

# Tractable Query Answering over Ontologies with Datalog<sup>±</sup>

Andrea Cali<sup>2,1</sup>, Georg Gottlob<sup>1,2</sup>, and Thomas Lukasiewicz<sup>1,‡</sup>

<sup>1</sup> Computing Laboratory, University of Oxford, UK  
firstname.lastname@comlab.ox.ac.uk

<sup>2</sup> Oxford-Man Institute of Quantitative Finance, University of Oxford, UK

**Abstract.** We present a family of expressive extensions of Datalog, called Datalog<sup>±</sup>, as a new paradigm for query answering over ontologies. The Datalog<sup>±</sup> family admits existentially quantified variables in rule heads, and has suitable restrictions to ensure highly efficient ontology querying. In particular, we show that query answering under so-called *guarded Datalog<sup>±</sup>* is PTIME-complete in data complexity, and that query answering under so-called *linear Datalog<sup>±</sup>* is in AC<sub>0</sub> in data complexity. We also show how negative constraints and a general class of key constraints can be added to Datalog<sup>±</sup> while keeping ontology querying tractable. We then show that *linear Datalog<sup>±</sup>*, enriched with a special class of key constraints, generalizes the well-known *DL-Lite* family of tractable description logics. Furthermore, the Datalog<sup>±</sup> family is of interest in its own right and can, moreover, be used in various contexts such as data integration and data exchange. This work is a short version of [8].

## 1 Introduction

Ontologies play a key role in the Semantic Web [4], data modeling, and information integration [19]. Recent trends in ontological reasoning have shifted from decidability issues to tractability ones, as e.g. reflected by the work on the *DL-Lite* family of tractable description logics (DLs) [12, 22]. An important result of these works is that the main variants of *DL-Lite* are not only decidable, but that answering (unions of) conjunctive queries for these logics is in LOGSPACE, and actually in AC<sub>0</sub>, in data complexity (i.e., the complexity where both the query and the constraints are fixed), and query answering in *DL-Lite* is FO-rewritable [12]. As observed in [21], the lack of value creation makes plain Datalog not very well suited for ontological reasoning with inclusion axioms either. It is thus natural to extend Datalog in order to nicely accommodate ontological axioms and constraints such as those expressible in the *DL-Lite* family.

This paper proposes and studies variants of Datalog that are suited for efficient ontological reasoning, and, in particular, for tractable ontology-based query answering. We introduce the Datalog<sup>±</sup> family of Datalog variants, which extend plain Datalog by the possibility of existential quantification in rule heads, and by a number of other features, and, at the same time, restrict the rule syntax in order to achieve tractability. In the Datalog<sup>±</sup> family, rules are a restricted form of *tuple-generating dependencies (TGDs)* and *equality-generating dependencies (EGDs)*. More specifically, in *guarded*

---

<sup>‡</sup> Alternative address: Institut für Informationssysteme, Technische Universität Wien, Austria.

$Datalog^\pm$ , rules are *guarded TGDs (GTGDs)*, i.e., TGDs with a single atom in the body that contains all universally quantified variables; in *linear Datalog<sup>±</sup>*, rules are *linear TGDs (LTGDs)*, i.e., TGDs that have a single atom in the body. We characterize the data complexity of query answering for both the above sublanguages of  $Datalog^\pm$ . We show that query answering is PTIME-complete and in  $AC_0$  (which is the complexity of evaluating fixed first-order formulas over a database or finite structure), when the query and the set of GTGDs and LTGDs are fixed, respectively. We then enrich  $Datalog^\pm$  by adding negative constraints, which are Horn clauses with a (not necessarily guarded) conjunction of atoms in their body and the truth constant *false*, denoted  $\perp$ , in the head. We show that the introduction of such constraints does not increase the complexity of query answering in  $Datalog^\pm$ . As a further extension, we add *non-conflicting keys*, which are special EGDs that do not interact with TGDs, and thus also do not increase the complexity of query answering in  $Datalog^\pm$ . We deal only with *keys*, since this suffices to capture the most common tractable ontology languages in the literature. The class of *non-conflicting keys* is a generalization of the one in [10].

**Further results.** We have several other results that, for space reasons, we do not include here. We refer the reader to [8] for more details. A central result is that the  $Datalog^\pm$  family is able to express the most common tractable ontology languages, in particular, the *DL-Lite* family [12] and F-Logic Lite [9]. In particular, it is interesting to see that linear  $Datalog^\pm$  enriched with non-conflicting keys is expressive enough to capture the expressive power of the description logics *DL-Lite<sub>F</sub>*, *DL-Lite<sub>R</sub>*, and *DL-Lite<sub>A</sub>*, which are highly suitable for ontological modeling and reasoning. Finally, we are able to deal with an extension of  $Datalog^\pm$  with stratified negation. We provide a canonical model and a perfect model semantics, and we show that they coincide. We thus provide a natural stratified negation for query answering over ontologies, which has been an open problem to date, since it is in general based on several strata of infinite models.

**Related work.** Note that the results of the present paper are related to but very different from the results in [7], where complexity issues of guarded TGDs as well as of another class, called *weakly guarded TGDs*, were first studied. There, it was shown that the combined complexity of query answering and fact inference with guarded TGDs is 2-EXPTIME complete, and it was noted that the data complexity is polynomial, which is now much refined by the present paper. Extensions of  $Datalog$  that allow the introduction of values not appearing in the active domain of the input database have been proposed in the literature [1, 5, 11, 10]; the introduction of such values is usually called *value invention*. Other works integrate  $Datalog$  with description logic knowledge bases [16, 6, 23], which is a different perspective.

## 2 Preliminaries

We briefly recall some basics on databases, queries, TGDs, and the chase.

**Databases and Queries.** We assume (i) an infinite universe of *data constants*  $\Delta$  (which constitute the “normal” domain of a database), (ii) an infinite set of (*labeled*) *nulls*  $\Delta_N$  (used as “fresh” Skolem terms, which are placeholders for unknown values, and can thus be seen as variables), and (iii) an infinite set of variables  $\mathcal{X}$  (used in dependencies and queries). Different constants represent different values (*unique name assumption*), while different nulls may represent the same value. We assume a lexicographic order

