

PUG: A Framework and Practical Implementation for Why & Why-Not Provenance

Seokki Lee · Bertram Ludäscher · Boris Glavic

Received: date / Accepted: date

Abstract Explaining why an answer is (or is not) returned by a query is important for many applications including auditing, debugging data and queries, and answering hypothetical questions about data. In this work, we present the first *practical* approach for answering such questions for queries with negation (first-order queries). Specifically, we introduce a graph-based provenance model that, while syntactic in nature, supports reverse reasoning and is proven to encode a wide range of provenance models from the literature. The implementation of this model in our PUG (Provenance Unification through Graphs) system takes a provenance question and Datalog query as an input and generates a Datalog program that computes an *explanation*, i.e., the part of the provenance that is relevant to answer the question. Furthermore, we demonstrate how a desirable factorization of provenance can be achieved by rewriting an input query. We experimentally evaluate our approach demonstrating its efficiency.

1 Introduction

Provenance for relational queries records how results of a query depend on the query's inputs. This type of information can be used to explain *why* (and *how*) a result is derived by a query over a given database. Recently, provenance-like techniques have been used to explain why a tuple (or a set of tuples described declaratively by a pattern) is *missing* from the query result (see [19]

✉ Seokki Lee, Boris Glavic
10 W 31st Street, Chicago, IL 60616
E-mail: slee195@hawk.iit.edu, E-mail: bglavic@iit.edu

Bertram Ludäscher
501 E. Daniel St, Champaign, IL 61820
E-mail: ludaesch@illinois.edu

$r_1 : Q(X, Y) :- \text{Train}(X, Z), \text{Train}(Z, Y), \neg \text{Train}(X, Y)$

Relation Train		
fromCity	toCity	$\mathbb{N}[X]$
seattle	seattle	p
seattle	chicago	q
chicago	seattle	r
washington dc	seattle	s
new york	washington dc	t
new york	chicago	u

Result of query Q	
X	Y
washington dc	chicago
new york	seattle
chicago	chicago

Fig. 1: Example train connection database and query

for a survey covering both provenance and missing answer techniques). However, the two problems have been treated mostly in isolation. Consider the following observation from [24]: asking why a tuple t is absent from the result of a query Q is equivalent to asking why t is present in $\neg Q$ (i.e., the complement of the result of Q wrt. the active domain). Thus, a unification of *why* and *why-not* provenance is naturally achieved by developing a provenance model for queries with negation. The approach for provenance and missing answers from [43] is based on the same observation.

In this paper, we introduce a graph model for provenance of first-order (FO) queries expressed as *non-recursive Datalog queries with negation*¹ (or *Datalog* for short) and an efficient method for explaining a (missing) answer using SQL. Our approach is based on the observation that typically only a part of the provenance, which we call *explanation* in this work, is actually relevant for answering the user's provenance question about the existence or absence of a result.

¹ or, equivalently, queries in full relational algebra (without aggregation), formulas in FO logic under the closed world assumption, and SPJUD-queries (select, project, join, union, difference).

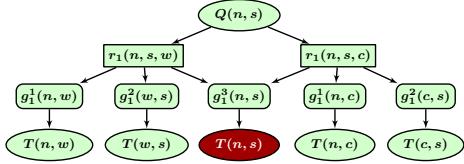


Fig. 2: Provenance graph explaining WHY $Q(n, s)$

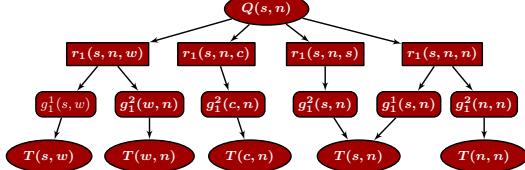


Fig. 3: Provenance graph explaining WHYNOT $Q(s, n)$

Example 1. Consider the relation `Train` in Fig. 1 that stores train connections. Datalog rule r_1 in Fig. 1 computes which cities can be reached with exactly one transfer, but not directly. We use the following abbreviations in provenance graphs: T = `Train`; n = New York; s = Seattle; w = Washington DC and c = Chicago. Given the result of this query, the user may be interested to know why he/she is able to reach Seattle from New York (WHY $Q(n, s)$) with one intermediate stop but not directly or why it is not possible to reach New York from Seattle in the same fashion (WHYNOT $Q(s, n)$).

An explanation for either type of question should justify the existence (absence) of a result as the success (failure) to derive the result through the rules of the query. Furthermore, it should explain how the existence (absence) of tuples in the database caused the derivation to succeed (fail). Provenance graphs providing this type of justification for WHY $Q(n, s)$ and WHYNOT $Q(s, n)$ are shown in Fig. 2 and Fig. 3, respectively. These graphs contain three types of nodes: *rule nodes* (boxes labeled with a rule identifier and the constant arguments of a rule derivation), *goal nodes* (rounded boxes labeled with a rule identifier and the goal's position in the rule's body), and *tuple nodes* (ovals). In these provenance graphs, nodes are either colored in *light green* (successful/existing) or *dark red* (failed/missing).

Example 2. Consider the explanation (provenance graph in Fig. 2) for question WHY $Q(n, s)$. Seattle can be reached from New York by stopping in Washington DC or Chicago and there is no direct connection between these two cities. These options correspond to two successful derivations using rule r_1 with $X=n$, $Y=s$, and $Z=w$ (or $Z=c$, respectively). In the provenance graph, there are two rule nodes representing these derivations of $Q(n, s)$ based on rule r_1 . A derivation is successful if all goals in the body evaluate to true, i.e., a successful rule node is connected to successful goal nodes (e.g., r_1 is connected to g_1^1 , the 1st goal in the rule's body). A positive (negated)

goal is successful if the corresponding tuple is (is not) in the database. Thus, a successful goal node is connected to the node corresponding to the existing (green) or missing (red) tuple justifying the goal, respectively.

Supporting negation and missing answers is challenging, because we need to enumerate all potential ways of deriving a missing answer (or intermediate result corresponding to a negated subgoal) and explain why each of these derivations has failed. For that, we have to decide how to bound the set of missing answers to be considered. Using the closed world assumption, only values that exist in the database or are postulated by the query are used to construct missing tuples. As is customary, we refer to this set of values as the active domain $adom(I)$ of a database instance I .

Example 3. Fig. 3 shows the explanation for WHYNOT $Q(s, n)$, i.e., why it is not true that New York is reachable from Seattle with exactly one transfer, but not directly. The tuple $Q(s, n)$ is missing from the result because all potential ways of deriving this tuple through rule r_1 have failed. In this example, $adom(I)=\{c, n, s, w\}$ and, thus, there exist four failed derivations of $Q(s, n)$ choosing either of these cities as the intermediate stop between Seattle and New York. A rule derivation fails if at least one goal in the body evaluates to false. Failed positive goals in the body of a failed rule are explained by missing tuples (red tuple nodes). For instance, we cannot reach New York from Seattle with an intermediate stop in Washington DC (the first failed rule derivation from the left in Fig. 3) because there exists no connection from Seattle to Washington DC (tuple node $T(s, w)$ in red), and Washington DC to New York (tuple node $T(w, n)$ in red). The successful goal $\neg T(s, n)$ (there is no direct connection from Seattle to New York) does not contribute to the failure of this derivation and, thus, is not part of the explanation.

Observe that nodes for missing tuples and successful rule derivations are conjunctive in nature (they depend on all their children) while existing tuples and failed rule derivations are disjunctive (they only require at least one of their children to be present).

Provenance Model. By recording which rule derivations justify the existence or absence of a query result, our model is suited well for debugging both data and queries. However, simpler provenance types, e.g., only tracking data dependencies, are sufficient for some applications. For example, assume that we record for each train connection from which webpage we retrieved information about this train connection. A user may be interested in knowing based on which webpages a query answer was computed. This question can be answered using a simpler provenance type called Lineage

(semiring $\text{Which}(X)$ [18]) which records the set of input tuples a result depends on. For such applications, we prefer simpler provenance types, because they are easier to interpret and more efficient to compute. Importantly, only minor modifications to our framework were required to support such provenance types.

Relationship to Other Provenance Models. In comparison to other provenance models, our model is more syntax-driven. We argue that this is actually a feature (not a bug). An important question is what is the semantic justification of our model, i.e., how do we know whether it correctly models Datalog query evaluation semantics and whether all (and only) relevant provenance is captured. First, we observe that our model encodes Datalog query semantics by construction. We justify that all relevant provenance is captured indirectly by demonstrating that our model captures sufficient information to derive provenance according to well-established models. Specifically, we demonstrate that our model is equivalent to provenance games [24] which also support FO queries. It was proven in [24] that provenance polynomials, the most general form of provenance in the semiring model [18, 22], for a result of a positive query can be “read out” from a provenance game. By being equivalent to provenance games, our provenance model also enjoys this property. We extend this result to queries with negation by relating our model to semiring provenance for FO model checking [13, 39, 43]. We prove that, for any FO formula φ , we can generate a query Q_φ such that the semiring provenance annotation of the formula $\pi(\varphi)$ according to a \mathcal{K} -interpretation π (annotation of positive and negated literals [13]) can be extracted efficiently from our provenance graph for Q_φ . Note that non-recursive Datalog queries with negation and FO formulas under the closed world assumption have the same expressive power and, thus, we use these languages interchangeably.

Reverse Reasoning (How-to Queries). For some applications, a user may not be interested in how a result was derived, but wants to understand how a result of interest can be achieved through updates to the database (see e.g., [30, 39, 31]). We extend our approach to support such “reverse reasoning” by introducing a third possible state of nodes in a provenance graph reserved for facts and derivations whose truth is undetermined. The provenance graph generated over an instance with undetermined facts represents a set of provenance graphs - one for each instance that is derived by assigning a truth value to each undetermined fact. We demonstrate that these graphs can be used to compute the semiring provenance of an FO formula under a provenance tracking interpretation as defined in [39].

Computing Explanations. We utilize Datalog to generate provenance graphs that explain a (missing) query result. Specifically, we instrument the input Datalog program to compute provenance bottom-up. Evaluated over an instance I , the instrumented program returns the edge relation of an explanation (provenance graph).

The main driver of our approach is a rewriting of Datalog rules (so-called *firing rules*) that captures successful and failed rule derivations. Firing rules for positive queries were first introduced in [23]. We have generalized this concept to negation and failed rule derivations. Firing rules provide sufficient information for constructing any of the provenance graph types we support. To make provenance capture efficient, we avoid capturing derivations that will not contribute to an explanation. We achieve this by propagating information from a user’s provenance question throughout the query to prune derivations that 1) do not agree with the constants of the question or 2) cannot be part of the explanation based on their success/failure status. For instance, in the running example, $Q(n, s)$ is only connected to derivations of rule r_1 with $X=n$ and $Y=s$.

We implemented our approach in *PUG* [26] (Provenance Unification through Graphs), an extension of our *GProM* [1] system. Using PUG, we compile rewritten Datalog programs into relational algebra, and translate such algebra expressions into SQL code that can be executed by a standard relational database backend.

Factorizing Provenance. Nodes in our provenance graphs are uniquely identified by their label. Thus, common subexpressions are shared leading to more compact graphs. For instance, observe that $g_1^3(n, s)$ in Fig. 2 is shared by two rule nodes. We exploit this fact by rewriting the input program to generate more concise, but equivalent, provenance. This is akin to factorization of provenance polynomials in the semiring model and utilizes factorized databases techniques [33, 34].

Contributions. This paper extends our previous work [26] in the following ways: we extend our model to support less informative, but more concise, provenance types; we extend our provenance model to support reverse reasoning [13] where the truth of some facts in the database is left undetermined; we demonstrate that our provenance graphs (explanations) are equivalent to provenance games [24] and how semiring provenance and its FO extension as presented in [39] can be extracted from our provenance model; we demonstrate how to rewrite an input program to generate a desirable (concise) factorization of provenance and evaluate the performance impact of this technique; finally, we present an experimental comparison with the language-integrated provenance techniques implemented in *Links* [9].

The remainder of this paper is organized as follows. We discuss related work in Sec. 2 and review Datalog in Sec. 3. We define our model in Sec. 4 and prove it to be equivalent to provenance games [24] in Sec. 5. We, then, show how our approach relates to semiring provenance for positive queries and FO model checking in Sec. 6 and 7, respectively. We present our approach for computing explanations in Sec. 8, and factorization in Sec. 9. Sec. 10 covers our implementation in *PUG* which we evaluate in Sec. 11. We conclude in Sec. 12.

2 Related Work

Our provenance graphs have strong connections to other provenance models for relational queries and to approaches for explaining missing answers.

Provenance Games. Provenance games [24] model the evaluation of a given query (input program) P over an instance I as a 2-player game in a way that resembles SLD(NF) resolution. By virtue of supporting negation, provenance games can uniformly answer why and why-not questions. We prove our approach to be equivalent to provenance games in Sec. 5. Köhler et al. [24] present an algorithm that computes the provenance game for a program P and database I . However, this approach requires instantiation of the full game graph (which enumerates all existing and missing tuples) and evaluation of a recursive Datalog[−] program over this graph using the well-founded semantics [10]. In contrast, our approach directly computes succinct explanations that contain only relevant provenance.

Database Provenance. Several provenance models for database queries have been introduced in related work, e.g., see [5, 22]. The semiring annotation framework generalizes these models for positive relational algebra (and, thus, positive non-recursive Datalog). An essential property of the \mathcal{K} -relational model is that the semiring of *provenance polynomials* $\mathbb{N}[X]$ generalizes all other semirings. It has been shown in [24] that provenance games generalize $\mathbb{N}[X]$ for positive queries. Since our graphs are equivalent to provenance games in the sense that there exist lossless transformations between both models (see Sec. 5), our graphs also encode $\mathbb{N}[X]$ and, thus, all other provenance models expressible as semirings (see Sec. 6.2). Provenance graphs which are similar to our graphs restricted to positive queries have been used as graph representations of semiring provenance (e.g., see [7, 8, 22]). Both our graphs and the boolean circuits representation of semiring provenance [8] explicitly share common subexpressions in the provenance. While these circuits support recursive queries, they do not support negation. Recently, extension of

circuits for semirings with monus (supporting set difference) have been discussed [38]. The semiring model has also been applied to record provenance of model checking for first-order (FO) logic formulas [39, 13, 43]. This work also supports missing answers using the observation made earlier in [24]. Support for negation relies on 1) translating formulas into negation normal form (*nnf*), i.e., pushing all negations down to literals, and 2) annotating both positive and negative literals using a separate set X and \bar{X} of indeterminates in provenance expressions where variables from X are reserved for positive literals and variables from \bar{X} for negated literals. This idea of using dual (positive and negative) indeterminates is an independent rediscovery of the approach from [6] which applied this idea for FO queries. The main differences between these approaches are 1) that the results from [6] were only shown for one particular semiring ($Bool(X \cup \bar{X})$, the semiring of boolean expressions over dual indeterminates) and 2) that [6] supports recursion in the form of well-founded Datalog and answer set programming (disjunctive Datalog). We prove that our model encompasses the model from [13]. The notion of causality is also closely related to provenance. Meliou et al. [29] compute causes for answers and non-answers. However, the approach requires the user to specify which missing inputs are considered as causes for a missing output. Roy et al. [36, 37] employ causality to compute explanations for high or low outcomes of aggregation queries as sets of input tuples which have a large impact on the result. Such sets of tuples are represented compactly through selection queries. A similar method was developed independently by Wu et al. [41].

Why-not and Missing Answers. Approaches for explaining missing answers are either based on the query [2, 3, 4, 40] (i.e., which operators filter out tuples that would have contributed to the missing answer) or based on the instance [20, 21] (i.e., what tuples need to be inserted into the database to turn the missing answer into an answer). The missing answer problem was first stated for query-based explanations in the seminal paper by Chapman et al. [4]. Huang et al. [21] first introduced an instance-based approach. Since then, several techniques have been developed to exclude spurious explanations, to support larger classes of queries [20], and to support distributed Datalog systems in Y! [42]. The approaches for instance-based explanations (with the exception of Y!) have in common that they treat the missing answer problem as a view update problem: the missing answer is a tuple that should be inserted into a view corresponding to the query and this insertion has to be translated as an insertion into the database instance. An explanation is then one particular solution to this view update problem. In contrast to these previous works,

our provenance graphs explain missing answers by enumerating all failed rule derivations that justify why the answer is not in the result. Thus, they are arguably a better fit for use cases such as debugging queries, where in addition to determining which missing inputs justify a missing answer, the user also needs to understand why derivations have failed. Furthermore, we do support queries with negation. Importantly, solutions for view update missing answer problems can be extracted from our provenance graphs. Thus, in a sense, provenance graphs with our approach generalize some of the previous approaches (for the class of queries supported, e.g., we do not support aggregation yet). Interestingly, recent work has shown that it may be possible to generate more concise summaries of provenance games [11, 35] and provenance graphs [28] that are particularly useful for negation and missing answers to deal with the potentially large size of the resulting provenance. Similarly, some missing answer approaches [20] use c-tables to compactly represent sets of missing answers. These approaches are complementary to our work.

Computing Provenance Declaratively. The concept of rewriting a Datalog program using firing rules to capture provenance as variable bindings of derivations was introduced by Köhler et al. [23]. They apply this idea for provenance-based debugging of positive Datalog. Firing rules are also similar to relational implementations of provenance capture in Perm [12], LogicBlox [16], Orchestra [17], and GProM [1]. Zhou et al. [44] leverage such rules for the distributed ExSPAN system using either full propagation or reference based provenance. The extension of firing rules for negation is the main enabler of our approach.

