Goal-Driven Query Answering for Existential Rules with Equality

Michael Benedikt and **Boris Motik**
University of Oxford

Efthymia Tsamoura
Alan Turing Institute & University of Oxford

Abstract

Inspired by the magic sets for Datalog, we present a novel goal-driven approach for answering queries over terminating existential rules with equality (aka TGDs and EGDs). Our technique improves the performance of query answering by pruning the consequences that are not relevant for the query. This is challenging in our setting because equalities can potentially affect all predicates in a dataset. We address this problem by combining the existing singularization technique with two new ingredients: an algorithm for identifying the rules relevant to a query and a new magic sets algorithm. We show empirically that our technique can significantly improve the performance of query answering, and that it can mean the difference between answering a query in a few seconds or not being able to process the query at all.

1 Introduction

Existential rules with equality, also known as *tuple- and equality generating dependencies* (TGDs and EGDs) or Datalog$^{\pm}$ rules, extend Datalog by allowing rule heads to contain existential quantifiers and the equality predicate $\approx$. Answering a conjunctive query Q over a set of existential rules Σ and a base instance B is key to dealing with incomplete information in information systems (Fagin et al. 2005). The problem is undecidable in general, but many decidable cases are known (Baget et al. 2011b; König et al. 2015; Baget et al. 2015a; Gottlob, Manna, and Pieris 2015; Leone et al. 2012). Systems such as Llunatic (Geerts et al. 2014), RDFox (Motik et al. 2014), DLV$^{\exists}$ (Leone et al. 2012), ChaseFUN (Bonifati, Ileana, and Linardi 2017), Ontop (Calvanese et al. 2017), and Graal (Baget et al. 2015b) implement various query answering techniques. One solution to this problem is to evaluate the query in a *universal model* of $\Sigma \cup B$, and a common and practically relevant case is when a finite universal model can be computed using a *chase* procedure. Many chase variants have been proposed. Although checking chase termination is undecidable for all variants (Deutsch, Nash, and Remmel 2008; Marnette 2009), numerous sufficient *acyclicity* conditions (Marnette 2009; Krötzsch and Rudolph 2011; Grau et al. 2013) guarantee termination of at least the oblivious Skolem chase; we call such $\Sigma \cup B$ *terminating*.

Computing a universal model in full when only a specific query is to be answered may be inefficient. We experimentally show that query answers often depend only on a small subset of the universal model, particularly for queries containing constants, so the chase may perform a lot of unnecessary work. Moreover, universal models sometimes cannot be computed due to their size. In such cases, *goal-driven* query answering techniques, which take the query into account, hold the key to efficient query answering.

One possibility, implemented in systems such as Ontop (Calvanese et al. 2017) and Graal (Baget et al. 2015b), is to *rewrite* the relevant rules into a new query that can be evaluated directly on the base instance. Rewriting into first-order queries is possible for DL-lite (Calvanese et al. 2007), linear TGDs (Calì, Gottlob, and Lukasiewicz 2012), and sticky TGDs (Gottlob, Orsi, and Pieris 2014; Calì, Gottlob, and Pieris 2010), among others. Frontier-guarded (Baget et al. 2011a), weakly-guarded (Gottlob, Rudolph, and Simkus 2014), and Horn-$\mathcal{SHIQ}$ (Eiter et al. 2012) rules can be rewritten into Datalog. Rewriting approaches, however, cannot handle common properties that can be handled via the chase, such as transitivity, and they typically support only "innocuous" equalities that do not affect query answers.

In Datalog and logic programming, the *magic sets* algorithm (Bancilhon et al. 1986; Beeri and Ramakrishnan 1991) annotates the rules with *magic* atoms, which ensure that bottom-up evaluation of the magic program simulates top-down query evaluation. This influential idea has been adapted to disjunctive (Alviano et al. 2012a) and finitely recursive (Calimeri et al. 2009) programs, programs with aggregates (Alviano, Greco, and Leone 2011), and *Shy* existential rules (Alviano et al. 2012b). These approaches, however, do not handle existential rules with equality.

In this paper we present what we believe to be the first goal-driven query answering technique for terminating existential rules *with* equality. Given a set of rules Σ and a query Q, we compute a logic program P such that, for each base instance B, the answers to Q on $\Sigma \cup B$ and $P \cup B$ coincide, but processing the latter is typically much more efficient. Our approach combines existing techniques such as *singularization* (Marnette 2009) with two new ingredients: a new *relevance analysis* algorithm that identifies irrelevant rules, and a new magic sets variant that handles existential rules with equality. These two techniques are complementary: the first one prunes rules whose consequences are irrelevant to the query, and the second one prunes the irrelevant consequences of the remaining rules. Since equalities can potentially affect any predicate, both techniques are needed to efficiently identify the relevant equalities.

We have empirically evaluated our technique on a recent benchmark that includes a diverse set of existential rules (Benedikt et al. 2017). Our results show that goal-driven query answering is generally more efficient than computing the chase in full. In fact, our approach can mean the difference between success and failure: even though the chase cannot be computed in several cases, we can answer the relevant queries in a few seconds. We also show that relevance analysis alone is very effective at eliminating irrelevant rules even without equalities. Finally, we show that magic sets alone can be less efficient on queries without constants, but it greatly benefits queries with constants. A combination of both techniques usually provides the best performance.

All proofs and further experimental results are given in the appendix.

2 Preliminaries

We use the standard first-order logic notions of *variables*, *constants*, *function symbols*, *predicates*, *arity*, *terms*, *atoms*, and *formulas*, and $\approx$ is the binary *equality* predicate. Atoms $\approx(s,t)$ are *equational* and are usually written as $s \approx t$, and all other atoms are *relational*. A *fact* is a variable-free atom, an *instance* is a (possibly infinite) set of facts, and a *base instance* is a finite, function-free instance. We consider two notions of entailment: $\models$ interprets $\approx$ as an "ordinary predicate" without any special semantics, whereas $\models_\approx$ interprets $\approx$ under the usual semantics of equality without the *unique name assumption* (UNA)—that is, distinct constants can be derived equal. A theory T *satisfies UNA* if no two distinct constants a and b exist such that $T \models_\approx a \approx b$. For example, let $\varphi = A(a) \wedge a \approx b$. Then, $\varphi \models_\approx A(b)$ and $\varphi \models_\approx b \approx a$, and φ does not satisfy UNA. In contrast, $\varphi \not\models A(b)$ and $\varphi \not\models b \approx a$. We often abbreviate a tuple $t_1, \ldots, t_n$ as $\mathbf{t}$, and we often treat $\mathbf{t}$ as a set and write $t_i \in \mathbf{t}$.

A term t *occurs* in a term, atom, tuple, or set X if X contains t possibly nested inside another term; $\mathsf{vars}(X)$ is the set of variables occurring in X; and X is *ground* if $\mathsf{vars}(X) = \emptyset$. For σ a mapping of variables and/or constants to terms, $\sigma(X)$ replaces each occurrence of a term t in X with $\sigma(t)$ if the latter is defined, and σ is a *substitution* if its domain is finite and contains only variables. For μ a mapping of ground terms to ground terms, $\mu[\![X]\!]$ replaces each occurrence of a term t not nested in a function symbol with $\mu(t)$ if the latter is defined. For example, let $A = R(f(x), g(a))$; then, $\sigma(A) = R(f(b), g(c))$ for $\sigma = \{x \mapsto b, a \mapsto c\}$, and $\mu[\![A]\!] = R(f(x), h(d))$ for $\mu = \{a \mapsto b, g(a) \mapsto h(d)\}$.

Existential rules are logical implications of two forms: $\forall \mathbf{x}.[\lambda(\mathbf{x}) \rightarrow \exists \mathbf{y}.\rho(\mathbf{x},\mathbf{y})]$ is a *tuple-generating dependency* (TGD), and $\forall \mathbf{x}.[\lambda(\mathbf{x}) \rightarrow t_1 \approx t_2]$ is an *equality-generating dependency* (EGD), where $\lambda(\mathbf{x})$ and $\rho(\mathbf{x},\mathbf{y})$ are conjunctions of relational, function-free atoms with variables in $\mathbf{x}$ and $\mathbf{x} \cup \mathbf{y}$, respectively, t_1 and t_2 are variables from $\mathbf{x}$ or constants, and each variable in $\mathbf{x}$ occurs in $\lambda(\mathbf{x})$. Quantifiers $\forall \mathbf{x}$ are commonly omitted. Conjunction $\lambda(\mathbf{x})$ is the *body* of a rule, and $\rho(\mathbf{x},\mathbf{y})$ and $t_1 \approx t_2$ are its *head*. We assume that queries are defined using a *query predicate* Q that does not occur in rule bodies or under existential quantifiers. A tuple $\mathbf{a}$ of constants is an *answer* to Q on a finite set of existential rules Σ and a base instance B iff $\Sigma \cup B \models_\approx \mathsf{Q}(\mathbf{a})$.

When treating $\approx$ as "ordinary," we allow rule bodies to contain equality atoms, and we can axiomatize the "true" semantics of $\approx$ for Σ as follows. Let $\mathsf{R}(\Sigma)$ and $\mathsf{C}(\Sigma)$ contain the *reflexivity* (1) and *congruence* (2) axioms, respectively, instantiated for each n-ary predicate R in Σ distinct from $\approx$ and each $1 \leq i \leq n$. Let ST contain the *symmetry* (3) and the *transitivity* (4) axioms. We assume that each base instance contains only the predicates of Σ, since the equality axioms are then determined only by Σ. Then, for each base instance B and tuple of constants $\mathbf{a}$, we have $\Sigma \cup B \models_\approx \mathsf{Q}(\mathbf{a})$ if and only if $\Sigma \cup \mathsf{R}(\Sigma) \cup \mathsf{C}(\Sigma) \cup \mathsf{ST} \cup B \models \mathsf{Q}(\mathbf{a})$.

$$R(\ldots, x_i, \ldots) \rightarrow x_i \approx x_i \tag{1}$$

$$R(\ldots, x_i, \ldots) \wedge x_i \approx x_i' \rightarrow R(\ldots, x_i', \ldots) \tag{2}$$

$$y \approx x \rightarrow x \approx y \tag{3}$$

$$x \approx y \wedge y \approx z \rightarrow x \approx z \tag{4}$$

Our algorithms use logic programming, which we define next. A *rule* r has the form $R(\mathbf{t}) \leftarrow R_1(\mathbf{t}_1) \wedge \cdots \wedge R_n(\mathbf{t}_n)$, where $R(\mathbf{t})$ and $R_i(\mathbf{t}_i)$ are atoms possibly containing function symbols. Each variable in r must occur in some $\mathbf{t}_i$. To distinguish existential from logic programming rules, we use $\rightarrow$ for the former and $\leftarrow$ for the latter. Conjunction $\mathsf{b}(r) = R_1(\mathbf{t}_1) \wedge \cdots \wedge R_n(\mathbf{t}_n)$ is the *body* of r and we often treat it as a set, and atom $\mathsf{h}(r) = R(\mathbf{t})$ is the *head* of r. Predicate $\approx$ is always ordinary in logic programming, so R and R_i can be $\approx$. A *(logic) program* P is a finite set of rules, and it is interpreted in first-order logic as usual. Again, we assume that a query in P is defined using the predicate Q not occurring in rule bodies. For I an instance, $T_P(I)$ is the result of extending I with $\sigma(\mathsf{h}(r))$ for each rule $r \in P$ and substitution σ such that $\sigma(\mathsf{b}(r)) \subseteq I$. Finally, for B a base instance, we inductively define a sequence of interpretations where $I_0 = B$ and $I_i = T_P(I_{i-1})$ for $i > 0$; then, the *least fixpoint* of P on B is $T_P^\infty(B) = \bigcup_{i \geq 0} I_i$. It is well known that $P \cup B \models F$ iff $F \in T_P^\infty(B)$ holds for each fact F.

Our algorithms reduce query answering over existential rules to reasoning in logic programming. We eliminate existential quantifiers by computing the *Skolemization* $\mathsf{sk}(\Sigma)$ of a set Σ of existential rules. Set $\mathsf{sk}(\Sigma)$ contains each EGD of Σ as a logic programming rule and, for each TGD $\tau = \lambda(\mathbf{x}) \rightarrow \exists \mathbf{y}.\rho(\mathbf{x},\mathbf{y}) \in \Sigma$ and each $R(\mathbf{t}) \in \rho(\mathbf{x},\mathbf{y})$, set $\mathsf{sk}(\Sigma)$ contains the rule $\sigma(R(\mathbf{t})) \leftarrow \lambda(\mathbf{x})$ where σ is a substitution mapping each variable $y \in \mathbf{y}$ to $f_{\tau,y}(\mathbf{x}')$ for $\mathbf{x}' = \mathsf{vars}(\lambda(\mathbf{x})) \cap \mathsf{vars}(\rho(\mathbf{x},\mathbf{y}))$ and $f_{\tau,y}$ a fresh function symbol unique for τ and y. Let $P = \mathsf{sk}(\Sigma)$; if Σ and B do not contain $\approx$, then for each predicate R and tuple $\mathbf{a}$ of constants, we have $\Sigma \cup B \models R(\mathbf{a})$ iff $P \cup B \models R(\mathbf{a})$. If Σ or B contains $\approx$, we can axiomatize equality using axioms $\mathsf{R}(P)$, $\mathsf{C}(P)$, and ST defined analogously to (1)–(4); then, $P' = P \cup \mathsf{R}(P) \cup \mathsf{C}(P) \cup \mathsf{ST}$ captures the intended semantics of $\approx$, and $\Sigma \cup B \models_\approx R(\mathbf{a})$ iff $P' \cup B \models R(\mathbf{a})$.

