

Acyclicity Conditions and Their Application to Query Answering in Description Logics

Bernardo Cuenca Grau, Ian Horrocks, Markus Krötzsch, Clemens Kupke,
Despoina Magka, Boris Motik, and Zhe Wang

Department of Computer Science, Oxford University
Oxford, United Kingdom

Abstract

Answering conjunctive queries (CQs) over a set of facts extended with existential rules is a key problem in knowledge representation and databases. This problem can be solved using the *chase* (aka *materialisation*) algorithm; however, CQ answering is undecidable for general existential rules, so the chase is not guaranteed to terminate. Several *acyclicity conditions* provide sufficient conditions for chase termination. In this paper, we present two novel such conditions—*model-faithful acyclicity* (MFA) and *model-summarising acyclicity* (MSA)—that generalise many of the acyclicity conditions known so far in the literature.

Materialisation provides the basis for several widely-used OWL 2 DL reasoners. In order to avoid termination problems, many of these systems handle only the OWL 2 RL profile of OWL 2 DL; furthermore, some systems go beyond OWL 2 RL, but they provide no termination guarantees. In this paper we investigate whether various acyclicity conditions 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 DL ontologies are MSA, and that the facts obtained via materialisation are not too large. Thus, our results suggest that principled extensions to materialisation-based OWL 2 DL reasoners may be practically feasible.

Introduction

Existential rules are positive, function-free first-order implications that may contain existentially quantified variables in the head. In databases, they are known as *dependencies* (Abiteboul, Hull, and Vianu 1995) and are used to capture a wide range of schema constraints; in particular, they are 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 et al. 2005). They also provide the foundation for several prominent knowledge representation formalisms, such as Datalog[±] (Calì, Gottlob, and Pieris 2010; Calì et al. 2010).

Answering *conjunctive queries* (CQs) over a set of facts extended with existential rules is a fundamental reasoning

problem in both database and KR settings. This problem is undecidable (Beeri and Vardi 1981) in general, and it can be characterised using *chase* (Johnson and Klug 1984; Maier, Mendelzon, and Sagiv 1979), a technique closely related to the hypertableau calculus (Motik, Shearer, and Horrocks 2009). The chase extends in a forward-chaining manner the original set of facts by introducing facts implied by the rules. The result of the chase is called the *universal model*, and an arbitrary conjunctive query can be answered over the original set of facts and the rules by simply evaluating the query in the universal model.

Rules with existentially quantified variables in the head—so-called *generating rules*—require the introduction of fresh individuals, and cyclic applications of generating rules may lead to non-termination; moreover, determining whether chase terminates on a set of rules and facts is undecidable. 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. 2011), guarded rules (Calì et al. 2010), and ‘sticky’ rules (Calì, Gottlob, and Pieris 2011). In the second group, one uses a sufficient (but not necessary) *acyclicity* condition that ensures chase termination. Roughly speaking, acyclicity conditions analyse information flow between the 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 *safety* (SF) (Meier, Schmidt, and Lausen 2009), *stratification* (ST) (Deutsch, Nash, and Rummel 2008), *acyclicity of a graph of rule dependencies* (aGRD) (Baget, Mugnier, and Thomazo 2011), *joint acyclicity* (JA) (Krötzsch and Rudolph 2011), and *super-weak acyclicity* (SWA) (Marnette 2009).

Acyclicity conditions are relevant for at least two reasons. First, unlike guarded rules, acyclic rules can axiomatise structures of arbitrary shapes, as long as these structures are bounded in size. Second, the chase result for acyclic rules can be stored and manipulated as if it were a database. This is important in data exchange, where the goal is to materialise the transformed database.

In this paper, we argue that acyclicity is also relevant for *description logics* (DLs), the KR formalisms underpin-

ning the OWL 2 DL ontology language (Cuenca Grau et al. 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 et al. 2007; Krötzsch, Rudolph, and Hitzler 2007; Glimm et al. 2008; Ortiz, Calvanese, and Eiter 2008; Lutz, Toman, and Wolter 2009; Pérez-Urbina, Motik, and Horrocks 2010; Rudolph and Glimm 2010; Kontchakov et al. 2011).

Answering CQs over expressive DLs, however, is quite technical and of high computational complexity. Therefore, practical applications often solve this problem using *materialisation*, in which ontology consequences are precomputed using forward-chaining and stored in a semantic data store; examples of such systems include Oracle’s Semantic Data Store (Wu et al. 2008), Sesame (Broekstra, Kampman, and van Harmelen 2002), Jena (Carroll et al. 2004), OWLim (Kiryakov, Ognyanov, and Manov 2005), and DLE-Jena (Meditkos and Bassiliades 2008). This approach is possible if (i) the ontology is *Horn* (Hustadt, Motik, and Sattler 2005), and (ii) forward-chaining is guaranteed to terminate. In practice, condition (ii) is achieved by computing the materialisation only w.r.t. the inference rules that correspond to the part of the ontology in the OWL 2 RL profile; this systematically excludes generating rules and is thus terminating, but incomplete in general. Generating rules are partially supported in systems such as OWLim (Bishop and Bojanov 2011) and Jena; however, such support is typically ad hoc and provides no completeness and/or termination guarantees. Acyclicity conditions can be used to address these issues: if a Horn ontology is acyclic, a complete materialisation can be computed without the risk of non-termination.

Motivated by the practical importance of chase termination, in this paper we present two new acyclicity conditions: *model-faithful acyclicity* (MFA) and *model-summarising acyclicity* (MSA). Roughly speaking, these acyclicity conditions 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 rules, which makes it a very general condition. 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 of MFA checking, MFA may be unsuitable for practical application, so we introduce MSA. Intuitively, MSA 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. Finally, we prove that MSA is strictly more general than SWA—one of the most general acyclicity conditions currently in use.

Both of these conditions can be applied to general existential rules *without* equality. Equality can be incorporated via *singularisation* (Marnette 2009)—a technique that transforms the rules to encode the effects of equality, but does not

prevent the evaluation of a conjunctive query in the universal model. Singularisation is orthogonal to acyclicity: after computing the transformed rules, one can use MFA, MSA, or in fact an arbitrary acyclicity notion to check whether the result is acyclic; if so, the chase of the original set of rules will terminate. Unfortunately, singularisation is non-deterministic: some ways of transforming the rules may produce acyclic rule sets, but not all ways are guaranteed to do so. Thus, we refine singularisation to obtain an upper and a lower bound for acyclicity. We also show that, when used with JA, the lower bound coincides with WA.