3 Datalog

A Datalog program P consists of a finite set of rules $r_i : R(\vec{X}) :- R_1(\vec{X}_1), \dots, R_n(\vec{X}_n)$ where \vec{X}_j denotes a tuple of variables and/or constants. We assume that the rules of a program are labeled r_1 to r_m . $R(\vec{X})$ is the *head* of the rule, denoted as $head(r_i)$, and $R_1(\vec{X}_1), \dots, R_n(\vec{X}_n)$ is the *body* (each $R_j(\vec{X}_j)$ is a *goal*). We use $vars(r_i)$ to denote the set of variables in r_i . In this paper, consider non-recursive Datalog with negation (FO queries), so goals $R_j(\vec{X}_j)$ in the body are *literals*, i.e., atoms $L(\vec{X}_j)$ or their negation $\neg L(\vec{X}_j)$, and recursion is not allowed. All rules r of a program have to be *safe*, i.e., every variable in r must occur positively in r 's body (thus, head variables and variables in negated goals must also occur in a positive goal). For example, Fig. 1 shows a Datalog query with a single rule r_1 . Here, $head(r_1)$ is $Q(X, Y)$ and $vars(r_1)$ is $\{X, Y, Z\}$. The rule is safe since the head

variables ($\{X, Y\}$) and the variables in the negated goal ($\{X, Y\}$) also occur positively in the body. The set of relations in the schema over which P is defined is referred to as the extensional database (EDB), while relations defined through rules in P form the intensional database (IDB), i.e., the IDB relations are those defined in the head of rules. We require that P has a distinguished IDB relation Q , called the *answer* relation. Given P and instance I , we use $P(I)$ to denote the result of P evaluated over I . Note that $P(I)$ includes the instance I , i.e., all EDB atoms that are true in I . For an EDB or IDB predicate R , we use $R(I)$ to denote the instance of R computed by P and $R(t) \in P(I)$ to denote that $t \in R(I)$ according to P .

We use $adom(I)$ to denote the active domain of instance I , i.e., the set of all constants that occur in I . Similarly, we use $adom(R.A)$ to denote the active domain of attribute A of relation R . In the following, we make use of the concept of a rule derivation. A *derivation* of a rule r is an assignment of variables in r to constants from $adom(I)$. For a rule with n variables, we use $r(c_1, \dots, c_n)$ to denote the derivation that is the result of binding $X_i = c_i$. We call a derivation *successful* wrt. an instance I if each atom in the body of the rule is true in I and *failed* otherwise.

4 Provenance Model

We now introduce our provenance model and formalize the problem addressed in this work: compute the subgraph of a provenance graph for a given query (input program) P and instance I that explains existence/absence of a tuple in/from the result of P .

4.1 Negation and Domains

To be able to explain why a tuple is missing, we have to enumerate all failed derivations of this tuple and, for each such derivation, explain why it failed. As mentioned in Sec. 1, we have to decide how to bound the set of missing answers. We propose a simple, yet general, solution by assuming that each attribute of an IDB or EDB relation has an associated domain.

Definition 1 (Domain Assignment). *Let $S = \{R_1, \dots, R_n\}$ be a database schema where each $R_i(A_1, \dots, A_m)$ is a relation. Given an instance I of S , a domain assignment dom is a function that associates with each attribute $R.A$ a domain of values. We require $dom(R.A) \supseteq adom(R.A)$.*

In our approach, the user specifies each $dom(R.A)$ as a query $dom_{R.A}$ that returns the set of admissible values for the domain of attribute $R.A$. These associated

domains fulfill two purposes: 1) to reduce the size of explanations and 2) to avoid semantically meaningless answers. For instance, if there exists another attribute `Price` in the relation `Train` in Fig. 1, then $adom(I)$ would also include all the values that appear in this attribute. Thus, some failed rule derivations for r_1 would assign prices as intermediate stops. Different attributes may represent the same type of entity (e.g., `fromCity` and `toCity` in our example) and, thus, it would make sense to use their combined domain values when constructing missing answers. For now, we leave it up to the user to specify attribute domains.

When defining provenance graphs in the following, we are only interested in rule derivations that use constants from the associated domains of attributes accessed by the rule. Given a rule r and variable X used in this rule, let $attrs(r, X)$ denote the set of attributes that variable X is bound to in the body of the rule. In Fig. 1, $attrs(r_1, Z) = \{\text{Train.fromCity}, \text{Train.toCity}\}$. We say a rule derivation $r(c_1, \dots, c_n)$ is *domain grounded* iff $c_i \in \bigcap_{A \in attrs(r, X_i)} \text{dom}(A)$ for all $i \in \{1, \dots, n\}$. For a relation $R(A_1, \dots, A_n)$, we use $\text{TUP}(R)$ to denote the set of all possible tuples for R , i.e., $\text{TUP}(R) = \text{dom}(R.A_1) \times \dots \times \text{dom}(R.A_n)$.

4.2 Provenance Graphs

Provenance graphs justify the existence (or absence) of a query result based on the success (or failure) to derive it using a query's rules. They also explain how the existence or absence of tuples in the database caused derivations to succeed or fail, respectively. Here, we present a constructive definition of provenance graphs that provide this type of justification. Nodes in these graphs carry two types of labels: 1) a label that determines the node type (tuple, rule, or goal) and additional information, e.g., the arguments and rule identifier of a derivation; 2) the success/failure status of nodes. Note that the first type of labels uniquely identifies nodes.

Definition 2 (Provenance Graph). *Let P be a Datalog program, I a database instance, dom a domain assignment for I , and \mathbb{L} the domain of strings. The provenance graph $\mathcal{PG}(P, I)$ is a graph $(V, E, \mathcal{L}, \mathcal{S})$ with nodes V , edges E , and node labelling functions $\mathcal{L} : V \rightarrow \mathbb{L}$ and $\mathcal{S} : V \rightarrow \{\text{True}, \text{False}\}$ (True for true/success and False for false/failure). We require that $\forall v, v' \in V : \mathcal{L}(v) = \mathcal{L}(v') \rightarrow v = v'$. The graph $\mathcal{PG}(P, I)$ is defined as follows:*

– **Tuple nodes:** For each n -ary EDB or IDB predicate R and tuple (c_1, \dots, c_n) of constants from the associated domains ($c_i \in \text{dom}(R.A_i)$), there exists a node v labeled $R(c_1, \dots, c_n)$. $\mathcal{S}(v) = \text{True}$ iff $R(c_1, \dots, c_n) \in P(I)$ and $\mathcal{S}(v) = \text{False}$ otherwise.

- **Rule nodes:** For every successful domain grounded derivation $r_i(c_1, \dots, c_n)$, there exists a node v in V labeled $r_i(c_1, \dots, c_n)$ with $\mathcal{S}(v) = \text{True}$. For every failed domain grounded derivation $r_i(c_1, \dots, c_n)$ where $\text{head}(r_i(c_1, \dots, c_n)) \notin P(I)$, there exists a node v as above but with $\mathcal{S}(v) = \text{False}$. In both cases, v is connected to the tuple node $\text{head}(r_i(c_1, \dots, c_n))$.
- **Goal nodes:** Let v be the node corresponding to a derivation $r_i(c_1, \dots, c_n)$ with m goals. If $\mathcal{S}(v) = \text{True}$, then for all $j \in \{1, \dots, m\}$, v is connected to a goal node v_j labeled g_i^j with $\mathcal{S}(v_j) = \text{True}$. If $\mathcal{S}(v) = \text{False}$, then for all $j \in \{1, \dots, m\}$, v is connected to a goal node v_j with $\mathcal{S}(v_j) = \text{False}$ if the j^{th} goal is failed in $r_i(c_1, \dots, c_n)$. Each goal is connected to the corresponding tuple node.

Our provenance graphs model query evaluation by construction. A tuple node $R(t)$ is successful in $\mathcal{PG}(P, I)$ iff $R(t) \in P(I)$. This is guaranteed, because each tuple built from values of the associated domain exists as a node v in the graph and its label $\mathcal{S}(v)$ is decided based on $R(t) \in P(I)$. Furthermore, there exists a successful rule node $r(\vec{c}) \in \mathcal{PG}(P, I)$ iff the derivation $r(\vec{c})$ succeeds for I . Likewise, a failed rule node $r(\vec{c})$ exists iff the derivation $r(\vec{c})$ is failed over I and $\text{head}(r(\vec{c})) \notin P(I)$. Fig. 2 and 3 show subgraphs of $\mathcal{PG}(P, I)$ for the query from Fig. 1. Since $Q(n, s) \in P(I)$ (Fig. 2), this tuple node is connected to all successful derivations with $Q(n, s)$ in the head which in turn are connected to goal nodes for each of the three goals of rule r_1 . $Q(s, n) \notin P(I)$ (Fig. 3) and, thus, its node is connected to all failed derivations with $Q(s, n)$ as a head. Here, we have assumed that all cities can be considered as start and end points of missing train connections, i.e., both $\text{dom}(\text{Train.fromCity})$ and $\text{dom}(\text{Train.toCity})$ are defined as $adom(\text{Train.fromCity}) \cup adom(\text{Train.toCity})$. Thus, we have considered derivations $r_1(s, n, Z)$ for $Z \in \{c, n, s, w\}$.

4.3 Provenance Questions and Explanations

Recall that the problem we address in this work is how to explain the existence or absence of (sets of) tuples using provenance graphs. Such a set of tuples specified as a pattern and paired with a qualifier (WHY /WHYNOT) is called a *provenance question* (PQ) in this paper. The two questions presented in Example 1 use constants only, but we also support provenance questions with variables, e.g., for a question WHYNOT $Q(n, X)$ we return all explanations for missing tuples where the first attribute is n , i.e., why it is not the case that a city X can be reached from New York with one transfer, but not directly. We say a tuple t' of constants *matches* a tuple t of variables and constants written as $t' \preceq t$ if

we can unify t' with t , i.e., we can equate t' with t by applying a valuation that substitutes variables in t with constants from t' .

Definition 3 (Provenance Question). *Let P be a query, I an instance, Q an IDB predicate, and dom a domain assignment for I . A provenance question ψ is of the form $\text{WHY } Q(t)$ or $\text{WHYNOT } Q(t)$ where $t = (v_1, \dots, v_n)$ consists of variables and domain constants $(\text{dom}(\text{Q.A}))$ for each attribute Q.A . We define:*

$$\begin{aligned}\text{PATTERN}(\psi) &= Q(t) \\ \text{MATCH}(\text{WHY } Q(t)) &= \{Q(t') \mid t' \in P(I) \wedge t' \preccurlyeq t\} \\ \text{MATCH}(\text{WHYNOT } Q(t)) &= \{Q(t') \mid t' \notin P(I) \wedge t' \preccurlyeq t \wedge t' \in \text{TUP}(Q)\}\end{aligned}$$

In Example 2 and 3, we have presented subgraphs of $\mathcal{PG}(P, I)$ as *explanations* for PQs, implicitly claiming that these subgraphs are sufficient for explaining these PQs. We now formally define this type of explanation.

Definition 4 (Explanation). *The explanation $\text{EXPL}(P, \psi, I, \text{dom})$ for a PQ ψ according to P , I , and dom is the subgraph of $\mathcal{PG}(P, I)$ containing only nodes that are connected to at least one node in $\text{MATCH}(\psi)$.*

In the following we will drop dom from $\text{EXPL}(P, \psi, I, \text{dom})$ if it is clear from the context or irrelevant for the discussion. Given this definition of explanation, note that 1) all nodes connected to a tuple node matching the PQ are relevant for computing this tuple and 2) only nodes connected to this node are relevant for the outcome. Consider $Q(t') \in \text{MATCH}(\psi)$ for a question $\text{WHY } Q(t)$. Since $Q(t') \in P(I)$, all successful derivations with head $Q(t')$ justify the existence of t' and these are precisely the rule nodes connected to $Q(t')$ in $\mathcal{PG}(P, I)$. For $\text{WHYNOT } Q(t)$ and matching $Q(t')$ we have $Q(t') \notin P(I)$ which is the case if all derivations with head $Q(t')$ have failed. In this case, all such derivations are connected to $Q(t')$ in the provenance graph. Each such derivation is connected to all of its failed goals which are responsible for the failure. Now, if a rule body references IDB predicates, then the same argument can be applied to reason that all rules directly connected to these tuples explain why they (do not) exist. Thus, by induction, the explanation contains all relevant tuple and rule nodes that explain the PQ.

5 Provenance Graphs and Provenance Games

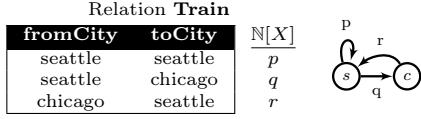
We now prove that provenance graphs according to Def. 2 are equivalent to provenance games. Thus, our model inherits the semantic foundation of provenance games. Specifically, provenance games were shown to encode Datalog query evaluation. Furthermore, the interpretation of provenance game graphs as 2-player games

provides a strong justification for why the nodes reachable from a tuple node justify the existence/absence of the tuple. We show how to transform a provenance game $\Gamma(P, \psi, I)$ into an explanation $\text{EXPL}(P, \psi, I)$ and vice versa to demonstrate that both are equivalent representations of provenance. We define a function $\text{TR}_{\Gamma \rightarrow \text{EXPL}}$ that maps provenance games to graphs and its inverse $\text{TR}_{\text{EXPL} \rightarrow \Gamma}$. Before that, we first give an overview of provenance games (see [24] for more details).

Provenance Games. Similar to our provenance graphs, provenance games are graphs that record successful and failed rule derivations. Provenance games consist of four types of nodes (e.g. Fig. 4d): rule nodes (boxes labeled with a rule identifier and the constant arguments of a rule derivation), goal nodes (boxes labeled with a rule identifier and the goal's position in the rule's body), tuple nodes (ovals), and EDB fact nodes (boxes labeled with an EDB relation name and the constants of a tuple). Every tuple node in a provenance game appears both positively and negatively, i.e., for every tuple node $R(t)$, there exists a tuple node $\neg R(t)$. Given a program P and database instance I , a provenance game is constructed by creating a positive and negative tuple node $R(c_1, \dots, c_n)$ for each n -ary predicate R and for all combinations of constants c_i from the active domain $\text{adom}(I)$. Similarly, nodes are created for rule derivations, i.e., a rule where variables have been replaced with constants from $\text{adom}(I)$ and each goal in the body of a rule (similar to Def. 2). In the game, a derivation of rule r for a vector of constants \vec{c} is labeled as $r(\vec{c})$, e.g., a derivation $\text{Q}_{\text{shop}}(s, s) := \text{T}(s, c), \text{T}(c, s), \text{T}(s, s)$ of r_2 in Fig. 4a is represented as a rule node labeled with $r_2(s, s, c, s)$. Finally, EDB fact nodes are added for each tuple in I , e.g., $r_T(s, s)$ for the tuple $(\text{seattle}, \text{seattle})$ in the Train relation (Fig. 4a). Tuple nodes are connected to the grounded rule nodes that derive them (have the tuple in their head), rule nodes to goal nodes for the grounded goals in their body, and goal nodes to negated tuple nodes corresponding to the goal (positive goals) or positive tuple nodes (negated goals). Such a game is interpreted as a 2-player game where the players argue for/against the existence of a tuple in the result of evaluating P over I . The existence of strategies for a player in this game determines tuple existence and success of rule derivations. A *solved game* is one where each node in the game graph is labeled as either won W (there exists a strategy for the player starting in this position) or lost L (no such strategy exists). A tuple node $R(t)$ is labeled as W iff the tuple $R(t)$ exists. A corollary of this is that a rule is labelled L if the corresponding derivation is successful and W otherwise.² Given such a solved

² This follows from the semantics of the type of 2-player game used here. The details are beyond the scope of this paper.

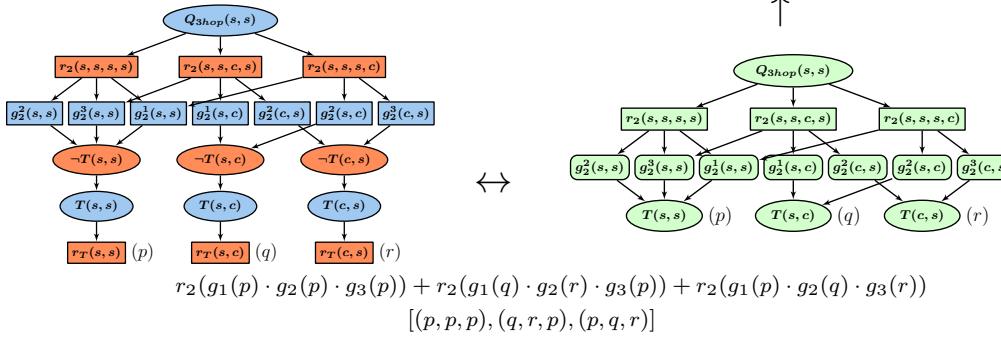
$$r_2 : Q_{3hop}(X, Y) := T(X, A), T(A, B), T(B, Y)$$



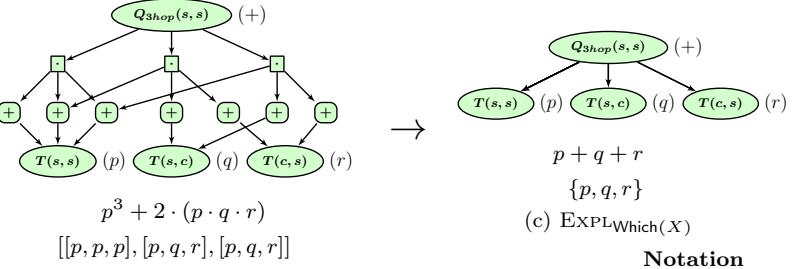
Result of query Q_{3hop}

X	Y	$\mathbb{N}[X]$
seattle	seattle	$p^3 + 2pqr$

(a) Example Train relation and query r_2



(d) Provenance game for $Q_{3hop}(s, s)$



(b) $EXPL_{\mathbb{N}[X]}$

$$p^3 + 2 \cdot (p \cdot q \cdot r)$$

$$[[p, p, p], [p, q, r], [p, q, r]]$$

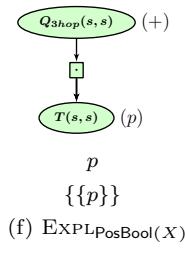
$$p + q + r$$

$$\{p, q, r\}$$

$$(c) EXPL_{\text{Which}(X)}$$

$$\text{Notation}$$

$$\begin{cases} \{ \} \text{ set} \\ [] \text{ bag} \\ () \text{ list} \end{cases}$$



(c) $EXPL_{\text{Which}(X)}$

$$\{p, q, r\}$$

$$(d) EXPL_{\text{PosBool}(X)}$$

$$\{ \{p\} \}$$

$$(e) EXPL_{\text{PosBool}(X)}$$

Fig. 4: Transformations exemplified using the provenance graph for $Q_{3hop}(s, s)$. For each graph, we show the structure of the provenance encoded by this graph and the corresponding semiring annotation where applicable.

game (denoted as $\Gamma(P, I)$), we can extract a subgraph rooted at an IDB tuple $Q(t)$ as the provenance of $Q(t)$. Similar to how we derive an explanation for a PQ with $\text{PATTERN}(\psi) = Q(t)$ where t may contain variables as the subgraph of the provenance graph $\mathcal{PG}(P, I)$ containing all IDB tuple nodes matching t and nodes reachable from these nodes, we can derive the corresponding subgraph in the provenance game $\Gamma(P, I)$ and denote it as $\Gamma(P, \psi, I)$ (we call such subgraphs *game explanations*).

