# Estimating the Cardinality of Conjunctive Queries over RDF Data Using Graph Summarisation 

Giorgio Stefanoni<br>Bloomberg L.P., London, UK

Boris Motik<br>University of Oxford, UK

Egor V. Kostylev<br>University of Oxford, UK


#### Abstract

Estimating the cardinality (i.e., the number of answers) of conjunctive queries is particularly difficult in RDF systems: queries over RDF data are navigational and thus tend to involve many joins. We present a new, principled cardinality estimation technique based on graph summarisation. We interpret a summary of an RDF graph using a possible world semantics and formalise the estimation problem as computing the expected cardinality over all RDF graphs represented by the summary, and we present a closed-form formula for computing the expectation of arbitrary queries. We also discuss approaches to RDF graph summarisation. Finally, we show empirically that our cardinality technique is more accurate and more consistent, often by orders of magnitude, than the state of the art.


## ACM Reference Format:

Giorgio Stefanoni, Boris Motik, and Egor V. Kostylev. 2018. Estimating the Cardinality of Conjunctive Queries over RDF Data Using Graph Summarisation. In Proceedings of The 2018 Web Conference (WWW 2018). ACM, New York, NY, USA, 17 pages. https://doi.org/10.1145/3178876.3186003

## 1 INTRODUCTION

The Resource Description Framework (RDF) [23] data model is often used in applications where developing a rigid schema is either infeasible or undesirable. An RDF graph is a set of triples of the form $\langle s, p, o\rangle$, where $s, p$, and $o$ are objects called resources; by viewing triples as directed labelled edges, an RDF graph corresponds to an ordinary directed labelled graph. SPARQL [18] is the standard RDF query language. Conjunctive queries (also known as basic graph patterns) are the basic type of SPARQL queries, and estimating their cardinality (i.e., the number of answers) is an important problem. For example, estimates are used to determine the cost of query plans, so estimation accuracy directly influences plan quality [26]. Thus, cardinality estimators are key components of most RDF systems.

Cardinality estimation has a long tradition in databases. It is usually solved by summarising a database to different kinds of synopses that can be used to accurately estimate the cardinality of certain queries. One-dimensional synopses, such as one-dimensional histograms [20,33] and wavelets [28], can summarise one attribute of one relation, so they can be used on queries involving one selection on a single relation. Multidimensional synopses, such as multidimensional histograms $[1,6,16,34]$, multidimensional wavelets [8, 13], discrete cosine transforms [24], and kernel methods [15], can summarise several attributes of one relation, so they can be

[^0]used on queries involving several selections on a single relation. Schema-level synopses, such as graphical models [14, 40], sampling [2, 41], TuG synopses [36], and statistical views [5, 11, 38] can summarise several attributes over different relations, but they often require a join schema that identifies the supported joins.

Queries whose cardinality cannot be estimated using the available synopses are typically broken into subqueries that can be estimated, and partial estimates are combined using ad hoc assumptions [4, 12]: the independence assumption assumes that each selection or join affects the query answers independently, the preservation assumption assumes that each attribute value of a relation participating in a join is present in the join result, and the containment assumption assumes that, for each pair of joined attributes, all values of one attribute are contained in the other attribute. However, these assumptions usually do not hold in practice and thus introduce estimation errors that compound exponentially with the number of joins [21]. Due to the graph-like nature of RDF, queries in RDF often navigate over paths and thus contain many (self-)joins (ten or more are not rare). Techniques that aim to address these specifics of RDF either just adapt relational approaches [19, 31, 37] or focus on particular types of queries (e.g., star or chain queries) [30] and fall back to the ad hoc assumptions for other types of queries. Hence, as we show in Section 6, accurately estimating the cardinality of complex RDF queries remains challenging.

To address these issues, in Section 3 we propose a new, principled technique for estimating the cardinality of conjunctive queries over RDF data. Our technique is based on graph summarisation-the process of compressing a graph by merging its vertices. Summaries have already been used for graph exploration [29, 39], fast approximate graph analytics [25, 35], and query processing [7]. In Section 3, we introduce a specific kind of summary that we use as a schemalevel synopsis for RDF data. Following LeFevre and Terzi [25], we interpret a summary using a possible world semantics as a family of RDF graphs represented by the summary. We formalise the cardinality estimation problem as computing the expectation of the query cardinality across this family, and so our approach does not require any ad hoc assumptions. Finally, we show how to determine the probability that the estimation error exceeds a given bound. To the best of our knowledge, this is the first approach that can provide such guarantees for arbitrary queries.

In Section 4 we present formulas for computing the expected cardinality and its variance. Our formulas can handle queries of the form SELECT * WHERE \{ BGP \} for BGP a basic graph pattern, which form the core of SPARQL. Note that, without DISTINCT, projecting variables in the SELECT clause does not affect the query cardinality.

To apply our framework in practice, effective graph summarisation algorithms are needed. Any summarisation algorithm can be used with our technique in principle, but some will fit our framework better than others: we intuitively expect that a good summary
should merge vertices participating in similar connections. Hence, algorithms for unlabelled graphs [25, 29, 35] are unlikely to work well on RDF, where links are inherently labelled. We experimented with an adaptation of the SNAP [39] technique to RDF, but it produced very large summaries. A recently proposed RDF-specific summarisation technique seems very promising, but its computational complexity seems very high [7]. Thus, in Section 5 we present a new graph summarisation approach. In the first step, we use a simple notion of similarity based on the labels on the edges that resources occur in. If the resulting summary is not sufficiently small, in the second step we further merge resources based on the similarity of the connections resources participate in.

In Section 6 we present the results of an extensive evaluation on LUBM, WatDiv, and DBLP benchmarks with 55 M to 109 M triples, and 15 to 103 queries. We compared the accuracy of our technique against RDF-3X [31] and the characteristic sets technique [30] (both specifically targeting RDF data), PostgreSQL (which uses one-dimensional histograms), and SystemX (a prominent commercial RDBMS that combines histograms with dynamic sampling); we considered both triple table and vertical partitioning for storing RDF in RDBMSs. We show that our technique is generally more accurate than the related approaches, often by orders of magnitude.

The proofs of all of our technical results are given in full in the appendices.

## 2 DEFINITIONS AND NOTATION

We now recall the relevant definitions and notation. RDF vocabulary consist of resources, which are URIs (i.e., abstract entities) or literals (i.e., concrete values described by datatypes such as $x s d$ :string or $x s d$ :integer). A term is a resource or a variable. An (RDF) atom is a triple of the form $\langle s, p, o\rangle$, where $s, p$, and $o$ are terms called the subject, predicate, and object, respectively. An (RDF) triple is a variable-free atom. An (RDF) graph is a finite set of triples.

A substitution $\pi$ is a mapping of variables to terms; the domain and range of $\pi$ are dom $(\pi)$ and $\operatorname{rng}(\pi)$, respectively; varrng $(\pi)$ is the set of all variables in $\mathrm{rng}(\pi)$; and $\pi$ is empty if $\operatorname{dom}(\pi)=\emptyset$. Let $Z$ be a term, an atom, or a set of atoms. Then, $\pi(Z)$ is obtained from $Z$ by replacing each occurrence of $x \in \operatorname{dom}(\pi)$ in $Z$ with $\pi(x)$; if $Z$ is a set of atoms, then $\pi(Z)$ is also a set and does not contain duplicates. Moreover, $\operatorname{res}(Z), \operatorname{var}(Z)$, and term $(Z)$ are the sets of resources, variables, and terms occurring in $Z$, respectively.

SPARQL [18] is a standard query language for RDF. Its syntax is very verbose, so we use a more compact notation that captures basic SPARQL queries of the form SELECT * WHERE \{ BGP \}. SPARQL variables are written as ? $x$, but we drop the question mark and reserve (possibly indexed) letters $x, y$, and $z$ for variables. A conjunctive query $(C Q) q$ is a finite set of atoms. A substitution $\pi$ is an answer to $q$ on an RDF graph $G$ if $\operatorname{dom}(\pi)=\operatorname{var}(q)$ and $\pi(q) \subseteq G$; moreover, $\llbracket q \rrbracket_{G}$ is the set of all answers to $q$ on $G$, and the cardinality of $q$ on $G$ is the size of $\llbracket q \rrbracket_{G}$. For technical reasons, we allow $q$ to be empty, in which case $\llbracket q \rrbracket_{G}$ contains just one empty answer. These definitions are compatible with the ones by Pérez et al. [32]; however, since we consider only CQs without variable projection, we simplify the notation and treat the query as a finite set of atoms.

The q-error is often used to measure precision of cardinality estimates [30]: if a query returns $N$ answers and its cardinality is
estimated as $E$, the $q$-error of the estimate is $e=\max \left(N^{\prime} / E^{\prime}, E^{\prime} / N^{\prime}\right)$, where $N^{\prime}=\max (N, 1)$ and $E^{\prime}=\max (E, 1)$. Intuitively, the q-error shows the difference in the orders of magnitude between a real cardinality $N$ and its estimation $E$, regardless of whether $E$ overor undershoots $N$. We use $N^{\prime}$ and $E^{\prime}$ instead of $N$ and $E$ in the definition of $q$-error in order to avoid division by zero, as well as to prevent artificially high q-error values when $E<1$.

Let $n \geq m>0$ be integers. Then, $\binom{n}{m}=\frac{n!}{m!(n-m)!}$ is the binomial coefficient, $(n)_{m}=n(n-1) \ldots(n-m+1)$ is the falling factorial, and $[m, n]$ is the set of all integers between $m$ and $n$ (inclusive).

## 3 CARDINALITY ESTIMATION FRAMEWORK

To estimate the cardinality of a query without ad hoc assumptions, we must 'compress' our input RDF graph while preserving as much of its structure as possible. Graph summarisation has already been used to this end in related settings [7, 25, 29, 35, 39], and in our work we use it to estimate CQ cardinality in a principled way.

Consider the input graph $G$ in Figure 1 where employees are connected to their cars and managers, and resources are assigned to classes (using the $r d f$ :type predicate) as shown in the legend. A summary $S$ of $G$ is obtained by merging the resources of $G$ into buckets as shown using dotted boxes; for example, employees $e_{1}$ and $e_{2}$ are replaced by bucket $b_{1}$, and $e_{3}$ and $e_{4}$ by another bucket $b_{3}$. In contrast to the existing approaches, in our approach each edge in $S$ is labelled with the number of edges of $G$ that collapse due to merging. For example, the manages edge from $e_{1}$ to $e_{2}$ collapses into a self-loop on $b_{1}$ of weight 1 , and the edges from $e_{1}$ to $e_{3}$ and from $e_{2}$ to $e_{4}$ collapse into an edge from $b_{1}$ to $b_{3}$ of weight 2 .

Our objective is to estimate the cardinality of a CQ $q$ on $G$ using $S$, rather than $G$. To this end, we interpret $S$ using the 'possible worlds' semantics [25]: we assume that $S$ represents with equal probability each graph $G^{\prime}$ that summarises as $S$-that is, if we replace the resources of $G^{\prime}$ with buckets in the same way as when we constructed $S$ from $G$, we obtain $S$. Thus, $S$ represents the input graph $G$, but, due to the loss of information in summarisation, it represents other graphs as well. In Figure 1, in addition to $G$, summary $S$ also represents graphs $G_{1}$ and $G_{2}$. We then estimate the cardinality of $q$ on $G$ as the expected cardinality of $q$ over all graphs represented by $S$. Instead of the assumptions mentioned in Section 1, we thus use a consistent semantic interpretation of $S$.

We next formalise our notion of a summary. We do not talk of a 'summary of an input graph $G$ ': a summary is a synopsis that represents a family of graphs.

Definition 3.1. A summary $S=\langle H, w, \mu\rangle$ is a triple where $H$ is an $R D F$ graph called the summarisation graph, $w: H \rightarrow \mathbb{N}$ is a weight function assigning a positive natural number to each triple in $H$, and $\mu$ is a surjective summarisation function from a finite set of resources dom $(\mu)$ to res $(H)$. The resources of $\mathrm{res}(H)$ are called buckets.

For $Z$ a term, an atom, or a set of atoms, $\mu(Z)$ is obtained from $Z$ by replacing each resource $d \in \operatorname{res}(Z) \cap \operatorname{dom}(\mu)$ with $\mu(d)$; if $Z$ is a set, then $\mu(Z)$ is a set as well (i.e., duplicates are removed). The $S$-size of $a$ bucket $b \in \operatorname{res}(H)$ is defined as $s[b]=|\{d \in \operatorname{dom}(\mu) \mid \mu(d)=b\}|$, and the $S$-size of a triple $\left\langle b_{1}, b_{2}, b_{3}\right\rangle \in \operatorname{res}(H) \times \operatorname{res}(H) \times \operatorname{res}(H)$ is defined as $s\left[\left\langle b_{1}, b_{2}, b_{3}\right\rangle\right]=s\left[b_{1}\right] \times s\left[b_{2}\right] \times s\left[b_{3}\right]$.

While buckets in histograms record value ranges, in our approach $\mu$ explicitly associates buckets with resources; this is needed


Figure 1: An example RDF graph $G$ (a), its summary $S$ (b), and example graphs $G_{1}$ and $G_{2}$ that also summarise as $S$ (c)
as there is no natural notion of a range of RDF resources. In Figure 1, we have $\mu\left(e_{1}\right)=\mu\left(e_{2}\right)=b_{1}, \mu\left(c_{1}\right)=\mu\left(c_{2}\right)=b_{2}, \mu\left(e_{3}\right)=\mu\left(e_{4}\right)=b_{3}$, $\mu\left(c_{3}\right)=\mu\left(c_{4}\right)=b_{4}$, and $\mu$ is identity on all remaining resources, all of which occur in $G$ as a predicate or the object of a triple whose predicate is $r d f: t y p e$ (such as owns, rdf:type, or Van).

The size of a bucket $b$ is the number of resources mapped to $b$; in our example, $s\left[b_{i}\right]=2$ for $i \in[1,4]$, but $s[r d f:$ type $]=s[$ Van $]=1$. For $f \in G$ and $h \in H$ where $\mu(f)=h$, we say that $f$ is summarised as $h$, and that $h$ can be expanded into $f$. Clearly, $h$ can be expanded into at most $s[h]$ triples. In our example, we have

$$
\begin{aligned}
s\left[\left\langle b_{1}, \text { manages, } b_{1}\right\rangle\right] & =s\left[b_{1}\right] \times s[\text { manages }] \times s\left[b_{1}\right]=4, \\
s\left[\left\langle b_{4}, r d f: t y p e, \text { Van }\right\rangle\right] & =s\left[b_{4}\right] \times s[r d f: \text { type }] \times s[\text { Van }]=2 .
\end{aligned}
$$

Definition 3.2 formalises the 'possible worlds' semantics of $S$.
Definition 3.2. A summary $S=\langle H, w, \mu\rangle$ represents a graph $G$ if (i) $\operatorname{res}(G) \subseteq \operatorname{dom}(\mu)$, (ii) $H=\mu(G)$, and (iii) for each $h \in H$, it is the case that $w[h]=|\{f \in G \mid \mu(f)=h\}|$. Moreover, $R_{S}$ is the set of all graphs that $S$ represents. Finally, summary $S$ is consistent if $R_{S} \neq \emptyset$.

At most $s[h]$ triples can be summarised as $h \in H$, so $w[h] \leq s[h]$ must hold for $S$ to be consistent. Also, a graph in $R_{S}$ is obtained by expanding each $h \in H$ arbitrarily into $w[h]$ triples if $w[h] \leq s[h]$ holds. Thus, we can check the consistency of $S$ as in Proposition 3.3.

Proposition 3.3. A summary $S=\langle H, w, \mu\rangle$ is consistent if and only if $w[h] \leq s[h]$ holds for each $h \in H$.

We now estimate the cardinality of a CQ $q$ as the expected cardinality of $q$ over all graphs represented by $S$. We also introduce the related, natural notion of cardinality variance.

Definition 3.4. The expectation $E_{q, S}$ and the variance $\sigma_{q, S}^{2}$ of the cardinality of a CQ q on a summary $S$ are defined as

$$
E_{q, S}=\sum_{G \in R_{S}} \frac{\left|\llbracket q \rrbracket_{G}\right|}{\left|R_{S}\right|} \text { and } \sigma_{q, S}^{2}=\sum_{G \in R_{S}} \frac{\left(\| q q \rrbracket_{G} \mid-E_{q, S}\right)^{2}}{\left|R_{S}\right|}
$$

Please note that $E_{q, S}$ is not necessarily a natural number. For example, if $q$ is a variable-free query, then $\left|\llbracket q \rrbracket_{G}\right|$ is equal to the number of graphs in $R_{S}$ that contain $q$; since $\left|\llbracket q \rrbracket_{G}\right| \leq\left|R_{S}\right|$ is always satisfied, $E_{q, S}$ is a rational number satisfying $0 \leq E_{q, S} \leq 1$.

