

# Dyadic Existential Rules

GEORG GOTTLÖB

*Department of Computer Science, University of Oxford, Oxford OX1 3QG, UK*  
*Faculty of Informatics, TU Wien, Vienna, Austria*  
(e-mail: [georg.gottlob@cs.ox.ac.uk](mailto:georg.gottlob@cs.ox.ac.uk))

MARCO MANNA and CINZIA MARTE

*Department of Mathematics and Computer Science, University of Calabria, Arcavacata, Italy*  
(e-mails: [marco.manna@unical.it](mailto:marco.manna@unical.it), [cinzia.marte@unical.it](mailto:cinzia.marte@unical.it))

*submitted 29 January 2023; revised 4 July 2023; accepted 25 July 2023*

---

## Abstract

Existential rules form an expressive Datalog-based language to specify ontological knowledge. The presence of existential quantification in rule-heads, however, makes the main reasoning tasks undecidable. To overcome this limitation, in the last two decades, a number of classes of existential rules guaranteeing the decidability of query answering have been proposed. Unfortunately, only some of these classes fully encompass Datalog and, often, this comes at the price of higher computational complexity. Moreover, expressive classes are typically unable to exploit tools developed for classes exhibiting lower expressiveness. To mitigate these shortcomings, this paper introduces a novel general syntactic condition that allows us to define, systematically and in a uniform way, from any decidable class  $\mathcal{C}$  of existential rules, a new class called Dyadic- $\mathcal{C}$  enjoying the following properties: (i) it is decidable; (ii) it generalizes Datalog; (iii) it generalizes  $\mathcal{C}$ ; (iv) it can effectively exploit any reasoner for query answering over  $\mathcal{C}$ ; and (v) its computational complexity does not exceed the highest between the one of  $\mathcal{C}$  and the one of Datalog.

**KEYWORDS:** existential rules, Datalog, ontology-based query answering, tuple-generating dependencies, computational complexity

---

## 1 Introduction

In ontology-based query answering, a conjunctive query is typically evaluated over a logical theory consisting of a relational database paired with an ontology. Description logics (Baader *et al.* 2003) and existential rules – a.k.a. tuple generating dependencies, or Datalog<sup>±</sup> rules – (Baget *et al.* 2011) are the main languages used to specify ontologies. In particular, the latter are essentially classical datalog rules (Abiteboul *et al.* 1995) extended with existential quantified variables in rule-heads. The presence of existential quantification in the head of rules, however, makes query answering undecidable in the general case. To overcome this limitation, in the last two decades, a number of classes of existential rules – based on both semantic and syntactic conditions – that guarantee the decidability of query answering have been proposed. Concerning the semantic conditions, we recall *finite expansions sets*, *finite treewidth sets*, *finite unification sets*, and *strongly parsimonious sets* (Baget *et al.* 2009; Baget *et al.* 2011; Leone *et al.* 2019). Each of

Table 1. Computational complexity of query answering

Class $\mathcal{C}$	Data complexity	Combined complexity
Weakly-(Fr)-Guarded <sup>[1],[2]</sup>	<b>EXPTIME</b> -complete	<b>2EXPTIME</b> -complete
(Fr)-Guarded <sup>[2]</sup>	<b>PTIME</b> -complete	<b>2EXPTIME</b> -complete
Weakly-Acyclic <sup>[3]</sup>	<b>PTIME</b> -complete	<b>2EXPTIME</b> -complete
Jointly-Acyclic <sup>[4]</sup>	<b>PTIME</b> -complete	<b>2EXPTIME</b> -complete
Datalog <sup>[5]</sup>	<b>PTIME</b> -complete	<b>EXPTIME</b> -complete
Shy <sup>[6]</sup>	<b>PTIME</b> -complete	<b>EXPTIME</b> -complete
Ward <sup>[7]</sup>	<b>PTIME</b> -complete	<b>EXPTIME</b> -complete
Protected <sup>[8]</sup>	<b>PTIME</b> -complete	<b>EXPTIME</b> -complete
Sticky-(Join) <sup>[9]</sup>	<b>AC</b> <sub>0</sub>	<b>EXPTIME</b> -complete
Linear <sup>[10]</sup>	<b>AC</b> <sub>0</sub>	<b>PSPACE</b> -complete
Joinless <sup>[11]</sup>	<b>AC</b> <sub>0</sub>	<b>PSPACE</b> -complete
Inclusion-Dependencies <sup>[12]</sup>	<b>AC</b> <sub>0</sub>	<b>PSPACE</b> -complete

[1] Baget et al. (2011); [2] Cali et al. (2013); [3] Fagin et al. (2005); [4] Krötzsch and Rudolph (2011); [5] Ceri et al. (1989); [6] Leone et al. (2019); [7] Gottlob and Pieris (2015); [8] Baldazzi et al. (2022); [9] Cali et al. (2012b); [10] Cali et al. (2012a); [11] Gogacz and Marcinkowski (2017); [12] Johnson and Klug (1984).

these classes encompasses a number of concrete classes based on syntactic conditions (Baget et al. 2011; Cali et al. 2013; Fagin et al. 2005; Krötzsch and Rudolph 2011; Ceri et al. 1989; Leone et al. 2019; Gottlob and Pieris 2015; Baldazzi et al. 2022; Cali et al. 2012b; Cali et al. 2012a; Gogacz and Marcinkowski 2017; Johnson and Klug 1984). Table 1 summarizes these classes and their computational complexity with respect to query answering, by distinguishing between *combined complexity* (the input consists of a database, an ontology, a conjunctive query, and a tuple of constants) and *data complexity* (only a database is given as input, whereas the remaining parameters are considered fixed).

Unfortunately, on the one side, despite the fact that existential rules generalize datalog rules, only some of these syntactic classes fully encompass Datalog and, in some cases, this even comes at the price of higher computational complexity of query answering. Moreover, on the other side, expressive classes typically need *ad hoc* reasoners without being able to exploit mature tools developed for classes exhibiting lower expressiveness.

With the aim of mitigating the two aforementioned shortcomings, this paper introduces a novel general syntactic condition that allows to define, systematically and in a uniform way, from any decidable class  $\mathcal{C}$  of existential rules, a new class called **Dyadic- $\mathcal{C}$**  that enjoys the following properties: (i) it is decidable; (ii) it generalizes Datalog;<sup>1</sup> (iii) it generalizes  $\mathcal{C}$ ; and (iv) it can effectively exploit any reasoner for query answering over  $\mathcal{C}$ . In particular, let  $\mathbb{C}_d$  (resp.,  $\mathbb{C}_c$ ) be the data (resp., combined) complexity of query answering over  $\mathcal{C}$ ,

<sup>1</sup> Strictly speaking, to guarantee that **Dyadic- $\mathcal{C}$**  generalizes Datalog, one has to focus on any  $\mathcal{C} \supseteq \text{Af-Inds}$ , where **Af-Inds** is the very simple class of existential rules defined in Section 3.3 such that (a) rules are inclusion dependencies with no existential variable and (b) predicates in rule-heads do not appear in rule-bodies. Indeed, all known classes of existential rules based on semantic conditions as well as all concrete classes reported in Table 1 do encompass **Af-Inds**.

query answering over Dyadic- $\mathcal{C}$  is  $\mathbf{PTIME}^{\mathbb{C}_d}$  (resp.,  $\mathbf{EXPTIME}^{\mathbb{C}_c}$  provided that there is at least an exponential jump from  $\mathbb{C}_d$  to  $\mathbb{C}_c$ ). Since all the classes reported in Table 1 comply with the exponential jump assumption, we get the following: (a) whenever  $\mathbb{C}_d \supseteq \mathbf{PTIME}$  (entries 1–8 of Table 1), then query answering over Dyadic- $\mathcal{C}$  is complete for  $\mathbb{C}_d$  (resp.,  $\mathbb{C}_c$ ); (b) in all the remaining cases (entries 9–12 of Table 1), query answering over Dyadic- $\mathcal{C}$  is complete for  $\mathbf{PTIME}$  (resp.,  $\mathbf{EXPTIME}$ ), namely it has the same complexity of query answering over Datalog.

Concerning the key principle at the heart of this new general syntactic condition, basically, an ontology  $\Sigma$  belongs to Dyadic- $\mathcal{C}$  if one can easily construct a pair  $(\Sigma_{\text{HG}}, \Sigma_{\mathcal{C}})$  of ontologies, called *dyadic*, such that: (i)  $\Sigma_{\text{HG}} \cup \Sigma_{\mathcal{C}}$  is equivalent to  $\Sigma$  with respect to query answering; (ii)  $\Sigma_{\mathcal{C}} \in \mathcal{C}$ ; and (iii)  $\Sigma_{\text{HG}}$  is a set rules called *head-ground* with respect to  $\Sigma_{\text{HG}} \cup \Sigma_{\mathcal{C}}$  (Gottlob and Pieris 2015). Intuitively,  $\Sigma_{\text{HG}}$  satisfies the following properties: (1) it belongs to Datalog; (2) for each database  $D$ , the chase procedure (Deutsch et al. 2008) over  $D \cup \Sigma_{\text{HG}} \cup \Sigma_{\mathcal{C}}$  never generates atoms containing null-values via rules of  $\Sigma_{\text{HG}}$ ; (3) head-predicates of  $\Sigma_{\text{HG}}$  and body-predicates of  $\Sigma_{\text{HG}}$  are disjoint; and (4) head-predicates of  $\Sigma_{\text{HG}}$  and head-predicates of  $\Sigma_{\mathcal{C}}$  are disjoint. Finally, since Dyadic- $\mathcal{C}$  is well-defined even if  $\mathcal{C}$  is a class of existential rules based on some semantic conditions and, if so, since query answering is still decidable over Dyadic- $\mathcal{C}$ , then – in analogy with the existing semantic classes – the union of all the Dyadic- $\mathcal{C}$  classes are called *dyadic decomposable sets*.

The article is a revised version of an earlier workshop paper (Gottlob et al. 2022). Specifically, the content that was previously presented in a single preliminary section has been expanded and reorganized into two longer separate sections, namely Sections 2 and 3. These sections now contain the necessary background information, ensuring that the paper is self-contained. Furthermore, the previous notion of “dyadic decomposition” has evolved into the novel notion of a “Dyadic Pair of TGDs”, which is discussed in Section 4. This new notion captures the essential properties of dyadic decompositions and also generalizes the notion of ontology, providing new perspectives and insights. Additionally, in Section 5, the notion of “Dyadic Decomposable Sets” is now supported by a canonical concrete algorithm that produces a Dyadic Pair of TGDs from each Dyadic Decomposable Set. The revisions also lead to new results regarding decidability and complexity. First, if  $\mathcal{C}$  is an abstract (resp., concrete) and decidable class, then Dyadic- $\mathcal{C}$  is now also an abstract (resp., concrete) and decidable class. Second, the relationship between Datalog and any Dyadic- $\mathcal{C}$  is made explicit, emphasizing the low expressive power required for  $\mathcal{C}$  to ensure that Dyadic- $\mathcal{C}$  fully encompasses Datalog. Finally, the computational complexity analysis is completed in Section 6, where both data and combined complexity for any Dyadic- $\mathcal{C}$  class are systematically studied.

## 2 Preliminaries

In this section, we introduce the syntax and the semantics of the class of rules that generalizes Datalog with existential quantifiers in rule-heads. Regarding computational complexity, we assume the reader is familiar with the basic complexity classes used in the subsequent sections:  $\mathbf{AC}_0 \subseteq \mathbf{PTIME} \subseteq \mathbf{NP} \subseteq \mathbf{PSPACE} \subseteq \mathbf{EXPTIME} \subseteq \mathbf{2EXPTIME}$ . Moreover, for a complexity class  $\mathbb{C}$ , we denote by  $\mathbf{PTIME}^{\mathbb{C}}$  (resp.,  $\mathbf{EXPTIME}^{\mathbb{C}}$ ) the class of decision problems that can be solved by an oracle Turing machine operating in polynomial (resp., exponential) time with the aid of an oracle that decides a problem in  $\mathbb{C}$ .

## 2.1 Basics on relational structures

Fix three pairwise disjoint lexicographically enumerable infinite sets  $\mathbf{C}$  of *constants*,  $\mathbf{N}$  of *nulls* ( $\varphi, \varphi_0, \varphi_1, \dots$ ), and  $\mathbf{V}$  of *variables* ( $x, y, z$ , and variations thereof). Their union is denoted by  $\mathbf{T}$  and its elements are called *terms*. For any integer  $k \geq 0$ , we may write  $[k]$  for the set  $\{1, \dots, k\}$ ; in particular, as usual, if  $k = 0$ , then  $[k] = \emptyset$ .

An *atom*  $\underline{a}$  is an expression of the form  $P(\mathbf{t})$ , where  $\text{preds}(\underline{a}) = P$  is a (*relational predicate*),  $\mathbf{t} = t_1, \dots, t_k$  is a tuple of *terms*  $\text{arity}(\underline{a}) = \text{arity}(P) = k \geq 0$  is the *arity* of both  $\underline{a}$  and  $P$ , and  $\underline{a}[i]$  denotes the  $i$ -th term  $\mathbf{t}[i] = t_i$  of  $\underline{a}$ , for each  $i \in [k]$ . In particular, if  $k = 0$ , then  $\mathbf{t}$  is the empty tuple and  $\underline{a} = P()$ . By  $\text{consts}(\underline{a})$  and  $\text{vars}(\underline{a})$  we denote, respectively, the set of constants and variables occurring in  $\underline{a}$ . A *fact* is an atom that contains only constants.

