

Graph Summaries for Optimizing Graph Pattern Queries on RDF Databases

Angela Maduko¹, Kemafor Anyanwu², Amit Sheth³, John Miller¹, Paul Schliekelman⁴,

¹Department of Computer Science, University of Georgia

²Department of Computer Science, North Carolina State University

³Kno.e.sis Center, Wright State University

⁴Department of Statistics, University of Georgia

maduko@cs.uga.edu, kogan@ncsu.edu, amit.sheth@wright.edu, jam@cs.uga.edu, pdschlie@stat.uga.edu

Abstract. The adoption of the Resource Description Framework (RDF) as a metadata and semantic data representation standard is spurring the development of high-level mechanisms for storing and querying RDF data. A common approach for managing and querying RDF data is to build on Relational/Object Relational Database systems and translate queries in an RDF query language into queries in the native language of the underlying system, typically SQL. The standard query paradigm for RDF is graph pattern matching which matches a query graph against a data graph. When translated into SQL, such queries involve join operations. To process join operations, the database query optimizer attempts to determine an optimal join order using a cost model which employs the expected cardinality of join results as a key parameter. This parameter is estimated from a statistical summary of the data that is maintained in memory. One limitation that arises with this approach is that it could lead to the selection of a sub-optimal query plan due to estimation errors, from the fact that the data summarization technique employed by database systems are oblivious of the graph structure of RDF data. In this work, we present data summarization techniques that take cognizance of the graph structure of RDF data for providing estimates of RDF graph patterns. Our approach is to take an RDF query and formalize it as a query graph pattern which itself may be made up of smaller graph patterns. These patterns may be viewed as a graph template which can match actual subgraphs in the RDF data. Our approach is divided into two phases, a preprocessing phase in which we create the summaries and a processing phase in which we use the summaries. Summarizations are produced by counting or estimating the frequency with which subgraphs match these patterns. We present and evaluate two basic summarization techniques and two pruning techniques for situations where the summaries are too large to fit the available memory budget. Once created these summaries supply values to query optimizers enabling them to produce better plans. These summaries are based on edge sequences which are linearized edges of subgraphs. In the first technique, rather than storing all the complete graph patterns with their corresponding frequencies, we reduce the storage by factoring out common sub-patterns and storing the summary as a prefix tree (or P-Tree). We may condense the P-Tree by pruning small subgraphs based on a valuation scheme that blends information about their importance and estimation power. In the second technique, we assume there are situations in which the edge occurrences on edge sequences are position independent (i.e., there is little correlation between the frequencies of neighboring edges). The independence assumption allows us to reduce the size of the summary even further. When applying either technique during query processing, we assume conditional independence to estimate the frequencies of larger patterns. We present extensive experiments on real world and synthetic datasets which confirm the feasibility of our approach. Our experiments are geared towards showing that the estimates obtained from the proposed summaries are effective for optimizing queries posed over RDF graphs.

Keywords: Cardinality estimation, Graph summaries, Data summaries, RDF Join Queries

1 Introduction

The growing interest in the Resource Description Framework (RDF) [30] with its accompanying schema language RDFS [6] as a metadata and semantic data representation standard in the Semantic Web is spurring on the development of large-scale storage and high-level querying systems for RDF data. As observed at the 2008 Semantic Technology conference, a number of commercial applications such as those from Talis and Twine already use and apply very large RDF datasets, and continue to seek techniques to manage ever increasingly larger datasets. In order to exploit the maturity and wealth of research that has been invested in Database Management Systems (DBMS), some of the proposed systems for storing and querying RDF [8][29][10] employ a Relational/Object Relational Database backend (also see [49] for related survey) that shred the RDF graph into relations. The two most common techniques of shredding RDF graphs into relations are the schema-oblivious technique and the schema-aware technique. The schema oblivious technique stores all triples in one

relation, while the schema-aware technique has one relation for each property type. Many of the proposed query languages for querying RDF, support graph pattern matching as the primary query paradigm [28][41][42]. In this paradigm, a query is a definition of a graph pattern that is to be matched in the database and the result of the query is the list of all occurrences or *matchings* of the graph pattern. Query processing techniques vary depending on the underlying storage techniques. For example, for systems that are built on relational database systems, queries are translated from an RDF query language such as SPARQL to SQL and then pushed to the underlying database for processing. One implication of the graph shredding storage approach with respect to query processing is that reconstructing subgraphs from database relations is done using relational join operations. Depending on the complexity of the query, several join operations may be required. Since multiple joins are required to compute larger subgraphs from the edges in the relations, there are several possible query plans for processing such queries due to the different possible orderings of join operations. Thus, query processing often involves a cost-based optimization step in which the query optimizer is faced with the task of using a cost model to attempt to choose an optimal query plan from amongst several alternatives. A key parameter of the cost model is the expected cardinality of the intermediate results obtained after the execution of each operator in the plan. The estimation of the cardinality of intermediate join results is equivalent to estimating the frequency of sub-patterns of a query graph pattern. These cardinalities are estimated using a statistical summary of the data maintained by the optimizer. As an illustration, the SPARQL [41] query (namespaces are omitted) of **Figure 1a** (with graph pattern shown in **Figure 1b**) asks for professors employed in universities in the United States who direct a Semantic Web project. If the schema-aware storage model (one relation for each relationship type) [50] is used for storing the data into edge relations, three join operations (shown in **Figure 1c**) will be needed to process this query. To optimize the query, the optimizer will need accurate estimates of the frequency of subpatterns of **Figure 1c**, such as those shown in **Figure 1d, e and f**.

In this work, we focus on efficient techniques for estimating the frequency of patterns in a graph database. Noting that (1) the number of possible subgraphs in a graph database could be exponential and (2) it is more expedient if the estimates are computed without disk accesses, since they are needed at optimization time, we focus on techniques that summarize subgraphs in the graph database so as to fit in the available memory budget. Obviously, such a summary will be useful only if it captures the interactions among subgraphs. Our choice of subgraphs for capturing interactions amongst graphs is strengthened by the observations in [51], [58] and [59], where it is shown that subtrees and subgraphs perform better than paths in capturing interactions among graphs and trees, respectively. However, the number of unique subgraphs greatly exceeds the number of paths in a graph. It is thus infeasible to examine all subgraphs, and efficient pruning techniques will be crucial for efficient query optimization. We propose two summarization techniques that differ in their representation and pruning techniques. The Pattern Tree (P-Tree) and the Maximal Dependence Tree (MD-Tree).

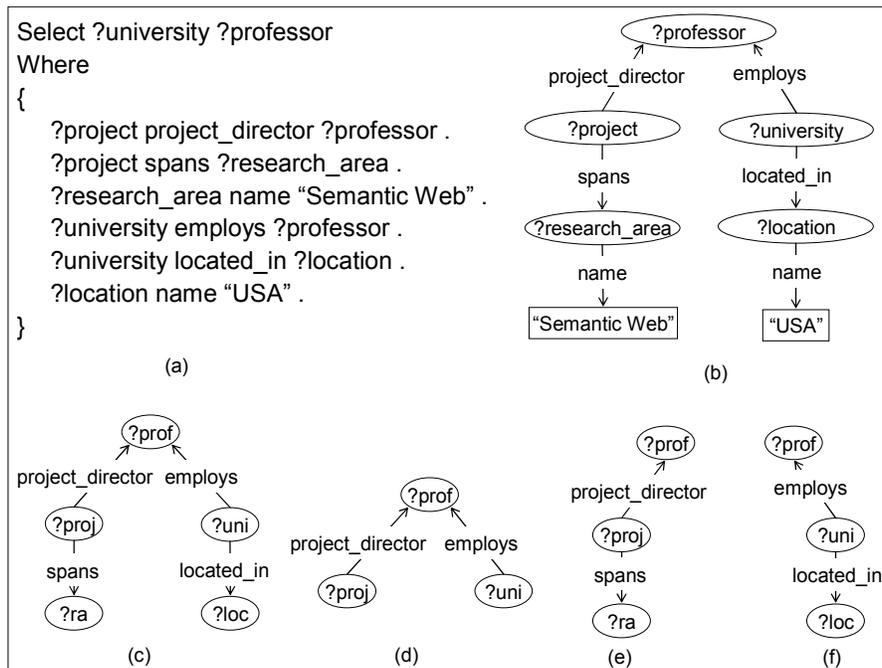


Figure 1: (a) A SPARQL query involving several join operations, (b) its corresponding graph pattern and (c), (d), (e), (f) sub-graph patterns.

Given a list of patterns and their frequencies, our P-Tree technique organizes the patterns in a prefix tree to save space needed for storing the patterns. The pruning technique of the P-Tree is based on two insights: (1) the frequency of a pattern may be estimated using a function of the frequency of its subpattern; and (2) information about the relative importance, as defined by the user, of subpatterns could lead to characterizing some as more worth while to estimate than others. For example, frequent subpatterns from a query workload are more important than infrequent ones for tuning purposes. We prune the P-Tree by blending the significance of patterns for estimation and for tuning purposes. The MD-Tree is based on the observation that edges in certain positions in patterns may largely determine their frequencies. The idea then, is to assess the edge position with the greatest influence on the frequencies of the patterns. If no position of great influence exists, we assume that edges occur independently at each position. The pruning technique of the MD-Tree is based on the observation that correlations or high-order statistical dependencies often exist among subgraphs. It may be prohibitive to keep all such dependencies; thus we attempt to capture the most informative dependencies in the given space. In particular, this work makes the following contributions:

- *Generic Representation for an RDF Graph Pattern.* With the multiple classification of resources allowed in the RDF data model, graph patterns may exist in the instance graph that are not explicitly represented in the schema. We create a *Semantic and Structural Summary* for RDF/S by adding structure information from the RDF instance graph to a Semantic Summary [3] created from the RDF schema.
- *Graph Sequentialization.* Using the *Semantic and Structural Summary*, we adapt a technique described in [57] for creating a partial ordering of edges in an undirected graph to obtain a variable length sequence for RDF graph patterns.
- *Summarization Framework.* Given that the number of unique graph patterns may be exponential, we develop heuristics for pruning patterns of size at most maxL (for an integer maxL) to fit a given memory budget.
- *Estimation Framework.* Using the proposed summaries, we show how to obtain estimates of the cardinality of both patterns that are represented in the summary and those that are not explicitly represented in the summary. The estimates obtained can be used to guide a query optimizer in choosing the optimal query plan.
- *Experimental Validation.* We conduct experiments on real-life and synthetic datasets which validate our approach. Our experiments show the efficiency of the proposed summaries in providing estimates that help in optimizing both unconstrained join queries and constrained join queries over RDF graphs.

The rest of this paper is structured as follows: In Section 2, we formally define the problem we address in this work and briefly discuss background work. Section 3 discusses the proposed summaries, while Section 4 presents the experimental evaluation of the summaries. In Section 5 we discuss related work and conclude the paper in Section 6.

2 Preliminaries

In this section, we give basic terminology used in this work and then we formulate the problem this work addresses.

Data Model. The RDF data model represents information in terms of “triples” of the form (entityA relationship1 entityB) where each triple captures a named binary relationship between two entities, e.g., entityA and entityB. For example, the triple (*X authorOf Z*) denotes a relationship between an entity represented by X (e.g., an author) and another entity or literal represented by Z (e.g., a publication). It is natural to view such triples using a graph model in which entities are represented as nodes and named relationships as labeled edges. Thus, an RDF database - a collection of triples – in graph theoretic terms is a directed labeled graph (possibly consisting of several connected components). The companion specification of RDF, RDFS, provides domain vocabulary for defining the types of entities, i.e., classes (e.g. Author) and relations (e.g. authorOf) in the domain. RDFS can also be represented as a directed labeled graph in which nodes are labeled with the names of classes and edges are labeled with names of relations. We now formalize this graph model.

Let C, P, I and L be the sets of class Universal Resource Identifiers (URIs), property URIs, instance URIs and literal values, respectively; the concepts of RDF schemas and instances can be formalized as follows.

Definition 1. RDF Schema Graph: An RDF schema graph $G_S = (V_S, E_S, \lambda_S, C, P)$ is a directed labeled graph where V_S is the set of nodes of G_S and E_S is the set of edges (i.e. subset of cartesian product of V_S) of G_S and λ_S :

$(V_S \cup E_S) \rightarrow C \cup P$ is a surjective labeling function that maps vertices and edges of G_S to class and property URIs, respectively, such that $\lambda_S(v) \in C$ and $\lambda_S(e) \in P$ for any $v \in V_S$ and $e \in E_S$.

Definition 2. RDF Instance Graph: An RDF instance graph $G_I = (V_I, E_I, \lambda_I, \tau, I, P, L)$ defined on a schema graph G_S is a directed labeled graph where V_I and E_I are sets of vertices and edges of G_I , respectively; $\lambda_I : (V_I \cup E_I) \rightarrow I \cup L \cup P$ is a surjective function that maps vertices and edges of G_I to instance URIs or literals or property URIs, respectively such that $\lambda_I(v) \in I \cup L$ for any $v \in V_I$ and $\lambda_I(e) \in P$ for any $e \in E_I$. On the other hand, $\tau : V_I \rightarrow 2^{V_S}$ is a function which maps nodes of G_I to sets of nodes of G_S . This typing of nodes is such that for any edge $e = (u, v)$ in G_I , if $\lambda_I(e) = p$, then there is an edge $e' = (u', v')$ in E_S for which $\lambda_S(e') = p$, $\tau(u) = \{u'\}$ and $\tau(v) = \{v'\}$.

Note that our model of the instance graph consists of only ground RDF graphs which do not contain reified statements (i.e., assertions about statements). Example 1 explains the schema and instance graphs given in Figure 2.

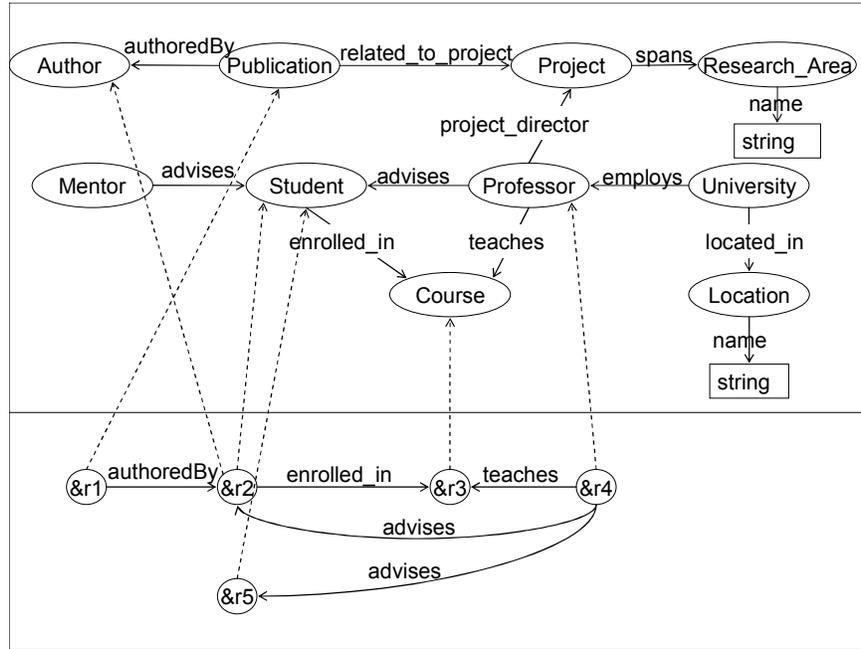


Figure 2: RDF Schema and instance graphs

Example 1. The top part of Figure 2 shows an RDF schema graph which describes the University domain. The class “Author” denotes the set of entities that are publication authors, while the class “Publication” denotes the set of entities that are publications. The relationship “authoredBy” asserts that instances of the class “Author” are related to instances of the class “Publication” by the “authoredBy” relationship. This is depicted in the instance graph shown in the bottom part of the figure. The dotted arrows denote a member of relationship from instances to classes so that entity &r2 is asserted to be an instance of the class “Author” as well as “Student”, while the entity &r1 is asserted to be a member of the class “Publication”. &r1 is also asserted to be related to &r2 by the “authorOf” relationship.

Definition 3. RDF Graph Pattern: An RDF graph pattern or simply a pattern $G_P = (V_P, E_P, \lambda_P, \alpha, P)$ is a connected edge-labeled directed graph where V_P is a set of nodes, $E_P \subseteq E_S$, $\lambda_P : E_P \rightarrow P$ and for $e \in E_P$, $\lambda_P(e) = \lambda_S(e)$ and $\alpha : V_P \rightarrow N$ maps nodes to the set of natural numbers.

Problem Formulation. This work focuses on exploiting information in an RDF graph to improve the statistics needed for cost-based query optimization of join queries posed over RDF data stores. In particular, we focus on improving the statistics needed for estimating the cardinality of equi-join operations. This amounts to obtaining information about the frequencies of patterns in an RDF graph.

Our approach is to model a query as a graph pattern defined by edge labels and topology (i.e., the nodes are considered variables and thus are not bound to any particular label).

Definition 4. Given an instance graph G_I and a query pattern G_P , we say that G_P has a *matching* in G_I if we can map every edge in G_P to an edge in G_I with the same label i.e., for $e \in E_P$, $\lambda_P(e) = \lambda_I(e')$ where $e' \in E_I$. This matching is such that the topology of the nodes and edges in G_P are preserved. Thus, a matching of a query pattern in an instance graph constitutes a result of the query.

G_{I1} and G_{I2} are two unique subgraphs in G_I matching G_P , if at least one node in G_{I1} is different (labeled with a different URI) from its corresponding node in G_{I2} . From the foregoing, it is clear that all unique matchings of a query pattern constitute the result of the query. Thus, our goal is to obtain the frequency or cardinality of all unique matchings of a particular pattern.

Example 2. *The pattern $1\text{---teaches}\text{---}2$, $1\text{---advises}\text{---}3$ has two unique matchings in the instance graph shown in the bottom part of Figure 2. The matchings are $\&r4\text{---teaches}\text{---}\&r3$, $\&r4\text{---advises}\text{---}\&r2$ and $\&r4\text{---teaches}\text{---}\&r3$, $\&r4\text{---advises}\text{---}\&r5$.*

Definition 5. The *frequency* of a query pattern is the number of unique matchings it has in the instance graph.

