

Modeling Datalog Assertion and Retraction in Linear Logic*

Edmund S. L. Lam and Iliano Cervesato

June 2012
CMU-CS-12-126
CMU-CS-QTR-113

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Carnegie Mellon University, Qatar campus.
The author can be reached at sllam@qatar.cmu.edu or iliano@cmu.edu.

Abstract

Practical algorithms have been proposed to efficiently recompute the logical consequences of a Datalog program after a new fact has been asserted or retracted. This is essential in a dynamic setting where facts are frequently added and removed. Yet while assertion is logically well understood as incremental inference, the monotonic nature of first-order logic is ill-suited to model retraction. As such, the traditional logical interpretation of Datalog offers at most an abstract specification of Datalog systems, but has tenuous relations to the algorithms that perform efficient assertions and retractions in practical implementations. This report proposes a logical interpretation of Datalog based on linear logic. It not only captures the meaning of Datalog updates, but also provides an operational model that underlies the dynamic changes of the set of inferable facts, all within the confines of logic. We prove the correctness of this interpretation with respect to its traditional counterpart.

* Funded by the Qatar National Research Fund as project NPRP 09-667-1-100 (Effective Programming for Large Distributed Ensembles)

Keywords: Datalog, Linear Logic, Retraction, Assertion, Dynamic updates

Contents

1	Introduction	1
2	Preliminaries	1
2.1	Notations	1
2.2	Datalog	2
2.3	Linear Logic: LV^{obs} and LV^{obs-}	3
3	Example	4
4	Modeling Datalog Assertion	5
4.1	Inference Rules and Interpretation	5
4.2	Quiescence and Complete Inference	7
4.3	Correctness	8
5	Modeling Datalog Retraction	12
5.1	Retraction and Absorption Rules	12
5.2	Quiescence and Complete Retraction	14
5.3	The llD^{Full} Interpretation	14
5.4	Cycles and Re-assertion	15
5.5	Correctness	16
6	Related Work	23
7	Conclusions and Future Work	24

List of Figures

1	Logical Specification of Datalog Updates	1
2	Syntax of Datalog	2
3	Fragment of the Intuitionistic Sequent Calculus, LJ	3
4	LV^{obs} and LV^{obs-} Intuitionistic Linear Logic	4
5	Example of Linear Logic Interpretation	5
6	LV^{obs-} derivations of an example \mathcal{LD}^{Assert} interpretation	6
7	\mathcal{LD}^{Assert} Interpretation of Datalog	6
8	\mathcal{LD}^{Assert} Inference Macro Rule	7
9	An Example of \mathcal{LD}^{Full-} Interpretation with Assertion/Retraction Derivations	8
10	An Example of \mathcal{LD}^{Full} Interpretation	12
11	\mathcal{LD}^{Full} Interpretation of Datalog	13
12	\mathcal{LD}^{Full} Inference Macro Rule	15
13	Assertion and Retraction, in \mathcal{LD}^{Full}	16
14	\mathcal{LD}^{Full} Retraction Macro Rule	17
15	\mathcal{LD}^{Full} Absorption Macro Rule	18
16	Example of Quiescence in the Presence of Cycles	19
17	Example of re-assertion during retraction	20

$$\begin{array}{l} \text{Assertion: } \mathcal{P}(\mathcal{B}) \xrightarrow{+a}_{\mathcal{P}} \mathcal{P}(\mathcal{B}, a) \qquad \text{Retraction: } \mathcal{P}(\mathcal{B}, a) \xrightarrow{-a}_{\mathcal{P}} \mathcal{P}(\mathcal{B}) \\ \text{where } \mathcal{P}(\mathcal{B}) = \{p(\vec{t}) \mid \mathcal{P}, \mathcal{B} \vdash p(\vec{t})\}. \end{array}$$

Figure 1: Logical Specification of Datalog Updates

1 Introduction

Datalog, a logic programming language defined in the early 1980s, was originally used to describe and implement deductive databases [GMN84]. Given a set \mathcal{B} representing the data in the database, a Datalog program \mathcal{P} specified recursive views that a user could then query. Logically, \mathcal{B} is a set of ground atomic formulas (facts) and \mathcal{P} a set of implicational formulas of a restricted form (Horn clauses). The semantics of such a Datalog system, as well as some actual implementations [CGT89], was given by the set of the atomic logical consequences of \mathcal{P} with respect to \mathcal{B} , in symbols $\{p(\vec{t}) \mid \mathcal{P}, \mathcal{B} \vdash p(\vec{t})\}$ — a beautifully pure use of logic.

In the last ten years, Datalog has had a resurgence in a variety of domains other than databases: to provide user authentication [LM03], to implement network protocols [GW10, LCG⁺06], to program cyber-physical systems [ARLG⁺09] and to support deductive spreadsheets [Cer07], just to cite a few. In these applications, the contents of the fact base \mathcal{B} is highly dynamic, at least compared to the relatively static databases of the 1980s. This requires efficient algorithms to assert new facts and to retract stale facts. While such algorithms have been proposed [ARLG⁺09, CARG⁺12, GMS93, LCG⁺06, NJLS11], the connection with logic has lagged behind, especially for retraction. Indeed, although it is easy to give a logical specification of assertion and retraction, as done in Figure 1, a direct implementation of these specifications is horribly inefficient as it involves recomputing logical consequences from scratch after each assertion or retraction.

In this report, we propose a logical interpretation of Datalog that supports updates internally. Like with the traditional reading, the semantics of a Datalog system $(\mathcal{P}, \mathcal{B})$ will be the set $\{p(\vec{t}) \mid \mathcal{P}, \mathcal{B} \vdash p(\vec{t})\}$ of its logical consequences (plus some bookkeeping information). However, the assertion or retraction of a fact a will trigger a small number of logical inferences to determine the new set of logical consequences rather than performing a full recalculation. In this way, the proposed encoding not only constitutes a logical specification of the static and dynamic aspects of a Datalog system, but it can also be viewed as an efficient algorithm to compute updates. Because retraction results in the removal of inferred facts, we base our encoding on a fragment of first-order linear logic rather than on the traditional predicate calculus.

The main contributions of this report are therefore 1) the definition of an encoding of Datalog which models the assertion and retraction of facts within logic, and 2) a proof of its correctness with respect to the traditional interpretation.

The rest of this report is organized as follows: Section 2 reviews Datalog and linear logic. Section 3 illustrates our approach with an example. Section 4 introduces an initial linear logic interpretation of Datalog that focuses just on assertion. Section 5 extends it to capture retraction as well. We discuss related work in Section 6, while Section 7 outlines directions of future work.

2 Preliminaries

This section provides a brief review of Datalog and linear logic. We first recall some common mathematical notions that we will use throughout this report.

2.1 Notations

In this report, we will refer to both sets and multisets. We write $\{a_1, \dots, a_n\}$ for a multiset with elements a_1, \dots, a_n (possibly repeated) and $S \uplus S'$ for the multiset union of S and S' . We use \emptyset to denote both the empty set and the empty multiset. Given a multiset S , we write $set(S)$ for the set of the elements of S .

$$\begin{array}{ll}
\text{Terms:} & t ::= x \mid c \\
\text{Facts:} & a, b, h ::= p(\vec{t}) \\
\text{Clause:} & D ::= r : \forall \vec{x}. B \supset h \\
\text{Clause bodies:} & B ::= b \wedge B \mid b \\
\text{Programs:} & \mathcal{P} ::= D, \mathcal{P} \mid \cdot \\
\text{Fact bases:} & \mathcal{B} ::= p(\vec{c}), \mathcal{B} \mid \cdot
\end{array}$$

Figure 2: Syntax of Datalog

(without duplicates). We write \bar{a} for a generic set or multiset, disambiguating as needed. Sequences will be written as (a_1, \dots, a_n) or \vec{a} , with ϵ for the empty sequence and S, S' for the concatenation of two sequences S and S' . A substitution $[t_1/x_1, \dots, t_n/x_n]$, denoted θ , is a mapping of variables x_i to terms t_i . Given a collection S (set, multiset or sequence) of formulas and $\theta = [t_1/x_1, \dots, t_n/x_n]$, we denote with $\theta(S)$ the simultaneous substitution of all free occurrences of x_i in S with t_i . We denote the set of free variables in S as $FV(S)$.

2.2 Datalog

Datalog is a declarative rule-based language, founded on traditional first-order logic. Its syntax is given in Figure 2. A *Datalog system* consists of a fact base, \mathcal{B} , and a program, \mathcal{P} . A *fact base* is a set of ground atomic formulas, called *facts*, of the form $p(\vec{c})$, where p is a predicate of fixed arity and \vec{c} is a sequence of constants drawn from a fixed and finite signature Σ . A *Datalog program* is a set of *clauses* of the form $r : \forall \vec{x}. b_1 \wedge \dots \wedge b_n \supset h$ where h and b_i are atomic formulas possibly mentioning variables among \vec{x} , and r is a unique label acting as the rule's name. We call h the *head* of r and $b_1 \wedge \dots \wedge b_n$ its *body*. Datalog clauses are a restricted form of Horn clauses: the body cannot be empty, the arguments of each predicate must be either constants or variables (function symbols are disallowed), and each variables in h must appear in the antecedent $b_1 \wedge \dots \wedge b_n$ (variables are *range-restricted*). Furthermore, to keep the presentation simple, we shall assume that no two body atoms have common instances.

Datalog classifies predicates p into *base predicates*, whose values are supplied by the user and can appear only in the body of a clause, and *inferred predicates* corresponding to views computed by the system. Clearly, all heads have inferred predicates. We classify the facts built from these predicates into *base facts* and *inferred facts*, respectively. The set of all facts that can be built from the signature Σ and the predicates mentioned in a program \mathcal{P} is traditionally called the *Herbrand base* of \mathcal{P} . We denote it $herbrand(\mathcal{P})$. Furthermore, we write $herbrand_B(\mathcal{P})$ and $herbrand_I(\mathcal{P})$ for the portions of $herbrand(\mathcal{P})$ that contain base and inferred facts, respectively. The common assumption that both Σ and \mathcal{P} are finite ensures that these sets are finite as well.

The semantics of a Datalog system $(\mathcal{P}, \mathcal{B})$ is defined as the set of all atomic logical consequences of \mathcal{P} and \mathcal{B} , formally $\{p(\vec{t}) \mid \mathcal{P}, \mathcal{B} \vdash p(\vec{t})\}$. This set, the *inferential closure* of \mathcal{P} and \mathcal{B} , consists of both base and inferred facts. We will denote it as $\mathcal{P}(\mathcal{B})$, thereby emphasizing that a fixed program \mathcal{P} can be viewed as a function that maps a fact base \mathcal{B} to the set of its logical consequences. This is particularly convenient given that this report focuses on the effect of dynamic updates to the fact base of a Datalog system. Figure 3 details a traditional logic proof system, the intuitionistic sequent calculus, *LJ*. We will make reference to *LJ* in our proofs in Section 4.3 and 5.5.

Given a fact base \mathcal{B} , the *assertion* of a fact a yields the updated fact base $\mathcal{B}' = \mathcal{B} \cup \{a\}$ (which we abbreviate “ \mathcal{B}, a ”). Similarly, the *retraction* of a from \mathcal{B} produces $\mathcal{B}' = \mathcal{B} \setminus \{a\}$. For convenience, we will assume that $a \notin \mathcal{B}$ in the first case and $a \in \mathcal{B}$ in the second. The effect of either type of update is to change the set $\mathcal{P}(\mathcal{B})$ to $\mathcal{P}(\mathcal{B}')$. We denote this as $\mathcal{P}(\mathcal{B}) \xrightarrow{\alpha} \mathcal{P}(\mathcal{B}')$ where the *action* α is $+a$ in the case of assertion and $-a$ for retraction. The logical meaning of both forms of update is summarized in Figure 1. Although it is possible to compute $\mathcal{P}(\mathcal{B}')$ from scratch on the basis of \mathcal{P} and \mathcal{B}' , it is generally much more efficient to do so starting from $\mathcal{P}(\mathcal{B})$. Although a number of algorithms have been proposed to do exactly this [ARLG⁺09, CARG⁺12, GMS93, LCG⁺06, NJLS11], they lack a logical justification, especially in the

$$\begin{array}{l}
\text{Formula } A, B, C \quad ::= \quad p(\vec{t}) \mid A \wedge B \mid A \supset B \mid \forall x.A \\
\text{Context } \Delta \quad ::= \quad \cdot \mid \Delta, A
\end{array}$$

$$\begin{array}{c}
\frac{}{A \vdash A} (Id) \qquad \frac{\Delta \vdash C}{\Delta, A \vdash C} (W_L) \qquad \frac{\Delta, A, A \vdash C}{\Delta, A \vdash C} (C_L) \qquad \frac{\Delta_1 \vdash A \quad \Delta_2, B \vdash C}{\Delta_1, \Delta_2, A \supset B \vdash C} (\supset_L) \\
\\
\frac{t \in \Sigma \quad \Delta, [t/x]A \vdash C}{\Delta, \forall x.A \vdash C} (\forall_L) \qquad \frac{\Delta, A_i \vdash C}{\Delta, A_1 \wedge A_2 \vdash C} (\wedge_{L_i}) \qquad \frac{\Delta \vdash A \quad \Delta \vdash B}{\Delta \vdash A \wedge B} (\wedge_R)
\end{array}$$

Figure 3: Fragment of the Intuitionistic Sequent Calculus, LJ

case of retraction. This report provides such a justification.

2.3 Linear Logic: LV^{obs} and LV^{obs-}

Linear logic [Gir87] is a substructural refinement of traditional logic. In traditional logic, assumptions grow monotonically during proof search. By restricting the use of the structural rules of contraction and weakening, linear logic enables the context to grow and shrink as inference rules are applied. This permits interpreting formulas as *resources* that are produced and consumed during the process of constructing a derivation. The result is a rich logic that is well-suited for modeling concurrent and dynamic systems [CS09].

In linear logic, the usage pattern of a formula is made explicit by replacing the connectives of traditional logic with a new set of operators. For instance, multiplicative conjunction or tensor ($A \otimes B$) denotes the simultaneous occurrence of A and B , and additive conjunction ($A \& B$) denotes the exclusive occurrence of either A or B . Hence \otimes and $\&$ models two understandings of “conjunction” on resources. Traditional logic implication $A \supset B$ is replaced by $A \multimap B$ which preserves that idea of deducing B from A , but differs in that A is consumed upon the application of $A \multimap B$. An atom $p(\vec{t})$ is defined as in Figure 2. The fragment of linear logic used in this report is shown in the top part of Figure 4. With the exception of the \multimap_a operator (discussed below), it is a subset of LV^{obs} [CS09], itself a restriction of linear logic that has been used to faithfully encode concurrent languages and rewriting systems [CS09].

We use the judgment $\Gamma; \Delta \rightarrow C$ to express that the linear logic formula C is derivable from the *unrestricted context* Γ (a set of formulas that can be used as often as needed) and the *linear context* Δ (a multiset of formulas that must be used exactly once) [Pfe95, CS09]. The proof system defining this judgment is given in the bottom part of Figure 4. Rules (*obs*), (\multimap_L) and (\multimap_L^-) make use of the *tensorial closure* $\otimes \Delta$ of a context Δ , defined as follows:

$$\begin{array}{l}
\otimes(\cdot) \quad = \quad 1 \\
\otimes(A, \Delta) \quad = \quad A \otimes (\otimes \Delta)
\end{array}$$

The language used in this report, LV^{obs-} , adds one construct to LV^{obs} (highlighted in Figure 4): *guarded implication*, written $A \multimap_a B$, where A and B are any formulas and a is an atomic formula. Its left rule (\multimap_L^-) differs from that of regular implication in that it is applicable only when the linear context *does not* contain the atomic formula a ; moreover, once applied, a is inserted into the context alongside the consequent B . It can be thought of as an implication with *negation-as-absence* on an atomic formula a (i.e., $\neg a \otimes A \multimap B \otimes a$).