A (*relational schema*)  $\mathbf{S}$  is a finite set of predicates, each with its own arity. The set of *positions* of  $\mathbf{S}$ , denoted by  $\text{pos}(\mathbf{S})$ , is defined as the set  $\{P[i] \mid P \in \mathbf{S} \wedge 1 \leq i \leq \text{arity}(P)\}$ , where each  $P[i]$  denotes the  $i$ -th *position* of  $P$ . A (*relational structure*) over  $\mathbf{S}$  is any (possibly infinite) set of atoms using only predicates from  $\mathbf{S}$ . The *domain* of a structure  $S$ , denoted by  $\text{dom}(S)$ , is the set of all the terms forming the atoms of  $S$ . An *instance* over  $\mathbf{S}$  is any structure  $I$  over  $\mathbf{S}$  such that  $\text{dom}(I) \subseteq \mathbf{C} \cup \mathbf{N}$ . A *database* over  $\mathbf{S}$  is any finite instance over  $\mathbf{S}$  containing only facts. The *active domain* of an instance  $I$ , denoted by  $\text{dom}(I)$ , is the set of all the terms occurring in  $I$ , whereas the *Herbrand Base* of  $I$ , denoted by  $\text{HB}(I)$ , is the set of all the atoms that can be formed using the predicate symbols of  $\mathbf{S}$  and terms of  $\text{dom}(I)$ .

Consider two sets of terms  $T_1$  and  $T_2$  and a map  $\mu : T_1 \rightarrow T_2$ . Given a set  $T$  of terms, the *restriction* of  $\mu$  with respect to  $T$  is the map  $\mu|_T = \{t \mapsto \mu(t) : t \in T_1 \cap T\}$ . An *extension* of  $\mu$  is any map  $\mu'$  between terms, denoted by  $\mu' \supseteq \mu$ , such that  $\mu'|_{T_1} = \mu$ . A *homomorphism* from a structure  $S_1$  to a structure  $S_2$  is any map  $h : \text{dom}(S_1) \rightarrow \text{dom}(S_2)$  such that both the following hold: (i) if  $t \in \mathbf{C} \cap \text{dom}(S_1)$ , then  $h(t) = t$ ; and (ii)  $h(S_1) = \{P(h(\mathbf{t})) : P(\mathbf{t}) \in S_1\} \subseteq S_2$ .

## 2.2 Conjunctive queries

A *conjunctive query* (CQ)  $q$  over a schema  $\mathbf{S}$  is a (first-order) formula of the form

$$\langle \mathbf{x} \rangle \leftarrow \exists \mathbf{y} \Phi(\mathbf{x}, \mathbf{y}), \quad (1)$$

where  $\mathbf{x}$  and  $\mathbf{y}$  are tuples (often seen as sets) of variables such that  $\mathbf{x} \cap \mathbf{y} = \emptyset$ , and  $\Phi(\mathbf{x}, \mathbf{y})$  is a conjunction (often seen as a set) of atoms using only predicates from  $\mathbf{S}$ . In particular,

- $\text{dom}(\Phi) \subseteq \mathbf{x} \cup \mathbf{y} \cup \mathbf{C}$ ,
- whenever a variable  $z$  belongs to  $\mathbf{x} \cup \mathbf{y}$ , then  $z$  occurs also in  $\Phi$ ,
- $\mathbf{x}$  are the *output* variables of  $q$ , and
- $\mathbf{y}$  are the *existential* variables of  $q$ .

To highlight the output variables, we may write  $q(\mathbf{x})$  instead of  $q$ . The *evaluation* of  $q$  over an instance  $I$  is the set  $q(I)$  of every tuple  $\mathbf{t}$  of constants admitting a homomorphism  $h_{\mathbf{t}}$  from  $\Phi(\mathbf{x}, \mathbf{y})$  to  $I$  such that  $h_{\mathbf{t}}(\mathbf{x}) = \mathbf{t}$ .

A *Boolean conjunctive query* (BCQ) is a CQ with no output variable, namely an expression of the form  $\langle \rangle \leftarrow \exists \mathbf{y} \Phi(\mathbf{y})$ . An instance  $I$  satisfies a BCQ  $q$ , denoted  $I \models q$ , if  $q(I)$  is nonempty, namely  $q(I)$  contains only the empty tuple  $\langle \rangle$ .

### 2.3 Tuple-generating dependencies

A *tuple-generating dependency* (TGD)  $\sigma$  – also known as (*existential*) *rule* – over a schema  $\mathbf{S}$  is a (first-order) formula of the form

$$\Phi(\mathbf{x}, \mathbf{y}) \rightarrow \exists \mathbf{z} \Psi(\mathbf{x}, \mathbf{z}), \tag{2}$$

where  $\mathbf{x}$ ,  $\mathbf{y}$ , and  $\mathbf{z}$  are pairwise disjoint tuples of variables, and both  $\Phi(\mathbf{x}, \mathbf{y})$  and  $\Psi(\mathbf{x}, \mathbf{z})$  are conjunctions (often seen as a sets) of atoms using only predicates from  $\mathbf{S}$ . In particular,

- $\Phi$  (resp.,  $\Psi$ ) contains all and only the variables in  $\mathbf{x} \cup \mathbf{y}$  (resp.,  $\mathbf{x} \cup \mathbf{z}$ ),
- constants (but not nulls) may also occur in  $\sigma$ ,
- $\mathbf{x} \cup \mathbf{y}$  are the *universal* variables of  $\sigma$  denoted by  $\text{vars}_{\forall}(\sigma)$ ,
- $\mathbf{z}$  are the *existential* variables of  $\sigma$  denoted by  $\text{vars}_{\exists}(\sigma)$ , and
- $\mathbf{x}$  are the *frontier* variables of  $\sigma$  denoted by  $\text{vars}_{\sim}(\sigma)$ .

We refer to  $\text{body}(\sigma) = \Phi$  and  $\text{head}(\sigma) = \Psi$  as the *body* and *head* of  $\sigma$ , respectively. If  $|\text{head}(\sigma)| = 1$ , the TGD is called *single-head*, otherwise it called *multi-head*. If  $\text{vars}_{\exists}(\sigma) = \emptyset$  and  $|\text{head}(\sigma)| = 1$ , then  $\sigma$  is called *datalog* rule. With  $h\text{-preds}(\sigma)$  (resp.,  $b\text{-preds}(\sigma)$ ) we denote the set of predicates in  $\text{head}(\sigma)$  (resp.,  $\text{body}(\sigma)$ ). An instance  $I$  satisfies  $\sigma$ , written  $I \models \sigma$ , if the existence of a homomorphism  $h$  from  $\Phi$  to  $I$  implies the existence of a homomorphism  $h' \supseteq h|_{\mathbf{x}}$  from  $\Psi$  to  $I$ .

An *ontology*  $\Sigma$  is a set of rules. An instance  $I$  satisfies  $\Sigma$ , written  $I \models \Sigma$ , if  $I \models \sigma$  for each  $\sigma \in \Sigma$ . Without loss of generality, we assume that  $\text{vars}(\sigma_1) \cap \text{vars}(\sigma_2) = \emptyset$ , for each pair  $\sigma_1, \sigma_2$  of rules in  $\Sigma$ . Operators  $\text{vars}_{\exists}$ ,  $h\text{-preds}$ , and  $b\text{-preds}$  naturally extend on ontologies.

A class  $\mathcal{C}$  of ontologies is any (typically infinite) set of TGDs fulfilling some syntactic or semantic conditions (see, e.g., the classes shown in Table 1, some of which will be formally defined in the subsequent sections). In particular, **Datalog** is the class of ontologies containing only datalog rules.

Finally, the *schema* of an ontology  $\Sigma$ , denoted  $\text{sch}(\Sigma)$ , is the subset of  $\mathbf{S}$  containing all and only the predicates occurring in  $\Sigma$ , whereas  $\text{arity}(\Sigma) = \max_{P \in \text{sch}(\Sigma)} \text{arity}(P)$ . For simplicity of exposition, we write  $\text{pos}(\Sigma)$  instead of  $\text{pos}(\text{sch}(\Sigma))$ .

### 2.4 Ontological query answering

Consider a database  $D$  and a set  $\Sigma$  of TGDs. A *model* of  $D$  and  $\Sigma$  is an instance  $I$  such that  $I \supseteq D$  and  $I \models \Sigma$ . Let  $\text{mods}(D, \Sigma)$  be the set of all models of  $D$  and  $\Sigma$ . The *certain answers* to a CQ  $q$  w.r.t.  $D$  and  $\Sigma$  are defined as the set of tuples  $\text{cert}(q, D, \Sigma) = \bigcap_{M \in \text{mods}(D, \Sigma)} q(M)$ . Accordingly, for any fixed schema  $\mathbf{S}$ , two ontologies  $\Sigma_1$  and  $\Sigma_2$  over  $\mathbf{S}$  are said to be  *$\mathbf{S}$ -equivalent* (in symbols  $\Sigma_1 \equiv_{\mathbf{S}} \Sigma_2$ ) if, for each  $D$  and  $q$  over  $\mathbf{S}$ , it holds that  $\text{cert}(q, D, \Sigma_1) = \text{cert}(q, D, \Sigma_2)$ . The pair  $D$  and  $\Sigma$  satisfies a BCQ  $q$ , written  $D \cup \Sigma \models q$ , if  $\text{cert}(q, D, \Sigma) = \langle \rangle$ , namely  $M \models q$  for each  $M \in \text{mods}(D, \Sigma)$ . Fix a class  $\mathcal{C}$  of ontologies. The computational problem studied in this work – called  $\text{CERT-EVAL}[\mathcal{C}]$  – can be schematized as follows:

PROBLEM:  $\text{CERT-EVAL}[\mathcal{C}]$ .  
 INPUT: A database  $D$ , an ontology  $\Sigma \in \mathcal{C}$ , a conjunctive query  $q(\mathbf{x})$ , and a tuple  $\mathbf{c} \in \mathbb{C}^{|\mathbf{x}|}$ .  
 QUESTION: Does  $\mathbf{c} \in \text{cert}(q, D, \Sigma)$  hold?

In what follows, with a slight abuse of terminology, whenever we say that  $\mathcal{C}$  is decidable, we mean that  $\text{CERT-EVAL}[\mathcal{C}]$  is decidable. Note that  $\mathbf{c} \in \text{cert}(q, D, \Sigma)$  if, and only if,  $D \cup \Sigma \models q(\mathbf{c})$ , where  $q(\mathbf{c})$  is the BCQ obtained from  $q(\mathbf{x})$  by replacing, for each  $i \in \{1, \dots, |\mathbf{x}|\}$ , every occurrence of the variable  $\mathbf{x}[i]$  with the constant  $\mathbf{c}[i]$ . Actually, the former problem is  $\mathbf{AC}_0$  reducible to the latter.

While considering the computational complexity of  $\text{CERT-EVAL}[\mathcal{C}]$ , we recall the following convention: (i) *combined complexity* means that  $D, \Sigma, q$ , and  $\mathbf{c}$  are given in input; and (ii) *data complexity* means that only  $D$  and  $\mathbf{c}$  are given in input, whereas  $\Sigma$  and  $q$  are considered fixed. Accordingly, we point out that complexity results reported in Table 1 refer to  $\text{CERT-EVAL}[\mathcal{C}]$  under this convention.

### 2.5 The chase procedure

The chase procedure (Deutsch et al. 2008) is a tool exploited for reasoning with TGDs. Consider a database  $D$  and a set  $\Sigma$  of TGDs. Given an instance  $I \supseteq D$ , a *trigger* for  $I$  is any pair  $\langle \sigma, h \rangle$ , where  $\sigma \in \Sigma$  is a rule as in equation (2) and  $h$  is a homomorphism from  $\text{body}(\sigma)$  to  $I$ . Let  $I' = I \cup h'(\text{head}(\sigma))$ , where  $h' \supseteq h|_{\mathbf{x}}$  maps each  $z \in \text{vars}_{\exists}(\sigma)$  to a “fresh” null  $h'(z)$  not occurring in  $I$  such that  $z_1 \neq z_2$  in  $\text{vars}_{\exists}(\sigma)$  implies  $h'(z_1) \neq h'(z_2)$ . Such an operation which constructs  $I'$  from  $I$  is called *chase step* and denoted  $\langle \sigma, h \rangle(I) = I'$ .

Without loss of generality, we assume that nulls introduced at each trigger functionally depend on the pair  $\langle \sigma, h \rangle$  that is involved in the trigger. For example, given a rule  $\sigma$  as in equation (2) and a homomorphism  $h$ , it is sufficient to pick  $\varphi_{\langle \mathbf{z}, h(\mathbf{x}, \mathbf{y}) \rangle}$  as the fresh null replacing  $\mathbf{z}$  when the chase produces the trigger  $\langle \sigma, h \rangle$ . Accordingly, the processing order of rules and triggers does not change the result of the chase, and hence  $\text{chase}(D, \Sigma)$  can be considered unique. The *chase procedure* of  $D \cup \Sigma$  is an exhaustive application of chase steps, starting from  $D$ , which produce a sequence  $I_0 = D \subset I_1 \subset I_2 \subset \dots \subset I_m \subset \dots$  of instances in such a way that: (i) for each  $i \geq 0$ ,  $I_{i+1} = \langle \sigma, h \rangle(I_i)$  is a chase step obtained via some trigger  $\langle \sigma, h \rangle$  for  $I_i$ ; (ii) for each  $i \geq 0$ , if there exists a trigger  $\langle \sigma, h \rangle$  for  $I_i$ , then there exists some  $j > i$  such that  $I_j = \langle \sigma, h \rangle(I_{j-1})$  is a chase step; and (iii) any *trigger*  $\langle \sigma, h \rangle$  is used only once. We define  $\text{chase}(D, \Sigma) = \cup_{i \geq 0} I_i$ .

