



HAL
open science

Efficient Enumeration of Recursive Plans in Transformation-based Query Optimizers

Amela Fejza, Pierre Genevès, Nabil Layaïda

► **To cite this version:**

Amela Fejza, Pierre Genevès, Nabil Layaïda. Efficient Enumeration of Recursive Plans in Transformation-based Query Optimizers. 2022. hal-03692274v2

HAL Id: hal-03692274

<https://inria.hal.science/hal-03692274v2>

Preprint submitted on 24 Oct 2022 (v2), last revised 3 Apr 2024 (v7)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Efficient Enumeration of Recursive Plans in Transformation-based Query Optimizers

Amela Fejza*, Pierre Genevès†, Nabil Layaida‡

Tyrex team, Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG

38000 Grenoble, France

Email: *amela.fejza@inria.fr, †pierre.geneves@inria.fr, ‡nabil.layaida@inria.fr,

Abstract—Query optimizers built on the transformation-based Volcano/Cascades framework are used in many database systems. Transformations proposed earlier on the logical query dag (LQDAG) data structure, which is key in such a framework, focus only on recursion-free queries. In this paper, we propose the recursive logical query dag (RLQDAG) which extends the LQDAG with the ability to capture and transform recursive queries, leveraging recent developments in recursive relational algebra. Specifically, this extension includes: (i) the ability of capturing and transforming sets of recursive relational terms thanks to (ii) annotated equivalence nodes used for guiding transformations that are more complex in the presence of recursion; and (iii) RLQDAG rewrite rules that transform sets of subterms in a grouped manner, instead of transforming individual terms in a sequential manner; and that (iv) incrementally update the necessary annotations. Core concepts of the RLQDAG are formalized using a syntax and formal semantics with a particular focus on subterm sharing and recursion. The result is a clean generalization of the LQDAG transformation-based approach, enabling more efficient explorations of plan spaces for recursive queries. An implementation of the proposed approach shows significant performance gains compared to the state-of-the-art.

I. INTRODUCTION

Recursive queries enable powerful information extraction, especially from linked data structures such as trees and graphs. However, important data management system components, such as the widely used Volcano framework [30], were designed for recursion-free queries. A typical transformation-based query optimizer operates by (i) translating a query into a relational algebraic term, (ii) applying algebraic transformations in order to search for equivalent yet more efficient evaluation plans, during a so-called *plan enumeration phase*, (iii) executing the query by running one of the explored plans.

Works on extending relational algebra (RA) with recursion [2], [11] resulted in powerful recursive relational algebras [5], [6], [35], capable of capturing queries with transitive closures [5] and even more general forms of recursion [6], [35]. This line of works recently led to μ -RA [35] which provides a rich set of rewrite rules for recursive terms enabling efficient recursive evaluation plans not available with earlier approaches.

The enumeration phase is crucial as it may produce terms which are drastically more efficient. It has been heavily researched for recursion-free queries. It is common practice to allocate a time budget for this enumeration phase, as it is notoriously known that exhaustive plan space explorations

may not be feasible in practice for certain queries. The faster we generate the space of equivalent plans, the more likely we will be able to find terms with more efficient evaluation.

With recursion, plan spaces are often significantly larger than in the non-recursive setting, due to new interplays between recursive and non-recursive operators. The efficiency of recursive plan enumeration becomes critical. The speed of plan enumeration directly determines whether query evaluation plans enabled by, e.g. μ -RA [35], are within range or theoretically reachable but still out of reach for a practical query optimizer. This motivates the search for efficient methods for enumerating recursive plans.

Contributions: We present the recursive logical query dag (RLQDAG), which extends Volcano’s LQDAG [30] and the μ -RA framework [35] for the purpose of efficiently enumerating recursive query plans. Contributions include the first extension of the LQDAG with the support of recursive terms; a formalization of important RLQDAG concepts in terms of formal syntax and semantics, with a particular focus on the sharing of common subterms in the presence of recursion; RLQDAG transformations that generalize rewrite rules for individual recursive terms of [35] and enable efficient and grouped transformations of sets of recursive terms. The RLQDAG relies on a concept of annotated equivalence nodes with incremental updates, used for guiding transformations of recursive subterms. Contributions also include a complete prototype implementation of the proposed approach; and an experimental assessment using third-party queries on synthetic and real datasets.

II. BACKGROUND AND RELATED WORK

A. Approaches in recursive query optimization

Recursion is considered in only a small fraction of the numerous works on query optimization. Three main lines of work with recursive queries can be identified:

- The Datalog line of works [1], [8], [17], [23], [33], [52], [56], [60] developed methods for optimizing recursive queries formulated in Datalog: magic-sets [10], [26], [50], demand transformations [57], automated reversals [45], and the FGH rule [61]. Although the syntax of Datalog greatly differs from RA, the effects of magic-sets [10], [26], [50] and of demand transformations [57] are comparable to pushing certain kinds of selections

and projections. These techniques are very sensitive to whether the Datalog program is written in a left-linear or right-linear manner, but one can use the automated reversal technique proposed in [45] to fully exploit them. The optimization framework proposed in [61] gathers magic-sets, semi-naive evaluation and proposes a new FGH rule for optimizing recursive Datalog programs with aggregations. Datalog engines do not explore plan spaces but use heuristics to find a good plan to evaluate queries. However, currently, no matter which combination of existing Datalog optimizations a Datalog engine implements, it will not be able to find plans where recursions have been merged automatically similar to those found by the μ -RA approach [35]. This is because, currently, in a Datalog program corresponding to the optimized translation of A^+/B^+ at least one of the two transitive closures A^+ or B^+ will be fully materialized (even if there is no solution to A^+/B^+). On real datasets, this can make Datalog query evaluation an order of magnitude slower than query evaluation with RA-based systems, as noticed in [35].

- The line of works based on relational algebra [5], [6], [15], [35], [36] attempts to extend relational algebra with operators to capture forms of recursion. For instance, α -RA [5] extends RA with an operator to capture transitive closures. LFP-RA [6] proposes a more general least fixpoint operator, and an extension of this work gave effective criteria for optimization in its presence [36]. Recently, μ -RA [35] proposed to extend RA with a μ operator which is also a fixpoint with an appropriate set of restrictions. This enables μ -RA to combine all earlier RA-based recursion optimization rules in the same framework [35], while adding new rewrite rules for recursive terms, in particular for merging recursions. This makes μ -RA the most advanced system for RA-based recursive query optimization as it can generate plans not reachable by other approaches. The approach that we propose in this paper extends μ -RA with a new method for enumerating recursive plans much more efficiently, thanks to a generalization of the rules presented in μ -RA so that the generalized rules (presented in Section IV) apply directly on a factorized representation on the recursive plan space.
- Ad-hoc optimizations for regular queries. Both Datalog-based and RA-based approaches can capture queries with expressive forms of recursion, going beyond regularity. Several works have focused on optimizing queries in which recursion is restricted to regular patterns, such as regular path queries (RPQs). Automata based approaches have been developed to answer RPQ queries [27], [37], [66]. A hybrid approach that combines finite state machines and α -RA is presented in [63] and extended in [3]. All these works are limited to RPQs and their unions or conjunctions. In comparison, the work we present in this paper supports more expressive forms of recursion, that may include non-regular patterns (see e.g. the experimental section V with queries of the form

$A^n B^n$ for instance).

B. Limits of the RA-based approach

The RA line of work offers several advantages including high expressivity and rich plan spaces. In addition, it can be seen as an interesting approach for extending the (non-recursive) RA approach already widely used in RDBMS implementations. One main limit however, is that the whole approach critically depends on the ability to quickly enumerate plans during the plan exploration phase. While plan enumeration has not been studied yet in the presence of recursion, it has been extensively studied for recursion-free queries.

C. Plan enumeration for non-recursive (SPJ) queries

Most of the works on plan enumeration focus on select-project-join (SPJ) queries. Techniques proposed for SPJ can be divided into two main groups: bottom-up and top-down approaches. Bottom-up approaches generate the plan space by starting from the leaves (initial relations) and going up in the tree of operators when progressively exploring alternate combinations of operators. In contrast, top-down approaches start from the root and recursively explore subbranches in search for possible alternatives.

Pioneering and seminal works on the bottom-up approach were proposed in [51] with refinements in execution strategies [41] and the Starburst plan generator [32], [40].

The top-down approach was introduced with EXODUS [29] and the seminal works on Volcano [30] and Cascades [28] that added memoization to increase efficiency. The purpose of memoization is to re-use information already found in order to avoid redundant calculations. An advantage of the top-down approach is that it makes it possible to use branch and bound pruning, as noticed in [54].

1) *Join enumeration:* With SPJ queries the crux of plan enumeration is join enumeration [46]. The join enumeration problem consists in finding the best join order for a query with n join operators. It is an NP-hard problem [34], and has been the focus of extensive research.

Ono and Lohman showed that the worst-case complexity of join enumeration is $O(3^n)$ (reached with the clique query structure) [47]. Join enumeration can be made as efficient as the lower bound of $O(3^n)$ as long as the generation of duplicates is avoided [48]; and this is implementable with a technique based on Volcano memoization [48].

Algorithmic improvements for the bottom-up approach have been researched in [43], [44], [59]. Top-down approaches have been shown to be competitive with bottom-up ones in [16]. This triggered further research which resulted in many algorithmic improvements of the top-down approach [16], [18]–[21], culminating with [53] which proposes an algorithm based on [21] for enumerating, à la Volcano, cross-product free trees of join operators in a complete manner.

2) *Limits of existing approaches and difficulties for extending them with recursion:* All the previous approaches focus mainly on SPJ queries, with some extensions to support outer joins [19], [20], [24], [25], [42] and aggregations [14].

Very few works consider other operators, as noticed in [13] and [53]. In particular, union has been largely neglected in the literature, with the notable exception of [65] that proposes a union pushdown technique for disjunctive queries. The idea is that union can be placed immediately after the join operators whenever the selectivity of the join operator is very low. To the best of our knowledge, plan enumeration for transformation-based query optimizers has not been studied yet in the presence of recursive operators.

The reason why join enumeration constitutes the crux of the problem of enumerating plans for SPJ queries is because with those queries that are union-free and recursion-free, selections and projections can be abstracted. Selections and projections can be pulled in the tree of operators of an SPJ, reaching a phase where only the remaining join operators are under scrutiny, and once the problem of finding a good join order is solved, selections and projections can be pushed back as close to the sources as possible. In the presence of recursion however, the situation is not comparable anymore, because selections and projections cannot be pulled and pushed back so easily. This makes recursion-free plan enumeration techniques hardly reusable in the presence of recursion.

