

On Separability of Ontological Constraints

Andrea Cali^{1,3}, Marco Console², and Riccardo Frosini²

¹Dept. of Computer Science and Inf. Systems, Birkbeck University of London, UK

²Dip. di Informatica e Sistemistica, Università di Roma “La Sapienza”, Italy

³Oxford-Man Institute of Quantitative Finance, University of Oxford, UK

andrea@dcs.bbk.ac.uk

{console.marco,frosini.riccardo}@gmail.com

Abstract. When data schemata are enriched with expressive constraints that aim at representing the domain of interest, in order to answer queries one needs to consider the logical theory consisting of both the data and the constraints. Query answering in such a context is called ontological query answering. Commonly adopted database constraints in this field are tuple-generating dependencies (TGDs) and equality-generating dependencies (EGDs). It is well known that their interaction leads to intractability or undecidability of query answering even in the case of simple subclasses. Several conditions have been found to guarantee *separability*, that is lack of interaction, between TGDs and EGDs. Separability makes EGDs (mostly) irrelevant for query answering and therefore often guarantees tractability, *as long as the theory is satisfiable*. In this paper we review the two notions of separability found in the literature, as well as several syntactic conditions that are sufficient to prove them. We then shed light on the issue of satisfiability checking, showing that under a sufficient condition called *deep separability* it can be done by considering the TGDs only. We show that, fortunately, in the case of TGDs and EGDs, separability implies deep separability. This result generalizes several analogous ones, proved *ad hoc* for particular classes of constraints. Applications include the class of *sticky TGDs* and EGDs, for which we provide a syntactic separability condition which extends the analogous one for *linear TGDs*; preliminary experiments show the feasibility of query answering in this case.

1 Introduction

When a database D is equipped with an *ontology* Σ , that is, a knowledge base consisting of constraints that express relevant properties of the underlying domain, queries are not answered merely against the database instance D , but against the logical theory $D \cup \Sigma$. Several languages have been proposed for ontologies, with different computational properties. The *DL-Lite* family [2, 15, 23] has the advantage of a low (AC_0 , which is contained in LOGSPACE) data complexity of conjunctive query answering and of knowledge base satisfiability. The well-known Entity-Relationship (ER) [17] model has recently gained importance in ontology specification, due to the fact that it is comprehensible to

theorists and practitioners, while having good expressive power. The ER^\pm family of ER-like languages [12], in particular, comprises several tractable (w.r.t. conjunctive query answering) ontology languages, which properly generalize the main languages of the DL-Lite family. Another relevant, more general class of ontology languages, is the *Datalog* $^\pm$ family, that is, a family of rule-based languages derived from Datalog (see, e.g., [9] and references therein) whose rules are (function-free) Horn rules, possibly with existentially quantified variables in the head, called *tuple-generating dependencies (TGDs)*, enriched with functionality constraints in the form of *equality-generating dependencies (EGDs)*, and *negative constraints*, a form of denial constraints. The simplest and least expressive Datalog $^\pm$ languages are able to properly capture the most prominent ontology languages, including most of the DL-lite family.

In this paper we focus on the interaction between TGDs and EGDs. A TGD is a first-order implication that forces the existence of tuples under certain conditions, and it is of the form $\forall \mathbf{X} \forall \mathbf{Y} \varphi(\mathbf{X}, \mathbf{Y}) \rightarrow \exists \mathbf{Z} \psi(\mathbf{X}, \mathbf{Z})$, where $\varphi(\mathbf{X}, \mathbf{Y})$ and $\psi(\mathbf{X}, \mathbf{Z})$ are conjunctions of atoms over a relational schema. An EGD forces equality of values under certain conditions, and it is of the form $\forall \mathbf{X} \varphi(\mathbf{X}) \rightarrow X_i = X_j$, where $\varphi(\mathbf{X})$ is a conjunction of atoms over a relational schema, and $\{X_i, X_j\} \subseteq \mathbf{X}$. To answer a query Q over an instance D and a set Σ of TGDs and EGDs, we could in principle expand D according to Σ , inferring all the entailed additional knowledge, and then evaluate Q against the obtained instance¹. In doing so, unknown values (those corresponding to the existentially-quantified variables of TGDs) will be represented by *labelled nulls*, which are a sort of placeholder. EGDs play a role in answering as they can enforce equalities; if an equality between two distinct constants is enforced, the process fails as the theory is unsatisfiable.

Example 1. Consider the following set Σ of TGDs and EGDs (we omit universal quantifiers to avoid clutter):

$$\begin{array}{ll} \sigma_1 : r_1(X) \rightarrow \exists Y r_2(X, Y) & \sigma_4 : r_4(X, Y) \rightarrow r_5(X, Y) \\ \sigma_2 : r_2(X, Y) \rightarrow r_3(X, Y) & \sigma_5 : r_5(X, Y) \rightarrow r_2(X, Y) \\ \sigma_3 : r_3(X, Y) \rightarrow r_4(X, Y) & \eta : r_3(X, Y), r_3(X, Z) \rightarrow Y = Z \end{array}$$

Notice that η is a key dependency, imposing that atoms in r_3 have unique values on their first attribute. Now, let us take $D = \{r_1(a), r_3(a, b)\}$ and the (ground) Boolean conjunctive query Q defined as $q \leftarrow r_2(a, b)$ (q is a propositional predicate), which simply asks whether $r_2(a, b)$ holds. Let us answer Q by expanding D according to Σ . We first obtain $r_2(a, \zeta)$ from σ_1 , where ζ is a null. From σ_3 we get $r_4(a, b)$. We then add $r_3(a, \zeta)$ from σ_2 ; at this point we need to apply η and enforce $\zeta = b$, so that $r_3(a, \zeta)$ becomes $r_3(a, b)$. Due to η , therefore, the query has positive answer, written $D \cup \Sigma \models Q$. However, even in the absence of η , Q would be answered positively. In fact, if we proceed with the expansion we get

¹ This is not a query answering procedure in general, as the expansion (which is called *chase*; see Section 2) can be infinite; however, the chase is a key tool for studying the query answering problem, as it will be clear in Section 2.