Using $E_{q, S}$ and $\sigma_{q, S}^{2}$, we can bound from above the probability $\mathrm{P}(e \geq \varepsilon)$ that the q -error of the estimate exceeds a given amount $\varepsilon$ as shown in Theorem 3.5. Intuitively, for a graph $G^{\prime}$ chosen from $R_{S}$ uniformly at random, $\mathrm{P}(e \geq \varepsilon)$ is the probability that $E_{q, S}$ over- or underestimates $\left|\llbracket q \rrbracket_{G^{\prime}}\right|$ by a factor larger than $\varepsilon$; we can equivalently express this by saying that $\mathrm{P}(e \geq \varepsilon) \cdot\left|R_{S}\right|$ is the number of such graphs $G^{\prime}$. We show in Section 6.6 that $\mathrm{P}(e \geq \varepsilon)$ is often small even
for $\varepsilon$ as low as 10 , so in such cases there is good chance that the cardinality of $q$ on our input graph is close to $E_{q, s}$.

Theorem 3.5. Let $q$ be a CQ and let $S$ be a summary such that $E_{q, S} \geq 1$. Then, for each $\varepsilon>1$ and fore the $q$-error of estimating the cardinality of $q$ as $E_{q, S}$, we have

$$
\mathrm{P}(e \geq \varepsilon) \leq\left(\frac{\varepsilon \cdot \sigma_{q, S}}{(\varepsilon-1) \cdot E_{q, S}}\right)^{2} .
$$

If $\varepsilon \geq E_{q, S}$ holds additionally, then the numerator $\varepsilon \cdot \sigma_{q, S}$ on the right-hand side of the inequality can be replaced with $\sigma_{q, s}$.

## 4 COMPUTING THE EXPECTATION

Given a CQ $q$ and a summary $S=\langle H, w, \mu\rangle$, we can compute $E_{q, S}$ by iterating over all graphs in $R_{S}$, but this is impractical since the number of such graphs is exponential in $|H|$. In this section, we present formulas that compute $E_{q, S}$ in time polynomial in $|H|$.

### 4.1 Intuition

Our formulas are quite complex, so we first discuss the intuitions using the summary $S$ for Figure 1 and several example queries.

We first explain how to compute $E_{q, S}$ for variable-free queries, which is the basic element of our approach. Let $q_{1}=\left\{f_{1}, f_{2}\right\}$ where $f_{1}=\left\langle e_{1}\right.$, manages, $\left.e_{3}\right\rangle$ and $f_{2}=\left\langle e_{3}\right.$, owns, $\left.c_{3}\right\rangle$. Since $q_{1}$ is variablefree, for each $G^{\prime} \in R_{S}$, the cardinality $\llbracket q_{1} \rrbracket_{G^{\prime}}$ of $q_{1}$ on $G^{\prime}$ is one if $q_{1} \subseteq G^{\prime}$, and it is zero otherwise; thus, the numerator of $E_{q_{1}, S}$ is equal to the number of graphs in $R_{S}$ that contain all atoms of $q_{1}$ (i.e., that contain $q_{1}$ as a subset). Triples $f_{1}$ and $f_{2}$ of the query are summarised as $h_{1}=\left\langle b_{1}\right.$, manages, $\left.b_{3}\right\rangle$ and $h_{2}=\left\langle b_{3}\right.$,owns, $\left.b_{4}\right\rangle$; for readability, let $w_{i}=w\left[h_{i}\right]$ and $s_{i}=s\left[h_{i}\right]$ for $i \in\{1,2\}$. Now every expansion of $h_{1}$ is determined by a choice of $w_{1}$ triples from $s_{1}$ possibilities, so there are a total of $\left(\begin{array}{c}\binom{s_{1}}{w_{1}} \text { expansions of } h_{1} \text {. Moreover, } \quad \text {, }{ }^{2}\end{array}\right.$ each expansion of $h_{1}$ that contains $f_{1}$ is obtained by choosing $f_{1}$ and the remaining $w_{1}-1$ triples from $s_{1}-1$ possibilities; thus, a total of $\binom{s_{1}-1}{w_{1}-1}$ expansions of $h_{1}$ contain $f_{1}$. Analogously, $h_{2}$ has $\binom{s_{2}}{w_{2}}$ expansions, of which $\binom{s_{2}-1}{w_{2}-1}$ contain $f_{2}$. Now a key observation is that $f_{1}$ and $f_{2}$ are summarised as distinct triples; consequently, the expansions of $h_{1}$ and $h_{2}$ are independent, and so the total number of expansions of $h_{1}$ and $h_{2}$, and the number of expansions containing both $f_{1}$ and $f_{2}$, are given by the product of the above factors. Finally, the expansions of each triple $h \in H \backslash\left\{h_{1}, h_{2}\right\}$ are independent from and thus irrelevant to $q_{1}$, so they do not affect $E_{q_{1}, s}$. Hence, we get

$$
E_{q_{1}, S}=\frac{\binom{s_{1}-1}{w_{1}-1}}{\binom{s_{1}}{w_{1}}} \cdot \frac{\binom{s_{2}-1}{w_{2}-1}}{\binom{s_{2}}{w_{2}}}=\frac{w_{1}}{s_{1}} \cdot \frac{w_{2}}{s_{2}}=\frac{2}{2 \cdot 2} \cdot \frac{2}{2 \cdot 2}=0.25 .
$$

Next, we demonstrate how to handle queries with variables. Let $q_{2}=\{\langle x$, manages, $y\rangle,\langle y$, owns, $z\rangle\}$. For each $G^{\prime} \in R_{S}$ and each $\pi \in \llbracket q_{2} \rrbracket_{G^{\prime}}$, query $\pi\left(q_{2}\right)$ does not contain variables, so we can compute $E_{\pi\left(q_{2}\right), S}$ just as for $q_{1}$. We can thus compute $E_{q_{2}, S}$ by summing the contribution of each such $\pi$, but enumerating all $\pi$ would be inefficient. However, for each $\pi \in \llbracket q_{2} \rrbracket_{G^{\prime}}$, there exists $\tau \in \llbracket \mu\left(q_{2}\right) \rrbracket_{H}$ such that $\mu(\pi(x))=\tau(x)$ for all $x \in \operatorname{dom}(\pi)$; we say that $\tau$ expands into $\pi$. Thus, for each $\tau$, we can compute $E_{\pi\left(q_{2}\right), S}$ for just one prototypical substitution $\pi$ that $\tau$ expands into, and multiply $E_{\pi\left(q_{2}\right), S}$ by the number of the expansions of $\tau$. On our example, evaluating $\mu\left(q_{2}\right)$ on $H$ produces substitutions $\tau_{1}=\left\{x \mapsto b_{1}, y \mapsto b_{3}, z \mapsto b_{4}\right\}$, $\tau_{2}=\left\{x \mapsto b_{1}, y \mapsto b_{1}, z \mapsto b_{2}\right\}, \quad \tau_{3}=\left\{x \mapsto b_{1}, y \mapsto b_{3}, z \mapsto b_{2}\right\}$. For $\tau_{1}$, we can select, say, $\pi=\left\{x \mapsto e_{1}, y \mapsto e_{3}, z \mapsto c_{3}\right\}$ as a prototypical expansion of $\tau_{1}$; then $\pi\left(q_{2}\right)=q_{1}$, and so $E_{\pi\left(q_{2}\right), S}=0.25$. Moreover, since $\tau_{1}$ maps both atoms of $q_{2}$ to distinct triples in $H$, each expansion $\pi$ of $\tau_{1}$ does the same; thus, we can obtain all such $\pi$ by expanding each variable in $\operatorname{dom}\left(\tau_{1}\right)$ independently, and so there are $s\left[\tau_{1}(x)\right] \cdot s\left[\tau_{1}(y)\right] \cdot s\left[\tau_{1}(z)\right]=2 \cdot 2 \cdot 2=8$ expansions of $\tau_{1}$. Thus, $\tau_{1}$ contributes to $E_{q_{2}}, S$ by $8 \cdot 0.25=2$. We compute the contributions of $\tau_{2}$ and $\tau_{3}$ analogously as $2^{3} \cdot \frac{1 \cdot 1}{4 \cdot 4}=0.5$ and $2^{3} \cdot \frac{2 \cdot 1}{4 \cdot 4}=1$. By summing all contributions, we get $E_{q_{2}, S}=2+0.5+1=3.5$.

Generalising to any CQ $q$, formula (1) sums the contribution of each $\tau \in \llbracket \mu(q) \rrbracket_{H}$ : the first factor counts the expansions of $\tau$ to a prototypical $\pi$, and the second one counts the contribution of $\pi$.

$$
\begin{equation*}
E_{q, S}=\sum_{\tau \in \llbracket \mu(q) \rrbracket_{H}} \prod_{x \in \operatorname{var}(q)} s[\tau(x)] \cdot \prod_{a \in q} \frac{w[\tau(\mu(a))]}{s[\tau(\mu(a))]} \tag{1}
\end{equation*}
$$

However, formula (1) is correct only if each answer to $\mu(q)$ in $H$ maps all atoms of $\mu(q)$ to distinct triples in $H$. Consider a query $q_{3}=\left\{\left\langle e_{3}, o w n s, x\right\rangle,\left\langle e_{3}, o w n s, y\right\rangle\right\}$, so $\llbracket \mu\left(q_{3}\right) \rrbracket_{H}$ contains answers

$$
\begin{array}{ll}
\tau_{1}=\left\{x \mapsto b_{2}, y \mapsto b_{4}\right\}, & \tau_{2}=\left\{x \mapsto b_{4}, y \mapsto b_{2}\right\}, \\
\tau_{3}=\left\{x \mapsto b_{4}, y \mapsto b_{4}\right\}, & \tau_{4}=\left\{x \mapsto b_{2}, y \mapsto b_{2}\right\} .
\end{array}
$$

Answers $\tau_{1}$ and $\tau_{2}$ map the atoms of $\mu\left(q_{3}\right)$ to distinct triples in $H$, so their contributions are given by (1) as $2 \cdot 2 \cdot \frac{2}{2 \cdot 2} \cdot \frac{1}{2 \cdot 2}=0.5$. In contrast, $\tau_{3}$ maps both atoms of $\mu\left(q_{3}\right)$ to $h_{3}=\left\langle b_{3}\right.$,owns, $\left.b_{4}\right\rangle$, so $\tau_{3}$ expands into two kinds of $\pi$. First, $\pi$ can map both atoms of $q_{3}$ to the same triple (e.g., $\left\langle e_{3}, o w n s, c_{4}\right\rangle \in G_{2}$ ). Then $\pi\left(q_{3}\right)$ contains just one atom, so $E_{\pi\left(q_{3}\right), S}=w\left[h_{3}\right] / s\left[h_{3}\right]=2 / 4=0.5$; also, $x$ and $y$ are mapped to the same value in each such $\pi$, so the number of such $\pi$ is $N_{1}=s\left[\tau_{3}(x)\right]=2$. Second, $\pi$ can map the atoms of $q_{3}$ to distinct triples. Then, each expansion of $h_{3}$ containing $\pi\left(q_{3}\right)$ corresponds to a choice of $w\left[h_{3}\right]-2$ triples out of $s\left[h_{3}\right]-2$ possibilities, so

$$
E_{\pi\left(q_{3}\right), S}=\frac{\binom{s\left[h_{3}\right]-2}{w\left[h_{3}\right]-2}}{\binom{s\left[h_{3}\right]}{w\left[h_{3}\right]}}=\frac{\left(w\left[h_{3}\right]\right)_{2}}{\left(s\left[h_{3}\right]\right)_{2}}=\frac{(2)_{2}}{(4)_{2}}=\frac{2 \cdot 1}{4 \cdot 3} \approx 0.167
$$

Moreover, there are $N_{2}=s\left[\tau_{3}(x)\right] \cdot s\left[\tau_{3}(y)\right]=2 \cdot 2=4$ expansions of $\tau_{3}$, and $N_{1}=2$ of these map $x$ and $y$ to the same resource; thus, $N_{2}-N_{1}=2$ expansions of $\tau_{3}$ map $x$ and $y$ to distinct resources; and so $\tau_{3}$ contributes to $E_{q_{3}, S}$ by $N_{1} \cdot 0.5+\left(N_{2}-N_{1}\right) \cdot 0.167=1.33$. Finally, answer $\tau_{4}$ maps both atoms of $\mu\left(q_{3}\right)$ to the same triple $h_{4}=\left\langle b_{3}, o w n s, b_{2}\right\rangle$, but no expansion of $\tau_{4}$ contains two triples since $w\left[h_{4}\right]=1$; thus, we consider only $s\left[\tau_{4}(x)\right]=2$ expansions $\pi$ of $\tau_{4}$, each mapping both atoms of $q_{3}$ to the same triple and thus contributing by $E_{\pi\left(q_{3}\right), S}=1 / 2=0.5$; hence, $\tau_{4}$ contributes to $E_{q_{3}, S}$ by $2 \cdot 0.5=1$. Finally, we get $E_{q_{3}, S} \approx 0.5+0.5+1.33+1=3.33$.

### 4.2 Formalisation

For this section, we fix a consistent summary $S=\langle H, w, \mu\rangle$, a CQ $q$ with $\operatorname{res}(q) \subseteq \operatorname{dom}(\mu)$, and an arbitrary total order $\leq$ on terms where $d \leq x$ holds for each resource $d$ and each variable $x$.

We first introduce a notion of a partition of $q$. Intuitively, a partition groups the atoms of $q$ into disjoint sets, and it describes the 'type' of answer to $q$ on graphs in $R_{S}$ where all atoms from each such group are mapped to the same triple.

Definition 4.1. A partition of $q$ is a set $P$ of mutually disjoint nonempty subsets $u_{i}$ of $q$ such that $q=\bigcup_{u_{i} \in P} u_{i}$.

The term graph $\mathcal{G}_{P}$ for $P$ contains all terms from term $(q)$ as vertices, and undirected edges between terms $s_{1}$ and $s_{2}, p_{1}$ and $p_{2}$, and $o_{1}$ and $o_{2}$ for each $u \in P$ and all atoms $\left\langle s_{1}, p_{1}, o_{1}\right\rangle$ and $\left\langle s_{2}, p_{2}, o_{2}\right\rangle$ in u. Partition $P$ is unifiable if no two distinct resources from res $(q)$ are reachable in $\mathcal{G}_{P}$, in which case substitution $\gamma_{P}$ maps each variable $x \in \operatorname{var}(q)$ to the $\leq$-least term reachable from $x$ in $\mathcal{G}_{P}$.

The partition base $B$ for $q$ is the set of all unifiable partitions of $q$. The partial order $\leq$ on $B$ is defined so that, for $P, P^{\prime} \in B$, relationship $P \leq P^{\prime}$ holds iff, for each $u \in P$, there exists $u^{\prime} \in P^{\prime}$ with $u \subseteq u^{\prime}$.

For $P, P^{\prime} \in B$ with $P \leq P^{\prime}, a$ chain from $P$ to $P^{\prime}$ of length $\ell \geq 0$ is a sequence $P=P_{0}, \ldots, P_{\ell}=P^{\prime}$ of partitions from $B$ where $P_{i-1}<P_{i}$ holds for $i \in[1, \ell]$; moreover, $C^{e}\left(P, P^{\prime}\right)$ and $C^{o}\left(P, P^{\prime}\right)$ are the numbers of chains from $P$ to $P^{\prime}$ of even and odd length, respectively; finally, $K\left(P, P^{\prime}\right)=C^{e}\left(P, P^{\prime}\right)-C^{o}\left(P, P^{\prime}\right)$.

Note that $\emptyset$ is the only unifying partition of $q=\emptyset$. Moreover, $C^{e}(P, P)=1$ and $C^{o}(P, P)=0$ for each $P \in B$, so $K(P, P)=1$.

Let $q_{3}=\left\{a_{1}, a_{2}\right\}$ for $a_{1}=\left\langle e_{3}, o w n s, x\right\rangle$ and $a_{2}=\left\langle e_{3}, o w n s, y\right\rangle$ be as in Section 4.1. Then, $P_{1}=\left\{\left\{a_{1}\right\},\left\{a_{2}\right\}\right\}$ and $P_{2}=\left\{\left\{a_{1}, a_{2}\right\}\right\}$ are the unifiable partitions of $q_{3}$ : partition $P_{1}$ represents the answers to $q_{3}$ on graphs in $R_{S}$ that map $a_{1}$ and $a_{2}$ to distinct triples, and $P_{2}$ represents the answers that map $a_{1}$ and $a_{2}$ to the same triple; the latter is captured by $\gamma_{P_{2}}=\{y \mapsto x\}$ (assuming $x \leq y$ ). Also, $P_{1}<P_{2}$ and $K\left(P_{1}, P_{2}\right)=-1$ show that the answers to $q_{3}$ satisfying $P_{2}$ are exactly the answers to $q_{3}$ minus the answers to $q_{3}$ that satisfy $P_{1}$.

