

Ontology-Based Query Answering over Datalog-Expressible Rule Sets is Undecidable

David Carral¹, Lucas Larroque², Michaël Thomazo²

¹LIRMM, Inria, University of Montpellier, CNRS, Montpellier, France

²Inria, DI ENS, ENS, CNRS, PSL University, Paris, France

{david.carral, lucas.larroque, michael.thomazo}@inria.fr,

Abstract

Ontology-based query answering is a problem that takes as input an ontology \mathcal{R} (typically expressed by existential rules), a set \mathcal{F} of facts, and a Boolean conjunctive query (CQ) q , and asks whether $\mathcal{R}, \mathcal{F} \models q$. This problem is undecidable in general, and a widely investigated approach to tackle it in some cases is query rewriting: given some “rule query” $\langle \mathcal{R}, q \rangle$, we compute a Boolean query $q_{\mathcal{R}}$ such that, for any fact set \mathcal{F} , it holds that $\mathcal{R}, \mathcal{F} \models q$ if and only if $\mathcal{F} \models q_{\mathcal{R}}$. Insofar, previous work has mostly focused on output queries $q_{\mathcal{R}}$ expressed as union of Boolean conjunctive queries (UCQs), and an effective algorithm that computes such a query $q_{\mathcal{R}}$ whenever it exists has been proposed in the literature. However, UCQ-rewritability is not a very general notion and many real-world interesting rule queries do not admit UCQ-rewritings. This raises the question whether such a generic algorithm can be designed for a more expressive target language, such as datalog. We solve this question by the negative, by studying the difference between datalog-expressibility and datalog-rewritability. More precisely, we show that query answering under datalog-expressible rule queries is undecidable.

1 Introduction

Efficiently accessing data is an important step in many real-world applications. Ontologies have been identified as an important tool to help a user to express their information needs, allowing them to use a vocabulary they are familiar with, while enabling a system to perform automated reasoning, leading to more complete answers. Ontology-based query answering (OBQA) is a core problem therein, where a set of facts is queried while taking into account the domain knowledge expressed in an ontology. These ontologies may be expressed in a variety of formalisms, such as Description Logics or existential rules. The OBQA problem is typically framed as, given a fact set \mathcal{F} , an ontology \mathcal{R} , and a Boolean CQ q , check if $\mathcal{F}, \mathcal{R} \models q$, where \models denotes the classical first-order logic entailment.

This problem is undecidable when the ontology can range over any set of existential rules. Thus, a lot of research has focused on finding decidable and even tractable classes of rule sets; see (Mugnier and Thomazo 2014) for an introduction to these. Particularly relevant to us are classes based on the so-called *query rewriting* approach. Given an ontology \mathcal{R} and a Boolean CQ q , one computes a Boolean UCQ $q_{\mathcal{R}}$ such that for any fact set \mathcal{F} , it holds that $\mathcal{F}, \mathcal{R} \models q$ if

and only if $\mathcal{F} \models q_{\mathcal{R}}$. As most data is stored in relational databases, which have been designed to efficiently process CQs, most research has focused on rewriting the output query $q_{\mathcal{R}}$ as a UCQ. A natural question is then, given an ontology \mathcal{R} and a BCQ q , is there a UCQ-rewriting $q_{\mathcal{R}}$ for \mathcal{R} and q ? In other terms, is that true that the rule query $\langle \mathcal{R}, q \rangle$ is *UCQ-expressible*? This does not always hold, and it is actually undecidable to check whether it is the case (Baget et al. 2011). However, there exists an effective procedure algorithm that *computes* a UCQ-rewriting when given as input a UCQ-expressible rule query (König et al. 2015). In other terms, the UCQ-expressibility of every rule query (the existence of a UCQ-rewriting) in a class and UCQ-rewritability of that class (the computability of UCQ-rewritings for all rule queries in that class) are two notions that coincide, which possibly explains why they have not been introduced separately in the literature.

Syntactic conditions such as linearity (Calì, Gottlob, and Kifer 2013) or stickiness (Calì, Gottlob, and Pieris 2010) guarantee the existence of UCQ-rewritings for any BCQ. Moreover, there is also DL-Lite (Artale et al. 2009), which is a widely used Description Logic that can be translated into existential rules. However, the expressivity of these languages is too limited for many real-world ontologies. A natural task is to consider a more expressive target query language for the rewritings. It is known that considering first-order queries instead of union of conjunctive queries does not allow covering more classes of existential rules (Rossman 2008). We then focus on another classical language, namely datalog. Note that all UCQ-expressible rule queries are also datalog-expressible but the converse is not true.

As discussed in Section 5, there are many known and interesting classes for which specific datalog-rewriting algorithms have been designed. However, no generic algorithm, such as in the case of UCQ-expressibility, is known so far. The contribution of this paper is to show that, unfortunately, no such algorithm exists. This is done by proving that the problem of checking $\mathcal{R}, \mathcal{F} \models q$ under the assumption that $\langle \mathcal{R}, q \rangle$ is datalog-expressible is undecidable, contradicting the existence of a rewriting algorithm for datalog-expressible queries. We prove the result by reduction from the halting problem of Turing machines to OBQA, where the difficulty lies in ensuring that rule queries produced by the reduction are datalog-expressible.

2 Preliminaries

We assume that the reader is familiar with first-order logic and basic concepts from computability theory. We only provide a brief recap of these topics in this section.

2.1 First-Order Logic

We define Preds, Funs, Cons, and Vars to be mutually disjoint and countably infinite sets of predicates, function symbols, constants, and variables, respectively. We associate every $s \in \text{Preds} \cup \text{Funs}$ with some *arity* $\text{ari}(s) \geq 0$. For every $n \geq 0$, the sets of all n -ary predicates and all n -ary function symbols are also countably infinite. The set Terms includes Cons, Vars, and the set of all nullary function symbols; and contains $f(t_1, \dots, t_n)$ for every $n \geq 1$, every n -ary $f \in \text{Funs}$, and every $t_1, \dots, t_n \in \text{Terms}$. We write lists t_1, \dots, t_n of terms as \vec{t} and often treat these lists as sets.

An *atom* is a first-order formula of the form $P(\vec{t})$ with \vec{t} a list of terms and P a $|\vec{t}|$ -ary predicate. An atom that features a predicate P is a P -*atom*. A *fact* is a function- and variable-free atom. For a nullary predicate P , we write P as a shortcut for the nullary atom $P()$. For a first-order formula F and a list \vec{x} of variables, we write $F[\vec{x}]$ to indicate that \vec{x} is the set of all free variables occurring in F .

Definition 1. An (existential) rule R is a function- and constant-free first-order formula of the form $\forall \vec{x}.(B[\vec{x}] \rightarrow \exists \vec{z}.H[\vec{y}, \vec{z}])$ where B and H are atom conjunctions, H is non-empty, and all variables in \vec{y} occur in \vec{x} . We refer to \vec{y} as the frontier of R , and to B and $\exists \vec{z}.H$ as the body and the head of R , respectively. Such a rule is *datalog* if it does not feature existential variables; that is, if \vec{z} is the empty list.

We omit universal quantifiers when writing rules. Also, we replace wedges with commas when writing atom conjunctions, which we often identify with sets.

A *substitution* π is a function from variables to terms. For an atom $P(\vec{t})$, let $\pi(P(\vec{t}))$ be the atom that results from replacing all occurrences of every variable x in $P(\vec{t})$ with $\pi(x)$ if the latter is defined. For some atom sets \mathcal{A} to \mathcal{B} , a *homomorphism* π from \mathcal{A} to \mathcal{B} is a substitution such that $\pi(\mathcal{A}) \subseteq \mathcal{B}$. Often, we abuse notation and consider substitutions (and thus homomorphisms) that map terms to terms. An *isomorphism* π from \mathcal{A} to \mathcal{B} is an injective homomorphism \mathcal{A} to \mathcal{B} such that $\pi(\mathcal{A}) = \mathcal{B}$ and $(\pi^{-1}(\mathcal{B}) = \mathcal{A})$.

A *Boolean conjunctive query* (BCQ) is a function-free first-order formula of the form $\exists \vec{z}.B[\vec{z}]$ with B a non-empty conjunction of atoms. A *union of BCQs* (UBCQs) is defined in the obvious manner. Under standard first-order semantics, a fact set \mathcal{F} entails a UBCQ $\bigvee_{i=1}^n \exists \vec{z}_i.B_i$ if there exists a homomorphism from B_k to \mathcal{F} for some $1 \leq k \leq n$.

A *knowledge base* is a pair $\langle \mathcal{R}, \mathcal{F} \rangle$ with \mathcal{R} a finite rule set and \mathcal{F} a finite fact set. Without loss of generality, we assume that (\dagger) existentially quantified variables do not re-occur across different rules in the same rule set. For a BCQ q , we write $\langle \mathcal{R}, \mathcal{F} \rangle \models q$ to indicate that $\mathcal{R} \cup \mathcal{F}$ entails q under standard first-order semantics.

2.2 The Chase Algorithm

We use the chase to define entailment procedurally. More precisely, we present a chase variant previously considered

by (Urbani et al. 2018) in which datalog rules are applied with higher priority, and non-datalog rules are applied in parallel and only if they are not already satisfied. The use of this specific variant, which produces a single chase sequence from an input knowledge base, simplifies some of the proofs presented later.