$r_5(a, b)$ from σ_4 and finally $r_2(a, b)$ from σ_5 . Therefore we would have had the same result without η . This can be shown to hold for every query Q ; therefore, *provided that the theory $D \cup \Sigma$ is satisfiable* (that is, the expansion does not fail), we can answer every query, for every D , by considering the TGDs only. This property is called *separability*, and it defines a form of lack of interaction (hence the name) between EGDs and TGDs in a certain set.

In this paper we provide the following contributions:

1. We illustrate the two notions of separability (the old and the new one) found in the literature, relating them to each other, and to the several syntactic conditions, sufficient to separability (one of the two notions), that have been proposed in the literature in the last decade or so.
2. We consider the problem of checking satisfiability of a theory consisting of an instance and a set of TGDs and EGDs. This can be done by answering suitable Boolean Conjunctive queries, and this has been shown in several cases. However, the proof that this technique works has always been *ad hoc*, depending on the particular class of TGDs and EGDs. Here, we show that this technique for the satisfiability check is sound and complete if a semantic condition, which we call *deep separability*, holds. We show that, fortunately, (new) separability implies deep separability. This way, we unify several known results under a single general theorem, and we also provide a useful tool for all separable classes that are yet to be discovered.
3. Finally, we mention, as an application, a syntactic condition which is sufficient for separability of *sticky sets of TGDs* [10] and EGDs, we hint at our preliminary experiment with a prototype system, and we discuss some open problems.

2 Preliminaries

In this section we recall some basics on databases, queries, tuple-generating dependencies, equality-generating dependencies, and the chase procedure.

General. We define the following pairwise disjoint (infinite) sets of symbols: (i) a set Γ of *constants* (constitute the “normal” domain of a database), (ii) a set Γ_N of *labeled nulls* (used as placeholders for unknown values, and thus can be also seen as variables), and (iii) a set Γ_V of *variables* (used in queries and dependencies). Different constants represent different values (*unique name assumption*), while different nulls may represent the same value. A lexicographic order is defined on $\Gamma \cup \Gamma_N$, such that every value in Γ_N follows all those in Γ . Throughout the paper, we denote by \mathbf{X} the sequence of variables X_1, \dots, X_k , where $k \geq 0$. Also, let $[n]$ be the set $\{1, \dots, n\}$, for any integer $n \geq 1$.

A *relational schema* \mathcal{R} (or simply *schema*) is a set of *relational symbols* (or *predicates*), each with its associated arity. A *position* $r[i]$ (in a schema \mathcal{R}) is identified by a predicate $r \in \mathcal{R}$ and its i -th argument (or attribute). A *term* t is a constant, null, or variable. An *atomic formula* (or simply *atom*) has the form $r(t_1, \dots, t_n)$, where r is an n -ary relation, and t_1, \dots, t_n are terms. Conjunctions

of atoms are often identified with the sets of their atoms. A *database (instance)* D for a schema \mathcal{R} is a (possibly infinite) set of atoms of the form $r(\mathbf{t})$ (a.k.a. *facts*), where r is an n -ary predicate of \mathcal{R} , and $\mathbf{t} \in (\Gamma \cup \Gamma_N)^n$. We denote as $r(D)$ the set $\{\mathbf{t} \mid r(\mathbf{t}) \in D\}$.

A *substitution* from one set of symbols S_1 to another set of symbols S_2 is a function $h : S_1 \rightarrow S_2$ defined as follows: (i) \emptyset is a substitution (empty substitution), (ii) if h is a substitution, then $h \cup \{X \rightarrow Y\}$ is a substitution, where $X \in S_1$ and $Y \in S_2$, and h does not already contain some $X \rightarrow Z$ with $Y \neq Z$. If $X \rightarrow Y \in h$ we write $h(X) = Y$. A *homomorphism* from a set of atoms A_1 to a set of atoms A_2 , both over the same schema \mathcal{R} , is a substitution $h : \Gamma \cup \Gamma_N \cup \Gamma_V \rightarrow \Gamma \cup \Gamma_N \cup \Gamma_V$ such that: (i) if $t \in \Gamma$, then $h(t) = t$, and (ii) if $r(t_1, \dots, t_n)$ is in A_1 , then $h(r(t_1, \dots, t_n)) = r(h(t_1), \dots, h(t_n))$ is in A_2 . If there are homomorphisms from A_1 to A_2 and vice-versa, then A_1 and A_2 are *homomorphically equivalent*. The notion of homomorphism naturally extends to conjunctions of atoms. Given a set of symbols S , two atoms \underline{a}_1 and \underline{a}_2 are *S -isomorphic* iff there exists a bijection h such that $h(\underline{a}_1) = \underline{a}_2$, $h^{-1}(\underline{a}_2) = \underline{a}_1$, and $h(X) = X$, for each $X \in S$.