Translating between Provenance Graphs and Provenance Games. The translation $\text{TR}_{\text{EXPL} \rightarrow \Gamma}$ of a provenance graph into the corresponding game and the reverse transformation $\text{TR}_{\Gamma \rightarrow \text{EXPL}}$ are straightforward. Thus, we only sketch $\text{TR}_{\text{EXPL} \rightarrow \Gamma}$ here. EDB tuple nodes are expanded to subgraphs $\neg R(t) \rightarrow R(t) \rightarrow r_R(t)$ for existing tuples and $\neg R(t) \rightarrow R(t)$ for missing tuples. IDB tuple nodes are always expanded to subgraphs of the later form. Rule and goal nodes and their interconnections are preserved. Goal nodes are connected to negated tuple nodes (positive goals) and to positive tuple nodes (negated goals). For positive tuple and goal nodes, we translate T to W (won) and F to L (lost). For negated tuple nodes and rule nodes, this mapping is reversed, i.e., T to L and F to W .

Theorem 1. Let P be a program, I a database instance, and ψ a PQ. We have:

$$\text{TR}_{\Gamma \rightarrow \text{EXPL}}(\Gamma(P, \psi, I)) = \text{EXPL}(P, \psi, I)$$

$$\text{TR}_{\text{EXPL} \rightarrow \Gamma}(\text{EXPL}(P, \psi, I)) = \Gamma(P, \psi, I)$$

Proof. See our accompanying technical report[27]. \square

Example 4. Consider rule r_2 in Fig. 4a computing which cities can be reached from another city through a path of length 3. The provenance game and provenance graph for $Q_{3hop}(s, s)$ are shown in Fig. 4d and Fig. 4e, respectively. In the provenance graph, goal nodes are directly connected to tuple nodes. In the game, they are represented as positive and negative tuple nodes and EDB fact nodes (the lower three levels). That is, every subgraph $\neg T(X, Y) \rightarrow T(X, Y) \rightarrow r_T(X, Y)$ in Fig. 4d is equivalently encoded as a single tuple node $T(X, Y)$ in Fig. 4e. Both graphs record the 3 paths of length 3 which start and end in Seattle: 1) $s \rightarrow s \rightarrow s \rightarrow s$, 2) $s \rightarrow c \rightarrow s \rightarrow s$, and 3) $s \rightarrow s \rightarrow c \rightarrow s$.

6 Semiring Provenance for Positive Queries

The semiring annotation model [22, 15, 18] is widely accepted as a provenance model for positive queries. An interesting question is how our model compares to provenance polynomials (semiring $\mathbb{N}[X]$), the most general form of annotation in the semiring model. It was shown in [24] that, for positive queries, the result of a query annotated with semiring $\mathbb{N}[X]$ can be extracted from the provenance game by applying a graph transformation. The equivalence shown in Sec. 5 extends this result to our provenance graph model. That being said, we develop simplified versions of our graph model to directly

support less informative provenance semirings such as Lineage which only tracks data-dependencies between input and output tuples. We now introduce the semiring annotation framework for positive queries and its use in provenance tracking and, then, explain our simplified provenance graph types.

6.1 \mathcal{K} -relations

In the semiring framework, relations are annotated with elements from a commutative semiring. A commutative semiring is a structure $\mathcal{K} = (K, +_{\mathcal{K}}, \cdot_{\mathcal{K}}, 0_{\mathcal{K}}, 1_{\mathcal{K}})$ over a set K where the addition and multiplication operations are associative, commutative, and have a neutral element ($0_{\mathcal{K}}$ and $1_{\mathcal{K}}$, respectively). Furthermore, multiplication with zero yields zero and multiplication distributes over addition. A relation annotated with the elements of a semiring \mathcal{K} is called a \mathcal{K} -relation. Operators of positive relational algebra (\mathcal{RA}^+) for \mathcal{K} -relations compute annotations for tuples in their output by combining annotations from their input using the operations of the semiring. Intuitively, multiplication represents conjunctive use of inputs (as in a join) whereas addition represents alternative use of inputs (as in a union or projection). We are interested in \mathcal{K} -relations, because it was shown that many provenance types can be expressed as semiring annotations.

Semiring homomorphisms are important for our purpose since they allow us to translate between different provenance semirings and understand their relative informativeness. A semiring homomorphism $h : \mathcal{K}_1 \rightarrow \mathcal{K}_2$ is a function from K_1 to K_2 that respects the operations of semirings, e.g., $h(k_1 +_{\mathcal{K}_1} k_2) = h(k_1) +_{\mathcal{K}_2} h(k_2)$. As shown in [14], if there exists a surjective homomorphism between one provenance semiring \mathcal{K}_1 and another semiring \mathcal{K}_2 , then \mathcal{K}_1 is more informative than \mathcal{K}_2 (see [18] for the technical details justifying this argument). We introduce several provenance semirings below and explain the homomorphisms that link the most informative semiring ($\mathbb{N}[X]$) to less informative semirings.

$(\mathbb{N}[X], +, \cdot, 0, 1)$: The elements of semiring $\mathbb{N}[X]$ are polynomials with natural number coefficients and exponents over a set of variables X representing tuples. Any polynomial can be written as a sum of products by applying the equational laws of semirings, e.g., the provenance polynomial for query result $Q_{3hop}(s, s)$ is $p^3 + 2pqr$ (Fig. 4a). An important property of $\mathbb{N}[X]$ is that there exist homomorphisms from $\mathbb{N}[X]$ to any other semiring.

$(PosBool(X), +, \cdot, 0, 1)$: The elements of $PosBool(X)$ are derived from $\mathbb{N}[X]$ by making both addition and multiplication idempotent and applying an additional equational law: $x + x \cdot y = x$. An element from $PosBool(X)$

can be encoded as a set of sets of variables with the restriction that every inner set k is minimal, i.e., there is no other inner set k' that is a subset of k . For example, the provenance polynomial $p^3 + 2pqr$ of $Q_{3hop}(s, s)$ is simplified as follows: $p^3 + 2pqr = p + pqr = p$.

$(Which(X), +, \cdot, 0, 1)$: In the $Which(X)$ semiring, addition is equivalent to multiplication: $x + y = x \cdot y$ for $x, y \notin \{0, 1\}$, and both addition and multiplication are idempotent. This semiring has sometimes also be called the Lineage semiring. Alternatively, the semiring can be defined over the powerset of the set of variables X [5].

Other semirings of interest are $(\mathbb{B}[X], +, \cdot, 0, 1)$ which is derived from $\mathbb{N}[X]$ by making addition idempotent ($x + x \equiv x$), semiring $(Trio(X), +, \cdot, 0, 1)$ where multiplication is idempotent ($x \cdot x \equiv x$), and $(Why(X), +, \cdot, 0, 1)$ where both addition and multiplication are idempotent.

6.2 \mathcal{K} -explanations

We now introduce simplified versions of our provenance graphs that each corresponds to a certain provenance semiring. Given a positive query P , $PQ \psi$, and database I , we use $EXPL_{\mathcal{K}}(P, \psi, I)$ to denote a \mathcal{K} -explanation for ψ . A \mathcal{K} -explanation is a provenance graph that encodes the \mathcal{K} -provenance of all query results from $MATCH(\psi)$, i.e., the set of answers the user is interested in. In the following, we first show how to extract $\mathbb{N}[X]$ from our provenance graph. Then, for each homomorphism implementing the derivation of a less informative provenance model from a more informative provenance model in the semiring framework, there is a corresponding graph transformation over our provenance graphs that maps $EXPL_{\mathbb{N}[X]}(P, \psi, I)$ to $EXPL_{\mathcal{K}}(P, \psi, I)$. The following theorem shows that we can reuse the existing mapping from provenance games to provenance polynomials by composing it with the mapping $TR_{EXPL \rightarrow \Gamma}$.

Theorem 2. *Let $TR_{EXPL \rightarrow \mathbb{N}[X]}$ denote the function $TR_{\Gamma \rightarrow \mathbb{N}[X]} \circ TR_{EXPL \rightarrow \Gamma}$. Given a positive input program P , database instance I , and tuple $t \in P(I)$, denote by $\mathbb{N}[X](P, I, t)$ the $\mathbb{N}[X]$ annotation of t over an abstractly tagged version of I (each tuple t is annotated with a unique variable x_t). Then,*

$$TR_{EXPL \rightarrow \mathbb{N}[X]}(EXPL(P, I, t)) = \mathbb{N}[X](P, I, t)$$

Proof. The proof is shown in [27]. \square

Consider the explanation for $WHY Q_{3hop}(s, s)$ shown in Fig. 4e. Recall that there are three options for reaching Seattle from Seattle with two intermediate stops corresponding to three derivations of Q_{3hop} using rule r_2 . These three derivations are shown in the provenance graph, e.g., $r_2(s, s, s, s)$ is the derivation that uses the

local train connection inside Seattle three times. Annotating the train connections with variables p , q , and r as shown in Fig. 4a and ignoring rule information encoded in the graph, the provenance encoded by our model is a bag (denoted as $[]$) of lists (denoted as $()$) of these variables. Each list corresponds to a rule derivation where variables are ordered according to the order of their occurrence in the body of the rule. For instance, (q, r, p) corresponds to taking a train from Seattle to Chicago (q), then from Chicago to Seattle (r), and finally a local connection inside Seattle (p). We now illustrate the graph transformations yielding \mathcal{K} -explanations from $\text{EXPL}_{\mathbb{N}[X]}$ based on this example.

Semiring $\mathbb{N}[X]$. In Fig. 4e, we (1) replace rule nodes with multiplication (i.e., $r_2(s, s, s, s) \rightarrow \cdot$) and (2) replace goal nodes with addition (e.g., $g_4^1(s, s) \rightarrow +$) to generate a graph that encodes $\mathbb{N}[X]$ as shown in Fig. 4b (denoted as $\text{EXPL}_{\mathbb{N}[X]}$). Applying this transformation, the rule instantiation $r_2(s, s, c, s)$ deriving result tuple $Q_{\text{shop}}(s, s)$ can no longer be distinguished from $r_2(s, s, s, c)$, because they are connected to the same tuple nodes. The only information retained is which arguments are used how often by a rule (labelled with \cdot). To extract $\mathbb{N}[X]$, we (1) replace labels of leaf nodes with their annotations from Fig. 4a (e.g., $T(s, s)$ is replaced with p) and (2) replace IDB tuple nodes with addition.

Semiring $\text{PosBool}(X)$. $\text{EXPL}_{\text{PosBool}(X)}$ (Fig. 4f) is computed from $\text{EXPL}_{\mathbb{N}[X]}$ by first collapsing rule nodes if the subgraphs rooted at these rule nodes are isomorphic and dropping all the goal nodes. Then, “.” nodes are removed if one or more subgraphs rooted at children of such a node is isomorphic to the subgraphs rooted at the children of another “.” node. Applying this process to our example, after the first step, two “.” nodes (one connects $Q_{\text{shop}}(s, s)$ to p and the other for each p , q , and r) exist in the graph corresponding to p and $p \cdot q \cdot r$. In the second step, $p \cdot q \cdot r$ is removed because it contains p ($T(s, s)$) as a subgraph.

Semiring $\text{Which}(X)$. The semiring $\text{Which}(X)$ (aka Lineage) is reached by collapsing all intermediate nodes and directly connecting tuple nodes (e.g., $Q_{\text{shop}}(s, s)$) with other tuple nodes (e.g., $T(s, s)$) as shown in Fig. 4c.

$\text{EXPL}_{\mathbb{B}[X]}$ and $\text{EXPL}_{\text{Trio}(X)}$ are derived from $\text{EXPL}_{\mathbb{N}[X]}$ by collapsing isomorphic subgraphs rooted at rule nodes and by dropping all the goal nodes, respectively. The combination of these transformations achieves $\text{EXPL}_{\text{Why}(X)}$.

7 Semiring Provenance for FO Model Checking

The semiring framework was recently extended for capturing provenance of first-order (FO) model checking [39, 13]. We now study the relationship of our model to

semiring provenance for FO queries. Based on the observation first stated in [24] (provenance for FO queries and, thus, also FO logic, naturally supports missing answers), the authors explain missing answers based on FO provenance [43]. Another interesting aspect of [13] is that it allows some facts to be left undetermined (their truth is undecided). This enables how-to queries [31], i.e., given an expected outcome, which possible world compatible with the undecided facts would produce this outcome. In this section, we first introduce the model from [13], then demonstrate how our approach can be extended to support undetermined truth values. Finally, we show how the annotation computed by the approach presented in [13] for a FO formula φ can be efficiently extracted from the provenance graph generated by our approach for a query Q_φ which is derived from φ through a translation $\text{TL}_{\varphi \rightarrow Q}$.

7.1 K-Interpretations and Dual Polynomials

In [39, 18], the authors define semiring provenance for formulas in FO logic. Let A be a domain of values. We use ν to denote an assignment of the free variables of φ to values from A . Given a so-called \mathcal{K} -interpretation π which is a function mapping positive and negative literals to annotations from \mathcal{K} , the annotation of a formula $\pi[\![\varphi]\!]_\nu$ for a given valuation ν is derived using the rules below. For sentences, i.e., formulas without free variables, we omit the valuation and write $\pi(\varphi)$ to denote $\pi[\![\varphi]\!]_\nu$ for the empty valuation ν . Furthermore, op is used to denote a comparison operator (either $=$ or \neq).

$$\begin{aligned} \pi[\![R(\mathbf{x})]\!]_\nu &= \pi(R(\nu(\mathbf{x}))) & \pi[\![\neg R(\mathbf{x})]\!]_\nu &= \pi(\neg R(\nu(\mathbf{x}))) \\ \pi[\![x \text{ op } y]\!]_\nu &= \text{if } \nu(x) \text{ op } \nu(y) \text{ then } 1 \text{ else } 0 & \pi[\![\neg\varphi]\!]_\nu &= \pi[\![\text{nnf}(\varphi)]]\!]_\nu \\ \pi[\![\varphi_1 \vee \varphi_2]\!]_\nu &= \pi[\![\varphi_1]\!]_\nu + \pi[\![\varphi_2]\!]_\nu & \pi[\![\varphi_1 \wedge \varphi_2]\!]_\nu &= \pi[\![\varphi_1]\!]_\nu \cdot \pi[\![\varphi_2]\!]_\nu \\ \pi[\![\exists x \varphi]\!]_\nu &= \sum_{a \in A} \pi[\![\varphi]\!]_{\nu[x \mapsto a]} & \pi[\![\forall x \varphi]\!]_\nu &= \prod_{a \in A} \pi[\![\varphi]\!]_{\nu[x \mapsto a]} \end{aligned}$$

Both conjunction and universal quantification correspond to multiplication, and annotations of positive and negative literals are read from π . This model deals with negation as follows. A negated formula is first translated into negation normal form (nnf). A formula in nnf does not contain negation except for negated literals. Any formula can be translated into this form using DeMorgan rules, e.g., $\neg(\forall x \varphi) \equiv (\exists x \neg \varphi)$. By pushing negation to the literal level using nnf , and annotating both positive and negative literals, the approach avoids extending the semiring structure with an explicit negation operation.

Provenance tracking for FO formula has to take into account the dual nature of the literals. The solution presented in [13, 39] is to use polynomials over two sets of variables: variables from X and \bar{X} are exclusively used

to annotate positive and negative literals, respectively. For any variable $x \in X$, there exists a corresponding variable $\bar{x} \in \bar{X}$ and vice versa. Furthermore, if x annotates $R(\mathbf{a})$, then \bar{x} can only annotate $\neg R(\mathbf{a})$ (and vice versa). The semiring of dual indeterminate polynomials is then defined as the structure generated by applying the congruence $x \cdot \bar{x} = 0$ to the polynomials from $\mathbb{N}[X \cup \bar{X}]$. The resulting structure is denoted by $\mathbb{N}[X, \bar{X}]$. Intuitively, this congruence encodes the standard logic equivalence $R(\mathbf{a}) \wedge \neg R(\mathbf{a}) \equiv \text{false}$. Importantly, in a $\mathbb{N}[X, \bar{X}]$ -interpretation π , we can decide which facts are true/false and whether to track provenance for these facts. Furthermore, we can leave the truth of some literals undetermined. Below, we show all feasible combinations for annotating $R(\mathbf{a})$ and $\neg R(\mathbf{a})$ in π and their meaning. For instance, if we annotate $R(\mathbf{a})$ with 1 or 0, this corresponds to asserting the fact $R(\mathbf{a})$, but not tracking provenance for it. By setting $R(\mathbf{a}) = x$ and $\neg R(\mathbf{a}) = \bar{x}$, we leave the truth of $R(\mathbf{a})$ undecided. Note that $R(\mathbf{a}) = 0$ and $\neg R(\mathbf{a}) = 0$ (as well as $R(\mathbf{a}) = 1$ and $\neg R(\mathbf{a}) = 1$) are not considered here since they lead to incompleteness (inconsistency).

$$\begin{aligned}
 \pi(R(\mathbf{a})) = 1 & \quad \pi(\neg R(\mathbf{a})) = 0 & (\text{true, no provenance}) \\
 \pi(R(\mathbf{a})) = 0 & \quad \pi(\neg R(\mathbf{a})) = 1 & (\text{false, no provenance}) \\
 \pi(R(\mathbf{a})) = x & \quad \pi(\neg R(\mathbf{a})) = 0 & (\text{true, track provenance}) \\
 \pi(R(\mathbf{a})) = 0 & \quad \pi(\neg R(\mathbf{a})) = \bar{x} & (\text{false, track provenance}) \\
 \pi(R(\mathbf{a})) = x & \quad \pi(\neg R(\mathbf{a})) = \bar{x} & (\text{undetermined})
 \end{aligned}$$

Consider a sentence φ .³ The annotation $\pi(\varphi)$ computed for φ over π with undetermined facts represents a set of possible models for φ . By choosing for each undetermined fact $R(\mathbf{a})$ in $\pi(\varphi)$ whether it is true or not, we “instantiate” one possible model for φ . By encoding a set of possible models, $\pi(\varphi)$ allows for reverse reasoning: we can find models that fulfill certain properties from the set of models encoded by $\pi(\varphi)$.

