



HAL
open science

A Journey to the Frontiers of Query Rewritability

Piotr Ostropolski-Nalewaja, Jerzy Marcinkowski, David Carral, Sebastian Rudolph

► **To cite this version:**

Piotr Ostropolski-Nalewaja, Jerzy Marcinkowski, David Carral, Sebastian Rudolph. A Journey to the Frontiers of Query Rewritability. PODS 2022 - 41st ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, Jun 2022, Philadelphia, United States. pp.359-367, 10.1145/3517804.3524163 . hal-03599601

HAL Id: hal-03599601

<https://hal.science/hal-03599601>

Submitted on 14 Mar 2022

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.

A Journey to the Frontiers of Query Rewritability (Extended Technical Report)

Piotr Ostropolski-Nalewaja
University of Wrocław, Institute of Computer Science

David Carral
LIRMM, Inria, University of Montpellier, CNRS

Jerzy Marcinkowski
University of Wrocław, Institute of Computer Science

Sebastian Rudolph
TU Dresden, Computational Logic Group

ABSTRACT

We consider (first-order) query rewritability in the context of theory-mediated query answering. The starting point of our journey is the FUS/FES conjecture, which states that any theory that is a finite expansion set (FES) and admits query rewriting (BDD, FUS) must be uniformly bounded. We show that this conjecture holds for a large class of BDD theories, which we call “local”. Upon investigating how “non-local” BDD theories can actually get, we discover unexpected phenomena that, we think, are at odds with prevailing intuitions about BDD theories.

1 INTRODUCTION

The scenario we consider in this paper has been studied extensively both in database theory and in knowledge representation: consider a database *instance* \mathbb{D} (also called *fact set*, or *structure*) and a *theory* \mathcal{T} (or *rule set*) that consists of *tuple generating dependencies* (or *rules*). For a given conjunctive query ϕ , we ask if \mathbb{D} and \mathcal{T} together logically entail ϕ , written: $\mathbb{D}, \mathcal{T} \models \phi$. This problem is also referred to as *ontology-mediated query answering*.

The chase. The notion of the *chase* is fundamental in this context. It denotes a structure obtained from \mathbb{D} via the *chase procedure*, which iteratively adds new terms and atoms in order to satisfy the constraints from \mathcal{T} , producing a growing sequence of structures $\mathbb{D} = Ch_0(\mathbb{D}, \mathcal{T}), Ch_1(\mathbb{D}, \mathcal{T}), Ch_2(\mathbb{D}, \mathcal{T}), \dots$. The chase $Ch(\mathbb{D}, \mathcal{T})$ is then obtained as the corresponding fixpoint $\bigcup_{i \in \mathbb{N}} Ch_i(\mathbb{D}, \mathcal{T})$.

It is known [1] that a conjunctive query is entailed if and only if it holds in the corresponding chase, which allows for reducing entailment to model checking:

$$\forall \mathcal{T} \forall \mathbb{D} \forall \phi \quad (Ch(\mathbb{D}, \mathcal{T}) \models \phi \Leftrightarrow \mathbb{D}, \mathcal{T} \models \phi).$$

Finite expansion sets. We say \mathcal{T} enjoys the *finite expansion set* property [2] (or simply *is FES*) if, for every \mathbb{D} , all conjunctive queries satisfied in $Ch(\mathbb{D}, \mathcal{T})$ are already jointly satisfied after finitely many chase steps. More precisely, \mathcal{T} is FES if:

$$\forall \mathbb{D} \exists i \in \mathbb{N} \forall \phi \quad (Ch(\mathbb{D}, \mathcal{T}) \models \phi \Leftrightarrow Ch_i(\mathbb{D}, \mathcal{T}) \models \phi) \quad (\text{FES})$$

This is an important property, since $Ch(\mathbb{D}, \mathcal{T})$ is typically an infinite structure, only existing as an abstract mathematical object, and impossible to query, whereas $Ch_i(\mathbb{D}, \mathcal{T})$ is always finite and so in principle it can be computed and queried.

Bounded derivation depth. An arguably even more beneficial property a theory can enjoy in this context is the following: We say \mathcal{T} has the *bounded derivation depth* property (or *is BDD*) if

$$\forall \phi \exists i \in \mathbb{N} \forall \mathbb{D} \quad (Ch(\mathbb{D}, \mathcal{T}) \models \phi \Leftrightarrow Ch_i(\mathbb{D}, \mathcal{T}) \models \phi), \quad (\text{BDD})$$

which means that, in order to evaluate ϕ , it is enough to run only the first i steps of the chase, with i depending on ϕ but not on \mathbb{D} .

As it turns out, BDD is equivalent [3] to FUS (*finite unification set*) [2]. FUS is the class ensuring that conjunctive queries always rewrite: for each ϕ one can compute a query $\phi_{\mathcal{T}}$, being a union of conjunctive queries, such that for each \mathbb{D} we have that $Ch(\mathbb{D}, \mathcal{T}) \models \phi$ exactly if $\mathbb{D} \models \phi_{\mathcal{T}}$ — this is known to be equivalent to the existence of an arbitrary first-order rewriting [4].

Behold the extreme usefulness of this property: instead of querying $Ch(\mathbb{D}, \mathcal{T})$, an elusive infinite structure, we can equivalently query \mathbb{D} , the only structure we have immediate access to.

No wonder the BDD/FUS property has been considered in literally hundreds of papers. Numerous classes of BDD theories have been identified and intensively studied, among them subclasses with decidable membership like:

- linear theories, where rules have at most one body atom;
- guarded BDD theories (while not all guarded theories are BDD, it is decidable to determine if a guarded theory is BDD [5, 6]), generalizing linear theories;
- sticky theories, defined by a reasonably natural syntactic restriction on the use of joins [7].

Apart from the decidable subclasses of BDD, there are also natural undecidable subclasses:

- bounded Datalog theories, already studied decades before the class BDD itself was discovered [8];
- binary BDD theories, where the arity of relation symbols is at most 2 (also studied in the context of description logics [9]);
- backward shy theories [10], a superclass of sticky theories.

As we expose in this paper, despite this extensive body of work, we still understand very little about the deeper mathematical properties of BDD theories. In particular, we are going to show that the intuition of BDD theories as being “local”, “only depending on the small pieces of \mathbb{D} ”, and “unable to look too far” while (more or less) correct for all aforementioned classes of BDD theories is blatantly incorrect for BDD theories in general.

The FUS/FES conjecture. There is a striking similarity between formulas (BDD) and (FES), inspiring a natural conjecture, which we call the *FUS/FES conjecture*: For any BDD theory that is also FES, it should be possible to choose the number i in a **uniform way**, independently from both \mathbb{D} and ϕ . In other words, the conjecture says that if \mathcal{T} is both BDD and FES then:

$$\exists i \in \mathbb{N} \forall \phi \forall \mathbb{D} \quad (Ch(\mathbb{D}, \mathcal{T}) \models \phi \Leftrightarrow Ch_i(\mathbb{D}, \mathcal{T}) \models \phi) \quad (\text{UBDD})$$

This conjecture was studied earlier [11] and a proof was proposed, which however turned out to be incorrect and was later withdrawn.

Later, it was shown that the conjecture would hold if the assumption that \mathcal{T} is FES were replaced by a stronger property [12].

Main results. We identify the new, generic class of *local* BDD theories. Informally, such theories enforce that the creation of every chase atom depends only on a constant number of facts from the database. Importantly, this new class not only includes most of the aforementioned subclasses of BDD (with sticky and backward shy being notable exceptions) but also all BDD theories over signatures with a maximum predicate arity of 2.

With this new class of theories defined (and the above inclusions proved), we present our three main results: **First** that the FUS/FES conjecture holds for local theories (Theorem 2). **Second** that the conjecture holds for every BDD theory over a binary signature (Corollary 1) – an immediate consequence of the fact that such theories are local (Theorem 1).

Yet, as least as interesting as these insights we find our **third** result (reflected in Theorem 3): the discovery of very much non-local (and not even what we call *bounded-degree local*) theories that are still BDD. Such theories not only defy many of the popular intuitions about the BDD class, but also shows that all previous investigations into that class have probably barely scratched its surface and that there is a lot of room for new decidable/syntactic classes of BDD theories, richer than all hitherto considered. One plausible reason why this new world exhibits counterintuitive phenomena and has gone entirely unnoticed is that it requires theories of arity higher than 2. Binary theories are much easier to imagine and they are mainly responsible for shaping our intuitions.

Summarizing, the **main message of this paper** is: even if a (finite) counterexample to the FUS/FES conjecture should exist, it is going to be found nowhere near the familiar avenues of the BDD class. But we also show that the known avenues only reach a small part of the BDD class and there is a lot of uncharted territory left.

Organisation of this paper. Apart from its preliminary sections (Sections 2–7), the paper is organized as follows:

In Section 8 we define *local* theories. Later we state our first result Theorem 1 and present a brief insight into its proof.

In Section 9 we state our second result – that the FUS/FES conjecture holds for local theories (Theorem 2) – and prove it. Also we note that the conjecture holds for theories over binary signatures (Corollary 1).

In Section 10, we notice that sticky theories, while BDD, are not always local. We define another, weaker, notion, of *bounded-degree-local* theories (or *bd-local*) and note that it covers sticky. Later we discuss properties of this newly defined class.

In Section 11 we examine the intuition that “BDD theories are unable to look too far”. We define the notion of *distancing* theories, and show that if a theory is local then it is also distancing. We also notice that backward shy theories are distancing, so that most of the previously known examples of BDD theories are indeed distancing. We show however, that there exists a BDD theory \mathcal{T}_d that is not distancing. As a corollary we get that, for this BDD theory, the rewriting $\phi_{\mathcal{T}_d}$ (which is a disjunction of conjunctive queries) of a query ϕ can require disjuncts of exponential size with respect to the size of ϕ . This is in stark contrast to any BDD theories previously considered.

Finally in Section 12, we conclude and discuss future work.

This is an extended technical report of an article with the same title published at PODS 2022. It includes the complete formal proofs for all results in the appendix.

2 PRELIMINARIES

Queries and TGDs. A *conjunctive query* (CQ) is a formula $\psi(\bar{y}) = \exists \bar{x} \beta(\bar{x}, \bar{y})$ with β being a non-empty conjunction of atomic formulas over some *signature* (or *schema*) Σ (which is a finite set of relation symbols) and over some set of variables and set of constants. So, for example $\exists x \text{Siblings}(\text{Abel}, x), \text{Female}(x)$ is a CQ.

A *boolean CQ* (short: BCQ) is a CQ with all variables quantified (as in the preceding example). We refer to β as the *body* of $\psi(\bar{y})$. By a *union of conjunctive queries* (UCQ) we mean a formula being a disjunction of CQs. By the *size* of a CQ, denoted $|\psi(\bar{y})|$, we mean the number of atomic formulas it is built of. By the *width* of a UCQ we will mean the size of its greatest disjunct. Also, we will sometimes treat UCQs as sets of CQs.

A *theory* or a *rule set* is a finite set of *tuple generating dependencies* (TGDs, often just referred to as *rules*). A *rule* is a first-order logic formula of the form $\forall \bar{x}, \bar{y} (\beta(\bar{x}, \bar{y}) \Rightarrow \exists \bar{w} \alpha(\bar{y}, \bar{w}))$, where \bar{x}, \bar{y} and \bar{w} are pairwise disjoint lists of variables, $\beta(\bar{x}, \bar{y})$ (the rule’s *body*) is a conjunction of atomic formulas and $\alpha(\bar{y}, \bar{w})$ (the rule’s *head*) is an atomic formula. The *frontier* \bar{y} of a rule ρ , denoted $\text{fr}(\rho)$, is the set of all variables that occur both in the body and the head of the rule. We omit universal quantifiers when writing rules and treat conjunctions of atoms (such as β) as sets of atoms.

Note that, in database theory terminology, our rules are “single head” TGDs. This is the only reasonable choice in this context, since we want to talk about theories over a binary signature: if we allowed rule heads to comprise several atoms, then rules with predicates of any arity could be easily simulated using only arity 2 predicates. On the other hand, this does not restrict our results in the cases where we do not assume that the signature is binary, since every multi-head theory can be rewritten into a single-head one, using higher-arity auxiliary predicates, and this rewriting does not affect the property of being FUS or FES.

Structures and entailment. A database *instance* (or *structure* or *fact set*) is a set of facts – atomic formulas over Σ . For a structure \mathbb{F} over Σ we let $\text{dom}(\mathbb{F})$ denote its *active domain* – the set of all terms that appear in the facts of \mathbb{F} . For $c, c' \in \text{dom}(\mathbb{F})$, we let $\text{dist}_{\mathbb{F}}(c, c')$ denote the distance between c and c' in the Gaifman graph of \mathbb{F} : the vertices of this graph are elements of $\text{dom}(\mathbb{F})$ and two vertices are connected by an edge if and only if they appear in the same fact. We define the *degree* of \mathbb{F} as the degree of \mathbb{F} ’s Gaifman graph.

\mathbb{F} is a model of \mathcal{T} (written: $\mathbb{F} \models \mathcal{T}$) if \mathbb{F} satisfies all TGDs from \mathcal{T} . For a pair \mathcal{T}, \mathbb{D} , a CQ $\phi(\bar{y})$, and a tuple $\bar{a} \in \text{dom}(\mathbb{D})^{|\bar{y}|}$, we write $\mathcal{T}, \mathbb{D} \models \phi(\bar{a})$ (or $\mathcal{T}, \mathbb{D}, \bar{a} \models \phi(\bar{y})$) to indicate that \mathcal{T} and \mathbb{D} jointly entail $\phi(\bar{a})$, which means that $\phi(\bar{a})$ holds in each structure \mathbb{F} satisfying $\mathbb{F} \models \mathcal{T}, \mathbb{D}$, which serves as a shortcut for $\mathbb{F} \models \mathcal{T} \wedge \mathbb{D} \subseteq \mathbb{F}$.

EXAMPLE 1. Consider the instance $\mathbb{D}_a = \{\text{Human}(\text{Abel})\}$ and the theory \mathcal{T}_a consisting of the following two rules:

$$\text{Human}(y) \Rightarrow \exists z \text{Mother}(y, z)$$

$$\text{Mother}(x, y) \Rightarrow \text{Human}(y)$$

Then $\mathcal{T}_a, \mathbb{D}_a \models \exists y, z \text{Mother}(\text{Abel}, y), \text{Mother}(y, z)$.

Homomorphisms and query containment. For structures \mathbb{D}, \mathbb{F} , a *homomorphism* from \mathbb{D} to \mathbb{F} is a function $h : \text{dom}(\mathbb{D}) \rightarrow \text{dom}(\mathbb{F})$

such that $A(h(\vec{x})) \in \mathbb{F}$ for each fact $A(\vec{x}) \in \mathbb{D}$ and that $h(a) = a$ for any constant $a \in \text{dom}(\mathbb{D})$. Given a fact $\alpha = A(\vec{x}) \in \mathbb{D}$, we will use $h(\alpha)$ to denote $A(h(\vec{x}))$.

For two CQs $\phi(\vec{y})$ and $\psi(\vec{y})$, with the same set of free variables, we say that $\phi(\vec{y})$ *contains* $\psi(\vec{y})$ if for every structure \mathbb{D} and for every tuple $\vec{a} \in \text{dom}(\mathbb{D})^{|\vec{y}|}$ if $\mathbb{D} \models \psi(\vec{a})$ then also $\mathbb{D} \models \phi(\vec{a})$. It is well known that $\phi(\vec{y})$ contains $\psi(\vec{y})$ if and only if there is a homomorphism¹ from $\phi(\vec{y})$ to $\psi(\vec{y})$ that is the identity on \vec{y} .

Core of a structure. A substructure \mathbb{H} of a finite structure \mathbb{G} is a *core* of \mathbb{G} (see [13]) if there exists a homomorphism $h : \mathbb{G} \rightarrow \mathbb{H}$ but there is no homomorphism from \mathbb{G} to \mathbb{H}' where \mathbb{H}' is a proper substructure of \mathbb{H} . Note that the definition of homomorphisms ensures that $h(a) = a$ for every constant $a \in \text{dom}(\mathbb{G})$.

It is well known that [13]:

- (1) Every finite structure has a core.
- (2) Cores of a finite structure are unique up to isomorphism.
- (3) If \mathbb{H} is a core (of any structure) then it is a core of itself.

Given a (finite) structure \mathbb{G} , we let $\text{Core}(\mathbb{G})$ denote a function that returns some induced substructure of \mathbb{G} that is a core.

Connected queries, rules and theories. For a CQ, one can define its Gaifman graph in the natural way: Variables are the vertices of this graph and two variables are connected by an edge if and only if they both appear in the same atomic formula. A conjunctive query is *connected* if its Gaifman graph is connected. A TGD is *connected* if its body is. A theory is *connected* if each of its rules is.

All the theories we consider in this paper are connected with the important exception of theories over a binary signature. Forcing theories to be connected will help us to better express the nuances of the BDD class in Sections 10–11. This assumption does not reduce the expressive power of such theories due to the following trivial trick: add a fresh variable as an additional, first variable in all the atoms appearing in the rules of the theory. This will make the theory connected, and it will obviously preserve its BDD and FES status. But it will increase the arity – so if we care about the arity we do not get connectivity for free.

Note that after applying the trivial trick to an instance \mathbb{F} the distance between each c and c' from $\text{dom}(\mathbb{F})$ will be at most 2. Also, applying this trick turns any instance with a Gaifman graph of a low (bounded) degree into one with a high degree Gaifman graph.

