

# Provenance for Large-scale Datalog

David Zhao <sup>†</sup>, Pavle Subotić<sup>\*</sup>, and Bernhard Scholz<sup>†</sup>

<sup>†</sup>School of Computer Science, The University of Sydney.

<sup>\*</sup>School of Computer Science, University College London.

July 11, 2019

## Abstract

Logic programming languages such as Datalog have become popular as Domain Specific Languages (DSLs) for solving large-scale, real-world problems, in particular, static program analysis and network analysis. The logic specifications which model analysis problems, process millions of tuples of data and contain hundreds of highly recursive rules. As a result, they are notoriously difficult to debug. While the database community has proposed several data provenance techniques that address the *Declarative Debugging Challenge* for Databases, in the cases of analysis problems, these state-of-the-art techniques do not scale.

In this paper, we introduce a novel bottom-up Datalog evaluation strategy for debugging: our provenance evaluation strategy relies on a new provenance lattice that includes proof annotations, and a new fixed-point semantics for semi-naïve evaluation. A debugging query mechanism allows arbitrary provenance queries, constructing partial proof trees of tuples with minimal height. We integrate our technique into Soufflé, a Datalog engine that synthesizes C++ code, and achieve high performance by using specialized parallel data structures. Experiments are conducted with DOOP/DaCapo, producing proof annotations for tens of millions of output tuples. We show that our method has a runtime overhead of  $1.27\times$  on average while being more flexible than existing state-of-the-art techniques.

## 1 Introduction

Datalog and other logic specification languages [24, 30, 5, 22] have seen a rise in popularity in recent years, being widely used to solve real-world problems including program analysis [24, 3], declarative networking [45, 23], security analysis [32] and business applications [5]. Logic programming provides declarative semantics for computation, resulting in succinct program representations and rapid-prototyping capabilities for scientific and industrial applications. Rather than prescribing the computational steps imperatively, logic specifications define the intended result declaratively and thus can express computations concisely. For example, logic programming has gained traction in the area of program analysis due to its flexibility in building custom program analyzers [24], points-to analyses for Java programs [8], and security analysis for smart contracts [21, 20].

Despite the numerous advantages, the declarative semantics of Datalog poses a debugging challenge. Strategies employed in debugging imperative programs such as inspecting variables at given points in the program execution do not translate to declarative programming. Logic specifications lack the notions of state and state transitions. Instead, they have relations that contain tuples. These relations may be input relations, such as those describing the instance of an analysis, intermediate relations, or output relations, such as those containing the results of an analysis. Relations can only be viewed in full, without any explanation of the origin or derivation of data, after the completion of a complicated evaluation strategy. Thus, the Datalog user will find the results alone of logic evaluation inconclusive for debugging purposes.

When debugging Datalog specifications, there are two main scenarios: (1) an unexpected output tuple appears, or (2) an expected output tuple does not appear. These may occur as a result of a fault in the input data, and/or a fault in the logic rules. Both scenarios call for mechanisms to explain how an output tuple is derived, or why the tuple cannot be derived from the input tuples. The standard mechanism for

these explanations is a *proof tree*. In the case of explaining the existence of an unexpected tuple, a proof tree describes formally the sequence of rule applications involved in generating the tuple. On the other hand, a failed proof tree, where at least one part of the proof tree doesn't hold, may explain why an expected tuple cannot be derived in the logic specification. These proof trees can be seen as a form of *data provenance* witness, that is, an explanation vehicle for the origins of data [9, 12].

In the presence of complex Datalog specifications and large datasets, Datalog debugging becomes an even bigger challenge. While recent developments in Datalog evaluation engines, such as Soufflé [24], have enabled the effective evaluation of complex Datalog specifications with large data using scalable bottom-up evaluation strategies [39, 34], unlike top-down evaluation, bottom-up evaluation does not have an explicit notion of a proof tree in its evaluation. Therefore, to facilitate debugging in bottom-up evaluation, state-of-the-art [27, 13, 26] techniques have been developed that rewrite the Datalog specification with provenance information. Using these techniques, Datalog users follow a *debugging cycle* which allows them to find anomalies in the input relations and/or the logic rules. In such setups, the typical debugging cycle comprises the phases of (1) defining an investigation query, (2) evaluating the logic specification to produce provenance witness, (3) investigating the faults based on the provenance information, and (4) fixing the faults. For complex Datalog specifications, the need for re-evaluation for each investigation is impractical. For example, DOOP [8] with a highly precise analysis setting may take multiple days to evaluate for medium to large-sized Java programs. Although state-of-the-art approaches scale for the database querying use cases, such approaches are not practical for industrial scale static analysis problems.

A further difficulty in developing debugging support for Datalog is providing understandable provenance witnesses. Use cases such as program analysis tend to produce proof trees of very large height. For example, investigations on medium sized program analyses in DOOP have minimal height proof trees of over 200 nodes. Therefore, a careful balance must be struck between enough information and readability in the debugging witnesses.

In this paper, we present a novel debugging approach that targets Datalog programs with characteristics of those found in static program analysis. Our approach scales to large dataset and ruleset sizes and provides succinct and interactively navigable provenance information.

The first aspect of our technique is a novel Datalog provenance evaluation strategy that augments the intensional database (IDB) with *Proof Annotations* and hence allows fast proof tree exploration for *all* debugging queries, *without the need for re-evaluation*. The exploration uses the proof annotations to construct proof trees for tuples *lazily*, i.e., a debugging query for a tuple produces the rule and the subproofs of the rule. The subproofs when expanded in consecutive debugging queries, will produce a *minimal* height proof tree for the given tuple. Our system also supports *non-existence* explanations of a tuple. In this case, proof annotations are not helpful since they cannot describe non-existent tuples. Thus, we adapt an approach from [28] to provide a *user-guided* procedure for explaining the non-existence of tuples.

We implement the provenance evaluation strategy in the synthesis framework of Soufflé [24], to produce *specialized* data structures and an interactive debugging query system for each logic specification. Our approach is tightly integrated into the Soufflé engine, thus achieving high performance and generalizability that no previous provenance approach is able to achieve. We demonstrate the feasibility of our technique through the complex Java points-to framework, DOOP, running the Java DaCapo benchmark suite, which produces tens of millions of output tuples. We demonstrate that the initial implementation of our novel provenance method incurs a runtime overhead of 1.27 $\times$ , and memory consumption overhead of 1.45 $\times$  on average. Thus, our provenance evaluation strategy is capable of processing large datasets with no performance disadvantage compared to existing techniques, while being more flexible for answering debugging queries.

Our contributions in this work are as follows:

- a provenance evaluation strategy for Datalog specifications, defining a new evaluation domain based on a provenance lattice which extends the standard Datalog subset lattice with proof annotations,
- the leveraging of parallel bottom-up evaluation to give minimal height proof trees, and provenance queries for constructing minimal height proof trees utilizing proof annotations, allowing effective bug investigation with a minimum number of user interactions,
- an efficient and scalable integration of the proof tree generator system into Soufflé, using specialized data structures for storing proof annotations, and

|                  |                 |  |
|------------------|-----------------|--|
| 11: a = new 0(); | new(a, l1).     | r1: vpt(Var, Obj) :- new(Var, Obj).  |
| 12: b = a;       | assign(b, a).   | r2: vpt(Var, Obj) :- assign(Var, Var2),<br>vpt(Var2, Obj).                               |
| 13: c = new P(); | new(c, l3).     | r3: vpt(Var, Obj) :- load(Var, Y, F),<br>store(P, F, Q),<br>vpt(Q, Obj),<br>alias(P, Y). |
| 14: d = new P(); | new(d, l4).     |  |
| 15: c.f = a;     | store(c, f, a). | r4: alias(Var1, Var2) :- vpt(Var1, Obj),<br>vpt(Var2, Obj),<br>Var1 != Var2.             |
| 16: e = d.f;     | load(e, d, f).  |  |
| 17: b = c.f;     | load(b, c, f).  |  |
| 18: a = b;       | assign(a, b).   |  |

(a) Input Program

(b) EDB Tuples

(c) Datalog Points-to Analysis

Figure 1: Program Analysis Datalog Setup

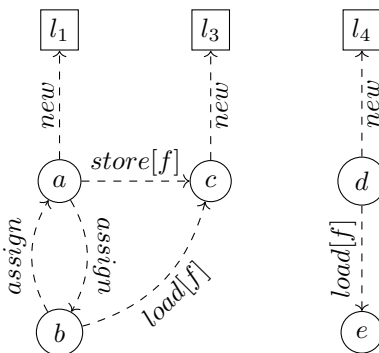


Figure 2: Points-to Input Diagram

- large-scale experiments using the DOOP program analysis framework with DaCapo benchmarks with tens of millions of tuples, measuring on average  $1.27\times$  overheads for runtime and  $1.45\times$  overheads for memory.

The paper is organized as follows: In Section 2 we motivate our provenance method and describe its use in a real-world program analysis use case. In Section 3 we detail the theoretical basis of our method with regards to the provenance evaluation strategy along with the provenance queries to construct proof trees for tuples. We also demonstrate the minimality properties and present the practical solution that results from this theory. In Section 4 we detail the implementation of our system in Soufflé. In Section 5 we present experiments that show the feasibility of our provenance system. In Section 6 we outline related work, and we conclude in Section 7.

## 2 Motivation and Problem Statement

A common approach to characterize the evaluation of a Datalog specification is through *proof trees*. A proof tree for a tuple describes the derivation of that tuple from input tuples and rules. During the debugging cycle, proof trees *explain* why an output tuple exists, and therefore are critical for any investigation into anomalies. Note that potentially there could be an infinite number of proof trees for the explanation of any given tuple. However, end users desire concise proof trees such that that the faulty behaviour of the logic specification is revealed quickly. In this section, we describe how proof trees can be used to debug a Datalog specification and an overview of our method for generating minimal proof trees for output tuples.

## 2.1 Use Case: Program Analysis

### 2.1.1 Points-To Analysis

We illustrate the utility of debugging via proof trees through a program analysis use case. Figure 1 illustrates a points-to analysis implemented in Datalog. The points-to analysis resembles a field-sensitive but flow-insensitive analysis [36]. The input relations (also known as EDB) of the logic specification are the relations *new*, *assign*, *load*, and *store* express the input program in relational form. The relation *new* represent the object-creation sites of the input program, the relation *assign* the assignments, and relations *load/store* the read and write accesses of objects via a field. Figure 1a shows an input program encoded in the form of input relations in Figure 1b. The graph in Figure 2 represents the input relations. The nodes represent either object-creation sites or variables. The edges are object-creation sites, assignments, and load/store instructions. The graph shows a clear separation between objects of  $l_4$  with objects  $l_1$  and  $l_3$ . The goal is to compute the var-points-to set in the form of the output relation *vpt*. The Datalog rules computing the var-points-to set are given in Figure 1c. The first rule makes the variable *Var* point to object *Obj* where *Obj* is the line number of the object-creation site as an abstraction for all possible objects that could be created by this object-creation site. The second rule shows the transfer of the var-points-to set from source *Var2* of the assignment to its destination *Var*. The third rule transfers the var-points-to set from the source of a store instruction *Q* to the destination of a load instruction *Var*. The transfer is conditional depending on whether field *F* of the load and store instructions match and whether the instance variable *Y* of the load and the instance variable *P* of the store instruction alias. The last rule expresses the alias relation between variables *Var1* and *Var2*, i.e., two variables alias if they share at least one object-creation site *Obj* in their var-points-to sets. The relations *vpt* and *alias* are the output of the analysis and are called the IDB of the Datalog specification.

### 2.1.2 Minimal Height Proof Trees

The analysis example in Figure 1 computes the output relation *alias* that captures the alias information of two variables. A user may investigate *why* a tuple  $(a, b)$  exists in the output relation *alias*, i.e., how the analysis derives  $alias(a, b)$  from the input data via the rules. Intuitively, this information is contained in the points-to input diagram (cf. Figure 2) showing that variables  $a$  and  $b$  may reach the same object. However, it is not an explanation, as a proof tree would be, for the tuple  $alias(a, b)$  as shown in Figure 3. The proof tree shows that  $alias(a, b)$  is derived by rule  $r_4$  using the facts  $vpt(a, l_1)$  and  $vpt(b, l_1)$ . This outcome is expected since it tells us that  $a$  and  $b$  point to the same object ( $l_1$  in this case), and thus they may alias.

$$\frac{\frac{new(a, l_1)}{vpt(a, l_1)} r_1 \quad \frac{assign(b, a) \quad \frac{new(a, l_1)}{vpt(a, l_1)} r_1}{vpt(b, l_1)} r_2 \quad a \neq b}{alias(a, b)} r_4$$

Figure 3: Full proof tree for  $alias(a, b)$

The importance of minimality of proof tree height is shown in Figure 3, which depicts the proof tree resulting from the assignment in line  $l_2$  in the input program. In the input program, there is a circular assignment in lines  $l_2$  and  $l_8$  caused by the flow-insensitivity of the input program, and thus the tuple  $vpt(b, l_1)$  could be derived in an arbitrary number of rule applications, as shown in Figure 4.

Thus, even for this small example, there are infinitely many valid proof trees for the tuple  $alias(a, b)$ . A provenance system ought to produce the most concise proof tree so that an end user can understand the derivation of a tuple with the least effort.