The *chase bottom* is the finite set of all null-free atoms in  $\text{chase}(D, \Sigma)$  and is defined as  $\text{chase}^{\perp}(D, \Sigma) = \text{chase}(D, \Sigma) \cap \text{HB}(D)$ .

It is well known that  $\text{chase}(D, \Sigma)$  is a *universal model* of  $D \cup \Sigma$ , that is, for each  $M \in \text{mods}(D, \Sigma)$  there is a homomorphism from  $\text{chase}(D, \Sigma)$  to  $M$ . Hence, given a BCQ  $q$  it holds that  $\text{chase}(D, \Sigma) \models q \Leftrightarrow D \cup \Sigma \models q$ .

We recall that  $\text{chase}(D, \Sigma)$  can be decomposed into *levels* (Calì et al. 2010): each atom of  $D$  has level  $\gamma = 0$ ; an atom of  $\text{chase}(D, \Sigma)$  has level  $\gamma + 1$  if, during its generation, the exploited trigger  $\langle \sigma, h \rangle$  maps the body of  $\sigma$  via  $h$  to atoms whose maximum level is  $\gamma$ . We refer to the part of the chase *up to level*  $\gamma$  as  $\text{chase}^{\gamma}(D, \Sigma)$ . Clearly,  $\text{chase}(D, \Sigma) = \cup_{\gamma \geq 0} \text{chase}^{\gamma}(D, \Sigma)$ . Finally, a trigger *involved* at a certain level  $j$  if it gives rise to an atom of level  $j$ .

## 3 Considered decidable classes of TGDs

In this section we provide an overview of the main existing decidable classes of TGDs. We recall both syntactic and semantic classes, where the first are based on a specific

syntactic condition that can be checked, while the latter are classes that do not come with a syntactic property that can be checked on rules and, hence, are not recognizable. Finally, we introduce a very simple new class of existential rules called Af-Inds. We will exploit the latter to sharpen our results presented in Sections 5 and 6.

### 3.1 Preliminary notions

We start fixing some basics notions. We have chosen to provide a uniform notation for the key existing notions of affected and invaded positions, such as attacked, protected, harmless, harmful, and dangerous variables (Leone et al. 2019; Cali et al. 2013; Krötzsch and Rudolph 2011; Berger et al. 2022; Gottlob et al. 2022). Basically, these notions serve to separate positions in which the chase can introduce only constants from those where nulls might appear.

*Definition 1 (S-affected positions)*

Consider an ontology  $\Sigma$  and a variable  $z \in vars_{\exists}(\Sigma)$ . A position  $\pi \in pos(\Sigma)$  is *z-affected* (or *invaded* by  $z$ ) if one of the following two properties holds: (i) there exists  $\sigma \in \Sigma$  such that  $z$  appears in the head of  $\sigma$  at position  $\pi$ ; (ii) there exist  $\sigma \in \Sigma$  and  $x \in front(\sigma)$  such that  $x$  occurs both in  $head(\sigma)$  at position  $\pi$  and in  $body(\sigma)$  at  $z$ -affected positions only. Moreover, a position  $\pi \in pos(\Sigma)$  is *S-affected*, where  $S \subseteq vars_{\exists}(\Sigma)$ , if: (i) for each  $z \in S$ ,  $\pi$  is  $z$ -affected; and (ii) for each  $z \in vars_{\exists}(\Sigma)$ , if  $\pi$  is  $z$ -affected, then  $z \in S$ . ■

Note that for every position  $\pi$  there exists a unique set  $S$  such that  $\pi$  is  $S$ -affected. We write  $aff(\pi)$  for this set  $S$ . Moreover,  $aff(\Sigma) = \{\pi \in pos(\Sigma) \mid aff(\pi) \neq \emptyset\}$ , and  $nonaff(\Sigma) = pos(\Sigma) \setminus aff(\Sigma)$ .

We point out that the notion above presented is a refined version of the classical notion of affected position (Cali et al. 2008). In particular, it holds that if a position  $\pi$  is  $S$ -affected, then  $\pi$  is also affected; whereas if  $\pi$  is affected, then  $\pi$  may not be  $S$ -affected. Moreover, the  $S$ -affected notion coincides with the one of *attacked positions by a variable* (Leone et al. 2019; Krötzsch and Rudolph 2011). We highlight that its key nature and properties are not modified by the notion of  $S$ -affected position introduced above. Hence, for simplicity of exposition, we give only this refined definition. In the same spirit, we classify variables occurring in a conjunction of atoms.

*Definition 2 (Variables classification)*

Let  $\sigma$  be a TGD and  $x \in vars(body(\sigma))$ . Then, (i) if  $x$  occurs at positions  $\pi_1, \dots, \pi_n$  and  $\bigcap_{i=1}^n aff(\pi_i) = \emptyset$ ,  $x$  is said to be *harmless*; (ii) if  $x$  is not harmless, with  $S = \bigcap_{i=1}^n aff(\pi_i)$ , it is said to be *S-harmful*; (iii) if  $x$  is  $S$ -harmful and belongs to  $vars_{\sim}(\sigma)$ ,  $x$  is *S-dangerous*. ■

Given a variable  $x$  that is  $S$ -dangerous, we write  $dang(x)$  for the set  $S$ . Hereinafter, for simplicity of exposition, the prefix  $S$ - is omitted when it is not necessary. Consider an ontology  $\Sigma$ . Given a rule  $\sigma \in \Sigma$ , we denote by  $dang(\sigma)$  (resp.,  $harmless(\sigma)$  and  $harmful(\sigma)$ ) the dangerous (resp., harmless and harmful) variables in  $\sigma$ . These sets of variables naturally extend to the whole  $\Sigma$  by taking, for each of them, the union over all the rules of  $\Sigma$ .

### 3.2 Decidable classes of existential rules

We now survey the 15 concrete classes reported in Table 1 as well as the known abstract classes based on semantic conditions. On the one side, we report some specific syntactic conditions whenever these are useful for the rest of the presentation; on the other side, for all of them (both concrete and abstract), we recall their containment relationships. For the rest of the section, fix a Datalog<sup>3</sup> ontology  $\Sigma$ .

The class FES (Baget et al. 2009) stands for finite expansions sets, which intuitively are sets of TGDs which ensure the termination of the chase procedure. The class BTS (Baget et al. 2009) stands for bounded treewidth sets, which intuitively are sets of TGDs which guarantee that the (possibly infinite) instance constructed by the chase procedure has bounded treewidth. The class FUS (Baget et al. 2011) stands for finite unification sets, which intuitively are sets of TGDs which guarantee the termination of (resolution-based) backward chaining procedures. The class SPS (Leone et al. 2019) stands for strongly parsimonious sets, which intuitively are sets of TGDs which guarantee that the parsimonious chase procedure can be reapplied a number of times that is linear in the size of the query.

The class Weakly-Acyclic (Fagin et al. 2005) is based on the *acyclicity* condition. To define the latter, we recall the *label graph* of  $\Sigma$ ,  $G(\Sigma) = \langle N, A \rangle$ , defined as follows: (i)  $N = \cup_{P \in \text{sch}(\Sigma)} \text{pos}(P)$ ; (ii)  $(\pi_1, \pi_2, \forall) \in A$  if there are  $\sigma \in \Sigma$  and  $x \in \text{vars}_{\wedge}(\sigma)$  such that  $x$  occurs both in  $\text{body}(\sigma)$  at position  $\pi_1$  and in  $\text{head}(\sigma)$  at position  $\pi_2$ ; and (iii)  $(\pi_1, \pi_2, \exists) \in A$  if there are  $\sigma \in \Sigma$ ,  $x \in \text{vars}_{\wedge}(\sigma)$ , and  $y \in \text{vars}_{\exists}(\sigma)$  such that both  $x$  occurs in  $\text{body}(\sigma)$  at position  $\pi_1$  and  $y$  occurs in  $\text{head}(\sigma)$  at position  $\pi_2$ . The *existential graph* of  $\Sigma$  is  $G_{\exists}(\Sigma) = \langle N, A \rangle$ , where: (i)  $N = \cup_{\sigma \in \Sigma} \text{vars}_{\exists}(\sigma)$ ; and (ii)  $(X, Y) \in A$  if the rule  $\sigma$  where  $y$  occurs contains a universal variable  $x$ -affected and occurring in  $\text{head}(\sigma)$ . Therefore,  $\Sigma$  belongs to Weakly-Acyclic (resp., Jointly-Acyclic) if  $G(\Sigma)$  (resp.,  $G_{\exists}(\Sigma)$ ) has no cycle going through an  $\exists$ -arc (resp., is acyclic).

We now recall the notion of marked variable, in order to define the class Sticky (Cali et al. 2012b). A variable  $x$  of  $\Sigma$  is *marked* if (i) there is  $\sigma \in \Sigma$  such that  $x$  occurs in  $\text{body}(\sigma)$  but not in  $\text{head}(\sigma)$ ; or (ii) there is  $\sigma \in \Sigma$  such that  $x$  occurs in  $\text{head}(\sigma)$  at position  $\pi$  together with some  $\sigma' \in \Sigma$  having a marked variable in its body at position  $\pi$ . Accordingly, the *stickiness* condition states that  $\Sigma$  is Sticky if, for each  $\sigma \in \Sigma$ ,  $x$  occurs multiple times in  $\text{body}(\sigma)$  implies  $x$  is not marked.

The class Linear (Cali et al. 2012a) is based on the *linearity* condition: an ontology  $\Sigma$  belongs to Linear if each rule contains at most one body atoms. This class generalizes the class Inclusion-Dependencies (Abiteboul et al. 1995; Johnson and Klug 1984) in which rules contain only one body atom and one head atom and the repetition of variables is not allowed neither in the body nor in the head.

The class Guarded (Cali et al. 2013) is based on the *guardedness* condition: an ontology  $\Sigma$  belongs to Guarded if for each rule  $\sigma \in \Sigma$  there is  $\underline{a}$  in  $\text{body}(\sigma)$  such that  $\text{vars}_{\forall}(\sigma) = \text{vars}_{\forall}(\underline{a})$ . In similar fashion,  $\Sigma$  belongs to Weakly-Guarded if, for each  $\sigma \in \Sigma$ , there is an atom of  $\text{body}(\sigma)$  containing all the affected variables of  $\sigma$ .

We recall the *shyness* condition underlying the class Shy (Leone et al. 2019). An ontology  $\Sigma$  is Shy if, for each  $\sigma \in \Sigma$  the following conditions both hold: (i) if a variable  $x$



occurs in more than one body atom, then  $x$  is harmless; (ii) for every pair of distinct dangerous variable  $z$  and  $w$  in different atoms,  $dang(z) \cap dang(w) = \emptyset$ .

The class Ward (Gottlob and Pieris 2015) is based on the *wardedness* condition:  $\Sigma \in \text{Ward}$  if, for each  $\sigma \in \Sigma$ , there are no dangerous variables in  $body(\sigma)$ , or there exists an atom  $\underline{a} \in body(\sigma)$ , called a *ward*, such that (i) all the dangerous variables in  $body(\sigma)$  occur in  $\underline{a}$ , and (ii) each variable of  $vars(\underline{a}) \cap vars(body(\sigma) \setminus \{\underline{a}\})$  is harmless.

Having finished with syntactic and semantic conditions, we close the section with a proposition stating their containment relationships (Baget et al. 2011; Krötzsch and Rudolph 2011; Leone et al. 2019; Baldazzi et al. 2022).

*Proposition 1*

The following classes are pairwise uncomparable, except for:

- Inclusion-Dependencies  $\subset$  Joinless, Inclusion-Dependencies  $\subset$  Linear;
- Joinless  $\subset$  Sticky  $\subset$  Sticky-Join  $\subset$  FUS;
- Linear  $\subset$  Guarded, Linear  $\subset$  Protected;
- Guarded  $\subset$  Weakly-Guarded, Guarded  $\subset$  Fr-Guarded;
- Weakly-Guarded  $\subset$  Weakly-Fr-Guarded  $\subset$  BTS;
- Fr-Guarded  $\subset$  Weakly-Fr-Guarded;
- Datalog  $\subset$  Weakly-Guarded, Datalog  $\subset$  Protected, Datalog  $\subset$  Weakly-Acyclic;
- Protected  $\subset$  Ward, Protected  $\subset$  Shy  $\subset$  SPS;
- Weakly-Acyclic  $\subset$  Jointly-Acyclic  $\subset$  FES.

Throughout the remainder of the paper, let  $\mathbb{E}_{syn}$  denote the set of all 15 decidable syntactic classes reported in Table 1. Analogously, let  $\mathbb{E}_{sem}$  denote the set of known decidable abstract classes considered in this paper, namely FES, FUS, BTS, and SPS.

### 3.3 Autonomous full inclusion dependencies

The aim of this section is to introduce a very simple new class of existential rules called Af-Inds. Additionally, we characterize the main properties of this class.