on  $\Delta \cup \Delta_N$ , with every symbol in  $\Delta_N$  following all symbols in  $\Delta$ . We denote by  $\mathbf{X}$  sequences of variables  $X_1, \dots, X_k$  with  $k \geq 0$ . We assume a *relational schema*  $\mathcal{R}$ , which is a finite set of *relation names* (or *predicates*). A *term*  $t$  is a constant, null, or variable. An *atomic formula* (or *atom*)  $\mathbf{a}$  has the form  $P(t_1, \dots, t_n)$ , where  $P$  is  $n$ -ary predicate, and  $t_1, \dots, t_n$  are terms. We denote by  $\text{dom}(\mathbf{a})$  the set of all its arguments. These notations naturally extend to sets and conjunctions of atoms. Conjunctions of atoms are often identified with the sets of their atoms. A *database (instance)*  $D$  for  $\mathcal{R}$  is a (possibly infinite) set of atoms with predicates from  $\mathcal{R}$  and arguments from  $\Delta \cup \Delta_N$ . Such  $D$  is *ground* iff it contains only atoms with arguments from  $\Delta$ . A *conjunctive query (CQ)* over  $\mathcal{R}$  has the form  $Q(\mathbf{X}) = \exists \mathbf{Y} \Phi(\mathbf{X}, \mathbf{Y})$ , where  $\Phi(\mathbf{X}, \mathbf{Y})$  is a conjunction of atoms with the variables  $\mathbf{X}$  and  $\mathbf{Y}$ . A *Boolean CQ (BCQ)* over  $\mathcal{R}$  is a CQ of the form  $Q()$ . Answers to CQs and BCQs are defined via *homomorphisms*, which are mappings  $\mu: \Delta \cup \Delta_N \cup \mathcal{X} \rightarrow \Delta \cup \Delta_N \cup \mathcal{X}$  such that (i)  $c \in \Delta$  implies  $\mu(c) = c$ , (ii)  $c \in \Delta_N$  implies  $\mu(c) \in \Delta \cup \Delta_N$ , and (iii)  $\mu$  is naturally extended to atoms, sets of atoms, and conjunctions of atoms. The set of all *answers* to a CQ  $Q(\mathbf{X}) = \exists \mathbf{Y} \Phi(\mathbf{X}, \mathbf{Y})$  over a database  $D$ , denoted  $Q(D)$ , is the set of all tuples  $\mathbf{t}$  over  $\Delta$  for which there exists a homomorphism  $\mu: \mathbf{X} \cup \mathbf{Y} \rightarrow \Delta \cup \Delta_N$  such that  $\mu(\Phi(\mathbf{X}, \mathbf{Y})) \subseteq D$  and  $\mu(\mathbf{X}) = \mathbf{t}$ . The *answer* to a BCQ  $Q()$  over  $D$  is *Yes*, denoted  $D \models Q$ , iff  $Q(D) \neq \emptyset$ .

**TGDs.** Given a relational schema  $\mathcal{R}$ , a *tuple-generating dependency (TGD)*  $\sigma$  is a first-order formula of the form  $\forall \mathbf{X} \forall \mathbf{Y} \Phi(\mathbf{X}, \mathbf{Y}) \rightarrow \exists \mathbf{Z} \Psi(\mathbf{X}, \mathbf{Z})$ , where  $\Phi(\mathbf{X}, \mathbf{Y})$  and  $\Psi(\mathbf{X}, \mathbf{Z})$  are conjunctions of atoms over  $\mathcal{R}$ , called the *body* and the *head* of  $\sigma$ , respectively. We usually omit the universal quantifiers in TGDs. Such  $\sigma$  is satisfied in a database  $D$  for  $\mathcal{R}$  iff, whenever there is a homomorphism  $h$  that maps the atoms of  $\Phi(\mathbf{X}, \mathbf{Y})$  to atoms of  $D$ , there exists an extension  $h'$  of  $h$  that maps the atoms of  $\Psi(\mathbf{X}, \mathbf{Z})$  to atoms of  $D$ . All sets of TGDs are finite here. The notion of *query answering* under TGDs is defined as follows. For a set of TGDs  $\Sigma$  on  $\mathcal{R}$ , and a database  $D$  for  $\mathcal{R}$ , the set of *models* of  $D$  given  $\Sigma$ , denoted  $\text{mods}(\Sigma, D)$ , is the set of all databases  $B$  such that  $B \models D \cup \Sigma$ . The set of *answers* to a CQ  $Q$  on  $D$  given  $\Sigma$ , denoted  $\text{ans}(Q, \Sigma, D)$ , is the set of all tuples  $\mathbf{a}$  such that  $\mathbf{a} \in Q(B)$  for all  $B \in \text{mods}(\Sigma, D)$ . The *answer* to a BCQ  $Q$  over  $D$  given  $\Sigma$  is *Yes*, denoted  $D \cup \Sigma \models Q$ , iff  $\text{ans}(Q, \Sigma, D) \neq \emptyset$ . Note that query answering under general TGDs is undecidable [3], even when the schema and TGDs are fixed [7]. Query answering under a certain class  $\mathcal{C}$  of TGDs is said to be *FO-rewritable* iff, for every given CQ  $Q$  and for every given set  $\Sigma$  of TGDs in  $\mathcal{C}$ , there exists a first order query  $Q_{FO}$  such that, for every instance  $D$ , it holds  $Q_{FO}(D) = \text{ans}(Q, \Sigma, D)$ . We recall that the two problems of CQ and BCQ evaluation under TGDs are LOGSPACE-equivalent [13, 18, 17, 15]. Moreover, it is easy to see that the query output tuple (QOT) problem (as a decision version of CQ evaluation) and BCQ evaluation are AC<sub>0</sub>-reducible to each other. Henceforth, we thus focus only on the BCQ evaluation problem. All complexity results carry over to the other problems. We also recall that query answering under TGDs is equivalent to query answering under TGDs with only singleton atoms in their heads [7]. In the sequel, w.l.o.g., every TGD has a singleton atom in its head.

**The TGD Chase.** The *chase* was introduced for checking the implication of dependencies [20], and later also for checking query containment [18]. It repairs a database relative to a set of dependencies, so that the result of the chase satisfies the dependencies. By “chase”, we refer both to the chase procedure and to its output. The TGD chase works on a database through so-called TGD *chase rules* (for an extended chase with also equality-generating dependencies (EGDs), see [8]). The TGD chase rule comes in

two flavors: *oblivious* and *restricted*, where the restricted one repairs TGDs only when not satisfied. We focus on the oblivious one, since it makes proofs technically simpler. The (*oblivious*) *TGD chase rule* defined below is the building block of the chase.

**TGD CHASE RULE.** Consider a database  $D$  for a relational schema  $\mathcal{R}$ , and a TGD  $\sigma$  on  $\mathcal{R}$  of the form  $\Phi(\mathbf{X}, \mathbf{Y}) \rightarrow \exists \mathbf{Z} \Psi(\mathbf{X}, \mathbf{Z})$ . Then,  $\sigma$  is *applicable* to  $D$  if there exists a homomorphism  $h$  that maps the atoms of  $\Phi(\mathbf{X}, \mathbf{Y})$  to atoms of  $D$ . Let  $\sigma$  be applicable, and  $h_1$  be a homomorphism that extends  $h$  as follows: for each  $X_i \in \mathbf{X}$ ,  $h_1(X_i) = h(X_i)$ ; for each  $Z_j \in \mathbf{Z}$ ,  $h_1(Z_j) = z_j$ , where  $z_j$  is a “fresh” null, i.e.,  $z_j \in \Delta_N$ ,  $z_j$  does not occur in  $D$ , and  $z_j$  lexicographically follows all other nulls already introduced. The application of  $\sigma$  adds to  $D$  the atom  $h_1(\Psi(\mathbf{X}, \mathbf{Z}))$  if not already in  $D$ . ■