### 2.1.3 Proof Tree Fragments for Debugging

Suppose a Datalog user discovers an unexpected tuple in the output, which indicates that a fault exists somewhere in the logic specification. The aim is to investigate the root cause of this fault. Since proof trees

$$\frac{\text{assign}(a, b)}{\frac{\text{assign}(b, a) \quad \frac{\dots}{\text{vpt}(b, l_1)} r_2}{\text{vpt}(a, l_1)} r_2} \text{vpt}(b, l_1)$$

Figure 4: Infinitely many derivations for  $\text{vpt}(b, l_1)$ , resulting from the circular assignment in lines  $l_2$  and  $l_8$  in the input program

provide explanations for the existence of a tuple, the proof tree of an unexpected tuple will help identify the fault in the logic specification.

An example fault could be if rule  $r_3$  was altered as follows,

```

r3: vpt(Var, Obj) :- load(Var, Y, F),
                    store(P, F, Q),
                    vpt(Q, Obj),
                    vpt(P, Obj1),
                    vpt(Y, Obj2).

```

Note the condition that objects P and Y must alias now no longer holds. A minor typo may have introduced this fault, and as a consequence of this typo, the analysis produces the extra tuple  $(a, e)$  in relation *alias*. This additional tuple becomes a *symptom* of the fault. To diagnose this fault, the proof tree of tuple *alias*( $a, e$ ) highlights the root cause of the fault.

However, in practice, a full proof tree may be too large to provide a meaningful explanation even if it is of minimal height, and as experiments in Section 5 show, proof trees for real-world program analyses (e.g., DOOP) can exceed heights of 200. Thus a Datalog user may want to explore only relevant *fragments* of it interactively. A fragment of a proof tree is a partial subtree, which consists of some number of levels. For instance, we may construct fragments comprising of 2 levels to explore only parts of the proof tree that are relevant.

We illustrate the exploration of fragments of the proof tree in Figure 5. In the figure, tuple  $t$  denotes the symptom of the fault, i.e.,  $t$  is an unexpected tuple in the output. The aim is to explore the proof tree for  $t$  to find the root cause for this fault. In our example, the user follows the scent of the fault by expanding proof tree fragments that show anomalies. This process produces a path of exploration in the proof tree. The path of exploration discovers the root cause of the fault efficiently, without constructing and displaying the full proof tree of an output tuple.

Concretely, we may wish to explain the tuple *alias*( $a, e$ ). Figure 6 illustrates the exploration of an explanation for *alias*( $a, e$ ) by generating proof tree fragments of 2 levels at a time. The user generates the first fragment and decides that  $\text{vpt}(e, l_1)$  is the most relevant explanation for the fault, and continues down this path. As a result, the root cause (for example, the erroneous rule  $r_3$ ) is discovered after two fragments. This interaction mechanism also justifies the choice to minimize the height of proof trees. By doing this, we

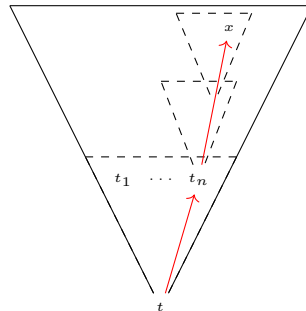


Figure 5: Interactive exploration of fragments of a proof tree for  $t$

$$\frac{\text{load}(e, d, f) \quad \text{store}(c, f, a) \quad \text{vpt}(a, l_1) \quad \text{vpt}(c, l_3) \quad \text{vpt}(d, l_4)}{\text{vpt}(e, l_1)} r_3$$

$$\frac{\text{vpt}(a, l_1) \quad \text{vpt}(e, l_1) \quad a \neq e}{\text{alias}(a, e)} r_4$$

Figure 6: Exploring the proof of  $\text{alias}(a, e)$  to find the erroneous rule  $r_3$

minimize the number of user interactions (i.e., proof tree fragments) required to discover the root cause for an anomaly.

Note that provenance queries for databases (cf. [13, 26, 27]) have in general shorter evaluation times, and, hence, Database approaches may re-evaluate for each provenance query the whole specification. However, for large and complex Datalog specifications, provenance queries as shown before are to be performed interactively and promptly, i.e., in a single debugging cycle. The Datalog user may want to issue multiple provenance queries in a single investigation phase of a debugging cycle. Hence, a re-evaluation of the whole logic specification for each provenance query becomes prohibitive, and a new provenance approach for Datalog is required.

## 2.2 Proof Trees and Problem Statement

We use standard terminology for Datalog, taken from [2]. A Datalog specification  $P$  consists of a set of *rules*, of the form  $R_0(X_0) :- R_1(X_1), \dots, R_n(X_n), \psi(X_1, \dots, X_n)$ . Each  $R_i(X_i)$  is a *predicate*, consisting of a *relation name*  $R_i$  and an argument  $X_i$  consisting of the correct number of variables and constants. The term  $\psi(X_1, \dots, X_n)$  denotes a conjunction of constraints on the variables in the rule. These constraints may include, for example, arithmetic constraints (such as less than), or negation of a predicate. A predicate can be *instantiated* to form a *tuple* where each variable is mapped to a constant. An instantiated rule is a rule with each predicate replaced by its instantiation such that the variable mappings are consistent between predicates and the constraints are satisfied. A set of tuples forms an instance  $I$ , and we denote the input instance to be  $EDB$ .

Given a Datalog specification  $P$ , an input instance  $EDB$  of  $P$ , and a tuple  $t$  produced by  $P$ , we want to find a *proof tree* of minimal height for  $t$ . We define a proof tree as follows:

**Definition 1** (Proof Tree). *Let  $P$  be a Datalog specification, and let  $EDB$  be an input instance. A proof tree  $\tau_t$  for a tuple  $t$  computed by  $P$  is a labeled tree where (1) each vertex is labeled with a tuple, (2) each leaf is labeled with an input tuple in  $EDB$ , (3) the root is labeled with  $t$ , and (4) for a vertex labeled with  $t_0$ , there is a valid instantiation  $t_0 :- t_1, \dots, t_n$  of a rule  $\rho$  in  $P$  such that the direct children of  $t_0$  are labeled with  $t_1, \dots, t_n$ . Moreover, the vertex is associated with  $\rho$ .*

A proof tree for  $t$  can be viewed as an *explanation* for the existence of  $t$ , by showing how it is derived from other tuples using the rules in the Datalog specification. Importantly, in the context of Datalog, there is a strong connection between a proof tree and the *model* of the specification. A model corresponds to a set of tuples that satisfy a Datalog specification when it is viewed as a constraint system. In a proof tree, each node corresponds to exactly one tuple from the model since each instantiation for the body of a rule generates exactly one tuple.

To formalize the problem statement, we need to characterize proof trees of minimal height. Note that the set of proof trees for a Datalog specification could be constructed inductively by the height of the trees. We denote  $\tau_t$  to be a proof tree for tuple  $t$ , and  $T^k$  to be the set of proof trees of height at most  $k$ . This construction leads to a convenient description of what it means for a proof tree to be of minimal height.

**Definition 2.** *We define the set of all proof trees inductively. Let  $T^0 = \{\tau_t \mid t \in EDB\}$  be the set of proof trees for tuples in the input instance. Then, define  $T^k$  in terms of  $T^{k-1}$ :  $T^k = \{\tau_t \mid t :- t_1, \dots, t_n \text{ is a valid instantiation and } \forall t_i : \exists \tau_{t_i} \in T^{k-1}\}$ . Then,  $T = \bigcup_{i \geq 0} T^i$  is the set of all proof trees produced by the specification  $P$ .*

Note that each  $T^k$  consists of proof trees of height at most  $k$  since if  $t :- t_1, \dots, t_n$  is an instantiation of a rule, then the height of the proof tree for  $t$  is equal to the maximum height of the proof trees for  $t_1, \dots, t_n$  plus 1. By defining the set of proof trees inductively, a proof tree of minimal height for a given tuple  $t$  has height given by

$$\min \{k \geq 0 \mid \exists \tau_t \in T^k\}.$$

Intuitively, this means that a proof tree for a tuple  $t$  is of minimal height if there does not exist another valid proof tree with a smaller height. We emphasize that a valid proof tree must exist since we have assumed that tuple is in the IDB of the Datalog specification and therefore its existence can be proved. Based on this inductive construction of proof trees, we reduce the problem of generating a fragment of a proof tree into the following incremental search problem.

**Problem statement:** Let  $P$  be a Datalog specification, and  $I$  be the instance computed by  $P$ . Then, given a tuple  $t \in I$ , find the tuples  $t_1, \dots, t_n$  that form the direct children of  $t$  in a minimal height proof tree for  $t$ , as depicted in Figure 7. Denote  $t_1, \dots, t_n$  to be a *configuration* of the body of the corresponding rule.

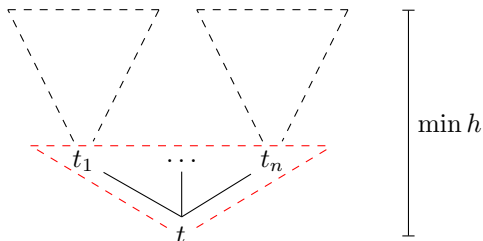


Figure 7: One level of a proof tree of minimal height for  $t$

Note that we could recursively construct the subtrees rooted at each  $t_i$ , which form valid proof trees for these tuples. Thus, this recursive construction solves the original problem of constructing a fragment of the proof tree of minimal height. Once a certain number of levels have been constructed, or if the only remaining leaves are in the EDB (characterized by having a proof tree consisting of only a single node), then the fragment is complete.

### 3 A New Provenance Method

A simple solution to this problem might be to generate a minimal height proof tree by brute-force searching for matching tuples and return the direct children of the root node. However, this is an unfeasible approach for real-world problems where there may be millions of tuples, and there are also no guarantees that the produced proof trees are of minimal height. Moreover, the two main evaluation strategies for Datalog, *bottom-up* and *top-down* are unsuitable for solving this problem on their own. Bottom-up evaluation is an efficient method for generating tuples but does not store any information related to proof trees. On the other hand, top-down evaluation does compute proof trees as part of its execution, but there are no guarantees for minimality of height. Additionally, to prove the existence of a particular tuple requires proving the existence of every intermediate tuple up to the input tuples, and thus the problem of generating only fragments of proof trees cannot be solved by top-down. Thus, we present a hybrid solution for generating proof trees, consisting of a provenance evaluation strategy based on bottom-up evaluation, plus a debugging query mechanism to construct proof trees.

We summarize the system in Figure 8. The Datalog specification and input tuples (EDB) are the input into the system. A pre-processing step of provenance Datalog evaluation generates a set of tuples (IDB) alongside proof annotations for these tuples. The annotated tuples form the input into the interactive proof tree generator system.

The proof tree generator is at the core of the interactive exploration of proofs for tuples. A user queries for a fragment of a proof tree, e.g., *two levels of a proof tree for  $vpt(b, l1)$* , and the system returns the

corresponding result. This system can answer any number of queries, and the user can query for *any* fragment of the proof tree for *any* tuple. As previously mentioned, this allows the user to interactively explore the proof for a tuple and find a meaningful explanation for a tuple.

The provenance evaluation strategy resembles a pre-computation step of the Datalog. The evaluation is performed only *once*, but the IDB with proof annotations can subsequently answer *any* debugging query using the same IDB resulting from evaluation. The ability to answer any debugging query without re-evaluation is an advantage over other selective provenance systems [13, 27], where the query is given prior to evaluation, which is then instrumented based on the query, and thus evaluation must be performed for each different query.

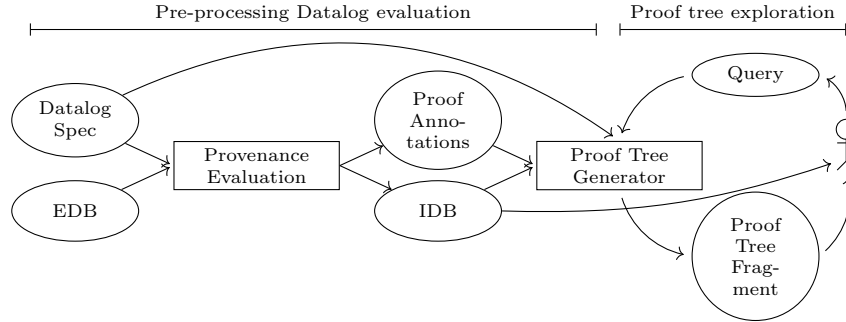


Figure 8: Synthesized Proof Tree Generator system

### 3.1 Standard Bottom-Up Evaluation

The basis of our approach is the standard bottom-up evaluation strategy for Datalog specification [2]. The computational domain of standard bottom-up evaluation is the subset lattice consisting of sets of tuples, denoted *instances*  $I$ . The *naïve* algorithm for evaluation is based on the immediate consequence operator,  $\Gamma_P$ , which generates new tuples by applying rules in the Datalog specification to tuples in the current instance.

$$\Gamma_P(I) = I \cup \{t \mid t :- t_1, \dots, t_n \text{ is a valid instantiation of a rule in } P \text{ with each } t_i \in I\}$$

The result of Datalog evaluation is attained when  $\Gamma_P$  reaches a fixpoint, i.e., when  $\Gamma_P(I) = I$ . Note that this evaluation appears closely related to the inductive construction of proof trees, and indeed the set of tuples represented by  $T^i$  is equal to the set of tuples generated by the  $i$ -th application of  $\Gamma_P$ .