Conjunctive Queries. A *conjunctive query (CQ)* Q of arity n over a schema \mathcal{R} has the form $q(\mathbf{X}) \leftarrow \varphi(\mathbf{X}, \mathbf{Y})$, where $\varphi(\mathbf{X}, \mathbf{Y})$ is a conjunction of atoms over \mathcal{R} , \mathbf{X} and \mathbf{Y} are sequences of variables or constants in Γ , and q is an n -ary predicate that does not occur in \mathcal{R} . $\varphi(\mathbf{X}, \mathbf{Y})$ is called the *body* of q , denoted as $body(q)$. A *Boolean CQ (BCQ)* is a CQ of zero arity.

The *answer* to an n -ary CQ Q of the form $q(\mathbf{X}) \leftarrow \varphi(\mathbf{X}, \mathbf{Y})$ over a database D , denoted as $Q(D)$, is the set of all n -tuples $\mathbf{t} \in \Gamma^n$ for which there exists a homomorphism $h : \mathbf{X} \cup \mathbf{Y} \rightarrow \Gamma \cup \Gamma_N$ such that $h(\varphi(\mathbf{X}, \mathbf{Y})) \subseteq D$ and $h(\mathbf{X}) = \mathbf{t}$. A BCQ has only the empty tuple $\langle \rangle$ as possible answer, in which case it is said that has positive answer. Formally, a BCQ Q has *positive* answer over D , denoted as $D \models Q$, iff $\langle \rangle \in Q(D)$, or, equivalently, $Q(D) \neq \emptyset$.

Dependencies. Given a schema \mathcal{R} , a *tuple-generating dependency (TGD)* σ over \mathcal{R} is a first-order formula $\forall \mathbf{X} \forall \mathbf{Y} \varphi(\mathbf{X}, \mathbf{Y}) \rightarrow \exists \mathbf{Z} \psi(\mathbf{X}, \mathbf{Z})$, where $\varphi(\mathbf{X}, \mathbf{Y})$ and $\psi(\mathbf{X}, \mathbf{Z})$ are conjunctions of atoms over \mathcal{R} , called the *body* and the *head* of σ , denoted as $body(\sigma)$ and $head(\sigma)$, respectively. Henceforth, to avoid notational clutter, we will omit the universal quantifiers in TGDs. Such σ is satisfied by a database D for \mathcal{R} iff, whenever there exists a homomorphism h such that $h(\varphi(\mathbf{X}, \mathbf{Y})) \subseteq D$, there exists an extension h' of h (i.e., $h' \supseteq h$) where $h'(\psi(\mathbf{X}, \mathbf{Z})) \subseteq D$.

Inclusion dependencies (IDs) are the simplest class of TGDs; they have just one body-atom and one head-atom, without repetition of variables neither in the body nor in the head. They are usually denoted as $r_1[\mathbf{A}] \subseteq r_2[\mathbf{B}]$, where \mathbf{A} and \mathbf{B} are sets of attributes (arguments) of r_1 and r_2 respectively. Such an ID expresses that for each r_1 -atom, its values in \mathbf{A} have to appear in \mathbf{B} some r_2 atom. How this is represented by a TGD is obvious.

An *equality-generating dependency (EGD)* η over \mathcal{R} is a first-order formula $\forall \mathbf{X} \varphi(\mathbf{X}) \rightarrow X_i = X_j$, where $\varphi(\mathbf{X})$ is a conjunction of atoms over \mathcal{R} , called the *body* and denoted as $body(\eta)$, and $X_i = X_j$ is an equality among variables of \mathbf{X} .

Henceforth, for brevity, we will omit the universal quantifiers in EGDs. Such η is satisfied by a database D for \mathcal{R} iff, whenever there exists a homomorphism h such that $h(\varphi(\mathbf{X})) \subseteq D$, then $h(X_i) = h(X_j)$. Special cases of EGDs are *functional dependencies (FDs)* of the form $r : \mathbf{A} \rightarrow \mathbf{B}$, where \mathbf{A} and \mathbf{B} are sets of attributes of r ; such a dependency is satisfied in an instance D if there are no two atoms in D that have the same values on \mathbf{A} but different values on \mathbf{B} . If the union of \mathbf{A} and \mathbf{B} is the whole set of attributes of r , the FD is said to be a *key dependency (KD)*. How a FD is expressed by EGDs is obvious.

CQ Answering under Dependencies. We now define the notion of *query answering* under TGDs and EGDs. Given a database D for \mathcal{R} , and a set Σ of TGDs and EGDs over \mathcal{R} , the *models* of D w.r.t. Σ , denoted as $\text{mods}(D, \Sigma)$, is the set of all databases B such that $B \models D \cup \Sigma$, i.e., $B \supseteq D$ and B satisfies Σ . The *answer* to a CQ Q w.r.t. D and Σ , denoted as $\text{ans}(Q, D, \Sigma)$, is the set $\{\mathbf{t} \mid \mathbf{t} \in Q(B), \text{ for each } B \in \text{mods}(D, \Sigma)\}$. The *answer* to a BCQ Q w.r.t. D and Σ is *positive*, denoted as $D \cup \Sigma \models Q$, iff $\text{ans}(Q, D, \Sigma) \neq \emptyset$. Note that query answering under general TGDs and EGDs is undecidable. In fact, this is true even in extremely simple cases such as that of IDs and keys [16].

We recall that the two problems of CQ and BCQ answering under TGDs and EGDs are LOGSPACE-equivalent [6]. Moreover, it is easy to see that the query output tuple problem (as a decision version of CQ answering) and BCQ evaluation are AC_0 -reducible to each other. Thus, we henceforth focus only on the BCQ answering problem.

