



HAL
open science

On the k -Boundedness for Existential Rules

Stathis Delivorias, Michel Leclère, Marie-Laure Mugnier, Federico Ulliana

► **To cite this version:**

Stathis Delivorias, Michel Leclère, Marie-Laure Mugnier, Federico Ulliana. On the k -Boundedness for Existential Rules. RuleML+RR, Sep 2018, Luxembourg, Luxembourg. pp.48-64, 10.1007/978-3-319-99906-7_4. lirmm-01921140

HAL Id: lirmm-01921140

<https://hal-lirmm.ccsd.cnrs.fr/lirmm-01921140v1>

Submitted on 13 Nov 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the k -Boundedness for Existential Rules

Stathis Delivorias, Michel Leclère, Marie-Laure Mugnier, and Federico Ulliana

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

Abstract. The chase is a fundamental tool for existential rules. Several chase variants are known, which differ on how they handle redundancies possibly caused by the introduction of nulls. Given a chase variant, the halting problem takes as input a set of existential rules and asks if this set of rules ensures the termination of the chase for any factbase. It is well-known that this problem is undecidable for all known chase variants. The related problem of boundedness asks if a given set of existential rules is bounded, i.e., whether there is a predefined upper bound on the number of (breadth-first) steps of the chase, independently from any factbase. This problem is already undecidable in the specific case of datalog rules. However, knowing that a set of rules is bounded for some chase variant does not help much in practice if the bound is unknown. Hence, in this paper, we investigate the decidability of the k -boundedness problem, which asks whether a given set of rules is bounded by an integer k . We prove that k -boundedness is decidable for three chase variants, namely the oblivious, semi-oblivious and restricted chase.

This report is a revised version of the paper published at RuleML+RR 2018.

1 Introduction

Existential rules (see [CGK08,BLMS09,CGL09] for the first papers and [GOPS12,MT14] for introductory courses) are a positive fragment of first-order logic that generalizes the deductive database query language Datalog and knowledge representation formalisms such as Horn description logics (see e.g. [CGL⁺05,KRH07,LTW09]). These rules offer the possibility to model the existence of unknown individuals by means of existentially quantified variables in rule heads, which enables reasoning on incomplete data with the open-domain assumption. Existential rules have the same logical form as database constraints known as tuple-generating dependencies, which have long been investigated [AHV95]. Reborn under the names of existential rules, Datalog[∃] or Datalog⁺, they have raised significant interest in the last years as ontological languages, especially for the ontology-mediated query-answering and data-integration issues.

A knowledge base (KB) is composed of a set of existential rules, which typically encodes ontological knowledge, and a factbase, which contains factual data. The forward chaining, also known as the *chase* in databases, is a fundamental tool for reasoning on rule-based knowledge bases and a considerable literature has been devoted to its analysis. Its ubiquity in different domains comes from the fact it allows one to compute a

universal model of the knowledge base, i.e., a model that maps by homomorphism to any other model of the knowledge base. This has a major implication in problems like answering queries with ontologies since it follows that a (Boolean) conjunctive query is entailed by a KB if and only if it maps by homomorphism to a universal model.

Several variants of the chase have been defined: *oblivious* or naive chase (e.g. [CGK08]), *skolem* chase [Mar09], *semi-oblivious* chase [Mar09], *restricted* or standard chase [FKMP05], *core* chase [DNR08] (and its variant, the equivalent chase [Roc16]). All these chase variants compute logically equivalent results.¹ Nevertheless, they differ on their ability to detect the redundancies that are possibly caused by the introduction of unknown individuals (often called *nulls*). Note that, since redundancies can only be due to nulls, all chase variants output exactly the same results on rules without existential variables (i.e., Datalog rules, also called range-restricted rules [AHV95]). Then, for rules with existential variables the chase produces iteratively new information until no new rule application is possible. The (re-)applicability of rules is depending on the ability of each chase variant to detect redundancies. Evidently this has a direct impact on the termination. Of course, if a KB has no finite universal model then none of the chase variants will terminate. This is illustrated by Example 1.

Example 1. Take the KB $\mathcal{K} = (F, \mathcal{R})$, where \mathcal{R} contains the rule $R = \forall x (Human(x) \rightarrow \exists y (parentOf(y, x) \wedge Human(y)))$ and $F = \{Human(Alice)\}$. The application of the rule R on the initial factbase F , entails the existence of a new (unknown) individual y_0 (a *null*) generated by the existential variable y in the rule. This yields the factbase $\{Human(Alice), parentOf(y_0, Alice), Human(y_0)\}$, which is logically translated into an existentially closed formula: $\exists y_0 (Human(Alice) \wedge parentOf(y_0, Alice) \wedge Human(y_0))$. Then, R can be applied again by mapping x to y_0 thereby creating a new individual y_1 . It is easy to see that in this case the forward chaining does not halt, as the generation of each new individual enables a novel rule application. This follows from the fact that the universal model of the knowledge base is infinite. \triangle

However, for the case of KBs which have a *finite* universal model, all chase variants can be totally ordered with respect to the inclusion of the sets of factbases on which they halt: $oblivious < semi-oblivious = skolem < restricted < core$. Here, $X_1 < X_2$ means that when X_1 halts on a KB, so does X_2 , and there are KBs for which the reciprocal is false. The oblivious chase is the most redundant kind of the chase as it performs all possible rule applications, without checking for redundancies. The core chase is the less redundant chase as it computes a minimal universal model by reducing every intermediate factbase to its core. In between, we find the semi-oblivious chase (equivalent to the skolem-chase) and the restricted chase. The first one does not consider isomorphic facts that would be generated by consecutive applications of a rule according to the same mapping of its frontier variables (i.e, variables shared by the rule body and head). The second one discards all rule applications that produce “locally redundant” facts. The chase variants are illustrated by Example 2 (for better presentation, universal quantifiers of rules will be omitted in the examples):

¹ In addition, the *parsimonious chase* was introduced in [LMTV12]. However, this chase variant, aimed towards responding at atomic queries, does not compute a universal model of the KB, hence it is outside the family of chase variants studied here.

Example 2. Consider the knowledge bases $\mathcal{K}_1 = (F, \{R_1\})$, $\mathcal{K}_2 = (F, \{R_2\})$, and $\mathcal{K}_3 = (F', \{R_3\})$ built from the facts $F = \{p(a, a)\}$ and $F' = \{\exists w p(a, w)\}$ and the rules $R_1 = p(x, y) \rightarrow \exists z p(x, z)$, $R_2 = p(x, y) \rightarrow \exists z p(y, z)$ and $R_3 = p(x, y) \rightarrow \exists z (p(x, x) \wedge p(y, z))$. Then, the oblivious chase does not halt on \mathcal{K}_1 while the semi-oblivious chase does. Indeed, there are infinitely many different rule applications on the atoms $p(a, z_0), p(a, z_1), \dots$ that can be generated with R_1 ; yet, all rule applications map the frontier variable x to the same constant a , and are therefore filtered by the semi-oblivious chase. In turn, the semi-oblivious chase does not halt on \mathcal{K}_2 while the restricted chase does. Here again, there are infinitely many rule applications on the atoms $p(a, z_0), p(z_0, z_1), \dots$ that can be generated with R_2 ; since each of them maps the frontier variables to new existentials, they are all performed by the semi-oblivious chase. However, all generated atoms are redundant with the initial atom $p(a, a)$ and the restricted chase deems the first (and then all successive) rule applications as redundant. On the other hand, the restricted chase does not halt on \mathcal{K}_3 while the core chase does. In this case, the first rule application yields $\exists w \exists z_0 (p(a, w) \wedge p(a, a) \wedge p(w, z_0))$. This is logically equivalent to $p(a, a)$ i.e., its core, which leads to the core-chase termination at the next step. However, the restricted chase checks only for redundancy of the newly added atoms with respect to the previous factbase, and does not take into account that the addition of new atoms can cause redundancies elsewhere in the factbase (in this example, the previous atom $p(a, w)$ together with the new atom $p(w, z_0)$ are redundant with respect to the new atom $p(a, a)$). So with the restricted chase, R_3 will be always applicable. Finally, note that $p(a, a)$ is a (finite) universal model for all knowledge bases $\mathcal{K}_1, \mathcal{K}_2$, and \mathcal{K}_3 . \triangle

The termination problem, which asks whether for a given set of rules the chase will terminate on any factbase, is undecidable for all chase variants [DNR08, BLM10, GM14]. Following previous work on Datalog, we study the related problem of *boundedness* in a *breadth-first* setting, i.e., the chase performs rule applications that correspond to a certain breadth-first level before any rule application that corresponds to a higher breadth-first level. Then, given a chase variant X , we call a set of rules X -*bounded* if there is k (called the bound) such that, for any factbase, the X -chase stops after at most k breadth-first steps. Of course, since chase variants differ with respect to termination, they also differ with respect to boundedness.