Union and recursion greatly extend the expressive power of SPJ queries. However, they not only make plan enumeration significantly more complex, but they also generate significantly larger plan spaces. This is because their addition generates many new possible combinations to be explored due to new interplays, for instance between unions and joins (e.g. distributivity of natural join over union) or between recursions and joins. This worsens the combinatorics of plans to be enumerated and motivates even more the interest of finding efficient techniques for enumerating recursive plans.

D. The logical query dag (LQDAG)

The method that we propose extends a key component used in the top-down enumeration approach: the logical query dag (LQDAG). The LQDAG is a directed acyclic graph data structure used to represent and generate the logical plan space. It was introduced in [30] and improved in [28] by the same authors. It is also well described as the AND-OR-DAG in [49], [53] where it is used for detecting and unifying common subexpressions for multi-query optimization [49]; and for generating the space of cross-product free join trees [53].

The LQDAG contains nodes of two different types: *equivalence* nodes and *operation* nodes. Equivalence nodes can only have operation nodes as children and vice versa: operation nodes can only have equivalence nodes as children. The purpose of an equivalence node is to explicitly regroup equivalent subterms. An operation node corresponds to an algebraic operation like: join (\bowtie), filter (σ_θ) etc. The LQDAG can be seen as a factorized representation of a set of terms.

Inspired from [49], Figure 1 illustrates a sample LQDAG, and Figure 2 depicts its expansion obtained after the application of commutativity and associativity rules for the join operator.




Fig. 1. An initial LQDAG.




Fig. 2. An expanded LQDAG.

III. THE RLQDAG

The RLQDAG extends the LQDAG with the ability to capture and transform sets of recursive terms.

A. Syntax

First of all, we propose a syntax for the RLQDAG. The purpose is to be able to syntactically express a term that denotes a (potentially very large) set of recursive algebraic terms. This makes it possible to develop transformations of sets of terms formally (i.e. with a high level of precision), and express them as rewrite rules that transform one RLQDAG term d into another RLQDAG term d' . The syntax of RLQDAG terms, given in Fig. 3, focuses on formalizing the concepts of equivalence nodes, operation nodes, sharing of common subterms, and recursion.

γ	::=	$\begin{array}{ l} \alpha \\ Y \end{array}$	Pointer to equivalence node Equivalence node Reference
α	::=	$\begin{array}{ l} d \\ d, \alpha \\ \text{let } Y = \alpha_1 \text{ in } \alpha_2 \end{array}$	Equivalence node internals Operation node Operation nodes Reference binder
d	::=	$\begin{array}{ l} X \\ \rho_a^b(\gamma) \\ \sigma_f(\gamma) \\ \gamma \bowtie \gamma' \\ \gamma \triangleright \gamma' \\ \gamma \cup \gamma' \\ \tilde{\pi}_a(\gamma) \\ \mu X. \gamma \cup \alpha_{rec} \end{array}$	Operation node Relation variable Rename Filter Join Antijoin Union Antiprojection Fixpoint (recursion)
α_{rec}	::=	$[\alpha]_{\mathfrak{D}}$	Annotated equivalence node

Fig. 3. Syntax of RLQDAG terms.

An equivalence node is a node that can have several operation nodes d as children, possibly with binders. The binder construct “let $Y = \alpha_1$ in α_2 ” enables the explicit sharing of a common equivalence node α_1 within the branches of another equivalence node α_2 . For that purpose, it assigns a new fresh reference name Y to α_1 , and allows Y to be used multiple times in α_2 as a reference to α_1 . Hence, the general definition of γ is either an equivalence node $[\alpha]$ or a reference Y to an existing equivalence node.

Operation nodes are defined by the variable d in the abstract syntax of Fig. 3. They include the main algebraic operations of recursive relational algebra. Each operand of an operation node d points in turn to an equivalence node (through γ). The rename operator $\rho_a^b(\gamma)$ renames column a into column b in the equivalence node γ . The filter operator $\sigma_f(\gamma)$ applies the filtering expression f to the equivalence node γ . The antiprojection operator $\tilde{\pi}_a(\gamma)$ removes column a from the equivalence node γ .

Recursive terms can be expressed using fixpoint operation nodes. The principle, inspired from earlier works in recursive relational algebras [6], [35], consists in the introduction of a least-fixpoint binder operation node (μ) that binds a fresh variable X to some expression in which the variable X can appear, thus explicitly denoting recursion. In the RLQDAG, as defined in the abstract syntax of Fig. 3 and illustrated in Fig. 4, a fixpoint operator node is written $\mu X. \gamma \cup \alpha_{rec}$. The first operand γ is an equivalence node that models the constant part (the base case) of the recursion. X cannot occur within γ . The other equivalence node α_{rec} is the recursive part. An important aspect is that α_{rec} contains at least one free occurrence of the recursive variable X . This is a major difference between the fixpoint operation node and the other operation nodes. This aspect will lead to a number of new definitions and formal developments. This is because depending on how the recursive variable is used in that branch, the transformation and sharing of subterms in the RLQDAG may, or may not, be allowed.

For example, the following RLQDAG term corresponds to the transitive closure of some relation A :

$$\mu X. [A] \cup [[A] \bowtie [X]]_{\mathfrak{D}}^{\mathfrak{R}}$$

In this case, the equivalence node for the constant part contains the relation variable A , and the equivalence node for the recursive part is composed of the join between A and X . In a RLQDAG, the recursive equivalence node of each fixpoint operation node is annotated (with annotations \mathfrak{D} and \mathfrak{R}). These annotations will be useful for guiding the application of transformations and will be defined and explained in Section III-D.

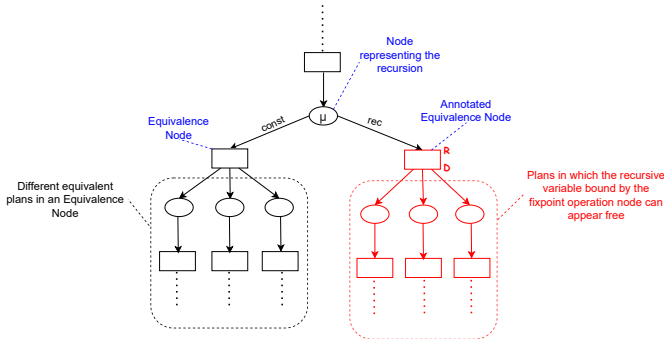


Fig. 4. Structure of a recursive term in the RLQDAG.

There are two reasons why we distinguish α_{rec} from a general equivalence node α in the abstract syntax. The first reason is that those equivalence nodes for recursive parts are

equipped with annotations (see Section III-E.). The second reason is that we want to allow a maximum level of sharing while preventing the sharing of subterms with free occurrences of a recursive variable. We thus forbid the use of the binding construct to share subterms with free variables.

We consider the following (usual) restrictions over the abstract syntax presented in Fig. 3: we consider only positive, linear and non mutually recursive RLQDAG terms: (i) positive means that recursive variables only appear in the left-hand operand of an antijoin operator node; (ii) linear means that one of the operands in a join or antijoin operator node is constant in the free variable; (iii) Non-mutually recursive terms means that fixpoint operator nodes are properly nested so that there is only one free variable in any subbranch of an annotated equivalence node (this variable may occur several times).

B. Semantics of RLQDAG terms

The interpretation of a RLQDAG term is the set of all recursive relational algebraic terms that it represents. Formally, the semantics of a RLQDAG $[\alpha]$ is a set of μ -RA terms as defined by the functions $S_\alpha[\llbracket \cdot \rrbracket]$ and $S_\gamma[\llbracket \cdot \rrbracket]$ presented in Fig. 5, where E denotes a variable environment used to keep track of the variable definitions introduced by binders for the sharing of subterms. The interpretation of a RLQDAG $[\alpha]$ is $S_\alpha[\llbracket \alpha \rrbracket]_\emptyset$.

$$\begin{aligned} S_\gamma[\llbracket [\alpha] \rrbracket]_E &= S_\alpha[\llbracket \alpha \rrbracket]_E \\ S_\gamma[\llbracket [Y] \rrbracket]_E &= S_\alpha[\llbracket E(Y) \rrbracket]_E \\ \\ S_\alpha[\llbracket [d] \rrbracket]_E &= S_d[\llbracket d \rrbracket] \\ S_\alpha[\llbracket [d, \alpha] \rrbracket]_E &= S_d[\llbracket d \rrbracket] \cup S_\alpha[\llbracket \alpha \rrbracket]_E \\ S_\alpha[\llbracket [\text{let } Y = \alpha_1 \text{ in } \alpha_2] \rrbracket]_E &= S_\alpha[\llbracket \alpha_2 \rrbracket]_{E \oplus \{y \mapsto \alpha_1\}} \\ \\ S_d[\llbracket [X] \rrbracket] &= X \\ S_d[\llbracket [\sigma_f(\gamma)] \rrbracket] &= \{ \sigma_f(t) \mid t \in S_\gamma[\llbracket \gamma \rrbracket]_E \} \\ S_d[\llbracket [\gamma_1 \bowtie \gamma_2] \rrbracket] &= \{ t \bowtie t' \mid t \in S_\gamma[\llbracket \gamma_1 \rrbracket]_E \wedge \\ &\quad t' \in S_\gamma[\llbracket \gamma_2 \rrbracket]_E \} \\ S_d[\llbracket [\gamma_1 \triangleright \gamma_2] \rrbracket] &= \{ t \triangleright t' \mid t \in S_\gamma[\llbracket \gamma_1 \rrbracket]_E \wedge \\ &\quad t' \in S_\gamma[\llbracket \gamma_2 \rrbracket]_E \} \\ S_d[\llbracket [\gamma_1 \cup \gamma_2] \rrbracket] &= \{ t \cup t' \mid t \in S_\gamma[\llbracket \gamma_1 \rrbracket]_E \wedge \\ &\quad t' \in S_\gamma[\llbracket \gamma_2 \rrbracket]_E \} \\ S_d[\llbracket [\rho_a^b(\gamma)] \rrbracket] &= \{ \rho_a^b(t) \mid t \in S_\gamma[\llbracket \gamma \rrbracket]_E \} \\ S_d[\llbracket [\tilde{\pi}_a(\gamma)] \rrbracket] &= \{ \tilde{\pi}_a(t) \mid t \in S_\gamma[\llbracket \gamma \rrbracket]_E \} \\ S_d[\llbracket [\mu X. \gamma \cup \alpha_{rec}] \rrbracket] &= \{ \mu X. t \cup t_{rec} \mid t \in S_\gamma[\llbracket \gamma \rrbracket]_E \\ &\quad \wedge t_{rec} \in S_\alpha[\llbracket \alpha_{rec} \rrbracket]_E \} \end{aligned}$$

Fig. 5. Formal semantics of RLQDAG terms.

a) *Notion of well-formedness:* A well-formed RLQDAG is a RLQDAG whose interpretation is a set of semantically equivalent terms:

Definition 1 (Well-formedness). A RLQDAG $[\alpha]$ is well-formed if and only if $\forall t, t' \in S_\alpha[\llbracket \alpha \rrbracket]_\emptyset, \text{eval}(t) = \text{eval}(t')$.

