

Acyclicity Notions for Existential Rules and Their Application to Query Answering in Ontologies

Bernardo Cuenca Grau

BERNARDO.CUENCA.GRAU@CS.OX.AC.UK

Ian Horrocks

IAN.HORROCKS@CS.OX.AC.UK

Markus Krötzsch

MARKUS.KROETZSCH@CS.OX.AC.UK

Clemens Kupke

CLEMENS.KUPKE@CS.OX.AC.UK

Despoina Magka

DESPOINA.MAGKA@CS.OX.AC.UK

Boris Motik

BORIS.MOTIK@CS.OX.AC.UK

Zhe Wang

ZHE.WANG@CS.OX.AC.UK

*Department of Computer Science, University of Oxford
Parks Road, Oxford OX1 3QD, United Kingdom*

Abstract

Answering conjunctive queries (CQs) over a set of facts extended with existential rules is a prominent problem in knowledge representation and databases. This problem can be solved using the *chase* algorithm, which extends the given set of facts with fresh facts in order to satisfy the rules. If the chase terminates, then CQs can be evaluated directly in the resulting set of facts. The chase, however, does not terminate necessarily, and checking whether the chase terminates on a given set of rules and facts is undecidable. Numerous *acyclicity notions* were proposed as sufficient conditions for chase termination. In this paper, we present two new acyclicity notions called *model-faithful acyclicity* (MFA) and *model-summarising acyclicity* (MSA). Furthermore, we investigate the landscape of the known acyclicity notions and establish a complete taxonomy of all notions known to us. Finally, we show that MFA and MSA generalise most of these notions.

Existential rules are closely related to the Horn fragments of the OWL 2 ontology language; furthermore, several prominent OWL 2 reasoners implement CQ answering by using the chase to *materialise* all relevant facts. In order to avoid termination problems, many of these systems handle only the OWL 2 RL profile of OWL 2; furthermore, some systems go beyond OWL 2 RL, but without any termination guarantees. In this paper we also investigate whether various acyclicity notions can provide a principled and practical solution to these problems. On the theoretical side, we show that query answering for acyclic ontologies is of lower complexity than for general ontologies. On the practical side, we show that many of the commonly used OWL 2 ontologies are MSA, and that the number of facts obtained by materialisation is not too large. Our results thus suggest that principled development of materialisation-based OWL 2 reasoners is practically feasible.

1. Introduction

Existential rules are first-order implications between conjunctions of function-free atoms that may contain existentially quantified variables in the implication's consequent (Baget, Leclère, Mugnier, & Salvat, 2011a; Cali, Gottlob, Lukasiewicz, Marnette, & Pieris, 2010a). Such rules are used in a variety of ways in databases, knowledge representation, and logic programming. In database theory, existential rules are known as *tuple-generating dependencies* (Abiteboul, Hull, & Vianu, 1995) and are used to capture a wide range of schema

constraints. Furthermore, they are also used as declarative data transformation rules in *data exchange*—the process of transforming a database structured according to a source schema into a database structured according to a target schema (Fagin, Kolaitis, Miller, & Popa, 2005). Existential rules also provide the foundation for several prominent knowledge representation formalisms, such as Datalog[±] (Calì, Gottlob, & Pieris, 2010b; Calì et al., 2010a), and they are also closely related to logic programs with function symbols in the head. Practical applications of existential rules range from bioinformatics (Mungall, 2009) to modelling complex structures of chemical compounds (Magka, Motik, & Horrocks, 2012; Hastings, Magka, Batchelor, Duan, Stevens, Ennis, & Steinbeck, 2012).

Answering *conjunctive queries* (CQs) over a set of facts extended with existential rules is a fundamental, yet undecidable (Beeri & Vardi, 1981) reasoning problem for existential rules. The problem can be characterised using *chase* (Johnson & Klug, 1984; Maier, Mendelzon, & Sagiv, 1979)—a technique closely related to the hypertableau calculus (Motik, Shearer, & Horrocks, 2009b; Baumgartner, Furbach, & Niemelä, 1996). In a forward-chaining manner, the chase extends the original set of facts with facts that can be derived using the rules. The result of the chase is a *universal* model, in the sense that an arbitrary CQ over the original facts and rules can be answered by evaluating the query in this model.

1.1 Chase Termination and Acyclicity Notions

Rules with existentially quantified variables in the head—so-called *generating rules*—require the introduction of fresh individuals. Cyclic applications of generating rules may prevent the chase from terminating, and in fact determining whether chase terminates on a set of rules and facts is undecidable (Deutsch, Nash, & Rummel, 2008). However, several decidable classes of existential rules have been identified, and the existing proposals can be classified into two main groups. In the first group, rules are restricted such that their possibly infinite universal models can be represented using finitary means. This group includes rules with universal models of bounded treewidth (Baget et al., 2011a), guarded rules (Calì et al., 2010a), and ‘sticky’ rules (Calì, Gottlob, & Pieris, 2011). In the second group, one uses a sufficient (but not necessary) *acyclicity* notion that ensures chase termination.

Roughly speaking, acyclicity notions analyse the information flow between rules to ensure that no cyclic applications of generating rules are possible. *Weak acyclicity* (WA) (Fagin et al., 2005) was one of the first such notions, and it was extended to notions such as *safety* (Meier, Schmidt, & Lausen, 2009), *stratification* (Deutsch et al., 2008), *acyclicity of a graph of rule dependencies* (aGRD) (Baget, Mugnier, & Thomazo, 2011b), *joint acyclicity* (JA) (Krötzsch & Rudolph, 2011), and *super-weak acyclicity* (SWA) (Marnette, 2009). Syntactic acyclicity criteria have also been investigated in the context of logic programs with function symbols in the rule heads, where the goal is to recognise logic programs with finite stable models. Several such notions have been implemented in state of the art logic programming engines, such as omega-restrictedness (Syrjänen, 2001) from the SMOLETS system (Syrjänen & Niemelä, 2001), lambda-restrictedness from the ASP grounder GRINGO (Gebser, Schaub, & Thiele, 2007), argument-restrictedness (Lierler & Lifschitz, 2009) from the DLV system (Leone, Pfeifer, Faber, Eiter, Gottlob, Perri, & Scarcello, 2006), and many others (Calimeri, Cozza, Ianni, & Leone, 2008; Greco, Spezzano, & Trubitsyna, 2012; De Schreye & Decorte, 1994).

1.2 Applications of Acyclicity Notions

Acyclicity notions are interesting for several reasons. First, unlike guarded rules, acyclic rules can axiomatise structures of arbitrary shapes, as long as these structures are bounded in size. Second, the result of the chase for acyclic rules can be stored and manipulated as if it were a database; this is important, for example, in data exchange, where the goal is to materialise the transformed database.

In this paper, we further argue that acyclicity notions are also relevant to *description logics* (DLs)—knowledge representation formalisms underpinning the OWL 2 ontology language (Cuenca Grau, Horrocks, Motik, Parsia, Patel-Schneider, & Sattler, 2008). CQ answering over DL ontologies is a key reasoning service in many DL applications, and the problem was studied for numerous different DLs (Calvanese, De Giacomo, Lembo, Lenzerini, & Rosati, 2007; Krötzsch, Rudolph, & Hitzler, 2007; Glimm, Horrocks, Lutz, & Sattler, 2008; Ortiz, Calvanese, & Eiter, 2008; Lutz, Toman, & Wolter, 2009; Pérez-Urbina, Motik, & Horrocks, 2009; Rudolph & Glimm, 2010; Kontchakov, Lutz, Toman, Wolter, & Zakharyashev, 2011). Answering CQs over ontologies, however, is quite technical and often of high computational complexity. Therefore, practical OWL 2 reasoners frequently solve this problem using *materialisation*—a reasoning technique in which the relevant consequences of the ontology are precomputed using chase, allowing queries to be directly evaluated in the materialised set of facts. Examples of materialisation-based systems include Oracle’s Semantic Data Store (Wu, Eadon, Das, Chong, Kolovski, Annamalai, & Srinivasan, 2008), Sesame (Broekstra, Kampman, & van Harmelen, 2002), OWLIM (Kiryakov, Ognyanov, & Manov, 2005), Jena (Carroll, Dickinson, Dollin, Reynolds, Seaborne, & Wilkinson, 2004), and DLE-Jena (Meditkos & Bassiliades, 2008). Such reasoning is possible if (i) the ontology is *Horn* (Hustadt, Motik, & Sattler, 2005) and thus does not require disjunctive reasoning, and (ii) the chase is guaranteed to terminate. To satisfy the second assumption, reasoners often consider only axioms in the OWL 2 RL profile (Motik, Cuenca Grau, Horrocks, Wu, Fokoue, & Lutz, 2009a); this systematically excludes generating rules and thus trivially ensures chase termination, but it also makes the approach incomplete. Generating rules are partially supported in systems such as OWLIM (Bishop & Bojanov, 2011) and Jena, but such support is typically ad hoc and provides no completeness and/or termination guarantees. Acyclicity notions can be used to address these issues: if an ontology is Horn and acyclic, a complete materialisation can be computed without the risk of non-termination.

1.3 Our Contributions

Motivated by the practical importance of chase termination, in this paper we present two new acyclicity notions: *model-faithful acyclicity* (MFA) and *model-summarising acyclicity* (MSA). Roughly speaking, these acyclicity notions use a particular model of the rules to analyse the implications between existential quantifiers, which is why we call them *model based*. In particular, MFA uses the actual ‘canonical’ model induced by the facts and the rules, which makes the notion very general. We prove that checking whether a set of existential rules is MFA is 2EXPTIME-complete, and it becomes EXPTIME-complete if the predicates in the rules are of bounded arity. Due to the high complexity, MFA may be unsuitable for practical application. Thus, we introduce MSA, which can be understood as MFA in which the analysis is performed over models that ‘summarise’ (or overestimate) the

actual models. Checking MSA of existential rules can be realised via checking entailment of ground atoms in datalog programs. We use this close connection between MSA and datalog to prove that checking MSA is EXPTIME-complete for general existential rules, and that it becomes coNP-complete if the arity of rule predicates is bounded.

We next conduct a detailed investigation of the landscape of known acyclicity notions, augmented with MFA and MSA. For the class of logic programs that correspond to existential rules with skolemised existential quantifiers, we show that MSA and MFA strictly subsume existing acyclicity notions known from logic programming. We also show that MSA is strictly more general than SWA—one of the most general acyclicity notions known in database theory. Furthermore, we investigate the relationship between the known notions and thus complete the picture with respect to their relative expressiveness.

Both MSA and MFA can be applied to general existential rules *without* equality. Equality can be incorporated via *singularisation*—a technique proposed by Marnette (2009) that transforms the rules to encode the effects of equality. Singularisation is orthogonal to acyclicity: after computing the transformed rules, one can use MFA, MSA, or in fact any notion to check whether the result is acyclic; if so, the chase of the singularised rules terminates, and the chase result can be used in a particular way to answer arbitrary CQs. Unfortunately, singularisation is nondeterministic: some ways of transforming the rules may produce acyclic rule sets, but not all ways are guaranteed to do so. In this paper, we refine singularisation to obtain practically useful upper and lower bounds for acyclicity. We also show that, when used with JA, our lower bound actually coincides with WA.

We next turn our attention to theoretical and practical issues of using acyclicity for materialisation-based CQ answering over ontologies. On the theoretical side, we show that checking MFA and MSA of Horn-*SROLF* ontologies is EXPTIME- and PTIME-complete, respectively, and that answering CQs over acyclic Horn-*SROLF* ontologies is EXPTIME-complete as well. Furthermore, we show that, for Horn-*SHIF* ontologies, the complexity of checking MFA and of answering CQs drops to PSPACE. Answering CQs is EXPTIME-complete for general (i.e., not acyclic) Horn-*SHIF* ontologies (Eiter, Gottlob, Ortiz, & Simkus, 2008; Ortiz, Rudolph, & Simkus, 2011), so acyclicity makes this problem easier. Furthermore, Horn ontologies can be extended with arbitrary SWRL rules (Horrocks & Patel-Schneider, 2004) without affecting decidability or worst-case complexity, provided that the union of the ontology and SWRL rules is acyclic; this is in contrast to the general case, where SWRL extensions of DLs easily lead to undecidability.

On the practical side, we explore the limits of reasoning with acyclic OWL 2 ontologies via materialisation. We checked MFA, MSA, and JA for 336 Horn ontologies; furthermore, to estimate the impact of materialisation, we compared the size of the materialisation with the number of facts in the original ontologies. Our experiments revealed that many ontologies are MSA, and that some complex ones are MSA but not JA; furthermore, the universal models obtained via materialisation are typically not too large. Thus, our results suggest that principled, materialisation-based reasoning for ontologies beyond the OWL 2 RL profile may be practically feasible.

This is an extended version of a paper by Cuenca Grau, Horrocks, Krötzsch, Kupke, Magka, Motik, and Wang (2012) published at KR 2012.

2. Preliminaries

In this section we introduce definitions and notation used in the rest of our paper.

2.1 First-Order Logic

We use the standard notions of constants, function symbols, and predicate symbols, where \approx is the equality predicate, \top is universal truth, and \perp is universal falsehood. Each function or predicate symbol is associated with a nonnegative integer arity. Variables, terms, substitutions, atoms, first-order formulae, sentences, interpretations (i.e., structures), and models are defined as usual. By a slight abuse of notation, we often identify a conjunction with the set of its conjuncts. Furthermore, we often abbreviate a vector of terms t_1, \dots, t_n as \vec{t} ; we define $|\vec{t}| = n$; and we often identify \vec{t} with the set of indexed terms $\{t_1, \dots, t_n\}$. With $\varphi(\vec{x})$ we stress that $\vec{x} = x_1, \dots, x_n$ are the free variables of a formula φ , and with $\varphi\sigma$ we denote the result of applying a substitution σ to φ . A term, atom, or formula is *ground* if it does not contain variables; a *fact* is a ground atom. The *depth* $\text{dep}(t)$ of a term t is defined as 0 if t is a constant or a variable, and $\text{dep}(t) = 1 + \max_{i=1}^n \text{dep}(t_i)$ if $t = f(t_1, \dots, t_n)$. A term t' is a *subterm* of a term t if $t' = t$ or $t = f(\vec{s})$ and t' is a subterm of some $s_i \in \vec{s}$; if additionally $t' \neq t$, then t' is a *proper subterm* of t . A term s is *contained* in an atom $P(\vec{t})$ if $s \in \vec{t}$, and s *occurs* in $P(\vec{t})$ if s is a subterm of some term $t_i \in \vec{t}$; thus, if s is contained in $P(\vec{t})$, then s also occurs in $P(\vec{t})$, but the converse may not hold. A term s is *contained* (resp. *occurs*) in a set of atoms I if s is *contained* (resp. *occurs*) in some atom in I .

In first-order logic, the equality predicate \approx is commonly assumed to have a predefined interpretation—that is, every first-order interpretation is required to interpret \approx as the smallest reflexive relation over the domain. Satisfaction of a sentence φ in an interpretation I where \approx is interpreted in this way is written $I \models \varphi$, and entailment of a sentence ψ from a sentence φ is written $\varphi \models \psi$. Unless otherwise stated, we use this standard interpretation of equality throughout this paper.

Equality, however, can also be treated as an ordinary predicate with an explicit axiomatisation. Let Σ be an arbitrary set of function-free first-order formulae. Then, $\Sigma_{\approx} = \emptyset$ if \approx does not occur in Σ ; otherwise, Σ_{\approx} contains formulae (1)–(3) and an instance of formula (4) for each n -ary predicate P occurring in Σ different from \approx , and for each $1 \leq i \leq n$. Note that all variables in all of these formulae are (implicitly) universally quantified.

$$\rightarrow x \approx x \tag{1}$$

$$x_1 \approx x_2 \rightarrow x_2 \approx x_1 \tag{2}$$

$$x_1 \approx x_2 \wedge x_2 \approx x_3 \rightarrow x_1 \approx x_3 \tag{3}$$

$$P(x_1, \dots, x_i, \dots, x_n) \wedge x_i \approx x'_i \rightarrow P(x_1, \dots, x'_i, \dots, x_n) \tag{4}$$

If \approx is treated as an ordinary predicate, satisfaction of a formula φ in a model I is written $I \models_{\approx} \varphi$, and entailment of a formula ψ from formula φ is written $\varphi \models_{\approx} \psi$. Please note that, according to our definitions, $I \models_{\approx} \varphi$ can hold even if interpretation I interprets predicate \approx in an arbitrary way; in contrast, $I \models \varphi$ can hold only if interpretation I interprets predicate \approx as the identity relation on the model's domain. The consequences of Σ w.r.t. \models and of $\Sigma \cup \Sigma_{\approx}$ w.r.t. \models_{\approx} coincide—that is, for each first-order sentence ψ constructed using the symbols from Σ , we have $\Sigma \models \psi$ if and only if $\Sigma \cup \Sigma_{\approx} \models_{\approx} \psi$.

2.2 Rules and Queries

An *instance* is a finite set of function-free facts. An *existential rule* (or just *rule*) is a function-free sentence of the form

$$\forall \vec{x} \forall \vec{z}. [\varphi(\vec{x}, \vec{z}) \rightarrow \exists \vec{y}. \psi(\vec{x}, \vec{y})] \quad (5)$$

where $\varphi(\vec{x}, \vec{z})$ and $\psi(\vec{x}, \vec{y})$ are conjunctions of atoms, and tuples of variables \vec{x} , \vec{y} , and \vec{z} are pairwise disjoint. Formula φ is the *body* and formula ψ is the *head* of the rule. For brevity, quantifiers $\forall \vec{x} \forall \vec{z}$ are often omitted. For convenience, we sometimes identify a rule body or head with the set of the respective conjuncts. A *datalog rule* is a rule where \vec{y} is empty. A rule is *equality-free* if it does not contain the equality predicate \approx . A term s *occurs* in an existential rule if s occurs in a head or body atom of the rule, and these definitions are extended to a set of rules in the obvious way; existential rules do not contain function symbols, so an analogous notion of s being *contained* in a rule coincides with this one. Two variables are *directly connected* in a rule if they occur together in a body atom of the rule; furthermore, *connected* is the transitive closure of *directly connected*; finally, a rule of the form (5) is *connected* if all pairs of variables $w, w' \in \vec{x} \cup \vec{z}$ are connected in the rule.

A *conjunctive query* (CQ) is a formula of the form $Q(\vec{x}) = \exists \vec{y}. \varphi(\vec{x}, \vec{y})$, where $\varphi(\vec{x}, \vec{y})$ is a conjunction of atoms; the query is *Boolean* if \vec{x} is empty. A substitution θ mapping \vec{x} to constants is an *answer* to $Q(\vec{x})$ w.r.t. a set of rules Σ and instance I if $\Sigma \cup I \models Q(\vec{x})\theta$. Answering CQs is the core reasoning problem in many applications of existential rules.

When answering a conjunctive query $Q(\vec{x})$ over a set of rules Σ and an instance I , in the rest of this paper we implicitly assume that $Q(\vec{x})$ and I contain only the predicates from Σ . This simplifies the presentation since it allows us to define various transformations of Σ without having to take into account possible predicates that occur in $Q(\vec{x})$ or I only. This assumption is w.l.o.g., as we can always extend Σ with tautological rules of the form $P(\vec{x}) \rightarrow P(\vec{x})$ for each predicate P occurring in $Q(\vec{x})$ or I but not in Σ .

Furthermore, we assume that \approx does not occur in the body of any rule in Σ or in the query $Q(\vec{x})$. This is w.l.o.g. since we can eliminate each atom of the form $x \approx t$ in a rule body and further replace x with t in the rest of the rule; furthermore, to eliminate body atoms of the form $a \approx b$ with a and b constants, we can introduce a fresh predicate O_a , add a new rule $\rightarrow O_a(a)$, replace each body atom $a \approx b$ with conjunction $O_a(x) \wedge x \approx b$ in which x is a fresh variable, and finally eliminate atom $x \approx b$ as before. Similarly, we do not provide an explicit support for the inequality predicate $\not\approx$. Inequality in rule heads can be simulated using an ordinary predicate: each atom of the form $s \not\approx t$ occurring in a rule head can be replaced with **NotEqual**(s, t), where **NotEqual** is a fresh ordinary predicate that is explicitly axiomatised as irreflexive; note that, if \approx is handled as a regular predicate explicitly axiomatised by rules (1)–(4), then the replacement axioms (4) must be instantiated for $P = \mathbf{NotEqual}$ as well. In contrast, atoms involving the inequality predicate occurring in rule bodies generally require disjunctive reasoning, which is not supported by existential rules.

Finally, we assume that conjunctions $\varphi(\vec{x}, \vec{z})$ and $\psi(\vec{x}, \vec{y})$ in each rule of the form (5) are both not empty. We also assume that \top and \perp are treated as ordinary unary predicates, and that the semantics of \top is captured explicitly in Σ by instantiating the following rule for each n -ary predicate P occurring in Σ :

$$P(x_1, \dots, x_n) \rightarrow \top(x_1) \wedge \dots \wedge \top(x_n) \quad (6)$$

These assumptions ensure that $I \cup \Sigma$ is always satisfiable, but that $\Sigma \cup I \models \exists y. \perp(y)$ if and only if $I \cup \Sigma$ is unsatisfiable w.r.t. the conventional treatment of \top and \perp . By allowing body atoms of the form $\top(x)$, without loss of generality we can require each existential rule to be *safe* (i.e., that each universally quantified variable occurring in a head atom also occurs in a body atom of the rule), which greatly simplifies many of our definitions.

In database theory, satisfaction and entailment are often considered only w.r.t. *finite* interpretations under the *unique name assumption* (UNA); the latter ensures that distinct constants are interpreted as distinct elements. In contrast, such assumptions are not customary in ontology-based KR. In this paper, we do not assume UNA, as UNA can be axiomatised explicitly if needed using the inequality predicate (or a simulation thereof). Furthermore, in this paper we investigate theories that are satisfiable in finite models (i.e., for which the chase is finite); thus, the difference between finite and infinite satisfiability is immaterial to our results.

We frequently use *skolemisation* to interpret rules in *Herbrand* interpretations, which are defined as possibly infinite sets of ground atoms. In particular, for each rule r of the form (5) and each variable $y_i \in \vec{y}$, let f_r^i be a function symbol globally unique for r and y_i of arity $|\vec{x}|$; furthermore, let θ_{sk} be the substitution such that $\theta_{\text{sk}}(y_i) = f_r^i(\vec{x})$ for each $y_i \in \vec{y}$. Then, the skolemisation $\text{sk}(r)$ of r is the following rule:

$$\varphi(\vec{x}, \vec{z}) \rightarrow \psi(\vec{x}, \vec{y})\theta_{\text{sk}} \quad (7)$$

The skolemisation $\text{sk}(\Sigma)$ of a set of rules Σ is obtained by skolemising each rule in Σ . Skolemisation does not affect the answers to CQs—that is, for each conjunctive query $Q(\vec{x})$ formed from only the predicates in Σ , each instance I , and each substitution σ , we have $\Sigma \cup I \models Q(\vec{x})\sigma$ if and only if $\text{sk}(\Sigma) \cup \Sigma_{\approx} \cup I \models_{\approx} Q(\vec{x})\sigma$.

2.3 The Skolem Chase

Answering CQs can be characterised using *chase*, and in this paper we use the *skolem chase* variant (Marnette, 2009). Let $r = \varphi \rightarrow \psi$ be a skolemised rule and let I be a set of ground atoms. A set of ground atoms S is a *consequence* of r on I if substitution σ exists mapping the variables in r to the terms occurring in I such that $\varphi\sigma \subseteq I$ and $S \subseteq \psi\sigma$. The result of *applying* r to I , written $r(I)$, is the union of all consequences of r on I . For Ω a set of skolemised rules, $\Omega(I) = \bigcup_{r \in \Omega} r(I)$. Let I be a finite set of ground atoms, let Σ be a set of rules, let $\Sigma' = \text{sk}(\Sigma) \cup \Sigma_{\approx}$, and let Σ'_f and Σ'_n be the subsets of Σ' containing rules with and without function symbols, respectively. The *chase sequence* for I and Σ is a sequence of sets of facts $I_{\Sigma}^0, I_{\Sigma}^1, \dots$ where $I_{\Sigma}^0 = I$ and, for each $i > 0$, set I_{Σ}^i is defined as follows:

- if $\Sigma'_n(I_{\Sigma}^{i-1}) \not\subseteq I_{\Sigma}^{i-1}$, then $I_{\Sigma}^i = I_{\Sigma}^{i-1} \cup \Sigma'_n(I_{\Sigma}^{i-1})$,
- otherwise $I_{\Sigma}^i = I_{\Sigma}^{i-1} \cup \Sigma'_f(I_{\Sigma}^{i-1})$.

The *chase* of I and Σ is defined as $I_{\Sigma}^{\infty} = \bigcup_i I_{\Sigma}^i$; note that I_{Σ}^{∞} can be infinite. The chase can be used as a ‘database’ for answering CQs: a substitution σ is an answer to Q over Σ and I if and only if $I_{\Sigma}^{\infty} \models_{\approx} Q\sigma$. The chase of I and Σ *terminates* if $i \geq 0$ exists such that $I_{\Sigma}^i = I_{\Sigma}^j$ for each $j \geq i$; the chase of Σ terminates *universally* if the chase of I and Σ terminates for each I . If the skolem chase of I and Σ terminates, then both the *nonoblivious chase* (Fagin et al., 2005) and the *core chase* (Deutsch et al., 2008) of I and Σ terminate as well.

The *critical instance* I_Σ^* for a set of rules Σ contains all facts that can be constructed using all predicates occurring in Σ , all constants occurring in the body of a rule in Σ , and one special fresh constant $*$. The skolem chase for I_Σ^* and Σ terminates if and only if the skolem chase of Σ terminates universally (Marnette, 2009).

2.4 Acyclicity Notions

Checking whether the skolem chase terminates on a given instance is undecidable, and checking universal skolem chase termination is conjectured to be undecidable as well. Consequently, various sufficient *acyclicity notions* have been proposed in the literature. Formally, an acyclicity notion X is a class of finite sets of rules; such a definition allows us to talk about (proper) containment between acyclicity notions. We sometimes write ‘ Σ is X ’, by which we mean ‘ $\Sigma \in X$ ’. We next introduce weak and joint acyclicity: the former is one of the first such notions considered in the literature; and as we show in Section 3, the latter notion is relatively powerful, yet still easy to understand. We use these two notions throughout the paper to present examples and state various technical claims. In Section 3 we present the definitions of many other acyclicity notions known in the literature.

In the following, let Σ be a set of rules where no variable occurs in more than one rule. A *position* is an expression of the form $P|_i$ where P is an n -ary predicate and i is an integer with $1 \leq i \leq n$. Given a rule r of the form (5) and a variable w occurring in r , the set $\text{Pos}_B(w)$ of *body positions* of w contains each position $P|_i$ such that $P(t_1, \dots, t_n) \in \varphi(\vec{x}, \vec{z})$ and $t_i = w$ for some vector \vec{t} of terms. The set $\text{Pos}_H(w)$ of *head positions* is defined analogously, but w.r.t. the head atoms of r . Note that, since each variable occurs in at most one rule in Σ , sets $\text{Pos}_B(w)$ and $\text{Pos}_H(w)$ are (indirectly) associated with the rule that contains w . In the rest of this paper, whenever we use notation such as $\text{Pos}_H(w)$ or $\text{Pos}_B(w)$, we silently assume that no variable occurs in more than one rule and so the notation is unambiguous. This is clearly w.l.o.g. as one can always arbitrarily rename variables in different rules.

Weak acyclicity (WA) (Fagin et al., 2005) can be applied to existential rules that contain the equality predicate. The *WA dependency graph* $WA(\Sigma)$ for Σ contains positions as vertices; furthermore, for each rule $r \in \Sigma$ of the form (5), each variable $x \in \vec{x}$, each position $P|_i \in \text{Pos}_B(x)$, and each variable $y \in \vec{y}$, graph $WA(\Sigma)$ contains

- a *regular edge* from $P|_i$ to each $Q|_j \in \text{Pos}_H(x)$ such that $Q \neq \approx$ and,
- a *special edge* from $P|_i$ to each $Q|_j \in \text{Pos}_H(y)$ such that $Q \neq \approx$.

Set Σ is WA if $WA(\Sigma)$ does not contain a cycle that involves a special edge. Equality atoms are effectively ignored by WA.

Joint acyclicity (JA) (Krötzsch & Rudolph, 2011) generalises WA, but it is applicable only to equality-free rules. For an existentially quantified variable y in Σ , let $\text{Move}(y)$ be the smallest set of positions such that

- $\text{Pos}_H(y) \subseteq \text{Move}(y)$, and
- for each existential rule $r \in \Sigma$ and each universally quantified variable x occurring in r , if $\text{Pos}_B(x) \subseteq \text{Move}(y)$, then $\text{Pos}_H(x) \subseteq \text{Move}(y)$.

The *JA dependency graph* $JA(\Sigma)$ of Σ is defined as follows. The vertices of $JA(\Sigma)$ are the existentially quantified variables occurring in Σ . Given arbitrary two such variables y_1 and y_2 , the JA dependency graph $JA(\Sigma)$ contains an edge from y_1 to y_2 whenever the rule that contains y_2 also contains a universally quantified variable x such that $\text{Pos}_H(x) \neq \emptyset$ and $\text{Pos}_B(x) \subseteq \text{Move}(y_1)$. Set Σ is JA if $JA(\Sigma)$ does not contain a cycle.

2.5 Rule Normalisation

Existential rules can often be transformed into other existential rules by replacing parts of the rule head or body with atoms involving fresh predicates. Such a transformation is called *normalisation*, and is often used as a preprocessing step to bring the rules into a suitable form. For example, Horn OWL 2 axioms can be translated into existential rules by using the well known transformations of first-order logic, and the latter can then be normalised to a form we describe in Section 6. In this section we introduce a definition of rule normalisation that captures all similar methods known to us.

Let r be a rule of the form (8), where φ_1 , φ_2 , ψ_1 , and ψ_2 are conjunctions of atoms satisfying $\vec{x}_1 \cup \vec{x}_2 = \vec{x}_3 \cup \vec{x}_4$, $\vec{z}_2 \cap \vec{z}_3 = \emptyset$, and $\vec{y}_2 \cap \vec{y}_3 = \emptyset$.

$$\varphi_1(\vec{x}_1, \vec{z}_1, \vec{z}_2) \wedge \varphi_2(\vec{x}_2, \vec{z}_1, \vec{z}_3) \rightarrow \exists \vec{y}_1, \vec{y}_2, \vec{y}_3. [\psi_1(\vec{x}_3, \vec{y}_1, \vec{y}_2) \wedge \psi_2(\vec{x}_4, \vec{y}_1, \vec{y}_3)] \quad (8)$$

A *normalisation step* replaces a conjunction in either the head or the body of the rule with an atom involving a fresh predicate. More precisely, a *head normalisation step* replaces $\psi_1(\vec{x}_3, \vec{y}_1, \vec{y}_2)$ with atom $Q(\vec{x}_3, \vec{y}_1)$ where Q is a fresh predicate, thus replacing r with rule (9), and it adds rule (10).

$$\varphi_1(\vec{x}_1, \vec{z}_1, \vec{z}_2) \wedge \varphi_2(\vec{x}_2, \vec{z}_1, \vec{z}_3) \rightarrow \exists \vec{y}_1, \vec{y}_3. [Q(\vec{x}_3, \vec{y}_1) \wedge \psi_2(\vec{x}_4, \vec{y}_1, \vec{y}_3)] \quad (9)$$

$$Q(\vec{x}_3, \vec{y}_1) \rightarrow \exists \vec{y}_2. \psi_1(\vec{x}_3, \vec{y}_1, \vec{y}_2) \quad (10)$$

Alternatively, a *body normalisation step* replaces $\varphi_1(\vec{x}_1, \vec{z}_1, \vec{z}_2)$ with atom $Q(\vec{x}_1, \vec{z}_1)$ where Q is a fresh predicate, thus replacing r with rule (11), and it adds rule (12).

$$Q(\vec{x}_1, \vec{z}_1) \wedge \varphi_2(\vec{x}_2, \vec{z}_1, \vec{z}_3) \rightarrow \exists \vec{y}_1, \vec{y}_2, \vec{y}_3. [\psi_1(\vec{x}_3, \vec{y}_1, \vec{y}_2) \wedge \psi_2(\vec{x}_4, \vec{y}_1, \vec{y}_3)] \quad (11)$$

$$\varphi_1(\vec{x}_1, \vec{z}_1, \vec{z}_2) \rightarrow Q(\vec{x}_1, \vec{z}_1) \quad (12)$$

Given a set of existential rules Σ , normalisation steps are often applied to Σ iteratively. If the predicate Q introduced in each step is always fresh, we call such normalisation *without structure sharing*. In contrast, normalisation *with structure sharing* allows the predicate Q to be reused across different normalisation steps. For example, once a predicate Q is introduced in a head normalisation step to replace $\varphi_1(\vec{x}_1, \vec{z}_1, \vec{z}_2)$, then a conjunction of the form $\varphi_1(\vec{x}'_1, \vec{z}'_1, \vec{z}'_2)$ where $\vec{x}'_1, \vec{z}'_1, \vec{z}'_2$ are renamings of $\vec{x}_1, \vec{z}_1, \vec{z}_2$ can be replaced with $Q(\vec{x}'_3, \vec{y}'_1)$ without introducing the corresponding rule (10). An analogous optimisation can be used in a body normalisation step.