*Definition 3 (Af-Inds)*

An ontology  $\Sigma$  belongs to Af-Inds (autonomous full inclusion dependencies) if  $\Sigma$  belongs to Inclusion-Dependencies and the following conditions are also satisfied: (1) head predicates do not appear in bodies (autonomous property); (2) rules have no existential variables (full property).

Now, we show that any class  $\mathcal{C}$  of TGDs in  $\mathbb{E}_{syn} \cup \mathbb{E}_{sem}$  includes the class just defined. Formally, it holds the following.

*Proposition 2*

Consider a class  $\mathcal{C} \in \mathbb{E}_{syn} \cup \mathbb{E}_{sem}$  of TGDs. Then, Af-Inds  $\subseteq \mathcal{C}$ .

*Proof*

Thanks to Proposition 1, the statement becomes equivalent to show that (i) Af-Inds  $\subseteq$  Inclusion-Dependencies and (ii) Af-Inds  $\subseteq$  Datalog. By Definition 3, the class Af-Inds contains all the rules that have only one body and head atom, without repetition of variables

neither in the body nor in the head, and that satisfy the autonomous property (head atom does not appear in bodies) and the full property (rules have only one head atom without existential variables). Accordingly, relation (i) and (ii) are trivially fulfilled.  $\square$

We conclude the section by providing the complexity of the class Af-Inds.

*Proposition 3*

CERT-EVAL[Af-Inds] is in  $\mathbf{AC}_0$  in data complexity and  $\mathbf{NP}$ -complete in combined complexity.

*Proof*

By Proposition 2, Af-Inds  $\subseteq$  Inclusion-Dependencies. Hence, the data complexity of the problem CERT-EVAL[Af-Inds] is inherit from that of CERT-EVAL[Inclusion-Dependencies], that is  $\mathbf{AC}_0$ . For the combined complexity, we first observe that the problem CERT-EVAL[Af-Inds] is  $\mathbf{NP}$ -hard, building upon the well-known fact that CERT-EVAL[ $\emptyset$ ] is already  $\mathbf{NP}$ -hard. The latter refers to the problem of evaluating a query against a database in the absence of an ontology. Secondly, to prove the completeness of the CERT-EVAL[Af-Inds] problem, we show that given a query  $q(\mathbf{x})$  and an ontology  $\Sigma$ , it is possible to construct in  $\mathbf{NP}$  a CQ  $q_\Sigma(\mathbf{x})$  such that  $c \in \text{cert}(D, \Sigma, q)$  iff  $c \in q_\Sigma(D)$ , with  $\mathbf{c}$  being a tuple in  $C^{|\mathbf{x}|}$ . To this aim, for each atom  $\underline{a} \in q(\mathbf{x})$ , we guess if leave  $\underline{a}$  unchanged, or “resolve”  $\underline{a}$  with the body of some rule  $\sigma$  in  $\Sigma$  such that  $\text{head}(\sigma)$  unify with  $\underline{a}$ . Accordingly,  $q_\Sigma(D)$  is polynomial with respect to the input and, finally, it is possible to guess in  $\mathbf{NP}$  an homomorphism to check if  $c \in q_\Sigma(D)$ .  $\square$

### 4 Dyadic pairs of TGDs

In this section we lay the groundwork for the main contribution of the paper, that is the definition of a new decidable class of TGDs called Dyadic- $\mathcal{C}$ . To this aim we first introduce some preliminary notions in order to define a *dyadic pair* and, then we conclude with some computational properties.

#### 4.1 Formal definition

We start introducing the concept of *head-ground* set of rules, being roughly “non-recursive” rules in which nulls are neither created nor propagated.

*Definition 4 (Head-ground rules)*

Consider an ontology  $\Sigma$ . A set  $\Sigma' \subseteq \Sigma$  is *head-ground* w.r.t.  $\Sigma$  if the following are true:

- (1)  $\Sigma' \in \text{Datalog}$ ; (2) each head atom of  $\Sigma'$  contains only harmless variables w.r.t.  $\Sigma$ ;
- (3)  $h\text{-preds}(\Sigma') \cap b\text{-preds}(\Sigma') = \emptyset$ ; and (4)  $h\text{-preds}(\Sigma') \cap h\text{-preds}(\Sigma \setminus \Sigma') = \emptyset$ .  $\blacksquare$

The following example is given to better understand the above definition.

*Example 1*

Consider the next set of rules:

$$\begin{aligned} \sigma_1 : & R(x_1, y_1), S(y_1, u_1), T(u_1, v_1) \rightarrow \exists z_1, w_1 Q(z_1, w_1) \\ \sigma_2 : & C(y_2), R(x_2, z_2) \rightarrow S(y_2, z_2) \\ \sigma_3 : & D(y_3, z_3), R(x_3, w_3) \rightarrow T(x_3, y_3) \end{aligned}$$

$$\begin{aligned} \sigma_4 : \quad & Q(x_4, y_4) \rightarrow \exists z_4 A(x_4, z_4) \\ \sigma_5 : \quad & A(x_5, z_5), D(y_5, z_5) \rightarrow Q(x_5, y_5) \end{aligned}$$

A subset of head-ground rule w.r.t.  $\Sigma$  is given by  $\Sigma_{\text{HG}} = \{\sigma_2, \sigma_3\}$ . In fact, *harmless*( $\Sigma$ ) is the set  $\{x_1, y_1, y_2, x_2, z_2, x_3, y_3, z_3, y_5, z_5\}$ ; hence, according to Definition 4, it is easy to check that (i)  $\sigma_2$  and  $\sigma_3$  are datalog rules; (ii) the head atoms of  $\sigma_2$  and  $\sigma_3$  contain only harmless variables; (iii) both predicates  $S$  and  $T$  do not occur in the body of any rule in  $\Sigma_{\text{HG}}$ , and (iv) both predicates  $S$  and  $T$  do not occur in the head of any rule in  $\{\sigma_1, \sigma_4, \sigma_5\}$ . On the contrary, none of the rules in  $\{\sigma_1, \sigma_4, \sigma_5\}$  can be part of any head-ground subset of  $\Sigma$ . Indeed, according to Definition 4, both  $\sigma_1$  and  $\sigma_2$  violate properties (1) and (2), whereas  $\sigma_5$  violates property (2). Hence, we observe that the set  $\Sigma_{\text{HG}}$  is also maximal. ■

Having in mind the notion of head-ground set of rules, we can now formally define what is a dyadic pair.

*Definition 5 (Dyadic pairs)*

Consider a class  $\mathcal{C}$  of TGDs. A pair  $\Pi = (\Sigma_{\text{HG}}, \Sigma_{\mathcal{C}})$  of TGDs is *dyadic* with respect to  $\mathcal{C}$  if the next hold: (1)  $\Sigma_{\text{HG}}$  is head-ground with respect to  $\Sigma_{\text{HG}} \cup \Sigma_{\mathcal{C}}$ ; and (2)  $\Sigma_{\mathcal{C}} \in \mathcal{C}$ . ■

Whenever the above definition applies, we also say, for short, that  $\Pi$  is a  $\mathcal{C}$ -dyadic pair. Consider the following example to more easily understand the concept of dyadic pair.

*Example 2*

Consider the following pair  $\Pi = (\Sigma_{\text{HG}}, \Sigma_{\mathcal{C}})$  of TGDs, where  $\Sigma_{\text{HG}}$  is:

$$\begin{aligned} P(x_1) &\rightarrow H_1(x_1) \\ P(x_2) &\rightarrow H_2(x_2) \\ Q(x_3) &\rightarrow H_3(x_3). \end{aligned}$$

and  $\Sigma_{\mathcal{C}}$  is:

$$\begin{aligned} H_1(x_1) &\rightarrow \exists y_1, z_1 R(y_1, z_1) \\ H_2(x_2) &\rightarrow \exists y_2 Q(y_2) \\ R(y_3, x_3), H_3(x_3) &\rightarrow S(y_3). \end{aligned}$$

In particular,  $\Pi$  is a dyadic pair with respect to any  $\mathcal{C} \in \{\text{Guarded}, \text{Shy}, \text{Ward}\}$ . To this aim, let  $\Sigma = \Sigma_{\text{HG}} \cup \Sigma_{\mathcal{C}}$ . It easy computable that  $\text{aff}(\Sigma) = \{R[1], R[2], Q[1], S[1]\}$ , where  $\text{aff}(R[1]) = \{y_1\}$ ,  $\text{aff}(R[2]) = \{z_1\}$ ,  $\text{aff}(Q[1]) = \{y_2\}$ , and  $\text{aff}(S[1]) = \{y_1\}$ . Accordingly,  $\text{harmless}(\Sigma) = \{x_1, x_2, x_3\}$ ,  $\text{harmful}(\Sigma) = \{y_3\}$  and  $\text{dang}(\Sigma) = \{y_3\}$ . To prove that  $\Pi$  is a dyadic pair, we have first to show that  $\Sigma_{\text{HG}}$  is an head-ground set of rules with respect to  $\Sigma$ . Clearly,  $\Sigma_{\text{HG}} \in \text{Datalog}$  and each head atom contains only harmless variables; moreover, the head predicates do not appear neither in body atoms of  $\Sigma_{\text{HG}}$  nor in head atoms of  $\Sigma_{\mathcal{C}}$ . Hence,  $\Sigma_{\text{HG}}$  is head-ground with respect to  $\Sigma$ . It remains to show that  $\Sigma_{\mathcal{C}} \in \mathcal{C}$ . We focus on the last rule of  $\Sigma_{\mathcal{C}}$ , since the first two rules are linear rules, and hence are trivially guarded, shy and ward rules. The last rule belongs to **Guarded** since the atom  $R(y_3, x_3)$  contains all the universal variables of the rule (guardedness condition); it belong to **Shy** since the variable  $x_3$  that occurs in two body atoms is harmless (shyness condition); finally, it belongs to **Ward** since atom  $R(y_3, x_3)$  is the ward that contains the dangerous variables ( $y_3$ ) and shares with the rest of the body only harmless variables ( $x_3$ ) (wardedness condition). ■

The next step is to extend the query answering problem – classically defined over an ontology – over a dyadic pair. Therefore, we extend both notions of chase and certain answers for a dyadic pair. Accordingly, given a dyadic pair  $\Pi = (\Sigma_{HG}, \Sigma_C)$ , we define

$$dp\text{-chase}(D, \Pi) = \text{chase}(D, \Sigma_{HG} \cup \Sigma_C) \tag{3}$$

and

$$dp\text{-cert}(q, D, \Pi) = \text{cert}(q, D, \Sigma_{HG} \cup \Sigma_C). \tag{4}$$

Now we can fix the problem studied in the rest of the paper.

PROBLEM: DP-CERT-EVAL[ $\mathcal{C}$ ].  
 INPUT: A database  $D$ , a  $\mathcal{C}$ -dyadic pair  $\Pi = (\Sigma_{HG}, \Sigma_C)$  of TGDs, a CQ  $q(\mathbf{x})$ , and a tuple  $\mathbf{c} \in \mathcal{C}^{|\mathbf{x}|}$ .  
 QUESTION: Does  $\mathbf{c} \in dp\text{-cert}(q, D, \Pi)$  holds?

### 4.2 Computational properties

For the rest of the section, fix a decidable class  $\mathcal{C}$  of TGDs. Given a database  $D$  and a  $\mathcal{C}$ -dyadic pair  $\Pi$  of TGDs, we define the following set of ground atoms:

$$gra(D, \Pi) = \{\underline{a} \in \text{chase}(D, \Pi) \mid \Pi = (\Sigma_{HG}, \Sigma_C) \wedge \text{preds}(\underline{a}) \in h\text{-preds}(\Sigma_{HG})\}. \tag{5}$$

Our idea is to reduce query answering over a dyadic pair  $\Pi$  to query answering over  $\mathcal{C}$ , the latter being decidable by assumption.

*Theorem 1*

Consider a database  $D$ , a  $\mathcal{C}$ -dyadic pair  $\Pi = (\Sigma_{HG}, \Sigma_C)$  of TGDs, and a conjunctive query  $q(\mathbf{x})$ . Let  $D^+ = D \cup gra(D, \Pi)$ . It holds that  $dp\text{-cert}(q, D, \Pi) = \text{cert}(q, D^+, \Sigma_C)$ .

*Proof*

Consider  $\Pi = (\Sigma_{HG}, \Sigma_C)$ . By equation (4),  $dp\text{-cert}(q, D, \Pi) = \text{cert}(q, D, \Sigma_{HG} \cup \Sigma_C)$ . Moreover, by fixing any arbitrary  $|\mathbf{x}|$ -ary tuple  $\mathbf{c}$  of constants, it holds that  $\mathbf{c} \in \text{cert}(q, D, \Sigma_{HG} \cup \Sigma_C) \Leftrightarrow D \cup \Sigma_{HG} \cup \Sigma_C \models q(\mathbf{c})$  and  $\mathbf{c} \in \text{cert}(q, D^+, \Sigma_C) \Leftrightarrow D^+ \cup \Sigma_C \models q(\mathbf{c})$ . Let  $q' = q(\mathbf{c})$ . Accordingly, the thesis boils down to showing that

$$D \cup \Sigma_{HG} \cup \Sigma_C \models q' \Leftrightarrow D^+ \cup \Sigma_C \models q'.$$

[ $\Rightarrow$ ] Assume that  $D \cup \Sigma_{HG} \cup \Sigma_C \models q'$  holds. Hence,  $\text{chase}(D, \Sigma_{HG} \cup \Sigma_C) \models q'$ . Given that  $gra(D, \Pi) \subseteq \text{chase}(D, \Sigma_{HG} \cup \Sigma_C)$ , it holds that  $\text{chase}(D \cup gra(D, \Pi), \Sigma_{HG} \cup \Sigma_C) \models q'$ . Moreover, since  $gra(D, \Pi)$  contains all the auxiliary ground consequences of  $\Sigma_{HG}$ , the latter becomes equivalent to  $\text{chase}(D \cup gra(D, \Pi), \Sigma_C) \models q'$ . Hence,  $D \cup gra(D, \Pi) \cup \Sigma_C \models q'$ , that is  $D^+ \cup \Sigma_C \models q'$ .