In this definition $\text{eval}(t)$ returns the interpretation of the individual recursive relational algebraic term t (i.e. the set of mappings returned by t ; or in other terms the resulting relational table of t), as defined by e.g. the formal semantics function $\llbracket \cdot \rrbracket_V$ defined for μ -RA terms in [35].

b) *Types*: The *type* of an operation node d is the set of column names obtained in the result of the evaluation of any subbranch of d . In a well-formed RLQDAG, all d_i under the same equivalence node are semantically equivalent, and thus have the same type. For this reason, we also define the *type of an equivalence node*: $type(\gamma)$ as the type of one of its operation nodes. Notice that the type of an annotated equivalence node of some fixpoint operation node d corresponds to the type of the equivalence node of the constant part of d . For a given filter operation node $\sigma_f(\gamma)$, we denote by $filt(f) \subseteq type(\sigma_f(\gamma))$ the subset of column names used in the filtering function f .

C. Recursive terms and rule applicability

A significant novelty introduced by recursion when compared to the classical setting of non-recursive algebraic terms resides in the criteria used to trigger rewrite rules. In the classical non-recursive setting these criteria are trivial in the sense that they only depend on top-level operators. For example, when applying join distributivity over union, the applicability of the rewrite rule $A \bowtie (B \cup C) \longrightarrow (A \bowtie B) \cup (A \bowtie C)$ can be determined by examining only the combination of the two top-most operators, i.e. the top-level (\bowtie) with the operator immediately underneath (i.e. \cup). For some other rewrite rules, applicability criteria may also include some additional verifications such as (non)-interaction between e.g. the set of columns being filtered, or the columns being removed (in the cases of filter and antiprojection, respectively). In any case of the non-recursive setting, these verifications need only to look at the two top-most algebraic operators under scrutiny (and their potential immediate parameters). No further traversal of subtrees of operators are required. For instance, in the previous example of join distributivity over union, B and C do not need to be traversed at all when determining rule applicability.

In sharp contrast, rules for transforming recursive terms rely on criteria that are significantly more sophisticated as they sometimes require a whole traversal of the recursive part of a fixpoint term. This is because opportunities for rule application with recursive terms depend on how the recursive variable is used within the recursive parts of fixpoints. It is known since the works of [6], [35], [36] that criteria for rule application are significantly more sophisticated in the presence of recursion as they need to examine how the recursive variables are used. We show that it is still possible to apply rules over sets of recursive terms at once, using the concept of *annotated equivalence nodes*, that we define in § III-E, and for which we need some preliminary definitions.

D. Preliminary definitions of auxiliary functions in RLQDAG

Definition 2 (Unfolding). *Let α be an equivalence node. The unfolding of α , denoted $unfold(\alpha)$ is α in which all occurrences of equivalence node variable names Y are replaced by their definitions (binders are simply unfolded).*

We now define two auxiliary functions over RLQDAG terms, that generalize the definitions for individual algebraic terms found in [35]. These functions will be used for defining annotated equivalence nodes in Section III-E.

a) Notion of destabilizer in RLQDAG:

We define the destabilizer ($destab()$) as the set of columns that can be changed during an iteration of the fixpoint. $destab()$ can be seen as the complement of the stabilizer calculated for individual terms in μ -RA [35], which is generalized for RLQDAG terms (i.e. sets of μ -RA terms). The calculation of $destab()$ relies on a set of derivations, and it needs to traverse term subtrees in order to analyze how the occurrence of free variables are used.

Definition 3. *For a fixpoint operation node $\mu X. \gamma \cup \alpha_{rec}$ we consider $\alpha' = unfold(\alpha_{rec})$ and we define $destab(\alpha', X)$ as the following set of column names: $destab(\alpha', X) = \{c \in \mathfrak{C} \mid \exists p \in d(\alpha', X) \ p(c) \neq c\}$ where $d(\cdot, \cdot)$ computes the set of derivations [35] over a RLQDAG term:*

$$\begin{aligned}
d((d, \alpha), X) &= d(d, X) \\
d([\alpha_1] \cup [\alpha_2], X) &= d(\alpha_1, X) \cup d(\alpha_2, X) \\
d([\alpha_1] \triangleright [\alpha_2], X) &= d(\alpha_1, X) \\
d([\alpha_1] \bowtie [\alpha_2], X) &= d(\alpha_1, X) \cup d(\alpha_2, X) \\
d(\rho_a^b([\alpha]), X) &= \{p \circ (b \rightarrow a, a \rightarrow \perp) \mid p \in d(\alpha, X)\} \\
d(\tilde{\pi}_a([\alpha]), X) &= \{p \circ (a \rightarrow \perp) \mid p \in d(\alpha, X)\} \\
d(\sigma_f([\alpha]), X) &= d(\alpha, X) \\
d(\mu(Z. \gamma \cup [\alpha]_{\mathfrak{D}}^{\mathfrak{R}}), X) &= \emptyset \\
d(X, X) &= \{\emptyset\} \text{ (a singleton identity)} \\
d(R, X) &= \emptyset
\end{aligned}$$

and where \circ represents the composition and $(a_1 \rightarrow b_1, \dots, a_n \rightarrow b_n)$ denotes the function that maps each a_i to its b_i and every other column name to itself.

b) *Notion of rigidity in RLQDAG*: We define a function $rigid()$ that computes the set of columns that cannot be added nor removed from a fixpoint operation node. Intuitively, $rigid()$ can be thought as the “complement” of the boolean predicate add (introduced for individual terms in μ -RA [35]) but which is generalized for a RLQDAG term, and returns a set of columns (instead of a boolean). A column $c \in \mathfrak{C}$ cannot be added nor removed from an annotated equivalence node α_{rec} (recursive in X) when $c \in rigid(unfold(\alpha), X)$ and $rigid()$ is defined as follows:

Definition 4 (Rigidity).

$$\begin{aligned}
rigid((d, \alpha), X) &= rigid(d, X) \\
rigid([\alpha_1] \cup [\alpha_2], X) &= rigid(\alpha_1, X) \cup rigid(\alpha_2, X) \\
rigid([\alpha_1] \bowtie [\alpha_2], X) &= rigid(\alpha_1, X) \cup rigid(\alpha_2, X) \\
rigid([\alpha_1] \triangleright [\alpha_2], X) &= rigid(\alpha_1, X) \cup rigid(\alpha_2, X) \\
rigid(\rho_a^b([\alpha]), X) &= rigid(\alpha, X) \cup \{a, b\} \\
rigid(\tilde{\pi}_a([\alpha]), X) &= \emptyset \text{ when } X \notin free(\alpha) \\
&= rigid(\alpha, X) \cup \{a\} \text{ otherwise} \\
rigid(\sigma_f([\alpha]), X) &= rigid(\alpha, X) \cup filt(f) \\
rigid(\mu(Z. \gamma \cup [\alpha]_{\mathfrak{D}}^{\mathfrak{R}}), X) &= rigid(\alpha, X) \cup rigid(\gamma, X) \\
rigid(R, X) &= type(R) \text{ when } X \neq R \\
rigid(X, X) &= \emptyset
\end{aligned}$$

E. Annotated equivalence node

An annotated equivalence node (α_{rec} in the abstract syntax of RLQDAG terms given in Fig. 3) is an equivalence node of a recursive part of a fixpoint, which is annotated with information that characterize how the recursive variable is used. Specifically:

Definition 5. Given a RLQDAG operation node $d = \mu X. \gamma \cup \alpha_{rec}$, the annotated equivalence node α_{rec} is defined as:

$$[\alpha]_{\mathfrak{D}}^{\mathfrak{R}}$$

where $\mathfrak{D} = destab(\alpha, X)$ and $\mathfrak{R} = rigid(\alpha, X)$.

Annotations \mathfrak{D} and \mathfrak{R} are intended to characterize all the subterms of the annotated equivalence node (thanks to definition 6). The goal of annotated equivalence nodes is to guide and maximize the grouped application of transformations, while also maximizing the sharing of common subterms.

In the sequel, we detail RLQDAG transformations, by introducing new rewrite rules that are capable of transforming sets of subterms at once.

a) *Notion of consistency:* Intuitively, a RLQDAG is *consistent* iff it is well-formed and in addition, for any of its annotated equivalence node $[\alpha_2]_{\mathfrak{D}}^{\mathfrak{R}}$, the annotations \mathfrak{D} and \mathfrak{R} are the same for every operation nodes directly underneath (no matter on which subbranch of the equivalence node they are computed, they coincide). More formally:

Definition 6 (Consistency). A RLQDAG $[\alpha]$ is consistent iff:

- 1) it is well-formed;
- 2) for all fixpoint operator node $\mu X. \gamma \cup [\alpha_2]_{\mathfrak{D}}^{\mathfrak{R}}$ occurring in α , we have $cons(\alpha_2, X)_{\mathfrak{D}}^{\mathfrak{R}}$ where:

$$cons([d], X)_{\mathfrak{D}}^{\mathfrak{R}} = destab(d, X) = \mathfrak{D} \wedge rigid(d, X) = \mathfrak{R}$$

$$cons([d, \alpha], X)_{\mathfrak{D}}^{\mathfrak{R}} = destab(d, X) = \mathfrak{D} \wedge rigid(d, X) = \mathfrak{R} \wedge cons([\alpha], X)_{\mathfrak{D}}^{\mathfrak{R}}$$

In the remaining, we only consider consistent RLQDAGs.

IV. RLQDAG TRANSFORMATIONS

We now propose RLQDAG transformations whose purpose is to efficiently build the space of equivalent recursive plans. RLQDAG transformations generalize the rewrite rules proposed in [35] for individual μ -RA terms so that they can operate on the RLQDAG representation: i.e. on sets of μ -RA terms at once (instead of successively on individual terms). RLQDAG transformations leverage annotated equivalence nodes to guide the transformation of subterms. They also update the RLQDAG structure with new annotations when needed, in an incremental manner. The incremental aspect for updating annotations is important as it avoids numerous subterm traversals, thus enabling more efficient grouped transformations.

For each fixpoint transformation operation, we describe which parts are shared, how new generated terms are attached and what happens with the other plans already present in the equivalence node. The creation of new combinations of

operation nodes may in turn generate more opportunities for transformations that are also explored.

