Towards Elastic Incrementalization for Datalog

David Zhao
University of Sydney
dzha3983@uni.sydney.edu.au

Mukund Raghothaman
University of Southern California
raghotha@usc.edu

Pavle Subotić
Microsoft
pavlesubotic@microsoft.com

Bernhard Scholz
University of Sydney
bernhard.scholz@sydney.edu.au

ABSTRACT

Various incremental evaluation strategies for Datalog have been developed that reuse computations for small input changes. These methods assume that incrementalization is always a better strategy than recomputation. However, in real-world applications such as static program analysis, recomputation can be cheaper than incrementalization for large updates.

This work introduces an elastic incremental approach with two strategies that can be selected based on the impact of the input change. The first strategy is a *Bootstrap* strategy that recomputes the entire result for high-impact changes. The second is an *Update* strategy that performs an incremental update for low-impact changes. Our approach allows for a lightweight Bootstrap strategy suitable for high-impact changes, with the trade-off that Update may require more work for small changes. We demonstrate our approach using real-world applications and compare our elastic incremental approach to existing methods.

CCS CONCEPTS

• **Software and its engineering** → **Constraint and logic languages**; *Incremental compilers*; • **Information systems** → *Relational database query languages*; • **Theory of computation** → *Interactive computation*; *Streaming models*; *Program analysis*; *Data provenance*.

KEYWORDS

Datalog, incremental evaluation, Datalog compilers, provenance

1 INTRODUCTION

Logic languages such as Datalog have seen widespread adoption in recent years in areas such as static program analysis [2, 8, 12, 14], declarative networking [4, 20, 47], security analysis [33], business applications [3] and machine learning [27]. The main reasons for the widespread adoption have been the availability of high-performance logic engines [3, 17, 22] and the ease of expressing programs declaratively, i.e., computations can be expressed succinctly, providing means for rapid prototyping of scientific and industrial applications.

The standard Datalog evaluation (which we refer to as *batch-mode*), computes the Intensional Database (IDB, or the *output* facts), given a set of rules, and an Extensional Database (EDB, or the *input* facts). However, many real-world applications recompute most of their IDB with slight variations of the EDB [26, 31]. Hence, several state-of-the-art Datalog engines have proposed incremental evaluation techniques [24, 26, 29] to facilitate streaming, i.e., the evaluation reuses the IDB from the previous computation to compute the new IDB given some changes to the EDB.

State-of-the-art incremental evaluation approaches operate on several assumptions: (1) that the impact, i.e., number of overall tuple changes, is proportional to the update size, and (2) that the use cases exhibit a continuous stream of small impact updates. Indeed for several use cases [31] these assumptions tend to hold. However, for other notable use cases such as program analysis in a continuous integration/continuous delivery (CI/CD) setup [10, 11, 42] these assumptions do not hold. For example, static analyses written in Datalog can consist of hundreds or thousands of highly recursive rules and relations [8, 13]. Due to the complexity of the ruleset, one can no longer assume that update size is proportional to the impact size. Our experimental evaluation on the Doop program analysis framework found large variability in the *impact* of updates due to the connectivity of points-to analyses, where even small program changes may substantially change pointer sets of variables. Another concern is that static analyzers are deployed in CI/CD pipelines where state-of-the-art incremental evaluation gives no guarantee that updates will be structurally small. For instance, when the code base is updated, the initial change is often a refactor or new feature implementation. Such code changes typically result in large changes to the input of an analysis. These changes may then be followed by smaller changes resulting from minor review suggestions, but as we show, even these smaller input changes do not necessarily result in small impacts. Thus, we argue the success of incremental evaluation techniques for such use cases requires minimizing the overhead of evaluating large impact updates.

(a) Batch-mode   **(b) Incremental**

Figure 1: Batch-mode vs. Incremental Evaluation

(a) Bootstrap Strategy   **(b) Update Strategy**

(c) Elastic Incremental, E_4 causes a restart in the stream using Bootstrap

Figure 2: Elastic Incremental Evaluation

Consider Fig. 1 that illustrates a generic incremental computation setup. Fig. 1a shows a *batch-mode* evaluation for a Datalog program P and EDBs E_1, E_2, and E_3. The batch-mode evaluation runs the program P with each EDB separately to produce the IDBs I_1, I_2, and I_3. However, if only a small portion of the EDB and IDB changes between runs, many computations in the batch-mode evaluation may have been repeated. An incremental evaluation ΔP (illustrated in Fig. 1b) can reuse computations from a previous run, which we call an *epoch*. A *Computational State* σ_1 encodes the previous computations for I_1 in a special format so that the next run can reuse the computations. With state σ_1 and the change in input ΔE_2, the incremental evaluation produces the output I_2 and the new computational state σ_2. This process may be repeated, with a series of updates to the EDB being provided via ΔE_i. For the first epoch, we use the empty state as computational state and E_1 as ΔE_1 to produce I_1 and the state σ_1. Any subsequent change ΔE_i in the EDB is processed by using the previous computational state σ_{i-1} to generate I_i.

State-of-the-art incremental evaluation strategies [26, 28, 29] represent their computational state exhaustively to perform small updates efficiently. Because of their exhaustive representation, however, initiating a stream with the state-of-the-art approaches can be prohibitively slow and cannot be used to react to large updates. For example, when a static program analysis seeks to reuse previous computations for a large code refactoring, significant portions of the control flow graph may have been replaced. In such a use case, an incremental evaluation will essentially perform two computations, one to delete the old control flow graph, and one to compute the new control flow graph with additional overheads caused by the incrementalization. Therefore, these heavyweight updates are better served by an evaluation strategy that is closer to standard batch-mode evaluation augmented with state for the future updates to be performed incrementally.

In this work, we demonstrate that both fully non-incremental and fully incremental strategies are not effective in some scenarios. Therefore, we propose an *elastic* incremental evaluation scheme called *Bootstrap-Update*, which is a hybrid approach. Our approach has two distinct strategies to evaluate an update: a specialized *Bootstrap* denoted as P_b (Fig. 2a) and *Update* denoted as P_u (Fig. 2b). The specialized Bootstrap resembles an augmented batch-mode evaluation that produces the computational state from scratch to allow subsequent updates, whereas Update is an incremental evaluation strategy.

Our approach proposes a novel *sparse* encoding that maintains a *lightweight* state σ. Our state exhibits a space complexity of $O(|I|)$ (i.e. linear in the size of the output) whereas existing incremental encodings [26, 29] have a worst-case space complexity of $O(m|I|)$ where m is the number of fixpoint iterations in the semi-naive evaluation algorithm [1]. Our lightweight state allows for an accelerated Bootstrap algorithm that can handle high-impact updates by efficiently recomputing the state from scratch, with the trade-off that the Update strategy may require more work for smaller updates. Furthermore, we provide a simple heuristic for choosing the appropriate strategy: we re-run the bootstrap when the incremental update takes more than a fraction (as a *switching parameter*) of the last bootstrap's runtime. This switching parameter typically depends on the behavior of each application and the typical update characteristics for that application. Our solution operates under the insight that if we have comparable performance with batch-mode Datalog evaluation on large impact updates and a small slow down on low impact updates, we will have an overall net gain by selective application of incremental evaluation.

We have integrated our elastic Bootstrap-Update incremental evaluation in the open-source, high-performance Datalog engine Soufflé [23]. We have performed an extensive evaluation on a number of use cases that shows our approach's utility compared to existing techniques on both large and small updates. We also provide a discussion of the practical considerations for building incremental evaluation in Soufflé that include relational data structures, parallelization, and scheduling strategies.

In summary, we make the following contributions in this paper:

(1) We present a new problem - that incremental evaluation should be *elastic*, i.e., it should be sensitive to the impact of an update.

(2) We present a novel incremental evaluation using a sparse derivation counting encoding, exhibiting superior performance and lower memory overhead for elastic use cases.

(3) We extend Soufflé, an open-source Datalog evaluation engine for elastic incremental evaluation, and propose several engine optimizations for superior performance.

(4) We provide an extensive experimental evaluation validating the utility of our contribution.

2 BACKGROUND

In this section, we provide an example to explain the background of Datalog evaluation.

2.1 Example: Datalog Pointer Analysis

We present a program analysis example written in Datalog that computes pairs of variables that may *alias* in a source program.

Figure 3a shows a fragment of object-oriented source code encoded in the form of relations in Figure 3b and represented diagrammatically in Figure 4. From this encoded relational representation of the source program, a (field sensitive but flow insensitive [41]) pointer analysis is written in Datalog, in Figure 3c. In this analysis, the input relations (also known as Extensional Database, or EDB) are new, assign, load, and store, each of which represents a particular type of operation in the source program. During the analysis, the Datalog specification computes output relations (also known as Intensional Database, or IDB) vpt, which relates variables and the objects that they point to, and alias, which relates pairs of variables that may point to the same object.

This logic specification consists of four *rules* (here labeled r1 through r4). Each rule is a Horn clause consisting of two parts: the *predicate* on the left of the implication sign (:-) is the *head*, and the set of predicates on the right is the *body*. Each predicate consists of a *relation name* and a sequence of constants and variables of appropriate arity. For example, the rule

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

has the predicate vpt(Var,Obj) as the head, and the two predicates assign(Var,Var2) and vpt(Var2,Obj) as the body. Negation and constraints are omitted for now but are discussed in more detail in Section 3.3.

A predicate may be *instantiated*, where all its variables are mapped to constants to form a *tuple*. An instantiated rule is a rule where each predicate is instantiated, such that the variable mappings are compatible between all the predicates. A Datalog rule is read from right to left as a universally quantified implication: "for all rule instantiations, if every tuple in the body is derivable, then the corresponding tuple for the head is also derivable".

2.2 Semi-Naïve Evaluation

To evaluate a Datalog specification, modern engines use a *bottom-up* approach, which begins from the input tuples, and in each step attempts to derive more tuples using an *immediate consequence operator* $\Gamma_P(I) = I \cup \{t \mid r = t :\text{-} t_1, \ldots, t_n, \text{ each } t_i \in I\}$ such that r is a valid instantiation of a rule in P with each $t_i \in I$. The evaluation ends when a *fixpoint* is reached. Many Datalog solvers improve on this bottom-up strategy by utilizing *semi-naïve* evaluation. Semi-naïve evaluation proposes two main optimizations: (1) Stratification: the Datalog specification is split into *strata*. Firstly, a precedence graph of relations is computed, where there is an edge from relation R_{body} to R_{head} if R_{body} appears in the body of a rule with R_{head} in the head. 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 relations in 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 at least one tuple generated in the previous iteration, avoiding the recomputation of tuples already computed in prior iterations.

The standard semi-naïve evaluation is presented in Algorithm 1 for a single stratum. The inputs for the algorithm are E, the input set of tuples (since this is a single stratum, the input may be EDB tuples or tuples from earlier strata), and P, the set of Datalog rules forming the stratum.

Algorithm 1 Semi-Naïve(E, P)

1: $\Delta_0 \leftarrow E$
2: **for all** $k \in \{1, 2, \ldots\}$ **do**
3: $\quad I_{k-1} \leftarrow \bigcup_{0 \leq i < k} \Delta_i$
4: $\quad \Delta_k \leftarrow \Pi_P[I_{k-1} \mid \Delta_{k-1}] \setminus I_{k-1}$
5: $\quad$ **if** $\Delta_k = \emptyset$ **then**
6: $\quad\quad$ **return** I_{k-1}

This algorithm begins by initializing the delta and the full set of tuples from the input (line 1). In the fixpoint loop, line 4 is the critical line, evaluating the Datalog rules. This line uses notation adapted from [29], which introduces a *rule evaluation operator*, Π, where

$$\Pi_P[I \mid \Delta] = \left\{ t \;\middle|\; \begin{array}{l} t :\text{-} t_1, \ldots, t_n \text{ is an instantiated rule in } P \\ \text{where } \{t_1, \ldots, t_n\} \subseteq I \text{ and } \{t_1, \ldots, t_n\} \cap \Delta \neq \emptyset \end{array} \right\}$$

Here, Π_P finds the head tuples of all rules in P instantiated from tuples in I, where at least one body tuple also exists in Δ. For the rest of this paper, the program P is omitted from Π_P where it is clear. The dependence on Δ is the new knowledge optimization in semi-naïve evaluation. By requiring that at least one body tuple for each rule derivation is contained in Δ_{k-1}, the algorithm ensures that new tuples are only generated if at least one body tuple was new in the previous iteration.