The Chase Procedure. The *chase procedure* (or simply *chase*) is a fundamental algorithmic tool introduced for checking implication of dependencies [21], and later for checking query containment [20]. Informally, the chase is a process of repairing a database w.r.t. a set of dependencies so that the resulted database satisfies the dependencies. We shall use the term chase interchangeably for both the procedure and its result. The chase works on an instance through the so-called TGD and EGD *chase rules*. The TGD chase rule comes in two different equivalent fashions: *oblivious* and *restricted* [6], where the restricted one repairs TGDs only when they are not satisfied. In the sequel, we focus on the oblivious one for better technical clarity. The chase rules follow.

TGD CHASE RULE. Consider a database D for a schema \mathcal{R} , and a TGD σ of the form $\varphi(\mathbf{X}, \mathbf{Y}) \rightarrow \exists \mathbf{Z} \psi(\mathbf{X}, \mathbf{Z})$ over \mathcal{R} . If σ is *applicable* to D , i.e., there exists a homomorphism h such that $h(\varphi(\mathbf{X}, \mathbf{Y})) \subseteq D$, then: (i) define $h' \supseteq h$ such that $h'(Z_i) = z_i$, for each $Z_i \in \mathbf{Z}$, where $z_i \in \Gamma_N$ is a “fresh” labeled null not introduced before, and following lexicographically all those introduced so far, and (ii) add to D the set of atoms $h'(\psi(\mathbf{X}, \mathbf{Z}))$, if not already in D .

EGD CHASE RULE. Consider a database D for a schema \mathcal{R} , and an EGD η of the form $\varphi(\mathbf{X}) \rightarrow X_i = X_j$ over \mathcal{R} . If η is *applicable* to D , i.e., there exists a homomorphism h such that $h(\varphi(\mathbf{X})) \subseteq D$ and $h(X_i) \neq h(X_j)$, then: (i) if $h(X_i)$ and $h(X_j)$ are both constants of Γ , then there is a *hard violation* of η , and the chase *fails*, otherwise (ii) replace each occurrence of $h(X_j)$ with $h(X_i)$, if $h(X_i)$ precedes $h(X_j)$ in the lexicographic order, or vice-versa otherwise.

Given a database D and a set of dependencies $\Sigma = \Sigma_T \cup \Sigma_E$, where Σ_T are TGDs and Σ_E are EGDs, the chase algorithm for D and Σ consists of an exhaustive application of the chase rules in a breadth-first fashion, which leads to a (possibly infinite) database. Roughly, the chase of D w.r.t. Σ , denoted as $\text{chase}(D, \Sigma)$, is the (possibly infinite) instance constructed by iteratively applying (i) the TGD chase rule once, and (ii) the EGD chase rule as long as it is applicable (i.e., until a fixed point is reached). A formal definition of the chase algorithm is given, e.g., in [10].

Example 2 (from [14]). Let $\mathcal{R} = \{r, s\}$. Consider the set Σ of TGDs and EGDs over \mathcal{R} constituted by the TGDs $\sigma_1 = r(X, Y) \rightarrow \exists Z r(Z, X), s(Z)$ and $\sigma_2 = r(X, Y) \rightarrow r(Y, X)$, and the EGD $\eta = r(X, Y), r(X', Y) \rightarrow X = X'$. Let D be the database for \mathcal{R} consisting of the single atom $r(a, b)$. During the construction of $\text{chase}(D, \Sigma)$ we first apply σ_1 , and we add the atoms $r(z_1, a), s(z_1)$, where z_1 is a “fresh” null of Γ_N . Moreover, σ_2 is applicable and we add the atom $r(b, a)$. Now, the EGD η is applicable and we replace each occurrence of z_1 with the constant b ; thus, we get the atom $s(b)$. We continue by applying exhaustively the chase rules as explained above. ■

The (possibly infinite) chase of D w.r.t. Σ is a *universal model* of D w.r.t. Σ , i.e., for each database $B \in \text{mods}(D, \Sigma)$, there exists a homomorphism from $\text{chase}(D, \Sigma)$ to B [19, 18]. Using this fact it can be shown that the chase is a formal tool for query answering under TGDs and EGDs. In particular, given a BCQ Q , $D \cup \Sigma \models Q$ iff $\text{chase}(D, \Sigma) \models Q$, providing that the chase does not fail. If the chase fails, then the set of models of D w.r.t. Σ is empty, and $D \cup \Sigma \models Q$ trivially.

3 What is separability?

In this section we provide a brief review of the interaction of TGDs and EGDs in ontological query answering.

It is well known that the interaction of TGDs and EGDs leads to undecidability of query answering; this happens even in simple sub-classes of constraints, such as key and inclusion dependencies [4, 13]. It is therefore useful to identify classes of constraints which do not suffer from this “harmful” interaction. To this aim, a key condition is that of *separability* [12, 11], which we give below.

Definition 1 (Separability). *Consider a set Σ_T of TGDs over a schema \mathcal{R} , and a set Σ_E of EGDs over \mathcal{R} . We say that the set $\Sigma = \Sigma_T \cup \Sigma_E$ is separable if, for every database D for \mathcal{R} , either $\text{chase}(D, \Sigma)$ fails, or, $\text{chase}(D, \Sigma) \models Q$ iff $\text{chase}(D, \Sigma_T) \models Q$, for every BCQ Q over \mathcal{R} .*

Notably, there is another definition of separability in the literature, which is adopted in [4, 13, 8]. Such a definition is similar to the above Definition 1, but with a difference which makes it stronger. For the sake of completeness, we give such a definition below.

