

Finite Model Reasoning in Horn-*SHIQ*

Yazmín Ibañez-García¹, Carsten Lutz², and Thomas Schneider²

¹ KRDB Research Centre, Free Univ. of Bozen-Bolzano, Italy, {ibanezgarcia@inf.unibz.it}

² Univ. Bremen, Germany, {clu, tschneider}@informatik.uni-bremen.de

Abstract. Finite model reasoning in expressive DLs such as *ALCQI* and *SHIQ* requires non-trivial algorithmic approaches that are substantially different from algorithms used for reasoning about unrestricted models. In contrast, finite model reasoning in the inexpressive fragment *DL-Lite_F* of *ALCQI* and *SHIQ* is algorithmically rather simple: using a TBox completion procedure that reverses certain terminological cycles, one can reduce finite subsumption to unrestricted subsumption. In this paper, we show that this useful technique extends all the way to the popular and much more expressive Horn-*SHIQ* fragment of *SHIQ*.

1 Introduction

Description logics (DLs) that include inverse roles and some form of counting such as functionality restrictions lack the finite model property (FMP) and, consequently, reasoning w.r.t. the class of finite models (*finite model reasoning*) does not coincide with reasoning w.r.t. the class of all models (*unrestricted reasoning*). On the one hand, this distinction is becoming increasingly important since DLs are nowadays regularly used in database applications, where models are generally assumed to be finite. On the other hand, finite model reasoning is rarely used in practice, mainly because for many popular DLs that lack the FMP, no algorithmic approaches to finite model reasoning are known that lend themselves towards efficient implementation.

Typical examples include the expressive DLs *ALCQI* and *SHIQ*, which are both a fragment of the OWL2 DL ontology language. While finite model reasoning in *ALCQI* and *SHIQ* are known to have the same complexity as unrestricted reasoning, namely EXPTIME-complete [9], the algorithmic approaches to the two cases are rather different. For unrestricted reasoning, there is a wide range of applicable algorithms such as tableau and resolution calculi, which often perform rather well in practical implementations. For finite model reasoning, all known approaches rely on the construction of some system of inequalities [3,9], and then solve this system over the integers; the crux is that the system of inequalities is of exponential size in the best case, and consequently it is far from obvious how to come up with efficient implementations. Note that the same is true for the two-variable fragment of first-order logic with counting quantifiers (C2), into which DLs such as *ALCQI* and *SHIQ* can be embedded [12,13], that is, all known approaches to finite model reasoning in C2 rely on solving (at least) exponentially large systems of inequalities.

Interestingly, the situation is quite different on the other end of the expressive power spectrum. *DL-Lite_F* is a very inexpressive DL that is used in database applications, but lacks the FMP because it still includes inverse roles and functionality restrictions.

Building on a technique that was developed in a database context by Cosmadakis, Kanellakis, and Vardi to decide the implication of inclusion dependencies and functional dependencies in the finite [4], Rosati has shown that finite model reasoning in $DL-Lite_{\mathcal{F}}$ can be reduced (in polynomial time) to unrestricted reasoning in $DL-Lite_{\mathcal{F}}$ [14]. In fact, the reduction is conceptually simple and relies on completing the TBox by finding certain cyclic inclusions and ‘reversing’ them. For example, the cycle

$$\exists r^- \sqsubseteq \exists s \quad \exists s^- \sqsubseteq \exists r \quad (\text{funct } r^-) \quad (\text{funct } s^-)$$

that consists of existential restrictions in the ‘forward direction’ and functionality statements in the ‘backwards direction’ would lead to the addition of the reversed cycle

$$\exists s \sqsubseteq \exists r^- \quad \exists r \sqsubseteq \exists s^- \quad (\text{funct } r) \quad (\text{funct } s).$$

As a consequence, finite model reasoning in $DL-Lite_{\mathcal{F}}$ does not require any new algorithmic techniques and can be implemented as efficiently as unrestricted reasoning. Given that $DL-Lite_{\mathcal{F}}$ is a very small fragment of \mathcal{ALCQI} and \mathcal{SHIQ} , these observations raise the question whether the cycle reversion technique extends also to larger fragments of these DLs. In particular, $DL-Lite_{\mathcal{F}}$ is a ‘Horn DL’, and such logics are well-known to be algorithmically more well-behaved than non-Horn DLs such as \mathcal{ALCQI} and \mathcal{SHIQ} . Maybe this is the reason why cycle reversion works for $DL-Lite_{\mathcal{F}}$?