Algorithm 1 continues by merging the newly discovered tuples into the full relation (line 3). Then, if a fixpoint has been reached (i.e., no new tuples are generated), the evaluation ends (line 5).

As a concrete example of semi-naïve evaluation, consider the recursive stratum containing vpt in the running example. In the initialization phase, the algorithm simply copies the inputs. Therefore, in iteration 0,

$$\Delta_0 = I_0 = \left\{ \begin{array}{l} \mathsf{new(a,L1), new(c,L3), new(d,L4),} \\ \mathsf{assign(a,b), assign(b,a), store(c,f,a),} \\ \mathsf{load(e,d,f), load(b,c,f)} \end{array} \right\}$$

In iteration 1, note that the vpt relation is empty in I_0. Therefore, only the non-recursive rule r1 can be applied, generating

$$\Delta_1 = \{\mathsf{vpt(a,L1), vpt(c,L3), vpt(d,L4)}\}$$
$$I_1 = I_0 \cup \Delta_1$$

Using I_1 and Δ_1, the algorithm can now apply the recursive rules of vpt as well. Rule r1 no longer applies since there are no tuples from relation new in Δ_1. From rule r2, we can derive vpt(b,L1) from the instantiation vpt(b,L1) :- assign(b,a), vpt(a,L1). From rule r3, we can again derive vpt(b,L1), from vpt(b,L1) :-

```
L1: a = new O();      new(a, L1).      vpt(Var, Obj) :- new(Var, Obj).      //r1
L2: b = a;            assign(b, a).    vpt(Var, Obj) :- assign(Var, Var2),
                                                        vpt(Var2, Obj).      //r2
L3: c = new P();      new(c, L3).      vpt(Var, Obj) :- load(Var, Y, F),
L4: d = new P();      new(d, L4).                       store(P, F, Q),
                                                        vpt(Q, Obj),
L5: c.f = a;          store(c, f, a).                   vpt(P, Obj2),
L6: e = d.f;          load(e, d, f).                    vpt(Y, Obj2).        //r3
L7: b = c.f;          load(b, c, f).   alias(Var1, Var2) :- vpt(Var1, Obj),
L8: a = b;            assign(a, b).                        vpt(Var2, Obj). //r4
```

(a) Input Program   (b) EDB Tuples   (c) Datalog Pointer Analysis

Figure 3: Program Analysis Datalog Setup

Figure 4: Pointer Input Diagram

load(b,c,f), store(c,f,a), vpt(a,L1), vpt(c,L3), vpt(c,L3). Therefore, these two derivations generate the same tuple, and so,

$$\Delta_2 = \{vpt(b, L1)\}$$
$$I_2 = I_1 \cup \Delta_2$$

In iteration 3, rule r2 can generate vpt(a,L1). However, this tuple is already contained in I_2, and therefore $I_3 = I_2$ and a fixpoint is reached.

2.3 Incremental Datalog Evaluation

Incremental evaluation refers to a procedure to *update* the result of the Datalog computation given some changes in the input without performing a full recomputation. An incremental evaluation proceeds in *epochs*, where each epoch represents one round of updates, i.e., inserting/deleting tuples from the input and computing the new result and state. We refer to the inserted and deleted tuples as the *diff*. For the workflow in Fig. 2c, each I_k represents the result of epoch k, and each ΔE_k represents the corresponding diff. To summarize, the central problem of incremental evaluation is as follows:

Definition 2.1 (Incremental Evaluation). Given a Datalog program P, an input data set E, the result $P(E)$, an insertion set E^+ and a deletion set E^-, compute the result $P((E \cup E^+) \setminus E^-)$.

The cost of an update can be measured by its impact. Typically, high impact changes result in greater computation overhead.

Definition 2.2 (Incremental Update Impact). The impact of an update is the number of IDB tuples changed as a consequence of the update, i.e., ΔI.

We note that while the state-of-the-art incremental evaluation strategies, such as DRed [16], its related strategies [18, 19, 29], and counting-based algorithms [26, 28] have proven worthwhile for applications where each update has a small impact on the computed result, we have observed that this assumption does not hold in general for all incremental workloads. For a concrete example, consider our running example. We may remove the line L6 in Figure 3a as part of an update to the software. This removed line would result in the input tuple load(e,d,f) being removed. From the graph in Figure 4, this only affects a single edge, and does not affect the connected component containing a, b, and c. Therefore, computing the result after performing this update should take advantage of this separation, and this update has *small* impact. However, imagine also removing the line L1 as part of the same software update. Then, the input tuple new(a,L1) would be deleted, and both connected components in Figure 4 would be affected. This results in an update with *large* impact, where half of the tuples in vpt are deleted, and all of the tuples in alias are deleted. In these situations, where both small and large updates may be present, state-of-the-art incremental evaluation strategies may not be effective.

3 ELASTIC INCREMENTAL EVALUATION

This section describes our encoding and algorithms for elastic incremental evaluation. Recall from Fig. 2 that we have two evaluation procedures, one to initialize the computation state and one to incrementally update it. We call these evaluations *Bootstrap* and *Update* strategies, respectively (see Fig. 2c). Our Bootstrap strategy mimics a standard semi-naïve evaluation and also computes the computational state to allow subsequent updates. The Bootstrap strategy either initiates the streaming or is a restart strategy for large updates. Recall that the update strategy needs a computational state σ, which is maintained in each epoch. In previous approaches [26, 28], this computational state involves a vector of numbers per tuple in the IDB. Each number in the vector represents a count in some fixpoint iteration. In the worst case, the length of the vector is determined by the number of iterations m in the fixpoint computation. Hence, the state may exhibit a worst-case space complexity of $O(m|I|)$ where $|I|$ is the size of the output.

In contrast, our approach maintains a lightweight computational state consisting of two numbers per tuple. The first number is a *derivation count*, and the second number is the iteration in which the tuple is first derived. The derivation count represents the number of ways that the tuple can be derived in the iteration when it is first deduced and allows the reuse of computation in the next epoch. Furthermore, our encoding is a sparse version of the vector of numbers in previous approaches, keeping only the first iteration rather than the whole vector. Therefore, the worst-case space complexity of this encoding is $O(|I|)$.

With our lightweight computational state σ, we can switch between Bootstrap and Update to adapt to lightweight and heavyweight updates accordingly. When given an incremental update, we provide a heuristic for switching between both strategies. We first attempt the Update strategy. If it times out (the timeout is set to some fraction using a *switching parameter* of the previous Bootstrap's runtime strategy), we discard its partial state and produce the output and computational state from scratch using Bootstrap. The timeout is dependent on the application and needs to be fine-tuned appropriately. In contrast, previous approaches have a single strategy and cannot adapt to light and heavy updates.

We introduce some notation for describing our approach. We define a sequence of sets $\langle D_1, D_2, \ldots \rangle$ where set D_k denotes the set of *rule instantiations*. Set $D_k = \{(t :\!- t_1, \ldots, t_n)\}$ contains all the rule instantiations that are computed in iteration k. The derivation count of tuple t in iteration k is the number of rule instantiations $(t :\!- t_1, \ldots, t_n)$ whose head is t.

For the sake of simplicity, we define $\mathcal{N}^{\#}$ as a sequence of counting multisets for describing the derivation counts of tuples. We use the standard definition of multisets, where each $\mathcal{N}_k^{\#} = \{(t \mapsto c)\}$ denotes the number of rule instantiations $t :\!- t_1, \ldots, t_n$ for tuple t in D_k. For notational convenience, we will express the elements with multiplicities $t \mapsto c$ as t^c, and we use N_k to denote the set projection of $\mathcal{N}_k^{\#}$.

3.1 Bootstrap Algorithm

The Bootstrap algorithm is a specialized counting algorithm that efficiently computes the sequence of multisets from scratch, mimicking a semi-naive evaluation while also producing the incremental computation state.

For example, consider our running example. In the initial phase, the input E becomes iteration 0, where the counting semantics mean that every tuple has a count of 1. Therefore,

$$\mathcal{N}_0^{\#} = \left\{ \begin{array}{l} \mathsf{new(a, L1)}^1, \mathsf{new(c, L3)}^1, \mathsf{new(d, L4)}^1, \\ \mathsf{assign(a, b)}^1, \mathsf{assign(b, a)}^1, \mathsf{store(c, f, a)}^1, \\ \mathsf{load(e, d, f)}^1, \mathsf{load(b, c, f)}^1 \end{array} \right\}$$

In iteration 1, we apply the non-recursive rule r1. In this case, all tuples have a count of 1:

$$\mathcal{N}_1^{\#} = \left\{ \mathsf{vpt(a, L1)}^1, \mathsf{vpt(c, L3)}^1, \mathsf{vpt(d, L4)}^1 \right\}$$

In iteration 2, however, the counting semantics causes a divergence from the standard semi-naïve evaluation. Recall that the tuple vpt(b,L1) is derivable from two rules:

(1) vpt(b, L1) :- assign(b, a), vpt(a, L1), and
(2) vpt(b, L1) :- load(b, c, f), store(c, f, a),
        vpt(a, L1), vpt(c, L3), vpt(c, L3)

Therefore, vpt(b, L1) has a count of 2 in iteration 2:

$$\mathcal{N}_2^{\#} = \left\{ \mathsf{vpt(b, L1)}^2 \right\}$$

In iteration 3, no new tuples are derivable. Therefore, a fixpoint has been reached, and the Datalog evaluation ends.

The bootstrap algorithm is illustrated in Algorithm 2. The main extension from the standard semi-naïve evaluation (Algorithm algorithm generates $\mathcal{N}^{\#}$, the sequence of multisets, in contrast to the standard sets in standard semi-naïve. To compute these multisets, we first introduce a version of the rule evaluation operator that computes sets of *rule instantiations*:

$$\Pi_P^D[I \mid I_{\mathrm{in}}] =$$
$$\left\{ (t :\!- t_1, \ldots, t_n) \left| \begin{array}{l} t :\!- t_1, \ldots, t_n \text{ in } P \text{ where } \{t_1, \ldots, t_n\} \subseteq I \\ \text{and } \{t_1, \ldots, t_n\} \cap I_{\mathrm{in}} \neq \emptyset \end{array} \right. \right\}$$

From the Π_P^D operator, we can define a counting version:

$$\Pi_P^{\#}[I \mid I_{\mathrm{in}}] =$$
$$\left\{ t^v \mid v = \#\text{rule instantiations } (t :\!- t_1, \ldots, t_n) \in \Pi_P^D[I \mid I_{\mathrm{in}}] \right\}$$

where v is the number of ways that the tuple t can be derived.

Algorithm 2 presents the lightweight bootstrap algorithm for a single stratum. Its structure is almost identical to the standard semi-naïve evaluation algorithm. The main difference for Bootstrap is that it maintains a separate sequence of multisets $\mathcal{N}^{\#}$, where each $\mathcal{N}_k^{\#}$ is similar to Δ_k of semi-naïve, and contains all the new tuples computed in iteration k. The algorithm begins by initializing $\mathcal{N}_0^{\#}$ to be equal to E (line 1). In the fixpoint loop, the algorithm first creates a set projection of the current iteration's multiset (line 3), where the operator taking the support of a multiset is defined as $\mathrm{Supp}(\mathcal{N}_{k-1}^{\#}) = \{t \mid (t^c) \in \mathcal{N}_{k-1}^{\#} \text{ and } c > 0\}$. In other words, $\mathrm{Supp}(\mathcal{N})$ is the set projection of tuples in $\mathcal{N}_0$. The algorithm also computes the full state of the relations up to iteration $k - 1$ (line 4), in the same way as the semi-naïve algorithm. These two auxiliary sets, N_{k-1} and I_{k-1}, are used in the rule evaluation on line 5. This rule evaluation computes all tuples that are new in the current iteration, and excludes any tuples that were computed in earlier iterations. By excluding existing tuples, the algorithm maintains the sparsification property, exhibiting a space complexity of $O(|I|)$. The algorithm exits and returns the evaluation state $(E, \mathcal{N}^{\#})$ (line 6) if no new tuples are generated in the current iteration, which is checked via the emptiness of the set projection of $\mathcal{N}_k^{\#}$.