However, this naïve evaluation will repeat computations since a tuple computed in some iteration will then be recomputed in every subsequent iteration. Therefore, the standard implementation of bottom-up evaluation in real systems such as [24, 40] is *semi-naïve*. Semi-naïve evaluation contains two main optimizations over naïve bottom-up evaluation:

1. **Precedence graph optimization:** the Datalog specification is split into *strata*. Firstly, a precedence graph of relations is computed, then each strongly connected component of the precedence graph forms a stratum. Each stratum is evaluated in a bottom-up fashion as a separate fixpoint computation in order based on the topological order of SCCs. The input to a particular stratum is the output of the previous strata in the precedence graph.
2. **New knowledge optimization:** within a single stratum, the evaluation is optimized in each iteration by considering the new tuples generated in the previous iteration. A new tuple is generated in the current iteration only if it directly depends on tuples generated in the previous iteration, therefore avoiding the recomputation of tuples already computed in prior iterations. We describe this process in further detail in Section 4.1

With these two optimizations, semi-naïve performs less repeated computations than the naïve algorithm, however, our method for generating proof trees must now be tailored to semi-naïve evaluation.



$$vpt(b, l_1) \quad \leftrightarrow \quad (vpt(b, l_1), 2) \quad \leftrightarrow \quad \frac{\text{assign}(b, a) \quad \frac{\text{new}(a, l_1)}{vpt(a, l_1)} r_1}{vpt(b, l_1)} r_2 \quad \Bigg] 2$$

Figure 9: Connecting a tuple to a proof tree via a height annotation

Another essential extension of Datalog is negation, and the standard semantics for negated Datalog is *stratified negation* [2]. A negated predicate is denoted with a ! symbol, for example  $!vpt(Var, Obj)$  denotes the negation of  $vpt(Var, Obj)$ . Semantically, a negated predicate evaluates to true if no matching tuples exist in the instance. With stratified negation semantics, a negated predicate is only allowed if the contained variables exist in positive predicates elsewhere in the body of the rule (a condition also known as *groundedness*), and if the corresponding relation does not appear in a cycle in the precedence graph. During evaluation, the stratification of the precedence graph is carried out in a way such that the negated relations can be treated as input into a stratum, and a negated predicate is treated as a constraint, which holds if no corresponding tuple exists in the input instance.

### 3.2 Provenance Evaluation Strategy

These standard bottom-up evaluation semantics are extended to compute a minimal height proof tree for each tuple. Our extended semantics store *proof annotations* alongside the original tuples. In particular, for each tuple, the annotations are the *height* of the minimal height proof tree, and a number denoting the *rule* which generated the tuple. By using this extra information, we can efficiently generate minimal height proof trees to answer provenance queries (see Section 3.3).

In the context of semi-naïve evaluation, and in particular the precedence graph optimization, we describe the provenance evaluation strategy here for a single fixpoint computation (i.e., a single stratum). The resulting correctness properties translate directly to the evaluation of the full Datalog specification since correctness holds for every stratum in the evaluation.

The rule number annotation is easily computed during bottom-up evaluation. With bottom-up evaluation, a new tuple  $t$  is generated if there is a rule  $\rho_k : R(X) :- R_1(X_1), \dots, R_n(X_n), \psi(X_1, \dots, X_n)$  and a set of tuples  $t_1, \dots, t_n$  such that  $t :- t_1, \dots, t_n$  forms a valid instantiation of the above rule. If this is the case, we say that the rule firing of  $\rho_k$  generates  $t$ , and thus the identifier  $\rho(t) = k$  is stored as the rule number annotation for  $t$ . In this way, for each tuple, we track which rule is fired to generate that tuple.

However, the height annotations are more involved and relate closely to the semantics of bottom-up evaluation. Thus, we must develop a formalism for the height annotations, to ensure that it correctly computes the height of the minimal height proof tree for each tuple. To formalize tuples with height annotations, we define a *provenance lattice* as our domain of computation, which extends the standard subset lattice with proof annotations. An element of the provenance lattice is a provenance instance.

**Definition 3** (Provenance Instance). *A provenance instance is an instance of tuples  $I$  along with a function*

$$h : I \rightarrow \mathbb{N}$$

*which provides a height annotation of each tuple in the instance. We denote a provenance instance to be the pair  $(I, h)$ .*

The aim of these height annotations is to connect a tuple to its proof tree, as depicted in Figure 9. The middle value is a tuple along with its height annotation, which is an example of an augmented tuple in a provenance instance. The corresponding proof tree on the right has height matching this annotation.

Similar to the subset lattice of standard bottom-up evaluation, the domain of provenance evaluation should also form a lattice, in our case, based on the subset lattice of standard bottom-up evaluation, but with elements being provenance instances rather than instances. We denote this to be the *provenance lattice*  $\mathcal{L}$ , where elements are provenance instances. The ordering  $\sqsubseteq$  of elements in the lattice is defined by:

$$(I_1, h_1) \sqsubseteq (I_2, h_2) \iff I_1 \subseteq I_2 \text{ and } \forall t \in I_1 : h_1(t) \geq h_2(t)$$

Intuitively, this ordering specifies that an augmented instance  $(I_1, h_1)$  is ‘less than’ another augmented instance  $(I_2, h_2)$  if all tuples in  $I_1$  also appear in  $I_2$ , with larger or equal height annotation. Therefore, moving ‘up’ the lattice towards the top element results in augmented instances with more tuples and smaller height annotations. This property guarantees the minimality of these height annotations since a bottom-up Datalog evaluation is equivalent to applying a monotone function to move ‘up’ a lattice.

The property that  $\sqsubseteq$  a valid partial order is essential to demonstrate that standard properties of Datalog evaluation hold.

**Lemma 1.**  $\sqsubseteq$  is a partial order

*Proof.* We show that each of the three properties of a partial order hold.

- Reflexivity:  $(I, h) \sqsubseteq (I, h)$  since  $I \subseteq I$  and  $\forall t \in I : h(t) = h(t)$
- Anti-symmetry: If  $(I_1, h_1) \sqsubseteq (I_2, h_2)$  and  $(I_2, h_2) \sqsubseteq (I_1, h_1)$ , then  $I_1 \subseteq I_2$  and  $I_2 \subseteq I_1$  (so  $I_1 = I_2$ ), and  $\forall t \in I_1 : h_1(t) \geq h_2(t)$  and  $\forall t \in I_2 : h_2(t) \geq h_1(t)$ . Since  $I_1 = I_2$ , we conclude that  $\forall t \in I_1 : h_1(t) = h_2(t)$ .

Thus,  $(I_1, h_1) = (I_2, h_2)$ .

- Transitivity: If  $(I_1, h_1) \sqsubseteq (I_2, h_2)$  and  $(I_2, h_2) \sqsubseteq (I_3, h_3)$ , then we have  $I_1 \subseteq I_2 \subseteq I_3$ . Also, we have  $\forall t \in I_1 : h_1(t) \geq h_2(t)$ , and  $\forall t \in I_2 : h_2(t) \geq h_3(t)$ . Since  $I_1 \subseteq I_2$ , we have  $\forall t \in I_1 : h_1(t) \geq h_2(t) \geq h_3(t)$ .

Therefore, we have  $I_1 \subseteq I_3$  and  $\forall t \in I_1 : h_1(t) \geq h_3(t)$ , and thus  $(I_1, h_1) \sqsubseteq (I_3, h_3)$ .

Thus, this demonstrates that  $\sqsubseteq$  is a partial order.  $\square$

In a similar fashion to the immediate consequence operator  $\Gamma_P$  operating on the subset lattice of Datalog instances, provenance evaluation is achieved with a consequence operator  $\mathcal{T}_P$  operating on the provenance lattice. The result of evaluation is reached when  $\mathcal{T}_P$  reaches a fixpoint, i.e., when  $\mathcal{T}_P((I, h)) = (I, h)$ . The main property  $\mathcal{T}_P$  is that once a fixpoint has been reached, the proof tree height annotations are *minimal*, and they correspond to the heights of the smallest height proof trees.

The consequence operator  $\mathcal{T}_P$  is defined in terms of the  $\Gamma_P$  operator:

**Definition 4** (Consequence operator). *The consequence operator,  $\mathcal{T}_P$ , generates a new provenance instance:*

$$\mathcal{T}_P((I, h)) = (\Gamma_P(I), h')$$

where  $h'$  is defined as follows. For any tuple  $t \in \Gamma_P(I)$ , let

$$G_t = \{(t_1, \dots, t_n) \mid t :- t_1, \dots, t_n \text{ is a valid rule instantiation with each } t_i \in \Gamma_P(I)\}$$

be the set of all configurations of rule bodies generating  $t$ . Note this may be empty in the case of EDB tuples. Then,

$$h'(t) = \begin{cases} h(t) & \text{if } G_t = \emptyset \\ \min_{g \in G_t} \{\max_{t_i \in g} \{h(t_i)\} + 1\} & \text{otherwise} \end{cases}$$

The generation of new tuples behaves in the same way as  $\Gamma_P$ . To illustrate the height annotations, consider the rule instantiation  $vpt(b, l_1) :- (assign(b, a), 0), (vpt(a, l_1), 1)$ , with height annotations written alongside body tuples for convenience. From this rule instantiation, we would generate the tuple  $vpt(b, l_1)$  with height annotation  $\max(0, 1) + 1 = 2$ . However, the instantiation  $vpt(b, l_1) :- (load(b, c, f), 0), (store(c, f, a), 0), (vpt(a, l_1), 1), (alias(c, \dots), 1)$ , would also generate  $vpt(b, l_1)$ , but with height annotation  $\max(0, 0, 1, 2) + 1 = 3$ . The resulting instance after applying  $\mathcal{T}_P$  will contain only the smaller annotation, and thus the resulting provenance tuple is  $(vpt(b, l_1), 2)$ .

Also, note that this semantics allow for the *update* of the height annotation for a tuple  $t \in I$ . If  $\mathcal{T}_P(I, h) = (\Gamma_P(I), h')$  results in  $h'(t) < h(t)$ , then we say that the height annotation of  $t$  has been updated. An update may happen if  $\mathcal{T}_P$  generates new tuples which form a valid configuration of a rule body generating  $t$ , with lower height annotations than a previous derivation.

We illustrate this definition of provenance evaluation strategy using the running example. We denote  $(t, h)$  to be a tuple  $t$  with height annotation  $h$ . Assume, for the sake of illustration, that the input instance is computed from a previous fixpoint, and thus has different height annotations, which may be the case with a single fixpoint in semi-naïve evaluation. Note that for the full Datalog specification, all input tuples will have a height annotation of 0. Figure 10 shows the derived relations under the fixpoint computation with the provenance evaluation strategy. Importantly, in iteration 3, the height annotation for  $vpt(b, l_1)$  is updated as a result of a new derivation using  $\mathbf{b} = \mathbf{c.f}$ ;  $\mathbf{c.f} = \mathbf{a}$ ; is computed, with lower height annotation than the previous derivation which used  $\mathbf{b} = \mathbf{a}$ ;. This demonstrates that annotations of tuples may be updated after they are initially computed, which is essential to maintain minimality.

**Input:**

$$\{(new(a, l_1), 0), (assign(b, a), 6), (new(c, l_3), 0), (new(d, l_4), 0), \\ (store(c, f, a), 0), (load(e, d, f), 0), (load(b, c, f), 0), (assign(a, b), 0)\}$$

**Fixpoint iterations:**

$$\begin{aligned} i_0 &: \emptyset \\ i_1 &: \{(vpt(a, l_1), 1), (vpt(c, l_3), 1), (vpt(d, l_4), 1)\} \\ i_2 &: \{(vpt(a, l_1), 1), (vpt(c, l_3), 1), (vpt(d, l_4), 1), (vpt(b, l_1), 7)\} \\ i_3 &: \{(vpt(a, l_1), 1), (vpt(c, l_3), 1), (vpt(d, l_4), 1), (vpt(b, l_1), 3)\} \\ i_4 &: \{(vpt(a, l_1), 1), (vpt(c, l_3), 1), (vpt(d, l_4), 1), (vpt(b, l_1), 3)\} \end{aligned}$$

Figure 10: IDB relation  $vpt$  in each iteration of the fixpoint computation for the example Datalog specification

It remains to be shown that the provenance evaluation strategy is correct, i.e., that  $\mathcal{T}_P$  terminates and results in the same set of tuples as  $\Gamma_P$ . Additionally, we must show that the height annotations resulting from provenance evaluation strategy is minimal.

**Lemma 2.**  $\mathcal{T}_P$  computes the same tuples as  $\Gamma_P$  at fixpoint, i.e.

1.  $\exists k$  s.t.  $\mathcal{T}_P(\mathcal{T}_P^k((I, h))) = \mathcal{T}_P^k(I, h)$ , and
2.  $\mathcal{T}_P^k(I, h) = (\Gamma_P^k(I), h^k)$  for some level annotation function  $h^k$

*Proof.* By definition,  $\mathcal{T}_P$  generates tuples in the same fashion as  $\Gamma_P$ . Since  $\Gamma_P$  always reaches a fixpoint, say after  $l$  iterations, i.e.,  $\Gamma_P(\Gamma_P^l(I)) = \Gamma_P^l(I)$ , we have

$$\mathcal{T}_P^l((I, h)) = (\Gamma_P^l(I), h^l)$$

Any further applications of  $\mathcal{T}_P$  do not change the set of tuples since  $\Gamma_P$  has already reached a fixpoint. Thus, after  $l$  iterations,  $\mathcal{T}_P$  computes the same tuples as  $\Gamma_P$ .