Let Σ' be a set of rules obtained via normalisation (with or without structure sharing) from Σ . It is well known that Σ' is a conservative extension of Σ . Consequently, for each instance I and each BCQ Q that does not use the freshly introduced predicates, we have $\Sigma \cup I \models Q$ if and only if $\Sigma' \cup I \models Q$.

3. Novel Acyclicity Notions

Weak acyclicity has considerably influenced the field of data exchange in databases, but it is a rather strict notion and so it may not be sufficient in many applications of existential rules. Joint acyclicity significantly relaxes weak acyclicity and was developed mainly for rule based knowledge representation applications.

In Section 3.1 we show that even joint acyclicity—one of the most general acyclicity notions developed so far—does not capture rules corresponding to axioms commonly found in ontologies for which the chase terminates universally. To address this important limitation, we propose in Section 3.2 *model-faithful acyclicity* (MFA)—a novel, very general, notion that can be used to successfully ensure chase termination for many ontologies used in practice. The computational cost of checking MFA is, however, rather high; hence, in Section 3.3 we introduce *model-summarising acyclicity* (MSA)—a more strict notion that is easier to check and produces the same results as MFA on most existing ontologies.

3.1 Limitations of Existing Acyclicity Notions

To motivate our new acyclicity notions, we first present an example that shows how known acyclicity notions, such as JA, are not satisfied by rules that are equivalent to very simple axioms that abound in OWL ontologies.

Example 1. Let Σ be the set of rules (13)–(17).

$$r_1 = A(x_1) \rightarrow \exists y_1. R(x_1, y_1) \wedge B(y_1) \quad (13)$$

$$r_2 = R(x_2, z_1) \wedge B(z_1) \rightarrow A(x_2) \quad (14)$$

$$r_3 = B(x_3) \rightarrow \exists y_2. R(x_3, y_2) \wedge C(y_2) \quad (15)$$

$$r_4 = C(x_4) \rightarrow D(x_4) \quad (16)$$

$$r_5 = R(x_5, z_2) \wedge D(z_2) \rightarrow B(x_5) \quad (17)$$

Rules r_1 and r_2 correspond to the description logic axiom $A \equiv \exists R.B$, rule r_3 corresponds to axiom $B \sqsubseteq \exists R.C$, rule r_4 corresponds to axiom $C \sqsubseteq D$, and rule r_5 corresponds to axiom $\exists R.D \sqsubseteq B$. Such axioms are very common in OWL ontologies.

By the definition of JA from Section 2, we have $\text{Move}(y_1) = \{R|_2, B|_1, R|_1, A|_1\}$. Thus, the JA dependency graph contains an edge from y_1 to itself, so the set of axioms Σ is not JA. In contrast, the following table shows the chase sequence for I_Σ^* and Σ .

$A(*)$	$R(*, f(*))$	$R(f(*), g(f(*)))$	$D(g(f(*)))$
$B(*)$	$B(f(*))$	$C(g(f(*)))$	
$C(*)$	$R(*, g(*))$	$D(g(*))$	
$D(*)$	$C(g(*))$		
$R(*, *)$			

Rule r_2 is not applicable to $R(f(*), g(f(*)))$ since I_3 does not contain the fact $B(g(f(*)))$ necessary to match the atom $B(z_1)$ from the rule. Thus, the chase terminates. \diamond

All existing acyclicity notions essentially try to estimate whether an application of a rule can produce facts that can (possibly by applying chase to other rules) repeatedly

trigger the same rule in an infinite manner. The key difference between various notions is how rule applicability is determined. In particular, JA considers each variable in a rule in isolation and does not check satisfaction of all body atoms at once; for example, rule (14) is not applicable to the facts generated by rule (15), but this can be determined only by considering variables x_2 and z_1 in rule (14) simultaneously. These notions thus overestimate rule applicability and, as a result, they can fail to detect chase termination.

3.2 Model-Faithful Acyclicity (MFA)

Our main intuition for addressing this problem is that more precise chase termination guarantees can be obtained by tracking rule applicability more ‘faithfully’. A simple solution is to be completely precise about rule applicability: one can run the skolem chase and then use sufficient checks to identify cyclic computations. Since no sufficient, necessary, and computable test can be given for the latter, we must adopt a practical approach. For example, we can ‘raise the alarm’ and stop the process if the chase derives a ‘cyclic’ term $f(\vec{t})$, where f occurs in \vec{t} . This idea can be further refined; for example, one could stop only if f occurs nested in a term some fixed number of times. The choice of the appropriate test thus depends on an application; however, as our experiments show, checking only for one level of nesting suffices in many cases. In particular, no term $f(\vec{t})$ with f occurring in \vec{t} is generated in the chase of the set of rules Σ from Example 1.

Definition 2. *A term t is cyclic if a function symbol f exists such that some term $f(\vec{s})$ is a subterm of t , and some term $f(\vec{u})$ is a proper subterm of $f(\vec{s})$.*

Our notion of acyclicity is declarative: the given set of rules Σ is transformed into a new set of rules Σ' that tracks rule dependencies using fresh predicates; then, Σ is identified as being acyclic if Σ' does not entail a special nullary predicate C . Since acyclicity is defined via entailment, it can be decided using any theorem proving procedure for existential rules that is sound and complete. Acyclicity guarantees termination of the skolem chase, which also guarantees termination of nonoblivious chase and core chase. We call our notion *model-faithful acyclicity* because it estimates rule application precisely, by examining the actual structure of the universal model of Σ .

Definition 3. *For each rule $r = \varphi(\vec{x}, \vec{z}) \rightarrow \exists \vec{y}. \psi(\vec{x}, \vec{y})$ and each variable $y_i \in \vec{y}$, let F_r^i be a fresh unary predicate unique for r and y_i ; furthermore, let S and D be fresh binary predicates, and let C be a fresh nullary predicate. Then, $\text{MFA}(r)$ is the following rule:*

$$\varphi(\vec{x}, \vec{z}) \rightarrow \exists \vec{y}. \left[\psi(\vec{x}, \vec{y}) \wedge \bigwedge_{y_i \in \vec{y}} \left[F_r^i(y_i) \wedge \bigwedge_{x_j \in \vec{x}} S(x_j, y_i) \right] \right]$$

For a set Σ of rules, $\text{MFA}(\Sigma)$ is the smallest set that contains $\text{MFA}(r)$ for each rule $r \in \Sigma$, rules (18)–(19), and rule (20) instantiated for each F_r^i corresponding to some $r \in \Sigma$:

$$S(x_1, x_2) \rightarrow D(x_1, x_2) \tag{18}$$

$$D(x_1, x_2) \wedge S(x_2, x_3) \rightarrow D(x_1, x_3) \tag{19}$$

$$F_r^i(x_1) \wedge D(x_1, x_2) \wedge F_r^i(x_2) \rightarrow C \tag{20}$$

The set Σ is model-faithful acyclic (MFA) w.r.t. an instance I if $I \cup \text{MFA}(\Sigma) \not\models C$; furthermore, Σ is universally MFA¹ if Σ is MFA w.r.t. I_Σ^* .

Example 4. Let Σ be the set of rules from Example 1. Then, $\text{MFA}(r_1)$ and $\text{MFA}(r_3)$ are given by (21) and (22), respectively; since r_1 and r_3 contain a single existentially quantified variable each, we omit the superscripts in F_{r_1} and F_{r_3} for the sake of clarity. Thus, $\text{MFA}(\Sigma)$ consists of rules (14), (16), and (17), rules (21)–(22), rules (18)–(19), and rule (20) instantiated for F_{r_1} and F_{r_3} .

$$A(x_1) \rightarrow \exists y_1. R(x_1, y_1) \wedge B(y_1) \wedge F_{r_1}(y_1) \wedge S(x_1, y_1) \quad (21)$$

$$B(x_3) \rightarrow \exists y_2. R(x_3, y_2) \wedge C(y_2) \wedge F_{r_3}(y_2) \wedge S(x_3, y_2) \quad (22)$$

It is straightforward to see that the chase of I_Σ^* and $\text{MFA}(\Sigma)$ consists of the facts presented in Example 1, augmented with the following facts:

$S(*, f(*))$	$D(*, f(*))$	$D(f(*), g(f(*)))$
$S(*, g(*))$	$D(*, g(*))$	$D(*, g(f(*)))$
$F_{r_1}(f(*))$	$S(f(*), g(f(*)))$	
$F_{r_3}(g(*))$	$F_{r_3}(g(f(*)))$	

The chase of I_Σ^* and $\text{MFA}(\Sigma)$ does not contain C , which implies that $I \cup \text{MFA}(\Sigma) \not\models C$. As a result, Σ is universally MFA. \diamond

MFA is formulated as a semantic, rather than a syntactic notion, and is thus mainly independent from algorithmic details: entailment $I \cup \text{MFA}(\Sigma) \models C$ can be checked using an arbitrary sound and complete first-order calculus. In Section 4 we discuss the relationship between MFA and existing notions, and we show that MFA generalises most of them.

The following proposition shows that MFA characterises the derivations of the skolem chase in which no cyclic terms occur.

Proposition 5. A set Σ of rules is not MFA w.r.t. an instance I if and only if $I_{\text{MFA}(\Sigma)}^\infty$ contains a cyclic term.

Proof. Let $\Sigma' = \text{MFA}(\Sigma)$, and let $I_{\Sigma'}^0, I_{\Sigma'}^1, \dots$ be the chase sequence for I and Σ' . Moreover, let f_r^i be the function symbol used to skolemise the i -th existentially quantified variable in rule r , as defined in Section 2.2. We next prove that the following claims hold for all terms t and t' occurring in $I_{\Sigma'}^k$, each rule r , each integer i , and each integer k , as well as $k = \infty$.

1. Term t is of the form $f_r^i(\vec{u})$ if and only if $F_r^i(t) \in I_{\Sigma'}^k$.
2. Term t is of the form $f_r^i(\vec{u})$ and $t' \in \vec{u}$ if and only if $S(t', t) \in I_{\Sigma'}^k$.
3. If t' is a proper subterm of t , then $D(t', t) \in I_{\Sigma'}^{k+2}$; furthermore, $D(t', t) \in I_{\Sigma'}^\infty$ if and only if t' is a proper subterm of t .

1. In the rest of this paper we often omit ‘universally’; furthermore, when used as an acyclicity notion, MFA means ‘universally MFA’.

(Claims 1 and 2, direction \Rightarrow) The proof is by induction on k . Set $I_{\Sigma'}^0$ does not contain functional terms, and so it clearly satisfies both claims. For the induction step, assume that both claims hold for $I_{\Sigma'}^{k-1}$ and consider $I_{\Sigma'}^k$. Since $I_{\Sigma'}^{k-1} \subseteq I_{\Sigma'}^k$, both claims clearly hold for each term t that occurs in $I_{\Sigma'}^{k-1}$. Consider an arbitrary term t of the form $f_r^i(\vec{u})$ that does not occur in $I_{\Sigma'}^{k-1}$, and an arbitrary term $t' \in \vec{u}$. Clearly, t is introduced into $I_{\Sigma'}^k$ by an application of the skolemisation of $\text{MFA}(r)$ for some rule $r \in \Sigma$. Since the head of $\text{MFA}(r)$ contains atoms $F_r^i(y_i)$ and $S(x_j, y_i)$ for each $x_j \in \vec{x}$, we have $F_r^i(t) \in I_{\Sigma'}^k$ and $S(t', t) \in I_{\Sigma'}^k$ for each $t' \in \vec{u}$, and so we have $F_r^i(t) \in I_{\Sigma'}^\infty$ and $S(t', t) \in I_{\Sigma'}^\infty$ for each $t' \in \vec{u}$ as well. Finally, since $I_{\Sigma'}^\infty = \bigcup_k I_{\Sigma'}^k$, these claims clearly hold for $k = \infty$.

(Claims 1 and 2, direction \Leftarrow) Predicate S and each predicate F_r^i occur in Σ' only in head atoms of the form $F_r^i(y_i)$ and $S(x_j, y_i)$; hence, the skolemised rules contain these predicates only in head atoms of the form $F_r^i(f_r^i(\vec{x}))$ and $S(x_j, f_r^i(\vec{x}))$, which clearly implies our claim.

(Claim 3, the first part for $k \neq \infty$) The proof is by induction on k . The base case holds vacuously since $I_{\Sigma'}^0$ does not contain functional terms. Assume now that the claim holds for some $k - 1$, and consider an arbitrary term $t = f_r^i(\vec{u})$ occurring in $I_{\Sigma'}^k$ such that t' is a subterm of some $t_i \in \vec{u}$. By Claim 2, we have $S(t_i, t) \in I_{\Sigma'}^k$; furthermore, t_i occurs in $I_{\Sigma'}^{k-1}$, so by the induction assumption we have $D(t', t_i) \in I_{\Sigma'}^{k+1}$. Finally, the rules without functional terms are applied before the rules with functional terms; hence, by rule (19) we have $D(t', t) \in I_{\Sigma'}^{k+2}$, as required.

(Claim 3, the second part) The ‘proper subterm’ relation is transitive, and rules (18) and (19) effectively define D as the transitive closure of S , which clearly implies this claim.

Assume now that $I_{\Sigma'}^\infty$ contains a cyclic term t . Then, some term $t_1 = f_r^i(\vec{s})$ is a subterm of t and some term $t_2 = f_r^i(\vec{u})$ is a proper subterm of t_1 . By Claims 1 and 3, then we have $\{F_r^i(t_2), D(t_2, t_1), F_r^i(t_1)\} \subseteq I_{\Sigma'}^\infty$. But then, since Σ' contains rule (20), we have $C \in I_{\Sigma'}^\infty$, so Σ is not MFA. For the converse claim, assume that Σ is not MFA w.r.t. an instance I . Then, by Definition 3 we have that $I \cup \text{MFA}(\Sigma) \models C$. Since the special nullary predicate C occurs only on the right-hand side of rule (20), there exist terms t_1 and t_2 , a rule $r \in \Sigma$, and a predicate F_r^i such that $\{F_r^i(t_1), D(t_1, t_2), F_r^i(t_2)\} \subseteq I_{\Sigma'}^\infty$. Since $F_r^i(t_1)$ and $F_r^i(t_2)$ are contained in $I_{\Sigma'}^\infty$, Claim 1 implies that t_1 and t_2 are of the form $t_1 = f_r^i(\vec{u}_1)$ and $t_2 = f_r^i(\vec{u}_2)$, respectively. Finally, $D(t_1, t_2) \in I_{\Sigma'}^\infty$ and Claim 3 imply that t_1 is a proper subterm of t_2 , so $I_{\Sigma'}^\infty$ contains a cyclic term. \square

This characterisation implies termination of skolem chase of MFA rules Σ in 2EXPTIME. In particular, a term t derived by the skolem chase of $\Sigma' = \text{MFA}(\Sigma)$ cannot be cyclic by Proposition 5; such t can then be seen as a tree with branching factor bounded by the maximum arity of a function symbol in $\text{sk}(\Sigma')$ and with depth bounded by the number of function symbols in $\text{sk}(\Sigma')$. The chase can thus generate at most a doubly exponential number of different terms and atoms. The 2EXPTIME bound already holds if the rules are WA (Calì et al., 2010b), so CQ answering for MFA rules is not harder than for WA rules.

Proposition 6. *If a set of rules Σ is MFA w.r.t. an instance I , then the skolem chase for I and Σ terminates in double exponential time.*

Proof. Let $\Sigma' = \text{MFA}(\Sigma)$, let c , f , and p be the number of constants, function symbols, and predicate symbols, respectively, occurring in $\text{sk}(\Sigma')$, let ℓ be the maximum arity of a function symbol, and let a be the maximum arity of a predicate symbol in $\text{sk}(\Sigma')$. Consider

now an arbitrary term t occurring in $I_{\Sigma'}^{\infty}$; clearly, t can be seen as a tree with branching factor ℓ containing constants in the leaf nodes and function symbols in the internal nodes; furthermore, since t is not cyclic, $\text{dep}(t) \leq f$, the number of leaves is bounded by ℓ^f , and the total number of nodes is bounded by $f \cdot \ell^f$. Each node is assigned a constant or a function symbol, so the number of different terms occurring in $I_{\Sigma'}^{\infty}$ is bounded by $\wp = (c + f)^{f \cdot \ell^f}$, and the number of different atoms in $I_{\Sigma'}^{\infty}$ is bounded by $p \cdot \wp^a$, which is clearly doubly exponential in Σ and I . Consequently, the size of $I_{\Sigma'}^{\infty}$ is at most doubly exponential in Σ and I . Furthermore, for an arbitrary set of facts I' and rule r , the set $r(I')$ can be computed by examining all mappings of the variables in r to the terms occurring in I' , which requires exponential time in the size of r and polynomial time in the size of I' . Consequently, $I_{\Sigma'}^{\infty}$ can be computed in time that is double exponential in I and Σ . Finally, it is straightforward to see that $I_{\Sigma}^{\infty} \subseteq I_{\Sigma'}^{\infty}$, so I_{Σ}^{∞} can be computed in double exponential time as well. \square

By Proposition 6, answering a BCQ over MFA rules is in 2EXPTIME. We next prove that checking MFA w.r.t. a specific instance I is also in 2EXPTIME, and that checking universal MFA is 2EXPTIME-hard. This provides tight complexity bounds for both problems. Towards this goal, we first establish in Lemma 7 a relationship between answering certain kinds of queries over certain kinds of rules and checking whether a related set of rules is universally MFA; we use this relationship in several hardness proofs in the rest of this paper. Then, in Theorem 8 we present our main complexity result.

Lemma 7. *Let Σ be a set of weakly acyclic, constant-free, equality-free, and connected rules with predicates of nonzero arity, let A and B be unary predicates, let R be a fresh binary predicate, let a be a constant, and let Ω be Σ extended with rule (23).*

$$R(z, x) \wedge B(x) \rightarrow \exists y. [R(x, y) \wedge A(y)] \quad (23)$$

Then, we have $\{A(a)\} \cup \Sigma \not\models B(a)$ if and only if Ω is universally MFA.

Proof. Let $I = \{A(a)\}$, and let $I_{\Sigma}^0, I_{\Sigma}^1, \dots$ be the chase sequence for I and Σ . Furthermore, let $\Omega' = \text{MFA}(\Omega)$, let $J = I_{\Omega}^*$, let $J_{\Omega'}^0, J_{\Omega'}^1, \dots$ be the chase sequence for J and Ω' , and let f be the function symbol used to skolemise the existential quantifier in rule (23). Set Σ is constant-free, so a is the only constant occurring in each set I_{Σ}^i .

We next show that the facts in $J_{\Omega'}^j$ are of a certain form. To this end, for each $\ell \geq 0$, let $t_{\ell} = f(\dots f(*) \dots)$ where the function symbol f is repeated ℓ times (by this definition, we have $t_0 = *$); also, each term or fact obtained from t_{ℓ} by zero or more applications of predicates or function symbols not in $\{f, D, S, C, R\}$ is of *level* ℓ . By induction on the chase sequence for J and Ω' , we next prove that the sequence satisfies the following property (\blacklozenge):

for each fact $F \in J_{\Omega'}^j$, some integer ℓ exists such that F is of the form $R(*, *)$ or $R(t_{\ell}, t_{\ell+1})$, or the predicate of F is contained in $\{D, S, C\}$, or F is an ℓ -level fact and the predicate of F is not contained in $\{D, S, C, R\}$.

Set $J_{\Omega'}^0 = J$ clearly satisfies property (\blacklozenge) since each fact in it is clearly of level 0. Now assume that $J_{\Omega'}^j$ satisfies property (\blacklozenge) for some j , and consider an application of a rule $r \in \Omega'$. If r corresponds to rule (18), (19), (20), or (23), then the result of the rule application clearly satisfies property (\blacklozenge). Otherwise, r is safe and no body atom contains a predicate in

$\{D, S, C, R\}$; by induction assumption, then some atom is matched to a fact of some level ℓ ; the body atoms of r are connected, so all body atoms are matched to facts of the same level; finally, the head atoms of r contain function symbols different from f , but no constants or predicates of zero arity, so each fact derived by an atom in the head of r either contains predicate S or is of level ℓ .

We next show that the chase sequences for I and Σ , and for J and Ω' are related by the following property (\diamond) :

for each fact F' of level 1 and the fact F obtained by replacing each t_1 in F' with a , we have $F \in I_\Sigma^i$ for some i if and only if $F' \in J_{\Omega'}^j \setminus J$ for some j .

The proof of (\diamond) is straightforward: J contains $R(*, *)$ and $B(*)$, so $J_{\Omega'}^1$ contains $R(*, f(*))$ and $A(f(*))$; moreover, due to (\blacklozenge) , term t_1 plays in the chase sequence for J and Ω' the ‘same role’ as constant a in the chase sequence for I and Σ , so the rule applications to facts of level 1 in the former chase sequence correspond one-to-one with rule applications in the latter chase sequence. We omit the formal details for the sake brevity.

Now assume that $\{A(a)\} \cup \Sigma \models B(a)$. Then, $B(a) \in I_\Sigma^i$ holds for some i . By property (\diamond) , then integer j exists such that $B(f(*)) \in J_{\Omega'}^j$. But then, due to rule (23), some $\ell \geq j$ exists such that $A(f(f(*))) \in J_{\Omega'}^\ell$. By Proposition 5, then Ω is not universally MFA.

Conversely, assume that $\{A(a)\} \cup \Sigma \not\models B(a)$. Since Σ is weakly acyclic and equality-free, Σ is super-weakly acyclic (Marnette, 2009); as we will show in Section 4 (see Theorem 19), Σ is then MFA as well. Now consider an arbitrary integer j and fact $F \in J_{\Omega'}^j$. If F is of level 0 or 1, since Σ is MFA, fact F does not contain a cyclic term. Furthermore, $B(a) \notin I_\Sigma^\infty$ so, by property (\diamond) , fact F is not of the form $B(f(*))$; thus, rule (23) does not ‘fire’ to introduce facts of level greater than 1. Consequently, F does not contain a cyclic term, and so, by Proposition 5, the set Ω is universally MFA. \square

Theorem 8. *Given a set of rules Σ , deciding whether Σ is MFA w.r.t. an instance I is in 2EXPTIME, and deciding whether Σ is universally MFA is 2EXPTIME-hard. Both results hold even if the arity of predicates in Σ is bounded.*

Proof. (Membership) Let $\Sigma' = \text{MFA}(\Sigma)$, let $I_{\Sigma'}^0, I_{\Sigma'}^1, \dots$ be the chase sequence for I and Σ' , and let \wp, p , and a be as stated in the proof of Proposition 6. The number of different atoms that can be constructed from \wp terms is bounded by $k = p \cdot \wp^a$; note that this is double exponential even if a is bounded. Let $k' = k + 4$; we next show that whether Σ is MFA w.r.t. I can be decided by constructing $I_{\Sigma'}^{k'}$ and then checking whether $C \in I_{\Sigma'}^{k'}$. As in the proof of Proposition 6, the latter can be done in double exponential time.

If $I_{\Sigma'}^k = I_{\Sigma'}^{k'}$, then $I_{\Sigma'}^\infty = I_{\Sigma'}^k$, so Σ is not MFA if and only if $C \in I_{\Sigma'}^k$. Otherwise, we have $I_{\Sigma'}^k \subsetneq I_{\Sigma'}^{k'}$; but then, $I_{\Sigma'}^{k'+1}$ clearly contains at least one cyclic term $t = f_r^i(\vec{t})$ such that $t' = f_r^i(\vec{s})$ is a subterm of some $t_i \in \vec{t}$. Since $I_{\Sigma'}^{k'+1}$ satisfies Claims 1–3 from the proof of Proposition 5, we have $D(t_i, t) \in I_{\Sigma'}^{k+3}$; by rule (20) and the fact that rules without functional terms are applied before rules with functional terms, we have $C \in I_{\Sigma'}^{k'}$; thus, $C \in I_{\Sigma'}^\infty$, so Σ is not MFA by Proposition 5.

(Hardness) We prove the claim by a reduction from the problem of checking $I \cup \Sigma \models Q$, where Σ is a *weakly acyclic* set of equality-free rules and with predicates of bounded arity,

I is an instance, and $Q = \exists \vec{y}. \xi(\vec{y})$ is a Boolean conjunctive query. Cali et al. (2010b) show that, for such I , Σ , and Q , deciding $I \cup \Sigma \models Q$ is 2EXPTIME-complete. We next transform I , Σ , and Q so that we can apply Lemma 7, which proves our claim.

Let $\Sigma_1 = \Sigma \cup \{\xi(\vec{y}) \rightarrow B\}$ where B is a fresh predicate of zero arity; clearly, $I \cup \Sigma \models Q$ holds if and only if $I \cup \Sigma_1 \models B$ holds.

Let Σ_2 and I_2 be obtained by eliminating constants from the rules in Σ_1 ; that is, we initially set $I_2 = I$ and then, for each rule $r \in \Sigma$ and each constant c occurring in r , we replace all occurrences in c with a fresh variable w_c , add an atom $O_c(w_c)$ to the body of r where O_c is a fresh predicate uniquely associated with c , and add a fact $O_c(c)$ to I_2 . It is straightforward to see that $I \cup \Sigma_1 \models B$ if and only if $I_2 \cup \Sigma_2 \models B$.

Finally, to transform Σ_2 and I_2 into Σ_3 , we define some notation. Let \hat{P} be a fresh $n+1$ -ary predicate unique for each n -ary predicate P , and let w be a fresh variable not occurring in Σ_2 . For a conjunction of atoms φ , let $\hat{\varphi} = \bigwedge_{P(\vec{t}) \in \varphi} \hat{P}(\vec{t}, w)$. Rule (24) is obtained from I_2 as specified below, where A is a fresh unary predicate, each constant c occurring in I_2 is associated with a distinct, fresh variable v_c , and \vec{v}_c is the vector of all such variables:

$$A(w) \rightarrow \exists \vec{v}_c. \bigwedge_{P(c_1, \dots, c_k) \in I_2} \hat{P}(v_{c_1}, \dots, v_{c_k}, w) \quad (24)$$

Finally, the set Σ_3 contains rule (24) and a rule

$$\hat{\varphi}(\vec{x}, \vec{z}, w) \rightarrow \exists \vec{y}. \hat{\psi}(\vec{x}, \vec{y}, w) \quad \text{for each rule} \quad \varphi(\vec{x}, \vec{z}) \rightarrow \exists \vec{y}. \psi(\vec{x}, \vec{y}) \text{ in } \Sigma_2. \quad (25)$$

Clearly, all predicates in Σ_3 are of nonzero arity; all rules in Σ_3 are constant-free and connected; and A occurs only in the body of rule (24) and Σ_2 is WA, so Σ_3 is WA as well. Finally, let $I_3 = \{A(a)\}$ where a is a fresh constant; by induction on the chase sequences for Σ_2 and I_2 , and Σ_3 and I_3 , it is straightforward to show that, for each integer i and each fact $P(c_1, \dots, c_k)$, we have $P(c_1, \dots, c_k) \in (I_2)_{\Sigma_2}^i$ if and only if $\hat{P}(f_{c_1}(a), \dots, f_{c_k}(a), a) \in (I_3)_{\Sigma_3}^{i+1}$, where f_{c_1}, \dots, f_{c_k} are the skolem functions used to skolemise v_{c_1}, \dots, v_{c_k} in rule (24). Thus, $I_2 \cup \Sigma_2 \models B$ if and only if $I_3 \cup \Sigma_3 \models \hat{B}(a)$, which by Lemma 7 implies our claim. \square

The results of Theorem 8 are somewhat discouraging: known acyclicity notions can typically be checked in PTIME or in NP. We consider MFA to be an ‘upper bound’ of practically useful acyclicity notions. We see two possibilities for improving these results. In Section 3.3 we introduce an approximation of MFA that is easier to check; our experiments (see Section 7) show that this notion often coincides with MFA in practice. Furthermore, we show next that the complexity is lower for rules of the following shape.

Definition 9. *A rule r of the form (5) is an \exists -1 rule if \vec{y} is empty or \vec{x} contains at most one variable.*

As we discuss in the following sections, \exists -1 rules capture (extensions of) Horn DLs. We next show that BCQ answering and MFA checking for \exists -1 rules is easier than for general rules. Intuitively, if Σ is MFA and contains only \exists -1 rules, then all functional terms in $\text{sk}(\text{MFA}(\Sigma))$ are unary and hence the number of different terms and atoms derivable by chase becomes exponentially bounded, as shown by the following theorem.

Theorem 10. *Let Σ be a set of \exists -1 rules, and let I be an instance. Checking whether Σ is MFA w.r.t. I is in EXPTIME, and checking whether Σ is universally MFA is EXPTIME-hard. Moreover, if Σ is MFA w.r.t. I , then answering a BCQ over Σ and I is EXPTIME-complete.*

Proof. We defer the proof of both hardness claims to Section 6, which deals with an even smaller class of rules that correspond to Horn description logic ontologies. In particular, we prove hardness of BCQ answering in Lemma 59, and hardness of checking whether Σ is MFA w.r.t. I in Lemma 60. In the rest of this proof, we show both membership results.

Let $\Sigma' = \text{MFA}(\Sigma)$; let c be the number of constants in an instance; and let f be the number of function symbols in the rules. Since Σ contains only \exists -1 rules, Σ' also contains only \exists -1 rules; consequently, all functional terms in $\text{sk}(\Sigma')$ are of arity 1. Hence, each noncyclic term can be understood as a sequence of at most f function symbols, so the total number of different noncyclic terms is bounded by $\wp = c \cdot (f + 1)^f$. The total number of atoms is bounded by $p \cdot \wp^a$, where p is the number of predicates and a is the maximum arity of a predicate in Σ' . Note that this is exponential even if a is fixed. As in the proof of Proposition 6, we can now show that either the chase for Σ' and I terminates or a cyclic term is derived in exponential time, which proves that the complexity of checking whether Σ is MFA w.r.t. I is in EXPTIME.