Algorithm 2 Bootstrap(E)

1: $\mathcal{N}_0^{\#} \leftarrow \{(t^1) \mid t \in E\}$
2: **for all** $k \in \{1, 2, \ldots\}$ **do**
3: $\quad N_{k-1} \leftarrow \mathrm{Supp}(\mathcal{N}_{k-1}^{\#})$
4: $\quad I_{k-1} \leftarrow \cup_{0 \leq i \leq k-1} N_i$
5: $\quad \mathcal{N}_k^{\#} \leftarrow \{(t^v) \in \Pi^{\#}[I_{k-1} \mid N_{k-1}] \mid t \notin I_{k-1}\}$
6: $\quad$ **if** $\mathrm{Supp}(\mathcal{N}_k^{\#}) = \emptyset$ **then**
7: $\quad\quad$ **return** $(E, \mathcal{N}^{\#})$

Correctness. To demonstrate the correctness of Algorithm 2, we need to show that it computes the same resulting set of tuples as standard semi-naïve evaluation (Algorithm 1). To do this, we need to demonstrate two basic properties: *(a)* each N_k of Bootstrap is equal to Δ_k of semi-naïve, and *(b)* both Bootstrap and semi-naïve evaluation terminate after the same number of iterations.

To show this, we introduce the following lemma:

LEMMA 3.1. *Given a Datalog program P, for all A, B such that $B \subseteq A$, $\mathrm{Supp}(\Pi_P^{\#}[A \mid B]) = \Pi_P[A \mid B]$.*

This property can be shown since a tuple $t \in \Pi_P[A \mid B]$ if and only if there is a rule instantiation that computes it. If this is the case, then the same rule instantiation also fits $\Pi_P^{\#}[A \mid B]$ with a count of at least one. As a corollary, we can show that Bootstrap and semi-naïve both produce the same set of tuples in each iteration.

LEMMA 3.2. *Given a Datalog program P and an input set E, each I_{k-1} of Bootstrap is equal to I_{k-1} of semi-naïve.*

The proof is by induction over k, since $\mathrm{Supp}(\mathcal{N}_i^{\#}) = \Delta_i$ (from Lemma 3.1) for each iteration i, then each iteration's result must be identical to semi-naïve. Note that both Bootstrap and semi-naïve terminate after the same number of iterations since $\mathrm{Supp}(\mathcal{N}_i^{\#}) = \Delta_i$ for every iteration i, and therefore $\mathrm{Supp}(\mathcal{N}_i^{\#}) = \emptyset$ if and only if $\Delta_i = \emptyset$. Therefore, both algorithms terminate after the same number of iterations and thus produce the same set of resulting tuples.

3.2 Incremental Update Algorithm

The Update algorithm is a procedure that takes a computational state, either computed by Bootstrap or by a previous Update, and a set of changes to the inputs. The algorithm returns the updated computational state after applying the input changes.

The Update algorithm produces a computational state $\mathcal{N}_k^{\#}$ from the computational state $\mathcal{N}_k^{o}$ of the previous epoch, following the iterations of the previous epoch's fixpoint. In each iteration, the algorithm applies the insertions and deletions resulting from the given changes to the input.

For example, consider the running example, where we remove $\mathsf{assign}(b, a)$ and insert $\mathsf{store}(d, f, c)$. Iteration 0 reflects these updates:

$$\mathcal{N}_0^{\#} = \left\{ \begin{array}{l} \mathsf{new}(a, L1)^1, \mathsf{new}(c, L3)^1, \mathsf{new}(d, L4)^1, \\ \mathsf{assign}(a, b)^1, \mathsf{store}(d, f, c)^1, \mathsf{store}(c, f, a)^1, \\ \mathsf{load}(e, d, f)^1, \mathsf{load}(b, c, f)^1 \end{array} \right\}$$

In iteration 1, the inputs for the non-recursive rule (i.e., the relation new) don't change, thus

$$\mathcal{N}_1^{\#} = \left\{ \mathsf{vpt}(a, L1)^1, \mathsf{vpt}(c, L3)^1, \mathsf{vpt}(d, L4)^1 \right\}$$

In iteration 2, however, the deletion of $\mathsf{assign}(b, a)$ means that $\mathsf{vpt}(b, L1)$ is no longer derivable from the rule

```
vpt(b,L1) :- assign(b,a), vpt(a,L1)
```

Meanwhile, the insertion of $\mathsf{store}(d, f, c)$ means that a new tuple, $\mathsf{vpt}(e, L3)$, can now be derived from

```
vpt(e,L3) :- load(e,d,f), store(d,f,c),
       vpt(c,L3), vpt(d,L4), vpt(d,L4)
```

Therefore, the *diff* in iteration 2 can be expressed as

$$\left\{ \mathsf{vpt}(b, L1)^{-1}, \mathsf{vpt}(e, L3)^{+1} \right\}$$

The result for iteration 2 after applying this diff is

$$\mathcal{N}_2^{\#} = \left\{ \mathsf{vpt}(b, L1)^1, \mathsf{vpt}(e, L3)^1 \right\}$$

At this point, no new tuples can be further derived, either by tuples that already existed previously or by tuples that are newly inserted. Therefore, a fixpoint is reached, and the evaluation terminates.

Our novel sparse computational state requires some notion of *re-discovery* since a tuple only exists in the first iteration. If a tuple is deleted in its first iteration, it may still be derivable in a later iteration. In this case, the Update algorithm re-discovers whether the tuple is either derived in a later iteration or is truly deleted from the IDB. This re-discovery process is a notion of provenance [9, 46], where we find derivations for tuples that are deleted in earlier iterations.

The incremental update algorithm introduces extended notation over Bootstrap. The rule evaluation notation, $\Pi_P[I \mid I_{\mathrm{in}}]$, denotes tuples resulting from rules in P instantiated from I, with at least one body tuple also in I_{in}. This is extended with

$$\Pi_P^{\#}[I \mid I_1 \mid I_2] =$$
$$\left\{ t^v \middle| \begin{array}{l} v = \text{number of rule instantiations } t :\text{-} t_1, \ldots, t_n \\ \text{in } P \text{ where } \{t_1, \ldots, t_n\} \subseteq I \text{ and } \{t_1, \ldots, t_n\} \cap I_1 \neq \emptyset \\ \text{and } \{t_1, \ldots, t_n\} \cap I_2 \neq \emptyset \end{array} \right\}$$

This notation derives tuples from rule instantiations where at least one body tuple is from I_1, and also at least one body tuple is from I_2. In Update, I_1 and I_2 would be the deltas from semi-naïve evaluation and the diffs from the incremental update, respectively, allowing the rule evaluation to compute newly changed tuples in the current iteration of the current epoch. Additionally, the algorithm uses $\oplus$ and $\ominus$, the standard multiset addition and subtraction operators for operations involving multisets.

The update algorithm computes the updates to the sequence of multisets $\mathcal{N}^{\#}$, which result from applying the insertions and deletions to the input. The algorithm also makes use of a number of auxiliary sets: I_k^o and I_k maintain the full sets of tuples up to iteration k for the previous and current epoch respectively, I_k^- and I_k^+ maintain the tuples that are deleted and inserted respectively up to iteration k, and N_k^o and N_k are the set projections of $\mathcal{N}_k^{\#o}$ and $\mathcal{N}_k^{\#}$, storing the tuples that are new in iteration k in the previous and current epoch, respectively.

Algorithm 3 is presented for a single stratum and takes the state of the previous epoch $(E, \mathcal{N}^{\#o})$, and the incremental update (E^-, E^+) consisting of a set of tuples to be deleted and a set of tuples to be inserted, respectively. Note that $\mathcal{N}^{\#}$ may be the IDB sequence from the bootstrap stage or the result of a previous incremental update. The algorithm begins by initializing the input state by applying E^- and E^+, and storing the result in $\mathcal{N}_0^{\#}$ (line 3). Then, the algorithm initializes the sets I_0^- and I_0^+ to be the updates in iteration 0.

In the fixpoint loop, the rule evaluation on line 10 is the core part of this algorithm. This step starts with the multiset of tuples from the previous epoch and applies deletions and insertions resulting from applying Datalog rules with the insertions and deletions for the current epoch. This step is split into three terms: the deletion term, the insertion term, and the re-discovery term. The deletion term, $\Pi^{\#}[I_{k-1}^o \mid N_{k-1}^o \mid I_{k-1}^-] \setminus I_{k-1}^o$, computes tuples that are

---

Algorithm 3 Update($(E, \mathcal{N}^{\#o}), (E^-, E^+)$)

Ensure: $E^- \subseteq E, E \cap E^+ = \emptyset$

1: $\mathcal{N}_0^\# \leftarrow E \setminus E^- \cup E^+$
2: $N_{k-1} \leftarrow \text{Supp}(\mathcal{N}_{k-1}^\#)$
3: $N_{k-1}^o \leftarrow \text{Supp}(\mathcal{N}_{k-1}^{\#o})$
4: $I_0^- \leftarrow E^-$
5: $I_0^+ \leftarrow E^+$
6: **for all** $k \in \{1, 2, \dots\}$ **do**
7: $\quad I_{k-1} \leftarrow \cup_{0 \le i \le k-1} N_i$
8: $\quad I_{k-1}^o \leftarrow \cup_{0 \le 1 \le k-1} N_i^o$
9: $\quad \triangleright$ Here, $\mathcal{A} \setminus B$ denotes $\{(t^v) \in \mathcal{A} \mid t \notin B\}$
10: $\quad \mathcal{N}_k^\# \leftarrow \mathcal{N}_k^{\#o} \ominus (\Pi^\#[I_{k-1}^o \mid N_{k-1}^o \mid I_{k-1}^-] \setminus I_{k-1}^o)$
$\qquad\qquad \oplus (\Pi^\#[I_{k-1} \mid N_{k-1} \mid I_{k-1}^+] \setminus I_{k-1})$
$\qquad\qquad \oplus (I_{k-1}^- \cap \Pi^\#[I_{k-1}^o \cap I_{k-1} \mid N_{k-1}] \setminus I_{k-1})$
11: $\quad \mathcal{N}_k^\# \leftarrow \{(t^v) \in \mathcal{N}_k^\# \mid t \notin I_{k-1}^+\}$
12: $\quad N_k \leftarrow \text{Supp}(\mathcal{N}_k^\#)$
13: $\quad N_k^o \leftarrow \text{Supp}(\mathcal{N}_k^{\#o})$
14: $\quad I_k^- \leftarrow (I_{k-1}^- \setminus N_k) \cup (N_k^o \setminus I_k)$
15: $\quad I_k^+ \leftarrow (I_{k-1}^+ \setminus N_k^o) \cup (N_k \setminus I_k^o)$
16: $\quad$ **if** $\text{Supp}(\mathcal{N}_k^\#) = \emptyset$ **then**
17: $\qquad$ **return** $(E \setminus E^- \cup E^+, \mathcal{N}^\#)$

---

deleted in the current iteration as a result of a derivation where the body contains both a tuple in the delta (N_{k-1}^o) and a deleted tuple (I_{k-1}^-). The set minus notation excludes tuples that were in earlier iterations in the previous epoch, preventing over-deletion since the tuples would not be present in the current iteration due to sparsification. The insertion term, $\Pi^\#[I_{k-1} \mid N_{k-1} \mid I_{k-1}^+] \setminus I_{k-1}$, computes tuples that are inserted as a result of the body of a derivation containing an inserted tuple. Tuples that already exist in previous iterations (i.e., tuples that are contained in I_{k-1}) are excluded to maintain the sparsification invariant. The re-discovery term, $I_{k-1}^- \cap \Pi[I_{k-1}^o \cap I_{k-1}^n \mid N_{k-1}]$, computes tuples that are deleted in previous iterations I_{k-1}^-, but where an alternative derivation exists in the current iteration. This re-discovery rule applies in the situation where a tuple is deleted from some iteration but can still be derived in a later iteration. In this case, the re-discovery term computes this later derivation.

The sparsification term (line 11) does not perform any rule evaluation but excludes tuples from iteration k that were inserted in an earlier iteration (as a result of a new derivation). These tuples should be deleted to maintain the sparsification invariant that a tuple is only present in a single iteration in any given epoch.