If there exists a  $k \geq l$  such that  $\mathcal{T}_P$  reaches fixpoint after  $k$  iterations, then the theorem is proved. Consider applying  $\mathcal{T}_P$  to  $(\Gamma_P^l(I), h^l)$ . The set of tuples will not change. For any tuple  $t \in \Gamma_P^l(I)$ , the height annotation can only decrease as a result of applying  $\mathcal{T}_P$  since  $\mathcal{T}_P$  takes the minimum height over all rule configurations generating  $t$  and  $h^l(t)$  also must result from such a configuration.

The height annotation is bounded from below by 0 since EDB tuples have non-negative annotations, and each subsequently generated tuple has increasing annotation. Therefore, applying  $\mathcal{T}_P$  monotonically decreases the height annotation of  $t$ , which is bounded from below, so eventually, a fixpoint must be reached. Since this holds for all tuples in  $\Gamma_P^l(I)$ ,  $\mathcal{T}_P$  must reach a fixpoint after  $k \geq l$  iterations.  $\square$

We have established that the provenance evaluation strategy terminates and computes the same set of tuples as standard bottom-up evaluation. It remains to be shown that the proof height annotations are minimal, i.e., that they reflect the real height of the minimal height proof tree for each tuple, and also that

they correspond to real proof trees. The property of minimal height annotations is the major result of this section since it demonstrates that our method generates proof trees of minimal height.

**Theorem 1.** *Let  $\mathcal{T}_P^k((I, h)) = (\Gamma_P^k(I), h^k)$  be the resulting instance at fixpoint of  $\mathcal{T}_P$ . Then, for any arbitrary tuple  $t \in \Gamma_P^k(I)$ ,*

1. *there does not exist any sequence of tuples  $t_1, \dots, t_n$  such that  $t :- t_1, \dots, t_n$  is a valid instantiation of a rule in  $P$  with each  $t_i \in \Gamma_P^k(I)$  and  $h(t) > \max\{h(t_1), \dots, h(t_n)\} + 1$ , and*
2. *there is a valid proof tree for  $t$  with height  $h^k(t)$*

*Proof.* The proof for part 1 is by contradiction. Assume such a sequence of tuples  $t_1, \dots, t_n$  exists. Consider applying  $\mathcal{T}_P$  to the instance.

$$\mathcal{T}_P(\Gamma_P^k(I), h^k) = (\Gamma_P^k(I), h^{k+1})$$

with  $h^{k+1}(t) = \min_{g \in G_t} \{\max_{t_i \in g} \{h^k(t_i)\} + 1\}$  by definition of  $\mathcal{T}_P$ . The set of tuples does not change since we assume that a fixpoint of  $\Gamma_P$  has already been reached.

Since the sequence  $t_1, \dots, t_n$  is a valid rule body configuration generating  $t$ , it is an element of  $G_t$ , and therefore is considered when updating the height annotation of  $t$ . Since the height annotation resulting from this sequence is lower than  $h^k(t)$ , the update will happen, and thus a fixpoint has not yet been reached.

Thus, we have a contradiction, and so such a sequence producing a lower height annotation cannot exist.

The proof for part 2 is by induction on the height annotation of  $t$ . Let  $h = h^k(t)$  for simplicity.

If  $h = 0$ , then  $t$  is in the EDB. In this case, the proof tree with a single node corresponding to  $t$  is a valid proof tree. Otherwise, for  $h > 0$ , assume the hypothesis is true for all tuples with height annotation less than  $h$ . By definition of  $\mathcal{T}_P$ , there exists a sequence  $t :- t_1, \dots, t_n$  such that

$$h = \max(h^k(t_1), \dots, h^k(t_n)) + 1$$

By the assumption, there are valid proof trees for each  $t_i$  of height  $h^k(t_i)$ . We can generate a proof tree as follows:

$$\frac{\frac{\dots}{t_1} \quad \dots \quad \frac{\dots}{t_n}}{t}$$

where each  $\dots$  represents the subtree forming a valid proof tree for each  $t_i$ . This resulting proof tree has height

$$\max(h^k(t_1), \dots, h^k(t_n)) + 1$$

which equals  $h$ . This forms a valid proof tree for  $t$  of height  $h^k(t)$ . □

We have shown the correctness and minimal height annotations of the provenance evaluation strategy for a single fixpoint computation. To evaluate a stratified Datalog specification in a semi-naïve fashion, each stratum is evaluated as a separate fixpoint using the provenance evaluation strategy. The correctness of the evaluation of a full Datalog specification follows from the correctness of each fixpoint evaluation.

### 3.2.1 Complexity of Provenance Evaluation Strategy

In this section, we discuss the complexity of the provenance evaluation strategy. We characterize this complexity by the number of rule firings during evaluation. With standard bottom-up evaluation, we say that a rule is fired if it generates a new tuple. Therefore, for each tuple generated by the Datalog specification, there is exactly 1 rule firing. However, with the provenance evaluation strategy, a rule is also fired if it results in an update for the height annotation of a tuple. Therefore, we consider the number of updates performed during evaluation of the specification as a characterization of the extra amount of work done by provenance evaluation compared to standard bottom-up evaluation.

**Theorem 2.** *An upper bound for the number of updates performed is  $\mathcal{O}(n \times \max h)$ , where  $\max h$  denotes the maximum attained height annotation for any tuple during evaluation and  $n$  the number of tuples generated by the specification.*

*Proof.* To prove this, we need to show two things: (1) that it is a true upper bound, and (2) that it is a tight bound.

To prove (1), consider a tuple  $t$  attaining a height annotation of  $\max h$ . Its annotation may only be updated if there is a valid derivation for  $t$  with a lower height. In the worst case, in each update, we reduce the annotation by 1, and thus we must perform  $\max h$  updates to  $t$ . Considering all tuples produced by the specification, we may update all tuples in this way in the worst case, and therefore, we have  $\mathcal{O}(n \times \max h)$  updates.

To prove (2), we show an example attaining the upper bound, in Figure 11. In this example, the maximum height annotation is  $2k$ , and the tuple  $\text{reach}(a, e)$  will be updated  $k$  times as new derivations are computed using nodes in the bottom chain. Furthermore, each tuple  $\text{reach}(a, x)$  corresponding to nodes  $x$  in the ‘leg’ must be updated  $\mathcal{O}(k)$  times as the tuple  $\text{reach}(a, e)$  is updated. Since there are  $k$  nodes in the leg, each of which is updated  $\mathcal{O}(k)$  times, we have in total  $\mathcal{O}(k^2)$  updates, which coincides with the upper bound. Therefore, this upper bound is tight.  $\square$

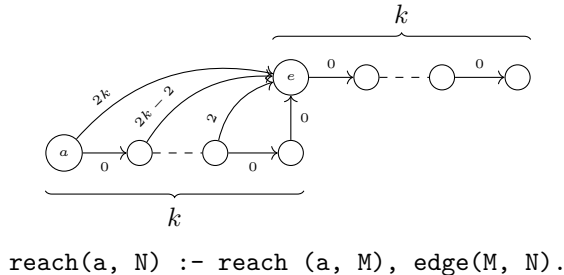


Figure 11: Example Datalog specification demonstrating the upper bound is tight. The label on each edge denotes the height annotation of that tuple.

We also note that  $\max h$  cannot exceed  $n$  since in each iteration of  $\mathcal{T}_P$  where a new tuple is generated, the height annotation of that tuple cannot exceed the maximum height annotation in the previous iteration, plus 1. If no more tuples are generated, then the height annotation for any tuple may not increase. Therefore, by generating a new tuple, we increase  $\max h$  by at most 1, and therefore this value is at most the total number of tuples generated.

Therefore, in the worst case, the provenance evaluation strategy may have to do a quadratic amount of extra work compared to standard bottom-up evaluation. However, as real-world examples (see Section 5) show, such instances rarely occur, and scalability is maintained in most real-world cases.

### 3.3 Proof Tree Construction by Provenance Queries

Given a provenance instance  $(I, h)$  computed by the provenance evaluation strategy, and a tuple  $t \in I$ , the aim is to construct one level of a minimal height proof tree for  $t$ . We utilize the height annotations  $h$  and rule number annotations that are stored alongside the instance during bottom-up evaluation. We use a top-down approach for proof tree construction, starting from a query tuple and recursively finding tuples that form a valid instantiation of a rule generating the query tuple. Denote  $h(t)$  to be the height annotation, and  $\rho(t)$  to be the rule corresponding with the rule annotation for  $t$ .

The result of this search would be a sequence  $t_1, \dots, t_n$  such that  $t \text{ :- } t_1, \dots, t_n$  is a valid instantiation of  $\rho(t)$  leading to a minimal height proof tree. A pre-requisite is that the provenance instance  $(I, h)$  is the result of bottom-up evaluation, and since all possible tuples are computed during this evaluation, we know that each  $t_1, \dots, t_n$  exists in  $I$ . Thus, this problem would be solved by searching for tuples in the already computed instance  $I$ .

However, we must constrain this search such that the result is part of a proof tree of minimal height since there may be multiple valid configurations for the body of  $\rho(t)$ , and some configurations may not lead to minimal height proof trees. These constraints result from the annotations from bottom-up evaluation. From Theorem 1, there exists a configuration for the body which leads to a minimal height annotation for the head, and the height annotation for tuple  $t$  is generated as

$$h(t) = \max(h(t_1), \dots, h(t_n)) + 1$$

by the consequence operator. Therefore, a configuration leading to the minimal height proof tree is  $t_1, \dots, t_n$  where  $h(t_i) < h(t)$  for each  $t_i$ . Note that there may be multiple configurations leading to a proof tree of minimal height, and any of these configurations is a valid result for the problem.

The problem can be phrased as the following goal search. Given a tuple  $t$ , and a rule  $\rho(t) : R(X) :- R_1(X_1), \dots, R_n(X_n), \psi(X_1, \dots, X_n)$  generating  $t$ , we want to find tuples  $t_1, \dots, t_n \in I$  such that  $t :- t_1, \dots, t_n$  is a valid instantiation of  $\rho(t)$ , with proof annotations of each  $t_i$  satisfying the former constraints.

$$? :- R_1(X_1), \dots, R_n(X_n), \psi(X_1, \dots, X_n), \text{matches}(t, X_1, \dots, X_n), h(R_1(X_1)) < h(t), \dots, h(R_n(X_n)) < h(t)$$

The condition  $\text{matches}(t, X_1, \dots, X_n)$  denotes that for a result  $t_1, \dots, t_n$ , the variable mapping from each  $X_i$  to  $t_i$  is consistent with the variable mapping from  $X$  to  $t$ . This is related to the problem of unification in Prolog, and in our context is crucial to ensure that the resulting configuration forms a valid instantiation of  $\rho$ .

**Example:** We illustrate this construction using the running example. The query is for the tuple  $\text{alias}(a, b)$ . From the initial bottom-up evaluation, the height annotation is  $h(\text{alias}(a, b)) = 4$ , and the generating rule is  $r_4 : \text{alias}(\text{Var1}, \text{Var2}) :- \text{vpt}(\text{Var1}, \text{Obj}), \text{vpt}(\text{Var2}, \text{Obj})$ .

The search is for tuples forming a configuration for the body of  $r_4$ ,  $\text{vpt}(\text{Var1}, \text{Obj}), \text{vpt}(\text{Var2}, \text{Obj})$  satisfying the constraints

$$\begin{aligned} \text{Var1} &= a, \\ \text{Var2} &= b, \\ h(\text{vpt}(\text{Var1}, \text{Obj})) &< 4, \\ h(\text{vpt}(\text{Var2}, \text{Obj})) &< 4 \end{aligned}$$

In this example, the first two constraints corresponds with  $\text{matches}(t, X_1, \dots, X_n)$ , and the last two constraints enforce the conditions for proof height annotations. Therefore, the goal search is

$$\begin{aligned} ? :- \text{vpt}(\text{Var1}, \text{Obj}), \text{vpt}(\text{Var2}, \text{Obj}), \text{Var1} \neq \text{Var2}, \text{Var1} = a, \text{Var2} = b, \\ h(\text{vpt}(\text{Var1}, \text{Obj})) < 4, h(\text{vpt}(\text{Var2}, \text{Obj})) < 4 \end{aligned}$$

In this case, we find the tuples  $\text{vpt}(a, l_1), \text{vpt}(b, l_1)$ , which form the next level of the proof tree:

$$\frac{\text{vpt}(a, l_1) \quad \text{vpt}(b, l_1) \quad a \neq b}{\text{alias}(a, b)} r_4$$

This goal search can be seen as a special case of Datalog, denoted **0-IDB** Datalog. The properties of 0-IDB Datalog are:

- There are no IDB relations, and every possible tuple is in EDB
- There is a single goal search query
- The evaluation terminates as soon as the first solution for the goal search is found

Our method for searching for body tuples fits into this framework. The IDB resulting from the provenance evaluation strategy forms the EDB of our 0-IDB program, and the search for matching body tuples is the single goal search query.

The property that the evaluation terminates as soon as the first solution is found is desirable for efficiency reasons. Our goal search contains constraints to ensure that *any* result forms a valid configuration for a minimal height generator, and therefore the first solution found is sufficient to solve our problem.

The complexity of the goal search depends highly on the data structures used in the implementation. We assume a fully (B-Tree) indexed nested loop joins. Therefore searching for a tuple for a rule with an  $m$  nested join, requires  $\mathcal{O}(\log^m n)$  time. Given a proof tree height of  $k$ , we need  $\mathcal{O}(k \log^m n) \equiv \mathcal{O}(\log^m n)$  to traverse a single branch.