Boundedness ensures several semantic properties. Indeed, if a set of rules is X -bounded with k the bound, then, for any factbase F , the saturation of F at rank k (i.e., the factbase obtained from F after k X -chase breadth-first steps) is a universal model of the KB; the reciprocal also holds true for the core chase. Moreover, boundedness also ensures the UCQ-rewritability property (also called the finite unification set property [BLMS11]): any (Boolean) conjunctive query q can be rewritten using the set of rules \mathcal{R} into a (Boolean) union of conjunctive queries Q such that for any factbase F , q is entailed by (F, \mathcal{R}) if and only if Q is entailed by F . It follows that many interesting static analysis problems such as query containment under existential rules become decidable when a ruleset is bounded. Note that the conjunctive query rewriting procedure can be designed in a such a way that it terminates within k breadth-first steps with k the bound for the core chase [LMU16]. Finally, from a practical viewpoint, the degree of boundedness can be seen as a measure of the recursivity of a ruleset, and most likely,

this is reflected in the actual number of breadth-first steps required by the chase for a given factbase or the query rewriting process for a given query, which is expected to be much smaller than the theoretical bound.

As illustrated by Example 1, the presence of existential variables in the rules can make the universal model of a knowledge base infinite and so the ruleset unbounded, even for the core chase. However, the importance of the boundedness problem has been recognized already for rules without existential variables. Indeed, the problem has been first posed and studied for Datalog, where it has been shown to be undecidable [HKMV95,Mar99]. Example 3 illustrates some cases of bounded and unbounded rulesets in this setting.

Example 3. Consider the rulesets $\mathcal{R}_1 = \{R\}$ and $\mathcal{R}_2 = \{R, R'\}$ where $R = p(x, y) \wedge p(y, z) \rightarrow p(x, z)$ and $R' = p(x, y) \wedge p(u, z) \rightarrow p(x, z)$. The set \mathcal{R}_1 contains a single transitivity rule for the predicate p . This set is clearly unbounded as for any integer k there exists a factbase $F = \{p(a_i, a_{i+1}) \mid 0 \leq i < 2^k\}$ that requires k chase steps. On the other hand, \mathcal{R}_2 also contains a rule that joins individuals on disconnected atoms. In this case, we have that *i*) if R generates some facts then R' generates these same facts as well and *ii*) R' needs to be applied only at the first step, for any F , as it does not produce any new atom at a later step. Therefore, \mathcal{R}_2 is bounded with the bound $k = 1$. Note that since these examples are in Datalog, the specificities of the chase variants do not play any role. \triangle

Finally, the next example illustrates boundedness for non-Datalog rules.

Example 4. Consider the ruleset $\mathcal{R} = \{p(x, y) \rightarrow \exists z(p(y, z) \wedge p(z, y))\}$ and the fact $F = \{p(a, b)\}$. With all variants, the first chase step yields $F_1 = \{p(a, b), p(b, z_0), p(z_0, b)\}$. Then, two new rule applications are possible, which map $p(x, y)$ to $p(b, z_0)$ and $p(z_0, b)$, respectively. The oblivious and semi-oblivious chases will perform these rule applications and go on forever. Hence, the chase on \mathcal{R} is not bounded for these two variants. On the other hand, the restricted chase does terminate. It will not perform any of these rule applications on F_1 . Indeed, the first application would add the facts $\{p(z_0, z_1), p(z_1, z_0)\}$, which can “folded” into F_1 by a homomorphism that maps z_1 to b (while leaving z_0 fixed), and this is similar for the second rule application. We can check that actually the restricted chase will stop on any factbase, and is bounded with $k = 1$. The same holds here for the core chase. \triangle

Despite the relatively negative results on boundedness, knowing that a set of rules is bounded for some chase variant does not help much in practice anyway, if the bound is unknown or even very large. Hence, the goal of this paper is to investigate the k -boundedness problem, which asks, for a given chase variant, whether for any factbase, the chase stops after at most k breadth-first steps.

Our main contribution is to show that k -boundedness is indeed decidable for the oblivious, semi-oblivious and restricted chases. Actually, we obtain a stronger result by exhibiting a property that a chase variant may fulfill, namely consistent heredity, and prove that k -boundedness is decidable as soon as this property is satisfied. We show

that it is the case for all the known chase variants except for the core chase. Hence, the decidability of k -boundedness for the core chase remains an open question.

2 Preliminaries

We consider a first-order setting with constants but no other function symbols. A term is either a constant or a variable. An atom is of the form $r(t_1, \dots, t_n)$ where r is a predicate of arity n and the t_i are terms. Given a set of atoms A , we denote by $\text{vars}(A)$ and $\text{terms}(A)$ the set of its variables and terms. A *factbase* is a set of atoms, logically interpreted as the existentially closed conjunction of these atoms. A *homomorphism* from a set of atoms A to a set of atoms B (notation: $\pi : A \rightarrow B$), is a substitution $\pi : \text{vars}(A) \rightarrow \text{terms}(B)$ such that $\pi(A) \subseteq B$. In this case, we also say that A maps to B (by π). A homomorphism from A to B is an *isomorphism* if its inverse is also a homomorphism. A set of atoms A is a *core* if there is no homomorphism from A to one of its strict subsets. We denote by \models the classical logical consequence and by \equiv the logical equivalence. It is well-known that, given sets of atoms A and B seen as existentially closed conjunctions, there is a homomorphism from A to B if and only if $B \models A$. When A and B are cores, $A \equiv B$ if and only if there is an isomorphism from A to B .

An *existential rule* (or simply rule), denoted by R , is a formula $\forall \bar{x} \forall \bar{y} (B(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} H(\bar{x}, \bar{z}))$ where B and H , called the body and the head of the rule, are conjunctions of atoms, \bar{x} and \bar{y} are sets of universally quantified variables, and \bar{z} is a set of existentially quantified variables. We call *frontier* the variables shared by the body and head of the rule, that is $\text{frontier}(R) = \bar{x}$. In the following we will refer to a rule as a pair of sets of atoms (B, H) by interpreting their common variables as the frontier. A *knowledge base* (KB) $\mathcal{K} = (F, \mathcal{R})$ is a pair where F is a factbase and \mathcal{R} is a set of existential rules. We implicitly assume that all the rules as well as the factbase employ disjoint sets of variables, even if, for convenience, we reuse variable names in examples.

Let F be a factbase and $R = (B, H)$ be an existential rule. We say that R is *applicable* on F via π if there exists a homomorphism π from its body B to F . We call the pair (R, π) a *trigger*. We denote by π^s a *safe* extension of π which maps all existentially quantified variables in H to fresh variables as follows : for each existential variable z we have that $\pi^s(z) = z_{(R, \pi)}$ ². The factbase $F \cup \pi^s(H)$ is called an *immediate derivation* from F through (R, π) . Given a factbase F and a ruleset \mathcal{R} we define a *derivation* from F and \mathcal{R} , denoted by \mathcal{D} , as a (possibly infinite) sequence of triples $D_0 = (\emptyset, \emptyset, F_0), D_1 = (R_1, \pi_1, F_1), D_2 = (R_2, \pi_2, F_2), \dots$ where $F_0 = F$ and every F_i ($i > 0$) is an *immediate derivation* from F_{i-1} through a new trigger (R_i, π_i) , that is, $(R_i, \pi_i) \neq (R_j, \pi_j)$ for all $i \neq j$. The *sequence of rule applications associated* with a derivation is simply the sequence of its triggers $(R_1, \pi_1), (R_2, \pi_2), \dots$. A *subderiva-*

² This fixed way to choose a new fresh variable allows us to always produce the same atoms for a given trigger and that is without loss of generality since each trigger appears at most once on a derivation.

tion of a derivation \mathcal{D} is any derivation \mathcal{D}' whose sequence of rule applications is a subsequence³ of the sequence of rule applications associated with \mathcal{D} .