Example 5. Reconsider query r_1 from Fig. 1. Assume that we want to determine what effect building a direct train connection from New York to Seattle would have on the query result $Q(n, s)$. Thus, we make the assumption that the database instance is as in Fig. 1 with the exception that we keep $T(n, s)$ undetermined. In first-order logic, r_1 is expressed as: $\text{only2hop}(x, y) \equiv \exists z(T(x, z) \wedge T(z, y) \wedge \neg T(x, y))$ and fact $Q(n, s)$ as: $\varphi \equiv \text{only2hop}(n, s)$. The database when keeping $T(n, s)$ undetermined is encoded as a $\mathbb{N}[X, \bar{X}]$ -interpretations π which assigns variables to positive literals as shown in Fig. 1 (the corresponding negated literals are annotated with 0). $\pi(T(n, s)) = v$, $\pi(\neg T(n, s)) = \bar{v}$, and we annotate all remaining positive literals with 0 and negative

³ We only restrict the discussion to sentences for simplicity. The arguments here also hold for formulas with free variables.

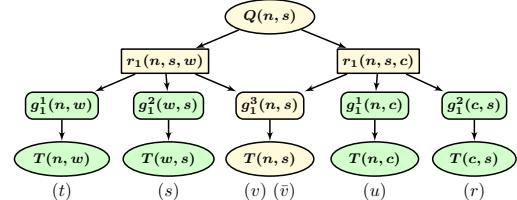


Fig. 5: Provenance graph for WHY $Q(n, s)$ when $T(n, s)$ is left undetermined

literals with 1. Computing $\pi(\varphi)$ using the rules above, we get $(t \cdot s \cdot \bar{v}) + (u \cdot r \cdot \bar{v})$. There are two ways of deriving the query result $Q(n, s)$ which both depend on the absence of a direct train connection from New York to Seattle (\bar{v}). Now if we decide to introduce such a connection, we can evaluate the effect of this choice by setting $\bar{v} = 0$ in the provenance polynomial above (the absence of this connection has been refuted), i.e., we get $(t \cdot s \cdot 0) + (u \cdot r \cdot 0) = 0$. Thus, if we were to introduce such a connection, then $Q(n, s)$ would no longer be a result.

7.2 Supporting Undeterminism in Provenance Graphs

Supporting undetermined facts in our provenance model is surprisingly straightforward. We introduce a new label U which is used to label nodes whose success/failure (existence/absence) is undetermined. To account for this new label, we amend the rules for determining connectivity and node labeling as follows:

- For a goal node v_g (no matter whether positive or negative) that is connected to a tuple node v_t with $\mathcal{S}(v_t) = \text{U}$, we set $\mathcal{S}(v_g) = \text{U}$ (goals corresponding to undetermined tuples are undetermined).
- A rule node is successful (G) if all its goals are successful, a rule node is failed if some of its goals are failed (F), and finally a rule node is undetermined (U) if at least one of its goals is undetermined and none of its goals are failed. Successful rule nodes are connected to all goals, failed rule nodes to failed and undetermined goals (these may provide further justification for the failure), and undetermined rule nodes to all goals (successful and undetermined goals will determine success of the rule nodes under choices).
- An IDB tuple exists (G) if at least one of its rule derivations is successful. It is connected to all successful and undetermined rule derivations (these may provide additional justifications under certain choices for undetermined facts). An IDB tuple is absent (F) if all of its rule derivations fail. Finally, an IDB tuple is undetermined (U) if at least one of its rule derivations is undetermined and none is successful.

$$\begin{array}{c}
\frac{\varphi := \exists x : \varphi_1}{Q_\varphi(\text{FREE}(\varphi)) := \text{Dom}(x), Q_{\varphi_1}(\text{FREE}(\varphi_1))} \quad (1) \quad \frac{\varphi := \neg R(\mathbf{x}), \text{FREE}(\varphi) = \{x_1, \dots, x_n\}}{Q_\varphi(\text{FREE}(\varphi)) := \text{Dom}(x_1), \dots, \text{Dom}(x_n), \neg R(\mathbf{x})} \quad (2) \quad \frac{\varphi := R(\mathbf{x})}{Q_\varphi(\text{FREE}(\varphi)) := R(\mathbf{x})} \quad (3) \\
\frac{\varphi := \varphi_1 \vee \varphi_2, \text{FREE}(\varphi_1) = \{x_1, \dots, x_n, y_1, \dots, y_m\}, \text{FREE}(\varphi_2) = \{x_1, \dots, x_n, z_1, \dots, z_l\}}{Q_\varphi(\text{FREE}(\varphi)) := \text{Dom}(z_1), \dots, \text{Dom}(z_k), Q_{\varphi_1}(\text{FREE}(\varphi_1)), Q_{\varphi_2}(\text{FREE}(\varphi_2))} \quad (4) \quad \frac{\varphi := x \mathbf{op} y}{Q_\varphi(x, y) := \text{Dom}(x), \text{Dom}(y), x \mathbf{op} y} \quad (5) \\
\frac{\varphi := \forall x : \varphi_1, \text{FREE}(\varphi) = \{x_1, \dots, x_n\}}{Q_\varphi(\text{FREE}(\varphi)) := \text{Dom}(x_1), \dots, \text{Dom}(x_k), \neg Q_{\varphi_1}(\text{FREE}(\varphi))} \quad (6) \quad \frac{\varphi := \varphi_1 \wedge \varphi_2}{Q_\varphi(\text{FREE}(\varphi)) := Q_{\varphi_1}(\text{FREE}(\varphi_1)), Q_{\varphi_2}(\text{FREE}(\varphi_2))} \quad (7)
\end{array}$$

Fig. 6: Translating a first-order formula φ into a first-order query Q_φ

Undetermined tuple nodes are connected to all their rule derivations (failed ones may be additional justifications for absence while undetermined ones may justify either existence or absence).

Example 6. Consider Example 5 in our extended provenance graph model. Let $v_{n,s}$ be the node corresponding to $\text{T}(n, s)$. If we set $\mathcal{S}(v_{n,s}) = \text{U}$ to indicate that $\text{T}(n, s)$ should be considered as undetermined, then we get the provenance graph in Fig. 5. Our approach correctly determines that under this assumption the truth of $\text{Q}(n, s)$ is undetermined and that there are two potential derivations of this result which also are undetermined, because they depend on existing tuples as well as the undetermined tuple $\text{T}(n, s)$. To evaluate the effect of choosing $\text{T}(n, s)$ to be true or false, we would set $\mathcal{S}(v_{n,s}) = \text{T}$ or $\mathcal{S}(v_{n,s}) = \text{F}$ and propagate the effect of this change bottom-up through-out the provenance graph.

Importantly, the provenance graph captured for an instance with undetermined facts is sufficient for evaluating the effect of setting any of these undetermined facts to false or true. That is, just like it is not necessary to reevaluate the semiring annotation of a formula to evaluate the impact of such a choice, it is also not necessary to recapture provenance in our model to evaluate a choice. For lack of space, we are not discussing the details of a corresponding extension for provenance games, but still would like to remark that undetermined facts correspond to draws in the game (neither player has a winning strategy). In the type of two-player games employed in provenance games, draws are caused by cycles in the game graph. To leave the existence of an EDB tuple undetermined, we introduce an EDB fact node for the tuple and add a self-edge to this node which causes the tuple node to be a draw in the game.

7.3 From First-order Formulas to Datalog

We now present a translation $\text{TL}_{\varphi \rightarrow Q}$ from FO formulas φ to boolean Datalog queries Q_φ . The query generated based on a formula φ is equivalent to the formula in the following sense: if φ evaluates to true for a model, then

$Q_\varphi(I)$ returns true. Here, I is the instance that contains precisely the tuples corresponding to literals that are true in the model. We assume that the free variables of a formula (the variables not bound by any quantifier) are distinct from variable names bound by quantifiers and that no two quantifiers bind a variable of the same name. This can be achieved by renaming variables in a formula that does not fulfill this condition. For example, $(\forall x R(x, y)) \wedge (\exists y S(y))$ does not fulfill this condition, but the equivalent formula $(\forall x R(x, y)) \wedge (\exists z S(z))$ does. We also assume an arbitrary, but fixed, total order $<_{Var}$ over variables that appear in formulas. We use $\text{FREE}(\varphi)$ to denote the list of free variables of a formula φ ordered increasingly by $<_{Var}$. For instance, for $\varphi := \forall x : R(x, y)$ we have $\text{FREE}(\varphi) = \{y\}$. Our translation $\text{TL}_{\varphi \rightarrow Q}$ takes as input a formula φ and outputs a Datalog program with an answer predicate Q_φ . The translation rules are shown in Fig. 6. Each rule translates one construct (e.g., a quantifier) and outputs one or more Datalog rules. The Datalog program generated by the translation for an input φ is the set of Datalog rules generated by applying the rules from Fig. 6 to all sub-formulas of φ . Here, we assume the existence of a unary predicate Dom whose extension is the domain A . Most translation rules are straightforward and standard. Logical operators are translated into their obvious counterpart in Datalog, e.g., a conjunction $\varphi_1 \wedge \varphi_2$ is translated into a rule with two body atoms Q_{φ_1} and Q_{φ_2} . The rules generated for a formula φ return the formula's free variables to make them available to formulas that use φ . For instance, since Datalog does not support universal quantification directly, we have to simulate it using double negation ($\forall x \varphi$ is rewritten as $\neg \exists x \neg \varphi$). Disjunctions are turned into unions. The complexity of the rule for disjunction stems from the fact that, in $\varphi_1 \vee \varphi_2$, the sets of free variables for φ_1 and φ_2 may not be the same. To make them union compatible, use $\text{FREE}(\varphi)$ as the arguments of the heads of the rules for both φ_1 and φ_2 , and add additional goals Dom to ensure that these rules are safe.

Example 7. Consider a directed graph encoded as its edge relation R . The formula $\varphi := \forall x \exists y \text{R}(x, y)$ checks

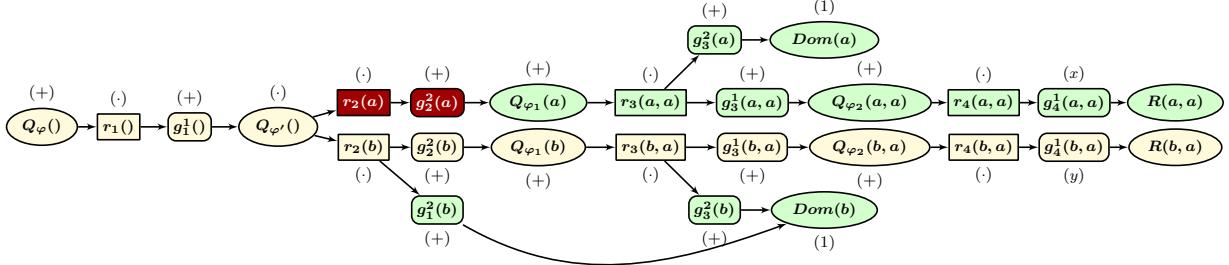


Fig. 7: Provenance graph for query Q_φ based on $\varphi := \forall x \exists y R(x, y)$ when $R(a, a)$ is true ($\pi(R(a, a)) = x$) and $\pi(\neg R(a, a)) = 0$ and $R(b, a)$ is left undetermined ($\pi(R(b, a)) = y$ and $\pi(\neg R(b, a)) = \bar{y}$). Note that $\varphi_1 := \exists y R(x, y)$ and $\varphi_2 := R(x, y)$. The result of $\text{TR}_{\text{EXPL} \rightarrow \mathbb{N}[X, \bar{X}]}$ shown besides the nodes encodes the dual polynomial $\pi(\varphi) = x \cdot y$.

whether all nodes in the graph have outgoing edges. Let $\varphi_1 := \exists y R(x, y)$ and $\varphi_2 := R(x, y)$. Translating this formula, we get:

$$\begin{aligned} Q_\varphi() &:= \neg Q_{\varphi_1}() \quad Q_{\varphi_1}() &:= \text{Dom}(x), \neg Q_{\varphi_1}(x) \\ Q_{\varphi_1}(x) &:= \text{Dom}(y), Q_{\varphi_2}(x, y) \quad Q_{\varphi_2}(x, y) &:= R(x, y) \end{aligned}$$

7.4 From Graphs to FO Semiring Provenance

Given a formula φ in negation normal form (nnf) and a $\mathbb{N}[X, \bar{X}]$ -interpretation π , we now demonstrate how to extract $\pi[\varphi]_\nu$ from the subgraph of the provenance graph generated based on π over $\text{TL}_{\varphi \rightarrow Q}(\varphi)$ rooted at the tuple node $Q_\varphi(\nu(\text{FREE}(\varphi)))$. First, we apply $\text{TL}_{\varphi \rightarrow Q}(\varphi)$ to compute Q_φ . Then, we generate an instance I_π where the existence of a tuple corresponding to a literal $R(\mathbf{a})$ is determined based on the truth value of this literal encoded by its annotation $\pi(R(\mathbf{a}))$. A tuple $R(\mathbf{a})$ exists in I_π if $\pi(R(\mathbf{a})) = x$ or $\pi(R(\mathbf{a})) = 1$ and $\pi(\neg R(\mathbf{a})) = 0$, the tuple is missing if $\pi(\neg R(\mathbf{a})) = \bar{x}$ or $\pi(\neg R(\mathbf{a})) = 1$ and $\pi(R(\mathbf{a})) = 0$, and the tuple's existence is undetermined if $\pi(R(\mathbf{a})) = x$ and $\pi(\neg R(\mathbf{a})) = \bar{x}$. Note that this corresponds to the truth value according to the 5 cases we have discerned in Sec. 7.1.

Next, we generate the provenance graph $\mathcal{PG}(Q_\varphi, I_\pi)$. If the formula has free variables, then the provenance graph will contain multiple tuple nodes $Q_\varphi(\nu(\text{FREE}(\varphi)))$, one for each valuation ν of the free variables, and the subgraph rooted at one such tuple node encodes $\pi[\varphi]_\nu$. By applying a function $\text{TR}_{\text{EXPL} \rightarrow \mathbb{N}[X, \bar{X}]}$ (defined in the following), we translate the subgraph rooted at tuple $Q_\varphi(\nu(\text{FREE}(\varphi)))$ in $\mathcal{PG}(Q_\varphi, I_\pi)$ into $\pi[\varphi]_\nu$.

The function $\text{TR}_{\text{EXPL} \rightarrow \mathbb{N}[X, \bar{X}]}$ replaces nodes in the provenance graph with nodes labeled as “+”, “.”, and annotations of literals. The polynomial $\pi[\varphi]_\nu$ can then be read from the graph generated by $\text{TR}_{\text{EXPL} \rightarrow \mathbb{N}[X, \bar{X}]}$ through a top-down traversal. Intuitively, the translation can be explained as follows. Datalog rules are a conjunction of atoms and, thus, are replaced with multiplication. There may exist multiple ways that derive

an IDB tuple through the rules of query. That is, IDB tuple nodes represent addition. The exception is IDB tuples that are used in a negated fashion which are replaced with multiplication, because, for the goal to succeed, all derivations of the tuple have to fail. Note that a tuple is used in a negated way if there is an odd number of negated goals on the path between the root of the provenance graph and IDB tuple node. In a program produced by our translation rules, this can only be the case for tuples that correspond to head predicate of a rule computing the $\neg \exists$ part of the translation of a universal quantification.

$\text{TR}_{\text{EXPL} \rightarrow \mathbb{N}[X, \bar{X}]}$ consists of the following steps:

1. Replace tuple nodes $\text{Dom}(\mathbf{x})$ with 1.
2. A goal node connected to an EDB tuple node representing a literal $R(\mathbf{a})$ is replaced by $\pi(R(\mathbf{a}))$ if the goal is positive and $\pi(\neg R(\mathbf{a}))$ otherwise.
3. Next, all EDB tuple nodes are removed leaving the goal nodes formerly connected to EDB tuple nodes to be the new leaves of the graph.
4. Rule nodes are replaced with multiplication (.)
5. Next all remaining goal nodes are replaced with addition (+).
6. Finally, nodes v_t corresponding to IDB tuples are replaced with addition with the exception of IDB tuples corresponding to the head predicate (Q_φ) of the second rule of a translated universal quantification which are replaced with multiplication (.)

Example 8. Consider the formula and query from Example 7. Assume that $A = \{a, b\}$ and consider that interpretation π which tracks provenance for edge (a, a) , keeps $R(b, a)$ undetermined, and sets all other positive literals to false without provenance tracking, i.e., $\pi(R(a, a)) = x$, $\pi(\neg R(a, a)) = 0$, $\pi(R(b, a)) = y$, $\pi(\neg R(b, a)) = \bar{y}$, and for all other $R(\mathbf{a})$ we have $\pi(R(\mathbf{a})) = 0$ and $\pi(\neg R(\mathbf{a})) = 1$. That is, in I_π , tuple $R(a, a)$ exists, tuple $R(b, a)$'s existence is undetermined, and all other tuples are missing. Fig. 7 shows the provenance graph $\mathcal{PG}(Q_\varphi, I_\pi)$. The truth of the universal quantification in φ is undetermined, because while there exists no $a \in A$

such that $\neg\varphi_1$ for $\nu := (x = a)$ is true (there is an outgoing edge starting at a), the truth of $\neg\varphi_1$ is undetermined for $\nu := (x = b)$ (the existence of edge $R(b, a)$ is undetermined). The truth of $\exists y R(a, y)$ and $\exists y R(b, y)$ is justified by the existing tuples $R(a, a)$ and $R(b, a)$, respectively. Applying the translation $\text{TR}_{\text{EXPL} \rightarrow \mathbb{N}[X, \bar{X}]}$, we get graph with node labels shown in Fig. 7 which corresponds to the polynomial $1 \cdot x \cdot 1 \cdot 1 \cdot y = x \cdot y = \pi(\varphi)$.