Definition 2 (Old separability). Consider a set Σ_T of TGDs over a schema \mathcal{R} , and a set Σ_E of EGDs over \mathcal{R} . We say that the set $\Sigma = \Sigma_T \cup \Sigma_E$ is separable if, for every database D for \mathcal{R} , (i) if the chase fails, then $D \not\models \Sigma_E$ (D does not satisfy Σ), and (ii) if the chase does not fail, we have $\text{chase}(D, \Sigma) \models Q$ iff $\text{chase}(D, \Sigma_T) \models Q$, for every BCQ Q over \mathcal{R} .

The old separability is a special case of the new separability as it enforces condition (i) (Definition 2), which we reformulate below, calling it *EGD-stability*.

Definition 3 (EGD-stability). Consider a set Σ_T of TGDs over a schema \mathcal{R} , and a set Σ_E of EGDs over \mathcal{R} . We say that the set $\Sigma = \Sigma_T \cup \Sigma_E$ is EGD-stable if, for every instance D for \mathcal{R} , $D \models \Sigma_E$ implies that $\text{chase}(D, \Sigma)$ does not fail.

EGD stability guarantees that the satisfiability of $D \cup \Sigma$ (i.e., the existence, or non-failure, of $\text{chase}(D, \Sigma)$) can be determined by merely checking whether $D \models \Sigma_E$. In fact, the difference between the old and the new definitions of separability resides in the fact that under the new separability condition, a chase failure does not imply that the initial database D violates some EGD in Σ_E , while this holds under the old definition. The satisfiability check is a fundamental step in query answering (see Section 3.2), and EGD stability guarantees it can be easily done. However, with a more sophisticated version of the satisfiability check (see Section 4), the old separability can be relaxed to the new separability, giving raise to the discovery of new (syntactic) classes that enjoy (new) separability. Section 3.1 below provides an overview of the most relevant syntactic classes in the literature.

3.1 Related Work

The notion of (old) separability was first introduced in [4, 13], in the context of inclusion dependencies (IDs) and key dependencies (KDs) (see e.g. [1]). The general idea is to define a *sufficient syntactic condition* for separability, which can be efficiently checked. This was done while extending an early class of IDs and KDs (see e.g. [1]), called *key-based*. Key-based IDs and KDs were proposed in the milestone work by Johnson and Klug [20], and they are in fact separable, though not defined explicitly as such. Key-based IDs and KDs are defined as follows²: (i) for each relational predicate r , there is *only one* KD defined on it; (ii) for each ID $r[\mathbf{X}] \subseteq s[\mathbf{Y}]$, where \mathbf{X} and \mathbf{Y} are set of attributes of r and s respectively (see [1] for the notation), (ii.a) \mathbf{X} is *disjoint* from any key set of attributes for r , and (ii.b) \mathbf{Y} is *properly contained* in the set of key attributes for s .

The more general class of *non-key-conflicting (NKC) IDs*, with respect to a set of KDs, is defined as follows: (i) for each relational predicate r , there is *only one* KD defined on it; (ii) for every ID $r[\mathbf{X}] \subseteq s[\mathbf{Y}]$, \mathbf{Y} is *not a proper* superset of the set of key attributes for s (if such set exists). The class of KDs and NKC

² The definition in the original paper is different but equivalent; we choose a definition which is more clear in the context of this paper.

IDs is separable, and it properly captures the well known class of *foreign-key dependencies*.

The class of KDs and non-key conflicting IDs was generalized in [8] to general TGDs, which are assumed to have a single atom in the head, without loss of generality. The idea is analogous, and the condition is as follows: (i) for each relational predicate r , there is *only one KD* defined on it; (ii) each existentially quantified variable in the head of a TGD must occur only once; (iii) for each TGD $\sigma = \varphi(\mathbf{X}, \mathbf{Y}) \rightarrow \exists \mathbf{Z} r(\mathbf{X}, \mathbf{Z})$, the set of \mathbf{X} -attributes of $r(\mathbf{X}, \mathbf{Z})$ is not a proper superset of the set of key attributes of r .

In [10], the class of non-key-conflicting TGDs was straightforwardly extended to treat functional dependencies (FDs) rather than keys.

The literature so far discussed deals with sufficient *syntactic* conditions that guarantee, in fact, EGD-stability. However, there are classes of TGDs and EGDs such that EGDs are triggered in the chase, but which enjoy separability.

Example 3. Consider the set of TGDs and EGDs in Example 1. It is separable, but it is not EGD-stable. In fact, if we consider the instance $D' = \{r_3(a, b), r_2(a, c)\}$, we have that $D' \models \Sigma_E$ but $D \cup \Sigma_E$ is unsatisfiable because the chase fails due to a hard violation.

This leads in [12] to the definition of separability as in Definition 1 in this paper. This work deals with IDs and KDs which express an expressive variant of the Entity-Relationship model [17]. By means of graph-related properties, *necessary and sufficient* syntactic conditions for separability are provided, thus defining useful tractable classes of constraints.

The case of *linear TGDs* (TGDs with a single body-atom) and KDs was later considered in [14, 11]. In these works, sufficient syntactic conditions are proposed that guarantee separability, without imposing non-egd-triggerability. The conditions are quite involved and they make use of backward-resolution. Interestingly, the complexity of checking the syntactic condition, called *non-conflict* condition, is the same as that of query answering, that is PSPACE-complete in combined complexity.

At this point, we considered separability but not the problem of satisfiability. While separability allows us to ignore EGDs in the case of satisfiable theories, the problem of deciding the satisfiability remains, as well as that of determining its complexity. This will be the subject of next section.