In this paper, we show that the cycle reversion technique extends all the way to the expressive Horn- \mathcal{SHIQ} fragment of \mathcal{SHIQ} , which is rather popular in database applications [6,11,5,2] and properly extends $DL-Lite_{\mathcal{F}}$ and other relevant Horn fragments such as \mathcal{ELLF} . In particular, we show that finite satisfiability in Horn- \mathcal{SHIQ} can be reduced to unrestricted satisfiability in Horn- \mathcal{SHIQ} by completing the TBox with reversed cycles in the style of Cosmadakis et al. and of Rosati. While the reduction technique is essentially the same as for $DL-Lite_{\mathcal{F}}$, the construction of a finite model in the correctness proof is much more subtle and demanding. Another crucial difference to the $DL-Lite_{\mathcal{F}}$ case is that, when completing Horn- \mathcal{SHIQ} TBoxes, the cycles that have to be considered can be of exponential length, and thus the reduction is not polynomial. On first glance, this of course casts a doubt on the practical relevance of the proposed reduction. Still, we are confident that our approach will lead to implementable algorithms for finite model reasoning in Horn- \mathcal{SHIQ} . Specifically, the state-of-the-art of efficient reasoning in Horn description logics is to use so-called consequence based calculi, as introduced for Horn- \mathcal{SHIQ} in [7] and implemented for example in the reasoners CEL, CB, and ELK [1,7,8]. Instead of first completing the TBox and then handing over the completed TBox to a reasoner, it seems well possible to integrate the reversion of cycles directly as an inference rule into such a calculus. This avoids the detection of cycles by uninformed, brute-force search, and instead searches for cycles in the consequences that have already been computed by the calculus, anyway. Since the efficiency of consequence-based calculi are largely due to the fact that, for typical inputs, the set of derived consequences is relatively small, we expect that this will work well in practical applications. For now, though, we leave it as future work to work out the details of such a calculus.

Some proof details are deferred to the appendix of the long version of the paper, to be found at <http://www.informatik.uni-bremen.de/tdki/research/papers.html>

2 Preliminaries

The original definition of Horn- \mathcal{SHIQ} is based on a notion of polarity and somewhat unwieldy [6]; alternative definitions have been proposed later, see for example [10]. For brevity, we directly introduce Horn- \mathcal{SHIQ} TBoxes in a certain normal form similar to the one used in [7].

Let \mathbb{N}_C and \mathbb{N}_R be countably infinite and disjoint sets of concept names and role names. A Horn- \mathcal{SHIQ} TBox \mathcal{T} is a set of *concept inclusions* (CIs) that can take the following forms:

$$K \sqsubseteq A \quad K \sqsubseteq \perp \quad K \sqsubseteq \exists r.K' \quad K \sqsubseteq \forall r.K' \quad K \sqsubseteq (\leq 1 r K') \quad K \sqsubseteq (\geq n r K')$$

where K and K' denote a conjunction of concept names, A a concept name, r a (potentially inverse) role, and $n \geq 1$. Throughout the paper, we will deliberately confuse conjunctions of concept names and sets of concept names. The empty conjunction is abbreviated to \top .

It was observed in [7] that, for the purposes of deciding unrestricted satisfiability, the above form can be assumed without loss of generality; that is, every Horn- \mathcal{SHIQ} TBox \mathcal{T} conformant with the original definition in [6] can be converted in polynomial time into a TBox \mathcal{T}' in the above form such that for all concept names A in \mathcal{T} , we have that A is satisfiable w.r.t. \mathcal{T} iff A is satisfiable w.r.t. \mathcal{T}' . It is straightforward to verify that all necessary transformations, such as coding out role hierarchies and transitive roles do not rely on unrestricted models to be available and thus, the introduced TBox normal form can be assumed w.l.o.g. also for finite satisfiability.

The semantics for Horn- \mathcal{SHIQ} is given as usual in terms of interpretations \mathcal{I} . For a given TBox \mathcal{T} and a concept inclusion $C \sqsubseteq D$, we write $\mathcal{T} \models C \sqsubseteq D$ if $\mathcal{I} \models C \sqsubseteq D$ for all models \mathcal{I} of \mathcal{T} , and $\mathcal{T} \models_{\text{fin}} C \sqsubseteq D$ if the same holds for all finite models. We recall that, in Horn- \mathcal{SHIQ} , (un)satisfiability and subsumption can be mutually reduced to each other in polynomial time, and that this also holds in the finite. The following examples show that, in Horn- \mathcal{SHIQ} , finite and unrestricted reasoning do not coincide.