A formula of the form $A \multimap B \otimes A$ is called *reassertive* in that applying it consumes A from the linear context Δ but immediately restores it there. If $A \multimap B \otimes A$ is in the unrestricted context Γ , it can then be applied again and again, producing arbitrarily many copies of B . By contrast, the *guarded reassertive formula* $A \multimap_a B \otimes A$ can be used at most once and produce a single B , even if it is unrestricted. Indeed, it is enabled only if $a \notin \Delta$ but its application adds a to Δ thereby preventing additional uses. Therefore, a

$$\begin{array}{ll}
\text{Atomic formulas:} & a, b ::= p(\vec{t}) \\
\text{Formulas:} & A, B, C ::= a \mid 1 \mid A \otimes B \mid A \& B \mid A \multimap B \mid \forall x. A \mid \boxed{A \multimap_a B} \\
\text{Unrestricted contexts:} & \Gamma ::= . \mid \Gamma, A \\
\text{Linear contexts:} & \Delta ::= . \mid \Delta, A
\end{array}$$

$$\begin{array}{ccc}
\frac{}{\Gamma; \Delta \longrightarrow \otimes \Delta} \text{ (obs)} & \frac{\Gamma, A; \Delta, A \longrightarrow C}{\Gamma, A; \Delta \longrightarrow C} \text{ (clone)} & \frac{\Gamma; \Delta \longrightarrow C}{\Gamma; \Delta, 1 \longrightarrow C} \text{ (1}_L\text{)} \\
\frac{\Gamma; \Delta, A_1, A_2 \longrightarrow C}{\Gamma; \Delta, A_1 \otimes A_2 \longrightarrow C} \text{ (}\otimes_L\text{)} & \frac{\Gamma; \Delta, A_i \longrightarrow C}{\Gamma; \Delta, A_1 \& A_2 \longrightarrow C} \text{ (}\&_{L_i}\text{)} & \frac{t \in \Sigma \quad \Gamma; \Delta, [t/x]A \longrightarrow C}{\Gamma; \Delta, \forall x. A \longrightarrow C} \text{ (}\forall_L\text{)} \\
\frac{\Gamma; \Delta, B \longrightarrow C}{\Gamma; \Delta, \Delta', (\otimes \Delta') \multimap B \longrightarrow C} \text{ (}\multimap_L\text{)} & \boxed{\frac{a \notin \Delta \quad \Gamma; \Delta, B, a \longrightarrow C}{\Gamma; \Delta, \Delta', (\otimes \Delta') \multimap_a B \longrightarrow C} \text{ (}\multimap_{L'}\text{)}}
\end{array}$$

Figure 4: LV^{obs} and LV^{obs-} Intuitionistic Linear Logic

can act both as a guard and a witness of the application of this rule. We will make abundant use of guarded reassertive formulas and choose the witness a in such a way that distinct instances can be used at most once in a derivation.

Given unrestricted context Γ and linear context Δ , the pair (Γ, Δ) is *quiescent* if whenever $\Gamma; \Delta \longrightarrow \otimes \Delta'$ it must be the case that $\Delta = \Delta'$. In the sequel, we will be encoding Datalog clauses as guarded reassertive implications in Γ and facts as atomic formulas in Δ . Because the rules in Figure 4 operate exclusively on the left-hand side of a sequent, the inferential closure of a fact base will be obtained when these contexts reach quiescence. The primary purpose of guarded implications is indeed to allow us to reach quiescence. This is crucial to our results in Section 4.

3 Example

Consider the Datalog program \mathcal{P} shown in Figure 5, together with a simplified version $\llbracket \mathcal{P} \rrbracket'$ of its linear logic interpretation. This program computes the paths in a graph: the base predicate $E(x, y)$ means that there is an edge between nodes x and y , while the inferred predicate $P(x, y)$ indicates that there is a path between them. Each Datalog clause is represented by an unrestricted linear implication called an *inference rule* (\mathcal{I}_1 and \mathcal{I}_2). Intuitively, Datalog facts are not quite *linear* in that they should not be consumed when applied to inference rules, yet we do not want to model them as *unrestricted* formulas in Γ because they are subject to retraction. Hence we model Datalog facts as *linear* atomic formulas in Δ which can be consumed, and Datalog clauses as reassertive implications. In Figure 5, we model the retraction of a fact $E(x, y)$ by defining an equal and opposite fact $\tilde{E}(x, y)$ that will consume a copy of the atomic formula $E(x, y)$ ¹. Each inference rule, modeling the inference of the head, embeds one or more scoped linear implications called *retraction rules* (highlighted in boxes), that deals with the retraction of this same fact.

Embedded retraction rules prepare for the retraction of the head of the corresponding clause in the event of the retraction of one of its body elements. For instance, the retraction rule in \mathcal{I}_1 states that if $E(x, y)$ is retracted — represented by $\tilde{E}(x, y)$, we shall retract $P(x, y)$ — represented by $\tilde{P}(x, y)$. Similarly, inference rule \mathcal{I}_2 contains two embedded retraction rules, each preparing for the retraction of its consequent $\tilde{P}(x, z)$ in the event of the retraction of either of its antecedents ($\tilde{E}(x, y)$ or $\tilde{P}(y, z)$, respectively). These two retraction rules are connected by the $\&$ operator, thereby strictly enforcing the exclusive use of either of the rules and

¹ \tilde{E} and \tilde{P} are ordinary predicates. The notation is meant to indicate that they are related to E and P , respectively.

$$\mathcal{P} = \begin{cases} r_1 : \forall x, y. E(x, y) \supset P(x, y) \\ r_2 : \forall x, y, z. E(x, y) \wedge P(y, z) \supset P(x, z) \end{cases}$$

Linear Logic Interpretation (Simplified):

$$\llbracket \mathcal{P} \rrbracket' = \begin{cases} \forall x, y. E(x, y) \\ \multimap P(x, y) \otimes E(x, y) \otimes \\ \boxed{(\tilde{E}(x, y) \multimap \tilde{P}(x, y) \otimes \tilde{E}(x, y))} & - \mathcal{I}_1 \\ \forall x, y, z. E(x, y) \otimes P(y, z) \\ \multimap P(x, z) \otimes E(x, y) \otimes P(y, z) \otimes \\ \boxed{(\tilde{E}(x, y) \multimap \tilde{P}(x, z) \otimes \tilde{E}(x, y))} \& & - \mathcal{I}_2 \\ \boxed{(\tilde{P}(y, z) \multimap \tilde{P}(x, z) \otimes \tilde{P}(y, z))} \end{cases}$$

$$\mathcal{A}_{\mathcal{P}} = \begin{cases} \forall x, y. E(x, y) \otimes \tilde{E}(x, y) \multimap 1 & - \mathcal{A}_1 \\ \forall x, y. P(x, y) \otimes \tilde{P}(x, y) \multimap 1 & - \mathcal{A}_2 \end{cases}$$

Figure 5: Example of Linear Logic Interpretation

hence ensuring that we introduce at most a single $\tilde{P}(x, z)$ in the event of the retraction of either of $P(x, z)$'s antecedents. Retraction rules enter a derivation as *linear* formulas, which means that they disappear the moment they are used. The retraction of some base fact $E(a, b)$ will trigger a cascade of applications of retraction rules, effectively retracting all paths $P(c, d)$ that are causally dependent on $E(a, b)$. Notice that retraction rules are reassertive (their antecedents appear as in their consequent). This is because a single retraction atom $\tilde{E}(x, y)$ might be responsible for producing more than one retraction atoms (e.g. $\tilde{E}(1, 2)$ might need to produce $\tilde{P}(1, 2)$ and $\tilde{P}(1, 3)$) hence the actual consumption of a retraction atom $\tilde{E}(x, y)$ cannot be the responsibility of a retraction rule, but is pushed to an *absorption rule* (e.g., \mathcal{A}_1 and \mathcal{A}_2), which models the annihilation of equal and opposite atoms.

The interpretation $\llbracket \mathcal{P} \rrbracket'$ illustrated in Figure 5 is a simplified form of our approach. Indeed, some additional bookkeeping information needs to be included to guarantee quiescence and correctness. Sections 4 and 5 will describe refinements that will provide such guarantees.

4 Modeling Datalog Assertion

In this section, we define a linear logic interpretation of Datalog program which focuses only on assertions. We call this interpretation llD^{Assert} . It maps a Datalog assertions inference $\mathcal{P}(\mathcal{B}) \xrightarrow{+a}_{\mathcal{P}} \mathcal{P}(\mathcal{B}, a)$ to a derivation of an LV^{obs-} sequent $\Gamma; \Delta \longrightarrow \otimes \Delta'$ where Γ encodes \mathcal{P} , Δ encodes $\mathcal{P}(\mathcal{B})$ together with a , and Δ' encodes $\mathcal{P}(\mathcal{B}, a)$.

4.1 Inference Rules and Interpretation

We model a Datalog clause as a guarded implication in the unrestricted context and a Datalog fact as a linear atomic formula. Given a Datalog clause $r : \forall \vec{x}. b_1 \wedge \dots \wedge b_n \supset h$, we interpret it as a linear logic implication $\forall \vec{x}. b_1 \otimes \dots \otimes b_n \multimap_{r^\sharp(\vec{x})} h \otimes b_1 \otimes \dots \otimes b_n$ which we call an *inference rule* (denoted by \mathcal{I}). The witness $r^\sharp(\vec{x})$ serves the purpose of reaching quiescence — it is discussed in detail below. While inference rules are unrestricted formulas that can be applied as many times as needed (more precisely, as many times as the (\multimap_L^-) rule permits), Datalog facts are linear atomic formulas. The responsibility of modeling their non-linearity is

$$\begin{aligned}
\mathcal{P} &= I_1 : \forall x, y. E(x, y) \supset P(x, y) \quad , \quad I_2 : \forall x, y, z. E(x, y) \wedge P(y, z) \supset P(x, z) \\
[\mathcal{P}] &= \forall x, y. \mathcal{I}_1^{(x,y)}, \forall x, y, z. \mathcal{I}_2^{(x,y,z)} \\
\mathcal{I}_1^{(x,y)} &= E(x, y) \multimap_{I_1^\#(x,y)} P(x, y) \otimes E(x, y) \\
\mathcal{I}_2^{(x,y,z)} &= E(x, y) \otimes P(y, z) \multimap_{I_2^\#(x,y,z)} P(x, z) \otimes E(x, y) \otimes P(y, z)
\end{aligned}$$

Assertion of $E(1, 2)$: $\emptyset \xrightarrow{+E(1,2)}_{[\mathcal{P}]}^{LL} \Delta$ **where** $\Delta = \{P(1, 2), E(1, 2), I_1^\#(1, 2)\}$

$$\begin{array}{c}
\overline{[\mathcal{P}]; P(1, 2), E(1, 2), I_1^\#(1, 2)} \longrightarrow \boxed{\otimes \Delta} \quad (obs) \\
\mathcal{I}_1^{(1,2)} \overline{\overline{[\mathcal{P}]; E(1, 2)} \longrightarrow \otimes \Delta} \quad (\mathcal{I}_L)
\end{array}$$

Assertion of $E(2, 3)$:

$\Delta \xrightarrow{+E(2,3)}_{[\mathcal{P}]}^{LL} \Delta'$ **where** $\Delta' = \{P(1, 2), P(2, 3), P(1, 3), E(1, 2), E(2, 3), I_1^\#(1, 2), I_1^\#(2, 3), I_2^\#(1, 2, 3)\}$

$$\begin{array}{c}
\overline{[\mathcal{P}]; P(1, 2), P(2, 3), P(1, 3), E(1, 2), E(2, 3), I_1^\#(1, 2), I_1^\#(2, 3), I_2^\#(1, 2, 3)} \longrightarrow \boxed{\otimes \Delta'} \quad (obs) \\
\mathcal{I}_2^{(1,2,3)} \overline{\overline{[\mathcal{P}]; P(1, 2), \overline{P(2, 3), E(1, 2)}, E(2, 3), I_1^\#(1, 2), I_1^\#(2, 3)} \longrightarrow \otimes \Delta'} \quad (\mathcal{I}_L) \\
\mathcal{I}_1^{(2,3)} \overline{\overline{[\mathcal{P}]; P(1, 2), E(1, 2), \overline{E(2, 3)}, I_1^\#(1, 2)} \longrightarrow \otimes \Delta'} \quad (\mathcal{I}_L)
\end{array}$$

Figure 6: LV^{obs-} derivations of an example llD^{Assert} interpretation

achieved by making inference rule reassertive. This choice is taken in anticipation to extending our logic interpretation with retraction in Section 5.

Figure 7 formalizes the encoding function $[-]$ which takes a Datalog program \mathcal{P} and returns the corresponding linear logic interpretation $[\mathcal{P}]$. Datalog facts are mapped to atomic formulas, and multisets of facts are mapped to linear contexts of atomic formulas. A Datalog clause is mapped to a guarded reassertive implication.

Inference rules are guarded by an *atomic formula witness* $r^\#(\vec{x})$ where r is the unique name of the clause and \vec{x} are all the variables of the clause. Note that $r^\#$ is not a Datalog predicate but an artifact of the encoding. In particular $r^\#(\vec{c}) \notin \text{herbrand}(\mathcal{P})$ for any \vec{c} . We write $\Delta^\#$ for a context fragment consisting only of rule witnesses. An llD^{Assert} state is a multiset of facts and witnesses. Specifically, given a Datalog program

$$\begin{array}{l}
\text{Fact} \quad [p(\vec{t})] = p(\vec{t}) \\
\text{Program} \quad \left\{ \begin{array}{l} [D, \mathcal{P}] = [D], [\mathcal{P}] \\ [\cdot] = \cdot \end{array} \right. \\
\text{Conjunction} \quad [b \wedge B] = [b] \otimes [B] \\
\text{Clause} \quad [r : \forall \vec{x}. B \supset h] = \forall \vec{x}. [B] \multimap_{r^\#(\vec{x})} [h] \otimes [B] \\
\text{Fact base} \quad \left\{ \begin{array}{l} [\mathcal{B}, a] = [\mathcal{B}], [a] \\ [\emptyset] = \cdot \end{array} \right.
\end{array}$$

Figure 7: llD^{Assert} Interpretation of Datalog

$$\begin{array}{l}
\text{Let } \lceil \mathcal{P} \rceil \equiv \Gamma, \mathcal{I} \\
\mathcal{I} \equiv \forall \vec{x}. \otimes \Delta \multimap_{r^\sharp(\vec{x})} h \otimes \otimes \Delta \\
\\
\frac{r^\sharp(\vec{t}) \notin \Delta' \quad \lceil \mathcal{P} \rceil; [\vec{t}/\vec{x}]h, [\vec{t}/\vec{x}]\Delta, r^\sharp(\vec{t}), \Delta' \longrightarrow C}{\lceil \mathcal{P} \rceil; [\vec{t}/\vec{x}]\Delta, \Delta' \longrightarrow C} (\mathcal{I}_L)
\end{array}$$

which decomposes into:

$$\frac{r^\sharp(\vec{t}) \notin \Delta' \quad \frac{\frac{\lceil \mathcal{P} \rceil; [\vec{t}/\vec{x}]h, [\vec{t}/\vec{x}]\Delta, r^\sharp(\vec{t}), \Delta' \longrightarrow C}{\lceil \mathcal{P} \rceil; [\vec{t}/\vec{x}](h \otimes \otimes \Delta), r^\sharp(\vec{t}), \Delta' \longrightarrow C} (\otimes_L)^*}{\lceil \mathcal{P} \rceil; [\vec{t}/\vec{x}]\mathcal{I}, [\vec{t}/\vec{x}]\Delta, \Delta' \longrightarrow C} (\forall_L)^*}{\lceil \mathcal{P} \rceil; \mathcal{I}, [\vec{t}/\vec{x}]\Delta, \Delta' \longrightarrow C} (\text{clone})}{\lceil \mathcal{P} \rceil; [\vec{t}/\vec{x}]\Delta, \Delta' \longrightarrow C} (\multimap_L^-)$$

Figure 8: llD^{Assert} Inference Macro Rule

\mathcal{P} , a multiset of facts $F \subseteq \text{herbrand}(\mathcal{P})$ and a multiset of rule witnesses Δ^\sharp , an llD^{Assert} state is defined as the set of linear formula $\lceil F \rceil, \Delta^\sharp$. Moreover, given base facts $\mathcal{B} \subseteq \text{herbrand}_B(\mathcal{P})$, the state $\lceil F \rceil, \Delta^\sharp$ is *reachable* from \mathcal{B} if $\lceil \mathcal{P} \rceil; \lceil \mathcal{B} \rceil \longrightarrow \otimes \lceil F \rceil \otimes \otimes \Delta^\sharp$ is derivable. Note that $\lceil \mathcal{B} \rceil$ is always reachable from \mathcal{B} — it is called an *initial state*.