Finally, we consider various theoretical and practical issues surrounding the use of acyclicity for CQ answering over DL ontologies. On the theoretical side, we show that checking MFA and MSA of Horn-*SHIQ* ontologies is PSPACE- and PTIME-complete, respectively, and that answering CQs over acyclic Horn-*SHIQ* ontologies is in PSPACE. The latter problem is EXPTIME-hard for general (i.e., not acyclic) Horn-*SHIQ* ontologies (Ortiz, Rudolph, and Simkus 2011), so acyclicity makes the problem easier. Furthermore, Horn ontologies can be extended with arbitrary SWRL rules (Horrocks and Patel-Schneider 2004) without affecting neither decidability nor 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 lead to the undecidability.

On the practical side, we explore the limits of reasoning with acyclic OWL 2 DL ontologies via materialisation. We checked MFA, MSA, and JA for 149 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 paper is accompanied with an online technical report that contains all proofs.¹

Preliminaries

We use the standard notions of constants, function symbols, and predicate symbols, where $\approx$ is the equality predicate. Each function or predicate symbol is associated with a nonnegative integer arity. Variables, terms, substitutions, atoms, and first-order formulae, sentences, interpretations (i.e., structures), and models are defined as usual. We abbreviate with $\vec{t}$ a vector of terms $t_1, \dots, t_n$ and define $|\vec{t}| = n$. 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 $t = f(\vec{t})$. Satisfaction of a sentence φ in an interpretation I (written $I \models \varphi$), and entailment of a sentence ψ from a sentence φ (written $\varphi \models \psi$), are defined as usual. By a slight

¹<http://www.cs.ox.ac.uk/isg/TR/acyclicity.pdf>

abuse of notation, we identify a conjunction of atoms with a set of atoms. A term t' is a *subterm* of a term t if $t' = t$ or $t = f(\vec{t})$ and t' is a subterm of some $t_i \in \vec{t}$; if additionally $t' \neq t$, then t' is a *proper* subterm of t . An atom $P(\vec{t})$ *contains* a term t if $t \in \vec{t}$, and a set of atoms I *contains* t if some atom in I contains t .

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

$$\forall \vec{x} \forall \vec{y} \vec{z}. [\varphi(\vec{x}, \vec{z}) \rightarrow \exists \vec{y}. \psi(\vec{x}, \vec{y})] \quad (1)$$

where $\varphi(\vec{x}, \vec{z})$ and $\psi(\vec{x}, \vec{y})$ are conjunctions of atoms, and the 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 $\forall \vec{x} \forall \vec{y} \vec{z}$ are often omitted. If $\vec{y}$ is empty, the rule is called a *datalog* rule. In database theory, satisfaction and entailment are often considered only w.r.t. *finite* interpretations under the *unique name assumption* (UNA), where distinct constants are interpreted as distinct elements; in contrast, such assumptions are not customary in KR. Since we study rules that can be satisfied in finite models, the restriction to finite satisfiability is immaterial; also, we do not assume UNA, which can be axiomatised if needed.

A *conjunctive query* (CQ) is a formula $Q(\vec{x})$ of the form $\exists \vec{y}. \varphi(\vec{x}, \vec{y})$, 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 a fundamental reasoning problem in many applications of existential rules.

In first-order logic, the equality predicate $\approx$ is commonly assumed to have a predefined interpretation. The semantics of $\approx$, however, can be axiomatised explicitly. Let Σ be an arbitrary set of rules; w.l.o.g. we assume that no rule in Σ contains $\approx$ in the body. Then, $\Sigma_{\approx} = \emptyset$ if $\approx$ does not occur in Σ ; otherwise, $\Sigma_{\approx}$ contains rules (2)–(4) and an instance of rule (5) for each predicate P occurring in Σ .

$$\rightarrow x \approx x \quad (2)$$

$$x_1 \approx x_2 \rightarrow x_2 \approx x_1 \quad (3)$$

$$x_1 \approx x_2 \wedge x_2 \approx x_3 \rightarrow x_1 \approx x_3 \quad (4)$$

$$P(\vec{x}) \wedge x_i \approx x'_i \rightarrow P(x_1, \dots, x'_i, \dots, x_n) \quad (5)$$

The consequences of Σ (where $\approx$ is treated as having a well-defined interpretation) and $\Sigma \cup \Sigma_{\approx}$ (where $\approx$ is treated as an ordinary predicate) coincide.

Sometimes we use *skolemisation* to interpret rules in *Herbrand* interpretations—possibly infinite sets of ground atoms. In particular, for each rule r of the form (1) and each variable $y_i \in \vec{y}$, let f_r^i be a function symbol that is globally unique for r and y_i ; furthermore, let θ be the substitution such that $\theta(y_i) = f_r^i(\vec{x})$ for each $y_i \in \vec{y}$. Then, the skolemisation $\text{sk}(r)$ of r is the rule

$$\varphi(\vec{x}, \vec{z}) \rightarrow \psi(\vec{x}, \vec{y})\theta \quad (6)$$

The skolemisation $\text{sk}(\Sigma)$ of a set of rules Σ is obtained by skolemising each rule in Σ . For each CQ $Q(\vec{x})$, instance I , and substitution σ , we have $\Sigma \cup I \models Q(\vec{x})\sigma$ if and only if $\text{sk}(\Sigma) \cup \Sigma_{\approx} \cup I \models Q(\vec{x})\sigma$.

Answering CQs can be characterised using *chase*, and we use the *skolem chase* variant (Marnette 2009). The result of applying a skolemised rule $r = \varphi \rightarrow \psi$ to a set of ground atoms I is the smallest set $r(I)$ that contains $\psi\sigma$ for each substitution σ from the variables in r to the terms occurring in I such that $\varphi\sigma \subseteq I$; furthermore, 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 I_{Σ}^i for $i > 0$ 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. Chase can be used as a ‘database’ for answering CQs: a substitution σ is an answer to Q over Σ and I iff $I_{\Sigma}^{\infty} \models Q\sigma$. Chase of I and Σ *terminates* if $i \geq 0$ exists such that $I_{\Sigma}^i = I_{\Sigma}^j$ for each $j \geq i$; 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 the nonoblivious chase (Fagin et al. 2005), and the core chase (Deutsch, Nash, and Rummel 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 $*$. If the skolem chase for I_{Σ}^* and Σ terminates, then the skolem chase of Σ terminates universally (Marnette 2009).

Universal chase termination is undecidable, and various sufficient *acyclicity* conditions have been proposed. In the following, let Σ be a set of rules where w.l.o.g. 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 $1 \leq i \leq n$. Given a rule r of the form (1) and a variable w , the set $\text{Pos}_B(w)$ of *body positions* of w consists of all positions $P|_i$ for which $P(t_1, \dots, t_n) \in \varphi(\vec{x}, \vec{z})$ exists with $t_i = w$. The set $\text{Pos}_H(w)$ is defined analogously.