Consider a rule $\forall \vec{x}. \forall \vec{y}. B[\vec{x}] \rightarrow \exists \vec{z}. H[\vec{y}, \vec{z}]$. For every $z \in \vec{z}$, let f_z be a fresh $|\vec{y}|$ -ary function symbol unique for z . Moreover, let $\text{Sk}(\exists \vec{z}. H)$ be the atom set that results from replacing all occurrences of every $z \in \vec{z}$ in H with the term $f_z(\vec{y})$; note that, because of (\dagger) , this symbol is also unique for a rule within a given rule set. In the following, we simply write z as a shortcut for a (nullary) function symbol of the form f_z . Moreover, we write $w(t)$ or $w^1(t)$ as a shortcut for a unary term such as $f_w(t)$, $w^2(t)$ as a shortcut for $f_w(f_w(t))$, and so on.

A *trigger* t is a pair $\langle R, \pi \rangle$ where $R = \forall \vec{x}. \forall \vec{y}. B \rightarrow \exists \vec{z}. H$ is a rule and π is a homomorphism with domain $\vec{x} \cup \vec{y}$. Let $\text{support}(t) = \pi(B)$ and $\text{output}(t) = \pi(\text{Sk}(\exists \vec{z}. H))$. The trigger t is *applicable* to an atom set \mathcal{A} if $\text{support}(t) \subseteq \mathcal{A}$ and \mathcal{A} does not include $\hat{\pi}(H)$ for all extensions $\hat{\pi}$ of π . A trigger is *datalog* if it features a datalog rule.

For a rule R and an atom set \mathcal{A} , let $R(\mathcal{A})$ be the set that includes $\text{output}(t)$ for every trigger t with R that is applicable to \mathcal{A} . For a rule set \mathcal{R} , let $\mathcal{R}(\mathcal{A}) = \bigcup_{R \in \mathcal{R}} R(\mathcal{A}) \cup \mathcal{A}$. Moreover, let \mathcal{R}^\forall and \mathcal{R}^\exists be the sets of all datalog and non-datalog rules in \mathcal{R} , respectively. The *datalog closure* of \mathcal{A} and \mathcal{R} is the minimal superset \mathcal{A}' of \mathcal{A} such that $\mathcal{R}^\forall(\mathcal{A}') = \mathcal{A}'$; that is, \mathcal{A}' is the (unique) minimal superset of \mathcal{A} that satisfies all of the rules in \mathcal{R}^\forall under standard first-order semantics. Note that an atom set satisfies a rule if no trigger with the latter is applicable to the former.

Definition 2. For a knowledge base $\mathcal{K} = \langle \mathcal{R}, \mathcal{F} \rangle$; let $\text{Ch}_1(\mathcal{K}) = \mathcal{F}$, let $\text{Ch}_i(\mathcal{K})$ be the datalog closure of \mathcal{R} and $\text{Ch}_{i-1}(\mathcal{K})$ for every even $i \geq 1$, and let $\text{Ch}_i(\mathcal{K}) = \mathcal{R}^\exists(\text{Ch}_{i-1}(\mathcal{K}))$ for every odd $i \geq 2$. Also, let $\text{Ch}(\mathcal{K}) = \bigcup_{i \geq 1} \text{Ch}_i(\mathcal{K})$ be the chase of \mathcal{K} .

The chase is a handy tool to answer conjunctive queries over a knowledge base.

Proposition 3. A knowledge base \mathcal{K} entails a UBCQ q if and only if $\text{Ch}(\mathcal{K})$ entails q .

The above holds because $\text{Ch}(\mathcal{K})$ is a universal model for $\mathcal{K} = \langle \mathcal{R}, \mathcal{F} \rangle$. That is, there is a homomorphism from the chase of \mathcal{K} to every model of this first-order theory.

2.3 Computability Theory

Definition 4. A transition function for a set Q of states is a (total) function from $(Q \setminus \{q_f\}) \times \{0, 1, B\}$ to $(Q \setminus \{q_s\}) \times \{0, 1\} \times \{L, R\}$ where q_s and q_f are two different states. A (Turing) machine is a pair $\langle Q, \delta \rangle$ where Q is a set of states that contains q_s and q_f , and δ is a transition function for Q .

As per our definition, all machines reuse the same *initial* q_s and *final* q_f states, as well as the same binary alphabet $\{0, 1, B\}$. Moreover, machines do not write blanks, may not reenter the starting state after the initial configuration, and halt if they reach the final state.

Definition 5. A configuration is a tuple $\langle n, t, p, q \rangle$ where $n \geq 1$ is a natural number, t is a function from $\{1, \dots, n\}$ to $\{0, 1, B\}$ that maps n to B , p is some number in $\{1, \dots, n\}$, and q is a state. The starting configuration on some word $w_1, \dots, w_n \in \{0, 1\}^*$ is the tuple $\langle n + 1, t, 1, q_s \rangle$ where t is the function that maps every $i \in \{1, \dots, n\}$ to w_i , (and $n + 1$ to B).

Given some configuration $\langle n, t, p, q \rangle$, we use t to encode the contents of the tape at each position; moreover, we use p and q to encode the position of the head and the current state of the machine, respectively.

Definition 6. Consider a machine $M = \langle Q, \delta \rangle$, a configuration $C = \langle n, t, p, q \rangle$ with $q \in Q \setminus \{q_f\}$, and the tuple $\delta(t(p), q) = \langle r, a, D \rangle$. Then, let $\text{Next}_M(C)$ be the configuration $\langle n + 1, t', p', r \rangle$ such that:

- For every $1 \leq i \leq n$ with $i \neq p$, let $t'(i) = t(i)$. Moreover, let $t'(p) = a$ and $t'(n + 1) = B$.
- If $D = L$; then $p' = p - 1$ if $p \geq 2$, and $p' = p$ otherwise. If $D = R$, then $p' = p + 1$.

As described above, any given machine defines a function that maps non-final configurations to configurations. An exhaustive iteration through these consecutive configurations that begins with a starting configuration yields a run.

Definition 7. Consider a machine M and a word \vec{w} .

- Define $\text{Run}_M(\vec{w})$ as the (possibly infinite) sequence C_0, \dots, C_n, \dots of configurations such that C_0 is the start configuration on \vec{w} , and $C_{i+1} = \text{Next}_M(C_i)$ for all i such that C_i does not feature the final state.
- The machine M halts on \vec{w} if $\text{Run}_M(\vec{w})$ is finite.

Note that, as per our definition, a machine halts on some word if and only if it reaches the final configuration. We do not discuss acceptance or rejection of a word by a machine; this is unnecessary for our purposes since the halting problem is already undecidable.

Proposition 8. The problem of checking if a machine halts on the empty word is undecidable.

3 Query Languages and Expressivity

The focus of this paper is on the query rewriting approach. To address it, we must first introduce the notions of rule query and of a \mathcal{L} -rewriting for some query language \mathcal{L} .

Definition 9. A rule query is a pair $\langle \mathcal{R}, q \rangle$ where \mathcal{R} is a rule set and q is a BCQ. A datalog query is a rule query with a datalog rule set.

The expressivity of rule queries has been studied from both a complexity viewpoint and a model-theoretic one (Rudolph and Thomazo 2015; Bourgaux et al. 2021). A classical and practically important question is whether a given rule query can be expressed within a less expressive query language. This question is formalized with the notion of \mathcal{L} -rewriting of a rule query $\langle \mathcal{R}, q \rangle$, for a query language \mathcal{L} (such as UCQ, datalog, first-order logic,...). We will consider only Boolean queries, hence we regard each query q as a set of fact sets, and we denote $\mathcal{F} \models q$ if $\mathcal{F} \in q$.

Definition 10 (\mathcal{L} -rewriting). Let \mathcal{L} be a query language. An \mathcal{L} -rewriting of a rule query $\langle \mathcal{R}, q \rangle$ over a set of predicates \mathbf{P}_e is a query $q' \in \mathcal{L}$ such that for any finite fact set \mathcal{F} over \mathbf{P}_e , $\langle \mathcal{R}, \mathcal{F} \rangle \models q$ if and only if $\mathcal{F} \models q'$.

Unless specified otherwise, we set $\mathbf{P}_e = \text{Preds}$, and discuss the other cases in the related work section. This means we consider rewritings that preserve BCQ entailment with respect to the initial rule query over any fact set.

Definition 11 (\mathcal{L} -Expressibility). Given some query language \mathcal{L} , a rule query $\langle \mathcal{R}, q \rangle$ is \mathcal{L} -expressible if it admits an \mathcal{L} -rewriting.

Related to \mathcal{L} -expressibility, but possibly a stronger requirement, is the notion of \mathcal{L} -rewritability: an \mathcal{L} -rewriting must not only exist, but we should be able to compute it.

Definition 12 (\mathcal{L} -Rewritability). Let \mathcal{L} be a query language. A set \mathcal{Q} of rule queries is \mathcal{L} -rewritable if there exists a procedure, which given q , computes an \mathcal{L} -rewriting q' of q if $q \in \mathcal{Q}$, and whose behavior is not specified (and may even not terminate) if $q \notin \mathcal{Q}$.