For an answer $\tau$ to $\mu(q)$ on $H$ and an appropriate partition $P$, we define coefficients $N_{\tau}(P)$ and $F_{\tau}(P)$ as follows.

Definition 4.2. Let $\tau$ be an answer to $\mu(q)$ on $H$. A partition $P \in B$ is satisfied by $\tau$ if, for each $x \in \operatorname{var}(q)$, (i) $\gamma_{P}(x) \in \operatorname{var}(q)$ implies $\tau(x)=\tau\left(\gamma_{P}(x)\right)$, and (ii) $\gamma_{P}(x) \in \operatorname{res}(q)$ implies $\tau(x)=\mu\left(\gamma_{P}(x)\right)$. Set $B_{\tau}$ contains all partitions of $B$ satisfied by $\tau$. For $P \in B_{\tau}$, let

$$
\begin{aligned}
& N_{\tau}(P)=\prod_{x \in \operatorname{varrng(\gamma _{P})} \text { s }} s[\tau(x)] \\
& F_{\tau}(P)= \begin{cases}\prod_{h \in \tau(\mu(q))} \frac{\text { and }}{} \frac{(w[h])_{H_{h}}}{(s[h])_{*_{h}}} & \text { if\#} \#_{h} \leq w[h] \forall h \in \tau(\mu(q)) \\
0 & \text { otherwise }\end{cases} \\
& \text { where } \#_{h}=|\{u \in P \mid \tau(\mu(u))=\{h\}\}| .
\end{aligned}
$$

In our example, $\gamma_{P_{2}}(x)=\gamma_{P_{2}}(y)$ and $\tau_{1}(x) \neq \tau_{1}(y)$ imply that $P_{2}$ is not satisfied by $\tau_{1}$, which intuitively means that no expansion of $\tau_{1}$ can match the atoms of $q_{3}$ as required by $P_{2}$. Moreover, $N_{\tau_{3}}\left(P_{1}\right)=s\left[\tau_{3}(x)\right] \cdot s\left[\tau_{3}(y)\right]=4$ and $N_{\tau_{3}}\left(P_{2}\right)=s\left[\tau_{3}(x)\right]=2$ count the expansions of $\tau_{3}$ that comply with $\gamma_{P_{1}}$ and $\gamma_{P_{2}}$, respectively; thus, the number of expansions of $\tau_{3}$ complying with just $P_{1}$ is

$$
K\left(P_{1}, P_{1}\right) \cdot N_{\tau_{3}}\left(P_{1}\right)+K\left(P_{1}, P_{2}\right) \cdot N_{\tau_{3}}\left(P_{2}\right)=1 \cdot 4-1 \cdot 2=2 .
$$

The contribution of each expansion of $\tau_{3}$ matching the atoms of $q_{3}$ to distinct triples is $F_{\tau_{3}}\left(P_{1}\right)$ : set $\tau_{3}(\mu(u))$ contains just $h_{3}$ for each $u \in P_{1}$, so $\#_{h_{3}}=2$, and $F_{\tau_{3}}\left(P_{1}\right) \approx 0.167$.

Thus, to compute $E_{q, S}$, we evaluate $\mu(q)$ in $H$ and, for each answer $\tau$, we combine $F_{\tau}(P)$ and $N_{\tau}\left(P^{\prime}\right)$ as in Theorem 4.3. Enumerating all $\tau$ corresponds to query answering and is thus polynomial in $|H|$, and the size of $B_{\tau}$ does not depend on $H$; hence $E_{q, S}$ can be computed in time polynomial in $|H|$. Moreover, Theorem 4.4 shows how to compute the variance of the cardinality of $q$.

Theorem 4.3. The following identity holds:

$$
E_{q, S}=\sum_{\tau \in \llbracket \mu(q) \mathbb{\rrbracket}_{H}} \sum_{P \in B_{\tau}} F_{\tau}(P) \cdot \sum_{P^{\prime} \in B_{\tau}, P \leq P^{\prime}} K\left(P, P^{\prime}\right) \cdot N_{\tau}\left(P^{\prime}\right) .
$$

Theorem 4.4. Identity $\sigma_{q, S}^{2}=E_{q \cup \rho(q), S}-E_{q, S}^{2}$ holds, where $\rho$ is a substitution mapping each variable in $\operatorname{var}(q)$ to a fresh variable.

We next introduce the notion of $\mu$-unification-free queries, where no two distinct atoms can ever be mapped to the same triple. Proposition 4.6 shows that, for such queries, the simpler formula (1) correctly computes the expectation $E_{q, S}$.

Definition 4.5. Atoms $a_{1}$ and $a_{2}$ are unifiable if a substitution $\kappa$ exists such that $\kappa\left(a_{1}\right)=\kappa\left(a_{2}\right)$. Query q is $\mu$-unification-free if no two distinct atoms of $\mu(q)$ are unifiable; otherwise, query $q$ is $\mu$-unifiable.

Proposition 4.6. Formula (1) correctly computes $E_{q, S}$ when $q$ is a $\mu$-unification-free query.

Queries $q_{1}$ and $q_{2}$ from Section 4.1 are $\mu$-unification-free, while $q_{3}$ is not. Moreover, query $q_{4}=\left\{\left\langle e_{3}, o w n s, x\right\rangle,\left\langle e_{4}, o w n s, y\right\rangle\right\}$ is also not $\mu$-unification-free: its atoms unify after applying $\mu$ to $q_{4}$. Note that queries used in practice often consist of atoms of the form $\left\langle t, r, t^{\prime}\right\rangle$ and $\langle t, r d f: t y p e, r\rangle$ where $t$ and $t^{\prime}$ are terms and $r$ is a distinct resource; then, such queries are unification-free whenever each such resource $r$ is assigned to a distinct bucket.

If $q$ is $\mu$-unification-free, then $P=\left\{\left\{a_{i}\right\} \mid a_{i} \in q\right\}$ is the only unifiable partition of $q$, so the estimation formula simplifies to (1). In particular, $P^{\prime}$ in the last sum in Theorem 4.3 can in such cases only be $P$, so the summation reduces to $N_{\tau}(P)$-the first product of (1). Moreover, the atoms of $\mu(q)$ cannot be equated, so $\#_{h}=1$ for each $h \in \tau(\mu(q))$; thus, $F_{\tau}(P)$ reduces to the second product of (1).

## 5 SUMMARISING RDF GRAPHS

As we have mentioned in Section 1, reusing existing graph summarisation techniques proved challenging: some produced summaries that did not lead to accurate cardinality estimates, and others could not be efficiently applied to large graphs. We thus developed a new summarisation approach consisting of two steps we describe next.

### 5.1 Typed Summary

To obtain a coarse measure of resource similarity, we assign to each resource $d$ in an RDF graph $G$ a type $\mathrm{t}_{d}=\langle C(d), O(d), I(d), P(d)\rangle$, which is a tuple consisting of four components described next.

The most important component of $\mathrm{t}_{d}$ is the class type $C(d)$ of $d$, which is the set defined as follows:

$$
C(d)= \begin{cases}\{o \in \operatorname{res}(G) \mid\langle d, r d f: t y p e, o\rangle \in G\} & \text { if } d \text { is a URI, } \\ \{\text { the datatype of } d\} & \text { if } d \text { is a literal. }\end{cases}
$$

One can intuitively expect that a summary whose buckets contain, say, both people and cars, is unlikely to provide good cardinality estimates. Thus, our summarisation approaches will put resources $d_{1}$ and $d_{2}$ into the same bucket only if $C\left(d_{1}\right)=C\left(d_{2}\right)$ holds. A similar idea has been used in query answering over ontologies [9].

Sets $O(d)$ and $I(d)$ describe the outgoing and incoming connections of $d$. Let $\left\langle p_{1}, \ldots, p_{k}\right\rangle$ be an arbitrary ordering of the predicates in $G$ different from $r d f: t y p e$; then, we could capture the outgoing connections of $d$ by a vector $\left\langle n_{1}, \ldots, n_{k}\right\rangle$, where each $n_{i}$ is the number of the outgoing connections from $d$ for predicate $p_{i}$. This, however, is likely to create too many types: resources in an RDF graph are unlikely to have exactly the same outgoing connections with the same cardinalities. We thus represent the cardinality information using histograms. For each predicate $p_{i}$, we compute the set out $p_{p_{i}}$ that, for each resource $d_{j}$ occurring in $G$ in subject position, contains a pair $\left\langle d_{j},\right|\left\{o \in \operatorname{res}(G)\right.$ such that $\left.\left\langle d_{j}, p_{i}, o\right\rangle \in G\right\}\rangle$ associating $d_{j}$ with the number of resources that $d_{j}$ is connected to via property $p_{i}$. We sort the pairs in out $p_{i}$ by the value of their second component into a list $\left\langle d_{1}, n_{1}\right\rangle, \ldots,\left\langle d_{N}, n_{N}\right\rangle$. Finally, given a predetermined number $J_{i}$ of histogram buckets (not to be confused with summary buckets) for $p_{i}$, we create a histogram for out $p_{i}$ : for $w=\left\lceil N / J_{i}\right\rceil$, the first bucket contains resources $d_{1}, \ldots, d_{w}$ and has ID 1, the second bucket contains resources $d_{w+1}, \ldots, d_{2 w}$ and has ID 2, and so on. Thus, each bucket contains resources with similar outgoing frequencies for $p_{i}$, and the difference between bucket IDs is indicative of the bucket similarity. Then, the outgoing type $O(d)$ of a resource $d$ is the vector $\left\langle O_{1}, \ldots, O_{k}\right\rangle$, where each $O_{i}$ is the ID of the bucket that $d$ was assigned to in the histogram for $p_{i}$. The incoming type $I(d)$ of $d$ is defined analogously. The numbers of buckets can be selected independently for each predicate and direction: it is reasonable to use a larger number of buckets if the relevant frequencies vary significantly across resources.

Finally, we also identify highly connected groups of resources: one can intuitively expect that such resources should be assigned to the same bucket. To this end, we divide the resources into a predetermined number of partitions of roughly the same size while minimising the number of edges between partitions using the algorithms by Karypis and Kumar [22]. The partition type $P(d)$ of a resource $d$ is the partition to which $d$ was assigned to.

Thus, if the types $\mathrm{t}_{d_{1}}$ and $\mathrm{t}_{d_{2}}$ of resources $d_{1}$ and $d_{2}$ are the same, then $d_{1}$ and $d_{2}$ occur in the same classes (or datatypes), have comparable outgoing and incoming frequencies for all predicates, and belong to the same highly connected component of the RDF graph; hence, we can expect $d_{1}$ and $d_{2}$ to be similar. Thus, we define the typed summary $S=\langle H, w, \mu\rangle$ of an RDF $G$ as follows: for each resource $d \in \operatorname{res}(G)$, if $d$ occurs in $G$ in triples of the form $\langle s, d, o\rangle$ or $\langle s, r d f: t y p e, d\rangle$, we let $\mu(d)=d$; otherwise, we let $\mu(d)=b_{\mathrm{t}_{d}}$, where $b_{\mathrm{t}_{d}}$ is a distinct bucket uniquely associated with the type $\mathrm{t}_{d}$ of $d$. Finally, we define $H=\mu(G)$ and adjust the function $w$ accordingly.

### 5.2 Summary Refinement by MinHashing

If a typed summary is large, we compress it by merging similar buckets and types. To identify similar buckets, we define the vicinity $V_{S}(b)$ of a bucket $b \in \operatorname{res}(H)$ as the union of the following sets:

$$
\begin{aligned}
V_{S}^{\text {out }}(b) & =\left\{\left\langle b_{p}, b_{o}\right\rangle \mid\left\langle b, b_{p}, b_{o}\right\rangle \in H \wedge b_{p} \neq \mu(r d f: t y p e)\right\}, \\
V_{S}^{\text {in }}(b) & =\left\{\left\langle b_{i}, b_{p}\right\rangle \mid\left\langle b_{i}, b_{p}, b\right\rangle \in H \wedge b_{p} \neq \mu(r d f: t y p e)\right\} .
\end{aligned}
$$

Since vicinity $V_{S}(b)$ describes the connections of $b$, given buckets $b_{1}, b_{2} \in \operatorname{res}(H)$, the degree of commonality between $V_{S}\left(b_{1}\right)$ and $V_{S}\left(b_{2}\right)$ is indicative of the similarity of $b_{1}$ and $b_{2}$. This is captured by the faccard index [27] of $b_{1}$ and $b_{2}$, which is defined as

$$
\mathcal{J}_{S}\left(b_{1}, b_{2}\right)=\frac{\left|V_{S}\left(b_{1}\right) \cap V_{S}\left(b_{2}\right)\right|}{\left|V_{S}\left(b_{1}\right) \cup V_{S}\left(b_{2}\right)\right|}
$$

if $V_{S}\left(b_{1}\right) \cup V_{S}\left(b_{2}\right) \neq \emptyset$, and otherwise $\mathcal{J}_{S}\left(b_{1}, b_{2}\right)=1$. Thus, we can naïvely reduce the size of $S$ by repeatedly merging the pair of buckets $b_{1}$ and $b_{2}$ with maximal $\mathcal{J}_{S}\left(b_{1}, b_{2}\right)$, but this is impractical for two reasons. First, computing $\mathcal{J}_{S}\left(b_{1}, b_{2}\right)$ requires iterating over $V_{S}\left(b_{1}\right)$ and $V_{S}\left(b_{2}\right)$, which can be large. Second, the number of pairs of $b_{1}$ and $b_{2}$ can be large in even moderately sized summaries.

We address the first problems using MinHashing [27], which uses fixed-sized signatures to approximate $\mathcal{J}_{S}\left(b_{1}, b_{2}\right)$. Given two integer parameters $m$ and $n$, we generate a MinHash scheme $\mathcal{F}$ of size $m \times n$, which is an $m \times n$ matrix of hash functions chosen uniformly at random with replacement. Each $\mathcal{F}[i, j]$ maps elements of $V_{S}(b)$ (i.e., pairs of resources) to natural numbers. The signature of a bucket $b \in \operatorname{res}(H)$ on $\mathcal{F}$ and $S$ is the $m \times n$ matrix $\mathcal{M}^{b}$ where

$$
\mathcal{M}^{b}[i, j]=\min _{\alpha \in V_{S}(b)} \mathcal{F}[i, j](\alpha)
$$

if $V_{S}(b) \neq \emptyset$, and otherwise $\mathcal{M}^{b}[i, j]=\infty$. It is known that

$$
\hat{\mathcal{J}}_{S}\left(b_{1}, b_{2}\right)=\frac{\left|\left\{\langle i, j\rangle \mid \mathcal{M}^{b_{1}}[i, j]=\mathcal{M}^{b_{2}}[i, j]\right\}\right|}{m \cdot n}
$$

is a good approximation of the Jaccard index of buckets $b_{1}$ and $b_{2}$ [27]. For $i \in[1, m]$, we write $\mathcal{M}^{b}[i]$ for the $i$-th row of $\mathcal{M}^{b}$.

We address the second problem using locality sensitive hashing [27]: we generate a hash function LSH for $\mathcal{M}^{b}[i]$ and use it to assign each bucket $b$ to a $\operatorname{bin} \operatorname{Bins}\left[\operatorname{Lsh}\left(\mathcal{M}^{b}[i]\right)\right]$. For each $\operatorname{Bin} \in \operatorname{Bins}$ and all buckets $b_{1}, b_{2} \in \operatorname{Bin}$, it is known that $\hat{\mathcal{J}}_{S}\left(b_{1}, b_{2}\right)$ is likely to be high, so all pairs of buckets in Bin are likely to be similar [27].