[ $\Leftarrow$ ] Assume that  $D^+ \cup \Sigma_C \models q'$ , hence  $\text{chase}(D^+, \Sigma_C) \models q'$ . Since  $\Sigma_C \subseteq \Sigma_{HG} \cup \Sigma_C$ , it holds that  $\text{chase}(D^+, \Sigma_{HG} \cup \Sigma_C) \models q'$ . By hypothesis,  $gra(D, \Pi) \subseteq \text{chase}(D, \Sigma_{HG} \cup \Sigma_C)$ ; hence  $\text{chase}(D, \Sigma_{HG} \cup \Sigma_C) \models q'$ , that is  $D \cup \Sigma_{HG} \cup \Sigma_C \models q'$ . □

According to Theorem 1, to solve DP-CERT-EVAL[ $\mathcal{C}$ ] one can first “complete”  $D$  and then performing classical query evaluation. To this aim we design Algorithms 1 and 2. The

---

**Algorithm 1:**  $\text{DpCertEval}_{[C]}(q, D, \Pi, \mathbf{c})$

---

**Input:** A database  $D$ , a dyadic pair  $\Pi$ , a CQ  $q(\mathbf{x})$ , and a tuple  $\mathbf{c} \in C^{|\mathbf{x}|}$

- 1  $D^+ := \text{Complete}_{[C]}(D, \Pi)$ ;
- 2 **return**  $\mathbf{c} \in \text{cert}(q, D^+, \Sigma_C)$ ;

---

correctness of the latter will be proved in Proposition 4. Consequently, the correctness of the former is guaranteed by Theorem 1.

In particular, given a database  $D$  and a dyadic pair  $\Pi = (\Sigma_{\text{HG}}, \Sigma_C)$ , Algorithm 2 iteratively constructs the set  $D^+ = D \cup \text{gra}(D, \Pi)$ , with  $\text{gra}(D, \Pi)$  being the set defined by equation (5). Roughly speaking, the first two instructions are required, respectively, to add  $D$  to  $D^+$  and to initialize a temporary set  $\tilde{D}$  used to store ground consequences derived from  $\Sigma_{\text{HG}}$ . The rest of the algorithm is an iterative procedure that computes the certain answers (instruction 5) to the queries constructed from the rules of  $\Sigma_{\text{HG}}$  (instruction 4) and completes the initial database  $D$  (instruction 7) until no more auxiliary ground atoms can be produced (instruction 6). We point out that, in general,  $\tilde{D} \subseteq \text{gra}(D, \Pi)$  holds; in particular,  $\tilde{D} = \text{gra}(D, \Pi)$  holds in the last execution of instruction 7 or, equivalently, when the condition  $D \cup \tilde{D} \supset D^+$  examined at instruction 6 is false, since all the auxiliary ground atoms have been added to  $D^+$ .

Before we prove that Algorithm 2 always terminates and correctly constructs  $D^+$ , we show the following preliminary lemma.

---

**Algorithm 2:**  $\text{Complete}_{[C]}(D, \Pi)$

---

**Input:** A database  $D$  and a  $\mathcal{C}$ -dyadic pair  $\Pi = (\Sigma_{\text{HG}}, \Sigma_C)$

**Output:** The set  $D^+$  of ground atoms

- 1  $D^+ := D$
- 2  $\tilde{D} := \emptyset$
- 3 **for** each rule of the form  $\Phi(\mathbf{x}, \mathbf{y}) \rightarrow \text{H}_i(\mathbf{x})$  in  $\Sigma_{\text{HG}}$  **do**
- 4      $q := \langle \mathbf{x} \rangle \leftarrow \Phi(\mathbf{x}, \mathbf{y})$
- 5      $\tilde{D} = \tilde{D} \cup \{\text{H}_i(\mathbf{t}) \mid \mathbf{t} \in \text{cert}(q, D^+, \Sigma_C)\}$
- 6 **if**  $(D \cup \tilde{D} \supset D^+)$  **then**
- 7      $D^+ := D \cup \tilde{D}$
- 8     **go to** instruction 2
- 9 **return**  $D^+$

---

*Lemma 1*

Consider a database  $D$  and a set  $\Sigma$  of TGDs. Let  $\Sigma' \subseteq \Sigma$  and  $X \subseteq \text{chase}^1(D, \Sigma) \setminus D$ . Then,  $\text{chase}(D \cup X, \Sigma') \subseteq \text{chase}(D, \Sigma)$ .

*Proof*

Let  $X = \{\underline{a}_1, \dots, \underline{a}_n\}$ . For each  $j \in [n]$ , let  $\text{lev}(\underline{a}_i)$  be the level of  $\underline{a}_i$  inside  $\text{chase}(D, \Sigma)$ . Let  $p = \max_{j \in [n]} \{\text{lev}(\underline{a}_i)\}$ . The proof proceeds by induction on the level  $i$  of  $\text{chase}(D \cup X, \Sigma')$ .

*Base case:*  $i = 1$ . We want to prove that  $\text{chase}^1(D \cup X, \Sigma') \subseteq \text{chase}(D, \Sigma)$ . Let  $\underline{a}$  be an atom of  $\text{chase}^1(D \cup X, \Sigma')$  generated exactly at level  $i = 1$ . By definition,  $\underline{a}$  is obtained

due to some trigger  $\langle \sigma, h \rangle$  such that  $h$  maps  $body(\sigma)$  to  $D \cup X$ . If  $\underline{a} \in chase^p(D, \Sigma)$ , then the claim holds trivially. Otherwise, we can show that  $\underline{a} \in chase^{p+1}(D, \Sigma)$ . Indeed, since  $h$  maps  $body(\sigma)$  to  $D \cup X$ , then  $h$  is also a trigger involved at level  $p + 1$  since it maps  $body(\sigma)$  to  $chase^p(D, \Sigma)$ . In particular,  $h$  maps at least one atom of  $body(\sigma)$  to some atom  $\underline{a}_k \in X$  such that  $lev(\underline{a}_k) = p$ . Since, by definition, nulls introduced during the chase functionally depend on the involved triggers, then  $\underline{a}$  necessarily belongs to  $chase^{p+1}(D, \Sigma)$ .

*Induction step:  $i = \ell$ .* Given that for every level  $i \leq \ell - 1$ ,  $chase^i(D \cup X, \Sigma') \subseteq chase(D, \Sigma)$  (*induction hypothesis*), we prove that  $chase^\ell(D \cup X, \Sigma') \subseteq chase(D, \Sigma)$  holds, too. Let  $\beta$  be an atom of  $chase^\ell(D \cup X, \Sigma')$  generated exactly at level  $i = \ell$ . By definition,  $\beta$  is obtained via some trigger  $\langle \sigma', h' \rangle$  such that  $h'$  maps  $body(\sigma')$  to atoms with level at most  $\ell - 1$ . Accordingly, by induction hypothesis,  $h'$  maps  $body(\sigma')$  also to  $chase(D, \Sigma)$ . Hence, since the processing order of rules and triggers does not change the result of the chase and nulls functionally depend on the involved triggers, it follows that also  $\beta \in chase(D, \Sigma)$ .  $\square$

With the next proposition, we prove that Algorithm 2 always terminates and correctly constructs  $D^+$ .

*Proposition 4*

Consider a database  $D$  and a  $\mathcal{C}$ -dyadic pair  $\Pi$  of TGDs. It holds that Algorithm 2 both terminates and computes  $D^+ = D \cup gra(D, \Pi)$ .

*Proof*

Let  $\Pi = (\Sigma_{HG}, \Sigma_C)$ . We proceed by proving first the termination of Algorithm 2 and then its correctness.

**Termination.** To prove the termination of Algorithm 2, it suffices to show that each instruction alone always terminates and that the overall procedure never falls into an infinite loop. First, observe that  $|gra(D, \Pi)| \leq |h\text{-preds}(\Sigma_{HG})| \cdot d^\mu$ , where  $d = |const(D)|$  and  $\mu = \max_{P \in h\text{-preds}(\Sigma_{HG})} \text{arity}(P)$ . Instructions 1, 2, 4, 8, and 9 trivially terminate. Instructions 6 and 7 both terminate, since  $\tilde{D} \subseteq gra(D, \Pi)$  always holds (see correctness below). Each time instruction 3 is reached, the **for**-loop simply scans the set  $\Sigma_{HG}$ , which is finite by definition. Concerning instruction 5, it suffices to observe that its termination relies on the termination of  $\text{CERT-EVAL}[\mathcal{C}]$  – which is true by hypothesis – and on the fact that, for each query  $q$ , to construct the set  $\{H_i(\mathbf{t}) \mid \mathbf{t} \in cert(q, D^+, \Sigma_C)\}$ , the problem  $\text{CERT-EVAL}[\mathcal{C}]$  must be solved at most  $d^\mu$  times, where  $d^\mu$  is the maximum number of tuples  $\mathbf{t}$  for which the check  $\mathbf{t} \in cert(q, D^+, \Sigma_C)$  has to be performed. Since each instruction alone terminates, it remains to analyze the overall procedure. It contains two loops. The first, namely the **for**-loop at instruction 3, is not problematic; indeed, we shown that it locally terminates. The second one, namely the **go to**-loop, depends on the evaluation of the **if**-instruction, which can be executed at most  $|gra(D, \Pi)|$  times. Thus, also the **go to**-loop does the same.

**Correctness.** We now claim that Algorithm 2 correctly completes the database. Let  $D^+$  be the output of Algorithm 2. Our claim is that  $D^+ = D \cup gra(D, \Pi)$ .

Inclusion 1 ( $D^+ \supseteq D \cup \text{gra}(D, \Pi)$ ). Assume, by contradiction, that  $D \cup \text{gra}(D, \Pi)$  contains some atom that does not belong to  $D^+$ . This means that there exists some  $j > 0$  such that both  $\bar{D} = ((D \cup \text{gra}(D, \Pi)) \cap \text{chase}^{j-1}(D, \Sigma_{\text{HG}} \cup \Sigma_{\mathcal{C}})) \subseteq D^+$  and  $((D \cup \text{gra}(D, \Pi)) \cap \text{chase}^j(D, \Sigma_{\text{HG}} \cup \Sigma_{\mathcal{C}})) \setminus D^+ \neq \emptyset$  hold. Thus, there exists some  $\underline{a} \in \text{chase}^j(D, \Sigma_{\text{HG}} \cup \Sigma_{\mathcal{C}})$  whose level is exactly  $j$  and that does not belong to  $D^+$ . Let  $\langle \sigma, h \rangle$  be the trigger used by the chase to generate  $\underline{a}$ , where  $\sigma$  is of the form  $\Phi(\mathbf{x}, \mathbf{y}) \rightarrow \mathbf{H}(\mathbf{x})$ . Clearly,  $h$  maps  $\Phi(\mathbf{x}, \mathbf{y})$  to  $\text{chase}^{j-1}(D, \Sigma_{\text{HG}} \cup \Sigma_{\mathcal{C}})$ , and we also have that  $\underline{a} = \mathbf{H}(h(\mathbf{x}))$ . Consider now the query  $q = \langle \mathbf{x} \rangle \leftarrow \Phi(\mathbf{x}, \mathbf{y})$  constructed from  $\sigma$  by Algorithm 2 at instruction 4. Thus,  $\text{chase}^{j-1}(D, \Sigma_{\text{HG}} \cup \Sigma_{\mathcal{C}}) \models q(h(\mathbf{x}))$  holds. Since  $\bar{D} \subseteq D^+$ , we have that  $\text{chase}^{j-1}(D, \Sigma_{\text{HG}} \cup \Sigma_{\mathcal{C}}) \subseteq \text{chase}(\bar{D}, \Sigma_{\mathcal{C}}) \subseteq \text{chase}(D^+, \Sigma_{\mathcal{C}})$ . Hence,  $\text{chase}(D^+, \Sigma_{\mathcal{C}}) \models q(h(\mathbf{x}))$ , namely  $h(\mathbf{x}) \in \text{cert}(q, D^+, \Sigma_{\mathcal{C}})$  and, thus,  $\underline{a} \in D^+$ , which is a contradiction.

Inclusion 2 ( $D^+ \subseteq D \cup \text{gra}(D, \Pi)$ ). Let  $D^+$  be the set produced by Algorithm 2. Let  $\ell$  be the number of time instruction 7 of Algorithm 2 is executed. At each execution  $i \in [\ell]$  of instruction 7, the algorithm computes the set  $\tilde{D}_i$  containing only auxiliary ground atoms, and produces the set  $D_i^+ = D \cup \tilde{D}_i$ . By construction,  $D^+ = D \cup \tilde{D}_\ell$ . Let  $\tilde{D}_0 = \emptyset$ ,  $D_0^+ = D \cup \tilde{D}_0$ , and  $I_i = \tilde{D}_i \setminus \tilde{D}_{i-1}$ , for each  $i \in [\ell]$ . We show that  $D \cup \tilde{D}_\ell \subseteq D \cup \text{gra}(D, \Pi)$ , that is  $\tilde{D}_\ell \subseteq \text{gra}(D, \Pi)$ . We proceed by induction on the number  $\ell$  of iterations.