Note that the absence of condition on the behavior of the algorithm in the case where $q \notin \mathcal{L}$ allows us to consider query languages whose membership is not decidable – this is in particular the case of UBCQ-expressible rule sets (Baget et al. 2011). In that case, \mathcal{L} -expressibility and \mathcal{L} -rewritability are identical notions: this is a direct consequence of the algorithm presented in (König et al. 2015) (see Algorithm 1 and Theorem 7).

Theorem 13. The class of all UBCQ-expressible rule queries is UBCQ-rewritable.

From this theorem, we can derive a useful corollary:

Corollary 14. There is a procedure to check if a knowledge base $\langle \mathcal{R}, \mathcal{F} \rangle$ entails a BCQ q that is sound, complete, and terminates if $\langle \mathcal{R}, q \rangle$ is UBCQ-expressible.

Proof. First, we compute a UBCQ-rewriting q' of $\langle \mathcal{R}, q \rangle$ using the procedure described by (König et al. 2015), which is guaranteed to terminate if $\langle \mathcal{R}, q \rangle$ is UBCQ-expressible. Then, we can simply check if $\mathcal{F} \models q'$; the result of this check indicates if $\langle \mathcal{R}, \mathcal{F} \rangle \models q$. \square

The very same reasoning procedure described in the proof of Corollary 14 can be applied off-the-shelf for any class of rule queries that is UBCQ-expressible such as linear (Calì, Gottlob, and Kifer 2013), sticky (Calì, Gottlob, and Pieris 2010), non-local rule sets (Ostropolski-Nalewaja et al. 2022), DL-Lite (Artale et al. 2009), or any other class of UCQ-expressible rule queries yet to be defined.

In this paper, we investigate whether analogous generic procedures can be developed for more expressive query languages. As it is known that BFO-rewritability and UBCQ-rewritability coincide (Rossman 2008) for rule queries, we set \mathcal{L} to datalog and turn our attention to the relationships between datalog-expressibility and datalog-rewritability. As datalog is a major query language, its expressivity has been studied (Feder and Vardi 2003; Dawar and Kreutzer 2008; Rudolph and Thomazo 2016), and it has served as a target query language for numerous

classes of rule queries (see related work). So far, however, the relationship between datalog-expressibility and datalog-rewritability remains unclear. We prove that:

Theorem 15. *The class of all datalog-expressible rule queries is not datalog-rewritable.*

Proof. If the above result does not hold, then we can define a procedure that solves BCQ entailment for datalog-rewritable inputs analogous to the one discussed in Corollary 14. This clashes with Theorem 16 to be proven in the next section, and hence, this theorem follows by contradiction. \square

4 Entailment for Datalog-Expressible Queries is Undecidable

Our only goal in this section is to show that:

Theorem 16. *There is no procedure to check if a knowledge base $\langle \mathcal{R}, \mathcal{F} \rangle$ entails a BCQ q that is sound, complete, and terminates if $\langle \mathcal{R}, q \rangle$ is datalog-expressible.*

Here is our high-level strategy to prove this result:

- In Definition 18, we introduce a reduction that takes a machine M as input and produces the rule set \mathcal{R}_M .
- We show that a machine M halts on the empty word ε if and only if $\langle \mathcal{R}_M, \emptyset \rangle \models \text{Halt}$ where Halt is a nullary predicate occurring in \mathcal{R}_M ; see Lemma 22. The knowledge base $\langle \mathcal{R}_M, \emptyset \rangle$ will be denoted by \mathcal{K}_M .
- We show that, if a machine M does not halt on ε , then the rule query $\langle \mathcal{R}_M^\forall, \text{Halt} \rangle$ is a datalog-rewriting of $\langle \mathcal{R}_M, \text{Halt} \rangle$;¹ see Lemma 24.

Once we establish all the above claims, we can readily prove the main result in this section:

Proof of Theorem 16. First, we prove that the rule query $\langle \mathcal{R}_M, \text{Halt} \rangle$ is datalog-expressible for any given machine M . If M halts on ε , then $\langle \mathcal{R}_M, \emptyset \rangle \models \text{Halt}$ by □. Therefore, $\langle \mathcal{R}_M, \mathcal{F} \rangle \models \text{Halt}$ for every fact set \mathcal{F} since first-order logic entailment is monotonic, and the rule query $\langle \{\rightarrow \text{Halt}\}, \text{Halt} \rangle$ is a valid datalog-rewriting of $\langle \mathcal{R}_M, \text{Halt} \rangle$ in this case. Otherwise, we conclude that $\langle \mathcal{R}_M, \text{Halt} \rangle$ is datalog-expressible by □.

Now suppose for a contradiction that there is a procedure such as the one discussed in Theorem 16. Given a machine M , we can first compute \mathcal{R}_M and then use such a procedure to effectively decide if $\langle \mathcal{R}_M, \emptyset \rangle \models \text{Halt}$ since $\langle \mathcal{R}_M, \text{Halt} \rangle$ is datalog-expressible. Moreover, we can use the result of this check to verify if M halts on ε by □. Note the clash with Proposition 8. \square

We address Points □, □, and □ separately in each of followin subsections.

¹Remember that \mathcal{R}_M^\forall is the set of all datalog rules in \mathcal{R}_M . We introduced this notation in Section 2.2.

$$\begin{array}{ll}
 \rightarrow \exists s. \mathbf{q}_s(s, s) & (R^{\mathbf{q}_s}) \\
 \mathbf{q}_s(x, x) \rightarrow \mathbf{RS}(x) & (R^{\mathbf{RS}}_{\mathbf{q}_s}) \\
 \mathbf{N}(\ell_i, \ell_{i+1}) \rightarrow \mathbf{N}^+(\ell_i, \ell_{i+1}) & (R^{\mathbf{N}^+}_{\text{init}}) \\
 \mathbf{N}(\ell_i, \ell_{i+1}), \mathbf{N}^+(\ell_{i+1}, \ell_j) \rightarrow \mathbf{N}^+(\ell_i, \ell_j) & (R^{\mathbf{N}^+}_{\text{tr}}) \\
 \mathbf{q}_s(\ell_1, \ell_1), \mathbf{N}^+(\ell_1, \ell_i) \rightarrow \mathbf{RS}(\ell_i) & (R^{\mathbf{RS}}) \\
 \mathbf{RS}(x) \rightarrow \mathbf{B}(x, x) & (R^{\mathbf{B}}_{\mathbf{RS}}) \\
 \mathbf{q}_f(p, c) \rightarrow \mathbf{Halt} & (R^{\mathbf{Halt}}_{\text{final}}) \\
 \\
 \mathbf{Eq}(x, y) \rightarrow \mathbf{Eq}(y, x) & (R^{\mathbf{Eq}}_{\text{sym}}) \\
 \mathbf{Eq}(x, y), \mathbf{Eq}(y, z) \rightarrow \mathbf{Eq}(x, z) & (R^{\mathbf{Eq}}_{\text{tr}}) \\
 \\
 \mathbf{Chaos} \rightarrow \mathbf{Halt} & (R^{\mathbf{Halt}}_{\text{Chaos}}) \\
 \mathbf{N}^+(\ell_i, \ell_i) \rightarrow \mathbf{Chaos} & (R^{\mathbf{Ch}}_{\text{loop}}) \\
 \mathbf{N}^+(\ell_{-i}, \ell_1), \mathbf{q}_s(\ell_1, \ell_1) \rightarrow \mathbf{Chaos} & (R^{\mathbf{Ch}}_{\mathbf{N}^+ \cdot \mathbf{q}_s}) \\
 \mathbf{q}_f(p, c_i), \mathbf{N}^+(c_i, c_j) \rightarrow \mathbf{Chaos} & (R^{\mathbf{Ch}}_{\mathbf{q}_f \cdot \mathbf{N}^+}) \\
 \mathbf{N}(\ell_i, \ell_{i+1}), \mathbf{N}(\ell_i, \ell'_{i+1}) \rightarrow \mathbf{Eq}(\ell_{i+1}, \ell'_{i+1}) & (R^{\mathbf{Eq}}_{<2\mathbf{N}}) \\
 \mathbf{N}(\ell_i, \ell_{i+1}), \mathbf{N}(\ell'_i, \ell_{i+1}) \rightarrow \mathbf{Eq}(\ell_i, \ell'_i) & (R^{\mathbf{Eq}}_{<2\mathbf{N}^-}) \\
 \mathbf{q}_s(\ell_1, \ell'_1), \mathbf{q}_s(\ell''_1, \ell'''_1) \rightarrow \mathbf{Eq}(\ell_1, \ell'''_1) & (R^{\mathbf{Eq}}_{<2\mathbf{q}_s})
 \end{array}$$

Figure 1: Some of the Rules in \mathcal{R}_M

4.1 Point □ : The Reduction

Before presenting our reduction in Definition 18, we define the signature of the output rule set:

Definition 17. *For a machine $M = \langle Q, \delta \rangle$; let $\text{Preds}(M)$ be the set of predicates that contains the nullary predicates Halt and Chaos ; the unary predicate \mathbf{RS} ; the binary predicates \mathbf{N} , \mathbf{N}^+ , \mathbf{O} , $\mathbf{1}$, \mathbf{B} , and \mathbf{Eq} ; and a fresh binary predicate $\mathbf{q} \in Q$.²*

Before reading ahead, remember that rules are function- and constant-free by Definition 1. Hence, all terms occurring in these formulas are variables.