We will introduce four chase variants, namely oblivious (**O**), semi-oblivious (**SO**), restricted (**R**), equivalent chase (**E**). As explained later, some pairs of chase variants introduced in the literature have similar behavior, in which case we chose to focus on one of the two. All the chase variants are derivations that comply with some condition of applicability of the triggers.

Definition 1. Let \mathcal{D} be a derivation of length n from a factbase F and a ruleset \mathcal{R} , and F_n the factbase obtained after the n rule applications of \mathcal{D} . A trigger (R, π) is called:

1. **O-applicable** on \mathcal{D} if R is applicable on F_n via π .
2. **SO-applicable** on \mathcal{D} if R is applicable on F_n via π and for every trigger (R, π') in the sequence of triggers associated with \mathcal{D} , the restrictions of π and π' to the frontier of R are not equal.
3. **R-applicable** on \mathcal{D} if $R = (B, H)$ is applicable on F_n via π and π cannot be extended to a homomorphism $\pi' : B \cup H \rightarrow F_n$.
4. **E-applicable** on \mathcal{D} if $R = (B, H)$ is applicable on F_n via π and it does not hold that $F_n \equiv F_n \cup \pi^s(H)$. ⊣

Note that for $X \in \{\mathbf{O}, \mathbf{R}, \mathbf{E}\}$, the applicability of the trigger only depends on F_n (hence we can also say the trigger is X -applicable on F_n), while for the **SO**-chase we have to take into account the previous triggers. Note also that the definitions of **O**- and **SO**-trigger applicability allow one to extend a derivation with a rule application that does not add any atom, i.e., $F_{n+1} = F_n$; however, this is not troublesome since no derivation can contain twice the same triggers.

Given a derivation \mathcal{D} , we define the *rank* of an atom as follows: $\text{rank}(A) = 0$ if $A \in F_0$, otherwise let $R = (B, H)$ and (R, π) be the first trigger in the sequence \mathcal{D} such that $A \in \pi^s(H)$, then $\text{rank}(A) = 1 + \max_{A' \in \pi(B)} \{\text{rank}(A')\}$. When we consider a breadth-first chase, the rank of an atom intuitively corresponds to the chase step at which it has been generated. This notion is naturally extended to triggers: $\text{rank}((R, \pi)) = 1 + \max_{A' \in \pi(B)} \{\text{rank}(A')\}$.

The *depth* of a finite derivation is the maximal rank of one of its atoms. Finally, a derivation \mathcal{D} is X -breadth-first (where $X \in \{\mathbf{O}, \mathbf{SO}, \mathbf{R}, \mathbf{E}\}$) if it satisfies the following two properties:

- (1) *rank compatibility*: for all elements D_i and D_j in \mathcal{D} with $i < j$, the rank of the trigger of D_i is smaller or equal to the rank of the trigger of D_j , and
- (2) *rank exhaustiveness*: for every rank k of a trigger in \mathcal{D} , let $D_i = (R_i, \pi_i, F_i)$ be the last element in \mathcal{D} such that $\text{rank}((R_i, \pi_i)) = k$. Then, every trigger which is X -applicable on the subderivation D_1, \dots, D_i is of rank $k + 1$.

Definition 2 (Chase variants). Let F be a factbase and \mathcal{R} be a ruleset. We define four variants of the chase:

³ A sequence S is a subsequence of a sequence S' if S' can be obtained from S by inserting some (or no) elements in S .

1. An **oblivious chase** is any derivation \mathcal{D} from F and \mathcal{R} .
2. A **semi-oblivious chase** is any derivation \mathcal{D} from F and \mathcal{R} such that for every element $D_i = (R_i, \pi_i, F_i)$ of \mathcal{D} , the trigger (R_i, π_i) is **SO**-applicable on the subderivation D_0, D_1, \dots, D_{i-1} of \mathcal{D} .
3. A **restricted chase** is any derivation \mathcal{D} from F and \mathcal{R} such that for every element $D_i = (R_i, \pi_i, F_i)$ of \mathcal{D} , the trigger (R_i, π_i) is **R**-applicable to on the subderivation D_0, D_1, \dots, D_{i-1} of \mathcal{D} .
4. An **equivalent chase** is any **E**-breadth-first derivation \mathcal{D} from F and \mathcal{R} such that for every element $D_i = (R_i, \pi_i, F_i)$ of \mathcal{D} , the trigger (R_i, π_i) is **E**-applicable on the subderivation D_0, D_1, \dots, D_{i-1} of \mathcal{D} . \dashv

We will abbreviate the above chase variants with **O**-chase, **SO**-chase, **R**-chase, and **E**-chase, respectively. Unless otherwise specified, when we use the term X-chase derivation, we will be referring to any of the four chase variants. Furthermore, with *breadth-first* X-chase derivation we will always imply X-breadth-first X-chase derivation.

An X-chase derivation \mathcal{D} from F and \mathcal{R} is *exhaustive* if for all $i \geq 0$, if a trigger (R, π) is X-applicable on the subderivation D_1, \dots, D_i , then there is a $k \geq i$ such that one of the two following holds:

1. $D_k = (R, \pi, F_k)$ or
2. (R, π) is not X-applicable on D_1, \dots, D_k .

Exhaustivity is also known as *fairness*. An X-chase derivation is *terminating* if it is both exhaustive and finite.

It is well-known that for $X \in \{\mathbf{O}, \mathbf{SO}, \mathbf{E}\}$, if there exists a terminating derivation for a given KB, then all exhaustive derivations on this KB are terminating. This does not hold for the restricted chase, because the order in which rules are applied matters, as illustrated by the next example:

Example 5. We assume two rules $R_1 = p(x, y) \rightarrow \exists z p(y, z)$ and $R_2 = p(x, y) \rightarrow p(y, y)$ and $F = \{p(a, b)\}$. Let $\pi = \{x \mapsto a, y \mapsto b\}$. Then (R_1, π) and (R_2, π) are both **R**-applicable. If (R_2, π) is applied first, then the derivation is terminating. However if we apply (R_1, π) first, and (R_2, π) second we produce the factbase $F_2 = \{p(a, b), p(b, z_{(R_1, \pi)}), p(b, b)\}$ and with $\pi' = \{x \mapsto b, y \mapsto z_{(R_1, \pi)}\}$ we have that (R_1, π') as well as (R_2, π') are again both **R**-applicable. Consequently, if we always choose to apply R_1 before R_2 then the corresponding derivation will be infinite. \triangle

Let us now link the four previous chase variants to some other known chase variants. The semi-oblivious and skolem chases, both defined in [Mar09], lead to similar derivations. Briefly, the skolem chase consists of first skolemizing the rules (by replacing existentially quantified variables with skolem functions whose arguments are the frontier variables) then running the oblivious chase. Both chase variants yield isomorphic results, in the sense that they generate exactly the same sets of atoms, up to a bijective renaming of nulls by skolem terms. Therefore, we chose to focus on one of the two, namely the semi-oblivious chase. The core chase [DNR08] and the equivalent chase

[Roc16] have similar behaviors as well. We remind that a core of a set of atoms is one of its minimal equivalent subsets, and that two equivalent sets of atoms have isomorphic cores. The core chase proceeds in a breadth-first manner and, at each step, performs in parallel all rule applications according to the restricted chase criterion, then computes a core of the resulting factbase. Hence, the core chase may remove at some step atoms that were introduced at a former step. After i breadth-first steps, the equivalent chase and the core chase yield logically equivalent factbases, and they terminate on the same inputs. This follows from the facts that computing the core after each rule application or after a sequence of rule applications gives isomomorphic results, and that $F_i \equiv F_{i+1}$ if and only if $\text{core}(F_i)$ is isomorphic to $\text{core}(F_{i+1})$. However, it is sometimes more convenient to handle the equivalent chase from a formal point of view because of its monotonicity (in the sense that within a derivation $F_i \subseteq F_{i+1}$).