Unfortunately, it is prohibitive to count all matchings of a pattern from G_I at query optimization time so that the use of a summary of G_I is necessary. Thus, the problem we address in this work can be stated more precisely as follows:

Given an RDF schema graph, a corresponding instance graph, and a memory budget B , create a summary of size at most B , from which one can obtain accurate estimates of the number of unique matchings of a pattern from the instance graph.

2.1 Canonical Labeling of Graphs

In order to count the frequencies of patterns, we need an efficient way to enumerate all unique patterns. In other words, we need a canonical labeling scheme for patterns. In this section, we discuss the DFS Coding canonical labeling and the gSpan algorithm [57] which we adopt in this work for labeling and counting the frequencies of matchings of a pattern.

With minimal modifications, efficient pattern mining algorithms, such as gSpan [57] and [52] (which uses disk-based indexes for limited memory settings) can be used to discover subgraphs matching patterns and count their frequencies. In general, these techniques first develop a canonical label (i.e., unique code) for a graph; then subgraph frequencies are computed based on the canonical label. In the next paragraph, we briefly review the *minimum Depth First Search (DFS) code* [57] of a graph as its canonical label, which we adopt in this work.

DFS Coding. This technique uses a Depth First Search (DFS) traversal of a graph to transform it into an edge sequence called a DFS code. Each edge (u, v) in the graph is represented by a 5-tuple $\langle i, j, l_i, l_{(i, j)}, l_j \rangle$, where i and j are integers denoting the DFS discovery times of nodes u and v and l_i , l_j and $l_{(i, j)}$ are the labels of u , v , and the edge (u, v) respectively. The edges are then ordered by listing those in the DFS tree (tree edges) in the order in which they were discovered, then inserting the remaining edges into the ordered list as follows: Given a tree edge (u, v) , all non-tree edges from v are listed immediately after (u, v) ; if (u_i, v_j) and (u_i, v_k) are two such non-tree edges, (u_i, v_j) is listed before (u_i, v_k) only if $j < k$. Since a graph can have multiple DFS codes, the minimum, obtained based on a linear ordering of all its DFS codes is chosen as its canonical label.

gSpan. Given an undirected graph G and an integer F , the gSpan algorithm allows for enumerating all subgraphs of G whose frequencies are at least F , where the minimal DFS code of each subgraph is its canonical label. The algorithm iteratively generates and counts all unique subgraphs of length $i+1$ from those of length i , whose frequencies are at least F (the length of a subgraph is defined as the number of edges it contains). Each generated subgraph of length $i+1$ with frequencies less than F is pruned on the assumption that it cannot lead to the generation of any new frequent subgraph (i.e., with frequency at least F). Details of the algorithm can be found in [57].

3 Approach

In this section, we present our proposed techniques for obtaining more accurate estimates of intermediate result sizes for use by query optimizers by pre-computing summaries of the instance graph, as well as, developing

formulas to combine this summary data. We first discuss the Semantic and Structural Summary in section 3.1, then the Pattern Tree summary in section 3.2 and then in section 3.3, we discuss the Maximal Dependence Tree summary.

3.1 Semantic and Structural summary

Recall that at least a node URI distinguishes two unique matchings of a query pattern in the instance graph so that to summarize the matchings of patterns in the instance graph, we need a general representation that captures all unique matchings of any given query pattern. We achieve this by refining the RDF Schema graph to include all patterns that exist in the RDF instance but are not represented in the schema (we give explanation of this later) graph and derive our representation from the refined schema graph. In a previous work [3], we introduced the RDF *Semantic Summary* as a data structure that captures all possible paths which can be obtained from the RDF schema based on either explicit definitions in or deductions from the schema.

Definition 6. Given an RDF Schema graph $G_S = (V_S, E_S, \lambda_S, C, P)$, a Semantic Summary $G_{SS} = (V_{SS}, E_{SS}, \lambda_{SS})$ is a directed labeled graph where E_{SS} the set of edges of the graph is given by

$$E_{SS} = \{e_i \mid e_i \in E_S, \lambda_S(e_i) \neq \lambda_S(e_j), i \neq j\}.$$

V_{SS} the set of nodes of the graph is given by

$$V_{SS} = \{\text{domain}(u) \cup \text{domain}(v) \mid u \in E_S, v \in E_S, \lambda_S(u) = \lambda_S(v)\} \cup \{\text{range}(u) \cup \text{range}(v) \mid u \in E_S, v \in E_S, \lambda_S(u) = \lambda_S(v)\}$$

$$\lambda_{SS}(e) = \lambda_S(e) \text{ for } e \in E_S$$

In other words, if multiple edges have the same label, these edges are merged into one by merging their respective source and destination nodes. If u and v are two nodes of G_S merged into w and if u' and v' are two nodes of the instance graph that are types of u and v , u' and v' are types of w in G_{SS} . The edges emergent from or incident on a node w in G_{SS} is the union of the respective edges emergent from or incident on all nodes in G_S that were merged into w .

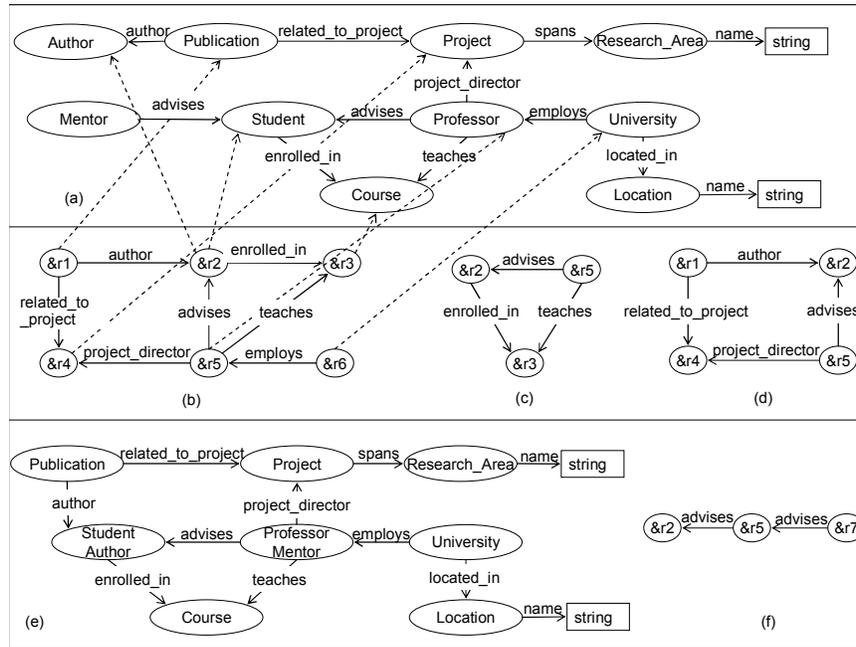


Figure 3: (a) Sample RDF Schema Graph, (b) Instance graph, (c) and (d) Graph Patterns, (e) Semantic and Structural Summary (f) Spurious Graph Pattern.

Although G_{SS} captures all possible paths defined in the RDF schema, it may not describe the structure of all patterns in the instance graph. This stems from the multiple classifications of resources allowed in the RDF data model. We enhance G_{SS} with this capability by merging nodes in G_{SS} which have any instance nodes in common. The resultant nodes are then annotated with unique integer ids. Each edge in G_{SS} is also annotated with a unique integer id. We refer to this data structure as an RDF *Semantic and Structural Summary*. Essentially, a

node in the Semantic and Structural Summary represents either a single node or a collection of nodes in the Semantic Summary. Example 3 illustrates the merging of nodes in the Semantic and Structural Summary to reflect all subgraphs in the instance graph.

Example 3. In the RDF schema graph shown in Figure 3a above, the property type “advises” is defined on the domains “Mentor” and “Professor”. A Semantic Summary of Figure 3a represents these two classes by the same schema object with respect to the property “advises” as shown in Figure 3e. Figure 3c and Figure 3d show two subgraphs drawn from the instance graph shown in Figure 3b. While the former subgraph is reflected in the schema, the latter is not because the resource &r2 is multiply classified as an “Author” and a “Student”. A Semantic and Structural Summary of Figure 3a will incorporate this subgraph by representing these two classes with the same schema object as shown in Figure 3e.

Formally, if $extent(u)$ is a function that gives the nodes in the instance graph G_I which are mapped to a node u in G_{SS} , we define the RDF Semantic and Structural Summary as follows.

Definition 7. RDF Semantic and Structural Summary: Given an RDF Semantic Summary $G_{SS} = (V_{SS}, E_{SS}, \lambda_{SS})$, an RDF Semantic and Structural Summary is a directed labelled graph $G_{SSS} = (V_{SSS}, E_{SSS}, \alpha, \lambda_{SSS})$ where $\lambda_{SSS} : E_{SSS} \rightarrow P$ is a labelling function that maps elements of E_{SSS} to the set of property URIs, V_{SSS} the set of nodes of the graph is given by

$$V_{SSS} = \{ v \in 2^{V_{SS}} \mid |v| > 1 \Rightarrow \bigcap_{i=1}^{|v|} extent(u_i) \neq \emptyset, u_i \in v \}$$

E_{SSS} the set of edges of the graph is given by

$$E_{SSS} = \{ e = (u, v) \mid u, v \in V_{SSS}, \text{ and } \exists e' = (u', v'), u' \in u, v' \in v \text{ and } \lambda_{SSS}(e) = \lambda_{SS}(e') \}.$$

$\alpha : (V_{SSS} \cup E_{SSS}) \rightarrow N$ is a numbering function such that $\alpha : V_{SSS} \rightarrow N$ and $\alpha : E_{SSS} \rightarrow N$ are injections.

In other words, if two nodes of the Semantic Summary have a common member in the instance graph, the Semantic and Structural Summary merges the two nodes. Also every edge of the Semantic and Structural Summary is either an edge of the Semantic Summary or an edge resulting from the merging of nodes of the Semantic Summary.

The Semantic and Structural Summary differs from the structural summaries described in [1][18][39] in the sense that it may contain properties defined in the RDF schema which do not exist at all in the instance base. We allow the representation of such properties so that updates to the instance graph that eventually introduce them can be accommodated gracefully. Furthermore, the Semantic and Structural Summary may introduce spurious paths and cycles in the data summary. For example, if resource &r5 in Figure 3b is also defined as an instance of the class *Author*, incorporating this in the Semantic and Structural Summary will cause the classes *Student*, *Author*, *Professor* and *Mentor* to be merged. The merged node will have the edge *advises* as a self loop so that the spurious pattern shown in Figure 3f is seen as valid. We will investigate possible solutions to this problem in the future.

Having discussed the Semantic and Structural Summary, we now direct our attention to our proposed MD-Tree and P-Tree summaries. To create each of the proposed summary, we begin by generating all patterns of length at most $maxL$ (where the length of a pattern is the number of edges it contains) and counting their frequencies using a slight modification of the gSpan algorithm with input graph $G_I = (V_I, E_I, \lambda_I, \tau, I, P, L)$ and the frequency threshold set to one. However, we represent each edge $e = (u, v)$ in G_I by a 5-tuple $\langle i, j, \lambda_i(e), \tau(u), \tau(v) \rangle$, where i and j are integers that denote the DFS discovery times of nodes u and v and label λ_i and type τ are the functions defined in section 2, where $\lambda_i(\cdot)$ and $\tau(\cdot)$ (i.e., the ranges of λ_i and τ) are integers in accordance with the RDF Semantic and Structural Summary. The sequence of edges/quintuples obtained after the algorithm is run represents the structure of patterns of G_I . While gSpan discusses the minimum DFS code in the context of undirected labeled graphs, for directed labeled graphs, we ignore edge directions during DFS traversal so as to maintain the connectivity of the graph, however, the directions are kept implicitly in the quintuples. Example 4 explains how to obtain the minimal DFS code of a pattern.

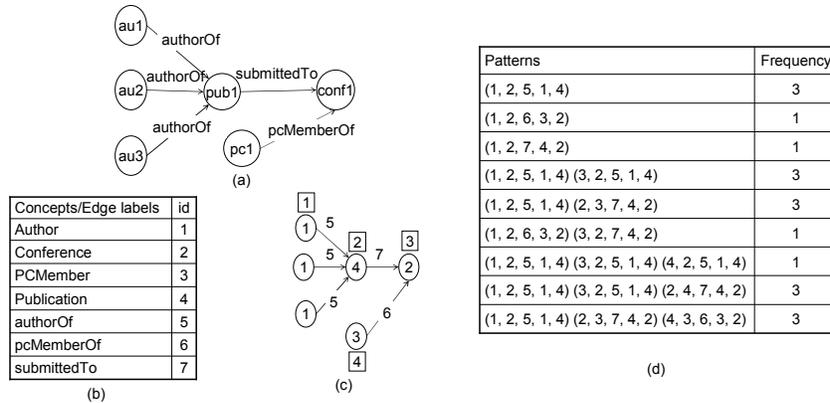


Figure 4: (a) Sample instance graph (b) Unique ids of concepts and edge labels (c) Modified instance graph reflecting unique ids of concepts and edges (d) Unique Edge Sequences for Patterns and their frequencies

Example 4. Figure 4a shows a directed labeled graph that models conference information. All publication authors are classified as members of the class “Author”, publications as “Publication” and so on. Concepts and edge labels are assigned integer ids as shown in Figure 4b. To obtain the edge sequence for the subgraph $au1$ $authorOf$ $pub1$, $pub1$ $submittedTo$ $conf1$, $pc1$ $pcpMember$ of $conf1$, we begin DFS with the edge $authorOf$ as it is lexicographically the smallest label. Traversal proceeds as shown by the boxed subscripts associated with the nodes to yield the pattern $(1,2,5,1,4)$ $(2,3,7,4,2)$ $(4,3,6,3,2)$. Note that the direction of the edge labeled “ $pcMemberOf$ ” is implicit in the sequence. Figure 4d shows all patterns of length at most 3 and their frequencies from the graph in Figure 4c

Note that, in our directed graph model, an edge, for example $(5, 1, 4)$ in Figure 4, may appear in a pattern of length at least two, in one of three possible directions: forward, as in $(1, 2, 5, 1, 4)$; backward, as in $(3, 2, 5, 1, 4)$; or potentially in a self-loop, as in $(2, 2, 5, 1, 4)$.

At this point, we have the necessary summary data to make accurate estimates for the sizes of joins (intermediate or final). If the number of joins is less than or equal to $maxL$, the estimate may be exact (and from empirical testing is always very close). Note, estimates are not always exact, since the $gSpan$ algorithm provides an efficient way to determine frequency counts, they may differ slightly from a more straightforward, though very inefficient, technique of pre-computing all joins. When the number of joins is greater than $maxL$, formulas must be used to provide somewhat less accurate estimates.

3.2 Pattern Tree (P-Tree)

Given the set of patterns and their frequencies, the P-Tree is a prefix tree representation of the patterns that achieves more compactness than the list of patterns. Its nodes are labeled with edge patterns (patterns of length 1) such that any pattern in P can be obtained by a concatenation of node labels on a path from the root. Also, each node is associated with the frequency of the pattern it represents.

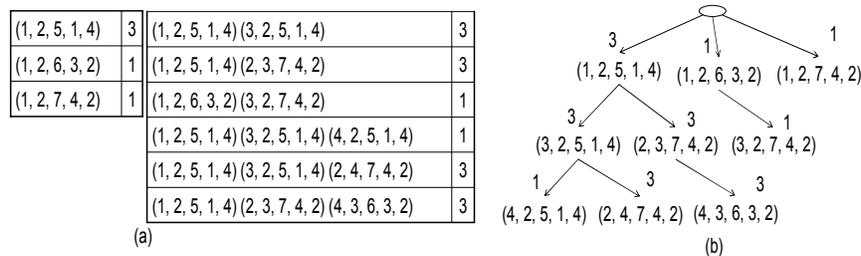


Figure 5: (a) Patterns and their frequencies (b) Pattern Tree

Having discussed the P-Tree, we now discuss how to prune the P-Tree to fit a memory budget. Given a pattern, we define the function *sequence* on patterns that gives the minimal DFS code of the pattern, i.e., the edge sequence of the pattern.

Definition 8. Given two patterns P_i and P_j with $\text{sequence}(P_i) = e_1, e_2, \dots, e_k$ and $\text{sequence}(P_j) = e'_1, e'_2, \dots, e'_m$ of lengths k and m , respectively, we say that P_i is a sub-pattern of P_j and respectively P_j is a super-pattern of P_i if $m > k$ and $e_i = e'_i$ for $1 \leq i \leq k$. Further, we say that P_i is a maximal sub-pattern of P_j and respectively P_j is a minimal super-pattern of P_i if k equals $m-1$.

Lemma 1. Given a set P of patterns of length at most maxL , every pattern P in P of length m , $m \geq 2$ has exactly one unique maximal sub-pattern in P .

Proof. The canonical representation of patterns is a linearization of the patterns so that every pattern of length at least 2 has only one maximal sub-pattern.

By Lemma 1, pattern P_i is the parent of pattern P_j in the P-Tree, if P_i is its maximal sub-pattern. The root of the tree is an empty node whose children are patterns of length 1. We illustrate this with an example.

Example 5. Figure 5b shows the P-Tree for the patterns in Figure 5a, which are the same patterns introduced in Figure 4a. Except for the root, each node is associated with an edge pattern and the frequency of the pattern it represents. Thus, the frequency of the pattern (1,2,5,1,4) (3,2,5,1,4) (4,2,5,1,4) can be obtained by traversing the leftmost branch of the tree.

Pruning the Pattern Tree.

To motivate the pruning of the P-Tree, we observe that given two patterns P_i and P_j which have almost the same edge patterns, the frequency of P_i may be within δ of that of P_j , where δ is a small non-zero positive integer. The idea of our P-Tree approach then is to identify sets of patterns which have almost the same edge patterns such that for a set say P , the frequencies of patterns in P are within δ of at least one pattern in P say P_i . Thus, given P_i , the frequencies of all other patterns in P can be estimated within δ error so that we can safely eliminate all other patterns in P from the summary. In this section, we discuss how we construct a P-Tree that fits the given summary budget.

If the size of the P-Tree exceeds the budget, it has to be systematically pruned so as to avoid a large increase in its overall estimation error. The question then is: *which nodes are to be pruned and in what order?* To answer this question, we develop the concepts of *preference* and *estimation* values of patterns. We begin by introducing some notations.

Notations.