Example 1. Let $\mathcal{T} = \{ A \sqsubseteq \exists r.B, B \sqsubseteq \exists r.B, B \sqsubseteq (\leq 1 r^- \top), A \sqcap B \sqsubseteq \perp \}$. Then A is satisfiable w.r.t. \mathcal{T} , but not finitely satisfiable. In fact, when $d \in A^{\mathcal{I}}$ in some model \mathcal{I} of \mathcal{T} , then the CI $B \sqsubseteq \exists r.B$ and functionality assertion on r^- enforces an infinite chain $r(d, d_1), r(d_1, d_2), \dots$ with $d \in A^{\mathcal{I}}, d \notin B^{\mathcal{I}}$ and $d_2, d_3, \dots \in B^{\mathcal{I}}$.

Let $\mathcal{T}' = \{ A_1 \sqsubseteq \exists r.A_2, A_2 \sqsubseteq \exists r.(A_1 \sqcap B), \top \sqsubseteq (\leq 1 r^- \top) \}$. The reader might want to convince herself that $\mathcal{T}' \not\models A_1 \sqsubseteq B$, but $\mathcal{T}' \models_{\text{fin}} A_1 \sqsubseteq B$.

Eliminating At-Least Restrictions

The usual normal form for Horn- \mathcal{SHIQ} does not comprise at-least restrictions, that is, CIs of the form $K \sqsubseteq (\geq n r K')$ are not allowed. This is achieved by replacing each such CI with

$$K \sqsubseteq \exists r.B_i, \quad B_i \sqsubseteq K', \quad B_i \sqcap B_j \sqsubseteq \perp \quad 1 \leq i < j \leq n \quad (1)$$

where each B_i is a fresh concept name. If infinite models are admitted, it is quite easy to see that this translation preserves the satisfiability of all concept names in \mathcal{T} , exploiting

the tree model property of Horn- \mathcal{SHIQ} . For finite satisfiability, the same construction works, but a more refined argument is needed to show this.

Proposition 2. *Let \mathcal{T}' be obtained from \mathcal{T} by replacing the CI $K \sqsubseteq (\geq n r K')$ with the CIs (1), and let A be a concept name from \mathcal{T} . Then A is finitely satisfiable w.r.t. \mathcal{T} iff A is finitely satisfiable w.r.t. \mathcal{T}' .*

Proof. The “if” direction is trivial since every model of \mathcal{T}' is also a model of \mathcal{T} . For the “only if” direction, let \mathcal{I} be a finite model of \mathcal{T} with $A^{\mathcal{I}} \neq \emptyset$. We construct a finite model \mathcal{J} of \mathcal{T}' with $A^{\mathcal{J}} \neq \emptyset$ by taking n copies of \mathcal{I} and ‘rewiring’ all role edges across the concept names B_i can be interpreted in a non-conflicting way.

Specifically, since \mathcal{I} satisfies $K \sqsubseteq (\geq n r K')$ we can choose a function $\text{succ} : K^{\mathcal{I}} \times \{0, \dots, n-1\} \rightarrow \Delta^{\mathcal{I}}$ such that the following conditions are satisfied:

- for all $d \in K^{\mathcal{I}}$ and $i < n$: $(d, \text{succ}(d, i)) \in r^{\mathcal{I}}$ and $\text{succ}(d, i) \in (K')^{\mathcal{I}}$;
- for all $d \in K^{\mathcal{I}}$ and $i < j < n$: $\text{succ}(d, i) \neq \text{succ}(d, j)$.

Then define the desired interpretation \mathcal{J} by setting

$$\begin{aligned} \Delta^{\mathcal{J}} &= \{d_i \mid d \in \Delta^{\mathcal{I}} \text{ and } i < n\} \\ E^{\mathcal{J}} &= \{d_i \mid d \in E^{\mathcal{I}} \text{ and } i < n\} \quad \text{for all } E \in \mathbf{N}_C \setminus \{B_0, \dots, B_{n-1}\} \\ B_i^{\mathcal{J}} &= \{d_i \mid d \in \Delta^{\mathcal{I}}\} \quad \text{for all } i < n \\ s^{\mathcal{J}} &= \{(d_i, e_i) \mid (d, e) \in s^{\mathcal{I}} \text{ and } i < n\} \quad \text{for all } s \in \mathbf{N}_R \setminus \{r\} \\ r^{\mathcal{J}} &= \{(d_i, e_i) \mid (d, e) \in r^{\mathcal{I}}, i < n, \text{ and } d \notin K^{\mathcal{I}} \text{ or } e \neq \text{succ}(d, j) \text{ for any } j\} \\ &\quad \cup \{(d_i, e_{(i+j) \bmod n}) \mid (d, e) \in r^{\mathcal{I}}, i, j < n, \text{ and } e = \text{succ}(d, j)\} \end{aligned}$$