The important notion of the (*derivation*) *level* of an atom in a TGD chase is defined as follows. Let  $D$  be the *initial* database from which the chase is constructed. Then: (1) The atoms in  $D$  have level 0. (2) Let a TGD  $\Phi(\mathbf{X}, \mathbf{Y}) \rightarrow \exists \mathbf{Z} \Psi(\mathbf{X}, \mathbf{Z})$  be applied at some point in the construction of the chase, and let  $h$  and  $h_1$  be as in the TGD chase rule. If the atom with highest level among those in  $h_1(\Phi(\mathbf{X}, \mathbf{Y}))$  has level  $k$ , then the added atom  $h_1(\Psi(\mathbf{X}, \mathbf{Z}))$  has level  $k + 1$ . The chase algorithm for  $\Sigma$  and  $D$  consists of an exhaustive application of the TGD chase rule in a breadth-first (level-saturating) fashion, which leads as result to a (possibly infinite) *chase* for  $\Sigma$  and  $D$ , denoted  $\text{chase}(\Sigma, D)$ . The *chase of level up to*  $k \geq 0$  for  $\Sigma$  and  $D$ , denoted  $\text{chase}^k(\Sigma, D)$ , is the set of all atoms in  $\text{chase}(\Sigma, D)$  of level at most  $k$ . The (possibly infinite) chase relative to TGDs is a *universal model*, i.e., for every  $B \in \text{mods}(\Sigma, D)$ , there exists a homomorphism that maps  $\text{chase}(\Sigma, D)$  onto  $B$  [15, 7], and thus  $Q(\text{chase}(\Sigma, D)) = \text{ans}(Q, \Sigma, D)$ .

### 3 Guarded Datalog<sup>±</sup>

We now introduce guarded Datalog<sup>±</sup> as a class of special TGDs that show computational data tractability, while being at the same time expressive enough to model ontologies. BCQs relative to such TGDs can be evaluated on a finite part of the chase of constant size in the data complexity.

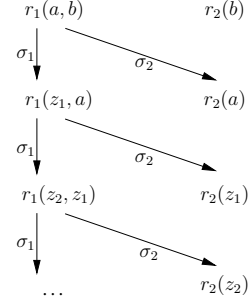
A TGD  $\sigma$  is *guarded* iff it contains an atom in its body that contains all universally quantified variables of  $\sigma$ . The leftmost such atom is the *guard atom* (or *guard*) of  $\sigma$ . The non-guard atoms in the body of  $\sigma$  are the *side atoms* of  $\sigma$ .

*Example 3.1.* The TGD  $r(X, Y), s(Y, X, Z) \rightarrow \exists W s(Z, X, W)$  is guarded (via the guard  $s(Y, X, Z)$ ), while the TGD  $r(X, Y), r(Y, Z) \rightarrow r(X, Z)$  is not guarded.

Sets of guarded TGDs (with single-atom heads) are theories in the guarded fragment of first-order logic [2]. In the sequel, let  $\mathcal{R}$  be a relational schema, let  $D$  be a database for  $\mathcal{R}$ , and let  $\Sigma$  be a set of guarded TGDs on  $\mathcal{R}$ . We first give some preliminary definitions as follows. The *chase graph* for  $\Sigma$  and  $D$  is the directed graph consisting of  $\text{chase}(\Sigma, D)$  as the set of nodes and having an arc from  $\mathbf{a}$  to  $\mathbf{b}$  iff  $\mathbf{b}$  is obtained from  $\mathbf{a}$  and possibly other atoms by a one-step application of a TGD  $\sigma \in \Sigma$ . Here, we mark  $\mathbf{a}$  as *guard* iff  $\mathbf{a}$  is the guard of  $\sigma$ . The *guarded chase forest* for  $\Sigma$  and  $D$  is the restriction of the chase graph for  $\Sigma$  and  $D$  to all atoms marked as guards and their children. The *subtree* of an atom  $\mathbf{a}$  in this forest, denoted  $\text{subtree}(\mathbf{a})$ , is the restriction of the forest to all successors of  $\mathbf{a}$ . The *type* of an atom  $\mathbf{a}$ , denoted  $\text{type}(\mathbf{a})$ , is the set of all atoms  $\mathbf{b}$  in  $\text{chase}(\Sigma, D)$  that have only constants and nulls from  $\mathbf{a}$  as arguments. Informally, the type of  $\mathbf{a}$  is the set of all atoms that determine the subtree of  $\mathbf{a}$  in the forest.

*Example 3.2.* Consider the TGDs  $\sigma_1 : r_1(X, Y), r_2(Y) \rightarrow \exists Z r_1(Z, X)$  and  $\sigma_2 : r_1(X, Y) \rightarrow r_2(X)$ , applied on an instance  $D = \{r_1(a, b), r_2(b)\}$ . The first part of the (infinite) guarded chase forest for  $\{\sigma_1, \sigma_2\}$  and  $D$  is shown in Fig. 3, where every arc is labeled with the applied TGD.

Given a finite set  $S \subseteq \Delta \cup \Delta_N$ , two sets of atoms  $A_1$  and  $A_2$  are  $S$ -isomorphic (or isomorphic if  $S = \emptyset$ ) iff there exists a bijection  $\beta: A_1 \cup \text{dom}(A_1) \rightarrow A_2 \cup \text{dom}(A_2)$  such that (i)  $\beta$  and  $\beta^{-1}$  are homomorphisms, and (ii)  $\beta(c) = c = \beta^{-1}(c)$  for all  $c \in S$ . Two atoms  $a_1$  and  $a_2$  are  $S$ -isomorphic (or isomorphic if  $S = \emptyset$ ) iff  $\{a_1\}$  and  $\{a_2\}$  are  $S$ -isomorphic. The notion of  $S$ -isomorphism (or isomorphism if  $S = \emptyset$ ) is naturally extended to more complex structures, such as pairs of two subtrees  $(V_1, E_1)$  and  $(V_2, E_2)$  of the guarded chase forest, and two pairs  $(\mathbf{b}_1, S_1)$  and  $(\mathbf{b}_2, S_2)$ , where  $\mathbf{b}_1$  and  $\mathbf{b}_2$  are atoms, and  $S_1$  and  $S_2$  are sets of atoms. The following lemma shows that if two atoms have  $S$ -isomorphic types, then their subtrees are also  $S$ -isomorphic.



**Fig. 1.** Guarded chase forest in Example 3.2.

**Lemma 3.1.** *Let  $S$  be a set of constants and nulls, and  $\mathbf{a}_1$  and  $\mathbf{a}_2$  be atoms from  $\text{chase}(\Sigma, D)$  with  $S$ -isomorphic types. Then, the subtree of  $\mathbf{a}_1$  is  $S$ -isomorphic to the subtree of  $\mathbf{a}_2$ .*

The next lemma provides an upper bound for the number of all non- $T$ -isomorphic pairs consisting of an atom and a type.