Algorithm 1 uses these ideas to compute a summary of a graph $G$. It takes as arguments a set of types T covering all resources of $G$. The algorithm first converts $G$ into a trivial summary $S$ (line 3) where $H=G$ and $\mu(d)=d$ for each $d \in \operatorname{res}(G)$, and then it computes the sets $\mathcal{B}_{\mathrm{t}}$ of buckets of type t for each $\mathrm{t} \in \mathrm{T}$ (line 4). The algorithm next enters a loop (lines 5-21) that merges buckets of the same type. It returns $S$ if the number of edges in $H$ is in the required space budget (line 6); otherwise, it considers each type $t \in T$ (lines 8-16) and computes the signature of each bucket $b \in \mathcal{B}_{\mathrm{t}}$ (line 9). For each row $i$ of the signature (lines $10-16$ ), the algorithm applies locality sensitive hashing to each bucket $b \in \mathcal{B}_{\mathrm{t}}$ (line 12) and adds $b$ to a bin. The number of bins is not predetermined, which reduces collisions. Finally, for all bins containing at least two buckets, all buckets in the bin are scheduled to be merged into one designated bucket $b$ (lines 13-16). Locality sensitive hashing ensures that similar buckets are likely to end in the same bin, but this is not guaranteed so the standard algorithm would merge $b^{\prime}$ and $b$ only if the estimate $\hat{\mathcal{J}}_{S}\left(b, b^{\prime}\right)$ is above a particular threshold. We skip this step for two reasons: the quadratic step of computing $\hat{\mathcal{J}}_{S}\left(b, b^{\prime}\right)$ for all pairs of buckets in a bin can be prohibitive, and the chance that dissimilar buckets are assigned to the same bin is small since we do not limit the number of bins. After all signature rows have been considered, the merges are applied to the summary (line 18). Finally, if the size of $S$ has not been reduced sufficiently, then the types in T are too

```
Algorithm 1 MinHash summarisation algorithm
    Inputs: \(G\) : an RDF graph
            T : a set of types
            \(m \times n\) : size of the MinHash scheme
            target : target size of the summary
    Generate a MinHash scheme \(\mathcal{F}\) of size \(m \times n\)
    Generate a locality sensitive hash LSH : \(\mathbb{N}^{n} \rightarrow \mathbb{N}\)
    Let \(S:=\langle H, w, \mu\rangle\) be the trivial summary of \(G\)
    for all \(\mathrm{t} \in \mathrm{T}\) do \(\mathcal{B}_{\mathrm{t}}:=\{b \in \operatorname{res}(H) \mid b\) is of type t\(\}\)
    loop
        if \(|H| \leq\) target then return \(S\)
        queue \(:=\emptyset\)
        for all \(t \in T\) do
            for all \(b \in \mathcal{B}_{\mathrm{t}}\) do Compute \(\mathcal{M}^{b}\) using \(\mathcal{F}\) and \(S\)
            for \(i:=1\) to \(m\) do
            Bins := \(\emptyset\)
            for all \(b \in \mathcal{B}_{\mathrm{t}}\) do Add \(b\) to \(\operatorname{Bins}\left[\operatorname{LsH}\left(\mathcal{M}^{b}[i]\right)\right]\)
            for Bin \(\in\) Bins such that \(\mid\) Bin \(\mid \geq 2\) do
                Choose some \(b \in \operatorname{Bin}\)
                for \(b^{\prime} \in \operatorname{Bin}\) with \(b \neq b^{\prime}\) do
                    Add \(\left\langle b^{\prime}, b\right\rangle\) to queue and remove \(b^{\prime}\) from \(\mathcal{B}_{\mathrm{t}}\)
        sizeBefore \(:=|H|\)
        for all \(\left\langle b^{\prime}, b\right\rangle \in\) queue do Merge \(b^{\prime}\) into \(b\) in \(S\)
        if sizeBefore \(-|H| \leq(|H|-\) target \() \cdot 0.01\) then
            Merge similar types in \(T\)
            if no merge has happened then return \(S\)
```

specific to allow further merges, so the algorithm merges similar types in T (line 20) to allow further reducing in the size of $S$.

Let t and $\mathrm{t}^{\prime}$ be types whose outgoing and incoming types are $\left\langle O_{1}, \ldots, O_{k}\right\rangle$ and $\left\langle O_{1}^{\prime}, \ldots, O_{k}^{\prime}\right\rangle$, and $\left\langle I_{1}, \ldots, I_{k}\right\rangle$ and $\left\langle I_{1}^{\prime}, \ldots, I_{k}^{\prime}\right\rangle$. Then, t and $\mathrm{t}^{\prime}$ can be similar only if their class and partition types coincide, and their similarity $\mathcal{J}\left(\mathrm{t}, \mathrm{t}^{\prime}\right)$ is the average of the generalised Jaccard indexes [27] of their outgoing and incoming types:

$$
\mathcal{J}\left(\mathrm{t}, \mathrm{t}^{\prime}\right)=\frac{1}{2}\left(\sum_{i} \frac{\min \left(O_{i}, O_{i}^{\prime}\right)}{\max \left(O_{i}, O_{i}^{\prime}\right)}+\sum_{i} \frac{\min \left(I_{i}, I_{i}^{\prime}\right)}{\max \left(I_{i}, I_{i}^{\prime}\right)}\right) .
$$

Merging t and $\mathrm{t}^{\prime}$ produces $\mathrm{t}^{\prime \prime}$ with the class and partition types same as t and $\mathrm{t}^{\prime}$, outgoing type $\left\langle\left(O_{1}+O_{1}^{\prime}\right) / 2, \ldots,\left(O_{k}+O_{k}^{\prime}\right) / 2\right\rangle$, and an analogous incoming type; we also set $\mathcal{B}_{\mathrm{t}^{\prime \prime}}=\mathcal{B}_{\mathrm{t}} \cup \mathcal{B}_{\mathrm{t}^{\prime}}$.

To merge the types in line 20 , the algorithm first groups the types by their class and partition type. Ideally, we would merge the most similar pairs within each such group, but computing the similarity of each pair in the group would be prohibitively expensive. Instead, the algorithm selects 500 pairs at random, computes their similarity, and merges the most similar pair if their similarity is at least $50 \%$. This process is repeated until the size of T is reduced by $20 \%$ (if possible), which then allows for further merging of buckets.

### 5.3 Choosing the Parameters

To recapitulate, our algorithm is parametrised by (i) the number of buckets for the incoming and outgoing histograms per predicate, (ii) the number of partitions, (iii) the target size of the summary,
and (iv) the size of the MinHash scheme $m \times n$ (if used). Standard techniques can be used to choose the number of buckets, such as the Freedman-Diaconis rule [10]. We found it reasonable to partition a graph into 10 partitions, and fall back to just one partition if more than $20 \%$ of the edges are cut. We also identified $m=20, n=2$, and target sizes in the order of tens of thousands as reasonable defaults.

## 6 EXPERIMENTAL EVALUATION

We implemented our techniques in a prototype system called SumRDF. ${ }^{1}$ The system is written in Java. It evaluates queries over graphs stored in RAM using left-deep index nested loop joins. We evaluated our system with three distinct objectives. First, we compared the accuracy of the estimates produced by SumRDF and several state-of-the-art RDF and relational systems. Second, we investigated the impact of the correct handling of $\mu$-unifiable queries; these correspond to self-joins and are difficult to estimate. Third, we evaluated the usefulness of the bounds provided by Theorem 3.5. We ran our experiments on a Dell computer with 96 GB of RAM, Fedora 22, and two Xeon X5667 processors with 16 physical cores.

### 6.1 Test Data

We used three well-known RDF benchmarks shown in Table 1. We group queries by their shape as linear, star, snowflake (joined stars), and complex (e.g., cyclic). For each group, we show the number of all queries, the number of $\mu$-unifiable queries, and the minimal and maximal numbers of atoms in a query. Each query is assigned an ID. All datasets and queries are available on the SumRDF Web site.
LUBM [17] is a synthetic benchmark comprising a data generator, an OWL 2 DL ontology, and 14 test queries. The ontology must be taken into account for the queries to produce results, so we generated an RDF graph with 500 universities and extended it with triples implied by the ontology, thus obtaining 91.1 M triples. The 14 queries are relatively simple, and are all $\mu$-unification-free, so we additionally handcrafted 24 queries, five of which are $\mu$-unifiable.
WatDiv [3] is another synthetic benchmark designed to stress-test RDF systems. We used its data generator to obtain an RDF graph with 108.9 M triples. The benchmark also includes three queries and 17 query templates with one parameter, and a generator that replaces the parameter with resources from the RDF graph. We instantiated each query template into five concrete queries and, if the parameter was not in the class position, we obtained one more query by replacing the parameter with a fresh variable. We thus obtained 103 queries overall, all of which are $\mu$-unification-free.
DBLP is a bibliography database that was converted into an RDF graph of 55.5 M triples. We used six queries already been considered in the literature [42], and we handcrafted nine additional queries. We thus obtained 15 queries, eight of which are $\mu$-unifiable.

### 6.2 Constructing the Summaries

LUBM can be easily partitioned into ten partitions, and the numbers of connections do not vary much across resources so we used just one bucket per histogram. The typed summary for LUBM contained fewer than 15,000 edges, so the refinement step was not needed. In contrast, WatDiv and DBLP could not be effectively partitioned, so

[^1]we used just one partition type, and we manually selected between one and 64 histogram buckets per predicate. The typed summaries were large so we further refined them using $m=20, n=2$, and target of 45,000 for WatDiv and 30,000 for DBLP. Table 1 shows that each graph was compressed by a factor more than 1,000 .

### 6.3 Comparison Systems

We evaluated our system against a portfolio of the following four systems, which represent the state of the art in open-source and commercial (RDF) data management.
RDF-3X [31] v0.3.7 is a state of the art RDF system whose cardinality estimator was specifically tailored to RDF and was shown to outperform the related approaches.
C-SET is the characteristic sets technique [30] that specifically targets RDF and pays particular attention to star queries. There is no publicly available implementation of this technique, so we reimplemented it ourselves. It was shown to outperform the approach by Stocker et al. [37], so we did not test the latter separately.
PostgreSQL v9.4.5 uses a conventional cardinality estimator with one-dimensional histograms. Statistics collection was automatic.
SystemX is a commercial RDBMS. ${ }^{2}$ Its highly tuned cardinality estimator operates in two modes: in the static mode the system uses statistics collected automatically prior to estimation, and in the dynamic mode it can sample the data during estimation if the estimate accuracy is low. We used the maximum sampling level.
To store RDF data into RDBMSs, we replaced all resources with unique integers, and then we considered two storage modes. First, we stored all triples in one ternary triple table, so all queries amount to self-joins over this table. Second, we vertically partitioned the data by storing $\langle s, r d f:$ type, $o\rangle$ as tuple $\langle s\rangle$ in a unary relation $o$, and $\langle s, p, o\rangle$ with $p \neq r d f$ :type as tuple $\langle s, o\rangle$ in a binary relation $p$. Using multiple tables generally allows for more accurate statistics collection. We thus obtained a total of eight comparison systems: RDF-3X, C-SET, two (P-T and P-V) PostgreSQL variants, and two static (XS-T and XS-V) and two dynamic (XD-T and XD-V) SystemX variants. For systems other than C-SET, we obtained their cardinality estimates using their EXPLAIN facility.

### 6.4 Accuracy Experiments

Figure 2 shows the q-error distribution per dataset and system, thus summarising the results of our accuracy experiments on all 156 test queries. Each bar chart shows the percentages of the queries with q-error in the intervals from the legend. Each box plot groups the q -errors per system: the box shows the lower and upper quartiles of the $q$-error, the line dividing the box shows the median, the line whiskers show the minimum and maximum, and the diamond mark shows the average. Please note that the maximum q-errors sometimes fall beyond the shown portion of the $y$ axis.

As one can see from the figure, SumRDF outperforms the other systems by several orders of magnitude on LUBM and WatDiv. On DBLP, it exhibits the best maximum and average $q$-errors, and SystemX with dynamic sampling is the only system that provides comparable performance. We conjecture the latter to be due to

[^2]Table 1: Dataset, summary, and query statistics

| Benchmark | Original graph $G$ |  | Summary $S$ |  |  |  | Queries by type |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Nr | Nr | Nr | Nr | Reduction | Construction | Linear |  |  |  | Star |  |  |  | Snowflake |  |  |  | Complex |  |  |  |
|  | resources | triples | buckets | triples | factor | time | total | $\mu$-un. | min | max | total | $\mu$-un. | min | max | total | $\mu$-un. | min | max | total | $\mu$-un. | min | max |
| LUBM | 16.4 M | 91.1 M | 313 | 14,926 | 6.1 k | 6 min | 16 | 3 | 1 | 6 | 2 | 0 | 4 | 10 | 6 | 0 | 4 | 10 | 14 | 2 | 4 | 9 |
| WatDiv | 10.2 M | 108.9 M | 5,504 | 43,350 | 2.5 k | 2 h | 30 | 0 | 2 | 3 | 40 | 0 | 2 | 9 | 30 | 0 | 5 | 9 | 3 | 0 | 6 | 10 |
| DBLP | 25.4 M | 55.5 M | 4,765 | 28,631 | 1.9 k | 45 min | 4 | 0 | 1 | 4 | 1 | 1 | 3 | 3 | 2 | 1 | 6 | 7 | 8 | 6 | 5 | 11 |

Table 2: Estimates confidence for varying q-errors

| Q-Error $\varepsilon$ | 10 |  |  |  |  | 100 |  |  |  |  | 1000 |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| The \% bound on $\mathrm{P}(e \geq \varepsilon)$ from Theorem 3.5 | [0,1] | (1, 5] | $(5,10]$ | (10, 25] | > 25 | [0,1] | (1, 5] | $(5,10]$ | $(10,25]$ | > 25 | [0,1] | (1, 5] | (5, 10] | $(10,25]$ | $>25$ |
| LUBM (29) | 20 (2) | 0 (0) | 6 (1) | 2 (0) | 1 (0) | 25 (3) | 0 (0) | 4 (0) | 0 (0) | 0 (0) | 28 (1) | 1 (0) | 0 (0) | 0 (0) | 0 (0) |
| WatDiv (48) | 16 (1) | 6 (1) | 1 (0) | 10 (2) | 15 (3) | 42 (0) | 1 (0) | 0 (0) | 0 (0) | 5 (0) | 43 (0) | 0 (0) | 0 (0) | 0 (0) | 5 (0) |
| DBLP (3) | 2 (0) | 1 (0) | 0 (0) | 0 (0) | 0 (0) | 3 (0) | 0 (0) | 0 (0) | 0 (0) | 0 (0) | 3 (0) | 0 (0) | 0 (0) | 0 (0) | 0 (0) |



Figure 2: Histograms and box plots of the distribution of the q-error on test datasets

Table 3：The q－error for $\mu$－unifiable queries

| $\theta$ |  |  |  | $$ | ち | i | $\stackrel{\ominus}{\grave{x}}$ | $\begin{aligned} & \overrightarrow{\grave{x}} \\ & \vec{x} \end{aligned}$ | ث | 交 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\sum_{0}$ <br> U <br> U 2 <br> U 2 <br> U <br> U 4 <br> $\mathrm{U5}$ | 1.4 k111.0 k1.11.21.1 | $\begin{array}{\|r\|} \hline 769.8 \\ 926.3 \\ 1.4 \\ 1.8 \\ 1.0 \end{array}$ | $\begin{array}{\|c\|} \hline 71.3 \\ 111.0 \mathrm{k} \\ 254.5 \\ 455.4 \\ 5.9 \end{array}$ | 205.2 k | 3.5 T | 122.5 | 26.3 G | 2.6 k | 17.5 k | 3.4 |
|  |  |  |  | 111.0 k | 7.9 T | 111.0 k | 12.3 M | 111.0 k | 982.0 | 111.0 k |
|  |  |  |  | 67.5 k | 25 G | 2.7 | 87.3 M | 1.2 | 8 k | 1.0 |
|  |  |  |  | 763.9 k | 650 G | 9.8 | 49.2 M | 1.1 | 2.7 k | 1.2 |
|  |  |  |  | 1.3 | 11.2 k | 1.6 | 10.2 k | 1.5 | 12.2 k | 1.5 |
| D 1 <br> D 2 <br> D 3 <br> D 4 <br> D 4 <br> D 5 <br> D 6 <br> D 7 <br> D 8 <br> D | 1.2 | 1.2 | 9.9 k | 440.8 k | 106.3 | 3.4 | 220.4 k | 1.4 | 2.8 | 2.8 |
|  | 2.1 | 2.1 | 1.0 M | 1.0 M | 23.2 k | 40.0 | 1.0 M | 8.8 | 15.2 | 3.2 |
|  | 238.0 | 19.9 | 11.9 k | 238.0 | 45.4 | 238.0 | 238.0 | 238.0 | 79.3 | 238.0 |
|  | 320.2 | 320.2 | 3.9 k | 29.7 k | 432.1 | 1.0 k | 213.5 k | 5.5 k | 6.3 k | 55.8 |
|  | 721.3 | ． 3 | 313.3 | 4.1 k | 3.9 | 344.4 | 4.1 k | 4.1 | 2.0 | 4.1 k |
|  | 57.0 | 57.0 | 57.0 | 57.0 | 2.7 | 57.0 | 57.0 | 57.0 | 57.0 | 57.0 |
|  | 210.3 | 210.3 | 10.1 k | 10.1 k | 10.1 k | 561.3 | 10.1 k | 76.5 | 10.1 k | 8.2 |
|  | 817.0 | 817.0 | 817.0 | 817.0 | 38.9 | 817.0 | 817.0 | 817.0 | 817.0 | 817.0 |
| min | 1.1 | 1.0 | 5.9 | 1.3 | 2.7 | 1.6 | 57.0 | 1.1 | 2.0 | 1.0 |
| med | 210.3 | 57.0 | 817.0 | 29.7 k | 10.1 k | 122.5 | 213.5 k | 76.6 | 982.1 | 55.8 |
| avg | 8.8 k | 296.1 | 90.1 k | 204.3 k | 940.8 G | 8.8 k | 31.8 M | 9.5 k | 4.3 k | 9.2 k |
| max | 111.0 k | 926.2 | 1.0 M | 1.0 M | 7.9 T | 111.0 k | 263.4 M | 111.0 k | 17.6 k | 111.0 k |

selections in queries：if a query contains a resource that has not been captured accurately using histograms，more accurate statistics are obtained on demand．Overall，our technique considerably improves the state of the art in cardinality estimation of RDF queries．