Definition 18. *For a machine $M = \langle Q, \delta \rangle$, let \mathcal{R}_M be the rule set over the predicates in $\text{Preds}(M)$ that includes $\mathcal{R}_M^{\text{Em}}$, $\mathcal{R}_M^{\text{Eq}}$, and $\mathcal{R}_M^{\text{Ch}}$, which are defined below. In these definitions, we write $\langle x, y \rangle_{\mathbf{RS}}^+$ for some variables x and y as a shortcut for the conjunction $\mathbf{RS}(x), \mathbf{N}^+(x, y), \mathbf{RS}(y)$.*

The rule set $\mathcal{R}_M^{\text{Em}}$ contains the first seven rules in Figure 1 plus all of the following:

- For every $q \in Q \setminus \{q_f\}$, we add:

$$\mathbf{q}(p, c), \mathbf{RS}(c) \rightarrow \exists n. \mathbf{N}(c, n) \quad (R^{\mathbf{N}}_q)$$

²Predicate names \mathbf{RS} , \mathbf{N} , and \mathbf{Eq} are shortcuts for “reachable from start state”, “next”, and “equality”, respectively.

- For every $\langle q, a \rangle \mapsto \langle r, b, L \rangle \in \delta$, we add:

$$\begin{aligned} & \mathbf{q}(p_i, c_j), \mathbf{a}(p_i, c_j), \mathbf{N}(p_{i-1}, p_i), \mathbf{N}(c_j, c_{j+1}), \\ & \langle p_i, c_{j+1} \rangle_{\text{RS}}^+ \rightarrow \mathbf{r}(p_{i-1}, c_{j+1}), \mathbf{b}(p_i, c_{j+1}) \quad (R_{q,a}^{r,b,L}) \\ & \mathbf{q}(p_1, c_j), \mathbf{a}(p_1, c_j), \mathbf{q}_s(p_1, p_1), \mathbf{N}(c_j, c_{j+1}), \\ & \langle p_1, c_{j+1} \rangle_{\text{RS}}^+ \rightarrow \mathbf{r}(p_1, c_{j+1}), \mathbf{b}(p_1, c_{j+1}) \quad (R_{q,a}^{r,b,S}) \end{aligned}$$

- For every $\langle q, a \rangle \mapsto \langle r, b, R \rangle \in \delta$, we add:

$$\begin{aligned} & \mathbf{q}(p_i, c_j), \mathbf{a}(p_i, c_j), \mathbf{N}(p_i, p_{i+1}), \mathbf{N}(c_j, c_{j+1}), \\ & \langle p_{i+1}, c_{j+1} \rangle_{\text{RS}}^+ \rightarrow \mathbf{r}(p_{i+1}, c_{j+1}), \mathbf{b}(p_i, c_{j+1}) \quad (R_{q,a}^{r,b,R}) \\ & \mathbf{q}(\ell_i, \ell_i), \mathbf{a}(\ell_i, \ell_i), \mathbf{N}(\ell_i, \ell_{i+1}), \mathbf{RS}(\ell_{i+1}) \\ & \rightarrow \mathbf{r}(\ell_{i+1}, \ell_{i+1}), \mathbf{b}(\ell_i, \ell_{i+1}) \quad (S_{q,a}^{r,b,R}) \end{aligned}$$

- For every $q \in Q \setminus \{q_f\}$ and $a \in \{0, 1, B\}$, we add:

$$\begin{aligned} & \mathbf{a}(p_i, c_j), \mathbf{q}(p_k, c_j), \mathbf{N}(c_j, c_{j+1}), \mathbf{N}^+(p_i, p_k), \\ & \langle p_k, c_{j+1} \rangle_{\text{RS}}^+, \langle p_i, c_{j+1} \rangle_{\text{RS}}^+ \rightarrow \mathbf{a}(p_i, c_{j+1}) \quad (R_{\text{copy}}^{a,L}) \\ & \mathbf{q}(p_i, c_j), \mathbf{a}(p_k, c_j), \mathbf{N}(c_j, c_{j+1}), \mathbf{N}^+(p_i, p_k), \\ & \langle p_k, c_{j+1} \rangle_{\text{RS}}^+, \langle p_i, c_{j+1} \rangle_{\text{RS}}^+ \rightarrow \mathbf{a}(p_k, c_{j+1}) \quad (R_{\text{copy}}^{a,R}) \end{aligned}$$

The rule set $\mathcal{R}_M^{\text{Eq}}$ contains the two rules in the middle of Figure 1 plus the following for every $P \in \text{Preds}(M) \setminus \{\text{Eq}, \text{Chaos}, \text{Halt}\}$ and $1 \leq i \leq \text{ar}(P)$:

$$\begin{aligned} & P(x_1, \dots, x_n) \rightarrow \text{Eq}(x_i, x_i) \quad (R_{P,i}^{\text{Eq}}) \\ & P(x_1, \dots, x_n), \text{Eq}(x_i, y_i) \\ & \rightarrow P(x_1, \dots, x_{i-1}, y_i, x_{i+1}, \dots, x_n) \quad (R_{\text{cg}}^{P,i}) \end{aligned}$$

The rule set $\mathcal{R}_M^{\text{Ch}}$ contains the last seven rules in Figure 1, and includes all of the following rule sets:

$$\begin{aligned} & \{\mathbf{q}(p, c), \mathbf{q}(p', c) \rightarrow \text{Eq}(p, p') \mid q \in Q\} \quad (R_q^{\text{Eq}}) \\ & \{\mathbf{q}(p, c), \mathbf{r}(p', c) \rightarrow \text{Chaos} \mid q \neq r \text{ in } Q\} \quad (R_{q \neq r}^{\text{Ch}}) \\ & \{\mathbf{s}(p, c), \mathbf{N}^+(c, p) \rightarrow \text{Chaos} \mid s \in \{0, 1, B\} \cup Q\} \quad (R_{s, \text{N}^+}^{\text{Ch}}) \\ & \{\mathbf{a}(p, c), \mathbf{b}(p, c) \rightarrow \text{Chaos} \mid a \neq b \text{ in } \{0, 1, B\}\} \quad (R_{a \neq b}^{\text{Ch}}) \end{aligned}$$

To ease the understanding, we first provide a high-level explanation for each of the rule sets defined in Definition 3; more detailed explanations specific to each rule will follow in Sections 4.2 and 4.3.

For a machine M , the rule set $\mathcal{R}_M^{\text{Eq}}$ ensures that the result of the chase on $\langle \mathcal{R}_M, \emptyset \rangle$ encodes the information in the run of M on the empty word ε . Put differently, this rule set ensures that the chase on $\langle \mathcal{R}_M, \emptyset \rangle$ becomes a procedure that closely emulates the computation of M on this specific input. This rule set is prominently used in the following section to show Lemma 22.

The rule set $\mathcal{R}_M^{\text{Eq}}$ ensures that Eq , which is just a regular binary predicate, behaves as an axiomatization of first-order equality; that is, of the special predicate \approx . In particular, note the rules of type $R_{\text{cg}}^{P,i}$, which ensure that Eq defines congruent relation in the result of the chase. The other rules in this rule set ensure that it defines an equivalence relation.

The rule set $\mathcal{R}_M^{\text{Ch}}$ allows us to derive Halt in many cases only using datalog rules. For example, using these rules

we can readily show that Halt is in the result of the chase on a knowledge base $\langle \mathcal{R}_M^{\forall}, \mathcal{F} \rangle$ if \mathcal{F} is a fact set that includes an N -cycle; that is, a non-empty set of facts of the form $\{\mathbf{N}(d_1, d_2), \dots, \mathbf{N}(d_{n-1}, d_n), \mathbf{N}(d_n, d_1)\}$. Thus, we can show that $\langle \mathcal{R}_M, \mathcal{F} \rangle \models \text{Halt}$ if and only if $\langle \mathcal{R}_M^{\forall}, \mathcal{F} \rangle \models \text{Halt}$ for such a fact set \mathcal{F} . In turn, this greatly simplifies the proof of Lemma 24.

4.2 Point \square : Machine Emulation

Our only goal in this subsection is to show that \mathcal{R}_M correctly simulates M on the empty word (Lemma 22). Let us clarify this intuition with the following example.

Example 19. Consider a state q and a machine $M = \langle Q, \delta \rangle$ such that $Q = \{q_s, q_f, q\}$, $\delta(q_s, B) = \langle q, 1, R \rangle$, and $\delta(q, B) = \langle q_f, 0, L \rangle$. We depict $\text{Run}_M(\varepsilon)$ and the chase of $\langle \mathcal{R}_M, \emptyset \rangle$ in Figure 2; note how the latter structure encodes the information in the former.


For instance, the alphabet symbol 1 occurs in the second position of the tape in the third configuration of $\text{Run}_M(\varepsilon)$; this information is encoded in $\text{Ch}(\mathcal{K}_M)$ with the atom $1(n(s), n^2(s))$.³ The alphabet symbol 1 is encoded using the corresponding binary predicate 1; the fact that this symbol occurs in the second position of the tape in the third configuration is encoded using the second $n(s)$ and third $n^2(s)$ elements of the N -chain, respectively. In other words, the first and second arguments of an 1-atom encode the information about the position and configuration of some occurrence of the alphabet symbol 1, respectively. Note that this intuition is reflected in the naming of the variables appearing in rules, where p stands for position and c stands for configuration.

