

Oblivious and Semi-Oblivious Boundedness for Existential Rules

Pierre Bourhis^{1,3,4}, Michel Leclère^{2,4}, Marie-Laure Mugnier^{2,4},
 Sophie Tison^{3,4}, Federico Ulliana^{2,4} and Lily Galois^{3,4}

¹ CNRS, France

² Univ. Montpellier, LIRMM, France

³ Univ. Lille, CRISTAL, France

⁴ Inria, France

{firstname.lastname}@inria.fr

Abstract

We study the notion of boundedness in the context of positive existential rules, that is, whether there exists an upper bound to the depth of the chase procedure, that is independent from the initial instance. By focussing our attention on the oblivious and the semi-oblivious chase variants, we give a characterization of boundedness in terms of FO-rewritability and chase termination. We show that it is decidable to recognize if a set of rules is bounded for several classes and outline the complexity of the problem.

1 Introduction

We consider the setting of ontology-based query answering (OBQA) in which answers to conjunctive queries are logically entailed from a knowledge base constituted of a set of facts and an ontology. Existential rules, also known as Tuple Generating Dependencies (TGDs) in database theory, are an expressive knowledge representation language well studied in the OBQA setting [Cali *et al.*, 2009; Baget *et al.*, 2011; Cali *et al.*, 2013]. These rules generalize function-free Horn rules (like those of datalog) with existentially quantified variables in the rule heads, which allow one to assert the existence of unknown individuals, and hence to reason in open domains. Beside datalog, existential rules generalize the Semantic Web language RDF Schema, as well as most Description Logics used in the OBQA context, namely Horn description logics, in particular those at the core of the tractable profiles of the ontological language OWL 2.

The two main approaches developed to answer conjunctive queries on existential rules knowledge bases are *materialization* and *query rewriting*. Both can be seen as ways of reducing query answering to a classical database query evaluation problem. Materialization relies on a forward chaining technique, called the *chase*, that consists in expanding the factbase with the facts entailed by rules until fixpoint. In contrast, query rewriting is a backward chaining mechanism that consists in rewriting an input query using relevant rules, so that its answers on the knowledge base are exactly the answers of the rewritten query on the factbase alone. Query answering being undecidable for existential rules, both materialization and query rewriting may not terminate.

This led to intensive research aiming at characterizing decidable and tractable classes of existential rules. Several syntactic restrictions were proposed to ensure chase termination (e.g. weak-acyclicity [Fagin *et al.*, 2005]) or the existence of a (finite) first-order rewriting of a conjunctive query, a property referred as *FO-rewritability* [Calvanese *et al.*, 2007]. Nevertheless, the interactions between chase termination and FO-rewritability have been little investigated so far, and not much is known for existential rules on which both hold. What are the relationships between these two properties?

Answering this question leads us to another fundamental problem, which has been extensively studied for datalog, namely (*uniform*) *boundedness* [Hillebrand *et al.*, 1995]. Boundedness concerns the recursivity of rules, and asks whether there is an upper bound on the depth of the chase, which is independent from any factbase. The property is key for practical optimization of reasoning as it implies that the ruleset is essentially non-recursive (although syntactic conditions may fail to capture this). It is known that boundedness and FO-rewritability are equivalent in the case of datalog [Ajtai and Gurevich, 1994], but this does not hold for existential rules. Moreover, in the setting of existential rules, the notion of boundedness also depends on the considered chase variant as every chase behaves differently with respect to termination.

We focus our attention on the *oblivious* and *semi-oblivious* (a.k.a. Skolem) chase [Marnette, 2009]. As a matter of fact, almost all known sufficient conditions for chase termination fall within these chase variants (from the simplest ones: rich-acyclicity [Hernich and Schweikardt, 2007], weak-acyclicity [Fagin *et al.*, 2005] and acyclic-GRD [Baget *et al.*, 2011] to the more general MFA [Grau *et al.*, 2013]), at the exception of the recent work of [Carral *et al.*, 2017] which applies to the restricted chase variant. Importantly, we consider a *breadth-first* version for both variants, which ensures the minimal depth of the chase [Delivorias *et al.*, 2018].

Our main contribution is a characterization of boundedness in terms of chase termination and FO-rewritability. This means that a set of rules is bounded if and only if it ensures both chase termination for any instance and FO-rewritability for any conjunctive query. We show this by proving two orthogonal results. The first is a bound on the depth of existential variables when the chase terminates on all instances. The second is a bound on the (breadth-first) rank at which facts using terms of a given depth are inferred.

This connection reveals important differences between the two variants. For the oblivious case we show that, when chase termination holds, FO-rewritability on *full-atomic queries* (queries with a single atom and only answer variables) is equivalent with FO-rewritability. Moreover, for the case of *fully-existential rules* (rules where all head atoms have at least one existential variable), we show that chase termination is equivalent to boundedness and so it implies FO-rewritability. None of these properties hold for the semi-oblivious chase.

Recognizing if a set of existential rules is bounded is undecidable already for datalog [Hillebrand *et al.*, 1995]. However, we show the decidability of the problem for major classes of existential rules as direct corollaries of our characterizations and existing results from the literature. Precisely, the problem is PSpace-complete for linear and sticky rules and in 2Exptime for guarded rules. Finally, we consider the *k-boundedness* problem (i.e., whether the chase terminates in k steps on all instances), which was recently proven decidable for several chase variants, including those investigated here [Delivorias *et al.*, 2018]. We show that deciding if a rule-set is k -bounded is in 2Exptime for the breadth-first (semi-)oblivious chase and co-NExptime-complete for datalog.