We noticed no pattern in the systems＇behaviour on queries with q－error above 10：each system undershoots on some，but overshoots on other queries．Different variants of the same system often differ greatly on the same query；for example，P－T might overshoot while P－V would undershoot，and similarly for SystemX．Also，for each dataset and system，there is a query where the system is most accurate，but SumRDF is never the least accurate．

Atoms of the form $\langle x, r d f: t y p e, C\rangle$ ，whose presence in the query does not affect the query answers，were a common source of estima－ tion errors．In C－SET，P－T，and XS－T，such atoms impose a selection on the rdf：type property that is combined using independence as－ sumption，which usually leads to a significant underestimation．

We further analysed difficult queries，where SumRDF exhibited a q－error above 10．Five queries of LUBM have q－errors between 750 and 1,000 ；three are triangular and two are cyclic．Nine queries of WatDiv have $q$－error between 10 to 20 ，and all of these contain one or two selections．Finally，on DBLP，three queries have q－errors between 10 and 20，and additional five have $q$－errors above 20 ； all are either cyclic or have up to three selections．Thus，queries containing cycles and／or selections are sometimes hard for SuM－ RDF．However，most systems exhibited a q－error above 10 on these queries as well．SystemX with dynamic sampling performed better on acyclic queries with selections：the system would additionally sample whenever the statistics were insufficiently precise．

## 6．5 Impact of $\mu$－Unification

Formula（1）for $\mu$－unification－free queries does not require comput－ ing the partition base for the query，so the formula should be easier to evaluate than the general formula from Theorem 4．3．Hence，
we investigated whether using the general formula improves es－ timation accuracy．Table 3 compares the $q$－error of the estimate computed using formula（1）with the $q$－errors of all other systems （including SumRDF）for all $\mu$－unifiable queries．WatDiv does not occur in the table since all of its queries are $\mu$－unification－free．

As one can see，$\mu$－unifiable queries tend to be difficult：for each query，at least one system exhibits a q－error above 900 ．However， handling $\mu$－unifiable queries correctly reduces the maximum error by two orders of magnitude，and it also considerably improves the median and the average．Finally，SumRDF considerably outperforms all other systems on all aggregate metrics，apart from the median of XD－V，whereas the aggregate metrics of formula（1）seem sim－ ilar to $\mathrm{P}-\mathrm{V}$ ．Thus，our principled approach to $\mu$－unifiable queries contributes substantially to estimation accuracy．

## 6．6 Probabilistic Error Bounds

For $\varepsilon$ equal to 10,100 ，and 1,000 ，Table 2 shows the bound on the probability of q －error exceeding $\varepsilon$ computed as stated in Theo－ rem 3．5．The probabilities（expressed as percentages）are grouped into five intervals．Each table cell shows the number of queries in the interval，and in parentheses the number of queries where the q －error indeed exceeds $\varepsilon$ ．For example，on LUBM，there are 20 queries for which the probability of $q$－error being larger than 10 is less than $1 \%$ ，and for two of these the $q$－error is larger than 10 ．

We could not compute $\sigma_{q, s}$ in several cases：query $q \cup \rho(q)$ from Theorem 4.4 is always $\mu$－unifiable and tends to be very complex． Moreover，Theorem 3.5 does not apply to queries with $E_{q, S}<1$ ． The numbers in parentheses next to each dataset name show the number of queries on which we could compute $\mathrm{P}(e \geq \varepsilon)$ ．

Already for $\varepsilon=10$ ，we can correctly bound with $99 \%$ certainty the estimates for 18 out of 29 LUBM queries， 15 out of 48 WatDiv queries，and two out of three DBLP queries．With $\varepsilon=100$ ，we cover $75 \%$ of LUBM， $87 \%$ of WatDiv，and all of the DBLP queries．These results suggest that the confidence bounds provided by Theorem 3.5 are indeed very useful．In our future work，we shall incorporate them into query planning algorithms so that a planner can prefer plans with higher confidence cardinality estimates．

## 7 CONCLUSION AND OUTLOOK

We presented a new approach for estimating the cardinality of CQs on RDF graphs．Our approach is based on graph summarisation and it formalises the estimation problem based on a precise＇possible worlds＇semantics．Our technique can also provide statistical con－ fidence for an estimate，even for arbitrarily shaped，complex CQs． Finally，we showed experimentally that our approach outperforms state of the art RDF and relational cardinality estimators．

We see many exciting opportunities for future work．On the the－ oretical side，supporting range queries should be easy，but adding DISTINCT and aggregation is likely to be more involved．We shall also try to incorporate further information about the data and thus reduce the number of graphs represented by a summary．On the practical side，we shall further analyse the sources of errors and develop ways of computing more precise summaries，possibly by incorporating variants of dynamic sampling．We shall also inves－ tigate how to update the summary without recomputing it from scratch when the underlying graph is updated．

## ACKNOWLEDGMENTS

The research presented in this paper was supported by the EPSRC projects $\mathrm{MaSI}^{3}$ (EP/P025943/1), DBOnto (EP/L012138/1), and ED ${ }^{3}$ (EP/N014359/1).

## REFERENCES

[1] Ashraf Aboulnaga and Surajit Chaudhuri. 1999. Self-tuning Histograms: Building Histograms Without Looking at Data. In Proc. of the ACM SIGMOD Int. Conf. on Management of Data (SIGMOD 1999). ACM, Philadelphia, PA, USA, 181-192.
[2] Swarup Acharya, Phillip B. Gibbons, Viswanath Poosala, and Sridhar Ramaswamy. 1999. Join Synopses for Approximate Query Answering. SIGMOD Record 28, 2 (1999), 275-286.
[3] Günes Aluç, Olaf Hartig, M. Tamer Özsu, and Khuzaima Daudjee. 2014. Diversified Stress Testing of RDF Data Management Systems. In Proc. of the 13th Int. Semantic Web Conference (ISWC 2014) (LNCS), Vol. 8796. Springer, Riva del Garda, Italy, 197-212.
[4] Nicolas Bruno. 2003. Statistics on Query Expressions in Relational Database Management Systems. Ph.D. Dissertation. Columbia University.
[5] Nicolas Bruno and Surajit Chaudhuri. 2004. Conditional Selectivity for Statistics on Query Expressions. In Proc. of the ACM SIGMOD Int. Conf. on Management of Data (SIGMOD 2004). ACM, Paris, France, 311-322.
[6] Nicolas Bruno, Surajit Chaudhuri, and Luis Gravano. 2001. STHoles: A Multidimensional Workload-aware Histogram. SIGMOD Record 30, 2 (2001), 211-222.
[7] Šejla Čebirić, François Goasdoué, and Ioana Manolescu. 2015. Query-oriented Summarization of RDF Graphs. P-VLDB 8, 12 (2015), 2012-2015.
[8] Kaushik Chakrabarti, Minos Garofalakis, Rajeev Rastogi, and Kyuseok Shim. 2001. Approximate query processing using wavelets. VLDB fournal 10, 2 (2001), 199-223.
[9] Julian Dolby, Achille Fokoue, Aditya Kalyanpur, Aaron Kershenbaum, Edith Schonberg, Kavitha Srinivas, and Li Ma. 2007. Scalable Semantic Retrieval through Summarization and Refinement. In Proc. of the 22nd AAAI Conf. on Artificial Intelligence (AAAI 2007). AAAI Press, Vancouver, BC, Canada, 299-304.
[10] David Freedman and Persi Diaconis. 1981. On the Histogram as a Density Estimator: $L_{2}$ theory. Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete 57, 4 (1981), 453-476.
[11] César A. Galindo-Legaria, Milind Joshi, Florian Waas, and Ming-Chuan Wu. 2003. Statistics on Views. In Proc. of the 29th Int. Conf. on Very Large Databases (VLDB 2003). Morgan Kaufmann, Berlin, Germany, 952-962.
[12] Hector Garcia-Molina, Jeffrey D. Ullman, and Jennifer Widom. 2000. Database System Implementation. Prentice-Hall, Upper Saddle River, NJ, USA.
[13] Minos Garofalakis and Phillip B. Gibbons. 2002. Wavelet Synopses with Error Guarantees. In Proc. of the ACM SIGMOD Int. Conf. on Management of Data (SIGMOD 2002). ACM, Madison, WI, USA, 476-487.
[14] Lise Getoor, Benjamin Taskar, and Daphne Koller. 2001. Selectivity Estimation Using Probabilistic Models. In Proc. of the ACM SIGMOD Int. Conf. on Management of Data (SIGMOD 2001). ACM, Santa Barbara, CA, USA, 461-472.
[15] Dimitrios Gunopulos, George Kollios, Vassilis J. Tsotras, and Carlotta Domeniconi. 2000. Approximating Multi-dimensional Aggregate Range Queries over Real Attributes. In Proc. of the ACM SIGMOD Int. Conf. on Management of Data (SIGMOD 2000). ACM, Dallas, TX, USA, 463-474.
[16] Dimitrios Gunopulos, George Kollios, Vassilis J. Tsotras, and Carlotta Domeniconi. 2005. Selectivity estimators for multidimensional range queries over real attributes. VLDB Journal 14, 2 (2005), 137-154.
[17] Y. Guo, Z. Pan, and J. Heflin. 2005. LUBM: A benchmark for OWL knowledge base systems. Journal of Web Semantics 3, 2-3 (2005), 158-182.
[18] Steve Harris and Andy Seaborne. 2013. SPARQL 1.1 Query Language, W3C Recommendation. (March 21 2013).
[19] Hai Huang and Chengfei Liu. 2011. Estimating Selectivity for Joined RDF Triple Patterns. In Proc. of the 20th ACM Conf. on Information and Knowledge Management (CIKM 2011). ACM, Glasgow, United Kingdom, 1435-1444.
[20] Yannis Ioannidis. 2003. The History of Histograms (Abridged). In Proc. of the 29th Int. Conf. on Very Large Databases (VLDB 2003). Morgan Kaufmann, Berlin, Germany, 19-30.

21] Yannis E. Ioannidis and Stavros Christodoulakis. 1991. On the Propagation of Errors in the Size of Join Results. In Proc. of the ACM SIGMOD Int. Conf. on Management of Data (SIGMOD 1991). ACM, Denver, CO, USA, 268-277.
[22] George Karypis and Vipin Kumar. 1998. A Fast and High Quality Multilevel Scheme for Partitioning Irregular Graphs. SIAM fournal on Scientific Computing 20, 1 (1998), 359-392.
[23] Graham Klyne, Jeremy J. Carroll, and Brian McBride. 2014. RDF 1.1: Concepts and Abstract Syntax.
https://www.w3.org/TR/rdf11-concepts/. (February 25 2014).
[24] Ju-Hong Lee, Deok-Hwan Kim, and Chin-Wan Chung. 1999. Multi-dimensional Selectivity Estimation Using Compressed Histogram Information. In Proc. of the ACM SIGMOD Int. Conf. on Management of Data (SIGMOD 1999). ACM, Philadelphia, PA, USA, 205-214.
[25] Kristen LeFevre and Evimaria Terzi. 2010. GraSS: Graph Structure Summarization. In Proc. of the SIAM Int. Conf. on Data Mining (SDM 2010). SIAM, Columbus, OH, USA, 454-465.
[26] Viktor Leis, Andrey Gubichev, Atanas Mirchev, Peter Boncz, Alfons Kemper, and Thomas Neumann. 2015. How Good Are Query Optimizers, Really? P-VLDB 9, 3 (2015), 204-215.
[27] Jure Leskovec, Anand Rajaraman, and Jeffrey D. Ullman. 2014. Mining of Massive Datasets, 2nd Ed. Cambridge University Press, New York, NY, USA.
[28] Yossi Matias, Jeffrey Scott Vitter, and Min Wang. 1998. Wavelet-based Histograms for Selectivity Estimation. In Proc. of the ACM SIGMOD Int. Conf. on Management of Data (SIGMOD 1998). ACM, Seattle, WA, USA, 448-459.
[29] Saket Navlakha, Rajeev Rastogi, and Nisheeth Shrivastava. 2008. Graph Summarization with Bounded Error. In Proc. of the ACM SIGMOD Int. Conf. on Management of Data (SIGMOD 2008). ACM, Vancouver, BC, Canada, 419-432.
[30] Thomas Neumann and Guido Moerkotte. 2011. Characteristic Sets: Accurate Cardinality Estimation for RDF Queries with Multiple Joins. In Proc. of the 27th IEEE Int. Conf. on Data Engineering (ICDE 2011). IEEE Computer Society, Hannover, Germany, 984-994.
[31] Thomas Neumann and Gerhard Weikum. 2010. The RDF-3X engine for scalable management of RDF data. VLDB fournal 19, 1 (2010), 91-113.
[32] J. Pérez, M. Arenas, and C. Gutierrez. 2009. Semantics and complexity of SPARQL. ACM Transactions on Database Systems 34, 3 (2009), 1-45.
[33] Viswanath Poosala, Peter J. Haas, Yannis E. Ioannidis, and Eugene J. Shekita. 1996. Improved Histograms for Selectivity Estimation of Range Predicates. In Proc. of the ACM SIGMOD Int. Conf. on Management of Data (SIGMOD 1996). ACM, Montreal, QC, Canada, 294-305.
[34] Viswanath Poosala and Yannis E. Ioannidis. 1997. Selectivity Estimation Without the Attribute Value Independence Assumption. In Proc. of the 23rd Int. Conf. on Very Large Databases (VLDB 1997). Morgan Kaufmann, Athens, Greece, 486-495.
[35] Matteo Riondato, David García-Soriano, and Francesco Bonchi. 2017. Graph summarization with quality guarantees. Data Mining and Knowledge Discovery 31, 2 (2017), 314-349.
[36] Joshua Spiegel and Neoklis Polyzotis. 2009. TuG Synopses for Approximate Query Answering. ACM Transactions on Database Systems 34, 1 (2009), 3:1-3:56.
[37] Markus Stocker, Andy Seaborne, Abraham Bernstein, Christoph Kiefer, and Dave Reynolds. 2008. SPARQL Basic Graph Pattern Optimization Using Selectivity Estimation. In Proc. of the 17th Int. Conf. on World Wide Web (WWW 2008). ACM, Beijing, China, 595-604.
[38] Pawel Terlecki, Hardik Bati, César A. Galindo-Legaria, and Peter Zabback. 2009. Filtered statistics. In Proc. of the ACM SIGMOD Int. Conf. on Management of Data (SIGMOD 2009). ACM, Providence, RI, USA, 897-904.
[39] Yuanyuan Tian, Richard A. Hankins, and Jignesh M. Patel. 2008. Efficient Aggregation for Graph Summarization. In Proc. of the ACM SIGMOD Int. Conf. on Management of Data (SIGMOD 2008). ACM, Vancouver, BC, Canada, 567-580.
[40] Kostas Tzoumas, Amol Deshpande, and Christian S. Jensen. 2013. Efficiently Adapting Graphical Models for Selectivity Estimation. VLDB fournal 22, 1 (2013), 3-27.
[41] Feng Yu, Wen-Chi Hou, Cheng Luo, Dunren Che, and Mengxia Zhu. 2013. CS2: A New Database Synopsis for Query Estimation. In Proc. of the ACM SIGMOD Int. Conf. on Management of Data (SIGMOD 2013). ACM, New York, NY, USA, 469-480.
[42] Lei Zou, M. Tamer Özsu, Lei Chen, Xuchuan Shen, Ruizhe Huang, and Dongyan Zhao. 2014. gStore: A Graph-based SPARQL Query Engine. VLDB Journal 23, 4 (2014), 565-590.

## A PROOF FOR SECTION 3

In this appendix, we prove Proposition 3.3 and Theorem 3.5.