We denote a set of patterns of length k by P_k and we define the function **freq** whose domain is the power set of patterns of length at most maxL and whose range is the set of positive integers, such that if X is a single pattern, $\text{freq}(X)$ maps to the frequency of the pattern. If however, X contains more than one pattern, $\text{freq}(X)$ maps to the sum of the frequencies of all patterns in X . Formally, we have that

$$\text{freq}(X) = \begin{cases} \text{frequency of } p \text{ in } G_1, p \in X & \text{if } |X| = 1 \\ \sum_{p \in X} \text{freq}(p) & \text{if } |X| > 1 \end{cases}$$

We define function **children** whose domain and range are the set of patterns of length at most maxL and the power set of patterns of length at most maxL , respectively, such that for a pattern P_i , $\text{children}(P_i)$ maps to the set of children of P_i in the P-Tree.

To prune the P-Tree, we consider that in some scenarios, certain patterns may be considered more important than others. For example, frequent patterns in a query workload, for tuning the summary. Our preference value for a pattern captures this notion.

Definition 9. Preference Value of a Pattern. Given a set of patterns $P = \{p_1, p_2, \dots, p_m\}$ with frequencies $(\text{freq}(p_1), \text{freq}(p_2), \dots, \text{freq}(p_m))$. Let $P_{PI} = (p_{PI1}, p_{PI2}, \dots, p_{PIm})$ be a vector such that $0 \leq p_{PIi} \leq 1$ for every $p_{PIi} \in p_{PI}$. If p_{PIi} defines the importance of pattern p_i , we define the preference value of p_i (p_{PV_i}) as the number of patterns which are less important than p_i , i.e., number of patterns p_j in P such that $p_{PIi} > p_{PIj}$.

In Definition 9 above, we do not assume any particular technique for computing the importance of a pattern. However, for the purpose of tuning the summary to favor frequent patterns, the importance of a pattern can simply be computed as the ratio of its frequency to that of the most frequent. Our testing did not include preference values, so we will defer further treatment of them until the appendix.

To motivate the estimation value of patterns, we note that if there is a match for a pattern p where $\text{sequence}(p) = e_1, e_2, \dots, e_k$ in the tree, its frequency $\text{freq}(p)$, is the integer associated with the matched node labeled e_k . If e_k is pruned in the tree, we guess $\text{freq}(p)$ from $\text{freq}(p')$, where $\text{sequence}(p') = e_1, e_2, \dots, e_{k-1}$ is the parent of p in the tree, under the assumption that children of p' in the tree have a uniform frequency distribution. Thus if p' has m children with total frequency N , the frequency of each child is estimated as N/m . When the children of p' are contracted, we compute and associate the ratio $N/(m \times \text{freq}(p'))$ with p' . We keep this ratio (rounded to the nearest integer) and not N/m for ease of frequency propagation as we will discuss later. We refer to this ratio as the *growth rate* of p' , denoted by p'_{GR} . We compute one growth rate for p' for all its children, to avoid overly increasing the size of the tree as patterns are pruned.

Our estimate for a contracted pattern p will be inaccurate if the frequency distribution of the children of its parent p' is not uniform. To exploit the uniformity assumption, we attempt to prune the P-Tree by deleting the children of patterns for which the assumption holds. To do this, we let the probability of the occurrence of any child p of a pattern p' be its proportion with respect to the total frequency of all children of p' (i.e., $\text{freq}(p)/N$, where N is the sum of $\text{freq}(p_j)$ for all p_j in $\text{children}(p')$). If the random variable Y defines the occurrence of a child of p' , we can measure the evenness of the probability distribution of Y using its *entropy* [44] $H(\text{Pr}_Y)$ given by:

$$H(\text{Pr}_Y) = - \sum_j \text{Pr}_Y(p_j) \log_2(\text{Pr}_Y(p_j))$$

The entropy of a probability distribution gives a sense of the unpredictability of any value from the distribution. The entropy of a probability distribution is maximized if the distribution is uniform. Thus, to measure how uniformly distributed Y is, we normalize $H(\text{Pr}_Y)$ by dividing by its maximum entropy. We denote this ratio as p'_{ENT} for pattern P' in the tree. If p' has one child, we set p'_{ENT} to 1.

Definition 10. Estimation Value of a Pattern. Given a set of patterns $P = \{p_1, p_2, \dots, p_m\}$ with frequencies $(\text{freq}(p_1), \text{freq}(p_2), \dots, \text{freq}(p_m))$ and some $\delta \geq 0$, the estimation value of p_i (p_{ENTI}) is given by:

$$p_{ENTI} = \frac{\left(\left| \{p_j \mid p_j \in \text{children}(p_i) \text{ and } |(\text{freq}(p_i) p_{GRi}) - \text{freq}(p_j)| \leq \delta \} \right| \right)^h}{|\text{children}(p_i)|}$$

By definition, for a pattern p_i , p_{ENTI} is at most 1. It is 1, if the frequency distribution of the children of p_i is uniform. If the exponent h is set to 1, the second term of the product measures how closely p estimates all its children with at most δ error. Thus, if p_i and p_j both have three children, if p_i estimates only two within δ error, while p_j estimates just one also within δ , this value will be higher for p_i (2/3) than for p_j (1/3). However, if p_i has six children and estimates only two within δ error, then the value will be 1/3 for both p_i and p_j , although p_i estimates more its children outside δ than p_j . To cope with this, we set h to 1.5 so that the numerator will not overly dominate the denominator. To find the optimal value for δ , i.e., the smallest integer for which the pruned P-Tree fits the given budget, we recursively increase a small integer exponentially in powers of 2 until enough patterns can be pruned to meet the budget. Suppose this occurs at $\varepsilon = 2^i$, then we find the optimal δ which lies between 2^{i-1} and 2^i using binary search.

As we noted before, when the size of the unpruned P-Tree exceeds the budget, we reduce the size of the P-Tree by systematically selecting some nodes of the P-Tree which are to be pruned. We compute the value of each node of the P-Tree (the combination of its preference and estimation values, or simply the latter if the former is not given) and select nodes to be pruned as shown in the *Prune P-Tree* algorithm shown below.

Prune P-Tree

Input: P-Tree T , Budget B , constant c

Output: Pruned P-Tree T

01 $\delta \leftarrow 0$; $\text{inc} \leftarrow 0$; Set estimable $\leftarrow \emptyset$; done \leftarrow false

02 **while** done = false **do**

03 estimable $\leftarrow \emptyset$

```

04   for each internal node  $v_i$  in  $T$  in bottom – up order do
05        $P_{E_{v_i}} \leftarrow \text{compute\_estimation\_value}(\delta)$ 
06       if  $P_{E_{v_i}} > 0$  then
07           insert all the children of  $v_i$  into estimable
08       end if
09   end for
10   if  $\text{sizeof}(T) - \text{sizeof}(\text{estimable}) \leq B$  and  $\delta$  is optimal then
11       compute estimation values of patterns in  $T$ 
12       prune all patterns in estimable
13       done  $\leftarrow$  true
14   else
15       adjust  $\delta$ 
16   end if
17 end while

```

Listing 1: Pattern Tree Construction Algorithm

As lines 4-8 of the Prune-Tree algorithm show, estimation values are used to select the patterns to be pruned. The running time of the algorithm $O(Ld^{\max L} \log(\max_i \{\text{freq}(p_i)\}))$, is reasonable since logarithm is a slowly growing function and $\max L$ will typically be small. The running time stems from the loops in lines 4-9 and 2-17. Lines 4-9 run in $O(Ld^{\max L})$ time, where L and d are the numbers of unique edge labels and the maximum degree of nodes in the graph respectively. Recall that the root of the P-Tree has a child for each unique edge label in the graph while internal nodes have at most d children. Lines 2-17 will be executed at most $\log(\max_i \{\text{freq}(P_i)\})$ times, since all patterns are estimable at $\delta = \max_i \{\text{freq}(P_i)\}$.

We illustrate the pruning process with an example. In the example below, we assume that information about the preference values of patterns are not provided.

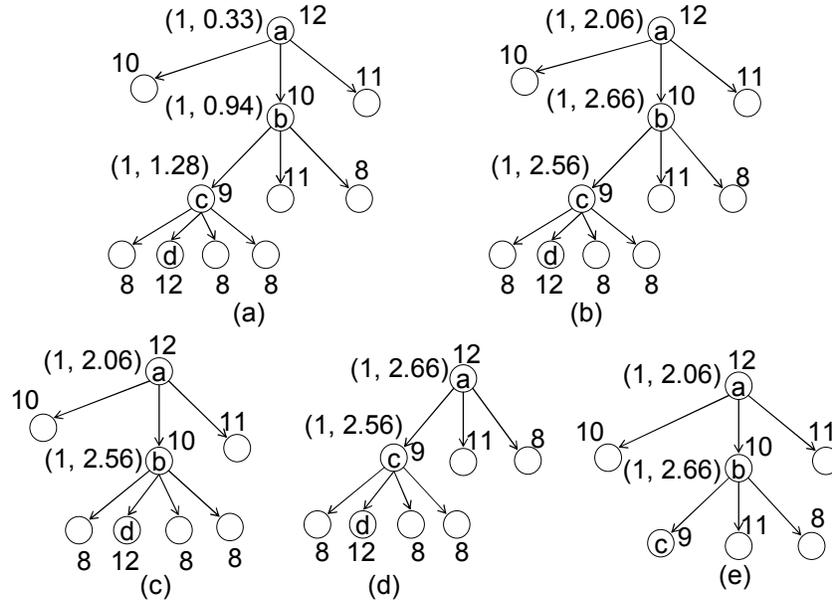


Figure 6: Pruning nodes of the Pattern Tree

Example 6. *Figure 6a* shows a subtree of a P-Tree. The 2-tuple (P_z, P_v) associated with each internal node is its growth rate and its value, computed at $\delta = 1$ and $c = 0$ and assuming no importance information is given so that $P_{P_{v_i}}$ is zero for all patterns. We show how the values are computed using node “a”. Its growth rate is given by $(10+10+11)/(3 \times 12)$, which rounds to 1, so its frequency (12) estimates that of one child (11). With exponent 1.5, the second term of the equation in Definition 10 is 0.333. The entropy of the frequency distribution of its children is given by $10/31 \times \log_2(31/10) + 10/31 \times \log_2(31/10) + 11/31 \times \log_2(31/11)$, or 1.583, with a maximum entropy ($\log_2(3)$) of 1.585 and ratio 0.999. Its estimation value is then $0.999(0.333) = 0.333$. In *Figure 6b*, the values are computed at $\delta = 2$. In *Figure 6c*, *d*, and *e*, the children of nodes *b*, *a*, and *c* have been pruned at $\delta =$

2, with total estimation errors of 4, 5, and 6. Our technique will result in the pruning of **Figure 6c** since *b* has the largest value.

We note that when a set of internal nodes children(P_i) with parent P_i are to be pruned, if the children of any node in children(P_i) have been pruned, the average of the growth rates of such nodes are computed and associated with P_i . Thus, a node P in the pruned P-Tree may have at most $\max L$ growth rates, ordered in increasing order of the original depths of their source in the P-Tree.

Frequency Estimation Using the Pattern Tree.

Given a pattern $G_p = (V_p, E_p, \lambda_p, P)$, we obtain its edge sequence $p = e_1, e_2, \dots, e_k$ and check that each pair of edges e_i and e_j in p is connected in the semantic and structural summary. If they are, we match p against the P-Tree. If we find a complete match for p in the Pattern Tree, we return the frequency of the matched node e_k in the P-Tree. If we find a partial match, we consider the last matched node v_j in the P-Tree. If it matches e_k , we return its associated frequency which is the exact frequency of p if no descendant of v_j was pruned. If it matches e_i $i < k$, we use its frequency to estimate that of the pruned node which originally matched e_k . We note that estimating the frequency of e_k requires estimating and propagating those of its $k-i-1$ immediate pruned ancestors. If $\xi_1, \xi_2, \dots, \xi_r$, $k \leq r \leq \max L$ are the growth rates associated with v_j and $\text{freq}(p_j)$ is the frequency associated with v_j , we estimate the frequency of e_k as:

$$\text{freq}(p_j) \times \prod_{r=1}^{k-i} \xi_r$$

For ease of exposition, we illustrate the estimation process with an example.

Example 7. To estimate the frequency of pattern (a, b, c, d) using the pruned P-Tree of **Figure 6c**, we find the partial match $(a, b, _, d)$, we return 12 since node *d* is matched. With the P-Tree of **Figure 6e**, we find the match $(a, b, c, _)$. We return 9, since the growth rate of *c* is 1.

3.3 Maximal Dependence Tree (MD-Tree)

To construct our MD-Tree summary, we adapted the Maximal Dependence Decomposition (MDD) technique that was proposed in [9]. The MDD was proposed as a technique that captures the most informative dependencies that exists amongst a set of DNA sequences of the same length, where it is most possible, due to limited data, to obtain a satisfactory estimate of all dependencies in the sequences. In our setting, we observe that even if it is possible to obtain accurate estimates of all dependencies that exist amongst a set of patterns of the same length, the space overhead that will be incurred in maintaining these dependencies will be very large. Our goal therefore, is to capture the most informative dependencies that exist amongst patterns of the same length, while incurring as small a space overhead as possible.

To motivate the MD-Tree, we observe that edges in certain positions in patterns may largely determine their frequencies. For example, the three patterns of length 3 in Figure 4d have the edge in the first (leftmost) position in common so that the edge in the second (middle) position, likely exerts a greater influence on the frequency of each of these patterns. The idea behind our adaptation of the Maximal Dependence Decomposition (MDD) approach [9] then, is to assess the edge position with the greatest influence on the frequencies of the patterns. If no position of great influence exists, we assume that edges occur independently at each position. Thus, given a set of patterns of length at most $\max L$, we capture the probabilities with which edges occur at any position $p \leq \max L$ in the patterns in a tree data structure, which we call the MD-Tree. The next paragraphs discuss how to construct the MD-Tree through an adaptation of the Maximal Dependence Decomposition [9] technique. First, we give notations used in this discussion.

Notations.

In addition to the notations introduced in section 3.2, we now introduce more notations used in the discussions in this section.

Given a instance graph $G_I = (V_I, E_I, \lambda_I, \tau, I, P, L)$, we let N_E denote the number of unique edge labels in G , i.e., N_E is size of the mapping $(\lambda_I(e), \tau(u), \tau(v))$, for each edge $e = (u, v)$ in E_I . Further, we let β be an integer in the range $[1, 3]$ such that: 1) β is 1 if all edges in E_I are forward edges (in a depth first search traversal of the graph) only, 2) β is 2 if in addition to forward edges, E_I also contains backward edges only, 3) β is 3 if in addition to forward edges, E_I also contains self loop edges only or both backward and self loop edges. This is based on the

assumption that a graph which contains a self loop edge is likely to contain a backward edge. We denote the probability with which edge e_i (its label and directionality) occurs at position j in the pattern by $\Pr(e_i, j)$.

Having established the notations used in the discussion, we begin the discussion of the MD-Tree by introducing our notion of *edge occurrence probability matrix*. Given a set P_K of patterns of length k , suppose it is known that edges occur independently at any position in the patterns, then we can estimate the frequency of any pattern P_j with edge sequence (e_1, e_2, \dots, e_k) as the product of the total frequency of patterns in P_K and the probability with which an edge e_i occurs at position j that is:

$$\text{freq}(P_j) = \text{freq}(P_K) \times \Pr(e_1, 1) \times \Pr(e_2, 2) \times \dots \times \Pr(e_k, k)$$

To this effect, we pre-compute and maintain these probabilities in a probability matrix where $\Pr(e_i, j)$ is given by the ratio of the total frequencies of patterns in P_K such that edge e_i occurs at position j and the total frequencies of all patterns in P_K , i.e.,

$$\Pr(e_i, j) = \frac{\text{freq}(\{p \in P_K \text{ and edge } e_i \text{ occurs at position } j\})}{\text{freq}(P_K)}$$

Definition 11. Edge Occurrence Probability Matrix. A probability matrix PM_K for a set P_K of patterns of length k with $k > 0$, is a $\beta N_E \times k$ matrix whose rows represent the possible edge sub-patterns that may appear in any pattern in P_K and whose columns represent the positions in which the edge sub-patterns may occur. The $(i, j)^{\text{th}}$ entry of PM_K contains the probability that edge e_i occurs at position j .

To construct the probability matrix PM_K for patterns in P_K , we first obtain the row indices by assigning unique integer ids to the possible edge sub-patterns that may appear in any pattern of length at least two, in multiples of β , corresponding with the possible edge directions (forward, backward and self-loop). To ascertain that edge patterns are uniquely identified, we assign integer x to edge sub-pattern e_i such that if x modulo β is zero, one or two then x identifies e_i in the forward direction, backward direction or self-loop, respectively. The column indices are simply the k positions in which edge patterns may occur. Next, if edge sub-pattern e_i is assigned integer x , and e_i occurs at position j , we store $\Pr(e_i, j)$ in cell (x, j) of PM_K . Example 8 explains how the probability matrix in Figure 7c is created.

Example 8. Figure 7a below shows the set P_2 of patterns of length two from Figure 4d. The dimension of the probability matrix PM_2 for patterns in P_2 is 6×2 (i.e. β is 2, since there are no self loop edges). To construct PM_2 , we assign integer ids to the edge types $(5, 1, 4)$, $(6, 3, 2)$ and $(7, 4, 2)$ as shown in Figure 7b. Next, we compute the entries for each cell (i, j) in PM_2 as shown in Figure 7c. Thus, cell $(1, 1)$ holds the probability that edge type $(5, 1, 4)$ occurs in a forward direction at position 1 in the patterns, i.e., $6/7$, cell $(2, 2)$ holds the probability that edge type $(5, 1, 4)$ occurs in a backward direction at position 2 of the patterns, i.e., $3/7$ etc. Under the independence assumption, the frequency of the pattern $(1, 2, 5, 1, 4)(3, 2, 5, 1, 4)$ is estimated as $7(6/7)(2/7)$, i.e., $12/7$, which rounds to 2.

(1, 2, 5, 1, 4) (3, 2, 5, 1, 4)	3
(1, 2, 5, 1, 4) (2, 3, 7, 4, 2)	3
(1, 2, 6, 3, 2) (3, 2, 7, 4, 2)	1

(a)

1	(5, 1, 4)
2	(5, 1, 4)
3	(6, 3, 2)
4	(6, 3, 2)
5	(7, 4, 2)
6	(7, 4, 2)

(b)

	1	2
1	6/7	0
2	0	3/7
3	1/7	0
4	0	0
5	0	3/7
6	0	1/7

(c)

Figure 7: (a) Patterns of length 2 (b) Integer ids of edge types in the forward and backward directions (c) A Probability Matrix for Patterns of Length 2

If the independence assumption holds for all sets of patterns of length k , $1 \leq k \leq \max L$, then probability matrices $PM_1, PM_2, \dots, PM_{\max L}$ will suffice for estimating the frequency of these patterns. This assumption gives rise to our notion of the *Base MD-Tree*.