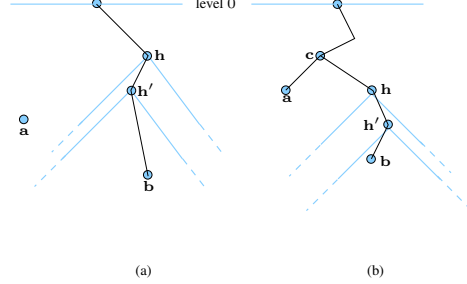
**Lemma 3.2.** *Let  $w$  be the maximal arity of a predicate in  $\mathcal{R}$ ,  $\delta = (2w)^w \cdot 2^{(2w)^{w \cdot |\mathcal{R}|}}$ , and  $\mathbf{a} \in \text{chase}(\Sigma, D)$ . Let  $P$  be a set of pairs  $(\mathbf{b}, S)$ , each consisting of an atom  $\mathbf{b}$  and a type  $S$  of atoms  $\mathbf{c}$  with arguments from  $\mathbf{a}$  and new nulls. If  $|P| > \delta$ , then  $P$  contains at least two  $\text{dom}(\mathbf{a})$ -isomorphic pairs.*

Using the above two lemmas, we are now ready to prove that the atoms in the type of an atom  $\mathbf{a}$  in the guarded chase forest cannot be generated at a depth of the guarded chase forest that exceeds the depth of the atom  $\mathbf{a}$  by a value that depends only on the TGDs. Here, the *guarded depth* of an atom  $\mathbf{a}$  in the guarded chase forest for  $\Sigma$  and  $D$ , denoted  $\text{depth}(\mathbf{a})$ , is the length of the path from  $D$  to  $\mathbf{a}$  in the forest. Note that this is generally different from the derivation level.

**Lemma 3.3.** *Let  $\mathbf{a}$  be a guard in the chase graph for  $\Sigma$  and  $D$ , and  $\mathbf{b} \in \text{type}(\mathbf{a})$ . Then,  $\text{depth}(\mathbf{b}) \leq \text{depth}(\mathbf{a}) + k$ , where  $k$  depends only on  $\Sigma$ .*

**Proof.** Let  $k = (2w)^w \cdot 2^{(2w)^{w \cdot |\mathcal{R}|}}$ , where  $w$  is the maximal arity of a predicate in  $\mathcal{R}$ . Suppose  $\text{depth}(\mathbf{b}) > \text{depth}(\mathbf{a}) + k$  for one atom  $\mathbf{b}$  in the type of  $\mathbf{a}$ . That is, the path  $P$  in the guarded chase forest leading to  $\mathbf{b}$  has a length greater than  $k$  from the depth of  $\mathbf{a}$ . Suppose first  $\mathbf{b}$  contains no nulls. By Lemma 3.2, there are two isomorphic atoms  $\mathbf{h}$  and  $\mathbf{h}'$  (in this order) on  $P$  with isomorphic types (see Fig. 3). By Lemma 3.1, the subtree of  $\mathbf{h}$  is isomorphic to the subtree of  $\mathbf{h}'$ . But then  $\mathbf{b}$  is also in the subtree of  $\mathbf{h}$ , on a path  $Q$  that is at least one edge shorter than  $P$ , which contradicts the assumption

that  $P$  is the path leading to  $b$ . Suppose next the set of nulls  $N$  in  $b$  is nonempty, and consider the common predecessor  $c$  of  $a$  and  $b$  of largest depth. By Lemma 3.2, there are two  $N$ -isomorphic atoms  $h$  and  $h'$  (in this order) on  $P$  with  $N$ -isomorphic types (see Fig. 3). Since  $b$  is in the type of  $a$ , it cannot contain new nulls compared to  $a$ .

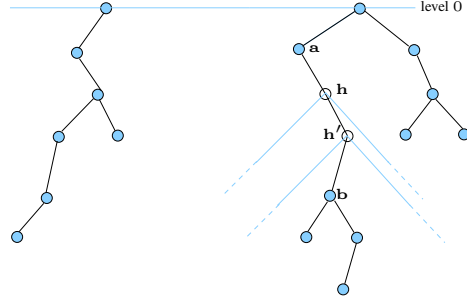


**Fig. 2.** Construction in Lemma 3.3's proof.

type of  $a$  can be obtained on paths of length at most  $k$ .  $\square$

We next show that BCQs can be evaluated using only a finite, initial portion of the guarded chase forest, whose size is determined by the TGDs only. Here, the *guarded chase* of level up to  $k \geq 0$  for  $\Sigma$  and  $D$ , denoted  $g\text{-chase}^k(\Sigma, D)$ , is the set of all atoms in the forest of depth at most  $k$ .

**Lemma 3.4.** *Let  $Q$  be a BCQ over  $\mathcal{R}$ . If there exists a homomorphism  $\mu$  that maps  $Q$  into  $\text{chase}(\Sigma, D)$ , then there exists a homomorphism  $\lambda$  that maps  $Q$  into  $g\text{-chase}^k(\Sigma, D)$ , where  $k$  depends only on  $Q$  and  $\mathcal{R}$ .*



**Fig. 3.** Construction in Lemma 3.4's proof.

$\mu(Q)$  is not contained in  $g\text{-chase}^k(\Sigma, D)$ , this tree must contain a path  $P$  of length greater than  $\delta$  of which all inner nodes (i.e., without start and end node) do not belong to  $\mu(Q)$  and have no branches (i.e., have exactly one outgoing edge). Let  $a$  be the start node of  $P$ . By Lemma 3.2, there are two  $\text{dom}(a)$ -isomorphic atoms  $h$  and  $h'$  on  $P$  with  $\text{dom}(a)$ -isomorphic types. By Lemma 3.1,  $\text{subtree}(h)$  is  $\text{dom}(a)$ -isomorphic to  $\text{subtree}(h')$ . Thus, we can remove  $h$  and the path to  $h'$ , obtaining a path that is at least one edge shorter. Let  $\iota$  be the homomorphism mapping  $\text{subtree}(h')$  to  $\text{subtree}(h)$ , and let  $\mu' = \mu \circ \iota$ . Then,  $\mu'$  is a homomorphism that maps  $Q$  into  $\text{chase}(\Sigma, D)$  such that  $\text{depth}(\mu') < \text{depth}(\mu)$ , which contradicts  $\mu$  being a homomorphism of this kind such that  $\text{depth}(\mu)$  is minimal. This shows that  $\mu(Q)$  is contained in  $g\text{-chase}^k(\Sigma, D)$ .  $\square$

Since  $c$  is the common predecessor of  $a$  and  $b$  of largest depth,  $b$  also cannot contain new nulls compared to  $c$ , and thus compared to  $h$  and  $h'$ . So,  $b$  is in the types of both  $h$  and  $h'$ . By Lemma 3.1, the subtree of  $h$  is  $N$ -isomorphic to the subtree of  $h'$ . But then  $b$  is also in the subtree of  $h$ , on a path  $Q$  that is at least one edge shorter than  $P$ , which contradicts the assumption that  $P$  is the path leading to  $b$ . In summary, all atoms in the

**Proof.** Let  $k = n \cdot \delta$ , where  $n = |Q|$ ,  $\delta = (2w)^w \cdot 2^{(2w)^w \cdot |\mathcal{R}|}$ , and  $w$  is the maximal arity of a predicate in  $\mathcal{R}$ . Suppose there exists a homomorphism that maps  $Q$  into  $\text{chase}(\Sigma, D)$ . Let  $\mu$  be a homomorphism of this kind such that  $\text{depth}(\mu) = \sum_{q \in Q} \text{depth}(\mu(q))$  is minimal. We now show that  $\mu(Q)$  is contained in  $g\text{-chase}^k(\Sigma, D)$ . Towards a contradiction, suppose the contrary. Consider the tree consisting of all atoms in  $\mu(Q)$  and their ancestors in the guarded chase forest for  $\Sigma$  and  $D$ . As

The following definition captures a general property, where also the whole derivation of the query atoms are contained in the finite, initial portion of the guarded chase forest, whose size is determined by the TGDs only.

