A Journey to the Frontiers of Query Rewritability

Piotr Ostropolski-Nalewaja
Jerzy Marcinkowski
postropolski@cs.uni.wroc.pl
jma@cs.uni.wroc.pl
University of Wrocław, Institute of Computer Science
Wrocław, Poland

David Carral
LIRMM, Inria, University of Montpellier, CNRS
Montpellier, France
david.carral@inria.fr

Sebastian Rudolph
TU Dresden, Computational Logic Group
Dresden, Germany
sebastian.rudolph@tu-dresden.de

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.

CCS CONCEPTS

• **Theory of computation → Logic**; **Automated reasoning**.

KEYWORDS

First-Order logic, Existential Rules, First-Order Rewritability

ACM Reference Format:
Piotr Ostropolski-Nalewaja, Jerzy Marcinkowski, David Carral, and Sebastian Rudolph. 2022. A Journey to the Frontiers of Query Rewritability. In *Proceedings of the 41st ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS '22), June 12–17, 2022, Philadelphia, PA, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3517804.3524163

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 $\mathbb{D}=Ch_0(\mathbb{D},\mathcal{T})$, $Ch_1(\mathbb{D},\mathcal{T})$, $Ch_2(\mathbb{D},\mathcal{T})$, ... 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 [11] that a conjunctive query is entailed if and only if it holds in the corresponding chase, which allows for reducing entailment to model checking:

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

Finite expansion sets. We say $\mathcal{T}$ enjoys the *finite expansion set* property [3] (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:

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

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

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

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 [6] to FUS (*finite unification set*) [3]. 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 [15].

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 [4, 9]), generalizing linear theories;
• sticky theories, defined by a reasonably natural syntactic restriction on the use of joins [8].

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 [12];
• binary BDD theories, where the arity of relation symbols is at

ACM ISBN 978-1-4503-9260-0/22/06...$15.00
https://doi.org/10.1145/3517804.3524163

most 2 (also studied in the context of description logics [1]);
• backward shy theories [18], 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:

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

This conjecture was studied earlier [10] 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 [5].

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.

The complete proofs of all results are included in the appendix of an extended technical report, which is available online [17].

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 $\exists x\ Siblings(Abel, x), Female(x)$ 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 $\forall \bar{x}, \bar{y}\big(\beta(\bar{x},\bar{y}) \Rightarrow \exists \bar{w}\ \alpha(\bar{y},\bar{w})\big)$, 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 $fr(\rho)$, 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 $dom(\mathbb{F})$ denote its *active domain* – the set of all terms that appear in the facts of $\mathbb{F}$. For $c, c' \in dom(\mathbb{F})$, we let $dist_{\mathbb{F}}(c, c')$ denote the distance between c and c' in the Gaifman graph of $\mathbb{F}$: the vertices of this graph are elements of $dom(\mathbb{F})$ 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 $\bar{a} \in dom(\mathbb{D})^{|\bar{y}|}$, 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 $\mathbb{D}_a = \{Human(Abel)\}$ and the theory $\mathcal{T}_a$ consisting of the following two rules:*

$$Human(y) \Rightarrow \exists z\; Mother(y, z)$$
$$Mother(x, y) \Rightarrow Human(y)$$

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

Homomorphisms and query containment. For structures $\mathbb{D}$, $\mathbb{F}$, a *homomorphism* from $\mathbb{D}$ to $\mathbb{F}$ is a function $h : dom(\mathbb{D}) \rightarrow dom(\mathbb{F})$ 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 $a \in dom(\mathbb{D})$. Given a fact $\alpha = A(\vec{x}) \in \mathbb{D}$, we will use $h(\alpha)$ to denote $A(h(\vec{x}))$.

For two CQs $\phi(\bar{y})$ and $\psi(\bar{y})$, with the same set of free variables, we say that $\phi(\bar{y})$ *contains* $\psi(\bar{y})$ if for every structure $\mathbb{D}$ and for every tuple $\bar{a} \in dom(\mathbb{D})^{|\bar{y}|}$ if $\mathbb{D} \models \psi(\bar{a})$ than also $\mathbb{D} \models \phi(\bar{a})$. It is well known that $\phi(\bar{y})$ contains $\psi(\bar{y})$ if and only if there is a homomorphism[1] from $\phi(\bar{y})$ to $\psi(\bar{y})$ that is the identity on $\bar{y}$.