*Base case:* Let  $i = 1$ . We claim that  $\tilde{D}_1 \subseteq \text{gra}(D, \Pi)$ . By construction, the set  $\tilde{D}_1 = \{\mathbf{H}_i(\mathbf{t}) \mid i \in [|\Sigma_{\text{HG}}|] \wedge \mathbf{t} \in \text{cert}(q, D_0^+, \Sigma_{\mathcal{C}})\}$ . Since  $D_0^+ = D$  and the component  $\Sigma_{\mathcal{C}}$  does not produce any atom in the first iteration of the algorithm, the latter is equal to  $\{\mathbf{H}(\mathbf{t}) \mid i \in [|\Sigma_{\text{HG}}|] \wedge \mathbf{t} \in \text{cert}(q, D, \emptyset)\} = \{\underline{a} \in \text{chase}^1(D, \Sigma_{\text{HG}}) \mid \text{preds}(\underline{a}) \in h\text{-preds}(\Sigma_{\text{HG}})\} \subseteq \{\underline{a} \in \text{chase}(D, \Sigma_{\text{HG}} \cup \Sigma_{\mathcal{C}}) \mid \text{preds}(\underline{a}) \in h\text{-preds}(\Sigma_{\text{HG}})\} = \text{gra}(D, \Pi)$ .

*Induction step:* Given that, for  $i = \ell - 1$ ,  $\tilde{D}_{\ell-1} \subseteq \text{gra}(D, \Pi)$  (*induction hypothesis*), we prove that  $\tilde{D}_\ell \subseteq \text{gra}(D, \Pi)$  holds, too. By construction  $\tilde{D}_\ell = \{\mathbf{H}_i(\mathbf{t}) \mid i \in [|\Sigma_{\text{HG}}|] \wedge \mathbf{t} \in \text{cert}(q, D_{\ell-1}^+, \Sigma_{\mathcal{C}})\} = \{\mathbf{H}_i(\mathbf{t}) \mid i \in [|\Sigma_{\text{HG}}|] \wedge \mathbf{t} \in \text{cert}(q, D \cup \tilde{D}_{\ell-1}, \Sigma_{\mathcal{C}})\} = \{\underline{a} \in \text{chase}(D \cup \tilde{D}_{\ell-1}, \Sigma_{\text{HG}} \cup \Sigma_{\mathcal{C}}) \mid \text{preds}(\underline{a}) \in h\text{-preds}(\Sigma_{\text{HG}})\}$ . Since the set  $D \cup \tilde{D}_{\ell-1}$  already contains all the ground consequences of  $\Sigma_{\text{HG}}$ , the latter is equivalent to  $\{\underline{a} \in \text{chase}(D \cup \tilde{D}_{\ell-1}, \Sigma_{\mathcal{C}}) \mid \text{preds}(\underline{a}) \in h\text{-preds}(\Sigma_{\text{HG}})\}$ . Applying Lemma 1 together with the induction hypothesis, the last one is a subset of  $\{\underline{a} \in \text{chase}(D, \Sigma_{\text{HG}} \cup \Sigma_{\mathcal{C}}) \mid \text{preds}(\underline{a}) \in h\text{-preds}(\Sigma_{\text{HG}})\} = \text{gra}(D, \Pi)$ .  $\square$

We conclude the section by proving the decidability of the problem DP-CERT-EVAL[ $\mathcal{C}$ ], under the assumption that CERT-EVAL[ $\mathcal{C}$ ] is decidable.

*Theorem 2*

If CERT-EVAL[ $\mathcal{C}$ ] is decidable, then DP-CERT-EVAL[ $\mathcal{C}$ ] is decidable.

*Proof*

To prove the decidability of DP-CERT-EVAL[ $\mathcal{C}$ ] we design Algorithm 1. Let  $D$  be a database,  $\Pi = (\Sigma_{\text{HG}}, \Sigma_{\mathcal{C}})$  a dyadic pair,  $q(\mathbf{x})$  a CQ, and  $\mathbf{c}$  a tuple in  $\mathbf{C}^{\mathbf{x}}$ . Clearly, step 1 always terminates since it recalls Algorithm 2 that, as shown in Proposition 4, always terminates and correctly constructs the set  $D^+$ . By Theorem 1, checking if  $\mathbf{c} \in dp\text{-cert}(q, D, \Pi)$  boils down to checking if  $\mathbf{c} \in \text{cert}(q, D^+, \Sigma_{\mathcal{C}})$ , that is decidable by hypothesis. Hence, step 2 never falls in a loop and Algorithm 1 correctly computes DP-CERT-EVAL[ $\mathcal{C}$ ].  $\square$

### 5 Dyadic decomposable sets

In this section we introduce a novel general condition that allows to define, from any decidable class  $\mathcal{C}$  of ontologies, a new decidable class called **Dyadic- $\mathcal{C}$**  enjoying desirable properties. The union of all the **Dyadic- $\mathcal{C}$**  classes, with  $\mathcal{C}$  being any decidable class of TGDs, forms what we call *dyadic decomposable sets*, which encompass and generalize any other existing decidable class, including those based on semantic conditions.

We start the section by providing a classification of atoms of the rule-body, according to where dangerous variables appear; then we define the class **Dyadic- $\mathcal{C}$**  proving that the query answer problem over this class is decidable.

*Definition 6 (Atoms classification)*

Consider a set  $\Sigma$  of TGDs and a rule  $\sigma \in \Sigma$ . An atom  $a$  of  $body(\sigma)$  is  $\sigma$ -*problematic* if (i)  $a$  contains a dangerous variable w.r.t.  $\Sigma$ , or (ii)  $a$  is connected to some  $\sigma$ -*problematic* atom via some harmful variable. The set of all the *problematic* atoms of  $\sigma$  is denoted by  $p\text{-atoms}(\sigma)$ , whereas  $s\text{-atoms}(\sigma) = body(\sigma) \setminus p\text{-atoms}(\sigma)$  denotes the set of all the *safe* atoms of  $\sigma$ . ■

We highlight that  $p\text{-atoms}$  and  $s\text{-atoms}$  can share only harmless variables. The next example is to clarify the above definition.

*Example 3*

Consider the database  $D = \{L(a), R(a, b)\}$ , and the following set  $\Sigma$  of TGDs:

$$\begin{aligned} \sigma_1 : & & L(x_1) & \rightarrow & \exists y_1 P(y_1, x_1) \\ \sigma_2 : & & P(x_2, y_2) & \rightarrow & \exists z_2 Q(y_2, z_2, x_2) \\ \sigma_3 : & & P(z_3, x_3), Q(x_3, y_3, z_3) & \rightarrow & S(z_3) \\ \sigma_4 : & & P(x_4, y_4), Q(y_4, z_4, w_4), S(w_4), R(u_4, v_4) & \rightarrow & T(x_4, z_4, u_4). \end{aligned}$$

We focus on rule  $\sigma_4$ . Taking into account that  $aff(\Sigma) = \{P[1], S[1], Q[2], Q[3]\}$ , it follows that  $dang(\Sigma) = \{x_2, x_4, z_3, z_4\}$ ,  $harmful(\Sigma) = \{x_2, x_4, y_3, z_3, z_4, w_4\}$ , and  $harmless(\Sigma) = \{x_1, x_3, y_2, y_4, u_4, v_4\}$ . Hence,  $p\text{-atoms}(\sigma_4) = \{P(x_4, y_4), Q(y_4, z_4, w_4), S(w_4)\}$ , whereas  $s\text{-atoms}(\sigma_4) = \{R(u_4, v_4)\}$ . ■

The second step consists in selecting variables shared by  $p\text{-atoms}$  and  $s\text{-atoms}$  together with the harmless frontier variables appearing in  $s\text{-atoms}$ . The latter can be expressed via the set  $hf(\sigma) = \{x \in vars(s\text{-atoms}(\sigma)) \mid x \in harmless(\sigma) \wedge x \in vars_{\sim}(\sigma)\}$ . Then, we define  $vars_{\star}(\sigma) = \{vars(p\text{-atoms}(\sigma)) \cap vars(s\text{-atoms}(\sigma))\} \cup hf(\sigma)$ . Trivially, it holds that  $vars_{\star}(\sigma) \subseteq vars(s\text{-atoms}(\sigma))$ . Finally, let

$$\begin{aligned} hg(\sigma) : s\text{-atoms}(\sigma) & \rightarrow \mathbf{Aux}_{\sigma}(vars_{\star}(\sigma)), \\ main(\sigma) : \mathbf{Aux}_{\sigma}(vars_{\star}(\sigma)), p\text{-atoms}(\sigma) & \rightarrow \exists vars_{\exists}(\sigma) head(\sigma), \\ hg(\Sigma) = \bigcup_{\sigma \in \Sigma} hg(\sigma) & \quad \text{and} \quad main(\Sigma) = \bigcup_{\sigma \in \Sigma} main(\sigma). \end{aligned}$$

By considering again Example 3, we have that  $vars_{\star}(\sigma_4) = \{u_4, y_4\}$  and also that  $main(\sigma_4) : \mathbf{Aux}_{\sigma_4}(u_4, y_4), P(x_4, y_4), Q(y_4, z_4, w_4), S(w_4, y_4) \rightarrow T(x_4, z_4, u_4)$ .

In the special case in which a variable  $x$  of  $vars_{\star}(\sigma)$  occurs  $n > 1$  times in  $head(\sigma)$ , then  $x$  also occurs  $n$  times in the head of  $hg(\sigma)$ . Accordingly,  $x$  occurs with different



names (e.g.,  $x_1, \dots, x_n$ ) both in the head and in the body of  $main(\sigma)$ . For example, if the ontology contains only the rule  $\sigma : P(x) \rightarrow R(x, x)$ , then  $hg(\sigma) : P(x) \rightarrow \mathbf{Aux}_\sigma(x, x)$  and  $main(\sigma) : \mathbf{Aux}_\sigma(x_1, x_2) \rightarrow R(x_1, x_2)$ . Clearly, the latter two rules together are equivalent to  $\sigma$  w.r.t. to the schema  $\{P, R\}$ . We prefer to keep the formal definition of  $hg(\Sigma)$  and  $main(\Sigma)$  light without formalizing such special cases.

We are now ready to formally introduce the class  $\mathbf{Dyadic-C}$ .

*Definition 7 (Dyadic-C)*

Consider a class  $\mathcal{C}$  of TGDs such that  $\mathbf{CERT-EVAL}[\mathcal{C}]$  is decidable. We say that  $\Sigma$  belongs to  $\mathbf{Dyadic-C}$  if  $\Sigma$  belongs to  $\mathcal{C}$  or if  $main(\Sigma)$  belongs to  $\mathcal{C}$ . ■

According to the previous definition, one can easily state the following property.

*Proposition 5*

Consider a class  $\mathcal{C}$  of TGDs. It holds that  $\mathcal{C} \subseteq \mathbf{Dyadic-C}$ .

By Definition 7, to check if an ontology  $\Sigma$  belongs to  $\mathbf{Dyadic-C}$ , one has to verify if  $\Sigma \in \mathcal{C}$ , or  $main(\Sigma) \in \mathcal{C}$ . We observe that the construction of the set  $main(\Sigma)$  explained above, is polynomial (indeed linear) with respect to the size of  $\Sigma$ . Hence, the following result holds.

*Theorem 3*

Consider a class  $\mathcal{C}$  of TGDs and assume that checking whether an ontology belongs to  $\mathcal{C}$  is doable in some complexity class  $\mathbb{C} \supseteq \mathbf{PTIME}$ . Then, checking whether an ontology belongs to  $\mathbf{Dyadic-C}$  is decidable and it belongs to  $\mathbb{C}$ .

*Proof*

We start recalling that, by Definition 7, an ontology  $\Sigma$  belongs to  $\mathbf{Dyadic-C}$  if (i)  $\Sigma \in \mathcal{C}$ , or (ii)  $main(\Sigma) \in \mathcal{C}$ . Accordingly, checking condition (i) is doable in some complexity class  $\mathbb{C} \supseteq \mathbf{PTIME}$ , by assumption; otherwise, the construction of the set  $main(\Sigma)$  is done by a procedure that is polynomial (indeed linear) with respect to the size of  $\Sigma$  and, hence, always terminates. Accordingly, checking condition (ii) is also decidable and doable in the complexity class  $\mathbb{C}$ . □

The next step is to prove the existence of a  $\mathcal{C}$ -dyadic pair for any ontology  $\Sigma \in \mathbf{Dyadic-C}$ .

*Theorem 4*

Consider a set  $\Sigma \in \mathbf{Dyadic-C}$ . There exists a  $\mathcal{C}$ -dyadic pair  $\Pi = (\Sigma_{\mathbf{HG}}, \Sigma_{\mathcal{C}})$  of TGDs such that  $\Sigma_{\mathbf{HG}} \cup \Sigma_{\mathcal{C}} \equiv_{sch(\Sigma)} \Sigma$ . In particular,

1. If  $\Sigma \in \mathcal{C}$ , then  $\Sigma_{\mathbf{HG}} = \emptyset$  and  $\Sigma_{\mathcal{C}} = \Sigma$ ;
2. If  $\Sigma \notin \mathcal{C}$ , then  $\Sigma_{\mathbf{HG}} = hg(\Sigma)$  and  $\Sigma_{\mathcal{C}} = main(\Sigma)$ .