To model the application of a Datalog clause D , we consider the sequence of linear logic derivation steps that corresponds to the cloning and instantiation of inference rule $\lceil D \rceil$, followed by the actual application of the inference rule instance and assembly of the tensor closure of consequents into atomic formulas. Figure 8 factors out the mapping of this sequence of linear logic derivation steps into a *macro rule* (\mathcal{I}_L) . Lemma 1 in Section 4.3 entails that any derivation on LV^{obs-} states has an equivalent derivation comprised of only (\mathcal{I}_L) and (obs) derivation steps.

We informally define a transition $\Delta \xrightarrow{+a}_{\lceil \mathcal{P} \rceil} \Delta'$ that maps a state Δ to another state Δ' such that Δ' has reached quiescence and $\lceil \mathcal{P} \rceil; \Delta, a \longrightarrow \otimes \Delta'$. We will extend it to retraction in Section 5 and formally define it in Section 5.3.

4.2 Quiescence and Complete Inference

Figure 6 illustrates an example LV^{obs-} derivation based on the llD^{Assert} interpretation of our example in Figure 5. It features two successive assertion steps from an empty state. For clarity, we highlight (using boxes) the formula fragments that are acted upon in each derivation step. The first derivation represents the assertion of $E(1, 2)$ from the empty state, during which the inference rule instance $\mathcal{I}_1^{(1,2)}$ is applied and we produce the state Δ . Note that without the guards $I_1^\sharp(x, y)$ and $I_2^\sharp(x, y, z)$, the inference rules would allow the derivation of goal formulas with arbitrarily many copies of the head of the clauses that have been applied. Instead, from $E(1, 2)$ we can only derive a single instance of $P(1, 2)$ because an application of $\mathcal{I}_1^{(x,y)}$ introduces the atomic formula $I_1^\sharp(1, 2)$ which, by rule (\multimap_L^-) , prevents further application of $\mathcal{I}_1^{(x,y)}$. In fact, derivations such as this will always reach quiescence. The second derivation represents the assertion of $E(2, 3)$ from the state Δ obtained in the previous derivation. Here, we reach quiescence with Δ' after applying inference rule instances $\mathcal{I}_1^{(2,3)}$ and $\mathcal{I}_2^{(1,2,3)}$.

Given a Datalog program \mathcal{P} , a state Δ reachable from fact base \mathcal{B} and some new fact a , to derive the inferential closure of (\mathcal{B}, a) , we construct the derivation of $\lceil \mathcal{P} \rceil; \Delta, a \longrightarrow \otimes \Delta'$ such that Δ' is quiescent with respect to $\lceil \mathcal{P} \rceil$.

$$\mathcal{P} = I : \forall x. A(x) \wedge B(x) \supset D(x)$$

$$\begin{array}{ll} \lceil \mathcal{P} \rceil' &= \forall x. \mathcal{I}_1^{(x)} & \mathcal{A}_{\mathcal{P}} &= \forall x. \mathcal{A}_A^{(x)}, \forall x. \mathcal{A}_B^{(x)}, \forall x. \mathcal{A}_D^{(x)} \\ \mathcal{I}_1^{(x)} &= A(x) \otimes B(x) \multimap_{I^\sharp(x)} D(x) \otimes A(x) \otimes B(x) \otimes \mathcal{R}_1^{(x)} & \mathcal{A}_A^{(x)} &= A(x) \otimes \tilde{A}(x) \multimap 1 \\ \mathcal{R}_1^{(x)} &= (\tilde{A}(x) \otimes I^\sharp(x) \multimap \tilde{D}(x) \otimes \tilde{A}(x)) \& & \mathcal{A}_B^{(x)} &= B(x) \otimes \tilde{B}(x) \multimap 1 \\ &(\tilde{B}(x) \otimes I^\sharp(x) \multimap \tilde{D}(x) \otimes \tilde{B}(x)) & \mathcal{A}_D^{(x)} &= D(x) \otimes \tilde{D}(x) \multimap 1 \end{array}$$

$$\begin{array}{c} \textbf{Assertion of } B(5) \\ \lceil A(5) \rceil \xrightarrow{+B(5)}_{\lceil \mathcal{P} \rceil}^{LL} \Delta' \text{ where } \Delta' = \lceil A(5), B(5), D(5), \mathcal{R}_1^{(5)}, I^\sharp(5) \rceil \\ \hline \mathcal{I}_1^{(5)} \frac{\lceil \mathcal{P} \rceil', \mathcal{A}_{\mathcal{P}}; \boxed{A(5), B(5), D(5), \mathcal{R}_1^{(5)}, I^\sharp(5)} \longrightarrow \boxed{\otimes \Delta'}}{\lceil \mathcal{P} \rceil', \mathcal{A}_{\mathcal{P}}; \boxed{A(5), B(5)} \longrightarrow \otimes \Delta'} \text{ (obs)} \\ \hline \textbf{Retraction of } A(5) \\ \Delta' \xrightarrow{-A(5)}_{\lceil \mathcal{P} \rceil}^{LL} \Delta'' \text{ where } \Delta'' = \lceil B(5) \rceil \\ \hline \mathcal{A}_A^{(5)} \frac{\lceil \mathcal{P} \rceil', \mathcal{A}_{\mathcal{P}}; \boxed{B(5)} \longrightarrow \boxed{\otimes \Delta''}}{\lceil \mathcal{P} \rceil', \mathcal{A}_{\mathcal{P}}; \boxed{A(5)}, B(5), \tilde{A}(5)} \longrightarrow \otimes \Delta'' \text{ (obs)} \\ \hline \mathcal{A}_D^{(5)} \frac{\lceil \mathcal{P} \rceil', \mathcal{A}_{\mathcal{P}}; \boxed{A(5)}, B(5), \boxed{D(5)}, \tilde{A}(5), \boxed{\tilde{D}(5)}} \longrightarrow \otimes \Delta''}{\lceil \mathcal{P} \rceil', \mathcal{A}_{\mathcal{P}}; \boxed{A(5), B(5), D(5)}, \mathcal{R}_1^{(5)}, I^\sharp(5), \tilde{A}(5)} \longrightarrow \otimes \Delta'' \\ \hline \mathcal{R}_1^{(5)} \frac{\lceil \mathcal{P} \rceil', \mathcal{A}_{\mathcal{P}}; \boxed{A(5), B(5), D(5)}, \mathcal{R}_1^{(5)}, I^\sharp(5), \tilde{A}(5)} \longrightarrow \otimes \Delta'' \end{array}$$

Figure 9: An Example of llD^{Full-} Interpretation with Assertion/Retraction Derivations

4.3 Correctness

In this section, we argue for the soundness and completeness of the llD^{Assert} linear logic interpretation of Datalog program.

As mentioned earlier, for LV^{obs-} sequents of the form we are interested in, we can always find a normalized LV^{obs-} derivation that consist just of (\mathcal{I}_L) and (obs) rules. This simplifies the proof of the correctness of our interpretation.

Lemma 1 (Normalized Derivations) *Given a Datalog program \mathcal{P} and states $\lceil F_1 \rceil, \Delta_1^\sharp$ and $\lceil F_2 \rceil, \Delta_2^\sharp$, a sequent $\lceil \mathcal{P} \rceil; \lceil F_1 \rceil, \Delta_1^\sharp \longrightarrow \otimes \lceil F_2 \rceil \otimes \otimes \Delta_2^\sharp$ has a derivation consisting of only applications of (\mathcal{I}_L) and (obs) derivation steps.*

Proof: The bulk of the proof of this Lemma follows from *permutability* results of linear logic [GPI94]. We need show that we can always disentangle a generic derivation of any sequent $\lceil \mathcal{P} \rceil; \lceil F_1 \rceil, \Delta_1^\sharp \longrightarrow \otimes \lceil F_2 \rceil \otimes \otimes \Delta_2^\sharp$, into an equivalent one which consists of an orderly sequence of LV^{obs-} derivation steps that corresponds to macro derivation steps (\mathcal{I}_L) described in Definition 8 and ending off with a (obs) derivation step. We need to show that all (\otimes_L) derivation steps applying on a formula $A \otimes B$ can be safely permuted downwards until right after the (\multimap_L^-) derivation step which infers the formula $A \otimes B$. Thanks to permutability results in [GPI94], we have this property for all derivation rules except for (\multimap_L^-) . To push the (\otimes_L) derivation step involving formula $A \otimes B$ safely downwards pass a (\multimap_L^-) step operating on disjoint set of formulas C and $C \multimap_a D$, we argue that for interpretation of well-formed Datalog programs, we always have $a \notin A$

and $a \notin B$ since rule witnesses a never occur in tensor formulas, hence such permutations are safe. Similarly, we permute upwards all occurrences of (\forall_L) derivation step applying to a inference rule $B \multimap_a h \otimes B$ until right after the (\multimap_L^-) derivation step which operates on $B \multimap_a h \otimes B \otimes a$. Lastly, occurrences of the *(clone)* derivation steps that operates on inference rules $\forall \vec{x}. B \multimap_a h \otimes B \otimes a$ in the unrestricted context are permuted upwards until right before the first occurrences of (\forall_L) that operates on the same inference rule. As such, we can replace each orderly sequence of *(clone)*, $(\forall_L)^*$, (\multimap_L^-) , $(\otimes_L)^*$ derivation steps by a (\mathcal{I}_L) . Hence we can obtain an lD^{Assert} normalized proof from any LV^{obs-} sequent $[\mathcal{P}]; [F_1], \Delta_1^\# \longrightarrow \otimes [F_2] \otimes \otimes \Delta_2^\#$. \square

Lemma 2 states the redundancy of having multiple copies of facts, with respect to Datalog inference. Essentially a consequence of the invertibility of the left contraction rule, this lemma allows us to safely go back and forth between the multisets of facts of LV^{obs-} and the sets of facts in traditional logic, in our soundness proof (Theorem 3). In the reverse direction of the proof, it allows us to focus on traditional logic proofs that are sufficiently *saturated*, meaning that multiple applications of the same instance of Horn clauses do not provide more information. This is crucial for our completeness proof (Theorem 4), given that in the lD^{Assert} interpretation of Datalog clauses, inference rules are guarded implications that will reach quiescence, preventing multiple applications of same implication rule instances.

Lemma 2 (Redundancy) *Given a Datalog program \mathcal{P} , states $\Delta_1 \equiv [F_1], \Delta_1^\#$ and $\Delta_2 \equiv [F_2], \Delta_2^\#$ and some fact $b \in \text{herbrand}(\mathcal{P})$, we have the following:*

1. $\mathcal{P}, b, b, F_1 \vdash \bigwedge F_2$ if and only if $\mathcal{P}, b, F_1 \vdash \bigwedge F_2$
2. $[\mathcal{P}]; [b], [b], \Delta_1 \longrightarrow [b] \otimes \otimes \Delta_2$ if and only if $[\mathcal{P}]; [b], \Delta_1 \longrightarrow \otimes \Delta_2$

Proof:

1. This follows immediately from the invertibility of the left contraction rules of the traditional intuitionistic logic proof system, *LJ*.
2. The proof proceeds by contraction as follows: Suppose there is some fact c such that $[\mathcal{P}]; [b], [b], \Delta_1 \longrightarrow [b] \otimes [c] \otimes \otimes \Delta_2$ and $[\mathcal{P}]; [b], \Delta_1 \not\rightarrow_{\Sigma} [b] \otimes [c] \otimes \otimes \Delta_2$, hence it must be the case that the extra $[b]$ has allowed the inference of $[c]$. We consider the 2 possible cases:
 - (*Case 1*) There exists some implication rule $\forall \vec{x} \otimes \Delta \multimap_a c \otimes \otimes \Delta$ in $[\mathcal{P}]$ such that there exists $b', b'' \in \Delta$ and we can find a unifier θ where $\theta(b') = b$ and $\theta(b'') = b$. This means that we can infer c from this rule which consumes two copies of b . Yet this is not possible, since we only consider Datalog clauses where no two body atoms b' and b'' have common instances.
 - (*Case 2*) There exists two implication rule $\otimes \Delta \multimap_{a1} h \otimes \otimes \Delta$ and $\otimes \Delta' \multimap_{a2} c \otimes \otimes \Delta'$ in $[\mathcal{P}]$ such that $b \in \Delta$ and $b \in \Delta'$, and hypothesize that applying the first implication consumes b and disallows application of the second rule if we only have one copy of b . However the inference rules in $[\mathcal{P}]$ are reassertive (antecedents Δ appear as consequents as well), hence the second rule is applicable regardless of whether the first rule applies before or after it.

Hence we can never find such a fact c . \square

Given a Datalog program and a multiset of facts, every valid LV^{obs-} derivation built from the lD^{Assert} interpretation has a corresponding valid traditional logic derivation. Therefore, our linear logic interpretation is sound with respect to the traditional semantics of Datalog.

Theorem 3 (Soundness) *Given a Datalog program \mathcal{P} , fact base \mathcal{B} and reachable lD^{Assert} state $[F], \Delta^\#$, if $[\mathcal{P}]; [\mathcal{B}] \longrightarrow \otimes [F] \otimes \otimes \Delta^\#$, then $\mathcal{P}, \mathcal{B} \vdash \bigwedge F$.*

Proof: For any sequent $[\mathcal{P}]; [\mathcal{B}] \longrightarrow \otimes [F] \otimes \otimes \Delta^\sharp$ constructed from our Datalog interpretation, we inductively define their corresponding normalized derivations as follows: the rule (*obs*) is the only base structural rule, while (\mathcal{I}_L) is the only inductive structural rule. By the form of the (\mathcal{I}_L) rule, we can conclude that all intermediate sequents of the derivation are of the form $[\mathcal{P}]; [F'], \Delta^\sharp \longrightarrow \otimes [F] \otimes \otimes \Delta^\sharp$, with initial sequents as a special case, such that $[F'] \equiv [\mathcal{B}]$ and $[\Delta'] \equiv \emptyset$. By Lemma 1, we can safely assert that no sequents other than those generated by the base (*obs*) and inductive (\mathcal{I}_L) structural rules are valid derivations of our Datalog interpretation. The inductive proof proceeds as follows:

- **Base (*obs*) structural rule:** We consider the sequent proof built from the (*obs*) derivation rule of LV^{obs-} . Hence for $[\mathcal{P}]; [F'], \Delta^\sharp \longrightarrow \otimes [F] \otimes \otimes \Delta^\sharp$ to be derivable for the (*obs*) rule, we must have $[F'] \equiv [F]$ and $\Delta^\sharp \equiv \Delta^\sharp$. As such we need to show that $\mathcal{P}, F' \vdash \bigwedge F$. This is proven in sequent calculus *LJ* simply by applications of the *left weakening* rule (W_L) to discard \mathcal{P} from the left, separating $\bigwedge F$ with applications of the right \wedge rule (\wedge_R) and finishing up with applications of the identity rule (*id*).
- **Inductive (\mathcal{I}_L) structural rule:** We consider a macro derivation step (\mathcal{I}_L). Induction hypothesis, we assume that for $[\mathcal{P}] \equiv \Gamma, (\forall \vec{x}. \otimes B \multimap_a h \otimes \otimes B)$, we have $[\mathcal{P}]; [\vec{t}/\vec{x}]h, [\vec{t}/\vec{x}]B, B', \Delta^\sharp, [\vec{t}/\vec{x}]a \longrightarrow \otimes [F] \otimes \otimes \Delta^\sharp$ where $[F'] \equiv [\vec{t}/\vec{x}]h, [\vec{t}/\vec{x}]B, B'$ and also $\mathcal{P}, F' \vdash F$. By applying the (\mathcal{I}_L) rule, we have $[\mathcal{P}]; [\vec{t}/\vec{x}]B, B', \Delta^\sharp \longrightarrow \otimes [F] \otimes \otimes \Delta^\sharp$. We now need to prove that we also have $\mathcal{P}, F'' \vdash \bigwedge F$ such that $[\vec{t}/\vec{x}]h, F'' \equiv F'$. We do this by considering a sequence of *LJ* derivations that mimics the (\mathcal{I}_L) rule, specifically we do the following: we apply *left contraction* (C_L) rules on each $[\vec{t}/\vec{x}]B \in F'$ then we apply the Datalog rule $D = \forall \vec{x}. B \supset h$ such that $[D] \equiv \forall \vec{x}. \otimes B \multimap_a h \otimes \otimes B$ to obtain $\mathcal{P}, [\vec{t}/\vec{x}]h, F'' \vdash F$ as the premise, which happens to be assumed valid by our induction hypothesis $\mathcal{P}, F' \vdash F$ (since $[\vec{t}/\vec{x}]h, F'' \equiv F'$).