We now introduce some notions that will be central for establishing our results on k -boundedness for the different chase variants.

Definition 3 (Restriction of a derivation). Let \mathcal{D} be a derivation from F and \mathcal{R} . For any $G \subseteq F$, the *restriction of \mathcal{D} induced by G* denoted by $\mathcal{D}|_G$, is the maximal derivation from G and \mathcal{R} obtained by a subsequence of the trigger sequence of \mathcal{D} . \dashv

The following example serves to demonstrate how a subset of the initial factbase induces the restriction of a derivation:

Example 6. Take $F = \{p(a, a), p(b, b)\}$, $R = p(x, y) \rightarrow \exists z p(y, z)$ and

$$\mathcal{D} = (\emptyset, \emptyset, F), (R, \pi_1, F_1), (R, \pi_2, F_2), (R, \pi_3, F_3), (R, \pi_4, F_4)$$

with $\pi_1 = \{x/y \mapsto a\}$, $\pi_2 = \{x/y \mapsto b\}$, $\pi_3 = \{x \mapsto a, y \mapsto z_{(R, \pi_1)}\}$, and $\pi_4 = \{x \mapsto z_{(R, \pi_1)}, y \mapsto z_{(R, \pi_3)}\}$.

The derivation \mathcal{D} produces the factbase

$$F_4 = F \cup \{p(a, z_{(R, \pi_1)}), p(b, z_{(R, \pi_2)}), p(z_{(R, \pi_1)}, z_{(R, \pi_3)}), p(z_{(R, \pi_3)}, z_{(R, \pi_4)})\}$$

Then, if $G = \{p(a, a)\}$, we have $\mathcal{D}|_G = (\emptyset, \emptyset, G), (R, \pi_1, G_1), (R, \pi_3, G_2), (R, \pi_4, G_3)$ is the restriction of \mathcal{D} induced by G where

$$G_3 = G \cup \{p(a, z_{(R, \pi_1)}), p(z_{(R, \pi_1)}, z_{(R, \pi_3)}), p(z_{(R, \pi_3)}, z_{(R, \pi_4)})\}$$

\triangle

Definition 4 (Ancestors). Let $D_i = (R_i, \pi_i, F_i)$ be an element of a derivation \mathcal{D} . Then every atom in $\pi_i(B_i)$ is called a *direct ancestor* of every atom in $(F_i \setminus F_{i-1})$. The (indirect) *ancestor* relation between atoms is defined as the transitive closure of the direct ancestor relation. The direct and indirect ancestor relations between atoms are extended to triggers: let $D_j = (R_j, \pi_j, F_j)$ where $j < i$. Then (R_j, π_j) is a *direct ancestor* of (R_i, π_i) if there is an atom in $(F_j \setminus F_{j-1})$ which is a direct ancestor of the atoms in $(F_i \setminus F_{i-1})$. We will denote the ancestors of sets of atoms and triggers as $\text{Anc}(F, \mathcal{D})$ and $\text{Anc}((R, \pi), \mathcal{D})$, respectively. The inverse of the ancestor relation is called the *descendant* relation. \dashv

There is an evident correspondence between the notion of ancestors and the notions of rank and depth. Suppose a ruleset with at most b atoms in the rules' bodies. The following lemma results from the fact that each atom has at most b direct ancestors and the length of a chain of ancestors cannot exceed the depth of a derivation.

Lemma 1 (The ancestor clue). Let \mathcal{D} be an X -chase derivation from F and \mathcal{R} . Then for any atom A of rank k in \mathcal{D} , $|F \cap \text{Anc}(A, \mathcal{D})| \leq b^k$; also for any trigger (R, π) of rank k in \mathcal{D} , $|F \cap \text{Anc}((R, \pi), \mathcal{D})| \leq b^k$.

This lemma will be instrumental for proving our results on k -boundedness as it allows one to characterize the maximal number of atoms that are needed to produce a new atom at a given chase step.

In the next section, we turn our attention to the properties of the derivations that are key to study k -boundedness.

3 Breadth-first Boundedness

As already mentioned, the concept of boundedness was first introduced for Datalog programs. A Datalog program is said to be bounded if the number of breadth-first steps of a bottom-up evaluation of the program is bounded independently from any database (this notion being more precisely called *uniform boundedness* to distinguish it from the notion of *program boundedness* that restricts the set of predicates that may occur in the database) [GMSV93, Abi89, GP94]. Applying this concept to the more general language of existential rules, and parametrizing it by the considered chase variant, X -boundedness can be specified as follows:

Definition 5. Let $X \in \{\mathbf{O}, \mathbf{SO}, \mathbf{R}, \mathbf{E}\}$. A ruleset \mathcal{R} is X -bounded if there is $k \in \mathbb{N}$ such that for every factbase F , every breadth-first X -chase derivation is of depth at most k . \dashv

This definition may seem natural, however it deserves some comments. First note that in Datalog all exhaustive derivations have the same length but not necessarily the same depth, as illustrated by the following example.

Example 7. Let $F = \{p(a)\}$ and $\mathcal{R} = \{R_1, R_2, R_3\}$ where $R_1 = p(x) \rightarrow q(x)$, $R_2 = q(x) \rightarrow r(x)$, $R_3 = p(x) \rightarrow r(x)$. Here are two exhaustive derivations:

$$\mathcal{D}_1 = (\emptyset, \emptyset, F), (R_1, \pi, F_1), (R_2, \pi, F_2), (R_3, \pi, F_2)$$

$$\mathcal{D}_2 = (\emptyset, \emptyset, F), (R_1, \pi, F_1), (R_3, \pi, F_2), (R_2, \pi, F_2)$$

where $\pi = \{x \mapsto a\}$. We can see that both derivations are exhaustive, however the depth of \mathcal{D}_1 is 2 whereas the depth of \mathcal{D}_2 is 1. \triangle

However, among all exhaustive derivations with Datalog rules, the class of breadth-first derivations are of minimal depth. This remains true for the oblivious and semi-oblivious chase derivations with existential rules:

Proposition 1. For each terminating **O**-chase derivation (resp. **SO**-chase derivation) from F and \mathcal{R} there exists a breadth-first terminating **O**-chase derivation (resp. **SO**-chase derivation) from F and \mathcal{R} of smaller or equal depth.

Proof: If \mathcal{D} is a terminating **O**-chase derivation, we can reorder the sequence of triggers associated with \mathcal{D} in such a way as to create a rank compatible **O**-chase derivation \mathcal{D}' (we know that the applicability condition is not affected if we perform some rule applications earlier). Then \mathcal{D}' is also exhaustive since the resulting factbase is the same. Moreover \mathcal{D}' has to be rank exhaustive, since if a trigger is **O**-applicable on a factbase at some step of the derivation, it is always **O**-applicable (unless it has already been applied). So \mathcal{D}' is breadth-first.

Let us now consider **SO**-chase derivations. For convenience in the following proof, given a trigger (R, π) , we slightly modify the definition of the safe extension π^s : for each existential variable z in H (the head of R), we define $\pi^s(z) = z_{f_R(\pi(x_1), \dots, \pi(x_n))}$ where f_R is a fresh symbol assigned to R , and (x_1, \dots, x_n) is a fixed ordering of the frontier variables in R . For brevity, we say that two triggers (R, π) and (R, π') such that π and π' have the same restriction to the frontier of R are “frontier-equal”. With the new definition, two frontier-equal triggers produce exactly the same set of atoms, i.e., $\pi^s(H) = \pi'^s(H)$. Since a **SO**-chase derivation does not have frontier-equal triggers, this modification of the names of fresh variables can be done without loss of generality.