3.2 Query Answering under Separable Constraints

In the case of a set $\Sigma = \Sigma_T \cup \Sigma_E$ of separable TGDs and EGDs, such that BCQ answering under Σ_T is decidable, given an instance D and a BCQ Q , to decide whether $D \cup \Sigma \models Q$, the following steps are needed:

1. Check whether the chase fails, that is, whether $D \cup \Sigma$ is satisfiable; if $D \cup \Sigma$ is unsatisfiable, then trivially $D \cup \Sigma \models Q$ (“*Ex falso quodlibet*”).
2. If $D \cup \Sigma$ is satisfiable, then by Definition 1 we know $\text{chase}(D, \Sigma) \models Q$ iff $\text{chase}(D, \Sigma_T) \models Q$, therefore we check whether $\text{chase}(D, \Sigma_T) \models Q$.

Apart from the complexity of the satisfiability check, we have that the complexity of query answering is the same as that of answering under TGDs only, which is a highly desirable property. In the case of EGD-stable constraints, the satisfiability check amounts simply to checking whether $D \models \Sigma_E$, which can be done in NP, and in PTIME (or better, in the even lower complexity class AC₀) if we consider Σ_E fixed. However, in the cases of separable but not EGD-stable constraints, the problem is to be addressed differently; this will be the subject of Section 4.

4 Separability and Satisfiability

In this section we address the problem of deciding whether, given an instance D and a set Σ of *separable* TGDs and EGDs, $D \cup \Sigma$ is satisfiable. As seen in Section 3, this preliminary check is needed, in the case of separability, before one proceeds to answer a BCQ Q by taking into account the TGDs only.

The satisfiability check is done as in [8, 12], by encoding *hard violations* of EGDs as a set \mathcal{Q}_V of Boolean conjunctive queries. Given a separable set $\Sigma = \Sigma_T \cup \Sigma_E$, and an instance D , we have that satisfiability holds if and only if, for each $Q \in \mathcal{Q}_V$ we have $D \cup \Sigma_T \not\models Q$ or, equivalently, if and only if all queries in \mathcal{Q}_V have negative answer when evaluated against D and Σ_T (TGDs only). The encoding is done as follows. First, we need to add auxiliary facts to D . For each pair of *distinct* constants c_1, c_2 appearing in D as arguments, we add the facts $neq(c_1, c_2)$ and $neq(c_2, c_1)$ to D , where $neq/2$ is an auxiliary predicate expressing that two constants are distinct. Then, for each EGD η of the form $\phi(\mathbf{X}) \rightarrow X_i = X_j$, with X_i and X_j in \mathbf{X} , we construct the BCQ Q_η as follows (quantifiers omitted for brevity): $Q_\eta : q \leftarrow \phi(\mathbf{X}), neq(X_i, X_j)$, where q is a propositional predicate. The set \mathcal{Q}_V encoding hard violations is then defined as the set of all Q_η so constructed, or equivalently $\mathcal{Q}_V = \bigcup_{\eta \in \Sigma_E} Q_\eta$. It is not difficult to prove that if $D \cup \Sigma \models Q_\eta$ then $chase(D, \Sigma)$ fails, and therefore $D \cup \Sigma$ is unsatisfiable. However, it still remains to determine whether $D \cup \Sigma \models Q_\eta$. Notice that in the case of non-egd-triggerable constraints we do not have this problem, as we merely need to check whether $D \models \Sigma_E$. In principle, separability (see Definition 1) does not tell us anything about the cases when the chase fails. However, we are still in luck because we can evaluate the (Boolean) conjunctive queries in \mathcal{Q}_V against $D \cup \Sigma_T$ rather than $D \cup \Sigma$. This is proved, for instance, in [12, 14], but *ad hoc*, by using the properties of the particular class of constraints involved. Here we show a general condition that allows us to perform the satisfiability check as above. We first need to introduce a preliminary condition, called *deep separability*, which is apparently more restrictive than separability (we shall then prove that it is implied by separability). We start by denoting by $chase^k(D, \Sigma)$ the result of the first k chase steps under Σ applied to an instance D , where a step is either an EGD or a TGD application.

Definition 4. *Let Σ be a set of TGDs and EGDs, with $\Sigma = \Sigma_T \cup \Sigma_E$, where the dependencies in Σ_T are TGDs and those in Σ_E are EGDs. We say that Σ is*

deeply separable if, for each integer $k \geq 0$, for each instance D with values in $\Gamma \cup \Gamma_N$, and for each Boolean conjunctive query Q , the following holds: if $\text{chase}^k(D, \Sigma)$ exists, then $\text{chase}^k(D, \Sigma) \models Q$ implies $\text{chase}(D, \Sigma_T) \models Q$.

The notion of *deep separability*, intuitively, guarantees that, at any step before a possible failure, the chase does not entail any atoms that are not entailed by the chase computed according to Σ_T only. We now come to the main result about deeply separable TGDs and EGDs. The result states that in the case of deep separability the satisfiability check done by answering suitable queries as above is correct and complete.

Theorem 1. *Let Σ be a set of deeply separable TGDs and EGDs, with $\Sigma = \Sigma_T \cup \Sigma_E$, where the dependencies in Σ_T are TGDs and those in Σ_E are EGDs. Let $\mathcal{Q}_V = \{Q_1, \dots, Q_m\}$ be the set of BCQs encoding violations of the EGDs in Σ_E as from the above construction. Then we have that $\text{chase}(D, \Sigma_T) \models Q_i$ for some $i \in \{1, \dots, m\}$ if and only if $\text{chase}(D, \Sigma)$ fails (or equivalently $D \cup \Sigma$ is unsatisfiable).*

PROOF (SKETCH). We prove the two directions of the implication separately.