3 THE SKOLEM CHASE.

The chase procedure is a standard algorithm, studied in a plethora of papers. It can be used to semi-decide whether $\mathcal{T}, \mathbb{D} \models \phi(\vec{a})$ for given theory \mathcal{T} , instance \mathbb{D} , CQ $\phi(\vec{y})$ and tuple $\vec{a} \in \text{dom}(\mathbb{D})^{|\vec{y}|}$. In fact the algorithm comes in many variants and flavors. The best way to present our results is by using a variant of the the *semi-oblivious Skolem chase*, which we define in this section.

We say that two CQs have the same *isomorphism type* when one can be obtained from the other by means of bijective renaming of variables. For a CQ Φ define $\tau(\Phi)$ as the isomorphism type of Φ .

For each possible isomorphism type $\tau = \tau(\phi(\vec{y}))$ of some *atomic* conjunctive query $\phi(\vec{y})$, and for each natural number $1 \leq i \leq \text{arity}(E)$, where E is the relation symbol of $\phi(\vec{y})$, let f_i^τ be a function

symbol, with arity equal to $|\vec{y}|$, that is the number of free variables in $\phi(\vec{y})$.

DEFINITION 2 (SKOLEMIZATION). For a given TGD ρ , of the form $\beta(\vec{x}, \vec{y}) \Rightarrow \exists \vec{w} \alpha(\vec{x}, \vec{w})$ by $sh(\rho)$ we denote the Skolemization of the head of ρ , that is the atom $\alpha(\vec{x}, \vec{w})$, with each variable $w \in \vec{w}$ replaced by the term $f_i^\tau(\vec{x})$, where i is the earliest position in $\alpha(\vec{x}, \vec{w})$ where the variable w occurs.

Let, for example ρ be $E(x, y, z), P(x) \Rightarrow \exists v R(y, v, z, v)$. Then $sh(\rho)$ will be the atom $R(y, f_2^\tau(y, z), z, f_2^\tau(y, z))$ where τ is the isomorphism type of $\exists v R(y, v, z, v)$. Notice that $sh(\rho)$ does not depend on the body of ρ , only on its head. In particular it does not depend on the non-frontier variables² of the body of ρ .

Now we can define the procedure of rule application. Parameters of this procedure are an instance \mathbb{F} , a rule ρ , and a mapping σ assigning elements of the active domain of \mathbb{F} to the variables that occur in the body of ρ :

DEFINITION 3 (RULE APPLICATION). Let ρ be a rule of the form $\beta(\vec{x}, \vec{y}) \Rightarrow \exists \vec{w} \alpha(\vec{y}, \vec{w})$, and let \mathbb{F} be a fact set.

- Define $\mathcal{H}om(\rho, \mathbb{F})$ as the set of all mappings σ from variables in $\vec{x} \cup \vec{y}$ to $\text{dom}(\mathbb{F})$ such that $\sigma(\beta(\vec{x}, \vec{y})) \subseteq \mathbb{F}$ (which means that all the atoms from β are in \mathbb{F} after we apply σ to them).
- For $\sigma \in \mathcal{H}om(\rho, \mathbb{F})$, define $appl(\rho, \sigma) = \sigma(sh(\rho))$.

With these notions in place, we can now define the chase procedure as a whole. Given an instance \mathbb{D} and a theory \mathcal{T} , it produces a sequence $(Ch_i(\mathcal{T}, \mathbb{D}))_{i \in \mathbb{N}}$ of instances and the structure $Ch(\mathcal{T}, \mathbb{D})$, according to the following definition.

DEFINITION 4 (SEMI-OBVIOUS SKOLEM CHASE PROCEDURE).

- $Ch_0(\mathcal{T}, \mathbb{D}) = \mathbb{D}$,
- $Ch_{i+1}(\mathcal{T}, \mathbb{D}) = Ch_i(\mathcal{T}, \mathbb{D}) \cup \{appl(\rho, \sigma) \mid \rho \in \mathcal{T}, \sigma \in \mathcal{H}om(\rho, Ch_i(\mathcal{T}, \mathbb{D}))\}$,
- $Ch(\mathcal{T}, \mathbb{D}) = \bigcup_{i \in \mathbb{N}} Ch_i(\mathcal{T}, \mathbb{D})$.

It is well known that $Ch(\mathcal{T}, \mathbb{D})$ is a universal model for \mathcal{T} and \mathbb{D} (i.e., a model that can be homomorphically mapped into any other model). Thus, this structure can be used to solve CQ entailment: for any theory \mathcal{T} , CQ $\phi(\vec{y})$, instance \mathbb{D} and $\vec{a} \in \text{dom}(\mathbb{D})^{|\vec{y}|}$, we have:

$$Ch(\mathcal{T}, \mathbb{D}) \models \phi(\vec{a}) \Leftrightarrow \mathbb{D}, \mathcal{T} \models \phi(\vec{a})$$

EXAMPLE 5. Let \mathcal{T}_a and \mathbb{D}_a be as in Example 1, then

- $Ch_0(\mathcal{T}_a, \mathbb{D}_a) = \mathbb{D}_a = \{Human(Abel)\}$,
- $Ch_1(\mathcal{T}_a, \mathbb{D}_a) = Ch_0(\mathcal{T}_a, \mathbb{D}_a) \cup \{Mother(Abel, mum(Abel))\}$,
- $Ch_2(\mathcal{T}_a, \mathbb{D}_a) = Ch_1(\mathcal{T}_a, \mathbb{D}_a) \cup \{Mother(mum(Abel), mum(mum(Abel)))\}$,

and so on (we use the function symbol “mum” as an alias for the ugly Skolem function symbol from Definition 2).

Now note that there is nothing in Definition 4 that could prevent us from taking $\mathbb{D} = Ch_2(\mathcal{T}_a, \mathbb{D}_a)$ and running the chase for such \mathbb{D} . It is easy to see that in that case we obtain $Ch(\mathcal{T}_a, \mathbb{D}) = Ch(\mathcal{T}_a, \mathbb{D}_a)$. This leads to the following easy insight:

OBSERVATION 6. If $\mathbb{D} \subseteq \mathbb{F} \subseteq Ch(\mathcal{T}, \mathbb{D})$ then $Ch(\mathcal{T}, \mathbb{F}) = Ch(\mathcal{T}, \mathbb{D})$.

²Including non-frontier variables as arguments of the functions f_i^τ (like x in the current example) would lead to the *oblivious chase*. Also note that the names of the terms do not identify ρ (their “rule of origin”). This is going to be important in the proof of Theorem 1.

¹The queries $\phi(\vec{y})$ and $\psi(\vec{y})$ are seen as structures here: the active domains of these structures are the sets of variables of $\phi(\vec{y})$ and $\psi(\vec{y})$.

Note that that equality in Observation 6 is to be understood literally (rather than “up to isomorphism”). This is crucial for our treatise and it constitutes the main reason why we use the Skolem naming convention.

Finally, let us state another property of the chase – a direct consequence of the fact that $Ch(\mathcal{T}, \mathbb{D})$ is a universal model.

PROPERTY 7. *Let \mathcal{T} be a theory and let \mathbb{D} and \mathbb{F} be fact sets satisfying $\mathbb{D} \subseteq \mathbb{F} \subseteq Ch(\mathcal{T}, \mathbb{D})$ as well as $\mathbb{F} \models \mathcal{T}$. Then there exists a homomorphism from $Ch(\mathcal{T}, \mathbb{D})$ to \mathbb{F} that is the identity on $dom(\mathbb{F})$.*

Frontier and birth atoms. Let α be an atom from $Ch(\mathcal{T}, \mathbb{D}) \setminus \mathbb{D}$, created as $appl(\rho, \sigma)$ for some $\rho \in \mathcal{T}$ and $\sigma \in Hom(\rho, \mathbb{D})$. We let the frontier of α (written: $fr(\alpha)$) denote the set of terms $\sigma(fr(\rho))$. Notice that there may be more than one rule application creating the same atom α , but:

OBSERVATION 8. *For every $\rho, \rho' \in \mathcal{T}$ and every $\sigma \in Hom(\rho, \mathbb{D})$ if $appl(\rho, \sigma) = appl(\rho', \sigma')$ then heads of ρ and ρ' when treated as CQs have the same isomorphism type.*

Clearly, for each $t \in dom(Ch(\mathcal{T}, \mathbb{D}))$ either $t \in dom(\mathbb{D})$ or t was created by the chase procedure as a Skolem term. Notice that, despite the name of the rule of \mathcal{T} that created t not being indicated in the name of t (as per our Skolem naming convention), the following still holds:

OBSERVATION 9. *Suppose $t \in dom(Ch(\mathcal{T}, \mathbb{D})) \setminus dom(\mathbb{D})$. Then there exists exactly one atom $\alpha \in Ch(\mathcal{T}, \mathbb{D})$ such that t appears in α , but $t \notin fr(\alpha)$.*

We will refer to such an atom α as the *birth atom* of t .

For the proof of the last observation notice that one can uniquely reconstruct α using only t . This is because t contains both the isomorphism type of α (so that one can reconstruct α up to bijective renaming of its terms) and its frontier terms.

4 THE THREE CLASSES

For a theory \mathcal{T} , an instance \mathbb{D} , a natural number n and a conjunctive query $\phi(y)$ we will write *Enough*($n, \phi(y), \mathbb{D}, \mathcal{T}$) as a shorthand for:

$$\forall \bar{a} \in dom(\mathbb{D})^{|\bar{y}|} \quad (Ch(\mathcal{T}, \mathbb{D}) \models \phi(\bar{a}) \Leftrightarrow Ch_n(\mathcal{T}, \mathbb{D}) \models \phi(\bar{a})).$$

Meaning that “it is *enough* to run n steps of the \mathcal{T} -chase on \mathbb{D} to fully evaluate ϕ ”.

Now we will provide definitions for three separate classes: BDD, FES, and UBDD [2, 14].

DEFINITION 10 (BDD). *A theory \mathcal{T} has the bounded derivation depth property if:*

$$\forall \Phi \exists n_{\Phi} \forall \mathbb{D} \text{ Enough}(n_{\Phi}, \Phi, \mathbb{D}, \mathcal{T})$$

DEFINITION 11 (FES). *A theory \mathcal{T} has the finite expansion set property if:*

$$\forall \mathbb{D} \exists n_{\mathbb{D}} \forall \Phi \text{ Enough}(n_{\mathbb{D}}, \Phi, \mathbb{D}, \mathcal{T})$$

DEFINITION 12 (UBDD). *Theory \mathcal{T} is uniformly BDD if:*

$$\exists c_{\mathcal{T}} \forall \mathbb{D} \forall \Phi \text{ Enough}(c_{\mathcal{T}}, \Phi, \mathbb{D}, \mathcal{T})$$

Notice the striking similarity between the above expressions. Connections between those classes will be discussed later in Section 7, but before we need to explore BDD and FES in greater detail.

5 BDD AND FINITE UNIFICATION SETS

As with many important notions, it happens that the same concept is defined by multiple communities using different properties. This is exactly the case with *finite unification set* property [2] and BDD.

DEFINITION 13 (FUS). *A theory \mathcal{T} has the finite unification set property if every conjunctive query $\psi(\bar{y})$ has a rewriting – a UCQ, denoted $\psi^{rew}(\bar{y})$, such that the following holds:*

For each instance \mathbb{D} and each tuple $\bar{a} \in dom(\mathbb{D})^{|\bar{y}|}$ we have

$$Ch(\mathcal{T}, \mathbb{D}), \bar{a} \models \psi(\bar{y}) \iff \mathbb{D}, \bar{a} \models rew(\psi(\bar{y})).$$

It is well known that \mathcal{T} is BDD if and only if it is FUS [3], so we will use the two terms interchangeably. For simplicity of our arguments and without loss of generality, we will require the set $rew(\psi(\bar{y}))$ to be *minimal*: If $\phi(\bar{y}) \neq \phi'(\bar{y})$ are two elements of $rew(\psi(\bar{y}))$ then $\phi(\bar{y})$ is not contained in $\phi'(\bar{y})$.

The BDD/FUS class admits several interesting properties. The following will be used later: facts about terms are produced by the chase soon after the terms are created (with only a constant delay).

OBSERVATION 14. *There exists a natural number n_{at} (depending only on \mathcal{T}) such that for any instance \mathbb{D} , for any $i \in \mathbb{N}$, for any tuple \bar{i} of domain elements from $dom(Ch_i(\mathcal{T}, \mathbb{D}))$ and for any $R \in \Sigma$, $Ch(\mathcal{T}, \mathbb{D}) \models R(\bar{i})$ implies $Ch_{i+n_{at}}(\mathcal{T}, \mathbb{D}) \models R(\bar{i})$.*

PROOF (SKETCH). For any query ϕ , let n_{ϕ} denote the constant from Definition 10. Note, that there is only a finite number of non-isomorphic atomic queries. Therefore, we can obtain n_{at} as $\max(\{n_{\phi} \mid \phi \text{ is an atomic query}\})$. \square

The BDD property – exercises. Now, we would like to encourage the reader to solve a few exercises. While they are not part of the actual proofs, we believe they might provide valuable insights.

EXERCISE 15. *Consider the theory \mathcal{T}_p over schema $\{E\}$ consisting of just one rule $E(x, y) \Rightarrow \exists z E(y, z)$. Show that this theory is BDD.*

Comment: This can be easily generalized: it is well known that all linear theories are BDD [14] (a theory is linear if each rule only has one atom in its body).

EXERCISE 16. *Show that, if \mathcal{T} is BDD and connected then there exists some $d \in \mathbb{N}$ such that for each \mathbb{D} and for each two terms $c, c' \in dom(\mathbb{D})$, if $dist_{Ch(\mathcal{T}, \mathbb{D})}(c, c') = 1$ then $dist_{\mathbb{D}}(c, c') \leq d$.*

Comment: the didactic purpose of Exercise 16 is to evoke or reinforce the intuition of BDD as a “locality” property: if terms from $dom(\mathbb{D})$ appear in one atom somewhere in $Ch(\mathcal{T}, \mathbb{D})$ then they could not possibly be far away from each other already in \mathbb{D} .

6 FES AND CORE TERMINATION

FES theories always (regardless of the initial structure \mathbb{D}) produce all the positive information present in $Ch(\mathcal{T}, \mathbb{D})$ already after a finite number of chase steps (this number can depend on \mathbb{D} though). This is, as well, the case with the *core termination* [1, 2]:

DEFINITION 17. *A theory \mathcal{T} is core-terminating if for each fact set \mathbb{D} there exists a $k \in \mathbb{N}$ such that for each $i \geq k$:*

$$Core(Ch_i(\mathcal{T}, \mathbb{D})) \text{ is isomorphic to } Core(Ch_{i+1}(\mathcal{T}, \mathbb{D})).$$

It is well known that \mathcal{T} is core-terminating if and only if it is FES [15], so we will use the two terms interchangeably. Moreover, the smallest numbers $n_{\mathbb{D}}$ and k satisfying Definitions 11 and 17 (respectively) are equal; from now on this number will be denoted with $c_{\mathcal{T},\mathbb{D}}$. It is also known [1] that $\text{Core}(\text{Ch}_{c_{\mathcal{T},\mathbb{D}}}(\mathcal{T}, \mathbb{D})) \models \mathcal{T}, \mathbb{D}$.

DEFINITION 18. *Given a FES theory \mathcal{T} and an instance \mathbb{D} , we let $\text{Core}(\mathcal{T}, \mathbb{D})$ denote $\text{Core}(\text{Ch}_{c_{\mathcal{T},\mathbb{D}}}(\mathcal{T}, \mathbb{D}))$.*

Exercises Also this section comes with a few exercises. Again, we expect them to provide valuable insight, but they are not required for our subsequent proofs.

EXERCISE 19. *Show that the theory from Exercise 15 is not FES.*

EXERCISE 20. *Show that the theory consisting of two rules $E(x, y) \Rightarrow \exists z E(y, z)$ and $E(x, x') \Rightarrow E(x, x)$ is FES.*

7 THE FUS/FES CONJECTURE

It is very easy to produce examples of BDD theories that are not UBDD (see Exercise 15). However, all examples we could produce are not core-terminating. Likewise, it is easy to produce examples of core-terminating theories that are not UBDD, but they are not BDD either (all unbounded Datalog theories will be very happy to serve as examples). This gives rise to the following conjecture.

CONJECTURE 1 (THE FUS/FES CONJECTURE). *Any theory that is both FUS and FES is also UBDD.*

This conjecture was studied in [11] where an incorrect proof was proposed, and in [12] where it was proved that it would hold true if the assumption that \mathcal{T} is FES was replaced by the significantly **stronger** assumption that it is *all-instances Skolem chase terminating*. Note that the conjecture would be false if infinite theories, over infinite (yet just binary) schemas were allowed:

EXAMPLE 21. *Suppose a relation symbol E_i for every $i \in \mathbb{N}$. Let the theory \mathcal{T}_{∞} consist of all rules of the form $E_i(x, y) \Rightarrow \exists z E_{i-1}(y, z)$ for $i \in \mathbb{N}^+$. Then \mathcal{T}_{∞} is BDD and core-terminating, but it is not UBDD. To see why this is the case, notice that only facts from a finite number of relations can appear in every given finite instance.*

8 LOCAL THEORIES.

We are now ready to introduce the central notion of *local theories*. As it will turn out, not only the FUS/FES conjecture holds for these (Theorem 2) but they also subsume all BDD theories over binary signatures (Theorem 1). Notwithstanding, as we will argue near the end of the paper, there exists an untapped potential within the BDD class beyond the veils of locality (Theorem 3).