2 Preliminary Definitions

We consider a relational vocabulary $\mathcal{V} = (\mathcal{P}, \mathcal{C})$ constituted of a finite set of predicates \mathcal{P} and a finite set of constants \mathcal{C} . A term v is a constant of \mathcal{C} or a variable. An atom is of the form $p(v_1 \dots v_k)$ where p is a predicate of arity k and the v_i are terms. We denote by $\text{terms}()$ the set of its terms and extend the notation to sets of atoms. An *embedding* φ from a set of atoms A to a set of atoms A' is a substitution of $\text{terms}(A)$ with $\text{terms}(A')$ such that $\varphi(A) \subseteq A'$. A *homomorphism* is an embedding which is the identity on constants.

An *instance* I is a conjunction of atoms on constants and (globally) existentially quantified variables. It is finite unless otherwise specified. Throughout this paper, we see an instance I as the set of its atoms and call *fact* any atom f that belongs to this set. Given a finite set \mathcal{P} of predicates, the *critical instance* I_a is composed of all facts built on \mathcal{P} and special constant a . Any instance I on \mathcal{P} can be embedded into I_a .

An *existential rule* σ is a closed formula $\forall \bar{x}\bar{y}(B[\bar{x}, \bar{y}] \rightarrow \exists \bar{z}H[\bar{x}, \bar{z}])$ where B and H are sets of atoms built on variables called the *body* and the *head* of the rule, also denoted by $\text{body}(\sigma)$ and $\text{head}(\sigma)$ respectively. The set of variables \bar{x} shared by B and H is called the *frontier* of the rule and is denoted by $\text{fr}(\sigma)$. The set of variables \bar{z} that belong to H only are called *existential variables* and are denoted by $\text{ex}(\sigma)$. Universal quantifiers will often be omitted in the remainder of the paper. A rule such that $\text{ex}(\sigma) = \emptyset$ is called *datalog*. A rule where all head atoms contain at least one existential variable is called *fully-existential* and denoted by *FE-rule*. We say that a rule σ is applicable on an instance I if there is a homomorphism π from $\text{body}(\sigma)$ to I and call the pair (σ, π) a *trigger* of I . Given a trigger (σ, π) , we denote by $\pi|_{\text{fr}(\sigma)} \subseteq \pi$ the restriction of π to $\text{fr}(\sigma)$.

A *knowledge base* (KB) is a pair (I, Σ) where I is an instance and Σ a set of existential rules. The chase is a fundamental tool for computing logical consequences from a KB

since, when it terminates, it computes a universal model of the KB, i.e., a model that maps by homomorphism to any other model of the KB (with a model being seen here as an instance). In this work, we focus our attention on the breadth-first oblivious (o-chase) and semi-oblivious (so-chase) variants. As discussed in Section 3.3, the breadth-first behavior is particularly interesting when studying boundedness.

Definition 1. Let (I, Σ) be a knowledge base and $\star \in \{\text{o}, \text{so}\}$ a chase variant. Then, the breadth-first \star -chase is defined as follows: $\star\text{-chase}^0(I, \Sigma) = I$ and for all saturation rank $i \geq 0$

$$\star\text{-chase}^{i+1}(I, \Sigma) = \star\text{-chase}^i(I, \Sigma) \cup \bigcup_{(\sigma, \pi)} \pi_\star(\text{head}(\sigma))$$

where (σ, π) is any trigger of $\star\text{-chase}^i(I, \Sigma)$ and $\pi_\star \supseteq \pi$ a substitution that replaces each existential variable $z \in \text{ex}(\sigma)$ with a fresh variable named as follows:

- $\pi_{\text{o}}(z) = z_{(\sigma, \pi)}$
- $\pi_{\text{so}}(z) = z_{(\sigma, \pi|_{\text{fr}(\sigma)})}$

Then, we define $\star\text{-chase}(I, \Sigma) = \bigcup_{i \geq 0} \star\text{-chase}^i(I, \Sigma)$. The \star -chase terminates on (I, Σ) if there is a rank k with $\star\text{-chase}(I, \Sigma) = \star\text{-chase}^k(I, \Sigma)$.

Note that for the o-chase fresh variables are named by the trigger from which they have been generated. Instead, for the so-chase the naming only depends on the frontier-restriction of the homomorphism of the trigger. This means that any two triggers having the same rule and agreeing on the image of its frontier variables produce *equal* results, hence only one of them is actually considered by the so-chase. The so-chase is very close to the *Skolem chase*, which relies on a skolemisation of the rules: first, each rule σ is transformed by replacing each occurrence of an existential variable z with a functional term $f_z^\sigma(\text{fr}(\sigma))$ on the frontier of σ ; then the o-chase is run on the skolemised rules. At each saturation rank, the Skolem chase produces a result isomorphic to that of the so-chase (up to the renaming of each Skolem term by the corresponding fresh variable), hence the forthcoming results on the so-chase also hold for the Skolem chase.

Example 1. Consider the rule $\sigma = p(x, y) \rightarrow \exists z p(x, z)$. Then o-chase(I, Σ) with $I = \{p(a, b)\}$ and $\Sigma = \{\sigma\}$ is infinite - as the chase does not terminate. The atom $p(a, z_{(\sigma, \pi_1)})$ with $\pi_1 = \{x \mapsto a, y \mapsto b\}$ is first inferred, then $p(a, z_{(\sigma, \pi_2)})$ with $\pi_2 = \{x \mapsto a, y \mapsto z_{(\sigma, \pi_1)}\}$, and so on. Here, each rule application enables a new trigger. In contrast, so-chase(I, Σ) is finite, in that only the first rule application will be performed, producing $p(a, z_{(\sigma, \{x \mapsto a\})})$, since all triggers map the frontier variable x to a . For the Skolem chase, σ is rewritten as $\sigma' = p(x, y) \rightarrow p(x, f_z^\sigma(x))$. The first rule application according to trigger (σ', π_1) produces $p(a, f_z^\sigma(a))$, then the chase halts as the same atom is produced by the next trigger.