It remains to verify that \mathcal{J} is indeed a model of \mathcal{T}' . Clearly, the CIs in (1) are satisfied. Moreover, it is not hard to verify that all concept inclusions in \mathcal{T} are satisfied by \mathcal{J} , using the fact that \mathcal{I} is a model of \mathcal{T} and the construction of \mathcal{J} . \square

From now on, we can thus safely assume that TBoxes do not contain at-least restrictions. Note that the above translation is polynomial only if the numbers n in at-least restrictions are coded in unary. The same is of course true in unrestricted reasoning with Horn- \mathcal{SHIQ} , where typically the same normal form is used.

3 Reduction to Unrestricted Satisfiability

We give a reduction of finite satisfiability to unrestricted satisfiability based on the completion of TBoxes with certain reversed cycles. Let \mathcal{T} be a Horn- \mathcal{SHIQ} TBox. A *finmod cycle* in \mathcal{T} is a sequence $K_1, r_1, K_2, r_2, \dots, r_{n-1}, K_n$, with K_1, \dots, K_n conjunctions of concept names and r_1, \dots, r_{n-1} (potentially inverse) roles that satisfies $K_n = K_1$ and

$$\mathcal{T} \models K_i \sqsubseteq \exists r_i. K_{i+1} \text{ and } \mathcal{T} \models K_{i+1} \sqsubseteq (\leq 1 r_i^- K_i) \quad \text{for } 1 \leq i < n.$$

By *reversing* a finmod cycle $K_1, r_1, K_2, r_2, \dots, r_{n-1}, K_n$ in a TBox \mathcal{T} , we mean to extend \mathcal{T} with the concept inclusions

$$K_{j+1} \sqsubseteq \exists r_j^- . K_j \text{ and } K_j \sqsubseteq (\leq 1 r_j K_{j+1}) \quad \text{for } 1 \leq j < n.$$

The completion \mathcal{T}_f of a TBox \mathcal{T} is obtained from \mathcal{T} by exhaustively reversing finmod cycles.

Example 3. The TBox \mathcal{T}' from Example 1 entails (in unrestricted models)

$$A_1 \sqcap B \sqsubseteq \exists r. A_2, \quad A_2 \sqsubseteq \exists r. (A_1 \sqcap B), \quad A_2 \sqsubseteq (\leq 1 r^- A_1 \sqcap B), \quad A_2 \sqcap B \sqsubseteq (\leq 1 r^- A_1).$$

Thus, $(A_1 \sqcap B), r, A_2, r, (A_1 \sqcap B)$ is a finmod cycle in \mathcal{T}' , which is reversed to

$$A_2 \sqsubseteq \exists r^- . (A_1 \sqcap B), \quad A_1 \sqcap B \sqsubseteq \exists r^- . A_2, \quad A_1 \sqcap B \sqsubseteq (\leq 1 r A_2), \quad A_1 \sqsubseteq (\leq 1 r A_2 \sqcap B).$$

Another finmod cycle in \mathcal{T}' is A_1, r, A_2, r, A_1 , reversed to

$$A_2 \sqsubseteq \exists r^- . A_1, \quad A_1 \sqsubseteq \exists r^- . A_2, \quad A_2 \sqsubseteq (\leq 1 r A_1), \quad A_1 \sqsubseteq (\leq 1 r A_2).$$

Note that \mathcal{T}'_f contains $A_1 \sqsubseteq \exists r^- . A_2$, $A_2 \sqsubseteq \exists r. (A_1 \sqcap B)$, and $A_2 \sqsubseteq (\leq 1 r A_1)$. Consequently $\mathcal{T}'_f \models A_1 \sqsubseteq B$, in correspondence with $\mathcal{T}' \models_{\text{fin}} A_1 \sqsubseteq B$.

The following result is the main result of this paper. It shows that TBox completion indeed provides a reduction from finite satisfiability to unrestricted satisfiability.