DEFINITION 22. *A theory \mathcal{T} is local if there exists some number $l_{\mathcal{T}} \in \mathbb{N}$ such that for every instance \mathbb{D} the following holds:*

$$\text{Ch}(\mathcal{T}, \mathbb{D}) = \bigcup_{\mathbb{F} \subseteq \mathbb{D}, |\mathbb{F}| \leq l_{\mathcal{T}}} \text{Ch}(\mathcal{T}, \mathbb{F})$$

Note that the Skolem naming convention is important here. Without it, it would be unclear or at least ambiguous what a union of chases is supposed to mean. Importantly, we obtain that locality implies the BDD property.

OBSERVATION 23. *If a theory is local then it is BDD.*

PROOF (SKETCH). Given a local theory \mathcal{T} and a CQ Φ , we show that there exists a natural number n_{Φ} satisfying Definition 10. Due to \mathcal{T} being local, we can, for every \mathbb{F} with $\mathbb{F} \models \Phi$, identify some $\mathbb{D} \subseteq \mathbb{F}$ with $\mathbb{D} \models \Phi$ and $|\mathbb{D}| \leq l_{\mathcal{T}}|\Phi|$. Therefore, we let $\mathcal{F} = \{\mathbb{D} \mid \text{Ch}(\mathcal{T}, \mathbb{D}) \models \Phi \wedge |\mathbb{D}| \leq l_{\mathcal{T}}|\Phi|\}$. For any $\mathbb{D} \in \mathcal{F}$, define $n_{\mathbb{D}}$ as the minimal natural number such that $\text{Ch}_{n_{\mathbb{D}}}(\mathcal{T}, \mathbb{D}) \models \Phi$. As \mathcal{F} contains just a finite number of non-isomorphic instances, picking $n_{\Phi} = \max\{n_{\mathbb{D}} \mid \mathbb{D} \in \mathcal{F}\}$ witnesses that Definition 10 applies. \square

We defer any further discussion about locality until later sections. And now, let us state our first result.

THEOREM 1. *Every BDD theory over binary signature is local.*

The detailed proof will be presented in Appendix A. Here, we just outline the proof idea. We start by observing that the atoms created in the chase by rules that contain existentially quantified variables form a forest. This is a crucial property of single-head rule sets over binary signatures as any such rule must have a frontier of size ≤ 1 . (Obviously, this property ceases to hold for signatures of higher arities.) Essentially, as the chase progresses, the existential atoms (those produced through “proper” existential rules) are created further and further away from the instance. If we were able to bound the number of each such atom’s *ancestors* — the atom set required for its creation — then we could finish the proof by slightly extending Observation 14. The forest shape, however, is not immediately sufficient to bound the ancestors of its atoms. To overcome this issue, we propose a normalization technique, that, given a BDD theory \mathcal{T} over binary signature, produces a new theory \mathcal{T}_{NF} , which might no longer be BDD but admits two important properties: First, the sets of existential atoms produced by \mathcal{T} and by \mathcal{T}_{NF} on any instance coincide (Lemma 48). Second, it is straightforward to find the mentioned bound on the number of ancestors for every atom produced by \mathcal{T}_{NF} (Lemma 57). These two properties grant us Theorem 1. But how does such a normalization work? In simple words, it relies on the fact that \mathcal{T} is BDD. In the absence of both disconnected and Datalog rules, the normalization would be rather simple: we could take any existential rule $\beta \Rightarrow \alpha \in \mathcal{T}$ and replace it with rules $\gamma \Rightarrow \alpha$ for all $\gamma \in \text{rew}(\beta)$. Dealing with disconnected bodies and Datalog rules complicates matters somewhat.

We are very confident that our proof can be generalized to all frontier-guarded BDD theories. As this is not yet spelled out in full detail, we prefer to be cautious and formulate it as conjecture.

CONJECTURE 2. *Every frontier-guarded BDD theory is local.*

9 THE FUS/FES CONJECTURE IS TRUE FOR LOCAL THEORIES

We proceed by presenting the second of our three main results.

THEOREM 2. *If a theory \mathcal{T} is FES and local then it is UBDD.*

This means that the FUS/FES conjecture holds for local theories. Before we start our proof, let us note that by Theorem 1 and Theorem 2, we can immediately conclude that Conjecture 1 holds for theories over binary signatures:

COROLLARY 1. *If a theory \mathcal{T} over a binary signature is both FES and BDD then it is UBDD.*

Also, Theorem 1 implies that any counterexample to the conjecture – should it exist – would have to be outside the realm of local classes. We explore this uncharted lands in Section 11.

Proof of Theorem 2

As any UBDD theory is core-terminating as well, it is easy to see that a theory \mathcal{T} is UBDD if and only if there exists some $c_{\mathcal{T}} \in \mathbb{N}$ such that $\text{Core}(\mathcal{T}, \mathbb{D}) \subseteq \text{Ch}_{c_{\mathcal{T}}}(\mathcal{T}, \mathbb{D})$ holds for any instance \mathbb{D} . Note that the numbers $c_{\mathcal{T}}$ here and in Definition 12 are equal. Thus we can reformulate Theorem 2 as follows:

THEOREM 2 (ALTERNATIVE). *Let \mathcal{T} be a core-terminating local theory. Then there exists a $c_{\mathcal{T}} \in \mathbb{N}$ such that $\text{Core}(\mathcal{T}, \mathbb{D}) \subseteq \text{Ch}_{c_{\mathcal{T}}}(\mathcal{T}, \mathbb{D})$ holds for any instance \mathbb{D} .*

Until the end of this section, we will consider \mathcal{T} a fixed theory that is both core-terminating and local (so also BDD). To simplify notation, $\text{Ch}(\mathcal{T}, \mathbb{D})$ will be shortened to $\text{Ch}(\mathbb{D})$.

DEFINITION 24. *For an instance \mathbb{D} define $\mathcal{I}_{\mathbb{D}}$ as the family of sets $\{\mathbb{F} \mid \mathbb{F} \subseteq \mathbb{D}, |\mathbb{F}| \leq l_{\mathcal{T}}\}$. Define $\mathbb{C}_{\mathbb{D}} = \bigcup_{\mathbb{F} \in \mathcal{I}_{\mathbb{D}}} \text{Core}(\mathcal{T}, \mathbb{F})$.*

LEMMA 25. *There exists a $k_{\mathcal{T}} \in \mathbb{N}$ depending only on \mathcal{T} (but not on \mathbb{D}), such that $\mathbb{C}_{\mathbb{D}} \subseteq \text{Ch}_{k_{\mathcal{T}}}(\mathbb{D})$.*

PROOF. The set $\mathcal{A} = \{\mathbb{F} \mid |\mathbb{F}| \leq l_{\mathcal{T}}\}$ of all instances (over Σ) of size at most $l_{\mathcal{T}}$ is finite (up to isomorphisms). Recall that \mathcal{T} is core-terminating and let $k_{\mathcal{T}} = \max\{c_{\mathcal{T}, \mathbb{D}} \mid \mathbb{D} \in \mathcal{A}\}$. \square

If we were able able to find a homomorphism $\tilde{h}_{\mathbb{D}}$ from $\text{Ch}(\mathbb{D})$ to $\mathbb{C}_{\mathbb{D}}$, the alternative formulation of Theorem 2 would be proved. Also, since \mathcal{T} is core-terminating, we know that for each $\mathbb{F} \in \mathcal{I}_{\mathbb{D}}$ there exists a homomorphism $h_{\mathbb{F}}$ from $\text{Ch}(\mathbb{F})$ to $\mathbb{C}_{\mathbb{D}}$ and we know that $\bigcup_{\mathbb{F} \in \mathcal{I}_{\mathbb{D}}} \text{Ch}(\mathbb{F}) = \text{Ch}(\mathbb{D})$. So can't we just define $\tilde{h}_{\mathbb{D}} = \bigcup_{\mathbb{F} \in \mathcal{I}_{\mathbb{D}}} h_{\mathbb{F}}$? Unfortunately not, because the domains of $h_{\mathbb{F}}$ and $h_{\mathbb{F}'}$ may overlap (for some $\mathbb{F} \neq \mathbb{F}'$) and there is no guarantee that $h_{\mathbb{F}}$ and $h_{\mathbb{F}'}$ will agree on the terms that are in both domains. If $\tilde{h}_{\mathbb{D}}$ could be produced this way, $\mathbb{C}_{\mathbb{D}} \models \mathcal{T}$ would always hold. Yet, we found an example (not included here) of a pair \mathbb{D}, \mathcal{T} for which $\mathbb{C}_{\mathbb{D}} \not\models \mathcal{T}$.

Luckily, the idea to build a global homomorphism $\tilde{h}_{\mathbb{D}}$ using the local homomorphisms $h_{\mathbb{F}}$ can be put to use in a different way, and the set of facts $\mathbb{C}_{\mathbb{D}}$ will indeed prove very useful in this context. The following lemma will be crucial in this endeavor:

LEMMA 26. *For any instance \mathbb{D} there exists a homomorphism $\tilde{h}_{\mathbb{D}}$ from $\text{Ch}(\mathbb{D})$ to $\text{Ch}(\mathbb{D})$ such that for each $t \in \text{dom}(\text{Ch}(\mathbb{D}))$ there is $\tilde{h}_{\mathbb{D}}(t) \in \text{dom}(\mathbb{C}_{\mathbb{D}})$.*

Let us first discuss how Theorem 2 can be concluded from Lemma 26. Suppose some \mathbb{D} is fixed and $\tilde{h}_{\mathbb{D}}$ is a homomorphism as in Lemma 26. We know that $t \in \text{dom}(\text{Ch}(\mathbb{D}))$ implies $\tilde{h}_{\mathbb{D}}(t) \in \text{dom}(\mathbb{C}_{\mathbb{D}})$ and we know that $\mathbb{C}_{\mathbb{D}} \subseteq \text{Ch}_{k_{\mathcal{T}}}(\mathbb{D})$. So one might be tempted to immediately conclude $\tilde{h}_{\mathbb{D}}(\text{Ch}(\mathbb{D})) \subseteq \text{Ch}_{k_{\mathcal{T}}}(\mathbb{D})$. But it is not quite that simple. Admittedly, Lemma 26 tells us that all the terms of $\tilde{h}_{\mathbb{D}}(\text{Ch}(\mathbb{D}))$ will indeed appear in $\text{Ch}_{k_{\mathcal{T}}}(\mathbb{D})$. But it says nothing like that about the atoms of $\tilde{h}_{\mathbb{D}}(\text{Ch}(\mathbb{D}))$. Rather, it might be that there are atoms in $\tilde{h}_{\mathbb{D}}(\text{Ch}(\mathbb{D}))$ that, despite having all their terms in $\text{dom}(\mathbb{C}_{\mathbb{D}})$ are not themselves in $\mathbb{C}_{\mathbb{D}}$. To overcome this little problem, we recall Observation 14 and let $c_{\mathcal{T}} = k_{\mathcal{T}} + n_{at}$. Then $\tilde{h}_{\mathbb{D}}(\text{Ch}(\mathbb{D})) \subseteq \text{Ch}_{c_{\mathcal{T}}}(\mathbb{D})$ follows as desired.

This means what remains to be presented in this section is the **proof of Lemma 26**:

DEFINITION 27. *Let \mathbb{D} be a set of facts and let $\mathbb{F} \subseteq \mathbb{D}$. We let $M_{\mathbb{F}}$ denote³ the substructure of $\text{Ch}(\mathbb{D})$ induced by the set of terms $\text{dom}(\text{Ch}(\mathbb{D})) \setminus (\text{dom}(\text{Ch}(\mathbb{F})) \setminus \text{dom}(\text{Core}(\mathcal{T}, \mathbb{F})))$.*

In the following, the terms of $\text{dom}(\text{Ch}(\mathbb{F})) \setminus \text{dom}(\text{Core}(\mathcal{T}, \mathbb{F}))$ will be referred to as *banned terms*.

LEMMA 28. *For any instance \mathbb{D} and for any $\mathbb{F} \subseteq \mathbb{D}$, the structure $M_{\mathbb{F}}$ is a model of \mathcal{T} and \mathbb{D} .*

PROOF. Clearly $M_{\mathbb{F}} \models \mathbb{D}$. In order to prove $M_{\mathbb{F}} \models \mathcal{T}$, consider any $\rho \in \mathcal{T}$ and any $\sigma \in \text{Hom}(\rho, M_{\mathbb{F}})$. Of course, $\text{appl}(\rho, \sigma) \in \text{Ch}(\mathbb{D})$, since $\text{Ch}(\mathbb{D})$ is by definition closed under rule applications.

In the following, the atom $\text{appl}(\rho, \sigma)$ will be mentioned often enough to deserve a shorter name, so we will call it α .

It is now sufficient (and necessary) to prove that there exists a homomorphism from α to some atom $\alpha' \in M_{\mathbb{F}}$, that is the identity on $\text{fr}(\alpha)$. In other words, we need to show that if the body of the rule ρ matches $M_{\mathbb{F}}$ (via mapping σ), then we can find an atom in $M_{\mathbb{F}}$ that witnesses satisfaction of ρ . Such an α' needs to have the same terms as α in the frontier positions and may have arbitrary terms in the positions of the existentially quantified variables in $\text{head}(\rho)$, except that if α had equal terms on two such positions then the respective terms in α' must also be equal.

If $\alpha \in M_{\mathbb{F}}$ then of course we pick $\alpha' = \alpha$. So, for the rest of the proof, assume $\alpha \notin M_{\mathbb{F}}$. Note that the only reason for α to be in $\text{Ch}(\mathbb{D})$ but not in $M_{\mathbb{F}}$ is that α contains some banned term t .

But $\sigma(\text{body}(\rho)) \subseteq M_{\mathbb{F}}$. Thus $\text{fr}(\alpha) \subseteq \text{dom}(M_{\mathbb{F}})$ and so $t \notin \text{fr}(\alpha)$. At this point, we can be sure that ρ is not a Datalog rule – atoms derived via a Datalog rule do not have non-frontier terms.

Term t being a non-frontier term of α means that α is the birth atom of t in $\text{Ch}(\mathbb{D})$. But $t \in \text{dom}(\text{Ch}(\mathbb{F}))$ so from Observation 9, we know that α is the birth atom of t in $\text{Ch}(\mathbb{F})$ and thus $\alpha \in \text{Ch}(\mathbb{F})$.

Note that $\text{fr}(\alpha) \subseteq \text{dom}(\text{Core}(\mathcal{T}, \mathbb{F}))$ (*), as $\text{fr}(\alpha) \subseteq \text{dom}(\text{Ch}(\mathbb{F}))$ and none of the terms in $\text{fr}(\alpha)$ are banned.

Now let $h_{\mathbb{F}}^* : \text{Ch}(\mathbb{F}) \rightarrow \text{Core}(\mathcal{T}, \mathbb{F})$ be a homomorphism as in Property 7. Since $\alpha \in \text{Ch}(\mathbb{F})$, we obtain $h_{\mathbb{F}}^*(\alpha) \in \text{Core}(\mathcal{T}, \mathbb{F})$ and thus $h_{\mathbb{F}}^*(\alpha) \in M_{\mathbb{F}}$. As $h_{\mathbb{F}}^*$ is a retraction, $h_{\mathbb{F}}^*(\text{fr}(\alpha)) = \text{fr}(\alpha)$ follows from (*). Hence, $h_{\mathbb{F}}^*(\alpha)$ can serve as our α' , concluding the proof. \square

LEMMA 29. *For any instance \mathbb{D} and any $\mathbb{F} \subseteq \mathbb{D}$ there exists a homomorphism $h_{M_{\mathbb{F}}}^*$ from $\text{Ch}(\mathbb{D})$ to itself that maps all terms to $\text{dom}(M_{\mathbb{F}})$ and is the identity on $\text{dom}(M_{\mathbb{F}})$.*

PROOF. Note that $M_{\mathbb{F}} \models \mathcal{T}$ (Lemma 28) and $\mathbb{D} \subseteq M_{\mathbb{F}}$. Then Property 7 ensures the existence of the claimed homomorphism. \square

Let $\mathcal{H}_{\mathbb{D}}$ be the set of all homomorphisms $h_{M_{\mathbb{F}}}^*$ for $\mathbb{F} \in \mathcal{I}_{\mathbb{D}}$.

Each $h_{M_{\mathbb{F}}}^* \in \mathcal{H}_{\mathbb{D}}$ has as its domain the set $\text{dom}(\text{Ch}(M_{\mathbb{F}}))$, that is equal to $\text{dom}(\text{Ch}(\mathbb{D}))$, and has as its image a subset of this domain. This means that one can compose such homomorphisms, and the resulting function will also be a homomorphism from $\text{Ch}(\mathbb{D})$ to $\text{Ch}(\mathbb{D})$ (and it will be the identity on $\text{dom}(\mathbb{D})$, since each $h_{M_{\mathbb{F}}}^*$ is). Now the rabbit is going to be pulled out of the hat: let us compose

³To be precise we should call this new structure $M_{\mathbb{D}, \mathbb{F}}$, but \mathbb{D} will be fixed and clear from the context.

all homomorphisms $h_{M_{\mathbb{F}}} \in \mathcal{H}_{\mathbb{D}}$, in any order. Call the resulting (“global”) homomorphism $\bar{h}_{\mathbb{D}}$.

Now recall that the proof of Lemma 26 (and thus also of Theorem 2) will be finished once we can show that $\bar{h}_{\mathbb{D}}(t) \in \text{dom}(\mathbb{C}_{\mathbb{D}})$ does indeed hold for each term $t \in \text{dom}(\text{Ch}(\mathbb{D}))$.