**Definition 3.1.** We say that  $\Sigma$  has the *bounded guard-depth property (BGDP)* iff, for every database  $D$  for  $\mathcal{R}$  and for every BCQ  $Q$ , whenever there exists a homomorphism  $\mu$  that maps  $Q$  into  $\text{chase}(\Sigma, D)$ , then there exists a homomorphism  $\lambda$  of this kind such that all ancestors of  $\lambda(Q)$  in the chase graph for  $\Sigma$  and  $D$  are contained in  $g\text{-chase}^{\gamma_g}(\Sigma, D)$ , where  $\gamma_g$  depends only on  $Q$  and  $\mathcal{R}$ .

It is not difficult to prove, based on Lemmas 3.3 and 3.4, that guarded TGDs enjoy the BGDP. Hence, BCQ answering under guarded TGDs is in PTIME in data complexity, since by the existence of the homomorphism  $\lambda$  in Lemma 3.4 and in Definition 3.1, we obtain  $Q(g\text{-chase}^{\gamma_g}(\Sigma, D)) = Q(\text{chase}(\Sigma, D)) = \text{ans}(Q, \Sigma, D)$ . The following theorem summarizes these and other results from [8].

**Theorem 3.1.** *Let  $D$  be a database for a relational schema  $\mathcal{R}$ ,  $\Sigma$  a finite set of guarded TGDs on  $\mathcal{R}$ , and  $Q$  a BCQ  $Q$  over  $\mathcal{R}$ . Then, deciding  $D \cup \Sigma \models Q$  is PTIME-complete in data complexity. It can be done in linear time in data complexity for atomic  $Q$ .*

## 4 Linear Datalog<sup>±</sup>

We now introduce linear Datalog<sup>±</sup> as a variant of guarded Datalog<sup>±</sup>, where we prove that query answering is even FO-rewritable (and thus in AC<sub>0</sub>) in the data complexity. Nonetheless, linear Datalog<sup>±</sup> is still expressive enough for representing ontologies [8] in many cases. A TGD is *linear* iff it contains only a singleton body atom. Notice that linear Datalog<sup>±</sup> generalizes the well-known class of *inclusion dependencies*.

We first define the bounded derivation-depth property, which is strictly stronger than the bounded guard-depth property.

**Definition 4.1.** A set of TGDs  $\Sigma$  has the *bounded derivation-depth property (BDDP)* iff, for every database  $D$  for  $\mathcal{R}$  and for every BCQ  $Q$  over  $\mathcal{R}$ , whenever  $D \cup \Sigma \models Q$ , then  $\text{chase}^{\gamma_d}(\Sigma, D) \models Q$ , where  $\gamma_d$  depends only on  $Q$  and  $\mathcal{R}$ .

Clearly, in the case of linear TGDs, for every  $\mathbf{a} \in \text{chase}(\Sigma, D)$ , the subtree of  $\mathbf{a}$  is determined only by  $\mathbf{a}$ , instead of  $\text{type}(\mathbf{a})$  (cf. Lemma 3.1). Therefore, for a single atom, its depth coincides with the number of applications of the TGD chase rule that are necessary to generate it. By this observation, as an immediate consequence of the fact that guarded TGDs enjoy the BGDP, we obtain that linear TGDs have the bounded derivation-depth property. The next result shows that queries relative to TGDs with the bounded derivation-depth property are FO-rewritable.

**Theorem 4.1.** *Let  $\Sigma$  be a set of TGDs over a relational schema  $\mathcal{R}$ ,  $D$  be a database for  $\mathcal{R}$ , and  $Q$  be a BCQ over  $\mathcal{R}$ . If  $\Sigma$  enjoys the BDDP, then  $Q$  is FO-rewritable.*

As an immediate consequence of the fact that linear TGDs enjoy the BDDP and of Theorem 4.1, we have that BCQs are FO-rewritable in the linear case.

**Corollary 4.1.** *Let  $\mathcal{R}$  be a relational schema,  $\Sigma$  be a set of linear TGDs over  $\mathcal{R}$ ,  $D$  be a database for  $\mathcal{R}$ , and  $Q$  be a BCQ over  $\mathcal{R}$ . Then,  $Q$  is FO-rewritable.*

## 5 Extensions

In this section, we extend Datalog<sup>±</sup> with negative constraints and EGDs, which are both important when representing ontologies.

**Negative Constraints.** A *negative constraint* (or *constraint*) is a first-order formula of the form  $\forall \mathbf{X} \Phi(\mathbf{X}) \rightarrow \perp$ , where  $\Phi(\mathbf{X})$  is a (not necessarily guarded) conjunction of atoms. It is often also written as  $\forall \mathbf{X} \Phi'(\mathbf{X}) \rightarrow \neg p(\mathbf{X})$ , where  $\Phi'(\mathbf{X})$  is obtained from  $\Phi(\mathbf{X})$  by removing the atom  $p(\mathbf{X})$ . We usually omit the universal quantifiers.

*Example 5.1.* If the unary predicates  $c$  and  $c'$  represent two classes, we may use the constraint  $c(X), c'(X) \rightarrow \perp$  to assert that the two classes have no common instances. Similarly, if additionally the binary predicate  $r$  represents a relationship, we may use  $c(X), r(X, Y) \rightarrow \perp$  to enforce that no member of the class  $c$  participates to  $r$ .

Query answering on a database  $D$  under TGDs  $\Sigma_T$  and constraints  $\Sigma_C$  can be done effortlessly by additionally checking that every constraint  $\sigma = \Phi(\mathbf{X}) \rightarrow \perp \in \Sigma_C$  is satisfied in  $D$  given  $\Sigma_T$ , each of which can be done by checking that the BCQ  $Q_\sigma = \Phi(\mathbf{X})$  evaluates to false in  $D$  given  $\Sigma_T$ . We write  $D \cup \Sigma_T \models \Sigma_C$  iff every  $\sigma \in \Sigma_C$  is false in  $D$  given  $\Sigma_T$ . We thus obtain immediately the following result. Here, a BCQ  $Q$  is true in  $D$  given  $\Sigma_T$  and  $\Sigma_C$ , denoted  $D \cup \Sigma_T \cup \Sigma_C \models Q$ , iff (i)  $D \cup \Sigma_T \models Q$  or (ii)  $D \cup \Sigma_T \not\models \Sigma_C$  (as usual in DLs).