*Proof*

We claim that for each ontology  $\Sigma \in \mathbf{Dyadic-C}$  it is possible to construct a  $\mathcal{C}$ -dyadic pair  $\Pi = (\Sigma_{\mathbf{HG}}, \Sigma_{\mathcal{C}})$  of TGDs such that  $\Sigma_{\mathbf{HG}} \cup \Sigma_{\mathcal{C}} \equiv_{sch(\Sigma)} \Sigma$ .

According to Definition 5, we recall that a pair  $\Pi = (\Sigma_{\mathbf{HG}}, \Sigma_{\mathcal{C}})$  is  $\mathcal{C}$ -dyadic if (i)  $\Sigma_{\mathbf{HG}}$  is head-ground with respect to  $\Sigma_{\mathbf{HG}} \cup \Sigma_{\mathcal{C}}$  and (ii)  $\Sigma_{\mathcal{C}} \in \mathcal{C}$ . Assume first that  $\Sigma \in \mathcal{C}$ . Then, trivially the pair  $(\emptyset, \Sigma)$  is a  $\mathcal{C}$ -dyadic pair. Assume now that  $\Sigma \notin \mathcal{C}$  and let  $\Pi = (hg(\Sigma), main(\Sigma))$ . Property (ii) is satisfied since by hypothesis  $\Sigma \in \mathbf{Dyadic-C}$ ; hence, it follows by definition that  $main(\Sigma) \in \mathcal{C}$ . It remains to show property (i). According to Definition 4, the set  $hg(\Sigma)$  has to satisfy four properties. Property 1 and 2 are trivially

**Algorithm 3:** CertEval<sub>[Dyadic-C]</sub>( $q, D, \Sigma, \mathbf{c}$ )

---

**Input:** A database  $D$ , a  $\Sigma \in \text{Dyadic-C}$ , a CQ  $q(\mathbf{x})$ , and a tuple  $\mathbf{c} \in C^{|\mathbf{x}|}$

- 1  $\Sigma_{\text{HG}} = hg(\Sigma)$ ;
- 2  $\Sigma_{\text{C}} = main(\Sigma)$ ;
- 3  $\Pi = (\Sigma_{\text{HG}}, \Sigma_{\text{C}})$ ;
- 4 **return** DpCertEval<sub>[C]</sub>( $q, D, \Pi, \mathbf{c}$ );

---

fulfilled since, by construction, for each  $\sigma \in \Sigma$ ,  $hg(\sigma)$  is a datalog rule and each head atom contains only harmless variables with respect to  $hg(\Sigma) \cup main(\Sigma)$ . Property 3 and 4 state that  $h\text{-preds}(hg(\Sigma)) \cap b\text{-preds}(hg(\Sigma)) = \emptyset$  and  $h\text{-preds}(hg(\Sigma)) \cap h\text{-preds}(main(\Sigma)) = \emptyset$ . These hold since, by construction,  $h\text{-preds}(hg(\Sigma)) = \{\text{Aux}_{\sigma} : \sigma \in \Sigma\}$ , where each  $\text{Aux}_{\sigma}$  is a predicate that does not occur neither in any body of  $hg(\Sigma)$  nor in any head of  $main(\Sigma)$ .

Concerning the equivalence between  $\Sigma_{\text{HG}} \cup \Sigma_{\text{C}}$  and  $\Sigma$ , we can observe that it easily comes from the shape of  $hg(\sigma)$  and  $main(\sigma)$  with respect to each original rule  $\sigma \in \Sigma$ . Indeed, first, the body of  $\sigma$  is first partitioned in  $s\text{-atoms}(\sigma)$  and  $p\text{-atoms}(\sigma)$ . Second, all the atoms if  $s\text{-atoms}(\sigma)$  form the body of  $hg(\sigma)$ . Then, all the variables of  $hg(\sigma)$  that are in join with  $p\text{-atoms}(\sigma)$  or are in the head of  $\sigma$  are collected in  $\text{Aux}_{\sigma}(vars_{*}(\sigma))$ . Finally,  $\text{Aux}_{\sigma}(vars_{*}(\sigma))$  is put in conjunction with  $p\text{-atoms}(\sigma)$  to form the body of  $main(\sigma)$ . Such a way of decomposing a rule  $\sigma$  is well-known to be correct for query answering purposes even when the variables in the auxiliary atom are harmful.  $\square$

It remains to show that Dyadic-C is decidable. We rely on Algorithm 3 together with Theorem 1 and Proposition 4 to state the following result.

*Theorem 5*

Consider a decidable class  $\mathcal{C}$  of TGDs. Then, CERT-EVAL[Dyadic-C] is decidable.

*Proof*

To prove the statement we provide the terminating Algorithm 3. Let  $\Sigma \in \text{Dyadic-C}$  an ontology. Instructions 1 and 2 of the algorithm are introduced in order to construct the components  $(hg(\Sigma), main(\Sigma))$  of a dyadic pair  $\Pi$ , which is successively initialized at instruction 3. Of course, the construction of  $\Pi$  is based on a polynomial procedure with respect to the size of the input  $\Sigma$ , hence these instructions always terminates. Finally, instruction 4 returns the result of the evaluation of the problem DP-CERT-EVAL[C]. To solve the latter, is invoked Algorithm 1, which in turn invokes Algorithm 2; their correctness is guaranteed by Theorem 1 and Proposition 4, respectively. Accordingly, CERT-EVAL[Dyadic-C] is decidable.  $\square$

Finally, we conclude the section by proving that, for any class  $\mathcal{C}$  including the class Af-Inds defined in Section 3.3, the class Dyadic-C includes Datalog.

*Theorem 6*

Consider a class  $\mathcal{C}$  of TGDs. If  $\mathcal{C} \supseteq \text{Af-Inds}$ , then Datalog  $\subseteq$  Dyadic-C.

*Proof*

Consider a datalog rule  $\sigma : \Phi(\mathbf{x}, \mathbf{y}) \rightarrow P(\mathbf{x})$ . Then, it can be decomposed into the following two rules:  $hg(\sigma) : \Phi(\mathbf{x}, \mathbf{y}) \rightarrow \text{Aux}_{\sigma}(\mathbf{x})$  and  $main(\sigma) : \text{Aux}_{\sigma}(\mathbf{x}) \rightarrow P(\mathbf{x})$ . Trivially,

$hg(\sigma) \in \Sigma_{HG}$ ; it remains to show that  $main(\sigma) \in \text{Af-Inds}$ . According to Definition 3, the full property immediately follows since datalog rules do not have existential variables, whereas the autonomous property holds since  $body(main(\sigma))$  contains only the fresh predicate  $H$ . In particular,  $main(\sigma) \in \mathcal{C}$ , since by assumption  $\mathcal{C} \supseteq \text{Af-Inds}$ ; hence, the thesis follows.  $\square$

### 6 Computational complexity of query answering

In this section we study the complexity of the CERT-EVAL problem over dyadic existential rules. We start analyzing the data complexity of the problem and then the combined complexity.

*Theorem 7*

Consider a class  $\mathcal{C}$  of TGDs. In data complexity, if  $\text{CERT-EVAL}[\mathcal{C}]$  belongs to some decidable complexity class  $\mathbb{C}$ , then the following hold:

1. If  $\mathbb{C} \subseteq \text{PTIME}$ , then  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}]$  is in **PTIME**;
2. If  $\mathbb{C} \supseteq \text{PTIME}$ , then  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}]$  is in  $\text{PTIME}^{\mathbb{C}}$ ;
3. If  $\mathbb{C} \supseteq \text{PTIME}$  is deterministic<sup>2</sup> and  $\text{CERT-EVAL}[\mathcal{C}]$  is  $\mathbb{C}$ -complete, then it holds that  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}]$  is  $\mathbb{C}$ -complete too;
4. If  $\mathbb{C} \supseteq \text{Af-Inds}$ , then  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}]$  is **PTIME**-hard.

*Proof*

To prove the theorem, we rely on the complexity of Algorithm 3. Let  $D$  be a database,  $\Sigma \in \text{Dyadic-}\mathcal{C}$  an ontology,  $q(\mathbf{x})$  a CQ, and  $\mathbf{c} \in \mathbb{C}^{|\mathbf{x}|}$  a tuple. Moreover, let  $d = |\text{consts}(D)|$  and  $\mu = \max_{P \in h\text{-preds}(\Sigma_{HG})} \text{arity}(P)$ . As shown in Theorem 5, Algorithm 3 terminates and correctly decides  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}]$ . More specifically, instructions 1 and 2 are introduced in order to construct, respectively, the first component  $hg(\Sigma)$  and the second component  $main(\Sigma)$  of a dyadic pair. The procedure used to build these sets is polynomial (indeed linear) with respect to the size of the ontology  $\Sigma$  in input. By neglecting the “trivial” instruction 3, the computational cost of the Algorithm mainly depends on instruction 4, that is on the invocation of Algorithm 1, which in turn invokes Algorithm 2 in order to compute the completed database  $D^+$ . As stated in Proposition 4, Algorithm 2 always terminates. In particular, the size of the set  $D^+$  is at most  $|h\text{-preds}(\Sigma_{HG})| \cdot d^\mu$ ; the problem  $\text{CERT-EVAL}[\mathcal{C}]$ , at instruction 5, is called  $d^\mu$  times, and the **for**-loop at instruction 3 is executed  $|\Sigma_{HG}|$  times. Therefore, by ignoring the computational costs of the oracle (i.e., checking whether  $\mathbf{t} \in \text{cert}(q, D^+, \Sigma_{\mathcal{C}})$ ), Algorithm 2 overall performs a number of steps that is linear in  $|\Sigma_{HG}| \cdot |h\text{-preds}(\Sigma_{HG})| \cdot d^{2\mu}$ . Indeed, this value is also an upper bound for the number of calls to the oracle. Since we are in data complexity, the following parameters are bounded: the maximum arity  $\mu$ , the size of the sets  $\Sigma_{HG}$  and  $\Sigma_{\mathcal{C}}$ , as well as the size and the number of each query  $q$  constructed at instruction 4 of Algorithm 2. Hence, the latter calls polynomially many times the problem  $\text{CERT-EVAL}[\mathcal{C}]$ . Accordingly, Algorithm 3 is polynomial and in turn it invokes polynomially many times an oracle to compute  $\text{CERT-EVAL}[\mathcal{C}]$ . Hence, if  $\mathbb{C} \subseteq \text{PTIME}$ , trivially  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}] \in \text{PTIME}$ ;

<sup>2</sup> By “deterministic” we mean that  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}]$  can be solved by a deterministic Turing machine. We refer to the classes **PTIME**, **PSPACE**, **i-EXPTIME**, **i-EXPSPACE**, for all  $i \geq 1$ .

Table 2. *Data complexity comparison of CERT-EVAL[C] with CERT-EVAL[Dyadic-C]*

Class $\mathcal{C}$	$\mathcal{C}$	Dyadic- $\mathcal{C}$
Weakly-(Fr)-Guarded	<b>EXPTIME</b> -complete	<b>EXPTIME</b> -complete
(Fr)-Guarded	<b>PTIME</b> -complete	<b>PTIME</b> -complete
Weakly-Acyclic	<b>PTIME</b> -complete	<b>PTIME</b> -complete
Jointly-Acyclic	<b>PTIME</b> -complete	<b>PTIME</b> -complete
Datalog	<b>PTIME</b> -complete	<b>PTIME</b> -complete
Shy	<b>PTIME</b> -complete	<b>PTIME</b> -complete
Ward	<b>PTIME</b> -complete	<b>PTIME</b> -complete
Protected	<b>PTIME</b> -complete	<b>PTIME</b> -complete
Sticky-(Join)	<b>AC</b> <sub>0</sub>	<b>PTIME</b> -complete
Linear	<b>AC</b> <sub>0</sub>	<b>PTIME</b> -complete
Joinless	<b>AC</b> <sub>0</sub>	<b>PTIME</b> -complete
Inclusion-Dependencies	<b>AC</b> <sub>0</sub>	<b>PTIME</b> -complete

whereas, if  $\mathbb{C} \supseteq \mathbf{PTIME}$ ,  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}] \in \mathbf{PTIME}^{\mathbb{C}}$ . To prove point 3 of the theorem, we observe that the membership follows from point 2 and from the fact that, for any deterministic class  $\mathbb{C} \supseteq \mathbf{PTIME}$ , it holds that  $\mathbf{PTIME}^{\mathbb{C}} = \mathbb{C}$ ; whereas, the hardness derives from Proposition 5, since Dyadic- $\mathcal{C}$  includes the class  $\mathcal{C}$ , that is  $\mathbb{C}$ -hard by assumption. Finally, to prove point 4, we recall that by Theorem 6,  $\text{Datalog} \subseteq \text{Dyadic-}\mathcal{C}$ ; hence, since  $\text{CERT-EVAL}[\text{Datalog}]$  is **PTIME**-hard, the thesis follows.  $\square$

Accordingly to the above theorem, immediately we get the following result.

*Corollary 1*

Complexity results in Table 2 do hold.