“Only if”. By contradiction, assume $\text{chase}(D, \Sigma_T) \models Q_i$ for some $i \in \{1, \dots, m\}$ but $\text{chase}(D, \Sigma)$ exists. It is not difficult to show that if $\text{chase}(D, \Sigma_T) \models Q_i$ then also $\text{chase}(D, \Sigma) \models Q_i$; this holds because $\text{chase}(D, \Sigma_T)$ is a universal model for $D \cup \Sigma_T$ and $\text{chase}(D, \Sigma)$ is a model (not necessarily universal) for $D \cup \Sigma_T$; therefore there exists a homomorphism from the former to the latter. If $\text{chase}(D, \Sigma) \models Q_i$ then the chase must necessarily fail. Contradiction.

“If”. Assume $\text{chase}(D, \Sigma)$ fails at step k by violation of the EGD $\eta \in \Sigma_E$. Then, $\text{chase}^{k-1}(D, \Sigma)$ exists; moreover, it is easily seen that $\text{chase}^{k-1}(D, \Sigma) \models Q_\eta$, where $Q_\eta \in \mathcal{Q}_V$ encodes the violation of η . By the hypothesis of deep separability we have $\text{chase}(D, \Sigma_T) \models Q_\eta$, hence the claim. \square

Notice that for the “If” direction we do not need deep separability nor separability; this direction of the implication holds for general TGDs and EGDs. Finally we show that, despite the appearances, deep separability is not a stricted condition than separability. In fact, separability implies deep separability.

Theorem 2. *Let Σ be a set of TGDs and EGDs. If Σ is separable, then it is deeply separable.*

PROOF (SKETCH). We distinguish two cases.

Case 1: $\text{chase}(D, \Sigma)$ exists. In this case the claim obviously holds.

Case 2: $\text{chase}(D, \Sigma)$ fails. Assume $\text{chase}(D, \Sigma)$ fails at step k ; therefore $\text{chase}^{k-1}(D, \Sigma)$ exists. Take D_F obtained from D by replacing each constant in Γ with a fresh null from Γ_N . Obviously $\text{chase}(D_F, \Sigma)$ exists and so does $\text{chase}^{k-1}(D_F, \Sigma)$. Take any BCQ Q such that $\text{chase}^{k-1}(D_F, \Sigma) \models Q$. It is straightforwardly seen that $\text{chase}(D_F, \Sigma) \models Q$ and, due to separability, $\text{chase}(D_F, \Sigma_T) \models Q$. Since $\text{chase}(D_F, \Sigma_T) \models Q$ is obtained from $\text{chase}(D_F, \Sigma)$

by the above renaming of constants into nulls, we have $\text{chase}(D, \Sigma_T) \models Q$. Since this holds for every step $\ell \leq k - 1$, the claim is proved. \square

The following result immediately follows from the above.

Corollary 1. *For every separable set Σ of TGDs and EGDs, for every instance D and for every BCQ Q :*

- *checking unsatisfiability of $D \cup \Sigma$ has the same complexity as query answering under Σ_T alone;*
- *checking whether $D \cup \Sigma \models Q$ has the same complexity as query answering under Σ_T alone.*

Notice that we mention *unsatisfiability* rather than satisfiability because Theorem 1 shows a reduction from failure (unsatisfiability) to BCQ answering under TGDs alone.

5 Preliminary Experiments

We show here some preliminary experiments on the separability check in ontologies. First of all, let us describe the ontology language we have adopted. In [11], a sufficient condition for separability (in the “new” meaning) is provided for the case of EGDs and linear TGDs (which, we remind the reader, are TGDs with exactly one head-atom and one body-atom). We have extended such condition, with a rather straightforward adaptation, to the class of *sticky sets of TGDs* [10], a relevant class that allows for a form of joins in the body. The condition requires, as in the case of [11], a significant number of containment checks under sticky sets of TGDs in practical cases. More specifically, there is a first step, like in [11], where the body of each EGD is expanded via a (terminating) variant of backward resolution, and then for each obtained subgoal, a suitable containment check (under TGDs alone) is performed. We have built a prototype system which performs such a separability check which requires – as is obvious from this paper – also the satisfiability check. In the first tests, we considered a few known ontologies, including the GALEN medical ontology and the LUBM ontology (benchmark from Lehigh University, which describes the organizational structure of universities). The tests have been run on an Intel QuadCore I7 at 2.0 GHz and with 4 Gb of RAM; the operating system was Windows 7 Professional carrying Java Standard Edition Runtime Environment 1.7.0. The considered ontologies do not include any expressive functionality constraints, therefore we designed some ourselves, according to the semantics. In such cases, separability checks ran in extremely short time, confirming our idea that real-world functionality constraints can be efficiently checked for separability. We therefore moved on to test the system on near-worst-case schemata. We “artificially” designed some ad-hoc ontologies that, to be checked for separability, require a high number of containment checks (double exponential in the maximum predicate arity). Some values are shown in Figure 1. Notice that

Max arity	Containment checks	Time (s)
4	18	0.372
5	61	1.635
6	186	7.603
7	540	67.797
8	1450	527.608
9	4748	6191.924

Fig. 1. Experimental results for the separability check in "near-worst-case" ontologies.

execution times are still reasonable for arities that are already too large to exist in practice. Given that the check is independent of queries, and therefore it is run only in case either the constraints or the instance change, we believe the first results are promising and the separability check is generally efficient in practice.