Definition 12. Base MD-Tree. Given the sets $P_1, P_2, \dots, P_{\max L}$ of patterns of length at most $\max L$, a base MD-Tree for the patterns in P_i $1 \leq i \leq \max L$ is a triple (r_T, V_T, E_T) where $r_T \in V_T$ is the root of the tree and V_T and E_T are the sets of nodes and edges of the tree, such that $|V_T - r_T| = |E_T| = \max L$. All nodes in $V_T - r_T$ are ordered children of r_T such that child i is associated with the probability matrix PM_i , for patterns in P_i . Each edge (r_T, i) is labeled with $\text{freq}(P_i)$, the total frequency of all patterns in P_i .

Example 9. Figure 8a shows the base MD-Tree for the patterns of Figure 4d.

Under the assumption that edge types appear independently in patterns, the base MD-Tree holds sufficient information for estimating the frequency of patterns. As we saw in Example 8, this assumption may not always hold. Thus, a refinement process on the base MD-Tree is required to capture dependency information that may exist among edge types with respect to the positions at which they occur in the patterns.

Given P_K , the set of all patterns of length k , suppose it is known that the occurrence of any edge at position j , $1 \leq i \leq k$ and $i \neq m$, depends on the edge at position m . We estimate the frequency of a pattern P_j with edge sequence (e_1, e_2, \dots, e_k) in P_K as:

$$\text{freq}(P_K) \times \Pr(e_m, m) \times \prod_{i=1, i \neq m}^k \Pr((e_i, i) | (e_m, m))$$

where $\Pr((e_i, i) | (e_m, m))$ is the conditional probability that e_i occurs at position i given that e_m occurred at position m . To refine the base MD-Tree to reflect this dependence, let node v_K be the node in the base MD-Tree that is associated with probability matrix PM_K . We create βN_E ordered children nodes rooted at v_K , where each child is associated with a new probability matrix PM_K of dimension $\beta N_E \times k-1$. Cell (a, b) of the probability matrix associated with child i of v_K contains the conditional probabilities that the edge type with integer id “a” occurs at position “b” given that the edge type with integer id “i” occurs at position m . The i th edge from v_K to its i th child is labeled with the probability that the edge type with integer id “i” occurs at position m in patterns

in P_K . After the refinement process, the probability matrix associated with node v_K is deleted and the integer “m” is associated with v_K .

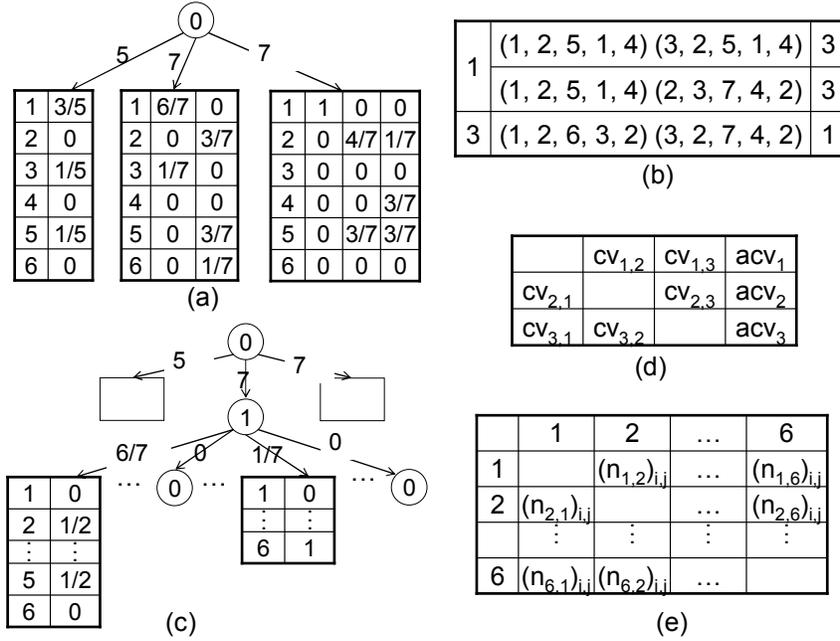


Figure 8: (a) The base MD-Tree to be refined (b) Patterns of length 2 with non-empty edge types at position 1 (c) The refined MD-Tree (d) Chi-square test statistic values (e) Computing the chi-square test statistic value

For ease of exposition, we illustrate the refinement process with an example.

Example 10. Given the base MD-Tree of Figure 8a, suppose we know that edge types occurring at position 1 for patterns of length 2 has the greatest influence on the edge types at position 2, we refine the second child node (which we denote as v_2 in this example) of the root of the base MD-Tree as follows. First, we create 6 new children nodes for node v_2 , one for each edge type. Next, we obtain the length 2 patterns which were used to create the probability matrix associated with v_2 i.e. the patterns shown in Figure 7a. Next, we partition these patterns with respect to the occurrence of the 6 edge types at position 1. As shown in Figure 8b, only the partitions for edge types 1 and 3 are non-empty. Using the patterns in the partition 1 and 3, we create two new probability matrices which we associate with child nodes 1 and 3 of v_2 respectively. Figure 8c shows the MD-Tree after the refinement process.

When one or more nodes of a base MD-Tree are refined as described above, we call the modified tree a *refined MD-Tree*.

Definition 13. Refined MD-Tree. A refined MD-Tree for the sets $P_1, P_2, \dots, P_{\max L}$ of patterns of length at most $\max L$, is a triple (r_T, V_T, E_T) where $r_T \in V_T$ is the root of the tree, V_T and E_T are the sets of nodes and edges of the tree, respectively. The set of nodes V_T can be partitioned into two disjoint non-empty sets V_{Tleaf} and $V_{Tnon-leaf}$ such that every node in $v \in V_{Tleaf}$ or $V_{Tnon-leaf}$ is a leaf node or non-leaf node, respectively. Further, a node v is associated with a probability matrix if and only if v is a leaf node and every non-leaf node $v, v \neq r_T$ has exactly βN_E children. The root node r_T has exactly $\max L$ children.

Finding the Position of Maximal Dependence.

As we noted earlier, we refine a node of the base MD-Tree by determining the position in patterns of length k that exerts the most influence on the others. We now discuss how we determine this position.

Given the set P_K of patterns of length k , we find the position that most influences the others by performing *chi-square association tests* for edge sub-patterns that occur at all pairs of positions i and $j, 1 \leq i, j \leq k, i \neq j$, using the chi-square test statistic given by:

$$\chi_0^2 = \sum_{m=1}^{\beta_{NE}} \sum_{n=1}^{\beta_{NE}} \frac{(O_{m,n} - E_{m,n})^2}{E_{m,n}}$$

In the equation above, $O_{m,n}$ is the sum of the frequency of patterns in P_K for which edge sub-patterns m and n occur at positions i and j , respectively and $E_{m,n}$, the expected mean of $O_{m,n}$ is given by:

$$E_{m,n} = \left(\sum_{a=1}^{\beta_{NE}} O_{m,a} \sum_{b=1}^{\beta_{NE}} O_{b,n} \right) / \sum_{a=1}^{\beta_{NE}} \sum_{b=1}^{\beta_{NE}} O_{a,b}$$

Thus, for the set P_K of patterns of length k , we find the position that most influences the others by performing chi-square association tests for edge sub-patterns which occur at all pairs of positions i and j , $1 \leq i, j \leq k$, $i \neq j$. For each position i , we sum the values obtained from the chi-square association tests for positions i and j , $1 \leq j \leq k$, $j \neq i$. We refer to this sum as the *aggregated chi-square value (ACV)* of position i . If m is the position which has the maximum ACV (MACV), then position m exerts the greatest influence on all other positions, if at least one of the chi-square values of positions m and j , $1 \leq j \leq k$, $j \neq m$, is statistically significant, i.e., if the probability that positions m and j are associated by chance is less than a given level of significance. We clarify this with an illustration.

Example 11. *Suppose we wish to find a position of maximal dependence for a set of patterns of length 3. First, we create a 3×4 matrix as shown in Figure 8d. To compute $CV_{1,2}$ for example, we create a 6×6 matrix as shown in Figure 8e, such that cell i,j contains the number of times the edge types with integer ids i and j occur at positions 1 and 2 in the patterns respectively. $CV_{1,2}$ is then the value of the chi-square test statistic for the 6×6 matrix. Next, we store the sum of all entries of each row of the 3×4 matrix to the fourth column, then we find the maximum of the sums over the three rows. Suppose the maximum sum is ACV_2 , we conclude that position 2 has the greatest influence on others but only if at least one of $CV_{2,j}$ is statistically significant*

Definition 14. Significant Node. Let v be a leaf node of a base or refined MD-Tree T , let PM_K be the probability matrix associated with v and let P_K be the set of patterns from which PM_K was created. We say that v is a significant node in T if there is a position j of maximal dependence in the patterns in P_K .

Definition 15. Complete MD-Tree. A complete MD-Tree is a refined MD-Tree T that has no significant leaf nodes.

Finding an Optimal MD-Tree.

From the foregoing discussion, it is obvious that the complete MD-Tree is the ideal choice amongst MD-Trees for obtaining accurate estimates of the frequency of patterns. However, the size of the complete MD-Tree may exceed the budget. Thus, given the sets of patterns $P_1, P_2, \dots, P_{\max L}$, our *optimal MD-Tree* is a refined MD-Tree whose size does not exceed B and from which the best estimates of the frequency of patterns in $P_1, P_2, \dots, P_{\max L}$ can be obtained. The idea is to choose a sub-tree of the complete MD-Tree that fits the budget and maximizes the MACV values of the refined nodes in the sub-tree. Our objective of selecting nodes with high MACV values is because a high MACV value indicates a strong position of maximal dependence. Given a complete MD-Tree (r_T, V_T, E_T) , let $S = (s_{v_1}, s_{v_2}, \dots, s_{v_m})$ be the size increment induced on the MD-Tree when node v_i was refined. Note that s_{v_i} is zero for the root and leaf nodes since they are not refined. Also let $I = (i_{v_1}, i_{v_2}, \dots, i_{v_m})$ be the impact of node v_i , given by $MACV_i / C_{v_i}$, rounded to the nearest integer, where C_v is the number of columns of the probability matrix associated with v_i . We normalize the MACV values to avoid favoring nodes of larger patterns. Note also that i_{v_j} is zero for the root and leaf nodes since they are not refined. The problem is to find a tree $T' = (V', E')$, $V' \subseteq V_T$ and $E' \subseteq E_T$ rooted at r_T , such that $\sum_j (s_{v_j}) \leq B$ and $\sum_j (i_{v_j})$ is maximized. This problem is an instance of the Tree Knapsack Problem (TKP) which is known to be NP-hard. Given x_j , an indicator variable with value 1 if v_j is selected as part of the optimal solution or 0 otherwise, TKP is formulated as the following integer programming problem: Maximize

$$\sum_j^m i_{v_j} x_j$$

constrained on

$$\sum_j^m s_{v_j} x_j \leq B, x_{\text{pred}(j)} \geq x_j$$

where $\text{pred}(j)$ denotes the predecessor (parent) of j in T' . With this reformulation, several pseudo-polynomial time solutions based on dynamic programming (DP) have been proposed such as [12] with $O(|V'|B)$ running time. We employ a greedy approximation with $O(|V'|)$ running time. Given the complete MD-Tree, the vectors S and I and the summary size budget B , our greedy approximation creates the tree $T' = (V', E')$ by choosing maximal impact subtrees of T' which fit the budget, as shown in the *Prune MD-Tree* algorithm.

Prune MD-Tree

Input: Tree Knapsack TK, Budget B

Output: Pruned Tree TK

```

01  $V = (v_1, v_2, \dots, v_m)$  //vector of nodes of TK in decreasing order of impacts
02 tree_capacity  $\leftarrow 0$ 
03 while  $V$  not empty do
04    $v \leftarrow$  extract first element of  $V$ 
05    $v\_ancestors\_size \leftarrow$  sizes of  $v$  and its ancestors
06   if tree_capacity +  $v\_ancestors\_size \leq B$  then
07     initialize node  $v'$  to  $v$ 
08     while  $u \neq r_T$  do
09       contract  $v'$  into  $u$  for edge  $(u, v')$ 
10       delete  $v'$  from  $V$ 
11       increment tree_capacity with  $v\_ancestors\_size$ 
12   else
13     prune subtree at  $v$ , delete its nodes from  $V$ 
14 end while

```

Listing 2: MD-Tree Construction Algorithm

Frequency Estimation Using the MD-Tree.

Given an optimal MD-Tree (r_T, V_T, E_T) , we define λ_v as a function that maps nodes in V_T to integers (for internal nodes) or probability matrices (for leaf nodes) with which they are associated. The integers associated with internal nodes denote the position at which the node was split. We also define λ_E as a function which maps edges in E_T to integers (for edges emanating from the root) or real numbers (for all other edges) with which they are associated. The integer on edge i emanating from the root denotes the total frequency of patterns of length i , while the real numbers on all other edges denote the conditional probability of the occurrence of its incident node given the edge pattern at the split position. We also define the function id on edge patterns that return the integer id assigned to the edge type of the pattern.

Let $P = e_1, e_2, \dots, e_k$ be the edge sequence of a graph pattern $G_P = (V_P, E_P, \lambda_P, P)$ of length k . To estimate the frequency of P , we check that each pair of edges e_i and e_j in P is connected in the semantic and structural summary. If so, beginning from the k^{th} child (say v) of r_T , we estimate $\text{freq}(P)$ as:

$$\text{freq}(P) = \lambda_E(r_T, v_1) \left(\prod_{i=1}^j \lambda_E(v_i, v_{i+1})_{id(e_{\lambda_v(v_i)})} \right) \left(\prod_{r=1, r \notin S}^k (\lambda_v(v_{j+1}))_{(id(e_r), r)} \right)$$

In the product above, given an edge $(v, v')_r$ subscripted with r , the subscript r denotes the r^{th} edge of node v . The integer j is the number of edges of the optimal MD-Tree found on the path from the root to a leaf node as defined by the subscripts on the edges, so that the node v_{j+1} is a leaf. The subscripts (r, r') are integer indices for accessing cell (r, r') of the probability matrix associated with node $\lambda_v(v_{j+1})$. The set S holds labels of all nodes on the path from R_T to v_{j+1} so that at v_{j+1} , any integer in the range $[1, \text{maxL}]$ not in the set S did not label any node on this path. We keep elements of S sorted as they are inserted, so that at the leaf v_{j+1} , we can check for elements of $[1, \text{maxL}]$ not contained in S in $O(\text{maxL})$ time. The depth of the MD-Tree is at most maxL ; thus the time complexity for estimating pattern frequencies is given by:

$$O(\text{maxL} \log(\text{maxL}))$$

We illustrate the estimation process by an example.

Example 12. To estimate the frequency of the pattern $P = e_1, e_2$ of length 2 given by $(1,2,5,1,4)$ $(3,2,5,1,4)$ from the MD-Tree of Figure 8c, we first access the second child node of the root. Let this node be denoted v_1 . Since $\lambda_V(v_1) = 1$, we insert 1 into set S and set the frequency of P ($\text{freq}(P)$) to $\text{freq}(P_2)$ which is 7. Recall from Figure 8b that $\text{id}(e_1) = \text{id}(1, 2, 5, 1, 4)$ is 1. So, we access the node on which the first edge of v_1 is incident. Let this node be v_2 . Next, we multiply $\text{freq}(P)$ by $\lambda_E(v_1, v_2)$ given by $6/7$, resulting in 6. Then, we obtain $\lambda_V(v_2)$ which yields the probability matrix PM_2 associated with v_2 which must have only one column. Since the set S contains the integer 1, the lone column of PM_2 must index position 2 of patterns in P_2 . Further, $\text{id}(e_2) = \text{id}(3,2,5,1,4)$ is 2, thus we access cell which represents the index $(2, 2)$ in PM_2 to obtain $1/2$. We then multiply $\text{freq}(P)$ which is currently 6 by $1/2$ to obtain 3, the final estimate for the frequency of P .

3.4 Estimating the Frequency of Large Patterns.

Given a large pattern whose length is greater than maxL , the idea is to use smaller patterns whose frequencies are known, to estimate the frequency of the large pattern. Thus we break up the large pattern into smaller patterns whose frequencies are known and then combine the frequencies of the small patterns to obtain an estimate for that of the large pattern. There are several possible ways of breaking up the large pattern starting from one in which the smaller patterns are disjoint to one in which the smaller patterns have overlaps in all but one edge. In [1], the authors demonstrated that breaking up a large path into non-disjoint smaller paths that intersect in all but one edge work well in estimating the frequency of large path. Further, the authors in [17] have also shown that breaking up a large twig into non-disjoint smaller twigs which intersect in all but one edge work well in estimating the frequency of the large twig. In the light of these works, we estimate the frequency of a large pattern by breaking it up into non-disjoint smaller patterns which intersect in all but one edge. We break up a large pattern by a traversal that visits all edges in the pattern as few times as possible. More formally, given a large pattern $G_P = (V_P, E_P, \lambda_P, P)$ with $|E_P| > \text{maxL}$, as always, we check that its edge sequence $p = e_1, e_2, \dots, e_k$ is connected in the semantic and structural summary. If so, we break up G_P into $G'_1, G'_2, \dots, G'_{|E|-\text{maxL}+1}$ non-disjoint connected patterns such that G'_i intersects G'_{i-1} in all but one edge. Let G''_i denote the intersecting edges of G'_i and G'_{i-1} . Next, we obtain the edge sequences $p'_1, p'_2, \dots, p'_{|E|-\text{maxL}+1}$ and $p''_2, p''_3, \dots, p''_{|E|-\text{maxL}+1}$ for the patterns $G'_1, G'_2, \dots, G'_{|E|-\text{maxL}+1}$ and $G''_2, G''_3, \dots, G''_{|E|-\text{maxL}+1}$, respectively. Like in [17], we assume conditional independence to estimate the frequency of G_P as follows:

$$\text{freq}(p) = \text{freq}(p'_1) \times \prod_{r=2}^{|E|+\text{maxL}-1} \frac{\text{freq}(p'_r)}{\text{freq}(p''_r)}$$

Since G_P may be broken up into $G'_1, G'_2, \dots, G'_{|E|-\text{maxL}+1}$ in several different ways, we select that for which frequency estimates of the patterns $p'_1, p'_2, \dots, p'_{|E|-\text{maxL}+1}$ are obtained along the deepest paths, i.e., paths which have the maximum total split nodes in the MD-Tree or along paths with the least pruned nodes in the P-Tree.