A. RLQDAG rewrite rules

We formalize all these ideas by introducing RLQDAG rewrite rules, based on the syntax of RLQDAG terms introduced in Section III. Specifically, RLQDAG rewrite rules are formalized as functions that take an equivalence node γ and return another equivalence node γ' obtained after applying transformations.

a) *Pushing filters into fixpoint operation nodes:* For pushing filters into sets of recursive terms, we introduce a function $pf()$, defined by considering all the syntactic decomposition cases of the input γ . Fig. 6 focuses on the two main cases that correspond to potential opportunities of pushing filters, i.e. the cases $pf([d])$ and $pf([d, \alpha])$ where d filters an equivalence node which contains a recursive subterm. For all the other cases, $pf()$ does not reorder operation nodes in the RLQDAG structure but simply traverses it in search for further transformation opportunities underneath¹. In Fig. 6, whenever the filter can be pushed through the fixpoint operation node, $pf()$ generates a new RLQDAG in which the filter operation node is put within the constant part of the fixpoint operation node. In the resulting equivalence node, the initial (suboptimal) term is replaced by the new term where filtering is done earlier (hence more efficient). For this transformation to happen, the criteria (on the last line of Fig. 6) must be satisfied. Whenever it is not the case, the RLQDAG structure is not reordered but simply recursively traversed in search for more transformation opportunities. The creation of new combinations of operation nodes may in turn provide new opportunities for other rewritings (for instance, new opportunities for pushing filters even further, or even other kinds of rewritings). This is the role of the $expand()$ function, formally defined in section IV-B. Intuitively, a call to $expand()$ on an equivalence node may further populate the equivalence node with new subterms. The $expand()$ function is in charge of exploring all opportunities for transformations. This is useful because other rewrite rules may apply, and the $expand()$ function basically triggers all possible applications of all rewrite rules.

Fig. 7 illustrates graphically a RLQDAG before transformation, and Fig. 8 depicts its updated structure obtained after the transformation defined in Fig. 6. New branches created by $pf()$ are represented in blue color (on both figures). The new term is added in the same equivalence node as the previous term, since they are semantically equivalent. Notice the incremental update of annotations performed by $pf()$ in Fig. 6: the annotations of the newly created term (in blue) are

¹For instance, two sample cases are the following:

$$pf(\mu X. \gamma \cup [\alpha]_{\mathfrak{D}}^{\mathfrak{R}}) = \mu X. expand(\gamma) \cup [expand(\alpha)]_{\mathfrak{D}}^{\mathfrak{R}}$$

$$pf([[A] \bowtie [B]]) = expand([A]) \bowtie expand([B])$$

where the $expand()$ function, defined in Section IV-B, simply traverses subterms in search for more transformation opportunities. $pf()$ is defined similarly for all the other syntactic cases of γ . Notice that when γ is a reference Y , there is no need to introduce a new binder, the reference name is used directly.

$$\text{pf}([\sigma_f([\mu X. \gamma \cup [\alpha]_{\mathcal{D}}^{\mathfrak{R}}])]) = \begin{cases} \bullet [\mu X'. \text{expand}([\sigma_f(\gamma)]) \cup [\text{expand}(\alpha_{\{X/X'\}})]_{\mathcal{D}}^{\mathfrak{R}'}] & \text{when } \text{filt}(f) \subseteq \mathcal{D} \\ \bullet [\sigma_f(\text{expand}([\mu X. \gamma \cup [\alpha]_{\mathcal{D}}^{\mathfrak{R}}])]) & \text{otherwise} \end{cases}$$

$$\text{pf}([\sigma_f([\mu X. \gamma \cup [\alpha]_{\mathcal{D}}^{\mathfrak{R}}, \alpha_2])]) = \begin{cases} \bullet [\mu X'. \text{expand}([\sigma_f(\gamma)]) \cup [\text{expand}(\alpha_{\{X/X'\}})]_{\mathcal{D}}^{\mathfrak{R}'}], & \text{expand}([\sigma_f([\alpha_2]])] \\ \text{when } \text{filt}(f) \subseteq \mathcal{D} & \\ \bullet [\sigma_f(\text{expand}([\mu X. \gamma \cup [\alpha]_{\mathcal{D}}^{\mathfrak{R}}, \alpha_2])]), & \text{expand}([\sigma_f([\alpha_2]])] & \text{otherwise} \end{cases}$$

where $\mathfrak{R}' = \mathfrak{R} \cup \text{filt}(f)$ and $\alpha_{\{X/X'\}}$ denotes α in which all occurrences of X are replaced by X' .

Fig. 6. Pushing a filter in an equivalence node containing a fixpoint operation node.

obtained from the annotations of the initial term. In that case \mathcal{D} is simply propagated whereas $\mathfrak{R}' = \mathfrak{R} \cup \text{filt}(f)$.

b) Pushing joins into fixpoint operation nodes: For pushing joins into sets of recursive terms we define a function $\text{pj}()$. $\text{pj}()$ takes an equivalence node γ as input and returns an expanded equivalence node γ' that contains all the subterms in γ with, in addition, all the subterms where all joins pushable in fixpoint operation nodes have been pushed. We define $\text{pj}()$ for all possible syntactic decompositions of a RLQDAG. Fig. 9 presents the definition of $\text{pj}()$ for the two main cases of interest. Again, other cases are defined without structure rearrangement but involving recursive calls to $\text{expand}()$ in search for further transformation opportunities underneath.

Notice that whenever a join can be pushed within a fixpoint operation node (see Fig. 9), it is possible to share the constant part of the fixpoint operation node. This is made explicit by the creation of the outermost “let” binder whose goal is to define and associate a name to the equivalence node, so that it can be referred multiple times (thus explicitly showing the sharing of subterms). Fig. 10 illustrates graphically the RLQDAG generated in Fig. 9. It shows that the newly created branch (in blue color) extends the set of semantically equivalent terms of the existing equivalence node.

c) Merging fixpoint operation nodes: The function $\text{mf}()$ defined in Fig. 11 takes an input γ and returns an equivalence node γ' containing all the subterms in γ with, in addition, all the terms in which recursions that can be merged are merged. A merging happens whenever (i) two recursions are joined and (ii) their annotated equivalence nodes allow them to be merged into a single recursion, as described in Fig. 11. The constant part of the new recursive term created is the join of the constant parts of the two initial fixpoints, and a new recursive part is created. Since the constant part has changed, a new recursive variable is introduced and recursive parts are also new (and cannot be shared²). New equivalence nodes are

²This does not prevent the sharing of potential subparts with no occurrence of a free variable.




Fig. 7. Initial RLQDAG before expansion with opportunity to apply $\text{pf}()$.




Fig. 8. Expansion of a RLQDAG by pushing a filter in a fixpoint operation node, with branches added by $\text{pf}()$ in blue color.

created. Fig. 12 illustrates the new recursive term produced, using subterm-sharing, after the transformation.

d) Pushing an antiprojection in a fixpoint operation node: For pushing antiprojections into fixpoint operation nodes we introduce a function $\text{pp}()$ that takes an equivalence node γ as input and returns an expanded equivalence node γ' where all pushable antiprojections have been pushed. $\text{pp}()$ is defined in Fig. 16.

The antiprojection is pushed when the criteria allows it, and this results in the creation of a new term which replaces the initial one in the existing equivalence node. Whenever the criteria is not satisfied, the initial term is left unchanged but traversed in search for more transformation opportunities.

Fig. 14 illustrates graphically the RLQDAG before the transformations. Fig. 15 shows the graphical representation of the RLQDAG obtained by the transformation of Fig. 13.

e) Pushing an antijoin in a fixpoint operation node: For pushing antijoins into fixpoints we introduce a function $\text{pa}()$ that takes an equivalence node γ as input and returns an expanded equivalence node γ' that contains all the subterms in γ with, in addition, all the subterms where all pushable antijoins have been pushed in fixpoint operation nodes. $\text{pa}()$ is defined in Fig. 16. Again, Fig. 16 focuses on the syntactic cases that correspond to when an antijoin might be pushed in a fixpoint. When criteria are satisfied, the antijoin is pushed in the constant part of the fixpoint operation node and the newly created term is added to the set of equivalent terms along with the rest. Fig. 17 illustrates graphically the RLQDAG created by the transformations of Fig. 16.

Notice that RLQDAG rules can be divided into two groups. The first group is composed of rewrite rules $\text{pj}()$, $\text{pa}()$ and $\text{mf}()$. The application of these rules always preserve preexisting terms in an equivalence node. These rules can only perform additions of new operation nodes in an equivalence node³. The second group is composed of the rules $\text{pf}()$ and $\text{pp}()$ that

³One operation node will then be selected using cost estimation heuristics depending on statistics on the cardinalities of the database instance.

$$\text{pj}([\beta \bowtie [\mu X. \gamma \cup [\alpha]_{\mathfrak{D}}^{\mathfrak{R}}]]) = \left\{ \begin{array}{l} \bullet \text{ let } \text{const} = \gamma \text{ in} \\ \quad [\text{expand}(\beta) \bowtie \text{expand}([\mu X. \text{const} \cup [\alpha]_{\mathfrak{D}}^{\mathfrak{R}}]) , \\ \quad \mu X'. \text{expand}([\beta \bowtie [\text{const}]] \cup [\text{expand}(\alpha_{\{X/X'\}})]_{\mathfrak{D}'}^{\mathfrak{R}'} } \\ \text{when } \text{type}(\beta) \cap \mathfrak{D} = \emptyset \text{ and } \text{type}(\beta) \setminus \text{type}(\gamma) \cap \mathfrak{R} = \emptyset \\ \bullet [\text{expand}(\beta) \bowtie \text{expand}([\mu X. \gamma \cup [\alpha]_{\mathfrak{D}}^{\mathfrak{R}}])] \text{ otherwise} \end{array} \right.$$

$$\text{pj}([\beta \bowtie [\mu X. \gamma \cup [\alpha]_{\mathfrak{D}}^{\mathfrak{R}}, \alpha_2]]) = \left\{ \begin{array}{l} \bullet \text{ let } \text{const} = \gamma \text{ in} \\ \quad [\text{expand}(\beta) \bowtie \text{expand}([\mu X. \text{const} \cup [\alpha]_{\mathfrak{D}}^{\mathfrak{R}}, \alpha_2]) , \\ \quad \mu X'. \text{expand}([\beta \bowtie [\text{const}]] \cup [\text{expand}(\alpha_{\{X/X'\}})]_{\mathfrak{D}'}^{\mathfrak{R}'} } \\ \quad [\text{expand}(\beta) \bowtie \text{expand}([\alpha_2])]] \\ \text{when } \text{type}(\beta) \cap \mathfrak{D} = \emptyset \text{ and } \text{type}(\beta) \setminus \text{type}(\gamma) \cap \mathfrak{R} = \emptyset \\ \bullet [[\text{expand}(\beta) \bowtie \text{expand}([\mu X. \gamma \cup [\alpha]_{\mathfrak{D}}^{\mathfrak{R}}, \alpha_2])] , \\ \quad [\text{expand}(\beta) \bowtie \text{expand}([\alpha_2])]] \text{ otherwise} \end{array} \right.$$

where we define^a $\mathfrak{D}' = \mathfrak{D} \cup \text{destab}(\beta, X)$ and $\mathfrak{R}' = \mathfrak{R} \cup \text{rigid}(\beta, X)$.