Theorem 4. *Let \mathcal{T} be a Horn-SHLQ TBox and A a concept name. Then A is finitely satisfiable w.r.t. \mathcal{T} iff A is satisfiable w.r.t. the completion \mathcal{T}_f of \mathcal{T} .*

The “only if” direction of Theorem 4 is an immediate consequence of the observation that all CIs added by the TBox completion are actually entailed by the original TBox in finite models.

Lemma 5. *Let $K_1, r_1, \dots, r_{n-1}, K_n$ a finmod cycle in \mathcal{T} , then for every $1 \leq i < n$, $\mathcal{T} \models_{\text{fin}} K_{i+1} \sqsubseteq \exists r_i^- . K_i$ and $\mathcal{T} \models_{\text{fin}} K_i \sqsubseteq (\leq 1 r_i K_{i+1})$.*

Proof. We have to show that, if $K_1, r_1, \dots, r_{n-1}, K_n$ is a finmod cycle in \mathcal{T} and \mathcal{I} is a finite model of \mathcal{T} , then $K_i^{\mathcal{I}} \sqsubseteq (\leq 1 r_i K_{i+1})^{\mathcal{I}}$ and $K_{i+1}^{\mathcal{I}} \sqsubseteq (\exists r_i^- . K_i)^{\mathcal{I}}$ for $1 \leq i < n$. We first note that, by the semantics of Horn-SHLQ, we must have $|K_1^{\mathcal{I}}| \leq \dots \leq |K_n^{\mathcal{I}}|$, thus $K_n = K_1$ yields $|K_1^{\mathcal{I}}| = \dots = |K_n^{\mathcal{I}}|$. Fix some i with $1 \leq i < n$. Using $|K_i^{\mathcal{I}}| = |K_{i+1}^{\mathcal{I}}|$, $K_i^{\mathcal{I}} \sqsubseteq (\exists r_i . K_{i+1})^{\mathcal{I}}$, and $K_{i+1}^{\mathcal{I}} \sqsubseteq (\leq 1 r_i^- K_i)^{\mathcal{I}}$, it is now easy to verify that $K_i^{\mathcal{I}} \sqsubseteq (\leq 1 r_i K_{i+1})^{\mathcal{I}}$ and $K_{i+1}^{\mathcal{I}} \sqsubseteq (\exists r_i^- . K_i)^{\mathcal{I}}$, as required. \square

The “if” direction of Theorem 4 is much more demanding to prove. It requires to construct a finite model of A and \mathcal{T} based on the assumption that there is a (possibly infinite) model of A and \mathcal{T}_f . Such a construction is presented in the next section.

4 Constructing Finite Models

We show that the completion \mathcal{T}_f of \mathcal{T} captures all finite entailments of \mathcal{T} , that is, we prove the “if” direction of Theorem 4 above.

Assume that the concept name A is satisfiable w.r.t. \mathcal{T}_f . Let $\text{CN}(\mathcal{T})$ denote the set of concept names used in \mathcal{T} . A subset $t \subseteq \text{CN}(\mathcal{T})$ is a *type for \mathcal{T}* if there is a (potentially infinite) model \mathcal{I} of \mathcal{T} and a $d \in \Delta^{\mathcal{I}}$ such that $\text{tp}_{\mathcal{I}}(d) := \{A \in \text{CN}(\mathcal{T}) \mid d \in A^{\mathcal{I}}\}$

is identical with t . We use $\text{TP}(\mathcal{T})$ to denote the set of all types for \mathcal{T} . Our aim is to construct a finite model \mathcal{I} of \mathcal{T}_f (and thus also of \mathcal{T}) that realizes all types in $\text{TP}(\mathcal{T})$. Note that since A is satisfiable w.r.t. \mathcal{T}_f , there is a type t for \mathcal{T} with $A \in t$. Since this type is realized in the finite model \mathcal{I} of \mathcal{T} that we construct, it follows that A is finitely satisfiable w.r.t. \mathcal{T} as desired.

Before we give details of the construction of \mathcal{I} , we introduce some relevant notation and preliminary results. For all $t, t' \in \text{TP}(\mathcal{T}_f)$ and roles r , we write

- $t \rightarrow_r t'$ if $\mathcal{T}_f \models t \sqsubseteq \exists r.t'$ and t' is maximal with this property;
- $t \rightarrow_r^1 t'$ if $t \rightarrow_r t'$ and $\mathcal{T}_f \models t' \sqsubseteq (\leq 1 r^- t)$;
- $t \overset{1}{\leftrightarrow}_r t'$ if $t \rightarrow_r^1 t'$ and $t' \rightarrow_{r^-}^1 t$.
- $t \Rightarrow_r^1 t'$ if $t \rightarrow_r^1 t'$ and there are $s \subseteq t$ and $s' \subseteq t'$ such that $s \overset{1}{\leftrightarrow}_r s'$, but $\mathcal{T}_f \models t' \sqsubseteq \exists r^- .t$ does not hold.