Recall our notion of banned terms. Now \mathbb{F} is no longer fixed, i.e., for each $\mathbb{F} \in \mathcal{I}_{\mathbb{D}}$ there is a set $\text{ban}_{\mathbb{F}}$ of terms that occur somewhere in $\text{Ch}(\mathbb{F})$ but not in $\text{Core}(\mathcal{T}, \mathbb{F})$. Each $h_{M_{\mathbb{F}}}^* \in \mathcal{H}_{\mathbb{D}}$ is the identity on all terms except those of $\text{ban}_{\mathbb{F}}$, and maps the terms from $\text{ban}_{\mathbb{F}}$ into $\text{dom}(\text{Core}(\mathcal{T}, \mathbb{F}))$, which means into $\text{dom}(\mathbb{C}_{\mathbb{D}})$.

Now suppose we apply $\bar{h}_{\mathbb{D}}$ to any $t \in \text{dom}(\text{Ch}(\mathbb{D}))$. If there is any $h_{M_{\mathbb{F}}}^* \in \mathcal{H}_{\mathbb{D}}$ with $h_{M_{\mathbb{F}}}^*(t) \neq t$ then of course $\bar{h}_{\mathbb{D}}(t) \in \text{dom}(\mathbb{C}_{\mathbb{D}})$. In case $h_{M_{\mathbb{F}}}^*(t) = t$ for each $h_{M_{\mathbb{F}}}^*$, consider any $\mathbb{F}_t \in \mathcal{I}_{\mathbb{D}}$ for which $t \in \text{dom}(\text{Ch}(\mathbb{F}_t))$. Then $t \in \text{dom}(\text{Core}(\mathcal{T}, \mathbb{F}_t)) \subseteq \text{dom}(\mathbb{C}_{\mathbb{D}})$.

10 SLIGHTLY BEYOND LOCALITY: STICKY THEORIES

Unfortunately, our notion of locality fails to characterize the entire BDD class, as demonstrated in the following example.

EXAMPLE 30. Let E be a relation of arity 4 and R one of arity 2. Read $E(a, b, b', c)$ as “ a sees an edge from b to b' colored with color c ” and $R(a, c)$ as “ a considers c a color”. The following one-rule sticky theory \mathcal{T} is not local:

$$E(x, y, y', t), R(x, t') \Rightarrow \exists y'' E(x, y', y'', t')$$

(meaning “if x sees an edge from y to y' and considers t' a color, then x must also see another edge from y' to some y'' of color t' ”).

To see that it is indeed not local, suppose it were and let $l_{\mathcal{T}}$ be the corresponding constant as in Definition 22. Now take an instance \mathbb{D} consisting of $l_{\mathcal{T}} + 1$ atoms: one atom $E(a, b_1, b_2, c_1)$ and atoms $R(a, c_i)$ for $1 \leq i \leq l_{\mathcal{T}}$. It is not hard to see that there are atoms in $\text{Ch}(\mathcal{T}, \mathbb{D})$ that require all the atoms from \mathbb{D} to be produced.

The only reason, however, for connected sticky theories to be non-local are high-degree vertices, like the a in the example. This leads to a natural generalization of the notion of locality:

DEFINITION 31. A theory \mathcal{T} will be called bounded-degree local (or *bd-local*) if for any $k \in \mathbb{N}$ there exists a constant $l_{\mathcal{T}}(k)$ such that for every instance \mathbb{D} having degree at most k , the following holds:

$$\bigcup_{\mathbb{F} \subseteq \mathbb{D}, |\mathbb{F}| \leq l_{\mathcal{T}}(k)} \text{Ch}(\mathcal{T}, \mathbb{F}) = \text{Ch}(\mathcal{T}, \mathbb{D})$$

As of yet, we have been unable to show that the FUS/FES conjecture holds for bounded-degree local theories, but we believe that with some additional effort, the ideas from Section 9 could probably be adapted to work also for such theories. And of course they do work if only instances of fixed degree are considered.

It is not hard to show that sticky theories are indeed *bd-local* (cf. Appendix E). Hence, in view of Conjecture 2, it seems that most known decidable BDD classes are bounded-degree local. Perhaps surprisingly, unlike local theories, not all bounded-degree local theories are BDD:

EXAMPLE 32. It is easy to see that the following single-rule theory is bounded-degree local but not BDD: $E(x, y, z), R(x, z) \Rightarrow R(y, z)$.

But even if not all *bd-local* theories are BDD, it is not straightforward to come up with a BDD theory that is not *bd-local*. So a

natural question arises: are there BDD theories that are not local in this generalized sense? We found it quite surprising to realize that the answer is positive:

EXAMPLE 33. The following BDD theory \mathcal{T}_c is not *bd-local*:

$$\begin{aligned} E(x, y) &\Rightarrow \exists x', y' R(x, y, x', y') \\ R(x, y, x', y'), E(y, z) &\Rightarrow \exists z' R(y, z, y', z') \end{aligned}$$

To prove that it is BDD one can notice that if $\mathcal{T}_c, \mathbb{D} \models \phi(\bar{a})$, for some \mathbb{D} and some $\bar{a} \in \text{dom}(\mathbb{D})^{|\bar{a}|}$ then $\text{Ch}_{|\phi(\bar{y})|} \models \phi(\bar{a})$. In order to prove that it is not *bd-local* consider, for each $n \in \mathbb{N}$, the instance \mathbb{D}_n consisting of atoms $E(a_1, a_2), E(a_2, a_3), \dots, E(a_n, a_1)$. The degree of this instance is 2. And there are atoms in $\text{Ch}_n(\mathcal{T}_c, \mathbb{D}_n)$ that are not in $\text{Ch}_n(\mathcal{T}_c, \mathbb{F})$ for any proper subset \mathbb{F} of \mathbb{D}_n .

We were, however, not able to find an example of a theory that would be *hereditary BDD* — a BDD theory such that every its subsets is BDD as well — but not *bd-local*. We think it reasonable to conjecture that there are no such theories.

11 FAR BEYOND LOCALITY: BDD THEORIES WITHOUT SMALL REWRITINGS

As we know, any local theory is also BDD. Additionally, local theories admit rewritings of linear width:

OBSERVATION 34. For each local theory \mathcal{T} and for each CQ Ψ , the size of the greatest disjunct in the rewriting is at most $l_{\mathcal{T}}|\Psi|$.

Clearly, the linear bound on the width of $\text{rew}(\Psi)$ gives us an immediate exponential upper bound on the number of its disjuncts. A matching lower bound is trivial to obtain:

OBSERVATION 35. Let \mathcal{T} consist of the two rules: $E(x, y) \Rightarrow R(x, y)$ and $E'(x, y) \Rightarrow R(x, y)$. Then \mathcal{T} is BDD and the number of disjuncts in $\text{rew}(\Psi)$ can be exponential in the size of Ψ .

Recall the notion of backwards shy theories [10] — these are BDD theories such that, for every query $\psi(\bar{y})$ if $\phi(\bar{y}) \in \text{rew}(\psi(\bar{y}))$ then only variables from \bar{y} can occur more than once in $\phi(\bar{y})$. Sticky theories are backward shy. It is easy to see that backward shy theories admit rewritings of linear width as well and, in consequence, also all sticky theories do. This is related to another notion of locality:

DEFINITION 36. We call a theory \mathcal{T} distancing if there is a $d_{\mathcal{T}} \in \mathbb{N}$ such that for any instance \mathbb{D} , any $c, c' \in \text{dom}(\mathbb{D})$, and any $n \in \mathbb{N}$ if $\text{dist}_{\text{Ch}(\mathcal{T}, \mathbb{D})}(c, c') \leq n$ then $\text{dist}_{\mathbb{D}}(c, c') \leq d_{\mathcal{T}}n$.

Is every BDD theory distancing? It might seem that this can be shown using Exercise 16. However, this is not the case, since the path from c to c' in $\text{Ch}(\mathcal{T}, \mathbb{D})$ might lead through atoms not containing any constants from the original \mathbb{D} . What can be proven is the following correspondency:

OBSERVATION 37. If a BDD theory admits rewritings of linear width, then it is distancing.

Assuming Conjecture 2, this implies that all theories from previously known BDD classes are distancing. The converse of Observation 37 does not hold, and such theories can be easily found:

OBSERVATION 38. The theory consisting of the single Datalog rule $A(x), E(x, y) \Rightarrow A(y)$ is distancing but not BDD.

So do there exist non-distancing BDD theories at all? Do there exist BDD theories that do not admit rewritings of linear width? The answer is given by Theorem 3, which constitutes the third main result of this paper:

THEOREM 3. *There exists a BDD theory that is non-distancing and does not even admit rewritings of polynomial width.*

To see how this is possible, consider the following theory:

DEFINITION 39. *Consider a signature with two binary predicates R and G . Let the theory \mathcal{T}_d consist of the following rules:*

$$\begin{aligned} (\text{loop}) \quad & \text{true} \Rightarrow \exists x R(x, x), G(x, x) \\ (\text{pins}) \quad & \forall x (\text{true} \Rightarrow \exists z, z' R(x, z), G(x, z')) \\ (\text{grid}) \quad & R(x, x'), G(x, u), G(u, u') \Rightarrow \exists z R(u', z), G(x', z) \end{aligned}$$

Note that the rules of \mathcal{T}_d are not single-head and some of them have empty bodies. One could easily reformulate them to avoid this at the cost of readability (see Appendix D).

We will think of instances over our signature (and of bodies of queries) as graphs with edges colored in red or green. For $n \in \mathbb{N}$, let $G^n(x_0, x_n)$ denote the CQ $\exists x_1 \dots \exists x_{n-1} G(x_0, x_1), \dots, G(x_{n-1}, x_n)$ and $R^n(x_0, x_n)$ likewise. Define conjunctive queries $\phi_R^n(x, y)$ by $\exists x', y' R^n(x, x'), R^n(y, y'), G(x', y')$ and let $\mathbb{G}^n(a, b)$ be a path of n green edges, with a as the first vertex and b as the last.

The following technical lemma substantiates Theorem 3:

LEMMA 40. (A) *The theory \mathcal{T}_d is BDD.*
 (B) *$G^{2^n}(x, y) \in \text{rew}_{\mathcal{T}_d}(\phi_R^n(x, y))$ holds for every $n \in \mathbb{N}$.*

Let us first prove claim (B) of the theorem, which implies that \mathcal{T}_d is not distancing. The claim follows once we notice that: (i) $\text{Ch}(\mathcal{T}_d, \mathbb{G}^{2^n}(a, b)) \models \phi_R^n(a, b)$ and (ii) if \mathbb{D} is a proper subset of $\mathbb{G}^{2^n}(a, b)$ then $\text{Ch}(\mathcal{T}_d, \mathbb{D}) \not\models \phi_R^n(a, b)$. Establishing (i) is immediate, as exemplified in Fig. 1 displaying the case $n = 3$. To show (ii), we note that if \mathbb{D} is a proper subset of $\mathbb{G}^{2^n}(a, b)$ then a and b are in two different connected components of \mathbb{D} and, since \mathcal{T}_d is connected, they are in two different connected components of $\text{Ch}(\mathcal{T}_d, \mathbb{D})$.




Figure 1: Fragment of $\text{Ch}(\mathcal{T}_d, \mathbb{G}^8(a_0, a_8))$ (print in colors!)

The proof of claim (A) is much harder (see Appendix B). It defines a rewriting procedure in the spirit of [4, 16], whose termination for any given query $\phi(y)$ is shown, via an invariant defined by a complicated multiset ordering.

As our final exercise illustrates, the reasons why \mathcal{T}_d is BDD are quite subtle indeed:

EXERCISE 41. *Show that without rule (loop), \mathcal{T}_d would not be BDD. Hint: Consider the CQ $\exists x, y R(x, y), G(x, y)$.*

A remark on Theorem 3. A folklore belief seems to be that the existence of BDD theories that enforce rewritings of unbounded width is a consequence of the fact that it is undecidable to check if a theory is BDD (see e.g. a recent stackexchange post

[17]). Our results call this belief into question, because being BDD is undecidable for theories with a binary signature, and yet such theories, if BDD, are local and thus admit rewritings of linear width.

Remark on distancing and linear width rewritings. One may ask whether distancing is the same as admitting rewritings of linear width. The answer is no. While Observation 37 shows one implication, the converse is not true. Consider theory \mathcal{T}_d from Definition 39, but with every predicate's arity increased by one, and the new variable r occurring in the last position in every atom. Then the new theory is distancing (unlike \mathcal{T}_d), but it still requires exponential size rewritings (and, of course, it will remain BDD).

12 CONCLUSIONS AND FUTURE WORK

Our major motivation to embark on this journey was the pending status of the FUS/FES conjecture. On our way, we realized that any progress in that direction requires to significantly advance our understanding of the BDD class, separating folklore beliefs from hard facts. To this end, we introduced several new notions, characterizing specific properties of theories, and investigated their correspondencies. Most notably, we defined *local theories*, a BDD subclass. Our major results are the following:

- We show that the FUS/FES conjecture holds for all local theories (Theorem 2), which include all theories over binary signatures (Corollary 1). If the conjecture holds in the general case, then our work may provide the basis for a complete proof. If it does not, we now know that we must look for counter-examples of higher arity to disprove it.
- We show that there are BDD theories that are non-distancing and even necessitate rewritings of exponential width (Theorem 3). This result highlights the limitations of existing BDD classes [7, 9, 10], which can only characterise rule sets that admit rewritings of polynomial width.

As for future work, we intend to explore the following:

- Study the relation between distancing and bd-local. More precisely, find out if there are theories that are BDD and bd-local but are not distancing.
- Extend the proof of Theorem 1 to show Conjecture 2, i.e., that all BDD frontier-guarded theories [2] are local and thus the FUS/FES conjecture holds for them. Also, show if the FUS/FES conjecture holds for bd-local theories and then, of course, try to show the conjecture in the general case!
- Define a class of BDD theories that contains rule sets such as the one from Definition 39. Also, define an expressive class that captures the intuitive notion of locality, contains all known BDD classes, and implies BDD membership.
- Even though we extend Theorem 3 in the appendix (see Lemma 40), we wonder if there is a theory that does not admit an elementary bound on the width of its rewritings.

Acknowledgements. Ostropolski-Nalewaja and Marcinkowski were supported by the Polish National Science Centre (NCN) grant 2016/23/B/ST6/01438. Carral was supported by the ANR project CQFD (ANR-18-CE23-0003), the DFG project 389792660 (TRR 248, Center for Perspicuous Systems) and by the BMBF in the Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI). Rudolph was supported by the European Research Council through the ERC Consolidator Grant DeciGUT (project number 771779).

REFERENCES

- [1] A. Deutsch, A. Nash, and J. B. Remmel, “The chase revisited,” in *Proc. of the 27th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2008*, (M. Lenzerini and D. Lembo, eds.), pp. 149–158, ACM, 2008.
- [2] J. Baget, M. Leclère, M. Mugnier, and E. Salvat, “On rules with existential variables: Walking the decidability line,” *Artif. Intell.*, vol. 175, no. 9-10, pp. 1620–1654, 2011.
- [3] A. Cali, G. Gottlob, and T. Lukasiewicz, “Datalog[±]: a unified approach to ontologies and integrity constraints,” in *Proc. of the 12th International Conference on Database Theory, ICDT 2009*, (R. Fagin, ed.), vol. 361 of *ACM International Conference Proceeding Series*, pp. 14–30, ACM, 2009.
- [4] M. König, M. Leclère, M. Mugnier, and M. Thomazo, “Sound, complete and minimal UCQ-rewriting for existential rules,” *Semantic Web*, vol. 6, no. 5, pp. 451–475, 2015.
- [5] P. Barceló, G. Berger, C. Lutz, and A. Pieris, “First-order rewritability of frontier-guarded ontology-mediated queries,” in *Proc. of the 27th International Joint Conference on Artificial Intelligence, IJCAI 2018*, pp. 1707–1713, 2018.
- [6] C. Civili and R. Rosati, “On the first-order rewritability of conjunctive queries over binary guarded existential rules,” in *Proc. of the 30th Italian Conference on Computational Logic*, (D. Ancona, M. Maratea, and V. Mascardi, eds.), vol. 1459 of *CEUR Workshop Proceedings*, pp. 25–30, CEUR-WS.org, 2015.
- [7] A. Cali, G. Gottlob, and A. Pieris, “Advanced processing for ontological queries,” *Proc. VLDB Endow.*, vol. 3, no. 1, pp. 554–565, 2010.
- [8] H. Gaifman, H. G. Mairson, Y. Sagiv, and M. Y. Vardi, “Undecidable optimization problems for database logic programs,” *J. ACM*, vol. 40, no. 3, pp. 683–713, 1993.
- [9] A. Artale, D. Calvanese, R. Kontchakov, and M. Zakharyashev, “The DL-Lite family and relations,” *J. Artif. Intell. Res.*, vol. 36, pp. 1–69, 2009.
- [10] M. Thomazo, *Conjunctive Query Answering Under Existential Rules - Decidability, Complexity, and Algorithms*. PhD thesis, Montpellier 2 University, France, 2013.
- [11] S. Delivrias, M. Leclère, M. Mugnier, and F. Ulliana, “On the k-boundedness for existential rules,” in *Proc. of 2nd International Joint Conference on Rules and Reasoning, RuleML+RR 2018*, (C. Benz Müller, F. Ricca, X. Parent, and D. Roman, eds.), vol. 11092 of *Lecture Notes in Computer Science*, pp. 48–64, Springer, 2018.
- [12] P. Bourhis, M. Leclère, M. Mugnier, S. Tison, F. Ulliana, and L. Gallois, “Oblivious and semi-oblivious boundedness for existential rules,” in *Proc. of the 28th International Joint Conference on Artificial Intelligence, IJCAI 2019*, (S. Kraus, ed.), pp. 1581–1587, ijcai.org, 2019.
- [13] P. Hell and J. Nešetřil, “The core of a graph,” *Discret. Math.*, vol. 109, no. 1-3, pp. 117–126, 1992.
- [14] A. Cali, G. Gottlob, and T. Lukasiewicz, “A general datalog-based framework for tractable query answering over ontologies,” *J. Web Semant.*, vol. 14, pp. 57–83, 2012.
- [15] J. Baget, F. Garreau, M. Mugnier, and S. Rocher, “Revisiting chase termination for existential rules and their extension to nonmonotonic negation,” *CoRR*, vol. abs/1405.1071, 2014.
- [16] G. Gottlob, G. Orsi, and A. Pieris, “Query rewriting and optimization for ontological databases,” vol. 39, Oct. 2014.
- [17] M. Krötzsch, “What is first-order rewritable (and fo-query)?” <https://cstheory.stackexchange.com/questions/4859/what-is-first-order-rewritable-and-fo-query>. Accessed: 2020-11-27.
- [18] J. Baget, M. Leclère, M. Mugnier, and E. Salvat, “Extending decidable cases for rules with existential variables,” in *Proc. of the 21st International Joint Conference on Artificial Intelligence, IJCAI 2009*, (C. Boutilier, ed.), pp. 677–682, 2009.
- [19] N. Dershowitz and Z. Manna, “Proving termination with multiset orderings,” *Commun. ACM*, vol. 22, no. 8, pp. 465–476, 1979.