Information pertaining to the position of the head as well as the state of the machine is analogously encoded using binary atoms over q , q_s , and q_f . For instance, the atom $q_f(s, n^2(s)) \in \text{Ch}(\mathcal{K}_M)$ is used to encode that, in the third configuration, the machine enters the final state q_f and its head is in the first position of the tape. Finally, note that the nullary fact Halt is in $\text{Ch}(\mathcal{K}_M)$ since this set contains an atom over the predicate q_f .

After the previous example, we can now elucidate the purpose of the rules in the rule set $\mathcal{R}_M^{\text{Eq}}$ for a given machine M . Roughly speaking, we can divide this rule set into five mutually disjoint partitions:

- First, we have the first six rules in Figure 1 as well as the rules of type R_q^{N} , which instantiate the N -chain in $\text{Ch}(\mathcal{K}_M)$. Moreover, these rules also materialize the atoms encoding the starting configuration on the empty word as well as every new blank symbol at the end of every subsequent configuration.
- Second, we have all rules of type $R_{q,a}^{r,b,L}$, $R_{q,a}^{r,b,S}$, $R_{q,a}^{r,b,R}$, or $S_{q,a}^{r,b,R}$. These rules instantiate the only rewritten alphabet symbol in every non-starting configuration, and update the state of the machine and move the head to the left ($R_{q,a}^{r,b,L}$ and $R_{q,a}^{r,b,S}$) or right ($R_{q,a}^{r,b,R}$ and $S_{q,a}^{r,b,R}$).

³Remember that $n^2(s)$ is a shortcut for the unary term $f_n(f_n(s))$. We introduced this notation in Section 2.2.


 Figure 2: The Sequence $\text{Run}_M(\varepsilon)$ and the Chase of $\langle \mathcal{R}_M, \emptyset \rangle$ where M is the machine from Example 19

- Fourth, the rules of type $R_{copy}^{a:L}$ and $R_{copy}^{a:R}$ take care of copying the alphabet symbols that are to the left and to the right of the head to every subsequent configuration, respectively. Note the use of the predicate N^+ to detect if a position for a given configuration is to the left or the right.
- Finally, Rule R_{final}^{Halt} , which produces Halt if the machine reaches the final configuration in the run.

All other rules in \mathcal{R}_M are mostly inactive during the computation of the chase on $\langle \mathcal{R}_M, \emptyset \rangle$. Namely, these rules solely derive facts of the form $\text{Eq}(t, t)$ and thus have no impact on deriving Halt when we only consider the knowledge base $\langle \mathcal{R}_M, \emptyset \rangle$ with the empty set.

We are ready to formalize the intuition presented in Example 19. Namely, for a machine M , we first define some atom sets that encode the information in every configuration in $\text{Run}_M(\varepsilon)$, and then show that the union of all of these sets is the result of the chase on $\langle \mathcal{R}_M, \emptyset \rangle$.

Definition 20. Consider a machine M , some $1 \leq k \leq |\text{Run}_M(\varepsilon)|$, and the k -th configuration $\langle k, t, p, q \rangle$ of $\text{Run}_M(\varepsilon)$. Let $\mathcal{C}_M(k)$ be the atom set that includes $\mathcal{C}_M(k-1)$ if $k > 1$, contains Halt if $q = q_f$, and includes

$$\begin{aligned} & \{N(n^i(s), n^{i+1}(s)) \mid 0 \leq i < k-1\} \cup \\ & \{N^+(n^i(s), n^j(s)) \mid 0 \leq i < j < k\} \cup \\ & \{\text{RS}(n^i(s)), \text{Eq}(n^i(s), n^i(s)) \mid 0 \leq i < k\} \cup \\ & \{\text{Pred}(t(i+1))(n^i(s), n^{k-1}(s)) \mid 0 \leq i < k\} \end{aligned}$$

where $\text{Pred}(0) = 0$, $\text{Pred}(1) = 1$, and $\text{Pred}(B) = B$. Moreover, let \mathcal{C}_M be the atom set that includes $\mathcal{C}_M(i)$ for every $1 \leq i \leq |\text{Run}_M(\varepsilon)|$.

Lemma 21. For a machine M , we have $\text{Ch}(\mathcal{K}_M) = \mathcal{C}_M$.

Sketch. First, we verify that $\text{Ch}_{2j+2}(\mathcal{R}_M) = \mathcal{C}_M(j)$ for every $1 \leq j \leq |\text{Run}_M(\varepsilon)|$ by induction on j . Furthermore, we show that, if M halts on the empty word ε , then $\text{Ch}(\mathcal{K}_M) = \text{Ch}_{2|\text{Run}_M(\varepsilon)|+2}(\langle \mathcal{R}_M, \emptyset \rangle)$. \square

For a machine M , the previous lemma establishes the formal correspondence between the information in $\text{Run}_M(\varepsilon)$ and the result of the chase of $\langle \mathcal{R}_M, \emptyset \rangle$. Applying this result, we can readily show the main result of this subsection.

Lemma 22. A machine M halts on the empty word if and only if the knowledge base $\langle \mathcal{R}_M, \emptyset \rangle$ entails Halt .

Proof. If $\langle \mathcal{R}_M, \emptyset \rangle$ entails Halt for some machine M , then the nullary atom Halt is in the result of the chase on $\langle \mathcal{R}_M, \emptyset \rangle$ by Proposition 3. Therefore, Halt is in $\mathcal{C}_M(k)$ for some $1 \leq k \leq |\text{Run}_M(\varepsilon)|$ by Lemma 21. By Definition 20, this is the case only if the k -th configuration in $\text{Run}_M(\varepsilon)$ features the final state q_f .

If a machine M halts on ε , then the nullary atom Halt is in $\mathcal{C}_M(|\text{Run}_M(\varepsilon)|)$ by Definition 20. Hence, Halt is in the result of the chase on $\langle \mathcal{R}_M, \emptyset \rangle$ by Lemma 21. \square

If our only goal was to show Lemma 22, we could reuse classical machine simulations relying on a grid structure, as done for instance in (Baget et al. 2011). However, to prove datalog-expressibility of $\langle \mathcal{R}_M, \text{Halt} \rangle$, it is convenient to have a simpler structure of Skolem terms: \mathcal{R}_M has only two existential rules, R_{qs} of empty frontier and R_q^n of frontier one. This facilitates the detection (through datalog rules) of fact sets that either do not encode a valid run of a machine.

4.3 Point \square : Datalog-Expressibility

For the remainder of the subsection, we fix a machine $M = \langle Q, \delta \rangle$ and a fact set \mathcal{F} . Moreover, we define another fact set that results from collapsing every set of constants in $\text{Ch}(\mathcal{R}_M^\forall, \mathcal{F})$ mutually interconnected by Eq .

Definition 23. Since rules $R_{p,i}^{\text{Eq}}$, R_{sym}^{Eq} , and R_{tr}^{Eq} are in \mathcal{R}_M^\forall , the Eq predicate defines an equivalence relation in $\text{Ch}(\mathcal{R}_M^\forall, \mathcal{F})$. For a constant c occurring in this atom set, let $[c]$ be a fresh constant unique for the equivalence class of c induced by Eq . Let $\mathcal{F}_M^{\text{Eq}} = \{P([c_1], \dots, [c_n]) \mid P(c_1, \dots, c_n) \in \text{Ch}(\mathcal{R}_M^\forall, \mathcal{F})\}$.

The fact set $\mathcal{F}_M^{\text{Eq}}$ is isomorphic to the chase of $\langle \mathcal{R}_M^\forall, \mathcal{F} \rangle$ if we had used first-order equality in \mathcal{R}_M instead of an axiomatization of this special predicate. If $\text{Chaos} \notin \mathcal{F}_M^{\text{Eq}}$, then this fact set has a clean structure that is quite useful in our proofs. After this short preamble, let us consider the main result of this subsection:

Lemma 24. If M does not halt on ε , then $\langle \mathcal{R}_M^\forall, \text{Halt} \rangle$ is a datalog-rewriting of $\langle \mathcal{R}_M, \text{Halt} \rangle$

Proof. Since \mathcal{R}_M includes \mathcal{R}_M^\forall and first-order entailment is monotonic, we conclude that $\langle \mathcal{R}_M^\forall, \mathcal{F} \rangle \models \text{Halt}$ implies $\langle \mathcal{R}_M, \mathcal{F} \rangle \models \text{Halt}$. If the premise of the lemma holds, then

we conclude that $\langle \mathcal{R}_M^\forall, \mathcal{F} \rangle \not\models \text{Halt}$ implies $\langle \mathcal{R}_M, \mathcal{F} \rangle \not\models \text{Halt}$ from Lemmas 25 and 33. \square

Because of the rules in $\mathcal{R}_M^{\text{Eq}}$, the predicate Eq behaves as an axiomatization of equality in the chase of $\langle \mathcal{R}_M^\forall, \mathcal{F} \rangle$. Hence, when we collapse constants interconnected via Eq to create $\mathcal{F}_M^{\text{Eq}}$, we preserve entailment of nullary facts and do not violate existing datalog rules.