For studying the combined complexity, we need to take into account the fact that the database returned by Algorithm 2 (namely,  $D^+$ ) is exponential with respect to the input one (namely,  $D$ ). Indeed, the check  $\mathbf{c} \in \text{cert}(q, D^+, \Sigma_{\mathcal{C}})$  performed by Algorithm 1 is done on an exponentially bigger database. Thus, in case  $\text{CERT-EVAL}[\mathcal{C}]$  would have the same data complexity and combined complexity, it might happen that the combined complexity of  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}]$  could be exponentially higher than the one of  $\text{CERT-EVAL}[\mathcal{C}]$ . Although all the considered classes in  $\mathbb{E}_{syn}^+$  do not suffer from this shortcoming, before stating our general result, we need to focus on “well-behaved” classes of TGDs. A class  $\mathcal{C}$  of TGDs enjoys the *dropping data-complexity* property if there is an exponential jump from the combined complexity of  $\text{CERT-EVAL}[\mathcal{C}]$  to the data complexity of  $\text{CERT-EVAL}[\mathcal{C}]$ .

*Proposition 6*

Each class in  $\mathbb{E}_{syn}$  enjoys the dropping data-complexity property.

We can now state the last result of the section, providing the combined complexity of problem  $\text{CERT-EVAL}$  over Dyadic- $\mathcal{C}$  sets of TGDs.

*Theorem 8*

Consider a class  $\mathcal{C}$  of TGDs. In combined complexity, if  $\text{CERT-EVAL}[\mathcal{C}]$  belongs to some decidable complexity class  $\mathbb{C}$  and  $\mathcal{C}$  enjoys the dropping data-complexity property, then the following hold:

1. If  $\mathbb{C} \subseteq \mathbf{EXPTIME}$ , then  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}]$  is in  $\mathbf{EXPTIME}$ ;
2. If  $\mathbb{C} \supseteq \mathbf{EXPTIME}$ , then  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}]$  is in  $\mathbf{EXPTIME}^{\mathbb{C}}$ ;
3. If  $\mathbb{C} \supseteq \mathbf{EXPTIME}$  is deterministic and  $\text{CERT-EVAL}[\mathcal{C}]$  is  $\mathbb{C}$ -complete, then it holds that  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}]$  is  $\mathbb{C}$ -complete too;
4. If  $\mathcal{C} \supseteq \text{Af-Ilds}$ , then  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}]$  is  $\mathbf{EXPTIME}$ -hard.

*Proof*

The argument proceeds similarly to proof of Theorem 7 by arguing on Algorithm 3 to determine the complexity of  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}]$ . Let  $D$  be a database,  $\Sigma \in \text{Dyadic-}\mathcal{C}$  an ontology,  $q(\mathbf{x})$  a CQ, and  $\mathbf{c} \in \mathbb{C}^{|\mathbf{x}|}$  a tuple. Moreover, let  $d = |\text{consts}(D)|$  and  $\mu = \max_{P \in h\text{-preds}(\Sigma_{\text{HG}})} \text{arity}(P)$ . As previously shown, Algorithm 3 invokes Algorithm 1, which in turn invokes Algorithm 2. Concerning the latter, by ignoring the computational costs of the oracle, it overall performs a number of step that is linear in  $|\Sigma_{\text{HG}}| \cdot |h\text{-preds}(\Sigma_{\text{HG}})| \cdot d^{2\mu}$ . Indeed, also in this case, this value is an upper bound for the number of calls to the oracle. This is enough to show point 2.

Concerning the memberships of point 1 and point 3, differently from the proof of Theorem 7, in combined complexity the maximum arity  $\mu$ , the size of the sets  $\Sigma_{\text{HG}}$  and  $\Sigma_{\mathcal{C}}$ , as well as the size and the number of each query constructed at instruction 4 of Algorithm 2 are not bounded. Accordingly, also the size of the completed database returned by Algorithm 2 (namely  $D^+$ ) may become exponential with respect to the input. More precisely,  $|\text{consts}(D^+)| = d$  and  $|D^+| \leq |h\text{-preds}(\Sigma_{\text{HG}})| \cdot d^\mu + |D|$ . Let  $n$  generically denote the size  $\|seq\|$  of any sequence  $seq$  of objects given in input to  $\text{CERT-EVAL}[\mathcal{C}]$ . We can now consider the cost function  $g(n)$  (resp.,  $f(n)$ ) of some algorithm/oracle that decides  $\text{CERT-EVAL}[\mathcal{C}]$  and shows that it belongs to  $\mathbb{C}$  (resp.,  $\mathbb{C}_d$ ) in combined (resp., data) complexity. According to the dropping data-complexity property, we know that  $g(n)$  grows at least exponentially faster than  $f(n)$ . Essentially, there is an exponential jump from  $\mathbb{C}_d$  to  $\mathbb{C}$  that does not depend on the size of the input database but only on the size of other parameters, namely the ontology, the query, and the tuple of constants. Consider now the query  $q' := \langle \mathbf{x} \rangle \leftarrow \Phi(\mathbf{x}, \mathbf{y})$  constructed at instruction 4 of Algorithm 2 (we call it  $q'$  to avoid confusion with  $q(\mathbf{x})$  mentioned at the beginning of this proof). At instruction 5 of the same algorithm, the oracle for  $\text{CERT-EVAL}[\mathcal{C}]$  checks whether  $\mathbf{t} \in \text{cert}(q', D^+, \Sigma_{\mathcal{C}})$  holds. Since  $g(n)$  grows at least exponentially faster than  $f(n)$ , we get that  $g(\|\mathbf{t}, q', D^+, \Sigma_{\mathcal{C}}\|)$  remains of the same exponential order of  $g(\|\mathbf{t}, q', D, \Sigma_{\mathcal{C}}\|)$ , although  $\|D^+\|$  may be exponentially larger than  $\|D\|$ .

Regarding point 1, if  $\mathbb{C} \subseteq \mathbf{EXPTIME}$ , then we know that a  $\mathbb{C}$ -oracle for  $\text{CERT-EVAL}[\mathcal{C}]$  works at most exponentially in  $\|\mathbf{t}, q', \Sigma_{\mathcal{C}}\|$  and at most polynomially (resp., exponentially) in  $\|D^+\|$  (resp.,  $\|D\|$ ); thus, in this case, the  $\mathbb{C}$ -oracle cannot reach double-exponential time but it remains exponential. Therefore,  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}] \in \mathbf{EXPTIME}$ .

For the memberships of point 3, we already know that  $\text{CERT-EVAL}[\text{Dyadic-}\mathcal{C}]$  is in  $\mathbf{EXPTIME}^{\mathbb{C}}$ . Consider now a  $\mathbb{C}$ -oracle  $O$  for  $\text{CERT-EVAL}[\mathcal{C}]$  characterized by the cost function  $g(n)$ . If  $\mathbb{C} \supseteq \mathbf{EXPTIME}$  is deterministic, then  $O$  works with respect to  $\|\mathbf{t}, q', \Sigma_{\mathcal{C}}\|$  in an exponentially faster way than with respect to  $\|D^+\|$ ; thus, also in this case,  $O$  cannot exceed the power of  $\mathbb{C}$ . Therefore,  $\mathbf{EXPTIME}^{\mathbb{C}}$ , in a sense, collapses to  $\mathbb{C}$ .

Finally, we conclude the proof by considering the hardness of points 3 and 4. In the first case, we observe that it derives from Proposition 5, since  $\text{Dyadic-}\mathcal{C}$  includes the

Table 3. Combined complexity comparison of CERT-EVAL[C] with CERT-EVAL[Dyadic-C]

Class $\mathcal{C}$	$\mathcal{C}$	Dyadic- $\mathcal{C}$
Weakly-(Fr)-Guarded	<b>2EXPTIME</b> -complete	<b>2EXPTIME</b> -complete
(Fr)-Guarded	<b>2EXPTIME</b> -complete	<b>2EXPTIME</b> -complete
Weakly-Acyclic	<b>2EXPTIME</b> -complete	<b>2EXPTIME</b> -complete
Jointly-Acyclic	<b>2EXPTIME</b> -complete	<b>2EXPTIME</b> -complete
Datalog	<b>EXPTIME</b> -complete	<b>EXPTIME</b> -complete
Shy	<b>EXPTIME</b> -complete	<b>EXPTIME</b> -complete
Ward	<b>EXPTIME</b> -complete	<b>EXPTIME</b> -complete
Protected	<b>EXPTIME</b> -complete	<b>EXPTIME</b> -complete
Sticky-(Join)	<b>EXPTIME</b> -complete	<b>EXPTIME</b> -complete
Linear	<b>PSPACE</b> -complete	<b>EXPTIME</b> -complete
Joinless	<b>PSPACE</b> -complete	<b>EXPTIME</b> -complete
Inclusion-Dependencies	<b>PSPACE</b> -complete	<b>EXPTIME</b> -complete

class  $\mathcal{C}$ , that is  $\mathbb{C}$ -hard by assumption. For point 4, we recall that by Theorem 6,  $\text{Datalog} \subseteq \text{Dyadic-}\mathcal{C}$ ; hence, since  $\text{CERT-EVAL}[\text{Datalog}]$  is **EXPTIME**-hard, it follows the thesis.  $\square$

The following immediately derives from above theorem.

#### Corollary 2

Complexity results in Table 3 do hold.

## 7 Conclusion

Dyadic decomposable sets form a novel decidable class of TGDs that encompasses and generalizes all the existing (syntactic and semantic) decidable classes of TGDs. In the near future, it would be interesting to implement a prototype for dyadic existential rules by exploiting different kinds of existing reasoners.

## Acknowledgments

Georg Gottlob is a Royal Society Research Professor and acknowledges support by the Royal Society under this scheme (Project “RAISON DATA” RP/R1/201074). Marco Manna was partially supported by the PNR projects “FAIR (PE00000013) - Spoke 9” and “Tech4You (ECS00000009) - Spoke 6”, under the NRRP MUR program funded by the NextGenerationEU. Cinzia Marte was partially supported by the PON project “Modelli, Sistemi e Competenze per l’Implementazione dell’Ufficio per il Processo - STARTUPP (H29J22000390006)”.

## References

- ABITEBOUL, S., HULL, R. AND VIANU, V. 1995. *Foundations of Databases*, vol. 8. Addison-Wesley Reading.
- BAADER, F., CALVANESE, D., MCGUINNESS, D. L., NARDI, D. AND PATEL-SCHNEIDER, P. F. Eds. 2003. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press.

- BAGET, J., LECLÈRE, M., MUGNIER, M. AND SALVAT, E. 2009. Extending decidable cases for rules with existential variables. In *IJCAI*, 677–682.
- BAGET, J., LECLÈRE, M., MUGNIER, M. AND SALVAT, E. 2011. On rules with existential variables: Walking the decidability line. *Artificial Intelligence* 175, 9–10, 1620–1654.
- BALDAZZI, T., BELLOMARINI, L., FAVORITO, M. AND SALLINGER, E. 2022. On the relationship between shy and warded datalog+/- . CoRR, abs/2202.06285.
- BERGER, G., GOTTLÖB, G., PIERIS, A. AND SALLINGER, E. 2022. *The Space-Efficient Core of Vatalog 2022*, vol. 47, 1:1–1:46.
- CALÌ, A., GOTTLÖB, G. AND KIFER, M. 2008. Taming the infinite chase: Query answering under expressive relational constraints. In *KR*. AAAI Press, 70–80.
- CALÌ, A., GOTTLÖB, G. AND KIFER, M. 2013. Taming the infinite chase: Query answering under expressive relational constraints. *Journal of Artificial Intelligence Research* 48, 115–174.
- CALÌ, A., GOTTLÖB, G. AND LUKASIEWICZ, T. 2012a. A general datalog-based framework for tractable query answering over ontologies. *Journal of Web Semantics* 14a, 57–83.
- CALÌ, A., GOTTLÖB, G. AND PIERIS, A. 2010. Advanced processing for ontological queries. *Proceedings of the VLDB Endowment* 3, 1–2, 554–565.
- CALÌ, A., GOTTLÖB, G. AND PIERIS, A. 2012b. Towards more expressive ontology languages: The query answering problem. *Artificial Intelligence* 193b, 87–128.
- CERI, S., GOTTLÖB, G. AND TANCA, L. 1989. What you always wanted to know about datalog (and never dared to ask). *IEEE Transactions on Knowledge and Data Engineering* 1, 1, 146–166.
- DEUTSCH, A., NASH, A. AND REMMEL, J. B. The chase revisited. In *PODS 2008*. ACM, 149–158.
- FAGIN, R., KOLAITIS, P. G., MILLER, R. J. AND POPA, L. 2005. Data exchange: Semantics and query answering. *Theoretical Computer Science* 336, 1, 89–124.
- GOGACZ, T. AND MARCINKOWSKI, J. *Converging to the Chase - A Tool for Finite Controllability*, vol. 83, 180–206.
- GOTTLÖB, G., MANNA, M. AND MARTE, C. Dyadic existential rules. In *Datalog*, CEUR Workshop Proceedings, vol. 3203. CEUR-WS.org, 83–96.
- GOTTLÖB, G. AND PIERIS, A. Beyond SPARQL under OWL 2 QL entailment regime: Rules to the rescue. In *IJCAI*. AAAI Press, 2999–3007.
- JOHNSON, D. S. AND KLUG, A. C. 1984. Testing containment of conjunctive queries under functional and inclusion dependencies. *Journal of Computer and System Sciences* 28, 1, 167–189.
- KRÖTZSCH, M. AND RUDOLPH, S. Extending decidable existential rules by joining acyclicity and guardedness. In *IJCAI*. IJCAI/AAAI, 963–968.
- LEONE, N., MANNA, M., TERRACINA, G. AND VELTRI, P. 2019. Fast query answering over existential rules. *ACM Transactions on Computational Logic* 20, 2, 12:1–12:48.