Finally, since $I_\Sigma^\infty \subseteq I_{\Sigma'}^\infty$, if Σ is MFA, then I_Σ^∞ can be computed in exponential time, so a BCQ over Σ and I can be answered in EXPTIME. \square

3.3 Model-Summarising Acyclicity (MSA)

The high cost of checking MFA of Σ arises because the arity of function symbols in $\text{sk}(\Sigma)$ is unbounded and the depth of cyclic terms can be linear in Σ . To obtain an acyclicity notion that is easier to check, we must coarsen the structure used for cycle analysis. We thus next introduce *model-summarising acyclicity*, which ‘summarises’ the models of Σ by reusing the same constant to satisfy an existential quantifier, instead of introducing ‘deep’ terms.

Definition 11. *Let S , D , and F_r^i be as specified in Definition 3; furthermore, for each rule $r = \varphi(\vec{x}, \vec{z}) \rightarrow \exists \vec{y}. \psi(\vec{x}, \vec{y})$ and each variable $y_i \in \vec{y}$, let c_r^i be a fresh constant unique for r and y_i . Then, $\text{MSA}(r)$ is the following rule, where θ_{MSA} is the substitution that maps each variable $y_i \in \vec{y}$ to c_r^i :*

$$\varphi(\vec{x}, \vec{z}) \rightarrow \psi(\vec{x}, \vec{y})\theta_{\text{MSA}} \wedge \bigwedge_{y_i \in \vec{y}} \left[F_r^i(y_i)\theta_{\text{MSA}} \wedge \bigwedge_{x_j \in \vec{x}} S(x_j, y_i)\theta_{\text{MSA}} \right]$$

For a set Σ of rules, $\text{MSA}(\Sigma)$ is the smallest set that contains $\text{MSA}(r)$ for each rule $r \in \Sigma$, rules (18)–(19), and rule (20) instantiated for each predicate F_r^i . Set Σ is model-summarising acyclic (MSA) w.r.t. an instance I if $I \cup \text{MSA}(\Sigma) \not\models C$; furthermore, Σ is universally MSA if Σ is MSA w.r.t. I_Σ^ .*

Example 12. *Consider again the set of rules Σ from Example 1. $\text{MSA}(r_1)$ and $\text{MSA}(r_3)$ are given by rules (26) and (27), respectively; since r_1 and r_3 contain a single existentially quantified variable each, we omit the superscripts in F_{r_1} , F_{r_3} , c_{r_1} , and c_{r_3} for the sake of*

clarity. Thus, $\text{MSA}(\Sigma)$ consists of rules (14), (16), and (17), rules (26)–(27), rules (18)–(19), and rule (20) instantiated for F_{r_1} and F_{r_3} .

$$A(x_1) \rightarrow R(x_1, c_{r_1}) \wedge B(c_{r_1}) \wedge F_{r_1}(c_{r_1}) \wedge S(x_1, c_{r_1}) \quad (26)$$

$$B(x_3) \rightarrow R(x_3, c_{r_3}) \wedge C(c_{r_3}) \wedge F_{r_3}(c_{r_3}) \wedge S(x_3, c_{r_3}) \quad (27)$$

The following table shows the chase sequence for I_Σ^* and $\text{MSA}(\Sigma)$.

$A(*)$	$R(*, c_{r_1})$	$R(c_{r_1}, c_{r_3})$	$D(c_{r_1}, c_{r_3})$
$B(*)$	$R(*, c_{r_3})$	$D(c_{r_3})$	
$C(*)$	$B(c_{r_1})$	$S(c_{r_1}, c_{r_3})$	
$D(*)$	$C(c_{r_3})$	$D(*, c_{r_3})$	
$R(*, *)$	$S(*, c_{r_1})$	$D(*, c_{r_1})$	
	$S(*, c_{r_3})$		
	$F_{r_1}(c_{r_1})$		
	$F_{r_3}(c_{r_3})$		

The result of the chase does not contain C , and so Σ is universally MSA. \diamond

Note that $\text{MSA}(\Sigma)$ is equivalent to a set of datalog rules: the only minor difference is that the rules in $\text{MSA}(\Sigma)$ can contain several head atoms, but such rules can clearly be transformed into equivalent datalog rules. Thus, MSA can be checked using a datalog reasoner. This connection with datalog and the complexity results by Dantsin, Eiter, Gottlob, and Voronkov (2001) for checking entailment of a ground atom in a datalog program provide us with the upper complexity bound for checking MSA in Theorem 13. The complexity of datalog reasoning is $O(r \cdot n^v)$ where r is the number of rules, v is the maximum number of variables in a rule, and n is the size of the set of facts that the rules are applied to; thus, checking MSA should be feasible if the rules in Σ are ‘short’ and so v is ‘small’.

Theorem 13. *For Σ a set of rules, deciding whether Σ is MSA w.r.t. an instance I is in EXPTIME, and deciding whether Σ is universally MSA is EXPTIME-hard. The two problems are in coNP and coNP-hard, respectively, if the arity of the predicates in Σ is bounded.*

Proof. (Membership) Let $\Sigma' = \text{MSA}(\Sigma)$, and note that Σ is MSA w.r.t. I if and only if $I \cup \Sigma' \not\models C$ if and only if $C \notin I_{\Sigma'}^\infty$. The total number of atoms occurring in $I_{\Sigma'}^\infty$ is $p \cdot c^a$, where p is the number of predicates, c is the number of constants, and a is the maximum arity of the predicates in Σ' ; this number is clearly exponential if a is not bounded. The rest of the proof is the same as in Theorem 8.

Assume now that a is bounded; then the number of ground atoms in $I_{\Sigma'}^\infty$ becomes polynomial. Furthermore, by the definition of the chase, $C \in I_{\Sigma'}^\infty$ if and only if there exist a sequence of rules r_1, \dots, r_n of the form $r_i = \varphi_i \rightarrow \psi_i$ and a sequence of substitutions $\sigma_1, \dots, \sigma_n$ such that $\varphi\sigma_i \subseteq I \cup \{\psi_j\sigma_j \mid j < i\} \subseteq I_{\Sigma'}^\infty$ for each $1 \leq i \leq n$ and $\psi_n\sigma_n = C$. Clearly, we can assume that $n \leq p \cdot c^a$, which is polynomial. Thus, we can guess the two sequences in non-deterministic polynomial time, and we can check the required property in polynomial time. Thus, $I \cup \Sigma' \models C$ can be checked in nondeterministic polynomial time, so checking whether Σ is MSA w.r.t. I is in coNP.

(Hardness) Let Σ be a set of datalog rules, let I be an instance, and let Q be ground atom. Checking whether $I \cup \Sigma \models Q$ is EXPTIME-complete in general (Dantsin et al., 2001). Furthermore, the problem is NP-hard if the arity of predicates is bounded: a rule in Σ can encode an arbitrary Boolean conjunctive query with atoms of bounded arity but arbitrarily many variables, for which answering is well known to be NP-hard.

Let Σ_4 and I_4 be obtained from Σ as in the proof of Theorem 8; then, $I \cup \Sigma \models Q$ if and only if $I_4 \cup \Sigma_4 \models B(a)$, and the set of rules Ω obtained from Σ_4 as specified in Lemma 7 is universally MFA if and only if $I_4 \cup \Sigma_4 \not\models B(a)$. Finally, the only existential variable in Ω occurs in a rule of the form (23), so it is straightforward to see that Ω is universally MFA if and only if \mathcal{T}' is universally MSA. \square

Before concluding this section, we present Theorem 14 and Example 15, which together show that MFA is strictly more general than MSA.

Theorem 14. *If a set of rules Σ is MSA (w.r.t. an instance I), then Σ is MFA (w.r.t. I) as well.*

Proof. Let $\Sigma_1 = \text{MFA}(\Sigma)$ and let $\Sigma_2 = \text{MSA}(\Sigma)$. Furthermore, let h be the mapping of ground terms to constants defined such that $h(t) = c_r^i$ if t is of the form $f_r^i(\dots)$, and $h(t) = t$ if t is a constant; for an atom $A = P(t_1, \dots, t_n)$, let $h(A) = P(h(t_1), \dots, h(t_n))$; and for an instance I , let $h(I) = \{h(A) \mid A \in I\}$. Finally, let $I_{\Sigma_1}^0, I_{\Sigma_1}^1, \dots$ be the chase sequence for I and Σ_1 , and let $I_{\Sigma_2}^0, I_{\Sigma_2}^1, \dots$ be the chase sequence for I and Σ_2 . Note that $\text{sk}(\Sigma_2) = \Sigma_2$ differs from $\text{sk}(\Sigma_1)$ only in that the former contains the constant c_r^i in place of each functional term $f_r^i(\vec{x})$. Please note that, although our definition of the chase applies rules with function symbols after rules without function symbols, one can clearly construct the chase of the function-free set of rules Σ_2 using any order of rule applications, including the one corresponding to the order of rule applications in the chase of Σ_1 . Assuming this slight modification, one can show by a straightforward induction on i that $h(I_{\Sigma_1}^i) \subseteq I_{\Sigma_2}^i$ for each i ; this implies $h(I_{\Sigma_1}^\infty) \subseteq I_{\Sigma_2}^\infty$. Consequently, $C \notin I_{\Sigma_2}^\infty$ clearly implies $C \notin I_{\Sigma_1}^\infty$; hence, if Σ is MSA, then Σ is MFA as well, as required. \square

Example 15. *Let Σ be the set of rules (28)–(31).*

$$r_1 = A(x) \rightarrow \exists y. R(x, y) \wedge B(y) \quad (28)$$

$$r_2 = B(x) \rightarrow \exists y. S(x, y) \wedge T(y, x) \quad (29)$$

$$r_3 = A(z) \wedge S(z, x) \rightarrow C(x) \quad (30)$$

$$r_4 = C(z) \wedge T(z, x) \rightarrow A(x) \quad (31)$$

It is straightforward to check that Σ is universally MFA, but not universally MSA. \diamond

3.4 Acyclicity Notions and Normalisation

As mentioned in Section 2.5, existential rules are often normalised into a particular form; however, this cannot destroy acyclicity: if a set of rules Σ is MFA, then each set of rules obtained from Σ by normalisation is MFA as well. This claim involves certain technical assumptions about the treatment of equality, which is why we postpone a formal proof of this statement until Section 5. Next, however, we show that normalisation can have a positive effect on acyclicity.

Example 16. Let Σ be the set containing only the following rule:

$$A(x) \rightarrow \exists y.[B(x) \wedge A(y)] \quad (32)$$

As specified in Section 2.2, this rule is skolemised as follows, which causes that the skolem chase of Σ and instance $I = \{A(a)\}$ does not terminate.

$$A(x) \rightarrow B(x) \wedge A(f(x)) \quad (33)$$

Note, however, that atoms $B(x)$ and $A(y)$ in the head of the rule do not share variables, so we can normalise this rule as follows, where Q is a fresh predicate of zero arity:

$$A(x) \rightarrow B(x) \wedge Q \quad (34)$$

$$Q \rightarrow \exists y.A(y) \quad (35)$$

It is straightforward to check that this normalised set of rules is MFA; in fact, the normalised set of rules is even JA. Intuitively, normalisation, as defined in Section 2.5, ensures that each functional symbol introduced during normalisation depends on as few variables in the rule as possible. \diamond

Normalisation, however, can have a negative effect on universal termination, as shown by the following example.

Example 17. Let Σ be the set containing only the following rule:

$$C(z) \wedge R(z, x) \wedge B(x) \rightarrow \exists y_1 \exists y_2.[R(x, y_1) \wedge R(y_1, y_2) \wedge B(y_2)] \quad (36)$$

One can readily check that Σ is universally MFA. Now let Σ' be the following set of rules, which is obtained by replacing conjunction $R(y_1, y_2) \wedge B(y_2)$ in the rule head with $Q(y_1)$:

$$C(z) \wedge R(z, x) \wedge B(x) \rightarrow \exists y_1.[R(x, y_1) \wedge Q(y_1)] \quad (37)$$

$$Q(y_1) \rightarrow \exists y_2.[R(y_1, y_2) \wedge B(y_2)] \quad (38)$$

Let f_1 and f_2 be function symbols used to skolemise the existential quantifier in rule (37) and (38), respectively. Since $Q(*) \in I_{\Sigma'}^*$, the chase of Σ' and $I_{\Sigma'}^*$ derives $R(*, f_2(*))$ and $B(f_2(*))$; but then, these facts, $C(*)$, and rule (37) derive $Q(f_1(f_2(*)))$, after which rule (38) derives $R(f_1(f_2(*)), f_2(f_1(f_2(*))))$ and $B(f_2(f_1(f_2(*))))$. The chase of Σ' and $I_{\Sigma'}^*$ thus contains a cyclic term, so Σ' is not universally MFA.

Intuitively, this problem occurs because the critical instance $I_{\Sigma'}^*$ for Σ' also instantiates the predicate Q introduced during normalisation. Such predicates, however, cannot occur in arbitrary input instances, so we can use the critical instance for Σ . Since $Q(*) \notin I_{\Sigma}^*$, the skolem chase of Σ' and I_{Σ}^* does not derive a cyclic term, from which we can conclude that the skolem chase of Σ' terminates on each instance I that contains facts constructed using only the predicates occurring in Σ . \diamond

4. Relationship with Known Acyclicity Notions

Many acyclicity notions have been proposed in the literature, but the relationships between them have been only partially investigated. We next investigate the relationship between MFA, MSA, and the acyclicity notions known to us, and we produce a detailed picture of their relative expressiveness. We show that MFA and MSA generalise most of these notions.

4.1 Acyclicity in Databases

Acyclicity notions have been considered in databases in data integration and data exchange scenarios. Weak acyclicity (Fagin et al., 2005) was one of the first such notions, and it has spurred on the research into more sophisticated notions for ensuring chase termination.

4.1.1 SUPER-WEAK ACYCLICITY

Marnette (2009) proposed *super-weak acyclicity* (SWA), which generalises weak acyclicity provided that the rules are equality-free. We next recapitulate the definition of SWA, and then we show that MSA and MFA are strictly more general than SWA.

Definition 18. Let Σ be a set of existential rules in which no variable occurs in more than one rule, and let θ_{sk} be the substitution used to skolemise the rules in Σ .² A place is a pair $\langle A, i \rangle$ where A is an n -ary atom occurring in a rule in Σ and $1 \leq i \leq n$. A set of places P' covers a set of places P if, for each place $\langle A, i \rangle \in P$, a place $\langle A', i' \rangle \in P'$ and substitutions σ and σ' exist such that $A\sigma = A'\sigma'$ and $i = i'$. Given a variable w occurring in a rule $r = \varphi \rightarrow \exists \vec{y}. \psi$, sets of places $\text{In}(w)$, $\text{Out}(w)$, and $\text{Move}(w)$ are defined as follows:

- set $\text{In}(w)$ contains each place $\langle R(\vec{t}), i \rangle$ such that $R(\vec{t}) \in \varphi$ and $t_i = w$;
- set $\text{Out}(w)$ contains each place $\langle R(\vec{t})\theta_{\text{sk}}, i \rangle$ such that $R(\vec{t}) \in \psi$ and $t_i = w$; and
- set $\text{Move}(w)$ is the smallest set of places such that
 - $\text{Out}(w) \subseteq \text{Move}(w)$ and
 - $\text{Out}(w') \subseteq \text{Move}(w)$ for each variable w' that is universally quantified in some rule in Σ such that $\text{Move}(w)$ covers $\text{In}(w')$.

The SWA dependency graph $\text{SWA}(\Sigma)$ of Σ contains a vertex for each rule of Σ , and an edge from a rule $r \in \Sigma$ to a rule $r' \in \Sigma$ if a variable x' occurring in both the body and the head of r' and an existentially quantified variable y occurring in the head of r exists such that $\text{Move}(y)$ covers $\text{In}(x')$. Set Σ is *super-weakly acyclic* (SWA) if $\text{SWA}(\Sigma)$ is *acyclic*.

Marnette (2009) uses a slightly different definition: the notation for places is the same as our notation for positions; a variable may occur in more than one rule so sets $\text{In}(w)$, $\text{Out}(w)$, and $\text{Move}(w)$ are defined w.r.t. a rule and a variable; and a rule trigger relation is used instead of the SWA dependency graph. For simplicity, Definition 18 introduces SWA in the same style as JA; however, both definitions capture the same class of rules.

The following theorem shows that MSA is more general than SWA. Furthermore, in Example 12 we argued that the set of rules Σ from Example 1 is MSA, and one can readily check that Σ is not SWA. Consequently, MSA is strictly more general than SWA.

Theorem 19. *If a set of rules Σ is SWA, then Σ is universally MSA.*

2. Substitution θ_{sk} is unique for each rule in Section 2.2; however, since each variable in Σ occurs in at most one rule, w.l.o.g. we can take θ_{sk} as the substitution used to skolemise all the rules in Σ .

Proof. Let $\Sigma' = \text{MSA}(\Sigma)$, let I^0, I^1, \dots be the chase sequence for I_Σ^* and Σ' , and let I^∞ be the chase of I_Σ^* and Σ' . Furthermore, let ρ be the function that maps constants to themselves and that maps ground functional terms as $\rho(f_r^i(\vec{t})) = c_r^i$, where f_r^i and c_r^i were introduced in Section 2.2 and Definition 11, respectively. Finally, let $\rho(P(t_1, \dots, t_n)) = P(\rho(t_1), \dots, \rho(t_n))$.

We next prove the following property (\blacklozenge): for each rule $r \in \Sigma$, each existentially quantified variable y_i occurring in r , each $P(\vec{t}) \in I^\infty$ where $P \notin \{\text{S}, \text{D}, \text{C}\}$, and each $t_j \in \vec{t}$ such that $t_j = c_r^i$, a substitution τ and a place $\langle A, j \rangle \in \text{Move}(y_i)$ exist such that $P(\vec{t}) = \rho(A\tau)$. The proof is by induction on the length of the chase. Since $I^0 = I_\Sigma^*$ does not contain a constant of the form c_r^i , property (\blacklozenge) holds vacuously for I^0 . Assume now that property (\blacklozenge) holds for some I^{k-1} , and consider an arbitrary rule $r \in \Sigma$, an existentially quantified variable y_i in r , a fact $P(\vec{t}) \in I^k \setminus I^{k-1}$ with $P \notin \{\text{S}, \text{D}, \text{C}\}$, and a term $t_j \in \vec{t}$ such that $t_j = c_r^i$. Fact $P(\vec{t})$ is derived in I^k from the head atom H of some rule $r^1 \in \text{MSA}(\Sigma)$. Let σ be the substitution used in the rule application; clearly, we have $H\sigma = P(\vec{t})$. Furthermore, let $r^2 \in \Sigma$ be the rule such that $r^1 = \text{MSA}(r^2)$, let $r^3 = \text{sk}(r^2)$, and let H^3 be the head atom of r^3 that corresponds to H ; clearly, we have $\rho(H^3\sigma) = P(\vec{t})$. Now if H has c_r^i in position j , then $r = r^1$ since r^1 is the only rule that contains c_r^i ; thus, $\langle H^3, j \rangle \in \text{Out}(y_i) \subseteq \text{Move}(y_i)$, so property (\blacklozenge) holds. Otherwise, H contains at position j a universally quantified variable x such that $\sigma(x) = c_r^i$. Let B_1, \dots, B_n be the body atoms of r^1 that contain x ; clearly, $\{B_1\sigma, \dots, B_n\sigma\} \subseteq I^{k-1}$. All these atoms satisfy the induction assumption, so for each $B_m \in \{B_1, \dots, B_n\}$ and each ℓ such that B_m contains variable x at position ℓ , a place $\langle B'_m, \ell \rangle \in \text{Move}(y_i)$ and substitution τ^m exist such that $B_m\sigma = \rho(B'_m\tau^m)$. Let σ' be the substitution obtained from σ by setting $\sigma'(w) = \tau^m(w)$ for each variable w for which $\tau^m(w)$ is a functional term; clearly, $B_m\sigma' = B'_m\tau^m$. But then, $\text{Move}(y_i)$ covers $\text{In}(x)$; hence, by the definition of Move , we have that $\langle H^3, j \rangle \in \text{Move}(y_i)$, so property (\blacklozenge) holds.

We additionally prove the following property (\blacklozenge): if $\text{S}(c_r^i, c_{r'}^{i'}) \in I^\infty$ for some i and i' , then $\text{SWA}(\Sigma)$ contains an edge from r to r' . Consider an arbitrary such fact, let y_i be the existentially quantified variable of r corresponding to c_r^i , and let k be the smallest integer such that $\text{S}(c_r^i, c_{r'}^{i'}) \in I^k$. Clearly, $\text{S}(c_r^i, c_{r'}^{i'})$ is derived in I^k from the head atom $\text{S}(x, c_{r'}^{i'})$ of rule r' . Let σ be the substitution used in the rule application; thus, $\sigma(x) = c_r^i$. Let B_1, \dots, B_n be the body atoms of r that contain x ; clearly, we have $\{B_1\sigma, \dots, B_n\sigma\} \subseteq I^{k-1}$. All these atoms satisfy property (\blacklozenge), so for each $B_m \in \{B_1, \dots, B_n\}$ and each ℓ such that B_m contains variable x at position ℓ , a place $\langle B'_m, \ell \rangle \in \text{Move}(y_i)$ and substitution τ^m exist such that $B_m\sigma = \rho(B'_m\tau^m)$. But then, as in the previous paragraph we have that $\text{Move}(y_i)$ covers $\text{In}(x)$, so $\text{SWA}(\Sigma)$ contains an edge from r to r' .

Assume now that Σ is not MSA, so $\text{C} \in I^\infty$; then $\{F_r^i(t), \text{D}(t, t'), F_{r'}^i(t')\} \subseteq I^\infty$ holds for some F_r^i due to rules (20). But then, since predicate F_r^i occurs in Σ' only in an atom $F_r^i(c_r^i)$, we have $t = t' = c_r^i$. Finally, since D is axiomatised in Σ' as the transitive closure of S , clearly $\text{SWA}(\Sigma)$ contains a path from r to itself, and so Σ is not SWA. \square

The rule set in Example 1 is MSA but not SWA. Furthermore, it is known that SWA is more general than JA, and the two notions differ only if at least one rule contains a body atom in which at least one variable occurs more than once (Krötzsch & Rudolph, 2013). The following example shows that SWA is strictly more general than JA.

Example 20. Let Σ be the set of the following rules:

$$r_1 = \quad A(x_1) \rightarrow \exists y. R(x_1, y) \wedge R(y, x_1) \wedge R(x_1, x_1) \quad (39)$$

$$r_2 = R(x_2, x_2) \rightarrow B(x_2) \quad (40)$$

$$r_3 = B(x_3) \rightarrow A(x_3) \quad (41)$$

One can readily verify that Σ is SWA, but not JA. \diamond

Theorem 19 holds even if Σ contains the equality predicate, but provided that the axiomatisation of equality (cf. Section 2) is taken as part of the input. On such rule sets, however, SWA, JA, MSA, and MFA are not strictly more general than WA. We discuss the underlying problems, as well as possible solutions, in Section 5.

4.1.2 ACYCLICITY BY REWRITING

Spezzano and Greco (2010) proposed an acyclicity notion called *Adn-WA*. Roughly speaking, one first rewrites a set of rules Σ into another set of rules Σ' by adorning the positions in the predicates that can contain infinitely many terms during the chase; then, one checks whether Σ' is WA. The rewriting algorithm is rather involved, so we do not recapitulate its definition; instead, we discuss it by means of an example. Spezzano and Greco used this example to show that Adn-WA is not subsumed by SWA, but the same example also shows that Adn-WA is not subsumed by MFA either.

Example 21. Let Σ be the set containing the following rules:

$$A(x) \rightarrow \exists y.R(x, y) \quad (42)$$

$$B(z) \wedge R(z, x) \rightarrow A(x) \quad (43)$$

The transformation by Spezzano and Greco (2010) produces a set Σ' that consists of three groups of rules. The first group contains rules (44)–(47).

$$A^b(x) \rightarrow \exists y.R^{bf}(x, y) \quad (44)$$

$$B^b(z) \wedge R^{bb}(z, x) \rightarrow A^b(x) \quad (45)$$

$$B^b(z) \wedge R^{bf}(z, x) \rightarrow A^f(x) \quad (46)$$

$$A^f(x) \rightarrow \exists y.R^{ff}(x, y) \quad (47)$$

For each n -ary predicate P , the transformation introduces predicates of the form P^m , where m is an adornment—a string of length n of letters b or f . Intuitively, if m contains letter b at position i , then during the chase construction the i -th position of P^m can contain only constants occurring in an instance. Rules (44)–(47) were derived as follows. Rule (44) is obtained from rule (42) by marking all positions of variable x with b , which effectively creates a variant of the rule whose body is applicable only to constants. Variable y in the head of rule (44) occurs under an existential quantifier, so the corresponding position is marked with f . Rule (45) is obtained from rule (43) in an analogous way. But then, since facts introduced by rule (44) can trigger an application of rule (43), the latter rule is marked as rule (46); predicate A^f in the head of rule (46) reflects the fact that variable x in the rule body is instantiated by atom $R^{bf}(z, x)$. Finally, facts derived by rule (46) can trigger an application of rule (42), so the latter rule is instantiated as (47). At this point the algorithm terminates: since no rule was instantiated with a marking B^f in the head, it is not possible to use predicate R^{ff} to mark the body of rule (43) in a consistent way.

The second group consists of rules (48)–(50), which ‘populate’ the adorned predicates with the contents of an instance.

$$R(x_1, x_2) \rightarrow R^{bb}(x_1, x_2) \quad (48)$$

$$A(x) \rightarrow A^b(x) \quad (49)$$

$$B(x) \rightarrow B^b(x) \quad (50)$$

The third group consists of rules (51)–(56), which ‘gather’ the content of each adorned predicate P^m into a fresh output predicates \hat{P} .

$$R^{bb}(x_1, x_2) \rightarrow \hat{R}(x_1, x_2) \quad (51)$$

$$R^{bf}(x_1, x_2) \rightarrow \hat{R}(x_1, x_2) \quad (52)$$

$$R^{ff}(x_1, x_2) \rightarrow \hat{R}(x_1, x_2) \quad (53)$$

$$A^b(x) \rightarrow \hat{A}(x) \quad (54)$$

$$A^f(x) \rightarrow \hat{A}(x) \quad (55)$$

$$B^b(x) \rightarrow \hat{B}(x) \quad (56)$$

It is straightforward to check that Σ is not MFA. In contrast, Σ' is WA; furthermore, Spezzano and Greco (2010) show that, for each instance I and each vector of ground terms \vec{t} , we have $\hat{P}(\vec{t}) \in I_{\Sigma'}^\infty$ if and only if $P(\vec{t}) \in I_\Sigma^\infty$. Since Σ' is WA, $I_{\Sigma'}^\infty$ is finite, and, by the previously mentioned property, I_Σ^∞ is finite as well. \diamond

The following example shows that MFA is not subsumed by Adn-WA, which indicates that MFA and Adn-WA are incomparable.

Example 22. Let Σ be the set containing the following rules:

$$r_1 = A(x) \rightarrow \exists y. R(x, y) \wedge B(y) \quad (57)$$

$$r_2 = S(z, x) \wedge B(x) \rightarrow \exists y. S(x, y) \quad (58)$$

The rules in the first group of the set Σ' obtained by the transformation are shown below; we do not show the rules in the second and the third group for the sake of brevity.

$$A^b(x) \rightarrow \exists y. R^{bf}(x, y) \wedge B^f(y) \quad (59)$$

$$S^{bb}(z, x) \wedge B^b(x) \rightarrow \exists y. S^{bf}(x, y) \quad (60)$$

$$S^{bf}(z, x) \wedge B^f(x) \rightarrow \exists y. S^{ff}(x, y) \quad (61)$$

$$S^{ff}(z, x) \wedge B^f(x) \rightarrow \exists y. S^{ff}(x, y) \quad (62)$$

The last rule ensures that the WA dependency graph for Σ' contains a special edge from position $S^{ff}|_2$ to itself; thus, Σ' is not WA, and therefore Σ is not Adn-WA. In contrast, one can readily verify that Σ is MFA. \diamond

Spezzano and Greco (2010) also proposed several optimisations of this transformation, the discussion of which is out of scope of this paper. All of them can be seen as ‘unfolding’ the rules in Σ up to a certain number of chase steps. This seems close to an idea by Baget

et al. (2011b), who propose to run the chase for some fixed number of steps before checking for potential cycles. A similar effect could be obtained by extending the notion of MFA to check for terms that contain a function symbol nested some fixed number of times.

Finally, note that the transformation by Spezzano and Greco (2010) is independent from the notion used to check the acyclicity of the transformed rule set; hence, given an arbitrary acyclicity notion X , one can define $\text{Adn-}X$ in the obvious way. Given arbitrary notions X and Y such that $X \subseteq Y$, it is obvious that $\text{Adn-}X \subseteq \text{Adn-}Y$; consequently, we have $\text{Adn-}X \not\subseteq \text{MFA}$ for each X such that $\text{WA} \subseteq X$. In contrast, however, it is not obvious whether the inclusion between $\text{Adn-}X$ and $\text{Adn-}Y$ is strict whenever the inclusion between X and Y is strict, or whether MFA is contained in $\text{Adn-}X$ for some X with $\text{WA} \subseteq X$. Finally, we conjecture that $X \subseteq \text{Adn-}X$ holds for an arbitrary notion X , but we do not have a formal proof of this conjecture. Due to the complex nature of the rewriting, we refrain from further analysis of these relationships.

4.1.3 MONITOR GRAPH

Meier et al. (2009) propose an idea that is similar in spirit to MFA. The idea is to track each chase step in an additional data structure called the *monitor graph*. If the chase is infinite, then the monitor graph contains cycles of arbitrary length; conversely, if one can show that the monitor graph does not contain a cycle of some fixed length, then the chase is guaranteed to terminate. While this idea is closely related to MFA, note that the definition of MFA is semantic; hence, one can use an arbitrary theorem proving technique to check whether $\text{MFA}(\Sigma) \models C$. In contrast, the notion of a monitor graph is specifically tied to the nonoblivious chase. It is well known that the result of the nonoblivious chase depends on the order in which the rules are applied; consequently, a set of rules can be identified as cyclic or acyclic depending on the selected rule application strategy. Because of this dependence, it is difficult to compare the monitor graph approach with other acyclicity notions.

4.2 Acyclicity in Knowledge Representation

Existential rules can capture knowledge representation formalisms such as Horn fragments of description logics (see Section 6), conceptual graphs (Baget, 2004; Baget et al., 2011a), and datalog[±] rules (Calì et al., 2010a), and so acyclicity notions allow for materialisation-based query answering over knowledge bases. In this context, Baget (2004) and Baget et al. (2011a) proposed the notion of *acyclic graph rule dependencies* (aGRD). Intuitively, aGRD introduces a rule dependency relation \prec for which $r_1 \prec r_2$ means that an application of rule r_1 on an instance I can subsequently trigger an application of rule r_2 . If the relation \prec is acyclic, then no rule can trigger itself so the skolem chase terminates on an arbitrary instance. This can be formalised as follows.

Definition 23. *The rule dependency relation $\prec \subseteq \Sigma \times \Sigma$ on a set of rules Σ is defined as follows. Let $r_1 = \varphi_1 \rightarrow \exists \vec{y}_1.\psi_1$ and $r_2 = \varphi_2 \rightarrow \exists \vec{y}_2.\psi_2$ be arbitrary rules in Σ , and let $\text{sk}(r_1) = \varphi_1 \rightarrow \psi'_1$ and $\text{sk}(r_2) = \varphi_2 \rightarrow \psi'_2$. Then, $r_1 \prec r_2$ if and only if there exist an instance I , a substitution σ_1 for all variables in $\text{sk}(r_1)$, and a substitution σ_2 for all variables in $\text{sk}(r_2)$ such that $\varphi_1\sigma_1 \subseteq I$, $\varphi_2\sigma_2 \not\subseteq I$, $\varphi_2\sigma_2 \subseteq I \cup \psi'_1\sigma_1$, and $\psi'_2\sigma_2 \not\subseteq I \cup \psi'_1\sigma_1$. Set Σ has an acyclic graph of rules dependencies (aGRD) if the relation \prec on Σ is acyclic.*

Definition 23 differs from the original definition by Baget (2004) in several ways. First, Baget uses fresh nulls to capture the effect of existential quantifiers, whereas Definition 23 uses skolem functions; however, this does not change the resulting relation \prec in any way. Second, Baget does not require $\psi'_2\sigma_2 \not\subseteq I \cup \psi'_1\sigma_1$. This condition intuitively ensures that an application of r_1 to I enables r_2 to derive something new; analogous optimisations were proposed by Deutsch et al. (2008) and Greco et al. (2012). It should be clear that Definition 23 is stronger than the one by Baget. To unify the notions used in various parts of this paper, we included this optimisation into Definition 23; however, we nevertheless call the resulting stronger notion aGRD.