Example 13. Figure 9a shows a large pattern of length 4. The exact matchings of this pattern in the sample graph database shown in Figure 9b is 4. Suppose maxL is 2, we estimate this pattern by subdividing it into the sub-patterns (a) $1\text{---head}\rightarrow 2, 2\text{---teacherOf}\rightarrow 3$, (b) $2\text{---teacherOf}\rightarrow 3, 2\text{---author}\rightarrow 4$, (c) $2\text{---author}\rightarrow 4, 4\text{---inJournal}\rightarrow 5$. Sub-pattern (a) has 3 matchings, (b) has 5 matchings and (c) also has 5 matchings. Sub-pattern (a) overlaps with (b) on the edge $2\text{---teacherOf}\rightarrow 3$ which has 3 matchings, while sub-pattern (b) overlaps with (c) on the edge $2\text{---author}\rightarrow 4$ which has 6 matchings. Therefore, the estimate for the pattern is $3(5/3)(5/6)$ which is 4.167 which rounds to 4.

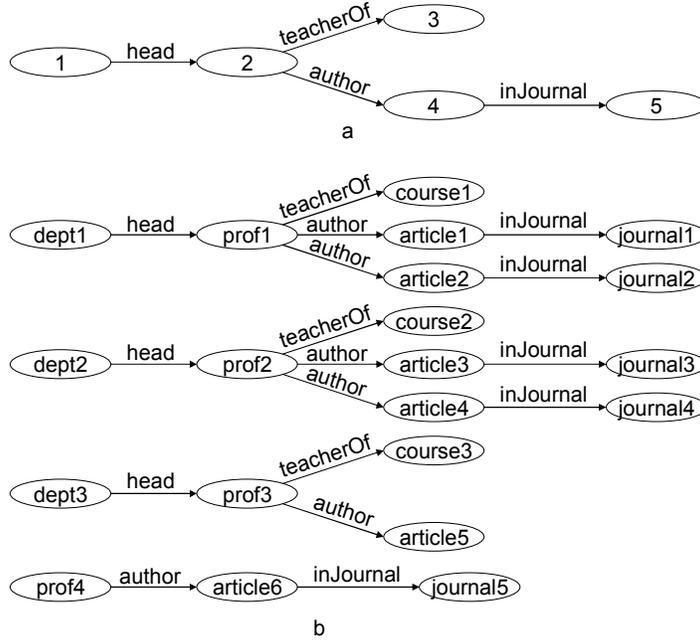


Figure 9: (a) A Large pattern and (b) A Sample Graph Database

Using the example shown in Figure 9, we compare our combining function to the traditional database join formula given by $N_A N_B / \max(V_A, V_B)$ where N_A is the number of tuples in relation A, N_B is the number of tuples in relation B, V_A and V_B are the numbers of distinct join attribute values in relation A and relation B, respectively.

Example 14. For the traditional database formula, there are several ways to subdivide the pattern of Figure 9a. We show only 2 ways here. One possible subdivision is $1\text{---}head\text{---}2, 2\text{---}teacherOf\text{---}3$ and $2\text{---}author\text{---}4, 4\text{---}inJournal\text{---}5$. The first sub-pattern has 3 matchings, while the second sub-pattern has 5 matchings and node 2 which is the join node (prof) has 3 distinct matchings for both. Applying the formula, we have $3(5/3)$ which is 5. Another possible subdivision is $1\text{---}head\text{---}2, 2\text{---}teacherOf\text{---}3$, $2\text{---}author\text{---}4$ and $4\text{---}inJournal\text{---}5$. The first sub-pattern has 3 matchings, the second (author relation) has 6 matchings and the third (inJournal relation) has 5 matchings. Node 2 (prof) which is the join node for the first and second sub-patterns has 3 matchings for the first and 4 matchings for the second. Node 4 (article) which is the join node for the second and third sub-patterns has 6 matchings for the second and 5 matchings for the third. Applying the formula recursively we have $3(6/4)(5/6)$ which is 3.75 which rounds to 4.

4 Experimental Evaluation

In this section, we present the results of the experimental study we conducted based on the proposed techniques described in this paper. The goal of our experimental evaluation is to show that the estimates obtained from the proposed summaries are indeed more accurate than commonly used techniques and more importantly effective for optimizing join queries posed over RDF graphs. We also demonstrate that the pruned summaries perform nearly as well as the unpruned summaries in terms of accuracy and query optimization. We show this by comparing the running time of several join queries optimized using our techniques to those optimized using other techniques.

4.1 Methodology

Techniques. We considered the following techniques in our experimental study, depending on how cardinality estimates are obtained for query optimization:

- **Proposed Graph Summaries (GSummaries):** We implemented the P-Tree and MD-Tree summaries proposed in this work. We implemented a main memory execution engine that uses the classic dynamic programming technique [43] for enumerating query plans. We used the schema-aware model [50] for shredding RDF graphs into relational tables.
- **Simple Greedy Technique (SGT):** We implemented a simple greedy technique that builds a left deep query plan by beginning with the relationship with the smallest cardinality and extending to its adjacent relationships using an algorithm similar to that for growing a minimum spanning tree. This is the technique used in the SPARLeR [31] engine for generating query plans.
- **Traditional Query Optimization Technique (TQOT):** We also implemented another execution engine that simulates a main memory relational database execution, using the same dynamic programming enumeration as the execution engine for our graph summaries. The difference however is in the way the result cardinality estimates are obtained. While for our proposed techniques, the graph summaries supply the cardinalities, for the main memory relational database, we used the well known uniformity assumptions employed by relational databases to compute the cardinality estimates. If n_1 and n_2 are the sizes of the two relations to be joined and if v_1 and v_2 are the number of distinct values of the join attributes, the cardinality of the join of the two relations is estimated as $n_1 n_2 / \max\{v_1, v_2\}$

Implementation Details. We implemented the proposed techniques (P-Tree and MD-Tree summaries) in C++ with experiments performed on a 1.8GHz Dual AMD Opteron processors and 16GB RAM. We created sparse matrices using sparseLib++ [23] libraries and used BRAHMS [21] to parse the graphs. We used the schema-aware model for shredding RDF graphs into tables and the B+Tree implementation of [20] for indexes. To compare the techniques, we implemented the joins using hash join algorithm with intermediate tables materialized in memory. Although our work is mainly targeted towards optimizing join queries without constraints, in our experiments, we show that with a combination of equi-depth [25] histograms and our proposed summaries, our techniques perform well for join queries with uri constraints or literal constraints. The histograms provide cardinality information for the constraints, based on which we compute the selectivity of the constraint, which we propagate to estimates of unconstrained joins from our graph summaries. If the value obtained from the histogram for a particular constrained relationship is f and the cardinality of that relationship when unconstrained is n , we compute the selectivity as f/n . If the estimate of an unconstrained join from our graph summaries is c , we then propagate the constraint to the unconstrained join by the product $c(f/n)$.

Datasets. For our experiments, we used one synthetic dataset and one real life dataset. The real life dataset which we used is the Mondial [24] dataset. The Mondial dataset is a rich compilation of geographical Web data sources on global statistics of world countries, cities, provinces, seas, and international organizations. We converted an OWL version of the Mondial dataset into RDF using protégé. For the synthetic dataset, we further enriched the Lehigh University Benchmark (LUBM) [19] dataset with new classes and relationships. To do so, we followed the model used by LUBM in determining the sizes of instances of classes and relationships such that the generated dataset models a real life university domain. For scalability experiments, we generated three LUBM datasets with 10, 15 and 20 universities. The table below shows the properties of these datasets.

Table 1: Dataset Properties

	Mondial	LUBM10	LUBM15	LUBM20
# Instance nodes	5841	248880	385530	528334
# Instance edges	18565	3021748	4677928	6416925
# Literal Nodes	12450	106300	164795	225209
# Literal edges	14002	429927	668110	915398
# Unique Instance edges	28	31	31	31

Summaries. Stored as a list, the size of all patterns of length up to three for the Mondial dataset is 378534 bytes and the size of the unpruned P-Tree and MD-Tree are 164792 and 128690 bytes, with the P-Tree giving about a 56% reduction in size and the MD-Tree giving about a 66% reduction in size. We pruned the P-Tree and MD-Tree, by constructing summaries which are 50% of the original unpruned summary sizes. **Figure 10** shows the sizes of the unpruned and pruned summaries for this dataset.

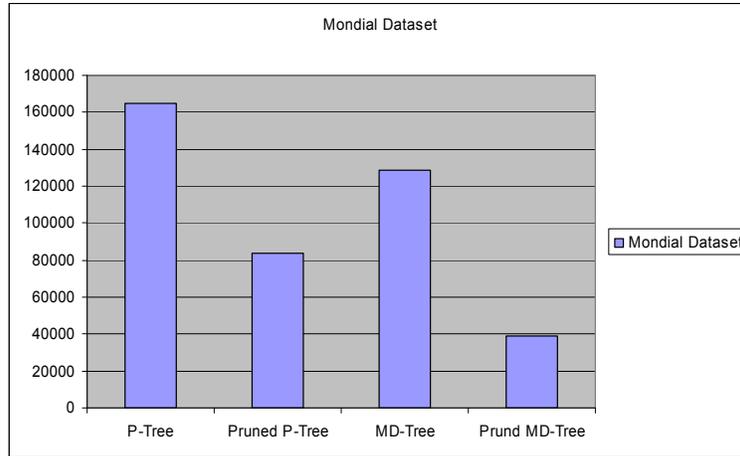


Figure 10: The sizes of the unpruned and pruned summaries for the Mondial dataset

On the other hand, the size (as a list) of all patterns of length two for the LUBM dataset is 9614 bytes. Its unpruned P-Tree and MD-Tree are 5582 and 9422 bytes, with the P-Tree giving about a 42% reduction in size and the MD-Tree giving a minimal reduction in size. Once again, we pruned the P-Tree and MD-Tree by constructing summaries which are 50% of the original unpruned summaries. For both datasets, we used a 5% significance level and a β value of 3 for constructing the MD-Tree summaries. Figure 11 shows the sizes of the unpruned and pruned summaries for the LUBM10 dataset.

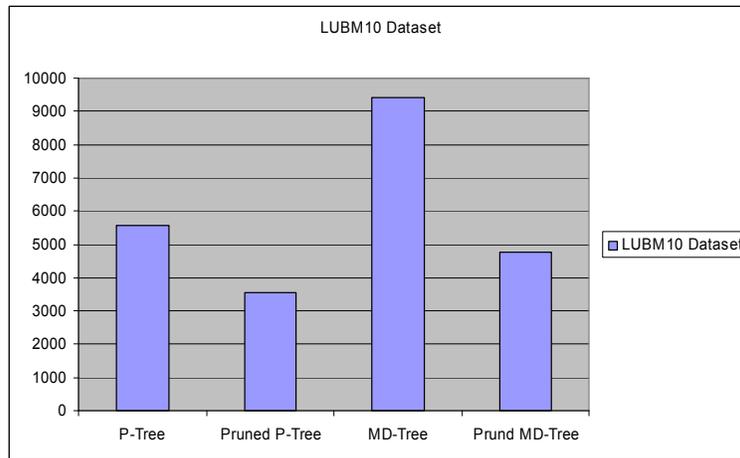


Figure 11: The sizes of the unpruned and pruned summaries for the LUBM10 dataset

Time Analysis. The time for discovering all patterns of length $\max L$ is the most time-consuming part of our approach. Fortunately, it is a preprocessing step and depends on the connectedness of the dataset. The time needed for constructing the P-Tree and MD-Tree summaries is in the order of tenths of seconds.

Queries. As we noted earlier, our proposed summaries are targeted towards join queries posed over RDF databases where relationships are specified in the joins but without constraints on the nodes. However, in our experiments, we also evaluated the efficiency of our techniques when combined with histograms for join queries which may have constraints on the nodes. To make our evaluations as realistic as possible, we avoided choosing queries which are a random concatenation of relationships in the data graph. While doing so could have led us to test queries with larger sizes, we made the choice of queries which are meaningful and are as close as possible to real life queries that users are likely to pose over the datasets. First, we chose queries which have no constraints on the nodes, then we also chose queries which have uri constraints on the nodes as well as those which have literal constraints on the nodes. For this reason, our queries can be grouped into three categories. 1) Join queries with no constraints at all, i.e., for which relationships are specified and all nodes are variables. 2) Join queries with uri constraints, i.e., those for which relationships are specified but some nodes are constrained to be of a particular uri. 3) Join queries with literal constraints, i.e., those for which relationships are specified but some nodes are constrained to be a particular literal value. We have included the queries in SPARQL in the appendix

of this paper. Queries 1 – 4 are posed over the Mondial dataset while queries 5 – 8 are posed over the LUBM dataset.

Error metric. We used the relative error metric $|freq(p) - freq(p^{\wedge})|/freq(p)$ to measure the estimation error where $freq(p)$ and $freq(p^{\wedge})$ are the true and estimated frequencies of p . Since the nature of the SGT is such that it does not provide estimates for patterns, we do not include it in our comparison of the estimation error of the techniques.

Results. We first show the results of queries over the Mondial dataset. Figure 12 shows the accuracy of the techniques for query 1, which is an unconstrained join query with three relationships. The cardinality of the result of this query is 4916. The unpruned P-Tree and the pruned P-Tree both estimate the cardinality of the result of this query as 4915 with an estimation error of 0.0002. The pruned P-Tree exhibits a good performance for this query. On the other hand, the unpruned MD-Tree estimates it as 4934 with an estimation error of 0.0037 while the pruned MD-Tree estimates it as 3276 with an estimation error of 0.3336. The TQOT estimates it as 3445 with an estimation error of 0.2992. The pruned MD-Tree estimates the cardinality of this query with the highest error amongst all the techniques. This is for the case where a maximal impact subtree is pruned because it does not fit the budget, leading to estimation using the independence assumption.

As Figure 13 shows, the unpruned P-Tree, pruned P-Tree, the unpruned MD-Tree and the TQOT all performed equally well while the pruned MD-Tree and SGT did not perform as well. We note that although the size of the pruned P-Tree is about half the size of the unpruned P-Tree, the pruned P-Tree has the same performance as the unpruned P-Tree. This is an indication of the efficiency of the pruning technique. As indicated by its estimation error, the performance of the pruned MD-Tree is not as good as its unpruned counterpart. As mentioned earlier, this is because of the pruning of a maximal impact subtree that does not fit the budget. For the TQOT, this is the case where the estimation error is not high enough to cause it to choose a sub-optimal plan.

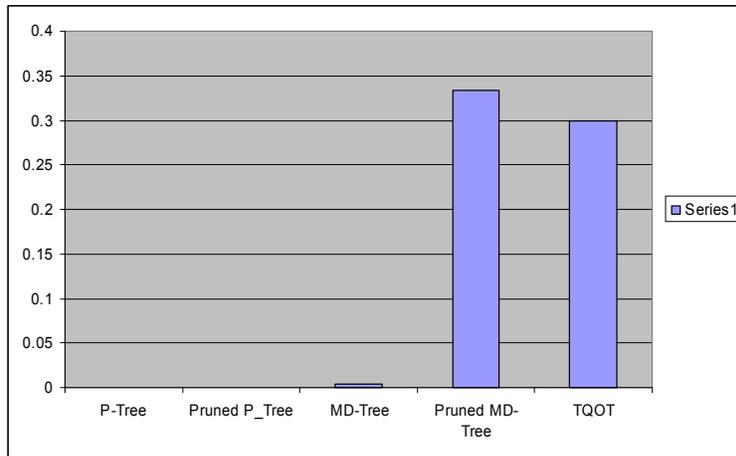


Figure 12: The estimation errors of the techniques for Query 1

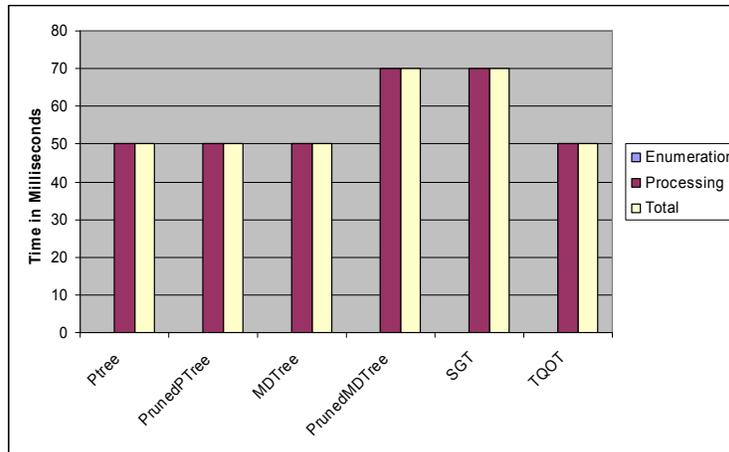


Figure 13: Result of Query 1: An Unconstrained Join Query over the Mondial Dataset

Figure 14, shows the estimation errors of the evaluated techniques for query 2, which is another unconstrained join query involving three relationships. For this query, the cardinality of the result is 877. The unpruned P-Tree and the pruned P-Tree both estimate the cardinality of the result as 865, with an estimation error of 0.0137. The effectiveness of the pruning technique of the P-Tree is shown by the fact that although the pruned P-Tree is about half the size of the unpruned P-Tree, its estimation error is the same as that of the unpruned P-Tree. The unpruned MD-Tree estimates it as 868 with an estimation error of 0.0103. The pruned MD-Tree estimates the result cardinality as 3276 with an estimation error of 2.7355, while the TQOT estimates the result cardinality as 2005 with an estimation error of 1.2862. Again, the unpruned P-Tree and pruned P-Tree have the best performance with the least estimation error, followed by the unpruned MD-Tree. The pruned MD-Tree has the worst performance.