### 3.4 Provenance for Non-Existence of Tuples via User Interaction

The provenance evaluation strategy of the previous section explains the existence of tuples in relations. However, the non-existence of tuples may also indicate faults in either the input relations and/or in the rules.

Therefore, we extend our approach explaining on why a tuple *cannot* be derived, i.e., if the user expects a tuple, but it does not appear in the IDB, the user may wish to investigate why the tuple is not produced. Alternatively, a user may want to understand why a negated body literal holds in a rule during the debugging process.

A non-existent tuple is characterized by *every* proof tree for the tuple failing to be constructed. The source of failure may be (1) tuples for the construction not being in the EDB/IDB, and/or (2) the constraints of rules not being satisfied. Given the potentially infinite number of failing proof trees, we avoid automatic procedures that represent a serious technical challenge and are not guaranteed to discover a failed proof tree containing the root cause of the fault. In practice, without a formal description of the root cause of the fault, the provenance system cannot decide which failed proof tree is most suitable<sup>1</sup>.

Hence, in our system, we take a pragmatic, semi-automated approach which is inspired by existing work such as [27, 28]. Our system leverages user domain knowledge and allows user interactions to incrementally guide the construction of a single failing proof tree. Each user interaction produces a failing *subproof*, or one level of the proof tree. This failing proof tree provides a succinct representation of valuable information for a Datalog user to discover on why an expected tuple is not being produced by the specification and does not burden the user with too much unnecessary information.

Formally, we define the problem as follows: given a provenance instance  $(I, h)$  computed by the provenance evaluation strategy, a tuple  $t \notin I$ , and a rule  $\rho : R(X) :- R_1(X_1), \dots, R_n(X_n), \psi(X_1, \dots, X_n)$  with head relation matching  $t$ , we aim to find a configuration  $t_1, \dots, t_n$  for the body of  $\rho$ , such that either: (1) at least one  $t_i \notin I$  or, (2) the constraints  $\psi(X_1, \dots, X_n)$  are not satisfied. Such a configuration forms a failing subproof, and recursively constructing subproofs results in a full failed proof tree. Note that it would be impossible to find a configuration where all tuples  $t_i \in I$  and constraints  $\psi(X_1, \dots, X_n)$  hold since the prior assumption is that  $t \notin I$ . If such an instantiation cannot be found, then the tuple  $t$  can be generated by the Datalog specification, and thus  $t \in I$ .

For showing the non-existence of a tuple, the provenance system supports the Datalog user in constructing the failing proof tree in stages. The debugging query for non-existence has three user interaction steps that are repeated until the root cause of the fault is found. The user interaction steps are as follows:

1. the user defines a query for the non-existence of a tuple,
2. the user selects a candidate rule from which the tuple may have been derived,
3. the user selects candidate variable values of unbound variables in the rule.

The system displays the rule application in the failing proof tree indicating the portions of the rule that fail (i.e., at least one literal / constraint must fail) and the portions of the rule that hold.

The Datalog user can continue the query with the newly found failing literals guiding the system to find the root cause of the fault. This process is semi-automated since the nature of the fault is known by the Datalog user only.

**Example:** Consider the example from Figure 1 for which we want to query the non-existence of the tuple  $vpt(b, l_4)$ . In the first user interaction step, the Datalog user queries for an explanation for the non-existence of the tuple  $vpt(b, l_4)$ . Then, the Datalog user selects an appropriate rule such as rule  $r_2$ .

<sup>1</sup>Proof annotations such as introduced in the previous section can only describe existent tuples in the IDB. It is impossible to consider such annotations for tuples that are *not* produced by the specification.

The system can then produce a partial instantiation for the body of the rule, where variables matching the head are replaced by concrete values from  $t$  such as,

$$vpt(b, l_4) \text{ :- } assign(b, Var_2), vpt(Var_2, l_4)$$

In the last user interaction step, the Datalog user selects instantiations for the remaining free variables in rule  $r_2$ . For example, the Datalog user may choose the value  $d$  for the free variable  $Var_2$ .

$$vpt(b, l_4) \text{ :- } assign(b, d), vpt(d, l_4)$$

Given the instantiated rule, the provenance system will evaluate which portions of the subproof fail and which portions hold. With that information, the Datalog user can continue the exploration of the failing portions to find the root cause of the fault. A simple colour labelling helps the Datalog user to indicate which portions fail and hold, respectively.

$$\frac{assign(b, d) \text{ X} \quad vpt(d, l_4) \text{ } \checkmark}{vpt(b, l_4)} r_2$$

In the above example, the red color and X denotes the non-existence of the tuple  $assign(b, d)$  in the IDB, i.e., a failing portion of the proof tree. The blue color with  $\checkmark$  indicates that  $vpt(d, l_4)$  holds.

In summary, our provenance system constructs a single failed subproof to explain the non-existence of a tuple. The construction of the failed subproof is guided by the Datalog user to ensure the answer contains a relevant explanation, given the infinitely many possible failed proof trees. The semi-automatic proof construction approach supports the Datalog user by highlighting which portions of the subproof hold and fail, respectively to guide the exploration.

### 3.5 Alternative Proof Tree Shapes

Our debugging strategy introduces an *interactive* system to explore fragments of proof trees to pinpoint faults in the Datalog specification. Therefore, we wish to minimize the number of user interactions required to find the fault. For this aim, minimal height proof trees are critical for reducing the number of user interactions in the fault investigation phase. The utility of this approach is backed by several user experiences in industrial-scale applications (see cf. Section 7.1.2 [38]).

While generating proof trees of minimal height is useful for users, in principle our framework is more general and can support a variety of metrics that may be beneficial in future applications. In this section, we outline general properties of proof tree metrics by having the following properties for function  $h$ :

1. The codomain of  $h$  must have a partial ordering  $\sqsubseteq$ , so that an update mechanism can be well defined. It is important that the annotation for a tuple can be updated if the same tuple is generated again with smaller (according to  $\sqsubseteq$ ) annotation. This ensures that the resulting annotations are always minimal since tuples will continue being updated with smaller annotations until a fixpoint with annotations is reached.
2. The metric must be *compositional*, i.e., if  $t$  is generated by a rule instantiation  $t \text{ :- } t_1, \dots, t_n$ , then  $h(t) = f(h(t_1), \dots, h(t_n))$ . The importance of this property is two-fold. Firstly, it ensures that the values of the annotations can be computed during evaluation of the Datalog specification, by encoding  $f$  as a functor in the transformed Datalog specification. For example, a rule may be transformed to be  $R(X, f(h_1, \dots, h_n)) \text{ :- } R_1(X_1, h_1), \dots, R_n(X_n, h_n)$ . to compute the value of the annotation.

Secondly, the compositional property is important for the reconstruction of the proof tree. In the backwards search for a body configuration that may produce the head tuple,  $f$  may be encoded as a constraint. For example, a backwards search may be

$$? \text{ :- } R_1(X_1), \dots, R_n(X_n), \psi(X_1, \dots, X_n), matches(t, X_1, \dots, X_n), h(t) = f(h(R_1(X_1)), \dots, h(R_n(X_n)))$$

where the last constraint ensures that the tuples found from the search correctly generate  $t$  with matching annotations.



3. The metric must be *monotone* and bounded, i.e.,  $h(t_i) \sqsubseteq h(t)$  for all  $1 \leq i \leq n$ , and *bounded*, i.e., there is a minimum value  $c$  such that  $c \sqsubseteq h(t)$  for any tuple  $t$ . This property ensures that the provenance evaluation strategy terminates. Monotonicity ensures that with each rule application, the annotation converges towards the minimum value  $c$ , and once it reaches  $c$ , then termination must occur.

If a given metric satisfies the above properties, then it can be used instead of proof tree height in our framework. Examples of such metrics could be the size of proof trees by number of nodes, or a sequence of  $k$  proof tree heights describing the smallest  $k$  proof trees for each tuple. Therefore, our approach could be applicable to wider applications than debugging. We leave the integration of other metrics in our provenance evaluation strategy as future work.

## 4 Implementation in Soufflé

In this section, we describe the implementation of our provenance system in Soufflé [24]. Soufflé [1] is an open-source system that is available under the UPL license and is implemented in C++. Soufflé is a parallel Datalog engine designed for shared memory, multi-core machines, synthesizing highly performant parallel C++ code from Datalog specifications.

Our provenance evaluation strategy and proof tree construction system are tightly integrated into the Soufflé engine<sup>2</sup>. Through this tight integration in Soufflé, we are able to achieve high parallel performance for the provenance evaluation, and enable a single evaluation phase to answer multiple debugging queries. In contrast, previous approaches [27, 13, 26] implement a Datalog re-writing scheme, and simply evaluate the re-written Datalog in an existing engine.

The Soufflé synthesizer performs a series of specialization steps based on Futamura projections [17], which synthesize a C++ program with the same semantics as the Datalog specification. The main specialization step is the compilation of Datalog into an intermediate representation called Relational Algebra Machine (RAM) which has imperative and relational algebra elements to perform simple relational algebra operations to compute fixed-points for semi-naïve evaluation. The RAM representation of a Datalog specification is in turn compiled into C++ code. The resulting C++ code implements a specialized semi-naïve algorithm for the rules in the Datalog specification that have similar performance to a hand-written program [24, 3]. In the following, we discuss the implementation of semi-naïve evaluation [2] in Soufflé and discuss subsequently how the synthesised semi-naïve evaluation is replaced by the provenance evaluation strategy. Semi-naïve evaluation avoids re-computations of tuples by assuming that new tuples can only be deduced from new tuples in the previous iteration. This is achieved by creating a *new* and *delta* version of each relation. The *new* and *delta* version of a relation store the tuples found in the current iteration and the previous iteration, respectively. For example, with a rule  $R_0(X_0) :- R_1(X_1), \dots, R_n(X_n), \psi(X_1, \dots, X_n)$ , each relation  $R_k$  is transformed to become a set of relations for each iteration  $i$ :

$$R_k^i, new_{R_k}^i, \Delta_{R_k}^i$$

where  $R_k^i$  stores all the tuples for relation  $R_k$  which are computed up until iteration  $i$ , while  $\Delta_{R_k}^i$  stores only the tuples in  $R_k$  computed in iteration  $i$ , without any tuples computed in previous iterations. The relation  $new_{R_k}^i$  is an intermediate relation used to compute the  $\Delta$  relations. The essential optimization, compared to naïve evaluation, is to realize that in iteration  $i + 1$ , a new tuple is only generated if it directly depends on a tuple generated in iteration  $i$ . If this condition doesn't hold, i.e., if it depends on knowledge generated in prior iterations, then the tuple would also have been generated in a previous iteration, and thus generating it again would be a redundant computation. Thus, a tuple is only generated in iteration  $i + 1$  if it depends on a  $\Delta^i$  relation. This constraint is enforced by transforming the original rule to a set of new Datalog rules

---

<sup>2</sup>Our provenance evaluation strategy is not specific to Soufflé – it can be integrated into any bottom-up evaluation Datalog engine.

which perform semi-naïve evaluation:

$$\begin{aligned}
new_{R_0}^{i+1}(X_0) &:- \Delta_{R_1}^i(X_1), R_2(X_2), \dots, R_n(X_n), \psi(X_1, \dots, X_n) \\
&\dots \\
new_{R_0}^{i+1}(X_0) &:- R_1(X_1), \dots, \Delta_{R_k}^i(X_k), \dots, R_n(X_n), \psi(X_1, \dots, X_n) \\
&\dots \\
new_{R_0}^{i+1}(X_0) &:- R_1(X_1), \dots, R_{n-1}(X_{n-1}), \Delta_{R_n}^i(X_n), \psi(X_1, \dots, X_n)
\end{aligned}$$

Thus,  $new_{R_0}^{i+1}$  contains tuples of  $R_0$  which depend directly on tuples generated in iteration  $i$ . The relation  $\Delta_{R_0}^{i+1}$  is computed as

$$\Delta_{R_0}^{i+1} = new_{R_0}^{i+1} - R_0^i$$

where the relations are viewed as sets of tuples and  $-$  denotes set minus. Thus,  $\Delta_{R_0}^{i+1}$  contains only tuples generated in iteration  $i$ , and no tuples generated in previous iterations. The relation  $R_0^{i+1}$  denotes all tuples generated in iterations  $0, \dots, i+1$ , and is computed as the union

$$R_0^{i+1} = R_0^i \cup \Delta_{R_0}^{i+1}$$

With these auxiliary relations, the final result for the relation  $R_0$  is the final  $R_0^i$  once a fix-point is reached, i.e., the result of the Datalog specification has stabilized. Note that Soufflé evaluates the  $\Delta_R^{i+1}$  relation by computing the tuples without an explicit set minus operation since an existence check determines whether the tuple already exists in  $R^i$  relation before it is inserted into  $\Delta_R^{i+1}$ . For example, Figure 12 depicts a snippet of a Soufflé RAM program, part of the semi-naïve evaluation of the rule  $r_2 : vpt(Var, Obj) :- assign(Var, Var2), vpt(Var2, Obj)$ . The join is performed via a loop nest iterating over tuples of relations efficiently via indexes [38]. Line 2 computes the new tuples to be added to the  $\Delta_{vpt}^{i+1}$  relation, where the NOT IN operation is an existence check to ensure the generated tuple does not already exist in the  $vpt^i$  relation.

```

1 SCAN assign AS t0
2 SEARCH @delta_vpt AS t1 ON INDEX t1.c0=t0.y WHERE (t0.x,t1.c1) NOT IN vpt
3 INSERT (t0.x, t1.c1) INTO @new_vpt

```

Figure 12: Existence check prior to inserting

## 4.1 Implementing Provenance Evaluation Strategy

The main challenge of integrating the provenance evaluation strategy is to allow the synthesis to be aware of proof annotations. In particular, the semi-naïve evaluation machinery must be replaced by the provenance evaluation strategy as described in Section 3.2 to handle the proof annotations. Another critical part of this machinery is the synthesis of data structures [38, 25] for relations that are specialized for the operations in the program. The synthesized data structures have to be extended for proof annotations as well, enabling an update semantics in Datalog for the annotations.