Weak acyclicity (WA) (Fagin et al. 2005) can be used with rules containing the equality predicate. The *WA dependency graph* D_{Σ} for Σ contains positions as vertices; furthermore, for each rule $r \in \Sigma$ of the form (1), each $x \in \vec{x}$, each $P|_i \in \text{Pos}_B(x)$, and each $y \in \vec{y}$, graph D_{Σ} 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 D_{Σ} does not contain a cycle going through a special edge. Equality atoms are effectively ignored by WA.

Joint acyclicity (JA) (Krötzsch and 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* of Σ is as follows. The vertices are the existentially quantified variables occurring in Σ . Given arbitrary two such variables y_1 and y_2 , the JA dependency graph 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 its JA dependency graph does not contain a cycle.

Super-weak acyclicity (SWA) (Marnette 2009) is a further generalisation of JA, and it can differ from JA only on rules in which a variable occurs more than once in some body atom. Since such rules are not obtained from DL knowledge bases, we omit the somewhat technical definition of SWA.

Spezzano and Greco (2010) suggest a rule rewriting framework for chase termination. A set of rules Σ is rewritten into a set of rules Σ' , and then Σ' is checked using an arbitrary acyclicity notion (e.g., weak acyclicity). This technique is thus independent from the actual acyclicity condition used, and so we do not discuss it any further.

Model-Faithful Acyclicity

As the following example shows, known acyclicity conditions, such as JA, do not guarantee chase termination on certain commonly occurring rules.

Example 1. Let Σ be the set of rules (7)–(9).

$$r_1 = A(x_1) \rightarrow \exists y_1. R(x_1, y_1) \wedge B(y_1) \quad (7)$$

$$r_2 = R(x_2, z) \wedge B(z) \rightarrow A(x_2) \quad (8)$$

$$r_3 = B(x_3) \rightarrow \exists y_2. R(x_3, y_2) \wedge C(y_2) \quad (9)$$

Note that $\text{Move}(y_1) = \{R|_2, B|_1, R|_1, A|_1\}$; hence, the JA dependency graph has a cyclic edge from y_1 to itself. The chase of Σ , however, terminates universally. Assume that f and g are used to skolemize r_1 and r_3 . Given a fact $A(a)$, rule r_1 derives $R(a, f(a))$ and $B(f(a))$, and rule r_3 derives $R(f(a), g(f(a)))$ and $C(g(f(a)))$; after this, rule r_2 is not applicable to $R(f(a), g(f(a)))$ since variable z in r_2 cannot be matched to $B(g(f(a)))$, and so the chase terminates. $\diamond$

Note that rules r_1 and r_2 in Example 1 encode the DL axiom $A \equiv \exists R.B$, and rule r_3 encodes $B \sqsubseteq \exists R.C$; such axioms abound in OWL ontologies. To enable applications of chase termination outlined in the introduction, we next propose a less restrictive acyclicity condition.

Acyclicity conditions try to estimate whether applying chase to a rule can produce facts that can (possibly by applying chase to other rules) repeatedly trigger the same rule. The key difference between various conditions is how rule applicability is determined. For example, JA and SWA consider each variable in a rule in isolation and do not check satisfaction of all body atoms at once; hence, they overestimate rule applicability. For example, rule (8) is not applicable to the facts generated by rule (9), but this can be determined only by considering variables x_2 and z in rule (8) simultaneously. 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. Clearly, no sufficient, necessary, and computable condition for the latter can be given, so we must adopt a practical approach; for example, we can ‘raise the alarm’ and stop the process if the chase derives a term $f(\vec{t})$ where f occurs in $\vec{t}$. This condition 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 condition is thus application dependent; 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 Σ from Example 1.

Meier, Schmidt, and Lausen (2009) proposed a related idea, where the chase is extended to keep track of a ‘monitor graph’, which is used to track rule dependencies and then stop the chase if certain conditions are satisfied. This approach uses a variant of the chase that is don’t-know non-deterministic: while all possible chase applications produce a model, not all applications will terminate, which can make acyclicity checking difficult.

In contrast, our notion of acyclicity is independent from any concrete notion of chase. The given rules Σ are transformed into a new set of rules Σ' , which 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 sound and complete theorem proving procedure for existential rules. Acyclicity guarantees termination of skolem chase, which then guarantees termination of nonoblivious and core chase as well. We call our notion *model-faithful acyclicity* because it estimates rule application precisely, by examining the actual model of Σ .

Definition 2. 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:

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

For Σ a set of rules, $\text{MFA}(\Sigma)$ is the smallest set that contains $\text{MFA}(r)$ for each rule $r \in \Sigma$, rules (10)–(11), and rule (12) instantiated for each predicate F_r^i :

$$S(x_1, x_2) \rightarrow D(x_1, x_2) \quad (10)$$

$$D(x_1, x_2) \wedge S(x_2, x_3) \rightarrow D(x_1, x_3) \quad (11)$$

$$F_r^i(x_1) \wedge D(x_1, x_2) \wedge F_r^i(x_2) \rightarrow C \quad (12)$$

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 MFA²* if Σ is MFA w.r.t. I_Σ^* .

Example 3. Let Σ be the set of rules from Example 1. Then, $\text{MFA}(r_1)$ and $\text{MFA}(r_3)$ are given by (13) and (14), resp.

$$A(x_1) \rightarrow \exists y_1. R(x_1, y_1) \wedge B(y_1) \wedge F_{r_1}^1(y_1) \wedge S(x_1, y_1) \quad (13)$$

²In the rest of this paper we typically omit ‘universally’.




Figure 1: Example MFA

$$B(x_3) \rightarrow \exists y_2. R(x_3, y_2) \wedge C(y_2) \wedge F_{r_3}^1(y_2) \wedge S(x_3, y_2) \quad (14)$$

Then, $\text{MFA}(\Sigma)$ consists of rules (8), rules (13)–(14), rules (10)–(11), and rule (12) instantiated for $F_{r_1}^1$ and $F_{r_3}^1$.

Set Σ is universally MFA: the interpretation shown in Figure 1 is a model of $I_\Sigma^* \cup \text{MFA}(\Sigma)$ that does not satisfy C . $\diamond$

MFA is defined as a semantic, rather than a syntactic condition, and entailment $I \cup \text{MFA}(\Sigma) \not\models C$ can be checked using sound and complete first-order calculus. In the following section we show that MFA is strictly more general than SWA. We next show that MFA characterises derivations of the skolem chase in a particular way.