Hence we have shown that Theorem 3 holds for all normalized derivations. \square

Given a Datalog program and a set of facts, every valid traditional logic derivation built from the traditional Horn clause interpretation of Datalog, has a corresponding LV^{obs-} derivation built from its llD^{Assert} interpretation. Our interpretation is therefore also complete.

Theorem 4 (Completeness) *Given a Datalog program \mathcal{P} , fact base \mathcal{B} and $F \subseteq \text{herbrand}(\mathcal{P})$, if $\mathcal{P}, \mathcal{B} \vdash \bigwedge F$, then there is Δ such that $[\mathcal{P}]; [\mathcal{B}] \longrightarrow \otimes [F] \otimes \otimes \Delta$.*

Proof: The proof of this theorem is similar to the proof of Theorem 3, except that we do structural induction on derivations of the corresponding fragment of *LJ* proof system (Figure 3), over well-formed Datalog programs and facts. We inductively define the *LJ* derivations we consider as follows: Base structural derivation steps built from *left weakening* (W_L) rules, (*id*) rules and *right conjunction* (\wedge_R) rules, while induction structural steps consist of an orderly sequence of derivation steps that correspond to application of Horn clauses. Specifically for a Datalog clause $\forall \vec{x}. B \supset h$, we instantiate the Horn clause by applying (\forall_L) rule on each variable \vec{x} , then left contraction rules (C_L) on each atomic formula in the context that matches to B , apply the implication rule (\supset_L). We rely on permutability results of *LJ* derivations [MNPS91] to safely restrict our attention to only sequents inductively generated by the above. Since we define our inductive steps in this manner, all intermediate sequents during the inductive steps are of the form $\mathcal{P}, F' \vdash \bigwedge F$ where $\mathcal{B} \subseteq F'$, with the initial sequent being a special case such that $F' \equiv \mathcal{B}$. We also exploit the results from Lemma 2 that states that multiple copies of facts are redundant, allowing us to omit *LJ* derivations that contains multiple applications of same Horn clause instances that produces the same facts (we cannot replicate this in LV^{obs-} guarded implications). The inductive proof proceeds as follows:

- **Base structural rule:** We consider the *LJ* sequent proof built from *left weakening* (W_L) rules, (*id*) rules and *right conjunction* (\wedge_R) rules. Hence for $\mathcal{P}, F' \vdash \bigwedge F$ to apply, we must have $F' \equiv F$. As such, from $\mathcal{P}, F' \vdash \bigwedge F$ we can use (W_L) rules to remove \mathcal{P} from the context, (\wedge_R) rules to break apart $\bigwedge F$ and (*id*) rules to match the individual atomic formulas in F . We need to show for some Δ^\sharp

and Δ that $\lceil \mathcal{P} \rceil; \lceil F' \rceil, \Delta^\sharp \longrightarrow \otimes \lceil F \rceil \otimes \Delta$. Since $F' \equiv F$, we have $\Delta^\sharp \equiv \Delta$ and apply the (*obs*) rule.

- **Inductive structural rule:** We consider the sequence of *LJ* derivation rules which involves the orderly sequence of derivation steps that results to the applications of Horn clauses as defined above. Induction hypothesis, we assume that given $\mathcal{P}, \lceil \vec{t}/\vec{x} \rceil h, \lceil \vec{t}/\vec{x} \rceil B, B' \vdash \bigwedge F$ and $\mathcal{P} \equiv \Gamma, (r : \forall \vec{x}. B \supset h)$ for some Γ , then we have $\lceil \mathcal{P} \rceil; \lceil \vec{t}, \vec{x} \rceil \lceil h \rceil, \lceil \vec{t}/\vec{x} \rceil \lceil B \rceil, \lceil B' \rceil, \Delta^\sharp \longrightarrow \otimes \lceil F \rceil \otimes \Delta$ for some Δ^\sharp and Δ . We show that we have the *LJ* sequent $\mathcal{P}, \lceil \vec{t}/\vec{x} \rceil B, B' \vdash \bigwedge F$ through the following *LJ* derivation steps: from $\mathcal{P}, \lceil \vec{t}/\vec{x} \rceil B, B' \vdash \bigwedge F$, given the Datalog clause $r : \forall \vec{x}. B \supset h$ we apply (\forall_L) derivation steps on each variable in \vec{x} , followed by (C_L) rules on each $\lceil \vec{t}/\vec{x} \rceil B$ in the context. Next, we apply this implication rule to obtain $\mathcal{P}, \lceil \vec{t}/\vec{x} \rceil h, \lceil \vec{t}/\vec{x} \rceil B, B' \vdash \bigwedge F$ as the premise of our derivation, which we assumed in our induction hypothesis. We now show that $\lceil \mathcal{P} \rceil; \lceil \vec{t}/\vec{x} \rceil \lceil B \rceil, \lceil B' \rceil, \Delta''^\sharp \longrightarrow \otimes \lceil F \rceil \otimes \Delta$ has a valid derivation for some Δ''^\sharp . We do this by applying the (\mathcal{I}_L) derivation step of the corresponding inference rule, i.e., $\forall \vec{x}. B \multimap_{r(\vec{x})} h \otimes B$, which yields $\lceil \mathcal{P} \rceil; \lceil \vec{t}, \vec{x} \rceil \lceil h \rceil, \lceil \vec{t}/\vec{x} \rceil \lceil B \rceil, \lceil B' \rceil, \Delta''^\sharp, r(\vec{t}) \longrightarrow \otimes \lceil F \rceil \otimes \Delta$ as the premise, and by having $\Delta^\sharp \equiv \Delta''^\sharp, r(\vec{t})$ we have our induction hypothesis. Hence we proved the inductive case.

Hence we have shown that Theorem 4 holds for all normalized derivations. \square

From Theorems 4 and 3, we have that the soundness and completeness. Given a fact base \mathcal{B} , we can always derive a state $\lceil F \rceil, \lceil \Delta^\sharp \rceil$ such that the set interpretation of $\lceil F \rceil$ (denoted $set(F)$) is equal to $\mathcal{P}(\mathcal{B})$ and this state is quiescent. Further more, there can only be one such quiescent state.

Theorem 5 (Correctness) *Let \mathcal{P} be a Datalog program and \mathcal{B} a collection of base facts.*

1. *There are $\lceil F \rceil$ and Δ^\sharp such that $\lceil \mathcal{P} \rceil; \lceil \mathcal{B} \rceil \longrightarrow \otimes \lceil F \rceil \otimes \otimes \Delta^\sharp$ and $Quiescent(\lceil \mathcal{P} \rceil, (\lceil F \rceil, \Delta^\sharp))$ and $set(F) = \mathcal{P}(\mathcal{B})$.*
2. *For any state Δ , if $\lceil \mathcal{P} \rceil; \lceil \mathcal{B} \rceil \longrightarrow \otimes \lceil \Delta \rceil$ and $Quiescent(\lceil \mathcal{P} \rceil, \Delta)$, then there is Δ^\sharp such that $\Delta = \lceil F \rceil, \Delta^\sharp$ and $set(F) = \mathcal{P}(\mathcal{B})$.*

Proof: We prove (1) and (2) as follows:

1. From Theorem 4, we can conclude that since $\mathcal{P}, \mathcal{B} \vdash \bigwedge \mathcal{P}(\mathcal{B})$, then we have $\lceil \mathcal{P} \rceil; \lceil \mathcal{B} \rceil \longrightarrow \otimes \lceil F \rceil \otimes \otimes \Delta^\sharp$ for some Δ^\sharp and $set(F) = \mathcal{P}(\mathcal{B})$. We prove that $Quiescent(\lceil \mathcal{P} \rceil, (\lceil F \rceil, \Delta^\sharp))$ by contradiction: Suppose that $Quiescent(\lceil \mathcal{P} \rceil, (\lceil F \rceil, \Delta^\sharp))$ does not hold, hence we can apply more derivation rules from that state. Yet given that our Datalog interpretation only has inference rule instances of the form $\otimes B \multimap_a h \otimes \otimes B$ which are reassertive, applying more such rules only produces an inferred fact h such that $h \notin \mathcal{P}(\mathcal{B})$. Yet we have a corresponding Datalog clause instance $B \supset h$ that should have been applied and captured in the inferential closure $\mathcal{P}(\mathcal{B})$ (i.e. $h \in \mathcal{P}(\mathcal{B})$), as such we have a contradiction. Hence, we must have $Quiescent(\lceil \mathcal{P} \rceil, (\lceil F \rceil, \Delta^\sharp))$.
2. Suppose that we have $\lceil \mathcal{P} \rceil; \lceil \mathcal{B} \rceil \longrightarrow \otimes \lceil \Delta \rceil$ and $Quiescent(\lceil \mathcal{P} \rceil, \Delta)$, but there is no Δ^\sharp such that $\Delta = \lceil F \rceil, \Delta^\sharp$ and $set(F) = \mathcal{P}(\mathcal{B})$. This implies that either we have $set(F) \subset \mathcal{P}(\mathcal{B})$ or $set(F) \supset \mathcal{P}(\mathcal{B})$. We cannot have $set(F) \supset \mathcal{P}(\mathcal{B})$ since by Theorem 3, we cannot infer more facts than $\mathcal{P}(\mathcal{B})$ (i.e. $\mathcal{P}, \mathcal{B} \vdash \bigwedge F$), but yet we cannot have $set(F) \subset \mathcal{P}(\mathcal{B})$ either, which would contradict $Quiescent(\lceil \mathcal{P} \rceil, \Delta)$. Hence our assumption that there is no Δ^\sharp such that $\Delta = \lceil F \rceil, \Delta^\sharp$ and $set(F) = \mathcal{P}(\mathcal{B})$ has reached a contradiction.

Hence we have shown that Theorem 5 holds \square

These results are incremental steps towards achieving the main goal of this report, which is to formalize an interpretation of both assertion and retraction in Datalog, done in Section 5.

$$\begin{aligned}
\mathcal{P} &= I : \forall x. A(x) \wedge B(x) \supset D(x) & \llbracket \mathcal{P} \rrbracket &= \forall x. \mathcal{I}_1^{(x)} \\
\mathcal{I}_1^{(x)} &= A(x) \otimes B(x) \multimap_{I^\sharp(x)} D(x) \otimes A(x) \otimes B(x) \otimes \mathcal{R}_1^{(x)} \otimes \boxed{A^\sharp(x) \otimes B^\sharp(x)} \\
\mathcal{R}_1^{(x)} &= (\tilde{A}(x) \otimes \boxed{A^\sharp(x) \otimes B^\sharp(x)}) \multimap \tilde{D}(x) \otimes \tilde{A}(x) \ \& \ (\tilde{B}(x) \otimes \boxed{A^\sharp(x) \otimes B^\sharp(x)}) \multimap \tilde{D}(x) \otimes \tilde{B}(x) \\
\mathcal{A}_{\mathcal{P}} &= \forall x. \mathcal{A}_A^{(x)}, \forall x. \mathcal{A}_B^{(x)}, \forall x. \mathcal{A}_D^{(x)} \\
\mathcal{A}_A^{(x)} &= A(x) \otimes \tilde{A}(x) \boxed{\multimap_{A^\sharp(x)}} (A^\sharp(x) \multimap 1) \\
\mathcal{A}_B^{(x)} &= B(x) \otimes \tilde{B}(x) \boxed{\multimap_{B^\sharp(x)}} (B^\sharp(x) \multimap 1) \\
\mathcal{A}_D^{(x)} &= D(x) \otimes \tilde{D}(x) \boxed{\multimap_{D^\sharp(x)}} (D^\sharp(x) \multimap 1)
\end{aligned}$$

Assertion of $B(5)$: $\lambda A(5) \int \xrightarrow{\llbracket \mathcal{P} \rrbracket}^{+B(5), LL} \Delta'$ where $\Delta' = \lambda A(5), B(5), D(5), \mathcal{R}_1^{(5)}, I^\sharp(5), A^\sharp(5), B^\sharp(5) \int$

$$\mathcal{I}_1^{(5)} \frac{\overline{\overline{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \boxed{A(5), B(5), D(5), \mathcal{R}_1^{(5)}, I^\sharp(5), A^\sharp(5), B^\sharp(5)} \longrightarrow \boxed{\otimes \Delta'}}}}{\overline{\overline{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \boxed{A(5), B(5)} \longrightarrow \otimes \Delta'}}}} \quad (obs)$$

Retraction of $A(5)$: $\Delta' \xrightarrow{\llbracket \mathcal{P} \rrbracket}^{-A(5), LL} \Delta''$ where $\Delta'' = \lambda B(5) \int$

$$\mathcal{R}_1^{(5)} \frac{\overline{\overline{\overline{\overline{\overline{\overline{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \boxed{B(5)} \longrightarrow \boxed{\otimes \Delta''}}}}}}}}}}{\overline{\overline{\overline{\overline{\overline{\overline{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \boxed{A(5)}, B(5), \boxed{\tilde{A}(5)} \longrightarrow \otimes \Delta''}}}}}}}}}}}}{\overline{\overline{\overline{\overline{\overline{\overline{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; A(5), B(5), \boxed{D(5)}, \tilde{A}(5), \boxed{\tilde{D}(5)} \longrightarrow \otimes \Delta''}}}}}}}}}}}}}}{\overline{\overline{\overline{\overline{\overline{\overline{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; A(5), B(5), D(5), \boxed{\mathcal{R}_1^{(5)}, I^\sharp(5), \tilde{A}(5), A^\sharp(5), B^\sharp(5)} \longrightarrow \otimes \Delta''}}}}}}}}}}}}}} \quad (obs)$$

Figure 10: An Example of llD^{Full} Interpretation

5 Modeling Datalog Retraction

In this section, we extend our linear logic interpretation of Datalog assertion to also account for the retraction of facts. We first illustrate our technique in Section 5.1 by informally introducing a simplified version called llD^{Full-} . Unfortunately, this simplified version admits LV^{obs-} sequent derivations that correspond to incomplete retraction of facts, meaning that when such a derivation reaches quiescence there is no guarantee that the state we obtain corresponds to a valid Datalog object. In Section, 5.2, we modify it into the llD^{Full} interpretation that provides guarantees of completeness when quiescence is achieved. Then in Section 5.3, we formally define the llD^{Full} interpretation that incorporates this modification and correctly models Datalog assertion and retraction.

5.1 Retraction and Absorption Rules

We model the removal of facts by means of *anti-predicates*. An anti-predicate \tilde{p} is an ordinary predicate that is meant to be *equal and opposite* to p . While p applied to terms \vec{t} forms a *fact*, \tilde{p} applied to terms form a *retraction fact*. For instance, recalling the reachability example from Figure 5, the predicates E and P will have equal and opposite anti-predicates \tilde{E} and \tilde{P} respectively. The intent is that an atom $\tilde{p}(\vec{t})$ will ultimately absorb its equal and opposite atom $p(\vec{t})$. To model this property, we define *absorption rules* of

$$\begin{array}{l}
\text{Fact} \quad \llbracket p(\vec{t}) \rrbracket = p(\vec{t}) \\
\text{Program} \quad \left\{ \begin{array}{l} \llbracket D, \mathcal{P} \rrbracket = \llbracket D \rrbracket, \llbracket \mathcal{P} \rrbracket \\ \llbracket \cdot \rrbracket = \cdot \end{array} \right. \\
\text{Conjunction} \quad \llbracket b \wedge B \rrbracket = \llbracket b \rrbracket \otimes \llbracket B \rrbracket \\
\text{Clause} \quad \llbracket r : \forall \vec{x}. B \supset h \rrbracket = \forall \vec{x}. \llbracket B \rrbracket \multimap_{r^\sharp(\vec{x})} \llbracket h \rrbracket \otimes \llbracket B \rrbracket \otimes \llbracket B^\sharp \rrbracket \otimes \text{Ret}(B, B, r^\sharp(\vec{x}), h) \\
\text{Fact base} \quad \left\{ \begin{array}{l} \llbracket \mathcal{B}, a \rrbracket = \llbracket \mathcal{B} \rrbracket, \llbracket a \rrbracket \\ \llbracket \emptyset \rrbracket = \cdot \end{array} \right.
\end{array}$$

where

$$\begin{array}{l}
\text{Ret}(b \wedge B', B, r^\sharp(\vec{x}), h) = (\llbracket \tilde{b} \rrbracket \otimes r^\sharp(\vec{x}) \otimes \llbracket B^\sharp \rrbracket \multimap \llbracket \tilde{h} \rrbracket \otimes \llbracket \tilde{b} \rrbracket) \& \text{Ret}(B', B, r^\sharp(\vec{x}), h) \\
\text{Ret}(b, B, r^\sharp(\vec{x}), h) = (\llbracket \tilde{b} \rrbracket \otimes r^\sharp(\vec{x}) \otimes \llbracket B^\sharp \rrbracket \multimap \llbracket \tilde{h} \rrbracket \otimes \llbracket \tilde{b} \rrbracket)
\end{array}$$

Figure 11: llD^{Full} Interpretation of Datalog

the form $\forall \vec{x}. p(\vec{x}) \otimes \tilde{p}(\vec{x}) \multimap 1$. We illustrate this by means of an example, displayed in Figure 9, which shows a Datalog program \mathcal{P} and its llD^{Full-} interpretation. We have three absorption rules $\mathcal{A}_-^{(x)}$, one for each predicate of this Datalog program. Also demonstrated here, we enrich the linear logic interpretation of Datalog programs with *retraction rules*: Each Datalog clause $r : \forall \vec{x}. p_1(\vec{t}_1) \wedge \dots \wedge p_n(\vec{t}_n) \supset p_h(\vec{t}_h)$ is modeled by an inference rule that contains an additional consequent, the additive conjunction of n retraction rules, each of which is uniquely associated with one of the antecedents $p_i(\vec{t}_i)$. For instance, inference rule $\mathcal{I}_1^{(x)}$ has an embedded retraction rule $\mathcal{R}_1^{(x)}$ in its consequent and $\mathcal{R}_1^{(x)}$ contains two linear implications, joined by the $\&$ operator. Each of these linear implications produces a retraction of $D(x)$ (i.e., $\tilde{D}(x)$) in the event that we have a retraction of $B(x)$ or $C(x)$ respectively.