We are now ready to state the main result of this section: our provenance graphs extended for undetermined facts can encode semiring provenance for first-order (FO) model checking. For simplicity, we only consider sentences, i.e., formulas φ without free variables, but the result also holds for formulas with free variables by only translating a subgraph of the provenance rooted at the IDB tuple node $Q_\varphi(\nu(\text{FREE}(\varphi)))$ which corresponds to the formula $\nu(\varphi)$.

Theorem 3. *Let φ be a formula, π a $\mathbb{N}[X, \bar{X}]$ -interpretation, $Q := \text{TL}_{\varphi \rightarrow Q}(\varphi)$, and I_π the instance corresponding to π as defined above. Then*

$$\text{TR}_{\text{EXPL} \rightarrow \mathbb{N}[X, \bar{X}]}(\mathcal{PG}(Q, I_\pi)) = \pi[\varphi]_\nu$$

Proof. We prove the theorem by induction over the structure of the input formula φ for a given valuation ν of $\text{FREE}(\varphi)$. For the full proof, see [27]. \square

8 Computing Explanations

We now present our approach for computing explanations using Datalog. Our approach generates a Datalog program $\mathbb{GP}_{P, \psi}$ by rewriting a given query (input program) P to return the edge relation of the explanation $\text{EXPL}(P, \psi, I)$ for a provenance question (PQ) ψ . Recall that a PQ is a pattern describing existing/missing outputs of interest and that an explanation for a PQ is a subgraph of the provenance which contains the provenance of all tuples described by the pattern.

Our approach for computing $\mathbb{GP}_{P, \psi}$ consists of the following steps that we describe in detail in the following subsections: 1) we unify the input program P with the PQ ψ by propagating constants from ψ top-down to prune derivations of outputs that do not match the PQ; 2) we determine for each IDB predicate whether the explanation may contain existing, missing, or both types of tuples from this predicate. Similarly, for each rule we determine whether successful, failed, all, or no derivations of this rule may occur in the provenance graph; 3) based on restricted and annotated version of the input program produced by the first two steps, we then generate *firing rules* which capture the variable bindings of successful and failed derivations of the input program's rules; 4) The result of the firing rules is

a superset of the set of relevant provenance fragments. We introduce additional rules that enforce connectivity to remove spurious fragments; 5) finally, we create rules that generate the edge relation of the explanation. This is the only step that depends on what provenance type (e.g., Fig. 4) is requested.

In the following, we will illustrate our approach using the provenance question $\psi_{n,s} = \text{WHY } Q(n, s)$ from Example 1, i.e., why New York is connected to Seattle via train with one intermediate stop, but there is no direct connection.

8.1 Unifying the Program with the PQ

The node $Q(n, s)$ in the provenance graph (Fig. 2) is only connected to derivations which return $Q(n, s)$. For instance, if variable X is bound to another city x (e.g., Chicago) in a derivation of the rule r_1 , then this rule cannot return the tuple (n, s) . This reasoning can be applied recursively to replace variables in rules with constants. That is, we unify the rules in the program top-down with the PQ. This process corresponds to selection push-down for relational algebra expressions. We may create multiple partially unified versions of a rule or predicate. For example, to explore successful derivations of $Q(n, s)$, we are interested in both train connections from New York to some city ($T(n, Z)$) and from any city to Seattle ($T(Z, s)$). Furthermore, we need to know whether there is a direct connection from New York to Seattle ($T(n, s)$). We store variable bindings as superscripts to distinguish multiple copies of a rule generated based on different bindings.

Example 9. *Given the question $\psi_{n,s}$, we unify the single rule r_1 using the assignment $(X=n, Y=s)$:*

$$r_1^{(X=n, Y=s)} : Q(n, s) :- T(n, Z), T(Z, s), \neg T(n, s)$$

This approach is correct because if we bind a variable in the head of rule, then only rule derivations that agree with this binding can derive tuples that agree with this binding. Based on this unification step, we know which bindings may produce fragments of $\mathcal{PG}(P, I)$ that are relevant for explaining the PQ (the pseudocode for the algorithm is presented in [25]). For an input P , we use P_{Unified} to denote the result of this unification.

8.2 Add Annotations based on Success/Failure

For $\text{WHY } Q(t)$ ($\text{WHYNOT } Q(t)$), we are only interested in subgraphs of the provenance rooted at existing (missing) tuple nodes for Q . With this information, we can infer restrictions for the success/failure state of nodes

in the provenance graph that are directly or indirectly connected to PQ node(s) (belong to the explanation). We store these restrictions as annotations $\textcolor{red}{T}$, $\textcolor{red}{F}$, and $\textcolor{red}{F}/\textcolor{red}{T}$ on heads and goals of rules and use these annotations to guide the generation of rules that capture derivations in step 3. Here, $\textcolor{red}{T}$ ($\textcolor{red}{F}$) indicates that we are only interested in successful (failed) nodes, and $\textcolor{red}{F}/\textcolor{red}{T}$ that we are interested in both.

Example 10. *Continuing with our running example question $\psi_{n,s}$, we know that $\text{Q}(n,s)$ is in the result (Fig. 1). This implies that only successful rule nodes and their successful goal nodes can be connected to this tuple node. Note that this annotation only indicates that it is sufficient to focus on successful rule derivations since failed ones cannot be connected to $\text{Q}(n,s)$.*

$r_1^{(X=n,Y=s),\textcolor{red}{T}} : \text{Q}(n,s) \textcolor{red}{T} :- \text{T}(n,Z) \textcolor{red}{T}, \text{T}(Z,s) \textcolor{red}{T}, \neg \text{T}(n,s) \textcolor{red}{T}$

We now propagate the annotations of the goals in r_1 throughout the program. That is, for any goal that is an IDB predicate, we propagate its annotation to the head of all rules deriving the goal's predicate and, then, propagate these annotations to the corresponding rule bodies. Note that the inverted annotation is propagated for negated goals (e.g., $\neg \text{T}(n,s) \textcolor{red}{T}$). For instance, if T would be an IDB predicate, then we would annotate the head of all rules deriving $\text{T}(n,s)$ with $\textcolor{red}{F}$, because $\text{Q}(n,s)$ can only exist if $\text{T}(n,s)$ does not exist.

Partially unified atoms such as $\text{T}(n,Z)$ may occur in both negative and positive goals. We annotate such atoms with $\textcolor{red}{F}/\textcolor{red}{T}$. The algorithm generating the annotation consists of the steps shown below (the pseudocode is presented in [25]). We use P_{Annot} to denote the result of this algorithm for $P_{Unified}$ (input to this step).

1. Annotate the head of all rules deriving tuples matching the question with $\textcolor{red}{T}$ (why) or $\textcolor{red}{F}$ (why-not).
2. Repeat the following steps until a fixpoint is reached:
 - (a) Propagate the annotation of a rule head to goals in the rule body as follows: propagate $\textcolor{red}{T}$ for $\textcolor{red}{T}$ annotated heads and $\textcolor{red}{F}/\textcolor{red}{T}$ for $\textcolor{red}{F}$ annotated heads.
 - (b) For each annotated positive goal in the rule body, we propagate its annotation ($\textcolor{red}{F}$, $\textcolor{red}{T}$, or $\textcolor{red}{F}/\textcolor{red}{T}$) to all rules that have this atom in the head. For negated goals, we propagate the inverted annotation (e.g., $\textcolor{red}{F}$ for $\textcolor{red}{T}$) unless the annotation is $\textcolor{red}{F}/\textcolor{red}{T}$ in which case we propagate $\textcolor{red}{F}/\textcolor{red}{T}$.

8.3 Creating Firing Rules

To compute the relevant subgraph of $\mathcal{PG}(P,I)$ (the explanation) for a PQ, we need to determine successful and/or failed rule derivations. Each derivation paired

```

 $\text{F}_{\text{Q},\textcolor{red}{T}}(n,s) :- \text{F}_{\text{r}_1,\textcolor{red}{T}}(n,s,Z)$ 
 $\text{F}_{\text{r}_1,\textcolor{red}{T}}(n,s,Z) :- \text{F}_{\text{T},\textcolor{red}{T}}(n,Z), \text{F}_{\text{T},\textcolor{red}{T}}(Z,s), \text{F}_{\text{T},\textcolor{red}{F}}(n,s)$ 
 $\text{F}_{\text{T},\textcolor{red}{T}}(n,Z) :- \text{T}(n,Z)$ 
 $\text{F}_{\text{T},\textcolor{red}{T}}(Z,s) :- \text{T}(Z,s)$ 
 $\text{F}_{\text{T},\textcolor{red}{F}}(n,s) :- \neg \text{T}(n,s)$ 

```

Fig. 8: Example firing rules for WHY $\text{Q}(n,s)$

with the information whether it is successful over the given database (and which goals are failed in case it is not successful) is sufficient for generating a fragment of $\mathcal{PG}(P,I)$. Successful derivations are always part of $\mathcal{PG}(P,I)$ for a given query (input program) P whereas failed rule derivations only appear if the tuple in the head failed, i.e., there are no successful derivations of any rule with this head. To capture the variable bindings of successful/failed rule derivations, we create “firing rules”. For successful rule derivations, a firing rule consists of the body of the rule (but using the firing version of each predicate in the body) and a new head predicate that contains all variables used in the rule. In this way, the firing rule captures all the variable bindings of a rule derivation. Furthermore, for each IDB predicate R that occurs as a head of a rule r , we create a firing rule that has the firing version of predicate R in the head and firing version of the rules r deriving the predicate in the body. For EDB predicates, we create firing rules that have the firing version of the predicate in the head and the EDB predicate in the body.

Example 11. *Consider the annotated program in Example 10 for the question $\psi_{n,s} = \text{WHY } \text{Q}(n,s)$. We generate the firing rules shown in Fig. 8. The firing rule for $r_1^{(X=n,Y=s),\textcolor{red}{T}}$ (the second rule from the top) is derived from the rule r_1 by adding Z (the only existential variable) to the head, renaming the head predicate as $\text{F}_{\text{r}_1,\textcolor{red}{T}}$, and replacing each goal with its firing version (e.g., $\text{F}_{\text{T},\textcolor{red}{T}}$ for the two positive goals and $\text{F}_{\text{T},\textcolor{red}{F}}$ for the negated goal). Note that negated goals are replaced with firing rules that have inverted annotations (e.g., the goal $\neg \text{T}(n,s) \textcolor{red}{T}$ is replaced with $\text{F}_{\text{T},\textcolor{red}{F}}(n,s)$). Furthermore, we introduce firing rules for EDB tuples (three rules at the bottom).*

We, now, extend firing rules to support queries with negation and capture missing answers. To construct a $\mathcal{PG}(P,I)$ fragment corresponding to a missing tuple, we need to find failed rule derivations with the tuple in the head and ensure that no successful derivations with this head exist (otherwise, we may capture irrelevant failed derivations of existing tuples). In addition, we need to determine which goals are failed for a failed rule derivation because only failed goals are connected to the node representing the failed rule derivation in the provenance graph. To capture this information, we add additional

```

 $F_{Q,F}(s, n) :- \neg F_{Q,T}(s, n)$ 
 $F_{Q,T}(s, n) :- F_{r1,T}(s, n, Z)$ 
 $F_{r1,F}(s, n, Z, V1, V2, \neg V3) :- F_{Q,F}(s, n), F_{T,F/T}(s, Z, V1),$ 
 $F_{T,F/T}(Z, n, V2), F_{T,F/T}(s, n, V3)$ 
 $F_{r1,T}(s, n, Z) :- F_{T,T}(s, Z), F_{T,T}(Z, n), F_{T,F}(s, n)$ 
 $F_{T,F/T}(s, Z, true) :- F_{T,T}(s, Z)$ 
 $F_{T,F/T}(s, Z, false) :- F_{T,F}(s, Z)$ 
 $F_{T,T}(s, Z) :- T(s, Z)$ 
 $F_{T,F}(s, Z) :- \text{dom}_{T.\text{toCity}}(Z), \neg T(s, Z)$ 

```

Fig. 9: Example firing rules for WHYNOT $Q(s, n)$

boolean variables — V_i for goal g^i — to the head of a firing rule that record for each goal whether it failed or not. The body of a firing rule for failed rule derivations is created by replacing every goal in the body with its F/T firing version, and adding the firing version of the negated head to the body (to ensure that only bindings for missing tuples are captured). Firing rules capturing failed derivations use the F/T firing versions of their goals because not all goals of a failed derivation have to be failed and the failure status determines whether the corresponding goal node is part of the explanation. A firing rule capturing missing tuples may not be safe, i.e., it may contain variables that only occur in negated goals. These variables should be restricted to the associated domains for the attributes the variables are bound to. Recall that associated domain $\text{dom}(R.A)$ for an attribute $R.A$ is given as an unary query $\text{dom}_{R.A}$. We use these queries in firing rules to restrict the values a variable is bound to. Thus, we ensure that only missing answers formed from the associated domains are considered and that firing rules are safe.

Example 12. Consider the question WHYNOT $Q(s, n)$ from Example 1. The firing rules generated for this question are in Fig. 9. We exclude the rules for the second goal $T(Z, n)$ and the negated goal $\neg T(s, n)$ which are analogous to the rules for the first goal $T(s, Z)$. New York cannot be reached from Seattle with exactly one transfer, i.e., $Q(s, n)$ is not in the result. Thus, we are only interested in failed derivations of rule r_1 with $X=s$ and $Y=n$. Furthermore, each rule node in the provenance graph corresponding to such a derivation will only be connected to failed subgoals. Thus, we need to capture which goals are successful or failed for each such failed derivation. We model this using boolean variables V_1 , V_2 , and V_3 (one for each goal) that are set to true iff the tuple corresponding to the goal exists. The firing version $F_{r1,F}(s, n, Z, V1, V2, \neg V3)$ of r_1 returns all variable bindings for derivations of r_1 such that $Q(s, n)$ is the head (i.e., guaranteed by adding $F_{Q,F}(s, n)$ to the body), the rule derivations are failed, and the tuple cor-

responding to the i^{th} goal exists for this binding iff V_i is true. The failure status of the i^{th} goal is V_i for positive goals and $\neg V_i$ for negated goals. To produce all these bindings, we need rules capturing successful and failed tuple nodes for each subgoal of the rule r_1 . We annotate such rules with F/T and use a boolean variable (true or false) to record whether a tuple exists (e.g., $F_{T,F/T}(s, Z, true) :- F_{T,T}(s, Z)$ is one of these rules). Similarly, $F_{T,F/T}(s, n, false)$ represents the fact that tuple $T(s, n)$ (connection from Seattle to New York) is missing. This causes the third goal of r_1 to succeed for any derivation where $X=s$ and $Y=n$. For each unified EDB atom annotated with F/T , we create four rules: one for existing tuples (e.g., $F_{T,T}(s, Z) :- T(s, Z)$), one for the failure case (e.g., $F_{T,F}(s, Z) :- \text{dom}_{T.\text{toCity}}(Z), \neg T(s, Z)$), and two for the F/T version. For the failure case, we use predicate $\text{dom}_{T.\text{toCity}}$ to only consider missing tuples (s, Z) where Z is a value from the associated domain.

Algorithm 1 takes as input the program P_{Annot} produced by step 2 and outputs a program P_{Fire} containing firing rules. The pseudocode for the subprocedures is presented in [25]. The algorithm maintains a queue $todo$ of annotated atoms that have to be processed which is initialized with $\text{PATTERN}(\psi)$, i.e., the provenance question atom. Furthermore, we maintain a set $done$ of atoms that have been processed already. Variables $todo$, $done$, and P_{Fire} are global variables that are shared with the subprocedures of this algorithm. For each atom $R(t)^\sigma$ (line 8) from the queue (here σ is the annotation of the atom, e.g., F), we mark the atom as done (line 9). We then consider two cases: R is an EDB atom or an IDB atom in which case we have to create firing rules for the predicate (relation) and the rules deriving it. The firing rules for EDB predicates check whether the tuples do or do not exist. These rules allow us to determine the success or failure of goals corresponding EDB predicates in rule derivations. For IDB predicates, we create firing rules that determine their existence based on successful or failed rule derivations captured by firing rules for the rules of the program. Consider a given program P with two rules: 1) $r_1 : Q(X) :- R(X, Y), Q_1(Y)$ and 2) $r_2 : Q_1(Y) :- S(Y, Z)$ where R and S are EDB relations and Q and Q_1 are IDB predicates. To capture provenance for the predicate $Q(X)$, we create firing rules for R and S to check existence or absence of tuples matching t in R and S . Moreover, we also generate firing rules for rules r_1 and r_2 to explain how derivations of $Q(X)$ through these rules have succeeded or failed. The firing rule for r_1 uses the firing rule for IDB predicate Q_1 which in turn uses the firing rule for r_2 since $\text{head}(r_2) = Q_1$. We describe these two cases in the following.