Fig. 9. Pushing join in an equivalence node containing a fixpoint operation node.

^aThis is because we need to recompute the destabilizer on the underlying set of derivations updated with $d(\beta, X)$ and we have $\text{destab}(A \cup B) = \text{destab}(A) \cup \text{destab}(B)$




Fig. 10. Expansion of a RLQDAG by pushing a join in a fixpoint operation node. The initial RLQDAG is in black color, and parts added by $\text{pj}()$ are in blue color.

$$\text{mf}([\mu X_1. \gamma_1 \cup [\alpha_1]_{\mathfrak{D}_1}^{\mathfrak{R}_1}] \bowtie [\mu X_2. \gamma_2 \cup [\alpha_2]_{\mathfrak{D}_2}^{\mathfrak{R}_2}]) = \left\{ \begin{array}{l} \bullet \text{ let } \text{const}_1 = \gamma_1 \text{ in} \\ \quad \text{let } \text{const}_2 = \gamma_2 \text{ in} \\ \quad [\text{expand}([\mu X_1. \text{const}_1 \cup [\alpha_1]_{\mathfrak{D}_1}^{\mathfrak{R}_1}]) \bowtie \text{expand}([\mu X_2. \text{const}_2 \cup [\alpha_2]_{\mathfrak{D}_2}^{\mathfrak{R}_2}]) , \\ \quad \mu X. \text{expand}([\text{const}_1 \bowtie \text{const}_2]) \cup \text{expand}([\alpha_1_{\{X_1/X\}} \cup \alpha_2_{\{X_2/X\}}]_{\mathfrak{D}}^{\mathfrak{R}})] \\ \text{when } (\text{type}(\gamma_1) \cap \text{type}(\gamma_2)) \cap (\mathfrak{D}_1 \cup \mathfrak{D}_2) = \emptyset \text{ and } \text{type}(\gamma_1) \setminus \text{type}(\gamma_2) \cap \mathfrak{R}_2 = \emptyset \text{ and } \text{type}(\gamma_2) \setminus \text{type}(\gamma_1) \cap \mathfrak{R}_1 = \emptyset \\ \bullet \text{expand}([\mu X_1. \text{const}_1 \cup [\alpha_1]_{\mathfrak{D}_1}^{\mathfrak{R}_1}]) \bowtie \text{expand}([\mu X_2. \text{const}_2 \cup [\alpha_2]_{\mathfrak{D}_2}^{\mathfrak{R}_2}]) \text{ otherwise} \end{array} \right.$$

$$\text{mf}([\mu X_1. \gamma_1 \cup [\alpha_1]_{\mathfrak{D}_1}^{\mathfrak{R}_1}, \alpha_3] \bowtie [\mu X_2. \gamma_2 \cup [\alpha_2]_{\mathfrak{D}_2}^{\mathfrak{R}_2}, \alpha_4]) = \left\{ \begin{array}{l} \bullet \text{ let } \text{const}_1 = \gamma_1 \text{ in} \\ \quad \text{let } \text{const}_2 = \gamma_2 \text{ in} \\ \quad [\text{expand}([\mu X_1. \text{const}_1 \cup [\alpha_1]_{\mathfrak{D}_1}^{\mathfrak{R}_1}]) \bowtie \text{expand}([\mu X_2. \text{const}_2 \cup [\alpha_2]_{\mathfrak{D}_2}^{\mathfrak{R}_2}]) , \\ \quad \mu X. \text{expand}([\text{const}_1 \bowtie \text{const}_2]) \cup \text{expand}([\alpha_1_{\{X_1/X\}} \cup \alpha_2_{\{X_2/X\}}]_{\mathfrak{D}}^{\mathfrak{R}}) , \\ \quad \text{expand}([\alpha_3]) \bowtie \text{expand}([\mu X_2. \gamma_2 \cup [\alpha_2]_{\mathfrak{D}_2}^{\mathfrak{R}_2}, \alpha_4]) , \\ \quad \text{expand}([\alpha_4]) \bowtie \text{expand}([\mu X_1. \gamma_1 \cup [\alpha_1]_{\mathfrak{D}_1}^{\mathfrak{R}_1}, \alpha_3])] \\ \text{when } (\text{type}(\gamma_1) \cap \text{type}(\gamma_2)) \cap (\mathfrak{D}_1 \cup \mathfrak{D}_2) = \emptyset \text{ and } \text{type}(\gamma_1) \setminus \text{type}(\gamma_2) \cap \mathfrak{R}_2 = \emptyset \text{ and } \text{type}(\gamma_2) \setminus \text{type}(\gamma_1) \cap \mathfrak{R}_1 = \emptyset \\ \bullet [\text{expand}([\mu X_1. \gamma_1 \cup [\alpha_1]_{\mathfrak{D}_1}^{\mathfrak{R}_1}, \alpha_3]) \bowtie \text{expand}([\mu X_2. \gamma_2 \cup [\alpha_2]_{\mathfrak{D}_2}^{\mathfrak{R}_2}, \alpha_4]) , \\ \quad \text{expand}([\alpha_3]) \bowtie \text{expand}([\mu X_2. \gamma_2 \cup [\alpha_2]_{\mathfrak{D}_2}^{\mathfrak{R}_2}, \alpha_4]) , \\ \quad \text{expand}([\alpha_4]) \bowtie \text{expand}([\mu X_1. \gamma_1 \cup [\alpha_1]_{\mathfrak{D}_1}^{\mathfrak{R}_1}, \alpha_3])] \text{ otherwise} \end{array} \right.$$

where^a $\mathfrak{D} = \mathfrak{D}_1 \cup \mathfrak{D}_2$ and $\mathfrak{R} = \mathfrak{R}_1 \cup \mathfrak{R}_2$ and $\alpha_{i\{X_i/X\}}$ denotes α_i in which all occurrences of X_i are replaced by X .

Fig. 11. Merging fixpoint operation nodes.

^aThis is because the only transformation of the recursive part that needs to be propagated to update the annotations is the union of the two former recursive parts and both destab and rigid are distributive over union.




Fig. 12. RLQDAG structure after merging fixpoint operation nodes.

perform logical optimizations. The application of these rules can replace terms in an equivalence node by more optimal variants (i.e. they discard logically suboptimal variants).

B. The overall expansion algorithm

We can now describe the overall RLQDAG expansion algorithm. We define the function $\text{expand}()$ that takes an equivalence node γ and returns the equivalence node γ' which contains all the terms obtained by transformations.

The expand function is simply defined as follows:

$$\begin{aligned} \text{expand}([d]) &= \text{applyAll}([d]) \\ \text{expand}([d, \alpha]) &= \text{applyAll}([d]) \cup \text{expand}(\alpha) \end{aligned}$$

where $\text{applyAll}()$ is in charge of applying all possible transformations on each operation node. This includes applying all rewrite rules defined in Section IV in combination with the more classical ones of relational algebra⁴:

$$\begin{aligned} \text{applyAll}([d]) &= \text{pfj}([d]) \cup \text{paj}([d]) \cup \text{pj}([d]) \cup \text{mf}([d]) \\ &\quad \cup \text{pp}([d]) \cup \text{allCodd}([d]) \end{aligned}$$

where $\text{allCodd}()$ applies all rewrite rules concerning classical (non-recursive) relational algebra adapted for RLQDAG. For example: $\text{pfj}()$ for pushing filters in join operation nodes, $\text{paj}()$ for pushing antiprojections in a join operation node, $\text{jassoc}()$ (for join associativity), $\text{dju}()$ (for distributivity of join over union operation nodes) etc.

$$\begin{aligned} \text{allCodd}([d]) &= \text{pfj}([d]) \cup \text{paj}([d]) \cup \\ &\quad \text{jassoc}([d]) \cup \dots \cup \text{dju}([d]) \end{aligned}$$

For instance, $\text{pfj}()$ is defined as shown in Fig. 18 for an RLQDAG in which a filter precedes an equivalence node which contains a join operation node. In other cases, $\text{pfj}()$ recursively traverses the structure with appropriate calls to

⁴ As a slight discrepancy between the theory and the implementation, the implementation of the RLQDAG expansion avoids some redundant calls to the $\text{expand}()$ function by implementing the $\text{applyAll}()$ function in an imperative-style double nested for loop, so as to implement a single case analysis and to trigger $\text{expand}()$ once only when needed.

$\text{expand}()$ in search for further transformation opportunities. Other rewrite rules of non-recursive relational algebra are also implemented in a similar way in the RLQDAG.