Definition 4. A term t is cyclic if some $f(\vec{s})$ is a subterm of t , and some $f(\vec{u})$ is a proper subterm of $f(\vec{s})$.

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

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 doubly exponentially many different terms and atoms. The 2EXPTIME bound already holds if the rules are WA; hence, CQ answering for MFA rules is not harder than for WA.

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.

Proposition 6 implies that answering a BCQ over MFA rules is in 2EXPTIME; furthermore, Cali, Gottlob, and Pieris (2010) provide the matching lower bound for WA rules. We next prove that checking MFA w.r.t. a specific instance I is in 2EXPTIME, and that checking universal MFA is 2EXPTIME-hard. These results provide tight complexity bounds for both problems.

Theorem 7. For Σ a set of rules, 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.

The results of Theorem 7 are somewhat discouraging: using the existing conditions, acyclicity can typically be checked in PTIME or in NP. We consider MFA to be an ‘upper bound’ of practically useful acyclicity conditions. We see two possibilities for improving these results. In the following section, we introduce an approximation of MFA that is easier to check; our evaluation shows that this condition often coincides with MFA in practice. Next, we show that the complexity is lower for rules of the following shape.

Definition 8. A rule r of the form (1) 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. The following theorem provides the upper bound; the lower bounds are given later on for a smaller class of rules that capture DLs.

Theorem 9. 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. Furthermore, if Σ is MFA, then answering a BCQ over Σ and I is in EXPTIME as well.

Model-Summarising Acyclicity

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

Definition 10. Let S , D , and F_r^i be as in Definition 2; 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 is the substitution that maps each variable $y_i \in \vec{y}$ to c_r^i :

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

For Σ a set of rules, $\text{MSA}(\Sigma)$ is the smallest set that contains $\text{MSA}(r)$ for each rule $r \in \Sigma$, rules (10)–(11), and rule (12) 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 11. Consider again the set of rules Σ in Example 1. $\text{MSA}(r_1)$ and $\text{MSA}(r_3)$ are given by the following rules (15) and (16), respectively:

$$A(x_1) \rightarrow R(x_1, c_{r_1}^1) \wedge B(c_{r_1}^1) \wedge F_{r_1}^1(c_{r_1}^1) \wedge S(x_1, c_{r_1}^1) \quad (15)$$

$$B(x_3) \rightarrow R(x_3, c_{r_3}^1) \wedge C(c_{r_3}^1) \wedge F_{r_3}^1(c_{r_3}^1) \wedge S(x_3, c_{r_3}^1) \quad (16)$$




Figure 2: Example MSA

Then, $\text{MSA}(\Sigma)$ consists of rules (8), rules (15)–(16), rules (10)–(11), and rule (12) instantiated for F_{r1}^1 and F_{r3}^1 .

Set Σ is universally MSA: the interpretation shown in Figure 2 is a model of $I_\Sigma^* \cup \text{MSA}(\Sigma)$ that does not satisfy C . $\diamond$

Note that $\text{MSA}(\Sigma)$ is a set of datalog rules, so MSA can be checked using a datalog reasoner. This connection with datalog provides the upper complexity bound for checking MSA: the following theorem follows from the well known complexity results of checking entailment of a ground atom in a datalog program (Dantsin et al. 2001). The complexity of datalog reasoning is $O(n^v)$ where 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 12. *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.*

We finish this section by proving strict inclusion relationships between MFA, MSA, and SWA. In particular, Theorem 13 and Example 14 show that MFA is strictly more general than MSA.

Theorem 13. *If a set of rules Σ is MSA (w.r.t. an instance I), then Σ is MFA (w.r.t. I) as well.*

Example 14. *Let Σ be the set of rules (17)–(20).*

$$A(x) \rightarrow \exists y. R(x, y) \wedge B(y) \quad (17)$$

$$B(x) \rightarrow \exists y. S(x, y) \wedge T(y, x) \quad (18)$$

$$A(z) \wedge S(z, x) \rightarrow C(x) \quad (19)$$

$$C(z) \wedge T(z, x) \rightarrow A(x) \quad (20)$$

Σ is universally MFA, but not universally MSA. $\diamond$

We now show that MSA is more general than SWA, and thus also more general than JA. The converse does not hold: the set Σ in Example 1 is MSA, but not SWA.

Theorem 15. *If a set of equality-free rules Σ is SWA, then Σ is universally MSA as well.*

Handling Equality via Singularisation

JA and SWA can be applied to rules with equality provided that the rule set contains rules (2)–(5). In both cases, however, rules (2) and (5) lead to a cycle as soon as the rule

set contains an existential quantifier. 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 (2)–(5) are not explicitly stated. On rules with equality, MFA and MSA are slightly more general than JA and SWA, but they still fail to capture certain relevant cases, as the following example demonstrates.

Example 16. *Consider the following set of rules.*

$$A(x) \wedge B(x) \rightarrow \exists y. [R(x, y) \wedge B(y)] \quad (21)$$

$$R(z, x_1) \wedge R(z, x_2) \rightarrow x_1 \approx x_2 \quad (22)$$

On (21)–(22), skolem chase derives the following infinite set of facts; but then, by Proposition 5 this set of rules is not MFA; by Theorem 13, it is not MSA either.

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

Example 16 shows that equalities between terms tend to proliferate during chase, which can lead to non-termination. Interestingly, the rules in Example 16 are WA. This is because WA is sufficient for termination of nonoblivious chase—a version of chase that expands an existential quantifier only if necessary. Already JA is more general than WA on rules without equality, so nonoblivious chase does not seem to provide an advantage over skolem chase w.r.t. termination on such rules; however, Example 16 shows that this is not the case for rules with equality.

Marnette (2009) proposed *singularisation* as a possible solution to this problem. The idea is to only partially axiomatise $\approx$ as being reflexive, symmetric, and transitive, but without the replacement property cf. rule (5). A set of rules Σ is modified in a way to take into account the lack of the replacement rules. This latter step is nondeterministic: there are many ways to modify Σ and, while some modifications will lead to chase termination, not all will do so.

We recapitulate the definition of singularisation. A *marking* M_r of a rule r of the form (1) is a mapping from each $w \in \vec{x} \cup \vec{z}$ to a single occurrence of w in φ ; all other variable occurrences are *unmarked* and all constants are also unmarked. A *marking* M of a set of rules Σ contains exactly one marking M_r for each $r \in \Sigma$. The *singularisation* of Σ under M is the set $\text{Sing}(\Sigma, M)$ containing

- for each $r \in \Sigma$, a rule obtained by replacing each unmarked occurrence of a body term t in r with a fresh variable z' and adding $t \approx z'$ to the body, and
- rules (2)–(4).