The algorithm continues by updating the I_k^- and I_k^+ sets (lines 14 and 15). Computing I_k^- (line 14) takes the deletion set from the previous iteration I_{k-1}^- and excludes the tuples that are newly computed in the current iteration N_k, along with tuples that are deleted in the current iteration ($N_k^o \setminus I_k^n$). Similarly, computing I_k^+ (line 15) takes the insertion set from the previous iteration and excludes tuples that already existed in the current iteration in the previous epoch (since these tuples already existed, so are not *newly* inserted in the current epoch), along with tuples that are inserted in the current iteration.

The algorithm exits if we have reached a fixpoint and the current iteration is identical to the previous iteration, i.e., if $\mathcal{N}_k^\#$ is empty (checked via emptiness of the set projection, in line 16).

Correctness. To show the correctness of our incremental update algorithm, we must show that it computes the same sequence of multisets as if we had applied Bootstrap to the altered input. In other words, we need to show that given a Datalog program P, an input set E, a deletion set E^- and an insertion set E^+, computing the result directly via Bootstrap($E^b = E \setminus E^- \cup E^+$) is equal to Update(Bootstrap(E), (E^-, E^+)). The central parts of the algorithm computing these results are lines 10 and 11. Before the final correctness proof, we need some intermediate properties of the N_k sets and the I^- and I^+ sets. The following important properties are that the validity properties of the E sets (i.e., that $E^+ \cap E = \emptyset$ and $E^- \subseteq E$) also hold for the I^o, I^-, and I^+ sets during the incremental update algorithm. Similar properties relating I^- and I^+ sets to the current epoch's I sets are also required. The eventual goal is to show that $I_k = I_k^o \setminus I_k^- \cup I_k^+$ for each iteration k, which is an important result for showing the correctness of the rule evaluations.

LEMMA 3.3. *For each iteration k, we have (1) $I_k^- \subseteq I_k^o$ and $I_k^- \cap I_k = \emptyset$, and (2) $I_k^+ \cap I_k^o = \emptyset$ and $I_k^+ \subseteq I_k$.*

To sketch the proof for this property, we perform an induction over the iterations. The base case holds because of the definition of I_0^- and I_0^+. Then, for each subsequent iteration, consider line 14 of Algorithm 3. Here, I_k^- takes the value of $(I_{k-1}^- \setminus N_k) \cup (N_k^o \setminus I_k)$. In the first part of the union, the property holds for I_{k-1}^- by the induction hypothesis. In the second part of the union, N_k^o is a subset of I_k^o by definition. Therefore, $I_k^- \subseteq I_k^o$. By similar arguments on line 15, $I_k^+ \subseteq I_k$. The second part of the property, i.e., that $I_k^- \cap I_k = \emptyset$ can be shown by a similar induction argument, again consider line 14.

As a corollary, we can show that the I^- and I^+ sets are correct.

COROLLARY 3.4. *For each iteration k, we have $I_k = I_k^o \setminus I_k^- \cup I_k^+$.*

It remains to be shown that Update is correct. Our criteria for correctness is that it computes the same sequence of multisets as if we had applied the bootstrap algorithm to the updated input, i.e., that the multisets $\mathcal{N}_i^\#$ as computed by Update and Bootstrap are the same for each iteration i. The following is the central theorem for our correctness proof.

THEOREM 3.5. *Given P, E, E^-, and E^+ as above, $\mathcal{N}_i^\#$ as computed by Update(Bootstrap(E), (E^-, E^+)) is equal to $\mathcal{N}_i^\#$ as computed by Bootstrap($E \setminus E^- \cup E^+$) for each iteration i.*

The proof of Theorem 3.5 is by induction over the iterations, and in each step, it considers all four parts of lines 10 and 11. By arguments over which sets each tuple is contained in, and careful consideration of the subset relationships between them, we can show that the counting multisets are the same as those produced by Bootstrap.

Sparsification. Another essential property of our elastic incremental evaluation strategy is the *sparsification invariant*.

LEMMA 3.6 (SPARSIFICATION INVARIANT). *For each iteration k, the sets N_k are disjoint.*

This property ensures that every tuple is only computed in a single iteration, with this iteration being the earliest one in which it is computed.

Re-discovery rules as a notion of provenance. The re-discovery term in the rule evaluation part of Algorithm 3 (the last term in line 10) is critical for maintaining the sparsification property of our algorithm. In particular, a tuple may be deleted in some iteration but still be derivable in a later iteration via a different rule or different body tuples. The re-discovery term allows the algorithm to recover these tuples in the later iteration.

The re-discovery is performed by the rule evaluation term $I_{k-1}^- \cap \Pi[I_{k-1}^o \cap I_{k-1}^n \mid N_{k-1}]$, which states that we compute tuples that were deleted in an earlier iteration (i.e., exist in I_{k-1}^-), but an alternative derivation exists in the current iteration ($\Pi[I_{k-1}^o \cap I_{k-1}^n \mid N_{k-1}]$).

Provenance can be defined as "discovering the derivations for a tuple". Similarly, the re-discovery term discovers derivations in the current iteration for tuples that were deleted in earlier iterations. In particular, for each tuple deleted in an earlier iteration, the re-discovery term finds derivations from $I_{k-1}^o \cap I_{k-1}^n$. Since this process resembles provenance, we adapt the backward rule evaluation techniques from [46] to compute the re-discovery rules.

3.3 Stratified Negation and Constraints

Our algorithms thus far have omitted any notion of negation or constraints. However, both negation and constraints are powerful and common extensions of Datalog. Constraints are a simpler case than negation, and may take the form of arithmetic constraints such as A < B or A != B where A and B are *grounded* variables (i.e., variables also occurring in a positive body predicate) or constants. In an instantiated rule, a constraint is satisfied if the instantiated arithmetic constraint is satisfied. For example,

```
alias(Var1,Var2) :- vpt(Var1,Obj), vpt(Var2,Obj),
                    Var1 != Var2
```

is a rule with arithmetic constraints, and an instantiation of the rule only derives a tuple if the inequality constraint is satisfied by the values given to Var1 and Var2.

Negation is more complicated than simple arithmetic constraints. Syntactically, negations are denoted as a negated predicate with the ! symbol. For example, the rule

```
path(X,Z) :- edge(X,Y), path(Y,Z), !edge(X,Z)
```

computes all the paths in a graph that are not direct edges. A negated predicate must contain only grounded variables or constants, and a negated predicate is satisfied if and only if the corresponding tuple (resulting from an instantiation) is not computable. In this work, we target the standard semantics for negation in Datalog: *stratified negation*. In this semantics, recursive negation is not permitted, and any negated predicates must be of a relation from either input or a previous stratum. With this semantics, a negated predicate is similar to a constraint, where a simple check of the input for a stratum is needed to determine whether the negation is satisfied. However, the truth value of a negation may change due to tuples being inserted or deleted from the negated relation, unlike constraints which do not change truth value after an incremental update. To adapt our Datalog evaluation algorithms to support stratified negation and constraints, the rule evaluation is extended to support these features.

The rule evaluation operator, $\Pi^\#$ is extended so that

$$\Pi_P^\#[I \mid I_{\text{in}}] =$$

$$\left\{ t^v \,\middle|\, \begin{array}{l} v = \#\text{instantiations } t \,\text{:-}\, t_1, \ldots, t_n, !t_{n+1}, \ldots, !t_{n+m}, \psi \\ \text{in } P \text{ where } \{t_1, \ldots, t_n\} \subseteq I, \{t_1, \ldots, t_n\} \cap I_{\text{in}} \neq \emptyset, \\ \{t_{n+1} \ldots t_{n+m}\} \cap I = \emptyset, \text{ and } \psi \text{ is satisfied} \end{array} \right\}$$

where ψ denotes the instantiated arithmetic constraints occurring in the rule. Replacing the rule evaluation operator in Bootstrap (Algorithm 2) with this extended version allows the algorithm to support stratified negation and constraints. However, the extension is more involved for Update since introducing negation also introduces new cases for deleting/inserting tuples. For example, consider the rule

```
path(X,Z) :- edge(X,Y), path(Y,Z), !edge(X,Z).
```

If we have a rule instantiation

```
path(a,c) :- edge(a,b), path(b,c), !edge(a,c)
```

where edge(a,c) is inserted as a result of an incremental update, then the head tuple path(a,c) must be deleted since the negation is no longer satisfied. The opposite situation may arise where the deletion of a tuple may lead to the consequent insertion of a tuple. Therefore, we further extend the rule evaluation operator so that

$$\Pi_P^\#[I \mid I_1 \mid I_2, I_2'] =$$

$$\left\{ t^v \,\middle|\, \begin{array}{l} v = \#\text{instantiations } t \,\text{:-}\, t_1, \ldots, t_n, !t_{n+1}, \ldots, !t_{n+m}, \psi \\ \text{in } P \text{ where } \{t_1, \ldots, t_n\} \subseteq I, \{t_1, \ldots, t_n\} \cap I_1 \neq \emptyset, \\ \{t_{n+1} \ldots t_{n+m}\} \cap I = \emptyset, \psi \text{ is satisfied, and} \\ (\{t_1, \ldots, t_n\} \cap I_2 \neq \emptyset \text{ or } \{t_{n+1} \ldots, t_{n+m}\} \cap I_2' \neq \emptyset) \end{array} \right\}$$

With this rule evaluation operator, a new tuple is derived if the rule instantiation contains body tuples from I, where at least one positive body tuple is also in I_1, and either there is a positive body tuple in I_2 or a negative body tuple in I_2'. Using this notation, the rule evaluation step of Algorithm 3 (line 10) becomes

$$\mathcal{N}_k^\# \leftarrow \mathcal{N}_k^\# \ominus (\Pi^\#[I_{k-1}^o \mid N_{k-1}^o \mid I_{k-1}^-, I_0^+] \setminus I_{k-1}^o)$$

$$\oplus (\Pi^\#[I_{k-1} \mid N_{k-1} \mid I_{k-1}^+, I_0^-] \setminus I_{k-1})$$

$$\oplus (I_{k-1}^- \cap \Pi[I_{k-1}^o \cap I_{k-1} \mid N_{k-1}])$$

where the first and second terms now handle stratified negation. The deletion term, $\Pi^\#[I_{k-1}^o \mid N_{k-1}^o \mid I_{k-1}^-, I_0^+] \setminus I_{k-1}^o$, now computes tuples that are deleted, either as a result of a deleted positive body tuple (I_{k-1}^-) or an inserted negated body tuple (I_0^+). We use iteration 0 for the negated tuples since stratified negation enforces that negations must be from the input of the current stratum. Similarly, the insertion term, $\Pi^\#[I_{k-1}^n \mid N_{k-1} \mid I_{k-1}^+, I_0^-] \setminus I_{k-1}^n$, computes tuples that are inserted either as a result of an inserted positive body tuple or a deleted negative body tuple. The other parts of the algorithms involve manipulating and merging relations and are independent of the Datalog rules. Therefore, no changes are needed to support negation or constraints. Hence, with the extensions to the rule evaluation presented above, our algorithms fully support Datalog with stratified negation and constraints.

4 INTEGRATION INTO SOUFFLÉ

In this section we outline how are approach is integrated in the Souflé Datalog engine, including several optimizations for incremental evaluation.

4.1 Core Implementation

Specialized data structures. Soufflé internally uses highly specialized, parallel B-tree data structures to store relations. For incremental evaluation, we associate each tuple with an iteration number and a count. Therefore, we must extend the internal data structures to allow for these auxiliary attributes. Importantly, these auxiliary attributes may be *updated*, e.g., if a new derivation is discovered, the count must be incremented. Thus, we implemented an update mechanism, and adapted the existing optimistic locking mechanism to support parallel operation.

Auxiliary relations. Auxiliary relations are necessary to represent the tuples that are inserted or deleted in Algorithm 3. These auxiliary relations are represented by separate instantiations of the original relations, with prefixes `diff_plus` and `diff_minus`, respectively. These `diff_plus` and `diff_minus` relations are not exact analogs of I^+ and I^-, since `diff_plus` and `diff_minus` may contain tuples where the derivation count is incremented/decremented, rather than only tuples that are fully inserted/deleted.