6 Discussion

In this paper we have given an overview of the problem of separability between TGDs and EGDs in the context of ontological query answering, where queries are (Boolean) conjunctive queries. We have reviewed the two main notions of separability found in the literature, the "old" one and the "new" one, and we have clarified the difference between them. We have then addressed the issue of checking satisfiability of a set of TGDs and EGDs together with an instance, and we have shown that this can be done by merely answering suitable queries in the case of deeply separable classes of constraints. We have shown the desirable property that all separable sets of constraints are also deeply separable. We have therefore clarified and proved formally a satisfiability check which was already employed in the literature, but proved on a case-by-case basis [11, 8, 12], depending on the class of constraints. We believe that our generalisation provides a useful tool for future studies, and a better insight into the satisfiability problem. Then, we have shown some preliminary experimental results on a prototype system we have built. While the experimentation is still at an initial stage, the first results are promising, even in "pathological" cases where we designed ontologies that behave, with respect to the separability check, as the (asymptotically) worst case.

Our research, given the broad variety of expressive ontology languages captured by the formalisms we considered, has broad applications in ontology-based data access, as we provide tools for the separability check. Notice that in the case of data exchange [19], where the chase is guaranteed to terminate for each instance, general EGDs can be easily treated, and the notion of separability has little significance.

Open problems. We currently have two main open problems at hand, which seem non-trivial. First, we would like to study the decidability of determining whether a given set of TGDs and EGDs is separable. This has been proved to be undecidable for the case of arbitrary TGDs and EGDs [22]; however, the

proof relies on the fact that query answering under arbitrary TGDs (without EGDs) is undecidable [3]. It is not clear whether undecidability holds also in the cases where query answering is undecidable under TGDs and EGDs together, but decidable under TGDs alone; relevant cases are, for instance, IDs and KDs, sticky sets of TGDs and EGDs, or guarded TGDs and KDs [7]. We conjecture that determining separability is undecidable for these classes, but a proof is still to be devised. The second problem is more technical, and is determining the complexity of checking the syntactic condition for the separability of sticky sets of TGDs and EGDs. We have an obvious EXPTIME lower bound, but the upper bound is currently unknown.

Acknowledgments. We would like to thank Georg Gottlob, Andreas Pieris and the anonymous reviewers for their insightful comments on this research.

References

1. S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
2. A. Artale, D. Calvanese, R. Kontchakov, and M. Zakharyashev. The DL-Lite family and relations. *J. Artif. Intell. Res.*, 36:1–69, 2009.
3. C. Beeri and M. Y. Vardi. The implication problem for data dependencies. In *Proc. of ICALP*, pages 73–85, 1981.
4. A. Cali. *Query Answering and Optimisation in Data Integration Systems*. PhD thesis, Università di Roma “La Sapienza”, 2003.
5. A. Cali, M. Console, R. Frosini, and A. Pieris. On equality constraints in ontological query answering. Technical report, University of London, Birkbeck College, 2012. Available from the authors.
6. A. Cali, G. Gottlob, and M. Kifer. Taming the infinite chase: Query answering under expressive relational constraints. In *Proc. of KR*, pages 70–80, 2008.
7. A. Cali, G. Gottlob, and T. Lukasiewicz. A general datalog-based framework for tractable query answering over ontologies. In *Proc. of PODS*, pages 77–86, 2009.
8. A. Cali, G. Gottlob, and T. Lukasiewicz. A general datalog-based framework for tractable query answering over ontologies. *J. Web Sem.*, 2012. to appear.
9. A. Cali, G. Gottlob, T. Lukasiewicz, B. Marnette, and A. Pieris. Datalog+/-: A family of logical knowledge representation and query languages for new applications. In *Proc. of LICS*, pages 228–242, 2010.
10. A. Cali, G. Gottlob, and A. Pieris. Advanced processing for ontological queries. *PVLDB*, 3(1):554–565, 2010.
11. A. Cali, G. Gottlob, and A. Pieris. Querying conceptual schemata with expressive equality constraints. In *Proc. of ER*, pages 161–174, 2011.
12. A. Cali, G. Gottlob, and A. Pieris. Ontological query answering under expressive Entity-Relationship schemata. *Inf. Syst.*, 37(4):320–335, 2012.
13. A. Cali, D. Lembo, and R. Rosati. On the decidability and complexity of query answering over inconsistent and incomplete databases. In *Proc. of PODS*, pages 260–271, 2003.
14. A. Cali and A. Pieris. On equality-generating dependencies in ontology querying – Preliminary report. In *Proc. of AMW*, 2011.
15. D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. Tractable reasoning and efficient query answering in description logics: The DL-lite family. *J. Autom. Reasoning*, 39(3):385–429, 2007.

16. A. K. Chandra and M. Y. Vardi. The implication problem for functional and inclusion dependencies. *SIAM Journal of Computing*, 14:671–677, 1985.
17. P. P. Chen. The entity-relationship model: towards a unified view of data. *ACM TODS*, 1(1):124–131, 1995.
18. A. Deutsch, A. Nash, and J. B. Remmel. The chase revisited. In *Proc. of PODS*, pages 149–158, 2008.
19. R. Fagin, P. G. Kolaitis, R. J. Miller, and L. Popa. Data exchange: Semantics and query answering. *Theor. Comput. Sci.*, 336(1):89–124, 2005.
20. D. S. Johnson and A. C. Klug. Testing containment of conjunctive queries under functional and inclusion dependencies. *J. Comput. Syst. Sci.*, 28(1):167–189, 1984.
21. D. Maier, A. O. Mendelzon, and Y. Sagiv. Testing implications of data dependencies. *ACM Trans. Database Syst.*, 4(4):455–469, 1979.
22. A. Pieris, 2012. Personal communication.
23. A. Poggi, D. Lembo, D. Calvanese, G. De Giacomo, M. Lenzerini, and R. Rosati. Linking data to ontologies. *J. Data Semantics*, 10:133–173, 2008.