Definition 2. The rank of a fact $f \in \star\text{-chase}(I, \Sigma)$, denoted by $\text{rank}(f)$, is 0 if $f \in I$ and $1 + \max\{\text{rank}(f') \mid f' \in \pi(\text{body}(\sigma))\}$ if f is produced by the trigger (σ, π) . This definition is naturally extended to terms and sets of facts. The rank of $\star\text{-chase}(I, \Sigma)$ is the smallest k such that $\star\text{-chase}(I, \Sigma) = \star\text{-chase}^k(I, \Sigma)$ if $\star\text{-chase}(I, \Sigma)$ terminates, and it is infinite otherwise.

Note that for the breadth-first chases we consider the above definition implies that $\text{rank}(f)$ is the smallest k such that $f \in \star\text{-chase}^k(I, \Sigma) \setminus \star\text{-chase}^{k-1}(I, \Sigma)$.

An *FO-query* $\phi(x_1, \dots, x_n)$ is a (function free) first-order formula whose free variables (called answer variables) are exactly $\{x_1, \dots, x_n\}$. A *conjunctive query* (CQ) is an FO-query which is an existentially quantified conjunction of atoms. An *atomic query* is a CQ with a single atom. A *full-atomic query* is an atomic query where all terms are free variables. A query is called *Boolean* if it does not have any free variable. As for instances, it will be handful to see CQs as sets of atoms, of course by distinguishing the answer variables. A *union of conjunctive queries* (UCQ) \mathcal{Q} is a disjunction of CQs with the same free variables, also seen as a set of CQs.

A tuple of constants $(a_1, \dots, a_n) \in \mathcal{C}^n$ is an *answer* to a CQ $Q(x_1, \dots, x_n)$ on an instance I if there is a homomorphism h from Q to I such that $h(x_i) = a_i$ for $1 \leq i \leq n$. Equivalently, $I \models Q[x_i \mapsto a_i]$, where \models denotes the classical logical consequence and $Q[x_i \mapsto a_i]$ is the Boolean query obtained from Q substituting each x_i with a_i . A tuple of constants $(a_1, \dots, a_n) \in \mathcal{C}^n$ is a *certain answer* to Q on a KB (I, Σ) if $I, \Sigma \models Q[x_i \mapsto a_i]$. This is equivalent to the existence of a saturation rank k such that $\star\text{-chase}^k(I, \Sigma) \models Q[x_i \mapsto a_i]$. In other words, the certain answers to Q on (I, Σ) are exactly its answers on the possibly infinite instance $\star\text{-chase}(I, \Sigma)$. The set of (certain) answers to a UCQ \mathcal{Q} is the union of the sets of (certain) answers to the CQs it contains.

2.1 Termination vs Boundedness

To begin our study, we need to present the relationships between chase termination and boundedness. Let $\star \in \{\text{o}, \text{so}\}$ be a chase variant, the \star -chase termination class, denoted by CT^\star , contains all rulesets Σ such that $\star\text{-chase}(I, \Sigma)$ terminates for all instances I . The \star -boundedness class, denoted by BN^\star , contains all bounded rulesets Σ , i.e., for which there exists an integer k such that $\star\text{-chase}^k(I, \Sigma) = \star\text{-chase}(I, \Sigma)$ for all instances I . Obviously, $\text{BN}^\star \subset \text{CT}^\star$.

Example 2. Let $\sigma_1 = p(x, y) \wedge p(y, z) \rightarrow p(x, z)$ and $\sigma_2 = p(x, y) \wedge p(w, z) \rightarrow p(x, z)$. Because both rules are datalog, $\{\sigma_1\} \in \text{CT}^\star$ and $\{\sigma_2\} \in \text{CT}^\star$. However, $\Sigma = \{\sigma_1\} \notin \text{BN}^\star$, since the rank of $\star\text{-chase}(I, \Sigma)$ depends on I . In contrast, $\{\sigma_2\} \in \text{BN}^\star$ and the bound is $k = 1$. Similarly, $\{\sigma_1, \sigma_2\} \in \text{BN}^\star$. Indeed, σ_2 produces at the first rank all atoms that can be produced by σ_1 at later ranks.

To get a better understanding of boundedness, it will be useful to decompose each rule of a set thereby distinguishing between its “datalog part” and its “existential part”. For instance, a rule of the form $p(x, y) \rightarrow \exists z p(x, z) \wedge q(x)$ can be decomposed into a datalog rule $p(x, y) \rightarrow q(x)$ and an FE-rule $p(x, y) \rightarrow \exists z p(x, z)$. Let σ be any existential rule of the form $B \rightarrow H_F \wedge H_D$ where B is the set of body atoms, H_F is the set of head atoms with at least one existential variable and H_D are the remaining head atoms. The datalog-fully existential decomposition of σ , denoted by $\text{DF}(\sigma)$, returns a set made of the FE-rule $B \rightarrow H_F$ together with a (single head) datalog rule of the form $B \rightarrow H_D^i$, for each $H_D^i \in H_D$. The definition is then extended to sets $\text{DF}(\Sigma) = \bigcup_{\sigma \in \Sigma} \text{DF}(\sigma)$. This decomposition preserves boundedness and termination.

Proposition 1. $\Sigma \in \text{CT}^\star$ iff $\text{DF}(\Sigma) \in \text{CT}^\star$ and $\Sigma \in \text{BN}^\star$ iff $\text{DF}(\Sigma) \in \text{BN}^\star$, when $\star \in \{\text{o}, \text{so}\}$.

3 Upper Bounding the Chase Depth

Our approach consists of defining a notion of *existential depth* for facts, proper to each chase, which is finite on a given instance if and only if the chase terminates on that instance. Then we show that for each chase, the existential depth of all facts produced by the chase for a given ruleset are bounded by those of the critical instance. This means that whenever the chase terminates on the critical instance there is an upper bound to the existential depth of the facts, for all instances. In the next section, with these results in hand, we use FO-rewritability to bound the rank at which any fact of a certain existential depth will be inferred. This will give us a characterization of boundedness for the oblivious and so-chase in terms of FO-rewritability and chase termination.