A *type partition* is a set $P \subseteq \text{TP}(\mathcal{T})$ that is minimal with the following conditions:

- P is non-empty;
- if $t \in P$ and $t \overset{1}{\leftrightarrow}_r t'$, then $t' \in P$.

We set $P \prec P'$ if there are $t \in P$ and $t' \in P'$ with $t' \subsetneq t$. We will later be referring to the strict partial order that is obtained by taking the transitive closure of \prec , denoted by \prec^+ . A proof of the following observation can be found in the appendix.

Lemma 6. \prec^+ is a strict partial order.

As a last bit of notation, if $\lambda = s \overset{1}{\leftrightarrow}_r s'$, then we use λ^- to denote $s' \overset{1}{\leftrightarrow}_{r^-} s$.

4.1 Constructing the Model

We construct \mathcal{I} by starting with an initial interpretation and then exhaustively applying four *completion steps* that we denote with (c1) to (c4). While constructing the sequence, we will make sure that the following invariants are satisfied:

- (i1) for each $d \in \Delta^{\mathcal{I}}$, we have $\text{tp}_{\mathcal{I}}(d) \in \text{TP}(\mathcal{T}_f)$;
- (i2) if $(d, d') \in r^{\mathcal{I}}$, then $\text{tp}_{\mathcal{I}}(d) \rightarrow_r \text{tp}_{\mathcal{I}}(d')$ or $\text{tp}_{\mathcal{I}}(d') \rightarrow_{r^-} \text{tp}_{\mathcal{I}}(d)$;
- (i3) if $\mathcal{T}_f \models K \sqsubseteq (\leq 1 r K')$, then $\mathcal{I} \models K \sqsubseteq (\leq 1 r K')$.

We shall prove in Section 4.2 that each of the steps (c1) to (c4) indeed preserves these invariants.

The initial interpretation \mathcal{I} is defined by introducing an element for every type, interpreting the concept names in the obvious way, and interpreting all role names as empty: $\Delta^{\mathcal{I}} = \text{TP}(\mathcal{T}_f)$; $A^{\mathcal{I}} = \{t \in \text{TP}(\mathcal{T}_f) \mid A \in t\}$; $r^{\mathcal{I}} = \emptyset$. The four completion steps are described in detail below. We prefer to apply rules with smaller numbers, that is, if completion steps (ci) and (cj) are both applicable and $i < j$, then we apply (ci) first.

- (c1) Select a $d \in \Delta^{\mathcal{I}}$ such that $\text{tp}_{\mathcal{I}}(d) \Rightarrow_r^1 t$, and $d \notin (\exists r.t)^{\mathcal{I}}$.
Add a fresh domain element e , and modify the extension of concept and role names such that $\text{tp}_{\mathcal{I}}(e) = t$ and $(d, e) \in r^{\mathcal{I}}$.

- (c2) Select a type partition P that is minimal w.r.t. the order \prec^+ , a $\lambda = s \stackrel{1}{\leftrightarrow}_r s'$ with $s \in P$, and an element $d \in \Delta^{\mathcal{I}}$ such that $d \in s^{\mathcal{I}}$ and $d \notin (\exists r.s')^{\mathcal{I}}$.

For each $s \in P$, set $n_s = |\{d \in \Delta^{\mathcal{I}} \mid d \in s^{\mathcal{I}}\}|$. Let $n_{\max} = \max\{n_s \mid s \in P\}$. Reserve fresh domain elements

$$\Delta := \{d_{s,i} \mid s \in P \text{ and } n_s < i \leq n_{\max}\}.$$

For each $s \in P$, define a set of s -instances

$$I_s = \{d \in \Delta^{\mathcal{I}} \mid d \in s^{\mathcal{I}}\} \cup \{d_{s,i} \mid n_s < i \leq n_{\max}\}.$$

To proceed, we treat each $\lambda = s \stackrel{1}{\leftrightarrow}_r s'$ with $s, s' \in P$ separately. Thus, fix such a λ . Define

$$R_\lambda := \{(d, e) \in r^{\mathcal{I}} \mid d \in s^{\mathcal{I}} \text{ and } e \in s'^{\mathcal{I}}\}.$$

We first note that it is a consequence of invariant (i3) that

(*) the relation R_λ is functional and inverse functional.