In the llD^{Full-} interpretation of Datalog programs, we can always find a valid LV^{obs-} derivation that represents each assertion or retraction of facts. Figure 9 illustrates two such derivations, starting from a state containing just the fact $A(5)$, the first derivation shows a derivation in the llD^{Full-} interpretation of this program, that models the assertion of a new fact $B(5)$. As such, we consider the derivation starting from the initial state, consisting of $A(5)$ and the newly asserted fact $B(5)$. We highlight (using boxes) the active fragments of the sequent which contributes to the application of the derivation step. Notice that when the inference rule instance $\mathcal{I}_1^{(5)}$ applies, we have not only the inferred fact $D(5)$, but we have the retraction rule instance $\mathcal{R}_1^{(5)}$ as well. This retraction rule is bookkeeping information that will facilitate the retraction of $D(5)$, in the event that either $A(5)$ or $B(5)$ is retracted. To illustrate this, the second derivation of this figure models the retraction of fact $A(5)$. We consider the state Δ' (obtained in the previous derivation) with the retraction atom $\tilde{A}(5)$ inserted. $\tilde{A}(5)$ and $I^\sharp(5)$ triggers the first additive conjunction component $\tilde{A}(5) \otimes I^\sharp(5) \multimap \tilde{D}(5) \otimes \tilde{A}(5)$ of the retraction rule instance $\mathcal{R}_1^{(5)}$ and produces $\tilde{D}(5)$ while retaining $\tilde{A}(5)$. The purpose of retaining $\tilde{A}(5)$ is that in the general case, $\tilde{A}(5)$ might be necessary for triggering the retraction of other atoms, yet we consume $I^\sharp(5)$ to allow future applications of inference rule instance $I^\sharp(5)$ should $A(5)$ be re-asserted by future updates. Joining the linear implications of retraction rule $\mathcal{R}_2^{(x)}$ with the $\&$ operator allows the creation of exactly *one* $\tilde{D}(5)$ atom from the retraction of $B(5)$ or $C(5)$. Hence we will not be left with a “dangling” copy of $\tilde{D}(5)$ should $C(5)$ be retracted as well (since $\mathcal{R}_1^{(5)}$ has completely been consumed by $\tilde{B}(5)$). We finally complete the derivation by applying the absorption of $D(5)$ and $A(5)$ by $\tilde{D}(5)$ and $\tilde{A}(5)$, via absorption rules $\mathcal{A}_D^{(5)}$ and $\mathcal{A}_A^{(5)}$ respectively. Note that Δ and Δ' are *quiescent* with respect to the llD^{Full-} interpretation of the Datalog program.

5.2 Quiescence and Complete Retraction

In this section we address a shortcoming of the llD^{Full-} interpretation: we cannot guarantee that reachable states derived by valid LV^{obs-} derivations model complete retraction of Datalog facts when they reach quiescence.

Refer to Figure 9 again. In the second derivation, which corresponds to the retraction of $A(5)$ starting from state $\Delta', \tilde{A}(5)$, rather than applying retraction rule $\mathcal{R}_1^{(5)}$, we could have instead immediately applied the absorption rule instance $\mathcal{A}_1^{(5)}$, yielding a state of quiescence and bypassing the retraction of $D(5)$. This state, however, would be an incorrect representation of the effect of retracting $A(5)$, hence such derivations should somehow be disallowed. In general, any pair of equal and opposite atoms $p(\vec{t})$ and $\tilde{p}(\vec{t})$ should be applied to (and consumed by) their corresponding absorption rule only when all retraction rules that can be applied to $p(\vec{t})$ have been applied. In other words, we must *inhibit* the application of absorption rules until all such retraction rules have been applied. To do so, we use the same technique with which we achieved quiescence in the presence of reassertive inference rules: guarded linear implications. Specifically we require that for each inference rule that is applied, for each antecedent $p(\vec{t})$ of this inference rule we produce a witness atom $p^\sharp(\vec{t})$; retraction rules additionally consume these witness $p^\sharp(\vec{t})$. We define absorption rules as guarded implications $p(\vec{x}) \otimes \tilde{p}(\vec{x}) \multimap_{p^\sharp(\vec{x})} (p^\sharp(\vec{x}) \multimap 1)$. The witnesses inhibit absorption rules until they are finally removed by retraction rules. Notice that an absorption rule produces a peculiar consequent $p^\sharp(\vec{x}) \multimap 1$. The purpose for this is to consume the atom $p^\sharp(\vec{x})$ produced as a consequent of applying the (\multimap_L^-) rule.

Figure 10 illustrates this idea on our last example. The boxes highlight the components added on top of the llD^{Full-} interpretation example of Figure 9. We call it the llD^{Full} interpretation, implemented by the function $\llbracket - \rrbracket$, defined in Section 5.3. Specifically inference rule $\mathcal{I}_1^{(x)}$ has two additional consequents $A^\sharp(x)$ and $B^\sharp(x)$, which are witness atoms for antecedents $A(x)$ and $B(x)$ respectively. Absorption rules are now guarded linear implications with witness atoms as inhibitors. For instance, $\mathcal{A}_A^{(x)}$ can only apply if there are no witnesses $A^\sharp(x)$ in the context. Finally, on top of consuming inference witness $I^\sharp(x)$ and producing retraction atom $\tilde{D}(x)$, the retraction rules in $\mathcal{R}_1^{(x)}$ have the additional responsibility of consuming witness atoms $A^\sharp(x)$ and $B^\sharp(x)$.

Figure 10 shows derivations for the llD^{Full} interpretation of our last example (Figure 9). The first derivation corresponds to the assertion of $B(5)$ starting from a state containing a single $A(5)$ and is similar to that in Figure 9 except we additionally infer witness atoms $A^\sharp(5)$ and $B^\sharp(5)$. The next derivation illustrates the subsequent retraction of $A(5)$. Now absorption rule $\mathcal{A}_A^{(5)}$ cannot be applied to $\tilde{A}(5)$ and $A(5)$ as there is a witness atom $A^\sharp(5)$ inhibiting $\mathcal{A}_A^{(5)}$. The only possible way forward is the application of retraction rule $\mathcal{R}_1^{(5)}$, which on top of producing retraction atom $\tilde{D}(5)$ and consuming inference witness $I^\sharp(5)$, consumes witness atoms $A^\sharp(5)$ and $B^\sharp(5)$, thereby unlocking the applicability of $\mathcal{A}_A^{(5)}$.

5.3 The llD^{Full} Interpretation

We now formalize the llD^{Full} interpretation of Datalog programs. For each atom representing a fact $p(\vec{t})$, we define its equal and opposite retraction atom $\tilde{p}(\vec{t})$. We will extend this notation to collections of atoms, hence for a multiset of atoms Δ , we denote a multiset of all equal and opposite retraction atoms of Δ is denoted $\tilde{\Delta}$. Beside the rule witnesses $r^\sharp(\vec{x})$ of Section 4.1, for each atom $p(\vec{x})$ we have another type of witnesses, the *fact witness* $p^\sharp(\vec{x})$. While a rule witness $r^\sharp(\vec{t})$ witnesses the application of an inference rule instance of r , a fact $p(\vec{t})$ witnesses a use of $p^\sharp(\vec{t})$ to infer some other fact. Given a Datalog program \mathcal{P} , for each predicate p of some arity n , we define the set of all absorption rules $\mathcal{A}_{\mathcal{P}}$ as follows:

$$\mathcal{A}_{\mathcal{P}} = \{ \forall \vec{x}. p(\vec{x}) \otimes \tilde{p}(\vec{x}) \multimap_{p^\sharp(\vec{x})} (p^\sharp(\vec{x}) \multimap 1) \mid p \text{ in } \mathcal{P} \}$$

Figure 11 formalizes the llD^{Full} interpretation of Datalog programs as the function $\llbracket - \rrbracket$. The main difference between this function and the function $\llbracket - \rrbracket$ of llD^{Full-} is that Datalog clauses are inference rules that include two additional types of consequents. Specifically, a retraction rule which is an additive

$$\begin{aligned}
\text{Let } \llbracket \mathcal{P} \rrbracket &\equiv \Gamma, \mathcal{I} \\
\mathcal{I} &\equiv \forall \vec{x}. \otimes \Delta \multimap_{r^\#(\vec{x})} h \otimes \otimes \Delta \otimes \mathcal{R}_r \otimes \otimes \Delta^\# \\
&\frac{r^\#(\vec{t}) \notin \Delta' \quad \llbracket \mathcal{P} \rrbracket, \mathcal{A}_\mathcal{P}; [\vec{t}/\vec{x}](h, \Delta, \mathcal{R}_r, \Delta^\#), r^\#(\vec{t}), \Delta' \longrightarrow C}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_\mathcal{P}; [\vec{t}/\vec{x}]\Delta, \Delta' \longrightarrow C} (\mathcal{I}_L)
\end{aligned}$$

which decomposes into:

$$\begin{aligned}
&\frac{r^\#(\vec{t}) \notin \Delta' \quad \frac{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_\mathcal{P}; [\vec{t}/\vec{x}](h, \Delta, \mathcal{R}_r), r^\#(\vec{t}), [\vec{t}/\vec{x}]\Delta^\#, \Delta' \longrightarrow C}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_\mathcal{P}; [\vec{t}/\vec{x}](h \otimes \otimes \Delta \otimes \mathcal{R}_r \otimes \otimes \Delta^\#), r^\#(\vec{t}), \Delta' \longrightarrow C} (\otimes_L)^*}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_\mathcal{P}; [\vec{t}/\vec{x}]\mathcal{I}, [\vec{t}/\vec{x}]\Delta, \Delta' \longrightarrow C} (-\circ_L^-) \\
&\frac{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_\mathcal{P}; [\vec{t}/\vec{x}]\mathcal{I}, [\vec{t}/\vec{x}]\Delta, \Delta' \longrightarrow C}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_\mathcal{P}; \mathcal{I}, [\vec{t}/\vec{x}]\Delta, \Delta' \longrightarrow C} (\forall_L)^* \\
&\frac{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_\mathcal{P}; \mathcal{I}, [\vec{t}/\vec{x}]\Delta, \Delta' \longrightarrow C}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_\mathcal{P}; [\vec{t}/\vec{x}]\Delta, \Delta' \longrightarrow C} (\text{clone})
\end{aligned}$$

Figure 12: llD^{Full} Inference Macro Rule

conjunction of linear implications defined by the function Ret and a tensor closure of inference fact witnesses derived from the antecedents of the inference rule. Altogether, given a Datalog clause $r : \forall \vec{x}. B \supset h$, its corresponding inference rule in llD^{Full} interpretation infers three types of bookkeeping information, namely, the retraction rule \mathcal{R}_r , the rule witness $r^\#(\vec{x})$ and fact witnesses $B^\#$.

Given a Datalog program \mathcal{P} , a multiset of facts $F \subseteq herbrand(\mathcal{P})$ and a multiset of retraction rules $\Delta^{\mathcal{R}}$ and a multiset of witness atoms $\Delta^\#$, a llD^{Full} state of this Datalog program is an LV^{obs-} linear context $\llbracket F \rrbracket, \Delta^{\mathcal{R}}, \Delta^\#$. Given base facts $\mathcal{B} \subseteq herbrand_B(\mathcal{P})$, the state $\llbracket F \rrbracket, \Delta^{\mathcal{R}}, \Delta^\#$ is *reachable* from \mathcal{B} if $\llbracket \mathcal{P} \rrbracket; [\mathcal{B}] \longrightarrow \otimes \llbracket F \rrbracket \otimes \otimes \Delta^{\mathcal{R}} \otimes \otimes \Delta^\#$ is derivable. Note that $\llbracket \mathcal{B} \rrbracket$ is always reachable from \mathcal{B} — it is called an *initial state*.

Similarly to Figure 8, Figure 12 defines the inference macro derivation rule for the llD^{Full} interpretation.

Figure 14 introduces the retraction macro step, which models the retraction of some fact h . As defined in Figure 11 by the function Ret , a retraction rule \mathcal{R}_r is an additive conjunction of linear implications of the form $\tilde{a}_i \otimes r(\vec{t}) \otimes \otimes \Delta^\# \multimap \tilde{a}_i \otimes \tilde{h}$, which are reassertive (antecedent \tilde{a}_i appears as consequent) and the only component that varies among the linear implications is the antecedent \tilde{a}_i . As such, retraction rule \mathcal{R}_r can be viewed as a formula that reacts to any \tilde{a}_i which is an antecedent of one of its linear implications and, when applied, consumes rule witness $r(\vec{t})$ and fact witnesses $\Delta^\#$, and produces \tilde{h} .

Figure 15 defines the last of the three macro derivation steps of the llD^{Full} interpretation. It models the absorption of an atom $p(\vec{t})$ and the equal and opposite $\tilde{p}(\vec{t})$, but only in the absence of any witness $p^\#(\vec{t})$.

Figure 13 illustrates the top-level formulation of the main goal of this report, defining $\Longrightarrow_{\llbracket \mathcal{P} \rrbracket}^{LL}$, our linear logic specification of Datalog assertion and retraction in terms of LV^{obs-} derivations. In contrast to the logical specification $\Longrightarrow_{\mathcal{P}}$ in Figure 1 which defines assertion and retraction based on inferential closures of base facts, $\Longrightarrow_{\llbracket \mathcal{P} \rrbracket}^{LL}$ is defined as follows: the assertion of a new base fact a to a state Δ maps to the derivation with fact a inserted into the original state, while the retraction an existing base fact a from state Δ maps to a derivation with \tilde{a} inserted. In both cases, we consider only resulting states which are quiescent. Theorem 13 in Section 5.5 states the correspondence of the llD^{Full} interpretation that formally prove the correctness this logical specification.