**Theorem 5.1.** *Let  $\mathcal{R}$  be a relational schema,  $\Sigma_T$  and  $\Sigma_C$  be sets of TGDs and constraints on  $\mathcal{R}$ , respectively,  $D$  be a database for  $\mathcal{R}$ , and  $Q$  be a BCQ on  $\mathcal{R}$ . Then,  $D \cup \Sigma_T \cup \Sigma_C \models Q$  iff (i)  $D \cup \Sigma_T \models Q$  or (ii)  $D \cup \Sigma_T \models Q_\sigma$  for some  $\sigma \in \Sigma_C$ .*

The next theorem shows that constraints do not increase the data and combined complexity of answering BCQs in the guarded and linear case. It follows from Theorem 5.1, by which the additional effort of deciding  $D \cup \Sigma_T \models \Sigma_C$  can be done by answering  $|\Sigma_C|$  BCQs  $Q_\sigma$  without constraints, where each query  $Q_\sigma$  has a size of  $\|\sigma\| \leq \|\Sigma_C\|$ , and in the data complexity,  $|\Sigma_C|$  and  $\|\sigma\|$  are constants. Here,  $|\Sigma_C|$  denotes the cardinality of  $\Sigma_C$ , while  $\|\Sigma_C\|$  denotes the input size of  $\Sigma_C$ . This additional effort does not increase the complexity of answering BCQs without constraints.

**Theorem 5.2.** *Let  $\mathcal{R}$  be a relational schema,  $\Sigma_T$  and  $\Sigma_C$  be sets of TGDs and constraints on  $\mathcal{R}$ , respectively,  $D$  be a database for  $\mathcal{R}$ , and  $Q$  be a BCQ. Then, deciding  $D \cup \Sigma_T \cup \Sigma_C \models Q$  in the guarded (resp., linear) case has the same data and the same combined complexity as deciding  $D \cup \Sigma_T \models Q$  in the guarded (resp., linear) case.*

**EGDs and Non-Conflicting Keys.** An *equality-generating dependency* (or *EGD*)  $\sigma$  is a first-order formula of the form  $\forall \mathbf{X} \Phi(\mathbf{X}) \rightarrow X_i = X_j$ , where  $\Phi(\mathbf{X})$ , called the *body* of  $\sigma$ , is a conjunction of atoms, and  $X_i$  and  $X_j$  are variables from  $\mathbf{X}$ . We call  $X_i = X_j$  the *head* of  $\sigma$ . Such  $\sigma$  is satisfied in a database  $D$  for  $\mathcal{R}$  iff, whenever there exists a homomorphism  $h$  such that  $h(\Phi(\mathbf{X}, \mathbf{Y})) \subseteq D$ , it holds that  $h(X_i) = h(X_j)$ .

*Example 5.2.* The EGD  $r(X, Y), r(X, Y') \rightarrow Y = Y'$  expresses that the second argument of the binary predicate  $r$  functionally depends on the first argument of  $r$ .



The chase is naturally extended to databases under EGDs in addition to TGDs: The *chase* of a database  $D$  in the presence of two sets  $\Sigma_T$  and  $\Sigma_E$  of TGDs and EGDs, respectively, denoted  $\text{chase}(\Sigma_T \cup \Sigma_E, D)$ , is computed by iteratively applying (1) a single TGD once, according to the order specified above, and (2) the EGDs, as long as applicable (i.e., until a fixpoint is reached), where EGDs are applied as follows:

**EGD CHASE RULE.** Consider a database  $D$  for a relational schema  $\mathcal{R}$ , and an EGD  $\sigma$  on  $\mathcal{R}$  of the form  $\Phi(\mathbf{X}) \rightarrow X_i = X_j$ . Such an EGD  $\sigma$  is *applicable* to  $D$  iff there exists a homomorphism  $\eta: \Phi(\mathbf{X}) \rightarrow D$  such that  $\eta(X_i)$  and  $\eta(X_j)$  are different and not both constants. If  $\eta(X_i)$  and  $\eta(X_j)$  are different constants, then there is a *hard violation* of  $\sigma$ , and the chase *fails*. Otherwise, the result of the application of  $\sigma$  to  $D$  is the database  $h(D)$  obtained from  $D$  by replacing every occurrence of a non-constant element  $e \in \{\eta(X_i), \eta(X_j)\}$  in  $D$  by the other element  $e'$  (if  $e$  and  $e'$  are both nulls, then  $e$  precedes  $e'$  in the lexicographic order). ■

While adding negative constraints is computationally effortless, adding EGDs is more problematic. The interaction of TGDs and EGDs leads to undecidability of query answering even in simple cases, such that of functional and inclusion dependencies [14], or keys and inclusion dependencies (see, e.g., [10], where the proof of undecidability is done in the style of Vardi as in [18]). It can even be seen that a *fixed* set of EGDs and guarded TGDs can simulate a universal Turing machine, and thus query answering and even propositional ground atom inference is undecidable with such fixed sets of dependencies. For this reason, we consider a restricted class of EGDs, namely, *non-conflicting key dependencies* (or *NC keys*), which show a controlled interaction with TGDs (and negative constraints), such that they do not increase the complexity of answering BCQs. Nonetheless, this class is sufficient for ontology modeling.

We now first concentrate on the semantic notion of separability for EGDs, which formulates a controlled interaction between EGDs and TGDs (and negative constraints), such that the EGDs do not increase the complexity of answering BCQs. We then provide a sufficient syntactic condition for the separability of EGDs, where we transfer a result by [10] about *non-key-conflicting* inclusion dependencies to the more general setting of Datalog<sup>±</sup>. In the context of description logics, general EGDs cannot be formulated, but only keys. Therefore, we mainly focus on keys here.

**Definition 5.1.** Let  $\mathcal{R}$  be a relational schema, and  $\Sigma_T$  and  $\Sigma_E$  be sets of TGDs and EGDs on  $\mathcal{R}$ , respectively. Then,  $\Sigma_E$  is *separable* from  $\Sigma_T$  iff for every database  $D$  for  $\mathcal{R}$ , the following conditions (i) and (ii) are both satisfied:

- (i) If there is a hard violation of an EGD in  $\text{chase}(\Sigma_T \cup \Sigma_E, D)$ , then there is also a hard violation of some EGD of  $\Sigma_E$ , when this EGD is directly applied to  $D$ .
- (ii) If there is no chase failure, then for every BCQ  $Q$ ,  $\text{chase}(\Sigma_T \cup \Sigma_E, D) \models Q$  iff  $\text{chase}(\Sigma_T, D) \models Q$ .