Core of a structure. A substructure $\mathbb{H}$ of a finite structure $\mathbb{G}$ is a *core* of $\mathbb{G}$ (see [14]) 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 $a \in dom(\mathbb{G})$.

It is well known that [14]:

(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 $Core(\mathbb{G})$ 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 $dom(\mathbb{F})$ 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 $\mathcal{T}, \mathbb{D} \models \phi(\bar{a})$ for given theory $\mathcal{T}$, instance $\mathbb{D}$, CQ $\phi(\bar{y})$ and tuple $\bar{a} \in dom(\mathbb{D})^{|\bar{y}|}$. 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 $\tau = \tau(\phi(\bar{y}))$ of some *atomic* conjunctive query $\phi(\bar{y})$, and for each natural number $1 \leq i \leq arity(E)$, where E is the relation symbol of $\phi(\bar{y})$, let f_i^{τ} be a function symbol, with arity equal to $|\bar{y}|$, that is the number of free variables in $\phi(\bar{y})$.

Definition 2 (Skolemization). *For a given TGD ρ, of the form $\beta(\bar{x}, \bar{y}) \Rightarrow \exists \bar{w}\; \alpha(\bar{x}, \bar{w})$ by $sh(\rho)$ we denote the* Skolemization of the head of *ρ, that is the atom $\alpha(\bar{x}, \bar{w})$, with each variable $w \in \bar{w}$ replaced by the term $f_i^{\tau}(\bar{x})$, where i is the earliest position in $\alpha(\bar{x}, \bar{w})$ 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 variables[2] 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 $\beta(\bar{x}, \bar{y}) \Rightarrow \exists \bar{w}\; \alpha(\bar{y}, \bar{w})$, 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 $\bar{x} \cup \bar{y}$ to $dom(\mathbb{F})$ such that $\sigma(\beta(\bar{x}, \bar{y})) \subseteq \mathbb{F}$ (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-oblivious Skolem chase procedure).

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

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

[2]Including non-frontier variables as arguments of the functions f_i^{τ} (like x in the current example) would lead to the *oblivious chase*. Note that the names of the terms do not identify ρ (their "rule of origin"), which is important in the proof of Theorem 1.

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 $\phi(\bar{y})$, instance $\mathbb{D}$ and $\bar{a} \in dom(\mathbb{D})^{|\bar{y}|}$, we have:

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

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 $Ch(\mathcal{T},\mathbb{F}) = Ch(\mathcal{T},\mathbb{D})$.*

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. *For some $t \in dom(Ch(\mathcal{T},\mathbb{D})) \setminus dom(\mathbb{D})$, there is exactly one atom $\alpha \in Ch(\mathcal{T},\mathbb{D})$ with t appears in α abd $t \notin fr(\alpha)$.*

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:

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

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 [3, 7].

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

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

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

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

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

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

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 [3] 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 [6], 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 $\bar{t}$ of domain elements from $dom(Ch_i(\mathcal{T},\mathbb{D}))$ and for any $R \in \Sigma$, $Ch(\mathcal{T},\mathbb{D}) \models R(\bar{t})$ implies $Ch_{i+n_{at}}(\mathcal{T},\mathbb{D}) \models R(\bar{t})$.*

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 $\max(\{n_\phi \mid \phi \text{ is an atomic query}\})$. □

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 [7] (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' of $dom(\mathbb{D})$, 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* [3, 11]:

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 [2], 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 [11] that $Core(Ch_{c_{\mathcal{T},\mathbb{D}}}(\mathcal{T}, \mathbb{D})) \models \mathcal{T}, \mathbb{D}$.