Rule evaluation. We extend the operations used in standard rule evaluation algorithms in Soufflé to support the extra operations required by the incremental evaluation algorithms. Soufflé uses nested loop joins for evaluating rules, which incorporate extra conditions and existence checks to ensure correctness. For incremental evaluation, further specialized existence checks are required, e.g., a tuple in `diff_minus` may not actually be deleted, and only one of its derivations is deleted. Therefore, we need a specialized existence check that uses its count in the full relation to determine if the tuple is fully deleted or not. The Datalog rules are then instrumented for incremental evaluation using these extra operations and auxiliary relations. Moreover, separate versions of rule instrumentation are required for the Bootstrap and Update algorithms.

Other operations. Other operations, such as merges between iterations and a cleanup operation between epochs, are also required, along with the rule evaluation extensions. In standard semi-naïve evaluation, at the end of each iteration, new tuples computed in that iteration are merged into the full relation, and this also becomes the delta for the following iteration. For incremental evaluation, further operations may take place, e.g., eager computation of the delta of the previous epoch, and eager computation of `diff_plus` and `diff_minus`. In between epochs, the incremental evaluation algorithms also require a cleanup stage, where the `diff_plus` and `diff_minus` relations are merged into the full relations to update the state in preparation for the following epoch.

4.2 Optimizations

Eager vs. lazy `diff_plus` *and* `diff_minus`. The `diff_plus` and `diff_minus` relations store tuples that are inserted and deleted in the current epoch, respectively. However, there is extra computation involved with the `diff_plus` and `diff_minus` relations in lines 14 and 15 of Algorithm 3. Here, a tuple in `diff_plus` may not actually be newly inserted - it may be a new derivation for a tuple that already existed. Similarly, a tuple in `diff_minus` may not actually be deleted - an alternative derivation may still hold. Thus, we need to check the full relation to determine if a tuple in `diff_plus` or `diff_minus` is actually inserted or deleted, respectively. This

check may be performed eagerly during the merge step in each iteration, with results stored in separate relations `actual_diff_plus` and `actual_diff_minus`, or lazily inside the rule evaluation. For the sake of clarity, our algorithms are presented with eager diff computations, which can be seen in lines 14 and 15. A lazy diff version would incorporate this computation directly in the rule evaluation. This design decision is a trade-off: eagerly computing `diff_plus` and `diff_minus` may result in wasted computation for tuples that are not considered in any rules, while lazy computation may mean the same check of the full relation is performed multiple times for a single tuple, if it occurs in multiple rule derivations. However, our experiments indicate that this trade-off generally favors eager diffs, where it can amortize the checks for tuples that occur in multiple rule derivations. For our benchmarks, the difference is generally within 15% in favor of eager diffs, but it can provide up to 4× speed up in some situations where tuples are frequently repeated in multiple rule derivations.

Filtering for re-discovery rules. The elastic algorithm includes the notion of *re-discovery*, which is required due to its sparsification. In the re-discovery rules, the algorithm finds all tuples which have been deleted in an earlier iteration but where an alternative derivation still exists for the current iteration. Naïvely, this could be done by instrumenting a rule to filter on `diff_minus`:

$$\text{R :- } \text{diff\_minus\_R}, \text{R}_1, \dots, \text{R}_k.$$

However, in some cases this can cause a problematic join, if there are few variables in common between the `diff_minus_R` atom and the remaining atoms. For example,

$$\text{R}(x, y, z) \text{ :- } \text{diff\_minus\_R}(x, y, z), \text{R}_1(x, a), \text{R}_2(y, a), \text{R}_3(z, a).$$

may cause duplication of work in $R_1(x, a)$ if there are many tuples in `diff_minus_R` with the same x value. Our solution is to divide the `diff_minus` relation so that it never causes extra work.

$$\text{R}(x, y, z) \text{ :- } \text{diff\_minus\_R}_x(x), \text{R}_1(x, a), \text{R}_2(y, a),$$
$$\text{diff\_minus\_R}_y(y), \text{R}_3(z, a), \text{diff\_minus\_R}_z(z).$$

Dividing the `diff_minus` relations ensures that each variable only acts as a filter and cannot multiply the work of the other atoms in the rule. Here, the x variable is scheduled first since we assume that `diff_minus_R`$_x(x)$ is smaller than $R_1(x, a)$. However, the other variables must be scheduled after their corresponding atom to prevent a cross-product with the previous atom.

This strategy of considering the variables in the filtering atom is inspired by worst-case optimal joins [32, 45]. These worst-case optimal join algorithms work by considering the *variables* in the atoms in some order, in contrast to traditional nested loop join algorithms that consider an atom order. For our re-discovery rules, this variable-based approach is used only for the filtering atom since the filtering atom is often performance-critical. Our benchmarks show that this technique is generally 2.5× faster than the naïve strategy, while in some situations it can be up to 15× faster.

Scheduling. Scheduling for join orders plays a vital role in the performance of Datalog rules [21, 38, 39]. With incremental evaluation, the assumption that the diffs are smaller than the full relations allows for better heuristics for automatic scheduling. Using this assumption, scheduling `diff_plus` or `diff_minus` first in a rule

evaluation generally improves performance by restricting the size of the search as early as possible. However, care must be taken to avoid cross-products. For example, consider the following rule:

$$R(a, d) \mathrel{:\!-} R_1(a, b), R_2(b, c), \texttt{diff\_minus\_R}_3(c, d).$$

In this case, moving $\texttt{diff\_minus\_R}_3(c, d)$ to the front of the rule would create a cross-product with $R_1(a, b)$ and may lead to worse performance than the original schedule. Hence, using simple automatic scheduling techniques, such as maximizing the number of bound variables in each atom, is crucial to maintain the performance of incremental evaluation.

5 EXPERIMENTAL EVALUATION

This experimental section aims to demonstrate the following claims:

- **Claim I:** Inviability of single incremental evaluations on variable update use cases.
- **Claim II:** The elastic incremental evaluation with a simple switch heuristic performs better compared to existing single strategy incremental evaluations, both in terms of runtime and memory usage, over a series of varying sized incremental updates.

Experimental Setup. Our experiments are run on an AMD Threadripper 2990WX machine with 128 GB memory, running Ubuntu 20.10 with GCC 10.2 used to generate all Soufflé executables. All experiments are run with 8 threads, and all I/O time is excluded from measurements.

We evaluate three versions of Soufflé: (1) Soufflé: Non-Incremental Soufflé engine. (2) Soufflé-counting: A baseline counting incremental algorithm, similar to DDLog, implemented and optimized for Soufflé. (3) Soufflé-elastic: The implementation of the technique presented in this paper. When necessary, we differentiate between elastic-update and elastic-bootstrap algorithms.

We also compare our approach to an industrial-strength incremental Datalog engine, Differential Datalog (DDLog) [36], which uses Differential Dataflow [26] as a backend. DDLog with Differential Dataflow is a state-of-the-art incremental engine that uses a variant of the counting algorithm.

We perform our evaluations using a set of dynamic Datalog use cases adapted by Frank McSherry [1] for benchmarking incremental Datalog engines. The use cases are described below:

(1) Doop [8]: a points-to program analysis framework for Java programs. This is a subset of the Doop program analysis library ported to DDLog. This use case contains a large number of rules and relations with complex recursion.

(2) CRDT: an implementation of a conflict-free replicated data type in Datalog. This use case represents an in-between ruleset with a medium number of rules, relations of moderate complexity, and arithmetic constraints.

(3) Galen [34]: a medical ontology inference task implemented in Datalog. This use case represents a typical ontological use case consisting of a small number of rules and relations with a simple recursive structure. However, the joins can be challenging in Galen.

---

[1]https://github.com/frankmcsherry/dynamic-datalog

Some basic statistics for the benchmarks are included in Table 1. To evaluate the performance of incremental evaluation algorithms, we generated update sets of varying sizes for each benchmark by randomly choosing a subset of EDB tuples that are incrementally deleted and inserted.

Table 1: Benchmark Statistics

Benchmark	Number of rules	EDB size	IDB size
Doop	90	11,014,960	41,665,029
CRDT	31	259,778	2,668,247
Galen	6	976,552	24,483,561

5.1 Single Strategy Incremental Evaluation

In this set of experiments we only consider single strategy evaluations, that is, we only include our Update (elastic-update) evaluation and thus **do not switch to Bootstrap**. In these experiments we do not establish the supremacy of any one technique. Rather, we show that single strategies are not viable compared to non-incremental evaluation. The results for the runtime of incremental updates for each evaluation implementation are shown in Fig. 5, while the impact sizes are in Table 2. These results are computed for one *cycle* of an update set. An update set is a randomly selected subset of EDB tuples. A cycle consists of one epoch where the update set is deleted, followed by one epoch where the update set is inserted. The horizontal line on each benchmark represents the runtime if non-incremental Soufflé performs the same task, i.e., running the whole benchmark twice from scratch. For each benchmark, there is a general trend that larger updates require more runtime. However, this performance is highly unpredictable, even if the size of the incremental update is constant.

Consider the performance of incremental evaluation for Doop in Fig. 5a. Here, there are five separate small update sets, which are each generated by randomly choosing 10 EDB tuples and running one cycle. These small updates all finished within two seconds, which is vastly faster than non-incremental Soufflé. For these small incremental updates all evaluations were very fast on average due to their very low impact, only affecting up to 25 of the IDB tuples. DDlog and elastic-update performed well and on par. Our general observation is that incremental evaluation is highly effective for these lightweight updates. For the 100 update size, the smallest impact was 88 IDB tuples, and the largest impact was 53,816 IDB tuples. As anticipated, this increased the variability of the results. Elastic-update exhibited large extremities, finishing within 5 seconds for the fastest, while more than 5,000 seconds for two update sets. DDLog also had high variance, with the fastest runtime being 5 seconds and the slowest being 213 seconds, well over the non-incremental engine time. Curiously, the fastest incremental update was also one of the higher impact ones, affecting 22,347 IDB tuples, while the slowest affected 140 IDB tuples, indicating that neither the size of the EDB updates nor the size of the impact is always helpful in predicting the runtime of the incremental update. While Soufflé-counting was generally faster than DDLog, it still exhibited a large variance, with runtimes ranging between 1.4 and 18 seconds. For the larger update sets, containing 400, 700, and 1000 tuples respectively, all evaluation strategies failed to compete with non-incremental Soufflé. For example, elastic-update was unable to

(a) Doop **(b) CRDT** **(c) Galen**

Figure 5: Incremental update size vs. runtime. The horizontal line in each figure is the runtime of non-incremental Soufflé on the respective benchmark, and the upwards arrows indicate timeouts.

Benchmark	Size and impacts of update				
Doop	10	100	400	700	1000
	7 - 25	88 - 53,816	26,701 - 5,257,937	49,175 - 3,589,404	87,424 - 6,559,713
CRDT	10	40	70	100	
	3,444 - 35,130	41,217 - 61,039	63,946 - 80,939	85,726 - 91,384	
Galen	10	10000	40000	70000	100000
	810 - 3,708	4,318,658 - 7,038,703	21,068,962 - 31,589,252	35,234,152 - 42,663,368	43,069,725 - 53,765,683

Table 2: The minimum and maximum impact for updates of each size; the impact is the overall number of IDB tuples inserted or deleted

complete any of the update sets within the time limit. These timed-out updates contained tuples deep in a complex recursive structure, indicating that the elastic-update algorithm does not handle these large impact updates well. Likewise, the counting algorithms implemented in both DDLog and in Soufflé exhibited generally poor performance compared to non-incremental Soufflé. Furthermore, these larger updates exhibited even greater variability, particularly for Soufflé-counting.