Lemma 25. *We have that $\langle \mathcal{R}_M, \mathcal{F}_M^{\text{Eq}} \rangle \not\models \text{Halt}$ if and only if $\langle \mathcal{R}_M, \mathcal{F} \rangle \not\models \text{Halt}$. Moreover, $\mathcal{F}_M^{\text{Eq}}$ satisfies all rules in \mathcal{R}_M^\forall (and therefore, $\text{Ch}(\mathcal{R}_M^\forall, \mathcal{F}_M^{\text{Eq}}) = \mathcal{F}_M^{\text{Eq}}$).*

Intuitively, the above lemma implies that we can disregard \mathcal{F} and focus on $\mathcal{F}_M^{\text{Eq}}$ for the remainder of the section. Now, let us explore the structure of the latter, which is not evident in the former. Namely, we show that, if $\text{Chaos} \notin \mathcal{F}_M^{\text{Eq}}$, then the set of all constants in $\mathcal{F}_M^{\text{Eq}}$ can be partitioned into a disjoint union of N-chains.

Definition 26. *An N-chain C is a non-empty (finite) list c_1, \dots, c_n of constants such that $c_i \neq c_j$ for every $1 \leq i < j \leq n$; $\text{N}(c_i, c_{i+1}) \in \mathcal{F}_M^{\text{Eq}}$ for every $1 \leq i \leq n-1$; and, if there is a fact of the form $\text{N}(d, e) \in \mathcal{F}_M^{\text{Eq}}$ such that either d or e occur in C , then $d = c_i$ and $e = c_{i+1}$ for some $1 \leq i \leq n-1$. A starting chain is an N-chain c_1, \dots, c_n such that $\text{qs}(c_1, c_1) \in \mathcal{F}_M^{\text{Eq}}$.*

Lemma 27. *If $\text{Chaos} \notin \mathcal{F}_M^{\text{Eq}}$; then the set of all constants occurring in $\mathcal{F}_M^{\text{Eq}}$ is a disjoint union of chains, there is at most one qs -atom in $\mathcal{F}_M^{\text{Eq}}$ (and hence is at most one starting chain), and there is one starting chain if and only if $\mathcal{F}_M^{\text{Eq}}$ contains some qs -atom.*

Sketch. The first, second, and third implications hold because \mathcal{R}_M^\forall includes $\{R_{\text{init}}^{\text{N}^+}, R_{\text{tr}}^{\text{N}^+}, R_{\text{loop}}^{\text{Ch}}, R_{<2\text{N}}^{\text{Eq}}, R_{<2\text{N}^-}^{\text{Eq}}\}$, $\{R_{<2\text{qs}}^{\text{Eq}}\}$, and $\{R_{\text{init}}^{\text{N}^+}, R_{\text{N}^+ \cdot \text{qs}}^{\text{Ch}}, R_{<2\text{qs}}^{\text{Eq}}\}$, respectively. Note that $\mathcal{F}_M^{\text{Eq}}$ satisfies all rules in \mathcal{R}_M^\forall by Lemma 25, and $c = d$ if $\text{Eq}(c, d) \in \mathcal{F}_M^{\text{Eq}}$ for some constants c and d . \square

Furthermore, we can show that, if there is a starting chain, then the fact set associated to this structure is isomorphic to a finite prefix of the chain in $\text{Ch}(\mathcal{K}_M)$.

Definition 28. *Let π_S be the (injective) substitution that maps s to the first element of (the starting chain) S , $n(s)$ to the second, $n^2(s)$ to the third, and so on. For a list \vec{t} of terms, let $\text{Facts}(\vec{t})$ be the set of all facts in $\mathcal{F}_M^{\text{Eq}}$ that can be defined using some non-nullary predicate and some terms occurring in \vec{t} .*

Remember that, since we slightly abuse notation, substitutions may be defined for non-variable terms.

Lemma 29. *If $\text{Chaos} \notin \mathcal{F}_M^{\text{Eq}}$ and there is a starting chain S , then the substitution π_S is an isomorphism from $\mathcal{C}_M(|S|)$ to $\text{Facts}(S)$.⁴*

⁴Remember that $\mathcal{C}_M(|S|)$ was introduced in Definition 20.

Sketch. We can use an argument analogous to the one in the proof of Lemma 21 to show that π_S is a homomorphism from $\mathcal{C}_M(|S|)$ to $\text{Facts}(S)$. Then, for the other direction, we use Rules $R_{\text{loop}}^{\text{Ch}}, R_{\text{N}^+ \cdot \text{qs}}^{\text{Ch}}, R_{\text{q}_f \cdot \text{N}^+}^{\text{Ch}}, R_{<2\text{N}}^{\text{Eq}}, R_{<2\text{N}^-}^{\text{Eq}}, R_{<2\text{qs}}^{\text{Eq}}, R_{\text{q} \neq \text{r}}^{\text{Ch}}, R_{\text{q}}^{\text{Eq}}, R_{\text{a} \neq \text{b}}^{\text{Ch}}$ and $R_{\text{s} \cdot \text{N}^+}^{\text{Ch}}$ to show that $\text{Facts}(S)$ only contains atoms in $\pi_S(\mathcal{C}_M(|S|))$. For instance, if an atom of the form $\text{N}^+(c, d)$ is in $\text{Facts}(S)$ but not in $\pi_S(\mathcal{C}_M(|S|))$, then either $c = d$, or c comes after d in the starting chain.

- If $c = d$, then $R_{\text{loop}}^{\text{Ch}}$ implies that $\text{Chaos} \in \text{Facts}(S)$, which contradicts the premise of the lemma.
- If c comes after d in the starting chain, then there is a list c_1, \dots, c_n of constants such that $c_1 = d$, $c_n = c$, and $\text{N}(c_i, c_{i+1}) \in \text{Facts}(S)$ for every $1 \leq i < n$. Then, $\text{N}^+(d, d) \in \text{Facts}(S)$ because of Rule $R_{\text{tr}}^{\text{N}^+}$, and $\text{Chaos} \in \text{Facts}(S)$ because of Rule $R_{\text{loop}}^{\text{Ch}}$, which contradicts the premise of the lemma.

Thus, there are no N^+ -atom in $\text{Facts}(S) \setminus \pi_S(\mathcal{C}_M(|S|))$. A careful analysis of all the other predicates then concludes the proof. \square

Now that we understand the structure and content of $\mathcal{F}_M^{\text{Eq}}$, we can declaratively describe how to extend this fact set to obtain the chase of $\langle \mathcal{R}_M, \mathcal{F}_M^{\text{Eq}} \rangle$.

Definition 30. *An N-chain $C = c_1, \dots, c_m$ is incomplete if the set $\mathcal{F}_M^{\text{Eq}}$ contains $\text{RS}(c_m)$ and some fact of the form $\text{q}(d, c_m)$ such that $\text{q} \neq \text{q}_f$. Given such an N-chain, let $\text{Compl}(C)$ be the atom set that contains $\text{N}(c_m, n(c_m))$, $\text{Eq}(n(c_m), n(c_m))$, and $\text{N}^+(c_i, n(c_m))$ for every $1 \leq i \leq m$. Let $\text{Compl}(\mathcal{F}_M^{\text{Eq}})$ be the atom set that includes $\text{Compl}(C)$ for every incomplete chain C .*

Lemma 31. *If $\text{Chaos} \notin \mathcal{F}_M^{\text{Eq}}$, then*

$$\text{Ch}(\mathcal{R}_M, \mathcal{F}_M^{\text{Eq}}) = \mathcal{F}_M^{\text{Eq}} \cup \text{Compl}(\mathcal{F}_M^{\text{Eq}}) \cup \pi_S(\text{Ch}(\mathcal{K}_M)).$$

Roughly speaking, if $\text{Chaos} \notin \mathcal{F}_M^{\text{Eq}}$, we can define the chase of $\langle \mathcal{R}_M, \mathcal{F}_M^{\text{Eq}} \rangle$ by taking $\mathcal{F}_M^{\text{Eq}}$, completing every incomplete chain, and appending the (possibly infinite) chain in $\text{Ch}(\mathcal{K}_M)$ to the starting chain in $\mathcal{F}_M^{\text{Eq}}$. To understand why this suffices, see the following example.

Example 32. *Consider a machine that has three states, q_s , q and q_f , that always goes right, and that goes to state q whenever it reads a blank, and goes to q_f whenever it reads a 0 or a 1. Obviously, this machine does not halt on the empty word, but does on every other word. The fact set on the left of Figure 3 contains two N-chains: the starting chain at the top, and another at the bottom. After one existential step and one datalog step, note that the atom $\text{q}_f(n(d), n(d))$ has not been derived: this is due to the fact that $\text{RS}(n(d))$ is not derivable, as $n(d)$ is not reachable through an N-chain from c_0 , the unique element for which $\text{qs}(c_0, c_0)$ holds. This blocks the application of Rule $S_{\text{q}, \text{a}}^{\text{r}, \text{b}, \text{R}}$.*

If ones try to circumvent that problem as on the right of Figure 3, with two N-chains connected by a q atom, the atom




Figure 3: Fact sets from Example 32 after one existential step and one closure under datalog rules. Dots represent constants occurring in the fact set, and crosses represent Skolem terms. Solid atoms are in the original fact set, dashed ones created by rule applications. Unlabelled arrows represent the N predicate. Predicates N^+ and Eq are not represented.

$N^+(n(d), n(c_2))$ cannot be mapped, hence Rule $R_{q,a}^{r,b,R}$ cannot be applied. In turn, this prevents the instantiation of the atom $q_{\mathcal{E}}(n(c_2), n(c_2))$.