5.4 Cycles and Re-assertion

In this section, we analyze two aspects of the llD^{Full} interpretation. Specifically, we discuss *cycles* in Datalog inference and the need for *re-assertion* when doing retraction.

$$\frac{a \notin \Delta \quad \llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \Delta, a \longrightarrow \otimes \Delta' \quad \text{Quiescent}(\Delta', (\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}))}{\Delta \xrightarrow{+a}^{\mathcal{L}\mathcal{L}} \llbracket \mathcal{P} \rrbracket \Delta'} \quad (\text{Infer})$$

$$\frac{a \in \Delta \quad \llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \Delta, \tilde{a} \longrightarrow \otimes \Delta' \quad \text{Quiescent}(\Delta', (\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}))}{\Delta \xrightarrow{-a}^{\mathcal{L}\mathcal{L}} \llbracket \mathcal{P} \rrbracket \Delta'} \quad (\text{Retract})$$

Figure 13: Assertion and Retraction, in $\mathit{llD}^{\text{Full}}$

Quiescence in Cycles: One important issue to consider for the $\mathit{llD}^{\text{Full}}$ interpretation, is whether quiescence can be achieved for Datalog programs with cycle. Figure 16 shows one such example modified from the graph example of Figure 6, along with its linear logic interpretation (we omit the absorption rules for brevity). This example attempts to generate reflexive paths $P(x, y)$ from base facts $E(x, y)$ that represent edges. Figure 16 also illustrates an $LV^{\text{obs-}}$ derivation which represents the assertion of $E(4, 5)$ which in fact, achieves quiescence: We infer the facts $P(4, 5)$ and $P(5, 4)$ through inference rule instances $\mathcal{I}_1^{(4,5)}$ and $\mathcal{I}_2^{(4,5)}$, as well as an additional copy of $P(4, 5)$ through $\mathcal{I}_2^{(5,4)}$, before we reach quiescence. We omit fact witnesses and retraction rules as they are inessential in this discussion. In general, as long as Datalog clauses are well-formed and the set Σ of all possible constants is finite, the $\mathit{llD}^{\text{Full}}$ interpretation of Datalog programs will reach quiescence under $LV^{\text{obs-}}$ derivations. This is because the possible combinations of rule instances are finite, and since for each rule instance we can apply it at most once in a derivation, all $LV^{\text{obs-}}$ derivations of well-formed Datalog programs reach quiescence.

Re-assertion during Retraction: In some cases, re-assertion of facts occurs during retraction while working towards quiescence. Figure 17 illustrates a Datalog program which exhibits such a behavior. From either base facts $A(x)$ or $B(x)$, we can infer $C(x)$. We consider an initial state Δ where $A(2)$ and $B(2)$ has been asserted. Hence, inference rule instances $\mathcal{I}_1^{(2)}$, $\mathcal{I}_2^{(2)}$ and $\mathcal{I}_3^{(2)}$ has been applied before reaching quiescence at state Δ . Note that while we have two copies of $C(2)$ (from $\mathcal{I}_1^{(2)}$ and $\mathcal{I}_2^{(2)}$), we only have one copy of $D(2)$, since rule instance $\mathcal{I}_3^{(2)}$ is only permitted to be applied once. Next we consider the situation when $A(2)$ is retracted, triggering the application of retraction rules $\mathcal{R}_1^{(2)}$ and $\mathcal{R}_3^{(2)}$ which ultimately absorbs $A(2)$, $C(2)$ and $D(2)$. This leads us to a seemingly inconsistent state, where $B(2)$ and $C(2)$ exists, yet $D(2)$ has been completely retracted as part of the retraction closure of $A(2)$. However, quiescence is only achieved when inference rule instance $\mathcal{I}_3^{(2)}$ is re-applied, re-asserting another copy of $D(2)$. While re-assertion guarantees the re-completion of assertion closures of the $\mathit{llD}^{\text{Full}}$ interpretation, it arguably adds overheads. One of our future objectives is to refine the $\mathit{llD}^{\text{Full}}$ interpretation to overcome the need for re-assertion during retraction.

5.5 Correctness

In this section, we highlight the proofs of correspondence of the $\mathit{llD}^{\text{Full}}$ linear logic interpretation of Datalog program.

For any $LV^{\text{obs-}}$ sequents constructed from the $\mathit{llD}^{\text{Full}}$ interpretations of Datalog assertion, we can always find a normalized $LV^{\text{obs-}}$ derivation that consists only of (\mathcal{I}_L) and (obs) derivation rules. Similarly, for sequents from $\mathit{llD}^{\text{Full}}$ interpretations of Datalog retraction, we can find normalized derivations that consist strictly only of (\mathcal{I}_L) , (\mathcal{R}_L) , (\mathcal{A}_L) and (obs) derivation rules. This will simplify the proof of the correctness of our interpretation.

Lemma 6 (Normalized Derivations) *Given Datalog program \mathcal{P} , $\mathit{llD}^{\text{Full}}$ states $\llbracket F_1 \rrbracket, \Delta_1^{\mathcal{R}}, \Delta_1^{\#}$ and $\llbracket F_2 \rrbracket, \Delta_2^{\mathcal{R}}, \Delta_2^{\#}$, a $LV^{\text{obs-}}$ derivation $\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket F_1 \rrbracket, \Delta_1^{\mathcal{R}}, \Delta_1^{\#} \longrightarrow \otimes \llbracket F_1 \rrbracket \otimes \otimes \Delta_2^{\mathcal{R}} \otimes \otimes \Delta_2^{\#}$ has a proof consisting of only applications of (\mathcal{I}_L) , (\mathcal{R}_L) , (\mathcal{A}_L) and (obs) derivation rules.*

Proof: The bulk of the proof of this Lemma is derived from *permutability* results of linear logic [GPI94].

$$\text{Let } \mathcal{R}_i \equiv \tilde{a}_i \otimes r(\vec{t}) \otimes \Delta^\# \multimap \tilde{a}_i \otimes \tilde{h}$$

$$\frac{\frac{\frac{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \tilde{a}_k, \tilde{h}, \Delta \longrightarrow C}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \tilde{a}_k \otimes \tilde{h}, \Delta \longrightarrow C} (\otimes_L)}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \tilde{a}_k, \mathcal{R}_k, r(\vec{t}), \Delta^\#, \Delta \longrightarrow C} (-\circ_L)}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \tilde{a}_k, \&_{i=1}^n \mathcal{R}_i, r(\vec{t}), \Delta^\#, \Delta \longrightarrow C} (\&_L)^*} (\mathcal{R}_L)$$

for some $k \geq 1$ and $k \leq n$, which decomposes into:

$$\frac{\frac{\frac{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \tilde{a}_k, \tilde{h}, \Delta \longrightarrow C}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \tilde{a}_k \otimes \tilde{h}, \Delta \longrightarrow C} (\otimes_L)}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \tilde{a}_k, \mathcal{R}_k, r(\vec{t}), \Delta^\#, \Delta \longrightarrow C} (-\circ_L)}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \tilde{a}_k, \&_{i=1}^n \mathcal{R}_i, r(\vec{t}), \Delta^\#, \Delta \longrightarrow C} (\&_L)^*} (\mathcal{R}_L)$$

Figure 14: lD^{Full} Retraction Macro Rule

Essentially, we need to disentangle a generic derivation of any sequent $\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket F_1 \rrbracket, \Delta_1^{\mathcal{R}}, \Delta_1^\# \longrightarrow \otimes \llbracket F_1 \rrbracket \otimes \otimes \Delta_2^{\mathcal{R}} \otimes \otimes \Delta_2^\#$, into one which comprises of an orderly sequence of LV^{obs-} derivation steps that correspond to macro derivation steps (\mathcal{I}_L) , (\mathcal{R}_L) and (\mathcal{A}_L) of Figure 12, 14 and 15, and ending off with a (obs) derivation step. We do this by showing that thanks to permutability results in [GPI94], we can safely rearrange derivation steps. Most of the proof is very similar to Lemma 1, except that derivation steps may have to be pushed pass the problematic $(\&)$ rule. Thankfully, retraction rules (the only formula where $\&$ operator appears in) are designed to be *confluent*: regardless of which $\&$ sub-formula is chosen, the same effect is achieved, for instance retraction rule $\&_{i=1}^n (\tilde{a}_i \otimes r(\vec{t}) \otimes \Delta^\# \multimap \tilde{a}_i \otimes \tilde{h})$ ultimately removes witnesses $r(\vec{t})$ and $\Delta^\#$, and introduces \tilde{h} regardless of which sub-formula is chosen. As such, all other derivation steps can be safely pushed pass it. We list the details of these permutation operation for each macro rule:

- (\mathcal{I}_L) rule: The proof is similar to Lemma 1, since the inference rule of lD^{Full} interpretation (Figure 12) is of the same form as lD^{Assert} inference rules (Figure 8), only the body contains additional witness atoms and retraction rules, which involves additional applications of (\otimes_L) and does not affect the results of Lemma 1.
- (\mathcal{R}_L) rule: For a retraction rule $\&_{i=1}^n (\tilde{a}_i \otimes r(\vec{t}) \otimes \Delta^\# \multimap \tilde{a}_i \otimes \tilde{h})$, since the top most (\otimes_L) rule that operates on the body of the retraction implication rule $\tilde{a}_k \otimes \tilde{h}$ can be pushed downwards until right before the $(-\circ_L)$ derivation step that applied on the retraction implication rule. Next, we have to push all $(\&_L)$ rules applied on the retraction rule upwards to the rule application $(-\circ_L)$ derivation step. This is possible because of the confluence result in retraction rules stated above.
- (\mathcal{A}_L) rule: We can push the top most (1_L) and $(-\circ_L)$ derivation steps to just right after the $(-\circ_L^-)$ derivation step of an application of the (\mathcal{A}_L) macro rule. Similarly to the (\mathcal{I}_L) macro rule, $(clone)$ and (\forall_L) derivation steps can be safely pushed upwards until right before the main $(-\circ_L^-)$ derivation step.

Hence we have shown that all LV^{obs-} derivations has a proof consisting of only the rules (\mathcal{I}_L) , (\mathcal{R}_L) , (\mathcal{A}_L) and (obs) . \square

For every lD^{Full} derivation that models the assertion of facts (excluding retraction), we have a corresponding lD^{Assert} derivation and vice versa. Both directions depend on the fact that witness facts and retraction rules have no effects on the assertion of new facts.

Lemma 7 (Completeness and Soundness of Assertion) *Given a Datalog Program \mathcal{P} , for any fact base \mathcal{B} , multiset of facts $F \subseteq \text{herbrand}(\mathcal{P})$, multiset of rule witnesses $\Delta_{\mathcal{R}}^\#$, multiset of fact witnesses $\Delta_{\mathcal{F}}^\#$ and*

Let $\forall \vec{x}. p(\vec{x}) \otimes \tilde{p}(\vec{x}) \multimap_{p^\sharp(\vec{x})} (p^\sharp(\vec{x}) \multimap 1) \in \mathcal{A}_{\mathcal{P}}$

$$\frac{p^\sharp(\vec{t}) \notin \Delta \quad \frac{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \Delta \longrightarrow C}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; p(\vec{t}), \tilde{p}(\vec{t}), \Delta \longrightarrow C} (\mathcal{A}_L)}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; p(\vec{t}), \tilde{p}(\vec{t}), \Delta \longrightarrow C}$$

which decomposes into:

$$\frac{\frac{\frac{\frac{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \Delta \longrightarrow C}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; 1, \Delta \longrightarrow C} (1_L)}{p^\sharp(\vec{t}) \notin \Delta \quad \frac{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; (p^\sharp(\vec{t}) \multimap 1), p^\sharp(\vec{t}), \Delta \longrightarrow C}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; (p(\vec{t}) \otimes \tilde{p}(\vec{t}) \multimap_{p^\sharp(\vec{t})} (p^\sharp(\vec{t}) \multimap 1)), p(\vec{t}), \tilde{p}(\vec{t}), \Delta \longrightarrow C} (\multimap_L)}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; (p(\vec{t}) \otimes \tilde{p}(\vec{t}) \multimap_{p^\sharp(\vec{t})} (p^\sharp(\vec{t}) \multimap 1)), p(\vec{t}), \tilde{p}(\vec{t}), \Delta \longrightarrow C} (\multimap_L^-)}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \mathcal{A}_{\mathcal{P}}, p(\vec{t}), \tilde{p}(\vec{t}), \Delta \longrightarrow C} (\forall_L)^*}{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; p(\vec{t}), \tilde{p}(\vec{t}), \Delta \longrightarrow C} (clone)$$

Figure 15: llD^{Full} Absorption Macro Rule

multiset of retraction rules $\Delta^{\mathcal{R}}$,

$$\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket \mathcal{B} \rrbracket \longrightarrow \otimes \llbracket F \rrbracket \otimes \otimes \Delta^{\mathcal{R}} \otimes \otimes \Delta_R^\sharp \otimes \otimes \Delta_F^\sharp \quad \text{if and only if} \quad \llbracket \mathcal{P} \rrbracket; \llbracket \mathcal{B} \rrbracket \longrightarrow \otimes \llbracket F \rrbracket \otimes \otimes \Delta_R^\sharp$$

Proof:

- (\Rightarrow) For any llD^{Full} derivation, we safely omit witnesses and retraction rules to obtain a llD^{Assert} derivation. This is because inference rules in both interpretations are the same, except that llD^{Full} inference rules have additional witnesses and retraction rules as consequents, which do not derive new facts.
- (\Leftarrow) We show that any llD^{Assert} normalized derivation has a corresponding llD^{Full} normalized derivation, i.e. in general, if $\llbracket \mathcal{P} \rrbracket; \llbracket F_1 \rrbracket, \Delta_{R_1}^\sharp \longrightarrow \otimes \llbracket F_2 \rrbracket \otimes \otimes \Delta_{R_2}^\sharp$ then $\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket F_1 \rrbracket, \Delta_{R_1}^{\mathcal{R}}, \Delta_{F_1}^\sharp, \Delta_{R_1}^\sharp \longrightarrow \otimes \llbracket F_2 \rrbracket \otimes \otimes \Delta_2^{\mathcal{R}} \otimes \otimes \Delta_{F_2}^\sharp \otimes \otimes \Delta_{R_2}^\sharp$ for some $\Delta_1^{\mathcal{R}}, \Delta_2^{\mathcal{R}}, \Delta_{F_1}^\sharp$ and $\Delta_{F_2}^\sharp$. We inductively define the llD^{Assert} derivation as follows: Base structural derivations steps are built from (*obs*) rule, while inductive structural derivation steps are built from (\mathcal{I}_L) macro rule of the llD^{Assert} interpretation. The inductive proof proceeds as follows:

- **Base structural rule:** We consider the application of the (*obs*) rule. Hence for $\llbracket \mathcal{P} \rrbracket; \llbracket F_1 \rrbracket, \Delta_{R_1}^\sharp \longrightarrow \otimes \llbracket F_2 \rrbracket \otimes \otimes \Delta_{R_2}^\sharp$ to result from the (*obs*) rule, we must have $\llbracket F_1 \rrbracket \equiv \llbracket F_2 \rrbracket$ and $\llbracket \Delta_{R_1} \rrbracket \equiv \llbracket \Delta_{R_2} \rrbracket$. We show that there exist some $\Delta_1^{\mathcal{R}}, \Delta_2^{\mathcal{R}}, \Delta_{F_1}^\sharp$ and $\Delta_{F_2}^\sharp$ such that the llD^{Full} sequent $\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket F_1 \rrbracket, \Delta_1^{\mathcal{R}}, \Delta_{F_1}^\sharp, \Delta_{R_1}^\sharp \longrightarrow \otimes \llbracket F_2 \rrbracket \otimes \otimes \Delta_2^{\mathcal{R}} \otimes \otimes \Delta_{F_2}^\sharp \otimes \otimes \Delta_{R_2}^\sharp$ has a derivation as well. For that we simply choose $\Delta_1^{\mathcal{R}} \equiv \Delta_2^{\mathcal{R}}$ and $\Delta_{F_1}^\sharp \equiv \Delta_{F_2}^\sharp$, hence the llD^{Full} sequent in question applies to the (*obs*) rule as well.
- **Inductive structural rule:** We consider the application of the llD^{Assert} (\mathcal{I}_L) macro derivation step. Induction hypothesis, we assume that if we have $\llbracket \mathcal{P} \rrbracket; [\vec{t}/\vec{x}][h], [\vec{t}/\vec{x}][B], \llbracket B' \rrbracket, \Delta_R^\sharp, r(\vec{t}) \longrightarrow \otimes \llbracket F_2 \rrbracket \otimes \otimes \Delta_{R_2}^\sharp$ such that $(\forall \vec{x}. \otimes B \multimap_{r(\vec{x})} h \otimes \otimes B) \in \llbracket \mathcal{P} \rrbracket$, then for some $\Delta_1^{\mathcal{R}}, \Delta_2^{\mathcal{R}}, \Delta_{F_1}^\sharp$ and $\Delta_{F_2}^\sharp$, we have $\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; [\vec{t}/\vec{x}][h], [\vec{t}/\vec{x}][B], \llbracket B' \rrbracket, \Delta_1^{\mathcal{R}}, \Delta_{R_1}^\sharp, \Delta_{F_1}^\sharp, r(\vec{t}) \longrightarrow \otimes \llbracket F_2 \rrbracket \otimes \otimes \Delta_2^{\mathcal{R}} \otimes \otimes \Delta_{R_2}^\sharp \otimes \otimes \Delta_{F_2}^\sharp$. For the llD^{Assert} sequent, we apply a llD^{Assert} (\mathcal{I}_L) macro derivation step to obtain $\llbracket \mathcal{P} \rrbracket; [\vec{t}/\vec{x}][B], \llbracket B' \rrbracket, \Delta_R^\sharp \longrightarrow \otimes \llbracket F_2 \rrbracket \otimes \otimes \Delta_{R_2}^\sharp$. We now need to show that $\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; [\vec{t}/\vec{x}][B], \llbracket B' \rrbracket, \Delta_3^{\mathcal{R}}, \Delta_{R_1}^\sharp, \Delta_{F_3}^\sharp \longrightarrow \otimes \llbracket F_2 \rrbracket \otimes \otimes \Delta_4^{\mathcal{R}} \otimes \otimes \Delta_{R_2}^\sharp \otimes \otimes \Delta_{F_4}^\sharp$ for some $\Delta_3^{\mathcal{R}}$,