Note that $\text{Sing}(\Sigma, M)$ is unique up to the renaming of the fresh variables. We identify the marked occurrence of a variable x in a rule as $\hat{x}$. The properties of singularisation can be summarised as follows: for an arbitrary set of rules Σ , a marking M for Σ , an instance I , and a fact $P(\vec{c})$, we have $I \cup \Sigma \cup \Sigma_{\approx} \models P(\vec{c})$ if and only if

$$I \cup \text{Sing}(\Sigma, M) \models \exists \vec{y}. [P(\vec{y}) \wedge \bigwedge_{y_i \in \vec{y}} y_i \approx c_i].$$

Example 17. Singularisation of the marked rule (23) produces rule (24).

$$A(\hat{x}) \wedge B(x) \wedge R(x, \hat{z}) \rightarrow C(x) \quad (23)$$

$$A(x) \wedge B(x_1) \wedge R(x_2, z) \wedge x \approx x_1 \wedge x \approx x_2 \rightarrow C(x) \quad (24)$$

Note that singularisation should be applied ‘globally’ to all rules, even to those without equality. $\diamond$

The absence of rules (5) often allows the skolem chase to terminate on $\text{Sing}(\Sigma, M)$; however, this may depend on the selected marking.

Example 18. Rule (21) from Example 16 admits the following two markings:

$$A(\hat{x}) \wedge B(x) \rightarrow \exists y.[R(x, y) \wedge B(y)] \quad (25)$$

$$A(x) \wedge B(\hat{x}) \rightarrow \exists y.[R(x, y) \wedge B(y)] \quad (26)$$

The skolem chase does not universally terminate for the singularisation obtained from (26); in contrast, the singularisation obtained from (25) is JA. $\diamond$

We use $\text{MFA}^\exists$ and $\text{MFA}^\forall$ to denote the classes of rule sets that are in MFA for *some* singularisation and for *all* singularisations, respectively; notions $\text{MSA}^\exists$, $\text{MSA}^\forall$, $\text{JA}^\exists$, and $\text{JA}^\forall$ are defined analogously. Clearly, $X^\forall \subseteq X^\exists$ for each $X \in \{\text{MFA}, \text{MSA}, \text{JA}\}$, and Example 18 shows this inclusion to be proper.

Theorem 19. $\text{JA}^\forall = \text{WA}$.

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 19 shows that $\text{JA}^\forall$ can be decided using WA. For the other cases, the following simple observation shows how to reduce the number of markings.

Proposition 20. Let M and M' be markings for Σ that agree on all variables that occur in both body and head, but not necessarily on the variables that occur only in the body of a rule. Then $\text{Sing}(\Sigma, M)$ is JA/MSA/MFA if and only if $\text{Sing}(\Sigma, M')$ is JA/MSA/MFA.

Despite this optimisation, the number of markings to check can still be exponential; hence, we next describe a useful approximation. Let $\text{Sing}_\cup(\Sigma) = \bigcup_{M \in \mathcal{M}} \text{Sing}(\Sigma, M)$, where $\mathcal{M}$ is a set of all markings for Σ that agree on all variables occurring only in the body of a rule in Σ . By Proposition 20, it is irrelevant how the markings of body variables are defined in $\mathcal{M}$. Let $\text{MFA}^\cup$ be the class containing rule sets Σ for which $\text{Sing}_\cup(\Sigma)$ is in MFA; $\text{MSA}^\cup$ and $\text{JA}^\cup$ are defined analogously. As the following theorem shows, $\text{Sing}_\cup(\Sigma)$ provides a ‘lower bound’ on the result attainable via singularisation.

Theorem 21. For each $X \in \{\text{MFA}, \text{MSA}, \text{JA}\}$, we have that $X^\cup \subseteq X^\forall$. The size of $\bigcup_{M \in \mathcal{M}} \text{Sing}(\Sigma, M)$ is exponential in the maximal number of variables that occur more than once in the body of any one rule in Σ , and it is linear in the number of rules in Σ .

This result is of particular interest when dealing with rules that are obtained from DLs, where each rule has at most one

variable that occurs in the head as well as multiple times in the body. On such rule sets, the size of $\text{Sing}_\cup(\Sigma)$ is linear in the size of Σ . For the general case, we can obtain the same complexity bounds despite the exponential increase in the number of rules.

Theorem 22. Deciding whether Σ is in $\text{MFA}^\cup$ ($\text{MFA}^\exists$, $\text{MFA}^\forall$) is 2EXPTIME-complete. Deciding whether Σ is in $\text{MSA}^\cup$ ($\text{MSA}^\exists$, $\text{MSA}^\forall$) is EXPTIME-complete.

Acyclicity of DL Ontologies

We now consider applying acyclicity conditions to DL ontologies. DLs are KR 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 R an atomic role, R^- is an *inverse role*; inverse roles can be used in atoms, and $R^-(t_1, t_2)$ is an abbreviation for $R(t_2, t_1)$. 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. DL ontologies consist of *axioms* about concepts and roles; these correspond to first-order sentences. For simplicity, we consider only normalised ontologies, where concepts are not nested. This is w.l.o.g. as each ontology can be normalised in linear time, and the normalised ontology is a conservative extension of the original one. In this paper, we consider only *Horn* DLs; ontologies in such DLs have at most one minimal Herbrand model, which is a prerequisite for materialisation-based reasoning—the main motivation for applying acyclicity to DLs.

A normalised Horn-*SRIQ* 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 , T are (not necessarily atomic) roles, and n is a positive integer. To guarantee decidability of reasoning, $\mathcal{T}$ must satisfy certain *global* conditions (Kutz, Horrocks, and Sattler 2006), which we omit for 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; the latter condition ensures that axioms of Type 6 can be represented using a nondeterministic finite automaton. Apart from Horn-*SRIQ*, we also consider Horn-*SRI* TBoxes, which do not contain rules of Type 2, as well as Horn-*SHIQ* TBoxes, where $R = S = T$ in all rules of Type 6; all Horn-*SHIQ* TBoxes are regular.

Each Horn-*SRIQ* axiom corresponds to an existential rule as shown in Table 1. A minor difference is that axioms in Table 1 can contain $\perp$ in the head, which can make a TBox $\mathcal{T}$ unsatisfiable w.r.t. an instance I . This can be handled by considering $\perp$ to be just another atomic concept, without special meaning. Technically, this ensures that $I \cup \mathcal{T}$ is satisfiable in the model constructed by the skolem chase; however, we consider $I \cup \mathcal{T}$ to be unsatisfiable if $I \cup \mathcal{T} \models \exists y. \perp(y)$. We consider a substitution θ to be an answer to a CQ $Q(\vec{x})$ w.r.t. a $\mathcal{T}$ and I if $I \cup \mathcal{T} \models \exists y. \perp(y)$ or $\mathcal{T} \cup I \models Q(\vec{x})\theta$. Due to this close correspondence between DL axioms and existential rules, in the rest of this paper we

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.	$R \circ S \sqsubseteq T$	$R(x_1, z) \wedge S(z, x_2) \rightarrow T(x_1, x_2)$

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

identify a TBox $\mathcal{T}$ with the corresponding set of rules.