Let \mathcal{D} be a terminating **SO**-chase derivation from a factbase F . We build a derivation \mathcal{D}_{bf} from \mathcal{D} by increasing rank as follows. Let $\mathcal{D}_0 = \mathcal{D} \setminus (\emptyset, \emptyset, F)$, $\mathcal{D}_{bf}^0 = (\emptyset, \emptyset, F)$. Starting from $i = 1$, we iteratively perform the following steps:

- 1) Let T be the set of all triggers (R, π) from \mathcal{D}_{i-1} such that there is a frontier-equal trigger (R, π') applicable on \mathcal{D}_{bf}^{i-1} , and let T' be the set composed of one trigger (R, π') for each (R, π) in T .
- 2) If $T = \emptyset$, $\mathcal{D}_{bf}^i = \mathcal{D}_{bf}^{i-1}$.
- 3) Otherwise, \mathcal{D}_{bf}^i is obtained by extending \mathcal{D}_{bf}^{i-1} with the triples corresponding to the triggers in T' (in any order), and \mathcal{D}_i is obtained from \mathcal{D}_{i-1} by removing the triples corresponding to the triggers in T .

We can easily check that the following conditions are fulfilled at each step of the algorithm: (a) $\mathcal{D}_{bf}^i \cdot \mathcal{D}_i$ is a well-formed derivation (b) there is a bijection between the triggers in \mathcal{D} and those in $\mathcal{D}_{bf}^i \cdot \mathcal{D}_i$, such that corresponding triggers are frontier-equal; (c) the depth of \mathcal{D}_{bf}^i is less or equal to the depth of \mathcal{D} ; (d) \mathcal{D}_{bf}^i is a breadth-first derivation. For Point (a), note that replacing (R, π) by (R, π') has no impact on the name of the obtained fresh variables, hence no impact on triggers that use atoms produced by (R, π) . For Point (d), note that \mathcal{D}_{bf}^i is rank-compatible by construction, and that it is rank-exhaustive: otherwise, there would be a trigger (R, π) still **SO**-applicable on \mathcal{D} , which is not possible since \mathcal{D} is terminating.

The algorithm terminates since the number of steps is upper bounded by the depth of \mathcal{D} . Let $i = d$ be the last step. Then, $\mathcal{D}_{d-1} = \emptyset$, hence, from (b), there is a bijection between the triggers in \mathcal{D} and those in $\mathcal{D}_{bf} = \mathcal{D}_{bf}^{d-1}$, such that corresponding triggers are frontier-equal. It follows that \mathcal{D}_{bf} is terminating. \square

The equivalent chase, which is inspired from the core chase, is breadth-first by definition. The case of the restricted chase is more complex, since, for a given factbase,

some exhaustive derivations may terminate, while others may not. It may happen that all breadth-first derivations terminate (with depth less than a predefined number k), but there is an exhaustive non-breadth-first derivation that does not terminate. It may also be the case that no breadth-first derivation terminates, but there is a non-breadth-first derivation that terminates (with predefined depth less than k), as illustrated by the next example.

Example 8. Let $F = \{p(a, b)\}$ and $\mathcal{R} = \{R_1, R_2, R_3\}$ with $R_1 = p(x, y) \rightarrow \exists z p(y, z)$, $R_2 = p(x, y) \rightarrow \exists z q(y, z)$ and $R_3 = q(y, z) \rightarrow p(y, y)$. It is easy to see that a breadth-first \mathbf{R} -chase derivation in this knowledge base cannot be terminating. However by applying only R_2 on F and then R_3 on the new atom, we obtain a terminating \mathbf{R} -chase derivation. Note also that, for any factbase, there is a terminating \mathbf{R} -chase derivation of depth at most 2. \triangle

Hence, in the case of the restricted chase, breadth-first derivations may not be derivations of minimal depth. More generally, one cannot exclude that other classes of derivations behave better with respect to depth. Moreover, it would be interesting to parametrize boundedness with respect to a specific kind of derivation that would be computed by some restricted chase algorithm. Therefore, a more general definition of boundedness could be based on the maximal depth of a class of derivations of interest. Then, boundedness based on breadth-first settings, as studied in this paper, could be seen as depth-based boundedness applied to breadth-first X -chase variants.

Finally, the following property gives more insight on the relationships between \mathbf{R} -chase derivations and rank-compatible \mathbf{R} -chase derivations (we recall that breadth-first derivations are rank-compatible derivations that are moreover rank-exhaustive).

Proposition 2. For each terminating \mathbf{R} -chase derivation from F and \mathcal{R} there exists a terminating rank-compatible \mathbf{R} -chase derivation from F and \mathcal{R} of smaller or equal depth.

Proof: Let \mathcal{D} be a terminating \mathbf{R} -chase derivation from F and \mathcal{R} . Let $\mathcal{T}_{\mathcal{D}}$ be its sequence of associated triggers and let \mathcal{T} be a sorting of $\mathcal{T}_{\mathcal{D}}$ such that the rank of each element is greater or equal to the rank of its predecessors. Note that \mathcal{T} contains exactly the same triggers as $\mathcal{T}_{\mathcal{D}}$, only the order has changed. Let \mathcal{D}' be the derivation defined by applying, when \mathbf{R} -applicable, the triggers using the order of \mathcal{T} . Because of the reordering, some of the triggers in \mathcal{T} may no longer be \mathbf{R} -applicable in \mathcal{D}' . However, \mathcal{D}' respects the rank compatibility property. We will show that it is a terminating \mathbf{R} -chase derivation. Suppose that there is a new trigger (R, π) (not present in \mathcal{T}) which is \mathbf{R} -applicable on \mathcal{D}' (with $R = (B, H)$). Let \hat{F} be the resulting factbase from \mathcal{D}' . So we can say that (R, π) is \mathbf{R} -applicable on \hat{F} . Let \tilde{F} be the resulting factbase from \mathcal{D} . Then, since $\hat{F} \subseteq \tilde{F}$, we have that (R, π) is \mathbf{O} -applicable on \tilde{F} . But because \mathcal{D} is a terminating \mathbf{R} -chase derivation, we know that (R, π) is not \mathbf{R} -applicable on \tilde{F} . Let $(R_1, \pi_1), \dots, (R_m, \pi_m)$ be the triggers of $\mathcal{T}_{\mathcal{D}}$ that do not appear in \mathcal{D}' (i.e., were not \mathbf{R} -applicable when constructing \mathcal{D}'). So

$$\tilde{F} = \hat{F} \cup \pi_1^s(H_1) \cup \dots \cup \pi_m^s(H_m) \quad (1)$$

where H_1, \dots, H_m are the heads of the rules R_1, \dots, R_m respectively. Since (R, π) is not \mathbf{R} -applicable on \tilde{F} we conclude that there is a homomorphism from $\pi^s(H)$ to \tilde{F} , i.e., a substitution $\sigma : \text{vars}(\pi^s(H)) \rightarrow \text{terms}(\tilde{F})$ such that $\sigma(\pi^s(H)) \subseteq \tilde{F}$, while σ is the identity on $\pi(B)$. Since $(R_1, \pi_1), \dots, (R_m, \pi_m)$ are not \mathbf{R} -applicable in \mathcal{D}' we know that there are substitutions $\sigma_1, \dots, \sigma_m$ such that for every $i \in \{1, \dots, m\}$ we have $\sigma_i : \text{vars}(\pi_i^s(H_i)) \rightarrow \text{terms}(\hat{F})$ and $\sigma_i(\pi_i^s(H_i)) \subseteq \hat{F}$ (i.e., homomorphisms from $\pi_i^s(H_i)$ to \hat{F}), where σ_i is the identity on $\pi_i(B_i)$. Since with $\sigma_1, \dots, \sigma_m$, only new variables are mapped to different terms (and all other variables are mapped to themselves), we can define the substitution $\dot{\sigma} = \bigcup_{i=1}^m \sigma_i$ which has the property that

$$\dot{\sigma}(\hat{F} \cup \pi_1^s(H_1) \cup \dots \cup \pi_m^s(H_m)) = \hat{F} \quad (2)$$

Moreover, the set of variables that are not identically mapped from $\dot{\sigma}$ is disjoint with the variable set $\text{vars}(\hat{F})$, because the new variables created from $(R_1, \pi_1), \dots, (R_m, \pi_m)$ are not present in \hat{F} . Therefore the composition $\dot{\sigma} \circ \sigma$ retains the set of new variables in $\pi^s(H)$ as its set of variables mapped to different terms. So by 1 and $\sigma(\pi^s(H)) \subseteq \tilde{F}$ we can write

$$\dot{\sigma} \circ \sigma(\pi^s(H)) \subseteq \dot{\sigma}(\hat{F} \cup \pi_1^s(H_1) \cup \dots \cup \pi_m^s(H_m))$$

which with 2 becomes

$$\dot{\sigma} \circ \sigma(\pi^s(H)) \subseteq \tilde{F}$$

which implies that (R, π) is not \mathbf{R} -applicable on \mathcal{D}' . That is a contradiction, which leads us to conclude that no such (R, π) exists, therefore \mathcal{D}' is a terminating \mathbf{R} -chase derivation. \square

As already mentioned, boundedness is shown to be undecidable for classes of existential rules like Datalog. However, the practical interest of this notion lies more on whether we can find the particular bound k , rather than knowing that there exists one and thus the ruleset is bounded. Because even if we cannot know whether a ruleset is bounded or not, it can be useful to be able to check a particular bound k . To this aim, we define the notion of k -boundedness where the bound is known, and we prove its decidability for three of the four chase variants.