Now let P and P' be as in the previous paragraph. We could answer queries over such P' by computing $T_{P'}^\infty(B)$

and evaluating Q on it, but this is inefficient even when P is just a Datalog program (Motik et al. 2015) since firing congruence rules can be prohibitively expensive. The *chase for logic programs* offers a more efficient method for reasoning with $P' \cup B$ by efficiently computing a representation of $T^{\infty}_{P'}(B)$. It is applicable if P does not contain constants, function symbols, or $\approx$ in the rule bodies. The algorithm constructs a sequence of pairs $\langle I_i, \mu_i \rangle$, $i \geq 0$, where I_i is an instance and μ_i maps ground terms to ground terms. The algorithm initializes I_0 to a normalized version of B where all constants reachable by $\approx$ in B are replaced by a representative, and it records these replacements in μ_0. For each $i > 0$, the chase selects a rule $r \in P$ and a substitution σ with $\sigma(\mathsf{b}(r)) \subseteq I_{i-1}$ and (i) if $\sigma(\mathsf{h}(r))$ is of the form $s \approx t$ and $s \neq t$, then one term, say s, is selected as the *representative*, and I_i and μ_i are obtained from I_{i-1} and μ_{i-1} by replacing t with s and setting $\mu_i(t) = s$; and (ii) if $\sigma(\mathsf{h}(r))$ does not contain $\approx$ and $\sigma(\mathsf{h}(r)) \notin I_{i-1}$, then $\mu_i = \mu_{i-1}$ and $I_i = I_{i-1} \cup \{\mu_i[\![\sigma(\mathsf{h}(r))]\!]\}$. The computation proceeds until no rule is applicable and then returns the final pair $\langle I_n, \mu_n \rangle$. If the representatives are always chosen as smallest in an arbitrary, but fixed well-founded order on ground terms, then the result is unique for P and B and it is called the *chase* of P on B, written $\mathsf{chase}(P, B)$. The following properties of the chase are well known (Benedikt et al. 2017).

Proposition 1. *For each program P, base instance B, $\mathsf{chase}(P, B) = \langle I, \mu \rangle$, and $P' = P \cup \mathsf{R}(P) \cup \mathsf{C}(P) \cup \mathsf{ST}$, (i) $\mu(t_1) = \mu(t_2)$ if and only if $P' \cup B \models t_1 \approx t_2$, for all ground terms t_1 and t_2, and (ii) $P' \cup B \models R(\mathbf{t})$ if and only if $R(\mu[\![\mathbf{t}]\!]) \in I$, for each ground relational atom $R(\mathbf{t})$.*

Intuitively, $\mu(t)$ is a unique *representative* of each ground term t, and the chase maintains and propagates facts only among the representative facts of $T^{\infty}_{P'}(B)$, instead of naïvely firing congruence rules. The algorithm is used in systems such as Llunatic and RDFox (Benedikt et al. 2017).

When reasoning with equality, an important question is whether programs are allowed to equate constants. We say that P and B *satisfy UNA* if $P' \cup B \models a \approx b$ implies $a = b$. Our algorithms do not require UNA to be satisfied, but certain steps can be optimized if we know that UNA is satisfied.

3 Motivation and Overview

To understand the challenges of goal-driven query answering over existential rules with $\approx$, let Σ^{ex} consist of (5)–(9).

$$A(x) \wedge R(x, y) \rightarrow \mathsf{Q}(x) \tag{5}$$

$$S(x, z) \rightarrow \exists y.R(x, y) \tag{6}$$

$$R(x, y) \wedge S(x, x') \wedge R(x', y') \rightarrow y \approx y' \tag{7}$$

$$B(x) \rightarrow \exists y.T(x, y) \wedge A(y) \tag{8}$$

$$T(x, y) \rightarrow x \approx y \tag{9}$$

Let $B^{ex} = \{B(a_1)\} \cup \{S(a_{i-1}, a_i) \mid 1 < i \leq n\}$. One can check that $\Sigma^{ex} \cup B^{ex} \models \mathsf{Q}(a_i)$ holds only for $i = 1$; however, all bottom-up techniques known to us will "fire" (6) and (7) for all a_i. In logic programming, goal-driven or top-down approaches, such as SLD resolution, start from the query and search for proofs backwards. The magic sets algorithm transforms a program so that evaluating the result bottom-up mimics top-down evaluation. These approaches are not directly applicable to existential rules, but we can apply them to the program $P' = P \cup \mathsf{R}(P) \cup \mathsf{C}(P) \cup \mathsf{ST}$, obtained by Skolemizing TGDs as $P = \mathsf{sk}(\Sigma^{ex})$ and then axiomatizing equality. This, however, is inefficient since the congruence axioms introduce many redundant proofs. In particular, Skolemizing (6) produces $R(x, f(x)) \leftarrow S(x, z)$. By rule (7), we have $P' \models f(a_{i-1}) \approx f(a_i)$ for $1 < i \leq n$ so, by the reflexivity, symmetry, and transitivity axioms for $\approx$, we have $P' \models f(a_i) \approx f(a_j)$ for $1 \leq i, j \leq n$. Hence, by the congruence axioms, we have $P' \models R(a_i, f(a_j))$. Thus, $P' \models \mathsf{Q}(a_1)$ has (at least) n proofs, where the first step uses a ground rule instance $\mathsf{Q}(a_1) \leftarrow A(a_1) \wedge R(a_1, f(a_i))$ for each $1 \leq i \leq n$. The magic sets algorithm will explore all of these proofs, which is very expensive. In contrast, our technique can answer the query by considering this rule instantiated only for $i = 1$ (see Example 5).

We present an approach that gives the benefits of top-down approaches, while radically pruning the set of considered proofs. Our approach has additional benefits. It does not require UNA, but certain steps can be optimized if $\Sigma \cup B$ satisfies UNA (e.g., if an earlier UNA check succeeded). Moreover, it preserves chase termination, and it includes an optimized magic set transformation using the symmetry of equality to greatly reducing the number of output rules.

Algorithm 1 Compute the answers to query Q over a finite set of existential rules Σ and a base instance B

1: $\Sigma_1 := \mathsf{sg}(\Sigma)$
2: $P_2 := \mathsf{sk}(\Sigma_1)$
3: $P_3 := \mathsf{relevance}(P_2, B)$
4: $P_4 := \mathsf{magic}(P_3)$
5: $P_5 := \mathsf{defun}(P_4)$
6: $P_6 := \mathsf{desg}(P_5)$
7: $\langle I, \mu \rangle := \mathsf{chase}(P_6, B)$
8: **for each** $\mathsf{Q}(\mathbf{a}) \in I$ where $\mathbf{a}$ are constants **do**
9:     **output** each tuple of constants $\mathbf{b}$ with $\mu[\![\mathbf{b}]\!] = \mathbf{a}$

Our technique is presented in the pipeline shown in Algorithm 1. Instead of axiomatizing equality, we first apply *singularization* (line 1), a well-known transformation that makes all relevant equalities explicit (Marnette 2009; ten Cate et al. 2009), and then we convert the result to a logic program using Skolemization (line 2). Next, we apply a relevance analysis algorithm (line 3) that identifies the rules relevant to the query. We next apply the magic sets transformation optimized for $\approx$ (line 4); the removal of irrelevant equality atoms during relevance analysis ensures that this step produces a smaller program. Finally, we remove the function symbols (line 5) and equalities (line 6) from rule bodies, obtaining a program that can be safely evaluated using the chase for logic programs (lines 7–9). We explain the components in detail in the following sections.

4 Singularization

Singularization is an alternative to congruence axioms.

Definition 1. *A* singularization *of an existential rule τ is obtained from τ by exhaustively (i) replacing each occurrence* *of a constant c in a relational body atom with a fresh variable x and adding atom $x \approx c$ to the body, and (ii) for each variable x occurring at least twice in (not necessarily distinct) relational body atoms, replacing one such occurrence with a fresh variable x' and adding atom $x' \approx x$ to the body.*

A singularization *of a set of existential rules Σ defining the query predicate* Q *is obtained by replacing each TGD of the form $\varphi \rightarrow \mathsf{Q}(x_1, \dots, x_n)$ with* (10) *for $x'_1, \dots, x'_n$ fresh variables, and then singularizing all existential rules.*

$$\varphi \wedge \bigwedge_{i=1}^{n} x_i \approx x'_i \rightarrow \mathsf{Q}(x'_1, \dots, x'_n) \tag{10}$$

Example 1. *On Σ^{ex}, singularization leaves* (6), (8)*, and* (9) *intact since their bodies do not contain repeated variables. Rules* (5) *and* (7) *are singularized as* (11) *and* (12).

$$A(x'') \wedge x \approx x'' \wedge R(x, y) \wedge x \approx x' \rightarrow \mathsf{Q}(x') \tag{11}$$

$$\begin{aligned} R(x, y) \wedge x \approx x'' \wedge \\ S(x'', x') \wedge x' \approx x''' \wedge R(x''', y') \rightarrow y \approx y' \end{aligned} \tag{12}$$

The result of singularization is not unique: (5) could also produce $A(x) \wedge x \approx x'' \wedge R(x'', y) \wedge x \approx x' \rightarrow \mathsf{Q}(x')$. In our approach, we let $\mathsf{sg}(\Sigma)$ be any singularization of Σ.

Singularization highlights the relevant equalities originating from joins, which in turn preserves all query answers without relying on congruence axioms: for each base instance B and tuple $\mathbf{a}$ of constants, we have $\Sigma \cup B \models_{\approx} \mathsf{Q}(\mathbf{a})$ if and only if $\mathsf{sg}(\Sigma) \cup \mathsf{R}(\Sigma) \cup \mathsf{ST} \cup B \models \mathsf{Q}(\mathbf{a})$. Singularization still relies on reflexivity axioms. As an important optimization, we prove that these do not need to be analyzed in the remaining steps of our pipeline, which ensures that our pipeline produces smaller, more efficient programs. To achieve this, we show that our transformations produce rules satisfying the following condition.

Definition 2. *A rule r is* $\approx$-safe *if, for each equality atom $A \in \mathsf{b}(r)$, the atom is of the form $x \approx y$ or $x \approx s$ for s a ground term, and $\mathsf{vars}(A) \cap \mathbf{t}_i \neq \emptyset$ for some relational atom $R_i(\mathbf{t}_i) \in \mathsf{b}(r)$. A program is* $\approx$-safe *all its rules are $\approx$-safe.*

Intuitively, $\approx$-safety ensures that each fact $t \approx t$ matching a body atom of a rule r can be derived from another relational body atom of r. Thus, we do not need to pass the reflexivity axiom as input to the steps of our pipeline in order to determine which of these are pertinent to Q. Instead, the pertinent reflexivity axioms are determined directly by the predicates occurring in the result of each pipeline step.

We apply Skolemization after singularization to eliminate existential quantifiers.

Example 2. *In our running example, only* (6) *and* (8) *contain existential quantifiers, so they are replaced by* (14)*, and* (16) *and* (17)*; all other rules are reinterpreted as logic programming rules. Program P_2 contains rules* (13)–(18).

$$\mathsf{Q}(x') \leftarrow A(x'') \wedge x \approx x'' \wedge R(x, y) \wedge x \approx x' \tag{13}$$

$$R(x, f(x)) \leftarrow S(x, z) \tag{14}$$

$$\begin{aligned} y \approx y' \leftarrow R(x, y) \wedge x \approx x'' \wedge S(x'', x') \wedge \\ x' \approx x''' \wedge R(x''', y') \end{aligned} \tag{15}$$

$$T(x, g(x)) \leftarrow B(x) \tag{16}$$

$$A(g(x)) \leftarrow B(x) \tag{17}$$

$$x \approx y \leftarrow T(x, y) \tag{18}$$

The answers to Q *on $\Sigma \cup B$ and $P_2 \cup \mathsf{R}(P_2) \cup \mathsf{ST} \cup B$ coincide on each base instance B, but the absence of congruence axioms considerably reduces the number of proofs: EGD* (7) *still ensures $P_2 \models f(a_i) \approx f(a_j)$ for $1 \leq i, j \leq n$, but $P_2 \models R(a_i, f(a_j))$ holds only for $i = j$. Thus, the only remaining proof of $P_2 \models \mathsf{Q}(a_1)$ is via the ground instance $\mathsf{Q}(a_1) \leftarrow C(a_1) \wedge R(a_1, f(a_1))$, which benefits all goal-driven techniques, including magic sets. Also, none of the facts derived by EGD* (7) *contribute to a proof of $\mathsf{Q}(a_1)$, so singularization makes EGD redundant; in Section 5 we present way to detect and eliminate such rules.*

5 Relevance Analysis

The next step of the pipeline eliminates rules all of whose consequences are irrelevant to Q. The idea is to homomorphically embed B into a much smaller instance B' called an *abstraction* of B. If B' is sufficiently small, we can analyze ways to derive answers to Q on B'; since homomorphism composition is a homomorphism, this will uncover all ways to derive an answer to Q on the original base instance B.

Definition 3. *A base instance B' is an* abstraction *of a base instance B w.r.t. a program P if there exists a homomorphism η from B to B' preserving the constants in P—that is, η maps constants to constants such that $\eta(B) \subseteq B'$ and $\eta(c) = c$ for each constant c occurring in P.*

To abstract B into B', we can use the *critical instance* for B: for C the set of constants of P and $*$ a fresh constant, we let B' contain $R(\mathbf{a})$ for each n-ary predicate R occurring in B and each $\mathbf{a} \in (C \cup \{*\})^n$. We can further refine the abstraction if predicates are *sorted* (i.e., each predicate position is associated with a sort such as strings or integers): we introduce a distinct fresh constant $*_i$ per sort, and we form B' as above while honoring the sorting requirements.

Algorithm 2 takes an $\approx$-safe program P and a base instance B, and it returns the rules relevant to answering Q on B. It selects an abstraction B' of B w.r.t. P (line 1), computes the consequences I of P on B' (line 2), and identifies the rules of P contributing to the answers of Q on B' by a form of backward chaining. It initializes the "ToDo" set $\mathcal{T}$ to all homomorphic images of the answers to Q on B' (line 3) and then iteratively explores $\mathcal{T}$ (lines 5–12). In each iteration, it extracts a fact F from $\mathcal{T}$ (line 6) and then identifies each rule r and substitution ν matching the head of r to F and the body of r in I (line 7). Such ν captures ways of deriving a fact represented by F from B via r, so r is added to the set $\mathcal{R}$ of relevant rules if such μ exists (line 8). Finally, the matched body atoms must be derivable as well, so they are all added to $\mathcal{T}$ (line 11). The "done" set $\mathcal{D}$ ensures that each fact is added to $\mathcal{T}$ just once, which ensures termination.