We next investigate the complexity of BCQ answering over acyclic DL TBoxes. For the membership, note that all rules in Table 1 are $\exists$ -1 rules; thus, Theorem 9 gives us an EXPTIME upper bound. We next prove a matching lower bound for WA Horn-*SRIQ* rules. Intuitively, axioms of Type 6 allow us to axiomatise 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, which may require one to examine all nodes in an exponential model of a Horn-*SRIQ* TBox. Furthermore, by Theorem 9, if we extend acyclic Horn-*SRIQ* rules Σ_1 with arbitrary SWRL rules Σ_2 , reasoning stays EXPTIME-complete, provided that $\Sigma_1 \cup \Sigma_2$ is acyclic; this is in contrast to general TBoxes for which SWRL extensions lead to undecidability. Thus, applications that need expressivity beyond what is available in OWL can benefit from the required expressivity without running into undecidability as long as the resulting ontology is acyclic.

Theorem 23. *Let $\mathcal{T}$ be a WA Horn-*SRIQ* TBox, let I be an instance, and let F be a fact. Then, checking whether $I \cup \mathcal{T} \models F$ is EXPTIME-hard.*

Note that Theorem 23 applies to Horn-*SRIQ* and thus does not rely on a particular treatment of equality.

The proof of Theorem 23 can be adapted to obtain the lower bound for checking MFA of Horn-*SRIQ* rules.

Proposition 24. *Checking whether a Horn-*SRIQ* TBox is universally MFA is EXPTIME-hard.*

As we show next, however, the complexity of query answering drops to PSPACE for MFA Horn-*SHIQ* ontologies. In contrast, checking entailment of a single fact is EXPTIME-hard in the general (i.e., not acyclic) case (Krötzsch, Rudolph, and Hitzler 2012). This drop in complexity is due to the fact that, if $R = S = T$ in all rules of Type 6, the automaton describing roles is of polynomial size, and the elimination of role inclusions by Demri and de Nivelle (2005) becomes polynomial. Thus, although acyclic TBoxes can axiomatise existence of polynomially deep and exponentially large structures, these structures are tree-like, which allows us to explore the structures one path at a time using the well-known tracing technique (Baader et al. 2007). The main difficulty in the membership proof of the following theorem is due to the fact that queries can contain transitive roles, so one cannot roll a query up into a concept. Since the TBox is Horn, however, one can guess places in the model that the query maps to. Given one such guess, one can ground the query and check entailment of each ground query atom individually, while taking transitive roles into

account. Furthermore, note that PSPACE-hardness proof of concept satisfiability checking by Baader et al. (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.

Theorem 25. *Let $\mathcal{T}$ be Horn-*SHIQ* TBox, let I be an instance such that $\mathcal{T}$ is MFA w.r.t. I , and let Q be a BCQ. Then, deciding $I \cup \mathcal{T} \models Q$ is PSPACE-complete.*

Although the proof of Theorem 25 takes into account ontology rules with equality (i.e., rules of Type 2), it assumes that equality is axiomatised by $\Sigma_{\approx}$ and hence it does not directly apply to *singularised* Horn-*SHIQ* rules. We conjecture, however, that the result in the theorem holds even if singularisation is applied.

The restriction to Horn-*SHIQ* rules also makes checking MFA w.r.t. an instance easier: this task can be accomplished by a minor variation of the query answering algorithm.

Theorem 26. *Let $\mathcal{T}$ be Horn-*SHIQ* 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.*

Finally, MSA provides us with a tractable condition for Horn-*SHIQ* 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. This result also holds for singularised rules.

Theorem 27. *Let $\mathcal{T}$ be Horn-*SHIQ* TBox, and let I be an instance. Then, deciding whether $\mathcal{T}$ is MSA w.r.t. I is in PTIME, and deciding whether $\mathcal{T}$ is universally MSA is PTIME-hard.*

Experiments

We have evaluated the applicability of various acyclicity conditions in practice. First, we implemented MFA, MSA, JA, and WA checkers, and used them to check acyclicity of a large corpus of Horn ontologies. Our goal was to determine whether a substantial portion of these ontologies are acyclic and could thus be used with (suitably extended) materialisation-based reasoners. Second, we computed the materialisation of the acyclic Horn ontologies and compared the size of the materialisation with the size of the original ABox. The goal of these tests was to see whether materialisation-based reasoning is practically feasible.

Tests were performed on the Oxford Super Computer HAL system with 8 2.8GHz processors and 16GB RAM. We

used a repository of 149 OWL ontologies whose TBox axioms can be transformed into existential rules. These ontologies include many of those in the Gardiner corpus (Gardiner, Tsarkov, and Horrocks 2006), the LUBM ontology, and a number of ontologies from the Open Biomedical Ontology (OBO) corpus. None of our test ontologies, however, has been obtained from typical conceptual models (i.e., ER or UML diagrams); due to the specific modelling patterns used in conceptual modelling, such ontologies are more likely to be cyclic. All test ontologies are available online.³

Acyclicity Tests

We implemented all acyclicity checks by adapting the Hermit reasoner. Hermit was used to transform an ontology into DL-clauses—formulae quite close to existential rules. DL-clauses were then preprocessed: at-least number restrictions in rule heads were replaced with existential quantification, atoms involving datatypes were eliminated, and DL-clauses with empty head were removed; datatypes and empty heads merely cause inconsistencies, and do not contribute to chase non-termination. If the DL-clauses contained equality, we check $X^\cup$ instead of X for each $X \in \{\text{MFA}, \text{MSA}, \text{JA}\}$ as a ‘lower-bound’ for acyclicity. To obtain an ‘upper bound’ for acyclicity, we checked whether the ontology was already cyclic when ignoring the rules containing equality. These steps produced a set of existential rules, which were further modified as required to encode the desired acyclicity check. Finally, Hermit was used to test universal acyclicity of the ontology by checking logical entailment w.r.t. the critical instance.