4 Decidability of k -boundedness for some chase variants

Definition 6 (*k -boundedness*). Given a chase variant X , a ruleset \mathcal{R} is X - k -bounded if for every factbase F , every breadth-first X -chase derivation is terminating with depth at most k . \dashv

Note that a ruleset which is k -bounded is also bounded, but the converse is not true. Our approach for testing k -boundedness is to construct a finite set of factbases whose size depends solely on k and \mathcal{R} , that acts as representative of *all* factbases for the boundedness problem. From this one could obtain the decidability of k -boundedness. Indeed,

for each representative factbase one can compute all breadth-first derivations of depth k and check if they are terminating.

For analogy, it is well-known that the oblivious chase terminates on all factbases if and only if it terminates on the so-called critical instance (i.e., the instance that contains all possible atoms on the constants occurring in rule bodies, with a special constant being chosen if the rule bodies have only variables) [Mar09]. However, it can be easily checked that the critical instance does not provide oblivious chase derivations of maximal depth, hence is not suitable for our purpose of testing k -boundedness. Also, to the best of our knowledge, no representative sets of all factbases are known for the termination of the other chase variants.

In this section, we prove that k -boundedness is decidable for the oblivious, semi-oblivious (skolem) and restricted chase variants by exhibiting such representative factbases. A common property of these three chase variants is that redundancies can be checked “locally” within the scope of a rank, while in the equivalent chase, redundancies may be “global”, in the sense by adding an atom we can suddenly make redundant atoms added by previous ranks.

Following this intuition, we define the notion of hereditary chase.

Definition 7. The X -chase is said to be *hereditary* if, for any X -chase derivation \mathcal{D} from F and \mathcal{R} , the restriction of \mathcal{D} induced by $F' \subseteq F$ is an X -chase derivation. \dashv

A chase is hereditary if by restricting a derivation on a subset of a factbase we still get a derivation with no redundancies. This captures the fact that redundancies can be tested “locally”. This property is fulfilled by the oblivious, semi-oblivious and restricted chase variants; a counter-example for the equivalent chase is given as the end of this section.

Proposition 3. The X -chase is hereditary for $X \in \{\mathbf{O}, \mathbf{SO}, \mathbf{R}\}$.

Proof: We assume that \mathcal{D} is an X -chase derivation from F and \mathcal{R} , and $\mathcal{D}|_{F'}$ is the restriction of \mathcal{D} induced by $F' \subseteq F$.

Case O By definition, an \mathbf{O} -chase derivation is any sequence of immediate derivations with distinct triggers, so the restriction of a derivation from a subfact of F is an \mathbf{O} -chase derivation.

Case SO The condition for \mathbf{SO} -applicability is that we do not have two triggers which map frontier variables in the same way. As \mathcal{D} fulfills this condition its subsequence $\mathcal{D}|_{F'}$ also fulfills it.

Case R The condition for \mathbf{R} -applicability imposes that for a trigger (R, π) there is no extension of π that maps the head of R to F . Since $\mathcal{D}|_{F'}$ generates a factbase included in the factbase generated by \mathcal{D} we conclude that \mathbf{R} -applicability is preserved. \square

Note however that when \mathcal{D} is breadth-first, it is not ensured that its restriction induced by F' is still breadth-first (because the rank exhaustiveness might not be satisfied). It is actually the case for the oblivious chase (since all triggers are always applied), but not for the other variants since some rule applications that would be possible from F' have not been performed in \mathcal{D} because they were redundant in \mathcal{D} given the whole F . The next examples illustrate these cases.

Example 9 (Semi-oblivious chase).

Let $F = \{p(a, b), r(a, c)\}$ and $\mathcal{R} = \{R_1 = p(x, y) \rightarrow r(x, y); R_2 = r(x, y) \rightarrow \exists z q(x, z); R_3 = r(x, y) \rightarrow t(y)\}$. Let \mathcal{D} be the (non terminating) breadth-first derivation of depth 2 from F whose sequence of associated triggers is $(R_1, \pi_1), (R_3, \pi_2), (R_2, \pi_2), (R_3, \pi_1)$ with $\pi_1 = \{x \mapsto a, y \mapsto b\}$ and $\pi_2 = \{x \mapsto a, y \mapsto c\}$ which produces $r(a, b), t(c), q(a, z_{(R_2, \pi_2)}), t(b)$; the trigger (R_2, π_1) is then **O**-applicable but not **SO**-applicable, as it maps equally the frontier variables as (R_2, π_2) . Let $F' = \{p(a, b)\}$. The restriction of \mathcal{D} induced by F' includes only $(R_1, \pi_1), (R_3, \pi_1)$ and is a **SO**-chase derivation of depth 2, however it is not breadth-first since now (R_2, π_1) is **SO**-applicable at rank 2 (thus the rank exhaustiveness is not satisfied). \triangle

Example 10 (Restricted chase).

Let $F = \{p(a, b), q(a, c)\}$ and $\mathcal{R} = \{R_1 = p(x, y) \rightarrow r(x, y); R_2 = r(x, y) \rightarrow \exists z q(x, z); R_3 = r(x, y) \rightarrow t(x)\}$. Let \mathcal{D} be the (terminating) breadth-first derivation of depth 2 from F whose sequence of associated triggers is $(R_1, \pi), (R_3, \pi)$ with $\pi = \{x \mapsto a, y \mapsto b\}$ which produces $\{p(a, b), q(a, c), r(a, b), t(a)\}$; note that the trigger (R_2, π) is **SO**-applicable but not **R**-applicable because of the presence of $q(a, c)$ in F . Let $F' = \{p(a, b)\}$. The restriction of \mathcal{D} induced by F' is a restricted chase derivation of depth 2, however it is not breadth-first since now (R_2, π) is **R**-applicable at rank 2 and thus has to be applied (to ensure the rank exhaustiveness of a breadth-first derivation). \triangle

Previous examples illustrate the need for a more appropriate property focusing on breadth-first derivations. Hence, we define another property, namely *consistent heredity*, which ensures that the restriction of a breadth-first derivation \mathcal{D} induced by F' can be extended to a breadth-first derivation (still from F'). When we consider breadth-first **X**-chases, heredity implies consistent heredity.

Definition 8. The **X**-chase is said to be *consistently hereditary* if for any factbase F and any breadth-first **X**-chase derivation \mathcal{D} from F and \mathcal{R} , the restriction of \mathcal{D} induced by $F' \subseteq F$ is a subderivation of a breadth-first **X**-chase derivation \mathcal{D}' from F' and \mathcal{R} . \dashv

Proposition 4. The **X**-chase is consistently hereditary for $X \in \{\mathbf{O}, \mathbf{SO}, \mathbf{R}\}$.

Proof: Let \mathcal{D} be a breadth-first **X**-chase derivation from F and \mathcal{R} and $\mathcal{D}_{|F'}$ the restriction of \mathcal{D} induced by $F' \subseteq F$.

Case O Since \mathcal{D} is breadth-first, it is rank compatible, and since the ordering of triggers is preserved in $\mathcal{D}_{|F'}$ we get that $\mathcal{D}_{|F'}$ is rank compatible. Similarly by the rank exhaustiveness of \mathcal{D} , all triggers which are descendants of F' appear in \mathcal{D} , so $\mathcal{D}_{|F'}$ is also rank exhaustive. Hence $\mathcal{D}_{|F'}$ is breadth-first.