The following result shows that adding separable EGDs to TGDs and constraints does not increase the data and combined complexity of answering BCQs in the guarded and linear case. It follows immediately from the fact that the separability implies that chase failure can be directly evaluated on  $D$ . For fixed (resp., variable)  $\Sigma_E$ , this can be done by evaluating a first-order formula on  $D$  (resp., in polynomial time in the size of  $D$  and  $\Sigma_E$ ), which clearly does not increase the data (resp., combined) complexity of answering BCQs in the three cases.

**Theorem 5.3.** *Let  $\mathcal{R}$  be a relational schema,  $\Sigma_T$  and  $\Sigma_C$  be sets of TGDs and constraints on  $\mathcal{R}$ , respectively, and  $D$  be a database for  $\mathcal{R}$ . Let  $\Sigma_E$  be a set of EGDs that is separable from  $\Sigma_T$ , and  $Q$  be a BCQ. Then, deciding  $D \cup \Sigma_T \cup \Sigma_C \cup \Sigma_E \models Q$  in the guarded (resp., linear) case has the same data complexity and also the same combined complexity as deciding  $D \cup \Sigma_T \models Q$  in the guarded (resp., linear) case.*

We next provide a sufficient syntactic condition for the separability of EGDs. We focus on *key dependencies (KDs, or simply keys)*: a key  $\kappa$  on a predicate  $r$  specifies a set  $\mathbf{K}$  of key attribute positions of  $r$ ; it is satisfied in a database  $D$  if no two atoms in  $D$  of the form  $r(\mathbf{t})$  have the same values in all positions in  $\mathbf{K}$ . Clearly, KDs are special types of EGDs. The EGD chase rule in case of a KD applies to pairs of tuples that violate the KD. For example, a KD  $\kappa$  asserting that the key attribute of a ternary predicate  $r$  is the first one is applicable to the instance  $D = \{r(a, b, z_3), r(a, z_2, z_1)\}$  (where the  $z_i$ 's are nulls), and its application yields the new instance  $\{r(a, b, z_1)\}$ . The following definition generalizes the notion of “non-key-conflicting” dependency relative to a set of keys, introduced in [10], to the context of arbitrary TGDs.

**Definition 5.2.** Let  $\kappa$  be a key, and  $\sigma$  be a TGD of the form  $\Phi(\mathbf{X}, \mathbf{Y}) \rightarrow \exists \mathbf{Z} r(\mathbf{X}, \mathbf{Z})$ . Then,  $\kappa$  is *non-conflicting (NC)* with  $\sigma$  iff either (i) the key does not apply to the predicate  $r$  in the head of  $\sigma$  or (ii) the positions of  $\kappa$  in  $r$  are not a proper subset of the  $\mathbf{X}$ -positions in  $r$  in the head of  $\sigma$ , and every existentially quantified variable in  $\sigma$  appears only once in the head of  $\sigma$ . We say  $\kappa$  is *non-conflicting (NC)* with a set of TGDs  $\Sigma_T$  iff  $\kappa$  is NC with every  $\sigma \in \Sigma_T$ . We say a set of keys  $\Sigma_K$  is *non-conflicting (NC)* with  $\Sigma_T$  iff every  $\kappa \in \Sigma_K$  is NC with  $\Sigma_T$ .

*Example 5.3.* Consider the four keys  $\kappa_1, \kappa_2, \kappa_3$ , and  $\kappa_4$  defined by the key attribute sets  $\mathbf{K}_1 = \{r[1], r[2]\}$ ,  $\mathbf{K}_2 = \{r[1], r[3]\}$ ,  $\mathbf{K}_3 = \{r[3]\}$ , and  $\mathbf{K}_4 = \{r[1]\}$ , respectively, and the TGD  $\sigma = p(X, Y) \rightarrow \exists Z r(X, Y, Z)$ . Then, the head predicate of  $\sigma$  is  $r$ , and the set of positions in  $r$  with universally quantified variables is  $\mathbf{H} = \{r[1], r[2]\}$ . Observe that all keys but  $\kappa_4$  are NC with  $\sigma$ , since only  $\mathbf{K}_4 \subset \mathbf{H}$ . Roughly, every atom added in a chase by applying  $\sigma$  would have a fresh null in some position in  $\mathbf{K}_1, \mathbf{K}_2$ , and  $\mathbf{K}_3$ , thus never firing  $\kappa_1, \kappa_2$ , and  $\kappa_3$ , respectively.

The following theorem shows that the NC property between keys and TGDs implies their separability. This generalizes a useful result of [10] on inclusion dependencies to the much larger class of all TGDs. The main idea behind the proof is roughly as follows. The NC condition between a key  $\kappa$  and a TGD  $\sigma$  assures that either (a) the application of  $\sigma$  in the chase generates an atom with a fresh null in a position of  $\kappa$ , and so the fact does not violate  $\kappa$  (see also Example 5.3), or (b) the  $\mathbf{X}$ -positions in the predicate  $r$  in the head of  $\sigma$  coincide with the key positions of  $\kappa$  in  $r$ , and thus any newly generated atom must have fresh distinct nulls in all but the key position, and may eventually be eliminated without violation. It then follows that the full chase does not fail. Since the new nulls are all distinct, it also contains a homomorphic image of the TGD chase. Therefore, the full chase is in fact homomorphically equivalent to the TGD chase.

**Theorem 5.4.** *Let  $\mathcal{R}$  be a relational schema,  $\Sigma_T$  and  $\Sigma_K$  be sets of TGDs and keys, respectively, such that  $\Sigma_K$  is NC with  $\Sigma_T$ . Then,  $\Sigma_K$  is separable from  $\Sigma_T$ .*

We conclude this section by stating that in the NC case, keys do not increase the data and the combined complexity of answering BCQs under guarded (resp., linear) TGDs and constraints. This result is immediate by Theorems 5.4 and 5.3.

**Corollary 5.1.** *Let  $\mathcal{R}$  be a relational schema,  $\Sigma_T$  and  $\Sigma_C$  be sets of TGDs and constraints on  $\mathcal{R}$ , respectively, and  $D$  be a database for  $\mathcal{R}$ . Let  $\Sigma_E$  be a set of EGDs that is NC with  $\Sigma_T$ , and  $Q$  be a BCQ. Then, deciding  $D \cup \Sigma_T \cup \Sigma_C \cup \Sigma_E \models Q$  in the guarded (resp., linear) case has the same data complexity and also the same combined complexity as deciding  $D \cup \Sigma_T \models Q$  in the guarded (resp., linear) case.*

## 6 Ontology Querying

In [8], we have provided a translation of the DLs  $DL-Lite_{\mathcal{F}}$ ,  $DL-Lite_{\mathcal{R}}$ , and  $DL-Lite_{\mathcal{A}}$  of the  $DL-Lite$  family [12] into linear Datalog<sup>±</sup> with negative constraints and non-conflicting keys, called Datalog<sub>0</sub><sup>±</sup>, and shown that the former are strictly less expressive than the latter. The other DLs of the  $DL-Lite$  family (including  $DL-Lite_{\mathcal{F},\square}$  and  $DL-Lite_{\mathcal{R},\square}$ ) can be similarly translated to Datalog<sub>0</sub><sup>±</sup>. We now illustrate the translation.