3.1 The Oblivious Case

Intuitively, the notion of existential depth of a term measures the number of fresh variable generation steps that led to the creation of this term.

Definition 3. The *existential depth* (or *simply depth*) of a term v that belongs to $\text{o-chase}(I, \Sigma)$ is

$$\text{depth}_\exists(v) = \begin{cases} 0 & \text{if } v \in \text{terms}(I) \\ 1 + \max\{\text{depth}_\exists(v_B)\} & \text{otherwise} \end{cases}$$

where v_B is any term in $\pi(\text{body}(\sigma))$ used by a trigger (σ, π) which generates v . The *existential depth of a fact f* is the *maximum existential depth of its terms*. The *existential depth of $\text{o-chase}(I, \Sigma)$* is the *maximum existential depth of its facts if it is finite and is infinite otherwise*.

To illustrate the definition, consider Example 1. The existential depth of terms in $\text{o-chase}(I, \{\sigma\})$ is unbounded, which is in line with the non-termination of the o-chase on $(I, \{\sigma\})$. The rule σ_1 in Example 2. shows the difference between rank and existential depth. For any I , the existential depth of terms (hence facts) is 0 because σ_1 is datalog, however their rank depends on I . More generally, for any term v and fact f in $\text{o-chase}(I, \Sigma)$ it holds that $\text{depth}_\exists(v) \leq \text{rank}(v)$ and $\text{depth}_\exists(f) \leq \text{rank}(f)$. Hence, if $\text{o-chase}(I, \Sigma)$ terminates, its existential depth is finite. Reciprocally, when the existential depth of $\text{o-chase}(I, \Sigma)$ is finite, so it is the number of its terms, and $\text{o-chase}(I, \Sigma)$ terminates. We point out that when dealing with sets of *FE-rules* the notions of rank and existential depth coincide, as illustrated by Example 1.

Proposition 2. If Σ is a set of FE-rules then, for all instance I and term v in $\text{o-chase}(I, \Sigma)$, holds that $\text{depth}_\exists(v) = \text{rank}(v)$.

It should be clear that, for a given ruleset, the o-chase may have unbounded rank even when it terminates on all instances (see for instance Example 2). Nevertheless, when a ruleset is in CT° , our goal is to show that there exists a bound on the *existential depth* of its terms, which holds for all instances. Aiming at this, we present a lemma stating that existential depth of terms are preserved by embeddings.

Lemma 3. For any embedding φ from I to I' and any $i \geq 0$, there exists an embedding $\varphi' \supseteq \varphi$ from $\text{o-chase}^i(I, \Sigma)$ to $\text{o-chase}^i(I', \Sigma)$ which preserves the existential depth of terms, i.e., for every term v in $\text{o-chase}(I, \Sigma)$ it holds that $\text{depth}_{\exists}(v) = \text{depth}_{\exists}(\varphi'(v))$.

It is well-known that the o-chase terminates on all instances if and only if it terminates on the critical instance [Marnette, 2009]. We leverage this property to compute a bound on the existential depth under chase termination.

Theorem 4. When $\Sigma \in \text{CT}^\circ$ there exists a constant k_d such that for every instance I , the existential depth of a term in $\text{o-chase}(I, \Sigma)$ is bounded by k_d .

Proof. Because $\Sigma \in \text{CT}^\circ$, the o-chase terminates on the critical instance I_a . Let k_d be the largest rank such that $\text{terms}(\text{o-chase}^{k_d}(I_a, \Sigma)) \setminus \text{terms}(\text{o-chase}^{k_d-1}(I_a, \Sigma)) \neq \emptyset$. Every instance I can be embedded into I_a . By Lemma 3 the existential depth of the terms in $\text{o-chase}(I, \Sigma)$ is bounded by that of $\text{o-chase}(I_a, \Sigma)$, which is in turn bounded by k_d . \square

Chase termination is a necessary condition for boundedness as it bounds the existential depths of the variables generated by the chase - but not the rank (see the datalog case). Interestingly, for FE-rules, chase termination also becomes a sufficient condition for boundedness, because the notion of rank and existential depth coincide (Proposition 2).

Corollary 5. For Σ a set of FE-rules, $\Sigma \in \text{CT}^\circ$ iff $\Sigma \in \text{BN}^\circ$.

For general existential rules, we will later show that when a restricted form of FO-rewritability holds, one can also provide a bound to the rank of the o-chase (Theorem 14).

3.2 The Semi-Oblivious Case

When applied to the so-chase, the previous notion of existential depth is not preserved by embedding, which hinders the possibility of using the critical instance to bound the existential depth of terms. As illustrated below, this is due to the fact that the so-chase makes equal the result of two distinct triggers agreeing on a rule frontier.

Example 3. Consider $I = \{p(a, b)\}$, $I' = I \cup \{r(a, b)\}$ and $\Sigma = \{\sigma_1 : p(x, y) \rightarrow \exists z r(z, y) \quad \sigma_2 : r(x, y) \rightarrow \exists z s(y, z)\}$. Then, $\text{so-chase}^2(I, \Sigma) = I \cup \{r(z_{(\sigma_1, \pi)}, b) \cup s(b, z_{(\sigma_2, \pi)})\}$ with $\pi = \{y \mapsto b\}$. Also, $\text{so-chase}^2(I, \Sigma) \subseteq \text{so-chase}^1(I', \Sigma)$ because all triggers applied by the chase from I are already applicable on I' . The application of σ_2 on $r(a, b)$ and $r(z_{(\sigma_1, \pi)}, b)$ gives equal results, hence $\text{so-chase}^1(I', \Sigma) = \text{so-chase}^2(I', \Sigma)$. In the embeddings from $\text{so-chase}^2(I, \Sigma)$ to $\text{so-chase}^2(I', \Sigma)$, $z_{(\sigma_2, \pi)}$ is mapped to itself, but both occurrences have different existential depth (resp. 2 and 1).

It is therefore natural to turn to the following notion of depth, which accounts for frontier terms only.