$$\begin{aligned}
 \mathcal{P} &= I_1 : \forall x, y. E(x, y) \supset P(x, y) \quad , \quad I_2 : \forall x, y. P(x, y) \supset P(y, x) \\
 \llbracket \mathcal{P} \rrbracket &= \forall x, y. \mathcal{I}_1^{(x,y)}, \forall x, y. \mathcal{I}_2^{(x,y)} \\
 \mathcal{I}_1^{(x,y)} &= E(x, y) \multimap_{I_1^\#(x,y)} P(x, y) \otimes E(x, y) \otimes E^\#(x, y) \otimes \mathcal{R}_1^{(x,y)} \\
 \mathcal{R}_1^{(x,y)} &= \tilde{E}(x, y) \otimes I_1^\#(x, y) \otimes E^\#(x, y) \multimap \tilde{P}(x, y) \otimes \tilde{E}(x, y) \\
 \mathcal{I}_2^{(x,y)} &= P(x, y) \multimap_{I_2^\#(x,y)} P(y, x) \otimes P(x, y) \otimes P^\#(x, y) \otimes \mathcal{R}_2^{(x,y)} \\
 \mathcal{R}_2^{(x,y)} &= \tilde{P}(x, y) \otimes I_2^\#(x, y) \otimes P^\#(x, y) \multimap \tilde{P}(y, x) \otimes \tilde{P}(x, y)
 \end{aligned}$$

Assertion of $E(4, 5)$: $\emptyset \xrightarrow{+E(4,5)}_{\mathcal{P}} \Delta$ where $\Delta = \{P(4, 5), P(5, 4), P(4, 5), E(4, 5), I_1^\#(4, 5), I_2^\#(4, 5), I_2^\#(5, 4), \dots\}$

$$\begin{array}{c}
 \overline{\overline{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \boxed{P(4, 5), P(5, 4), P(4, 5), E(4, 5), I_1^\#(4, 5), I_2^\#(4, 5), I_2^\#(5, 4), \dots} \longrightarrow \boxed{\otimes \Delta}} \quad (\text{obs}) \\
 \mathcal{I}_2^{(5,4)} \overline{\overline{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \boxed{P(5, 4)}, P(4, 5), E(4, 5), I_1^\#(4, 5), I_2^\#(4, 5), \dots} \longrightarrow \otimes \Delta}} \quad (\mathcal{I}_L) \\
 \mathcal{I}_2^{(4,5)} \overline{\overline{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \boxed{P(4, 5)}, E(4, 5), I_1^\#(4, 5), \dots} \longrightarrow \otimes \Delta}} \quad (\mathcal{I}_L) \\
 \mathcal{I}_1^{(4,5)} \overline{\overline{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \boxed{P(4, 5)}, E(4, 5), I_1^\#(4, 5), \dots} \longrightarrow \otimes \Delta}} \quad (\mathcal{I}_L) \\
 \overline{\overline{\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \boxed{E(4, 5)}} \longrightarrow \otimes \Delta}}
 \end{array}$$

Figure 16: Example of Quiescence in the Presence of Cycles

$\Delta_4^{\mathcal{R}}$, $\Delta_{F_3}^\#$ and $\Delta_{F_4}^\#$. To do this, we apply the llD^{Full} (\mathcal{I}_L) macro derivation step corresponding to the llD^{Full} inference rule instance $(\forall \vec{x}. \otimes B \multimap_{r(\vec{x})} h \otimes \otimes B \otimes \otimes \Delta^\# \otimes \mathcal{R}) \in \llbracket \mathcal{P} \rrbracket$. We obtain a derivation step with premise $\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; [\vec{t}/\vec{x}] \llbracket h \rrbracket, [\vec{t}/\vec{x}] \llbracket B \rrbracket, \llbracket B' \rrbracket, [\vec{t}/\vec{x}], \mathcal{R}, \Delta_3^{\mathcal{R}}, r(\vec{t}), \Delta_{R_3}^\#, [\vec{t}/\vec{x}] \Delta^\#, \Delta_{F_3}^\# \longrightarrow \otimes \llbracket F_2 \rrbracket \otimes \otimes \Delta_4^{\mathcal{R}} \otimes \otimes \Delta_{R_2}^\# \otimes \otimes \Delta_{F_4}^\#$ where if we take $\Delta_1^{\mathcal{R}} \equiv \mathcal{R}, \Delta_3^{\mathcal{R}}, \Delta_{F_1}^\# \equiv [\vec{t}/\vec{x}] \Delta^\#, \Delta_{F_3}^\#, \Delta_2^{\mathcal{R}} \equiv \Delta_4^{\mathcal{R}}$ and $\Delta_{F_2}^\# \equiv \Delta_{F_4}^\#$, is exactly equivalent to our induction hypothesis.

Hence we have proven the completeness and soundness of assertions. \square

We show that for derivations that do not consist of anti-predicates (retraction of facts), we have a quiescent state in llD^{Full} corresponds to a quiescent state in llD^{Assert} .

Lemma 8 (Correspondence of Quiescence) *Let \mathcal{P} be a Datalog program and \mathcal{B} a collection of base facts, for any $F, \Delta^{\mathcal{R}}, \Delta^\#$ such that $\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket \mathcal{B} \rrbracket \longrightarrow \otimes \llbracket F \rrbracket \otimes \otimes \Delta^{\mathcal{R}} \otimes \otimes \Delta^\#$,*

$$\text{Quiescent}(\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}), (\llbracket F \rrbracket, \Delta^{\mathcal{R}}, \Delta^\#) \quad \text{if and only if} \quad \text{Quiescent}(\llbracket \mathcal{P} \rrbracket, (\llbracket F \rrbracket, \Delta^\#))$$

Proof: This follows from Lemma 7. We show by contradiction:

- (\Rightarrow) Suppose that $\text{Quiescent}(\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}), (\llbracket F \rrbracket, \Delta^{\mathcal{R}}, \Delta^\#)$ but $\text{Quiescent}(\llbracket \mathcal{P} \rrbracket, (\llbracket F \rrbracket, \Delta^\#))$ does not hold. Hence there exists some $F_2, \Delta_2^{\mathcal{R}}, \Delta_2^\#$ such that $\llbracket \mathcal{P} \rrbracket; \llbracket F \rrbracket, \Delta^\# \longrightarrow \otimes \llbracket F_2 \rrbracket \otimes \otimes \Delta_2^\#$. and $F \neq F_2$ and $\Delta^\# \neq \Delta_2^\#$. By Lemma 7, we have $\llbracket \mathcal{P} \rrbracket; \llbracket F \rrbracket, \Delta^{\mathcal{R}}, \Delta^\# \longrightarrow \otimes \llbracket F_2 \rrbracket \otimes \otimes \Delta_2^{\mathcal{R}} \otimes \otimes \Delta_2^\#$, contradicting $\text{Quiescent}(\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}), (\llbracket F \rrbracket, \Delta^{\mathcal{R}}, \Delta^\#)$. Hence it must be that $\text{Quiescent}(\llbracket \mathcal{P} \rrbracket, (\llbracket F \rrbracket, \Delta^\#))$.
- (\Leftarrow) Suppose that $\text{Quiescent}(\llbracket \mathcal{P} \rrbracket, (\llbracket F \rrbracket, \Delta^\#))$ but $\text{Quiescent}(\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}), (\llbracket F \rrbracket, \Delta^{\mathcal{R}}, \Delta^\#)$ does not hold, Hence there exists some $F_2, \Delta_2^{\mathcal{R}}, \Delta_2^\#$ such that $\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket F \rrbracket, \Delta^{\mathcal{R}}, \Delta^\# \longrightarrow \otimes \llbracket F_2 \rrbracket \otimes \otimes \Delta_2^{\mathcal{R}} \otimes \otimes \Delta_2^\#$. and $F \neq F_2$ and $\Delta^\# \neq \Delta_2^\#$. By Lemma 7, we have $\llbracket \mathcal{P} \rrbracket; \llbracket F \rrbracket, \Delta^\# \longrightarrow \otimes \llbracket F_2 \rrbracket \otimes \otimes \Delta_2^\#$, contradicting $\text{Quiescent}(\llbracket \mathcal{P} \rrbracket, (\llbracket F \rrbracket, \Delta^\#))$. Hence it must be that $\text{Quiescent}(\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}), (\llbracket F \rrbracket, \Delta^{\mathcal{R}}, \Delta^\#)$.

2. The proof proceeds by contradiction. Suppose that we have $\llbracket \mathcal{P} \rrbracket; \llbracket \mathcal{B} \rrbracket \longrightarrow \otimes \llbracket \Delta \rrbracket$ and $Quiescent((\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}), \Delta)$, but for any $\Delta^{\mathcal{R}}, \Delta^{\sharp}$ such that $\Delta = [F], \Delta^{\mathcal{R}}, \Delta^{\sharp}$, $set(F) \neq \mathcal{P}(\mathcal{B})$. However by Lemma 7 we can always find a llD^{Assert} derivation $\llbracket \mathcal{P} \rrbracket; \llbracket \mathcal{B} \rrbracket \longrightarrow \otimes \llbracket \Delta \rrbracket$ and Theorem 5 states that if $Quiescent(\llbracket \mathcal{P} \rrbracket, \Delta)$ then $set(F) = \mathcal{P}(\mathcal{B})$, contradicting our assumption.

Hence we have shown the correctness of assertions. \square

While (\mathcal{I}_L) macro derivation rules in llD^{Full} corresponds to llD^{Assert} , applications of the (\mathcal{I}_L) macro rules in llD^{Full} that models assertion of facts provides certain properties that facilitates the retraction of facts. Lemma 10 states these properties of assertion.

Lemma 10 (Properties of Assertion) *Given a Datalog Program \mathcal{P} , for any llD^{Full} state $\Delta, \Delta^{\mathcal{R}}, \Delta^{\sharp}$, and for some Δ' , if $\llbracket \mathcal{P} \rrbracket; \mathcal{A}_{\mathcal{P}}; \Delta, \Delta^{\mathcal{R}}, \Delta^{\sharp} \longrightarrow \otimes \Delta'$ and $Quiescent((\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}), \Delta')$ we have the following properties:*

1. For each $a \in \Delta'$ such that $a \in herbrand_I(\mathcal{P})$ (i.e. a is inferred), there must exist some retraction rule $\&_{i=1}^n(\tilde{a}_i \otimes r(\vec{t}) \otimes \otimes B^{\sharp} \multimap \tilde{a}_i \otimes \tilde{h}) \in \Delta'$ such that $a = h$.
2. For each retraction rule $\&_{i=1}^n(\tilde{a}_i \otimes r(\vec{t}) \otimes \otimes B^{\sharp} \multimap \tilde{a}_i \otimes \tilde{h}) \in \Delta'$, we must have for each \tilde{a}_i such that equal and opposite $a_i \in \Delta'$.
3. For each retraction rule $\mathcal{R} = \&_{i=1}^n(\tilde{a}_i \otimes r(\vec{t}) \otimes \otimes B^{\sharp} \multimap \tilde{a}_i \otimes \tilde{h})$, $\mathcal{R} \in \Delta'$ if and only if $B^{\sharp} \in \Delta'$ and $r(\vec{t}) \in \Delta'$ as well.
4. There always exists some Δ' such that $Quiescent((\mathcal{P}, \mathcal{A}_{\mathcal{P}}), \Delta')$.

Proof: The proof of these properties relies entirely on the definition of inference rules in $\llbracket \mathcal{P} \rrbracket$ (Figure 12). Specifically they are of the form: $\forall \vec{x}. \otimes B \multimap_{r(\vec{x})} h \otimes \otimes B \otimes \otimes B^{\sharp} \otimes \mathcal{R}$ such that $\mathcal{R} = \&_{i=1}^n(\tilde{a}_i \otimes r(\vec{x}) \otimes \otimes B^{\sharp} \multimap \tilde{a}_i \otimes \tilde{h})$. Details of the proof for each property proceeds as follows:

1. For any fact $a \in \Delta'$, if a is not a base fact then it must be the product of an application of an inference rule instance from $\llbracket \mathcal{P} \rrbracket$. As such, by the definition of inference rules, for every fact $h = a$ inferred, we have a corresponding retraction rule $\mathcal{R} = \&_{i=1}^n(\tilde{a}_i \otimes r(\vec{x}) \otimes \otimes B^{\sharp} \multimap \tilde{a}_i \otimes \tilde{a})$ as a by-product. Hence we have property (1).
2. Since retraction rules are only introduced by inference rules, by the definition of inference rules, each inference rule $\forall \vec{x}. \otimes B \multimap_{r(\vec{x})} h \otimes \otimes B \otimes \otimes B^{\sharp} \otimes \mathcal{R}$ introduces retraction rule \mathcal{R} such that each \tilde{a}_i that appears in each implication of the retraction rule is the equal and opposite of a unique atom in B . Hence we have property (2).
3. Similarly, since retraction rules are only introduced by inference rules, by the definition of inference rules, each retraction rule $\mathcal{R} = \&_{i=1}^n(\tilde{a}_i \otimes r(\vec{x}) \otimes \otimes B^{\sharp} \multimap \tilde{a}_i \otimes \tilde{h})$ is introduced explicitly with B^{\sharp} as accompanying consequents, while the formulation of (\multimap_L^-) rule dictates the introduction of inhibiting atom $r(\vec{t})$. Hence we have property (3).
4. We argue that all inference rules are guarded implication rules inhibited by a unique rule instance $r(\vec{t})$. Since all well-formed Datalog program we consider have finite number of clauses and finite constants, there are finite number of rule instances $r(\vec{t})$. As such, we can apply finite number of (\mathcal{I}_L) macro derivation steps and hence all derivations modeling assertion will reach quiescence.

Hence we have proven these properties of assertion. \square

Like inference rules, retraction rules and absorption rules are defined in a manner which provides some important properties that help to guarantee the correctness of retraction. Lemma 11 formally asserts these properties.

Lemma 11 (Properties of Retraction) *Given a Datalog Program \mathcal{P} , for any llD^{Full} state that represent an intermediate state of retraction $\Delta, \Delta^{\mathcal{R}}, \Delta^{\sharp}$ such that $\Delta = \llbracket F \rrbracket, \llbracket a \rrbracket, \llbracket F' \rrbracket, \llbracket \tilde{F}' \rrbracket, \llbracket \tilde{a} \rrbracket$ and $F, F' \subseteq \text{herbrand}(\mathcal{P})$ and $a \in \text{herbrand}(\mathcal{P})$, and for some Δ' , if $\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \Delta, \Delta^{\mathcal{R}}, \Delta^{\sharp} \longrightarrow \otimes \Delta'$, we have the following properties:*

1. For each retraction rule $\mathcal{R} = \&_{i=1}^n (\tilde{a}_i \otimes r(\vec{t}) \otimes \otimes B^{\sharp} \multimap \tilde{a}_i \otimes \tilde{h})$, $\mathcal{R} \in \Delta'$ if and only if $B^{\sharp} \in \Delta'$ and $r(\vec{t}) \in \Delta'$.
2. If $\tilde{a} \notin \Delta'$ then there not exists any retraction rule $\&_{i=1}^n (\tilde{a}_i \otimes r(\vec{t}) \otimes \otimes B^{\sharp} \multimap \tilde{a}_i \otimes \tilde{h}) \in \Delta'$, such that $\tilde{a} = \tilde{a}_i$ for some $i \geq 1$ and $i \leq n$.
3. There always exists some Δ' such that $\text{Quiescent}((\mathcal{P}, \mathcal{A}_{\mathcal{P}}), \Delta')$.
4. If $\text{Quiescent}((\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}), \Delta')$ then there not exists any $\tilde{b} \in \Delta'$.