## A. 1 Proof of Proposition 3.3

To prove Proposition 3.3, we start by defining the inverse of the summarisation function $\mu$ in a summary $S=\langle H, w, \mu\rangle$.
Definition A.1. For $S=\langle H, w, \mu\rangle$ a summary, the inverse function $\mu^{-}$of $\mu$ maps each bucket $b \in \operatorname{res}(H)$ and each triple $h \in H$ as follows:

$$
\begin{aligned}
& \mu^{-}(b)=\{d \in \operatorname{dom}(\mu) \mid \mu(d)=b\}, \\
& \mu^{-}(h)=\{f \in \operatorname{dom}(\mu) \times \operatorname{dom}(\mu) \times \operatorname{dom}(\mu) \mid \mu(f)=h\} .
\end{aligned}
$$

Please note that, by Definition 3.1 of size $s[\cdot]$, we have that $s[b]=\left|\mu^{-}(b)\right|$ and $s[h]=\left|\mu^{-}(h)\right|$, for each $b \in \operatorname{res}(H)$ and each $h \in H$. The next lemma shows that the consistency of a summary can be determined as described in Proposition 3.3; in addition, it shows how to count the number of graphs represented by a consistent summary.

Lemma A.2. For $S=\langle H, w, \mu\rangle$ a summary, $S$ is consistent if and only if $w[h] \leq s[h]$ for each $h \in H$. Furthermore, when $S$ is consistent, the following holds:

$$
\begin{equation*}
\left|R_{S}\right|=\prod_{h \in H}\binom{s[h]}{w[h]} . \tag{2}
\end{equation*}
$$

Proof. Let $S=\langle H, w, \mu\rangle$ be a summary. Then, by Definition 3.2 of a graph represented by $S$ and by Definition A. 1 of $\mu^{-}$, for each graph $G$, we have that $G \in R_{S}$ if and only if
(a) $G \subseteq \bigcup_{h \in H} \mu^{-}(h)$, and
(b) $w[h]=\left|\mu^{-}(h) \cap G\right|$ for each $h \in H$.

We next show that $S$ is consistent if and only if $w[h] \leq s[h]$ for each $h \in H$.
$(\Rightarrow)$ Assume that $S$ is consistent, and so $R_{S} \neq \emptyset$. Consider an arbitrary RDF graph $G \in R_{S}$ and an arbitrary triple $h \in H$. Then, $w[h]$ is the number of triples in $\mu^{-}(h)$ that occur in $G$. Since $s[h]=\left|\mu^{-}(h)\right|$, we have that $w[h] \leq s[h]$.
$(\Leftarrow)$ Assume that $w[h] \leq s[h]$ for each $h \in H$. Then, for each $h \in H$, let $G_{h}$ be an arbitrary, but fixed subset of $\mu^{-}(h)$ of size $w[h]$; since $w[h] \leq s[h]$, such $G_{h}$ exists. Finally, let $G=\bigcup_{h \in H} G_{h}$. By properties (a) and (b), we have that $G \in R_{S}$.

Therefore, $S$ is consistent if and only if $w[h] \leq s[h]$ for each $h \in H$. To conclude the proof, we assume that $S$ is consistent and show that Equation (2) holds. By properties (a) and (b), each graph $G \in R_{S}$ can be generated by independently choosing, for each triple $h \in H$, a set $G_{h}$ of atoms of size $w[h]$ such that $G_{h} \subseteq \mu^{-}(h)$. For each $h \in H$, there are precisely

$$
\binom{s[h]}{w[h]}
$$

distinct such sets $G_{h}$. So, we can count the total number of RDF graphs in $R_{S}$ as shown in Equation (2).

## A. 2 Proof of Theorem 3.5

Theorem 3.5. Let $q$ be a CQ and let $S$ be a summary such that $E_{q, S} \geq 1$. Then, for each $\varepsilon>1$ and fore the $q$-error of estimating the cardinality of $q$ as $E_{q, s}$, we have

$$
\mathrm{P}(e \geq \varepsilon) \leq\left(\frac{\varepsilon \cdot \sigma_{q, S}}{(\varepsilon-1) \cdot E_{q, S}}\right)^{2} .
$$

If $\varepsilon \geq E_{q, S}$ holds additionally, then the numerator $\varepsilon \cdot \sigma_{q, S}$ on the right-hand side of the inequality can be replaced with $\sigma_{q, S}$.
Proof. First, note that since $\varepsilon>1$,

$$
\frac{\varepsilon \cdot \sigma_{q, S}}{(\varepsilon-1) \cdot E_{q, S}}>\frac{\sigma_{q, S}}{(\varepsilon-1) \cdot E_{q, S}},
$$

so the additional inequality implies the main one. Therefore, it is enough to prove the first bound for the case when $E_{q, S}>\varepsilon>1$ and the second bound for the case when $\varepsilon \geq E_{q, S} \geq 1$. We do this in the rest of the proof.

Let $E_{q, S}>\varepsilon>1$. The inequality

$$
e=\max \left(\frac{\max \left(\left|\llbracket q \rrbracket_{G}\right|, 1\right)}{\max \left(E_{q, S}, 1\right)}, \frac{\max \left(E_{q, S}, 1\right)}{\max \left(\left|\llbracket q \rrbracket_{G}\right|, 1\right)}\right)<\varepsilon
$$

is equivalent to

$$
\frac{\max \left(E_{q, S}, 1\right)}{\varepsilon}<\max \left(\left|\llbracket q \rrbracket_{G}\right|, 1\right)<\varepsilon \cdot \max \left(E_{q, S}, 1\right)
$$

and, since $E_{q, S}>\varepsilon>1$, to

$$
\frac{E_{q, S}}{\varepsilon}<\left|\llbracket q \rrbracket_{G}\right|<\varepsilon \cdot E_{q, S} .
$$

Therefore, since $\varepsilon \cdot E_{q, S}-E_{q, S}>E_{q, S}-E_{q, S} / \varepsilon$,

$$
\begin{aligned}
& \mathrm{P}(e<\varepsilon)=\mathrm{P}\left(\frac{E_{q, S}}{\varepsilon}<\left|\llbracket q \rrbracket_{G}\right|<\varepsilon \cdot E_{q, S}\right)=\mathrm{P}\left(\frac{E_{q, S}}{\varepsilon}-E_{q, S}<\left|\llbracket q \rrbracket_{G}\right|-E_{q, S}<\varepsilon \cdot E_{q, S}-E_{q, S}\right)> \\
& \quad \mathrm{P}\left(\frac{E_{q, S}}{\varepsilon}-E_{q, S}<\left|\llbracket q \rrbracket_{G}\right|-E_{q, S}<E_{q, S}-\frac{E_{q, S}}{\varepsilon}\right)=\mathrm{P}\left(| | \llbracket q \rrbracket_{G}\left|-E_{q, S}\right|<E_{q, S}-\frac{E_{q, S}}{\varepsilon}\right)=\mathrm{P}\left(| | \llbracket q \rrbracket_{G}\left|-E_{q, S}\right|<\frac{(\varepsilon-1) \cdot E_{q, S}}{\varepsilon}\right)> \\
& 1-\left(\frac{\varepsilon \cdot \sigma_{q, S}}{(\varepsilon-1) \cdot E_{q, S}}\right)^{2} .
\end{aligned}
$$

The last step follows from Chebyshev's inequality $\mathrm{P}\left(\left|\left|\llbracket q \rrbracket_{G}\right|-E_{q, S}\right| \geq k \cdot \sigma_{q, S}\right) \leq \frac{1}{k^{2}}$ for $k=\frac{(\varepsilon-1) \cdot E_{q, S}}{\varepsilon \cdot \sigma_{q, S}}$. This immediately implies the required bound

$$
\mathrm{P}(e \geq \varepsilon) \leq\left(\frac{\varepsilon \cdot \sigma_{q, S}}{(\varepsilon-1) \cdot E_{q, S}}\right)^{2}
$$

Let now $\varepsilon \geq E_{q, S} \geq 1$. The inequality $e<\varepsilon$ is again equivalent to

$$
\frac{\max \left(E_{q, S}, 1\right)}{\varepsilon}<\max \left(\left|\llbracket q \rrbracket_{G}\right|, 1\right)<\varepsilon \cdot \max \left(E_{q, S}, 1\right)
$$

and, since $\varepsilon \geq E_{q, S} \geq 1$, to $\left|\llbracket q \rrbracket_{G}\right|<\varepsilon \cdot E_{q, s}$. Therefore,

$$
\begin{aligned}
& \mathrm{P}(e<\varepsilon)=\mathrm{P}\left(\left|\llbracket q \rrbracket_{G}\right|<\varepsilon \cdot E_{q, S}\right)=\mathrm{P}\left(\left|\llbracket q \rrbracket_{G}\right|-E_{q, S}<\varepsilon \cdot E_{q, S}-E_{q, S}\right) \geq \\
& \mathrm{P}\left(E_{q, S}-\varepsilon \cdot E_{q, S}<\left|\llbracket q \rrbracket_{G}\right|-E_{q, S}<\varepsilon \cdot E_{q, S}-E_{q, S}\right)=\mathrm{P}\left(| | \llbracket q \rrbracket_{G}\left|-E_{q, S}\right|<\varepsilon \cdot E_{q, S}-E_{q, S}\right)=\mathrm{P}\left(| | \llbracket q \rrbracket_{G}\left|-E_{q, S}\right|<(\varepsilon-1) \cdot E_{q, S}\right)> \\
& 1-\left(\frac{\sigma_{q, S}}{(\varepsilon-1) \cdot E_{q, S}}\right)^{2} .
\end{aligned}
$$

The last step follows from Chebyshev's inequality for $k=\frac{(\varepsilon-1) \cdot E_{q, S}}{\sigma_{q, S}}$. This immediately implies the required bound

$$
\mathrm{P}(e \geq \varepsilon) \leq\left(\frac{\sigma_{q, S}}{(\varepsilon-1) \cdot E_{q, S}}\right)^{2}
$$

This concludes the proof of the theorem.

## B PROOFS FOR SECTION 4

In this appendix, we prove Theorems 4.3 and 4.4, and Proposition 4.6. For the rest of this section, we fix an arbitrary consistent summary $S=\langle H, w, \mu\rangle$ and a CQ $q$ such that $\operatorname{res}(q) \subseteq \operatorname{dom}(\mu)$.

## B. 1 Properties of the Partition Base

We start by formalising the properties of the unifiable partitions of $q$ that were intuitively outlined in Section 4.1. To this end, we first define the notion of an expansion w.r.t. a partition $P$ of an answer $\tau$ to $\mu(q)$ over $H$.

Definition B.1. Let $\tau$ be an answer to $\mu(q)$ on $H$ and let $P \in B_{\tau}$ be a partition. A substitution $\pi$ with $\operatorname{dom}(\pi)=\operatorname{var}(q)$ and $\operatorname{rng}(\pi) \subseteq \operatorname{dom}(\mu)$ is a $\tau$-expansion to $P$ if, for each $x \in \operatorname{var}(q)$,
(1) $\pi(x) \in \mu^{-}(\tau(x))$,
(2) $\gamma_{P}(x) \in \operatorname{var}(q)$ implies that $\pi(x)=\pi\left(\gamma_{P}(x)\right)$,
(3) $\gamma_{P}(x) \in \operatorname{res}(q)$ implies that $\pi(x)=\gamma_{P}(x)$.

Then, $\exp _{\tau}(P)$ contains each $\tau$-expansion to $P$, and $\exp _{\tau}^{\prec}(P)$ contains each $\pi \in \exp _{\tau}(P)$ for which no $P^{\prime} \in B_{\tau}$ exists such that $P<P^{\prime}$ and $\pi \in \exp _{\tau}\left(P^{\prime}\right)$.

Proposition B.2. For each answer $\tau$ to $\mu(q)$ on $H$ and each partition $P \in B_{\tau}$, the following properties hold:
(P1) for each $\pi \in \exp _{\tau}(P)$, exactly one partition $P^{\prime} \in B_{\tau}$ exists such that $P \leq P^{\prime}$ and $\pi \in \exp _{\tau}^{<}\left(P^{\prime}\right)$,
(P2) for each $\pi \in \exp _{\tau}^{<}(P)$ and all atoms $a, a^{\prime} \in q$, there exists $u \in P$ with $a, a^{\prime} \in u$ if and only if $\pi(a)=\pi\left(a^{\prime}\right)$.

Proof. For the rest of this proof, we let $\tau$ be an arbitrary answer to $\mu(q)$ on $H$.
Property $(P 1)$ Since $B_{\tau}$ is finite and $\leq$ is a partial order on $B_{\tau}$, for each $P \in B_{\tau}$ and each $\pi \in \exp _{\tau}(P)$, there exists a partition $P^{\prime} \in B_{\tau}$ such that $P \leq P^{\prime}$ and $\pi \in \exp _{\tau}^{\prec}\left(P^{\prime}\right)$. Hence, we next show that $\exp _{\tau}^{\prec}\left(P_{1}\right) \cap \exp _{\tau}^{\prec}\left(P_{2}\right)=\emptyset$ for all distinct partitions $P_{1}, P_{2} \in B_{\tau}$.

Consider two arbitrary distinct partitions $P_{1}, P_{2} \in B_{\tau}$. By the definition of $\exp _{\tau}^{\prec}\left(P_{1}\right)$ and $\exp _{\tau}^{<}\left(P_{2}\right)$, if $P_{1}<P_{2}$ or $P_{2}<P_{1}$ holds, we then have $\exp _{\tau}^{\prec}\left(P_{1}\right) \cap \exp _{\tau}^{\prec}\left(P_{2}\right)=\emptyset$. Thus, in the rest of this proof, we consider the case in which $P_{1} \nprec P_{2}$ and $P_{2} \nprec P_{1}$.

Let $\approx$ be the smallest equivalence relation on (the atoms of) $q$ such that $a \approx a^{\prime}$ for all atoms $a, a^{\prime} \in q$ such that $a, a^{\prime} \in u$ for $u \in P_{i}$ with $i \in[1,2]$. Then, let $P$ be the set of all the equivalence classes of $\approx$. Since $P_{1}$ and $P_{2}$ are partitions of $q$ and $\approx$ is an equivalence relation on the atoms of $q$, we have that $P$ is a partition of $q$ as well. By the definition of $\approx$, for each $i \in[1,2]$ and each $u \in P_{i}$, there exists an equivalence class of $\approx$ that contains each atom in $u$; thus, we have that $P_{i} \leq P$. By the initial assumptions, we have that $P_{1} \neq P_{2}, P_{1} \nless P_{2}$, and $P_{2} \nless P_{1}$; therefore, we also have that $P_{1} \prec P$ and $P_{2} \prec P$. We next show that $P \in B_{\tau}$ and, for each $\pi \in \exp _{\tau}\left(P_{1}\right) \cap \exp _{\tau}\left(P_{2}\right)$, we have that $\pi \in \exp _{\tau}(P)$. Please note that this suffices to show that $\exp _{\tau}^{<}\left(P_{1}\right) \cap \exp _{\tau}^{<}\left(P_{2}\right)=\emptyset$ since $P_{1}<P$ and $P_{2}<P$.

Consider an arbitrary substitution $\pi$ with $\operatorname{dom}(\pi)=\operatorname{var}(q)$ and $\operatorname{rng}(\pi) \subseteq \operatorname{dom}(\mu)$ such that $\pi \in \exp _{\tau}\left(P_{1}\right) \cap \exp _{\tau}\left(P_{2}\right)$. Since $\pi$ is a $\tau$ expansion to $P_{1}$ and $P_{2}$, for each $x \in \operatorname{var}(q)$, we have that $\tau(x)=\mu(\pi(x))$. Next, let $\mathcal{G}_{P}$ be the term graph for $P$ defined in Definition 4.1. To prove that $\pi$ is a $\tau$-expansion to $P$, we first show that, for all terms $s, t \in \operatorname{term}(q)$ such that $s$ reaches $t$ in $\mathcal{G}_{P}$, we have that $\pi(s)=\pi(t)$. To this end, we show that the following property holds for all atoms $a, a^{\prime} \in q$ :
(A1) $a \approx a^{\prime}$ implies that $\pi(a)=\pi\left(a^{\prime}\right)$.
To prove auxiliary property (A1), we proceed by induction on the number of steps required to derive $a \approx a^{\prime}$. For the base case, the empty relation $\approx_{0}$ clearly satisfies the property. For the inductive step, consider an arbitrary relation $\approx_{n}$ obtained in $n$ steps that satisfies the property; we show that the same holds for all the equivalences derivable from $\approx_{n}$. Since the equality relation $=$ is reflexive, symmetric, and transitive, the derivation of $a \approx_{n+1} a^{\prime}$ due to reflexivity, symmetry, or transitivity clearly preserves the required property. So, we next consider an arbitrary $i \in[1,2]$, a set $u \in P_{i}$, and two atoms $a, a^{\prime} \in u$, so that we derive $a \approx_{n+1} a^{\prime}$. Let $a=\langle s, p, o\rangle$ and let $a^{\prime}=\left\langle s^{\prime}, p^{\prime}, o^{\prime}\right\rangle$. Then, by the definition of $\mathcal{G}_{P_{i}}$, terms $s$, $p$, and $o$ reach $s^{\prime}, p^{\prime}$, and $o^{\prime}$ in $\mathcal{G}_{P_{i}}$, respectively. Thus, by the definition of $\gamma_{P_{i}}$, we also have that $\gamma_{P_{i}}(s)=\gamma_{P_{i}}\left(s^{\prime}\right), \gamma_{P_{i}}(p)=\gamma_{P_{i}}\left(p^{\prime}\right)$, and $\gamma_{P_{i}}(o)=\gamma_{P_{i}}\left(o^{\prime}\right)$. By properties (2) and (3) in Definition B.1, we have that $\pi(a)=\pi\left(a^{\prime}\right)$, as required.