Definition 4. The frontier existential depth (or simply frontier depth) of a term v that belongs to $\text{so-chase}(I, \Sigma)$ is

$$\text{depth}_{\exists}^{\text{fr}}(v) = \begin{cases} 0 & \text{if } v \in \text{terms}(I) \\ 1 & \text{if } \text{fr}(\sigma) = \emptyset \\ 1 + \max\{\text{depth}_{\exists}^{\text{fr}}(v_B)\} & \text{otherwise} \end{cases}$$

where v_B is any term in $\pi(\text{fr}(\sigma))$ used by a trigger (σ, π) which generates v . Accordingly, the frontier depth of a fact f is the maximum frontier depth of its terms. The frontier depth of so-chase (I, Σ) is defined as the maximum frontier depth of its facts if it is finite and is infinite otherwise.

Note that frontier depth coincides with the (usual) depth of terms generated by the Skolem chase.

Clearly, $\text{depth}_{\exists}^{\text{fr}}(v) \leq \text{depth}_{\exists}(v)$ for all v in $\text{o-chase}(I, \Sigma)$. The following example illustrates the difference between the two notions of (existential) depth.

Example 4. Let $\Sigma = \{\sigma = p(x, y, u) \rightarrow \exists z p(y, x, z)\}$. Starting from $I = \{p(a, b, c)\}$, the o-chase generates an infinite number of fresh variables v with increasing $\text{depth}_{\exists}(v)$. The rank of the so-chase is instead 2 and for each fresh variable v , $\text{depth}_{\exists}^{\text{fr}}(v) = 1$ as all triggers map $\text{fr}(\sigma)$ to $\text{terms}(I)$.

It is worth noting that not only the oblivious notion of existential depth is not effective for studying the so-chase, but also that the frontier depth is not well characterizing the behavior of the o-chase either. The crux is that the finiteness of the frontier depth cannot be related with the termination of the o-chase, as illustrated by Example 4. Using such a notion to study the o-chase would impede us, for instance, to establish Corollary 5, which relies on the fact that rank and existential depth coincide for the oblivious-chase (Property 2).

We are now ready to show that the frontier depth is preserved by embeddings. The next lemma and theorem are the counter-parts of Lemma 3 and Theorem 4 for the so-chase.

Lemma 6. For any embedding φ from I to I' and any $i \geq 0$, there exists an embedding $\varphi' \supseteq \varphi$ from $\text{so-chase}^i(I, \Sigma)$ to $\text{so-chase}^i(I', \Sigma)$ which preserves the frontier depth of terms.

Theorem 7. When $\Sigma \in \text{CT}^{\text{so}}$ there exists a constant k_d such that for every instance I , the frontier depth of a term in $\text{so-chase}(I, \Sigma)$ is bounded by k_d .

3.3 On the Interest of the Breadth-First Chase

We conclude this section with some remarks on the interest of studying boundedness for breadth-first chases. We assume that the reader is familiar with the notion of chase sequence.¹ We define the rank of a chase sequence on (I, Σ) as the maximal rank of its facts if it is finite, and infinite otherwise.

For the (semi-)oblivious chase, it is well-known that there is a terminating chase sequence for (I, Σ) if and only if all chase sequences for (I, Σ) terminate. However, not all terminating chase sequences have the same rank, and the minimal rank is obtained with breadth-first sequences [Delivorias et al., 2018]. This makes the notion of boundedness we consider equivalent to studying whether there exists a bound such that, for all instance, there exists a terminating chasing sequence whose rank is within the bound. Hence, it characterizes the fact that the chase can indeed terminate within that bound, if a strategy ensuring a minimal sequence rank is followed.

¹A chase sequence is any sequence of triggers satisfying the applicability criterion of the chase variant. For the oblivious chase, the same trigger should not be applied twice. For the semi-oblivious chase a trigger is not applied if a trigger for the same rule assigning the same image for the frontier variables has been applied before.

It is therefore natural to consider breadth-first chases which achieve this property, like the (semi-)oblivious chase. Example 2 illustrates this concept and shows that, already for datalog, the rank of some chase sequences may be not bounded, while the rank of all breadth-first sequences is bounded. This happens for instance if all applications of the transitivity rule σ_2 are performed before the rule σ_1 .

In the special case of FE-rules, it is not hard to see that all oblivious chase sequences for (I, Σ) have the same rank. However, this does not hold for the semi-oblivious chase. Below, a variation of Example 2, where some dummy variables are introduced, illustrates this point.

Example 5. Let $\Sigma = \{\sigma_1, \sigma_2\}$, with $\sigma_1 = p(x, y, t) \wedge p(y, z, u) \rightarrow \exists v p(x, z, v)$ and $\sigma_2 = p(x, y, t) \wedge p(w, z, u) \rightarrow \exists v p(x, z, v)$. The rank of so-chase (I, Σ) is bounded by 2 for any I , while again performing all applications of σ_2 before σ_1 gives derivations of different ranks.

4 The Impact of First Order Rewritability

We now turn our attention to FO-rewritability and show that it yields a bound on the rank of specific (sets of) facts that share terms with the initial instance I . For the o-chase, we bound the rank of facts that have all their terms in I . For the so-chase, we consider triggers that map a rule frontier to terms of I : we do not bound the rank of facts that allow to fire such triggers, but we show that for each such trigger $t = (\sigma, \pi)$, there is a trigger $t' = (\sigma, \pi')$ that agrees with t on the mapping of $\text{fr}(\sigma)$ and that is fired at a bounded rank. In Section 5, we will leverage these results to show that FO-rewritability yields a bound on the rank of all facts with a certain existential depth. For the o-chase, a restricted version of FO-rewritability is sufficient to get these properties.