Algorithm 1 Create Firing Rules

```

1: procedure CREATEFIRINGRULES( $P_{Annot}$ ,  $\psi$ )
2:    $P_{Fire} \leftarrow []$ 
3:    $state \leftarrow \text{typeof}(\psi)$ 
4:    $Q(t) \leftarrow \text{PATTERN}(\psi)$ 
5:    $todo \leftarrow [Q(t)^{state}]$ 
6:    $done \leftarrow []$ 
7:   while  $todo \neq []$  do                                 $\triangleright$  create rules for a predicate
8:      $R(t)^\sigma \leftarrow \text{POP}(todo)$ 
9:      $\text{INSERT}(done, R(t)^\sigma)$ 
10:    if  $\text{isEDB}(R)$  then
11:      CREATEEDBFIRINGRULE( $P_{Fire}$ ,  $R(t)^\sigma$ )
12:    else
13:      CREATEIDBNEGRULE( $P_{Fire}$ ,  $R(t)^\sigma$ )
14:       $rules \leftarrow \text{GETRULES}(R(t)^\sigma)$ 
15:      for all  $r \in rules$  do                       $\triangleright$  create firing rule for  $r$ 
16:         $args \leftarrow args(\text{head}(r))$ 
17:         $args \leftarrow args :: (args(\text{body}(r)) - args(\text{head}(r)))$ 
18:        CREATEIDBPOSRULE( $P_{Fire}$ ,  $R(t)^\sigma$ ,  $r$ ,  $args$ )
19:        CREATEIDBFIRINGRULE( $P_{Fire}$ ,  $R(t)^\sigma$ ,  $r$ ,  $args$ )
20:    return  $P_{Fire}$ 

```

EDB atoms (line 13). For an EDB atom $R(t)^\text{G}$, we use procedure CREATEEDBFIRINGRULE to create one rule $F_{R,\text{G}}(t) :- R(t)$ that returns tuples from relation R that match t . For missing tuples $(R(t)^\text{F})$, we extract all variables from t (some arguments may be constants propagated during unification) and create a rule that returns all tuples that can be formed from values of the associated domains of the attributes these variables are bound to and do not exist in R . This is achieved by adding goals $\text{dom}(X_i)$ as explained in Example 12.

IDB atoms (lines 13-19). IDB atoms with F or F/G annotations are handled in the same way as EDB atoms with these annotations. If the atom is $R(t)^\text{F}$ (line 13), we create a rule with $\neg F_{R,\text{G}}(t)$ in the body using the associated domain queries to restrict variable bindings. Similarly, for $R(t)^\text{F}/\text{G}$, the procedure called in line 13 adds two additional rules as shown in Fig. 9 (5th and 6th rule) for EDB atoms. Both types of rules only use the positive firing version for $R(t)$ and domain predicates in their body. Thus, these rules are independent of which rules derive R . Now, for any R , we create positive firing rules that correspond to the derivation of R through one particular rule. For that, we iterate over the annotated versions of all rules deriving R (lines 14+15). For each rule r with head $R(t)$, we create a rule $F_{R,\text{G}}(t) :- F_{R,\text{G}}(\vec{X})$ where \vec{X} is the concatenation of t with all existential variables from the body of r .

Rules (line 15-19). Consider a rule $r : R(t) :- g_1(\vec{X}_1), \dots, g_n(\vec{X}_n)$. If the head of r is annotated with G , then we create a rule with head $F_{R,\text{G}}(\vec{X})$ where $\vec{X} = \text{vars}(r)$ (stored in variable $args$, lines 16+17) and the same body as r except that each goal is replaced with its firing version with appropriate annotation (e.g., G for

```

 $F_{Q,\text{G}}(n, s) :- F_{r_1,\text{G}}(n, s, Z)$ 
 $F_{r_1,\text{G}}(n, s, Z) :- F_{T,\text{G}}(n, Z), F_{T,\text{G}}(Z, s), F_{T,\text{F}}(n, s)$ 
 $FC_{r_2,r_1,\text{G}}(n, Z) :- T(n, Z), F_{r_1,\text{G}}(n, s, Z)$ 
 $FC_{r_2,r_1,\text{G}}(Z, s) :- T(Z, s), F_{r_1,\text{G}}(n, s, Z)$ 
 $FC_{r_2,r_1,\text{F}}(n, s) :- \neg T(n, s), F_{r_1,\text{G}}(n, s, Z)$ 

```

Fig. 10: Example firing rules with connectivity checks

positive goals). For rules annotated with F or F/G , we create one additional rule with head $F_{R,\text{F}}(\vec{X}, \vec{V})$ where \vec{X} is defined as above, and \vec{V} contains V_i if the i^{th} goal of r is positive and $\neg V_i$ otherwise. The body of this rule contains the F/G version of every goal in r 's body plus an additional goal $F_{R,\text{F}}$ to ensure that the head atom is failed. As an example for this type of rule, consider the third rule from the top in Fig. 9.

Theorem 4 (Correctness of Firing Rules). *Let P be an input program, r denote a rule of P with m goals, and P_{Fire} be the firing version of P . We use $r(t) \models P(I)$ to denote that the rule derivation $r(t)$ is successful in the evaluation of program P over I . The firing rules for P correctly determine existence of tuples, successful derivations, and failed derivations for missing answers:*

- $F_{R,\text{G}}(t) \in P_{Fire}(I) \leftrightarrow R(t) \in P(I)$
- $F_{R,\text{F}}(t) \in P_{Fire}(I) \leftrightarrow R(t) \notin P(I)$
- $F_{r,\text{G}}(t) \in P_{Fire}(I) \leftrightarrow r(t) \models P(I)$
- $F_{r,\text{F}}(t, \vec{V}) \in P_{Fire}(I) \leftrightarrow r(t) \not\models P(I) \wedge \text{head}(r(t)) \notin P(I)$ and for $i \in \{1, \dots, m\}$ we have that V_i is false iff i^{th} goal fails in $r(t)$.

Proof. We prove Theorem 4 by induction over the structure of a program. For the proof, see [26] or [27]. \square

8.4 Connectivity Joins

To be in the result of a firing rule is a necessary, but not sufficient, condition for the corresponding rule node to be connected to a node $Q(t') \in \text{MATCH}(\psi)$ in the explanation. Thus, we have to check connectivity of intermediate results explicitly.

Example 13. Consider the firing rules for $\psi_{n,s}$ shown in Fig. 8. The corresponding rules with connectivity checks are shown in Fig. 10. All rule nodes corresponding to $F_{r,\text{G}}(n, s, Z)$ are guaranteed to be connected to the node $Q(n, s)$ (corresponding to the only atom in $\text{MATCH}(\psi_{n,s})$). Note that connectivity joins are also required for negative firing rules (e.g., $F_{r_1,\text{F}}(s, n, Z, V_1, V_2, \neg V_3)$ in Fig. 9 is used for WHYNOT). For sake of example, assume that instead of using T , rule r_1 uses an IDB relation R which is computed using a rule $r_2 : R(X, Y) :- T(X, Y)$.

Consider the firing rule $F_{r_2, \text{green}}(n, Z) :- T(n, Z)$ created based on the 1st goal of r_1 . Some provenance fragments computed by this rule may not be connected to $Q(n, s)$. A tuple node $R(n, c)$ for a constant c is only connected to the node $Q(n, s)$ iff it is part of a successful binding of r_1 . That is, for the node $R(n, c)$, there has to exist a tuple $R(c, s)$. Connectivity is achieved by adding the head of the firing rule for r_1 to the body of the firing rule for r_2 as shown in Fig. 10 (the 3rd and 4th rule).

Our algorithm traverses the query’s rules starting from PQ atom(s) to find all combinations of rules r_i and r_j such that the head of r_j can be unified with a goal in r_i ’s body. For each such pair (r_i, r_j) where the head of r_j corresponds to the k^{th} goal in the body of r_i , we create a rule $FC_{r_j, r_i^k, \text{green}}(\vec{X})$ as follows. We unify the variables of the k^{th} goal in the firing rule for r_i with the head variables of the firing rule for r_j . All remaining variables of r_i are renamed to avoid name clashes. We add the unified head of r_i to the body of r_j . These rules check whether rule nodes in the provenance graph are connected to nodes in $\text{MATCH}(\psi)$.

8.5 Computing the Edge Relation

The program created so far captures sufficient information for generating the edge relation of the explanation for a PQ (which is used when rendering graphs). We make this step part of the program to offload this work to database backend. To compute the edge relation, we use Skolem functions to create node identifiers. An identifier records the type of the node (tuple, rule, or goal), variables assignments, and the success/failure status of the node, e.g., a tuple node $T(n, s)$ that is successful would be represented as $f_T^{\text{green}}(n, s)$. Each rule firing corresponds to a fragment of $\mathcal{PG}(P, I)$. For example, one such fragment is shown in Fig. 11 (left). Such a sub-structure is created through a set of rules:

- One rule creating edges between tuple nodes for the head predicate and rule nodes
- One rule for each goal connecting a rule node to that goal node (only failed goals for failed rules)
- One rule creating edges between each goal node and the corresponding EDB tuple node

Example 14. Consider the firing rules with connectivity joins from Example 13. Some of the rules for creating the edge relation of the explanation sought by the user are shown in Fig. 11 (right). For example, each edge connecting the tuple node $Q(n, s)$ to a successful rule node $r_1(n, s, Z)$ is created by the top-most rule, and the 2nd rule creates an edge between $r_1(n, s, Z)$ and

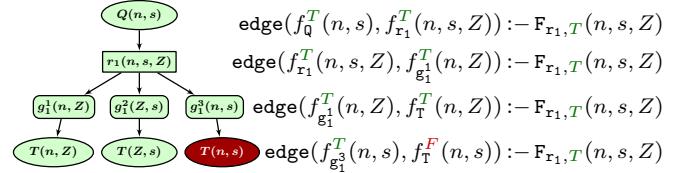


Fig. 11: Fragment of an explanation corresponding to a derivation of rule r_1 (left) and the rules generating the edge relation for such a fragment (right)

$g_1^1(n, Z)$. Edges for failed derivations are created by considering the corresponding node identifiers and a failure pattern (e.g., $F_{r1, \text{red}}(s, n, Z, V_1, V_2, \neg V_3)$).

8.6 \mathcal{K} -Explanations

To compute one of the \mathcal{K} -explanation types introduced in Sec. 6.2, we only have to adapt the rules generating the edge relation. As an example, we present the modifications for computing $\text{EXPL}^{\text{Which}(X)}$ (e.g., Fig. 4c). Recall that semiring $\text{Which}(X)$ models provenance as a set of contributing tuples and we encode this as a graph by connecting a head of a rule derivation to the atoms in its body. That is, for the $\text{EXPL}^{\text{Which}(X)}$, we create only one type of rule that connects tuple nodes for the head predicate to EDB tuple nodes. We use $\mathbb{GP}_{P, \psi}^{\text{Which}(X)}$ to denote the program generated in this way for an input program P , and a PQ ψ .

Example 15. Consider the graph fragment for r_1 in Fig. 11 (left) without rule and goal nodes. The rule that creates the edge between $Q(n, s)$ and $T(n, Z)$ is

$$\text{edge}(f_Q^{\text{green}}(n, s), f_T^{\text{green}}(n, Z)) :- F_{r1, \text{green}}(n, s, Z)$$

For each successful derivation of result $Q(n, s)$ using rule r_1 , a subgraph replacing Z with bindings from the derivation is included in $\text{EXPL}^{\text{Which}(X)}$.

8.7 Correctness

We now prove that our approach is correct.

Theorem 5. Let P be a program, I be a instance, and ψ a PQ. Program $\mathbb{GP}_{P, \psi}^{\text{Which}(X)}$ evaluated over I returns the edge relation of $\text{EXPL}(P, \psi, I)$.

Proof. To prove Theorem 5, we have to show that 1) only edges from $\mathcal{PG}(P, I)$ are in $\mathbb{GP}_{P, \psi}(I)$ and 2) the program returns precisely the set of edges of explanation $\text{EXPL}(P, \psi, I)$. The full proof is presented in [26] and in our accompanying report [27]. \square

Theorem 6. Let P be a positive program, I be a database instance, and ψ a PQ. The result of program $\mathbb{GP}_{P, \psi}^{\text{Which}(X)}$ is the edge relation of $\text{EXPL}^{\text{Which}(X)}(P, \psi, I)$.

Proof. We prove Theorem 6 by induction over the structure of a program as in the proof of Theorem 4. The full proof is presented in our technical report [27]. \square

9 Factorization

For provenance polynomials, we can exploit the distributivity law of semirings to generate factorizations of provenance [33] which are exponentially more concise in the best case. For instance, consider a query r_3 returning the end points of paths of length 2 evaluated over the edge-labelled graph in Fig. 12a. The provenance polynomial for the query result $Q_{2\text{hop}}(d)$ using the annotations from Fig. 12a is shown in Fig. 12d. Each monomial in the polynomial corresponds to one of the derivations of the result using r_3 . Each of these $2 \cdot (2^2)$ (we have two options as starting points and, for each hop, we have two options) derivations corresponds to one path of length 2 ending in d . When generating provenance graphs for provenance polynomials, we create “.” nodes for rule derivations and “+” nodes for IDB tuples. Fig. 12b is the factorized representation of this polynomial. We can exploit the fact that our approach shares common subexpressions to produce a particular factorization. This is achieved by rewriting the input program to partition a query by materializing joins and projections as new IDB relations which can then be shared. We first review f-trees and d-trees as introduced in [34] which encode possible nesting “schemas” for factorized representations of provenance (or query results), the size bounds for factorized representations based on d-trees proven in [34], and how to chose a d-tree for a query that results in the optimal worst-case size bound for the factorized representation of the provenance according to this d-tree. Then, we introduce a query transformation for conjunctive queries which, given an input query and the d-tree for this query, generates a rewritten query which returns a provenance graph factorized corresponding to this d-tree. We employ this rewriting to produce more concise provenance in PUG (experiments are shown in Sec. 11).

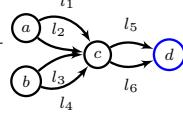
Factorized Representations. In [33, 34], a *factorized representation* (f-rep for short) of a relation is defined as an algebraic expression constructed using singleton relations (one tuple with one value) and the relational operators union and product. Any f-rep over a set of attributes from a schema S can be interpreted as a relation over S by evaluating the algebraic expression, e.g., $\{(a)\} \times (\{(b)\} \cup \{(c)\})$ is a factorized representation of the relation $\{(a, b), (a, c)\}$. Following the convention from [33], we denote a singleton $\{(a)\}$ as a . Factorization can be applied to compactly represent relations and

query results as well as provenance (e.g., Fig. 12b). We will factorize representations of provenance which encode variables of provenance polynomials as the tuples annotated by these variables and show how to extract provenance polynomials from provenance graphs generated in this way.

F-trees for F-reps. Olteanu et al. [34] introduce *f-trees* to encode the nesting structure of f-reps. At first, let us consider only f-trees which encode the nesting structure of a boolean query [33]. An f-tree for a boolean query Q (e.g., r_4 in Fig. 12a) is a rooted forest with one node for every variable of Q .⁴ An f-rep according to an f-tree \mathcal{T} nests values according to \mathcal{T} : a node labelled with X corresponds to a union of values from the attributes bound to X by the query. The values of attributes bound to children of a node X corresponding to a single value x bound to X are grouped under x . If a node has multiple children, then their f-reps are connected via \times . For example, consider an f-tree \mathcal{T} with root X and a single child Y for a query $Q() :- R(X, Y)$. An f-rep of Q according to \mathcal{T} would be of the form $x_1 \times (y_{1_1} \cup \dots \cup y_{1_m}) \cup \dots \cup x_m \times (y_{m_1} \cup \dots \cup y_{m_n})$, i.e., the Y values co-occurring with a given X value x are grouped as a union and then paired with x . An f-tree encodes (conditional) independence of the variables of a query in the sense that the values of one variable do not depend on the values of another variable. For instance, two siblings X and Y in an f-tree have to be independent since a union of X values is paired (cross-product) with a union of Y values. This is only correct if the values of X and Y are independent. The independence assumptions encoded in an f-tree may not hold for every possible query with the same schema as the f-tree. Thus, only some f-trees with a particular schema may be applicable for a query with this schema. It was shown in [34], that a query has an f-rep over an f-tree \mathcal{T} for any database iff for each relation in Q the variables assigned to attributes of this relation (these variables are called dependent) are on the same root-to-leaf path in the f-tree. This is called the *path condition*. Note that multiple references to the same relation in a query are considered as separate relations when checking this condition. For instance, consider the boolean query r_4 in Fig. 12a which checks if there are paths of length 2 ending in the node d . Fig. 12c shows two f-trees \mathcal{T}_1 and \mathcal{T}_2 for this query (ignore the sets on the side of nodes for now). An f-rep according to \mathcal{T}_2 for r_4 would encode a union of Y values paired (\times) with a union of Z values for this Y value. Each Z value nested under a Y value is then paired with a cross-product of L_1 and L_2 values.

⁴ In [34], relational algebra is used to express queries and nodes of f-trees represent equivalence classes of attributes which in Datalog correspond to query variables.

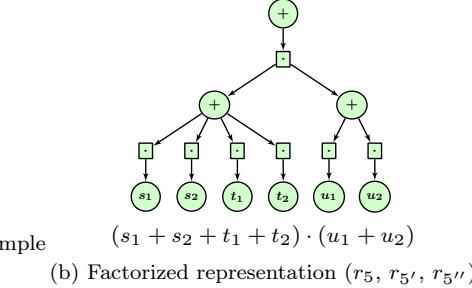
$r_3 : \text{Q}_{\text{2hop}}(X) :- \text{H}(Y, L_1, Z), \text{H}(Z, L_2, X)$
$r_4 : \text{Q}_{\text{2hop-d}}() :- \text{H}(Y, L_1, Z), \text{H}(Z, L_2, d)$
$r_5 : \text{Q}_{\text{2hop}}() :- \text{Q}_{L_1}(Z), \text{Q}_{L_2}(Z)$
$r_{5'} : \text{Q}_{L_1}(Z) :- \text{H}(Y, L_1, Z)$
$r_{5''} : \text{Q}_{L_2}(Z) :- \text{H}(Z, L_2, d)$



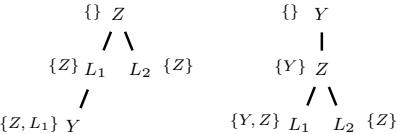
(a) 2hop queries (r_3 and r_4), rewriting (r_5 , r_5' , r_5'') according to d-tree \mathcal{T}_1 , and example database (graph)

Relation H

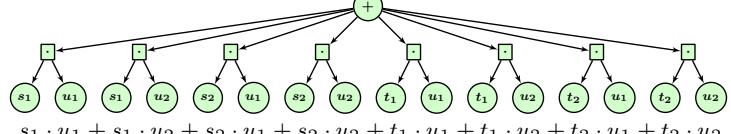
S	L	E
a	l_1	c
a	l_2	c
b	l_3	c
b	l_4	c
c	l_5	d
c	l_6	d



(b) Factorized representation (r_5 , r_5' , r_5'')



(c) Two d-trees of r_4 : \mathcal{T}_1 (left) and \mathcal{T}_2 (right)



(d) Flat representation (r_4)

Fig. 12: Factorized and flat provenance graphs ($\mathbb{N}[X]$) explaining WHY $\text{Q}_{\text{2hop}}(d)$ and two d-trees for r_4 .