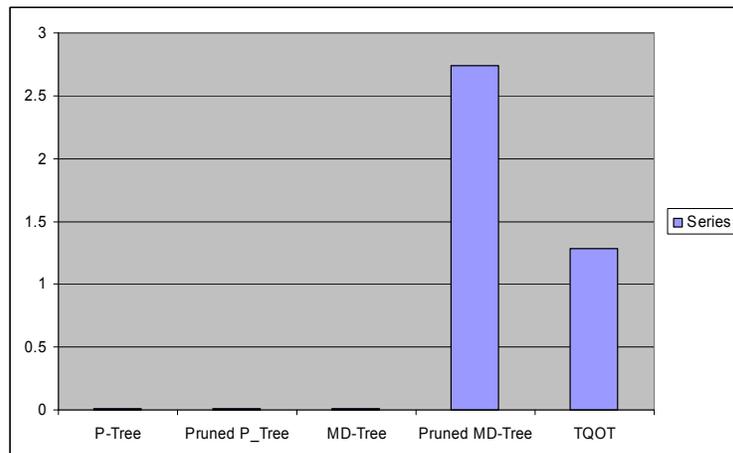


Figure 14: The estimation errors of the techniques for Query 2

As Figure 15 shows, the TQOT exhibits the worst performance, while all the other techniques exhibit the same performance. For the TQOT, this is the case where the join uniformity assumption does not hold and the estimation errors are large enough to cause it to choose a sub-optimal query plan. Although the estimation error of the pruned MD-Tree is larger than that of the unpruned MD-Tree as well as the TQOT, the pruned MD-Tree still performs as well as the unpruned MD-Tree and better than the TQOT. This is because the intermediate estimates obtained from the pruned MD-Tree are more accurate than the overall estimate, as such it still finds the optimal plan.

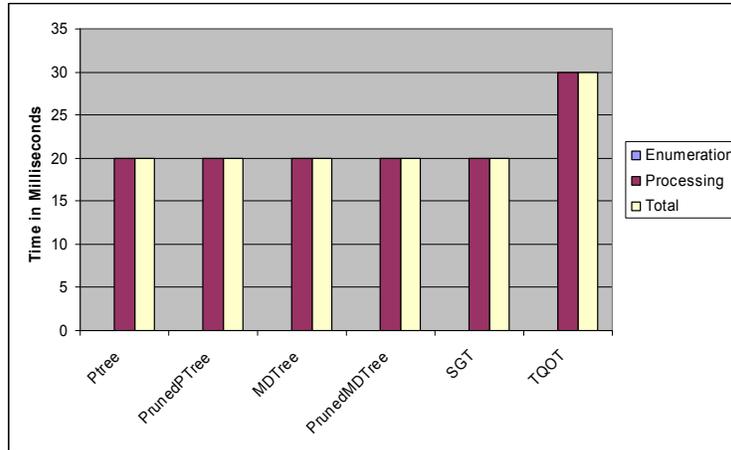


Figure 15: Result of Query 2: An Unconstrained Join Query over the Mondial Dataset

Figure 16 shows the estimation errors of the techniques for query 3, a join query involving three relationships, with a uri constraint. The cardinality of the result of this query is 1033. The unpruned P-Tree, pruned P-Tree and unpruned MD-Tree all estimate the cardinality of the result of this query as 429, with an estimation error of 0.5847. For this query, once again, the pruning technique for the P-Tree prunes effectively as the estimation error for the P-Tree is the same as that of the unpruned P-Tree. The pruned MD-Tree estimates the cardinality of this query as 0 with an estimation error of 1. The estimation error of the pruned MD-Tree is greater than that of the unpruned MD-Tree, reflecting the case when a maximal impact subtree of the MD-Tree is pruned in order to meet the budget. The TQOT estimates it as 289, with an estimation error of 0.7202.

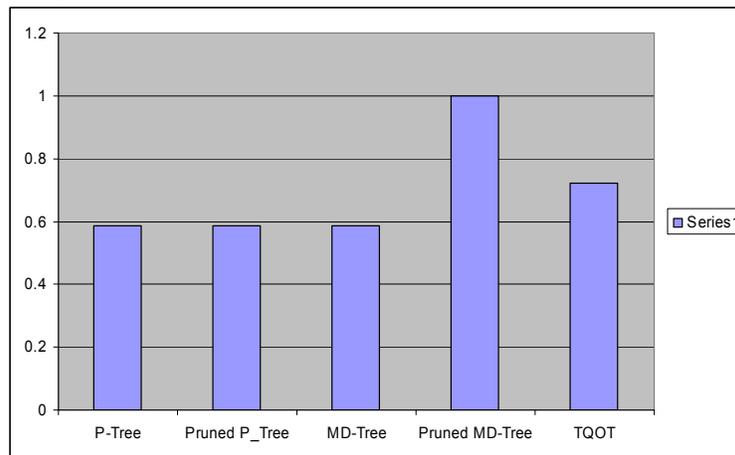


Figure 16: The estimation errors of the techniques for Query 3

Figure 17 shows the query processing performances of the techniques for query 3. For this query, all the techniques find the optimal query plan and thus exhibit the same performance. This is due to the constraint on the query which causes the techniques to find an optimal query plan, even though their estimation errors differ.

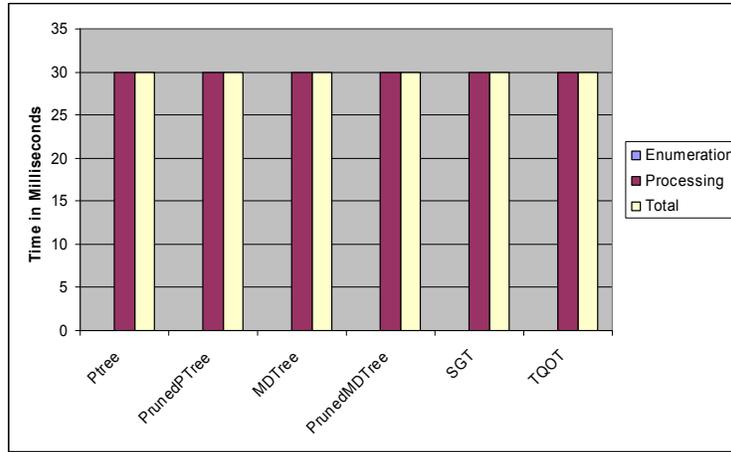


Figure 17: Result of Query 3: A Uri Constrained Join Query over the Mondial Dataset

Figure 18 shows the estimation errors of the evaluated techniques for query 4, a join query involving 4 relationships with a literal constraint. The cardinality of the result of this query is 162. The unpruned P-Tree, pruned P-Tree, unpruned MD-Tree and pruned MD-Tree all estimate the cardinality of the result as 1 with an estimation error of 0.9938. The TQOT estimates it as 1296 with an estimation error of 7. Recall that for queries with constraint, we propagate the selectivity of the constraint to the estimates of the joins. Our propagation assumes that a uniform number of patterns will be affected by the constraint. For this query, this uniform assumption does not hold and hence the larger errors in the estimates for the cardinality of the results.

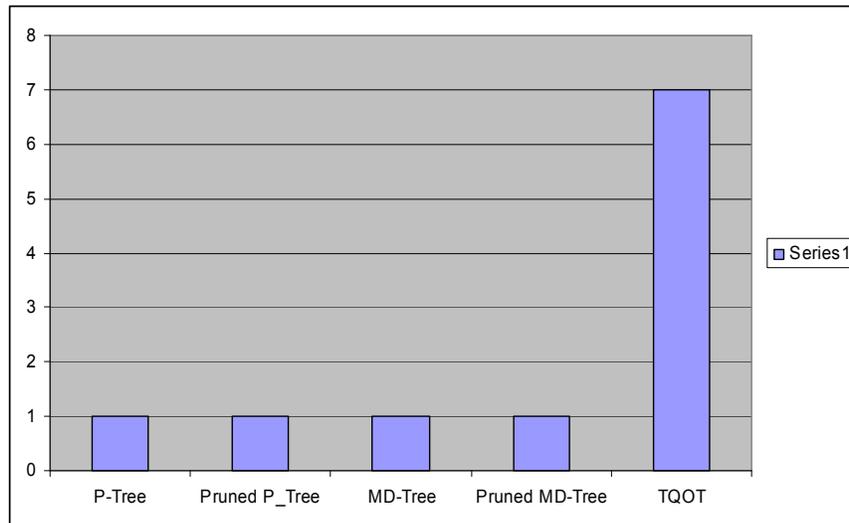


Figure 18: The estimation errors of the techniques for query 4

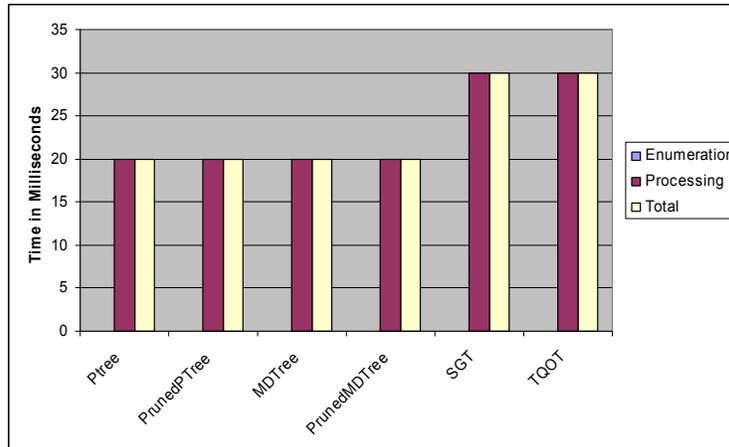


Figure 19: Result of Query 4: A Literal Constrained Join Query over the Mondial Dataset

As Figure 19 shows, for this query, both the SGT and TQOT techniques perform worse than our techniques. In conformity with the estimation errors of our techniques, all of our techniques exhibit the same performance in terms of the time spent for query processing. For this query, the effectiveness of our pruning techniques is once again shown by the fact that both the pruned P-Tree and pruned MD-Tree perform as well as their unpruned counterparts.

In summary for the Mondial dataset, our experimental results show that for the accuracy of estimates of unconstrained join queries, our techniques perform better than the Traditional Query Optimization Technique with the unpruned P-Tree, pruned P-Tree and the unpruned MD-Tree having an average error of 0.7% compared to 79% for the Traditional Query Optimization Technique. The pruned MD-Tree performs worst with an average error of 153%. For the constrained queries, all our techniques perform much better than the Traditional Query Optimization Technique with the unpruned P-Tree, pruned P-Tree and the unpruned MD-Tree having an average error of 79%, while the pruned MD-Tree has an average error of 100% compared to 386% average error of the Traditional Query Optimization Technique. Over both the unconstrained and constrained join queries, our techniques are superior with an average error of 40% for the unpruned P-Tree, pruned P-Tree and the unpruned MD-Tree and 127% average error of the pruned MD-Tree compared to 233% average error of the Traditional Query Optimization Technique.

As can be observed from the foregoing discussion, the accuracy of the pruned P-Tree was the same as that of the unpruned P-Tree for the queries 1 – 4. To assess the level of degradation of the accuracy of the pruned P-Tree with respect to the size of the pruned P-Tree, we pruned the P-Tree to fit a 10KB budget, 25KB budget and 50KB budget. Figure 20 shows the estimation errors of the pruned P-Trees of sizes 10KB, 25KB and 50KB for query 2. As the figure shows, the pruned P-Tree of size 10KB shows a large estimation error for this query, while the pruned P-Tree of Size 25KB and 50KB have the same minimal estimation error (0.0137) as the unpruned P-Tree (see Figure 14 above). As the figure shows, one can construct a summary that is as small as 25KB, without increasing the estimation error when compared to the unpruned P-Tree, for this query. As the figure also shows, it takes the same amount of time to process the query for each of the pruned summaries.

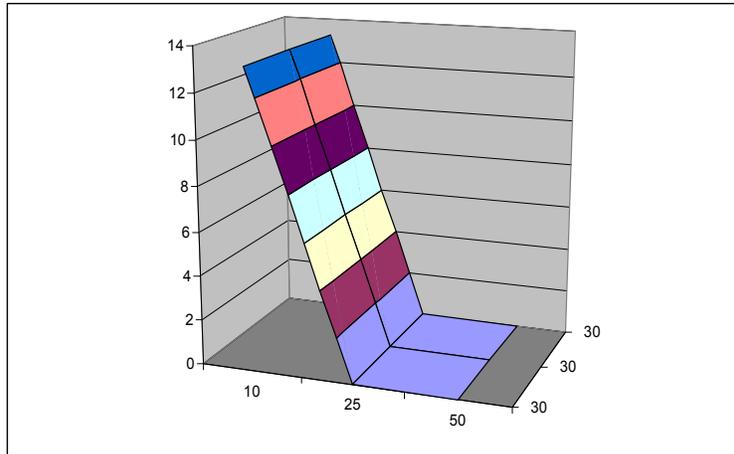


Figure 20: The Estimation Error for Query 2 for Summaries of Size 10KB, 25KB and 50KB

Figure 21 shows the estimation errors of the pruned P-Trees of sizes 10KB, 25KB and 50KB for query 3. Once again, the pruned P-Tree of size 10KB shows a large estimation error for this query, while the pruned P-Tree of Size 25KB and 50KB have the same minimal estimation error (0.5847) as the unpruned P-Tree (see Figure 16 above). Once again, as illustrated by this figure, for this query, we can construct a summary that is as small as 25KB, without increasing the estimation error when compared to the unpruned P-Tree. As Figure 20 and Figure 21 show, the estimation error increases as the size of the summary decreases. However, the estimation error of the summary does not overly increase for a reasonably small budget, in these cases a budget of size 25KB.

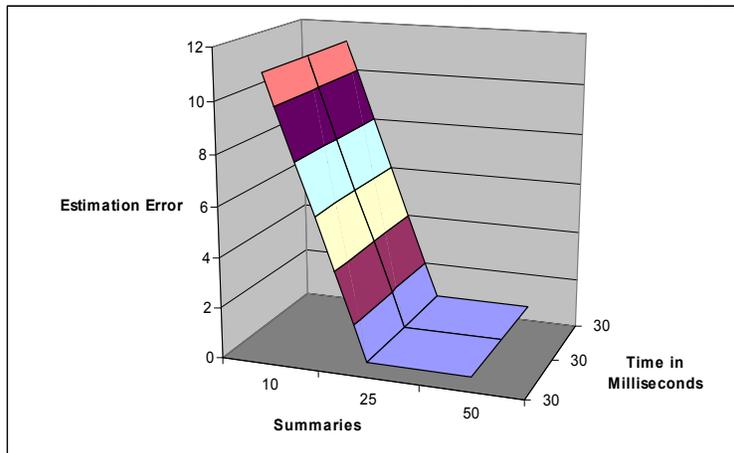


Figure 21: The Estimation Error for Query 3 for Summaries of Size 10KB, 25KB and 50KB

We now show the results over the LUBM dataset. For this dataset, once again we posed queries in categories 1 – 3. Figure 22 shows the estimation errors of the evaluated techniques for query 5, an unconstrained join query involving three relationships, over the LUBM10 dataset. The cardinality of the result of this query is 2912. The unpruned P-Tree, the pruned P-Tree and the unpruned MD-Tree all estimate it as 2906 with an estimation error of 0.0021 while the pruned MD-Tree estimates it as 1 with an estimation error of 0.9997. This is because maximal impact subtrees along the path of this query were aggressively pruned causing estimation to be based on the independence assumption which does not hold in this case. The TQOT estimates it as 2914 with an estimation error of 0.0007.

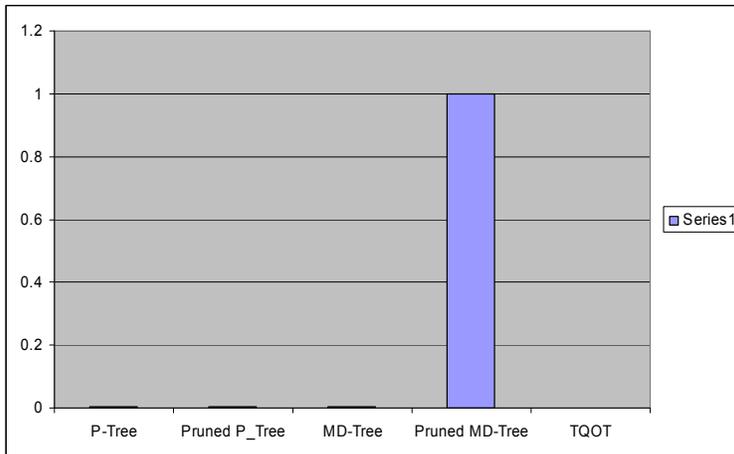


Figure 22: The estimation errors of the techniques for query 5

As Figure 23 shows, our pruned P-Tree and pruned MD-Tree show encouraging performances with both performing as well as their unpruned counterparts. Once again, the performance of the pruned P-Tree and pruned MD-Tree show the effectiveness of the pruning techniques. Our techniques have the same performance as the TQOT, while the SGT exhibits a worse performance than all the other techniques. Further, the time taken to enumerate the plans in all cases is very negligible. To show how the techniques scale with larger dataset, we posed the same query over LUBM15 and LUBM20 datasets and as Figure 24 shows, the techniques scale linearly with increase in dataset size and the performance of SGT gets even worse with increase in dataset size.