We can optimize the algorithm if $P \cup B$ is known to satisfy UNA (e.g., if an earlier UNA check was conducted). If ν matches an atom $x \approx t \in \mathsf{b}(r)$ as $c \approx c$, the corresponding derivation from P and B necessarily matches $x \approx t$ to $d \approx d$ for some constant d; due to $\approx$-safety, we can derive $d \approx d$ using reflexivity axioms, so we do not need to examine other proofs for $d \approx d$ (line 10). Moreover, if all matches

Algorithm 2 relevance(P, B)

1: **choose** an abstraction B' of B w.r.t. P
2: $I := T^{\infty}_{P'}(B')$ for $P' = P \cup \mathsf{R}(P) \cup \mathsf{ST}$
3: $\mathcal{D} := \mathcal{T} := \{\mathsf{Q}(\mathbf{a}) \in I \mid \mathbf{a}$ is a tuple of constants$\}$
4: $\mathcal{R} := \emptyset$ and $\mathcal{B} := \emptyset$
5: **while** $\mathcal{T} \neq \emptyset$ **do**
6:     **choose and remove** some fact F from $\mathcal{T}$
7:     **for each** $r \in P \cup \mathsf{ST}$ and substitution ν such that $\nu(\mathsf{h}(r)) = F$ and $\nu(\mathsf{b}(r)) \subseteq I$ **do**
8:         **if** $r \notin \mathcal{R} \cup \mathsf{ST}$ **then add** r to $\mathcal{R}$
9:         **for each** $G_i \in \nu(\mathsf{b}(r))$ **do**
10:             **if** G_i is not of the form $c \approx c$ with c a constant, or $P \cup B$ is not known to satisfy UNA **then**
11:                 **if** $G_i \notin \mathcal{D}$ **then add** G_i to $\mathcal{T}$ and $\mathcal{D}$
12:                 **if** G_i is an equality **then add** $\langle r, i\rangle$ to $\mathcal{B}$
13: **if** P and B and known to satisfy UNA **then**
14:     **for each** $r \in \mathcal{R}$ and each i-th atom of $\mathsf{b}(r)$ of the form $x \approx t$ with t a term and $\langle r, i\rangle \notin \mathcal{B}$ **do**
15:         **remove** $x \approx t$ from r and replace x with t in r
16: **return** $\mathcal{R}$

of $x \approx t$ are of such a form, then we can replace x with t in r and inform the subsequent magic sets transformation step that no equalities are relevant to this atom. To this end, Algorithm 2 maintains a set $\mathcal{B}$ of "blocked" body equality atoms that records all body equality atoms that can be matched to an equality not of the form $d \approx d$ (line 12). After considering all possibilities for deriving certain answers, body equality atoms that have not been blocked are removed (lines 13–15).

The computation of I in line 2 may not terminate in general, but we shall apply Algorithm 2 in our pipeline only in cases where termination is guaranteed. Theorem 1 shows that, in that case, the query answers remain preserved.

Theorem 1. *For each $\approx$-safe program P defining the query predicate Q and base instance B where line 2 of Algorithm 2 terminates, program $\mathcal{R} =$ relevance(P, B) is $\approx$-safe, and, for each tuple $\mathbf{a}$ of constants, $P \cup \mathsf{R}(P) \cup \mathsf{ST} \cup B \models \mathsf{Q}(\mathbf{a})$ if and only if $\mathcal{R} \cup \mathsf{R}(\mathcal{R}) \cup \mathsf{ST} \cup B \models \mathsf{Q}(\mathbf{a})$.*

If computing $T^{\infty}_{P'}(B')$ in line 1 is difficult (as is the case in some of our experiments), we can replace each term $f(\mathbf{x})$ in P with a fresh constant c_f unique for f. This does not affect the algorithm's correctness, since $T^{\infty}_{P'}(B)$ can still be homomorphically embedded into $T^{\infty}_{P'}(B')$. Moreover, the resulting program then does not contain function symbols, and so the computation in line 2 necessarily terminates.

Example 3. *In our running example, UNA holds on Σ^{ex} and B^{ex}, and we shall assume that this is known in advance. Moreover, we shall take B' to be the critical instance, containing facts $B(*)$ and $S(*, *)$. Computing the least fixpoint in line 2 of Algorithm 2 produces the following instance I.*

$$\begin{array}{cccc} B(*) & S(*,*) & R(*, f(*)) & f(*) \approx f(*) \\ T(*, g(*)) & A(g(*)) & * \approx * & * \approx g(*) \\ g(*) \approx * & g(*) \approx g(*) & \mathsf{Q}(*) & \mathsf{Q}(g(*)) \end{array}$$

Algorithm 2 starts by considering $\mathsf{Q}(*)$. Matching the fact to the head of (13) *and evaluating the body in I produces the ground rule instance*

$$\mathsf{Q}(*) \leftarrow A(g(*)) \wedge * \approx g(*) \wedge R(*, f(*)) \wedge * \approx *;$$

thus, rule (13) *is identified as relevant. UNA is known to hold, so atom $* \approx *$ is not considered any further, but the algorithm must consider the remaining body atoms. Matching $g(*) \approx f(*)$ to the head of* (18) *and evaluating the body in I produces the ground rule instance $* \approx g(*) \leftarrow T(*, g(*))$, so rule* (18) *is identified as relevant; moreover, matching $T(*, g(*))$ to the head of* (16) *produces the ground rule instance $T(*, g(*)) \leftarrow B(*)$, so rule* (16) *is identified as relevant as well. In contrast, matching $* \approx g(*)$ to the head of* (15) *produces query*

$$R(x, *) \wedge \approx x'' \wedge S(x'', x') \wedge x' \approx x''' \wedge R(x''', g(*)),$$

which has no matches in I; thus, rule (15) *is not added to $\mathcal{R}$. Next, matching $R(*, f(*))$ to the head of* (14) *and evaluating the body in I produces the ground rule instance $R(*, f(*)) \leftarrow S(*, *)$; thus, rule* (14) *is identified as relevant. Finally, $B(*)$ and $S(*, *)$ do not match any rule head, so the algorithm terminates. Thus, the algorithm returns all rules apart from* (15)*. Moreover, atom $x \approx x''$ in* (13) *is matched to $* \approx f(*)$ so it cannot be removed—that is, this equality is relevant. In contrast, atom $x \approx x'$ in* (13) *is matched* only *to $* \approx *$; since we know that UNA holds, this equality is irrelevant and it is removed in lines 13–15 of Algorithm 2. Consequently, the algorithm returns program P_3 consisting of rules* (19)–(23).

$$\mathsf{Q}(x) \leftarrow A(x'') \wedge x \approx x'' \wedge R(x, y) \tag{19}$$
$$R(x, f(x)) \leftarrow S(x, z) \tag{20}$$
$$T(x, g(x)) \leftarrow B(x) \tag{21}$$
$$A(g(x)) \leftarrow B(x) \tag{22}$$
$$x \approx y \leftarrow T(x, y) \tag{23}$$

6 Magic Sets for Existential Rules with $\approx$

We now present our variant of the magic sets transformation. Our technique also mimics top-down evaluation, but with a specialized treatment of equality that takes the symmetry of $\approx$ into account and further prunes the set of proofs for facts of the form $t \approx t$. In particular, singularization critically relies on reflexivity axioms, and passing these to the magic sets algorithm would significantly blow up the resulting rule set. The $\approx$-safety of the rules implies that our magic transformation does not need to be applied to the reflexivity rules, which results in a much more efficient magic program.

We follow Beeri and Ramakrishnan (1991) in defining adornments and magic predicates for predicates other than $\approx$, but, as we discuss shortly, we optimize these notions for $\approx$. Intuitively, an adornment identifies which arguments of an atom will be bound by sideways information passing, and a magic predicate will "collect" the passed arguments.

Definition 4. *An* adornment *for an n-ary predicate R other than $\approx$ is a string α of length n over alphabet* b *("bound") and* f *("free"), and m^{α}_{R} is a fresh* magic *predicate unique for R and α with arity equal to the number of* b*-symbols in α. An* adornment *for $\approx$ has the form* bb, bf, *or* fb, *and $m^{\mathsf{bb}}_{\approx}$ and $m^{\mathsf{b}\wr\mathsf{f}}_{\approx}$ are fresh* magic *predicates for $\approx$ of arity two and one, respectively. For α an adornment of length n and $\mathbf{t}$ an n-tuple of terms, $\mathbf{t}^{\alpha}$ contains in the same relative order each $t_i \in \mathbf{t}$ for which the i-th element of α is* b.

Definition 4 takes into account that, if one argument of an equality atom is bound, the other argument will also be bound due to the symmetry of $\approx$. Thus, predicate $m_{\approx}^{\mathsf{b}\wr\mathsf{f}}$ is used for both bf and fb, where notation $\mathsf{b}\wr\mathsf{f}$ stresses that the positions of b and f are interchangeable. Moreover, at least one argument of an equality atom must be bound, so $\approx$ cannot be adorned by ff. Definition 5 introduces a sideways information passing strategy (SIPS), which determines how information is propagated through the rule bodies. Function reorder reorders the rule bodies to maximize information passing, and function adorn decides which arguments of an atom should be bound given a set of available bindings.

Definition 5. *A* sideways information passing strategy *consists of the following two functions.*

For φ a conjunction of atoms and T a set of terms, $\mathsf{reorder}(\varphi, T)$ *returns an ordering $\langle R_1(\mathbf{t}_1), \dots, R_n(\mathbf{t}_n)\rangle$ of the conjuncts of φ such that, for each equality atom $R_i(\mathbf{t}_i)$ of the form $x \approx y$ or $x \approx s$ where s is ground, $z \in \mathsf{vars}(\mathbf{t}_i)$ exists such that $z \in T$ or $z \in \mathbf{t}_j$ for some $j < i$ with $R_j \neq \approx$.*

For $R(t_1, \dots, t_k)$ an atom and V a set of variables, $\mathsf{adorn}(R(t_1, \dots, t_k), V)$ *returns an adornment α for R such that* $\mathsf{vars}(t_j) \subseteq V$ *if the j-th element of α is* b.

Algorithm 3 implements the magic sets transformation optimized for $\approx$. It initializes the "ToDo" set $\mathcal{T}$ with the magic predicate for the query (line 1) and processes $\mathcal{T}$ iteratively. For each magic predicate m_R^α in $\mathcal{T}$ (line 3), it identifies each rule r that can derive R (line 4). Adornment $\alpha = \mathsf{b}\wr\mathsf{f}$ is processed as both bf and fb (lines 6–7), and each other α is processed as is (line 9). In all cases, the algorithm produces the *modified rule* by restricting the body of r by the magic predicate corresponding to the head predicate (line 12), and then it reorders (line 13) and processes (lines 14–19) the body of r. For each body atom $R_i(\mathbf{t}_i)$ with R_i occurring in the head of P, the algorithm uses the SIPS to determine an adornment γ identifying the bound arguments (line 15), and it generates the *magic rule* that populates $m_{\approx}^{\mathsf{b}\wr\mathsf{f}}$ or $m_{R_i}^{\gamma}$ with the bindings for R_i (line 18). The algorithm takes into account that $m_{\approx}^{\mathsf{b}\wr\mathsf{f}}$ captures both bf and fb (line 16–17). The magic rule would not be $\approx$-safe if $R_i(\mathbf{t}_i)$ were an equality atom with no arguments bound, which, at the end of our pipeline, could produce a rule with variables occurring the head but not the body. Thus, Definition 4 does not introduce the ff adornment for $\approx$, and Definition 5 requires at least one argument of each equality atom to be bound. Finally, the magic predicate must also be processed (line 19), and the "done" set $\mathcal{D}$ ensures that this happens just once.

Theorem 2 shows that Algorithm 3 preserves $\approx$-safety and query answers, and Theorem 3 shows that it also preserves chase termination. The latter does not hold for all programs with function symbols; for example, transforming $A(x) \leftarrow A(f(x))$ can produce the nonterminating rule $m_A^f(f(x)) \leftarrow m_A^f(x)$. However, in Algorithm 1, the magic sets are applied in line 4 only to programs P_3 that do not contain function symbols in the body, which suffices to show that the heads of the magic rules are function-free and cannot derive terms of unbounded depth, and therefore the transformation does not affect termination.