The following example shows that aGRD, even in its weaker form as originally proposed by Baget (2004), is not contained in SWA.

Example 24. Let Σ be the set consisting of the following rule:

$$r = \quad A(z_1, x, z_2) \wedge B(z_2) \rightarrow \exists y_1 \exists y_2. A(x, y_1, y_2) \quad (63)$$

To see that $r \prec r$ does not hold, consider the skolemisation r' of r :

$$\text{sk}(r) = \quad A(z_1, x, z_2) \wedge B(z_2) \rightarrow A(x, f_1(x), f_2(x)) \quad (64)$$

Now let I be an arbitrary instance, and let σ_1 and σ_2 be arbitrary substitutions such that $\{A(z_1, x, z_2)\sigma_1, B(z_2)\sigma_1\} \subseteq I$ and $\{A(z_1, x, z_2)\sigma_2, B(z_2)\sigma_2\} \not\subseteq I$. Since instance I contains only constants, atom $A(x, f_1(x), f_2(x))\sigma_1$ is of the form $A(a, f_1(a), f_2(a))$; but then, for $\{A(z_1, x, z_2)\sigma_2, B(z_2)\sigma_2\} \subseteq I \cup \{A(a, f_1(a), f_2(a))\}$ to hold, it must be that $\sigma_2(z_2) = f_2(a)$; thus, $B(z_2)\sigma_2 = B(f_2(a))$ should be contained in I , which is impossible since I is an instance and thus does not contain functional terms. Note that the additional condition by Greco et al. (2012) plays no role here. Thus, we have $r \not\prec r$, so Σ is aGRD even in the weaker form by Baget (2004). However, one can easily check that Σ is not SWA. \diamond

However, aGRD seems to be a rather weak notion: as the following example shows, even a set of rules without existential quantifiers can be cyclic according to this criterion.

Example 25. Let Σ be the set consisting of the following rules:

$$r_1 = \quad A(x) \rightarrow B(x) \quad (65)$$

$$r_2 = \quad B(x) \rightarrow C(x) \quad (66)$$

$$r_3 = \quad C(x) \rightarrow A(x) \quad (67)$$

To see that $r_1 \prec r_2$, let $I = \{A(a)\}$, let $\sigma = \{x \mapsto a\}$, and note that $A(x)\sigma \in I$, $B(x)\sigma \notin I$, $B(x)\sigma \in I \cup \{B(x)\sigma\}$, and $C(x)\sigma \notin I \cup \{B(x)\sigma\}$. Analogously, by taking $I = \{B(a)\}$ we get $r_2 \prec r_3$, and by taking $I = \{C(a)\}$ we get $r_3 \prec r_1$. Consequently, Σ is not aGRD. However, Σ is obviously WA since it does not contain existentially quantified variables. \diamond

Baget et al. (2011a) suggested that rule dependencies become more powerful if they are combined with an arbitrary acyclicity notion X . Intuitively, the main idea is to use \prec to partition a set of rules into strongly connected components, and then check whether each component is X ; we call this notion X^\prec . This idea can be formalised as follows.

Definition 26. Let Σ be a set of existential rules, and let \prec be the rule dependency relation on Σ . Relation \prec is extended to arbitrary sets $C \subseteq \Sigma$ and $C' \subseteq \Sigma$ such that $C \prec C'$ if and only if rules $r \in C$ and $r' \in C'$ exist such that $r \prec r'$. A dependency partition of Σ is a sequence of sets $\Sigma_1, \dots, \Sigma_n$ such that $\Sigma = \bigcup_{i=1}^n \Sigma_i$, each Σ_i is a strongly connected component of \prec , and $\Sigma_j \not\prec \Sigma_i$ for all i and j such that $1 \leq i < j \leq n$.

Let X be an arbitrary acyclicity notion. Then, $\Sigma \in X^\prec$ if a dependency partition $\Sigma_1, \dots, \Sigma_n$ of Σ exists such that, for each $1 \leq i \leq n$, we have $\Sigma_i \in X$, or Σ_i consists of a single rule r_i such that $r_i \not\prec r_i$.

If Σ is aGRD, then each strongly connected component Σ_i contains a single rule r_i such that $r_i \not\prec r_i$. Now if Definition 26 did not consider the special case where Σ_i consists of a single rule that does not depend on itself, then SWA^\prec would not extend aGRD; for example, the rule in Example 24 would not be in SWA^\prec . The extra condition in Definition 26 thus ensures that aGRD is contained in X^\prec regardless of the choice of X , and that aGRD can be understood as \emptyset^\prec —the acyclicity notion obtained by extending the empty notion (i.e., the notion under which no rule set is acyclic) with rule dependencies.

We next present two simple results. Proposition 27 precludes inclusions between certain acyclicity notions and will thus help us establish proper inclusions between many acyclicity notions. Furthermore, Proposition 28 shows that combining an acyclicity notion contained in SWA with rule dependencies creates a strictly stronger acyclicity notion; note that this holds even for the weaker form of rule dependencies originally proposed by Baget (2004).

Proposition 27. Let X and Y be acyclicity notions such that $X \subseteq Y$. Then, $X^\prec \subseteq Y^\prec$. Furthermore, if there exists a set $\Sigma \in Y \setminus X$ whose rule dependency relation has a cycle containing all the rules from Σ , then $Y \not\subseteq X^\prec$, $Y^\prec \not\subseteq X^\prec$, and $X^\prec \subsetneq Y^\prec$.

Proof. Relationship $X^\prec \subseteq Y^\prec$ is immediate from Definition 26. Assume now that there exists a set of rules $\Sigma \in Y \setminus X$ whose rule dependency relation has a cycle containing all the rules from Σ . By Definition 26, $\Sigma \notin X$ implies $\Sigma \notin X^\prec$, and $\Sigma \in Y$ implies $\Sigma \in Y^\prec$. But then, we clearly have $Y \not\subseteq X^\prec$ and $Y^\prec \not\subseteq X^\prec$, and the latter clearly implies $X^\prec \subsetneq Y^\prec$. \square

Proposition 28. For each acyclicity notion X such that $X \subseteq \text{SWA}$, we have $X \subsetneq X^\prec$ and $\text{aGRD} \not\subseteq X$.

Proof. Set Σ from Example 24 is in aGRD and thus in X^\prec ; however, Σ is not in SWA and hence not in X either. \square

MSA also does not contain aGRD; however, unlike for SWA, our claim depends on the optimisation in Definition 23. An analysis of the relationship between MSA and the version of rule dependencies originally proposed by Baget (2004) is out of scope of this paper.

Example 29. Let Σ be the set consisting of the following rules:

$$r_1 = R(x_1, x_1) \wedge U(x_1, z) \wedge U(x_2, z) \rightarrow R(x_1, x_2) \quad (68)$$

$$r_2 = R(z, x) \rightarrow \exists y. T(x, y) \quad (69)$$

$$r_3 = T(z, x) \rightarrow \exists y. U(x, y) \quad (70)$$

It is obvious that $r_1 \prec r_2$, $r_1 \not\prec r_3$, $r_2 \not\prec r_1$, $r_2 \not\prec r_2$, $r_2 \prec r_3$, $r_3 \not\prec r_2$, and $r_3 \not\prec r_3$. We next argue that $r_1 \not\prec r_1$ and $r_3 \not\prec r_1$, which implies that Σ is aGRD.

To see that $r_1 \not\prec r_1$, assume that an application of r_1 to an instance I produces an atom of the form $R(a, b)$; due to atom $R(x_1, x_1)$ in the body of r_1 , we have $R(a, a) \in I$. Now let $I' = I \cup \{R(a, b)\}$; since $R(a, a) \in I$, the rule application derives ‘something new’ only if $a \neq b$. Now assume that a substitution σ_2 exists that makes r_1 applicable to I' but not to I ; this rule application must ‘use’ the fact $R(a, b)$, which implies that $R(x_1, x_1)\sigma_2 = R(a, b)$; however, this is impossible since $a \neq b$. Consequently, we have $r_1 \not\prec r_1$, and this holds even for the version of rule dependencies by Baget (2004).

Furthermore, to see that $r_3 \not\prec r_1$, assume that r_3 is applicable to an instance I , and that the rule application derives a fact of the form $U(a, f(a))$. Now let $I' = I \cup \{U(a, f(a))\}$, and assume that a substitution σ_2 exists that makes r_1 applicable to I' but not to I ; this rule application must ‘use’ the fact $U(a, f(a))$, which implies that $\sigma_2(x_1) = \sigma_2(x_2) = a$ and $\sigma_2(z) = f(a)$. Furthermore, rule r_1 is applicable only if $R(a, a) \in I$; but then, the rule application does not derive ‘something new’ since $R(x_1, x_2)\sigma_2 = R(a, a)$. Consequently, we have $r_3 \not\prec r_1$; however, unlike in the previous paragraph, this claim depends on the optimisation in Definition 23.

Consider now the chase of I_Σ^* and $\text{MSA}(\Sigma)$ as shown below (facts involving the predicates D , F_{r_2} , and F_{r_3} are omitted for clarity). The chase result contains C , so Σ is not in MSA , and thus $\text{aGRD} \not\subseteq \text{MSA}$; as a corollary, we also get $\text{MSA} \subsetneq \text{MSA}^\prec$.

$R(*, *)$	$T(*, c_{r_2})$	$U(c_{r_2}, c_{r_3})$	$R(*, c_{r_2})$	$T(c_{r_2}, c_{r_2})$	C
$T(*, *)$	$U(*, c_{r_3})$	$S(c_{r_2}, c_{r_3})$		$S(c_{r_2}, c_{r_2})$	
$U(*, *)$	$S(*, c_{r_2})$				
	$S(*, c_{r_3})$				

Note that $R(*, c_{r_2})$ is derived from $R(*, *)$, $U(*, c_{r_3})$, and $U(c_{r_2}, c_{r_3})$, where the latter two facts are obtained from distinct instantiations of $\text{MSA}(r_3)$. Rule dependencies, however, analyse rule applicability w.r.t. $\text{sk}(r_3)$, which is closer to the actual skolem chase. \diamond

In contrast to this result, in Theorem 32 we will show that extending MFA with rule dependencies does not create a stronger notion: MFA^\prec coincides with MFA, which implies that $X^\prec \subseteq \text{MFA}$ for each notion X such that $X \subseteq \text{MFA}$. Towards this goal, we show in Lemma 30 that independent rule sets can be evaluated independently, and in Lemma 31 that a single rule that does not depend on itself can be applied only once.

Lemma 30. *Let Σ_1 and Σ_2 be sets of existential rules such that $\Sigma_2 \not\prec \Sigma_1$, and let F be a set of ground facts not containing a function symbol in $\text{sk}(\Sigma_2)$. Then, $F_{\Sigma_1 \cup \Sigma_2}^\infty = (F_{\Sigma_1}^\infty)_{\Sigma_2}^\infty$.*

Proof. Let $F_0 = F_{\Sigma_1}^\infty$; let F_0, F_1, \dots be the chase sequence for F_0 and Σ_2 where, for convenience, we assume each F_i to be obtained from F_{i-1} by a single rule application (this assumption is clearly w.l.o.g.); and let $F' = (F_0)_{\Sigma_2}^\infty$. By the definition of the skolem chase, we clearly have $F' \subseteq F_{\Sigma_1 \cup \Sigma_2}^\infty$. Furthermore, assume that $F_{\Sigma_1 \cup \Sigma_2}^\infty \not\subseteq F'$; then, a skolemised rule $r_1 \in \text{sk}(\Sigma_1)$ of the form $r_1 = \varphi_1(\vec{x}_1) \rightarrow \psi_1(\vec{x}_1)$ exists such that $F' \subsetneq r_1(F')$. Fix the smallest i such that $F_i \subsetneq r_1(F_i)$ (we clearly have $i > 0$), and let σ_1 be the substitution used in the application of r_1 . Furthermore, let $r_2 \in \text{sk}(\Sigma_2)$ be the skolemised rule of the form

$r_2 = \varphi_2(\vec{x}_2) \rightarrow \psi_2(\vec{x}_2)$ that is used to derive F_i from F_{i-1} , and let σ_2 be the substitution used in the application of r_2 . Now consider an arbitrary term $f(\vec{x}_2)$ in the head of r_2 and assume that $f(\vec{x}_2)\sigma_2$ occurs in F_{i-1} ; since the function symbol f is ‘private’ to r_2 , the head of r_2 must have been already instantiated for σ_2 ; but then, $\psi_2\sigma_2 \subseteq F_{i-1}$, which contradicts our assumption that $\psi_2\sigma_2 \in F_i \setminus F_{i-1}$. Thus, we have the following property (\star):

for each term $f(\vec{x}_2)$ occurring in the head of r_2 , ground term $f(\vec{x}_2)\sigma_2$ does not occur in F_{i-1} .

Finally, let δ be a function that maps each ground term in F_{i-1} to a fresh distinct constant; let $I = \delta(F_{i-1})$; let σ'_2 be the substitution defined by $\sigma'_2(w) = \delta(\sigma_2(w))$ for each variable w in r_2 ; and let σ'_1 be the substitution defined as follows for each variable w in r_1 :

- $\sigma'_1(w) = f(\delta(\vec{t}))$ if $\sigma_1(w) = f(\vec{t})$ for f a function symbol ‘private’ to r_2 ; and
- $\sigma'_1(w) = \delta(\sigma_1(w))$ otherwise.

We clearly have $\varphi_2\sigma'_2 \subseteq I$ and $\psi_2\sigma'_2 \not\subseteq I$; furthermore, by (\star), we also have $\varphi_1\sigma'_1 \subseteq I \cup \psi_2\sigma'_2$ and $\psi_1\sigma'_1 \not\subseteq I \cup \psi_2\sigma'_2$. Moreover, $\varphi_1\sigma'_1 \not\subseteq I$ follows from our assumption that i is the smallest integer such that $F_i \subsetneq r_1(F_i)$. But then, by Definition 23, we have $r_2 \prec r_1$ and, consequently, $\Sigma_2 \prec \Sigma_1$ as well, which is a contradiction. \square

Lemma 31. *Let $\Sigma = \{r\}$ be a singleton rule set such that $r \not\prec r$, and let F be a set of facts not containing a function symbol in $\text{sk}(\Sigma)$. Then, $F_\Sigma^\infty = \Sigma(F)$.*

Proof. Let $F_0 = F$, and let F_0, F_1, \dots be sets of facts such that each F_{i+1} is the union of F_i with the result of a distinct single application of r to F_0 ; clearly, $\bigcup_i F_i = \Sigma(F_0)$. Now assume that $\bigcup_i F_i \subsetneq \Sigma(\bigcup_i F_i)$; then analogously to the proof of Lemma 30, one can show that $r \prec r$, which is a contradiction; we omit the details for the sake of brevity. \square

Theorem 32. *Let Σ be an arbitrary set of rules and let I be an arbitrary instance. If Σ is MFA^\prec w.r.t. I , then Σ is also MFA w.r.t. I .*

Proof. Assume that Σ is in MFA^\prec ; let I be an arbitrary instance; let $\Sigma_1, \dots, \Sigma_n$ be a dependency partition of Σ ; let $\Upsilon_0 = \emptyset$ and $I_0 = I$; and, for each $1 \leq i \leq n$, let $\Upsilon_i = \bigcup_{\ell=1}^i \Sigma_\ell$ and $I_i = (I_{i-1})_{\Sigma_i}^\infty$. By the definition of dependency partitions, we have that $\Sigma_i \not\prec \Upsilon_{i-1}$ holds for each $1 \leq i \leq n$. We next show that, for each $0 \leq i \leq n$, the following two properties hold:

- (a) $I_i = (I_0)_{\Upsilon_i}^\infty$, and
- (b) I_i does not contain a cyclic term.

Set I_0 does not contain functional terms and hence it trivially satisfies (a) and (b). Now consider arbitrary $0 < i < n$ such that I_{i-1} satisfies (a) and (b). By the induction assumption, Lemma 30, $\Sigma_i \not\prec \Upsilon_{i-1}$, and $\Upsilon_i = \Sigma_i \cup \Upsilon_{i-1}$, we have that $(I_0)_{\Upsilon_i}^\infty = ((I_0)_{\Upsilon_{i-1}}^\infty)_{\Sigma_i}^\infty$; thus, I_i satisfies (a). To see that I_i satisfies (b), note that no function symbol used to skolemise the rules in Σ_i is used to skolemise the rules in Υ_{i-1} ; we call this property (\star). Now there are two ways to compute I_i .

- Assume that $\Sigma_i = \{r_i\}$ such that $r_i \not\prec r_i$. By Lemma 31, we have $I_i = r_i(I_{i-1})$; but then, I_i does not contain a cyclic term due to (\star).

- If Σ_i is MFA, then I_i does not contain a cyclic term due to (\star) and Proposition 5.

From the above claim we have that $I_n = I_{\Upsilon_n}^\infty = I_\Sigma^\infty$ and that I_n does not contain a cyclic term; but then, Σ is MFA w.r.t. I by Proposition 5. \square

Combinations of rule dependencies with acyclicity notions have also been considered in databases: Deutsch et al. (2008) proposed a notion of *stratification*, and Meier et al. (2009) further developed this idea and proposed a notion of *c-stratification*. Roughly speaking, each such notion checks whether all strongly connected components of a certain rule dependency graph are WA. The rule dependency notions, however, were developed for the nonoblivious chase and are thus different from Definition 23, as illustrated by the following rule:

$$r = \quad R(z, x) \rightarrow \exists y. R(x, y) \wedge R(y, y) \quad (71)$$

The skolem chase on the critical instance for r is infinite, and $r \prec r$ by Definition 23. In contrast, rule r does not pose problems for the nonoblivious chase. In particular, assume that the rule is matched to an atom $R(t_1, t_2)$, and that it derives $R(t_2, t_3)$ and $R(t_3, t_3)$. Then, rule r is not applicable to $R(t_2, t_3)$ or $R(t_3, t_3)$ since in either case the head atom is satisfied; hence, the rule dependency graphs by Deutsch et al. and Meier et al. are both empty. These results can be summarised as follows: if a rule set Σ satisfies the notion by Deutsch et al., then for each instance I there exists a finite nonoblivious chase sequence; furthermore, if Σ satisfies the notion by Meier et al., then for each instance I all chase sequences (regardless of the rule application strategy) are finite. Meier (2010) discusses in detail the subtle differences between these notions. Since these notions consider a different chase variant, we do not discuss them any further in this paper.

4.3 Acyclicity and Logic Programming

Acyclicity notions have also been considered in the context of disjunctive logic programs with function symbols under the answer set semantics, with the goal of ensuring that a given program has finitely many answer sets, all of which are finite. All of these notions must deal with disjunction and nonmonotonic negation, which is one of the main differences to the notions considered thus far. All notions from logic programming, however, are applicable to rules without disjunction and nonmonotonic negation, in which case they ensure termination of the skolem chase. Therefore, in this section we compare such specialisations of the acyclicity notions from logic programming with aGRD, WA, JA, SWA, MSA, and MFA. We simplify all definitions so that they apply only to skolemised existential rules—that is, we do not present parts of definitions that handle disjunctions in the head and nonmonotonic negation and function symbols in the body.

4.3.1 FINITE DOMAIN NOTION

Calimeri et al. (2008) proposed a finite domain (FD) notion. We next recapitulate this definition, but we do so in the style of Greco et al. (2012), which will come useful in Section 4.3.3 when we introduce Γ -acyclicity. Both approaches use an *argument graph* to determine possible ways for propagating ground terms between positions during chase. The definition of the argument graph is the same as that of the WA dependency graph (see Section 2.4),

but without the distinction between regular and special edges. To simplify the presentation, we consistently use the WA dependency graph instead of the argument graph.

Definition 33. Let Σ be a set of rules. A position $P|_i$ is Σ -recursive with a position $Q|_j$ if the WA dependency graph $WA(\Sigma)$ contains a cycle (consisting of regular and/or special edges) going through $P|_i$ and $Q|_j$. The set $\text{Pos}_{FD}(\Sigma)$ of finite domain positions of Σ is the largest set of positions in Σ such that, for each position $P|_i \in \text{Pos}_{FD}(\Sigma)$, each rule $r \in \Sigma$ of the form $r = \varphi(\vec{x}, \vec{z}) \rightarrow \exists \vec{y}.\psi(\vec{x}, \vec{y})$, and each head atom of r of the form $P(\vec{t})$, the following conditions are satisfied:

- if the i -th component of \vec{t} is a variable $x \in \vec{x}$, then $\text{Pos}_B(x) \cap \text{Pos}_{FD}(\Sigma) \neq \emptyset$; and
- if the i -th component of \vec{t} is a variable $y \in \vec{y}$, then, for each variable $x \in \vec{x}$, some position $Q|_j \in \text{Pos}_B(x) \cap \text{Pos}_{FD}(\Sigma)$ exists that is not Σ -recursive with $P|_i$.

Set Σ is FD if $\text{Pos}_{FD}(\Sigma)$ coincides with the set of all positions in Σ .

Note that the notion of Σ -recursive positions introduced above is symmetric: if $P|_i$ is Σ -recursive with $Q|_j$, then $Q|_j$ is also Σ -recursive with $P|_i$. Furthermore, note that Calimeri et al. (2008) defined FD as follows:

A set of rules Σ is FD if, for each rule $r = \varphi(\vec{x}, \vec{z}) \rightarrow \exists \vec{y}.\psi(\vec{x}, \vec{y})$ in Σ , each atom $Q(\vec{t})$ in the head of r , each j -th term of \vec{t} that is an existential variable y , and each variable $x \in \vec{x}$, there exists a position $P|_i \in \text{Pos}_B(x)$ such that $Q|_j$ is not Σ -recursive with $P|_i$.

Conditions in the above definition clearly correspond to the conditions in Definition 33; but then, since $\text{Pos}_{FD}(\Sigma)$ was defined as the maximal set satisfying these conditions, the two definitions of FD coincide.

We next show that WA is strictly contained in FD. To this end, we first prove that WA is contained in FD, and then we present an example showing that the inclusion is strict.

Proposition 34. If a set of rules Σ is WA, then Σ is FD.

Proof. Let Σ be a set of rules that is not FD. Then, there exist a rule $r \in \Sigma$, an atom $Q(\vec{t})$ in the head of r , a j -th term of \vec{t} equal to an existential variable y , and a variable $x \in \vec{x}$ such that each position $P|_i \in \text{Pos}_B(x)$ is Σ -recursive with $Q|_j$. The set $\text{Pos}_B(x)$ is not empty (\vec{x} contains precisely those variables occurring both in the body and the head of the rule), so choose an arbitrary position $P|_i \in \text{Pos}_B(x)$. The WA dependency graph $WA(\Sigma)$ then contains a special edge from $P|_i$ to $Q|_j$. Furthermore, since $Q|_j$ is Σ -recursive with $P|_i$, graph $WA(\Sigma)$ contains a cycle going through $P|_i$ and $Q|_j$. Thus, $WA(\Sigma)$ clearly contains a cycle containing a special edge, so Σ is not WA. \square

Example 35. Let Σ be the set containing rules (72) and (73).

$$r_1 = R(z, x) \wedge A(x) \rightarrow \exists y.S(x, y) \quad (72)$$

$$r_2 = S(x_1, x_2) \rightarrow R(x_1, x_2) \quad (73)$$

Set Σ is not WA since the WA dependency graph contains a special edge from $R|_2$ to $S|_2$ and a regular edge from $S|_2$ to $R|_2$. However, Σ is FD because position $S|_2$ is not Σ -recursive with $A|_1 \in \text{Pos}_B(x)$. Together with Proposition 34, we can conclude that $\text{WA} \subsetneq \text{FD}$.

In addition, we have $r_1 \prec r_2$ and $r_2 \prec r_1$. In Section 4.3.2 we will prove that $\text{FD} \subseteq \text{JA}$; hence, $\text{FD} \subsetneq \text{FD}^\prec$, $\text{WA}^\prec \subsetneq \text{FD}^\prec$, and $\text{FD} \not\subseteq \text{WA}^\prec$ from Propositions 27, 28, and 34. \diamond

4.3.2 ARGUMENT-RESTRICTED RULE SETS

Lierler and Lifschitz (2009) proposed the notion of argument-restricted rule sets, whose definition we summarise next.

Definition 36. An argument ranking for a set of rules Σ is a function α that assigns a nonnegative integer to each position in Σ such that the following conditions are satisfied for each rule $r \in \Sigma$, each universally quantified variable x in r , and each existentially quantified variable y in r :

1. for each $P|_i \in \text{Pos}_H(x)$, some $Q|_j \in \text{Pos}_B(x)$ exists such that $\alpha(P|_i) \geq \alpha(Q|_j)$; and
2. for each $P|_i \in \text{Pos}_H(y)$, some $Q|_j \in \text{Pos}_B(x)$ exists such that $\alpha(P|_i) > \alpha(Q|_j)$.

Set Σ is argument restricted (AR) if an argument ranking for Σ exists.

An argument-restricted set of rules has a finite skolem chase on an arbitrary instance: by a straightforward induction on the chase sequence, one can show that $\text{dep}(t_i) \leq \alpha(P|_i)$ for each ground fact $P(t_1, \dots, t_n)$ derived by the chase and each $1 \leq i \leq n$.

We next show that JA is strictly more general than AR. Towards this goal, we first prove an auxiliary lemma that establishes a relationship between the set Move from the definition of JA and an argument ranking; next, we use this lemma to prove that $\text{AR} \subseteq \text{JA}$; and finally we present an example that shows this inclusion to be proper.

Lemma 37. Let Σ be a set of rules, let α be an argument ranking for Σ , let y be an existentially quantified variable in Σ , and let $\text{Move}(y)$ be the set of positions used in the definition of JA. For each position $P|_i \in \text{Move}(y)$, some position $Q|_j \in \text{Pos}_H(y)$ exists such that $\alpha(P|_i) \geq \alpha(Q|_j)$ holds.

Proof. Let y be an existentially quantified variable occurring in some rule $r \in \Sigma$, and consider an arbitrary position $P|_i \in \text{Move}(y)$. We prove the claim by induction on the definition of $\text{Move}(y)$. The base case when $P|_i \in \text{Pos}_H(y)$ is trivial. Assume now that $P|_i \in \text{Pos}_H(x)$ for some variable x occurring in a rule $r' \in \Sigma$, and that $\text{Pos}_B(x) \subseteq \text{Move}(y)$, so $P|_i$ needs to be added to $\text{Move}(y)$. By the definition of an argument ranking and since $P|_i \in \text{Pos}_H(x)$, position $P'|_\ell \in \text{Pos}_B(x)$ exists such that $\alpha(P|_i) \geq \alpha(P'|_\ell)$. But then, since $P'|_\ell \in \text{Pos}_B(x) \subseteq \text{Move}(y)$, by the induction hypothesis we have that position $Q|_j \in \text{Pos}_H(y)$ exists such that $\alpha(P'|_\ell) \geq \alpha(Q|_j)$. Thus, $\alpha(P|_i) \geq \alpha(Q|_j)$ holds, as required. \square

Theorem 38. If a set of rules Σ is AR, then Σ is JA.

Proof. Assume that Σ is AR, let α be an argument ranking for Σ , and let $\text{JA}(\Sigma)$ be the JA dependency graph for Σ . We next prove the following claim: for each edge in $\text{JA}(\Sigma)$ from a variable y_1 to a variable y_2 , and for each position $Q|_j \in \text{Pos}_H(y_2)$, there exists a position

$P|_i \in \text{Pos}_H(y_1)$ such that $\alpha(P|_i) < \alpha(Q|_j)$. Consider an arbitrary edge from y_1 to y_2 in $JA(\Sigma)$ and an arbitrary position $Q|_j \in \text{Pos}_H(y_2)$. By the definition of the JA dependency graph, then the rule r that contains y_2 also contains a universally quantified variable x such that x occurs in the head of r and $\text{Pos}_B(x) \subseteq \text{Move}(y_1)$. Since α is an argument ranking for Σ , some position $P'|_\ell \in \text{Pos}_B(x)$ exists such that $\alpha(P'|_\ell) < \alpha(Q|_j)$. Since $P'|_\ell \in \text{Move}(y_1)$, by Lemma 37 position $P|_i \in \text{Pos}_H(y_1)$ exists such that $\alpha(P|_i) \leq \alpha(P'|_\ell)$. Thus, we have $\alpha(P|_i) < \alpha(Q|_j)$, and so our claim holds. But then, this claim clearly implies that the JA dependency graph $JA(\Sigma)$ is acyclic, and therefore Σ is JA. \square

Example 39. Let Σ be the set consisting of the following rules:

$$r_1 = R(z_1, x_1) \rightarrow \exists y_1. S(x_1, y_1) \quad (74)$$

$$r_2 = R(z_2, x_2) \rightarrow \exists y_2. S(y_2, x_2) \quad (75)$$

$$r_3 = S(x_3, x_4) \rightarrow T(x_3, x_4) \quad (76)$$

$$r_4 = T(x_5, x_6) \wedge T(x_6, x_5) \rightarrow R(x_5, x_6) \quad (77)$$

Let α be an argument ranking for Σ . Then, $\alpha(R|_2) < \alpha(S|_2)$ due to (74); $\alpha(R|_2) < \alpha(S|_1)$ due to (75); $\alpha(S|_1) \leq \alpha(T|_1)$ and $\alpha(S|_2) \leq \alpha(T|_2)$ due to (76); and $\alpha(T|_2) \leq \alpha(R|_2)$ or $\alpha(T|_1) \leq \alpha(R|_2)$ due to (77). Together, these observations are contradictory, so such α cannot exist and Σ is not AR. In contrast, $\text{Move}(y_1) = \{S|_2, T|_2\}$ and $\text{Move}(y_2) = \{S|_1, T|_1\}$, and so Σ is JA.

In addition, we have $r_1 \prec r_3$, $r_2 \prec r_3$, $r_3 \prec r_4$, $r_4 \prec r_1$, and $r_4 \prec r_2$; hence, we have $AR \subsetneq AR^\prec$, $AR^\prec \subsetneq JA^\prec$, and $JA \not\subseteq AR^\prec$ from Theorem 38 and Propositions 27 and 28. \diamond

Lierler and Lifschitz (2009, Thm. 4) proved that AR is strictly more general than FD. We next present an example that shows $FD \subsetneq AR$, but that also settles the relationships between FD^\prec and AR^\prec .