In fact, let $(d, e_1), (d, e_2) \in R_\lambda$. Then $(d, e_1), (d, e_2) \in r^{\mathcal{I}}$, $d \in s^{\mathcal{I}}$, and $e_1, e_2 \in s'^{\mathcal{I}}$. By λ , we have $\mathcal{T}_f \models s \sqsubseteq (\leq 1 r s')$. Thus, (i3) yields $e_1 = e_2$. Inverse functionality can be shown analogously.

Let R_λ^1 be the domain of R_λ , and let R_λ^2 be the range. By (*), we have $|R_\lambda^1| = |R_\lambda^2|$. By definition of Δ , we have $|I_s| = |I_{s'}|$. Moreover, $R_\lambda^1 \subseteq I_s$ and $R_\lambda^2 \subseteq I_{s'}$. We can thus choose a bijection π_λ between $I_s \setminus R_\lambda^1$ and $I_{s'} \setminus R_\lambda^2$. Now extend \mathcal{I} as follows:

- add all domain elements in Δ ;
- extend $r^{\mathcal{I}}$ with π_λ , for each $\lambda = s \stackrel{1}{\leftrightarrow}_r s'$;
- extend $r^{\mathcal{I}}$ with the converse of π_λ , for each $\lambda = s \stackrel{1}{\leftrightarrow}_{r^-} s'$;
- interpret concept names so that $\text{tp}_{\mathcal{I}}(d_{s,i}) = s$ for all $d_{s,i} \in \Delta$.

- (c3) Select a $d \in \Delta^{\mathcal{I}}$ such that $\text{tp}_{\mathcal{I}}(d) \rightarrow_r^1 t$ and $d \notin (\exists r.t)^{\mathcal{I}}$. Add a fresh domain element e , and modify the extension of concept and role names such that $\text{tp}_{\mathcal{I}}(e) = t$ and $(d, e) \in r^{\mathcal{I}}$.
- (c4) Select a $d \in \Delta^{\mathcal{I}}$ such that $\text{tp}_{\mathcal{I}}(d) \rightarrow_r t$ and $d \notin (\exists r.t)^{\mathcal{I}}$. Add the edge (d, t) to $r^{\mathcal{I}}$, where t is the element for the type t introduced in the initial interpretation \mathcal{I} .

We briefly discuss the main effects of prioritizing the completion steps. It is important to prefer (c1) over (c2): together with the preference of type partitions that are minimal w.r.t. \prec^+ in (c2), this ensures that when (c2) is executed on type partition P , $\lambda = s \stackrel{1}{\leftrightarrow}_{r^-} s'$ with $s, s' \in P$, and $d \in I_s \setminus R_\lambda^1$, then not only $\text{tp}_{\mathcal{I}}(d) \supseteq s$, but actually $\text{tp}_{\mathcal{I}}(d) = s$. This central property, put as Lemma 7 below, is essential to guarantee preservation of invariants (i2) and (i3) by execution of (c2). Preferring (c1) and (c2) over (c3) ensures that, when (c3) is executed, then there are no $s \sqsubseteq \text{tp}_{\mathcal{I}}(d)$ and $s' \sqsubseteq t$ such that $s \stackrel{1}{\leftrightarrow}_r s'$; and preferring (c3) over (c4) ensures that, when (c4) is executed, then we have $\mathcal{T}_f \not\models \text{tp}_{\mathcal{I}}(d) \sqsubseteq (\leq 1 r t)$. These statements are provided here only to help in understanding the model construction. A formal proof is omitted at this point, but it can be recovered from the proofs given in the subsequent sections.