We say that a pair (Q, Σ) is *FO-rewritable* (resp. UCQ-rewritable) if there is an FO-query (resp. a UCQ) \mathcal{Q} such that, for all I , the certain answers to Q on (I, Σ) are exactly the answers to \mathcal{Q} on I . It is known that FO-rewritability is equivalent to UCQ-rewritability². A set of rules Σ is *FO-rewritable* (or equivalently, UCQ-rewritable) if (Q, Σ) is FO-rewritable for every CQ Q . We denote by FO-R the class of FO-rewritable rulesets. We will also consider specific classes of CQs. Given a class of CQs \mathcal{C} , we say that a ruleset Σ is FO-rewritable with respect to \mathcal{C} if (Q, Σ) is FO-rewritable for all $Q \in \mathcal{C}$. We denote by FO-R ^{\mathcal{C}} the corresponding class. We first point out that FO-rewritability with respect to full-atomic queries, denoted by FO-R^{AF}, is a strictly weaker property than FO-rewritability.

Proposition 8. FO-R^{AF} \supset FO-R

Proof. The inclusion holds by definition, and to see that it is strict consider $\Sigma = \{\sigma = p(x, x_1), p(x_1, x_2), p(x_2, z) \rightarrow \exists y p(x, y), p(y, z)\}$. Σ is not FO-rewritable as for the Boolean query $Q = \{p(a, u), p(u, b)\}$, where a and b are constants, (Q, Σ) is not FO-rewritable (we would need an infinite union of Boolean CQs of the form $\{p(a, u_0), \dots, p(u_{i-1}, u_i), p(u_i, b)\}$, none of these queries being contained in another). However, $\Sigma \in \text{FO-R}^{\text{AF}}$ as (Q, Σ)

²It follows from the Homomorphism preservation theorem, a classical result in model theory.

is FO-rewritable for any $Q \in \text{AF}$. Indeed, σ cannot bring any answer to such query (in more technical terms, an existential variable of σ cannot be unified with an answer variable). \square

Note also that since full-atomic queries have only answer variables, they cannot be rewritten by means of FE-rules. Thus, every set of FE-rules is trivially in FO-R^{AF}. More interestingly, to check if $\Sigma \in \text{FO-R}^{\text{AF}}$ one can restrict the full-atomic queries of interest to those corresponding to the heads of the datalog rules yielded by the DF-decomposition of Σ .

Proposition 9. Let Σ be a ruleset and HD_Σ be the full-atomic queries given by heads of the datalog rules in $\text{DF}(\Sigma)$. Then, $\Sigma \in \text{FO-R}^{\text{AF}}$ if and only if $\Sigma \in \text{FO-R}^{\text{HD}_\Sigma}$.

The following lemma upper bounds the rank of all facts with terms in I for sets of rules enjoying FO-rewritability on full-atomic queries.

Lemma 10. If $\Sigma \in \text{FO-R}^{\text{AF}}$ there is a constant k_{AF} such that, for any instance I and fact f such that $\text{terms}(f) \subseteq \text{terms}(I)$, when $f \in \text{o-chase}(I, \Sigma)$ it holds that $\text{rank}(f) \leq k_{\text{AF}}$.

Proof. The number of (non-isomorphic) full-atomic queries to be considered is finite, as for Proposition 9. We take for k_{AF} the maximal number of breadth-first rewriting steps necessary to obtain a UCQ-rewriting of a full-atomic query (we refer here to the breadth-first rewriting based on aggregated piece-unifiers, see [König et al., 2013]). \square

The previous lemma also holds for the so-chase, however we want to derive a bound on the rank of facts with a certain frontier depth, and for that full-atomic rewritability is not enough. To illustrate, consider $\Sigma = \{\sigma = p(x, y, u), p(y, z, v) \rightarrow \exists w p(x, z, w)\}$. Here $\Sigma \in \text{FO-R}^{\text{AF}}$ (the only rewriting of a full-atomic query is the query itself because of the existential variable w). For any instance I , the frontier depth of facts in the so-chase is bounded by 1, however there is no bound on their rank (although the so-chase terminates). Therefore, we give a different property for the so-chase, which requires the power of FO-rewritability.

Lemma 11. If $\Sigma \in \text{FO-R}$ there is a constant k_{FO} such that, for any instance I and any trigger (σ, π) from so-chase (I, Σ) with $\pi(\text{fr}(\sigma)) \subseteq \text{terms}(I)$, there is also a trigger (σ, π') from so-chase (I, Σ) such that $\pi'_{|\text{fr}(\sigma)} = \pi_{|\text{fr}(\sigma)}$ and $\text{rank}(f) \leq k_{\text{FO}}$ for all $f \in \pi'(\text{body}(\sigma))$.

Proof. Similar to the proof of Lemma 10 but considering CQs of the form $Q_{\text{body}(\sigma)}$ whose atoms correspond to the atoms of $\text{body}(\sigma)$, for $\sigma \in \Sigma$, and all variables are existentially quantified except for those in $\text{fr}(\sigma)$. The number of such queries is bounded by the cardinal of Σ . We take for k_{FO} the maximal number of breadth-first rewriting steps necessary to obtain a UCQ-rewriting from any $Q_{\text{body}(\sigma)}$ query. The proof actually shows that FO-rewritability with respect to rule body queries is sufficient to derive the lemma. \square

5 Boundedness: Linking Depth and Rank

We can finally establish a connection between the rank and depth of a fact when the chase is run on FO-rewritable sets of rules. This will immediately lead us to a characterization of boundedness for the oblivious and semi-oblivious chases.

Theorem 12. *If $\Sigma \in \text{FO-R}^{\text{AF}}$ then for all instance I and fact $f \in \text{o-chase}(I, \Sigma)$ we have that $\text{rank}(f) \leq \text{depth}_{\exists}^{\text{fr}}(f) \times (k_{\text{AF}} + 1) + k_{\text{AF}}$ with k_{AF} the bound provided by Lemma 10.*

Theorem 13. *If $\Sigma \in \text{FO-R}$ then for all instance I and fact $f \in \text{so-chase}(I, \Sigma)$ we have that $\text{rank}(f) \leq \text{depth}_{\exists}^{\text{fr}}(f) \times (k_{\text{FO}} + 1) + k_{\text{FO}}$ with k_{FO} the bound provided by Lemma 11.*

For the o-chase, boundedness is exactly termination and FO-rewritability on full-atomic queries. Furthermore, for rulesets in CT° , the notions of FO-R and FO-R^{AF} coincide.