These two examples illustrate the use of predicates RS and N^+ , allowing to block the development of simulations that do not correspond to the empty word.

Leveraging the contingent definition of the chase of $\langle \mathcal{R}_M, \mathcal{F}_M^{Eq} \rangle$ given in Lemma 31, we can better understand when Halt is not entailed by this knowledge base:

Lemma 33. *If $\text{Halt} \notin \mathcal{F}_M^{Eq}$ and M does not halt on ε , then $\langle \mathcal{R}_M, \mathcal{F}_M^{Eq} \rangle$ does not entail Halt .*

Proof. By Proposition 3, we have that $\langle \mathcal{R}_M, \mathcal{F}_M^{Eq} \rangle$ does not entail Halt if Halt is not in the chase of $\langle \mathcal{R}_M, \mathcal{F}_M^{Eq} \rangle$. By Lemma 31, this is the case if Halt is not in \mathcal{F}_M^{Eq} , in $\text{Compl}(\mathcal{F}_M^{Eq})$, or in $\pi_S(\text{Ch}(\mathcal{K}_M))$. Indeed, Halt is not in the first set by the premise of the lemma, it is not in the second by Definition 30, and it is not in the third by the premise of the lemma and Lemma 22. Note that the premise of the lemma implies that $\langle \mathcal{R}_M, \mathcal{F}_M^{Eq} \rangle$ does not entail Chaos since the Rule $R_{\text{Chaos}}^{\text{Halt}}$ is in \mathcal{R}_M . \square

5 Related Work

Let us first point out that there are several variations around the notion of datalog-rewriting. A first dimension is that, rather than rewriting a specific rule query $\langle \mathcal{R}, q \rangle$, one can wish to rewrite \mathcal{R} into a datalog rule set \mathcal{R}' , and use \mathcal{R}' to compute answers for a class of queries. Another variation is focused on the fact sets on which the rewriting should output the same answer as the original rule query. In this paper, we consider the strong version where answers should be the same on every fact set over Pred . Quite often, defined datalog rewritings only preserve answers over fact sets on the original signature – this allows to introduce fresh predicates which are known not to belong to fact sets on which the datalog program is to be evaluated. There are cases for which there exists datalog-rewritings of rule queries for this relaxed definition but not for our restricted one; this is a consequence of Theorem 3 in (Krötzsch 2011). Our undecidability result implies undecidability of this more relaxed notion.

The use of datalog as a target language for rewritings has been studied over the last 15 years. The goal was to reduce reasoning task over expressive ontologies towards query answering over datalog, for which optimization techniques have been developed in the database community. This is even more pregnant today, as a variety of efficient datalog reasoners have been implemented (Nenov et al. 2015; Urbani et al. 2018). Such an approach has been proposed for providing disjunctive⁵ datalog rewritings for \mathcal{SHIQ} for fact entailment over the original signature (Hustadt, Motik, and Sattler 2007), later generalized to \mathcal{SHIQ}_{b_S} (Rudolph, Krötzsch, and Hitzler 2012). More recently, such reductions for Horn description logics have been implemented and evaluated (Carral, Dragoste, and Krötzsch 2018; Carral, González, and Koopmann 2019). Such datalog rewritings have also been studied for existential rules, for guarded (Benedikt et al. 2022), nearly guarded (Gottlob, Rudolph, and Simkus 2014), warded (Berger et al. 2022) and shy (Leone et al. 2019) rule sets.

Beyond these fragment specific reductions, the limits of datalog-rewritability have been explored. In (Marnette 2012), it is shown that whenever rule queries have bounded depth (meaning that if they are entailed, they are entailed by a portion of the chase that uses only Skolem terms of bounded depth), they are datalog-rewritable. This result applies for all syntactic fragment for which the chase is known to terminate (Grau et al. 2013), but datalog rewritability is not guaranteed (and not always possible) for rule sets having terminating restricted chase (Krötzsch, Marx, and Rudolph 2019) – this is proven by a data complexity argument.

Another question of interest is the size of the obtained rewritings. In (Ahmetaj, Ortiz, and Simkus 2018), the authors provide polynomial (disjunctive) datalog rewritings for (disjunctive) guarded rules queries. Non-recursive datalog has also been studied: while it does not increase the expressivity with respect to UCQs, the re-use of predicates allows to significantly reduce the size of rewritings, reaching polynomiality in some cases (Gottlob and Schwenck 2012; Gottlob et al. 2014).

All of these contributions are summarised in Table 1.

⁵For disjunctive existential rules and datalog, the reader is invited to consult (Deutsch and Tannen 2003)

Source Language	Target Language	Arbitrary Query Sig.	Implemented	Reference
<i>SHIQ</i>	disj. dat.	×	✓	(Hustadt, Motik, and Sattler 2007)
<i>SHIQ_{BS}</i>	disj. dat.	×	×	(Rudolph, Krötzsch, and Hitzler 2012)
Horn- <i>ALCHOIQ</i>	datalog	×	✓	(Carral, Dragoste, and Krötzsch 2018)
Horn- <i>SRIQ</i>	datalog	×	✓	(Carral, González, and Koopmann 2019)
Bounded Depth rules	datalog	×	×	(Marnette 2012)
Frontier guarded rules	datalog	✓	×	(Bárány, Benedikt, and ten Cate 2013)
Nearly Guarded Rules	datalog	✓	×	(Gottlob, Rudolph, and Simkus 2014)
Guarded disj. rules	disj. datalog	×	×	(Ahmetaj, Ortiz, and Simkus 2018)
Guarded rules	datalog	✓	✓	(Benedikt et al. 2022)
Warded rules	datalog	✓	✓	(Berger et al. 2022)
Linear	non rec. dat.	×	×	(Gottlob and Schwentick 2012)
Sticky(-join)	non rec. dat.	×	×	(Gottlob and Schwentick 2012)

Table 1: Summary of datalog rewriting approaches applicable for fact entailment

6 Future Work

To conclude, we discuss two distinct lines for future research, which naturally follow from our work.

Answer Expressible Rule Sets We intend to study an even more restrictive class of rewritable rule sets: a rule set \mathcal{R} is *answer datalog-expressible* if, for every CQ $q[\vec{x}]$, the rule query $\langle \mathcal{R}, q[\vec{x}] \rangle$ admits some *datalog-rewriting that preserves all answers*. That is, a datalog query $\langle \mathcal{R}', q'[\vec{x}] \rangle$ such that, for every fact set \mathcal{F} and every list \vec{a} of constants occurring in \mathcal{F} , we have that $\langle \mathcal{R}, \mathcal{F} \rangle \models q[\vec{x}/\vec{a}]$ if and only if $\langle \mathcal{R}', \mathcal{F} \rangle \models q'[\vec{x}/\vec{a}]$. With this definition in place, we wonder about the following problem:

Open Problem 34. *Consider a knowledge base $\mathcal{K} = \langle \mathcal{R}, \mathcal{F} \rangle$, a CQ q , and a list \vec{a} of constants occurring in \mathcal{F} . Is there a procedure to check if \vec{a} is an answer of q with respect to $\langle \mathcal{R}, \mathcal{F} \rangle$ that is sound, complete, and terminating if \mathcal{R} is answer datalog-expressible?*

The above question is quite relevant for our field of research, where we often study the theoretical properties of classes of rule sets and not of rule queries.

At the moment, we believe that the answer to this open problem is negative, which would yield a result strictly stronger than Theorem 16. However, a different proof strategy is required to show this since there are rule sets in the range of the reduction described in Definition 18 that are not answer datalog-rewritable.

Alternative Rewriting Languages In this paper, our primary focus is on datalog; in the future, we plan to study alternative query languages for rewritings. For instance, one could consider unions of Boolean conjunctive regular path queries (UBCRPQs) (Florescu, Levy, and Suciu 1998) and then consider the following problem:

Open Problem 35. *Is the class of all UBCRPQ-expressible queries is UBCRPQ-rewritable? Is there a procedure to check if a knowledge base $\langle \mathcal{R}, \mathcal{F} \rangle$ entails a BCQ q that is sound, complete, and terminates if the rule query $\langle \mathcal{R}, q \rangle$ is UBCRPQ-expressible?*

We can instantiate different versions of this open problem by considering different output rewriting languages. For instance, we could consider as unions of (non-conjunctive) regular path queries, monadic datalog, query languages based on context-free grammars (Medeiros, Musicante, and da Costa 2022), or any of the query languages considered by (Bourhis, Krötzsch, and Rudolph 2015).

As a closing remark, note that the answers to the first and second questions in Open Problem 35 might be negative and positive, respectively. That is, it is possible that we can solve entailment for UBCRPQ-expressible rule queries even if we cannot effectively compute rewritings for these. This is an exciting possibility that may lead us to the discovery of a novel kind of reasoning procedure for this expressive class of rule queries. Or perhaps future research will just result in another undecidability result, which would be strictly stronger than Theorem 16. Either way, we look forward to researching (and hopefully settling!) these questions.