Case SO As in the **O** case, we can easily see that triggers in $\mathcal{D}_{|F'}$ are ordered by rank. Now, suppose that $\mathcal{D}_{|F'}$ is not rank exhaustive, i.e., there are rule applications (descendants of F') that were *skipped* in \mathcal{D} because they mapped the frontier variables of a rule R in the same way that earlier rule applications (using atoms from $F \setminus F'$) did. Then new triggers will be applicable in $\mathcal{D}_{|F'}$.

Let \mathcal{D}' be a derivation, called the *breadth first completion* of $\mathcal{D}_{|F'}$, constructed as follows: for every breadth-first level κ , after sequentially applying all triggers of $\mathcal{D}_{|F'}$ of

rank κ that are still **SO**-applicable, we complete this rank by applying all other possible **SO**-applicable triggers of rank κ (in any order).

By construction, \mathcal{D}' is a breadth-first **SO**-chase derivation. We will now show that it is actually a completion of $\mathcal{D}_{|F'}$, in the sense that $\mathcal{D}_{|F'}$ is a subderivation of \mathcal{D}' . Indeed, suppose that the addition of a new trigger (R, π) at rank κ in \mathcal{D}' cancels the **SO**-applicability of a trigger (R, π') at rank $\kappa' > \kappa$ in $\mathcal{D}_{|F'}$. So (R, π) is “frontier-equal” with (R, π') . Then, since (R, π) is not in \mathcal{D} , and \mathcal{D} is rank-exhaustive, there is a “frontier-equal” trigger (R, π_D) in \mathcal{D} at rank $\kappa_D \leq \kappa$; this is not possible since (R, π_D) would also be frontier-equal to (R, π') , which would both belong to \mathcal{D} , which contradicts the fact that \mathcal{D} is a **SO**-chase derivation.

Case R Let \mathcal{D}' be the *breadth first completion* of $\mathcal{D}_{|F'}$ constructed similarly as in the previous case: for every breadth-first level κ , after sequentially applying all triggers of $\mathcal{D}_{|F'}$ of rank κ that are still **R**-applicable, we complete this rank by applying all other possible **R**-applicable triggers of rank κ (in any order). By construction, \mathcal{D}' is a breadth-first **R**-chase derivation.

We will also show that $\mathcal{D}_{|F'}$ is a subderivation of \mathcal{D}' . We do so by contradiction. Let (R, π) be the first trigger of $\mathcal{D}_{|F'}$ that does not appear in \mathcal{D}' .

We denote by \hat{F}' the resulting factbase after applying all the triggers that precede (R, π) in $\mathcal{D}_{|F'}$ and by G the resulting factbase after applying all triggers of \mathcal{D}' up to (R, π) (excluding (R, π)). Let $(R_1, \pi_1), \dots, (R_m, \pi_m)$ be the triggers that were not **R**-applicable in \mathcal{D} but were **R**-applicable in \mathcal{D}' and added before (R, π) . It holds that $G = \hat{F}' \cup \pi_1^s(H_1) \cup \dots \cup \pi_m^s(H_m)$.

Now, we have assumed that (R, π) is not **R**-applicable on G , hence not present in \mathcal{D}' . So, by the condition of **R**-applicability, there exists a homomorphism $\sigma : \pi^s(H) \rightarrow G$ (so also $\sigma(\pi^s(H)) \subseteq G$), which behaves as the identity on $\pi(B)$. We denote with F_i the factbase produced just before applying (R, π) on \mathcal{D} . We have that $\hat{F}' \subseteq F_i$, hence we get that $G \subseteq F_i \cup \pi_1^s(H_1) \cup \dots \cup \pi_m^s(H_m)$ and therefore we also have

$$\sigma(\pi^s(H)) \subseteq F_i \cup \pi_1^s(H_1) \cup \dots \cup \pi_m^s(H_m) \quad (3)$$

Now, because $(R_1, \pi_1), \dots, (R_m, \pi_m)$ were not **R**-applicable in \mathcal{D} we know that there exist respective homomorphisms $\sigma_j : \pi_j^s(H_j) \rightarrow F_i$ (so also $\sigma_j(\pi_j^s(H_j)) \subseteq F_i$), that behave as the identity on $\pi_j(B_j)$, for all $j \in \{1, \dots, m\}$. As the domains of all σ_j restricted to existential variables are disjoint, and σ_j are the identity on non-existential

variables, we can define the substitution $\dot{\sigma} := \bigcup_{i=1}^m \sigma_i$. By applying $\dot{\sigma}$ to both sides of (3) we get

$$\dot{\sigma} \circ \sigma(\pi^s(H)) \subseteq \dot{\sigma}(F_i \cup \pi_1^s(H_1) \cup \dots \cup \pi_m^s(H_m)) \quad (4)$$

which, considering that $\dot{\sigma}(F_i \cup \pi_1^s(H_1) \cup \dots \cup \pi_m^s(H_m)) \subseteq F_i$, yields

$$\dot{\sigma} \circ \sigma(\pi^s(H)) \subseteq F_i \quad (5)$$

The homomorphism $\dot{\sigma} \circ \sigma$ can only substitute the set of newly created variables in $\pi^s(H)$, hence qualifies as an extension of π , and from (5) we conclude that (R, π) is

not \mathbf{R} -applicable in \mathcal{D} . That is a contradiction, hence it must be the case that (R, π) is indeed \mathbf{R} -applicable in \mathcal{D}' . Therefore we have shown that all triggers of $\mathcal{D}_{|F'}$ appear in \mathcal{D}' , so indeed $\mathcal{D}_{|F'}$ is a subderivation of a breadth-first \mathbf{R} -chase derivation from F' . \square

The next property exploits the notion of consistent heredity to bound the size of the factbases that have to be considered.

Proposition 5. Let b be the maximum number of atoms in the bodies of the rules of a ruleset \mathcal{R} . Let X be any consistently hereditary chase. If there exist an F and a breadth-first X -chase \mathcal{R} -derivation from F that is of depth at least k , then there exist an F' of size $|F'| \leq b^k$ and a breadth-first X -chase \mathcal{R} -derivation from F' with depth at least k .

Proof: Let \mathcal{D} be a breadth-first X -chase derivation from F and \mathcal{R} of depth k . Let (R, π) be a trigger of \mathcal{D} of depth k . Let F' be the set of ancestors of (R, π) in F , and by Lemma 1 we know that $|F'| \leq b^k$. Since the X -chase is consistently hereditary, the restriction $\mathcal{D}_{|F'}$ (which trivially includes (R, π)) is a subderivation of a breadth-first X -chase derivation \mathcal{D}' from F' and \mathcal{R} . According to the proof of proposition 4, for all three consistently hereditary chase variants, \mathcal{D}' was constructed as a breadth-first completion of \mathcal{D} , therefore the ranks of common triggers are preserved from \mathcal{D} to $\mathcal{D}_{|F'}$ and \mathcal{D}' . And since \mathcal{D}' includes (R, π) in its sequence of associated rule applications, we have that (R, π) has also rank k in \mathcal{D}' , hence \mathcal{D}' is of depth at least k . \square

We are now ready to prove the main result.

Theorem 1. Determining if a set of rules is X - k -bounded is decidable for any consistently hereditary chase variant X . This is in particular the case for the oblivious, semi-oblivious and restricted chase variants.

Proof: By Proposition 5, to check if all breadth-first X -chase derivations from \mathcal{R} (with any factbase) are of depth at most k , it suffices to verify this property on all factbases of size less or equal to b^k . For a given factbase F , there is a finite number of (breadth-first) X -chase derivations from F and \mathcal{R} of depth at most k , hence we can effectively compute these derivations, and check if one of them can be extended to a derivation of depth $k + 1$. \square

Finally, the following example shows that the \mathbf{E} -chase (hence the core chase as well) is not consistently hereditary (hence not hereditary, as it the \mathbf{E} -chase is breadth-first).