Example 40. Let Σ be the set consisting of the following rules:

$$r_1 = A(x) \rightarrow \exists y. R(x, y) \quad (78)$$

$$r_2 = R(x_1, x_2) \rightarrow S(x_1, x_2) \quad (79)$$

$$r_3 = S(z, x) \wedge B(x) \rightarrow A(x) \quad (80)$$

The WA dependency graph for Σ contains a special edge from $A|_1$ to $R|_2$, as well as regular edges from $R|_2$ to $S|_2$ and from $S|_2$ to $A|_1$; thus, $R|_2$ is Σ -recursive with $A|_1$. Consequently, rule (78) cannot satisfy the conditions in Definition 33, so we have $R|_2 \notin \text{Pos}_{FD}(\Sigma)$, and thus Σ is not FD. In contrast, Σ is AR, as evidenced by the following argument ranking:

$$\alpha = \{A|_1 \mapsto 0, B|_1 \mapsto 0, R|_1 \mapsto 0, R|_2 \mapsto 1, S|_1 \mapsto 0, S|_2 \mapsto 1\}$$

In addition, we have $r_1 \prec r_2$, $r_2 \prec r_3$, and $r_3 \prec r_1$; hence, $FD \subsetneq FD^\prec$, $FD^\prec \subsetneq AR^\prec$, and $AR \not\subseteq FD^\prec$ from Propositions 27 and 28. \diamond

Finally, we note that λ -restricted programs by Gebser et al. (2007) and ω -restricted programs by Syrj nen (2001) are both included in FD and AR; thus, when restricted to skolemised existential rules, these notions are also included in JA.

4.3.3 Γ -ACYCLICITY

Greco et al. (2012) recently proposed the notion of Γ -acyclicity for logic programs with function symbols. The original definition of Γ -acyclicity is rather complex, so we next present a simplified version of Γ -acyclicity that is applicable to existential rules. To unify the naming style for the notions in this paper, we often write Γ -acyclicity as ΓA .

Greco et al. (2012) introduce a notion of an *activation graph*, which tracks whether a rule can trigger another rule. This notion is closely related to the notion of rule dependencies from Definition 23, but with the requirement that I is an arbitrary finite set of ground facts (possibly containing functional terms). To understand why the latter is needed in logic programming, consider the following logic program:

$$r_1 = \quad A(x) \wedge B(x) \rightarrow A(f(x)) \quad (81)$$

$$r_2 = \quad A(x) \wedge B(x) \rightarrow B(f(x)) \quad (82)$$

If we restrict the set I in Definition 23 to be an instance, then $r_1 \not\prec r_2$ and $r_2 \not\prec r_1$; however, the skolem chase of r_1 , r_2 , and facts $A(a)$ and $B(a)$ is infinite. Intuitively, r_1 and r_2 contain the same function symbol f , so to determine whether an application of r_1 can trigger an application of r_2 , we must allow the set I in Definition 23 to contain facts such as $B(f(a))$. In our setting, however, function symbols are introduced by skolemisation and are thus ‘private’ to each rule, which allows us to restrict the set I in Definition 23 to facts without functional terms. Thus, in the rest of this section, we simply reuse the rule dependency relation \prec from Definition 23, which gives us a slightly stronger version of ΓA for existential rules than the one proposed by Greco et al. (2012).

Furthermore, Greco et al. (2012) handle logic programming rules with functional terms in the body. Such rules, however, are not considered in this paper, which allows us to omit the definition of a *labelled argument graph* and simplify the notion of a *propagation graph* to a subset of the WA dependency graph.

We are now ready to present a simplified version of Γ -acyclicity that is applicable to existential rules.

Definition 41. *Let Σ be a set of rules. The rule dependency relation \prec is taken from Definition 23, and the set of finite domain positions $\text{Pos}_{FD}(\Sigma)$ is taken from Definition 33.*

The set of safe positions of Σ , written $\text{Pos}_S(\Sigma)$, is the least set of the positions of Σ such that $\text{Pos}_{FD}(\Sigma) \subseteq \text{Pos}_S(\Sigma)$, and $P|_i \in \text{Pos}_S(\Sigma)$ if and only if, for each rule $r \in \Sigma$, at least one of the following conditions is satisfied:

- *if P occurs in the head of r , then \prec does not contain a cycle going through r , or*
- *for each atom $P(\vec{t})$ in the head of $\text{sk}(r)$ and each variable x that occurs in i -th component of \vec{t} , we have $\text{Pos}_B(x) \cap \text{Pos}_S(\Sigma) \neq \emptyset$.*

A position is affected if it is not safe. The propagation graph $PG(\Sigma)$ for Σ has the affected positions of Σ as vertices, and the edges of $PG(\Sigma)$ are defined as in weak acyclicity, but restricted to affected positions. The set Σ is Γ -acyclic (ΓA) if $PG(\Sigma)$ does not contain a cycle that involves a special edge.

In order to relate ΓA to the notions considered thus far, we first establish some containment relationships. It is obvious from Definition 41 that $FD \subseteq \Gamma A$: if all positions in Σ are finite domain, then they are also safe and so the propagation graph is empty. Furthermore, the set of rules in Example 40 is actually ΓA (all positions are safe), but not FD ; hence, by Proposition 27, we have that $FD \subsetneq \Gamma A$, $FD \prec \subsetneq \Gamma A \prec$, and $\Gamma A \not\subseteq FD \prec$. Next, Proposition 42 observes that $aGRD$ is contained in ΓA , and Theorem 43 shows that, perhaps somewhat surprisingly, $\Gamma A \prec$ is contained in $AR \prec$.

Proposition 42. *If a set of rules Σ is aGRD, then Σ is ΓA .*

Proof. If the rule dependency relation \prec on Σ is acyclic, then by the first safety condition in Definition 41 all positions in Σ are safe; but then, $PG(\Sigma)$ is empty, and so Σ is ΓA . \square

Theorem 43. *If a set of rules Σ is $\Gamma A \prec$, then Σ is $AR \prec$.*

Proof. The claim clearly follows from the following property: if the rule dependency relation \prec for Σ has just one strongly connected component and Σ is ΓA , then Σ is AR . Thus, assume that each rule $r \in \Sigma$ occurs on a cycle of \prec . We next construct a mapping α that assigns a nonnegative integer to each position in Σ , and then we show that α is an argument ranking for Σ . In the rest of this proof, we write $p_1 \rightsquigarrow p_2$ if $WA(\Sigma)$ (see Section 2.4) contains a path (consisting of regular and/or special edges) from position p_1 to position p_2 .

Due to our assumption on Σ , the first item in Definition 41 never applies. Furthermore, let Ψ be the function that maps a set S of positions into another set of positions as follows:

$$\Psi(S) = S \cup \{P|_i \mid \text{Pos}_B(x) \cap S \neq \emptyset \text{ for each } r \in \Sigma, \text{ each atom } P(\vec{t}) \text{ in the head of } \text{sk}(r), \text{ and each variable } x \text{ occurring in the } i\text{-th component of } \vec{t}\}$$

Let $\Psi^0(S) = S$, $\Psi^k(S) = \Psi(\Psi^{k-1}(S))$ for each $k > 0$, and $\Psi^\infty(S) = \bigcup \Psi^k(S)$. From Definition 41 it is obvious that $\text{Pos}_S(\Sigma) = \Psi^\infty(\text{Pos}_{FD}(\Sigma))$.

We next define the mapping α . In the rest of this proof, let Y be the set containing each position $p \in \text{Pos}_{FD}(\Sigma)$ for which an existentially quantified variable y in Σ exists such that $p \in \text{Pos}_H(y)$. Furthermore, we use a convention that $\max \emptyset = 0$.

- For each position $p \in \text{Pos}_{FD}(\Sigma)$, we define $\alpha(p)$ as follows:

$$\alpha(p) = \begin{cases} |\{p' \in Y \mid p' \rightsquigarrow p \text{ and } p' \neq p\}| + 1 & \text{if } p \in Y \\ |\{p' \in Y \mid p' \rightsquigarrow p\}| & \text{if } p \notin Y \end{cases}$$

- For each position $p \in \text{Pos}_S(\Sigma) \setminus \text{Pos}_{FD}(\Sigma)$, we define $\alpha(p)$ as follows:

$$\alpha(p) = \left[\min\{k \mid p \in \Psi^k(\text{Pos}_{FD}(\Sigma))\} \right] + [\max\{\alpha(q) \mid q \in \text{Pos}_{FD}(\Sigma)\}]$$

- For each position p in Σ with $p \notin \text{Pos}_S(\Sigma)$, we define $\alpha(p)$ as follows, where $m(p)$ is the maximum number of special edges occurring in $PG(\Sigma)$ on a path ending at p :

$$\alpha(p) = m(p) + 1 + [\max\{\alpha(q) \mid q \in \text{Pos}_S(\Sigma)\}]$$

Since Σ is $\Gamma A \prec$, $PG(\Sigma)$ does not contain a cycle involving a special edge, so $m(p)$ is always a nonnegative integer and $\alpha(p)$ is correctly defined.

We next show that α is an argument ranking—that is, that it satisfies all conditions of Definition 36. To this end, consider an arbitrary rule $r \in \Sigma$, an arbitrary existentially quantified variable y in r , an arbitrary universally quantified variable x in r , and an arbitrary position $P|_i \in \text{Pos}_H(y)$; we have the following cases.

- $P|_i \in \text{Pos}_{FD}(\Sigma)$. By Definition 33, position $Q|_j \in \text{Pos}_B(x) \cap \text{Pos}_{FD}(\Sigma)$ exists that is not Σ -recursive with $P|_i$. Thus, we have $P|_i \not\rightsquigarrow Q|_j$; furthermore, $Q|_j \rightsquigarrow P|_i$ by the definition of $WA(\Sigma)$. Together, the latter two properties imply the following:

$$\{p' \in Y \mid p' \rightsquigarrow Q|_j \text{ and } p' \neq Q|_j\} \subseteq \{p' \in Y \mid p' \rightsquigarrow P|_i \text{ and } p' \neq P|_i\}$$

If $Q|_j \in Y$, this inclusion is strict since $Q|_j$ is contained in the set on the right-hand side, but not in the set on the left-hand side; thus, $\alpha(Q|_j) < \alpha(P|_i)$ holds, as required. If $Q|_j \notin Y$, then $\alpha(Q|_j) < \alpha(P|_i)$ holds since the definition of α ensures that $\alpha(P|_i) - \alpha(Q|_j)$ is at least 1.

- $P|_i \in \text{Pos}_S(\Sigma) \setminus \text{Pos}_{FD}(\Sigma)$. Let k be the smallest number with $P|_i \in \Psi^k(\text{Pos}_{FD}(\Sigma))$. By the definition of Ψ , there exists a position $Q|_j \in \text{Pos}_B(x) \cap \Psi^{k-1}(\text{Pos}_{FD}(\Sigma))$, so $k - 1 \geq \alpha(Q|_j)$ by the definition of α . Thus, $\alpha(P|_i) > \alpha(Q|_j)$ holds, as required.
- $P|_i \notin \text{Pos}_S(\Sigma)$. The first possibility is that some position $Q|_j \in \text{Pos}_B(x) \cap \text{Pos}_S(\Sigma)$ exists; but then, by the definition of α , we have $\alpha(Q|_j) < \alpha(P|_i)$, as required. The second possibility is that there exists some affected position $Q|_j \in \text{Pos}_B(x)$; but then, $Q|_j$ has at least one less incoming special edge in $PG(\Sigma)$ than $P|_i$; thus, we also have $\alpha(Q|_j) < \alpha(P|_i)$, as required.

To complete the proof, we must also consider an arbitrary position $P|_i \in \text{Pos}_H(x)$; however, the cases are analogous as above, so we omit them for the sake of brevity. \square

To place ΓA precisely in the landscape of acyclicity notions, we present three examples that disprove relevant containment relationships. Greco et al. (2012) stated that AR is strictly contained in ΓA , but we were unable to find a formal proof of that statement; in fact, Example 44 shows that this is not the case, and that actually $\Gamma A \prec \subsetneq AR \prec$ holds. Moreover, Example 45 shows that $\Gamma A \not\subseteq MSA$. Finally, Example 46 shows that $WA \prec \not\subseteq \Gamma A$.

Example 44. Let Σ be the set consisting of the following rules:

$$r_1 = A(x) \rightarrow \exists y.R(x, y) \quad (83)$$

$$r_2 = R(x_1, x_2) \rightarrow S(x_1, x_2) \quad (84)$$

$$r_3 = S(z, x) \wedge B(x) \rightarrow A(x) \quad (85)$$

$$r_4 = R(z, x) \rightarrow T(x, x) \quad (86)$$

$$r_5 = T(x, z) \rightarrow R(x, x) \quad (87)$$

$$r_6 = T(z_1, x) \wedge R(z_2, x) \rightarrow \exists y.T(x, y) \quad (88)$$

One can readily verify that the following mapping of positions to nonnegative integers is an argument ranking for Σ :

$$\alpha = \{A|_1 \mapsto 0, B|_1 \mapsto 0, R|_1 \mapsto 1, R|_2 \mapsto 1, S|_1 \mapsto 1, S|_2 \mapsto 1, T|_1 \mapsto 1, T|_2 \mapsto 2\}$$

We next argue that Σ is not ΓA . First, the rule dependency relation in Σ holds (at least) between the pairs of rules shown below. Thus, each rule in Σ occurs in \prec on a cycle, and so Σ is the only strongly connected component of \prec .

$$r_1 \prec r_2 \quad r_2 \prec r_3 \quad r_3 \prec r_1 \quad r_1 \prec r_4 \quad r_4 \prec r_5 \quad r_5 \prec r_2 \quad r_5 \prec r_6 \quad r_6 \prec r_5$$

Second, the WA dependency graph for Σ contains a special edge from $A|_1$ to $R|_2$ due to rule r_1 , a regular edge from $R|_2$ to $S|_2$ due to rule r_2 , and a regular edge from $S|_2$ to $A|_1$ due to rule r_3 ; consequently, $R|_2$ is Σ -recursive with $A|_1$; but then, rule r_1 does not satisfy the conditions in Definition 33, and so $R|_2 \notin \text{Pos}_{FD}(\Sigma)$. Furthermore, due to rule r_4 , we have $T|_1 \notin \text{Pos}_{FD}(\Sigma)$ and $T|_2 \notin \text{Pos}_{FD}(\Sigma)$ as well. Finally, $R|_1 \notin \text{Pos}_{FD}(\Sigma)$ due to rule r_5 . Consequently, the set of finite domain positions is given by $\text{Pos}_{FD}(\Sigma) = \{A|_1, B|_1, S|_1\}$.

Third, we argue that $\text{Pos}_S(\Sigma) = \text{Pos}_{FD}(\Sigma)$. In particular, there is no need to extend $\text{Pos}_S(\Sigma)$ with $R|_2$: position $R|_2$ occurs in the head of rule r_5 , but since $T|_2$ is not a finite domain position and r_5 occurs on a cycle of \prec , neither condition from Definition 41 holds. Analogously, positions $T|_1$ and $T|_2$ do not need to be added to $\text{Pos}_S(\Sigma)$ either.

Fourth, since positions $R|_2$, $T|_1$, and $T|_2$ are all affected, the propagation graph $PG(\Sigma)$ contains a special edge from $T|_2$ to itself due to rule r_6 . Consequently, Σ is not ΓA .

Finally, since Σ is the only strongly connected component of \prec , this example also shows that $AR \not\subseteq \Gamma A^\prec$ and $AR^\prec \not\subseteq \Gamma A^\prec$; but then, by Theorem 43, we have $\Gamma A^\prec \subsetneq AR^\prec$. \diamond

Example 45. Let Σ be the set of rules from Example 29. As explained in the example, Σ is aGRD, but not MSA and thus also not JA, AR, or FD. By Proposition 42, Σ is ΓA , which implies $\Gamma A \not\subseteq \text{MSA}$, and thus $\Gamma A \not\subseteq \text{SWA}$, $\Gamma A \not\subseteq \text{JA}$, $\Gamma A \not\subseteq \text{AR}$, and $\Gamma A \not\subseteq \text{FD}$. \diamond

Example 46. Let Σ be the set consisting of the following rules:

$$r_1 = R(x_1, x_1) \rightarrow \exists y_1 \exists y_2. [A(x_1) \wedge S(y_1, x_1) \wedge S(x_1, y_2)] \quad (89)$$

$$r_2 = A(x_2) \rightarrow B(x_2) \quad (90)$$

$$r_3 = B(x_3) \rightarrow R(x_3, x_3) \quad (91)$$

$$r_4 = S(x_4, x_4) \rightarrow \exists y_3 \exists y_4. [C(x_4) \wedge R(y_3, x_4) \wedge R(x_4, y_4)] \quad (92)$$

$$r_5 = C(x_5) \rightarrow D(x_5) \quad (93)$$

$$r_6 = D(x_6) \rightarrow S(x_6, x_6) \quad (94)$$

Note that $r_1 \not\prec r_4$ and $r_4 \not\prec r_1$, so the rule dependency relation \prec in Σ has two strongly connected components: the first one consists of r_1 , r_2 , and r_3 , and the second one consists of r_4 , r_5 , and r_6 . Moreover, each strongly connected component is WA, so Σ is WA^\prec .

In contrast, each position in Σ is Σ -recursive with itself, so $\text{Pos}_{FD}(\Sigma) = \emptyset$. Moreover, each position in Σ occurs in the head of a rule that (i) appears in \prec in a cycle and (ii) does not satisfy the second safety condition in Definition 41; hence, $\text{Pos}_S(\Sigma) = \emptyset$, and all positions are affected. But then, $PG(\Sigma) = WA(\Sigma)$, and so Σ is not ΓA . \diamond

It may seem counterintuitive that $AR \not\subseteq \Gamma A$, but $\Gamma A^\prec \subsetneq AR^\prec$. Intuitively, the notion of safe positions from Definition 41 uses the rule dependency relation, which allows us to construct an example that is in ΓA but not in AR. In ΓA^\prec , this extra condition is always applied to rules that occur in \prec on a cycle; thus, the notion of safe positions collapses to a notion weaker than AR, which in turn allows ΓA^\prec to be subsumed by AR^\prec .

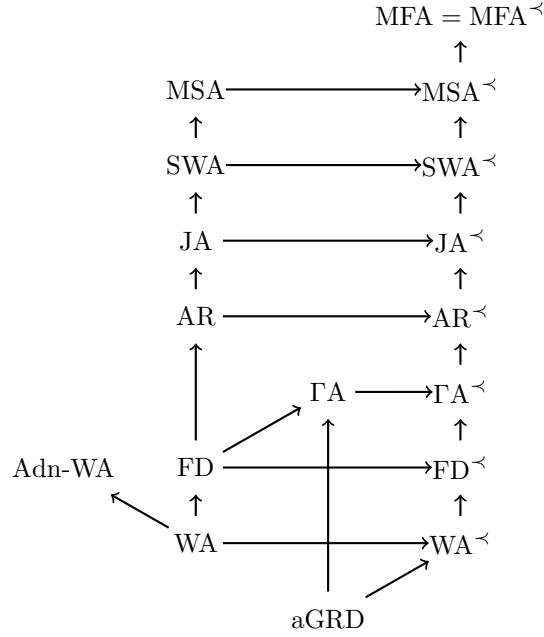


Figure 1: The Landscape of Acyclicity Notions

4.4 The Landscape of Acyclicity Notions

To obtain a complete picture of the relative expressiveness of the acyclicity notions considered in this paper, we make the following observations.

- The rule set from Example 15 is MFA but not MSA, and one can readily verify that $r_1 \prec r_2 \prec r_3 \prec r_4 \prec r_1$; but then, $\text{MSA}^< \subsetneq \text{MFA}^< = \text{MFA}$ by Proposition 27.
- The rule set from Example 20 is SWA but not JA, and one can readily verify that $r_1 \prec r_2 \prec r_3 \prec r_1$; but then, $\text{JA}^< \subsetneq \text{SWA}^<$ and $\text{SWA} \not\subseteq \text{JA}^<$ by Proposition 27.
- The rule set from Example 1 is MSA but not SWA, and one can readily verify that $r_1 \prec r_3 \prec r_4 \prec r_5 \prec r_2 \prec r_1$; but then, $\text{SWA}^< \subsetneq \text{MSA}^<$ and $\text{MSA} \not\subseteq \text{SWA}^<$ by Proposition 27.
- The rule set from Example 22 is aGRD: we have $r_1 \not\prec r_1$, $r_1 \not\prec r_2$, $r_2 \not\prec r_1$, and $r_2 \not\prec r_2$. Thus, $\text{aGRD} \not\subseteq \text{Adn-WA}$.
- The rule set from Example 22 is FD: we can assume all positions in the rule set to be finite domain without violating conditions of Definition 33. Thus, $\text{FD} \not\subseteq \text{Adn-WA}$.

The landscape of the acyclicity notions considered in this paper is shown in Figure 1. All inclusions between notions shown in the figure are strict: if a notion X is reachable from a notion Y via one or more (directed) arcs, then X is strictly more general than Y . Furthermore, all inclusions are also complete: if a notion X is not reachable from a notion Y via one or more (directed) arcs, then X does not contain Y .

5. Handling Equality via Singularisation

Most acyclicity notions presented so far provide no special provision for the equality predicate. If a set of rules Σ contains the equality predicate, one can always axiomatise equality explicitly and then check acyclicity. More precisely, the acyclicity of $\Sigma \cup \Sigma_{\approx}$ (under any notion introduced thus far) guarantees termination of the skolem chase of Σ . Furthermore, note that MFA and MSA are defined as entailment checks in first-order logic with equality, which effectively incorporates the rules of equality into these checks even if rules (1)–(4) are not explicitly given; however, the effect of such a definition is the same.

While handling equality explicitly may be simple, such an approach does not ensure termination of the skolem chase in many practically relevant cases. In particular, the following example shows that the equalities between terms tend to proliferate during skolem chase, which can lead to non-termination.

Example 47. Consider the set of rules Σ containing rules (95)–(96).

$$A(x) \wedge B(x) \rightarrow \exists y.[R(x, y) \wedge B(y)] \quad (95)$$

$$R(z, x_1) \wedge R(z, x_2) \rightarrow x_1 \approx x_2 \quad (96)$$

The skolem chase of I_{Σ}^* and Σ derives the following infinite set of facts:

$$\begin{array}{llll} R(*, f(*)) & B(f(*)) & * \approx f(*) & A(f(*)) \\ R(f(*), f(f(*))) & B(f(f(*))) & \dots & \end{array}$$

Thus, Σ is not universally MFA by Proposition 5, and by Theorem 14 it is not universally MSA either. \diamond

It is worth noticing that in the presence of equality WA is no longer subsumed by MFA and hence both notions become incomparable. As explained in Section 2.4, WA can be applied to rules containing the equality predicate (and without an explicit axiomatisation of equality). Under such a treatment, the rules in Example 47 are WA. This, however, does not contradict the results from Section 4: WA does not require an explicit axiomatisation of equality because it ensures termination of nonoblivious chase—an optimised chase variant that expands existential quantifiers only if necessary and that handles equality by replacing equal terms with canonical representatives. In contrast, the results in Section 4 ensure termination of the skolem chase; since this chase variant uses an explicit axiomatisation of equality, all of our results hold only for equality-free rules (or, equivalently, for the rules containing an explicit axiomatisation of equality). The rules in Example 47 are not WA if equality is axiomatised explicitly, which explains the apparent mismatch with Section 4.

In order to use the skolem chase with rule sets such as the ones in Example 47, Marnette (2009) proposed the *singularisation* technique. Roughly speaking, singularisation replaces the equality predicate \approx with a fresh binary predicate **Eq** to clarify that the two are to be treated differently; furthermore, it axiomatises **Eq** as reflexive, symmetric, and transitive, but it does not introduce replacement rules analogous to (4); finally, it modifies the rules in Σ to take the lack of the replacement rules into account. The chase of the transformed rule set is not a model of Σ , but it can be used to answer queries over Σ in a particular well-defined way. The modification of Σ , however, is nondeterministic: there are many ways to modify Σ and, while some may ensure termination of the skolem chase, not all are required to do so. We next recapitulate the definition of singularisation by Marnette (2009).

Definition 48. A marking M_r of a rule r of the form (5) is a mapping that assigns to each variable $w \in \vec{x} \cup \vec{z}$ a single occurrence of w in φ ; the marked occurrence of w in a rule is written w^\diamond . All other occurrences of w are unmarked, and all occurrences of constants are unmarked as well. For Σ a set of rules, a marking M of Σ contains exactly one marking M_r for each $r \in \Sigma$. Let Eq be a fresh binary predicate not occurring in Σ . The singularisation of Σ under M is the set $\text{Sing}(\Sigma, M)$ that contain rules

$$\rightarrow \text{Eq}(x, x) \quad (97)$$

$$\text{Eq}(x_1, x_2) \rightarrow \text{Eq}(x_2, x_1) \quad (98)$$

$$\text{Eq}(x_1, x_2) \wedge \text{Eq}(x_2, x_3) \rightarrow \text{Eq}(x_1, x_3) \quad (99)$$

and, for each rule $r \in \Sigma$, the rule obtained from r by replacing each atom $s \approx t$ with atom $\text{Eq}(s, t)$, and by replacing each unmarked occurrence of a term t in a body atom with a fresh variable z' and then adding atom $\text{Eq}(t, z')$ to the rule body.

Note that $\text{Sing}(\Sigma, M)$ is unique up to the renaming of the fresh variables. Furthermore, note that rule (97) can be transformed into a safe rule as explained in Section 2.2. Finally, note that $\text{Sing}(\Sigma, M)$ is equality-free (since \approx and Eq are different predicates); therefore, no specific treatment of equality is needed when computing its chase or checking its acyclicity.

Example 49. Singularisation of the marked rule (100) produces rule (101).

$$A(x^\diamond) \wedge B(x) \wedge R(x, z^\diamond) \rightarrow C(x) \quad (100)$$

$$A(x) \wedge B(x_1) \wedge R(x_2, z) \wedge \text{Eq}(x, x_1) \wedge \text{Eq}(x, x_2) \rightarrow C(x) \quad (101)$$

Note that singularisation should be applied ‘globally’ to all rules, including the ones that do not contain the equality predicate. \diamond

The properties of singularisation can be summarised as follows. Let Σ be a set of rules, let I be an instance, and let M be a marking for Σ . Furthermore, let $\Sigma' = \text{Sing}(\Sigma, M)$, and let $I' = I_{\Sigma'}^\infty$ be the chase of I and Σ' . Finally, note that predicate Eq is interpreted in I' as an equivalence relation, so let ρ be a function that maps each term t occurring in I' to an arbitrarily chosen representative from the equivalence class of t . The first-order interpretation $\rho(I')$ is defined as follows, where $\text{rng}(\rho)$ is the range of the mapping ρ , the set $\Delta^{\rho(I')}$ is the universe of $\rho(I')$, and $(P)^{\rho(I')}$ is the interpretation of a predicate P :

$$\begin{aligned} \Delta^{\rho(I')} &= \text{rng}(\rho) \\ (P)^{\rho(I')} &= \{ \langle \rho(t_1), \dots, \rho(t_n) \rangle \mid P(t_1, \dots, t_n) \in I \} \text{ for each } P \text{ different from } \text{Eq} \\ (\text{Eq})^{\rho(I')} &= \{ \langle x, x \rangle \mid x \in \Delta^{\rho(I')} \} \end{aligned}$$

Note that $\rho(I')$ interprets \approx as true equality—that is, each term t is interpreted in $\rho(I')$ as a representative of the equivalence class that contains t ; hence, $\rho(I')$ is not a Herbrand interpretation. Marnette (2010) showed that, for an arbitrary ρ , interpretation $\rho(I')$ is a *universal model* of Σ and I —that is, $\rho(I')$ can be homomorphically embedded into an arbitrary model of Σ and I . Thus, $\rho(I)$ can be used for query answering: for a Boolean conjunctive query Q , we have $I \cup \Sigma \models Q$ if and only if $\rho(I') \models Q$.

This result can be reformulated as follows. Let Σ , I , M , and I' be as specified above, and let us assume that Q is of the form $Q = \exists \vec{y}. \varphi(\vec{y})$. Furthermore, let r be the following rule, and let M' be an arbitrary marking of r :

$$r = \varphi(\vec{y}) \rightarrow H \quad (102)$$

Then, the above characterisation of singularisation implies that

$$\begin{aligned} I \cup \Sigma &\models Q && \text{if and only if} \\ I \cup \text{Sing}(\Sigma \cup \{r\}, M \cup M') &\models H && \text{if and only if} \\ I' \cup \text{Sing}(\{r\}, M') &\models H. \end{aligned}$$

Hence, we can answer Q w.r.t. Σ and I by evaluating $\text{Sing}(\{r\}, M')$ in the chase of I and $\text{Sing}(\Sigma, M)$. It is straightforward to generalise this approach to non-Boolean queries.

The absence of replacement rules (4) often allows the skolem chase to terminate on $\text{Sing}(\Sigma, M)$, but this may depend on the selected marking.

Example 50. Rule (95) from Example 47 admits the following two markings:

$$A(x^\diamond) \wedge B(x) \rightarrow \exists y. [R(x, y) \wedge B(y)] \quad (103)$$

$$A(x) \wedge B(x^\diamond) \rightarrow \exists y. [R(x, y) \wedge B(y)] \quad (104)$$

The skolem chase does not universally terminate for the singularisation obtained from (104) and (96). In contrast, the singularisation obtained from (103) and (96) is JA. \diamond

Definition 51. For $X \in \{\text{MFA}, \text{MSA}, \text{JA}\}$, acyclicity notion X^\exists (resp. X^\forall) contains each finite set of rules Σ such that $\text{Sing}(\Sigma, M) \in X$ for some (resp. each) marking M of Σ .

Clearly, $X^\forall \subseteq X^\exists$ for each $X \in \{\text{MFA}, \text{MSA}, \text{JA}\}$, and Example 50 shows this inclusion to be proper. We next show that JA^\forall actually coincides with WA.

Theorem 52. For Σ an arbitrary finite set of rules, Σ is JA^\forall if and only if Σ is WA.

Proof. ($\text{JA}^\forall \subseteq \text{WA}$) We prove the contrapositive, so let Σ be an arbitrary set of rules that is not WA; w.l.o.g. we assume that each variable in Σ occurs in at most one rule. We consider each edge from p to q in the WA dependency graph $\text{WA}(\Sigma)$ to be a triple $e = \langle p, q, t \rangle$, where $t = \cdot$ if the edge is regular and $t = *$ if the edge is special. By the definition of WA, for each such e , a rule $r \in \Sigma$ and universally quantified variable x occurring in the head and the body of r exist such that $p \in \text{Pos}_B(x)$, so let x_e be one such arbitrarily chosen but fixed variable; furthermore, if edge e is special, then an existentially quantified variable y exists such that $q \in \text{Pos}_H(y)$, so let y_e be one such arbitrarily chosen but fixed variable.

A *cycle* in $\text{WA}(\Sigma)$ is a sequence of edges e_1, \dots, e_n of the form $e_i = \langle p_i, q_i, t_i \rangle$ such that $q_i = p_{i+1}$ for each $1 \leq i < n$ and $q_n = p_1$. Such a cycle is *dangerous* if an edge e_k exists that is special; and such a cycle is *simple* if $x_{e_i} \neq x_{e_j}$ for all $1 \leq i < j \leq n$.³

Now let $\Pi' = e_1, \dots, e_n$ be an arbitrary dangerous cycle in $\text{WA}(\Sigma)$. If Π' is not simple, we show how to transform Π' to a shorter dangerous cycle. Towards this goal, assume that Π' contains edges $e_i = \langle p_i, q_i, t_i \rangle$ and $e_j = \langle p_j, q_j, t_j \rangle$ such that $1 \leq i < j \leq n$ and $x_{e_i} = x_{e_j}$; hence, some rule $r \in \Sigma$ contains body atoms in which x_{e_i} occurs at positions p_i and p_j . Furthermore, let e_k be an arbitrarily chosen, but fixed special edge in Π' ; such e_k exists since Π' is dangerous. We have the following possibilities.