Please note that $G_{P}$ is obtained by adding an undirected edge between $s_{1}$ and $s_{2}, p_{1}$ and $p_{2}$, and $o_{1}$ and $o_{2}$ for each $u \in P$ and all atoms $\left\langle s_{1}, p_{1}, o_{1}\right\rangle,\left\langle s_{2}, p_{2}, o_{2}\right\rangle \in u$. By the definition of $P$, we have that $\left\langle s_{1}, p_{1}, o_{1}\right\rangle \approx\left\langle s_{2}, p_{2}, o_{2}\right\rangle$ for each such pair of atoms. So, because the equality relation $=$ is reflexive, symmetric, and transitive, property (A1) implies that $\pi(s)=\pi(t)$ for all terms that are reachable in $\mathcal{G}_{P}$.

Since $\pi(s)=\pi(t)$ for all $s, t \in \operatorname{term}(q)$ such that $s$ reaches $t$ in $\mathcal{G}_{P}$, partition $P$ is unifiable, and so $P \in B$. We next show that $P$ is satisfied by $\tau$ and that $\pi$ satisfies conditions (2) and (3) in Definition B.1. For $x \in \operatorname{var}(q)$ a variable, we next consider two cases.

- Assume that $\gamma_{P}(x)$ is a variable $y \in \operatorname{var}(q)$. By the definition of $\gamma_{P}$, we have that $x$ reaches $y$ in $\mathcal{G}_{P}$, and so $\pi(x)=\pi(y)$. Since $\tau(x)=\mu(\pi(x))$ and $\tau(y)=\mu(\pi(y))$, we conclude that $\tau(x)=\tau(y)$.
- Assume that $\gamma_{P}(x)$ is a resource $d \in \operatorname{res}(q)$. By the definition of $\gamma_{P}$, we have that $x$ reaches $d$ in $G_{P}$, and so $\pi(x)=d$. Since $\tau(x)=\mu(\pi(x))$, we conclude that $\tau(x)=\mu(d)$.
Therefore, $P$ is satisfied by $\tau$ and $\pi$ satisfies conditions (2) and (3) in Definition B.1; thus, $\pi$ is a $\tau$-expansion to $P$, as required.
Property (P2) Let $P \in B_{\tau}$ be a partition and let $\pi \in \exp _{\tau}^{<}(P)$ be a $\tau$-expansion; we show that, for all $a, a^{\prime} \in q$, there exists $u \in P$ with $a, a^{\prime} \in u$ if and only if $\pi(a)=\pi\left(a^{\prime}\right)$.
$(\Rightarrow)$ By properties (2) and (3) in Definition B. 1 and by Definition 4.1 of $\gamma_{P}$, we have that $\pi(a)=\pi\left(a^{\prime}\right)$ for each $u \in P$ and all $a, a^{\prime} \in u$, as required.
$(\Leftarrow)$ For the sake of a contradiction, assume that two atoms $a, a^{\prime} \in q$ exist such that $\pi(a)=\pi\left(a^{\prime}\right)$, but no $u \in P$ exists such that $a, a^{\prime} \in u$. Then, let $P^{\prime}$ be the partition of $q$ such that, for all $a_{1}, a_{2} \in q$, there exists $u^{\prime} \in P^{\prime}$ such that $a_{1}, a_{2} \in u^{\prime}$ if and only if $\pi\left(a_{1}\right)=\pi\left(a_{2}\right)$. Please note that, by the definition of $P^{\prime}$, a set $u^{\prime} \in P^{\prime}$ exists such that $a, a^{\prime} \in u^{\prime}$.

We first show that $P<P^{\prime}$. To this end, consider an arbitrary $u \in P$ and two arbitrary atoms $a_{1}, a_{2} \in u$. Because $\pi$ is a $\tau$-expansion to $P$, we have that $\pi(a)=\pi\left(a^{\prime}\right)$, and so a set $u^{\prime} \in P^{\prime}$ exists such that $a_{1}, a_{2} \in u^{\prime}$. Thus, we have that $P \leq P^{\prime}$. Furthermore, by the initial assumption, there exists no $u \in P$ such that $a, a^{\prime} \in u$; hence, we also have that $P<P^{\prime}$.

By the construction of $P^{\prime}$, we have that $\pi(s)=\pi(t)$ for all terms $s, t \in \operatorname{term}(q)$ such that $s$ reaches $t$ in $\mathcal{G}_{P^{\prime}}$. So, $P^{\prime}$ is a unifiable partition of $q$ and satisfies properties (2) and (3) of Definition B.1. Since $\pi$ is a $\tau$-expansion to $P$, we have that $\tau(x)=\mu(\pi(x))$ for each $x \in \operatorname{var}(q)$; thus, $P^{\prime}$ is satisfied by $\tau$ and $\pi$ is a $\tau$-expansion to $P^{\prime}$. This is a contradiction, since $P<P^{\prime}$ and $\pi \in \exp _{\tau}^{<}(P)$.

## B. 2 Counting Expansions

We next show how one can effectively count, given an answer $\tau$ to $\mu(q)$ over $H$ and a partition $P \in B_{\tau}$, the number of $\tau$-expansions to $P$. To this end, we first an auxiliary lemma.

Lemma B.3. For all partitions $P, P^{\prime} \in B$ with $P<P^{\prime}$,

$$
\begin{equation*}
-K\left(P, P^{\prime}\right)=\sum_{P<P^{\prime \prime} \leq P^{\prime}} K\left(P^{\prime \prime}, P^{\prime}\right) \tag{3}
\end{equation*}
$$

Proof. Let $P$ and $P^{\prime}$ be as specified in the lemma. By Definition 4.1, we have that

$$
-K\left(P, P^{\prime}\right)=-\left(C^{e}\left(P, P^{\prime}\right)-C^{o}\left(P, P^{\prime}\right)\right)=C^{o}\left(P, P^{\prime}\right)-C^{e}\left(P, P^{\prime}\right) .
$$

Because $P$ and $P^{\prime}$ are distinct partitions, by Definition 4.1 of a chain from $P$ to $P^{\prime}$, the number of even chains from $P$ to $P^{\prime}$ in $B$ can be computed by summing, for each $P^{\prime \prime} \in B$ with $P<P^{\prime \prime}$ and $P^{\prime \prime} \leq P^{\prime}$, the number of odd chains from $P^{\prime \prime}$ to $P^{\prime}$ in $B$. Similarly, the number of odd chains from $P$ to $P^{\prime}$ in $B$ can be computed by summing, for each $P^{\prime \prime} \in B$ with $P<P^{\prime \prime}$ and $P^{\prime \prime} \leq P^{\prime}$, the number of even chains from $P^{\prime \prime}$ to $P^{\prime}$ in $B$. Hence, we have that

$$
\begin{aligned}
& C^{e}\left(P, P^{\prime}\right)=\sum_{P<P^{\prime \prime} \leq P^{\prime}} C^{o}\left(P^{\prime \prime}, P^{\prime}\right), \text { and } \\
& C^{o}\left(P, P^{\prime}\right)=\sum_{P<P^{\prime \prime} \leq P^{\prime}} C^{e}\left(P^{\prime \prime}, P^{\prime}\right) .
\end{aligned}
$$

Therefore, we have that

$$
\begin{aligned}
-K\left(P, P^{\prime}\right) & =C^{o}\left(P, P^{\prime}\right)-C^{e}\left(P, P^{\prime}\right)=\left(\sum_{P<P^{\prime \prime} \leq P^{\prime}} C^{e}\left(P^{\prime \prime}, P^{\prime}\right)\right)-\left(\sum_{P<P^{\prime \prime} \leq P^{\prime}} C^{o}\left(P^{\prime \prime}, P^{\prime}\right)\right) \\
& =\sum_{P<P^{\prime \prime} \leq P^{\prime}}\left(C^{e}\left(P^{\prime \prime}, P^{\prime}\right)-C^{o}\left(P^{\prime \prime}, P^{\prime}\right)\right)=\sum_{P<P^{\prime \prime} \leq P^{\prime}} K\left(P^{\prime \prime}, P^{\prime}\right) .
\end{aligned}
$$

Proposition B.4. The following equations hold for each answer $\tau$ to $\mu(q)$ on $\operatorname{res}(H)$ and each $P \in B_{\tau}$ :

$$
\begin{align*}
\exp _{\tau}(P) \mid & =N_{\tau}(P),  \tag{4}\\
\left|\exp _{\tau}^{<}(P)\right| & =\sum_{P^{\prime} \in B_{\tau}, P \leq P^{\prime}} K\left(P, P^{\prime}\right) \cdot N_{\tau}\left(P^{\prime}\right) . \tag{5}
\end{align*}
$$

Proof. Let $\tau$ be as specified in the lemma, and let $\leq_{\tau}$ be the restriction of $\leq$ to $B_{\tau}$.
We first prove Equation (4). To this end, consider a partition $P \in B_{\tau}$. By Definition B. 1 of a $\tau$-expansion to $P$, each $\pi \in \exp _{\tau}(P)$ can vary only on the variables occurring in the range of $\gamma_{P}$. Furthermore, for each $x \in \operatorname{varrng}\left(\gamma_{P}\right)$, we have that $\pi(x) \in \mu^{-}(\tau(x))$ and $\left|\mu^{-}(\tau(x))\right|=s[\tau(x)]$; thus, there are $N_{\tau}(P)=\prod_{x \in \operatorname{varrng}\left(\gamma_{P}\right)} s[\tau(x)]$ distinct $\tau$-expansions to $P$.

We next prove Equation (5) by induction on $\leq_{\tau}$.
Base case. Consider an arbitrary partition $P \in B_{\tau}$ that is $\leq_{\tau}$-maximal. We then have that $\exp _{\tau}{ }_{\tau}(P)=\exp _{\tau}(P)$; furthermore, by Equation (4), we have that $\left|\exp _{\tau}(P)\right|=N_{\tau}(P)$. In addition, we have that $\sum_{P \leq_{\tau} P^{\prime}} K\left(P, P^{\prime}\right) \cdot N_{\tau}\left(P^{\prime}\right)=K(P, P) \cdot N_{\tau}(P)$. Finally, by Definition 4.1 of $K$, we have that $K(P, P)=1$; so, Equation (5) holds.

Inductive Step. Consider an arbitrary $P \in B_{\tau}$ and assume that, for each $P^{\prime \prime} \in B_{\tau}$ with $P<_{\tau} P^{\prime \prime}$, we have that

$$
\left|\exp _{\tau}^{<}\left(P^{\prime \prime}\right)\right|=\sum_{P^{\prime \prime} \leq \leq_{\tau} P^{\prime}} K\left(P^{\prime \prime}, P^{\prime}\right) \cdot N_{\tau}\left(P^{\prime}\right) ;
$$

we show that the same holds for $P$. By the definition of $\exp _{\tau}^{<}(P)$, we have that $\exp _{\tau}{ }_{\tau}(P)=\exp _{\tau}(P) \backslash \bigcup_{P{ }_{\tau}} P^{\prime \prime} \exp _{\tau}\left(P^{\prime \prime}\right)$. By Proposition B.2, for each $\pi \in \bigcup_{P<_{\tau} P^{\prime \prime}} \exp _{\tau}\left(P^{\prime \prime}\right)$, exactly one partition $P^{\prime \prime} \in B_{\tau}$ exists such that $P<_{\tau} P^{\prime \prime}$ and $\pi \in \exp _{\tau}^{\prec}\left(P^{\prime \prime}\right)$. Therefore, we have that

$$
\begin{aligned}
\exp _{\tau}^{<}(P) & =\exp _{\tau}(P) \backslash \bigcup_{P<_{\tau} P^{\prime \prime}} \exp _{\tau}^{<}\left(P^{\prime \prime}\right), \text { and } \\
\left|\exp _{\tau}^{<}(P)\right| & =\left|\exp _{\tau}(P)\right|-\sum_{P<_{\tau} P^{\prime \prime}}\left|\exp _{\tau}^{<}\left(P^{\prime \prime}\right)\right|
\end{aligned}
$$

By Equation (4), we have that $\left|\exp _{\tau}(P)\right|=N_{\tau}(P)$. So, by the inductive hypothesis, we can compute $\left|\exp _{\tau}{ }_{\tau}(P)\right|$ as follows.

$$
\left|\exp _{\tau}^{<}(P)\right|=N_{\tau}(P)-\sum_{P<_{\tau} P^{\prime \prime}}\left(\sum_{P^{\prime \prime} \leq \tau} K P^{\prime}\left(P^{\prime \prime}, P^{\prime}\right) \cdot N_{\tau}\left(P^{\prime}\right)\right)=N_{\tau}(P)-\sum_{P<_{\tau} P^{\prime}} N_{\tau}\left(P^{\prime}\right) \cdot\left(\sum_{P<_{\tau} P^{\prime \prime} \leq_{\tau} P^{\prime}} K\left(P^{\prime \prime}, P^{\prime}\right)\right) .
$$

By Lemma B.3, for each $P^{\prime} \in B_{\tau}$ with $P<_{\tau} P^{\prime}$, we have that $\sum_{P<_{\tau} P^{\prime \prime} \leq_{\tau} P^{\prime}} K\left(P^{\prime \prime}, P^{\prime}\right)=-K\left(P, P^{\prime}\right)$; thus, we have that

$$
\left|\exp _{\tau}^{<}(P)\right|=N_{\tau}(P)-\sum_{P<_{\tau} P^{\prime}} N_{\tau}\left(P^{\prime}\right) \cdot\left(\sum_{P<_{\tau} P^{\prime \prime} \leq_{\tau} P^{\prime}} K\left(P^{\prime \prime}, P^{\prime}\right)\right)
$$

$$
=N_{\tau}(P)-\left(\sum_{P<_{\tau} P^{\prime}}-N_{\tau}\left(P^{\prime}\right) \cdot K\left(P, P^{\prime}\right)\right)=N_{\tau}(P)+\left(\sum_{P<_{\tau} P^{\prime}} N_{\tau}\left(P^{\prime}\right) \cdot K\left(P, P^{\prime}\right)\right) .
$$

Finally, by Definition 4.1 of $K$, we have that $K(P, P)$ is 1 , so the following holds:

$$
\begin{aligned}
\left|\exp _{\tau}^{\prec}(P)\right| & =N_{\tau}(P)+\left(\sum_{P<_{\tau} P^{\prime}} N_{\tau}\left(P^{\prime}\right) \cdot K\left(P, P^{\prime}\right)\right) \\
& =K(P, P) \cdot N_{\tau}(P)+\left(\sum_{P<_{\tau} P^{\prime}} K\left(P, P^{\prime}\right) \cdot N_{\tau}\left(P^{\prime}\right)\right)=\sum_{P \leq_{\tau} P^{\prime}} K\left(P, P^{\prime}\right) \cdot N_{\tau}\left(P^{\prime}\right) .
\end{aligned}
$$

## B. 3 Proofs of Theorem 4.3 and Proposition 4.6

We are now ready to prove the correctness of the formulae for computing $E_{q, S}$ given in Theorem 4.3 and Proposition 4.6.
Lemma B.5. For all $n, k, m \in \mathbb{N}$ with $n \geq m \geq k>0$,

$$
\frac{\binom{n-k}{m-k}}{\binom{n}{m}}=\frac{(m)_{k}}{(n)_{k}}
$$

Proof. Let $n, k, m$ be natural numbers as specified in the lemma. By the definition of binomial coefficient, we have that

$$
\begin{aligned}
\binom{n}{m} & =\frac{n!}{m!\cdot(n-m)!}, \text { and } \\
\binom{n-k}{m-k} & =\frac{(n-k)!}{(m-k)!\cdot(n-k-(m-k))!}=\frac{(n-k)!}{(m-k)!\cdot(n-m)!} .
\end{aligned}
$$

By using the standard laws of arithmetics, we then have that

$$
\frac{\binom{n-k}{m-k}}{\binom{n}{m}}=\frac{\frac{(n-k)!}{(m-k)!\cdot(n-m)!}}{\frac{n!}{m!\cdot(n-m)!}}=\frac{(n-k)!\cdot m!}{n!\cdot(m-k)!}=\frac{m \cdots(m-k+1)}{n \cdots(n-k+1)}=\frac{(m)_{k}}{(n)_{k}} .
$$

Lemma B.6. Let $\tau$ be an answer to $\mu(q)$ over $H$ and let $P \in B_{\tau}$ be a partition. The following holds for each $\pi \in \exp _{\tau}^{<}(P)$ :

$$
\begin{equation*}
F_{\tau}(P)=E_{\pi(q), S} \tag{6}
\end{equation*}
$$

Proof. Let $\tau$ and $P$ be as specified in the lemma. Consider an arbitrary substitution $\pi \in \exp _{\tau}^{<}(P)$ and let $q_{\pi}=\pi(q)$; we show that $F_{\tau}(P)=E_{q_{\pi}, s}$. Note that, for each $a \in q_{\pi}$, we have that $\operatorname{var}(a)=\emptyset$ and $\operatorname{res}(a) \subseteq \operatorname{dom}(\mu)$.