Example 11 (Equivalent chase). Let $F = \{s(b), p(a, a), p(a, b), p(b, c)\}$ and \mathcal{R} the following set of rules:

$$\begin{aligned} R_1 &= s(y) \wedge p(y, z) \wedge p(w, z) \wedge r(w) \rightarrow q(w) \\ R_2 &= p(x, y) \wedge p(y, z) \rightarrow t(y) \\ R_3 &= p(x, x) \wedge p(x, y) \wedge p(y, z) \rightarrow \exists w(p(w, z) \wedge r(w)) \\ R_4 &= t(y) \rightarrow r(y) \\ R_5 &= p(x, y) \rightarrow \exists u p(u, x) \end{aligned}$$

Here we can verify that any exhaustive \mathbf{E} -chase derivation from F and \mathcal{R} is of depth 3. Consider such a derivation \mathcal{D} that adds atoms in the following specific order at each breadth-first level (for clarity, we do not use standardized names for the nulls):


- 0 : $s(b), p(a, a), p(a, b), p(b, c)$
- 1 : $t(a), t(b), p(w_1, c), r(w_1), p(w_2, b), r(w_2), p(w_3, a), r(w_3)$
- 2 : $q(w_1), r(a), r(b), p(u_1, w_1)$
- 3 : $q(b)$

Below is a graphical representation of this derivation, where nodes are atoms and edges are colored according to different triggers:



At step 1, R_2 is applied twice, producing $t(a)$ and $t(b)$, and R_3 is applied three times, producing $p(w_1, c), r(w_1), p(w_2, b), r(w_2), p(w_3, a)$ and $r(w_3)$. Note that R_1 and R_4 are not applicable, and R_5 is not **E**-applicable because it would produce redundant atoms. At step 2, R_1 is applied once (producing $q(w_1)$), R_2 and R_3 are not **E**-applicable, R_4 is applied twice, and R_5 is applied once (producing $p(u_1, w_1)$). Finally, at step 3, R_1 is applied, which makes all further triggers redundant, hence no other rule is **E**-applicable.

Let $F' = F \setminus \{s(b)\}$. Let $\mathcal{D}_{F'}$ be the restriction of \mathcal{D} induced by F' . Here is a graphical representation of $\mathcal{D}_{F'}$:



At level 2, $\mathcal{D}_{F'}$ still produces $r(a)$, $r(b)$ and $p(u_1, w_1)$ but not $q(w_1)$, and there is no step 3 because R_1 is not applicable. We can see that $\mathcal{D}_{F'}$ is not an **E**-chase derivation because the application of R_5 at step 2 (which produces $p(u_1, w_1)$) is now redundant (this is due to the absence of $q(w_1)$). This already shows that the **E**-chase is not hereditary. Moreover, we can check on $\mathcal{D}_{F'}$ that no rule application before the application of R_5 is able to add information on w_1 that would make R_5 **E**-applicable at step 2. Hence, $\mathcal{D}_{F'}$ is not contained in any **E**-chase derivation from F' , which shows that the **E**-chase is not consistently hereditary. Note also that any exhaustive **E**-chase derivation from F' is of depth 2 and not 3 as from F . \triangle

5 Conclusion

In this paper, we investigated the problem of determining whether a ruleset is k -bounded, that is when the chase always halts within a predefined number of steps independently of the factbase. After discussing the concept of boundedness in breadth-first derivations, we have shown that k -boundedness is decidable for some important chase variants by establishing a common property that ensures decidability, namely “consistent heredity”. The complexity of the problem is independent from any data since the size of the factbases to be checked depends only on k and the size of the rule bodies. Our results indicate an EXPTIME upper bound for checking k -boundedness for both the **O**-chase and the **SO**-chase. For the **R**-chase, as the order of the rule applications matters, one needs to check all possible derivations. This leads to a 2-EXPTIME upper bound for the **R**-chase. We leave for further work the study of the precise lower complexity bound according to each kind of chase. Finally, we leave open the question of the decidability of the k -boundedness for the core (or equivalent) chase.

References

- .Abi89 Serge Abiteboul. Boundedness is undecidable for datalog programs with a single recursive rule. *Inf. Process. Lett.*, 32(6):281–287, 1989.
- .AHV95 Serge Abiteboul, Richard Hull, and Victor Vianu, editors. *Foundations of Databases: The Logical Level*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1st edition, 1995.
- .BLM10 Jean-François Baget, Michel Leclère, and Marie-Laure Mugnier. Walking the decidability line for rules with existential variables. In *KR 2010*, 2010.
- .BLMS09 Jean-François Baget, Michel Leclère, Marie-Laure Mugnier, and Eric Salvat. Extending decidable cases for rules with existential variables. In *IJCAI 2009*, pages 677–682, 2009.
- .BLMS11 J.-F. Baget, M. Leclère, M.-L. Mugnier, and E. Salvat. On Rules with Existential Variables: Walking the Decidability Line. *Artificial Intelligence*, 175(9-10):1620–1654, 2011.
- .CGK08 A. Cali, G. Gottlob, and M. Kifer. Taming the infinite chase: Query answering under expressive relational constraints. In *KR’08*, pages 70–80, 2008.
- .CGL⁺05 D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. DL-Lite: Tractable description logics for ontologies. In *AAAI*, pages 602–607, 2005.
- .CGL09 Andrea Cali, Georg Gottlob, and Thomas Lukasiewicz. A general datalog-based framework for tractable query answering over ontologies. In *PODS 2009*, pages 77–86, 2009.
- .DNR08 A. Deutsch, A. Nash, and J. B. Remmel. The chase revisited. In *PODS*, pages 149–158, 2008.
- .FKMP05 R. Fagin, P. G. Kolaitis, R. J. Miller, and L. Popa. Data Exchange: Semantics and Query Answering. *Theor. Comput. Sci.*, 336(1):89–124, 2005.
- .GM14 Tomasz Gogacz and Jerzy Marcinkowski. All-instances termination of chase is undecidable. In *ICALP 2014 Proceedings, Part II*, pages 293–304, 2014.
- .GMSV93 Haim Gaifman, Harry G. Mairson, Yehoshua Sagiv, and Moshe Y. Vardi. Undecidable optimization problems for database logic programs. *J. ACM*, 40(3):683–713, 1993.
- .GOPS12 Georg Gottlob, Giorgio Orsi, Andreas Pieris, and Mantas Simkus. Datalog and its extensions for semantic web databases. In *Reasoning Web*, pages 54–77, 2012.
- .GP94 Irène Guessarian and Marcos Veloso Peixoto. About boundedness for some datalog and datalog^{neg} programs. *J. Log. Comput.*, 4(4):375–403, 1994.
- .HKMV95 Gerd G. Hillebrand, Paris C. Kanellakis, Harry G. Mairson, and Moshe Y. Vardi. Undecidable boundedness problems for datalog programs. *J. Log. Program.*, 25(2):163–190, 1995.
- .KRH07 M. Krötzsch, S. Rudolph, and P. Hitzler. Complexity boundaries for Horn description logics. In *Proc. of AAAI*, pages 452–457. AAAI Press, 2007.
- .LMTV12 Nicola Leone, Marco Manna, Giorgio Terracina, and Pierfrancesco Veltri. Efficiently computable Datalog \exists programs. In *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning*, KR’12, pages 13–23. AAAI Press, 2012.
- .LMU16 Michel Leclère, Marie-Laure Mugnier, and Federico Ulliana. On bounded positive existential rules. In *Proceedings of the 29th International Workshop on Description Logics*, 2016.
- .LTW09 C. Lutz, D. Toman, and F. Wolter. Conjunctive query answering in the description logic \mathcal{EL} using a relational database system. In *Proc. of IJCAI*, pages 2070–2075, 2009.

- .Mar99 Jerzy Marcinkowski. Achilles, turtle, and undecidable boundedness problems for small datalog programs. *SIAM J. Comput.*, 29(1):231–257, September 1999.
- .Mar09 B. Marnette. Generalized schema-mappings: from termination to tractability. In *PODS*, pages 13–22, 2009.
- .MT14 Marie-Laure Mugnier and Michaël Thomazo. An introduction to ontology-based query answering with existential rules. In *Reasoning Web 2014*, pages 245–278, 2014.
- .Roc16 Swan Rocher. *Querying Existential Rule Knowledge Bases: Decidability and Complexity. (Interrogation de Bases de Connaissances avec Règles Existentielles : Décidabilité et Complexité)*. PhD thesis, University of Montpellier, France, 2016.