C. Correctness and completeness

Proposition 1 (Correctness). *Let $[\alpha]$ be a consistent RLQDAG, and $\alpha' = \text{expand}([\alpha])$, then we have $S_\gamma \llbracket [\alpha] \rrbracket_\emptyset \subseteq S_\gamma \llbracket [\alpha'] \rrbracket_\emptyset$ and α' is consistent.*

Proof Sketch. The proof of correctness is done by induction on the structure of α' , using a case-by-case analysis of syntactic decompositions (i.e. by considering separately each RLQDAG subcase and the corresponding case of the generalized rewrite rule that applies to it). The complete formal proof is too large to be listed here (with many straightforward subcases). We summarize the main principles as well as useful auxiliary properties. When proving correctness for each RLQDAG rewrite rule, (i) we focus on each newly created recursive term added by the expansion, and (ii) use the consistency hypothesis on the recursive term before transformation to apply the theorems for individual terms of [35] to each rearrangement of fixpoint operation node. We also need to show that the incremental updates of annotations performed by the RLQDAG rewrite rules preserve consistency. This is done for each update by leveraging the properties that both $\text{destab}()$ and $\text{rigid}()$ are distributive over union. \square

Proposition 2 (Completeness properties). *Let R be a set of RLQDAG rewrite rules such that R contains the 5 RLQDAG rewrite rules for recursive terms presented in Section IV, we consider $[\alpha] = \text{unfold}(\text{expand}_R([\alpha']))$ where α' is a consistent RLQDAG, and $\text{expand}_R()$ is the $\text{expand}()$ function in which rules in R are activated. The following properties hold:*

Property 1 (No pushable filter left unpushed).

$$\forall \sigma_f(\gamma) \in [\alpha], \nexists d \in \gamma \mid d = \mu X. [\kappa] \cup [\alpha_2]_{\mathfrak{D}}^{\mathfrak{R}} \text{ and } \text{filt}(f) \subseteq \mathfrak{D}.$$

Property 2 (No pushable antiprojection left unpushed).

$$\forall \tilde{\pi}_a(\gamma) \in [\alpha], \nexists d \in \gamma \mid d = \mu X. [\kappa] \cup [\alpha_2]_{\mathfrak{D}}^{\mathfrak{R}} \text{ and } a \notin \mathfrak{R}.$$

Property 3 (All pushable joins have been pushed).

$$\forall (\beta \bowtie \gamma) \in [\alpha], \text{ if } \exists d \in \gamma \text{ such that } d = \mu X. [\kappa] \cup [\alpha_2]_{\mathfrak{D}}^{\mathfrak{R}} \text{ and } \text{type}(\beta) \cap \mathfrak{D} = \emptyset \text{ and } \text{type}(\beta) \setminus \text{type}(\gamma) \cap \mathfrak{R} = \emptyset \text{ then } \exists d' \in [\alpha] \text{ such that } d' = \mu X'. [[\beta] \bowtie [\kappa]] \cup [\alpha_2 \{X/X'\}]_{\mathfrak{D}}^{\mathfrak{R}}.$$

Property 4 (All mergeable fixpoints have been merged).

$$\forall (\gamma_1 \bowtie \gamma_2) \in [\alpha], \text{ if } \mu X_1. [\kappa_1] \cup [\alpha_1]_{\mathfrak{D}_1}^{\mathfrak{R}_1} \in \gamma_1 \text{ and } \mu X_2. [\kappa_2] \cup [\alpha_2]_{\mathfrak{D}_2}^{\mathfrak{R}_2} \in \gamma_2 \text{ and we have } \gamma_1 \neq \gamma_2 \text{ and } (\text{type}(\gamma_1) \cap \text{type}(\gamma_2)) \cap (\mathfrak{D}_1 \cup \mathfrak{D}_2) = \emptyset \text{ and } \text{type}(\gamma_1) \setminus \text{type}(\gamma_2) \cap \mathfrak{R}_2 = \emptyset \text{ and } \text{type}(\gamma_2) \setminus \text{type}(\gamma_1) \cap \mathfrak{R}_1 = \emptyset \text{ then there exists } d \in [\alpha] \text{ such that } d = \mu X. [[\kappa_1] \bowtie [\kappa_2]] \cup [[\alpha_1] \cup [\alpha_2]]_{\mathfrak{D}}^{\mathfrak{R}}.$$

Property 5 (All pushable antijoins have been pushed).

$$\forall ([\mu X. \gamma \cup [\alpha]_{\mathfrak{D}}^{\mathfrak{R}}] \triangleright \beta) \in [\alpha], \text{ if } \text{type}(\beta) \cap \mathfrak{D} = \emptyset \text{ then } \exists d \in [\alpha] \text{ such that } d = [\mu X'. [\gamma \triangleright \beta] \cup [\alpha \{X/X'\}]_{\mathfrak{D}}^{\mathfrak{R}}]$$

Proof Sketch. These properties are proved by contradiction: (i) assuming the existence of a missed transformation opportunity

$$\text{pp}([\tilde{\pi}_a([\mu X. \gamma \cup [\alpha]_{\mathcal{D}}^{\mathfrak{R}}])]) = \begin{cases} \bullet [\mu X'. \text{expand}([\tilde{\pi}_a(\gamma)]) \cup [\text{expand}(\alpha_{\{X/X'\}})]_{\mathfrak{D}'}^{\mathfrak{R}'}] \\ \text{when } a \notin \mathfrak{R} \\ \bullet [\tilde{\pi}_a(\text{expand}([\mu X. \gamma \cup [\alpha]_{\mathcal{D}}^{\mathfrak{R}}]))] \text{ otherwise} \end{cases}$$

$$\text{pp}([\tilde{\pi}_a([\mu X. \gamma \cup [\alpha]_{\mathcal{D}}^{\mathfrak{R}}, \alpha_2])]) = \begin{cases} \bullet [\mu X'. \text{expand}([\tilde{\pi}_a(\gamma)]) \cup [\text{expand}(\alpha_{\{X/X'\}})]_{\mathfrak{D}'}^{\mathfrak{R}'}], \\ \text{expand}([\tilde{\pi}_a([\alpha_2]])] \text{ when } a \notin \mathfrak{R} \\ \bullet [\tilde{\pi}_a(\text{expand}([\mu X. \gamma \cup [\alpha]_{\mathcal{D}}^{\mathfrak{R}}, \alpha_2])), \\ \text{expand}([\tilde{\pi}_a([\alpha_2]])] \text{ otherwise} \end{cases}$$

where $\mathfrak{D}' = \mathfrak{D} \cup \{a\}$ and $\mathfrak{R}' = \mathfrak{R} \cup \{a\}$.

Fig. 13. Pushing antiprojection in an equivalence node containing at least one fixpoint operation node.

$$\text{pa}([\mu X. \gamma \cup [\alpha]_{\mathcal{D}}^{\mathfrak{R}} \triangleright \beta]) = \begin{cases} \bullet \text{let const} = \gamma \text{ in} \\ \bullet [\text{expand}([\mu X. \text{const} \cup [\alpha]_{\mathcal{D}}^{\mathfrak{R}}] \triangleright \text{expand}(\beta)), \\ \mu X'. \text{expand}([\text{const} \triangleright \beta]) \cup [\text{expand}(\alpha_{\{X/X'\}})]_{\mathfrak{D}'}^{\mathfrak{R}'}] \\ \text{when } \text{type}(\beta) \cap \mathfrak{D} = \emptyset \\ \bullet [\text{expand}([\mu X. \gamma \cup [\alpha]_{\mathcal{D}}^{\mathfrak{R}}] \triangleright \text{expand}(\beta))] \text{ otherwise} \end{cases}$$

$$\text{pa}([\mu X. \gamma \cup [\alpha]_{\mathcal{D}}^{\mathfrak{R}}, \alpha_2 \triangleright \beta]) = \begin{cases} \bullet \text{let const} = \gamma \text{ in} \\ \bullet [\text{expand}([\mu X. \text{const} \cup [\alpha]_{\mathcal{D}}^{\mathfrak{R}}, \alpha_2] \triangleright \text{expand}(\beta)), \\ \mu X'. \text{expand}([\text{const} \triangleright \beta]) \cup [\text{expand}(\alpha_{\{X/X'\}})]_{\mathfrak{D}'}^{\mathfrak{R}'}], \\ [\text{expand}([\alpha_2] \triangleright \text{expand}(\beta))] \\ \text{when } \text{type}(\beta) \cap \mathfrak{D} = \emptyset \\ \bullet [\text{expand}([\mu X. \gamma \cup [\alpha]_{\mathcal{D}}^{\mathfrak{R}}, \alpha_2] \triangleright \text{expand}(\beta)), \\ [\text{expand}([\alpha_2] \triangleright \text{expand}(\beta))] \text{ otherwise} \end{cases}$$

where $\mathfrak{R}' = \mathfrak{R} \cup \text{rigid}(\beta, X)$.

Fig. 16. Pushing antijoin in an equivalence node containing a fixpoint operation node.

$$\text{pfj}([\sigma_f([\gamma_1 \bowtie \gamma_2, \alpha])]) = \begin{cases} \bullet [\exp([\sigma_f(\gamma_1)]) \bowtie \exp(\gamma_2), \exp([\sigma_f(\alpha)])] \\ \text{when } \text{filt}(f) \subseteq \text{type}(\gamma_1) \wedge \text{filt}(f) \not\subseteq \text{type}(\gamma_2) \\ \bullet [\exp(\gamma_1) \bowtie \exp([\sigma_f(\gamma_2)]), \exp([\sigma_f(\alpha)])] \\ \text{when } \text{filt}(f) \not\subseteq \text{type}(\gamma_1) \wedge \text{filt}(f) \subseteq \text{type}(\gamma_2) \\ \bullet [\exp([\sigma_f(\gamma_1)]) \bowtie \exp([\sigma_f(\gamma_2)]), \exp([\sigma_f(\alpha)])] \\ \text{when } \text{filt}(f) \subseteq \text{type}(\gamma_1) \wedge \text{filt}(f) \subseteq \text{type}(\gamma_2) \\ \bullet [\sigma_f(\exp([\varphi \bowtie \psi])), \exp([\sigma_f(\alpha)])] \text{ otherwise} \end{cases}$$

Fig. 18. Pushing a filter in an equivalence node composed of at least one join operation node and other operation nodes ($\text{exp}()$ stands for $\text{expand}()$).

in the expansion, which (ii) necessarily implies some unrealized rule application (whereas the rule was applicable), and (iii) showing that the systematic structure traversal performed by $\text{expand}_R()$ leaves no room for such a missed opportunity, thus (iv) leading to a contradiction.




Fig. 14. Initial RLQDAG before expansion with opportunity to apply $\text{pp}()$.

Fig. 15. Expansion of a RLQDAG by pushing an antiprojection in a fixpoint operation node. The part added by $\text{pp}()$ is in blue color.

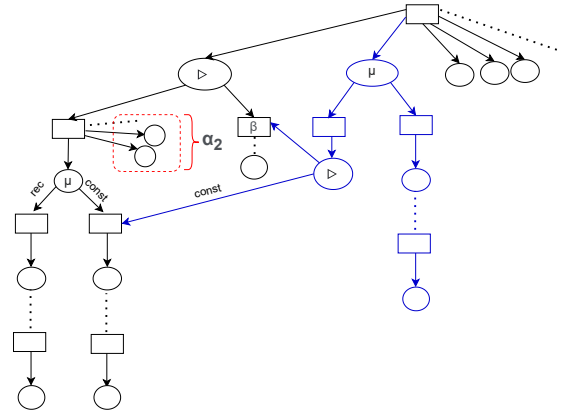


Fig. 17. Expansion of a RLQDAG by pushing an antijoin in a fixpoint operation node. The initial RLQDAG is in black color, and parts added by $\text{pa}()$ are in blue color.