Each acyclicity test was given a 500s timeout. The MSA test exceeded this limit on 2 ontologies, whereas the MFA test exceeded the limit on 26 ontologies. Of the 149 ontologies tested, 124 (83%) were MSA. Moreover, MFA and MSA are indistinguishable w.r.t. the test ontologies—that is, all MFA ontologies were found to be MSA as well (the converse holds per Theorem 13). Results are shown in Table 2. Given the large number of test ontologies, we cannot show results for each ontology. Instead, ontologies are grouped by number of generating rules (G-rules); for each group, the table shows the number of ontologies (Total) and the number of ontologies found to be MSA, JA, and WA.

Note that seven large OBO ontologies were MSA but not JA; thus, MSA may be especially useful on large and complex ontologies. Table 3 shows for each of these ontologies the number of generating rules (G-rules), if it uses equality (Eq), expressivity (DL), and the number of classes (C), properties (P), and axioms (A).

Of the 25 non-MFA ontologies, only two ontologies are in $\mathcal{ELH}^r$ or certain versions of DL-Lite; this is interesting since combined approaches (Lutz, Toman, and Wolter 2009; Kontchakov et al. 2011) can be used to support CQ answering over these ontologies.

Materialisation Tests

To estimate the practicability of materialisation in acyclic ontologies, we measured the maximal depth of function

G-rules	Total	MSA	JA	WA
ontologies without equality				
< 100	21	19	19	19
100–1K	33	30	30	23
1K–5K	18	14	14	12
5K–12K	9	8	6	6
12K–160K	7	5	3	3
ontologies with equality				
< 100	49	45	45	45
100–1K	0	0	0	0
1K–5K	2	0	0	0
5K–12K	5	3	0	0
12K–160K	5	0	0	0

Table 2: Results of acyclicity tests

G-rules	Eq	DL	C	P	A
biological_process_xp_self.imports.owl					
10980	yes	$\mathcal{SRLI}$	22375	183	47454
go_xp_regulation.owl					
11187	no	$\mathcal{SH}$	27883	5	50941
biological_process_xp_cell.imports.owl					
11274	yes	$\mathcal{SRLI}$	24309	293	50386
cellular_component_xp_go.imports.owl					
11473	no	$\mathcal{SR}$	35236	8	64026
biological_..._cellular_component.imports.owl					
11798	yes	$\mathcal{SRLI}$	25337	187	52759
go_xp_regulation.imports.owl					
23844	no	$\mathcal{SR}$	34293	8	104473
biological_..._multi_organism_process.imports.owl					
24678	no	$\mathcal{SR}$	34410	21	104873

Table 3: MSA but not JA ontologies

symbol nesting in terms generated by skolem chase. This measure, which we call *ontology depth*, is of interest as it can be used to establish a bound on the size of the chase. Our tests revealed that most ontologies have small depths: out of the 124 MSA ontologies, 83 (66.9%) have depths less than 5; 13 (10.5%) have depths from 5 to 9; 24 (19.4%) have depths from 10 to 19; 2 (1.6%) have depths from 20 to 49; and 2 (1.6%) have depths from 50 to 80.

We also computed the materialisation of several acyclic ontologies. Since our implementation is only prototypical, our primary goal was not to evaluate the performance of materialisation, but rather to estimate the increase in ABox size. Although this increase may not be perfectly linear, we believe that it can be estimated by examining moderately-sized ABoxes. Most of our test ontologies, however, do not have substantial ABoxes; ontologies are often publicly available as general vocabularies, whereas ABoxes are application-specific and are thus usually not publicly available. Because of this problem, we conducted two kinds of experiments.

First, we computed the materialisation of two ontologies with nontrivial ABoxes: LUBM with one university and the

³<http://hermit-reasoner.com/2011/acyclicity/TestCorpus.zip>

Depth	#	time		gen. size		mat. size	
		max	avg	max	avg	max	avg
< 5	82	69	0.9	27	2	35	5
5–9	13	68	11	37	11	41	13
10–80	14	549	101	281	51	283	53

Table 4: Materialisation times (in seconds) and sizes

‘kmi-basic-portal’ ontology. The TBox of LUBM contains 8 generating rules and has depth 1; the ABox before materialisation contains 100,543 facts. Materialisation took only 1 second, and it produced 150,530 new facts, of which 47,798 were added by generating rules. The ‘kmi-basic-portal’ ontology has 10 generating rules and has depth 2; the ABox contains 179 facts. Materialisation took only 0.01 seconds, and it added 975 new facts, of which 151 were added by generating rules.

Second, for each of the 124 ontologies identified as MSA, we computed an ABox by instantiating each class and property with fresh individuals. We then computed the materialisation and measured the *generated size* (number of facts introduced by generating rules, divided by the facts in the initial ABox), the *materialisation size* (facts in the materialisation, divided by facts in the initial ABox), and the materialisation time. Since most generating rules in these ontologies had 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 ABox size caused by materialisation. Of the 124 ontologies tested, 15 exceeded the 1,000s time limit for materialisation. The results for the other 109 ontologies are shown in Table 4. Ontologies are grouped by depth; each group shows the number of ontologies (#), and materialisation times, generated sizes, and materialisation sizes.

Thus, materialisation seems practically feasible for many ontologies: for 82 ontologies with depth less than 5, materialisation increases the ontology size by a factor of 5. 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.

Conclusion

In this paper, we have studied the problem of CQ answering over acyclic existential rules. We have proposed two novel acyclicity conditions that are sufficient to ensure chase termination and that generalise many of the existing acyclicity conditions known thus far.

We have then studied the problem of CQ answering over acyclic DL ontologies. Acyclicity provides several compelling benefits for DL query answering. First, the CQ answering problem over Horn ontologies is computationally easier than for general ontologies. Second, under acyclicity conditions it is possible to extend Horn ontologies with arbitrary SWRL rules while preserving both the decidability and the worst-case complexity of the formalism. Third, acyclicity enables principled extensions of ontology materialisation-based reasoners. Fourth, ontologies used in

practice are often acyclic, so our results open the door to practical CQ answering beyond the OWL 2 RL profile.

Acknowledgments

This work was supported by the EU FP7 project SEALS and by the EPSRC projects ConDOR, ExODA, and LogMap. B. Cuenca Grau is supported by a Royal Society University Research Fellowship.