Algorithm 3 $\mathsf{magic}(P)$

1: $\mathcal{D} := \mathcal{T} := \{m_{\mathsf{Q}}^{\alpha}\}$ and $\mathcal{R} := \{m_{\mathsf{Q}}^{\alpha} \leftarrow\}$, $\quad \alpha = \mathsf{f} \cdots \mathsf{f}$
2: **while** $\mathcal{T} \neq \emptyset$ **do**
3:     **choose and remove** some m_R^α from $\mathcal{T}$
4:     **for each** $r \in P \cup \mathsf{ST}$ such that $\mathsf{h}(r) = R(\mathbf{t})$ **do**
5:         **if** $R = \approx$ and $\alpha = \mathsf{b}\wr\mathsf{f}$ **then**
6:             $\mathsf{process}(r, \alpha, \mathsf{bf})$
7:             $\mathsf{process}(r, \alpha, \mathsf{fb})$
8:         **else**
9:             $\mathsf{process}(r, \alpha, \alpha)$
10: **return** $\mathcal{R}$

11: **procedure** $\mathsf{process}(r, \alpha, \beta)$ where $\mathsf{h}(r) = R(\mathbf{t})$
12:     **if** $r \notin \mathsf{ST}$ **then add** $\mathsf{h}(r) \leftarrow m_R^\alpha(\mathbf{t}^\beta) \wedge \mathsf{b}(r)$ to $\mathcal{R}$
13:     $\langle R_1(\mathbf{t}_1), \dots, R_n(\mathbf{t}_n)\rangle := \mathsf{reorder}(\mathsf{b}(r), \mathbf{t}^\beta)$
14:     **for** $1 \leq i \leq n$ s.t. R_i is $\approx$ or it occurs in P in a head **do**
15:         $\gamma := \mathsf{adorn}(R_i(\mathbf{t}_i), \mathsf{vars}(\mathbf{t}^\beta) \cup \bigcup_{j=1}^{i-1} \mathsf{vars}(\mathbf{t}_j))$
16:         **if** $R_i = \approx$ and $\gamma \in \{\mathsf{bf}, \mathsf{fb}\}$ **then** $S := m_{\approx}^{\mathsf{b}\wr\mathsf{f}}$
17:         **else** $S := m_{R_i}^{\gamma}$
18:         **add** $S(\mathbf{t}_i^\gamma) \leftarrow m_R^\alpha(\mathbf{t}^\beta) \wedge \bigwedge_{j=1}^{i-1} R_j(\mathbf{t}_j)$ to $\mathcal{R}$
19:         **if** $S \notin \mathcal{D}$ **then add** S to $\mathcal{T}$ and $\mathcal{D}$

Note: please remember that $t_1 \approx t_2$ can also be written as $\approx(t_1, t_2)$, so $R(\mathbf{t})$ and $R_i(\mathbf{t}_i)$ can be equality atoms.

Theorem 2. *For each $\approx$-safe program P defining the query predicate* Q*, program $\mathcal{R} = \mathsf{magic}(P)$ is $\approx$-safe; moreover, for each base instance B and each tuple $\mathbf{t}$ of ground terms,*

$$P \cup \mathsf{R}(P) \cup \mathsf{ST} \cup B \models \mathsf{Q}(\mathbf{t}) \quad \textit{iff}$$
$$\mathcal{R} \cup \mathsf{R}(\mathcal{R}) \cup \mathsf{ST} \cup B \models \mathsf{Q}(\mathbf{t}).$$

Theorem 3. *Let P be a program where the body of each rule is function-free, and let $P_1 = P \cup \mathsf{R}(P) \cup \mathsf{ST}$ and $P_2 = \mathcal{R} \cup \mathsf{R}(\mathcal{R}) \cup \mathsf{ST}$ for $\mathcal{R} = \mathsf{magic}(P)$. For each base instance B, if $T_{P_1}^{\infty}(B)$ is finite, then $T_{P_2}^{\infty}(B)$ is finite as well.*

Example 4. *Applying Algorithm 3 to P_3 produces program P_4 consisting of rules* (24)–(36)*. Horizontal lines separate the rules produced in each invocation of* process*. Please note that* (23) *produces* (30) *and* (31) *when* $\mathsf{b}\wr\mathsf{f}$ *is interpreted as* bf*, and* (30) *and* (31) *when* $\mathsf{b}\wr\mathsf{f}$ *is interpreted as* bf.

$$m_{\mathsf{Q}}^{\mathsf{f}} \leftarrow \tag{24}$$

$$\mathsf{Q}(x) \leftarrow m_{\mathsf{Q}}^{\mathsf{f}} \wedge A(x'') \wedge x \approx x'' \wedge R(x, y) \tag{25}$$
$$m_A^{\mathsf{f}} \leftarrow m_{\mathsf{Q}}^{\mathsf{f}} \tag{26}$$
$$m_{\approx}^{\mathsf{b}\wr\mathsf{f}}(x'') \leftarrow m_{\mathsf{Q}}^{\mathsf{f}} \wedge A(x'') \tag{27}$$
$$m_R^{\mathsf{bf}}(x) \leftarrow m_{\mathsf{Q}}^{\mathsf{f}} \wedge A(x'') \wedge x \approx x'' \tag{28}$$

$$A(g(x)) \leftarrow m_A^{\mathsf{f}} \wedge B(x) \tag{29}$$

$$x \approx y \leftarrow m_{\approx}^{\mathsf{b}\wr\mathsf{f}}(x) \wedge T(x, y) \tag{30}$$
$$m_T^{\mathsf{bf}}(x) \leftarrow m_{\approx}^{\mathsf{b}\wr\mathsf{f}}(x) \tag{31}$$

$$x \approx y \leftarrow m_{\approx}^{\mathsf{b}\wr\mathsf{f}}(y) \wedge T(x, y) \tag{32}$$
$$m_T^{\mathsf{fb}}(y) \leftarrow m_{\approx}^{\mathsf{b}\wr\mathsf{f}}(y) \tag{33}$$

$$T(x, g(x)) \leftarrow m_T^{\mathsf{bf}}(x) \wedge B(x) \tag{34}$$

$$T(x, g(x)) \leftarrow m_T^{\mathsf{fb}}(g(x)) \wedge B(x) \tag{35}$$

$$R(x, f(x)) \leftarrow m_R^{\mathsf{bf}}(x) \wedge S(x, z) \tag{36}$$

7 Final Transformations

The final steps ensure that the resulting program can be evaluated efficiently using the chase for logic programs, which, as explained in Section 2, can handle only programs with no constants, function symbols, and $\approx$ in the rule bodies. The magic sets transformation can introduce body atoms with function symbols, so Definition 6 removes these by introducing fresh predicates. Proposition 2 shows the query answers remain preserved since the fresh predicates can always be interpreted to reflect the structure of the ground functional terms encountered during the chase.

Definition 6. *Program* $\mathsf{defun}(P)$ *is obtained from a program* P *by exhaustively applying the following steps.*

1. *In the body of each rule, replace each occurrence of a constant* c *with a fresh variable* z_c *unique for* c*, add atom* $\mathsf{F}_c(z_t)$ *to the body, and add the rule* $\mathsf{F}_c(c) \leftarrow$*.*
2. *In the body of each rule, replace each occurrence of a term* t *of the form* $f(\mathbf{s})$ *with a fresh variable* z_t *unique for* t*, and add atom* $\mathsf{F}_f(\mathbf{s}, z_t)$ *to the body.*
3. *For each rule* r *and each term of the form* $f(\mathbf{s})$ *occurring in* $\mathsf{h}(r)$ *with* f *a function symbol considered in the second step, add the rule* $\mathsf{F}_f(\mathbf{s}, f(\mathbf{s})) \leftarrow \mathsf{b}(r)$*.*

Proposition 2. *For each program* P *and* $P' = \mathsf{defun}(P)$*, base instance* B*, predicate* R *not of the form* F_f*, and tuple* $\mathbf{t}$ *of ground terms,* $P \cup \mathsf{R}(P) \cup \mathsf{ST} \cup B \models R(\mathbf{t})$ *if and only if* $P' \cup \mathsf{R}(P') \cup \mathsf{ST} \cup B \models R(\mathbf{t})$*.*

Definition 7 reverses the effects of singularization and removes all body equality atoms. As a side-effect, this reduces the number of rule variables, which simplifies rule matching.

Definition 7. *The* desingularization *of a rule is obtained by repeatedly removing each body atom of the form* $x \approx t$ *while replacing* x *with* t *everywhere in the rule. For* P *a program,* $\mathsf{desg}(P)$ *contains a desingularization of each rule of* P*.*

We evaluate the final program using the chase for logic programs, which captures the effects of congruence axioms. Note, however, that program P_5 from Algorithm 1 contains fresh predicates introduced by the magic sets transformation and the elimination of function symbols, to which the chase will (implicitly) apply congruence axioms as well. Theorem 4 shows that this preserves the query answers, and its proof is not trivial: adding congruence axioms to a program produces new consequences, so the proof depends on the fact program P_5 was obtained as shown in Algorithm 1.

Theorem 4. *For each finite set of existential rules* Σ *defining the query predicate* Q*, each base instance* B*, each tuple of constants* $\mathbf{a}$*, and program* P_6 *obtained from* Σ *and* B *by applying Algorithm 1,* $\Sigma \cup B \models_{\approx} \mathsf{Q}(\mathbf{a})$ *if and only if* $P_6 \cup \mathsf{R}(P_6) \cup \mathsf{C}(P_6) \cup B \models \mathsf{Q}(\mathbf{a})$*.*

Theorem 5 shows that our entire pipeline is correct. Note that, if the Skolem chase of $\mathsf{sg}(\Sigma)$ terminates on every base instance, then line 2 of Algorithm 2 necessarily terminates.

Theorem 5. *For each finite set of existential rules* Σ *defining the query predicate* Q *such that the chase of* $\mathsf{sg}(\Sigma)$ *terminates on all base instances, and for each base instance* B*, Algorithm 1 outputs precisely all answers to* Q *on* $\Sigma \cup B$ *and then terminates.*

Example 5. *In our running example, program* P_6 *contains rules* (37)–(52)*, where* (34) *produces* (48) *and* (49)*, and* (35) *produces* (50) *and* (51)*.*

$$m^{\mathsf{f}}_{\mathsf{Q}} \leftarrow \tag{37}$$

$$\mathsf{Q}(x) \leftarrow m^{\mathsf{f}}_{\mathsf{Q}} \wedge A(x) \wedge R(x,y) \tag{38}$$

$$m^{\mathsf{f}}_{A} \leftarrow m^{\mathsf{f}}_{\mathsf{Q}} \tag{39}$$

$$m^{\mathsf{b}\wr\mathsf{f}}_{\approx}(x'') \leftarrow m^{\mathsf{f}}_{\mathsf{Q}} \wedge A(x'') \tag{40}$$

$$m^{\mathsf{bf}}_{R}(x'') \leftarrow m^{\mathsf{f}}_{\mathsf{Q}} \wedge A(x'') \tag{41}$$

$$A(g(x)) \leftarrow m^{\mathsf{f}}_{A} \wedge B(x) \tag{42}$$

$$\mathsf{F}_g(x, g(x)) \leftarrow m^{\mathsf{f}}_{A} \wedge B(x) \tag{43}$$

$$x \approx y \leftarrow m^{\mathsf{b}\wr\mathsf{f}}_{\approx}(x) \wedge T(x,y) \tag{44}$$

$$m^{\mathsf{bf}}_{T}(x) \leftarrow m^{\mathsf{b}\wr\mathsf{f}}_{\approx}(x) \tag{45}$$

$$x \approx y \leftarrow m^{\mathsf{b}\wr\mathsf{f}}_{\approx}(y) \wedge T(x,y) \tag{46}$$

$$m^{\mathsf{fb}}_{T}(y) \leftarrow m^{\mathsf{b}\wr\mathsf{f}}_{\approx}(y) \tag{47}$$

$$T(x, g(x)) \leftarrow m^{\mathsf{bf}}_{T}(x) \wedge B(x) \tag{48}$$

$$\mathsf{F}_g(x, g(x)) \leftarrow m^{\mathsf{bf}}_{T}(x) \wedge B(x) \tag{49}$$

$$T(x, g(x)) \leftarrow \mathsf{F}_g(x, z_{g(x)}) \wedge m^{\mathsf{fb}}_{T}(z_{g(x)}) \wedge B(x) \tag{50}$$

$$\mathsf{F}_g(x, g(x)) \leftarrow \mathsf{F}_g(x, z_{g(x)}) \wedge m^{\mathsf{fb}}_{T}(z_{g(x)}) \wedge B(x) \tag{51}$$

$$R(x, f(x)) \leftarrow m^{\mathsf{bf}}_{R}(x) \wedge S(x,z) \tag{52}$$

Program P_6 *contains no constants, function symbols, or equality atoms in the body, so we can answer* Q *on* P_6 *and* B^{ex} *using the chase for logic programs and thus avoid computing the least fixpoint for a program that explicitly axiomatizes equality. Doing so derives the following facts.*

$$m^{\mathsf{f}}_{\mathsf{Q}} \quad m^{\mathsf{f}}_{A} \quad A(g(a_1)) \quad \mathsf{F}_g(a_1, g(a_1)) \quad m^{\mathsf{b}\wr\mathsf{f}}_{\approx}(g(a_1)) \quad m^{\mathsf{fb}}_{T}(g(a_1)) \quad T(a_1, g(a_1))$$

Next, rule (46) *derives* $a_1 \approx g(a_1)$*, so the chase for logic programs takes* a_1 *as the representative of* $g(a_1)$ *and replaces* $g(a_1)$ *with* a_1*, thus deriving the following facts.*

$$m^{\mathsf{f}}_{\mathsf{Q}} \quad m^{\mathsf{f}}_{A} \quad A(a_1) \quad \mathsf{F}_g(a_1, a_1) \quad m^{\mathsf{b}\wr\mathsf{f}}_{\approx}(a_1) \quad m^{\mathsf{fb}}_{T}(a_1) \quad T(a_1, a_1)$$

After this, the chase further derives the following facts.

$$m^{\mathsf{bf}}_{T}(a_1) \quad m^{\mathsf{bf}}_{R}(a_1) \quad R(x, f(a_1)) \quad \mathsf{Q}(a_1)$$

Note that no facts involving a_i *with* $i \geq 2$ *are derived, as these are irrelevant to answering* Q*. Thus, evaluating* P_6 *on* B^{ex} *produces far fewer facts than applying the chase for logic programs to* $\mathsf{sk}(\Sigma^{ex})$*.*

8 Empirical Evaluation

We evaluated our technique using ChaseBench (Benedikt et al. 2017), a recent benchmark offering a mix of scenarios that simulate data exchange and ontology reasoning applications. We selected the scenarios summarized in Table 1, each comprising a set of existential rules, a base instance, and several queries. LUBM-100 and LUBM-1K are derived from the well-known Semantic Web LUBM (Guo, Pan, and Heflin 2011) benchmark; Deep300 is a "stress test" scenario; Doctors-1M simulates data exchange between

	TGDs	EGDs	Facts	Queries free	Queries const.
LUBM-100	136	0	12 M	4	10
LUBM-1K	136	0	120 M	4	10
DEEP300	1,300	0	1 k	20	0
DOCTORS-1M	5	4	1 M	7	11
STB-128	199	93	1 M	24	26
ONT-256	529	348	2 M	34	8

Note: "const." and "free" are the numbers of queries with and without constants, respectively.