By the definition of $\tau$-expansion, we have that $\tau(x)=\mu(\pi(x))$ for each $x \in \operatorname{var}(q)$, and so $\mu\left(q_{\pi}\right)=\tau(\mu(q))$. Because $P$ is satisfied by $\tau$, for each $u \in P$ and all atoms $a, a^{\prime} \in u$, we have that $\tau(\mu(a))=\tau\left(\mu\left(a^{\prime}\right)\right)$; so, $\tau(\mu(u))$ is a set consisting of a single triple from $H$. In the following, for each triple $h \in \tau(\mu(q))$, let $\#_{h} \in \mathbb{N}$ be as specified in Definition 4.2.

By property (P2) of Proposition B.2, for all atoms $a, a^{\prime} \in q$, we have that $\pi(a)=\pi\left(a^{\prime}\right)$ if and only if there exists $u \in P$ with $a, a^{\prime} \in u$. That is, there is a one-to-one correspondence between the triples in $q_{\pi}$ and the sets in $P$. So, for each triple $h \in \tau(\mu(q)$ ), we have that

$$
\begin{equation*}
\#_{h}=|\{u \in P \mid \tau(\mu(u))=\{h\}\}|=\left|\left\{f \in q_{\pi} \mid \mu(f)=h\right\}\right| . \tag{7}
\end{equation*}
$$

By Definition 3.4 of expectation, we have that

$$
\begin{equation*}
E_{q_{\pi}, S}=\sum_{G \in R_{S}} \frac{\left|\llbracket q_{\pi} \rrbracket_{G}\right|}{\left|R_{S}\right|} \tag{8}
\end{equation*}
$$

By the initial assumption, $S$ is consistent; so, $\left|R_{S}\right|>0$ and $E_{q_{\pi}, S}$ is well-defined. Next, let $\mathcal{N}=\sum_{G \in R_{S}}\left|\llbracket q_{\pi} \rrbracket_{G}\right|$. Since $q_{\pi}$ is a variable-free, $\left|\llbracket q_{\pi} \rrbracket_{G}\right|$ is either 0 or 1 ; thus, $\mathcal{N}$ is the number of graphs $G \in R_{S}$ with $q_{\pi} \subseteq G$. By Definition 3.2 of a graph represented by $S$, for each graph $G$, we have that $G \in R_{S}$ and $q_{\pi} \subseteq G$ if and only if
(A) $G \subseteq \bigcup_{h \in H} \mu^{-}(h)$,
(B) for each $h \in H$, the set $\mu^{-}(h) \cap G$ contains each $f \in q_{\pi}$ with $\mu(f)=h$, and
(C) $w[h]=\left|\mu^{-}(h) \cap G\right|$ for each $h \in H$.

To prove the lemma, we next consider two alternative cases.
(Case 1) Assume that a triple $h \in \tau(\mu(q))$ exists such that $\left|\#_{h}\right|>w[h]$; so $F_{\tau}(P)=0$. By Equation (7), \# ${ }_{h}$ is the number of triples in $q_{\pi}$ that $\mu$ maps onto $h$. So, there is no graph $G$ that satisfies both properties (B) and (C) above, and thus $\mathcal{N}=0$ and $E_{q_{\pi}, S}=0$. Therefore, $F_{\tau}(P)=E_{q_{\pi}, S}$ and the lemma holds.
(Case 2) Assume that $\#_{h} \leq w[h]$ for each $h \in \tau(\mu(q))$; so

$$
F_{\tau}(P)=\prod_{h \in \tau(\mu(q))} \frac{(w[h])_{\#_{h}}}{(s[h])_{\#_{h}}}
$$

Each graph $G \in R_{S}$ with $q_{\pi} \subseteq G$ can be generated by independently choosing, for each $h \in H$, a subset $G_{h}$ of $\mu^{-}(h)$ such that $\left|G_{h}\right|=w[h]$ and $G_{h}$ contains each $f \in q_{\pi}$ with $\mu(f)=h$. Please recall that $s[h]$ is the size of $\mu^{-}(h)$. Next, we show how to compute, for each $h \in H$, the number $\mathcal{N}_{h}$ of distinct such sets by considering two alternative cases.

- Assume that $h \in \tau(\mu(q))$. By Equation (7), there are precisely $\#_{h}$ distinct triples $f \in q_{\pi}$ such that $\mu(f)=h$, and so

$$
\mathcal{N}_{h}=\binom{s[h]-\#_{h}}{w[h]-\#_{h}}
$$

- Assume that $h \notin \tau(\mu(q))$. Hence, no triple $f \in q_{\pi}$ exists such that $\mu(f)=h$. So, we have that

$$
\mathcal{N}_{h}=\binom{s[h]}{w[h]}
$$

Therefore, we can compute $\mathcal{N}$ as follows:

$$
\mathcal{N}=\sum_{G \in R_{S}}\left|\llbracket q_{\pi} \rrbracket_{G}\right|=\prod_{h \in H} \mathcal{N}_{h}=\prod_{h \in \tau(\mu(q))}\binom{s[h]-\#_{h}}{w[h]-\#_{h}} \cdot \prod_{h \in(H \backslash \tau(\mu(q)))}\binom{s[h]}{w[h]} .
$$

By the definition of expected answer, we then have that

$$
E_{q_{\pi}, S}=\sum_{G \in R_{S}} \frac{\left|\llbracket q_{\pi} \rrbracket_{G}\right|}{\left|R_{S}\right|}=\frac{\mathcal{N}}{\left|R_{S}\right|}=\frac{\prod_{h \in \tau(\mu(q))}\binom{s[h]-\#_{h}}{w[h]-\#_{h}} \cdot \prod_{h \in(H \backslash \tau(\mu(q)))}\binom{s[h]}{w[h]}}{\left|R_{S}\right|}
$$

By using the formula for computing $\left|R_{S}\right|$ from Lemma A.2, we obtain the following:

$$
\begin{aligned}
& =\frac{\prod_{h \in \tau(\mu(q))}\binom{s[h]-\#_{h}}{w[h]-\#_{h}} \cdot \prod_{h \in(H \backslash \tau(\mu(q)))}\binom{s[h]}{w[h]}}{\prod_{h \in \tau(\mu(q))}\binom{s[h]}{w[h]} \cdot \prod_{h \in(H \backslash \tau(\mu(q)))\binom{s[h]}{w[h]}}^{\prod_{h \in \tau(\mu(q))}} \frac{\binom{s[h]-\#_{h}}{w[h]-\#_{h}}}{\binom{s[h]}{w[h]}} .}
\end{aligned}
$$

Finally, by Lemma B.5, we have that

$$
E_{q_{\pi}, S}=\prod_{h \in \tau(\mu(q))} \frac{\binom{s[h]-\#_{h}}{w[h]-\#_{h}}}{\binom{s[h]}{w[h]}}=\prod_{h \in \tau(\mu(q))} \frac{(w[h])_{\#_{h}}}{(s[h])_{\#_{h}}}=F_{\tau}(P)
$$

We are now ready to prove Theorem 4.3.
Theorem 4.3. The following identity holds:

$$
E_{q, S}=\sum_{\tau \in \llbracket \mu(q) \rrbracket_{H}} \sum_{P \in B_{\tau}} F_{\tau}(P) \cdot \sum_{P^{\prime} \in B_{\tau}, P \leq P^{\prime}} K\left(P, P^{\prime}\right) \cdot N_{\tau}\left(P^{\prime}\right)
$$

Proof. In the rest of this proof, for each substitution $\pi$ with $\operatorname{dom}(\pi)=\operatorname{var}(q)$ and $\operatorname{rng}(\pi) \subseteq \operatorname{dom}(\mu)$, we let $q_{\pi}=\pi(q)$; furthermore, let $P_{0}=\{\{a\} \mid a \in q\}$ be a partition of $q$. By Definition 4.1 of partition base, we have that $P_{0} \in B$; in addition, for each answer $\tau \mu(q)$ on $H$, we have that $P_{0} \in B_{\tau}$; finally, for each $P \in B$, we have that $P_{0} \leq P$.

Consider an arbitrary $G \in R_{S}$; because $S$ is consistent, such $G$ exists. Then, $\| \llbracket q \rrbracket_{G} \mid$ is the number of substitutions $\pi$ with dom $(\pi)=\operatorname{var}(q)$ and $\operatorname{rng}(\pi) \subseteq \operatorname{dom}(\mu)$ such that $q_{\pi} \subseteq G$. Since $G \in R_{S}$, for each such $\pi$, the substitution $\tau$ such that $\tau(x)=\mu(\pi(x))$, for each $x \in \operatorname{dom}(\pi)$, is an answer to $\mu(q)$ on $H$. In addition, we have that $\pi \in \exp _{\tau}\left(P_{0}\right)$. So, by Definition 3.4 of expected number of answers, the following holds:

$$
E_{q, S}=\sum_{G \in R_{S}} \frac{\left|\llbracket q \rrbracket_{G}\right|}{\left|R_{S}\right|}=\sum_{G \in R_{S}} \sum_{\tau \in \llbracket \mu(q) \rrbracket_{H}} \sum_{\pi \in \exp _{\tau}\left(P_{0}\right)} \frac{\left|\llbracket q \pi \rrbracket_{G}\right|}{\left|R_{S}\right|}
$$



$$
E_{q, S}=\sum_{G \in R_{S}} \sum_{\tau \in \llbracket \mu(q) \rrbracket_{H}} \sum_{\pi \in \exp _{\tau}\left(P_{0}\right)} \frac{\left|\llbracket q_{\pi} \rrbracket_{G}\right|}{\left|R_{S}\right|}=\sum_{G \in R_{S}} \sum_{\tau \in \llbracket \mu(q) \rrbracket_{H}} \sum_{P \in B_{\tau}} \sum_{\pi \in \exp _{\tau}(P)} \frac{\left|\llbracket q q_{\pi} \rrbracket_{G}\right|}{\left|R_{S}\right|} .
$$

So, by the definition of $E_{q_{\pi}}, S$, we can compute $E_{q, S}$ as follows:

$$
E_{q, S}=\sum_{\tau \in \llbracket \mu(q) \rrbracket_{H}} \sum_{P \in B_{\tau}} \sum_{\pi \in \exp _{\tau}^{〔}(P)} \sum_{G \in R_{S}} \frac{\left|\llbracket q_{\pi} \rrbracket_{G}\right|}{\left|R_{S}\right|}=\sum_{\tau \in \llbracket \mu(q) \rrbracket_{H}} \sum_{P \in B_{\tau}} \sum_{\pi \in \exp _{\tau}^{〔}(P)} E_{q_{\pi}}, S
$$

By Lemma B.6, we then have that

$$
E_{q, S}=\sum_{\tau \in \llbracket \mu(q) \rrbracket_{H}} \sum_{P \in B_{\tau}} \sum_{\pi \in \exp _{\tau}^{<}(P)} F_{\tau}(P) .
$$

Finally, Proposition B. 4 shows how to count the number of $\leq$-maximal $\tau$-expansions, so we can rewrite $E_{q, S}$ as follows:

$$
E_{q, S}=\sum_{\tau \in \llbracket \mu(q) \rrbracket_{H}} \sum_{P \in B_{\tau}} F_{\tau}(P) \cdot \sum_{P^{\prime} \in B_{\tau}, P \leq P^{\prime}} K\left(P, P^{\prime}\right) \cdot N_{\tau}\left(P^{\prime}\right) .
$$

As pointed out in Section 4.2, Proposition 4.6 follows immediately from Theorem 4.3.
Proposition 4.6. Formula (1) correctly computes $E_{q, S}$ when $q$ is a $\mu$-unification-free query.

## B.4 Proof of Theorem 4.4

We next show that our method for computing $\sigma_{q, S}$ is correct.
Theorem 4.4. Identity $\sigma_{q, S}^{2}=E_{q \cup \rho(q), S}-E_{q, S}^{2}$ holds, where $\rho$ is a substitution mapping each variable in $\operatorname{var}(q)$ to a fresh variable.
Proof. Let $\rho$ be a substitution mapping each variable in $\operatorname{var}(q)$ to a fresh variable, and let $q^{\prime}=\rho(q)$. We first show that equation (9) holds for each RDF graph $G \in R_{S}$.

$$
\begin{equation*}
\left|\llbracket q \rrbracket_{G}\right|^{2}=\left|\llbracket q \cup q^{\prime} \rrbracket_{G}\right| \tag{9}
\end{equation*}
$$

To this end, consider an arbitrary graph $G \in R_{S}$. Then, $\left|\llbracket q \rrbracket_{G}\right|^{2}$ is the cardinality of the set Pairs that contains each pair $\left\langle\pi, \pi^{\prime}\right\rangle$ of substitutions mapping the variables of $q$ to dom $(\mu)$ such that $\left\{\pi(q), \pi^{\prime}(q)\right\} \subseteq G$. But then, Equation (9) holds because the set of answers to $q \cup q^{\prime}$ over $G$ contains precisely, for each pair $\left\langle\pi, \pi^{\prime}\right\rangle \in$ Pairs, one distinct substitution $\kappa_{\pi, \pi^{\prime}}$ from $\operatorname{var}(q) \cup \operatorname{var}\left(q^{\prime}\right)$ to dom $(\mu)$. Such $\kappa_{\pi, \pi^{\prime}}$ can be obtained from $\pi$ and $\pi^{\prime}$ by setting, for each $x \in \operatorname{var}(q)$,

$$
\kappa_{\pi, \pi^{\prime}}(x)=\pi(x) \text { and } \kappa_{\pi, \pi^{\prime}}(\rho(x))=\pi^{\prime}(x)
$$

We are now ready to prove the theorem. By Definition 3.4 of variance, we have that

$$
\sigma_{q, S}^{2}=\sum_{G \in R_{S}} \frac{\left.\left|\llbracket q \rrbracket_{G}\right|^{2}-2 \cdot\left|\llbracket q \rrbracket_{G}\right| \cdot E_{q, S}+E_{q, S}\right)^{2}}{\left|R_{S}\right|}
$$

By using Equation (9) and by manipulating summations in the usual way, we can express $\sigma_{q, S}^{2}$ as follows.

$$
\sum_{G \in R_{S}}\left[\frac{\left|\llbracket q \cup q^{\prime} \rrbracket_{G}\right|}{\left|R_{S}\right|}\right]-2 \cdot E_{q, S} \cdot \sum_{G \in R_{S}}\left[\frac{\left|\llbracket q \rrbracket_{G}\right|}{\left|R_{S}\right|}\right]+E_{q, S}^{2}
$$

By Definition 3.4 of expectation, we then have that

$$
\sigma_{q, S}^{2}=E_{q \cup q^{\prime}, S}-2 \cdot E_{q, S}^{2}+E_{q, S}^{2}=E_{q \cup q^{\prime}, S}-\left(E_{q, S}\right)^{2},
$$

and so the theorem holds.


[^0]:    This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.
    WWW 2018, April 23-27, 2018, Lyon, France
    © 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.
    ACM ISBN 978-1-4503-5639-8/18/04.
    https://doi.org/10.1145/3178876.3186003

[^1]:    ${ }^{1}$ http://www.cs.ox.ac.uk/isg/tools/SumRDF/

[^2]:    ${ }^{2}$ Publishing benchmarks with the system named is prohibited by the system's license.