V. EXPERIMENTAL RESULTS

In order to assess the efficiency of the RLQDAG approach in practice, we implemented a prototype of the RLQDAG in the μ -RA system [35], which is, as described in Section II, one of the most advanced relational-based system for recursive query optimization; and the system with the richest plan spaces for recursive terms, that are explored using a dynamic programming strategy [35]. We compare the performance of our RLQDAG prototype implementation with the performance of the unmodified μ -RA system. We conducted extensive experiments and report on the main lessons learned.

A. Experimental setup

a) *Datasets*: We consider various datasets, graphs and trees, real and generated, as described in Table I.

Dataset	#nodes & #edges	Type	Nature
Yago [62]	42M & 62M	Knowledge graph	Real
Bahamas Leaks [9]	202K & 249K	Property graph	Real
Airbnb [7]	24K & 14K	Property graph	Real
LDBC [12]	908K & 1.9M	Property graph	Synthetic
Wikitree [22]	9.1M & 1.3M	Tree	Real
Academic tree [39]	765K & 1.5M	Tree	Real

TABLE I
DATASETS (AVAILABLE FROM [58]).




Fig. 19. Yago queries.

b) *Queries*: We consider a variety of recursive queries formulated against these datasets. Queries for Yago are mainly third-party regular path queries already considered in earlier papers in the literature [4], [31], [35], [64], and chosen because they are representative of the variety of possible recursive optimizations that can apply to them. Queries over the Airbnb dataset are inspired from [55]. We added more queries formulated over the Bahamas and LDBC datasets. We also considered non-regular queries (variants of same generation, and $a^n b^n$) for the Wikitree and Academic Tree datasets. Queries and datasets used in experiments are available at [58].

B. Results

First, we measure the enumeration capability of the RLQDAG in terms of the number of plans explored per second for each query. Figures 19, 20, 21, 22, and Fig. 23 respectively show the results obtained for the queries over each dataset. In these figures, the y axis (in log scale) indicates the number of plans per second explored by each approach for a given query (on the x axis). Results suggest that the RLQDAG approach always enumerates plans much faster (up to two orders of magnitude) when compared to the μ -RA system⁵.

Now, we set a time budget t (in seconds) for the plan space exploration and let the two systems generate plan spaces for that time budget. This means that after t elapsed seconds we

⁵Cases where histogram bars look very similar correspond to situations where both systems completed the plan space exploration in a very short time duration, which makes the difference hardly visible with the log scale.

stop the two explorations and look at the portions of plan spaces obtained by the two systems. Fig. 24 shows the results obtained for a time budget of $t = 0.5$ seconds. The y axis (in log scale) indicates the number of plans found. Fig. 24 also indicates the size of the complete (exhaustive) plan space obtained without any time restriction for the plan exploration ($t = \infty$). For example, for query Q31, the complete plan space contains more than 21.4 million plans. In 0.5 seconds, the RLQDAG prototype explored 1,019,026 plans whereas the μ -RA system explored only 5,751 plans. This is because although μ -RA uses dynamic programming techniques, it is not capable of benefitting from the RLQDAG’s grouping effect when applying complex rewrite rules on sets of recursive terms at once, thus rules are significantly more costly to apply. Seen from another perspective, this means that the RLQDAG’s approach is more effective in avoiding redundant subcomputations. We have conducted extensive experiments and overall results indicate that, for a given time budget, the RLQDAG prototype explores many more terms in comparison to the μ -RA system, in all cases. In some cases, the RLQDAG generates a number of plans which is greater by up to two orders of magnitude (for the same time budget). Such a speedup sometimes enables a complete exploration of the whole plan space in some cases (as shown e.g. for Q32, Q51, Q53, Q54 in Fig. 24). Fig. 25 illustrates the time spent in evaluating the best estimated plan taken from each of the explored plan space portions of Fig. 24. Notice that the two systems use the same heuristics for cost estimations [38], thus making relevant a head to head comparison. Results are presented in Fig. 25. They show the benefits of exploring a much larger plan space: the RLQDAG approach always provides similar or better performance, which is a direct consequence of the availability of more efficient recursive plans in the larger explored portion of the plan space.

We now report on experiments of exploring plan spaces with varying and increasing time budgets for the same query. For instance, Fig. 26 presents the number of plans explored (on the y axis) for different time budgets shown on the x axis for query Q31 and query 53 (2 queries from 2 different datasets). Fig. 27 shows the time spent in evaluating the best estimated plan selected from each plan space explored with the time budget shown on the x axis. Again, results shown in Fig. 26 indicate that the RLQDAG explores significantly more plans than the other approach for all considered time budgets (the y axis is in log scale). We can also observe that the difference between the amount of plans explored by each system stays of the same order, even when exploration time increases.

Results shown in Fig. 27 illustrate the benefits of exploring larger plan spaces since the best estimated term selected from the larger plan space is indeed more efficiently evaluated. This confirms the interest of efficiently exploring recursive plan spaces, since this shows that, in practice, larger plan spaces are very prone to contain more efficient recursive plans.




Fig. 20. Bahamas Leaks queries.




Fig. 21. Airbnb queries.




Fig. 22. LDBC queries.

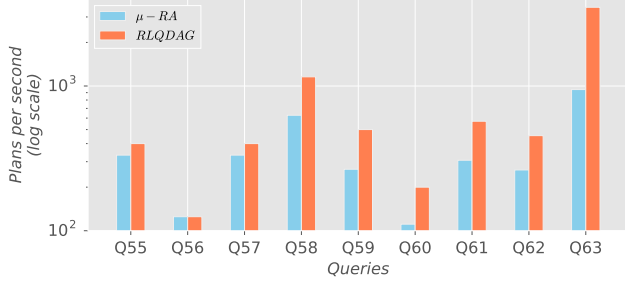


Fig. 23. Non-regular queries.

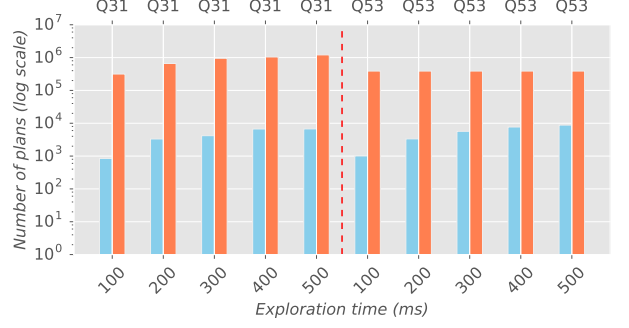


Fig. 26. Number of plans explored for different time budgets for Q31 and Q53.




Fig. 24. Plan spaces portions explored in 500ms.



Fig. 25. Evaluation of the best estimated plan found in the explored space portions for the given queries.

VI. CONCLUSION AND PERSPECTIVES

We propose the RLQDAG for capturing and efficiently transforming sets of recursive relational terms. This is done by introducing annotated equivalence nodes, and a formal syntax and semantics for RLQDAG terms that enable the development of RLQDAG rewrite rules on a solid ground.

RLQDAG rewrite rules transform sets of recursive terms while precisely describing how new subterms are created, attached, shared, and how new structural annotations are obtained with incremental updates.

The proposed formalisation of the RLQDAG in terms of syntax and semantics provided a convenient – if not instrumental – means to develop transformations. It helped in detecting and fixing intricate transformational issues (due to misplaced recursive calls) in an earlier version of the RLQDAG. We believe that this formalization can also serve in other and further (R)LQDAG extensions, thus contributing to the extensibility of the top-down transformational approach.

Practical experiments with a prototype implementation of the RLQDAG show that it efficiently generates the rich plan spaces introduced in the μ -recursive relational algebra. Experiments show significant performance gains compared to the state-of-the-art μ -RA system. This suggests that the RLQDAG represents a promising approach for the efficient enumeration of recursive relational query plans.