Table 1: Summary of the test scenarios

medical databases; and STB-128 and ONT-256 were produced using the IBENCH and TOXGENE rule and instance generators. The first three scenarios contain only TGDs, and the remaining ones contain EGDs as well. All rules are weakly acyclic, so the chase always terminates. Finally, UNA was known to hold in all cases (Benedikt et al. 2017).

To compute the chase of the final program (line 7 of Algorithm 1), we used the RAM-based RDFox system written in C++.[1] We implemented our technique in Java on top of the CHASEBENCH (Benedikt et al. 2017) library. We used just one thread while computing the chase. Our system and the test data are available online.[2]

No existing goal-driven query answering techniques can handle these scenarios, as explained in the introduction. Thus, we primarily compared the performance of computing the chase with no optimizations (MAT), with the pipeline that uses just the magic sets and omits the relevance analysis (MAG), with the pipeline that uses just the relevance analysis and omits the magic sets (REL), and the entire pipeline with both relevance analysis and magic sets (REL+MAG). The MAT variant thus provides us with a baseline, and the remaining tests allow us to identify the relative contribution of various steps of our pipeline. Note that, in the presence of EGDs, we always used singularization and the other relevant steps of our pipeline, rather than the congruence axioms. On TGDs only, our algorithm becomes equivalent to the classical algorithm of Beeri and Ramakrishnan (1991).

In each test run, we computed the program P_6 from Algorithm 1 (skipping the relevance analysis and/or magic sets, as required for the test type), computed $\mathsf{chase}(P_6, B)$, and output the certain answers of Q as shown in line 9 of Algorithm 1. We recorded the wall-clock time of each run (without the loading times) and the number of facts derived by the chase; the latter provides an implementation-independent measure of the work needed to answer a query. In line 1 of Algorithm 2, we abstracted the base instance using the typed critical instance; however, computing the least fixpoint of such an abstraction was infeasible on DEEP300 so, in this case only, we used the optimization from Section 5.

Figure 1 summarizes the query times and the numbers of derived facts for our 158 test queries. The whiskers of each box plot show the minimum and maximum values, the box shows the lower quartile, the median, and the upper quartile, and the diamond shows the average. The distributions of these values are shown in more detail in Appendix A. Table 2 shows the times for computing the least fixpoint of the abstraction in line 2 of Algorithm 2, which are insignificant in all cases apart from DEEP300. On REL and REL+MAG, one query of DEEP300 and nine queries of STB-128 could not be processed by the relevance analysis for reasons we discuss shortly. Moreover, all queries of LUBM-1K and DEEP300 on MAT, three queries of LUBM-1K and one of DEEP300 on MAG, all queries of LUBM-1K and one of DEEP300 on REL, and three queries of LUBM-1K on REL+MAG could not be processed due to memory exhaustion while computing the chase.

Figure 1 clearly shows that our technique is generally very effective and can mean the difference between success and failure: the chase for LUBM-1K and DEEP300 could not be computed on our test machine, whereas REL+MAG can answer 11 out of 14 queries on LUBM-1K in at most 45 s, and 19 out of 20 queries on DEEP300 in at most 18 s. The upper quartile of query times for REL+MAG is at least an order of magnitude below the times for MAT in all cases apart from DOCTORS-1M, where this holds for the median. Overall, REL+MAG achieves the best performance.

In addition, relevance analysis alone can lead to significant improvements: the query times for REL and REL+MAG are almost identical on DEEP300 and ONT-256, suggesting that the improvements are due to relevance analysis, rather than magic sets. On some queries of STB-128 and ONT-256, the relevance analysis eliminates all rules and thus proves that queries have no answers. The benefits of relevance analysis are marginal only on LUBM, mainly because its TGDs contain few existential quantifiers.

However, relevance analysis also has its pitfalls: we could not run it on one query of DEEP300 with eight body atoms, and on nine queries of STB-128 containing between 11 and 19 output variables that, after singularization, have between 19 and 22 body atoms. Atoms of these queries match to many facts in the least fixpoint of the abstraction, so query evaluation explodes either in line 2 or line 7 of Algorithm 2. One additional query of DEEP300 exhibited similar issues, which we addressed by (manually) tree-decomposing the query and thus reducing the number of matches.

To investigate the cases in which magic sets are particularly beneficial, Table 3 shows the minimum, maximum, and median of the query times and the numbers of derived facts for queries without and with constants. The maximum numbers of derived facts are particularly telling: without constants, REL+MAG derives more facts compared to REL; and with constants, the numbers for REL+MAG are several orders of magnitude smaller compared to REL. Programs P_3 in line 3 of Algorithm 1 are the same in both cases so this improvement is clearly due to magic sets. In contrast, the impact of constants is insignificant for REL, highlighting the different strengths of relevance analysis and magic sets.

Unfortunately, magic sets are not "free". For example, MAT and REL are faster than both MAG and REL+MAG on seven constant-free queries of DOCTORS-1M, and they outperform MAG on 17 queries of STB-128 and on 39 queries of ONT-256. In all these cases, the magic sets transforma-

[1] http://www.cs.ox.ac.uk/isg/tools/RDFox/
[2] http://github.com/tsamoura/chaseGoal

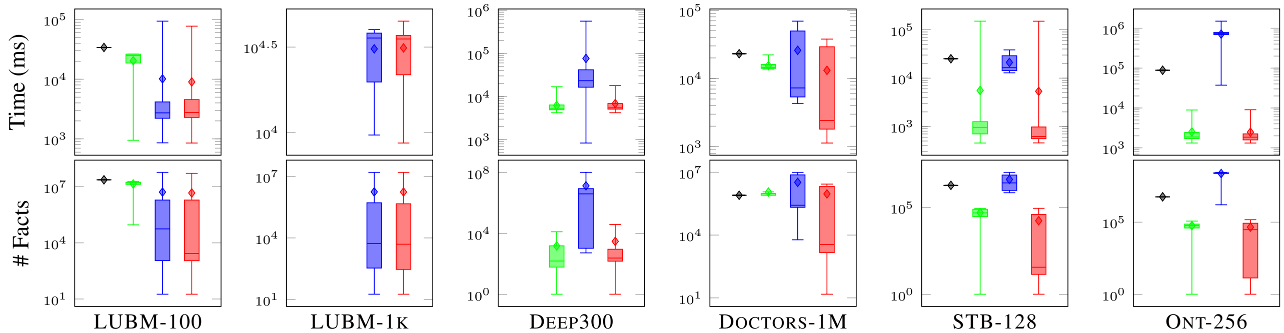

Figure 1: Times and the numbers of derived facts. Black, green, blue and red are MAT, REL, MAG and REL+MAG, respectively.

	LUBM	DEEP300	DOCTORS-1M	STB-128	ONT-256
Time	10 ms	4700 ms	10 ms	200 ms	650 ms

Table 2: Times for computing the least fixpoints of the base instance abstractions

tion increases the number of rules by one or two orders of magnitude, which introduces considerable overhead during chase computation. This can be particularly significant on queries without constants: such queries tend to have more answers and thus require exploring larger proofs.

Both relevance analysis and magic sets try to identify proofs deriving a goal. Magic sets do this “at runtime”: the transformed rules derive only facts that would be explored by top-down reasoning. In contrast, relevance analysis identifies rules that can participate in such proofs “offline” by checking whether all body atoms of a rule are derivable. Combining the two optimizations is particularly effective at reducing the overheads described in the previous paragraph.

9 Conclusion & Outlook

We presented a novel approach for goal-driven query answering over terminating existential rules. Our empirical results clearly show that our technique can lead to significant performance improvements and can mean the difference between success and failure to answer a query. In the future, we shall investigate the use of magic sets for query answering on nonterminating, but decidable classes of existential rules (e.g., guarded, linear, or sticky). We shall also consider adding optimizations such as tabling and subsumption.

Acknowledgments

This research was funded by the EPSRC projects DBOnto, ED^3, $MaSI^3$, PDQ, and the Alan Turing Institute.

References

Alviano, M.; Faber, W.; Greco, G.; and Leone, N. 2012a. Magic Sets for disjunctive Datalog programs. *Artificial Intelligence* 187:156–192.

Alviano, M.; Leone, N.; Manna, M.; Terracina, G.; and Veltri, P. 2012b. Magic-Sets for Datalog with Existential Quantifiers. In *Datalog 2.0*.

Alviano, M.; Greco, G.; and Leone, N. 2011. Dynamic Magic Sets for Programs with Monotone Recursive Aggregates. In *LPNMR*.

Baget, J.-F.; Mugnier, M.-L.; Rudolph, S.; and Thomazo, M. 2011a. Walking the complexity lines for generalized guarded existential rules. In *IJCAI*.

Baget, J.; Leclère, M.; Mugnier, M.; and Salvat, E. 2011b. On rules with existential variables: Walking the decidability line. *Artificial Intelligence* 175(9–10):1620–1654.

Baget, J.; Bienvenu, M.; Mugnier, M.; and Rocher, S. 2015a. Combining Existential Rules and Transitivity: Next Steps. In *IJCAI*.

Baget, J.; Leclère, M.; Mugnier, M.; Rocher, S.; and Sipieter, C. 2015b. Graal: A Toolkit for Query Answering with Existential Rules. In *RuleML*.

Bancilhon, F.; Maier, D.; Sagiv, Y.; and Ullman, J. D. 1986. Magic Sets and Other Strange Ways to Implement Logic Programs. In *PODS*.

Beeri, C., and Ramakrishnan, R. 1991. On the Power of Magic. *Journal of Logic Programming* 10(3&4):255–299.

Benedikt, M.; Konstantinidis, G.; Mecca, G.; Motik, B.; Papotti, P.; Santoro, D.; and Tsamoura, E. 2017. Benchmarking the Chase. In *PODS*.

Bonifati, A.; Ileana, I.; and Linardi, M. 2017. ChaseFUN: a Data Exchange Engine for Functional Dependencies at Scale. In *EDBT*.

Calì, A.; Gottlob, G.; and Lukasiewicz, T. 2012. A general Datalog-based framework for tractable query answering over ontologies. *Journal of Web Semantics* 14:57–83.

Calì, A.; Gottlob, G.; and Pieris, A. 2010. Query Answering under Non-guarded Rules in Datalog$^{\pm}$. In *RR*.

Calimeri, F.; Cozza, S.; Ianni, G.; and Leone, N. 2009. Magic Sets for the Bottom-Up Evaluation of Finitely Recursive Programs. In *LPNMR*.

Calvanese, D.; De Giacomo, G.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2007. Tractable Reasoning and Efficient

		Query times (seconds)								
		MAG			REL			REL+MAG		
		min	max	median	min	max	median	min	max	median
LUBM-100	N	1.62	93.56	11.43	1.69	26.24	25.29	1.59	77.14	11.30
	Y	0.86	2.98	2.58	0.95	26.29	24.92	0.85	3.22	2.68
LUBM-1K	N	19.74	19.74	19.74	N/A	N/A	N/A	19.74	19.74	19.74
	Y	9.64	40.10	36.19	N/A	N/A	N/A	8.66	44.93	35.79
DOCTORS-1M	N	47.27	69.21	54.86	14.05	22.21	16.11	25.57	37.81	29.81
	Y	4.28	8.51	5.66	13.75	15.57	14.27	1.15	2.70	2.10
STB-128	N	14.65	38.35	30.07	0.45	150.57	0.74	0.46	150.21	0.75
	Y	12.87	17.92	14.35	0.75	18.47	1.10	0.51	18.18	0.59
ONT-256	N	60.36	1.49 k	745.17	1.32	8.91	1.86	1.32	9.04	1.99
	Y	37.27	703.26	690.82	1.69	2.69	1.83	1.51	2.28	1.55

		The numbers of derived facts								
		MAG			REL			REL+MAG		
		min	max	median	min	max	median	min	max	median
LUBM-100	N	1.59 M	59.23 M	5.58 M	1.59 M	18.42 M	1.74 M	1.59 M	52.41 M	5.33 M
	Y	18	51.94 k	4.70 k	92.34 k	19.52 M	17.37 M	18.00	46.40 k	4.40 k
LUBM-1K	N	15.84 M	15.84 M	15.84 M	N/A	N/A	N/A	15.84 M	15.84 M	15.84 M
	Y	18	2.77 M	4.40 k	N/A	N/A	N/A	18.00	2.77 M	4.40 k
DOCTORS-1M	N	6.65 M	9.82 M	7.44 M	953.50 k	1.74 M	1.74 M	1.90 M	2.69 M	2.69 M
	Y	5.94 k	264.30 k	263.16 k	952.50 k	953.63 k	952.50 k	15.00	4.08 k	2.96 k
STB-128	N	1.88 M	10.80 M	7.76 M	0.00	90.05 k	39.47 k	1.00	90.64 k	40.00 k
	Y	723.19 k	2.92 M	1.05 M	30.00 k	90.00 k	65.00 k	9.00	51.00	20.00
ONT-256	N	6.35 M	283.68 M	254.05 M	0.00	118.09 k	59.00 k	1.00	147.49 k	45.91 k
	Y	1.59 M	224.23 M	224.22 M	29.36 k	118.07 k	54.12 k	9.00	41.00	11.50

Table 3: Running times and the numbers of facts for queries without ("N") and with ("Y") constants

Query Answering in Description Logics: The DL-Lite Family. *Journal of Automated Reasoning* 39(3):385–429.

Calvanese, D.; Cogrel, B.; Komla-Ebri, S.; Kontchakov, R.; Lanti, D.; Rezk, M.; Rodriguez-Muro, M.; and Xiao, G. 2017. Ontop: Answering SPARQL queries over relational databases. *Semantic Web* 8(3):471–487.