The results for CRDT, in Fig. 5b tell a similar story. Here, even small updates consisting of 10 EDB tuples exhibit unpredictable and poor performance. In comparison to Doop, the small updates for CRDT have a much larger impact, affecting between 3,444 and 35,130 IDB tuples. However, even this larger impact is around 1% of the IDB, and even with these overall small impacts, the runtime of incremental update is considerably slower than re-running the computation from scratch in Soufflé. Similar to Doop, the performance for larger updates only gets worse. For updates containing 40 EDB tuples, the runtimes varied between 9 and 13 seconds. While this variation is smaller than for Doop, the result still indicates that the performance of incremental evaluation is unpredictable. For larger updates containing 70 and 100 EDB tuples, DDLog was around 5× slower than non-incremental Soufflé, despite the update being only around 0.04% of the EDB and impacting only up to 3.4% of the IDB tuples. Update and Soufflé-counting were both more performant, however still slower than non-incremental Soufflé. The poor performance of incremental update algorithms may be due to the structure of the Datalog rules in CRDT. The rules contain several arithmetic inequality constraints, which cannot be indexed, and are checked after the corresponding value is known in the join. Therefore, incremental strategies that use indices to limit the

computation to updated tuples are ineffective in the presence of performance-critical inequalities. It is also interesting to note that the impact on the IDB tuples was much more consistent for CRDT when compared with Doop. For example, with updates containing 100 EDB tuples, the impact on IDB tuples ranged between 85,726 and 91,384 tuples. This may be due to the much simpler structure of the CRDT application, which contains a larger pre-processing stage followed by a very small recursive stratum.

On the other hand, Galen performed far better with DDLog for incremental evaluation. One reason for this is that Galen has a simple ruleset consisting of only 6 Datalog rules but with challenging join characteristics. DDLog is better optimized for these joins, and can outperform Soufflé for these incremental workloads. For small updates consisting of 10 EDB tuples, an incremental update takes between 0.1 and 0.2 seconds, providing far superior performance compared to a non-incremental engine. Even for medium sized updates consisting of 10,000 EDB tuples, DDLog's incremental update performance is generally faster than non-incremental Soufflé. Only when we consider larger updates of 40,000, 70,000, and 100,000 EDB tuples, or 4%, 7%, and 10% respectively, does the performance of incremental evaluation slow down considerably compared to non-incremental Soufflé. The impact of these larger updates on the IDB is up to 53M tuples, which is almost double the original IDB size. This impact indicates that not only are most of the IDB tuples affected but they are even affected in multiple iterations. Given this large impact, it is no surprise that the runtime for such an incremental update is slower than simply recomputing the result from scratch. Overall, DDLog performs well on this benchmark compared to non-incremental Soufflé. For Galen, the Soufflé incremental strategies do not perform as well. Soufflé-counting is

generally an order of magnitude slower for updates than DDLog, due to unfavorable join orderings. Elastic-update fares even worse, timing out for the larger updates above 10,000 EDB tuples. This is a result of these updates impacting tuples across multiple iterations, which the sparsification of the elastic strategy does not handle well.

These results indicate that state-of-the-art single strategy incremental evaluation algorithms perform well on small impact updates. However, they may be outperformed by a standard non-incremental Datalog engine for more complex applications or high-impact changes. Overall, **we demonstrate Claim I by highlighting the unpredictability and tendency for degraded performance of single strategy evaluations on large impact updates compared to non-incremental Soufflé.**

5.2 Elastic Incremental Evaluation

In this section, we evaluate the performance of our elastic incremental evaluation strategy. That is, **we evaluate the combination of the Update algorithm with Bootstrap.** We use an empirically determined switching parameter of 20% to determine when to use the Update and when to switch to Bootstrap. That is, if the update time is more than **20**% of the previous bootstrap time, we restart using Bootstrap.

For this experiment, we use example workloads for incremental evaluation, which consists of 13 epochs. The first epoch is the initial evaluation, then the following 6 epochs are small updates (containing 10 tuples), with alternating deletion and insertions. These are followed by one large update (1,000 for Doop, 100 for CRDT, and 100,000 for Galen) in epoch 7, then followed by another 4 small updates, with a large update as the final epoch. We note that these patterns may appear in all three of these benchmarks. For Doop, there is a common pattern of software updates consisting of a large refactor, followed by several smaller commits addressing minor comments. For CRDT, an application commonly used for collaborative online text editing, a large update may result from a large portion of text being moved around, while a smaller update may result from smaller additions or deletions from the text. For Galen, a medical ontology application associated with patient diagnosis, a large update may result from a medical test result being updated, while a smaller update may result from a minor symptom change.

For Doop, in Fig. 6a, all of the incremental evaluation strategies are able to effectively incrementalize for the small updates. However, the main differences across the full workload result from the bootstrap strategy, with both the initial evaluation and the large updates being faster or on-par with the state-of-the-art counting strategy. As a result, the elastic incremental strategy can complete this workload in 245 seconds, compared to 284 seconds for Soufflé-counting and 467 seconds for DDLog. In comparison, non-incremental Soufflé, which evaluates each epoch from scratch, achieves 304 seconds for this workload. Thus, this use case demonstrates that **an elastic incremental evaluation is effective for the complex Doop benchmark. Overall, we demonstrate an amortized net gain compared to non-incremental Soufflé as well as single strategy evaluations.**

For CRDT, in Fig. 6b, none of the incremental evaluation strategies are effective for reasons illustrated in Section 5.1, even for the small updates. Here, the elastic strategy hits the 20% heuristic

threshold for all updates, despite Update strategy actually being slightly faster than Bootstrap if it were allowed to run to completion. For this workload, non-incremental Soufflé completes all epochs in 19 seconds, followed by 25 seconds for the Soufflé-counting, 31 seconds for Soufflé-elastic, and 56 seconds for DDLog. **For this particular application, we conclude that incremental evaluation is ineffective in general.**

For Galen, in Fig. 6c, the incremental evaluation strategies were able to perform reasonably well. For epochs 1 and 5, the elastic update strategy reached the 20% heuristic threshold, thus triggering a bootstrap. If this threshold were not in place, the elastic update would have been faster for these small updates. Despite this, Soufflé-elastic is still highly competitive compared to the other incremental evaluation strategies, being able to finish the workload in 384 seconds, compared to 445 seconds for DDLog. Soufflé-counting was ineffective for the large updates for Galen and times out overall. In comparison, non-incremental Soufflé required 370 seconds for this workload. **The results demonstrate that our elastic evaluation is competitive for the Galen use case.**

Overall, **the experimental evaluation has validated Claim II by showing a performance improvement compared to single strategy approaches.** The limited overhead of our Bootstrap evaluation makes up for any cost induced by the Update evaluation. We believe with improved heuristics and tuning, this improvement can be further maximized.

Table 3: Memory usage for each engine, showing the minimum, average, and maximum memory usage across all of the update sets

Bench.	Engine	Min (MB)	Avg (MB)	Max (MB)
Doop	Soufflé	1,759	1,762	1,764
	Soufflé-elastic	7,473	7,492	7,505
	Soufflé-counting	9,106	9,449	11,387
	DDLog	17,381	23,352	27,851
CRDT	Soufflé	42	42	42
	Soufflé-elastic	335	346	352
	Soufflé-counting	328	337	344
	DDLog	786	829	858
Galen	Soufflé	901	931	960
	Soufflé-elastic	5,641	5,672	5,698
	Soufflé-counting	14,588	17,974	21,034
	DDLog	15,333	20,862	26,461

Along with runtime, another aspect of performance is memory usage. For example, in large program analysis use cases memory has been shown to be a limiting factor [23]. Table 3 shows the minimum, average, and maximum memory usage across all the update sets for each benchmark. These results show that non-incremental Soufflé uses the least memory by far since it does not need to keep the extra state that incremental evaluation requires. Among the incremental engines, Soufflé-elastic performs best since it only keeps the counts for one iteration for each tuple. On the other hand, the counting algorithm, both in Soufflé and in DDLog, requires keeping the count of each tuple for *every* iteration it is generated in, thus using extra memory to maintain this additional state.

6 RELATED WORK

There is a large corpus of incremental algorithms in related fields, including Databases [5], Logic-programming [37], Compilers [35],

(a) Doop

(b) CRDT

(c) Galen

Figure 6: Runtimes for an elastic workload. For each benchmark, the first epoch is an initial evaluation, followed by 6 epochs of small updates, then one large update, then 4 epochs of small updates, then one large update.

Model-checking [40] and SAT solving [25]. In this section, we focus exclusively on Datalog evaluation. The main body of work in incremental Datalog evaluation is related to the Delete-Rederive (DRed) algorithm [16]. The main weakness of this approach concerns *over-deletion*. This is resolved by re-deriving tuples that are over-deleted. The Counting algorithm presented in [16] is applicable only for non-recursive Datalog programs. For this approach, each tuple is associated with a *count* of the number of different derivations for that tuple. When removing or inserting a new tuple, that count is decremented or incremented respectively, and a tuple may be removed if the count reaches 0. However, with recursive Datalog programs, deleting tuples may cause the recursive decrement of the count, thus again leading to over-deletion. More recent developments include the Backward/Forward algorithm [29] and DRed^c [18], which are both optimizations of the DRed algorithm. These approaches aim to reduce the approximation induced by the over-deletion step. Backward/Forward uses a form of *backward evaluation* to eagerly check if over-deleted tuples still have a proof from the remaining input, while DRed^c maintains separate recursive and non-recursive counters to track the number of derivations of each tuple. While these approaches indeed reduce the over-deletion of DRed, they are still approximations and worst-case scenarios may exhibit large runtime overheads. Techniques like [15] use provenance information in the form of Boolean formulae for each tuple to determine if a deleted tuple has proof support. Systems such as [7] are a formally certified implementation of algorithms inspired by DRed for a limited subset of non-recursive Datalog.

The Differential Dataflow (DDF) system [26] implements incremental evaluation for Dataflow programming. The approach is similar to the counting algorithm, with each tuple being associated with a count for the number of derivations for that tuple. However, DDF permits recursive programs by storing a count *per iteration* of recursive evaluation. Its advantage is that it computes a precise result for an incremental update. However, the setting of Dataflow programming is different from Datalog, and more similar to stream programming, where programs tend to be less complex with smaller updates. Differential Datalog [36] is a Datalog engine built on top of DDF. Other systems, such as RDFox [18, 19, 28] and IncA/Ladder [43, 44], implement variations of DRed and DDF algorithms,

specialized to their respective domains. In comparison with existing approaches, our elastic evaluation is unique in that it has two evaluation phases, recognizing the importance of specializing the Bootstrap phase to initialize the computation. Our algorithms form a sparsified variation of the counting algorithm, allowing the efficient Bootstrap phase, and lowering the space overhead per tuple.

In addition to incremental evaluation, multiset semantics for Datalog have also been employed for other uses. [30] uses multiset semantics to enable magic set transformation for group-by aggregates. Meanwhile, [6] uses multiset semantics to better model replicated data that exist in many database applications.

7 CONCLUSION

This paper has demonstrated the pitfalls of existing incremental evaluation algorithms for use cases with varying sized updates. We have proposed the use of an elastic approach for incremental evaluation. We switch between a low overhead Bootstrap strategy that targets high impact updates and an Update strategy that targets low impact updates. We propose a simple heuristic for switching between the two strategies. Using this setup, we have shown that the elastic approach is effective in use cases where single strategy incremental evaluation struggles to perform adequately compared to regular Datalog evaluation.

For future work, one potential avenue is to use worst-case optimal join algorithms for Update rules. These rules typically involve large intermediate results with small final output, where worst-case optimal joins may be helpful. Other possibilities could include further investigating the characteristics of incremental evaluation algorithms to better understand when to switch between Bootstrap and Update.

REFERENCES

[1] Serge Abiteboul, Richard Hull, and Victor Vianu. 1995. *Foundations of Databases.* Addison-Wesley Publishing Company.

[2] Nicholas Allen, Bernhard Scholz, and Padmanabhan Krishnan. 2015. *Staged Points-to Analysis for Large Code Bases.* Springer Berlin Heidelberg, 131–150. https://doi.org/10.1007/978-3-662-46663-6_7

[3] Molham Aref, Balder ten Cate, Todd J. Green, Benny Kimelfeld, Dan Olteanu, Emir Pasalic, Todd L. Veldhuizen, and Geoffrey Washburn. 2015. Design and Implementation of the LogicBlox System. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data* (Melbourne, Victoria, Australia) *(SIGMOD '15).* ACM, New York, NY, USA, 1371–1382. https://doi.org/10.1145/2723372.2742796

[4] John Backes, Sam Bayless, Byron Cook, Catherine Dodge, Andrew Gacek, Alan J. Hu, Temesghen Kahsai, Bill Kocik, Evgenii Kotelnikov, Jure Kukovec, Sean McLaughlin, Jason Reed, Neha Rungta, John Sizemore, Mark A. Stalzer, Preethi Srinivasan, Pavle Subotic, Carsten Varming, and Blake Whaley. 2019. Reachability Analysis for AWS-Based Networks. In *Computer Aided Verification - 31st International Conference, CAV 2019, New York City, NY, USA, July 15-18, 2019, Proceedings, Part II.* 231–241.