For the provenance evaluation strategy, we need to amend relations by extra attributes to contain the proof annotations. We utilize the synthesis pipeline of Soufflé by introducing two provenance attributes for each relation. The first attribute represents the rule number of the rule which generated the tuple, and the second attribute represents the proof tree height. These two new attributes are introduced for each relation at the syntactic level in Soufflé. A predicate  $R(X)$  is transformed into  $R(X, @rule, @height)$ . For the sake of readability in this text, we distinguish between original and provenance tuples, where an original tuple is a provenance tuple without proof annotations. We rewrite all logic rules at the syntactic level to take account of the two provenance attributes constituting the proof annotation for our system, and to compute the value of the annotations. The proof annotation instrumentation is performed as follows where a rule

$$\rho_k : R(X) :- R_1(X_1), \dots, R_n(X_n), \psi(X_1, \dots, X_n)$$

is transformed into:

$$\rho_k : R(X, k, \max(@\text{height}_1, \dots, @\text{height}_n) + 1) :- \\ R_1(X_1, -, @\text{height}_1), \dots, R_n(X_n, -, @\text{height}_n), \psi(X_1, \dots, X_n)$$

The transformed provenance rule computes level and height annotations for a new tuple, according to the semantics in Section 3.2. Since the rules are known statically, the rule number annotation  $k$  can be assigned a constant value for each rule in the transformed specification. The rule numbers of the body predicates are ignored by using `_` in each body predicate since they do not influence the head predicate.

The transformed provenance rule syntactically represents the computation of proof annotations during rule evaluation. However, the actual execution of provenance rules differs from a standard semi-naïve evaluation as presented in [2]. The reason is the update mechanism: a newly discovered provenance tuple may overwrite an existing provenance tuple if they are the same original tuple, but the new tuple has a smaller height annotation.

The provenance evaluation strategy extends the semi-naïve algorithm by updating the rule number and the height annotation of a tuple (as defined by  $\mathcal{T}_P$ ) if the original tuple already exists and the newly generated tuple has smaller height annotation. In other words, if a smaller proof tree could be found in a subsequent iteration for the same tuple, then an update occurs. Otherwise, if the original tuple doesn't already exist, the provenance tuple is inserted into the relation as is. Thus, the rule computing  $\Delta_{R_0}^{i+1}$  in the semi-naïve evaluation is modified to accommodate the possibility of updates, i.e.,

$$\Delta_{R_0}^{i+1} = (\text{new}_{R_0}^{i+1} - R_0^i) \cup \{t \in R_0^i \mid h^i(t) > h^{i+1}(t)\}$$

where  $h^i$  denotes the height annotations in iteration  $i$ .

Therefore, with our provenance evaluation strategy, we integrate the possibility of an annotation update into the data structure. During an insertion operation, if the same original tuple is discovered, with a larger height annotation than the current tuple, then an update occurs. Therefore, we wish to call the insert operation if *either* the tuple does not already exist *or* the existing tuple has the larger proof annotation. A specialized existence check is implemented, implicitly implementing the semantics of the provenance  $\Delta_R^{i+1}$  relation. Similar to the standard existence check, the special existence check is implemented as part of the data structure that has been specialized for each relation.

The result is the RAM snippet in Figure 13. The obvious differences compared to Figure 12 are in lines 3-4, where the level annotation is computed within the loop nest, as part of the insertion. Furthermore, the `PROV NOT IN vpt` operation in line 3 denotes the special provenance existence check, which allows the `INSERT` to proceed if either the tuple does not already exist, or exists with a larger proof height annotation. Thus, this implements the update semantics discussed above.

```

1 SCAN assign AS t0
2 SEARCH @delta_vpt AS t1 ON INDEX t1.x=t0.y
3 IF (t0.x,t1.y,_,(max(t0.@height,t1.@height)+number(1))) PROV NOT IN vpt
4 INSERT (t0.x, t1.y, number(2), (max(t0.@height,t1.@height)+number(1))) INTO @new_vpt

```

Figure 13: Provenance version of RAM loop nest

However, the specializations in the data structures still remain to be discussed. Soufflé employs a highly specialized parallel B-Tree data structure [25], with index orderings for the attributes generated automatically via an optimization problem [38]. The B-Tree employs a special optimistic read/write lock for each node, allowing high throughput for parallel insertion. During an `insert` operation, a thread may obtain a read lease for each node as it checks whether the tuple to be inserted already exists. If an insertion is required, it checks if the lease has changed, and restarts the whole procedure if it has. Otherwise, it upgrades to a write lease and inserts the tuple into the correct position in the B-Tree. The data structure also takes advantage of Soufflé's Datalog evaluation setting, where a single relation is either read from, or written to, but never at the same time. Therefore, there are no interleaved reads and writes, and so read operations are not synchronized.

With the proof annotations, we modify these specializations so that they can take into account the provenance semantics. The important step is the *update* semantics, and thus we integrate an update mechanism

into the `insert` operation, without requiring to delete and then re-insert. The provenance evaluation strategy requires two main modifications to our B-Tree data structure. Firstly, the existence check for insertion should consider only the original tuple, and ignore annotations. This ensures that Datalog set semantics are preserved and that no duplicate original tuples can exist. However, note that we still need to retrieve the full tuple, including its proof annotations. This is important for the proof tree construction, discussed in the next section. To address this concern, we use different lexicographical orderings of indices for the `insert` and `retrieve` operations. The `insert` index order does not include the attributes storing the proof annotations (so that annotations are not considered when checking existence), while the `retrieve` index order does. We also need to ensure that updating an annotation does not change the ordering of tuples according to the index; otherwise subsequent index supported searches will fail. Therefore, the `retrieve` index order requires the annotation attributes to be at the end, as this guarantees that an update to the annotations does not affect tuple ordering.

Secondly, we must have a mechanism to update existing tuples to implement the update semantics in Section 3.2. To achieve this, we modified the `insert` operation so that it may also update any existing tuple with a smaller proof annotation. This insertion first requires to check if the original tuple exists in the B-Tree. If it does, rather than aborting (as it would with standard Datalog evaluation), the insertion then checks the annotation. If the height annotation of the existing tuple is larger than the tuple to be inserted, then the annotations of the existing annotation are updated with new values. Note that during an update, a read lease also needs to be validated and upgraded to a write lease. The integration of the update into the `insert` operation avoids any need to delete and re-insert tuples, thus improving the efficiency of the provenance evaluation strategy. These modifications to the B-Tree reflect the desired insertion semantics, and updates are handled directly in the `insert` operation. All other retrieval operations for the B-Tree are not modified, and tuples can be retrieved as normal, including their proof annotations.

## 4.2 Implementing a Proof Tree Construction User Interface

After the provenance evaluation strategy is completed, the proof tree construction stage is driven by the user. It is critical that this process is also fully parallelized and highly performant.

The user interface is implemented as a command line, where the user can enter queries to explain the existence and non-existence of a tuple. For example, the query `explain alias("a", "b")` results in the proof tree in Figure 14. Explaining the non-existence of a tuple, i.e., the query `explainnegation vpt("b", "14")`, results in the interaction in Figure 15. The user may also select the size of proof tree fragments to display, i.e., `setdepth 6` instructs the system to construct 6 levels of the proof tree in the next query. For each debugging query, the system invokes the relevant procedure to construct a proof tree fragment.

```

> explain alias("a", "b")
                                new("a", "11")
                                -----(R1)
new("a", "11") assign("b", "a") vpt("a", "11")
----- (R1) ----- (R2)
vpt("a", "11")          vpt("b", "11")          "a" != "b"
----- (R1)
                                alias("a", "b")

```

Figure 14: Explaining the tuple  $alias(a, b)$

It is critical that the proof tree construction procedures are highly performant since the constructed IDB may be very large, and we may need to search through many tuples to construct a proof tree fragment. Therefore, the proof tree construction procedures must be tightly integrated into the Soufflé system to enable a high-performance, parallel search. We integrate these procedures into the Soufflé RAM, utilizing the existing translation from RAM to parallel C++. Moreover, since the provenance evaluation strategy uses specialized B-Tree data structures, the proof tree construction can also utilize index supported searches to find relevant tuples.

Recall that the proof tree construction is facilitated by searches for subproofs. Therefore, we require a specialized framework in the Soufflé RAM to implement a subproof search. We term this framework a *subroutine* framework. Each subproof search can be implemented as a subroutine, thus integrating with the

```

Enter command > explainnegation vpt("b", "l4")
1: vpt(Var,Obj) :-
    new(Var,Obj).

2: vpt(Var,Obj) :-
    assign(Var,Var2),
    vpt(Var2,Obj).

3: vpt(Var,Obj) :-
    load(Var,Y,F),
    store(P,F,Q),
    alias(P,Y),
    vpt(Q,Obj).

Pick a rule number: 2
Pick a value for Var2: d
assign("b", "d") X vpt("d", "l4") ✓
------(R2)
                vpt("b", "l4")

```

Figure 15: Explaining the non-existence of the tuple  $vpt("b", "l4")$

Soufflé RAM.

To explain the existence of a tuple, a subproof search is required to search the body of a rule for matching body tuples, satisfying the constraint that proof tree height is lower than the current tuple. This backwards search for a single rule is implemented as a subroutine. For example, the rule  $r_2 : vpt(Var, Obj) :- assign(Var, Var2), vpt(Var2, Obj)$  is implemented as the subroutine in Figure 16.

```

1 SUBROUTINE vpt_2_subproof
2   SCAN assign AS t0 WHERE t0.x = argument(0) AND t0.@level_number < argument(2)
3   SEARCH vpt AS t1 ON INDEX t1.x=t0.y AND t1.y=argument(1) WHERE t1.@level_number < argument(2)
4   RETURN (t0.x, t0.y, t0.@rule_number, t0.@level_number, t0.y,
5           t1.y, t1.@rule_number, t1.@level_number)

```

Figure 16: Subroutine for example program

Lines 2-5 represent a search through a database that is already constructed by the initial bottom-up evaluation, to find tuples which satisfy the constraints required for the construction of a proof tree fragment. The values of `argument(0)` and `argument(1)` are the values in the head tuple, and `argument(2)` is the height annotation of the head tuple. Therefore, this subproof search is parameterized by the head tuple. The relations of the body atoms, `assign` and `vpt` are searched to find tuples `t0` and `t1` which match the body of the rule. Importantly, the constraints for the level number are encoded in lines 3-4, ensuring that the resulting tuples have level number annotations lower than the query tuple. As shown in Section 3.3, being able to apply this operation recursively allows us to generate the full proof tree.

Similarly, to generate a failed subproof to explain the non-existence of a tuple, the search for failing and holding parts of a subproof is implemented as a subroutine. Given an instantiated rule (which is produced via user interaction), a subroutine returns whether each body tuple is in the IDB and whether each constraint is satisfied.

## 5 Experiments

In this section, we conduct experiments with the provenance evaluation strategy and provenance queries implemented in Soufflé (see Section 4). The experiments are conducted for large-scale Datalog specifications. We have the following experimental research claims:

**Claim-I: Scalable Provenance Evaluation Strategy.** Our provenance evaluation strategy only has a minor impact on the runtime performance, i.e., our provenance evaluation strategy remains scalable for realistic datasets and rulesets.

| Benchmark | context-insensitive |            | 1-obj, 1-heap |            |
|-----------|---------------------|------------|---------------|------------|
|           | # EDB               | # IDB      | # EDB         | # IDB      |
| antlr     | 8,319,095           | 21,832,232 | 8,319,095     | 24,145,648 |
| bloat     | 4,468,277           | 13,104,020 | 4,468,277     | 15,417,516 |
| chart     | 8,743,770           | 22,975,742 | 8,743,729     | 25,289,200 |
| eclipse   | 4,389,770           | 13,076,265 | 4,389,799     | 15,389,708 |
| fop       | 8,769,583           | 22,970,533 | 8,769,572     | 25,283,913 |
| hsqldb    | 9,007,087           | 24,561,921 | 9,007,087     | 26,875,437 |
| kython    | 5,203,400           | 17,158,375 | 5,203,400     | 19,471,797 |
| luindex   | 4,396,394           | 13,415,336 | 4,396,394     | 15,728,788 |
| lusearch  | 4,396,415           | 13,415,390 | 4,396,394     | 15,728,788 |
| pmd       | 8,388,202           | 22,853,676 | 8,388,202     | 25,167,134 |
| xalan     | 8,670,980           | 23,488,951 | 8,670,966     | 25,802,385 |

Table 1: Statistics for DOOP benchmarks

**Claim-II: Scalable Proof Tree Construction.** The provenance queries for exploring proof paths scales to large proof tree fragments, allowing efficient interactive exploration of proof trees.

**Claim-III: Need for Proof Tree Exploration.** For realistic benchmarks, minimal proof trees are still very large substantiating the need for interactive proof tree exploration.