Deutsch, A.; Nash, A.; and Remmel, J. B. 2008. The chase revisited. In *PODS*.

Eiter, T.; Ortiz, M.; Simkus, M.; Tran, T.; and Xiao, G. 2012. Query Rewriting for Horn-$\mathcal{SHIQ}$ Plus Rules. In *AAAI*.

Fagin, R.; Kolaitis, P. G.; Miller, R. J.; and Popa, L. 2005. Data exchange: semantics and query answering. *Theoretical Computer Science* 336(1):89–124.

Geerts, F.; Mecca, G.; Papotti, P.; and Santoro, D. 2014. That's All Folks! LLUNATIC Goes Open Source. *PVLDB* 7(13):1565–1568.

Gottlob, G.; Manna, M.; and Pieris, A. 2015. Polynomial Rewritings for Linear Existential Rules. In *IJCAI*.

Gottlob, G.; Orsi, G.; and Pieris, A. 2014. Query Rewriting and Optimization for Ontological Databases. *ACM TODS* 39(3):25:1–25:46.

Gottlob, G.; Rudolph, S.; and Simkus, M. 2014. Expressiveness of guarded existential rule languages. In *PODS*.

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. *JAIR* 47:741–808.

Guo, Y.; Pan, Z.; and Heflin, J. 2011. LUBM: A benchmark for OWL knowledge base systems. *Web Semantics: Science, Services and Agents on the World Wide Web* 3(2-3).

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., and Rudolph, S. 2011. Extending Decidable Existential Rules by Joining Acyclicity and Guardedness. In *IJCAI*.

Leone, N.; Manna, M.; Terracina, G.; and Veltri, P. 2012. Efficiently Computable Datalog$^{\exists}$ Programs. In *KR*.

Marnette, B. 2009. Generalized schema-mappings: from termination to tractability. In *PODS*.

Motik, B.; Nenov, Y.; Piro, R.; Horrocks, I.; and Olteanu, D. 2014. Parallel Materialisation of Datalog Programs in Centralised, Main-Memory RDF Systems. In *AAAI*.

Motik, B.; Nenov, Y.; Piro, R.; and Horrocks, I. 2015. Handling owl:sameAs via Rewriting. In *AAAI*.

ten Cate, B.; Chiticariu, L.; Kolaitis, P. G.; and Tan, W. C. 2009. Laconic Schema Mappings: Computing the Core with SQL Queries. *PVLDB* 2(1):1006–1017.

A Full Experimental Results

We conducted all experiments on a MacBook Pro laptop with a 3.1 GHz Intel Core i7 processor, 16 GB of DDR3 RAM, and an 500 GB SSD, running macOS Sierra Version 10.12.5.

Figure 2 shows a *cumulative distribution* for the query times, the number of derived facts, and the number of rules in program P_6 from Algorithm 1. For example, the graphs in the first column show how many test queries can be answered in a given time; thus, in the LUBM-100 scenario, REL+MAG answers 11 test queries in under 3 s, and that this covers most test queries. Lines not reaching the top of a graph show queries that could not be answered.

B Proofs

We now prove Theorem 5, which shows that Algorithm 1 correctly computes all query answers in a finite amount of time. Towards this goal, we prove correctness of all intermediate steps as well.

B.1 Properties of Singularization

Recall that our pipeline in Algorithm 1 begins with the singularization transformation, Step 1. This introduces explicit equality atoms in the body of rules, as described in Section 4. Proposition B.1 summarizes the properties of singularization. ten Cate et al. (2009) proved the first two claims when UNA holds, but extending their argument in the absence of UNA is straightforward. The third property follows from the first two and the adjustments in (10) to the existential rules defining the query predicate Q.

Proposition B.1. *For each set Σ of existential rules, base instance B, predicate R, tuples of constants $\mathbf{a}$ and variables $\mathbf{x}$, constants b and c, and $\Sigma' = \mathsf{sg}(\Sigma) \cup \mathsf{R}(\Sigma) \cup \mathsf{ST}$,*

- $\Sigma \cup B \models_\approx R(\mathbf{a})$ *iff* $\Sigma' \cup B \models \exists \mathbf{x}.R(\mathbf{x}) \wedge \bigwedge_{i=1}^n x_i \approx a_i$*;*
- $\Sigma \cup B \models_\approx b \approx c$ *iff* $\Sigma' \cup B \models b \approx c$*; and*
- $\Sigma \cup B \models_\approx \mathsf{Q}(\mathbf{a})$ *iff* $\Sigma' \cup B \models \mathsf{Q}(\mathbf{a})$.

B.2 Correctness of the Relevance Analysis

In Algorithm 1, singularization is followed by the standard Skolemization step. Next, relevance analysis is applied to prune the set of rules, as described in Algorithm 2 of Section 5. Our relevance algorithm is new, and hence we will need to prove its correctness, which is captured in the following theorem.

Theorem 1. *For each $\approx$-safe program P defining the query predicate Q and base instance B where line 2 of Algorithm 2 terminates, program $\mathcal{R} = \mathsf{relevance}(P, B)$ is $\approx$-safe, and, for each tuple $\mathbf{a}$ of constants, $P \cup \mathsf{R}(P) \cup \mathsf{ST} \cup B \models \mathsf{Q}(\mathbf{a})$ if and only if $\mathcal{R} \cup \mathsf{R}(\mathcal{R}) \cup \mathsf{ST} \cup B \models \mathsf{Q}(\mathbf{a})$.*

Proof. Let B' be as chosen in line 1; let η be any mapping of constants to constants such that $\eta(B) \subseteq B'$ and $\eta(c) = c$ for each constant c occurring in P; let I and P' be as in line 2; let $\mathcal{R}_{\mathsf{int}}$ be the intermediate set of rules as computed just before line 13; and let $\mathcal{R}_{\mathsf{fin}}$ and $\mathcal{D}_{\mathsf{fin}}$ be the final sets. Note that $\mathsf{R}(\mathcal{R}_{\mathsf{int}}) = \mathsf{R}(\mathcal{R}_{\mathsf{fin}})$. Furthermore, $\mathcal{R}_{\mathsf{int}} \subseteq P$ so $\mathcal{R}_{\mathsf{int}}$ is $\approx$-safe, and the transformation in line 15 clearly preserves $\approx$-safety.

($\Leftarrow$) For $P_1 = \mathcal{R}_{\mathsf{fin}} \cup \mathsf{R}(\mathcal{R}_{\mathsf{fin}}) \cup \mathsf{ST}$, let $I_0, I_1, \ldots$ be the sequence of instances used to compute $T_{P_1}^\infty(B)$ as defined in Section 2. We show by induction i that, for each fact $F \in I_i$, we have $P \cup \mathsf{R}(P) \cup \mathsf{ST} \cup B \models F$. The base case is trivial. For the induction step, assume that the claim holds for I_i, and consider applying a rule $r \in \mathcal{R}_{\mathsf{fin}}$ with a substitution σ where $\sigma(\mathsf{h}(r)) = F$. Thus, we have $\sigma(\mathsf{b}(r)) \subseteq I_i$, and so $P \cup \mathsf{R}(P) \cup \mathsf{ST} \cup B \models \sigma(\mathsf{b}(r))$ holds by the inductive assumption. Rule $r \in \mathcal{R}_{\mathsf{fin}}$ is obtained from some rule $r' \in \mathcal{R}_{\mathsf{int}} \subseteq P$ by transformations in line 15. Let σ' be the extension of σ such that, for each variable x that was replaced with a term t, we set $\sigma'(x) = \sigma(t)$. Each body atom $A' \in \mathsf{b}(r')$ that was not removed in line 15 corresponds to some $A \in \mathsf{b}(r)$ and $\sigma'(A') = \sigma(A)$ holds. Moreover, since P is $\approx$-safe, each body atom A' of r' that was removed in line 15 is of the form $x \approx y$ or $x \approx t$, and variable x occurs in some atom $R_i(\ldots, x, \ldots) \in \mathsf{b}(r')$; but then, the inductive assumption and the fact that $\mathsf{R}(\mathcal{R}_{\mathsf{fin}})$ contains the reflexivity rules for B ensure $P \cup \mathsf{R}(P) \cup \mathsf{ST} \cup B \models \sigma'(A')$. Summarizing, we have $P \cup \mathsf{R}(P) \cup \mathsf{ST} \cup B \models \sigma'(\mathsf{b}(r'))$, and so $P \cup \mathsf{R}(P) \cup \mathsf{ST} \cup B \models \sigma'(\mathsf{h}(r'))$ holds, as required.

($\Rightarrow$) For $P_2 = P \cup \mathsf{R}(P) \cup \mathsf{ST}$, let $I_0, I_1, \ldots$ be the sequence of instances used to compute $T_{P_2}^\infty(B)$ as defined in Section 2. We prove the claim in two steps.

First, we show by induction on i that $\eta(I_i) \subseteq I$ holds. For the base case, we have $\eta(I_0) = B' \subseteq I$, as required. Now assume that $\eta(I_i) \subseteq I$ holds, and consider an arbitrary fact $F \in I_{i+1}$ derived by a rule $r \in P'$ and substitution σ; hence, $\sigma(\mathsf{b}(r)) \subseteq I_i$ holds. By the inductive assumption, we have $\eta(\sigma(\mathsf{b}(r))) \subseteq \eta(I_i) \subseteq I$. Moreover, η does not affect the constants occurring in P, so $\sigma'(\mathsf{b}(r)) \subseteq \eta(I_i) \subseteq I$ holds where σ' is the substitution defined by $\sigma'(x) = \eta(\sigma(x))$ for each variable x from the domain of σ; consequently, we have $\sigma'(\mathsf{h}(r)) \in I$. Finally, since η does not affect the constants occurring in P, we have $\sigma'(\mathsf{h}(r)) = \eta(F)$, as required.

Second, we show by induction in i that, for each fact $F \in I_i$ that is not of the form $t \approx t$ and that satisfies $\eta(F) \in \mathcal{D}_{\mathsf{fin}}$, we have $\mathcal{R}_{\mathsf{fin}} \cup \mathsf{R}(\mathcal{R}_{\mathsf{fin}}) \cup \mathsf{ST} \cup B \models F$. Line 3 of the algorithm ensures that $\eta(\mathsf{Q}(\mathbf{a})) \in \mathcal{D}_{\mathsf{fin}}$ holds for each $\mathbf{a}$ with $\mathcal{R}_{\mathsf{fin}} \cup \mathsf{R}(\mathcal{R}_{\mathsf{fin}}) \cup \mathsf{ST} \cup B \models \mathsf{Q}(\mathbf{a})$, so our claim holds.