3. Note that a cycle of length one is always simple.

- If $i \leq k < j$, let $\Pi'' = e, e_{i+1}, \dots, e_{j-1}$ where $e = \langle p_j, q_i, t_i \rangle$. If e_i is regular, then x_{e_i} occurs in a head atom of r at position q_i ; furthermore, if e_i is special, then some head atom of r contains an existentially quantified variable at position q_i . Either way, e is an edge of $WA(\Sigma)$, so Π'' is a cycle in $WA(\Sigma)$. Furthermore, e is special if $k = i$, and Π'' contains e_k otherwise; hence, Π'' is dangerous.
- Otherwise, let $\Pi'' = e_1, \dots, e_{i-1}, e, e_{j+1}, \dots, e_n$ where $e = \langle p_i, q_j, t_j \rangle$. Edges e and e_j are of the same type, so e is an edge of $WA(\Sigma)$ and Π'' is a cycle in $WA(\Sigma)$. Furthermore, e is special if $k = j$, and Π'' contains e_k otherwise; hence, Π'' is dangerous.

In both cases, Π'' contains at least one edge less than Π' . Thus, we can iteratively transform an arbitrary dangerous cycle in $WA(\Sigma)$ to a simple dangerous cycle Π .

Now let M be a marking for Σ that marks each variable w occurring in the body of a rule $r \in \Sigma$ as follows.

- If an edge $e = \langle p, q, t \rangle$ in Π exists such that $w = x_e$, then M marks an occurrence of w in r at position p (if there are multiple such occurrences, one is chosen arbitrarily). Since Π is simple, edge e is unique, and so M is correctly defined.
- Otherwise, M marks an arbitrarily chosen occurrence of w in r .

Let $\Sigma' = \text{Sing}(\Sigma, M)$, and let $JA(\Sigma')$ be the JA dependency graph for Σ' . To show that $JA(\Sigma')$ contains a cycle, we first prove the following property (\star).

For each subpath e_1, \dots, e_k of Π where edge e_1 is special and each edge e_i with $1 < i \leq k$ is regular, we have $\{q_i, \text{Eq}|_1\} \subseteq \text{Move}(y_{e_1})$ for each $1 \leq i \leq k$.⁴

Since Σ' contains rule (97), we clearly have $\text{Eq}|_1 \in \text{Move}(y_{e_1})$. We next prove (\star) by induction on k . For the base case $k = 1$, we have $q_1 \in \text{Move}(y_{e_1})$ by the definition of JA. For the induction step, assume that the claim holds for all subpaths of length k , and consider a subpath e_1, \dots, e_k, e_{k+1} . By the induction assumption and the fact that $q_k = p_{k+1}$, we have $p_{k+1} \in \text{Move}(y_{e_1})$. Furthermore, variable $x_{e_{k+1}}$ occurs in the body and the head atom of some rule $r \in \Sigma'$ at positions p_{k+1} and q_{k+1} , respectively. Finally, by the definition of M and the properties of singularisation, we have that $\text{Pos}_B(x_{e_{k+1}})$ contains p_{k+1} and possibly $\text{Eq}|_1$. But then, by the definition of JA, we have $q_{k+1} \in \text{Move}(y_{e_1})$, as required.

To complete the proof, consider now an arbitrary subpath e_1, \dots, e_ℓ of Π where edges e_1 and e_ℓ are special and each edge e_i with $1 < i < \ell$ is regular. By (\star) and the fact that $q_{\ell-1} = p_\ell$, we have $\{p_\ell, \text{Eq}|_1\} \subseteq \text{Move}(y_{e_1})$. Furthermore, as in the previous paragraph, $\text{Pos}_B(x_{e_\ell})$ contains p_ℓ and possibly $\text{Eq}|_1$; but then, $JA(\Sigma')$ contains an edge from y_{e_1} to y_{e_ℓ} . Since Π is a cycle, $JA(\Sigma')$ clearly contains a cycle, so Σ' is not JA, as required.

($JA^\forall \supseteq WA$) Assume that $\Sigma \notin JA^\forall$, so there exists a marking M for Σ such that $\Sigma' = \text{Sing}(\Sigma, M)$ is not JA. We assume that Σ does not contain an existentially quantified variable that occurs in an equality atom; this is w.l.o.g. as we can always replace each equality atom $y \approx t$ with an atom $R(x, t)$ and add a rule $R(x_1, x_2) \rightarrow x_1 \approx x_2$ for R a fresh binary predicate, and such a transformation clearly does not affect the membership of the

4. The notion of a subpath is defined in the obvious way; however, please note that, although Π is defined as a sequence of edges, subpaths of Π can ‘wrap around’ this sequence as Π is a cycle.

rule set in JA^\forall and WA . Now consider an arbitrary existentially quantified variable y , and arbitrary positions $p \in \text{Pos}_H(y)$ and $q \in \text{Move}(y)$ that do not involve Eq (both sets are w.r.t. Σ'); by induction on the construction of $\text{Move}(v)$, one can prove that $\text{WA}(\Sigma)$ then contains a sequence of regular edges from p to q . The proof is straightforward, and we omit the details for the sake of brevity. Similarly, consider an arbitrary edge from y_1 to y_2 in $\text{JA}(\Sigma')$, and arbitrary positions $p \in \text{Pos}_H(y_1)$ and $q \in \text{Pos}_H(y_2)$ that do not involve Eq ; by the definition of JA , a variable x occurring in the rule of y_2 and a position s not involving Eq exist such that $s \in \text{Move}(y_1)$ and $s \in \text{Pos}_B(x)$. But then $\text{WA}(\Sigma)$ contains a path consisting of regular edges from p to s , as well as a special edge from s to q . Since $\text{JA}(\Sigma')$ is cyclic, $\text{WA}(\Sigma)$ clearly contains a cycle involving a special edge. \square

Checking all possible markings may be infeasible: the number of candidates is exponential in the total number of variables that occur more than once in a rule body. Theorem 52 shows that JA^\forall can be decided using WA . For the other cases, the following simple observation shows how to reduce the number of markings.

Definition 53. *A variable x is relevant for a rule $r \in \Sigma$ if x occurs more than once in the body of r , and the head of r contains an atom $P(\vec{t})$ such that $x \in \vec{t}$ and P is not \approx .*

Proposition 54. *Let M and M' be markings for Σ such that, for each rule $r \in \Sigma$, the markings for r in M and M' coincide on each relevant variable in r . Then, for each instance I , the result of the skolem chase for I and $\text{Sing}(\Sigma, M)$ coincides with the result of the skolem chase for I and $\text{Sing}(\Sigma, M')$; furthermore, $\text{Sing}(\Sigma, M)$ is $\text{JA}/\text{MSA}/\text{MFA}$ if and only if $\text{Sing}(\Sigma, M')$ is $\text{JA}/\text{MSA}/\text{MFA}$.*

Proof. Consider an arbitrary rule $r \in \Sigma$. If a variable x occurs only in the body of r , then marking various occurrences of x in r clearly produces rules equivalent up to the renaming of variables. Furthermore, assume that a variable x occurs in the head of r only in an equality atom of the form $x \approx t$, and that the markings of x differ. Then, the rules obtained from r by singularisation will all have the same body (up to the renaming of variables); furthermore, the bodies contain atoms $\text{Eq}(x_i, x)$, and the rule heads are of the form $\text{Eq}(x, t)$. Since $\text{Sing}(\Sigma, M)$ and $\text{Sing}(\Sigma, M')$ contain rules (97)–(99), the skolem chase for I and $\text{Sing}(\Sigma, M)$ clearly derives the same ground atoms as the skolem chase for I and $\text{Sing}(\Sigma, M')$. \square

Despite this optimisation, the number of markings to check can still be exponential in the size of Σ , so we next describe a useful approximation. Let \mathcal{M} be a maximal set of markings for Σ such that, for all $M_1, M_2 \in \mathcal{M}$, each rule $r \in \Sigma$, and each variable x that is not relevant in r , the markings of x in r under M_1 and M_2 coincide. Intuitively, such \mathcal{M} contains all possible markings of the relevant variables, but the markings of all other variables coincide. By Proposition 54 it is clear that, given two such sets \mathcal{M}_1 and \mathcal{M}_2 , the skolem chase of $\bigcup_{M \in \mathcal{M}_1} \text{Sing}(\Sigma, M)$ and $\bigcup_{M \in \mathcal{M}_2} \text{Sing}(\Sigma, M)$ coincides for an arbitrary instance I ; thus, let \mathcal{M} be one arbitrarily chosen such set of markings. Also, let $\text{Sing}_\cup(\Sigma) = \bigcup_{M \in \mathcal{M}} \text{Sing}(\Sigma, M)$, let MFA^\cup be the class containing each rule set Σ such that $\text{Sing}_\cup(\Sigma) \in \text{MFA}$, and let MSA^\cup and JA^\cup be defined analogously. As the following proposition shows, $\text{Sing}_\cup(\Sigma)$ provides a ‘lower bound’ on acyclicity that can be obtained via singularisation.

Proposition 55. *For each $X \in \{MFA, MSA, JA\}$, we have that $X^\cup \subseteq X^\forall$. Furthermore, the size of $\text{Sing}_\cup(\Sigma)$ is exponential in the maximal number of relevant variables in a rule in Σ , and it is linear in the number of rules in Σ .*

Proof. The first claim follows from the fact that all considered notions of acyclicity are monotone in the sense that every subset of an acyclic rule set is also acyclic. The second claim follows from the fact that, if a rule r exists that contains k relevant variables and each variable occurs m times in r , then \mathcal{M} contains m^k different markings for r . \square

This result is interesting when dealing with rules that are obtained from DLs, where each rule has at most one relevant variable: on such rule sets, the size of $\text{Sing}_\cup(\Sigma)$ is linear in the size of Σ . For the general case, the complexity of acyclicity checking does not increase despite the exponential increase in the number of rules.

Theorem 56. *Deciding whether Σ is in MFA^\cup (MFA^\exists , MFA^\forall) is 2EXPTIME-complete. Deciding whether Σ is in MSA^\cup (MSA^\exists , MSA^\forall) is EXPTIME-complete.*

Proof. If Σ contains no equality, it is easy to see that Σ is in MFA^\cup (MFA^\exists , MFA^\forall) if and only if it is in MFA. The same can be observed for MSA. Thus, hardness follows from Theorems 8 and 13.

For membership, we first consider the cases of MFA^\exists , MFA^\forall , MSA^\exists , and MSA^\forall . Each of these properties can be decided by considering all of the at most exponentially many markings. Since $\text{Sing}(\Sigma, M)$ is linear in the size of Σ , the property can be checked for each marking for MFA in 2EXPTIME (cf. Theorem 8) and for MSA in EXPTIME (cf. Theorem 13). This yields the required bound since an exponential factor is not significant for the considered complexity classes.

For MFA^\cup and MSA^\cup , membership follows by observing that the membership of MFA and MSA in 2EXPTIME and EXPTIME, respectively, is obtained from the double/single exponential bound on the number of ground facts that potentially need to be derived in order to decide the required property. While $\text{Sing}_\cup(\Sigma)$ is exponentially larger than Σ , the maximal number of relevant ground facts is still the same since no new predicates or constant symbols are introduced. The increased number of rules leads to an exponential increase of the time to check applicability of all rules in each of the doubly/singly exponentially many steps, but this exponential factor does not affect membership of the decision problem in 2EXPTIME/EXPTIME. \square

We finish this section by examining the interaction between rule normalisation and singularisation. Note that normalisation reduces the number of variables in a rule, which at least at first sight suggests that normalisation could prevent one from finding a marking that ensures acyclicity of the singularised rules. We next show that this cannot happen if normalisation is used without structure sharing: if the original set of rules is MFA w.r.t. some set of markings, then the transformed set of rules is MFA w.r.t. a set of markings as well. Furthermore, we show that this does not hold if normalisation is used with structure sharing; hence, normalisation should be applied with care when used with singularisation.

Theorem 57. *Let Σ be a set of existential rules, let Σ' be obtained from Σ by applying a single normalisation step without structure sharing, and let I be an instance. Then each*

marking M of Σ for which $\text{Sing}(\Sigma, M)$ is MFA w.r.t. I can be extended to a marking M' of Σ' such that $\text{Sing}(\Sigma', M')$ is MFA w.r.t. I .

Proof. Let M be a marking of Σ such that $\text{Sing}(\Sigma, M)$ is MFA, let $r \in \Sigma$ be the rule of the form (8) to which the normalisation step is applied, and let Σ' be the set of rules obtained from Σ after the application of a normalisation step to r . We next prove that the claim holds for both a head and a body normalisation step.

(*Head Normalisation*) Assume that the set of rules Σ' is obtained by replacing a rule $r \in \Sigma$ with rules r_1 and r_2 of the following forms, where $\vec{x} = \vec{x}_3 \cup \vec{x}_4$:

$$\begin{aligned} r &= \varphi(\vec{x}, \vec{z}) \rightarrow \exists \vec{y}_1, \vec{y}_2, \vec{y}_3. [\psi_1(\vec{x}_3, \vec{y}_1, \vec{y}_2) \wedge \psi_2(\vec{x}_4, \vec{y}_1, \vec{y}_3)] \\ r_1 &= \varphi(\vec{x}, \vec{z}) \rightarrow \exists \vec{y}_1, \vec{y}_3. [Q(\vec{x}_3, \vec{y}_1) \wedge \psi_2(\vec{x}_4, \vec{y}_1, \vec{y}_3)] \\ r_2 &= Q(\vec{x}_3, \vec{y}_1) \rightarrow \exists \vec{y}_2. \psi_1(\vec{x}_3, \vec{y}_1, \vec{y}_2) \end{aligned}$$

Let M' be a marking that coincides with M on all rules different from r , that marks r_1 in the same way as M marks r , and that marks r_2 in the only possible way (note that the body of this rule does not contain repeated occurrences of variables); furthermore, let $\Omega = \text{Sing}(\Sigma, M)$ and $\Upsilon = \text{Sing}(\Sigma', M')$. We assume that rule r is skolemised by replacing each variable $y \in \vec{y}_1$ with $g_1^y(\vec{x})$, each variable $y \in \vec{y}_2$ with $g_2^y(\vec{x})$, and each variable $y \in \vec{y}_3$ with $g_3^y(\vec{x})$; rule r_1 is skolemised as r ; and rule r_2 is skolemised by replacing each variable $y \in \vec{y}_2$ with $h^y(\vec{x}_3, \vec{y}_1)$. Thus, the skolemised and singularised rules have the following form; formula φ' is a singularisation of φ , and all freshly introduced variables are contained in \vec{z}_1 :

$$\begin{aligned} \varphi'(\vec{x}, \vec{z}_1) &\rightarrow \psi_1(\vec{x}_3, \vec{g}_1(\vec{x}), \vec{g}_2(\vec{x})) \wedge \psi_2(\vec{x}_4, \vec{g}_1(\vec{x}), \vec{g}_3(\vec{x})) \\ \varphi'(\vec{x}, \vec{z}_1) &\rightarrow Q(\vec{x}_3, \vec{g}_1(\vec{x})) \wedge \psi_2(\vec{x}_4, \vec{g}_1(\vec{x}), \vec{g}_3(\vec{x})) \\ Q(\vec{x}_3, \vec{y}_1) &\rightarrow \psi_1(\vec{x}_3, \vec{y}_1, \vec{h}(\vec{x}_3, \vec{y}_1)) \end{aligned}$$

Finally, we inductively define a partial mapping μ from terms to terms as follows:

- $\mu(c) = c$ for each constant c ,
- $\mu(f(\vec{t})) = f(\mu(\vec{t}))$ for each function symbol f not of the form h^y or g_1^y and all terms \vec{t} such that $\mu(\vec{t})$ is defined, and
- $\mu(h^y(\vec{s}, \vec{g}_1(\vec{s}, \vec{t}))) = g_2^y(\mu(\vec{s}), \mu(\vec{t}))$ for each function symbols of the form h^y , the corresponding symbol g_2^y , and all terms \vec{s} and \vec{t} such that $\mu(\vec{s})$ and $\mu(\vec{t})$ are defined.

We next show the following property (\star): for each $A(\vec{t}) \in I_\Upsilon^\infty$ where A is a predicate occurring in Ω (i.e., A was not introduced by the normalisation step), $\mu(\vec{t})$ is defined and $A(\mu(\vec{t})) \in I_\Omega^\infty$. The proof is by induction on the chase sequence $I_\Upsilon^0, I_\Upsilon^1, \dots$ for I and Υ . The base case holds trivially. Furthermore, since Ω and Υ coincide on all rules apart from r , r_1 , and r_2 , the proof of the claim is trivial for each conclusion of a rule different from r_1 or r_2 . For the remaining cases, we can assume w.l.o.g. that I_Υ^{i+1} is obtained from I_Υ^i by a single application of r_1 of substitution σ and an application of r_2 to the result; thus, the rules together derive the following facts:

$$\begin{aligned} &Q(\vec{x}_3\sigma, \vec{g}_1(\vec{x}\sigma)) \\ &\psi_1(\vec{x}_3\sigma, \vec{g}_1(\vec{x}\sigma), \vec{h}(\vec{x}_3\sigma, \vec{g}_1(\vec{x}\sigma))) \\ &\psi_2(\vec{x}_4\sigma, \vec{g}_1(\vec{x}\sigma), \vec{g}_3(\vec{x}\sigma)) \end{aligned}$$

By the induction assumption, $\varphi'(\mu(\vec{x}\sigma), \mu(\vec{z}\sigma)) \subseteq I_\Omega^\infty$, and so I_Ω^∞ contains the following facts:

$$\begin{aligned} &\psi_1(\mu(\vec{x}_3\sigma), \vec{g}_1(\mu(\vec{x}\sigma)), \vec{g}_2(\mu(\vec{x}\sigma))) \\ &\psi_2(\mu(\vec{x}_4\sigma), \vec{g}_1(\mu(\vec{x}\sigma)), \vec{g}_3(\mu(\vec{x}\sigma))) \end{aligned}$$

Clearly, each term $h^y(\vec{x}_3\sigma, \vec{g}_1(\vec{x}\sigma))$ is of the form $h^y(\vec{s}, \vec{g}_1(\vec{s}, \vec{t}))$, so the mapping μ is defined on the term. Furthermore, by the definition of μ , it is clear that property (\star) holds.

The proof of (\star) also reveals that function symbols h^y occur in I_Γ^∞ always in (sub)terms of the form $h^y(\vec{s}, \vec{g}_1(\vec{s}, \vec{t}))$, and that $\mu(u)$ is defined for each term u occurring in I_Γ^∞ . This observation and the following property of μ clearly imply the claim of this theorem: if u is a cyclic term and $\mu(u)$ is defined, then $\mu(u)$ is cyclic as well. To prove the latter, it suffices to consider the following two cases.

- Assume that u is cyclic due to the repetition of a function symbol f not of the form h^y . Thus, u contains a subterm of the form $f(\vec{s})$, and some $s_i \in \vec{s}$ contains a subterm of the form $f(\vec{t})$. By the definition of μ , then $\mu(u)$ contains a term of the form $f(\mu(\vec{s}))$, and some $s'_i \in \mu(\vec{s})$ contains a subterm of the form $f(\mu(\vec{t}))$. Clearly, $\mu(u)$ is cyclic.
- Assume that u is cyclic due to the repetition of a function symbol of the form h^y . By the above observation, then u contains a subterm of the form $h^y(\vec{s}, \vec{g}_1(\vec{s}, \vec{t}))$, and some $s_i \in \vec{s} \cup \vec{t}$ contains a subterm of the form $h^y(\vec{v}, \vec{g}_1(\vec{v}, \vec{w}))$. By the definition of μ , then $\mu(u)$ contains a subterm of the form $g_2^y(\mu(\vec{s}), \mu(\vec{t}))$, and some $s'_i \in \mu(\vec{s}) \cup \mu(\vec{t})$ contains a subterm of the form $g_2^y(\mu(\vec{v}), \mu(\vec{w}))$. Clearly, u is cyclic.

(*Body Normalisation*) Assume that the set of rules Σ' is obtained by replacing a rule $r \in \Sigma$ with rules r_1 and r_2 of the following forms, where $\vec{x} = \vec{x}_1 \cup \vec{x}_2 \cup \vec{x}_3$, and $\vec{x}_1, \vec{x}_2, \vec{x}_3, \vec{z}_1, \vec{z}_2$, and \vec{z}_3 are all pairwise disjoint:

$$\begin{aligned} r &= \varphi_1(\vec{x}_1, \vec{x}_2, \vec{z}_1, \vec{z}_2) \wedge \varphi_2(\vec{x}_1, \vec{x}_3, \vec{z}_1, \vec{z}_3) \rightarrow \exists \vec{y}. \psi(\vec{x}, \vec{y}) \\ r_1 &= \varphi_1(\vec{x}_1, \vec{x}_2, \vec{z}_1, \vec{z}_2) \rightarrow Q(\vec{x}_1, \vec{x}_2, \vec{z}_1) \\ r_2 &= Q(\vec{x}_1, \vec{x}_2, \vec{z}_1) \wedge \varphi_2(\vec{x}_1, \vec{x}_3, \vec{z}_1, \vec{z}_3) \rightarrow \exists \vec{y}. \psi(\vec{x}, \vec{y}) \end{aligned}$$

For each marked variable v , let \vec{u}_v be the variables used to replace v in singularisation. Then, the singularised rule r can be represented as follows, where for clarity we do not show the free variables of various formulae, φ'_1 and φ'_2 do not contain atoms with predicate **Eq**, and Γ_1 and Γ_2 are the conjunctions of atoms with predicate **Eq** obtained by renaming unmarked occurrences of the variables in $\varphi_1(\vec{x}_1, \vec{x}_2, \vec{z}_1, \vec{z}_2)$ and $\varphi_2(\vec{x}_1, \vec{x}_3, \vec{z}_1, \vec{z}_3)$, respectively:

$$\varphi'_1 \wedge \varphi'_2 \wedge \Gamma_1 \wedge \Gamma_2 \rightarrow \exists \vec{y}. \psi(\vec{x}, \vec{y})$$

Now let M' be a marking that coincides with M on all rules different from r , and that, for each marked occurrence of a variable $w \in \vec{x}_1 \cup \vec{x}_2 \cup \vec{z}_1$ in r , marks r_1 and r_2 as follows.

- If the marked occurrence of w appears in $\varphi_1(\vec{x}_1, \vec{x}_2, \vec{z}_1, \vec{z}_2)$, then the corresponding occurrence of w is marked in r_1 ; in addition, if $w \in \vec{x}_1 \cup \vec{x}_2 \cup \vec{z}_1$, then the occurrence of w in atom $Q(\vec{x}_1, \vec{x}_2, \vec{z}_1)$ is marked in r_2 .

- If the marked occurrence of w appears in $\varphi_2(\vec{x}_1, \vec{x}_3, \vec{z}_1, \vec{z}_3)$, then the corresponding occurrence of w is marked in r_2 ; in addition, if $w \in \vec{x}_1 \cup \vec{x}_2 \cup \vec{z}_1$, then an arbitrary occurrence of w is marked in r_1 .

Since there is no structure sharing, Σ does not contain r_1 , so the above definition is well-formed. The singularisation of r_1 and r_2 under M' can be represented as follows:

$$\begin{aligned}\varphi_1'' \wedge \Gamma_1'' &\rightarrow Q(\vec{x}_1, \vec{x}_2, \vec{z}_1) \\ Q(\vec{x}'_1, \vec{x}_2, \vec{z}'_1) \wedge \varphi'_2 \wedge \Gamma'_2 &\rightarrow \exists \vec{y}. \psi(\vec{x}, \vec{y})\end{aligned}$$

By the definition of M' , it should be clear that $\varphi_1'' \wedge \Gamma_1''$ is isomorphic to a subset of $\varphi'_1 \wedge \Gamma'_1$. Based on this observation, it is now routine to prove that, if $A(\vec{t}) \in I_{\text{Sing}(\Sigma, M)}^\infty$ and A is different from the newly introduced predicate Q , then $A(\vec{t}) \in I_{\text{Sing}(\Sigma', M')}^\infty$, which clearly implies our claim. \square

In contrast to Theorem 57, the following example shows that normalisation with structure sharing can prevent one from finding a marking that makes the normalised rules acyclic. This example shows that normalisation must be used with care in applications that use singularisation to deal with equality.

Example 58. Let Σ be the following set of rules marked by a marking M shown below.

$$A(x) \wedge T(x^\diamond, z) \wedge B(z^\diamond) \rightarrow \exists y. [R(x, y) \wedge A(y)] \quad (105)$$

$$A(x^\diamond) \wedge T(x, z) \wedge C(z^\diamond) \rightarrow \exists y_1 \exists y_2. [S(x, y_1) \wedge T(y_1, y_2)] \quad (106)$$

$$R(z^\diamond, x_1^\diamond) \wedge R(z, x_2^\diamond) \rightarrow x_1 \approx x_2 \quad (107)$$

$$S(z^\diamond, x_1^\diamond) \wedge S(z, x_2^\diamond) \rightarrow x_1 \approx x_2 \quad (108)$$

$$T(z^\diamond, x_1^\diamond) \wedge T(z, x_2^\diamond) \rightarrow x_1 \approx x_2 \quad (109)$$

One can show that $\text{Sing}(\Sigma, M)$ is MFA w.r.t. the instance I given below.

$$I = \{A(a), R(a, a), T(a, b), B(b), A(a'), S(a', a'), T(a', b'), C(b')\}$$

Furthermore, let M_1 be a marking identical to M but which marks $A(x^\diamond)$ in rule (105), and let M_2 be a marking identical to M but which marks $T(x^\diamond, z)$ in rule (106). One can show that neither $\text{Sing}(\Sigma, M_1)$ nor $\text{Sing}(\Sigma, M_2)$ is MFA w.r.t. I .

Now let Σ' be obtained from Σ by applying normalisation with structure sharing to rules (105) and (106); thus, rules (105) and (106) are replaced with the following rules:

$$Q(x, z) \wedge B(z) \rightarrow \exists y. [R(x, y) \wedge A(y)] \quad (110)$$

$$Q(x, z) \wedge C(z) \rightarrow \exists y_1 \exists y_2. [S(x, y_1) \wedge T(y_1, y_2)] \quad (111)$$

$$A(x) \wedge T(x, z) \rightarrow Q(x, z) \quad (112)$$

Note that conjunction $A(x) \wedge T(x, z)$ occurs in Σ' only in rule (112); therefore, variable x in this conjunction can be marked in only one way. This, however, has the same effect as choosing M_1 or M_2 for Σ : no possible marking M' will make $\text{Sing}(\Sigma', M')$ MFA w.r.t. I . Intuitively, normalisation with structure sharing reduces the space of available markings, due to which it may be impossible to find a marking that makes the rules acyclic. \diamond

6. Applying Acyclicity to Horn Description Logics

In this section we apply various acyclicity notions to reasoning problems in *description logic* (DL) ontologies. Description logics are knowledge representation formalisms that underpin the Web Ontology Language (OWL). DL ontologies are constructed from *atomic concepts* (i.e., unary predicates), *atomic roles* (i.e., binary predicates), and *individuals* (i.e., constants). Special atomic concepts \top and \perp denote universal truth and falsehood, respectively. For each atomic role R , expression R^- is an *inverse role*; furthermore, a *role* is an atomic or an inverse role. DLs provide a rich set of *constructors* for building *concepts* (first-order formulae with one free variable) from atomic concepts and roles. A description logic *TBox* is a set of *axioms*, which correspond to first-order sentences. In this paper we consider only *Horn* description logics, in which TBoxes can be translated into existential rules. Furthermore, in this paper we will consider only normalised TBoxes, in which concepts do not occur nested in other concepts. The latter assumption is without loss of generality as each Horn description logic TBox can be normalised in linear time, and the normalised ontology is a model-conservative extension of the original one.

In this paper we consider several logics all of which are fragments of the description logic Horn-*SRQIF*, which provides the formal underpinning for a prominent subset of OWL. A normalised Horn-*SRQIF* TBox \mathcal{T} consists of axioms shown on the left-hand side of Table 1; in the table, A , B , and C are atomic concepts (including possibly \top and \perp), R , S , and T are (not necessarily atomic) roles, and a is an individual. To guarantee decidability of reasoning, \mathcal{T} must satisfy certain *global* conditions (Kutz, Horrocks, & Sattler, 2006), which we omit for the sake of brevity. Roughly speaking, only so-called *simple* roles are allowed to occur in axioms of Type 2, and axioms of Type 6 must be *regular* according to a particular condition that allows such axioms to be represented using a nondeterministic finite automaton. We also consider the following fragments of Horn-*SRQIF*.

- Horn-*SRI* TBoxes are not allowed to contain axioms of Type 2 or 7.
- Horn-*SHIF* TBoxes are not allowed to contain axioms of Type 7, and all axioms of Type 6 satisfy $R = S = T$. Note that all Horn-*SHIF* TBoxes are regular.
- Horn-*SHI* TBoxes inherit the restrictions from Horn-*SHIF* and are further not allowed to contain axioms of type 2.

To simplify the presentation, we do not consider general at-least number restrictions—that is, concepts of the form $\geq n R.A$ with $n > 1$. The translation of such concepts into rules would require an explicit inequality predicate. As explained in Section 2.2, the inequality predicate can be simulated using an ordinary predicate, and so the extension of our results to general at-least number restrictions is straightforward.

In the rest of this paper we allow inverse roles to occur in atoms, so we take an atom of the form $R^-(t_1, t_2)$ with R an atomic role as an abbreviation for $R(t_2, t_1)$. Then, each Horn-*SRQIF* axiom corresponds to an existential rule as shown in Table 1. As explained in Section 2.2, we treat \top and \perp as ordinary unary predicates where \top is explicitly axiomatised. Thus, we can take a substitution θ to be an answer to a CQ $Q(\vec{x})$ w.r.t. a \mathcal{T} and I if $\mathcal{T} \cup I \models Q(\vec{x})\theta$ or $I \cup \mathcal{T} \models \exists y. \perp(y)$; the latter condition takes into account that an unsatisfiable theory entails all possible formulae. Due to this close correspondence between

1.	$A \sqsubseteq \exists R.B$	$A(x) \rightarrow \exists y.[R(x, y) \wedge B(y)]$
2.	$A \sqsubseteq \leq 1 R.B$	$A(z) \wedge R(z, x_1) \wedge B(x_1) \wedge R(z, x_2) \wedge B(x_2) \rightarrow x_1 \approx x_2$
3.	$A \sqcap B \sqsubseteq C$	$A(x) \wedge B(x) \rightarrow C(x)$
4.	$A \sqsubseteq \forall R.B$	$A(z) \wedge R(z, x) \rightarrow B(x)$
5.	$R \sqsubseteq S$	$R(x_1, x_2) \rightarrow S(x_1, x_2)$
6.	$R \circ S \sqsubseteq T$	$R(x_1, z) \wedge S(z, x_2) \rightarrow T(x_1, x_2)$
7.	$A \sqsubseteq \{a\}$	$A(x) \rightarrow x \approx a$

 Table 1: Axioms of normalised Horn-*SRQIF* ontologies and corresponding rules

description logic axioms and existential rules, in the rest of this paper we identify a TBox \mathcal{T} with the corresponding set of rules.

The complexity of answering Boolean conjunctive queries over general (i.e., not acyclic) DL TBoxes is 2EXPTIME- and EXPTIME-complete for Horn-*SRQIF* (Ortiz et al., 2011) and Horn-*SHIF* (Eiter et al., 2008), respectively. In the rest of this section we investigate the complexity of this problem on acyclic ontologies, as well as the complexity of acyclicity checking. In particular, in Section 6.1 we consider the case when the TBox is expressed in Horn-*SRQIF*, for which we show that both BCQ answering and MFA checking are EXPTIME-complete. Then, in Section 6.2 we consider Horn-*SHIF* TBoxes, for which we show that the complexity of these problems drops to PSPACE.

6.1 Acyclic Horn-*SRQIF* TBoxes

We start by showing that BCQ answering for WA Horn-*SRI* TBoxes is EXPTIME-hard. Intuitively, this is due to the axioms of Type 6, which can be used to axiomatise existence of non-tree-like structures. Although regularity ensures that axioms of Type 6 can be represented by a nondeterministic finite automaton, this automaton can be exponential; as a consequence, axioms of Type 6 can axiomatise exponential non-tree-like structures, which is the main source of complexity.