Appendices

Appendix A PROOF OF THEOREM 1

Let us now fix a binary signature Σ and a BDD theory \mathcal{T} over Σ . By \mathcal{T}_{DL} we will denote the Datalog rules of \mathcal{T} and by \mathcal{T}_{\exists} its proper existential rules.

First we will need to distinguish, among all elements of \mathcal{T}_{\exists} , *detached*⁴ rules that are of the form $\phi(\bar{x}) \rightarrow \exists y, z \psi(y, z)$ or $\phi(\bar{x}) \rightarrow \exists y \psi(y)$, that is, rules having empty frontier. Note that when firing a detached rule, the newly created atom, has no common terms with the rest of the chase. Notice that, since we only consider binary schemas, the non-empty frontier of an existential rule always consists of exactly one variable⁵. Rules from \mathcal{T}_{\exists} which are not detached will be called *sensible*

Clearly, whatever fact set \mathbb{D} we consider, the structure $Ch(\mathcal{T}, \mathbb{D})$ is a disjoint union of three sets of atoms. One set consists of the original facts from \mathbb{D} . Second contains *existential atoms*, that is facts created in the process of the chase, by rules of \mathcal{T}_{\exists} . The third set consists of atoms that are created by the rules of \mathcal{T}_{DL} which are called *Datalog atoms*. We will denote the set of existential atoms of $Ch(\mathcal{T}, \mathbb{D})$ together with atoms of \mathbb{D} with $Ch_{\exists}(\mathcal{T}, \mathbb{D})$.

Let us now concentrate on the structure of $Ch_{\exists}(\mathcal{T}, \mathbb{D})$. Notice that there are again two kinds of atoms there: *detached* atoms, created by detached rules and *sensible* atoms, created by sensible rules. Notice also that our taxonomy of atoms implies a taxonomy of the terms of $Ch_{\exists}(\mathcal{T}, \mathbb{D})$ (that is the elements of $dom(Ch(\mathcal{T}, \mathbb{D})) \setminus dom(\mathbb{D})$): there are sensible terms, created by sensible rules and detached terms, created by detached rules. The set of detached terms will be called $det(Ch(\mathcal{T}, \mathbb{D}))$.

OBSERVATION 42. *The graph whose vertices are the terms from $dom(Ch(\mathcal{T}, \mathbb{D}))$ and whose edges are sensible atoms of $Ch_{\exists}(\mathcal{T}, \mathbb{D})$ is a forest. The set of the roots of the trees of this forest is equal to $dom(\mathbb{D}) \cup det(Ch(\mathcal{T}, \mathbb{D}))$. The number of children of any of the vertices of this forest is bounded by the number of existential rules in \mathcal{T} .*

For any $a \in dom(\mathbb{D}) \cup det(Ch(\mathcal{T}, \mathbb{D}))$ let $\mathbb{S}(a)$ be the set of all atoms of $Ch_{\exists}(\mathcal{T}, \mathbb{D})$ which are edges of the tree rooted in a . Following the naming convention, we call trees that are rooted in detached terms *detached trees*.

A.1 First (failed) attempt at the Crucial Lemma

For any given set of facts \mathbb{D} let a *parent function* $par_{\mathcal{T}}$ be any function from $Ch(\mathcal{T}, \mathbb{D}) \setminus \mathbb{D}$ to the power set of $Ch(\mathcal{T}, \mathbb{D})$ such that for any atom $\alpha \in Ch(\mathcal{T}, \mathbb{D})$ there exists a rule ρ and a mapping σ satisfying:

- $\alpha = appl(\rho, \sigma)$,
- $\sigma(body(\rho)) = par_{\mathcal{T}}(\alpha)$.

The parent function points to some (arbitrarily chosen) set of atoms that leads to the creation of α . Note that there may be more

⁴Such rules are called *disconnected* in [18], however we think that calling those rules *detached* might help the reader to distinguish those from rules that have disconnected bodies.

⁵Actually, the assumption we really use in the proof of Theorem 1 is not that relations are at most binary, but that the existential rules are “frontier one”.

than one such function as α could be created in more than one way during the chase.

Let \mathbb{D} be any set of facts and let $par_{\mathcal{T}}$ be some parent function. Then we define an *ancestor function* $anc_{\mathcal{T}}$ as follows:

- $anc_{\mathcal{T}}(\alpha) = \{\alpha\}$ for an atom $\alpha \in \mathbb{D}$,
- $anc_{\mathcal{T}}(\alpha) = \bigcup_{\alpha' \in par_{\mathcal{T}}(\alpha)} anc_{\mathcal{T}}(\alpha')$ for atoms of $Ch(\mathcal{T}, \mathbb{D}) \setminus \mathbb{D}$.

Intuitively the set $anc_{\mathcal{T}}(\alpha)$ consists of facts from \mathbb{D} which were used during the chase, to prove α . Of course there might be more than one ancestor function for any given set of facts as that function is strictly associated with a particular parent function. This freedom in taking parents, and so in picking ancestors, leads to some problems, as we will soon discover.

We would be one step from proving Theorem 1 if we had:

LEMMA 43 (CRUCIAL LEMMA, FIRST ATTEMPT, FALSE). *There is a natural number M , depending on the rule set \mathcal{T} but not on the fact set \mathbb{D} , such that for every ancestor function $anc_{\mathcal{T}}$, for each constant and any detached term t in $Ch(\mathcal{T}, \mathbb{D})$ holds:*

$$\left| \bigcup_{\alpha \in \mathbb{S}(t)} anc_{\mathcal{T}}(\alpha) \right| \leq M$$

But that lemma is unfortunately not true. For a counterexample see:

EXAMPLE 44. *Let \mathcal{T} consist of two rules:*

- $E(x, y), R(z, y) \rightarrow \exists v E(y, v)$
- $E(x, y), P(z) \rightarrow R(z, y)$

Suppose that M as in the lemma exists and that \mathbb{D} consists of atom $E(a_0, a_1)$ and atoms $P(b_i)$ for each $1 \leq i \leq M$.

Then $Ch(\mathcal{T}, \mathbb{D})$ will consist of an infinite number of new facts $E(a_1, a_2), E(a_2, a_3), E(a_3, a_4) \dots$. In order to build them, however, some facts about relation R will need to be proven by the second rule using a number of P atoms from \mathbb{D} . And it might happen (Skolem/semi-oblivious chase is non-deterministic in that aspect) that the proven R -facts will be $R(b_1, a_1), R(b_2, a_2), \dots, R(b_M, a_M)$, meaning that $\mathbb{S}(a_1)$ uses during its creation all the $M + 1$ facts of \mathbb{D} . But of course this is for irrelevant reasons: the same chase could be built if $P(b_1)$ was used each time some P was needed.

A.2 The normalization of \mathcal{T}

In order to circumvent the problems highlighted by Example 44, we will now transform the rule set \mathcal{T} into another rule set \mathcal{T}_{NF} which, apart from some other useful properties, will satisfy the equality (*) $Ch_{\exists}(\mathcal{T}_{NF}, \mathbb{D}) = Ch_{\exists}(\mathcal{T}, \mathbb{D})$.

First we will define the signature of \mathcal{T}_{NF} . Let us take a fresh set of nullary predicates $\mathbf{M} = \{M_{\phi} \mid \phi \text{ is a boolean CQ over } \Sigma\}$. Then our new signature⁶ Σ' is $\Sigma \cup \mathbf{M}$.

Two procedures will be used during the normalization: *body rewriting* and *body separation*.

DEFINITION 45 (BODY REWRITING). *Let ρ be some rule with body $\beta(\bar{x}, \bar{y})$ over Σ and a head $\gamma(\bar{y})$ which is an atom from Σ' possibly preceded with the existential quantifier (or two). Then by $Rew(\rho)$ we denote the set:*

$$\{\beta'(\bar{x}, \bar{y}) \Rightarrow \gamma(\bar{y}) : \beta'(\bar{x}, \bar{y}) \in rew_{\mathcal{T}}(\exists \bar{x} \beta(\bar{x}, \bar{y}))\}$$

⁶We briefly forget here about our promise that signatures would be finite.

While body rewriting can be applied both to existential rules and Datalog rules, the second procedure will only be applied to existential rules. It separates the disconnected fragment of body of given existential rule and “encapsulates” that fragment in a single nullary predicate from \mathbf{M} . That is, given a rule, the procedure returns a pair of rules. One being almost the original rule, but with its body changed to consist of nullary predicate and a connected conjunction of atoms. The second being a rule whose job is to prove the aforementioned nullary predicate.

DEFINITION 46 (BODY SEPARATION). *If ρ is an existential rule of the form $\beta(\bar{x}, \bar{y}) \wedge \phi(\bar{z}) \Rightarrow \exists \bar{u} \gamma(\bar{y}, \bar{u})$ such that:*

- $(\bar{x} \cup \bar{y}) \cap \bar{z} = \emptyset$,
- $\beta(\bar{x}, \bar{y})$ is connected,

then

- $sep_{cc}(\rho) = \beta(\bar{x}, \bar{y}) \wedge M_{\phi} \Rightarrow \exists \bar{u} \gamma(\bar{y}, \bar{u})$
- $sep_M(\rho) = \phi(\bar{z}) \Rightarrow M_{\phi}$

If the body of ρ is connected, we assume that $\phi(\bar{z})$ is empty. And we have a nullary predicate $M_{\emptyset} \in \Sigma'$ for this occasion.

The normalization algorithm is performed in three steps:

NORMALIZATION ALGORITHM

STEP ONE: $\mathcal{T}_I = \bigcup_{\rho \in \mathcal{T}_{\exists}} Rew(\rho)$

STEP TWO: $\mathcal{T}_{II} = \{sep_{cc}(\rho) : \rho \in \mathcal{T}_I\}$

STEP THREE: $\mathcal{T}_{III} = \bigcup_{\rho \in \mathcal{T}_I} Rew(sep_M(\rho))$

RETURN: $\mathcal{T}_{NF} = \mathcal{T}_{II} \cup \mathcal{T}_{III}$

The normalization allows us to attack the source of the problem highlighted in the previous section by separating the “disconnected ancestors” required by existential rules and encapsulating those ancestors within rules producing nullary predicates.

OBSERVATION 47. *Let ρ be a rule that creates a detached atom in $Ch(\mathcal{T}_{NF}, \mathbb{D})$ for some fact set \mathbb{D} . Then ρ is a rule from \mathcal{T}_{II} and its body consists of a single nullary atom.*

LEMMA 48. *For any set of facts \mathbb{D} over Σ :*

$$Ch_{\exists}(\mathcal{T}, \mathbb{D}) = Ch_{\exists}(\mathcal{T}_{NF}, \mathbb{D})$$

Here, our choice of the Skolem naming convention can be explained. We want to show equality of the structures in Lemma 48 and so if we choose our Skolem naming convention to be dependent on rule names then Lemma 48 would obviously be false.

A.3 Proof of Lemma 48

This entire subsection is devoted to the proof of Lemma 48, which will follow directly from Lemma 52 and Lemma 55. But first, as a warm-up, notice that:

EXERCISE 49. *Let \mathbb{D} be any instance and let $\bar{a} \in dom(Ch(\mathcal{T}, \mathbb{D}))^{|\bar{y}|}$. Let $\phi(\bar{y}) \in rew(\psi(\bar{y}))$ (where $\psi(\bar{y})$ is some CQ) and suppose $Ch(\mathcal{T}, \mathbb{D}) \models \phi(\bar{a})$. Then $Ch(\mathcal{T}, \mathbb{D}) \models \psi(\bar{a})$.*

Hint: Recall that $Ch(\mathcal{T}, Ch(\mathcal{T}, \mathbb{D})) = Ch(\mathcal{T}, \mathbb{D})$.

EXERCISE 50. Let $\psi(\bar{y})$ be any CQ and $\phi(\bar{y}) \in \text{rew}(\psi(\bar{y}))$. Let \mathbb{D} be any instance and $\bar{a} \in \text{dom}(\mathbb{D})^{|\bar{y}|}$. Suppose $Ch(\mathcal{T}, \mathbb{D}) \models \phi(\bar{a})$. Then there exists query $\phi'(\bar{y}) \in \text{rew}(\psi(\bar{y}))$ such that $\mathbb{D} \models \phi'(\bar{a})$.

Hint: Notice that $Ch(\mathcal{T}, Ch(\mathcal{T}, \mathbb{D})) = Ch(\mathcal{T}, \mathbb{D})$.

EXERCISE 51. Suppose $(\exists \bar{t}, \bar{u} \phi(\bar{t}, \bar{y}) \wedge \gamma(\bar{u})) \in \text{rew}_{\mathcal{T}}(\exists \bar{x} \beta(\bar{x}, \bar{y}))$ for some CQs ϕ, γ and β , where the tuples of variables $\bar{t} \cup \bar{y}$ and \bar{u} are disjoint. And suppose $\exists \bar{v} \zeta(\bar{v}) \in \text{rew}_{\mathcal{T}}(\exists \bar{u} \gamma(\bar{u}))$. Then there exists a query $\exists \bar{z} \varrho(\bar{z}, \bar{y}) \in \text{rew}_{\mathcal{T}}(\exists \bar{x} \beta(\bar{x}, \bar{y}))$ which is contained in $(\exists \bar{t}, \bar{v} \phi(\bar{t}, \bar{y}) \wedge \zeta(\bar{v}))$.

Hint: Let database \mathbb{D} be a frozen body of $\exists \bar{t}, \bar{v} \phi(\bar{t}, \bar{y}) \wedge \zeta(\bar{v})$ and observe that $\mathbb{D} \models \text{rew}_{\mathcal{T}}(\exists \bar{x} \beta(\bar{x}, \bar{y}))$

Notice that the only existential rules in \mathcal{T}_{NF} are the ones in \mathcal{T}_{II} . Notice also that the only Datalog rules in \mathcal{T}_{NF} are the ones in \mathcal{T}_{III} and hence the only atoms which are in $Ch(\mathcal{T}_{NF}, \mathbb{D})$ but not in $Ch_{\exists}(\mathcal{T}_{NF}, \mathbb{D})$ are the nullary atoms from the set \mathbf{M} .

By $Ch_{i,\exists}(\mathcal{T}, \mathbb{D})$ we denote the intersection of $Ch_i(\mathcal{T}, \mathbb{D})$ and $Ch_{\exists}(\mathcal{T}, \mathbb{D})$.

LEMMA 52. For each fact set \mathbb{D} over Σ and each $k \in \mathbb{N}$:

$$Ch_{k,\exists}(\mathcal{T}_{NF}, \mathbb{D}) \subseteq Ch_{\exists}(\mathcal{T}, \mathbb{D}).$$

PROOF. We will show this by induction on k . The case for $k = 0$ is trivial. Assume the claim is true for some $k \in \mathbb{N}$. We are going to show that it is also true for $k + 1$.

Let α be an atom from $Ch_{k+1,\exists}(\mathcal{T}_{NF}, \mathbb{D})$, which is not in $Ch_{k,\exists}(\mathcal{T}_{NF}, \mathbb{D})$. Then $\alpha = \text{appl}(\rho_{\alpha}, \sigma_{\alpha})$ for some:

- $\rho_{\alpha} \in \mathcal{T}_{II}$, the rule which actually created α , of the form $\phi(\bar{t}, \bar{y}) \wedge M_{\gamma} \Rightarrow \exists \bar{z} \alpha_0(\bar{y}, \bar{z})$,
- σ_{α} such that $\sigma_{\alpha}(\phi(\bar{t}, \bar{y})) \subseteq Ch_{k,\exists}(\mathcal{T}_{NF}, \mathbb{D})$,
- M_{γ} such that $Ch_{k,\exists}(\mathcal{T}_{NF}, \mathbb{D}) \models M_{\gamma}$.

It follows from the construction of \mathcal{T}_{NF} that there must exist:

- (1) A rule $\rho' \in \mathcal{T}_I$ of the form:

$$\phi(\bar{t}, \bar{y}) \wedge \gamma(\bar{u}) \Rightarrow \exists \bar{z} \alpha_0(\bar{y}, \bar{z})$$

which, by Step II of the normalization algorithm, led to the creation of ρ_{α} , such that the tuple \bar{u} of variables is disjoint with \bar{y} and with \bar{t} , and that $\phi(\bar{t}, \bar{y})$ is a connected query.