Proof: The proof of these properties relies on guarantees provided by assertion (Lemma 10) and the definition of retraction rules and absorption rules (Figure 14 and 15 respectively). Details of the proof for each property proceeds as follows:

1. Since all retraction rules are only introduced by inference rules and that state $\Delta, \Delta^{\mathcal{R}}, \Delta^{\sharp}$ is a reachable state, property 3 of Lemma 10 guarantees that initially, each retraction rule in $\&_{i=1}^n (\tilde{a}_i \otimes r(\vec{t}) \otimes \otimes B^{\sharp} \multimap \tilde{a}_i \otimes \tilde{h}) \in \Delta^{\mathcal{R}}$ has a complete set of witnesses in Δ^{\sharp} (i.e. $r(\vec{t}), B^{\sharp} \subseteq \Delta^{\sharp}$). During retraction, each application of a retraction \mathcal{R} consumes the retraction rule itself as well as the exact set of witnesses B^{\sharp} and $r(\vec{t})$ (and nothing more) which was introduced together with \mathcal{R} in the same inference rule instance. Hence, the property that any remaining retraction rules has its full set of witnesses in the context is preserved throughout intermediate retraction states.
2. If $\tilde{a} \notin \Delta'$, absorption rule $a \otimes \tilde{a} \multimap_{a^{\sharp}} (a^{\sharp} \multimap 1)$ must have been applied, since it is the only way to consume an \tilde{a} . In order for the absorption rule to apply, the context must not have any occurrences of a^{\sharp} , and the only way to have achieved this is that all retraction rules $\&_{i=1}^n (\tilde{a}_i \otimes r(\vec{t}) \otimes \otimes B^{\sharp} \multimap \tilde{a}_i \otimes \tilde{h})$ such that $\tilde{a} = \tilde{a}_i$ for some $i \geq 1$ and $i \leq n$ and $a^{\sharp} \in B^{\sharp}$, were applied to consume all instance of a^{\sharp} and hence allowing the absorption rule to be applied.
3. We argue that retraction rules are linear, hence there always exists a derivation which involves exhaustive and finite application of the retraction macro derivation (\mathcal{R}_L). Since we can apply finite number of (\mathcal{R}_L) derivations, we can generate finite numbers of retraction atoms \tilde{a} . As such, we can apply finite number of absorption (\mathcal{A}_L) macro derivation steps as well. Yet in retraction, it is possible that we need to reassert facts via (\mathcal{I}_L). By property (4) of Lemma 10 we can conclude that there are also finite number of reassertion steps. Hence we can always reach quiescence.
4. From (1), we argue that all retraction rules at any intermediate reachable states always have their accompanying witnesses B^{\sharp} and $r(\vec{t})$. As such we can exhaustively apply retraction rules to remove all witnesses inhibiting absorption rules from applying, after which all $\tilde{b} \in \Delta$ will be consumed. Hence if Δ' is quiescent, we must have removed all retraction atom \tilde{b} .

Hence we have proven these properties of retraction. \square

Inserting a set of retraction atoms $\tilde{\mathcal{B}}$ in a llD^{Full} state $\llbracket F \rrbracket, \llbracket \mathcal{B} \rrbracket, \Delta^{\mathcal{R}}, \Delta^{\sharp}$ will ultimately result in the derivation of the state Δ_1 which excludes $\llbracket \mathcal{B} \rrbracket$. Note that we disallow the insertion of a retraction atom \tilde{b} if b does not exist in the state. Theorem 12 states that (1) we can always find a quiescent state in retraction of some b and this state corresponds to the state with b and all its consequents (facts, retraction rules and witnesses) removed. (2) This quiescent state is unique.

Theorem 12 (Correctness of Retraction) *Given a Datalog Program \mathcal{P} and a base fact $b \in \text{herbrand}_B(\mathcal{P})$ and llD^{Full} states Δ_1 ,*

1. There exists some state $\llbracket F \rrbracket, \Delta^{\mathcal{R}}, \Delta^{\sharp}$, and Δ_1, Δ_2 , such that

$$\begin{aligned} & \llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket F \rrbracket, \llbracket b \rrbracket, \Delta^{\mathcal{R}}, \Delta^{\sharp} \longrightarrow \otimes \Delta_1 \otimes \otimes \Delta_2 \text{ and } \text{Quiescent}(\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}), (\Delta_1, \Delta_2) \text{ and} \\ & \llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket F \rrbracket, \Delta^{\mathcal{R}}, \Delta^{\sharp} \longrightarrow \otimes \Delta_1 \text{ and } \text{Quiescent}(\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}), \Delta_1 \text{ and} \\ & \llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket F \rrbracket, \llbracket b \rrbracket, \Delta^{\mathcal{R}}, \Delta^{\sharp}, \llbracket \tilde{b} \rrbracket \longrightarrow \otimes \Delta_1 \end{aligned}$$

2. For all states $\llbracket F \rrbracket, \Delta^{\mathcal{R}}, \Delta^{\sharp}$, there exists some Δ_1, Δ_2 , such that

$$\begin{aligned} \text{if } & \llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket F \rrbracket, \llbracket b \rrbracket, \Delta^{\mathcal{R}}, \Delta^{\sharp}, \llbracket \tilde{b} \rrbracket \longrightarrow \otimes \Delta_1 \text{ and } \text{Quiescent}(\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}), \Delta_1 \\ \text{then } & \llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket F \rrbracket, \llbracket b \rrbracket, \Delta^{\mathcal{R}}, \Delta^{\sharp} \longrightarrow \otimes \Delta_1 \otimes \otimes \Delta_2 \text{ and } \text{Quiescent}(\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}), (\Delta_1, \Delta_2) \\ & \text{and } \llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket F \rrbracket, \Delta^{\mathcal{R}}, \Delta^{\sharp} \longrightarrow \otimes \Delta_1 \end{aligned}$$

Proof: First, by relying on property 4 of Lemma 10, we can conclude that there exists an assertion state Δ_1, Δ_2 such that $\text{Quiescent}(\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}), (\Delta_1, \Delta_2)$ and by property 3 of Lemma 11 we can conclude that there exists a retraction state Δ_1 such that $\text{Quiescent}(\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}), \Delta_1$. We now need to show that there exist such quiescent assertion and retraction states where $\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket F \rrbracket, \llbracket b \rrbracket, \Delta^{\mathcal{R}}, \Delta^{\sharp}, \llbracket \tilde{b} \rrbracket \longrightarrow \otimes \Delta_1$. We assume that $\llbracket \mathcal{P} \rrbracket, \mathcal{A}_{\mathcal{P}}; \llbracket F \rrbracket, \llbracket b \rrbracket, \Delta^{\mathcal{R}}, \Delta^{\sharp}, \llbracket \tilde{b} \rrbracket \longrightarrow \otimes \Delta_3$ for some Δ_3 , then proceed to show that if $\Delta_1 = \Delta_3$, our assumption still holds. Property 4 of Lemma 11 states that quiescent retraction states have no retraction atoms, hence we know for sure Δ_3 contains no retraction atoms. By property 2 of Lemma 11, since $\tilde{b} \notin \Delta_3$, we can safely assert that no retraction rule $\mathcal{R} = \&_{i=1}^n(\tilde{a}_i \otimes r(\vec{t}) \otimes \otimes B^{\sharp} \multimap \tilde{a}_i \otimes \tilde{h}) \in \Delta^{\mathcal{R}}$ such that $\tilde{a}_i = \tilde{b}$ for some i is in Δ_3 as well. With all such retraction rule removed, property 1 of Lemma 11 further states that corresponding witnesses of these retraction rules are not present in Δ_3 either. Finally, we know that to remove all retraction atoms \tilde{a} from the context, Δ_3 was the result of several (\mathcal{A}_L) macro derivation steps that removed b and its consequent facts. Hence, we can conclude that Δ_3 can be intermediately obtained by removing b and all its consequents (facts, retraction rules and witnesses) from $\llbracket F \rrbracket, \Delta^{\mathcal{R}}, \Delta^{\sharp}$, yielding some Δ_4 which contains no retraction atoms. From Δ_4 we can apply (\mathcal{I}_L) macro steps to reassert facts common to b 's consequents and by Theorem 5 we can guarantee that we can find a unique quiescent state Δ_3 since this quiescent state must be unique, hence $\Delta_1 = \Delta_3$. \square

Corollary 13 formally asserts the bridge between our lLD^{Full} interpretation of Datalog in linear logic LV^{obs-} and the traditional interpretation (Figure 1 and 13).

Corollary 13 (Correctness) *Given a Datalog Program \mathcal{P} , for reachable states $\Delta_1, \Delta_1^{\mathcal{R}}, \Delta_1^{\sharp}$ and $\Delta_2, \Delta_2^{\mathcal{R}}, \Delta_2^{\sharp}$ such that $\Delta_1 = \llbracket \mathcal{P}(\mathcal{B}_1) \rrbracket$ and $\Delta_2 = \llbracket \mathcal{P}(\mathcal{B}_2) \rrbracket$, then we have the following:*

$$(\Delta_1, \Delta_1^{\mathcal{R}}, \Delta_1^{\sharp}) \xrightarrow{\alpha}_{\llbracket \mathcal{P} \rrbracket}^{LL} (\Delta_2, \Delta_2^{\mathcal{R}}, \Delta_2^{\sharp}) \text{ iff } \mathcal{P}(\mathcal{B}_1) \xrightarrow{\alpha}_{\mathcal{P}} \mathcal{P}(\mathcal{B}_2)$$

Note if $\alpha = +a$ we have $\mathcal{B}_2 = \mathcal{B}_1, a$. If $\alpha = -a$ we have $\mathcal{B}_1 = \mathcal{B}_2, a$.

Proof: Follows from Theorems 9 and 12. \square

6 Related Work

Early work on Datalog focused on efficient support for querying the logical consequences of a static fact base [CGT89] although some of it also supported assertions. Examples include the semi-naive strategy [BR87], QSQ [Vie86], and magic sets [BMSU86]. Recently, there has been a trend in the parallel and distributed programming research community to develop high-level declarative programming languages which are based on (or inspired by) Datalog. For instance, NetLog [GW10] and Network Datalog ($DNLog$) [LCG⁺06] target the development of networking protocols, Meld [ARLG⁺09] for programming of

large ensembles of embedded robotic devices, and the system in [LM03] enables distributed user authentication. All of these works to varying extents, credit the declarative nature of Datalog as a main motivation, mitigating the notorious complexities of parallel and distributed programming. These new application areas are characterized by frequent assertions *and* retractions. For example, the Nexcel language of [Cer07] uses a dialect of Datalog to compute derived data in a spreadsheet as the user types.

By defining formal operational semantics and effective algorithms, these efforts contributed efficient implementations to efficiently compute inferential closures incrementally as facts are dynamically added or removed [ARLG⁺09, CARG⁺12, GMS93, LCG⁺06, NJLS11]. However we are unaware of any work that bridges these algorithms back to formal logic. As such, our efforts here complement these developments.

7 Conclusions and Future Work

In this report, we have given a logical interpretation of Datalog in linear logic, which captures assertion and retraction of facts within the confines of logic. It directly maps assertion and retraction steps to derivations in linear logic, thereby functioning not only as an abstract specification, but as an operational blueprint for practical implementations of Datalog systems. Indeed, we prove the llD^{Full} interpretation corresponds with the traditional logic interpretation of Datalog programs.

In the future, we plan to implement a Datalog system based on llD^{Full} and exploiting ideas in [CS09]. We believe that this can be a foundationally strong approach for deriving expressive and powerful concurrent logic programming languages to be applied to highly parallel systems such as Claytronics [ARLG⁺09]. We also wish to further investigate re-assertion during retraction, briefly discussed in Section 5.4. Specifically, we are interested in exploring the possibility of deriving a linear logic interpretation of Datalog programs that removes the need for re-assertion during fact retraction, thereby providing a more streamlined and efficient operational semantics for practical implementation. We intend to explore the use of other variants of linear logic, in the attempt to streamline our interpretation of Datalog. For instance, representing Datalog facts as consumable subexponentials [NM09] and using polarized presentations of linear logic.

References

- [ARLG⁺09] M. P. Ashley-Rollman, P. Lee, S. C. Goldstein, P. Pillai, and J. D. Campbell. A language for large ensembles of independently executing nodes. In *Proc. of ICLP'09*, volume 5649/2009, pages 265–280. Springer-Verlag, 2009.
- [BMSU86] F. Bancilhon, D. Maier, Y. Sagiv, and J. D. Ullman. Magic sets and other strange ways to implement logic programs (extended abstract). In *Proc. of PODS'86*, pages 1–15. ACM, 1986.
- [BR87] I. Balbin and K. Ramamohanarao. A generalization of the differential approach to recursive query evaluation. *J. Log. Program.*, 4(3):259–262, 1987.
- [CARG⁺12] F. Cruz, M. P. Ashley-Rollman, S. C. Goldstein, Ricardo Rocha, and F. Pfenning. Bottom-up logic programming for multicores. In Vítor Santos Costa, editor, *Proc. of DAMP 2012 - Short Papers*. ACM Digital Library, January 2012.
- [Cer07] I. Cervesato. NEXCEL, a Deductive Spreadsheet. *The Knowledge Engineering Review*, 22:221–236, 2007.
- [CGT89] S. Ceri, G. Gottlob, and L. Tanca. What you always wanted to know about datalog (and never dared to ask). *IEEE Trans. on Knowl. and Data Eng.*, 1(1):146–166, March 1989.
- [CS09] I. Cervesato and A. Scedrov. Relating state-based and process-based concurrency through linear logic. *Inf. Comput.*, 207(10):1044–1077, 2009.

REFERENCES

- [Gir87] J. Y. Girard. Linear logic. *Theor. Comput. Sci.*, 50:1–102, 1987.
- [GMN84] H. Gallaire, J. Minker, and J. M. Nicolas. Logic and databases: A deductive approach. *ACM Comput. Surv.*, 16(2):153–185, June 1984.
- [GMS93] A. Gupta, I. S. Mumick, and V. S. Subrahmanian. Maintaining views incrementally. In *Proc. of SIGMOD'93*, pages 157–166. ACM, 1993.
- [GPI94] D. Galmiche, G. Perrier, and C. C. Inria. On proof normalization in linear logic. *Theor. Comput. Sci.*, 135:135–1, 1994.
- [GW10] S. Grumbach and F. Wang. Netlog, a rule-based language for distributed programming. In *Proc. of PADL'10*, volume 5937/2010, pages 88–103. Springer-Verlag, 2010.
- [LCG⁺06] B. T. Loo, T. Condie, M. Garofalakis, D. E. Gay, J. M. Hellerstein, P. Maniatis, R. Ramakrishnan, T. Roscoe, and I. Stoica. Declarative networking: language, execution and optimization. In *Proc. of SIGMOD '06*, pages 97–108. ACM, 2006.
- [LM03] Ninghui Li and John C. Mitchell. Datalog with constraints: A foundation for trust-management languages. In *Proc. of PADL'03*, volume 2562/2003, pages 58–73. Springer-Verlag, 2003.
- [MNPS91] D. Miller, G. Nadathur, F. Pfenning, and A. Scedrov. Uniform proofs as a foundation for logic programming. *Proc. of APAL'91*, 51:125–157, 1991.
- [NJLS11] V. Nigam, L. Jia, B. T. Loo, and A. Scedrov. Maintaining distributed logic programs incrementally. In *Proc. of PPDP'11*, pages 125–136. ACM, 2011.
- [NM09] V. Nigam and D. Miller. Algorithmic specifications in linear logic with subexponentials. In *Proc. of PPDP'09*, pages 129–140. ACM, 2009.
- [Pfe95] F. Pfenning. Structural cut elimination. In *Proc. of LICS'95*, pages 156–166. IEEE Press, 1995.
- [Vie86] L. Vieille. Recursive axioms in deductive databases: The query/subquery approach. In *Expert Database Conf.*, pages 253–267. Benjamin Cummings, 1986.