D-trees for D-reps. The size of a factorized representation can be further reduced by allowing subexpressions to be shared through definitions, i.e., using algebra graphs instead of trees. In [34], such representations are called *d-representations* (d-rep). Analogous to how f-trees define the structure of f-reps, d-trees were introduced to define the structure of d-reps. A d-tree is an f-tree where each node X is annotated with a set $\text{key}(X)$, a subset of its ancestors in the f-tree on which the node and any of its dependents depend on. The d-rep of the subtree rooted in X is unique for each combination of values from $\text{key}(X)$. That is, if $\text{key}(X)$ is a strict subset of the ancestors of X , then the same d-rep for the subtree at X can be shared by multiple ancestors, reducing the size of the representation. In Fig. 12c, the set key is shown beside each node, e.g., in \mathcal{T}_2 , the variable L_2 depends only on Z , but not on Y . An important result proven in [34] is that, for a given d-tree \mathcal{T} for a query Q , the size of d-rep of Q over a database I is bound by $|I|^{s^\uparrow(\mathcal{T})}$ where $s^\uparrow(\mathcal{T})$ is a rational number computed based on \mathcal{T} alone (see [34] for details of how to compute $s^\uparrow(\mathcal{T})$). This bound can be used to determine the d-tree for a query Q which will yield the d-rep of worst-case optimal size by enumerating the valid d-trees for Q and, then, choosing the one with the lowest value of s^\uparrow .

Example 16. Consider the d-rep for r_4 (Fig. 12a) over the example instance of relation H (Fig. 12a) according to d-tree \mathcal{T}_2 (Fig. 12c). Variable Y at the root of \mathcal{T}_2 is bound to the attribute S from the first reference of H, i.e., the starting point of paths of length 2 ending in d . There are two such starting points a and b . Now each of these are paired with the only valid intermediate node c on these paths (variable Z). Finally, for this node, we compute the cross-product of the L_1 and L_2 values connected to c . Since the L_2 values only depend on Z ,

we share these values when the same Z value is paired with multiple Y values. The final result is $(a \times c \times (l_1 \cup l_2) \times l^\uparrow) + (b \times c \times (l_3 \cup l_4) \times l^\uparrow)$ where $l^\uparrow := (l_5 \cup l_6)$.

Factorization of Provenance. For the provenance of a conjunctive query Q that is not a boolean query, i.e., it has one or more variables in the head (e.g., r_3 in Fig. 12a), we have to compute a provenance polynomial for each result of Q . We would like the factorization of the provenance of Q to clearly associate the provenance polynomial of a result t with the tuple t . That is, we want to avoid factorizations where head variables of Q are nested below variables that store provenance (appear only in the body) since reconstructing the provenance polynomial for t would require enumeration of the full provenance from the factorized representation in the worst case. For example, consider a query with head variable X and body variable Y . If Y is the root of a d-tree \mathcal{T} , then the d-rep of Q according to \mathcal{T} would be of the form $y_1 \times (x_{11} + \dots + x_{n1}) + \dots + y_m \times (x_{1m} + \dots + x_{nm})$. To extract the provenance polynomial for a result x_i , we may have to traverse all y values since there is no indication, for which y values, x_i appears in the sum $x_{1i} + \dots + x_{ni}$. We ensure this by constructing d-trees which do not include the head variables, but treat those as ancestors of every node in the d-tree when computing key for the nodes. For instance, to make \mathcal{T}_1 (Fig. 12c) a valid d-tree for capturing the provenance of r_3 (Fig. 12a), we treat the head variable X as a virtual ancestor of all nodes and get $\text{key}(Z) = \{X\}$ and $\text{key}(L_2) = \{Z, X\}$. Furthermore, if we are computing an explanation to a provenance question (PQ) ψ that binds one or more head variables to constants, then we can propagate these bindings before constructing a d-tree for the query. For example, to explain $\text{Q}_{\text{2hop}}(d)$, we would propagate the binding $X = d$ resulting in rule r_4 (Fig. 12a). Thus, any d-tree for r_4 can be used to cre-

ate a factorized $\text{EXPL}_{\mathbb{N}[X]}$ graph for the user question $\text{WHY}(\mathbb{Q}_{2\text{hop}}(d))$.

Rewriting Queries for Factorization. We now explain how, given a d-tree \mathcal{T} for a conjunctive query Q and positive PQ $\psi := \text{WHY } Q(t)$, to generate a Datalog query Q_{rewr} such that, for any database I , we have that $\text{EXPL}_{\mathbb{N}[X]}(Q_{\text{rewr}}, \psi, I)$ encodes $\mathbb{N}[X](Q_{\text{rewr}}, I, t)$ for each $t \in \text{MATCH}(\psi)$ factorized according to \mathcal{T} . We first unify the query with the PQ as described in Sec. 8.1. Given a unified input query Q and a d-tree \mathcal{T} , we compute Q_{rewr} as follows.

1. Assume a total order among the variables of Q (e.g., the lexicographical order). For every node X with children Y_1, \dots, Y_n in the d-tree \mathcal{T} , we generate

$$r_X : \mathbb{Q}_X(\text{key}(X)) :- \mathbb{Q}_{Y_1}(\text{key}(Y_1)), \dots, \mathbb{Q}_{Y_n}(\text{key}(Y_n))$$

2. Now for every atom $\mathbb{R}(Z_1, \dots, Z_m)$ in the body of Q , we find the shortest path starting in a root node that contains all nodes Z_1 to Z_m . Let $Y = Z_i$ for some i be the last node on this path. Then, we add atom $\mathbb{R}(Z_1, \dots, Z_m)$ to the body of rule r_Y created in the previous step.
3. Let X_1, \dots, X_n be the roots of the d-tree \mathcal{T} (being a forest, a d-tree may have multiple roots). Furthermore, let Y_1, \dots, Y_m denote the head variables of the unified input query Q with the PQ. We create

$$r_Q : \mathbb{Q}(Y_1, \dots, Y_m) :- \mathbb{Q}_{X_1}(\text{key}(X_1)), \dots, \mathbb{Q}_{X_n}(\text{key}(X_n))$$

The rewriting above creates a factorization according to a d-tree \mathcal{T} . However, it may contain rules which cannot potentially lead to reuse and, thus, result in overhead that could be avoided if we were able to identify such rules. We now present an optimization that removes such rules to further reduce the size of the generated provenance graphs. Consider two nodes X and Y in a d-tree where Y is the only child of X , i.e., $\text{key}(Y) = \text{key}(X) \cup \{X\}$. We would generate rules

$$r_X : \mathbb{Q}_X(\text{key}(X)) :- \mathbb{Q}_Y(X \cup \text{key}(X))$$

$$r_Y : \mathbb{Q}_Y(X \cup \text{key}(X)) :- \dots$$

In this case, the intermediate result \mathbb{Q}_Y does not lead to further factorization (we have a union of unions). Thus, we can merge the rules by substituting the atom $\mathbb{Q}_Y(X \cup \text{key}(X))$ in r_X with the body of r_Y . A similar situation may arise with the rule r_Q deriving the final query result. In general, we can merge any rule of the form $\mathbb{Q}_1(X_1, \dots, X_n) :- \mathbb{Q}_2(X_1, \dots, X_n)$ with the rule deriving \mathbb{Q}_2 (in our translation, there will be exactly one rule with head \mathbb{Q}_2).

Example 17. Consider the question $\text{WHY} \mathbb{Q}_{2\text{hop}}(d)$ over the query r_3 from Fig. 12a. Unifying the query with

this question yields r_4 (below r_3 in the same figure). To rewrite the query according to the d-tree \mathcal{T}_1 from Fig. 12c, we apply the above algorithm to create rules:

$$\begin{aligned} r_{Q_{2\text{hop}}} &: \mathbb{Q}_{2\text{hop}}() :- \mathbb{Q}_Z() & r_Z &: \mathbb{Q}_Z() :- \mathbb{Q}_{L_1}(Z), \mathbb{Q}_{L_2}(Z) \\ r_{L_1} &: \mathbb{Q}_{L_1}(Z) :- \mathbb{Q}_Y(Z, L_1) & r_Y &: \mathbb{Q}_Y(Z, L_1) :- \mathbb{H}(Y, L_1, Z) \\ r_{L_2} &: \mathbb{Q}_{L_2}(Z) :- \mathbb{H}(Z, L_2, d) \end{aligned}$$

Applying the optimizations introduced above, we merge the rules $r_{Q_{2\text{hop}}}$ with r_Z (the head \mathbb{Q}_Z is the body of $r_{Q_{2\text{hop}}}$). Since $\text{key}(Y) = \text{key}(L_1) \cup \{L_1\}$ and L_1 has only one child, we merge r_Y into r_{L_1} . The resulting program is shown as rules r_5 , $r_{5'}$ and $r_{5''}$ in Fig. 12a.

Factorized Explanations. To generate a concise factorization of provenance for a PQ ψ over a conjunctive query Q , we first find a d-tree \mathcal{T} with minimal s^\uparrow among all d-trees for Q (such a d-tree \mathcal{T} guarantees worst-case optimal size bounds for the generated factorization). Then, we rewrite the input query according to \mathcal{T} (explained above) and use the approach in Sec. 8 to generate $\text{EXPL}_{\mathbb{N}[X]}(Q_{\text{rewr}}, \psi, I)$ encoding the d-rep of $\mathbb{N}[X](Q_{\text{rewr}}, I, t)$ for each $t \in \text{MATCH}(\psi)$.

Example 18. Continuing with Example 17, assume we compute the $\mathbb{N}[X]$ explanation using the rewritten query (r_5 , $r_{5'}$, and $r_{5''}$). The result over the example database is shown in Fig. 12b. The top-most addition and multiplication correspond to the successful derivation using rule r_5 (using c as an intermediate hop from some node to d). The left branch below the multiplication encodes the four possible derivations of $\mathbb{Q}_{L_1}(c)$ ($s_1 + s_2 + t_1 + t_2$) and the right branch corresponds to the two derivations of $\mathbb{Q}_{L_2}(c)$ ($u_1 + u_2$). The polynomial captured by this graph is $(s_1 + s_2 + t_1 + t_2) \cdot (u_1 + u_2)$. That is, there are 4 ways to reach c from any starting node and two ways of reaching d from c leading to a total of $4 \cdot 2 = 8$ paths of length two ending in the node d .

10 Implementation

We have implemented the approach presented in this paper in a system called *PUG* (Provenance Unification through Graphs). PUG is an extension of GProM [1], a middleware that executes provenance requests using a relational database backend (shown in Fig. 13). We have extended the system to support Datalog enriched with syntax for stating provenance questions. The user provides a why or why-not question and the corresponding Datalog query as an input. Our system parses and semantically analyzes this input. Schema information is gathered by querying the catalog of the backend database (e.g., to determine whether an EDB predicate exists). Modules for accessing schema information are already part of the GProM system, but a

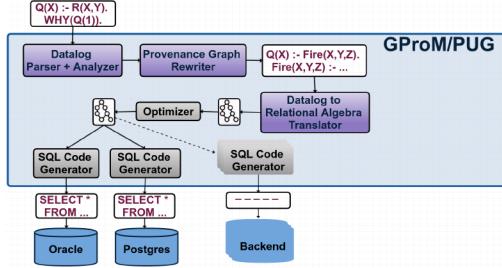


Fig. 13: PUG implementation in GProM

new semantic analysis component had to be developed to support Datalog. The algorithms presented in Sec. 8 are applied to create the program $\mathbb{G}\mathbb{P}_{P,\psi}$ for the input program P and the provenance question ψ which computes $\text{EXPL}(P, \psi, I)$ (analogously, $\text{EXPL}_{\mathcal{K}}(P, \psi, I)$ for $\mathbb{G}\mathbb{P}_{P,\psi}^{\mathcal{K}}$). This program is then translated into relational algebra (\mathcal{RA}). The resulting algebra expression is translated into SQL and sent to the backend database to compute the edge relation of the explanation for the question. Based on this edge relation, we render a provenance graph. For examples and installation guidelines see: <https://github.com/IITDBGroup/PUG>. While it would certainly be possible to directly translate the Datalog program into SQL without the intermediate translation into \mathcal{RA} , we choose to introduce this step to be able to leverage the existing heuristic and cost-based optimizations for \mathcal{RA} expressions provided by GProM [32] and use its library of \mathcal{RA} to SQL translators. Our translation of first-order queries (a program with a distinguished answer relation) to \mathcal{RA} is mostly standard. See [26] for details and an example.

11 Experiments

We evaluate the performance of our solution over a co-author graph relation extracted from DBLP (<http://www dblp org>) as well as over the TPC-H benchmark dataset (<http://www tpc org/tpch/default asp>). We mainly evaluate three aspects; 1) we compare our approach for computing explanations (EXPL) with the approach introduced for provenance games [24]. We call the provenance game approach **Direct Method** (DM), because it directly constructs the full provenance graph; 2) we compare our approach for Lineage (EXPL_{WHICH(X)}) to the language-integrated approach developed for the *Links* programming language [9]; 3) we evaluate the performance impact of rewriting queries to produce factorized provenance (Sec. 9). We have created subsets of the DBLP dataset with 100, 1K, 10K, 100K, 1M, and 8M co-author pairs (tuples). For the TPC-H benchmark, we used database sizes 10MB, 100MB, 1GB, and 10GB. All experiments were run on a machine with 2

$r_1 : \text{only2hop}(X, Y) :- \text{DBLP}(X, Z), \text{DBLP}(Z, Y), \neg \text{DBLP}(X, Y)$
$r_2 : \text{XwithYnotZ}(X, Y) :- \text{DBLP}(X, Y), \neg Q_1(X)$
$r_{2'} : Q_1(X) :- \text{DBLP}(X, 'Svein Johannessen')$
$r_3 : \text{only3hop}(X, Y) :- \text{DBLP}(X, A), \text{DBLP}(A, B), \text{DBLP}(B, Y), \neg E_1(X), \neg E_2(X)$
$r_{3'} : E_1(X) :- \text{DBLP}(X, Y)$
$r_{3''} : E_2(X) :- \text{DBLP}(X, A), \text{DBLP}(A, Y)$
$r_4 : \text{ordPriority}(X, Y) :- \text{CUSTOMER}(A, X, B, C, D, E, F, G), \text{ORDERS}(H, A, I, J, K, Y, M, N, O)$
$r_5 : \text{ordDisc}(X, Y) :- \text{CUSTOMER}(A, X, B, C, D, E, F, G), \text{ORDERS}(H, A, I, J, K, L, M, O, P), \text{LINEITEM}(H, Q, R, S, T, U, V, Y, W, Z, A', B', C', D', E', F')$
$r_6 : \text{partNotAsia}(X) :- \text{PART}(A, X, B, C, D, E, F, G, H), \text{PARTSUPP}(A, I, J, K, L), \text{SUPPLIER}(I, M, N, O, P, Q, R), \text{NATION}(O, S, T, U), \neg R_1(T, 'ASIA')$
$r_{6'} : R_1(T, Z) :- \text{REGION}(T, Z, V)$
$r_7 : \text{suppCust}(N) :- \text{SUPPLIER}(A, B, C, N, D, E, F), \text{CUSTOMER}(G, H, I, N, J, K, L, M)$

Fig. 14: DBLP and TPC-H queries for experiments

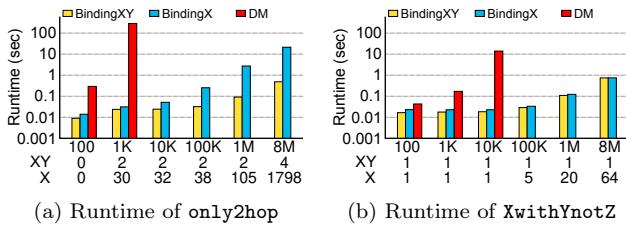
x 3.3Ghz AMD Opteron 4238 CPUs (12 cores in total) and 128GB RAM running Oracle Linux 6.4. We use the commercial DBMS X (name omitted due to licensing restrictions) and Postgres as a backend (DBMS X is the default). Unless stated otherwise, each experiment was repeated 100 times (we stopped executions that ran longer than 10 minutes) and we report the median runtime. Computations that did not finish within the allocated time are omitted from the graphs.

Workloads. We compute explanations for the queries in Fig. 14. For DBLP datasets, we consider: **only2hop** (r_1) which is our running example query in this paper; **XwithYnotZ** (r_2) that returns authors that are direct co-authors of a certain person Y , but not of “Svein Johannessen”; **only3hop** (r_3) that returns pairs of authors (X, Y) that are connected via a path of length 3 in the co-author graph where X is not a co-author or indirect co-author (2 hops) of Y . For TPC-H, we consider: **ordPriority** (r_4) which returns for each customer the priorities of her/his orders; **ordDisc** (r_5) which returns customers and the discount rates of items in their orders; **partNotAsia** (r_6) which finds parts that can be supplied from a country that is not in Asia; **suppCust** (r_7) returns nations having both suppliers and customers.