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

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 [10] where an incorrect proof was proposed, and in [5] 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 $\mathcal{T}_\infty$ 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 $\mathcal{T}_\infty$ 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:*

$$Ch(\mathcal{T}, \mathbb{D}) \quad = \quad \bigcup\nolimits_{\mathbb{F} \subseteq \mathbb{D}, |\mathbb{F}| \leq l_{\mathcal{T}}} \ Ch(\mathcal{T}, \mathbb{F})$$

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 $\mathcal{F} = \{\mathbb{D} \mid Ch(\mathcal{T}, \mathbb{D}) \models \Phi \wedge |\mathbb{D}| \leq l_{\mathcal{T}}|\Phi|\}$. For any $\mathbb{D} \in \mathcal{F}$, define $n_{\mathbb{D}}$ as the minimal natural number such that $Ch_{n_{\mathbb{D}}}(\mathcal{T}, \mathbb{D}) \models \Phi$. 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. □

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 is presented in Appendix A in [17]. 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 in [17]). Second, it is straightforward to find the mentioned bound on the number of ancestors for every atom produced by $\mathcal{T}_{NF}$ (Lemma 57 in [17]). 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 $\gamma \in rew(\beta)$. 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 $Core(\mathcal{T}, \mathbb{D}) \subseteq Ch_{c_{\mathcal{T}}}(\mathcal{T}, \mathbb{D})$ 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 $Core(\mathcal{T}, \mathbb{D}) \subseteq Ch_{c_{\mathcal{T}}}(\mathcal{T}, \mathbb{D})$ 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, $Ch(\mathcal{T}, \mathbb{D})$ will be shortened to $Ch(\mathbb{D})$.

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 $\mathbb{C}_{\mathbb{D}} = \bigcup_{\mathbb{F} \in \mathcal{I}_{\mathbb{D}}} Core(\mathcal{T}, \mathbb{F})$.*

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

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 $k_{\mathcal{T}} = \max\{c_{\mathcal{T},\mathbb{D}} \mid \mathbb{D} \in \mathcal{A}\}$. □

If we were able able to find a homomorphism $\bar{h}_{\mathbb{D}}$ from $Ch(\mathbb{D})$ 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 $Ch(\mathbb{F})$ to $\mathbb{C}_{\mathbb{D}}$ and we know that $\bigcup_{\mathbb{F} \in \mathcal{I}_{\mathbb{D}}} Ch(\mathbb{F}) = Ch(\mathbb{D})$. So can't we just define $\bar{h}_{\mathbb{D}} = \bigcup_{\mathbb{F} \in \mathcal{I}_{\mathbb{D}}} h_{\mathbb{F}}$? 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 $\bar{h}_{\mathbb{D}}$ 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 $\bar{h}_{\mathbb{D}}$ 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 $\bar{h}_{\mathbb{D}}$ from $Ch(\mathbb{D})$ to $Ch(\mathbb{D})$ such that for each $t \in dom(Ch(\mathbb{D}))$ there is $\bar{h}_{\mathbb{D}}(t) \in dom(\mathbb{C}_{\mathbb{D}})$.*

Let us first discuss how Theorem 2 can be concluded from Lemma 26. Suppose some $\mathbb{D}$ is fixed and $\bar{h}_{\mathbb{D}}$ is a homomorphism as in Lemma 26. We know that $t \in dom(Ch(\mathbb{D}))$ implies $\bar{h}_{\mathbb{D}}(t) \in dom(\mathbb{C}_{\mathbb{D}})$ and we know that $\mathbb{C}_{\mathbb{D}} \subseteq Ch_{k_{\mathcal{T}}}(\mathbb{D})$. So one might be tempted to immediately conclude $\bar{h}_{\mathbb{D}}(Ch(\mathbb{D})) \subseteq Ch_{k_{\mathcal{T}}}(\mathbb{D})$. But it is not quite that simple. Admittedly, Lemma 26 tells us that all the terms of $\bar{h}_{\mathbb{D}}(Ch(\mathbb{D}))$ will indeed appear in $Ch_{k_{\mathcal{T}}}(\mathbb{D})$. But it says nothing like that about the atoms of $\bar{h}_{\mathbb{D}}(Ch(\mathbb{D}))$. Rather, it might be that there are atoms in $\bar{h}_{\mathbb{D}}(Ch(\mathbb{D}))$ that, despite having all their terms in $dom(\mathbb{C}_{\mathbb{D}})$ 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 $\bar{h}_{\mathbb{D}}(Ch(\mathbb{D})) \subseteq Ch_{c_{\mathcal{T}}}(\mathbb{D})$ 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 $M_{\mathbb{F}}$ denote*[3] *the substructure of $Ch(\mathbb{D})$ induced by the set of terms $dom(Ch(\mathbb{D})) \setminus \big(dom(Ch(\mathbb{F})) \setminus dom(Core(\mathcal{T}, \mathbb{F}))\big)$.*