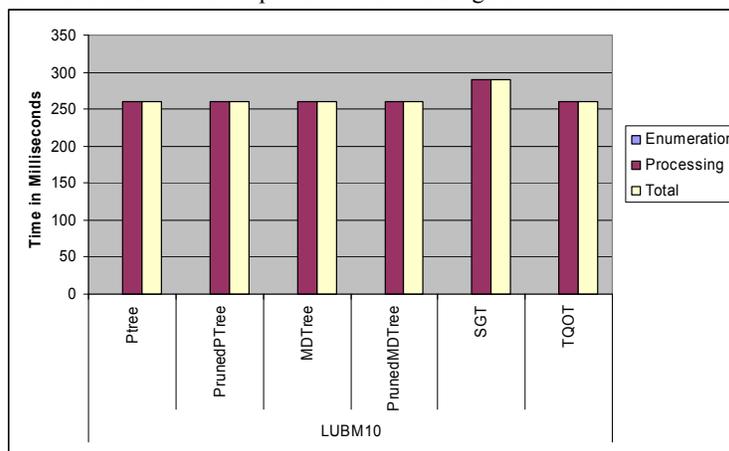


Figure 23: Result of Query 5: An Unconstrained Join Query over the LUBM10 dataset

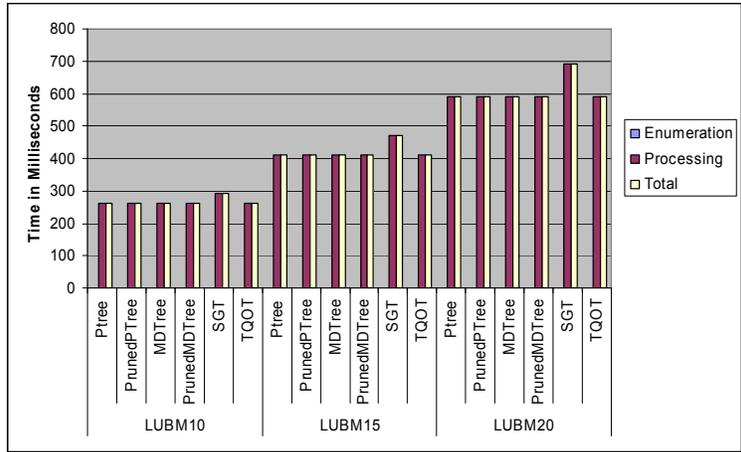


Figure 24: Result of Query 5: An Unconstrained Join Query over the LUBM10, LUBM15 and LUBM20 datasets

Figure 25 shows the estimation errors of the techniques for query 6, another unconstrained join query involving four relationships, over the LUBM10 dataset. The cardinality of the result of this query is 559. The unpruned P-Tree, the pruned P-Tree and the unpruned MD-Tree all perfectly estimate the cardinality of the result as 559 with no estimation error. However, the pruned MD-Tree estimates it as 1 with an estimation error of 0.998. Once again, this is due to the aggressive pruning of maximal impact subtrees along the path of this query with estimation using the independence assumption which does not hold. The TQOT estimates the cardinality of the result as 3790 with an estimation error of 5.78.

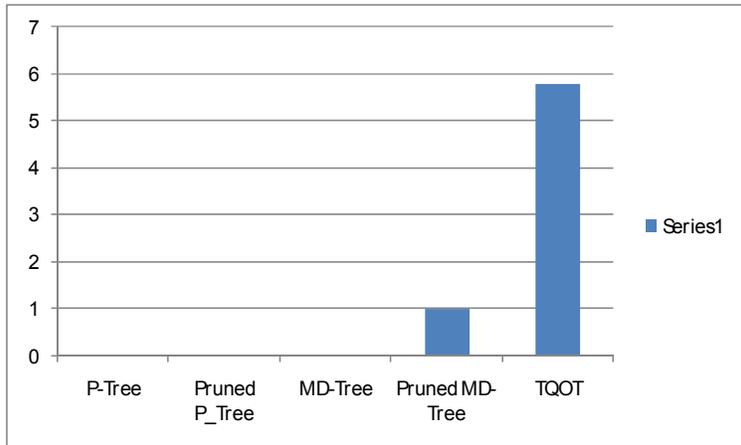


Figure 25: The estimation errors of the techniques for query 6

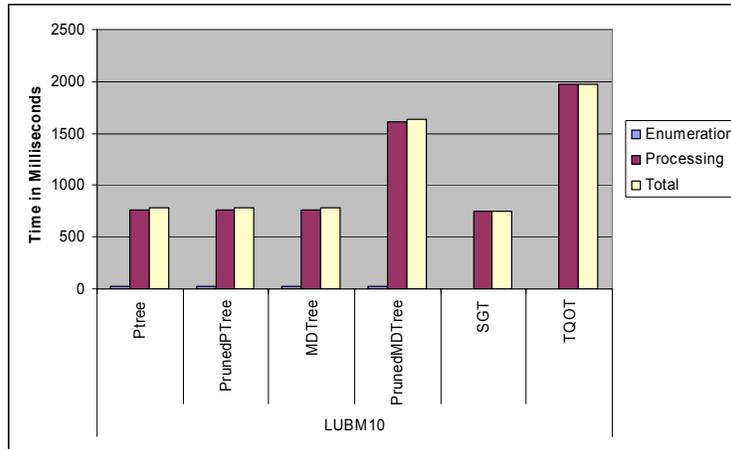


Figure 26: Result of Query 6: An Unconstrained Join Query over LUBM10 dataset

For this query, as shown in Figure 26, in conformity with their estimation errors, our pruned P-Tree performed as well as its unpruned counterpart while the pruned MD-Tree shows a worse performance than its unpruned counterpart. Our techniques all perform better than the TQOT. This is due to the fact that the uniformity assumptions on which it is based do not hold and the estimation errors are large enough to cause it to choose a worse plan than the P-Tree and MD-Tree techniques. The SGT also exhibits the same performance as our techniques.

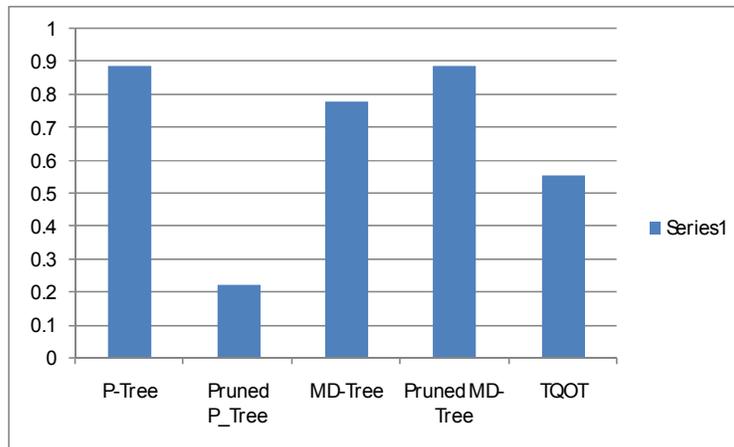


Figure 27: The estimation errors of the techniques for query 7

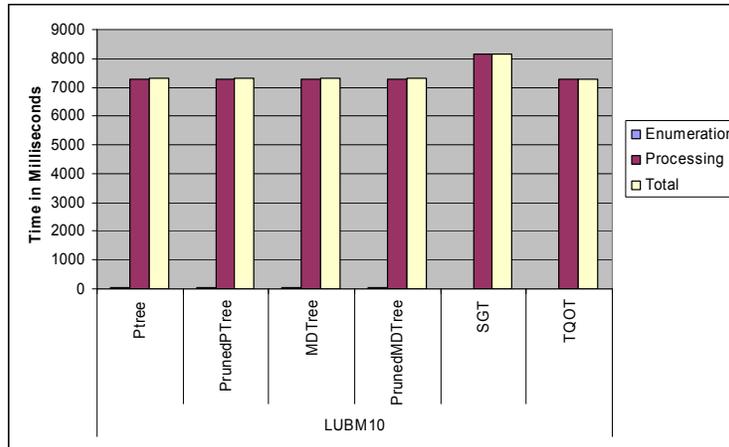


Figure 28: Result of Query 7: A Uri Constrained Join Query over LUBM10 dataset

Figure 27 shows the estimation errors of the techniques for query 7, a uri constrained join query involving three relationships over the LUBM10 dataset. The cardinality of the result of this query is 9. The unpruned P-Tree estimates it as 1 with an estimation error of 0.8889, the pruned P-Tree estimates it as 7 with an estimation error of 0.2222. For this query, the estimate obtained from the pruned P-Tree is more accurate than that obtained from the unpruned P-Tree. This is as a result of the combination function that is used to obtain the estimate of a pattern which is larger than the size of patterns maintained in the P-Tree. The unpruned MD-Tree estimates it as 2 with an estimation error of 0.7778, while the pruned MD-Tree estimates it as 1 with an estimation error of 0.8889. The TQOT estimates it as 4 with an estimation error of 0.5556. As shown in Figure 28, all techniques except the SGT came up with the same query plan and thus have the same performance. This is because their absolute estimation errors are small and thus they all find an optimal plan.

Figure 29 shows the estimation errors of the evaluated techniques for query 8, a literal constrained join query involving four relationships over the LUBM10 dataset. The cardinality of the result of this query is 93. The unpruned P-Tree, the pruned P-Tree, the unpruned MD-Tree and the pruned MD-Tree all estimate the cardinality of the result as 1 with an estimation error of 0.9892. Once again, this is because the uniform assumption based upon which the selectivity of the constraint is propagated to the estimates of the joins does not hold. The TQOT estimates the cardinality of the result as 341 with an estimation error of 2.6667.

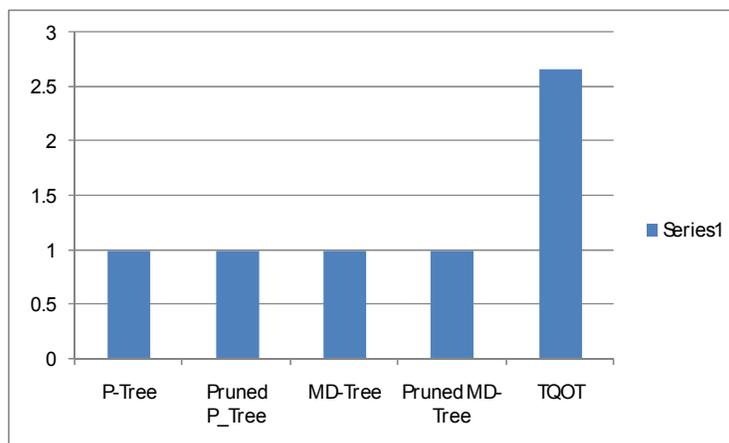


Figure 29: The estimation errors of the techniques for query 8

Although our techniques have a smaller overall estimation error for this query than the TQOT, the TQOT performs slightly better than our techniques, as Figure 30 shows. This is because the estimates obtained for intermediate patterns from our techniques with the selectivity propagation, caused it to choose a worse plan than the TQOT. The SGT shows a slightly worse performance than all our techniques.

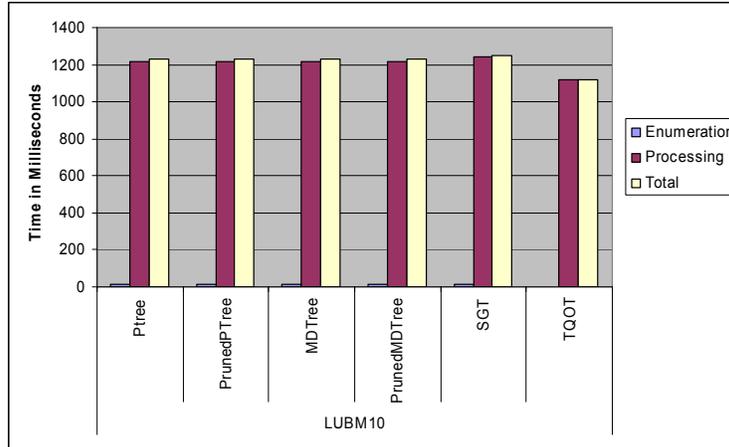


Figure 30: Result of Query 8: A Literal Constrained Join Query over LUBM10 dataset

In summary for the LUBM dataset, our experimental results show that our techniques perform better on the average in terms of the accuracy of the estimates. For the unconstrained join queries, the unpruned P-Tree, pruned P-Tree and unpruned MD-Tree all had an average error of 0.11%, while the pruned MD-Tree had an average error of 100% compared to 289% average error of the Traditional Query Optimization Technique. For the constrained join queries, the unpruned P-Tree had an average estimation error of 94%, the pruned P-Tree had an average error of 61%, the unpruned MD-Tree had an average error of 88%, while the pruned MD-Tree had an average error of 94% compared to 161% average error of the Traditional Query Optimization Technique. Over both the unconstrained and constrained join queries, the unpruned P-Tree had an average error of 47%, the pruned P-Tree had an average error of 30%, the unpruned MD-Tree had an average error of 44%, while the pruned MD-Tree had an average error of 97% compared to 225% average error of the Traditional Query Optimization Technique.

5 Related Work

An optimization approach for RDF Graph patterns that is based on the selectivity of RDF triples was proposed in [48]. To determine the selectivity of resources, they simply adapted the techniques developed for Relational Databases to RDF. Based on the selectivity of resources, they define the selectivity for a triple pattern and then a graph pattern. Similar to our approach, they determine the selectivity of graph patterns of size 2 (i.e., two joined triple patterns) by enumerating and storing all possible graph patterns of size 2 and their frequencies. However, unlike in our approach, they only consider joined patterns that are obtainable from the schema and as such, may miss some patterns. Furthermore, they do not consider approaches for pruning the summary such that it fits a given memory budget as well as techniques for tuning the summary to favor certain patterns.

The optimization unit of an implementation of the SPARQLer query language [31], adopts a greedy approach for optimizing supported graph pattern queries. Given the graph pattern for a query such that edges in the graph pattern are annotated with the expected cardinality of triples that satisfy the triple pattern it represents, a query plan is constructed by beginning with the edge which has the least cardinality and extending to its adjacent edges until all edges of the graph have been visited, using an algorithm similar to that used for growing a minimum spanning tree. In contrast to our approach, this does not provide for nor utilize the estimates of the cardinality of larger patterns in constructing the query plan. As such, this technique may result in sub-optimal plans. To the best of our knowledge, ours is the first work on summarizing RDF graphs for subgraph frequency estimation.

The Tuple Graph Synopses (TUGS) was proposed in [46] as a graph-based summary for cardinality estimation for relational databases. Given a schema graph which defines the join relationships amongst relations, this technique constructs a graph of the data contained in all relations and their join dependencies with respect to the schema graph. The TUGS synopsis for the data is then constructed by a systematic summarization of the instance graph. However, it is unclear how this approach may be applied to RDF graph stored using the schema-oblivious approach, nor RDF graph patterns which require self-joins, since the TUGS model does not consider self-joins of a relation. A technique that seamlessly copes with RDF patterns, regardless of the storage model used to map the data into relations is desirable.

For the XML data model, work most closely related are techniques that summarize XML data for selectivity estimation for path expressions [1] and twigs[38][39][51]. A fundamental difference between these techniques and ours is the data and query models. All the techniques except [39] assume tree-structured XML data. More important, all the techniques are proposed for either path or twig queries so that it is unclear how they apply to arbitrary graph-structured queries. Our estimation value for path patterns is similar in spirit to the notion of δ -derivable twigs introduced in [51] for pruning twigs whose estimated frequencies are within δ error of their true frequencies. However, the technique of [51] may blindly prune a pattern that, if left unpruned, may have caused more twigs to be pruned, thereby reducing the summary size further. In contrast, our value-based approach makes a more informed choice of patterns to be pruned.

Several efforts have been made in using graph-indexing schemes to reduce the cost of processing graph containment queries, over a collection of many disconnected graphs. In these approaches, a graph containment query is processed in two steps. The first step retrieves a candidate set of graphs that contain the indexed features of the query graph. The second step uses subgraph isomorphism to validate each candidate graph. GraphGrep [45] uses a path-based indexing approach that selects all paths of up to length l_p as the indexing feature. The size of the candidate set obtained in the first step could be large since paths do not keep graph structure. To cope with this, GIndex [58] uses frequent graph fragments as the indexing feature. To reduce the large (potentially exponential) number of frequent fragments, only discriminative frequent fragments are kept. Noting that the set of frequent graph fragments contain many more tree than non-tree structures, Tree+ Δ [59] indexes frequent trees, reducing the large index construction time of GIndex due to graph-mining. On demand, Tree+ Δ further reduces the size of the candidate set by selecting a small portion of discriminative non-tree features related to query graphs only. In complement, our work allows for optimizing the subgraph isomorphism tests, in the second step, using estimates of the cardinalities of both indexed and non-indexed fragments of the query. In addition, our technique can also be applied to a large connected graph.

6 Conclusions

Graph pattern querying is important for eliciting information from graph databases such as RDF data stores. Optimizing graph pattern queries requires estimating the frequency of subgraphs in a query graph pattern. In this work, we presented two techniques for summarizing the structure of graphs and we showed how to prune the summary to fit a given space budget. As our experiments showed, for unconstrained join queries, the proposed P-Tree and MD-Tree exhibited encouraging performance when compared to the traditional query optimization techniques used in relational database systems and the simple greedy technique used by the execution engine of the SPARQLer system. In terms of the accuracy of the estimates, our techniques showed superior performance over the Traditional Query Optimization technique for both unconstrained and constrained join queries. The pruned P-Tree is relatively stable for all datasets, but performs best when graph patterns that share a common sub-graph pattern co-occur. The pruned MD-Tree should perform best when single points of dependence exist among subgraphs, however, no such cases showed in our experimentation. Our experiments also show that the time taken to enumerate plans is often negligible when compared to the actual query processing time. Although the Dynamic Programming approach arguably may take some time for enumeration as the number of relationships grows larger, the authors in [32] have shown a new class of enumeration algorithms that work to reduce the enumeration time without a large decrease in performance when compared to the dynamic programming enumeration.

There are several directions we hope to explore in the future. First, we will look into a summarization and estimation framework for RDF graphs which may have subpropertyOf hierarchies on the edges. Next we will extend our graph summaries to better cope with constraints so as to avoid the less accurate estimates obtained when propagating the selectivity of a constraint to the estimates of the patterns. Further we will investigate techniques for gracefully accommodating updates to the data graph into our summaries without a complete reconstruction of the summaries.

Acknowledgments. This work was funded by NSF-ITR-IDM Award #0325464 and #071444.