- (2) A rule $\rho \in \mathcal{T}_{\exists}$ of the form:

$$\beta(\bar{x}, \bar{y}) \Rightarrow \exists \bar{z} \alpha_0(\bar{y}, \bar{z})$$

which, by Step I of the normalization algorithm, led to the creation of ρ' , such that:

$$(\exists \bar{t}, \bar{u} \phi(\bar{t}, \bar{y}) \wedge \gamma(\bar{u})) \in \text{rew}_{\mathcal{T}}(\exists \bar{x} \beta(\bar{x}, \bar{y})) \quad (\spadesuit)$$

- (3) A rule $\rho_{M_{\gamma}}$ in \mathcal{T}_{III} of the form $\zeta(\bar{u}) \Rightarrow M_{\gamma}$ such that $\exists \bar{u} \zeta(\bar{u}) \in \text{rew}_{\mathcal{T}}(\exists \bar{u} \gamma(\bar{u}))$ and a substitution σ_{ζ} from the variables in \bar{u} to $\text{dom}(Ch_{k-1,\exists}(\mathcal{T}_{NF}, \mathbb{D}))$ such that $\sigma_{\zeta}(\zeta(\bar{u})) \subseteq Ch_{k-1,\exists}(\mathcal{T}_{NF}, \mathbb{D})$. This is because, for ρ_{α} to be applicable in $Ch_{k,\exists}(\mathcal{T}_{NF}, \mathbb{D})$, there must be $Ch_{k,\exists}(\mathcal{T}_{NF}, \mathbb{D}) \models M_{\gamma}$, and since $\mathbb{D} \not\models M_{\gamma}$, a rule able to produce M_{γ} must have earlier been applied.

In order to complete the induction step, we need to show that $\alpha \in Ch_{\exists}(\mathcal{T}, \mathbb{D})$. Since ρ is a rule of \mathcal{T} this claim will follow once we can prove that:

$$Ch_{\exists}(\mathcal{T}, \mathbb{D}) \models \exists \bar{x} \beta(\bar{x}, \sigma_{\alpha}(\bar{y})) \quad (\diamond)$$

And (\diamond) will follow (using Exercise 49) once we can show that for some query $\varrho(\bar{y}) \in \text{rew}_{\mathcal{T}}(\exists \bar{x} \beta(\bar{x}, \bar{y}))$ there is:

$$Ch_{\exists}(\mathcal{T}, \mathbb{D}) \models \varrho(\sigma_{\alpha}(\bar{y})) \quad (\clubsuit)$$

Recall that from (\heartsuit) we know that:

$$Ch_{k,\exists}(\mathcal{T}_{NF}, \mathbb{D}) \models \exists \bar{t} \phi(\bar{t}, \sigma_{\alpha}(\bar{y})).$$

By hypothesis this implies that:

$$Ch_{\exists}(\mathcal{T}, \mathbb{D}) \models \exists \bar{t} \phi(\bar{t}, \sigma_{\alpha}(\bar{y})).$$

The claim (\clubsuit) would be now proven, using (\spadesuit) , if we could also show that:

$$Ch_{\exists}(\mathcal{T}, \mathbb{D}) \models \exists \bar{u} \gamma(\bar{u}),$$

But this may not be the case. All we know is that:

$$Ch_{k,\exists}(\mathcal{T}_{NF}, \mathbb{D}) \models \exists \bar{u} \zeta(\bar{u})$$

and therefore (using the hypothesis again) that:

$$Ch_{\exists}(\mathcal{T}, \mathbb{D}) \models \exists \bar{u} \zeta(\bar{u}).$$

So we get that $Ch_{\exists}(\mathcal{T}, \mathbb{D}) \models \exists \bar{s}, \bar{t} \phi(\bar{t}, \sigma_{\alpha}(\bar{y})) \wedge \zeta(\bar{s})$. Now use Exercise 51 to get (\clubsuit) . \square

The proof of the following lemma seems overly complicated. Why doesn't it just follow from Exercise 50? This is because, while we assume that \mathcal{T} is BDD, we never proved that \mathcal{T}_{NF} is BDD too.

LEMMA 53. Suppose $Ch(\mathcal{T}_{NF}, \mathbb{D}) \models M_{\phi}$ for some nullary fact M_{ϕ} . Then $Ch_1(\mathcal{T}_{NF}, \mathbb{D}) \models M_{\phi}$.

PROOF. Suppose $Ch(\mathcal{T}_{NF}, \mathbb{D}) \models M_{\phi}$ and let $\beta(\bar{x})$ be the body of the rule from \mathcal{T}_{III} that created M_{ϕ} .

Recall that $\beta(\bar{x})$ does not contain nullary predicates and that \mathcal{T}_{NF} consists only of existential rules (from \mathcal{T}_{II}) and of Datalog rules that prove nullary facts (from \mathcal{T}_{III}). This means that $Ch_{\exists}(\mathcal{T}_{NF}, \mathbb{D}) \models \exists \bar{x} \beta(\bar{x})$.

Now we can use Lemma 52 and get: $Ch_{\exists}(\mathcal{T}, \mathbb{D}) \models \exists \bar{x} \beta(\bar{x})$.

Now recall that $\exists \bar{x} \beta(\bar{x}) \in \text{rew}_{\mathcal{T}}(\phi)$. From Exercise 50 we get that there exists another boolean query β_0 such that $\beta_0 \in \text{rew}_{\mathcal{T}}(\phi)$ and that $\mathbb{D} \models \beta_0$. And from this we can immediately conclude that $Ch_1(\mathcal{T}_{NF}, \mathbb{D}) \models M_{\phi}$. \square

Combining Lemma 53 and Observation 47 we get:

COROLLARY 54. If $Ch(\mathcal{T}_{NF}, \mathbb{D}) \models \alpha$ for some detached atom α then $Ch_2(\mathcal{T}_{NF}, \mathbb{D}) \models \alpha$.

The last lemma we need for the proof of Lemma 48 is:

LEMMA 55. $Ch_{i,\exists}(\mathcal{T}, \mathbb{D}) \subseteq Ch_{i+2,\exists}(\mathcal{T}_{NF}, \mathbb{D})$

PROOF. The induction base $\mathbb{D} \subseteq Ch_{2,\exists}(\mathcal{T}_{NF}, \mathbb{D})$ is clearly true. And so let us move to the induction step.

Let α be an atom produced by an existential rule during the $i + 1$ -th step of $Ch(\mathcal{T}, \mathbb{D})$. Let a rule $\rho \in \mathcal{T}$ and a mapping σ be such that $\alpha = \text{appl}(\rho, \sigma)$ and let $\gamma(\bar{x}, \bar{y})$ be a body of ρ where \bar{y} are the frontier variables.

While we know that $Ch_i(\mathcal{T}, \mathbb{D}) \models \sigma(\gamma(\bar{x}, \bar{y}))$ we are not sure if $Ch_{i,\exists}(\mathcal{T}, \mathbb{D}) \models \sigma(\gamma(\bar{x}, \bar{y}))$: some of the atoms in $\sigma(\gamma(\bar{x}, \bar{y}))$ could be produced by Datalog rules of \mathcal{T} .

But $Ch(\mathcal{T}, \mathbb{D}) = Ch(\mathcal{T}, Ch_{i,\exists}(\mathcal{T}, \mathbb{D}))$, so if $Ch_i(\mathcal{T}, \mathbb{D}) \models \sigma(\gamma(\bar{x}, \bar{y}))$ then we can be sure that there exists a query $\exists \bar{z} \beta(\bar{z}, \bar{y}) \in \text{rew}_{\mathcal{T}}(\exists \bar{x} \gamma(\bar{x}, \bar{y}))$ such that $Ch_{i,\exists}(\mathcal{T}, \mathbb{D}) \models \exists \bar{z} \beta(\bar{z}, \bar{y})$ or, in other words, $Ch_{i,\exists}(\mathcal{T}, \mathbb{D}) \models \beta(\sigma_{\beta}(\bar{z}), \bar{y})$ for some substitution σ_{β} .

Observe that from the induction hypothesis we know that $Ch_{i+2,\exists}(\mathcal{T}_{NF}, \mathbb{D}) \models \beta(\sigma_{\beta}(\bar{z}), \bar{y})$.

Let now ρ' be the rule from \mathcal{T}_{NF} (or, to be more precise, from \mathcal{T}_{II}), such that $\text{head}(\rho') = \text{head}(\rho)$ and that the body of ρ' is of the form $\beta_0(\bar{u}, \bar{y}) \wedge M_{\phi}$, where $\beta(\bar{z}, \bar{y}) = \beta_0(\bar{u}, \bar{y}) \wedge \phi(\bar{v})$ for some \bar{v} disjoint from $\bar{u} \cup \bar{y}$. It follows from the construction of \mathcal{T}_{NF} that such a ρ' exists.

Let us consider the second case. Note that $Ch_{i+2,\exists}(\mathcal{T}_{NF}, \mathbb{D}) \models \sigma_{\beta}(\beta'(\bar{x}, \bar{y}))$ and so all we need to show in order to complete our induction step is that $Ch_{i+2}(\mathcal{T}_{NF}, \mathbb{D}) \models M_{\phi}$.

Clearly $Ch_{i+2,\exists}(\mathcal{T}_{NF}, \mathbb{D}) \models \exists \bar{z} \phi(\bar{z})$ as we know that $\phi \subseteq \beta$. From this we can conclude that $\mathbb{D} \models \text{rew}_{\mathcal{T}}(\phi)$ and thus $M_{\phi} \in Ch_{2,\exists}(\mathcal{T}_{NF}, \mathbb{D})$.

This ends the proof of Lemma 55 and therefore of Lemma 48. \square

Note that this implies that the sensible and non-nullary atoms of $Ch(\mathcal{T}_{NF}, \mathbb{D})$ form a set of trees having a tree $\mathbb{S}(t)$ for each term t that is a constant of \mathbb{D} or detached term of $Ch_{\exists}(\mathcal{T}_{NF}, \mathbb{D})$. And that this set of trees is exactly this same set of trees as in the case of rule set \mathcal{T} . Also, note that having Lemma 48 we get a very important:

COROLLARY 56. *For any set of facts \mathbb{D} :*

$$Ch(\mathcal{T}_{DL}, Ch_{\exists}(\mathcal{T}_{NF}, \mathbb{D}) \cup \mathbb{D}) = Ch(\mathcal{T}, \mathbb{D})$$

A.4 The Crucial Lemma

In this section, we state and prove the Crucial Lemma. Recall, that we want to prove, that for any term t being a constant of \mathbb{D} or a detached term of $Ch_{\exists}(\mathcal{T}_{NF}, \mathbb{D})$, the tree $\mathbb{S}(t)$ rooted in t and consisting of sensible atoms of $Ch_{\exists}(\mathcal{T}_{NF}, \mathbb{D})$ requires only a small subset of \mathbb{D} to be built by the chase $Ch(\mathcal{T}_{NF}, \mathbb{D})$.

First let us distinguish, among the parents and ancestors of some atom in $Ch(\mathcal{T}_{NF}, \mathbb{D})$, its *connected parents* and *connected ancestors*: given some parent function $\text{par}_{\mathcal{T}_{NF}}$ for $Ch(\mathcal{T}_{NF}, \mathbb{D})$ and an atom α of $Ch(\mathcal{T}_{NF}, \mathbb{D})$ we define the set of connected parents $\text{cpar}(\alpha)$ as the set of all the non-nullary atoms of $\text{par}_{\mathcal{T}_{NF}}(\alpha)$. Then for atoms of $Ch(\mathcal{T}_{NF}, \mathbb{D})$ we define their respective sets of connected ancestors as follows:

- $\text{canc}(\alpha) = \{\alpha\}$ for an atom $\alpha \in \mathbb{D}$,
- $\text{canc}(\alpha) = \bigcup_{\alpha' \in \text{cpar}(\alpha)} \text{canc}(\alpha')$, otherwise.

Before proceeding, let us define a number of measures, depending on \mathcal{T}_{NF} .

- Let k be the number of nullary predicates in \mathcal{T}_{NF} .
- Let h be the maximal number of atoms in the body of a rule from \mathcal{T}_{NF} .
- Let n be the number of rules in \mathcal{T}_{NF} .
- Let N be the number of elements of a full n -ary tree of depth h .

- Let $M = Nh + kh$

Now having prepared everything we are going to prove the Crucial Lemma:

LEMMA 57 (CRUCIAL). *For every set of facts \mathbb{D} , every term t that is constant of \mathbb{D} or detached term of $Ch_{\exists}(\mathcal{T}_{NF}, \mathbb{D})$ and every ancestor function $\text{anc}_{\mathcal{T}_{NF}}$:*

$$\left| \bigcup_{\alpha \in \mathbb{S}(t)} \text{anc}_{\mathcal{T}_{NF}}(\alpha) \right| \leq M$$

Notice that we now count ancestors with respect to \mathcal{T}_{NF} instead of ancestors with respect to \mathcal{T} . This is, in fact, the only difference between this lemma and the (false) Lemma 43.

PROOF. The atoms of $\mathbb{S}(t)$ are all produced by the rules of \mathcal{T}_{II} . Recall that bodies of those rules consist of one connected CQ and one nullary atom. The number of atoms that are ancestors of nullary atoms can easily be bounded by kh using Lemma 53.

What is left to be bound is the number of connected ancestors of $\mathbb{S}(t)$. Recall that for $t \in \text{dom}(\mathbb{D}) \cup \text{det}(Ch(\mathcal{T}_{NF}, \mathbb{D}))$ the set of facts $\mathbb{S}(t)$ is a tree with t as its root. This gives us a natural notion of depth of atoms in $\mathbb{S}(t)$, with atoms containing t being at depth one.

First we will consider the easier case when $t \in \text{det}(Ch(\mathcal{T}, \mathbb{D}))$. Then $\mathbb{S}(t)$ is a detached tree, not connected by atoms of $Ch(\mathcal{T}_{NF}, \mathbb{D})$ to \mathbb{D} . Moreover the connected parents and connected ancestors are defined in such a way that every atom is in one connected component of $Ch(\mathcal{T}_{NF}, \mathbb{D})$ with all its connected ancestors. So, no atom in $\mathbb{S}(t)$ has any connected ancestors, and thus the entire $\mathbb{S}(t)$ has in total at most kh ancestors.

Let us now move to the case when $t \in \text{dom}(\mathbb{D})$. Clearly, to be a connected ancestor of someone in $\mathbb{S}(t)$ an atom in \mathbb{D} must be a connected parent of someone in $\mathbb{S}(t)$:

$$\bigcup_{\alpha \in \mathbb{S}(t)} \text{canc}(\alpha) = \left(\bigcup_{\alpha \in \mathbb{S}(t)} \text{cpar}(\alpha) \right) \cap \mathbb{D}.$$

Using the above equality, one can bound the number of ancestors of atoms in the connected ancestors of someone in $\mathbb{S}(t)$ in the following way:

$$\left| \left(\bigcup_{\alpha \in \mathbb{S}(t)} \text{cpar}(\alpha) \right) \cap \mathbb{D} \right| \leq \sum_{\alpha \in \mathbb{S}(t)} |\text{cpar}(\alpha) \cap \mathbb{D}|.$$

Now, we can easily bound $|\text{cpar}(\alpha) \cap \mathbb{D}|$ with h for any atom α of $Ch(\mathcal{T}_{NF}, \mathbb{D})$. Finally we are going to show that only a finite (and bounded) number of atoms of $\mathbb{S}(t)$ have any connected parents in \mathbb{D} . Observe that, if an application of an existential rule ρ created an atom at depth greater than h then the connected part of the body ρ , of size at most h , could not “reach” the atoms of \mathbb{D} . From this we get Nh as bound for the number of connected ancestors of $\mathbb{S}(t)$ which implies that M is a bound on the number of ancestors of $\mathbb{S}(t)$ for any t . \square

A.5 Proving Theorem 1

Now we can finally prove Theorem 1. The following is an easy corollary of Lemma 57 and Corollary 56:

COROLLARY 58. *For any term $t \in \text{dom}(\mathbb{D}) \cup \text{det}(Ch(\mathcal{T}_{NF}, \mathbb{D}))$ there exists a subset \mathbb{D}' of \mathbb{D} of size at most M such that $\mathbb{S}(t) \subseteq Ch(\mathcal{T}, \mathbb{D}')$.*

This is almost Theorem 1, but only for existential atoms of $Ch(\mathcal{T}, \mathbb{D})$. However, observe that:

OBSERVATION 59. *There exists a constant $d_{\mathcal{T}}$ such that for any set of facts \mathbb{D} and for any atom α of $Ch(\mathcal{T}_{DL}, \mathbb{D})$ there exists a subset \mathbb{D}' of \mathbb{D} such that $|\mathbb{D}'| < d_{\mathcal{T}}$ and $\alpha \in Ch(\mathcal{T}, \mathbb{D}')$.*

PROOF. From Observation 14 the size of \mathbb{D}' can easily be bounded by $h^{n_{at}}$ where h is the maximal number of atoms in a rule of \mathcal{T} and n_{at} is the constant from Observation 14. \square

Thus $Md_{\mathcal{T}}$ is a locality constant for Datalog atoms of $Ch(\mathcal{T}, \mathbb{D})$. This concludes the proof of Theorem 1.

Appendix B PROOF OF LEMMA 40 (A). THE HIGH-LEVEL VIEW.

Since \mathcal{T}_d is fixed, we will use the notation $Ch(\mathbb{F})$ instead of $Ch(\mathcal{T}_d, \mathbb{F})$.