[5] Andrey Balmin, Yannis Papakonstantinou, and Victor Vianu. 2004. Incremental Validation of XML Documents. *ACM Trans. Database Syst.* 29, 4 (Dec. 2004), 710–751. https://doi.org/10.1145/1042046.1042050

[6] Leopoldo Bertossi, Georg Gottlob, and Reinhard Pichler. 2018. Datalog: Bag semantics via set semantics. *arXiv preprint arXiv:1803.06445* (2018).

[7] Angela Bonifati, Stefania Dumbrava, and Emilio Jesús Gallego Arias. 2018. Certified graph view maintenance with regular datalog. *Theory and Practice of Logic Programming* 18, 3-4 (2018), 372–389.

[8] Martin Bravenboer and Yannis Smaragdakis. 2009. Strictly Declarative Specification of Sophisticated Points-to Analyses. *SIGPLAN Not.* 44, 10 (2009), 243–262. https://doi.org/10.1145/1639949.1640108

[9] James Cheney, Laura Chiticariu, and Wang-Chiew Tan. 2009. Provenance in Databases: Why, How, and Where. *Foundations and Trends in Databases* 1 (2009), 379–474. https://doi.org/10.1561/1900000006

[10] Nathan Chong, Byron Cook, Konstantinos Kallas, Kareem Khazem, Felipe R. Monteiro, Daniel Schwartz-Narbonne, Serdar Tasiran, Michael Tautschnig, and Mark R. Tuttle. 2020. Code-Level Model Checking in the Software Development Workflow. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering: Software Engineering in Practice* (Seoul, South Korea) *(ICSE-SEIP '20).* Association for Computing Machinery, New York, NY, USA, 11–20. https://doi.org/10.1145/3377813.3381347

[11] Dino Distefano, Manuel Fähndrich, Francesco Logozzo, and Peter W. O'Hearn. 2019. Scaling Static Analyses at Facebook. *Commun. ACM* 62, 8 (July 2019), 62–70.

[12] Neville Grech, Lexi Brent, Bernhard Scholz, and Yannis Smaragdakis. 2019. Gigahorse: Thorough, Declarative Decompilation of Smart Contracts. In *Proceedings of the 41th International Conference on Software Engineering, ICSE 2019.* ACM, Montreal, QC, Canada, (to appear).

[13] Neville Grech, Lexi Brent, Bernhard Scholz, and Yannis Smaragdakis. 2019. Gigahorse: thorough, declarative decompilation of smart contracts. In *Proceedings of the 41st International Conference on Software Engineering, ICSE 2019, Montreal, QC, Canada, May 25-31, 2019,* Joanne M. Atlee, Tevfik Bultan, and Jon Whittle (Eds.). IEEE / ACM, 1176–1186. https://doi.org/10.1109/ICSE.2019.00120

[14] Neville Grech, Michael Kong, Anton Jurisevic, Lexi Brent, Bernhard Scholz, and Yannis Smaragdakis. 2018. MadMax: Surviving Out-of-Gas Conditions in Ethereum Smart Contracts. In *SPLASH 2018 OOPSLA.*

[15] Todd J. Green, Grigoris Karvounarakis, Zachary G. Ives, and Val Tannen. 2007. Update Exchange with Mappings and Provenance. In *In Very Large Data Bases (VLDB.* 675–686.

[16] Ashish Gupta, Inderpal Singh Mumick, and Venkatramanan Siva Subrahmanian. 1993. Maintaining views incrementally. *ACM SIGMOD Record* 22, 2 (1993), 157–166.

[17] Kryštof Hoder, Nikolaj Bjørner, and Leonardo de Moura. 2011. μZ– An Efficient Engine for Fixed Points with Constraints. In *Computer Aided Verification,* Ganesh Gopalakrishnan and Shaz Qadeer (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 457–462.

[18] Pan Hu, Boris Motik, and Ian Horrocks. 2018. Optimised maintenance of datalog materialisations. In *Thirty-Second AAAI Conference on Artificial Intelligence.*

[19] Pan Hu, Boris Motik, and Ian Horrocks. 2019. Modular Materialisation of Datalog Programs. (2019).

[20] Shan Shan Huang, Todd Jeffrey Green, and Boon Thau Loo. 2011. Datalog and Emerging Applications: An Interactive Tutorial. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data* (Athens, Greece) *(SIGMOD '11).* ACM, 1213–1216. https://doi.org/10.1145/1989323.1989456

[21] Muhammad Imran, Gábor E Gévay, and Volker Markl. 2020. Distributed Graph Analytics with Datalog Queries in Flink. In *Software Foundations for Data Interoperability and Large Scale Graph Data Analytics.* Springer, 70–83.

[22] Herbert Jordan, Bernhard Scholz, and Pavle Subotić. 2016. Soufflé: On Synthesis of Program Analyzers. *Proceedings of Computer Aided Verification* 28 (2016), 422–430.

[23] Herbert Jordan, Bernhard Scholz, and Pavle Subotić. 2016. Soufflé: On synthesis of program analyzers. In *International Conference on Computer Aided Verification.* Springer, 422–430.

[24] Grigoris Karvounarakis, Todd J. Green, Zachary G. Ives, and Val Tannen. 2013. Collaborative Data Sharing via Update Exchange and Provenance. *ACM Trans. Database Syst.* 38, 3, Article 19 (Sept. 2013), 42 pages.

[25] Yousef Kilani, Mohammad Bsoul, Ayoub Alsarhan, and Ahmad Al-Khasawneh. 2013. A Survey of the Satisfiability-Problems Solving Algorithms. *Int. J. Adv. Intell. Paradigms* 5, 3 (Sept. 2013), 233–256. https://doi.org/10.1504/IJAIP.2013.056447

[26] Frank McSherry, Derek Gordon Murray, Rebecca Isaacs, and Michael Isard. 2013. Differential Dataflow.. In *CIDR.*

[27] Hongyuan Mei, Guanghui Qin, Minjie Xu, and Jason Eisner. 2020. Neural Datalog Through Time: Informed Temporal Modeling via Logical Specification. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119),* Hal Daumé III and Aarti Singh (Eds.). PMLR, 6808–6819.

[28] Boris Motik, Yavor Nenov, Robert Piro, and Ian Horrocks. 2019. Maintenance of datalog materialisations revisited. *Artificial Intelligence* 269 (2019), 76–136.

[29] Boris Motik, Yavor Nenov, Robert Edgar Felix Piro, and Ian Horrocks. 2015. Incremental update of datalog materialisation: the backward/forward algorithm. In *Twenty-Ninth AAAI Conference on Artificial Intelligence.*

[30] Inderpal Singh Mumick, Hamid Pirahesh, and Raghu Ramakrishnan. 1990. The magic of duplicates and aggregates. In *VLDB.* Citeseer.

[31] Derek Murray, Frank McSherry, Rebecca Isaacs, Michael Isard, Paul Barham, and Martin Abadi. 2013. Naiad: A Timely Dataflow System. In *Proceedings of the 24th ACM Symposium on Operating Systems Principles (SOSP)* (proceedings of the 24th acm symposium on operating systems principles (sosp) ed.). ACM.

[32] Hung Q Ngo, Ely Porat, Christopher Ré, and Atri Rudra. 2018. Worst-case optimal join algorithms. *Journal of the ACM (JACM)* 65, 3 (2018), 1–40.

[33] Xinming Ou, Sudhakar Govindavajhala, and Andrew W. Appel. 2005. MulVAL: A Logic-based Network Security Analyzer. In *Proceedings of the 14th Conference on USENIX Security Symposium - Volume 14* (Baltimore, MD) *(SSYM'05).* USENIX Association, Berkeley, CA, USA, 8–8. http://dl.acm.org/citation.cfm?id=1251398.1251406

[34] Alan L Rector, Jeremy E Rogers, and Pam Pole. 1996. The GALEN high level ontology. In *Medical Informatics Europe'96.* IOS Press, 174–178.

[35] Thomas Reps. 1982. Optimal-Time Incremental Semantic Analysis for Syntax-Directed Editors. In *Proceedings of the 9th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Albuquerque, New Mexico) *(POPL '82).* Association for Computing Machinery, New York, NY, USA, 169–176. https://doi.org/10.1145/582153.582172

[36] Leonid Ryzhyk and Mihai Budiu. 2019. Differential Datalog. *Datalog* 2 (2019), 4–5.

[37] Diptikalyan Saha and C. R. Ramakrishnan. 2006. Incremental Evaluation of Tabled Prolog: Beyond Pure Logic Programs. In *Practical Aspects of Declarative Languages,* Pascal Van Hentenryck (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 215–229.

[38] Jiwon Seo, Stephen Guo, and Monica S Lam. 2013. Socialite: Datalog extensions for efficient social network analysis. In *2013 IEEE 29th International Conference on Data Engineering (ICDE).* IEEE, 278–289.

[39] Alexander Shkapsky, Mohan Yang, Matteo Interlandi, Hsuan Chiu, Tyson Condie, and Carlo Zaniolo. 2016. Big data analytics with datalog queries on spark. In *Proceedings of the 2016 International Conference on Management of Data.* 1135–1149.

[40] Oleg V. Sokolsky and Scott A. Smolka. 1994. Incremental Model Checking in the Modal Mu-Calculus. In *IN CAV, VOLUME 818 OF LNCS.* Springer-Verlag, 351–363.

[41] Manu Sridharan, Denis Gopan, Lexin Shan, and Rastislav Bodík. 2005. 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 2005, October 16-20, 2005, San Diego, CA, USA,* Ralph E. Johnson and Richard P. Gabriel (Eds.). ACM, 59–76. https://doi.org/10.1145/1094811.1094817

[42] Pavle Subotić. 2020. Concise Explanations in Static Analysis Driven Code Reviews. (2020). https://www.youtube.com/watch?v=FPCZ2TIxrpg&t=8888s Infer Practitioners 2020.

[43] Tamás Szabó, Sebastian Erdweg, and Gábor Bergmann. 2021. Incremental Whole-Program Analysis in Datalog with Lattices. (2021).

[44] Tamás Szabó, Sebastian Erdweg, and Markus Voelter. 2016. Inca: A dsl for the definition of incremental program analyses. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering.* 320–331.

[45] Todd L Veldhuizen. 2012. Leapfrog triejoin: A simple, worst-case optimal join algorithm. *arXiv preprint arXiv:1210.0481* (2012).

[46] David Zhao, Pavle Subotić, and Bernhard Scholz. 2020. Debugging Large-scale Datalog: A Scalable Provenance Evaluation Strategy. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 42, 2 (2020), 1–35.

[47] Wenchao Zhou, Micah Sherr, Tao Tao, Xiaozhou Li, Boon Thau Loo, and Yun Mao. 2010. Efficient Querying and Maintenance of Network Provenance at Internet-Scale. *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data* (2010), 615–626.

A PROOFS FOR THEOREMS

A.1 Proof of Lemma 3.2

PROOF. This proof is by induction over the iterations k. For $k = 1$, the semi-naïve algorithm takes Δ_0, while the bootstrap algorithm takes $\text{Supp}(\mathcal{B}_0^{\#})$. Both of these sets are defined to be E, so are equal. For the induction hypothesis, assume that all I_{i-1} of bootstrap equal I_{i-1} of semi-naïve. Then, $\text{Supp}(\mathcal{B}_i^{\#}) = \Delta_i$ by Lemma 3.1. Therefore, adding $\text{Supp}(\mathcal{B}_i^{\#})$ or Δ_i to the union results in the same set. □

A.2 Proof of Lemma 3.3

PROOF. This proof is by induction over the iterations. For $k = 0$, $I_0^- = E^-$ and $I_0^+ = E^+$ by definition, so properties (1) and (2) hold.

The induction hypothesis is that for iteration $k - 1$, we have $I_{k-1}^- \subseteq I_{k-1}^o$, $I_{k-1}^- \cap I_{k-1} = \emptyset$, $I_{k-1}^+ \cap I_{k-1}^o = \emptyset$, and $I_{k-1}^+ \subseteq I_{k-1}$.