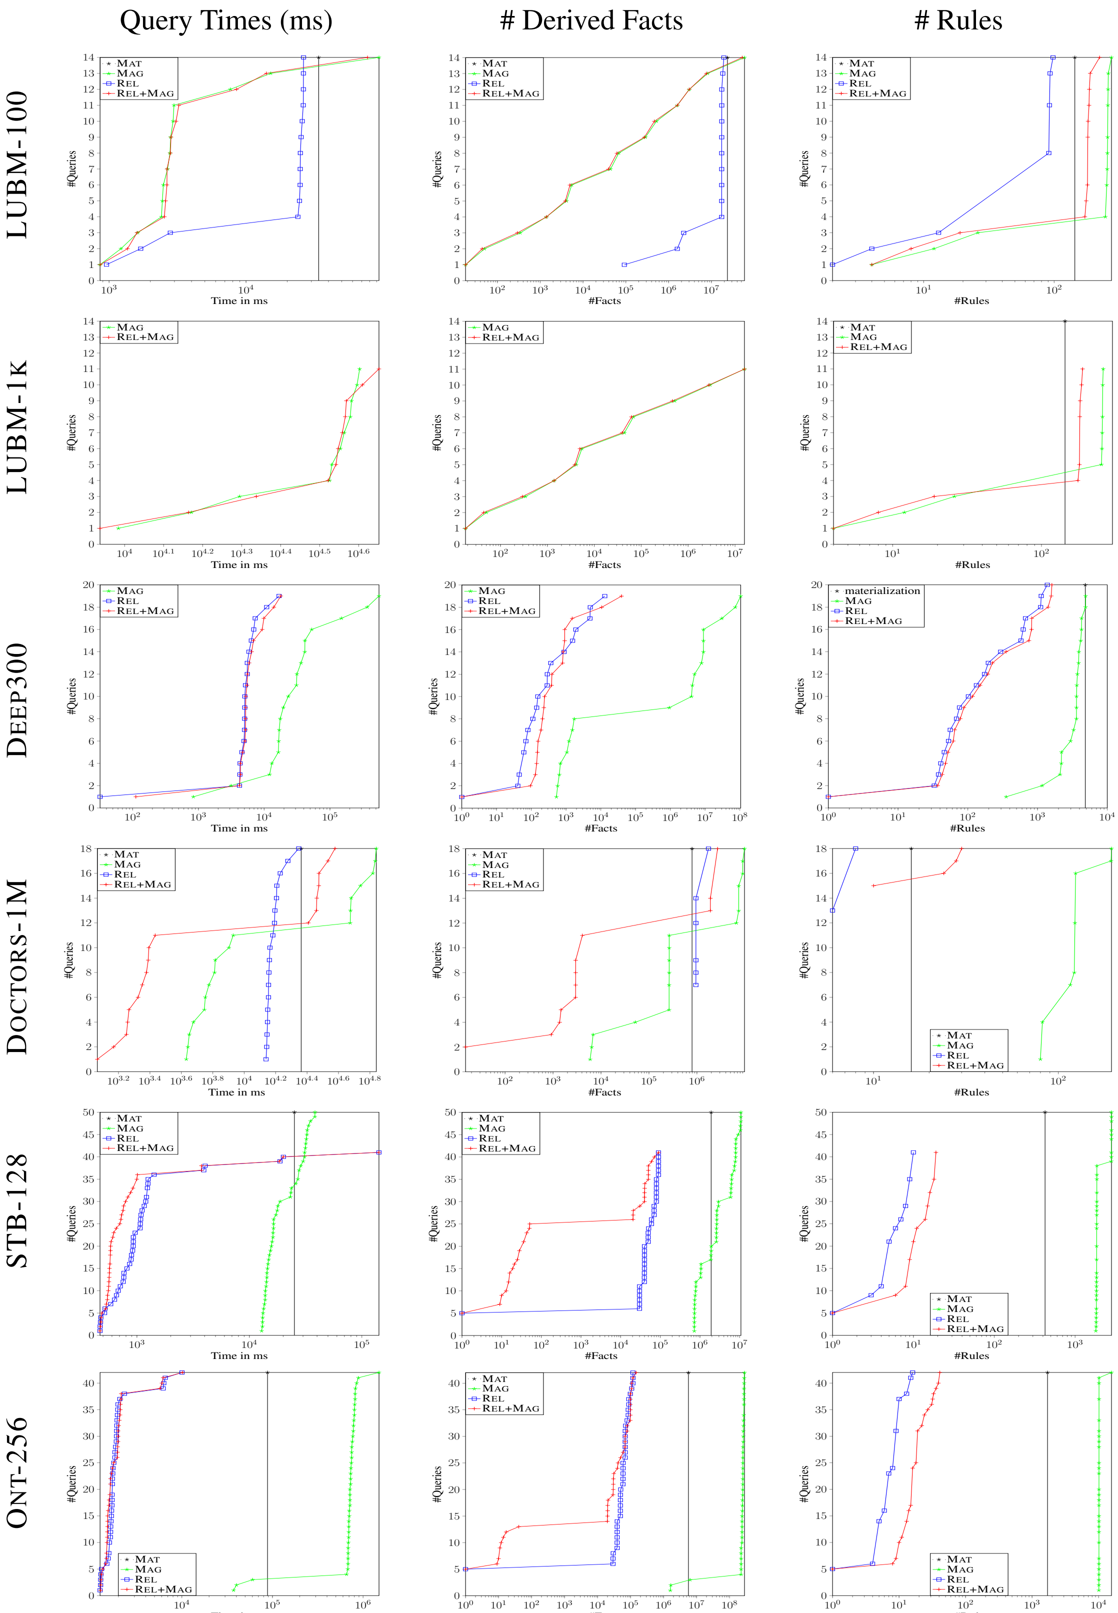

Figure 2: Times, numbers of derived facts, and total numbers of rules

The induction base holds trivially. Now assume that I_i satisfies this property and consider an arbitrary fact $F \in I_{i+1}$ derived by a rule $r \in P'$ and substitution σ; hence, $\sigma(\mathsf{b}(r)) \subseteq I_i$ holds. Moreover, assume that F is not of the form $t \approx t$; thus, $r \notin \mathsf{R}(P)$ and so $r \in P \cup \mathsf{ST}$. Also, assume that $\eta(F) \in \mathcal{D}_{\mathsf{fin}}$ holds; then $\eta(F)$ must have been added to $\mathcal{T}$ in line 11, so it must have been extracted from $\mathcal{T}$ at some point in line 6. Moreover, due to $r \in P \cup \mathsf{ST}$, rule r was considered in line 7. Finally, let $r' = \sigma(r)$; then $\eta(\mathsf{b}(r')) \subseteq I$ holds by the previous paragraph, which together with $\eta(\mathsf{h}(r')) = \eta(F)$ and the fact that η does not affect the constants in P ensures that substitution ν satisfying $\nu(r) = \eta(r')$ was considered in line 7. Consequently, r is added to $\mathcal{R}_{\mathsf{int}}$ in line 8. We next prove the following property:

$$\mathcal{R}_{\mathsf{fin}} \cup \mathsf{R}(\mathcal{R}_{\mathsf{fin}}) \cup \mathsf{ST} \cup B \models \mathsf{b}(r'). \tag{$*$}$$

First, consider a ground atom $F_i \in \mathsf{b}(r')$ not of the form $t \approx t$. If atom F_i is relational, or if F_i contains function symbols, or if $P \cup B$ does not satisfy UNA, then $\eta(F_i)$ is not of the form $c \approx c$ for c a constant; hence, $\eta(F_i)$ added to $\mathcal{D}_{\mathsf{fin}}$ in line 11, and so the inductive assumption ensures $\mathcal{R}_{\mathsf{fin}} \cup \mathsf{R}(\mathcal{R}_{\mathsf{fin}}) \cup \mathsf{ST} \cup B \models F_i$.

Second, consider a ground atom $t \approx t \in \mathsf{b}(r')$. Rule r is $\approx$-safe, so atom $t \approx t$ is obtained from some $A \in \mathsf{b}(r)$ of the form $x \approx y$ or $x \approx t$, and a relational atom $R_i(\mathbf{t}_i) \in \mathsf{b}(r)$ exists such that $\mathsf{vars}(A) \cap \mathbf{t}_i \neq \emptyset$. The previous paragraph ensures that $\mathcal{R}_{\mathsf{fin}} \cup \mathsf{R}(\mathcal{R}_{\mathsf{fin}}) \cup \mathsf{ST} \cup B \models \sigma(R_i(\mathbf{t}_i))$ holds; and $r \in \mathcal{R}_{\mathsf{int}}$ so $\mathsf{R}(\mathcal{R}_{\mathsf{fin}})$ contains all reflexivity rules for R_i; thus, $\mathcal{R}_{\mathsf{fin}} \cup \mathsf{R}(\mathcal{R}_{\mathsf{fin}}) \cup \mathsf{ST} \cup B \models t \approx t$ holds.

Thus, property $(*)$ holds. To complete the proof, we show that $\mathcal{R}_{\mathsf{fin}}$ contains a rule that derives F. In particular, $r \in \mathcal{R}_{\mathsf{int}}$ corresponds to some rule $r'' \in \mathcal{R}_{\mathsf{fin}}$ where the latter is obtained by transformations in line 15. Now consider an arbitrary body equality atom $F_i \in \mathsf{b}(r')$. If F_i is not of the form $c \approx c$ for some constant c, then either $P \cup B$ does not satisfy UNA or $\eta(F_i)$ is not of the form $c \approx c$ for c a constant; hence, $\langle r, i \rangle$ is added to $\mathcal{B}$ in line 12, and so the body atom of r'' corresponding to F_i is not eliminated in line 15. Thus, we have $\sigma(\mathsf{b}(r'')) \subseteq \mathsf{b}(r')$, so $\mathcal{R}_{\mathsf{fin}} \cup \mathsf{R}(\mathcal{R}_{\mathsf{fin}}) \cup \mathsf{ST} \cup B \models F$. ☐

B.3 Correctness of the Magic Transformation

In line 4 of Algorithm 1, we apply our variant of the magic sets algorithm to the result of the prior steps. Note that both the input and output of this step are rules which may contain equality atoms in both the head and the body. We need to prove not only that this preserves the semantics of the input, but that it also preserves $\approx$-safety. The latter guarantees that when, we remove the equalities in the final two steps of the pipeline, we are left with a domain-independent set of rules.

Theorem 2. *For each $\approx$-safe program P defining the query predicate Q, program $\mathcal{R} = \mathsf{magic}(P)$ is $\approx$-safe; moreover, for each base instance B and each tuple $\mathbf{t}$ of ground terms,*

$$\begin{gathered} P \cup \mathsf{R}(P) \cup \mathsf{ST} \cup B \models \mathsf{Q}(\mathbf{t}) \quad \textit{iff} \\ \mathcal{R} \cup \mathsf{R}(\mathcal{R}) \cup \mathsf{ST} \cup B \models \mathsf{Q}(\mathbf{t}). \end{gathered}$$

Proof. To see that $\mathcal{R}$ is $\approx$-safe, consider a rule $r \in P \cup \mathsf{ST}$ processed by the function $\mathsf{process}$. Since r is $\approx$-safe, the rule added to $\mathcal{R}$ in line 12 for r is clearly $\approx$-safe as well. Moreover, the reordered atoms in line 13 satisfy the condition from Definition 5:

- each rule $r \in P$ can be ordered in such a way since P is $\approx$-safe, and
- each rule $r \in \mathsf{ST}$ can be ordered in such a way since ff is not a valid adornment of $\approx$, so $\mathbf{t}^\beta$ in line 13 contains at least one variable from the body of r.

Now consider a rule $r' \in \mathcal{R}$ computed from r in line 18. For each atom $R_i(\mathbf{t}_i) \in \mathsf{b}(r') \subseteq \mathsf{b}(r)$ of the form $x \approx y$ or $x \approx s$, and for $z \in \mathsf{vars}(\mathbf{t}_i)$ the variable that satisfies the condition for $R_i(\mathbf{t}_i)$ in Definition 5, the $\approx$-safety of r' is ensured by $m_R^\alpha(\mathbf{t}^\beta) \in \mathsf{b}(r')$ if $z \in T$ for $T = \mathbf{t}^\beta$, and by $R_j(\mathbf{t}_j) \in \mathsf{b}(r')$ if $z \in \mathbf{t}_j$ for some $j < i$.

We now proceed to prove that query answers remain preserved.

($\Leftarrow$) Each rule $r' \in \mathcal{R}$ where $\mathsf{h}(r')$ does not contain a magic predicate is obtained from some rule $r \in P \cup \mathsf{ST}$ by appending an atom with a magic predicate to $\mathsf{b}(r)$. Thus, each derivation from $\mathcal{R} \cup \mathsf{R}(\mathcal{R}) \cup \mathsf{ST} \cup B$ of a fact not containing a magic predicate corresponds directly to a derivation of the fact from $P \cup \mathsf{R}(P) \cup \mathsf{ST} \cup B$.

($\Rightarrow$) Let $P_1 = P \cup \mathsf{R}(P) \cup \mathsf{ST}$ and $I = T_{P_1}^\infty(B)$, and let $I_0, I_1, \ldots$ be the sequence of instances used to compute I as defined in Section 2. We prove by induction on i that, for each fact $R(\mathbf{s}) \in I_i$ not of the form $s \approx s$, each magic predicate m_R^α occurring in $\mathcal{R}$, and each tuple of ground terms $\mathbf{s}$, the following two properties hold:

1. $\mathcal{R} \cup \mathsf{R}(\mathcal{R}) \cup \mathsf{ST} \cup B \cup \{m_\approx^{\mathsf{b}\wr\mathsf{f}}(\mathbf{s}^\beta)\} \models \approx(\mathbf{s})$ for each $\beta \in \{\mathsf{bf}, \mathsf{fb}\}$ if $R = \approx$ and $\alpha = \mathsf{b}\wr\mathsf{f}$;
2. $\mathcal{R} \cup \mathsf{R}(P') \cup \mathsf{ST} \cup B \cup \{m_R^\alpha(\mathbf{s}^\alpha)\} \models R(\mathbf{s})$ if R is distinct from $\approx$ or $\alpha \neq \mathsf{b}\wr\mathsf{f}$.

Line 1 ensures $m_{\mathsf{Q}}^\alpha \leftarrow \in \mathcal{R}$ for $\alpha = f \cdots f$, so fact m_{Q}^α is derivable from $\mathcal{R}$ and B, and therefore, for each ground tuple of terms $\mathbf{t}$ such that $P \cup \mathsf{R}(P) \cup \mathsf{ST} \cup B \models \mathsf{Q}(\mathbf{t})$, property 2 ensures $\mathcal{R} \cup \mathsf{R}(\mathcal{R}) \cup \mathsf{ST} \cup B \models \mathsf{Q}(\mathbf{t})$.

The base case holds trivially for each $R(\mathbf{s}) \in I_0 = B$. For the induction step, assume that I_i satisfies properties 1 and 2 and consider an arbitrary fact $R(\mathbf{s}) \in I_{i+1}$ not of the form $t \approx t$ derived by a rule $r \in P_1$ and substitution σ; hence, $\sigma(\mathsf{b}(r)) \subseteq I_i$

and $r \notin \mathsf{R}(P)$ hold. Moreover, consider an arbitrary magic predicate m_R^α occurring in $\mathcal{R}$. Then, m_R^α was extracted from $\mathcal{T}$ in line 3, and so rule $r \in P \cup \mathsf{ST}$ was considered in line 4. Let $R(\mathbf{t})$ be the head of r; thus, we clearly have $\sigma(\mathbf{t}) = \mathbf{s}$.

Now consider $\beta \in \{\mathsf{bf}, \mathsf{fb}\}$ if $R = \approx$ and $\alpha = \mathsf{b}\wr\mathsf{f}$, and let $\beta = \alpha$ otherwise. The algorithm calls $\mathsf{process}(r, \alpha, \beta)$ in lines 6–7 or line 9. Let $R_1(\mathbf{t}_1), \ldots, R_n(\mathbf{t}_n)$ be the reordered body or r from line 13; as we have already mentioned, such an ordering exists since P is $\approx$-safe. We next prove by another induction on $1 \leq j \leq n$ that the following property holds:

$$\mathcal{R} \cup \mathsf{R}(\mathcal{R}) \cup \mathsf{ST} \cup B \cup \{m_R^\alpha(\mathbf{s}^\beta)\} \models \sigma(R_j(\mathbf{t}_j)). \tag{$*$}$$

The base case is the same as the induction step. Consider $1 \leq j \leq n$ such that the claim holds for each $j' < j$.

First, assume $\sigma(R_j(\mathbf{t}_j))$ is of the form $t \approx t$. Program P is $\approx$-safe, so $R_j(\mathbf{t}_j)$ is of the form $x \approx y$ or $x \approx t$, and some $z \in \mathsf{vars}(\mathbf{t}_j)$ satisfies the condition of Definition 5.