Theorem 14. $\text{FO-R}^{\text{AF}} \cap \text{CT}^{\circ} = \text{BN}^{\circ} = \text{FO-R} \cap \text{CT}^{\circ}$

Proof. We start by showing that $\text{BN}^{\circ} \subseteq \text{FO-R} \cap \text{CT}^{\circ}$. By definition $\text{BN}^{\circ} \subseteq \text{CT}^{\circ}$. Then, $\text{BN}^{\circ} \subseteq \text{FO-R}$ follows from the equivalence between FO-R and the bounded-depth derivation property [Gottlob *et al.*, 2014]. Moreover, by Proposition 8 we have $\text{BN}^{\circ} \subseteq \text{FO-R}^{\text{AF}} \cap \text{CT}^{\circ}$. To conclude the proof, by Theorem 4 and 12 we have that $\text{FO-R}^{\text{AF}} \cap \text{CT}^{\circ} \subseteq \text{BN}^{\circ}$ and again by Proposition 8 follows $\text{FO-R} \cap \text{CT}^{\circ} \subseteq \text{BN}^{\circ}$. \square

For the so-chase, boundedness can be characterized again as termination and FO-rewritability by Theorem 7 and 13.

Theorem 15. $\text{BN}^{\text{so}} = \text{FO-R} \cap \text{CT}^{\text{so}}$

Summing up, we have the following differences between boundedness for o-chase and so-chase. o-chase-boundedness requires *i)* o-chase termination and full-atomic-rewritability and *ii)* is equivalent to o-chase termination for FE-rules. Intuitively, when a set of rules Σ is decomposed into $\text{DF}(\Sigma)$, the fully-existential part may cause non-termination of the o-chase, while the datalog part may cause non-FO-rewritability. Furthermore, the fully-atomic queries possibly leading to infinite rewritings in this case correspond to the heads of the datalog rules. Note however that this restricted form of FO-rewritability has still to be verified with respect to the whole set of rules. In contrast, so-chase-boundedness *i)* requires a stronger form of FO-rewritability and *ii)* FE-rules do not behave differently from general existential rules for this chase. Intuitively, for the so-chase, any existential rule (even an FE-rule) has an “underlying” datalog rule. This is illustrated by the following transformation. To each rule σ in Σ we assign a special predicate p_{σ} of arity $|\text{fr}(\sigma)|$. $\Psi(\Sigma)$ is obtained from Σ by replacing each rule $\sigma = B \rightarrow H$ with two rules: a datalog rule $B \rightarrow p_{\sigma}(\text{fr}(\sigma))$ and a rule $p_{\sigma}(\text{fr}(\sigma)) \rightarrow H$. It can be shown that $\Sigma \in \text{CT}^{\text{so}}$ iff $\Psi(\Sigma) \in \text{CT}^{\circ}$ and that $\Sigma \in \text{BN}^{\text{so}}$ iff $\Psi(\Sigma) \in \text{BN}^{\circ}$. This may also provide an alternative path to study so-chase boundedness by reducing it to o-chase boundedness.

6 Decidability and Complexity

From the undecidability of (uniform) boundedness of datalog [Hillebrand *et al.*, 1995], we immediately obtain the undecidability of membership to BN° and BN^{so} . A notable class of datalog rules with decidable boundedness (more precisely in linear time) is chain datalog [Guessarian and Peixoto, 1994]. We obtain that membership to BN^{so} , CT^{so} and FO-R remains undecidable for *FE-rules*, while the decidability of membership to BN° , hence to CT° , is still open.

Importantly, new decidability and complexity results about boundedness for specific existential rules studied in the literature can be obtained as direct corollaries of our results. This is in particular the case for classes known to be FO-rewritable.

Corollary 16. *For any class of existential rules $\mathcal{C} \in \text{FO-R}$, it holds that: $\mathcal{C} \in \text{BN}^{\circ}$ iff $\mathcal{C} \in \text{CT}^{\circ}$, and $\mathcal{C} \in \text{BN}^{\text{so}}$ iff $\mathcal{C} \in \text{CT}^{\text{so}}$.*

This implies that membership to BN° and BN^{so} is PSpace-complete for the two main classes of FO-rewritable existential rules, namely *linear* and *sticky*. Indeed, deciding CT° and CT^{so} is PSpace-complete for both [Calautti *et al.*, 2015; Calautti and Pieris, 2019]. We also get an upper bound on the complexity of membership to BN° and BN^{so} for a major class of existential rules, namely *guarded*. This class is neither CT^{so} nor FO-R. However, membership to CT° and CT^{so} for guarded rules is decidable in 2Exptime [Calautti *et al.*, 2015]. Then a careful reduction from [Barceló *et al.*, 2018] allows us to set the result. The paper shows that checking FO-rewritability for a single query under guarded rules is in 2Exptime. This suffices since by Lemma 10 and 11 we need to test only a polynomial number of queries.

We conclude by considering the k -boundedness problem, which asks whether the chase actually halts within k steps. The problem is decidable for the breadth-first (semi-)oblivious chase and any set of existential rules [Delivorias *et al.*, 2018]. Therefore, the k -boundedness question becomes interesting for dealing with fragments of existential rules where boundedness is undecidable. We study here the complexity of the following version of the problem. Given a ruleset Σ and a (unary encoded) integer k , does it hold that $\star\text{-chase}^k(I, \Sigma) = \star\text{-chase}(I, \Sigma)$ for all instance I ?