For property (1), we show $I_k^- \subseteq I_k^o$. Consider line 14 of Algorithm 3, where $I_k^- \leftarrow (I_{k-1}^- \setminus N_k) \cup (N_k^o \setminus I_k)$. In the first part of the union, $I_{k-1}^- \subseteq I_{k-1}^o$ by the induction hypothesis. Therefore, also $I_{k-1}^- \subseteq I_k^o$, since I_{k-1}^o grows monotonically. In the second part of the union, $N_k^o \subseteq I_k^o$ by definition of I_k^o. Therefore, $I_k^- \subseteq I_k^o$.

To show that $I_k^- \cap I_k = \emptyset$, consider the same line. In the first part of the union, $I_{k-1}^- \cap I_{k-1} = \emptyset$ by the induction hypothesis. We then exclude N_k, and since $I_k = I_{k-1} \cup N_k$ by definition, then $(I_{k-1}^- \setminus N_k) \cap I_k = \emptyset$. In the second part of the union, we exclude I_k. Therefore, $I_k^- \cap I_k = \emptyset$.

Property (2) holds by similar arguments on line 15. □

A.3 Proof of Lemma 3.4

PROOF. We first show that $I_k^o \setminus I_k = I_k^-$ by showing both directions of inclusion. The reverse direction, i.e., that $I_k^- \subseteq I_k^o \setminus I_k$ is a direct corollary of Lemma 3.3, that $I_k^- \subseteq I_k^o$ and $I_k^- \cap I_k = \emptyset$. For the forward direction, consider some tuple $t \in I_k^o \setminus I_k$. Then, t must be in some $N_i^o \setminus I_k$ for some $i \leq k$. Since $I_i \subseteq I_k$, t is also in $N_i^o \setminus I_i$. Therefore, $t \in I_i^-$. Also, t cannot be removed from I^- in a later iteration, since $t \notin I_k$, and therefore, $t \in I_k^-$.

We have shown both directions of inclusion, and therefore, $I_k^o \setminus I_k = I_k^-$. By a similar argument, $I_k \setminus I_k^o = I_k^+$. From these equalities:

$$I_k^o \setminus I_k^- \cup I_k^+ = I_k^o \setminus (I_k^o \setminus I_k) \cup (I_k \setminus I_k^o)$$
$$= (I_k^o \cap I_k) \cup (I_k \setminus I_k^o) = I_k$$

□

A.4 Proof of Theorem 3.5

PROOF. For this proof, we mainly consider the underlying sets of derivations rather than the counting multisets, since the counting multisets do not distinguish between different derivations. We introduce some new notation to convert between derivations and tuples: $\phi((t :\text{-} t_1, \ldots, t_n)) := t$ takes the head tuple of a derivation.

The proof is an induction over the iterations. The initial step, where $k = 0$, is true since both $\mathcal{B}_0^{\#}$ and $\mathcal{N}_0^{\#}$ take on the value of $E \setminus E^- \cup E^+$, where every tuple has a count of 1.

The induction hypothesis is that for all $0 \leq i < k$, we have $\mathcal{B}_i^{\#} = \mathcal{N}_i^{\#}$. We consider each of the four terms in lines 10 and 11. We first need to show that the sets of derivations computed by these lines are disjoint, so that the algorithm does not double count.

- For the deletion term (we label it (1)), we have the derivations $\{d \in \Pi^D[I_{k-1}^o \mid N_{k-1}^o \mid I_{k-1}^-] \mid \phi(d) \notin I_{k-1}^o\}$.
- For the insertion term (labelled (2)), we have derivations $\{d \in \Pi^D[I_{k-1} \mid N_{k-1} \mid I_{k-1}^+] \mid \phi(d) \notin I_{k-1}\}$. Since $I_{k-1}^+ \cap I_{k-1}^o = \emptyset$ (from Corollary 3.4), then $(2) \cap (1) = \emptyset$, since (1) takes derivations only from I_{k-1}^o.
- For the re-discovery term (labelled (3)), we have derivations $\{d \in \Pi^D[I_{k-1}^o \cap I_{k-1} \mid N_{k-1}] \mid \phi(d) \in I_{k-1}^-$ and $\phi(d) \notin I_{k-1}\}$. Since this takes derivations from I_{k-1}^o, and $I_{k-1}^o \cap I_{k-1}^+ = \emptyset$, then $(3) \cap (2) = \emptyset$. Also, since $I_{k-1}^- \subseteq I_{k-1}^o$ (from Corollary 3.4), we have $(3) \cap (1) = \emptyset$, since (1) excludes tuples from I_{k-1}^o.
- For the sparsification term (labelled (4)), we have $\mathcal{N}_k^{\#} \ominus (\mathcal{N}_k^{\#} \cap I_{k-1}^+)$. However, note that this term is processed after the three other terms. Therefore, it naturally excludes (1), and so $(4) \cap (1) = \emptyset$. Moreover, we have $I_{k-1}^+ \subseteq I_{k-1}$ (from Corollary 3.4), and so $(4) \cap (3) = \emptyset$ and $(4) \cap (2) = \emptyset$, since both (3) and (2) exclude I_{k-1}.

Since all 4 terms produce disjoint derivations, the algorithm does not double count when adding or removing any derivations. Next, we need to prove that for any derivation in $\mathcal{N}_k^{Do}$ and not in $\mathcal{N}_k^D$, it is removed by one of the four terms, and vice versa.

Consider a derivation $d \in \mathcal{N}_k^{Do} \setminus \mathcal{N}_k^D$. By definition, $d \in \{d \in \Pi^D[I_{k-1}^o \mid N_{k-1}^o] \mid \phi(d) \notin I_{k-1}^o\} \setminus \{d \in \Pi^D[I_{k-1} \mid N_{k-1}] \mid \phi(d) \notin I_{k-1}\}$. Then, there are two cases. The first case is that $\phi(d) \in I_{k-1}$. In this case, also $\phi(d) \notin I_{k-1}^o$, by our assumption, and so $\phi(d) \in I_{k-1}^+$ (by Corollary 3.4). Therefore, d would be removed by the sparsification term which removes all tuples that are in I_{k-1}^+. The second case is that $d \notin \Pi^D[I_{k-1} \mid N_{k-1}]$. In this case, one of the body tuples of d is in $I_{k-1}^o \setminus I_{k-1}$ (or in $N_{k-1}^o \setminus N_{k-1}$, which implies also that it is in $I_{k-1}^o \setminus I_{k-1}$), which equals I_{k-1}^- (by Corollary 3.4). Therefore, $d \in \Pi^D[I_{k-1}^o \mid N_{k-1}^o \mid I_{k-1}^-]$, and since $\phi(d) \notin I_{k-1}^o$ by assumption, it would be removed by the deletion term.

Now, for the opposite case, consider a derivation $d \in \mathcal{N}_k^D \setminus \mathcal{N}_k^{Do}$. We want to show that this derivation is inserted by one of the four terms. By definition, $d \in \{d \in (\Pi^D[I_{k-1} \mid N_{k-1}] \mid \phi(d) \notin I_{k-1}\} \setminus \{d \in \Pi^D[I_{k-1}^o \mid N_{k-1}^o] \mid \phi(d) \notin I_{k-1}^o\}$. Like the deletion case, there are two cases. The first is that at least one of the body tuples of d are in $I_{k-1} \setminus I_{k-1}^o$. Then, this tuple is in I_{k-1}^+, and therefore, $d \in \Pi^D[I_{k-1} \mid N_{k-1} \mid I_{k-1}^+]$. Since $\phi(d) \notin I_{k-1}$ by assumption, then d will be inserted by the insertion term. The second case is if $\phi(d) \in I_{k-1}^o$. Then, since $\phi(d) \notin I_{k-1}$, $\phi(d) \in I_{k-1}^o \setminus I_{k-1} = I_{k-1}^-$. If the first case doesn't hold, we know that all of the body tuples are not in $I_{k-1} \setminus I_{k-1}^o$, and therefore, they must all be in I_{k-1}^o. Therefore, $d \in \Pi^D[I_{k-1}^o \cap I_{k-1} \mid N_{k-1}]$. Since $\phi(d) \notin I_{k-1}$ by assumption, d would be inserted by the re-discovery term. □

B ADDITIONAL EXPERIMENTAL DATA

Benchmark	Engine	Updates	Epoch 0 (sec)	Epoch 1 (-) (sec)	Epoch 2 (+) (sec)	Memory (MB)
doop	Soufflé-elastic(update)	10	64.53	min 0.20 max 5.62	min 0.20 max 0.42	7486.5
		100	66.22	min 4.63 max 13624.27	min 0.38 max 2.98	7497.4
		400	-	-	-	-
		700	-		-	-
		1000	-	-	-	-
	Soufflé-counting	10	113.64	min 0.50 max 1.00	min 0.50 max 1.11	9116.2
		100	116.00	min 0.78 max 9.05	min 0.77 max 9.49	9131.4
		400	114.00	min 10.20 max 632.30	min 10.90 max 695.00	9540.6
		700	113.81	min 21.26 max 129.37	min 22.60 max 147.66	9539.8
		1000	116.92	min 14.40 max 428.29	min 16.70 max 500.99	9917.3
	DDLog	10	164.53	min 0.07 max 1.85	min 0.02 max 1.67	17495.4
		100	167.83	min 2.57 max 102.51	min 2.30 max 110.29	20682.4
		400	169.45	min 90.65 max 145.37	min 101.38 max 146.70	25002.1
		700	164.86	min 115.14 max 205.84	min 129.22 max 185.37	26350.2
		1000	167.35	min 149.71 max 232.32	min 148.30 max 205.08	27230.5
crdt	Soufflé-elastic(update)	10	1.99	min 1.64 max 1.74	min 1.45 max 1.51	338.7
		40	1.96	min 1.93 max 2.10	min 1.67 max 1.78	345.7
		70	1.96	min 2.32 max 2.68	min 2.01 max 2.58	349.3
		100	2.00	min 2.51 max 2.90	min 2.05 max 2.65	351.8
	Soufflé-counting	10	2.98	min 1.58 max 1.69	min 1.62 max 1.74	331.4
		40	3.02	min 1.83 max 1.99	min 1.80 max 2.00	336.6
		70	3.10	min 2.15 max 2.32	min 2.11 max 2.34	338.6
		100	2.92	min 2.21 max 2.47	min 2.19 max 2.49	341.6
	DDLog	10	8.52	min 0.71 max 5.97	min 0.69 max 5.74	804.5
		40	8.61	min 4.69 max 6.56	min 4.47 max 6.72	825.4
		70	8.59	min 7.02 max 7.36	min 6.78 max 7.33	833.1
		100	8.50	min 7.29 max 7.62	min 7.08 max 7.28	851.8
galen	Soufflé-elastic(update)	10	60.69	min 2.48 max 20.29	min 0.17 max 0.24	5666.9
		10000	59.64	min 37.16 max 94.22	min 0.29 max 0.62	5676.2
		40000	-	-	-	-
		70000	-	-	-	-
		100000	-	-	-	-
	Soufflé-counting	10	415.63	min 0.40 max 1.44	min 0.53 max 1.64	14595.4
		10000	415.05	min 147.14 max 203.44	min 146.88 max 239.68	15799.4
		40000	422.35	min 568.20 max 902.38	min 612.21 max 977.07	18776.6
		70000	411.63	min 996.40 max 1130.46	min 1025.11 max 1216.50	20027.0
		100000	414.45	min 1131.57 max 1602.21	min 21.03 max 1377.78	20671.6
	DDLog	10	152.09	min 0.09 max 0.20	min 0.06 max 0.09	15602.4
		10000	154.10	min 19.69 max 28.02	min 20.78 max 27.14	16771.1
		40000	152.70	min 74.72 max 96.72	min 77.23 max 99.10	22027.8
		70000	154.10	min 107.76 max 130.03	min 110.19 max 132.05	24195.4
		100000	157.96	min 135.58 max 160.32	min 137.96 max 168.43	25712.8

Table 4: Running times and memory usage for Dynamic Datalog Benchmarks, each min and max value denotes the minimum and maximum runtimes over 5 different datasets of the corresponding update size, - denotes timeout, and variations in Epoch 0 runtime are due to re-runs of the experiment