Lemma 59. *Let \mathcal{T} be a WA Horn-*SRI* TBox, let I be an instance, and let F be a fact. Then, checking whether $I \cup \mathcal{T} \models F$ is EXPTIME-hard.*

Proof. Let $\mathcal{M} = (\mathcal{S}, \mathcal{Q}, \delta, Q_0, Q_a)$ be a deterministic Turing machine, where \mathcal{S} is a finite set of symbols, \mathcal{Q} is a finite set of states, $\delta : \mathcal{Q} \times \mathcal{S} \rightarrow \mathcal{Q} \times \mathcal{S} \times \{\leftarrow, \rightarrow\}$ is a transition function, $Q_0 \in \mathcal{Q}$ is the initial state, and Q_a the accepting state. Furthermore, assume that an integer k exists such that \mathcal{M} halts on each input of length n in time 2^{n^k} . Given an arbitrary input S_{i_1}, \dots, S_{i_n} , we construct an MFA set of Horn-*SRI* rules \mathcal{T} and an instance I such that $I \cup \mathcal{T} \models Q_a(a)$ if and only if \mathcal{M} accepts the input. To simplify the presentation, we will use a slightly more general rule syntax than what is allowed by Table 1; however, all such rules can be brought into the required form by renaming parts of the rules with fresh predicates.

Let $\ell = n^k$; since k is a constant, ℓ is polynomial in n . Our construction uses a unary predicate for each symbol and state; for simplicity, we do not distinguish between the predicate and the symbol/state. In addition, the construction also uses binary predicates $L_i, R_i, T_i, U_i, D_i, H_i$, and V_i for $1 \leq i \leq \ell$, unary predicates A_i and B_i for $0 \leq i \leq \ell$, and unary predicates O_1, \dots, O_{n+1}, N_1 , and N_2 . Instance I contains only the fact $A_0(a)$. We

next present the rules of \mathcal{T} . The set \mathcal{T} will contain only Horn rules without empty heads, so it will be satisfiable in a minimal Herbrand model. For readability, we divide \mathcal{T} into groups of rules and prove for each group various facts about this minimal Herbrand model of $\mathcal{T} \cup I$, which is shown schematically in Figure 2. The construction of \mathcal{T} proceeds along the following lines.

- The first, the second, and the third group of rules construct the exponential grid shown at the bottom of Figure 2, whose edges are labelled with H_ℓ and V_ℓ . Each sequence of V_ℓ -edges will be used to encode the contents of the tape of the Turing machine at some point in time; furthermore, precisely one vertex in each such sequence will be labelled with a state, thus representing the position of the head. In contrast, H_ℓ -edges will connect different points in time and will be used to encode the transitions of the Turing machine.
- The fourth group labels the right-most V_ℓ -chain with $S_{i_1}, \dots, S_{i_n}, S_\square, S_\square, \dots, S_\square, S_\square$, where S_\square represents the empty tape symbol.
- The fifth and the sixth groups ensure that the symbols on the tape that are not modified by a move of a Turing machine are propagated between time points.
- The seventh and the eighth group encode the transitions of the Turing machine.
- The ninth group propagates the acceptance condition to the top of the figure by labelling the individual a with the accepting state Q_a .

We next present the rules of \mathcal{T} in detail.

The first group of rules in \mathcal{T} contains rules (113)–(115) for each $0 < i \leq \ell$, and rule (116) for each $1 < i \leq \ell$.

$$A_{i-1}(x) \rightarrow \exists y.[L_i(x, y) \wedge A_i(y)] \quad (113)$$

$$A_{i-1}(x) \rightarrow \exists y.[R_i(x, y) \wedge A_i(y)] \quad (114)$$

$$R_i(z, x) \wedge L_i(z, x') \rightarrow T_i(x, x') \quad (115)$$

$$L_i(z, x) \wedge T_{i-1}(z, z') \wedge R_i(z', x') \rightarrow T_i(x, x') \quad (116)$$

On I , these rules axiomatise existence of a triangular structure in the top part of Figure 2 containing T_i links.

The second group of rules in \mathcal{T} contains rule (117), rules (118)–(120) for each $0 < i \leq \ell$, and rule (121) for each $1 < i \leq \ell$.

$$A_\ell(x) \rightarrow B_0(x) \quad (117)$$

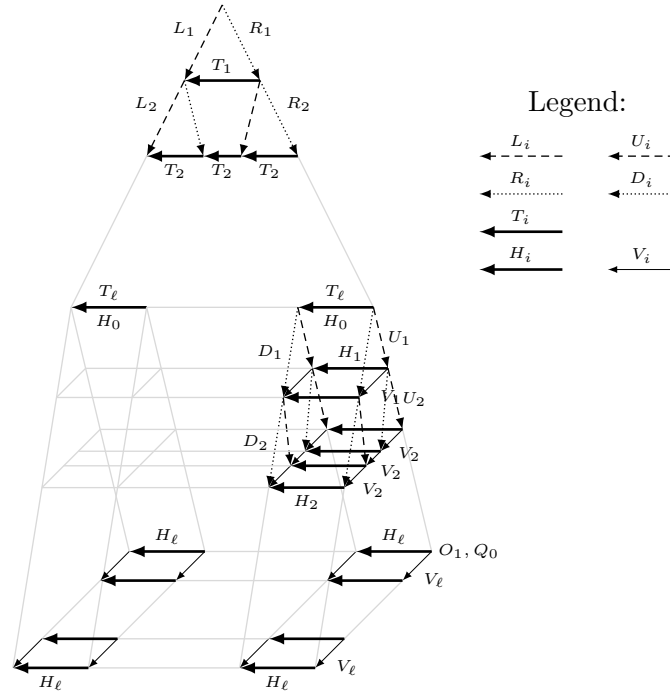
$$B_{i-1}(x) \rightarrow \exists y.[U_i(x, y) \wedge B_i(y)] \quad (118)$$

$$B_{i-1}(x) \rightarrow \exists y.[D_i(x, y) \wedge B_i(y)] \quad (119)$$

$$U_i(z, x) \wedge D_i(z, x') \rightarrow V_i(x, x') \quad (120)$$

$$D_i(z, x) \wedge V_{i-1}(z, z') \wedge U_i(z', x') \rightarrow V_i(x, x') \quad (121)$$

These rules axiomatise existence of triangular structures in the bottom part of Figure 2 containing V_i links.


 Figure 2: Grid Model of \mathcal{T}

The third group of rules in \mathcal{T} contains rule (122), and rules (123) and (124) for each $0 < i \leq \ell$.

$$T_\ell(x, x') \rightarrow H_0(x, x') \quad (122)$$

$$U_i(z, x) \wedge H_{i-1}(z, z') \wedge U_i(z', x') \rightarrow H_i(x, x') \quad (123)$$

$$D_i(z, x) \wedge H_{i-1}(z, z') \wedge D_i(z', x') \rightarrow H_i(x, x') \quad (124)$$

These rules axiomatise existence of H_i links, which with V_i links form a grid of size $2^i \times 2^i$ shown in Figure 2.

In the rest of this proof, for variables w_0 and w_ℓ , we use $R^\ell(w_0, w_\ell)$ as an abbreviation for $R_1(w_0, w_1) \wedge \dots \wedge R_\ell(w_{\ell-1}, w_\ell)$, where each w_i with $0 < i < m$ is a variable not occurring outside the conjunction. Furthermore, we analogously use $U^\ell(w_0, w_\ell)$ as an abbreviation for $U_1(w_0, w_1) \wedge \dots \wedge U_\ell(w_{\ell-1}, w_\ell)$.

The fourth group of rules in \mathcal{T} contains rule (125), rules (126) and (127) for each $1 \leq j \leq n$, and rules (128)–(129), where S_\square is the empty tape symbol. Remember that S_{i_1}, \dots, S_{i_n} encodes the input to \mathcal{M} .

$$A_0(z) \wedge R^\ell(z, z') \wedge U^\ell(z', x) \rightarrow O_1(x) \wedge Q_0(x) \quad (125)$$

$$O_j(z) \wedge V_\ell(z, x) \rightarrow O_{j+1}(x) \quad (126)$$

$$O_j(x) \rightarrow S_{i_j} \quad (127)$$

$$O_{n+1}(z) \wedge V_\ell(z, x) \rightarrow O_{n+1}(x) \quad (128)$$

$$O_{n+1}(x) \rightarrow S_\square(x) \quad (129)$$

Rule (125) labels the grid origin and sets the initial state as shown in Figure 2. Rules (126) ensure that the n subsequent nodes are labelled with O_2, \dots, O_{n+1} , and rule (128) propagates O_{n+1} to the rest of the V_ℓ -chain. Finally, rules (127) and (129) ensure that nodes labelled with O_j are also labelled with S_{i_j} , and that nodes labeled with O_{n+1} are labeled with S_\square . Thus, this group of rules in \mathcal{T} ensures that the right-most V_ℓ -chain in the grid contains the initial state of the tape of \mathcal{M} .

The fifth group of rules in \mathcal{T} contains rules (130)–(131) for each state $Q_k \in \mathcal{Q}$, and rules (132)–(133). These rules essentially ensure that all nodes before and after a node labelled with some state $Q_k \in \mathcal{Q}$ are labeled with N_1 and N_2 , respectively, thus indicating that the head is not above the node.

$$Q_k(z) \wedge V_\ell(x, z) \rightarrow N_1(x) \quad (130)$$

$$Q_k(z) \wedge V_\ell(z, x) \rightarrow N_2(x) \quad (131)$$

$$N_1(z) \wedge V_\ell(x, z) \rightarrow N_1(x) \quad (132)$$

$$N_2(z) \wedge V_\ell(z, x) \rightarrow N_2(x) \quad (133)$$

The sixth group of rules in \mathcal{T} contains rules (134)–(135) instantiated for each symbol $S_k \in \mathcal{S}$; these rules ensure that the contents of the tape is copied between successive time points for all points in the grid not containing the head.

$$N_1(z) \wedge S_k(z) \wedge H_\ell(z, x) \rightarrow S_k(x) \quad (134)$$

$$N_2(z) \wedge S_k(z) \wedge H_\ell(z, x) \rightarrow S_k(x) \quad (135)$$

The seventh group of rules in \mathcal{T} contains rules (136)–(137) instantiated for each symbol $S_k \in \mathcal{S}$ and each state $Q_k \in \mathcal{Q}$ such that $\delta(Q_k, S_k) = (Q_{k'}, S_{k'}, \leftarrow)$. These rules encode moves of \mathcal{M} where the head moves left.

$$Q_k(z) \wedge S_k(z) \wedge H_\ell(z, x) \rightarrow S_{k'}(x) \quad (136)$$

$$Q_k(z) \wedge S_k(z) \wedge H_\ell(z, z') \wedge V_\ell(x, z') \rightarrow Q_{k'}(x) \quad (137)$$

The eighth group of rules in \mathcal{T} contains rules (138)–(139) instantiated for each symbol $S_k \in \mathcal{S}$ and each state $Q_k \in \mathcal{Q}$ such that $\delta(Q_k, S_k) = (Q_{k'}, S_{k'}, \rightarrow)$. These rules encode moves of \mathcal{M} where the head moves right.

$$Q_k(z) \wedge S_k(z) \wedge H_\ell(z, x) \rightarrow S_{k'}(x) \quad (138)$$

$$Q_k(z) \wedge S_k(z) \wedge H_\ell(z, z') \wedge V_\ell(z', x) \rightarrow Q_{k'}(x) \quad (139)$$

The ninth group of rules in \mathcal{T} contains rules (140)–(143) for each $1 \leq i \leq \ell$; these rules simply ensure that acceptance is propagated back to the root of the upper tree.

$$Q_a(z) \wedge U_i(x, z) \rightarrow Q_a(x) \quad (140)$$

$$Q_a(z) \wedge D_i(x, z) \rightarrow Q_a(x) \quad (141)$$

$$Q_a(z) \wedge L_i(x, z) \rightarrow Q_a(x) \quad (142)$$

$$Q_a(z) \wedge R_i(x, z) \rightarrow Q_a(x) \quad (143)$$

The above discussion shows that labelling of the nodes in the grid shown in Figure 2 simulates the execution of \mathcal{M} on input S_{i_1}, \dots, S_{i_n} , where the contents of the tape at some time instant is represented by a V_ℓ -chain, and H_ℓ -links connect tape cells at successive time instants. Thus, $I \cup \mathcal{T} \models Q_a(a)$ if and only if \mathcal{M} accepts S_{i_1}, \dots, S_{i_n} in time 2^ℓ . It is straightforward to see that \mathcal{T} is WA, so the claim of this theorem holds. \square

The proof of Lemma 59 can be adapted to obtain the lower bound for checking MFA of Horn-*SRI* rules.

Lemma 60. *Checking whether a Horn-SRI TBox is universally MFA is EXPTIME-hard.*

Proof. Let \mathcal{M} be an arbitrary deterministic Turing machine and let S_{i_1}, \dots, S_{i_n} be an input string on which \mathcal{M} terminates in time 2^{n^k} . For such \mathcal{M} and S_{i_1}, \dots, S_{i_n} , let \mathcal{T} be as in the proof of Lemma 59. TBox \mathcal{T} is WA, it contains only constant-free, equality-free, and connected rules, and no predicate in \mathcal{T} is of zero arity; hence, by Lemma 7, a Horn-*SRI* TBox \mathcal{T}' exists such that \mathcal{M} accepts S_{i_1}, \dots, S_{i_n} if and only if \mathcal{T}' is not universally MFA. \square

Note that Lemmas 59 and 60 apply to Horn-*SRI* and thus do not rely on a particular treatment of equality. We can deal with the equality predicate in Horn-*SROIF* TBoxes using singularisation as described in Section 5, which leads us to the following result.

Theorem 61. *Let \mathcal{T} be a Horn-SROIF TBox, let M be a marking of \mathcal{T} , let I be an instance, and let Q be a BCQ. Then, checking whether $\text{Sing}(\mathcal{T}, M)$ is MFA w.r.t. I is EXPTIME-complete. Furthermore, if $\text{Sing}(\mathcal{T}, M)$ is MFA w.r.t. I , then checking whether $I \cup \mathcal{T} \models Q$ holds is EXPTIME-complete as well.*

Proof. Note that all rules in Table 1 are \exists -1 rules. Since all rules in $\text{Sing}(\mathcal{T}, M)$ are \exists -1 rules as well, Theorem 10 gives us an EXPTIME upper bound for both of our problems. The matching lower bounds follow from Lemmas 59 and 60 (note that every Horn-*SRI* TBox is also a Horn-*SROIF* TBox) and the fact that their proofs do not use predicate \approx . \square

In fact, Theorem 10 provides us with even stronger complexity bounds. In particular, even if \mathcal{T} does not satisfy all the required global conditions, and even if \mathcal{T} is extended with SWRL rules (Horrocks, Patel-Schneider, Bechhofer, & Tsarkov, 2005), the rules in \mathcal{T} are all still \exists -1 rules. Thus, one can decide whether such \mathcal{T} is MFA (universally or w.r.t. an instance) in EXPTIME, and if that is the case, one can answer BCQs in EXPTIME as well. Consequently, ontology-based applications can freely use the expressivity beyond what is currently available in OWL without an increase in the complexity of reasoning, assuming that the resulting TBox is acyclic.

We conclude this section by observing that MSA provides us with a tractable notion for Horn-*SROIF* rules. Intuitively, all rules in $\text{MSA}(\mathcal{T})$ have a bounded number of variables and all predicates in $\text{MSA}(\mathcal{T})$ are of bounded arity, which eliminates all sources of intractability in datalog reasoning. We prove the matching lower bound in Section 6.2 for the more specific case of Horn-*SHIF* ontologies.

Theorem 62. *Let \mathcal{T} be Horn-SROIF TBox, let M be a marking, and let I be an instance. Then, checking whether $\text{Sing}(\mathcal{T}, M)$ is MSA w.r.t. I is in PTIME.*

Proof. As one can see in Table 1, the rules in \mathcal{T} all contain a bounded number of variables and atoms in the body, and so the number of variables in the body of each rule in $\text{MSA}(\text{Sing}(\mathcal{T}, M))$ is bounded as well. Furthermore, the datalog program $\text{MSA}(\text{Sing}(\mathcal{T}, M))$ contains predicates of bounded arity, so its chase w.r.t. I is polynomial in size. Thus, the chase of I and $\text{MSA}(\mathcal{T})$ can be computed in polynomial time, which implies our claim. \square

6.2 Acyclic Horn- \mathcal{SHIF} TBoxes

The exponential lower bound of Lemmas 59 and 60 critically depend on axioms of Type 6, which can be used to encode exponential structures; furthermore, a combination of inverse roles and axioms of Types 2 and 7 (i.e., of inverse roles, number restrictions, and nominals) is also well known to be problematical (Horrocks & Sattler, 2007). In practice, however, TBoxes are often expressed in Horn- \mathcal{SHIF} , which disallows such axioms in TBoxes. We next show that, for such TBoxes, the complexity of both problems drops to PSPACE.

We first prove PSPACE-hardness for both problems. Note that the PSPACE-hardness proof of concept satisfiability checking by Baader, Calvanese, McGuinness, Nardi, and Patel-Schneider (2007) is not applicable to Horn ontologies since it uses disjunctive concepts. Nonetheless, PSPACE-hardness can be proved by a reduction from checking QBF validity.

Lemma 63. *Let \mathcal{T} be a WA Horn- \mathcal{SHI} TBox, let I be an instance, and let F be a fact. Then, checking whether $I \cup \mathcal{T} \models F$ is PSPACE-hard.*

Proof. Let $\varphi = Q_1x_1 \dots Q_nx_n.C_1 \wedge \dots \wedge C_k$ be an arbitrary quantified Boolean formula defined over variables x_1, \dots, x_n , where each $Q_i \in \{\exists, \forall\}$, $1 \leq i \leq n$ is a quantifier, and each C_j , $1 \leq j \leq k$ is a clause of the form $C_j = L_{j,1} \vee L_{j,2} \vee L_{j,3}$. Checking validity of φ is the canonical PSPACE-hard problem.

In the rest of this proof, for a binary predicate P and variables w_0 and w_m , we use $P^m(w_0, w_m)$ as an abbreviation for $P(w_0, w_1) \wedge \dots \wedge P(w_{m-1}, w_m)$, where each w_i with $0 < i < m$ is a variable not occurring outside the conjunction. Let \mathcal{T} be the Horn- \mathcal{SHI} TBox containing rules (144)–(147) for each $1 \leq i \leq n$, rule (148) for each clause C_j and each literal $L_{j,m} = x_\ell$ occurring in C_j , rule (149) for each clause C_j and each literal $L_{j,m} = \neg x_\ell$ occurring in C_j , rule (150), rule (151) for each $1 \leq i \leq n$ such that $Q_i = \exists$, and rule (152) for each $1 \leq i \leq n$ such that $Q_i = \forall$.

$$A_{i-1}(x) \rightarrow \exists y.[X_i^+(x, y) \wedge A_i(y)] \quad (144)$$

$$A_{i-1}(x) \rightarrow \exists y.[X_i^-(x, y) \wedge A_i(y)] \quad (145)$$

$$X_i^+(x, x') \rightarrow P(x, x') \quad (146)$$

$$X_i^-(x, x') \rightarrow P(x, x') \quad (147)$$

$$X_\ell^+(z', z) \wedge P^{n-\ell}(z, x) \wedge A_n(x) \rightarrow C_j(x) \quad (148)$$

$$X_\ell^-(z', z) \wedge P^{n-\ell}(z, x) \wedge A_n(x) \rightarrow C_j(x) \quad (149)$$

$$C_1(x) \wedge \dots \wedge C_k(x) \rightarrow F_n(x) \quad (150)$$

$$P(x, z) \wedge F_i(z) \rightarrow F_{i-1}(x) \quad (151)$$

$$X_i^+(x, z) \wedge F_i(z) \wedge X_i^-(x, z') \wedge F_i(z') \rightarrow F_{i-1}(x) \quad (152)$$

Strictly speaking, rules (148), (149), (150), and (152) are not Horn- \mathcal{SHI} rules, but they can be transformed into Horn- \mathcal{SHI} rules by replacing parts of their bodies with fresh concepts. It is straightforward to see that \mathcal{T} is WA.

Let $I = \{A_0(a)\}$, and let $I_{\mathcal{T}}^{\infty}$ be the chase of I and \mathcal{T} . Due to rules (144)–(145), $I_{\mathcal{T}}^{\infty}$ contains a binary tree of depth n in which each leaf node is reachable from a via a path that, for each $1 \leq i \leq n$, contains either X_i^+ or X_i^- . If we interpret the presence of X_i^+ and X_i^- as assigning variable x_i to t and f, respectively, then each leaf node corresponds to one possible assignment of x_1, \dots, x_n . Rules (148) and (149) then clearly label each leaf node with the clauses that are true in the node, and rule (150) labels each leaf node with F_n for which all clauses are true. Finally, rules (151) and (152) label each interior node of the tree with F_{i-1} according to the semantics of the appropriate quantifier of φ . Clearly, φ is valid if and only if $I \cup \mathcal{T} \models F_0(a)$, which implies our claim. \square

We next turn our attention to the upper bounds on the complexity of answering a BCQ over an MFA TBox, and checking whether a TBox is MFA. While in Section 6.1 we considered a TBox \mathcal{T} singularised according to some marking M , in this section we assume that equality in \mathcal{T} is handled by means of an explicit axiomatisation \mathcal{T}_{\approx} . As we explain next, this is because singularised rules are not ‘local’, which makes a PSPACE membership proof quite difficult. For example, consider the following singularised rule:

$$A(x) \wedge x \approx x' \wedge B(x') \rightarrow C(x) \quad (153)$$

Atoms $A(x)$ and $B(x')$ in the rule do not share variables and therefore need not be matched ‘locally’ in the chase of $\text{Sing}(\mathcal{T}, M)$ and I ; furthermore, the chase can be exponential in size, so it is not trivial to see how it can be explored using polynomial space. Nevertheless, we conjecture that it is possible to extend our proof to singularised rules as well; however, the details involved seem quite technical, without explaining much about the nature of BCQ answering under equality. Therefore, we leave this problem open and restrict ourselves to the technically simpler case when equality in \mathcal{T} is encoded explicitly using \mathcal{T}_{\approx} .

We next show that answering a BCQ Q over an MFA Horn- \mathcal{SHIF} TBox \mathcal{T} and an instance I can be performed in polynomial space. The proof uses the well-known tracing technique of inspecting a model of $\mathcal{T} \cup I$ using polynomial space. The key aspect of this result, however, is dealing with the transitive roles in the query, which allow the query to be embedded non-locally into the chase of \mathcal{T} and I . Note, however, that we can guess an embedding of Q into the result of $I_{\mathcal{T}}^{\infty}$ using nondeterministic polynomial time; furthermore, since $I_{\mathcal{T}}^{\infty}$ is a minimal Herbrand model of \mathcal{T} (i.e., since \mathcal{T} is Horn), we can check the entailment of each mapped atom of Q separately, and in doing so we can use the well-known encoding by Demri and de Nivelle (2005) to handle transitive roles.

Theorem 64. *Let \mathcal{T} be Horn- \mathcal{SHIF} TBox, let I be an instance such that \mathcal{T} is MFA w.r.t. I , and let Q be a BCQ. Then, checking whether $I \cup \mathcal{T} \models Q$ is PSPACE-complete.*

Proof. Hardness follows from Lemma 63. We next present a nondeterministic polynomial space algorithm that decides $I \cup \mathcal{T} \models Q$; by Savitch’s Theorem, this algorithm can be transformed into a deterministic polynomial space algorithm, which proves our claim.

Assume that BCQ Q is of the form $Q = \exists \vec{y}. B_1 \wedge \dots \wedge B_n$. Furthermore, let $\Upsilon = \text{sk}(\mathcal{T})$. Since \perp is just a regular atomic concept, $I \cup \mathcal{T}$ is always satisfiable in the chase $I_{\mathcal{T}}^{\infty}$ of I

and \mathcal{T} . Furthermore, $I \cup \mathcal{T} \models Q$ if and only if a substitution θ from the variables in \vec{y} to the terms in $I_{\mathcal{T}}^{\infty}$ exists such that $B_i\theta \in I_{\mathcal{T}}^{\infty}$ for each $1 \leq i \leq n$; the latter clearly holds if and only if $I \cup \Upsilon \models B_i\theta$. As shown in the proof of Theorem 10, each term in $I_{\mathcal{T}}^{\infty}$ is of the form $g_1(\dots g_{\ell}(a)\dots)$, where ℓ is less than or equal to the number of function symbols in Υ . Thus, the first step in deciding $I \cup \mathcal{T} \models Q$ is to examine all possible θ and then check $I \cup \Upsilon \models B_i\theta$ for each $1 \leq i \leq n$; this can clearly be done using a deterministic Turing machine that uses polynomial space to store θ , provided that each individual check $I \cup \Upsilon \models B_i\theta$ can also be decided in polynomial space.

If $B_i\theta$ is of the form $C(t)$, then let $\Upsilon' = \Upsilon$, and let $D = C$. Alternatively, if $B_i\theta$ is of the form $R(t', t)$, then let Υ' be Υ extended with the following rules, where D and E are fresh concepts not occurring in Υ and I :

$$\rightarrow E(t') \quad (154)$$

$$E(z) \wedge R(z, x) \rightarrow D(x) \quad (155)$$

It is straightforward to see that $I \cup \Upsilon \models R(t', t)$ if and only if $I \cup \Upsilon' \models D(t)$. Let Υ'' be obtained from Υ' by deleting each rule in Υ' of the form

$$R(x_1, z) \wedge R(z, x_2) \rightarrow R(x_1, x_2) \quad (156)$$

and, for each role R occurring in such a rule, replacing each rule of the form

$$A(z) \wedge R(z, x) \rightarrow B(x) \quad (157)$$

with the following rules, where $Q_{A,R,B}$ is a fresh concept unique for A , R , and B :

$$A(z) \wedge R(z, x) \rightarrow Q_{A,R,B}(x) \quad (158)$$

$$Q_{A,R,B}(z) \wedge R(z, x) \rightarrow Q_{A,R,B}(x) \quad (159)$$

$$Q_{A,R,B}(x) \rightarrow B(x) \quad (160)$$

This transformation corresponds to the well-known elimination of transitivity by Demri and de Nivelle (2005), so $I \cup \Upsilon' \models D(t)$ if and only if $I \cup \Upsilon'' \models D(t)$; the proof of this claim is straightforward and we omit it for the sake of brevity.

Let Ξ be Υ'' extended with the equality axioms (2) and (4). Since \approx does not occur in the body of the rules in Υ'' , we have that $I \cup \Upsilon'' \models D(t)$ if and only if $I \cup \Xi \models_{\approx} D(t)$. Let I_{Ξ}^{∞} be the chase for I and Ξ ; then $I \cup \Xi \models_{\approx} D(t)$ if and only if $D(t) \notin I_{\Xi}^{\infty}$. Note that Ξ contains rules of Types 1–5 from Table 1, rules (2) and (4), and possibly rules of the form $\rightarrow E(t_1)$. These facts can be used to show that each assertion in I_{Ξ}^{∞} is of one of the following forms, where a and b are constants, t is a constant or a term that contains only unary function symbols, f and g are unary function symbols, C is an atomic concept, and R is an atomic role:

- $C(t)$,
- $R(a, b)$, $R(a, f(b))$, $R(f(b), a)$, $R(t, f(t))$, $R(f(t), t)$, or
- $t \approx f(g(t))$, $f(t) \approx g(t)$, $a \approx b$, $a \approx f(b)$, or an equality symmetric to these ones.

The proof is by induction on the length of the chase sequence for I and Ξ , and the claim follows straightforwardly from the I_{Ξ}^{∞} form of rules of Types 1–5. Motik et al. (2009b) prove an analogous claim for a more general description logic, and their proof carries over to the above setting with only syntactic changes.

We say that x is the *central variable* in a rule of Type 1 or 3, and that z is the *central variable* in a rule of Type 2 or 4. W.l.o.g. we assume that the body of a rule of Type 5 does not contain inverse roles; then, x_1 is the *central variable* of a rule of Type 5. Finally, in the equality replacement rules (4), the *central variable* is the variable being replaced.

Clearly, $D(t) \notin I_{\Xi}^{\infty}$ if and only if a Herbrand interpretation J exists in which all assertions are of the form mentioned above, such that $I_{\Xi}^{\infty} \subseteq J$, $J \models_{\approx} \Xi$, and $D(t) \notin J$. We next show how to check the existence of such J using a nondeterministic Turing machine that runs in polynomial space.

Let f_1, \dots, f_m be all function symbols occurring in Ξ . We first guess a Herbrand interpretation J^0 over the constants of I satisfying $I \subseteq J^0$, and we check whether all rules in Ξ not of Type 1 are satisfied in J^0 . If that is the case, we consider each constant c in J_0 and call the following procedure for $s = c$ and $i = 1$:

1. If $i = m + 1$ return *true*.
2. Guess a Herbrand interpretation J^i such that each assertion in J^i is of a form as specified earlier and involves at least one term among $f_1(s), \dots, f_m(s)$.
3. If $D(t) \in J^i$, return *false*.
4. Check whether the equality symmetry rule (4) is satisfied in J^i ; if not, return *false*.
5. Check whether $J^i \cup J^{i-1} \cup \dots \cup J^0$ satisfies each rule in Ξ if the central variable of the rule is mapped to s ; if this is not the case for each rule, return *false*.
6. For each $1 \leq k \leq m$, recursively call this procedure for $f_k(s)$ and $i + 1$; if one of this calls returns *false*, return *false* as well.
7. Return *true*.

Assume that this procedure returns *true* for each constant c , and let J be the union of all J^i considered in the process. It is straightforward to see that $I \subseteq J$ and $D(t) \notin J$; furthermore, $J \models_{\approx} \Xi$ holds since the satisfaction of each rule $r \in \Xi$ in J can be ascertained ‘locally’, by inspecting the vicinity of the ground term that is mapped to the central variable of r . Furthermore, the recursion depth of our algorithm is m and at each recursion level we need to keep a polynomially sized interpretation J^i , so our algorithm can be implemented using a nondeterministic Turing machine that uses polynomial space. \square

Theorem 65. *Let \mathcal{T} be Horn- \mathcal{SHIF} TBox, and let I be an instance. Then, deciding whether \mathcal{T} is MFA w.r.t. I is in PSPACE, and deciding whether \mathcal{T} is universally MFA is PSPACE-hard.*

Proof. (Membership) Rules in $\text{MFA}(\mathcal{T})$ are ‘almost’ Horn- \mathcal{SHIF} rules: rule (19) can be made a Horn- \mathcal{SHIF} rule by replacing S in the body with D (which clearly does not affect the consequences of the rule), and the fact that rule (20) contains a nullary atom in the

head is immaterial. Thus, the claim can be proved by a straightforward adaptation of the membership proof of Theorem 64. The main difference in the algorithm is that, with n function symbols, we need to examine the models to depth $n+1$; however, such an algorithm still uses polynomial space.

(Hardness) Let φ be an arbitrary QBF, and let \mathcal{T} be as in the hardness proof of Lemma 63. TBox \mathcal{T} is WA; it contains only constant-free, equality-free, and connected rules; and it does not contain a predicate of zero arity. Hence, by Lemma 7, a Horn-*SHI* TBox \mathcal{T}' exists such that φ is valid if and only if \mathcal{T}' is not universally MFA. \square

We finish this section by proving that checking whether a set of Horn-*SHI* rules is universally MSA is PTIME-hard; in this way, we also obtain a matching lower bound for theorem Theorem 62 from Section 6.1.