Clearly, to prove that \mathcal{T}_d is BDD, it is enough to consider only connected queries, since a bound for any non-connected query can be derived from bounds of its connected components. Additionally, notice that due to the rule (loop), if ϕ is a boolean query and if \mathbb{D} is any instance then $Ch_1(\mathbb{D}) \models \phi$. So for the rest of the proof of Lemma 40 (A) **we will consider only connected non-boolean queries**. Let us start with:

DEFINITION 60 (MARKED QUERY). *We define a marked query as a pair $\langle \phi(\bar{y}), V \rangle$ where $\phi(\bar{y}) = \exists \bar{x} \beta(\bar{x}, \bar{y})$ is a CQ and V is a subset of variables of $\phi(\bar{y})$ such that $\bar{y} \subseteq V \subseteq \bar{y} \cup \bar{x}$.*

We say that the variables in V are *marked*. It will often be more convenient to refer to a marked query without (or before) specifying its components. The letter Q will be used for that. If $Q = \langle \phi(\bar{y}), V \rangle$ then by $V(Q)$ we mean V , and $q(Q)$ denotes $\phi(\bar{y})$. By $var(Q)$ we mean the set of all variables of $q(Q)$, and Q_R (Q_G) is the set of red (green) atoms in $q(Q)$.

For a marked query Q the intention behind the set of marked variables $V(Q)$ is that the variables of $V(Q)$ should not be mapped onto the chase-produced Skolem terms:

DEFINITION 61. *For a marked query $Q = \langle \phi(\bar{y}), V \rangle$, a fact set \mathbb{D} , and $\bar{a} \in dom(\mathbb{D})^{|\bar{y}|}$ we say that $Ch(\mathbb{D}) \models Q(\bar{a})$ if there exists a homomorphism $h : var(Q) \rightarrow dom(Ch(\mathbb{D}))$ witnessing $Ch(\mathbb{D}) \models \phi(\bar{a})$ such that for every v there is $h(v) \in dom(\mathbb{D})$ if and only if $v \in V$.*

Some marked queries are false – it follows directly from Definition 39 that they cannot be satisfied in any $Ch(\mathbb{D})$:

OBSERVATION 62. *For any instance \mathbb{D}*

- if $R(a, b)$ or $G(a, b)$ is an edge in $Ch(\mathbb{D})$ and if $b \in dom(\mathbb{D})$ then $a \in dom(\mathbb{D})$;
- if $C = \{E_0(a_0, a_1), E_1(a_1, a_2) \dots E_k(a_k, a_0)\}$ is a cycle in $Ch(\mathbb{D})$ (where each $E_i \in \{R, G\}$) then $C \subseteq \mathbb{D}$;
- if $E(a_1, b)$ and $E(a_2, b)$ are two atoms in $Ch(\mathbb{D})$ (where $E \in \{R, G\}$) and $a_1 \in dom(\mathbb{D})$ then also $a_2 \in dom(\mathbb{D})$.

Then it follows from Observation 62 that:

OBSERVATION 63. *Suppose $Ch(\mathbb{D}) \models Q$. Then:*

- (i) if $R(z, z')$ or $G(z, z')$ is an atom in $q(Q)$ and $z' \in V(Q)$ then also $z \in V(Q)$;

- (ii) if $\{E_0(z_0, z_1), E_1(z_1, z_2) \dots E_k(z_k, z_0)\}$ is a cycle in $q(Q)$ (where each $E_i \in \{R, G\}$) then $z_i \in V(Q)$ for each $i \in \{0, \dots, k\}$;

- (iii) if $E(z_1, u)$ and $E(z_2, u)$ are two atoms in $q(V)$ (where $E \in \{R, G\}$) and $z_1 \in V(Q)$ then also $z_2 \in V(Q)$.

Marked queries satisfying conditions (i)-(iii) of Observation 63 will be called *properly marked*. Marked queries whose all variables are marked will be called *totally marked*. Properly marked queries that are not totally marked will be called *live*. Notice that for a totally marked query $\langle \phi(\bar{a}), V \rangle$, an instance \mathbb{D} and a tuple $\bar{a} \in dom(\mathbb{D})^{|\bar{y}|}$ there is $Ch(\mathbb{D}) \models \langle \phi(\bar{a}), V \rangle$ if and only if $\mathbb{D} \models \phi(\bar{a})$.

B.1 High-level proof of claim (A). The process.

Now **we get some conjunctive query** $\phi(\bar{y})$, that will be fixed till the end of this proof, and we want to show that there exists a rewriting of ϕ , i.e. a finite set $rew(\phi(\bar{y}))$ of CQs such that for each \mathbb{D} and $\bar{a} \in dom(\mathbb{D})^{|\bar{y}|}$ there is $Ch(\mathbb{D}) \models \phi(\bar{a})$ if and only if there exists a $\psi(\bar{y}) \in rew(\phi(\bar{y}))$ such that $\mathbb{D} \models \psi(\bar{a})$.

Notice that $rew(\phi(\bar{y}))$ will be constructed if we can produce a finite set \mathcal{S} of marked queries, such that:

- (♣) for each \mathbb{D} and each $\bar{a} \in dom(\mathbb{D})^{|\bar{y}|}$ we have $Ch(\mathbb{D}) \models \phi(\bar{a})$ if and only if $Ch(\mathbb{D}) \models Q(\bar{a})$ holds for some $Q \in \mathcal{S}$;
- (♣) there are no live queries in \mathcal{S} .

\mathcal{S} will be constructed as the result of some *process*. As the starting point of this process let $\mathcal{S}_0 = \{Q : q(Q) = \phi(\bar{y})\}$, the set of all possible markings of $\phi(\bar{y})$. One can easily see that \mathcal{S}_0 satisfies Condition (♣) above, but there is no reason to think that it also satisfies Condition (♣).

Now, the plan is as follows. *Five operations* are going to be defined in Appendix B.3, called *cut-red*, *cut-green*, *fuse-red*, *fuse-green*, and *reduce*. Each of them will:

- take, as an input, some marked query Q ;
- remove from $q(Q)$, one variable and some atoms (operation *reduce* will remove one red atom and one green) ;
- keep the marking of the surviving variables unchanged;
- in one case (of operation *reduce*) add two new variables, one red atom, and two green atoms);
- return the resulting marked query (except for operation *reduce* which will return four marked queries: one for each of the four possible markings of the two new variables).

It will be shown in Appendix B.3 that:

LEMMA 64 (COMPLETENESS). *If a query is live then at least one of the five operations can be applied to it.*

At this point we can define our process, which is supposed to ultimately create \mathcal{S} . Start from \mathcal{S}_0 . Once \mathcal{S}_i is defined, which does not satisfy condition (♣), take any live query $Q \in \mathcal{S}_i$, apply one of the five operations to this query, and define \mathcal{S}_{i+1} as \mathcal{S}_i with Q replaced by the query (or queries) resulting from this application. Clearly, for the process to make sense we will need to prove (in Appendix B.4):

LEMMA 65 (SOUNDNESS). *Suppose an application of one of the five operations to some marked query Q_0 returns a set \mathcal{Q} consisting of one or several marked queries. Then, for any \mathbb{D} and any $\bar{a} \in dom(\mathbb{D})^{|\bar{a}|}$ we have $Ch(\mathbb{D}) \models Q_0(\bar{a})$ if and only if there exists a $Q \in \mathcal{Q}$ such that $Ch(\mathbb{D}) \models Q(\bar{a})$.*

It follows from the above lemma that each \mathcal{S}_i satisfies Condition (\clubsuit) . What we still need to show is that the process terminates: at some point we will get \mathcal{S}_i without live queries. To this end, we are going to use *ranks*.

B.2 High-level proof of claim (A). Termination.

For a marked query Q and an atom $\alpha \in Q_G$ the edge rank $erk(\alpha, Q)$ will be defined (Definition 75) as some natural number, reflecting “the minimal cost of hiking from a marked variable to α ”. Then we will prove that:

LEMMA 66. *Suppose a marked query Q' is returned as a result of applying one of the five operations to Q . Then:*

- (i) *If the operation is **cut-red** or **fuse-red** then $|Q_R| > |Q'_R|$.*
- (ii) *If the operation is **cut-green** then $|Q_R| = |Q'_R|$ and for each $\alpha \in Q_G$ there is $erk(\alpha, Q) \geq erk(\alpha, Q')$.*
- (iii) *If the operation is **fuse-green** then $|Q_R| > |Q'_R|$ or $|Q_R| \geq |Q'_R|$ and $erk(\alpha, Q) \geq erk(\alpha, Q')$ for each $\alpha \in Q_G$.*
- (iv) *If the operation is **reduce** and if $\alpha \in Q_G$ is the green atom removed by the operation and if $\alpha' \in q(Q')$ is any of the two green atoms added by the operation then:*
 - (a) $|Q_R| = |Q'_R|$;
 - (b) $erk(\alpha', Q') < erk(\alpha, Q)$;
 - (c) if $\beta \in Q_G \cap Q'_G$, then $erk(\beta, Q') \leq erk(\beta, Q)$.

Recall that all the operations except for *reduce* will decrease the number of variables and of atoms in Q .

Multisets. Using Lemma 66 we are going to prove that our process indeed terminates. To this end we borrow a technique from the term rewriting community. We will use the notation $\{\dots\}_m$ to denote a multiset. $M(A)$ will denote the family of all finite multisets with elements from A . By $<_m$ we will denote the (strict) multiset ordering on $M(\mathbb{N})$. By \mathcal{R} we will mean the set of all possible pairs $\langle k, A \rangle$ where $k \in \mathbb{N}$ and $A \in M(\mathbb{N})$. For $\langle k, A \rangle, \langle k', A' \rangle \in \mathcal{R}$ define $\langle k, A \rangle <_{\mathcal{R}} \langle k', A' \rangle$ if $k < k'$, or $k = k'$ and $A <_m A'$. Finally, let $<_M$ be the (strict) multiset ordering on $M(\mathcal{R})$. It is well known ([19]) that (\heartsuit) if A is well-ordered then the multiset ordering on $M(A)$ is also a well-ordering. So $<_m$ is a well-ordering. In consequence $<_{\mathcal{R}}$, which is the lexicographic ordering on the Cartesian product of two well-ordered sets is a well-ordering too. And, again using (\heartsuit) , we get that $<_M$ is a well-ordering.

- DEFINITION 67. • For a marked query Q define its rank $qrk(Q) \in \mathcal{R}$ as $\langle |Q_R|, \{erk(\alpha, Q) : \alpha \in Q_G\}_m \rangle$.
- For a set of marked queries \mathcal{S} define its rank $srk(\mathcal{S}) \in M(\mathcal{R})$ as the multiset $\{qrk(Q) : Q \in \mathcal{S}\}_m$.

Now, since the set $M(\mathcal{R})$ is well-ordered by $<_M$, to prove termination of our process it is enough to show that whenever it produces two subsequent sets \mathcal{S}_i and \mathcal{S}_{i+1} there must be $(*) srk(\mathcal{S}_{i+1}) <_M srk(\mathcal{S}_i)$. But recall that \mathcal{S}_{i+1} is \mathcal{S}_i with one marked query, call it Q , replaced by one of the five operations, with a set Q' consisting of one or several marked queries. So (this is how the multiset ordering works) in order to show $(*)$ it is enough to show that for each $Q' \in Q$ we have $Q' <_{\mathcal{R}} Q$.

But this follows immediately from Lemma 66, from the definition of the lexicographic ordering $<_{\mathcal{R}}$ (if the operation in question is

cut-red or *fuse-red*) and from the definition of the multiset ordering $<_m$ (for the remaining three operations).

B.3 Proof of Lemma 40 (A). Five operations.

It easily follows from Observation 63 that for every live query Q there must exist a *maximal* variable $x \in var(Q)$. By this we mean that $x \notin V$ and that no atom of the form $E(x, z)$ occurs in $q(Q)$ for $E \in \{G, R\}$. Note that:

LEMMA 68. *Let x be a maximal variable of a live query $\langle \phi(\bar{y}), V \rangle$. Then one of the following condition holds:*

- (i) *x occurs in exactly one atom $E(z, x)$, with $E \in \{G, R\}$;*
- (ii) *x occurs in exactly two atoms $R(x_r, x)$ and $G(x_g, x)$ for some x_r, x_g ;*
- (iii) *there exist at least two vertices $z \neq z'$ and $E \in \{G, R\}$ such that $E(z, x)$ and $E(z', x)$ are atoms of $\phi(\bar{y})$.*

As promised, now we can define the **five operations**. Suppose $Q = \langle \phi(\bar{y}), V \rangle$ is a live query and $x \in var(Q)$.

DEFINITION 69 (cut-red). *Suppose x is as in Lemma 68(i), with $E = R$. Define $cut-red(Q, x)$ as $\langle \phi'(\bar{y}), V(Q) \rangle$ where $\phi'(\bar{y})$ is created from $\phi(\bar{y})$ by removing the sole atom containing x .*

Operation cut-green is defined in an analogous way.

DEFINITION 70 (fuse-red). *Let x, z and z' be as in Lemma 68 (iii), with $E = R$. Then $fuse-red(Q, x, z, z') = \langle \phi'(\bar{y}), V \rangle$ where $\phi'(\bar{y})$ is $\phi(\bar{y})$ with all occurrences of z' renamed⁷ as z .*

Operation fuse-green is defined in an analogous way.

DEFINITION 71 (REDUCE). *Suppose x is as in Lemma 68 (ii). Let $\phi'(\bar{y})$ be a query obtained from $\phi(\bar{y})$ by removing atoms $R(x_r, x)$ and $G(x_g, x)$ and replacing them with atoms $G(x', x'')$, $G(x'', x_r)$, $R(x', x_g)$ where x' and x'' are fresh variables. Then define $reduce(Q, x)$ as the set of four marked queries $\langle \phi'(\bar{y}), V(Q) \rangle$, $\langle \phi'(\bar{y}), V(Q) \cup \{x'\} \rangle$, $\langle \phi'(\bar{y}), V(Q) \cup \{x', x''\} \rangle$ and $\langle \phi'(\bar{y}), V(Q) \cup \{x''\} \rangle$ ⁸.*

Now, Lemma 64 easily follows from Lemma 68 and from Definitions 69–71.

The next thing left to be proven in this section is Lemma 66. Notice that claim (i) of the Lemma is now obvious: cut-red just removes a single red edge and fuse-red merges two red edges into one. In order to prove claims (ii)-(iv) however one needs to work a little bit harder.

Ranks. We are now going to define the rank $erk(\alpha, Q)$ for a live query Q and an atom $\alpha \in Q_G$. To this end we consider paths, from some marked variable to α , traversing edges of $q(Q)$ (in both directions). To be more precise:

DEFINITION 72. *Given a live query Q an R -path is a finite sequence P such that:*

- *each of the elements of P is either $E(t, z)$ or $E^{-1}(z, t)$ for some $E(t, z)$ from $q(Q)$, where $E \in \{G, R\}$ (obviously, $E^{-1}(z, t)$ means that we traverse $E(t, z)$ backwards);*

⁷One could wonder what happens if one of the two variables we unify is in V and the other is not in V . But this is prohibited by Observation 63 (iii).

⁸Note that $\langle \phi'(\bar{y}), V(Q) \cup \{x''\} \rangle$ is not properly marked, and thus not live, and it will in no way contribute to our process any more.

- if $E(s, t)$ and $E'(u, z)$ are two consecutive elements of P then $t = u$ (where $E, E' \in \{G, R, G^{-1}, R^{-1}\}$);
- (\star) if $R(t, z)$ is an atom of $q(Q)$ then only one of $R(t, z)$ or $R^{-1}(z, t)$ can appear in P and it can appear at most once.

Notice that an atom $G(t, z)$ of $q(Q)$, as well as $G^{-1}(z, t)$, can appear any number of times in a R -path. Each R -path has its elevation and its cost⁹:

DEFINITION 73 (ELEVATION AND COST). For an empty R -path \emptyset we define $\text{cost}(\emptyset) = 0$ and $\text{elev}(\emptyset) = 3^{|Q_R|}$. For a path $P = P'E(x, z)$ we define:

- $\text{elev}(P) = \text{elev}(P')$ if $E \notin \{R, R^{-1}\}$
- $\text{elev}(P) = 3 \cdot \text{elev}(P')$ if $E = R$
- $\text{elev}(P) = \frac{1}{3} \cdot \text{elev}(P')$ if $E = R^{-1}$
- $\text{cost}(P) = \text{cost}(P') + \text{elev}(P')$ if $E \in \{G, G^{-1}\}$
- $\text{cost}(P) = \text{cost}(P')$ if $E \notin \{G, G^{-1}\}$

Notice that it follows from condition (\star) of Definition 72 that $\text{elev}(P)$ is always a positive natural number.

DEFINITION 74. For an atom $\alpha = G(u, u')$ in $q(Q)$ by an α -hike we mean an R -path such that:

- if $E(t, z)$ or $E^{-1}(t, z)$ is the first atom of P then $t \in V(Q)$;
- the last element of P is either $G(u, u')$ or $G^{-1}(u', u)$.

For $\alpha \in Q_G$ we denote the set of all α -hikes as $\text{hikes}(\alpha, Q)$.

DEFINITION 75. For an atom $\alpha \in Q_G$ its rank is defined as: $\text{erk}(\alpha, Q) = \min(\{\text{cost}(P) : P \in \text{hikes}(\alpha, Q)\})$.

It is easy to see that:

OBSERVATION 76. For a marked query Q and for $\alpha = G(u, u')$ in Q_G , if $P \in \text{hikes}(\alpha, Q)$ and $\text{cost}(P) = \text{erk}(\alpha, Q)$ then any of α or $G^{-1}(u', u)$ can occur only as the last atom of P .

Now we can finally prove Lemma 66.