References

- Ahmetaj, S.; Ortiz, M.; and Simkus, M. 2018. Rewriting guarded existential rules into small datalog programs. In Kimelfeld, B., and Amsterdamer, Y., eds., *21st International Conference on Database Theory, ICDT 2018, March 26-29, 2018, Vienna, Austria*, volume 98 of *LIPIcs*, 4:1–4:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Artale, A.; Calvanese, D.; Kontchakov, R.; and Zakharyashev, M. 2009. The dl-lite family and relations. *J. Artif. Intell. Res.* 36:1–69.
- Baget, J.; Leclère, M.; Mugnier, M.; and Salvat, E. 2011. On rules with existential variables: Walking the decidability line. *Artif. Intell.* 175(9-10):1620–1654.
- Bárány, V.; Benedikt, M.; and ten Cate, B. 2013. Rewriting guarded negation queries. In Chatterjee, K., and Sgall, J., eds., *Mathematical Foundations of Computer Science 2013 - 38th International Symposium, MFCS 2013, Klosterneuburg, Austria, August 26-30, 2013. Proceedings*, volume 8087 of *Lecture Notes in Computer Science*, 98–110. Springer.
- Benedikt, M.; Buron, M.; Germano, S.; Kappelmann, K.;

- and Motik, B. 2022. Rewriting the infinite chase. *Proc. VLDB Endow.* 15(11):3045–3057.
- Berger, G.; Gottlob, G.; Pieris, A.; and Sallinger, E. 2022. The space-efficient core of vatalog. *ACM Trans. Database Syst.* 47(1):1:1–1:46.
- Bourgau, C.; Carral, D.; Krötzsch, M.; Rudolph, S.; and Thomazo, M. 2021. Capturing homomorphism-closed decidable queries with existential rules. In Bienvenu, M.; Lakemeyer, G.; and Erdem, E., eds., *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021, Online event, November 3-12, 2021*, 141–150.
- Bourhis, P.; Krötzsch, M.; and Rudolph, S. 2015. Reasonable highly expressive query languages - IJCAI-15 distinguished paper (honorary mention). In Yang, Q., and Wooldridge, M. J., eds., *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, 2826–2832. AAAI Press.
- Calì, A.; Gottlob, G.; and Kifer, M. 2013. Taming the infinite chase: Query answering under expressive relational constraints. *J. Artif. Intell. Res.* 48:115–174.
- Calì, A.; Gottlob, G.; and Pieris, A. 2010. Query answering under non-guarded rules in datalog+/- . In Hitzler, P., and Lukasiewicz, T., eds., *Web Reasoning and Rule Systems - Fourth International Conference, RR 2010, Bressanone/Brixen, Italy, September 22-24, 2010. Proceedings*, volume 6333 of *Lecture Notes in Computer Science*, 1–17. Springer.
- Carral, D.; Dragoste, I.; and Krötzsch, M. 2018. The combined approach to query answering in horn-alchoiq. In Thielscher, M.; Toni, F.; and Wolter, F., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018*, 339–348. AAAI Press.
- Carral, D.; González, L.; and Koopmann, P. 2019. From horn-sriq to datalog: A data-independent transformation that preserves assertion entailment. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, 2736–2743. AAAI Press.
- Dawar, A., and Kreutzer, S. 2008. On datalog vs. LFP. In Aceto, L.; Damgård, I.; Goldberg, L. A.; Halldórsson, M. M.; Ingólfssdóttir, A.; and Walukiewicz, I., eds., *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, Part II - Track B: Logic, Semantics, and Theory of Programming & Track C: Security and Cryptography Foundations*, volume 5126 of *Lecture Notes in Computer Science*, 160–171. Springer.
- Deutsch, A., and Tannen, V. 2003. Reformulation of XML queries and constraints. In Calvanese, D.; Lenzerini, M.; and Motwani, R., eds., *Database Theory - ICDT 2003, 9th International Conference, Siena, Italy, January 8-10, 2003, Proceedings*, volume 2572 of *Lecture Notes in Computer Science*, 225–241. Springer.
- Feder, T., and Vardi, M. Y. 2003. Homomorphism closed vs. existential positive. In *Proceedings of the 18th IEEE Symposium on Logic in Computer Science (LICS 2003)*, 311–320.
- Florescu, D.; Levy, A. Y.; and Suciu, D. 1998. Query containment for conjunctive queries with regular expressions. In Mendelzon, A. O., and Paredaens, J., eds., *Proceedings of the Seventeenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 1-3, 1998, Seattle, Washington, USA*, 139–148. ACM Press.
- Gottlob, G., and Schwentick, T. 2012. Rewriting ontological queries into small nonrecursive datalog programs. In Brewka, G.; Eiter, T.; and McIlraith, S. A., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Thirteenth International Conference, KR 2012, Rome, Italy, June 10-14, 2012*. AAAI Press.
- Gottlob, G.; Kikot, S.; Kontchakov, R.; Podolskii, V. V.; Schwentick, T.; and Zakharyashev, M. 2014. The price of query rewriting in ontology-based data access. *Artif. Intell.* 213:42–59.
- Gottlob, G.; Rudolph, S.; and Simkus, M. 2014. Expressiveness of guarded existential rule languages. In Hull, R., and Grohe, M., eds., *Proceedings of the 33rd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS'14, Snowbird, UT, USA, June 22-27, 2014*, 27–38. ACM.
- Grau, B. C.; Horrocks, I.; Krötzsch, M.; Kupke, C.; Magka, D.; Motik, B.; and Wang, Z. 2013. Acyclicity notions for existential rules and their application to query answering in ontologies. *J. Artif. Intell. Res.* 47:741–808.
- Hustadt, U.; Motik, B.; and Sattler, U. 2007. Reasoning in description logics by a reduction to disjunctive datalog. *J. Autom. Reason.* 39(3):351–384.
- König, M.; Leclère, M.; Mugnier, M.; and Thomazo, M. 2015. Sound, complete and minimal ucq-rewriting for existential rules. *Semantic Web* 6(5):451–475.
- Krötzsch, M.; Marx, M.; and Rudolph, S. 2019. The power of the terminating chase (invited talk). In Barceló, P., and Calautti, M., eds., *22nd International Conference on Database Theory, ICDT 2019, March 26-28, 2019, Lisbon, Portugal*, volume 127 of *LIPICs*, 3:1–3:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Krötzsch, M. 2011. Efficient rule-based inferencing for OWL EL. In Walsh, T., ed., *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, 2668–2673. IJCAI/AAAI.
- Leone, N.; Manna, M.; Terracina, G.; and Veltri, P. 2019. Fast query answering over existential rules. *ACM Trans. Comput. Log.* 20(2):12:1–12:48.
- Marnette, B. 2012. Resolution and Datalog rewriting under value invention and equality constraints. *CoRR* abs/1212.0254.

- Medeiros, C. M.; Musicante, M. A.; and da Costa, U. S. 2022. Querying graph databases using context-free grammars. *J. Comput. Lang.* 68:101089.
- Mugnier, M., and Thomazo, M. 2014. An introduction to ontology-based query answering with existential rules. In Koubarakis, M.; Stamou, G. B.; Stoilos, G.; Horrocks, I.; Kolaitis, P. G.; Lausen, G.; and Weikum, G., eds., *Reasoning Web. Reasoning on the Web in the Big Data Era - 10th International Summer School 2014, Athens, Greece, September 8-13, 2014. Proceedings*, volume 8714 of *Lecture Notes in Computer Science*, 245–278. Springer.
- Nenov, Y.; Piro, R.; Motik, B.; Horrocks, I.; Wu, Z.; and Banerjee, J. 2015. Rdflox: A highly-scalable RDF store. In Arenas, M.; Corcho, Ó.; Simperl, E.; Strohmaier, M.; d’Aquin, M.; Srinivas, K.; Groth, P.; Dumontier, M.; Heflin, J.; Thirunarayan, K.; and Staab, S., eds., *The Semantic Web - ISWC 2015 - 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11-15, 2015, Proceedings, Part II*, volume 9367 of *Lecture Notes in Computer Science*, 3–20. Springer.
- Ostropolski-Nalewaja, P.; Marcinkowski, J.; Carral, D.; and Rudolph, S. 2022. A journey to the frontiers of query rewritability. In Libkin, L., and Barceló, P., eds., *PODS ’22: International Conference on Management of Data, Philadelphia, PA, USA, June 12 - 17, 2022*, 359–367. ACM.
- Rossmann, B. 2008. Homomorphism preservation theorems. *J. ACM* 55(3):15:1–15:53.
- Rudolph, S., and Thomazo, M. 2015. Characterization of the expressivity of existential rule queries. In Yang, Q., and Wooldridge, M. J., eds., *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, 3193–3199. AAAI Press.
- Rudolph, S., and Thomazo, M. 2016. Expressivity of datalog variants - completing the picture. In Kambhampati, S., ed., *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, 1230–1236. IJCAI/AAAI Press.
- Rudolph, S.; Krötzsch, M.; and Hitzler, P. 2012. Type-elimination-based reasoning for the description logic shiqbs using decision diagrams and disjunctive datalog. *Log. Methods Comput. Sci.* 8(1).
- Urbani, J.; Krötzsch, M.; Jacobs, C. J. H.; Dragoste, I.; and Carral, D. 2018. Efficient model construction for horn logic with vlog - system description. In Galmiche, D.; Schulz, S.; and Sebastiani, R., eds., *Automated Reasoning - 9th International Joint Conference, IJCAR 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings*, volume 10900 of *Lecture Notes in Computer Science*, 680–688. Springer.