As an experimental testbed, we use the DOOP [8] points-to analysis framework. We experiment with DOOP’s *context-insensitive* and *1-object-sensitive, 1-heap* (1-obj, 1-heap) analyses that exhibit different runtime complexities. As inputs for the points-to analyses, we compute the points-to sets for the DaCapo 2006 Java program benchmarks. Each analysis contains approx. 300 relations, 850 rules and produces up to approx. 26 million output tuples on the DaCapo benchmarks (See Table 1). Our experiments were performed on a computer with an Intel Xeon Gold 6130 CPU and 192 GB of memory, running Fedora 27. Soufflé executables were generated using GCC 7.3.1.

## 5.1 Performance of the Provenance Evaluation Strategy

In Table 2, we present the runtime and memory consumption of Soufflé with 8 threads, comparing standard Soufflé with our provenance evaluation strategy with proof annotations. We use the DaCapo benchmarks with both the context-insensitive and 1-obj, 1-heap analysis. As expected, Soufflé with proof annotations incurs an overhead during evaluation. This overhead for the provenance evaluation strategy is typically within a factor of 1.3 which is a small overhead to pay for being able to generate minimal proof trees for all possible tuples in the IDB. Hence, we demonstrate the viability of the provenance evaluation strategy for large-scale Datalog specifications, substantiating Claim I. We noticed that the runtime overhead for the context-insensitive analysis was smaller across all benchmarks than that of the 1-obj-1-heap analysis due to cache locality that was more prominent for smaller memory footprints. Note that the overhead for memory consumption is similar to performance overheads, at approximately  $1.45\times$ . This overhead results from the storage of extra proof annotations during evaluation.

In contrast, a naïve direct encoding approach (cf. Chapter 5, [44]), where each tuple is annotated with its full subproof (i.e., direct children in the proof tree), resulted in excessive memory usage (up to  $100\times$ ) on a simple transitive closure experiment with 2000 tuples, and thus cannot be deployed for large-scale Datalog specifications such as those found in DOOP.

Figures 17a and 17b show the total runtime and average memory usage for each of the DOOP DaCapo benchmarks with both DOOP (context-insensitive and 1-obj-1-heap) analyses, running with multiple threads. The figure demonstrates that the provenance evaluation strategy is scalable, in that the overhead is sustainable with an increasing number of threads. We observe that the overall runtime decreases for provenance and without provenance until 5 threads, and increased thereafter. This is caused by the synchronization of Soufflé’s rule evaluation system and is not specific to provenance. It is interesting to note that the runtime overhead is larger with fewer threads, being  $1.45\times$  for 1 thread while being  $1.23\times$  for 16 threads. Again, this is related to the underlying hardware architecture providing caches and memory lanes for each core. With more threads, the memory bandwidth to access the logical relations with proof annotations improves.

| Benchmark                  | Runtime (sec) |       |              | Memory (MB) |       |              |
|----------------------------|---------------|-------|--------------|-------------|-------|--------------|
|                            | No Prov.      | Prov. | ( $\times$ ) | No Prov.    | Prov. | ( $\times$ ) |
| <b>context-insensitive</b> |               |       |              |             |       |              |
| antlr                      | 9.73          | 12.29 | 1.26         | 595         | 900   | 1.51         |
| bloat                      | 9.54          | 12.25 | 1.28         | 596         | 900   | 1.51         |
| chart                      | 15.89         | 19.60 | 1.23         | 1,103       | 1,604 | 1.45         |
| eclipse                    | 9.64          | 11.76 | 1.22         | 593         | 898   | 1.51         |
| fop                        | 15.57         | 19.48 | 1.25         | 1,079       | 1,579 | 1.46         |
| hsqldb                     | 16.36         | 19.73 | 1.21         | 1,124       | 1,642 | 1.46         |
| ython                      | 11.00         | 13.62 | 1.24         | 731         | 1,090 | 1.49         |
| luindex                    | 9.62          | 12.00 | 1.25         | 594         | 905   | 1.52         |
| lusearch                   | 9.80          | 12.23 | 1.25         | 593         | 904   | 1.52         |
| pmd                        | 15.58         | 18.90 | 1.21         | 1,053       | 1,542 | 1.46         |
| xalan                      | 15.59         | 19.54 | 1.25         | 1,091       | 1,595 | 1.46         |
| geo-mean                   |               |       | 1.24         |             |       | 1.44         |
| <b>1-obj, 1-heap</b>       |               |       |              |             |       |              |
| antlr                      | 10.84         | 12.60 | 1.16         | 936         | 1,310 | 1.40         |
| bloat                      | 15.77         | 22.00 | 1.40         | 732         | 1,082 | 1.48         |
| chart                      | 21.84         | 28.13 | 1.29         | 1,242       | 1,788 | 1.44         |
| eclipse                    | 15.76         | 21.00 | 1.33         | 729         | 1,080 | 1.48         |
| fop                        | 22.21         | 29.63 | 1.33         | 1,216       | 1,756 | 1.44         |
| hsqldb                     | 23.01         | 29.43 | 1.28         | 1,256       | 1,823 | 1.45         |
| ython                      | 17.54         | 22.96 | 1.31         | 868         | 1,270 | 1.46         |
| luindex                    | 15.94         | 21.55 | 1.35         | 730         | 1,086 | 1.49         |
| lusearch                   | 15.95         | 21.25 | 1.33         | 731         | 1,087 | 1.49         |
| pmd                        | 21.58         | 28.21 | 1.31         | 1,190       | 1,725 | 1.45         |
| xalan                      | 22.09         | 28.68 | 1.30         | 1,224       | 1,773 | 1.45         |
| geo-mean                   |               |       | 1.31         |             |       | 1.46         |

Table 2: Runtime and memory usage overheads for Soufflé with and without proof annotations with 8 threads

|              | Runtime (sec) | Average Memory (MB) |
|--------------|---------------|---------------------|
| Our approach | 3:02.7        | 1,289               |
| Top- $k$     | 3:02.3        | 1,289               |

Table 3: Runtime and memory usage overheads for our provenance approach compared to [14]. Runtime is the total over all context-insensitive DaCapo analyses, and memory usage is the average over all context-insensitive DaCapo analyses.

The memory usage of the provenance evaluation strategy has a consistent overhead of  $1.45\times$ , which aligns with our expectations that there would be a reasonable overhead associated with storing the provenance annotations per tuple. Note that this overhead is constant over any number of threads since the amount of extra information stored overall does not change with the number of threads.

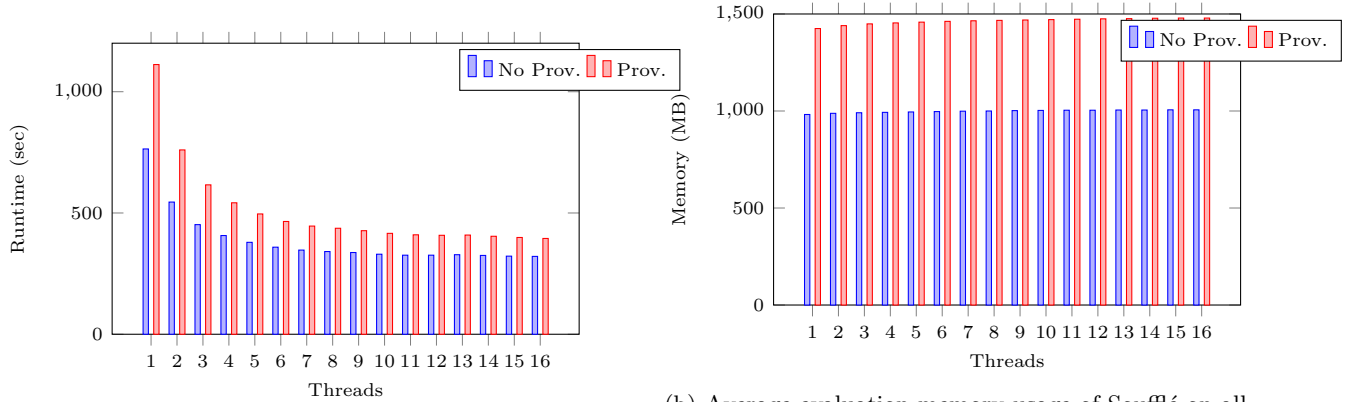
**Comparison with current approaches:** The current state of the art in tracking Datalog provenance is to instrument the specification with a given provenance query. The instrumented Datalog specification can then be evaluated using any Datalog engine. One example of this approach is the top- $k$  approach [13, 14], where Datalog specifications are instrumented based on a provenance query taking the form of a *derivation tree pattern*.

For our experiments, we implemented the instrumentation algorithm presented in [14], and evaluated the resulting Datalog using Soufflé, again using DOOP as the test Datalog specification. Since the instrumentation is done for a specific query, we chose an arbitrary tuple from the *VarPointsTo* relation in DOOP.

The results in Figure 3 showed that during evaluation time, the difference in both runtime and memory usage is less than 1%, demonstrating that our provenance encoding scheme is as scalable as the state-of-the-art. However, with our approach, we are able to answer *any* provenance query during proof construction time.

## 5.2 Proof Tree Construction

For the construction of proof trees, the performance of the provenance queries is instrumental. A debugging query constitutes a backward search for a rule (i.e., reverting the computational direction of a rule). The



(a) Total evaluation runtime of Soufflé on all DaCapo benchmarks with each DOOP (context-insensitive and 1-obj, 1-heap) analysis with and without provenance

(b) Average evaluation memory usage of Soufflé on all DaCapo benchmarks with each DOOP (context-insensitive and 1-obj, 1-heap) analysis with and without provenance

construction of the proof tree is performed level by level. The expansion of a node in the proof tree represents a single debugging query.

In Figure 18a, we show the time for proof tree fragments with levels up to a height of 20. We initiate the proof tree construction for randomly sampled output tuples in the DOOP DaCapo benchmarks. In the figure, we plot the runtime against the number of nodes in the proof tree fragment. Even for 20 levels, these proof trees contain over 15,000 nodes. Considering that full proof trees may have heights over 200, the corresponding full proof trees would be intractable to compute and understand due to exponential growth. However, this experiment shows that the construction of proof trees is approximately linear in the size of the tree. Therefore, provenance queries can be efficiently computed, and the method will scale well for interactive use. The interactive exploration of proof trees is scalable, with each debugging query on average taking less than 1 ms per node.

### 5.3 Characteristics of Proof Trees

In the following experiments, we demonstrate the difficulty of proof tree construction for Datalog specifications at large scale. Figure 18b shows the distribution of heights of full proof trees for the DaCapo benchmarks. The proof tree heights can be more than 300. While this may not seem prohibitive, the expected number of nodes in the proof tree is exponential in height. A non-linear regression performed on the sizes of actual proof trees, suggests that the branching factor of proof trees is approximately 1.466 for the DaCapo benchmarks. Therefore, since larger proof trees will have an exponential number of nodes, it is computationally intractable to construct full proof trees for these large specifications. Besides, there is a usability challenge in generating meaningful explanations for the existence of a tuple, which is addressed by the interactive exploration of fragments of a proof tree, with the user exploring only relevant fragments. This is in contrast to a full proof tree, where a user may have to interpret millions of nodes to find an explanation.

## 6 Related Work

Debugging for logic programming languages has a long history, with work having been done on algorithmic debugging strategies since the 1980s [16, 35]. These works present a framework for the algorithmic debugging method for Prolog programs, where a system asks the user questions about the intended model of the program to find buggy rules. However, they are based on the SLDNF resolution of Prolog which is not truly declarative, and thus the semantics of Datalog differ. Our method aligns practically with the interactive debugging frameworks presented here, but applied to the bottom-up evaluation of Datalog, and with sophisticated and efficient techniques to generate the debugging information.

**Debugging and Provenance.** Our method also fits into the established frameworks for provenance in Datalog [12] and debugging for Datalog [10]. The proof trees generated by our method are analogous to



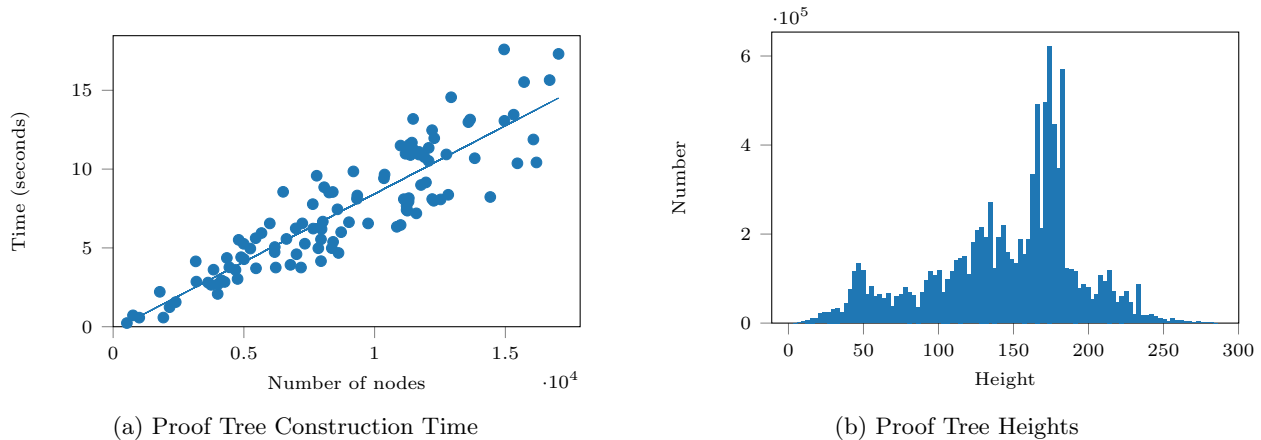


Figure 18: Proof Tree Construction and Statistics

the computation graphs presented in [10], and are equally effective for debugging. They can also be seen as an extension of *how*-provenance [9]. However, our hybrid method for generating proof trees is novel, i.e., it permits multiple debugging queries in a single debugging cycle. Hence, our method is especially useful for debugging large Datalog specifications.