APPENDIX A

Query 1

For each country, find the provinces, the province capitals and the rivers that flow through the country

Prefix protégé: <<http://protege.stanford.edu/kb#>>

```

Select ?country ?province ?province_capital ?river
Where
{
  ?country protégé:has-province ?province.
  ?province protégé:has-province-capital ?province_capital.
  ?river protégé:flows-through-country ?country
}

```

Query 2

For each country, find the provinces, the province capitals and the languages of the country

```

Prefix protégé: <http://protege.stanford.edu/kb#>
Select ?country ?province ?province_capital ?language
Where
{
  ?country protégé:has-province ?province.
  ?province protégé:has-province-capital ?province_capital.
  ?country protégé:languages ?language
}

```

Query 3

For each country that is a member of the World Trade Organization (WTrO), find the provinces and the province capitals

```

Prefix protégé: <http://protege.stanford.edu/kb#>
Select ?country ?province ?province_capital
Where
{
  ?country protégé:has-province ?province.
  ?province protégé:has-province-capital ?province_capital.
  protégé:WTrO protégé:member ?country
}

```

Query 4

For each country that is in Africa, find the provinces and the province capitals

```

Prefix protégé: <http://protege.stanford.edu/kb#>
Select ?country ?province ? province_capital ?continent
Where
{
  ?country protégé:has-province ?province.
  ?province protégé:has-province-capital ?province_capital.
  ?country protégé:encompassed ?continent.
  ?continent protégé:name "Africa"
}

```

Query 5

For each professor that directs a research group, find the course he teaches and the ta of the course

```

Prefix lubm: <univ-bench.owl#>
Select ?researchGroup ?professor ?course ?ta
Where
{
  ?researchGroup lubm:director ?professor.
  ?professor lubm:teacherOf ?course.
  ?ta lubm:teachingAssistantOf ?course
}

```

Query 6

For each professor that is a co pi of two projects, find the research group of the projects

```

Prefix lubm: <univ-bench.owl#>
Select ?project1 ?project2 ?professor ?researchGroup1 ?researchGroup2
Where
{
  ?project1 lubm:coPI ?professor.
  ?project2 lubm:coPI ?professor.
  ?researchGroup1 lubm:hasProject ?project1
  ?researchGroup2 lubm:hasProject ?project2
}

```

Query 7

For each project that has two sub projects, find the publications of the project authored by FullProfessor0 in department10, university0 and the citations of the publication

```

Prefix lubm: <univ-bench.owl#>
Prefix dept10_univ0: <http://www.Department10.University0.edu/>
Select ?project1 ?project2 ?project3 ?publication1 ?publication2
Where
{
  ?publication1 lubm:cites ?publication2.
  ?publication2 lubm:publicationAuthor dept10_univ0:FullProfessor0.
  ?project1 lubm:hasPublication ?publication1.
  ?project2 lubm:partProjectOf ?project1.
  ?project3 lubm:partProjectOf ?project1
}

```

Query 8

Find the journal publications of all head of departments whose name is FullProfessor2.

```

Prefix lubm: <univ-bench.owl#>
Select ?professor ?department ?publication ?journal
Where
{
  ?professor lubm:headOf ?department.
  ?professor lubm:name "FullProfessor2".
  ?publication lubm:publicationAuthor ?professor.
  ?publication lubm:publicationInJournal ?journal
}

```

APPENDIX B

We now show how the preference and estimation values of patterns are combined to obtain a single value with which the P-Tree is pruned.

Definition 16. Let $P = \{p_1, p_2, \dots, p_m\}$ be the set of patterns in the P-Tree and $p_{EV_{max}}$, the maximum expected value of patterns in P . Given a constant $c > 0$, the value of a pattern p_j is given by:

$$p_{V_j} = (1 + p_{EV_j})(1 + p_{PV_j}) + ip_{EV_{max}}$$

where i is an indicator variable whose value is 1 if $p_{PV_j} \geq c$ and 0 otherwise. The additive constants ensure that the value of a pattern is non-zero when either its preference or estimation value is zero. On the other hand, the second term allows for tuning the P-Tree by boosting the values of important patterns as defined by their preference values, to delay their pruning.

We note that when the size of the P-Tree exceeds the budget, we prune the tree systematically to meet the budget. To prune, we compute the value of each node of the P-Tree (the combination of its preference and estimation values, or simply the latter if the former is not given) and select nodes to be pruned as shown in the *Prune P-Tree* algorithm shown below.

Prune P-Tree

Input: P-Tree T, Budget B, constant c

Output: Pruned P-Tree T

```
01  $\delta \leftarrow 0$ ; inc  $\leftarrow 0$ ; Set estimable  $\leftarrow \emptyset$ ; done  $\leftarrow$  false
02 while done = false do
03     estimable  $\leftarrow \emptyset$ 
04     for each internal node v in T in bottom – up order do
05         compute estimation value given  $\delta$ 
06         if v's estimation value > 0 then
07             insert all its children into estimable
08         end if
09     end for
10     if  $sizeof(T) - sizeof(estimable) \leq B$  and  $\delta$  is optimal then
11         compute values of patterns in T
12         if c > 0
13             delete smaller valued patterns to meet B
14         else
15             prune all patterns in estimable
16         end if
17         done  $\leftarrow$  true
18     else
19         adjust  $\delta$ 
20     end if
21 end while
```

As lines 4-9 of the Prune-Tree algorithm show, estimation values are used to select the patterns to be pruned. However in lines 12-16, the patterns eventually pruned may differ when tuning the tree. The running time of the algorithm $O(Ld^{\max L} \log(\max_i \{freq(p_i)\}))$, is reasonable since logarithm is a slowly growing function and $\max L$ will typically be small. The running time stems from the loops in lines 4-9 and 2-21. Lines 4-9 run in $O(Ld^{\max L})$ time, where L and d are the numbers of unique edge labels and the maximum degree of nodes in the graph respectively. Recall that the root of the P-Tree has a child for each unique edge label in the graph while internal nodes have at most d children. Lines 2-21 will be executed at most $\log(\max_i \{freq(P_i)\})$ times, since all patterns are estimable at $\delta = \max_i \{freq(P_i)\}$

7 References

1. Abounaga, A., Alameldeen, A., Naughton, J.: Estimating the Selectivity of XML Path Expressions for Internet Scale Applications. In VLDB, 2001.
2. Aleman-Meza, B., Hakimpour, F., Arpinar, B. I., Sheth, A.: SwetoDblp Ontology of Computer Science Publications. Tech. Report, CS Department, University of Georgia, 2006
3. Anyanwu, K., Maduko, A., Sheth, A.: SemRank: Ranking Complex Relationship Search Results on the Semantic Web. In Proc. 14th WWW Conference, 2005.
4. Beckett, D.: RDF/XML Syntax Specification (Revised). W3C Recommendation. 10th February 2004. <http://www.w3.org/TR/rdf-syntax-grammar/>
5. Boulicaut, J., Bykowski, A., Rigotti, C.: Approximation of Frequency Queries By Means of Free-sets. In Proc. 4th European PKDD Conference, 2000.
6. Brickley, D., Guha, R. V.: RDF Vocabulary Description Language 1.0: RDF Schema. W3C Recommendation. 10th February, 2004. <http://www.w3.org/TR/rdf-schema>
7. Broder, A. Z., Kumar, R., Maghoul, F., Raghavan, P., Rajagopalan, S., Stata, R., Tomkins, T., Wiener, J.: Graph structure in the web. WWW9/Computer Networks, 2000.
8. Broekstra, J., Kampman, A., Harmelen, F. V.: Sesame: A Generic Architecture for Storing and Querying RDF and RDF Schema. In Proc. 1st ISWC Conference, 2002
9. Burge, C. Identification of Complete Gene Structures in Human Genomic DNA. Ph.D. Thesis, Stanford University, Stanford, CA. 1997
10. Carroll, J. et al.: Implementing the Semantic Web Recommendations. In Proc. 13th WWW Conference, 2004.
11. Chen, Z., Jagadish, H. V., Korn, F., Koudas, N., Muthukrishnan, S., Ng, Raymond., Srivastava, D.: Counting Twig Matches in a Tree. Proc. 17th ICDE Conference 2001.
12. Cho, G., Shaw, D. X. A Depth-First Dynamic Programming Algorithm for the Tree Knapsack Problem. INFORMS J. Computing 9(4) 1997, 431-438.
13. Dehaspe, L., Toivonen, H., King, R. D.: Finding Frequent Substructures in Chemical Compounds. In KDD, 1998.

14. Deshpande, A., Garofalakis, M., Rastogi, R.: Independence is Good: Dependency-Based Histogram Synopses for High-Dimensional Data. In Proc. 2001 ACM SIGMOD Conference, 2001.
15. Deshpande, M., Kuramochi, M., Wale, N.: Frequent Substructure-Based Approaches for Classifying Chemical Compounds. In TKDE. Vol. 17, No. 8. Aug. 05.
16. Freire, J., Haritsa, J., Ramanath, M., Roy, P., Simeon, J.: StatiX: Making XML Count. In Proc. 2002 ACM SIGMOD Conference, 2002.
17. Getoor, L., Taskar, B., Koller, D.: Selectivity Estimation using Probabilistic Models. In Proc. 2001 ACM SIGMOD Conference, 2001.
18. Goldman, R., Widom, J.: DataGuides: Enabling Query Formulation and Optimization in Semi-Structured Databases. In Proc. 23rd VLDB Conference, 1997.
19. Guo, Y., Pan, Z., Heflin, J.: LUBM: A Benchmark for OWL Knowledge Base Systems. Journal of Web Semantics 3(2), 2005, pp158-182.
20. <http://idlebox.net/2007/stx-btree/>
21. <http://lstdis.cs.uga.edu/projects/semdis/brahms>
22. <http://lstdis.cs.uga.edu/projects/semdis/swetodblp>
23. <http://math.nist.gov/sparselib++/>
24. <http://www.dbis.informatik.uni-goettingen.de/Mondial/>
25. Ioannidis, y.: The History of Histograms (abridged). In Proc. 29th VLDB Conference, 2003.
26. Jagadish, H., Koudas, N., Muthukrishnan, S.: Optimal Histograms with Quality Guarantees. In Proc. 24th VLDB Conference, 1998
27. Janik, M., Kochut, K.: BRAHMS: A WorkBench RDF Store and High Performance Memory System for Semantic Association Discovery. In Proc. 4th ISWC Conference, 2005.
28. Karvounarakis, G., Alexaki, S., Christophides, V., Plexousakis, D., Schol, M.: RQL: A Declarative Query Language for RDF. In WWW'02, Honolulu, Hawaii, USA, May7-11 2002
29. Karvounarakis, G., Christophides, V., Plexousakis, D., Alexaki, S.: Querying RDF descriptions for community web portals. In Proc. French BDA Conference, 2001.
30. Klyne, G., Carroll, J. J.: RDF Concepts and Abstract Syntax. W3C Recommendation. (Revised) February 2004. <http://www.w3.org/TR/rdf-syntax-grammar/>
31. Kochut, K., Janik, M.: SPARQLer: Extended SPARQL for Semantic Association Discovery. In ESWC 2007.
32. Kossmann, D., Stocker, K.: Iterative Dynamic Programming: A New Class of Query Optimization Algorithms. In ACM Transactions on Databases. Volume 25, Issue 1 (March 2000)
33. Lim, L., Wang, M., Padmanabhan, S., Vitter, J.S., Parr, R.: XPathLearner: An On-Line Self-Tuning Markov Histogram for XML Path Selectivity Estimation. In Proc. 28th VLDB Conference, 2002.
34. Magkanaraki, A., Karvounarakis, G., Christophides, V., Plexousakis, D., Anh, T.: Ontology Storage and Querying. Technical Report No 308, April 2002
35. McHugh, J., Widom, J.: Lore: Query Optimization for XML. In Proc. 25th VLDB Conference, 1999
36. Pei, J., Dong, G., Zou, W., Han, J.: On Computing Condensed Frequent Pattern Bases. In ICDM, 2002.
37. Perry, M. TontoGen: A Synthetic Data Set Generator for Semantic Web Applications. In SIGSEMIS Bulletin.
38. Polyzotis, N., Garofalakis, M., Ioannidis, Y.: Selectivity Estimation for XML Twigs. In ICDE, 2004.
39. Polyzotis, N., Garofalakis, M.: Statistical Synopses for Graph-Structured XML Databases. In SIGMOD, 2002.
40. Poosala, V., Ioannidis, E.: Selectivity Estimation Without the Attribute Value Independence Assumption. In Proc. 23rd VLDB Conference, 1997.
41. Prud'hommeaux, E., Seaborne, A.: SPARQL Query Language for RDF. W3C Working Draft. 19th April 2005. <http://www.w3.org/TR/rdf-sparql-query/>
42. Seaborne, A.: RDQL - A Query Language for RDF, WWW Consortium, Member Submission SUBM-RDQL-20040109, January 2004
43. Selinger, P. G., Astrahan, M. M., Lorie, R. A., Price, T. G.: Access Path Selection in a Relational Database Management System. In Proceedings of the ACM SIGMOD International Conference on Management of Data. 1979.
44. Shannon, C.E. A Mathematical Theory of Communication, Bell Syst. Tech. Journal 27, 379-423, 623-656. 1948.
45. Shasha, D., Wang, J. T. L., Giugno, R. Algorithmics and Applications of Tree and Graph Searching. In PODS, 2002
46. Spiegel, J., Polyzotis, N.: Graph-Based Synopses for Relational Selectivity Estimation. In Proc. 2006 ACM SIGMOD Conference, 2006.
47. Srivastava, J., Cooley, R., Deshpande, M., Tan, P. Web Usage Mining: Discovery and Applications of Usage Patterns from Web Data. ACM SIGKDD Explorations 2000.
48. Stocker, M., Seaborne, A., Bernstein, A., Kiefer, C., Reynolds, D.: SPARQL Basic Graph Pattern Optimization Using Selectivity Estimation. In WWW 2008.
49. SWAD-Europe Deliverable 10.2: Mapping Semantic Web Data with RDBMSes, 2003.
50. Theoharis, Y., Christophides, V., Karvounarakis, G.: Benchmarking Database Representations of RDF/S Stores. In Proc. 4th ISWC Conference, 2005.
51. Wang, C., Parthasarathy, S., Jin, R.: A Decomposition-Based Probabilistic Framework for Estimating the Selectivity of XML Twig Queries. In EDBT, 2006.
52. Wang, C., Wang, W., Pei, J., Zhu, Y., Shi, B. Scalable Mining of Large Disk-based Graph Databases. KDD 2004.
53. Wang, W., Jiang, h., Lu, H., Yu, j.: Bloom Histograms: Path Selectivity Estimation for XML Data with Updates. In Proc. 30th VLDB Conference, 2004.
54. Wilkinson, K.: Jena Property Table. In Proc. 3rd Jena User Conference, 2006.
55. Winer, F. B. Drug design and the CAS ONLINE Substructure Search System. Drug Inf J. 983;17(4):277-86.
56. Wu, J., Patel, J., Jagadish, H. V.: Estimating Answer Sizes for XML Queries. In Proc. 8th EDBT Conference, 2002.

57. Yan, X., Han, J.: gSpan: Graph-Based Substructure Pattern Mining. In Proc. 2nd ICDM Conference, 2002.
 58. Yan, X., Yu, P. S., Han, J. Graph Indexing: A Frequent Structure-based Approach. In SIGMOD, 2004.
 59. Zhao, P., Yu, J. X., Yu, P. S.: Graph Indexing: Tree + Delta \geq Graph. In VLDB, 2007.



Angela Maduko is a Ph.D candidate at the department of Computer Science in the University of Georgia. She received a B.S. degree in Computer Science with Statistics, from the University of Nigeria. She is a member of the [Large Scale Distributed Information Systems \(LSDIS\) lab](#) on the Semantic Discovery (SemDis) project. Her research interests are in the areas of Databases, Semantic Web, Information Retrieval, Data Mining and Statistical Methods in Bioinformatics.



In August 2007, Dr. Kemafor Anyanwu (Ogan) received a Ph.D degree in Computer Science from the University of Georgia and joined the faculty at the Department of Computer Science of North Carolina State University. She earlier received a B.Sc degree in Biochemistry from the University of Nigeria and worked for a few years at Biotech company (formerly called Serex Inc) in New Jersey. She worked as a graduate research assistant at the Large Scale Distributed Systems Lab at the University of Georgia on the Semantic Discovery (SemDis) Project. She also worked as a research intern with HP Labs, Advanced Database Group, Palo Alto in 2006. Her primary contributions have been in the area of querying Semantic Web databases. Her research interests are in the areas of Semantic Web, Databases, Data Mining, Information Retrieval, Services Computing, Intelligent Data Analytics and Bioinformatics.



[Amit Sheth](#) is an educator, researcher, and entrepreneur. He is the LexisNexis Ohio Eminent Scholar for Advanced Data Management and Analysis, an endowed faculty position funded by LexisNexis and the Ohio Board of Regents at Wright State University. He directs the [Kno.e.sis center](#) for Knowledge enabled Information & Services Science. Professor Sheth is an IEEE Fellow and has received such recognitions as the IBM Faculty award. He is on the editorial boards of several journals and is Editor-in-Chief of the International Journal on Semantic Web and Information Systems ([IJSWIS](#)). During 1994 and 2006, Prof. Sheth was a professor at the University of Georgia. He founded and directed the LSDIS lab at UGA, widely recognized as a

leading international research group in the areas of semantic Web, SOA, and workflows. Earlier, he served in R&D groups at Bellcore, Unisys, and Honeywell. He received his BE from BITS, Pilani in 1981 and MS and PhD from Ohio State University in 1983 and 1985, respectively. Professor Sheth's research has led to several commercial products and two companies, which he founded and managed in various executive roles (COB, CEO, CTO). Keys areas of his interests and contributions are: information integration, metadata, workflow management, Semantic Web, Web and knowledge services, and Services Science. Keys areas of applications of his research and technology development include: life sciences/biomedicine, health care, defense, intelligence, financial services, telecommunication and GIS.



John A. Miller is a Professor of Computer Science at the University of Georgia and has also been the Graduate Coordinator for the department for 9 years. His research interests include database systems, simulation, Web services and bioinformatics. Dr. Miller received a B.S. in Applied Mathematics from Northwestern University in 1980 and an M.S. and Ph.D. in Information and Computer Science from the Georgia Institute of Technology in 1982 and 1986, respectively. During his undergraduate education, he worked as a programmer at the Princeton Plasma Physics Laboratory. Dr. Miller is the author of over 130 publications covering all areas of his research interests. He is an Associate Editor for the ACM Transactions on Modeling and Computer Simulation and IEEE Transactions on Systems, Man and Cybernetics as well as an Editorial Board

Member for the Journal of Simulation and International Journal of Simulation and Process Modelling.



Paul Schliekelman grew up in Iowa. He received a B.S. in Physics from Iowa State University in 1992. After two years in the Peace Corps teaching physics and math in a rural high school in Cameroon, he entered graduate school in Biomathematics at North Carolina State University. His dissertation research focused on mathematical models that explored the effectiveness of newly proposed methods for driving desired genes into the insect pest populations using the mass release of genetically modified insects. After graduating in 2000, he took up a postdoctoral position in the Department of Integrative

Biology at UC Berkeley, working on problems in statistical genetics and evolutionary biology. He started a faculty position at the University of Georgia in 2001 and has been there since. His research interests are in

mathematical and statistical approaches to solving problems in genetics and evolutionary biology. His current research topics include developing methods for integrating genomewide expression information into gene mapping for complex traits, mathematical modeling of genome evolution, and mathematical modeling of the evolution of infectious disease resistance.