*Example 6.1.* Consider the following sets of atomic concepts, abstract roles, and individuals, which represent (i) the classes of scientists, articles, conference papers, and journal papers, (ii) the binary relations “has author”, “has first author”, and “is author of”, and (iii) two individuals, respectively:

$$\begin{aligned} \mathbf{A} &= \{Scientist, Article, ConPaper, JouPaper\}, \\ \mathbf{R}_A &= \{hasAuthor, hasFirstAuthor, isAuthorOf\}, \mathbf{I} = \{i_1, i_2\}. \end{aligned}$$

The following are concept inclusion axioms, which informally express that (i) conference and journal papers are articles, (ii) conference papers are not journal papers, (iii) every scientist has a publication, (iv)  $isAuthorOf$  relates scientists and articles:

$$\begin{aligned} ConPaper \sqsubseteq Article, JouPaper \sqsubseteq Article, ConPaper \sqsubseteq \neg JouPaper, \\ Scientist \sqsubseteq \exists isAuthorOf, \exists isAuthorOf \sqsubseteq Scientist, \exists isAuthorOf^- \sqsubseteq Article. \end{aligned}$$

The following role inclusion and functionality axioms express that (v)  $isAuthorOf$  is the inverse of  $hasAuthor$ , and (vi)  $hasFirstAuthor$  is a functional binary relationship:

$$isAuthorOf^- \sqsubseteq hasAuthor, hasAuthor^- \sqsubseteq isAuthorOf, (\text{funct } hasFirstAuthor).$$

Finally, the concept and role membership axioms  $Scientist(i_1)$ ,  $isAuthorOf(i_1, i_2)$ , and  $Article(i_2)$  express that the individual  $i_1$  is a scientist who authors the article  $i_2$ .

Then, the concept inclusion axioms are translated to the following TGDs and constraints (where we identify atomic concepts and roles with their predicates):

$$\begin{aligned} ConPaper(X) \rightarrow Article(X), JouPaper(X) \rightarrow Article(X), \\ ConPaper(X) \rightarrow \neg JouPaper(X), Scientist(X) \rightarrow \exists Z isAuthorOf(X, Z), \\ isAuthorOf(X, Y) \rightarrow Scientist(X), isAuthorOf(Y, X) \rightarrow Article(X). \end{aligned}$$

The role inclusion and functionality axioms are translated to these TGDs and EGDs:

$$\begin{aligned} isAuthorOf(Y, X) \rightarrow hasAuthor(X, Y), hasAuthor(Y, X) \rightarrow isAuthorOf(X, Y), \\ hasFirstAuthor(X, Y), hasFirstAuthor(X, Y') \rightarrow Y = Y'. \end{aligned}$$

Finally, the three concept and role membership axioms are translated to the three database atoms  $Scientist(i_1)$ ,  $isAuthorOf(i_1, i_2)$ , and  $Article(i_2)$ , respectively (where we also identify individuals with their constants).

**Acknowledgments.** The work of A. Calì and G. Gottlob was supported by the EPSRC grant Number EP/E010865/1 “Schema Mappings and Automated Services for Data Integration”. G. Gottlob, whose work was partially carried out at the Oxford-Man Institute of Quantitative Finance, gratefully acknowledges support from the Royal Society as the holder of a Royal Society-Wolfson Research Merit Award. T. Lukasiewicz’s work was supported by the German Research Foundation under the Heisenberg Programme.

## References

1. S. Abiteboul and V. Vianu. Datalog extensions for database queries and updates. *J. Comput. Syst. Sci.*, 43(1):62–124, 1991.
2. H. Andréka, I. Németi, and J. van Benthem. Modal languages and bounded fragments of predicate logic. *J. Philos. Logic*, 27(3):217–274, 1998.
3. C. Beeri and M. Y. Vardi. The implication problem for data dependencies. In *Proc. ICALP-1981, LNCS 115*, pp. 73–85. Springer, 1981.
4. T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Sci. Am.*, 284:34–43, 2001.
5. L. Cabibbo. The expressive power of stratified logic programs with value invention. *Inf. Comput.*, 147(1):22–56, 1998.
6. M. Cadoli, L. Palopoli, and M. Lenzerini. Datalog and description logics: Expressive power. In *Proc. DBPL-1997, LNCS 1369*, pp. 281–298. Springer, 1998.
7. A. Calì, G. Gottlob, and M. Kifer. Taming the infinite chase: Query answering under expressive relational constraints. In *Proc. KR-2008*, pp. 70–80. AAAI Press, 2008. Revised version: <http://benner.dbai.tuwien.ac.at/staff/gottlob/CGK.pdf>.
8. A. Calì, G. Gottlob, and T. Lukasiewicz. A general Datalog-based framework for tractable query answering over ontologies. In *Proc. PODS-2009*. ACM Press, 2009.
9. A. Calì and M. Kifer. Containment of conjunctive object meta-queries. In *Proc. VLDB-2006*.
10. A. Calì, D. Lembo, R. Rosati. On the decidability and complexity of query answering over inconsistent and incomplete databases. In *Proc. PODS-2003*, pp. 260–271. ACM Press, 2003.
11. F. Calimeri, S. Cozza, and G. Ianni. External sources of knowledge and value invention in logic programming. *Ann. Math. Artif. Intell.*, 50(3/4):333–361, 2007.
12. 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.
13. A. Chandra and P. Merlin. Optimal implementation of conjunctive queries in relational data bases. In *Proc. STOC-1977*, pp. 77–90. ACM Press, 1977.
14. A. K. Chandra and M. Y. Vardi. The implication problem for functional and inclusion dependencies is undecidable. *SIAM J. Comput.*, 14(3):671–677, 1985.
15. A. Deutsch, A. Nash, J. B. Remmel. The chase revisited. In *Proc. PODS-2008*, pp. 149–158.
16. F. M. Donini, M. Lenzerini, D. Nardi, and A. Schaerf.  $\mathcal{AL}$ -log: Integrating Datalog and description logics. *J. Intell. Inf. Syst.*, 10(3):227–252, 1998.
17. 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.
18. D. Johnson and A. Klug. Testing containment of conjunctive queries under functional and inclusion dependencies. *J. Comput. Syst. Sci.*, 28(1):167–189, 1984.
19. M. Lenzerini. Data integration: A theoretical perspective. In *Proc. PODS-2002*, pp. 233–246.
20. D. Maier, A. O. Mendelzon, and Y. Sagiv. Testing implications of data dependencies. *ACM Trans. Database Syst.*, 4(4):455–469, 1979.
21. P. F. Patel-Schneider and I. Horrocks. Position paper: A comparison of two modelling paradigms in the Semantic Web. In *Proc. WWW-2006*, pp. 3–12. ACM Press, 2006.
22. 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.
23. R. Rosati. On combining description logic ontologies and nonrecursive Datalog rules. In *Proc. RR-2008, LNCS 5341*, pp. 13–27. Springer, 2008.