References

- Abiteboul, S.; Hull, R.; and Vianu, V. 1995. *Foundations of Databases*. Addison-Wesley.
- Baader, F.; Calvanese, D.; McGuinness, D.; Nardi, D.; and Patel-Schneider, P. F., eds. 2007. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, 2nd edition.
- Baget, J.-F.; Leclère, M.; Mugnier, M.-L.; and Salvat, E. 2011. On rules with existential variables: Walking the decidability line. *Artificial Intelligence* 175(9–10):1620–1654.
- Baget, J.-F.; Mugnier, M.-L.; and Thomazo, M. 2011. Towards farsighted dependencies for existential rules. In *Proc. RR*, 30–45.
- Beeri, C., and Vardi, M. Y. 1981. The implication problem for data dependencies. In *Proc. ICALP*, 73–85.
- Bishop, B., and Bojanov, S. 2011. Implementing OWL 2 RL and OWL 2 QL rule-sets for OWLIM. In *Proc. OWLED*, volume 796 of *CEUR WS Proceedings*.
- Broekstra, J.; Kampman, A.; and van Harmelen, F. 2002. Sesame: A Generic Architecture for Storing and Querying RDF and RDF Schema. In *Proc. ISWC*, 54–68.
- Cali, A.; Gottlob, G.; Lukasiewicz, T.; Marnette, B.; and Pieris, A. 2010. Datalog+/-: A family of logical knowledge representation and query languages for new applications. In *Proc. LICS*, 228–242.
- Cali, A.; Gottlob, G.; and Pieris, A. 2010. Query answering under non-guarded rules in Datalog+/- . In *Proc. RR*, 1–17.
- Cali, A.; Gottlob, G.; and Pieris, A. 2011. New expressive languages for ontological query answering. In *Proc. AAAI*.
- Calvanese, D.; De Giacomo, G.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2007. Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *J. Automated Reasoning* 39(3):385–429.
- Carroll, J. J.; Dickinson, I.; Dollin, C.; Reynolds, D.; Seaborne, A.; and Wilkinson, K. 2004. Jena: Implementing the Semantic Web Recommendations. In *Proc. WWW—Alternate Track*, 74–83.
- Cuenca Grau, B.; Horrocks, I.; Motik, B.; Parsia, B.; Patel-Schneider, P.; and Sattler, U. 2008. OWL 2: The next step for OWL. *J. Web Semantics* 6(4):309–322.
- Dantsin, E.; Eiter, T.; Gottlob, G.; and Voronkov, A. 2001. Complexity and expressive power of logic programming. *ACM Comput. Surv.* 33(3):374–425.
- Demri, S., and de Nivelle, H. 2005. Deciding Regular Grammar Logics with Converse Through First-Order Logic. *J. Logic, Language and Information* 14(3):289–329.

- Deutsch, A.; Nash, A.; and Rummel, J. B. 2008. The chase revisited. In *Proc. PODS*, 149–158.
- Fagin, R.; Kolaitis, P. G.; Miller, R. J.; and Popa, L. 2005. Data exchange: semantics and query answering. *Theoretical Computer Science* 336(1):89–124.
- Gardiner, T.; Tsarkov, D.; and Horrocks, I. 2006. Framework for an automated comparison of description logic reasoners. In *Proc. ISWC*, 654–667.
- Glimm, B.; Lutz, C.; Horrocks, I.; and Sattler, U. 2008. Conjunctive query answering for the description logic *SHIQ*. *J. Artif. Intell. Res.* 31:157–204.
- Horrocks, I., and Patel-Schneider, P. F. 2004. A proposal for an OWL rules language. In *Proc. WWW*, 723–731.
- Hustadt, U.; Motik, B.; and Sattler, U. 2005. Data complexity of reasoning in very expressive description logics. In *Proc. IJCAI*, 466–471.
- Johnson, D. S., and Klug, A. C. 1984. Testing containment of conjunctive queries under functional and inclusion dependencies. *J. Comput. Syst. Sci.* 28(1):167–189.
- Kiryakov, A.; Ognjanov, D.; and Manov, D. 2005. OWLIM: A pragmatic semantic repository for OWL. In *WISE Workshops*, 182–192.
- Kontchakov, R.; Lutz, C.; Toman, D.; Wolter, F.; and Zakharyashev, M. 2011. The combined approach to ontology-based data access. In *Proc. IJCAI*, 2656–2661.
- Krötzsch, M., and Rudolph, S. 2011. Extending decidable existential rules by joining acyclicity and guardedness. In *Proc. IJCAI*, 963–968.
- Krötzsch, M.; Rudolph, S.; and Hitzler, P. 2007. Conjunctive queries for a tractable fragment of OWL 1.1. In *Proc. ISWC*, 310–323.
- Krötzsch, M.; Rudolph, S.; and Hitzler, P. 2012. Complexities of Horn description logics. *ACM Trans. Comp. Log.* To appear.
- Kutz, O.; Horrocks, I.; and Sattler, U. 2006. The Even More Irresistible *SR_QIQ*. In *Proc. KR*, 68–78.
- Lutz, C.; Toman, D.; and Wolter, F. 2009. Conjunctive query answering in the description logic *EL* using a relational database system. In *Proc. IJCAI*.
- Maier, D.; Mendelzon, A. O.; and Sagiv, Y. 1979. Testing implications of data dependencies. *ACM Trans. Database Syst.* 4(4):455–469.
- Marnette, B. 2009. Generalized schema-mappings: from termination to tractability. In *Proc. PODS*, 13–22.
- Meditskos, G., and Bassiliades, N. 2008. Combining a DL reasoner and a rule engine for improving entailment-based OWL reasoning. In *Proc. ISWC*, 277–292.
- Meier, M.; Schmidt, M.; and Lausen, G. 2009. On chase termination beyond stratification. *Proc. VLDB* 2(1):970–981.
- Motik, B.; Shearer, R.; and Horrocks, I. 2009. Hypertableau reasoning for description logics. *J. Artif. Intell. Res.* 36:165–228.
- Ortiz, M.; Calvanese, D.; and Eiter, T. 2008. Data complexity of query answering in expressive description logics via tableaux. *J. Automated Reasoning* 41(1):61–98.
- Ortiz, M.; Rudolph, S.; and Simkus, M. 2011. Query answering in the Horn fragments of the description logics *SH_QIQ* and *SR_QIQ*. In *Proc. IJCAI*, 1039–1044.
- Pérez-Urbina, H.; Motik, B.; and Horrocks, I. 2010. Tractable query answering and rewriting under description logic constraints. *J. Applied Logic* 8(2):186–209.
- Rudolph, S., and Glimm, B. 2010. Nominals, inverses, counting, and conjunctive queries or: Why infinity is your friend! *J. Artif. Intell. Res.* 39:429–481.
- Spezzano, F., and Greco, S. 2010. Chase termination: A constraints rewriting approach. *Proc. VLDB* 3(1):93–104.
- Wu, Z.; Eadon, G.; Das, S.; Chong, E. I.; Kolovski, V.; An-namalai, M.; and Srinivasan, J. 2008. Implementing an inference engine for RDFS/OWL constructs and user-defined rules in Oracle. In *Proc. ICDE*, 1239–1248.