In the following, the terms of $dom(Ch(\mathbb{F})) \setminus dom(Core(\mathcal{T}, \mathbb{F})$ 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 $\sigma \in \mathcal{H}om(\rho, M_{\mathbb{F}})$. Of course, $appl(\rho, \sigma) \in Ch(\mathbb{D})$, since $Ch(\mathbb{D})$ is by definition closed under rule applications.

In the following, the atom $appl(\rho, \sigma)$ 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 $fr(\alpha)$. 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 $head(\rho)$, 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 $Ch(\mathbb{D})$ but not in $M_{\mathbb{F}}$ is that α contains some banned term t.

But $\sigma(body(\rho)) \subseteq M_{\mathbb{F}}$. Thus $fr(\alpha) \subseteq dom(M_{\mathbb{F}})$ and so $t \notin fr(\alpha)$. 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 $Ch(\mathbb{D})$. But $t \in dom(Ch(\mathbb{F}))$ so from Observation 9, we know that α is the birth atom of t in $Ch(\mathbb{F})$ and thus $\alpha \in Ch(\mathbb{F})$.

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

Now let $h^*_{\mathbb{F}} : Ch(\mathbb{F}) \to Core(\mathcal{T}, \mathbb{F})$ be a homomorphism as in Property 7. Since $\alpha \in Ch(\mathbb{F})$, we obtain $h^*_{\mathbb{F}}(\alpha) \in Core(\mathcal{T}, \mathbb{F})$ and

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

thus $h^*_{\mathbb{F}}(\alpha) \in M_{\mathbb{F}}$. As $h^*_{\mathbb{F}}$ is a retraction, $h^*_{\mathbb{F}}(fr(\alpha)) = fr(\alpha)$ follows from (∗). Hence, $h^*_{\mathbb{F}}(\alpha)$ can serve as our α', concluding the proof. □

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

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. □

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 $dom(Ch(M_{\mathbb{F}}))$, that is equal to $dom(Ch(\mathbb{D}))$, 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 $Ch(\mathbb{D})$ to $Ch(\mathbb{D})$ (and it will be the identity on $dom(\mathbb{D})$, since each $h_{M_{\mathbb{F}}}$ is). Now the rabbit is going to be pulled out of the hat: let us compose **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 $\bar{h}_{\mathbb{D}}(t) \in dom(\mathbb{C}_{\mathbb{D}})$ does indeed hold for each term $t \in dom(Ch(\mathbb{D}))$.

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 $ban_{\mathbb{F}}$ of terms that occur somewhere in $Ch(\mathbb{F})$ but not in $Core(\mathcal{T}, \mathbb{F})$. Each $h^*_{M_{\mathbb{F}}} \in \mathcal{H}_{\mathbb{D}}$ is the identity on all terms except those of $ban_{\mathbb{F}}$, and maps the terms from $ban_{\mathbb{F}}$ into $dom(Core(\mathcal{T}, \mathbb{F}))$, which means into $dom(\mathbb{C}_{\mathbb{D}})$.