Proof of Lemma 66 (ii). Let $Q' = \text{cut-green}(Q, x)$. It follows directly from the construction that $|Q_R| = |Q'_R|$. Now take any $\alpha \in Q'_G$. We need to show that $\text{erk}(\alpha, Q) \geq \text{erk}(\alpha, Q')$. Let $z \in \text{var}(Q)$ be such that $G(z, x)$ is an atom of $q(Q)$. Consider any $P \in \text{hikes}(\alpha, Q)$. Let P' be a path obtained from P by deleting each occurrence of $G(z, x)$ and $G^{-1}(x, z)$. Then $P' \in \text{hikes}(\alpha, Q')$ and $\text{cost}(P') \leq \text{cost}(P)$. \square

Proof of Lemma 66 (iii). Let $Q' = \text{fuse-green}(Q, x, z, z')$.

If (\ast) there exists a variable $u \in \text{var}(Q)$ such that atoms $R(z, u)$ and $R(z', u)$ are in $q(Q)$, or that $R(u, z)$ and $R(u, z')$ are in $q(Q)$, then the two red edges merge in $q(Q')$ and $|Q_R| > |Q'_R|$. Also, if ($\ast\ast$) at least two of the atoms $R(z, z')$, $R(z', z)$, $R(z, z)$, and $R(z', z')$ are in $q(Q)$ then $|Q_R| > |Q'_R|$.

So suppose there is neither (\ast) nor ($\ast\ast$). Take any $\alpha \in Q'_G$ and $P \in \text{hikes}(\alpha, Q)$. Consider path P' obtained from P by replacing each occurrence of z' in P with z . Obviously, $\text{cost}(P) = \text{cost}(P')$. We now will show that $P' \in \text{hikes}(\alpha, Q')$. Clearly, one only needs to worry if condition (\star) of Definition 72 holds. Suppose towards contradiction that it does not. It can only happen when there exist $s, t \in \text{var}(Q')$ such that two atoms from $\{R(s, t), R^{-1}(t, s)\}$ appear in P .

⁹The cost of making a step depends on the current elevation. Current elevation depends on the difference between the total ascent and total descent.

Let us assume that $R(s, t)$ (other cases are analogous) appears twice in P' . But of course $R(s, t)$ could not appear twice in P , which is an R -path. So at least one occurrence of $R(s, t)$ in P' results from the unification of z and z' .

There are two cases: either (a) $s = t = z$ or (b) exactly one of s, t equals z . So suppose (b) happened and without loss of generality assume that $s = z$. We know that there was $R(z, t)$ and $R(z', t)$ in P and they both unified to $R(z, t)$ in P' . But this would imply (\ast), leading to a contradiction. The remaining case (a) is that $R(z, z)$ occurs twice in P' . But this would need ($\ast\ast$) to be true, which is a contradiction again.

So $P' \in \text{hikes}(\alpha, Q')$ and $\text{erk}(\alpha, Q) \geq \text{erk}(\alpha, Q')$. \square

Proof of Lemma 66 (iv). Let Q' be any of the marked queries in $\text{reduce}(Q, x)$. Let x_r and x_g be variables of Q such that $G(x_g, x)$ and $R(x_r, x)$ are atoms of $q(Q)$. Recall that $x', x'' \in \text{var}(Q')$ are two new variables added by *reduce*.

Clearly, $|Q'_R| = |Q_R|$, since when creating $q(Q')$ from $q(Q)$ we have just replaced one red edge with another.

Let us first show **claim (c)**. Take some $\beta \in Q_G \cap Q'_G$ and $P \in \text{hikes}(\beta, Q)$ such that $\text{cost}(P) = \text{erk}(\beta, Q)$. Our goal is to find $P' \in \text{hikes}(\beta, Q')$ such that $\text{cost}(P') \leq \text{cost}(P)$.

Obviously, one can assume that P does not contain, as a connected subsequence, $G(x_g, x)G^{-1}(x, x_g)$. If it did, we could remove such subsequence to get a new P with lower cost. And, as an R -path, P does not contain subsequence $R(x_r, x)R^{-1}(x, x_r)$ either.

If no atom of P contains x then $P \in \text{hikes}(\beta, Q')$ and thus we set $P' = P$. Otherwise it contains exactly one of $R(x_r, x)G^{-1}(x, x_g)$ or $G(x_g, x)R^{-1}(x, x_r)$ as a connected subsequence. Suppose $P = P_0R(x_r, x)G^{-1}(x, x_g)P_1$. Let $A = G^{-1}(x_r, x'')G^{-1}(x'', x')R(x', x_g)$ and consider path $P' = P_0AP_1$. Note that $P' \in \text{hikes}(\beta, Q)$. Now we need to show that $\text{cost}(P) > \text{cost}(P')$. But:

$$\text{cost}(P) = \text{cost}(P_0) + 3 \cdot \text{elev}(P_0) + 3 \cdot \text{elev}(P_0) \cdot \text{cost}(P_1)$$

$$\text{cost}(P') = \text{cost}(P_0) + 2 \cdot \text{elev}(P_0) + 3 \cdot \text{elev}(P_0) \cdot \text{cost}(P_1)$$

The case when $P = P_0G(x_g, x)R^{-1}(x, x_r)P_1$, is similar.

Let us now move to **claim (b)**. Take $P \in \text{hikes}(\alpha, Q)$ such that $\text{cost}(P) = \text{erk}(\alpha, Q)$. We are going to build a path $P' \in \text{hikes}(\alpha', Q')$ such that $\text{cost}(P') < \text{cost}(P)$.

There are two cases depending on the last atom of P :

In the first case P is $P_0G(x_g, x)$. Let then P' be $P_0R^{-1}(x_g, x')G(x', x'')$ or $P_0R^{-1}(x_g, x')G(x', x'')G(x'', x_r)$ (depending on whether α' is $G(x', x'')$ or $G(x'', x_r)$). Then of course $P' \in \text{hikes}(\alpha', Q)$ and $\text{cost}(P) = \text{cost}(P_0) + \text{elev}(P_0)$ while $\text{cost}(P') \leq \text{cost}(P_0) + \frac{2}{3} \cdot \text{elev}(P_0)$.

The second case is that $P = P_0R(x_r, x)G^{-1}(x, x_g)$. Here the argument is similar. Now we take $P' = P_0G^{-1}(x_r, x'')$ or $P' = P_0G^{-1}(x_r, x'')G^{-1}(x'', x')$ (again depending on whether α' is $G(x', x'')$ or $G(x'', x_r)$). Then $P' \in \text{hikes}(\alpha', Q)$ and $\text{cost}(P) = \text{cost}(P_0) + 3 \cdot \text{elev}(P_0)$ while $\text{cost}(P') \leq \text{cost}(P_0) + 2 \cdot \text{elev}(P_0)$. \square

B.4 Proof of Lemma 65

Let $Q = \langle \phi(\bar{y}), V \rangle$ be any live marked query.

Lemma 65 will follow directly from Lemmas 77–79:

LEMMA 77. For any set of facts \mathbb{D} , tuple $\bar{a} \in \text{dom}(\mathbb{D})^{|\bar{y}|}$ and x such that $Q' = \text{cut-green}(Q, x)$ (or $Q' = \text{cut-red}(Q, x)$) the following

holds:

$$\text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q(\bar{a}) \iff \text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q'(\bar{a}).$$

PROOF. Let $G(z, x)$ be the atom removed from $q(Q)$ by cut-green.

(\Rightarrow). It follows immediately as $q(Q')$ is a subset of $q(Q)$ and $V(Q') = V(Q)$.

(\Leftarrow). Let h' be a homomorphism witnessing that $\text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q'(\bar{a})$. Due to the rule (*pins*) of ruleset_d there must exist $t \in \text{dom}(\text{Ch}(\mathcal{T}_d, \mathbb{D}))$ such that $G(h'(z), a) \in \text{Ch}(\mathcal{T}_d, \mathbb{D})$. Define a new homomorphism h as $h' \cup \langle x, t \rangle$. Then h is a homomorphism witnessing that $\text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q(\bar{a})$.

The proof for cut-red is analogous. \square

LEMMA 78. For any set of facts \mathbb{D} , tuple $\bar{a} \in \text{dom}(\mathbb{D})^{|\bar{y}|}$ and x, z, z' such that $Q' = \text{fuse-green}(Q, x, z, z')$ (or $Q' = \text{fuse-red}(Q, x, z, z')$) the following holds:

$$\text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q(\bar{a}) \iff \text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q'(\bar{a}).$$

PROOF. Note that there exists a homomorphism (the one that identifies z' with z) from $q(Q)$ to $q(Q')$ which preserves markings of variables. Thus (\Leftarrow) is trivial.

(\Rightarrow). Let h be a homomorphism witnessing that $\text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q(\bar{a})$. If we can show that $h(z) = h(z')$ then h will witness that $\text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q'(\bar{a})$ as well. First recall that, since $x \notin V(Q)$ we can be sure that $h(x) \in \text{dom}(\text{Ch}(\mathcal{T}_d, \mathbb{D}) \setminus \text{dom}(\mathbb{D}))$. Then notice that it follows from the rules of \mathcal{T}_d that the in-degree, with respect to relation G , of any term of $\text{dom}(\text{Ch}(\mathcal{T}_d, \mathbb{D}) \setminus \text{dom}(\mathbb{D}))$ is at most one. Thus the only way for $\text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q(\bar{a})$ to happen is that $h(z) = h(z')$ thus $\text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q'(\bar{a})$. The proof for fuse-red is analogous. \square

LEMMA 79. For any set of facts \mathbb{D} , tuple $\bar{a} \in \text{dom}(\mathbb{D})^{|\bar{y}|}$ and x such that $Q = \text{reduce}(Q, x)$ the following holds:

$$\text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q(\bar{a}) \iff \exists Q' \in \mathcal{Q} \text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q'(\bar{a}).$$

PROOF. Let x', x'', x_r , and x_g be variables of $q(Q)$ such that $R(x_r, x), G(x_g, x) \in q(Q)$ and that $R(x', x_g), G(x', x''), G(x'', x_r) \in q(Q')$.

(\Leftarrow). Let $Q' \in \mathcal{Q}$ be such that $\text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q'(\bar{a})$ and let h' be a homomorphism witnessing that. We need to show that there exists a homomorphism h witnessing $\text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q(\bar{a})$.

Since $R(x', x_g), G(x', x''), G(x'', x_r)$ are atoms of $q(Q)$, and h' is a homomorphism, we know that $R(h'(x_g), h'(x')), G(h'(x'), h'(x''))$ and $G(h'(x''), h'(x_r))$ are atoms of $\text{Ch}(\mathcal{T}_d, \mathbb{D})$. But, since (*grid*) is a rule of \mathcal{T}_d , this implies that there exists an element $t \in \text{Ch}(\mathcal{T}_d, \mathbb{D})$, such that $G(h'(x_g), h'(t))$ and $R(h'(x_r), h'(t))$ are also in $\text{Ch}(\mathcal{T}_d, \mathbb{D})$.

Define h by:

- $h(u) = h'(u)$ for $u \in \text{var}(Q) \setminus \{x\}$.
- $h(x) = t$.

(\Rightarrow). Let h be a homomorphism witnessing that $\text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q(\bar{a})$. We will show that there exist $Q' \in \mathcal{Q}$ and a homomorphism h' such that h' is witnessing $\text{Ch}(\mathcal{T}_d, \mathbb{D}) \models Q'(\bar{a})$. Again recall that x is an unmarked variable and so let us find parents of $h(x)$ in $\text{Ch}(\mathcal{T}_d, \mathbb{D})$. Note that $h(x)$ could be created only by rule (*grid*) as it has an in-degree of two. Let σ be such that $\text{appl}(\text{grid}, \sigma) = R(h(x_r), h(x)), G(h(x_g), h(x))$. We set:

$$h' = h \setminus \{x, h(x)\} \cup \{x', \sigma(x), x'', \sigma(y)\}.$$

There are four possible picks for Q' from \mathcal{Q} . While homomorphism h' works for any of them as elements of Q differ only by markings, we need to make sure that the marking of Q' agrees with h' . Obviously for every variable of Q this is the case. However we have two new variables to consider, namely x' and x'' . Thus we need to take Q' such that it satisfies:

- $h'(x')$ is a constant of \mathbb{D} if and only if $x' \in V(Q')$.
- $h'(x'')$ is a constant of \mathbb{D} if and only if $x'' \in V(Q')$.

This is trivially possible due to Definition 71. \square

Appendix C GENERALIZATION OF LEMMA 40

For $K \in \mathbb{N}$ define \mathcal{T}_d^K as the theory, over $\Sigma_K = \{I_K, I_{K-1}, \dots, I_1\}$ (each I_k is a binary relation symbol), comprising, for each $1 \leq i < K$ and each $1 \leq k \leq K$ the following $2K + 1$ rules:

$$\begin{aligned} (\text{loop}) \quad & \text{true} \Rightarrow \exists x \ I_K(x, x), I_{K-1}(x, x), \dots, I_1(x, x) \\ (\text{pin}_k) \quad & \forall x \ (\text{true} \Rightarrow \exists z \ I_k(x, z)) \\ (\text{grid}_i) \quad & I_{i+1}(x, x'), I_i(x, u), I_i(u, u') \Rightarrow \exists z \ I_{i+1}(u', z), I_i(x', z) \end{aligned}$$

Using the ideas from Appendix B one can now show that:

LEMMA 80. For each $K \in \mathbb{N}$:

A. theory \mathcal{T}_d^K is BDD;

B. there is a query $\psi(y, y')$ such that $\text{rew}_{\mathcal{T}_d^K}(\psi(y, y'))$ contains a CQ of size $(K-1)$ -fold exponential in the size of ψ .

Like in the case of Lemma 40, claim **B** of Lemma 80 is relatively easy to prove.

For the proof of claim **A**, properly marked queries first need to be slightly redefined (but let us skip it here). Then the *five operations* need to be generalized in the natural way: we will now have K cut operations, K fuse operations, and $K-1$ reduce operations.

The non-obvious part is how to modify the ranks erk and qrk so that they do their job correctly in the new circumstances. For that I_i -paths need to be defined (for $1 \leq i < K$), analogous to R -paths in Definition 72. But now the condition (\star) will apply to atoms of the relation I_i (including I_i^1). Then i -elevation ($elev_i$) will be defined, like in Definition 73. And finally, we will need $cost_i$ of a path, calculated almost like the cost of the path in Definition 73: for a path $P = P'E(x, z)$ we have $cost_i(P) = cost(P') + elev_i(P')$ if $E = I_{i-1}$ or I_{i-1}^1 and $cost_i(P) = cost(P')$ otherwise. Notice that I_i is the new red and I_{i-1} is the new green. Note also that we have a new situation: E may now very well be neither “green” nor “red”. And that is fine, in such case it neither contributes to the elevation nor to the cost of the path.

Having the function $cost_i$, rank $erk(\alpha)$ of an atom α of the relation I_{i-1} is (like in Appendix B.3) defined as minimal $cost_i$ of an I_i -path from some marked variable to α , and rank $qrk_i(Q)$ is the multiset of all ranks $erk(\alpha)$ of all atoms α of I_{i-1} in Q . Finally, $qrk(Q)$ is the tuple:

$$\langle |Q_K|, qrk_K(Q), |Q_{K-1}|, qrk_{K-1}(Q), \dots, |Q_2|, qrk_2(Q) \rangle$$

where $|Q_i|$ is the number of the atoms of the relation I_i in Q . Clearly, the lexicographic ordering on the set of such ranks is a well ordering. Now, a careful case inspection shows that each of the $3K-1$ operations decreases the rank of a query.

Appendix D REMARKS ON \mathcal{T}_d

There are two technical problems one needs to deal with here:

- (1) There are conjunctions in the heads of rules (loop) and (grid) (notice that the rule (pins) can be simply split into two single-head rules).
- (2) we have predicate "true" in the bodies of rules (loop) and (pins)

To deal with the first problem we could introduce a new ternary predicate T , and rewrite the rules as:

$$\begin{array}{ll} \text{(loop)} & \text{true} \Rightarrow \exists x T(x, x, x) \\ \text{(grid)} & R(x, x'), G(x, u), G(u, u') \Rightarrow \exists z T(u', x', z) \end{array}$$

Then add Datalog rules:

$$\begin{array}{ll} T(x, y, z) & \Rightarrow R(x, z) \\ T(x, y, z) & \Rightarrow G(y, z) \\ R(x, z), G(y, z) & \Rightarrow T(x, y, z) \end{array}$$

In this way $T(x, y, z)$ is just a macro for $R(x, z)$ and $G(y, z)$ and clearly the transformation does not change the BDD status of the theory.

To deal with the second problem we need to replace the predicate "true" with something like "if anything at all exists". For rule (loop) it would mean that it should be replaced with three rules:

$$\begin{array}{ll} R(y, z) & \Rightarrow \exists x T(x, x, x) \\ G(y, z) & \Rightarrow \exists x T(x, x, x) \\ T(t, y, z) & \Rightarrow \exists x T(x, x, x) \end{array}$$

This transformation leads to a theory which is equivalent to our \mathcal{T}_d on all **nonempty** instances, but this is of course good enough.

An analogous (albeit slightly more complicated) trick works for (pins).

Appendix E STICKY AND BD-LOCAL

Here we prove the following observation:

OBSERVATION 81. *Every connected sticky theory is bd-local*

PROOF. Let \mathcal{T} be a connected sticky theory and \mathbb{D} be any instance with k being its degree. Let $A(\vec{a})$ be an atom of $Ch(\mathcal{T}, \mathbb{D})$ and \mathbb{F} be the minimal subset of \mathbb{D} such that $A(\vec{a}) \in Ch(\mathcal{T}, \mathbb{F})$. We show that $|\mathbb{F}|$ depends only on \mathcal{T} and k . As \mathcal{T} is connected and sticky we know that if $|\mathbb{F}| > 1$ then every fact of \mathbb{F} needs to contain some term of \vec{a} . From this we conclude that $|\mathbb{F}| \leq |\vec{a}|k$, thus we know that \mathcal{T} is local. \square