**Provenance in Datalog.** Methods for computing provenance in Datalog has been a well-explored field [6, 26, 15, 13, 11, 27], however with the caveat that all these previous approaches store the full provenance information during the evaluation of Datalog. [26], for example, stores the whole computation graph as an auxiliary relation during Datalog evaluation, which may be many times larger than the IDB itself in large analysis use cases. Approaches such as [13, 27] attempt to reduce the impact of provenance storage by only storing information relevant for a particular provenance query, which is given before the instrumentation and evaluation of Datalog. Thus, in these approaches, the Datalog specification needs to be re-evaluated for each different provenance query and therefore extends the investigation phase of the debugging cycle. The closest approach to ours is perhaps [15], where a boolean circuit representation for provenance is described, as well as an algorithm for generating such a boolean circuit during Datalog evaluation. However, there is no mechanism for exploring an understandable provenance representation, and no practical implementation of this work. Therefore, our approach is novel by minimizing the storage overhead of provenance information, allowing interactive exploration of proof trees, all while providing an effective integration into an existing semi-naïve Datalog engine.

**Other Applications for Datalog Provenance.** Debugging Datalog specifications is not the only use case for provenance, with user-guided approaches [42, 31, 33, 43] for program analysis also relying on tracking the origins of data. In [33, 42, 31], a user may tag certain static analysis alarms, to increase or decrease their importance in the next analysis cycle. In [43], the analysis system automatically generates an appropriate abstraction, by iteratively trying and refining failing abstractions. All these approaches rely on an annotation framework for Datalog: the user-guided systems require the user to add an annotation representing the importance of an alarm, and the abstraction refinement system requires the system to tag failing analyses with annotations. In any case, our provenance evaluation strategy would fit well into these systems, by providing an annotation framework at the Datalog engine level.

**Provenance in Databases.** Outside of Datalog, provenance has also been a focus of the database community, being useful for understanding the origins of data in large database systems. For instance, *Trio* [41, 7] and *Perm* [18, 19] are two such systems implementing provenance systems for relational databases. These systems focus on tracking the *lineage* of data in a database, rather than on debugging a query, and thus it is essential that the provenance information is stored directly alongside the data. It is important to note that in the context of databases, Datalog acts as a powerful query language rather than a specification logic. As a result, the Datalog specifications in these use cases typically consist of fewer rules [29], that do not exhibit complex patterns found in program analysis benchmarks [8]. Further study has also been undertaken in querying database provenance, with [37, 4] both presenting mechanisms to construct provenance information lazily after the database query is run. Thus, similarly to our approach, these allow the querying of arbitrary

parts of provenance information. However, both approaches are applied to database systems, with [37] reconstructing provenance information based on tracking file I/O and system calls during query evaluation, and [4] based on system logs produced by the database system. Therefore, with the highly complex recursive nature of real-world Datalog specifications, similar information may blow up drastically, and experiments of both works only show scalability up to 10,000 tuples. Moreover, these approaches are unsuitable for our setting of in-memory analysis workloads.

## 7 Conclusion

In this paper we have presented a novel provenance evaluation strategy for Datalog specifications. The provenance evaluation strategy extends tuples in the IDB with proof annotations. With the help of proof annotations, provenance queries can construct minimal proof trees incrementally. Our method has very small overheads at logic evaluation time in comparison to standard top-down evaluation or a naïve provenance methods that encode provenance information explicitly. Hence, our method enables debugging of large-scale logic specifications for the first time. We have implemented our provenance method in a high performance Datalog engine called Soufflé [24], and demonstrated its feasibility through the DOOP program analysis framework. We show that the runtime overheads of the provenance evaluation strategy are approximately  $1.27\times$ , and the memory overheads are  $1.45\times$ .

## References

- [1] souffle-lang/souffle: Soufflé is a variant of datalog for tool designers crafting analyses in horn clauses. soufflé synthesizes a native parallel c++ program from a logic specification. more information can be found here: <http://souffle-lang.github.io/>, 2017. Accessed: 19-10-2017.
- [2] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley Publishing Company, 1995.
- [3] N. Allen, B. Scholz, and P. Krishnan. *Staged Points-to Analysis for Large Code Bases*, pages 131–150. Springer Berlin Heidelberg, 2015.
- [4] B. S. Arab, D. Gawlick, V. Krishnaswamy, V. Radhakrishnan, and B. Glavic. Using reenactment to retroactively capture provenance for transactions. *IEEE Transactions on Knowledge and Data Engineering*, 30(3):599–612, March 2018.
- [5] M. Aref, B. ten Cate, T. J. Green, B. Kimelfeld, D. Olteanu, E. Pasalic, T. L. Veldhuizen, and G. Washburn. Design and implementation of the logicblox system. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, SIGMOD ’15, pages 1371–1382, New York, NY, USA, 2015. ACM.
- [6] T. Arora, R. Ramakrishnan, W. G. Roth, P. Seshadri, and D. Srivastava”. Explaining program execution in deductive systems. *Proceedings of Deductive and Object-Oriented Databases: Third International Conference*, pages 101–119, 1993.
- [7] O. Benjelloun, A. D. Sarma, C. Hayworth, and J. Widom. An introduction to uldbs and the trio system. *IEEE Data Engineering Bulletin*, 2006.
- [8] M. Bravenboer and Y. Smaragdakis. Strictly declarative specification of sophisticated points-to analyses. *SIGPLAN Not.*, 44(10):243–262, 2009.
- [9] P. Buneman, S. Khanna, and W.-C. Tan. Why and where: A characterization of data provenance. *Proceedings of the International Conference on Database Theory*, 1973:316–330, 2001.
- [10] R. Caballero, Y. García-Ruiz, and F. Sáenz-Pérez. A theoretical framework for the declarative debugging of datalog programs. In K.-D. Schewe and B. Thalheim, editors, *Semantics in Data and Knowledge Bases*, pages 143–159, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.

- [11] R. Caballero, Y. García-Ruiz, and F. Sáenz-Pérez. Debugging of wrong and missing answers for datalog programs with constraint handling rules. 07 2015.
- [12] J. Cheney, L. Chiticariu, and W.-C. Tan. Provenance in databases: Why, how, and where. *Foundations and Trends in Databases*, 1:379–474, 2009.
- [13] D. Deutch, A. Gilad, and Y. Moskovitch. Selective provenance for datalog programs using top-k queries. *Proceedings of the VLDB Endowment*, 8:1394–1405, 2015.
- [14] D. Deutch, A. Gilad, and Y. Moskovitch. Efficient provenance tracking for datalog using top-k queries. *The VLDB Journal*, 27(2):245–269, Apr 2018.
- [15] D. Deutch, T. Milo, S. Roy, and V. Tannen. Circuits for datalog provenance. *Conference on Database Theory*, 17:201–212, 2014.
- [16] W. Drabent, S. Nadjm-Tehrani, and J. Maluszyński. Meta-programming in logic programming. pages 501–521, 1989.
- [17] Y. Futamura. Partial evaluation of computation process – an approach to a compiler-compiler. *Higher Order Symbol. Comput.*, 12(4):381–391, Dec. 1999.
- [18] B. Glavic and G. Alonso. Perm: Processing provenance and data on the same data model through query rewriting. *IEEE International Conference on Data Engineering*, 25:174–185, 2009.
- [19] B. Glavic, R. J. Miller, and G. Alonso. Using sql for efficient generation and querying of provenance information. *Lecture Notes in Computer Science*, 8000:291–320, 2013.
- [20] N. Grech, L. Brent, B. Scholz, and Y. Smaragdakis. Gigahorse: Thorough, declarative decompilation of smart contracts. In *Proceedings of the 41th International Conference on Software Engineering, ICSE 2019*, page (to appear), Montreal, QC, Canada, May 2019. ACM.
- [21] N. Grech, M. Kong, A. Jurisevic, L. Brent, B. Scholz, and Y. Smaragdakis. Madmax: Surviving out-of-gas conditions in ethereum smart contracts. In *SPLASH 2018 OOPSLA*, 2018.
- [22] K. Hoder, N. Bjørner, and L. de Moura.  $\mu z$  – an efficient engine for fixed points with constraints. In G. Gopalakrishnan and S. Qadeer, editors, *Computer Aided Verification*, pages 457–462, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.
- [23] S. S. Huang, T. J. Green, and B. T. Loo. Datalog and emerging applications: An interactive tutorial. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*, SIGMOD ’11, pages 1213–1216. ACM, 2011.
- [24] H. Jordan, B. Scholz, and P. Subotic. Soufflé: On synthesis of program analyzers. *Proceedings of Computer Aided Verification*, 28:422–430, 2016.
- [25] H. Jordan, P. Subotić, D. Zhao, and B. Scholz. A specialized b-tree for concurrent datalog evaluation. In *Proceedings of the 24th Symposium on Principles and Practice of Parallel Programming*, PPOPP ’19, pages 327–339, New York, NY, USA, 2019. ACM.
- [26] S. Köhler, B. Ludäscher, and Y. Smaragdakis. Declarative datalog debugging for mere mortals. *Lecture Notes in Computer Science*, 7494:111–122, 2012.
- [27] S. Lee, S. Köhler, B. Ludäscher, and B. Glavic. Efficiently computing provenance graphs for queries with negation. *CoRR*, abs/1701.05699, 2017.
- [28] S. Lee, B. Ludäscher, and B. Glavic. Provenance summaries for answers and non-answers. *Proc. VLDB Endow.*, 11(12):1954–1957, Aug. 2018.
- [29] S. Liang, P. Fodor, H. Wan, and M. Kifer. Openrulebench: An analysis of the performance of rule engines. In *Proceedings of the 18th International Conference on World Wide Web*, WWW ’09, pages 601–610, 2009.

- [30] M. Madsen, M.-H. Yee, and O. Lhoták. From datalog to flix: A declarative language for fixed points on lattices. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '16, pages 194–208, New York, NY, USA, 2016. ACM.
- [31] R. Mangal, X. Zhang, A. V. Nori, and M. Naik. A user-guided approach to program analysis. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*, ESEC/FSE 2015, pages 462–473, New York, NY, USA, 2015. ACM.
- [32] X. Ou, S. Govindavajhala, and A. W. Appel. Mulval: A logic-based network security analyzer. In *Proceedings of the 14th Conference on USENIX Security Symposium - Volume 14*, SSYM'05, pages 8–8, Berkeley, CA, USA, 2005. USENIX Association.
- [33] M. Raghothaman, S. Kulkarni, K. Heo, and M. Naik. User-guided program reasoning using bayesian inference. pages 722–735, 06 2018.
- [34] R. Ramakrishnan and S. Sudarshan. Top-down vs. bottom-up revisited. In *In Proceedings of the International Logic Programming Symposium*, pages 321–336. MIT Press, 1991.
- [35] E. Y. Shapiro. *Algorithmic Program DeBugging*. MIT Press, Cambridge, MA, USA, 1983.
- [36] M. Sridharan, D. Gopan, L. Shan, and R. Bodík. Demand-driven points-to analysis for java. In *Proceedings of the 20th Annual ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications*, OOPSLA '05, pages 59–76, New York, NY, USA, 2005. ACM.
- [37] M. Stamatogiannakis, P. Groth, and H. Bos. Decoupling provenance capture and analysis from execution. In *7th USENIX Workshop on the Theory and Practice of Provenance (TaPP 15)*, Edinburgh, Scotland, 2015. USENIX Association.
- [38] P. Subotic, H. Jordan, L. Chang, A. Fekete, and B. Scholz. Automatic index selection for large-scale datalog computation. *PVLDB*, 12(2):141–153, 2018.
- [39] J. D. Ullman. Bottom-up beats top-down for datalog. In *Proceedings of the Eighth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, PODS '89, pages 140–149, New York, NY, USA, 1989. ACM.
- [40] J. Whaley, D. Avots, M. Carbin, and M. S. Lam. *Using Datalog with Binary Decision Diagrams for Program Analysis*, pages 97–118. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.
- [41] J. Widom. Trio: A system for integrated management of data, accuracy, and lineage. In *2nd Biennial Conference on Innovative Data Systems Research, CIDR 2005*, pages 262–276, 01 2005.
- [42] X. Zhang, R. Grigore, X. Si, and M. Naik. Effective interactive resolution of static analysis alarms. *Proc. ACM Program. Lang.*, 1(OOPSLA):57:1–57:30, Oct. 2017.
- [43] X. Zhang, R. Mangal, R. Grigore, M. Naik, and H. Yang. On abstraction refinement for program analyses in datalog. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '14, pages 239–248, New York, NY, USA, 2014. ACM.
- [44] D. Zhao. Honours thesis: Large-scale provenance for soufflé. 2017.
- [45] W. Zhou, M. Sherr, T. Tao, X. Li, B. T. Loo, and Y. Mao. Efficient querying and maintenance of network provenance at internet-scale. *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*, pages 615–626, 2010.