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

10 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 $Ch(\mathcal{T}, \mathbb{D})$ 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:*

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

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 in [17]). 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:*

$$E(x, y) \Rightarrow \exists x', y'\ R(x, y, x', y')$$

$$R(x, y, x', y'), E(y, z) \Rightarrow \exists z'\ R(y, z, y', z')$$

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 $\bar{a} \in dom(\mathbb{D})^{|\bar{a}|}$ then $Ch_{|\phi(\bar{y})|} \models \phi(\bar{a})$. 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 $E(a_1, a_2), E(a_2, a_3) \ldots, E(a_n, a_1)$. The degree of this instance is 2. And there are atoms in $Ch_n(\mathcal{T}_c, \mathbb{D}_n)$ that are not in $Ch_n(\mathcal{T}_c, \mathbb{F})$ 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 $rew(\Psi)$ 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 $rew(\Psi)$ can be exponential in the size of Ψ.*

Recall the notion of backwards shy theories [18] – these are BDD theories such that, for every query $\psi(\bar{y})$ if $\phi(\bar{y}) \in rew(\psi(\bar{y}))$ 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 $c, c' \in dom(\mathbb{D})$, and any $n \in \mathbb{N}$ if $dist_{Ch(\mathcal{T},\mathbb{D})}(c, c') \leq n$ then $dist_{\mathbb{D}}(c, c') \leq d_{\mathcal{T}} n$.*

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 $Ch(\mathcal{T}, \mathbb{D})$ might lead through atoms not containing any constants from the original $\mathbb{D}$. Nevertheless:

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.*

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

(loop) $\quad true \Rightarrow \exists x\; R(x,x), G(x,x)$

(pins) $\quad \forall x (true \Rightarrow \exists z, z'\; R(x,z), G(x,z'))$

(grid) $\quad R(x,x'), G(x,u), G(u,u') \Rightarrow \exists z\; R(u',z), G(x',z)$

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 in [17]).

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 $\exists x_1 \dots x_{n-1} G(x_0, x_1), \dots, G(x_{n-1}, x_n)$ and $R^n(x_0, x_n)$ likewise. Define conjunctive queries $\phi^n_R(x,y)$ 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) $G^{2^n}(x,y) \in rew_{\mathcal{T}_d}(\phi^n_R(x,y))$ *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) $Ch(\mathcal{T}_d, \mathbb{G}^{2^n}(a,b)) \models \phi^n_R(a,b)$ and (ii) if $\mathbb{D}$ is a proper subset of $\mathbb{G}^{2^n}(a,b)$ then $Ch(\mathcal{T}_d, \mathbb{D}) \not\models \phi^n_R(a,b)$. 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 $Ch(\mathcal{T}_d, \mathbb{D})$.

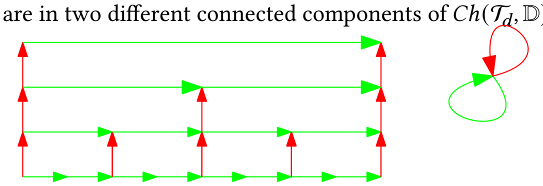

Figure 1: Fragment of $Ch(\mathcal{T}_d, \mathbb{G}^8(a_0, a_8))$ (print in colors!)

The proof of claim (A) is much harder (see Appendix B in [17]). It defines a rewriting procedure in the spirit of [13, 15], 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 the fact that it is undecidable to check if a theory is BDD (see e.g. a recent stackexchange post [16]). 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, seperating 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 [1, 8, 18], 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 [3] 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 and Ostropolski-Nalewaja was supported by the European Research Council through the ERC Consolidator Grant DeciGUT (project number 771779).

REFERENCES

[1] Alessandro Artale, Diego Calvanese, Roman Kontchakov, and Michael Zakharyaschev. 2009. The DL-Lite Family and Relations. *J. Artif. Intell. Res.* 36 (2009), 1–69.

[2] Jean-François Baget, Fabien Garreau, Marie-Laure Mugnier, and Swan Rocher. 2014. Revisiting Chase Termination for Existential Rules and their Extension to Nonmonotonic Negation. *CoRR* abs/1405.1071 (2014).

[3] Jean-François Baget, Michel Leclère, Marie-Laure Mugnier, and Eric Salvat. 2011. On rules with existential variables: Walking the decidability line. *Artif. Intell.* 175, 9-10 (2011), 1620–1654.