Theorem 17. *Deciding k -boundedness is in 2Exptime for existential rules for the o-chase and so-chase; co-NExptime-complete for datalog rules; in co-NExptime on FE-rules for the o-chase.*

Proof. (Sketch) The membership results rely on the decidability arguments from [Delivorias *et al.*, 2018]. Co-NExptime-hardness for datalog is by reduction from the co-NExptime-hard inclusion problem of non-recursive Boolean datalog queries [Benedikt and Gottlob, 2010]. \square

7 Outline and Perspectives

In this paper, we have characterized boundedness in terms of FO-rewritability and chase termination, for the oblivious and semi-oblivious chase variants. We conclude with a discussion on the extent of our results to more powerful chase variants (i.e., which terminate at least when the semi-oblivious chase terminate). Theorem 13 suggests that whenever $\Sigma \in \text{FO-R}$ if any such chase generates only terms of bounded *frontier depth* on all instances, then Σ is bounded. We leave open the question to determine if for other chase variants, like the restricted and the core chases, boundedness is again the intersection of chase termination and FO-rewritability.

Acknowledgements

Work supported by ANR projects CQFD (ANR-18-CE23-0003), DataCert (ANR-15-CE39-0009), DeLTA (ANR-16-CE40-0007) and CNRS-Momentum project Managing-Data.

References

- [Ajtai and Gurevich, 1994] Miklos Ajtai and Yuri Gurevich. Datalog vs first-order logic. 1994.
- [Baget *et al.*, 2011] Jean-François Baget, Michel Leclère, Marie-Laure Mugnier, and Eric Salvat. On rules with existential variables: Walking the decidability line. *Artif. Intell.*, 175(9-10):1620–1654, 2011.
- [Barceló *et al.*, 2018] Pablo Barceló, Gerald Berger, Carsten Lutz, and Andreas Pieris. First-order rewritability of frontier-guarded ontology-mediated queries. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden.*, pages 1707–1713, 2018.
- [Benedikt and Gottlob, 2010] Michael Benedikt and Georg Gottlob. The impact of virtual views on containment. *PVLDB*, 3(1):297–308, 2010.
- [Calautti and Pieris, 2019] Marco Calautti and Andreas Pieris. Oblivious chase termination: The sticky case. In *22nd International Conference on Database Theory, ICDT 2019, March 26-28, 2019, Lisbon, Portugal*, pages 17:1–17:18, 2019.
- [Calautti *et al.*, 2015] Marco Calautti, Georg Gottlob, and Andreas Pieris. Chase termination for guarded existential rules. In *Proceedings of the 34th ACM Symposium on Principles of Database Systems, PODS 2015, Melbourne, Victoria, Australia, May 31 - June 4, 2015*, pages 91–103, 2015.
- [Calì *et al.*, 2009] Andrea Calì, Georg Gottlob, and Thomas Lukasiewicz. Datalog extensions for tractable query answering over ontologies. In *Semantic Web Information Management - A Model-Based Perspective*, pages 249–279. 2009.
- [Calì *et al.*, 2013] Andrea Calì, Georg Gottlob, and Michael Kifer. Taming the infinite chase: Query answering under expressive relational constraints. *J. Artif. Intell. Res.*, 48:115–174, 2013.
- [Calvanese *et al.*, 2007] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family. *J. Autom. Reasoning*, 39(3):385–429, 2007.
- [Carral *et al.*, 2017] David Carral, Irina Dragoste, and Markus Krötzsch. Detecting chase (non)termination for existential rules with disjunctions. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 922–928, 2017.
- [Delivorias *et al.*, 2018] Stathis Delivorias, Michel Leclère, Marie-Laure Mugnier, and Federico Ulliana. On the *k*-boundedness for existential rules. In *International Joint Conference on Rules and Reasoning*, pages 48–64. Springer, 2018.
- [Fagin *et al.*, 2005] Ronald Fagin, Phokion G. Kolaitis, Renée J. Miller, and Lucian Popa. Data exchange: semantics and query answering. *Theor. Comput. Sci.*, 336(1):89–124, 2005.
- [Gottlob *et al.*, 2014] Georg Gottlob, Stanislav Kikot, Roman Kontchakov, Vladimir Podolskii, Thomas Schwentick, and Michael Zakharyashev. The price of query rewriting in ontology-based data access. *Artificial Intelligence*, 213:42–59, 2014.
- [Grau *et al.*, 2013] Bernardo Cuenca Grau, Ian Horrocks, Markus Krötzsch, Clemens Kupke, Despoina Magka, Boris Motik, and Zhe Wang. Acyclicity notions for existential rules and their application to query answering in ontologies. *J. Artif. Intell. Res.*, 47:741–808, 2013.
- [Guessarian and Peixoto, 1994] Irène Guessarian and Marcos Veloso Peixoto. About boundedness for some datalog and datalog^{neg} programs. *J. Log. Comput.*, 4(4):375–403, 1994.
- [Hernich and Schweikardt, 2007] André Hernich and Nicole Schweikardt. Cwa-solutions for data exchange settings with target dependencies. In *Proceedings of the Twenty-Sixth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 11-13, 2007, Beijing, China*, pages 113–122, 2007.
- [Hillebrand *et al.*, 1995] Gerd G Hillebrand, Paris C Kanelakis, Harry G Mairson, and Moshe Y Vardi. Undecidable boundedness problems for datalog programs. *The Journal of logic programming*, 25(2):163–190, 1995.
- [König *et al.*, 2013] Mélanie König, Michel Leclère, Marie-Laure Mugnier, and Michaël Thomazo. On the exploration of the query rewriting space with existential rules. In *Web Reasoning and Rule Systems - 7th International Conference, RR 2013, Mannheim, Germany, July 27-29, 2013. Proceedings*, pages 123–137, 2013.
- [Marnette, 2009] Bruno Marnette. Generalized schema-mappings: from termination to tractability. In *Proceedings of the Twenty-Eighth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2009, June 19 - July 1, 2009, Providence, Rhode Island, USA*, pages 13–22, 2009.