- If $z \in T$ for $T = \mathbf{t}^\beta$, then $z \in \mathbf{t}^\beta$ holds for the body atom of the form $m_R^\alpha(\mathbf{t}^\beta)$ added in line 12 for r, so $\sigma(z) = t \in \mathbf{s}^\beta$. Moreover, $m_R^\alpha(\mathbf{s}^\beta)$ holds by ($*$), and $\mathsf{R}(\mathcal{R})$ contains the reflexivity rules for m_R^α, and one of them derives $t \approx t$.
- If $z \in \mathbf{t}_{j'}$ for some $j' < j$ and body atom $R_{j'}(\mathbf{t}_{j'})$ of r, then ($*$) ensures that $\sigma(R_{j'}(\mathbf{t}_{j'}))$ is derived. Moreover, $\mathsf{R}(\mathcal{R})$ contains the reflexivity rules for $R_{j'}$, and one of them derives $t \approx t$.

Second, assume $\sigma(R_j(\mathbf{t}_j))$ is not of the form $t \approx t$ and that R_j does not occur in the head of a rule in P. The rules from $\mathsf{R}(\mathcal{R}) \cup \mathsf{ST}$ can derive only facts of the form $t \approx t$, so therefore R_j is different from $\approx$. But then we have $\sigma(R_j(\mathbf{t}_j)) \in B$, and thus ($*$) holds trivially.

Third, assume $\sigma(R_j(\mathbf{t}_j))$ is not of the form $t \approx t$ and that R_j is processed in lines 14–19. Let γ be the adornment for $R_j(\mathbf{t}_j)$ in line 15 and let S be the magic predicate in lines 16–17. The rule added to P in line 18 and ($*$) ensure $\mathcal{R} \cup \mathsf{R}(\mathcal{R}) \cup \mathsf{ST} \cup B \cup \{m_R^\alpha(\mathbf{s}^\beta)\} \models \sigma(S(\mathbf{t}_j^\gamma))$, so the inductive assumption for properties 1 and 2 ensures $\mathcal{R} \cup \mathsf{R}(\mathcal{R}) \cup \mathsf{ST} \cup B \cup \{m_R^\alpha(\mathbf{s}^\beta)\} \models \sigma(R_j(\mathbf{t}_j))$.

This completes the proof of ($*$). To complete the proof of this theorem, note that ($*$) and $r \in \mathsf{ST}$ trivially imply property 1. Moreover, if $R = \approx$ and $\alpha = \mathsf{b}\wr\mathsf{f}$, then ($*$) and the rule added to $\mathcal{R}$ in line 12 ensure property 1. Finally, in all other cases, ($*$) and the rule added to $\mathcal{R}$ in line 12 ensure property 2. □

B.4 Magic Transformation Preserves Chase Termination

We next show that the magic sets transformation preserves finiteness of the least fixpoint of the programs considered in our pipeline. Thus, applying the chase for logic programs to the result is guaranteed to terminate, as mentioned in Section 6.

Theorem 3. *Let P be a program where the body of each rule is function-free, and let $P_1 = P \cup \mathsf{R}(P) \cup \mathsf{ST}$ and $P_2 = \mathcal{R} \cup \mathsf{R}(\mathcal{R}) \cup \mathsf{ST}$ for $\mathcal{R} = \mathsf{magic}(P)$. For each base instance B, if $T_{P_1}^\infty(B)$ is finite, then $T_{P_2}^\infty(B)$ is finite as well.*

Proof. For F a fact, let $\mathsf{dep}(F)$ be the depth of F as usual (where constants have depth zero). Let M be the maximum depth of an atom in $T_{P_1}^\infty(B)$, and let $I_0, I_1, \ldots$ be the sequence of instances used to compute $T_{P_2}^\infty(B)$ as defined in Section 2. We show by induction on k that, for each fact $F \in I_k$, we have $\mathsf{dep}(F) \leq M$. The base case is trivial, so assume that the claim holds for some I_k and consider an application of a rule $r \in P_2$ to I_k with some σ. If $r \in \mathsf{R}(P) \cup \mathsf{ST}$ or r was added to P_2 in line 12, then P_1 contains a rule r' such that $\mathsf{b}(r') \subseteq \mathsf{b}(r)$; but then, since r' derives on P_1 and B an atom of depth at most M, so does r on P_2 and B. Finally, assume that r was added to P_2 in line 18 for some i. By the inductive assumption, we have $\mathsf{dep}(\sigma(S(\mathbf{t}_i^\gamma))) \leq M$ and $\mathsf{dep}(\sigma(R_j(\mathbf{t}_j))) \leq M$ for $1 \leq j < i$. Moreover, the body of r does not contain function symbols, and so the head atom $\mathsf{h}(r') = S(\mathbf{t}_i^\gamma)$ of r' does not contain a function symbol either, and all variables in $\mathsf{h}(r')$ occur in $\mathsf{b}(r')$. Clearly, we have $\mathsf{dep}(\sigma(\mathsf{h}(r'))) \leq M$, as required. □

B.5 Correctness of Eliminating Function Symbols from Bodies

Next, line 5 removes function symbols from the body Algorithm 1 and thus prepares the result so that it can be evaluated using the chase for logic programs. The correctness of this step is easily verified.

Proposition 2. *For each program P and $P' = \mathsf{defun}(P)$, base instance B, predicate R not of the form F_f, and tuple $\mathbf{t}$ of ground terms, $P \cup \mathsf{R}(P) \cup \mathsf{ST} \cup B \models R(\mathbf{t})$ if and only if $P' \cup \mathsf{R}(P') \cup \mathsf{ST} \cup B \models R(\mathbf{t})$.*

Proof. Let $I_1 = T_{P_1}^\infty(B)$ and $I_2 = T_{P_2}^\infty(B)$, where $P_1 = P \cup \mathsf{R}(P) \cup \mathsf{C}(P) \cup \mathsf{ST}$ and $P_2 = P' \cup \mathsf{R}(P') \cup \mathsf{ST}$. By routine inductions on the construction of I_1 and I_2, one can show that

$$I_2 = I_1 \cup \{\mathsf{F}_f(\mathbf{t}, f(\mathbf{t})) \mid f(\mathbf{t}) \text{ occurs in } I_1 \text{ and } \mathsf{F}_f \text{ occurs in } P'\}.$$

Indeed, each derivation of a fact F in I_1 by a rule $r \in P_1$ corresponds precisely to the derivation of F and $\mathsf{F}_f(\mathbf{t}, f(\mathbf{t}))$ for each term $f(\mathbf{t})$ occurring in F in I_2 by the rules obtained from r by transformations in Definition 6; moreover, each term occurring in some fact with predicate of the form F_f also occurs in a fact with a predicate not of the form F_f and so the reflexivity rules in $\mathsf{R}(P') \setminus \mathsf{R}(P)$ do not derive any new facts. Thus, $P_1 \cup B$ and $P_2 \cup B$ entail the same facts over predicates occurring in P. □

B.6 Correctness of Desingularization

Line 6 of Algorithm 1 desingularizes the rules from the preceding steps, as described in Section 7, by removing all equalities in the rule bodies. If a rule body γ is desingularizing as γ', then clearly γ' is a logical consequence of γ and the congruence axioms for equality (i.e., axioms (2) from Section 2). Since we wish to answer a query Q in the presence of congruence axioms, the desingularized rules clearly return all certain answers. The argument in the other direction (i.e., that adding congruence axioms does not introduce incorrect answers) is more subtle, and it requires analyzing the output of our pipeline as a whole.

Theorem 4. *For each finite set of existential rules Σ defining the query predicate Q, each base instance B, each tuple of constants $\mathbf{a}$, and program P_6 obtained from Σ and B by applying Algorithm 1, $\Sigma \cup B \models_{\approx} \mathsf{Q}(\mathbf{a})$ if and only if $P_6 \cup \mathsf{R}(P_6) \cup \mathsf{C}(P_6) \cup B \models \mathsf{Q}(\mathbf{a})$.*

Proof. Let Σ_1 and P_2–P_6 be as specified in Algorithm 3. Set Σ_1 and programs P_2–P_6 are all $\approx$-safe, and thus desingularization correctly produces a program where each variable in a rule also occurs in the body. Now let $P' = P_6 \cup \mathsf{R}(P_6) \cup \mathsf{C}(P_6) \cup \mathsf{ST}$ and consider a tuple of constants $\mathbf{a}$.

($\Rightarrow$) Assume $\Sigma \cup B \models_{\approx} \mathsf{Q}(\mathbf{a})$. Then, $P_5 \cup \mathsf{R}(P_5) \cup \mathsf{ST} \cup B \models \mathsf{Q}(\mathbf{a})$ holds by (i) Proposition B.1, capturing the properties of singularization, (ii) the properties of Skolemization in Section 2, (iii) Theorems 1 and 2, capturing the correctness of relevance analysis and the magic sets transformation, and (iv) Proposition 2, capturing the correctness of removal of function terms from rule bodies. Now let $P'' = P_5 \cup \mathsf{R}(P_6) \cup \mathsf{C}(P_6) \cup \mathsf{ST}$. Entailment in first-order logic is monotonic as rules are added; moreover, we clearly have $\mathsf{R}(P_5) = \mathsf{R}(P_6)$ and $\mathsf{C}(P_5) = \mathsf{C}(P_6)$, so $P'' \cup B \models \mathsf{Q}(\mathbf{a})$ holds. Let $I_0, I_1, \ldots$ be the sequence of instances used to compute $T^{\infty}_{P''}(B)$ as defined in Section 2. By induction on i, we show that $P' \cup B \models F$ holds for each fact $F \in I_i$, which implies our claim. The base case is obvious, so assume that this property holds for I_i and consider an application of a rule $r \in P''$ with substitution σ such that $\sigma(\mathsf{b}(r)) \subseteq I_i$. Let σ' be obtained from σ by setting $\sigma'(x) = \sigma(t)$ for each variable x replaced with term t when Definition 7 is applied to P_5. For each atom $R(\mathbf{t}) \in \mathsf{b}(r)$, the inductive assumption ensures $P' \cup B \models \sigma(R(\mathbf{t}))$, and the congruence rules in $\mathsf{C}(P_6)$ for R_i clearly ensure $P' \cup B \models \sigma'(R(\mathbf{t}))$. Thus, we have $P' \cup B \models \sigma'(\mathsf{h}(r))$. Finally, the transformation from Definition 7 does not affect variables in the head of r, so we have $\sigma(\mathsf{h}(r)) = \sigma'(\mathsf{h}(r))$.

($\Leftarrow$) Let $P'' = \mathsf{sk}(\Sigma) \cup \mathsf{R}(\mathsf{sk}(\Sigma)) \cup \mathsf{C}(\mathsf{sk}(\Sigma)) \cup \mathsf{ST}$ be the program obtained by Skolemizing Σ and then axiomatizing equality. By the properties of Skolemization from Section 2, we have $\Sigma \cup B \models_{\approx} \mathsf{Q}(\mathbf{a})$ if and only if $P'' \cup B \models \mathsf{Q}(\mathbf{a})$. Let $I_0, I_1, \ldots$ be the sequence of instances used to compute $T^{\infty}_{P'}(B)$ as defined in Section 2. By induction on i, we show that, for each fact $R(\mathbf{t}) \in I_i$ not of the form $s \approx s$ and where R occurs in P'' (i.e., R is not a magic predicate or of the form F_f), we have $P'' \cup B \models R(\mathbf{t})$. The base case holds trivially, so assume that this property holds for I_i and consider an application of a rule $r' \in P'$ with substitution σ such that $\sigma(\mathsf{b}(r')) \subseteq I_i$, the predicate of $\mathsf{h}(r')$ occurs in P'', and $\sigma(\mathsf{h}(r'))$ is not of the form $s \approx s$. Note that r' cannot be a reflexivity axiom (since such axioms always derive $s \approx s$). If r' is a congruence axiom, then we have $r' \in \mathsf{C}(\mathsf{sk}(\Sigma))$, so the property holds by inductive assumption. If $r' \in \mathsf{ST}$ and $\sigma(\mathsf{h}(r'))$ is not of the form $s \approx s$, then r' can derive a new fact only if $\sigma(\mathsf{b}(r'))$ does not contain a fact of such a form, so the property again holds by the inductive assumption. The only remaining possibility is $r' \in P_6$ with the predicate of $\mathsf{h}(r)$ occurring in P''. Rule r' is obtained by singularizing and then Skolemizing an existential rule $\tau \in \Sigma$, possibly desingularizing some equalities in the body during relevance analysis in line 15 of Algorithm 2, adding some magic atom $m^{\alpha}_{R}(\mathbf{t}^{\beta})$ to the body in line 12 of Algorithm 3, possibly transforming away the function symbols from this magic atom, and finally designgularizing the remaining equalities. W.l.o.g. we can assume that all desingularization steps "undo" the effects of the initial singularization (i.e., all results of desingularization are unique up to variable renaming). Moreover, Skolem terms contain only the variables occurring in the rule heads, and these are not affected by singularization and the other transformations, which modify only the bodies; thus, the heads of the rules are the same to what would be produced by Skolemizing τ before singularization. Hence, there exists a rule $r'' \in P''$ such that $\mathsf{h}(r'') = \mathsf{h}(r')$ and $\mathsf{b}(r'') \subseteq \mathsf{b}(r')$. Since the body of r'' does not contain equality atoms, the inductive assumption ensures $P'' \cup B \models \sigma(\mathsf{b}(r''))$, so we have $P'' \cup B \models \sigma(\mathsf{h}(r''))$, as required. □

B.7 Final correctness argument

The final step of Algorithm 1 is Step 7, which simply applies the chase for logic programs to the output of the prior steps.

Theorem 5. *For each finite set of existential rules Σ defining the query predicate Q such that the chase of $\mathsf{sg}(\Sigma)$ terminates on all base instances, and for each base instance B, Algorithm 1 outputs precisely all answers to Q on $\Sigma \cup B$ and then terminates.*

The correctness claim of Theorem 5 follows from combining Proposition 1, the properties of Skolemization form Section 2, and all the results proved in this section. For the termination claim, if the chase of $\mathsf{sg}(\Sigma)$ terminates on every base instance, then line 2 of Algorithm 2 necessarily terminates.