[4] Pablo Barceló, Gerald Berger, Carsten Lutz, and Andreas Pieris. 2018. First-Order Rewritability of Frontier-Guarded Ontology-Mediated Queries. In *Proc. of the 27th International Joint Conference on Artificial Intelligence, IJCAI 2018,*, Jérôme Lang (Ed.). ijcai.org, 1707–1713.

[5] Pierre Bourhis, Michel Leclère, Marie-Laure Mugnier, Sophie Tison, Federico Ulliana, and Lily Gallois. 2019. Oblivious and Semi-Oblivious Boundedness for Existential Rules. In *Proc. of the 28th International Joint Conference on Artificial Intelligence, IJCAI 2019,*, Sarit Kraus (Ed.). ijcai.org, 1581–1587.

[6] Andrea Calì, Georg Gottlob, and Thomas Lukasiewicz. 2009. Datalog$^{\pm}$: a unified approach to ontologies and integrity constraints. In *Proc. of the 12th International Conference on Database Theory, ICDT 2009, (ACM International Conference Proceeding Series)*, Ronald Fagin (Ed.), Vol. 361. ACM, 14–30.

[7] Andrea Calì, Georg Gottlob, and Thomas Lukasiewicz. 2012. A general Datalog-based framework for tractable query answering over ontologies. *J. Web Semant.* 14 (2012), 57–83.

[8] Andrea Calì, Georg Gottlob, and Andreas Pieris. 2010. Advanced Processing for Ontological Queries. *Proc. VLDB Endow.* 3, 1 (2010), 554–565.

[9] Cristina Civili and Riccardo Rosati. 2015. On the first-order rewritability of conjunctive queries over binary guarded existential rules. In *Proc. of the 30th Italian Conference on Computational Logic, (CEUR Workshop Proceedings)*, Davide Ancona, Marco Maratea, and Viviana Mascardi (Eds.), Vol. 1459. CEUR-WS.org, 25–30.

[10] Stathis Delivorias, Michel Leclère, Marie-Laure Mugnier, and Federico Ulliana. 2018. On the k-Boundedness for Existential Rules. In *Proc. of 2nd International Joint Conference on Rules and Reasoning, RuleML+RR 2018, (Lecture Notes in Computer Science)*, Christoph Benzmüller, Francesco Ricca, Xavier Parent, and Dumitru Roman (Eds.), Vol. 11092. Springer, 48–64.

[11] Alin Deutsch, Alan Nash, and Jeffrey B. Remmel. 2008. The chase revisited. In *Proc. of the 27th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2008,*, Maurizio Lenzerini and Domenico Lembo (Eds.). ACM, 149–158.

[12] Haim Gaifman, Harry G. Mairson, Yehoshua Sagiv, and Moshe Y. Vardi. 1993. Undecidable Optimization Problems for Database Logic Programs. *J. ACM* 40, 3 (1993), 683–713.

[13] Georg Gottlob, Giorgio Orsi, and Andreas Pieris. 2014. Query Rewriting and Optimization for Ontological Databases. *ACM Trans. Database Syst.* 39, 3 (2014), 25:1–25:46.

[14] Pavol Hell and Jaroslav Nesetril. 1992. The core of a graph. *Discret. Math.* 109, 1-3 (1992), 117–126.

[15] Mélanie König, Michel Leclère, Marie-Laure Mugnier, and Michaël Thomazo. 2015. Sound, complete and minimal UCQ-rewriting for existential rules. *Semantic Web* 6, 5 (2015), 451–475.

[16] Markus Krötzsch. 2020. 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.

[17] Piotr Ostropolski-Nalewaja, Jerzy Marcinkowski, David Carral, and Sebastian Rudolph. 2022. A Journey to the Frontiers of Query Rewritability (Extended Technical Report). In *https://hal.archives-ouvertes.fr/hal-03599601v1*.

[18] Michaël Thomazo. 2013. *Conjunctive Query Answering Under Existential Rules - Decidability, Complexity, and Algorithms.* Ph.D. Dissertation. Montpellier 2 University, France.