REFERENCES

- [1] *PODS'21: Proceedings of the 40th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, New York, NY, USA, 2021. Association for Computing Machinery.
- [2] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases: The Logical Level*. Addison-Wesley Longman Publishing Co., Inc., USA, 1st edition, 1995.
- [3] Z. Abul-Basher, N. Yakovets, P. Godfrey, S. Ghajar-Khosravi, and M. H. Chignell. Tasweet: optimizing disjunctive regular path queries in graph databases. In *EDBT/ICDT 2017 joint conference 20th international conference on extending database technology*. <https://doi.org/10.5441/002/edbt>, 2017.
- [4] Z. Abul-Basher, N. Yakovets, P. Godfrey, S. Ghajar-Khosravi, and M. H. Chignell. Tasweet: optimizing disjunctive regular path queries in graph databases. In *EDBT/ICDT 2017 joint conference 20th international conference on extending database technology*. <https://doi.org/10.5441/002/edbt>, 2017.
- [5] R. Agrawal. Alpha: an extension of relational algebra to express a class of recursive queries. *IEEE Transactions on Software Engineering*, 14(7):879–885, July 1988.
- [6] A. V. Aho and J. D. Ullman. Universality of data retrieval languages. In *Proceedings of the 6th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, POPL '79, pages 110–119, New York, NY, USA, 1979. ACM.
- [7] Airbnb. Airbnb. <http://insideairbnb.com/get-the-data.html>, 2022.
- [8] P. Alvaro, W. Marczak, N. Conway, J. Hellerstein, D. Maier, and R. Sears. Dedalus: Datalog in time and space. pages 262–281, 01 2010.
- [9] Bahamas-Leaks. Bahamas leaks. <https://www.kaggle.com/zusmani/paradisepanamapers>, 2016.
- [10] F. Bancilhon, D. Maier, Y. Sagiv, and J. D. Ullman. Magic sets and other strange ways to implement logic programs (extended abstract). In *Proceedings of the Fifth ACM SIGACT-SIGMOD Symposium on Principles of Database Systems*, PODS '86, page 115, New York, NY, USA, 1985. Association for Computing Machinery.
- [11] F. Bancilhon and R. Ramakrishnan. An amateur's introduction to recursive query processing strategies. In *Proceedings of the 1986 ACM SIGMOD International Conference on Management of Data*, SIGMOD '86, page 1652, New York, NY, USA, 1986. Association for Computing Machinery.
- [12] P. Boncz. Ldbc: Benchmarks for graph and rdf data management. In *Proceedings of the 17th International Database Engineering ; Applications Symposium*, IDEAS '13, page 12, New York, NY, USA, 2013. Association for Computing Machinery.
- [13] S. Chaudhuri. An overview of query optimization in relational systems. In *Proceedings of the Seventeenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, PODS '98, page 3443, New York, NY, USA, 1998. Association for Computing Machinery.
- [14] S. Chaudhuri and K. Shim. Including group-by in query optimization. In *VLDB*, volume 94, pages 12–15, 1994.
- [15] E. F. Codd. A relational model of data for large shared data banks. *Commun. ACM*, 13(6):377387, June 1970.
- [16] D. DeHaan and F. W. Tompa. Optimal top-down join enumeration. In *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, SIGMOD '07, page 785796, New York, NY, USA, 2007. Association for Computing Machinery.
- [17] Z. Fan, J. Zhu, Z. Zhang, A. Albarghouthi, P. Koutris, and J. M. Patel. Scaling-up in-memory datalog processing: Observations and techniques. *Proc. VLDB Endow.*, 12(6):695–708, 2019.
- [18] P. Fender and G. Moerkotte. A new, highly efficient, and easy to implement top-down join enumeration algorithm. pages 864 – 875, 05 2011.
- [19] P. Fender and G. Moerkotte. Counter strike: Generic top-down join enumeration for hypergraphs. *Proc. VLDB Endow.*, 6(14):18221833, sep 2013.
- [20] P. Fender and G. Moerkotte. Top down plan generation: From theory to practice. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*, pages 1105–1116, 2013.
- [21] P. Fender, G. Moerkotte, T. Neumann, and V. Leis. Effective and robust pruning for top-down join enumeration algorithms. In *2012 IEEE 28th International Conference on Data Engineering*, pages 414–425, 2012.
- [22] M. Fire and Y. Elovici. Data mining of online genealogy datasets for revealing lifespan patterns in human population. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 6(2):28, 2015.
- [23] M. Francis-Landau, T. Vieira, and J. Eisner. Evaluation of logic programs with built-ins and aggregation: A calculus for bag relations, 10 2020.
- [24] C. Galindo-Legaria and A. Rosenthal. How to extend a conventional optimizer to handle one-and two-sided outerjoin. In *1992 Eighth International Conference on Data Engineering*, pages 402–403. IEEE Computer Society, 1992.
- [25] C. Galindo-Legaria and A. Rosenthal. Outerjoin simplification and reordering for query optimization. *ACM Transactions on Database Systems (TODS)*, 22(1):43–74, 1997.
- [26] G. Gardarin. Magic functions: A technique to optimize extended datalog recursive programs. pages 21–30, 01 1987.
- [27] R. Goldman and J. Widom. Dataguides: Enabling query formulation and optimization in semistructured databases. Technical report, Stanford, 1997.
- [28] G. Graefe. The cascades framework for query optimization. *Data Engineering Bulletin*, 18, 1995.
- [29] G. Graefe and D. J. DeWitt. The exodus optimizer generator. In *Proceedings of the 1987 ACM SIGMOD International Conference on Management of Data*, SIGMOD '87, pages 160–172, New York, NY, USA, 1987. ACM.
- [30] G. Graefe and W. J. McKenna. The volcano optimizer generator: Extensibility and efficient search. In *Proceedings of the Ninth International Conference on Data Engineering*, pages 209–218, Washington, DC, USA, 1993. IEEE Computer Society.
- [31] A. Gubichev, S. J. Bedathur, and S. Seufert. Sparqling kleene: fast property paths in rdf-3x. In *First International Workshop on Graph Data Management Experiences and Systems*, pages 1–7, 2013.
- [32] L. M. Haas, J. C. Freytag, G. M. Lohman, and H. Pirahesh. Extensible query processing in starburst. In *Proceedings of the 1989 ACM SIGMOD International Conference on Management of Data*, SIGMOD 89, page 377388, New York, NY, USA, 1989. Association for Computing Machinery.
- [33] S. S. Huang, T. J. Green, and B. T. Loo. Datalog and emerging applications: An interactive tutorial. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*, SIGMOD '11, page 12131216, New York, NY, USA, 2011. Association for Computing Machinery.
- [34] T. Ibaraki and T. Kameda. On the optimal nesting order for computing n-relational joins. *ACM Transactions on Database Systems (TODS)*, 9(3):482–502, 1984.
- [35] L. Jachiet, P. Genevès, N. Gesbert, and N. Layaïda. On the optimization of recursive relational queries: Application to graph queries. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 681–697, 2020.
- [36] M. Kifer and E. L. Lozinskii. On compile-time query optimization in deductive databases by means of static filtering. *ACM Trans. Database Syst.*, 15(3):385426, sep 1990.
- [37] A. Koschmieder and U. Leser. Regular path queries on large graphs. In *International Conference on Scientific and Statistical Database Management*, pages 177–194. Springer, 2012.
- [38] M. Lawal, P. Genevès, and N. Layaïda. A cost estimation technique for recursive relational algebra. *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 2020.
- [39] J. F. Liénard, T. Achakulvisut, D. E. Acuna, and S. V. David. Intellectual synthesis in mentorship determines success in academic careers. *Nature communications*, 9(1):1–13, 2018.
- [40] G. M. Lohman. Grammar-like functional rules for representing query optimization alternatives. *SIGMOD Rec.*, 17(3):1827, June 1988.
- [41] G. M. Lohman, C. Mohan, L. M. Haas, D. Daniels, B. G. Lindsay, P. G. Selinger, and P. F. Wilms. *Query Processing in R**, pages 31–47. Springer Berlin Heidelberg, Berlin, Heidelberg, 1985.
- [42] G. Moerkotte, P. Fender, and M. Eich. On the correct and complete enumeration of the core search space. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, pages 493–504, 2013.
- [43] G. Moerkotte and T. Neumann. Analysis of two existing and one new dynamic programming algorithm for the generation of optimal bushy join trees without cross products. In *Proceedings of the 32nd International Conference on Very Large Data Bases*, VLDB '06, page 930941. VLDB Endowment, 2006.

- [44] G. Moerkotte and T. Neumann. Dynamic programming strikes back. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD '08, page 539552, New York, NY, USA, 2008. Association for Computing Machinery.
- [45] J. F. Naughton, R. Ramakrishnan, Y. Sagiv, and J. D. Ullman. Efficient evaluation of right-, left-, and multi-linear rules. In *Proceedings of the 1989 ACM SIGMOD International Conference on Management of Data*, SIGMOD '89, page 235242, New York, NY, USA, 1989. Association for Computing Machinery.
- [46] T. Neumann and B. Radke. Adaptive optimization of very large join queries. In *Proceedings of the 2018 International Conference on Management of Data*, SIGMOD '18, page 677692, New York, NY, USA, 2018. Association for Computing Machinery.
- [47] K. Ono and G. M. Lohman. Measuring the complexity of join enumeration in query optimization. In *Proceedings of the 16th International Conference on Very Large Data Bases*, VLDB '90, page 314325, San Francisco, CA, USA, 1990. Morgan Kaufmann Publishers Inc.
- [48] A. Pellenkoft, C. Galindo-Legaria, and M. Kersten. The complexity of transformation-based join enumeration. pages 306–315, 01 1997.
- [49] P. Roy, S. Seshadri, S. Sudarshan, and S. Bhobe. Efficient and extensible algorithms for multi query optimization. *SIGMOD Rec.*, 29(2):249260, May 2000.
- [50] D. Saccà and C. Zaniolo. On the implementation of a simple class of logic queries for databases. In *Proceedings of the Fifth ACM SIGACT-SIGMOD Symposium on Principles of Database Systems*, PODS '86, page 1623, New York, NY, USA, 1985. Association for Computing Machinery.
- [51] P. G. Selinger, M. M. Astrahan, D. D. Chamberlin, R. A. Lorie, and T. G. Price. Access path selection in a relational database management system. In *Proceedings of the 1979 ACM SIGMOD International Conference on Management of Data*, SIGMOD '79, pages 23–34, New York, NY, USA, 1979. ACM.
- [52] J. Seo, S. Guo, and M. S. Lam. Socialite: An efficient graph query language based on datalog. *IEEE Transactions on Knowledge and Data Engineering*, 27(7):1824–1837, 2015.
- [53] A. Shanbhag and S. Sudarshan. Optimizing join enumeration in transformation-based query optimizers. *Proc. VLDB Endow.*, 7(12):12431254, aug 2014.
- [54] L. Shapiro, D. Maier, P. Benninghoff, K. Billings, Y. Fan, K. Hatwal, Q. Wang, Y. Zhang, H.-M. Wu, and B. Vance. Exploiting upper and lower bounds in top-down query optimization. pages 20 – 33, 02 2001.
- [55] C. Sharma, R. Sinha, and K. Johnson. Practical and comprehensive formalisms for modelling contemporary graph query languages. *Information Systems*, 102:101816, 2021.
- [56] A. Shkapsky, M. Yang, M. Interlandi, H. Chiu, T. Condie, and C. Zaniolo. Big data analytics with datalog queries on spark. In *Proceedings of the 2016 International Conference on Management of Data*, SIGMOD '16, page 11351149, New York, NY, USA, 2016. Association for Computing Machinery.
- [57] K. T. Tekle and Y. A. Liu. More efficient datalog queries: Subsumptive tabling beats magic sets. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*, SIGMOD '11, page 661672, New York, NY, USA, 2011. Association for Computing Machinery.
- [58] Tyrex-repository. Datasets and queries used in experiments with the RLQDAG. <https://gitlab.inria.fr/tyrex-public/rlqdag>, 2022.
- [59] B. Vance and D. Maier. Rapid bushy join-order optimization with cartesian products. In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, SIGMOD '96, page 3546, New York, NY, USA, 1996. Association for Computing Machinery.
- [60] J. Wang, M. Balazinska, and D. Halperin. Asynchronous and fault-tolerant recursive datalog evaluation in shared-nothing engines. *Proc. VLDB Endow.*, 8(12):15421553, aug 2015.
- [61] Y. R. Wang, M. A. Khamis, H. Q. Ngo, R. Pichler, and D. Suciu. Optimizing recursive queries with program synthesis. *arXiv preprint arXiv:2202.10390*, 2022.
- [62] YAGO. Yago: A high-quality knowledge base. <https://www.mpi-inf.mpg.de/yago-naga/yago/>, july 2019.
- [63] N. Yakovets, P. Godfrey, and J. Gryz. WAVEGUIDE: Evaluating SPARQL property path queries. In *EDBT*, volume 2015, pages 525–528, 2015.
- [64] N. Yakovets, P. Godfrey, and J. Gryz. Waveguide: Evaluating SPARQL property path queries. In *EDBT*, volume 2015, pages 525–528, 2015.
- [65] J. young Chang and S. goo Lee. An optimization of disjunctive queries: union-pushdown. *Proceedings Twenty-First Annual International Computer Software and Applications Conference (COMPSAC'97)*, pages 356–361, 1997.
- [66] H. Zauner, B. Linse, T. Furche, and F. Bry. A RPL through RDF: Expressive navigation in RDF graphs. In *International Conference on Web Reasoning and Rule Systems*, pages 251–257. Springer, 2010.