Implementing DM. DM has to instantiate a graph with $\mathcal{O}(|\text{adom}(I)|^n)$ nodes where n is the maximal number of variables in a rule. We do not have a full implementation of DM, but compute a conservative lower bound for the runtime of the step constructing the game graph by executing a query (n -way cross-product over the active domain). Note that the actual runtime will be much higher because 1) several edges are created for each rule binding (we underestimate the number of

DBLP (#tuples)	100	1K	10K	100K
2 Variables (r_2)	0.043	0.171	14.016	-
3 Variables (r_1)	0.294	285.524	-	-
4 Variables (r_3)	56.070	-	-	-
TPC-H (Size)	10MB	100MB	1GB	10GB
> 10 Variables (r_4, r_5, r_6, r_7)	-	-	-	-

Fig. 15: Runtime of DM in seconds. For entries with ‘-’, the computation did not finish within 10 min.



(a) Runtime of only2hop

(b) Runtime of XwithYnotZ

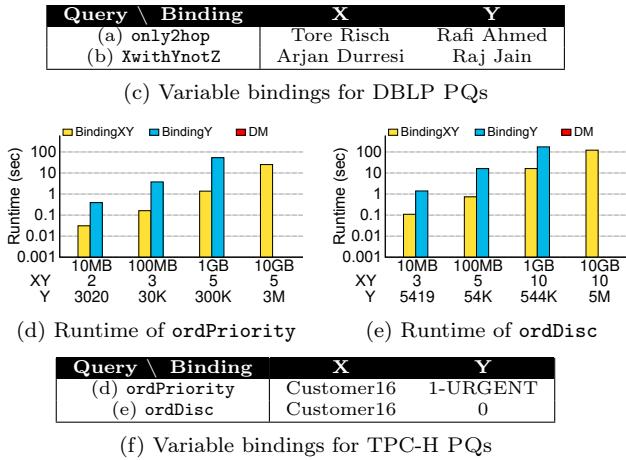
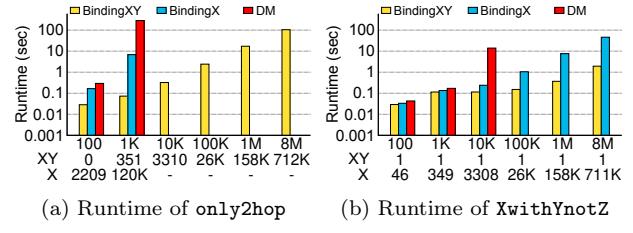


Fig. 16: Why questions: DBLP (top), TPC-H (bottom)

nodes of the constructed graph) and 2) recursive Datalog queries have to be evaluated over this graph using the well-founded semantics. The results for different instance sizes and number of variables are shown in Fig. 15. Even for only 2 variables, DM did not finish for datasets of more than 10K tuples within the allocated 10 min timeslot. For queries with more than 4 variables, DM did not even finish for the smallest dataset.

Why Questions. The runtime of generating explanations for why questions over the queries r_1, r_2, r_4 , and r_5 (Fig. 14) is shown in Fig. 16. For the evaluation, we consider the effect of different binding patterns on performance. Fig. 16c and 16f show which variables are bound by the provenance questions (PQs). Fig. 16a and 16b show the runtime for DBLP queries r_1 and r_2 , respectively. We also provide the number of rule nodes in the explanation for each binding pattern below the X axis. If only variable X is bound (BindingX), then the queries determine authors that occur together with the author we have bound to X in the query result. For instance, the explanation for `only2hop` with BindingX explains why persons are indirect, but not direct, co-authors of



(a) Runtime of only2hop

(b) Runtime of XwithYnotZ

Query \ Binding	X	Y
(a) only2hop	Tore Risch	Svein Johannessen
(b) XwithYnotZ	Tor Skeie	Joo-Ho Lee

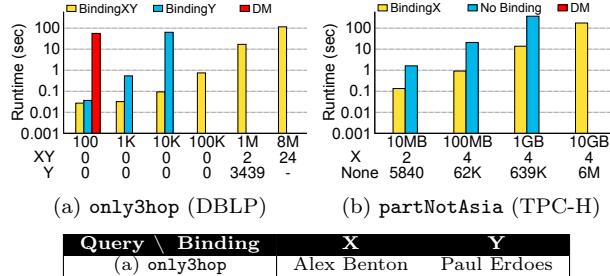
(c) Variable bindings for DBLP PQs

Fig. 17: Why-not questions over the DBLP dataset

“Tore Risch”. If both X and Y are bound (BindingXY), then the provenance for r_1 and r_2 is limited to a particular indirect and direct co-author, respectively. The runtime for generating explanations grows roughly linear in the dataset size and outperforms DM even for small instances. Furthermore, Fig. 16d and 16e (for r_4 and r_5 , respectively) show that our approach can handle queries with many variables (attributes in TPC-H) where DM times out even for the smallest dataset we have considered. Binding one variable (BindingY) in queries r_4 and r_5 expresses a condition, e.g., $Y = \text{'1-URGENT'}$ in r_4 requires the order priority to be urgent. If both variables are bound, then the PQ verifies the existence of orders for a certain customer (e.g., why “Customer16” has at least one urgent order). Runtimes exhibit the same trend as for the DBLP queries.

Why-not Provenance. We use queries r_1 and r_2 from Fig. 14 to evaluate the performance of computing explanations for failed derivations. When binding all variables in the PQ (BindingXY) using the bindings from Fig. 17c, these queries check if a particular set of authors do not appear together in the result. For instance, for `only2hop` (r_1), the query checks why “Tore Risch” is either not an indirect co-author or is a direct co-author of “Svein Johannessen”. The results for queries r_1 and r_2 (DBLP) are shown in Fig. 17a and 17b, respectively. The number of tuples produced by the provenance computation (the number of rule nodes is shown below the X axis) is quadratic in the database size resulting in a quadratic increase in runtime. DM only finishes within the allocated time for very small datasets while our approach scales to larger instances.

Queries with Negation. Recall that our approach also handles queries with negation. We choose rules r_3 (multiple negated goals) and r_6 (one negated goal) from Fig. 14 to evaluate the performance of answering why questions over such queries. We use the bindings shown in Fig. 18c. The results for r_3 and r_6 are shown in Fig. 18a and 18b, respectively. These results demonstrate that our approach efficiently computes explana-



¹ grcpi = ghost royal chocolate peach ivory

(c) Variable bindings for DBLP and TPC-H PQs

Fig. 18: Why questions for queries with negation

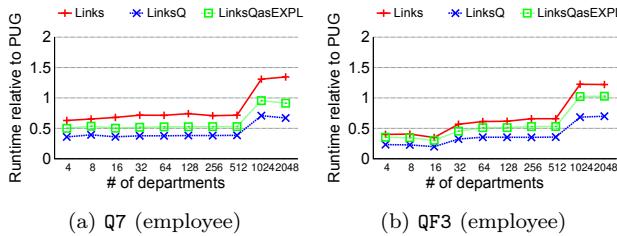
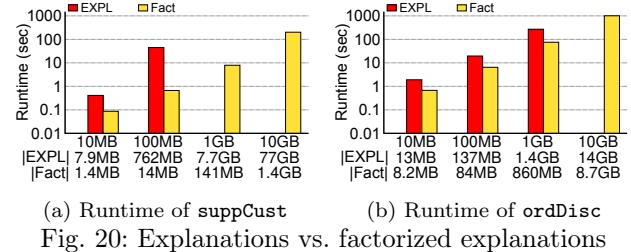


Fig. 19: Comparing Which(X) in PUG with Links

tions for such queries. When increasing the database size, the runtimes of PQs for these queries exhibit the same trend as observed for other why (why-not) questions and significantly outperform DM. For instance, the performance of partNotAsia (Fig. 18b), which contains many variables and negation exhibits the same trend as queries that have no negation (i.e., r_4 and r_5 in Fig. 16d and Fig. 16e, respectively).

Comparison with Links. In this experiment, we compare the runtime of computing EXPL $\text{Which}(X)$ (e.g., Fig. 4c) with computation of Lineage in Links^L from [9]. We show relative runtimes where PUG is normalized to 1. For this particular evaluation, we use Postgres as a backend since it is supported by both PUG and Links. Note that EXPL $\text{Which}(X)$ contains a full description of each tuple unlike Links^L which returns tuple identifiers (OIDs in Postgres). To get a nuanced understanding of the system’s performance, we show three runtimes for Links: 1) Links is the actual implementation in Links which computes Lineage (only OIDs) and where the runtime includes the construction of in-memory Links types from the provenance fetched from Postgres; 2) LinksQ is the runtime of the queries that Links uses to capture Lineage; and 3) LinksQasEXPL which joins the output of LinksQ with the base tables (i.e., as informative as EXPL $\text{Which}(X)$). We choose two queries from [9]. The query Q7 applies a range condition to the result of a two-way join. QF3 is a self-join on equality with an additional inequality condition (see [9] for more details). The queries are expressed over two tables `dept` and `emp`. The number of departments is varied from 4 to 2048 (by



(a) Runtime of suppCust (b) Runtime of ordDisc

Fig. 20: Explanations vs. factorized explanations

powers of 2 to replicate the setting from [9]), and each department has 100 employees on average. The relation `dept` consists of one attribute (department name), and `emp` has three attribute (department name, employee name, and salary). QF3 can be written in Datalog as:

$QF3(N1, N2) :- \text{emp}(_, D, N1, S), \text{emp}(_, D, N2, S), N1 \neq N2$

The runtimes of queries Q7 and QF3 are shown in Fig. 19a and 19b, respectively. Links performs better on smaller instances. The gap between Links and PUG shrinks with increasing dataset size. PUG outperforms Links and LinksQasEXPL on larger datasets.

Factorized Explanations. We now compare the performance of generating provenance for a query (EXPL) and a factorized representation of provenance (Fact) by rewriting the input query (Sec. 9). Factorization techniques perform best for many-to-many joins (e.g., the query r_7 in Fig. 14). The rewritten version of suppCust (r_7) producing factorized provenance is shown below.

```
r8 : suppCust(N) :- supp(N), cust(N)
r8' : supp(N) :- SUPPLIER(A, B, C, N, D, E, F)
r8'' : cust(N) :- CUSTOMER(G, H, I, N, J, K, L, M)
```

For this experiments, we use a 15 minute time-out. The runtimes for r_7 (yellow bars) and r_8 (red bars) are shown in Fig. 20a. We show the total result size in bytes below the X axis. The runtime of Fact grows roughly linear unlike EXPL whose growth is quadratic in dataset size. We also evaluate query r_5 which includes one-to-many joins to see how Fact performs for a query (Fig. 20b) where factorization only reduces size by a constant factor. This is confirmed by the measurements: the performance of Fact for r_5 is $\sim 30\%$ that of EXPL independent of dataset size.

12 Conclusions

We present a provenance model and unified framework for explaining answers and non-answers over first-order queries expressed in Datalog. Our efficient middleware implementation generates a Datalog program that computes the explanation for a provenance question and compiles this program into SQL. We prove that our model is expressive enough to encode a wide range of

provenance models from the literature and extend our approach to produce concise, factorized representations of provenance. In future work, we will investigate summarization of provenance (we did present a proof-of-concept in [28]) to deal with the large size of explanations for missing answers. We plan to also support query-based explanations [2, 3, 4, 40] and more expressive query languages (e.g., aggregation).

References

1. B. Arab, D. Gawlick, V. Radhakrishnan, H. Guo, and B. Glavic. A generic provenance middleware for database queries, updates, and transactions. In *TaPP*, 2014.
2. N. Bidoit, M. Herschel, and K. Tzompanaki. Immutably answering why-not questions for equivalent conjunctive queries. In *TaPP*, 2014.
3. N. Bidoit, M. Herschel, K. Tzompanaki, et al. Query-Based Why-Not Provenance with NedExplain. In *EDBT*, pages 145–156, 2014.
4. A. Chapman and H. V. Jagadish. Why Not? In *SIGMOD*, pages 523–534, 2009.
5. J. Cheney, L. Chiticariu, and W. Tan. Provenance in databases: Why, how, and where. *Foundations and Trends in Databases*, 1(4):379–474, 2009.
6. C. V. Damásio, A. Analyti, and G. Antoniou. Justifications for logic programming. In *Logic Programming and Non-monotonic Reasoning*, pages 530–542, 2013.
7. D. Deutch, A. Gilad, and Y. Moskovitch. Selective provenance for datalog programs using top-k queries. *VLDB*, 8(12):1394–1405, 2015.
8. D. Deutch, T. Milo, S. Roy, and V. Tannen. Circuits for datalog provenance. In *ICDT*, pages 201–212, 2014.
9. S. Fehrenbach and J. Cheney. Language-integrated provenance. *Science of Computer Programming*, 2017.
10. J. Flum, M. Kubierschky, and B. Ludäscher. Total and partial well-founded datalog coincide. In *ICDT*, pages 113–124, 1997.
11. B. Glavic, S. Köhler, S. Riddle, and B. Ludäscher. Towards constraint-based explanations for answers and non-answers. In *TaPP*, 2015.
12. B. Glavic, R. J. Miller, and G. Alonso. Using sql for efficient generation and querying of provenance information. In *In search of elegance in the theory and practice of computation*, pages 291–320. Springer, 2013.
13. E. Grädel and V. Tannen. Semiring provenance for first-order model checking. *arXiv preprint arXiv:1712.01980*, 2017.
14. T. Green. Containment of conjunctive queries on annotated relations. *Theory of Computing Systems*, 49(2):429–459, 2011.
15. T. Green, G. Karvounarakis, and V. Tannen. Provenance semirings. In *PODS*, pages 31–40, 2007.
16. T. J. Green, M. Aref, and G. Karvounarakis. Logicblox, platform and language: A tutorial. In *Datalog in Academia and Industry*, pages 1–8. Springer, 2012.
17. T. J. Green, G. Karvounarakis, Z. G. Ives, and V. Tannen. Update Exchange with Mappings and Provenance. In *VLDB*, pages 675–686, 2007.
18. T. J. Green and V. Tannen. The semiring framework for database provenance. In *PODS*, pages 93–99, 2017.
19. M. Herschel, R. Diestelkämper, and H. B. Lahmar. A survey on provenance: What for? what form? what from? *The VLDB Journal*, pages 1–26, 2017.
20. M. Herschel and M. Hernandez. Explaining Missing Answers to SPJUA Queries. *VLDB*, 3(1):185–196, 2010.
21. J. Huang, T. Chen, A. Doan, and J. Naughton. On the provenance of non-answers to queries over extracted data. In *VLDB*, pages 736–747, 2008.
22. G. Karvounarakis and T. J. Green. Semiring-annotated data: queries and provenance. *SIGMOD Record*, 41(3):5–14, 2012.
23. S. Köhler, B. Ludäscher, and Y. Smaragdakis. Declarative datalog debugging for mere mortals. In *Datalog 2.0: Datalog in Academia and Industry*, pages 111–122, 2012.
24. S. Köhler, B. Ludäscher, and D. Zinn. First-order provenance games. In *In Search of Elegance in the Theory and Practice of Computation*, pages 382–399. Springer, 2013.
25. S. Lee, S. Köhler, B. Ludäscher, and B. Glavic. Efficiently computing provenance graphs for queries with negation. Technical Report CoRR, arXiv:1701.05699, 2016.
26. S. Lee, S. Köhler, B. Ludäscher, and B. Glavic. A SQL-middleware unifying why and why-not provenance for first-order queries. In *ICDE*, pages 485–496, 2017.
27. S. Lee, B. Ludäscher, and B. Glavic. Pug: A framework and practical implementation for why and why-not provenance (extended version). Technical Report IIT/CS-DB-2018-02, Illinois Institute of Technology, 2018. <http://cs.iit.edu/%7edbgroup/pdfpubls/LL18.pdf>.
28. S. Lee, X. Niu, B. Ludäscher, and B. Glavic. Integrating Approximate Summarization with Provenance Capture. In *TaPP*, 2017.
29. A. Meliou, W. Gatterbauer, K. Moore, and D. Suciu. The Complexity of Causality and Responsibility for Query Answers and non-Answers. *VLDB*, 4(1):34–45, 2010.
30. A. Meliou, W. Gatterbauer, and D. Suciu. Reverse data management. *VLDB*, 4(12), 2011.
31. A. Meliou and D. Suciu. Tiresias: The database oracle for how-to queries. In *SIGMOD*, pages 337–348, 2012.
32. X. Niu, R. Kapoor, B. Glavic, D. Gawlick, Z. H. Liu, V. Krishnaswamy, and V. Radhakrishnan. Provenance-aware query optimization. In *ICDE*, pages 473–484, 2017.
33. D. Olteanu and J. Závodný. Factorised representations of query results: Size bounds and readability. In *ICDT*, pages 285–298. ACM, 2012.
34. D. Olteanu and J. Závodný. Size bounds for factorised representations of query results. *ACM Transactions on Database Systems (TODS)*, 40(1):2, 2015.
35. S. Riddle, S. Köhler, and B. Ludäscher. Towards constraint provenance games. In *TaPP*, 2014.
36. S. Roy, L. Orr, and D. Suciu. Explaining query answers with explanation-ready databases. *Proceedings of the VLDB Endowment*, 9(4):348–359, 2015.
37. S. Roy and D. Suciu. A formal approach to finding explanations for database queries. In *SIGMOD*, 2014.
38. P. Senellart. Provenance and probabilities in relational databases. *ACM SIGMOD Record*, 46(4):5–15, 2018.
39. V. Tannen. Provenance analysis for FOL model checking. *ACM SIGLOG News*, 4(1):24–36, 2017.
40. Q. T. Tran and C.-Y. Chan. How to conquer why-not questions. In *SIGMOD*, pages 15–26, 2010.
41. E. Wu and S. Madden. Scorpion: Explaining away outliers in aggregate queries. *VLDB*, 6(8):553–564, 2013.
42. Y. Wu, M. Zhao, A. Haeberlen, W. Zhou, and B. T. Loo. Diagnosing missing events in distributed systems with negative provenance. In *SIGCOMM*, pages 383–394, 2014.
43. J. Xu, W. Zhang, A. Alawini, and V. Tannen. Provenance analysis for missing answers and integrity repairs. *Data Engineering*, page 39, 2018.
44. W. Zhou, M. Sherr, T. Tao, X. Li, B. T. Loo, and Y. Mao. Efficient querying and maintenance of network provenance at internet-scale. In *SIGMOD*, pages 615–626, 2010.