Theorem 66. *Checking whether a Horn-*SHI* TBox \mathcal{T} is universally MSA is PTIME-hard.*

Proof. Let \mathcal{N} be a set of Horn propositional clauses of the form $\neg v_1 \vee \dots \vee \neg v_n \vee v_{n+1}$ and let v be a propositional variable; deciding $\mathcal{N} \models v$ is well known to be PTIME-hard. Let V_i be a concept uniquely associated with each propositional variable v_i ; let A be a fresh concept; and let \mathcal{T} be the TBox obtained by transforming each propositional clause in \mathcal{N} of the above form into rule (161).

$$A(x) \wedge V_1(x) \wedge \dots \wedge V_n(x) \rightarrow V_{n+1}(x) \quad (161)$$

Finally, let $I = \{A(a)\}$. Clearly, $\mathcal{N} \models v$ holds if and only if $I \cup \mathcal{T} \models V(a)$ holds. TBox \mathcal{T} is WA, it contains only constant-free, equality-free, and connected rules, and no predicate in \mathcal{T} is of zero arity; hence, by Lemma 7, a Horn-*SHI* TBox \mathcal{T}' exists such that $\mathcal{N} \not\models v$ holds if and only if \mathcal{T}' is universally MFA. Finally, the only existential variable in \mathcal{T}' occurs in a rule of the form (23), so it is straightforward to see that \mathcal{T}' is universally MFA if and only if \mathcal{T}' is universally MSA. \square

7. Experiments

To estimate the extent to which various acyclicity notions can be used in practice, we conducted two sets of experiments. First, we implemented MFA, MSA, JA, and WA checkers, and we used them to check acyclicity of a large corpus of Horn ontologies. Our goal was to see how many ontologies are acyclic and could thus be used with (suitably extended) materialisation-based OWL reasoners. Second, we computed the materialisation of the acyclic Horn ontologies and compared the number of facts before and after materialisation. The goal of these tests was to see whether materialisation-based reasoning with acyclic ontologies is practically feasible.

Tests were performed on a Windows R2 Server with two Intel Xeon 3.06GHz processors. We used a repository of 336 OWL ontologies whose TBox axioms can be transformed into existential rules where at least one rule contains an existential quantifier in the head. These ontologies include a large subset of the Gardiner ontology corpus (Gardiner, Tsarkov, & Horrocks, 2006), the LUBM ontology, several Phenoscape ontologies, and a number of ontologies from two versions of the Open Biomedical Ontology (OBO) corpus. Please note

that no test ontology has been obtained from conceptual models (e.g., the ER models or UML diagrams): due to the specific modelling patterns used in conceptual modelling, such ontologies are less likely to be acyclic. Each test ontology can be accessed online from our ontology repository by means of a unique ID.⁵ Each ID identifies one self-contained OWL ontology ‘frozen in time’ with all of its imports resolved at the time the ontology was added to the repository; furthermore, any possible future version of the ontology will be assigned a fresh ID. These measures should ensure that our experiments can be independently repeated at any point in the future.

7.1 Acyclicity Tests

We implemented all acyclicity checks by adapting the Hermit reasoner.⁶ Hermit was used to transform an ontology into DL-clauses—formulae quite close to existential rules. In the result, at-least number restrictions in head atoms were replaced with existential quantification, atoms involving datatypes were eliminated, and the DL-clauses with no head atoms were removed: datatypes and empty heads can cause inconsistencies, but they cannot prevent the skolem chase from terminating.

Each set of rules Σ obtained by the above preprocessing steps was considered in combination with each acyclicity notion $X \in \{\text{WA}, \text{JA}, \text{MSA}, \text{MFA}\}$ as follows. If Σ did not contain the equality predicate, we simply checked whether $\Sigma \in X$. If Σ contained the equality predicate, we checked whether $\Sigma \in X^\cup$, and we also checked whether $\Sigma' \in X$ for $\Sigma' \subseteq \Sigma$ the set of all rules of Σ that do not contain the equality predicate; these tests provided us with a ‘lower’ and an ‘upper’ bound for acyclicity, respectively. Each acyclicity test was performed by modifying Σ (or Σ') as required by X and then running Hermit to check for a particular logical entailment on the critical instance.

Our tests revealed MFA and MSA to be indistinguishable for all 336 test ontologies; that is, all MFA ontologies were found to be MSA as well (the converse holds per Theorem 14). A total of 213 (63.4%) ontologies were found to be MSA, including 43 of the 49 (87.8%) ontologies from the Gardiner corpus, 164 of the 208 (78.8%) OBO ontologies, and the LUBM ontology. In contrast, the GALEN ontology and its variants, the GO ontology and its extensions, and the 55 Phenoscope ontologies were found not to be MFA. These results are summarised in Table 2. Given the large number of ontologies tested, it would be impractical to present the results for each ontology individually. Instead, the ontologies are grouped by number of generating rules (G-rules), which are the rules containing an existential quantifier; for each group, Table 2 shows the total number of ontologies, as well as the numbers of ontologies found to be MSA, JA, and WA. Of the 123 ontologies that are not MFA, seven ontologies are in \mathcal{ELH}^r , so CQ answering over these ontologies can be realised using the combined approaches by Lutz et al. (2009) and Kontchakov et al. (2011).

The five older versions of OBO ontologies (IDs 00359, 00374, 00376, 00382, and 00486) are MSA, whereas their newer versions (IDs 00360, 00375, 00377, 00383, and 00487) are not MFA. In contrast, two older versions of OBO ontologies are not MFA (IDs 00432 and 00574), but their newer versions (IDs 00433 and 00575) are MSA.

5. URL <http://www.cs.ox.ac.uk/isg/ontologies/UID/xxxxx.owl> can be used to download an ontology that has been assigned ID ‘xxxxx’.

6. <http://www.hermit-reasoner.com/>

	ontologies without equality				ontologies with equality			
G-rules	Total	MSA	JA	WA	Total	MSA	JA	WA
< 100	44	42	42	41	48	39	39	39
100–1K	69	62	62	41	41	1	0	0
1K–5K	38	31	30	24	17	2	1	1
5K–12K	28	23	16	14	10	5	0	0
12K–160K	18	7	6	5	23	1	0	0

Table 2: Results of acyclicity tests

Finally, we found 15 large OBO ontologies (including different versions of the same ontologies) that are MSA but not JA. Thus, MSA seems to be particularly useful on complex ontologies since it analyses implications between existentially quantified variables more precisely than the previously known notions. Table 3 shows for each of these ontologies the number of generating rules (G-rules), whether the ontology uses the equality predicate (Eq), the ontology expressivity in the description logic family of languages (DL), and the number of classes (C), properties (P), and axioms (A) that the ontology contains. Different versions of the same ontology are distinguished in the table as ‘old’ and ‘new’. Two further ontologies (IDs 00762 and 00766) containing the equality predicate are MSA^\cup , but their status regarding joint acyclicity is unknown: they are JA when the rules involving the equality predicate are deleted, but are not JA^\cup .

7.2 Materialisation Tests

To estimate the practicability of materialisation in acyclic ontologies, we measured the maximal depth of function symbol nesting in terms generated by materialisation on critical instances. This measure, which we call *ontology depth*, is of interest as it provides us with a bound on the size of the materialisation. Out of the 213 MSA ontologies, our test succeeded on 207 of them (tests were aborted if they did not finish in three hours). On the latter ontologies, depth was distributed as follows:

- 123 (59.4%) ontologies have depth less than 5;
- 30 (14.5%) ontologies have depth between 5 and 9;
- 47 (22.7%) ontologies have depth between 10 to 19;
- 5 (2.4%) ontologies have depth between 20 and 49; and
- 2 (1.0%) ontologies have depth between 50 to 70.

These results leads us to believe that many (but clearly not all) ontologies have manageable depths, which should allow for successful materialisation-based query answering.

We also computed the materialisation for several acyclic ontologies. As our implementation is prototypical, our primary goal was not to evaluate the performance of computing the materialisation, but rather to estimate the blowup in the number of facts. We clearly do

Ontology ID	G-rules	Eq	DL	C	P	A
biological_process_xp_cell.imports-local.owl						
00371	7464	yes	\mathcal{SHIF}	17296	178	117925
biological_process_xp_cellular_component.imports-local.owl (old)						
00374	8270	yes	\mathcal{SHIF}	18673	186	126796
biological_process_xp_multi_organism_process.imports-local.owl (old)						
00382	8378	no	\mathcal{EL}^{++}	27900	18	295396
biological_process_xp_plant_anatomy.imports-local.owl (old)						
00386	7559	yes	\mathcal{SHIF}	19146	193	122062
biological_process_xp_plant_anatomy.imports-local.owl (new)						
00387	12025	yes	\mathcal{SRIF}	27412	215	213956
bp_xp_cell.imports-local.owl						
00398	7419	yes	\mathcal{SHIF}	17296	177	117881
bp_xp_cellular_component.imports-local.owl						
00400	7999	yes	\mathcal{SHIF}	18676	175	126540
cellular_component_xp_go.imports-local.owl (old)						
00415	7752	no	\mathcal{EL}^{++}	27890	8	210765
cellular_component_xp_go.imports-local.owl (new)						
00416	12269	no	\mathcal{EL}^{++}	37254	9	334762
fypo.owl						
00476	1834	no	\mathcal{EL}^{++}	1677	22	8027
go_xp_regulation.imports-local.owl (old)						
00486	7777	no	\mathcal{EL}^{++}	27891	5	295138
go_xp_regulation.owl (old)						
00488	7777	no	\mathcal{EL}^{++}	27883	5	214080
go_xp_regulation.owl (new)						
00489	9507	no	\mathcal{EL}^{++}	30170	6	238200
molecular_function_xp_regulators.imports-local.owl (old)						
00536	6762	no	\mathcal{EL}^{++}	25521	5	198170
molecular_function_xp_regulators.imports-local.owl (new)						
00537	11089	no	\mathcal{EL}^{++}	34135	8	316057

Table 3: MSA but not JA ontologies

not expect this blowup to depend linearly on size of the input number of facts; however, our results should provide us with a rough estimate of the performance of materialisation-based reasoning in practice. Most of our test ontologies, however, do not contain many facts: ontologies are often constructed as general vocabularies, while facts are often application-specific and are thus not publicly available. To overcome this problem, we conducted two kinds of experiments.

First, we computed the materialisation of two ontologies that contain facts: LUBM with one university (ID 00347), and the ‘kmi-basic-portal’ ontology (ID 00078). The TBox of LUBM has eight generating rules and depth one, and there are 100,543 facts before ma-

Depth	#	time		gen. size		mat. size	
		max	avg	max	avg	max	avg
< 5	123	38	0.7	65	2	89	8
5–9	30	71	7	122	30	132	38
10–70	54	9396	1807	1286	175	1297	189

Table 4: Materialisation times (in seconds) and sizes

terialisation. Materialisation took only 2 seconds, and it produced 231,200 new facts, of which 97,860 were added by the generating rules. The ‘kmi-basic-portal’ ontology has ten generating rules and depth two, and there were 198 facts before materialisation. Materialisation took only 0.03 seconds, and it produced 744 new facts, of which 145 were added by the generating rules.

Second, for each of the ontologies identified as MSA, we instantiated each class and each property with fresh individuals. We then computed the materialisation and measured the *generated size* (the number of facts introduced by the generating rules divided by the number of facts before materialisation), the *materialisation size* (the number of facts after materialisation divided by the number of facts before materialisation), and the time needed to compute the materialisation. Since most generating rules in these ontologies have singleton body atoms (i.e., they are of the form $A(x) \rightarrow \exists R.C(x)$), these measures should provide a reasonable estimate of the increase in the number of facts during materialisation. Table 4 summarises the results of our tests for the 207 ontologies on which the test succeeded. Ontologies are grouped by their depth, and each group shows the number of ontologies (#), and the maximal and average materialisation time, generated size, and materialisation size.

Thus, materialisation seems practically feasible for many ontologies: for 123 ontologies with depth less than 5, materialisation increases the ontology size by a factor of 8. This suggests that principled, materialisation-based reasoning for ontologies beyond the OWL 2 RL profile may be feasible, especially for ontologies with relatively small depths.

8. Conclusions

In this paper, we investigated acyclicity notions—sufficient conditions that ensure termination for skolem chase on existential rules. We proposed two novel notions, called MFA and MSA, for which we determined tight complexity bounds for membership checking, as well as for conjunctive query answering over acyclic existential rules.

We also conducted a thorough investigation of the acyclicity notions known in the literature, and we produced a complete taxonomy of their relative expressiveness. Our results show that MFA and MSA generalise most of the previously considered notions.

We next investigated ways to ensure acyclicity of existential rules that contain the equality predicate. To this end, we presented several optimisations of the singularisation technique by Marnette (2009). Our optimisations can often reduce the number of acyclicity checks needed, thus making the singularisation technique more suitable for practical use.

Finally, we studied the problem of answering conjunctive queries over acyclic DL ontologies. On the theoretical side, we showed that acyclicity can make this problem computation-

ally easier; furthermore, provided that the result is acyclic, one can extend Horn ontologies with arbitrary SWRL rules without affecting decidability and the worst-case complexity of query answering. On the practical side, we investigated the extent to which acyclicity notions enable principled extensions of materialisation-based ontology reasoners with support for existential quantification. Our tests show that many ontologies commonly used in practice are acyclic, and that the blowup in the number of facts due to materialisation is manageable. This suggests that principled extensions of materialisation-based ontology reasoners are practically feasible and useful.

An interesting topic for future work is to see whether our acyclicity notions can be used in a more general logic programming setting. We see several main sources of technical difficulties towards this goal. First, general logic programs can contain functional terms in body atoms. Such terms can ‘cancel out’ function symbols introduced by head atoms, and it is not clear how to take this into account in an acyclicity test. Second, logic programs can contain atoms under nonmonotonic negation, which are likely to need special treatment; Magka, Krötzsch, and Horrocks (2013) recently made a first step in that direction. Third, it might be desirable to modularise the ways in which these different concerns are handled and thus arbitrarily combine the approaches for handling function symbols in the body and/or the head with the approaches for dealing with nonmonotonic negation.

Acknowledgments

This work was supported by the Royal Society, the Seventh Framework Program (FP7) of the European Commission under Grant Agreement 318338, ‘Optique’, and the EPSRC projects ExODA, Score!, and MaSI³.

References

- Abiteboul, S., Hull, R., & Vianu, V. (1995). *Foundations of Databases*. Addison Wesley.
- Baader, F., Calvanese, D., McGuinness, D., Nardi, D., & Patel-Schneider, P. F. (Eds.). (2007). *The Description Logic Handbook: Theory, Implementation and Applications* (2nd edition). Cambridge University Press.
- Baget, J.-F. (2004). Improving the Forward Chaining Algorithm for Conceptual Graphs Rules. In Dubois, D., Welty, C. A., & Williams, M.-A. (Eds.), *Proc. of the 9th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR2004)*, pp. 407–414, Whistler, BC, Canada. AAAI Press.
- Baget, J.-F., Leclère, M., Mugnier, M.-L., & Salvat, E. (2011a). On rules with existential variables: Walking the decidability line. *Artificial Intelligence*, 175(9–10), 1620–1654.
- Baget, J.-F., Mugnier, M.-L., & Thomazo, M. (2011b). Towards Farsighted Dependencies for Existential Rules. In Rudolph, S., & Gutierrez, C. (Eds.), *Proc. of the 5th Int. Conf. on Web Reasoning and Rule Systems (RR 2011)*, Vol. 6902 of *LNC3*, pp. 30–45, Galway, Ireland. Springer.

- Baumgartner, P., Furbach, U., & Niemelä, I. (1996). Hyper Tableaux. In *Proc. of the European Workshop on Logics in Artificial Intelligence (JELIA '96)*, No. 1126 in LNAI, pp. 1–17, Évora, Portugal. Springer.
- Beeri, C., & Vardi, M. Y. (1981). The Implication Problem for Data Dependencies. In Even, S., & Kariv, O. (Eds.), *Proc. of the 8th Colloquium on Automata, Languages and Programming (ICALP 1981)*, Vol. 115 of *LNCS*, pp. 73–85, Acre (Akko), Israel. Springer.
- Bishop, B., & Bojanov, S. (2011). Implementing OWL 2 RL and OWL 2 QL rule-sets for OWLIM. In Dumontier, M., & Courtot, M. (Eds.), *Proc. of the OWL: Experiences and Directions Workshop (OWLED 2011)*, Vol. 796 of *CEUR WS Proceedings*, San Francisco, UCA, USA.
- Broekstra, J., Kampman, A., & van Harmelen, F. (2002). Sesame: A Generic Architecture for Storing and Querying RDF and RDF Schema. In Horrocks, I., & Hendler, J. A. (Eds.), *Proc. of the 1st Int. Semantic Web Conf. (ISWC 2002)*, Vol. 2342 of *LNCS*, pp. 54–68, Sardinia, Italy. Springer.
- Calì, A., Gottlob, G., Lukasiewicz, T., Marnette, B., & Pieris, A. (2010a). Datalog[±]: A Family of Logical Knowledge Representation and Query Languages for New Applications. In *Proc. of the 25th IEEE Symposium on Logic in Computer Science (LICS 2010)*, pp. 228–242, Edinburgh, United Kingdom. IEEE Computer Society.
- Calì, A., Gottlob, G., & Pieris, A. (2010b). Query Answering under Non-guarded Rules in Datalog[±]. In Hitzler, P., & Lukasiewicz, T. (Eds.), *Proc. of the 4th Int. Conf. on Web Reasoning and Rule Systems (RR 2010)*, Vol. 6333 of *LNCS*, pp. 1–17, Bressanone/Brixen, Italy. Springer.
- Calì, A., Gottlob, G., & Pieris, A. (2011). New Expressive Languages for Ontological Query Answering. In Burgard, W., & Roth, D. (Eds.), *Proc. of the 25th National Conference on Artificial Intelligence (AAAI 2011)*, pp. 1541–1546, San Francisco, CA, USA. AAAI Press.
- Calimeri, F., Cozza, S., Ianni, G., & Leone, N. (2008). Computable Functions in ASP: Theory and Implementation. In Garcia de la Banda, M., & Pontelli, E. (Eds.), *Proc. of the 24th Int. Conf. on Logic Programming (ICLP 2008)*, Vol. 5366 of *LNCS*, pp. 407–424, Udine, Italy. Springer.
- Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., & Rosati, R. (2007). Tractable Reasoning and Efficient Query Answering in Description Logics: The DL-Lite Family. *Journal of Automated Reasoning*, 9(3), 385–429.
- Carroll, J. J., Dickinson, I., Dollin, C., Reynolds, D., Seaborne, A., & Wilkinson, K. (2004). Jena: Implementing the Semantic Web Recommendations. In Feldman, S. I., Uretsky, M., Najork, M., & Wills, C. E. (Eds.), *Proc. of the 13th Int. Conf. on World Wide Web (WWW 2004)—Alternate Track*, pp. 74–83, New York, NY, USA. ACM.
- Cuenca Grau, B., Horrocks, I., Krötzsch, M., Kupke, C., Magka, D., Motik, B., & Wang, Z. (2012). Acyclicity Conditions and their Application to Query Answering in Description Logics. In Brewka, G., Eiter, T., & McIlraith, S. A. (Eds.), *Proc. of the 13th Int.*

- Conf. on the Principles of Knowledge Representation and Reasoning (KR 2012)*, pp. 243–253, Rome, Italy.
- Cuenca Grau, B., Horrocks, I., Motik, B., Parsia, B., Patel-Schneider, P., & Sattler, U. (2008). OWL 2: The next step for OWL. *Journal of Web Semantics: Science, Services and Agents on the World Wide Web*, 6(4), 309–322.
- Dantsin, E., Eiter, T., Gottlob, G., & Voronkov, A. (2001). Complexity and expressive power of logic programming. *ACM Computing Surveys*, 33(3), 374–425.
- De Schreye, D., & Decorte, S. (1994). Termination of Logic Programs: The Never-Ending Story. *Journal of Logic Programming*, 19–20, 199–260.
- Demri, S., & de Nivelle, H. (2005). Deciding Regular Grammar Logics with Converse Through First-Order Logic. *Journal of Logic, Language and Information*, 14(3), 289–329.
- Deutsch, A., Nash, A., & Rummel, J. B. (2008). The chase revisited. In Lenzerini, M., & Lembo, D. (Eds.), *Proc. of the 27th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS 2008)*, pp. 149–158, Vancouver, BC, Canada.
- Eiter, T., Gottlob, G., Ortiz, M., & Simkus, M. (2008). Query Answering in the Description Logic Horn-*SHIQ*. In Hölldobler, S., Lutz, C., & Wansing, H. (Eds.), *Proc. of the 11th European Conference on Logics in Artificial Intelligence (JELIA 2008)*, Vol. 5293 of *LNCIS*, pp. 166–179, Dresden, Germany. Springer.
- Fagin, R., Kolaitis, P. G., Miller, R. J., & Popa, L. (2005). Data exchange: semantics and query answering. *Theoretical Computer Science*, 336(1), 89–124.
- Gardiner, T., Tsarkov, D., & Horrocks, I. (2006). Framework for an Automated Comparison of Description Logic Reasoners. In Cruz, I. F., Decker, S., Allemang, D., Preist, C., Schwabe, D., Mika, P., Uschold, M., & Aroyo, L. (Eds.), *Proc. of the 5th Int. Semantic Web Conference (ISWC 2006)*, Vol. 4273 of *LNCIS*, pp. 654–667, Athens, GA, USA. Springer.
- Gebser, M., Schaub, T., & Thiele, S. (2007). Gringo: A New Grounder for Answer Set Programming. In Baral, C., Brewka, G., & Schlipf, J. S. (Eds.), *Proc. of the 9th Int. Conf. on Logic Programming and Nonmonotonic Reasoning (LPNMR 2007)*, Vol. 4483 of *LNCIS*, pp. 266–271, Tempe, AZ, USA.
- Glimm, B., Horrocks, I., Lutz, C., & Sattler, U. (2008). Conjunctive Query Answering for the Description Logic *SHIQ*. *Journal of Artificial Intelligence Research*, 31, 151–198.
- Greco, S., Spezzano, F., & Trubitsyna, I. (2012). On the Termination of Logic Programs with Function Symbols. In Dovier, A., & Santos Costa, V. (Eds.), *Proc. of the 8th Int. Conf. on Logic Programming (ICLP 2012)*, Vol. 17 of *Leibniz International Proceedings in Informatics*, pp. 323–333, Budapest, Hungary.
- Hastings, J., Magka, D., Batchelor, C., Duan, L., Stevens, R., Ennis, M., & Steinbeck, C. (2012). Structure-based classification and ontology in chemistry. *Journal of Cheminformatics*, 4(8).
- Horrocks, I., & Patel-Schneider, P. F. (2004). A Proposal for an OWL Rules Language. In *Proc. of the 13th Int. World Wide Web Conference (WWW 2004)*, pp. 723–731, New York, NY, USA. ACM Press.

- Horrocks, I., & Sattler, U. (2007). A Tableau Decision Procedure for *SHOIQ*. *Journal of Automated Reasoning*, 39(3), 249–276.
- Horrocks, I., Patel-Schneider, P. F., Bechhofer, S., & Tsarkov, D. (2005). Owl rules: A proposal and prototype implementation. *J. Web Sem.*, 3(1), 23–40.
- Hustadt, U., Motik, B., & Sattler, U. (2005). Data Complexity of Reasoning in Very Expressive Description Logics. In *Proc. of the 19th Int. Joint Conf. on Artificial Intelligence (IJCAI 2005)*, pp. 466–471, Edinburgh, UK. Morgan Kaufmann Publishers.
- Johnson, D. S., & Klug, A. C. (1984). Testing Containment of Conjunctive Queries under Functional and Inclusion Dependencies. *Journal of Computer and System Sciences*, 28(1), 167–189.
- Kiryakov, A., Ognyanov, D., & Manov, D. (2005). OWLIM – A Pragmatic Semantic Repository for OWL. In Dean, M., Guo, Y., Jun, W., Kaschek, R., Krishnaswamy, S., Pan, Z., & Sheng, Q. Z. (Eds.), *Proc. of the Int. Workshop on Web Information Systems Engineering (WISE 2005)*, Vol. 3807 of *LNCIS*, pp. 182–192, New York, NY, USA.
- Kontchakov, R., Lutz, C., Toman, D., Wolter, F., & Zakharyashev, M. (2011). The Combined Approach to Ontology-Based Data Access. In Walsh, T. (Ed.), *Proc. of the 22nd Int. Joint Conf. on Artificial Intelligence (IJCAI 2011)*, pp. 2656–2661, Barcelona, Spain.
- Krötzsch, M., & Rudolph, S. (2011). Extending Decidable Existential Rules by Joining Acyclicity and Guardedness. In Walsh, T. (Ed.), *Proc. of the 22nd Int. Joint Conf. on Artificial Intelligence (IJCAI 2011)*, pp. 963–968, Barcelona, Spain.
- Krötzsch, M., & Rudolph, S. (2013). On the Relationship of Joint Acyclicity and Super-Weak Acyclicity. Tech. rep. 3037, Institute AIFB, Karlsruhe Institute of Technology. Available online at <http://www.aifb.kit.edu/web/Techreport3013>.
- Krötzsch, M., Rudolph, S., & Hitzler, P. (2007). Conjunctive Queries for a Tractable Fragment of OWL 1.1. In Aberer, K., Choi, K.-S., Noy, N. F., Allemang, D., Lee, K.-I., Nixon, L. J. B., Golbeck, J., Mika, P., Maynard, D., Mizoguchi, R., Schreiber, G., & Cudré-Mauroux, P. (Eds.), *Proc. of the 6th Int. Semantic Web Conference (ISWC 2007)*, Vol. 4825 of *LNCIS*, pp. 310–323, Busan, Korea. Springer.
- Kutz, O., Horrocks, I., & Sattler, U. (2006). The Even More Irresistible *SHOIQ*. In Doherty, P., Mylopoulos, J., & Welty, C. A. (Eds.), *Proc. of the 10th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2006)*, pp. 68–78, Lake District, UK. AAAI Press.
- Leone, N., Pfeifer, G., Faber, W., Eiter, T., Gottlob, G., Perri, S., & Scarcello, F. (2006). The DLV system for knowledge representation and reasoning. *ACM Transactions on Computational Logic*, 7(3), 499–562.
- Lierler, Y., & Lifschitz, V. (2009). One More Decidable Class of Finitely Ground Programs. In Hill, P. M., & Warren, D. S. (Eds.), *Proc. of the 25th Int. Conf. on Logic Programming (ICLP 2009)*, Vol. 5649 of *LNCIS*, pp. 489–493, Pasadena, CA, USA. Springer.
- Lutz, C., Toman, D., & Wolter, F. (2009). Conjunctive Query Answering in the Description Logic \mathcal{EL} Using a Relational Database System. In Boutilier, C. (Ed.), *Proc. of the 21st*

- Int. Joint Conf. on Artificial Intelligence (IJCAI 2009)*, pp. 2070–2075, Pasadena, CA, USA.
- Magka, D., Krötzsch, M., & Horrocks, I. (2013). Computing Stable Models for Nonmonotonic Existential Rules. In *Proc. of the 23rd Int. Joint Conf. on Artificial Intelligence (IJCAI 2013)*. AAAI Press/IJCAI. To appear.
- Magka, D., Motik, B., & Horrocks, I. (2012). Modelling Structured Domains Using Description Graphs and Logic Programming. In Simperl, E., Cimiano, P., Polleres, A., Corcho, Ó., & Presutti, V. (Eds.), *Proc. of the 9th Extended Semantic Web Conference (ESWC 2012)*, Vol. 7295 of *LNCS*, pp. 330–344, Heraklion, Greece. Springer.
- Maier, D., Mendelzon, A. O., & Sagiv, Y. (1979). Testing Implications of Data Dependencies. *ACM Transactions on Database Systems*, 4(4), 455–469.
- Marnette, B. (2009). Generalized schema-mappings: from termination to tractability. In Paredaens, J., & Su, J. (Eds.), *Proc. of the 28th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS 2009)*, pp. 13–22, Providence, RI, USA.
- Marnette, B. (2010). *Tractable Schema Mappings Under Oblivious Termination*. Ph.D. thesis, University of Oxford, Oxford, UK.
- Meditskos, G., & Bassiliades, N. (2008). Combining a DL Reasoner and a Rule Engine for Improving Entailment-Based OWL Reasoning. In Sheth, A. P., Staab, S., Dean, M., Paolucci, M., Maynard, D., Finin, T. W., & Thirunarayan, K. (Eds.), *International Semantic Web Conference*, Vol. 5318 of *LNCS*, pp. 277–292, Karlsruhe, Germany.
- Meier, M. (2010). *On the Termination of the Chase Algorithm*. Ph.D. thesis, Universität Freiburg.
- Meier, M., Schmidt, M., & Lausen, G. (2009). On Chase Termination Beyond Stratification. *Proceedings of the VLDB Endowment*, 2(1), 970–981.
- Motik, B., Cuenca Grau, B., Horrocks, I., Wu, Z., Fokoue, A., & Lutz, C. (2009a). OWL 2 Web Ontology Language: Profiles, W3C Recommendation. <http://www.w3.org/TR/owl2-profiles/>.
- Motik, B., Shearer, R., & Horrocks, I. (2009b). Hypertableau Reasoning for Description Logics. *Journal of Artificial Intelligence Research*, 36, 165–228.
- Mungall, C. (2009). Experiences Using Logic Programming in Bioinformatics. In Hill, P. M., & Warren, D. S. (Eds.), *Proc. of the 25th Int. Conf. on Logic Programming (ICLP 2009)*, Vol. 5649 of *LNCS*, pp. 1–21, Pasadena, CA, USA. Springer.
- Ortiz, M., Calvanese, D., & Eiter, T. (2008). Data Complexity of Query Answering in Expressive Description Logics via Tableaux. *Journal of Automated Reasoning*, 41(1), 61–98.
- Ortiz, M., Rudolph, S., & Simkus, M. (2011). Query Answering in the Horn Fragments of the Description Logics *SHOIQ* and *SROIQ*. In Walsh, T. (Ed.), *Proc. of the 22nd Int. Joint Conf. on Artificial Intelligence (IJCAI 2011)*, pp. 1039–1044, Barcelona, Spain.

- Pérez-Urbina, H., Motik, B., & Horrocks, I. (2009). Tractable Query Answering and Rewriting under Description Logic Constraints. *Journal of Applied Logic*, 8(2), 151–232.
- Rudolph, S., & Glimm, B. (2010). Nominals, Inverses, Counting, and Conjunctive Queries or: Why Infinity is your Friend!. *Journal of Artificial Intelligence Research*, 39, 429–481.
- Spezzano, F., & Greco, S. (2010). Chase Termination: A Constraints Rewriting Approach. *Proceedings of the VLDB Endowment*, 3(1), 93–104.
- Syrjänen, T. (2001). Omega-Restricted Logic Programs. In Eiter, T., Faber, W., & Truszczyński, M. (Eds.), *Proc. of the 6th Int. Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2001)*, Vol. 2173 of *LNCS*, pp. 267–279, Vienna, Austria. Springer.
- Syrjänen, T., & Niemelä, I. (2001). The Smodels System. In Eiter, T., Faber, W., & Truszczyński, M. (Eds.), *Proc. of the 6th Int. Conf. on Logic Programming and Nonmonotonic Reasoning (LPNMR 2001)*, Vol. 2173 of *LNAI*, pp. 434–438, Vienna, Austria. Springer.
- Wu, Z., Eadon, G., Das, S., Chong, E. I., Kolovski, V., Annamalai, M., & Srinivasan, J. (2008). Implementing an Inference Engine for RDFS/OWL Constructs and User-Defined Rules in Oracle. In Alonso, G., Blakeley, J. A., & Chen, A. L. P. (Eds.), *Proc. of the 24th Int. Conf. on Data Engineering (ICDE 2008)*, pp. 1239–1248, Cancún, México. IEEE.