4.2 Satisfaction of Invariants

Application of (c1) preserves all invariants. It is obvious that the invariants (i1) and (i2) are preserved with each single application of (c1). We have to show that the same is true for (i3). Assume that completion processed $d \in \Delta^{\mathcal{I}}$ with $\text{tp}_{\mathcal{I}}(d) \Rightarrow_r^1 t$, and that e is the fresh domain element added. Assume to the contrary of what is to be shown that $\mathcal{T}_f \models K \sqsubseteq (\leq 1 r K')$ and there is a $e' \in \Delta^{\mathcal{I}}$ distinct from e such that $d \in K^{\mathcal{I}}$, $(d, e') \in r^{\mathcal{I}}$, and $e, e' \in K'^{\mathcal{I}}$. According to (i2), we distinguish the following cases:

- $\text{tp}_{\mathcal{I}}(d) \rightarrow_r \text{tp}_{\mathcal{I}}(e')$
Then $\mathcal{T}_f \models \text{tp}_{\mathcal{I}}(d) \sqsubseteq \exists r. \text{tp}_{\mathcal{I}}(e')$ and $\text{tp}_{\mathcal{I}}(e')$ is maximal with this property. Since $\text{tp}_{\mathcal{I}}(d) \rightarrow_r t$, we additionally have $\mathcal{T}_f \models \text{tp}_{\mathcal{I}}(d) \sqsubseteq \exists r. t$. Since $K \subseteq \text{tp}_{\mathcal{I}}(d)$ and $e, e' \in K'^{\mathcal{I}}$ implies $K' \subseteq \text{tp}_{\mathcal{I}}(e') \cap t$, a simple semantic argument shows that $\mathcal{T}_f \models K \sqsubseteq \exists r. (\text{tp}_{\mathcal{I}}(e') \cup t)$. The maximality of $\text{tp}_{\mathcal{I}}(e')$ thus implies $t \subseteq \text{tp}_{\mathcal{I}}(e')$, in contradiction to the fact that $d \notin (\exists r. t)^{\mathcal{I}}$ was true before the completion step.
- $\text{tp}_{\mathcal{I}}(e') \rightarrow_{r^-} \text{tp}_{\mathcal{I}}(d)$.
Then $\mathcal{T}_f \models \text{tp}_{\mathcal{I}}(e') \sqsubseteq \exists r^- . \text{tp}_{\mathcal{I}}(d)$ and, additionally, we have $\mathcal{T}_f \models \text{tp}_{\mathcal{I}}(d) \sqsubseteq \exists r. t$. Since $K \subseteq \text{tp}_{\mathcal{I}}(d)$ and $K' \subseteq \text{tp}_{\mathcal{I}}(e') \cap t$, a simple semantic argument shows that $\mathcal{T}_f \models \text{tp}_{\mathcal{I}}(e') \sqsubseteq t$. Since $\text{tp}_{\mathcal{I}}(e')$ is a type for \mathcal{T}_f by (i1), it follows that $t \subseteq \text{tp}_{\mathcal{I}}(e')$. This contradicts the fact that $d \notin (\exists r. t)^{\mathcal{I}}$ was true before the completion step.

Application of (c2) preserves all invariants. Invariant (i1) is clearly preserved by each single application of (c2). We have to prove that the same is true for (i2) and (i3). First, we show that the following property holds:

Lemma 7. *If $\lambda = s \overset{1}{\leftrightarrow}_r s'$ and $d \in I_s \setminus R_{\lambda}^1$, then $\text{tp}_{\mathcal{I}}(d) = s$.*

To show that (i2) is preserved by step (c2), consider an arbitrary pair $(d, d') \in r^{\mathcal{I}}$ that has been added in a step (c2). Hence $\pi_{\lambda}(d) = d'$, i.e., there is some $\lambda = s \overset{1}{\leftrightarrow}_r s'$ such that $d \in I_s \setminus R_{\lambda}^1$ and $d' \in I_{s'} \setminus R_{\lambda}^2$. From Lemma 7, we obtain $\text{tp}_{\mathcal{I}}(d) = s$. Analogously, considering $\lambda' = s' \overset{1}{\leftrightarrow}_{r^-} s$ and the fact $d' \in I_{s'} \setminus R_{\lambda}^2 = I_{s'} \setminus R_{\lambda'}^1$, we obtain from Lemma 7 that $\text{tp}_{\mathcal{I}}(d') = s'$. Consequently, $s \overset{1}{\leftrightarrow}_r s'$ yields $\text{tp}_{\mathcal{I}}(d) \rightarrow_r \text{tp}_{\mathcal{I}}(d')$ and $\text{tp}_{\mathcal{I}}(d') \rightarrow_{r^-} \text{tp}_{\mathcal{I}}(d)$.

We now show that (i3) is preserved by step (c2). Let $\mathcal{T}_f \models K \sqsubseteq (\leq 1 r K')$, and assume to the contrary of what is to be shown that, after some application of (c2), there are $(d, d_1), (d, d_2) \in r^{\mathcal{I}}$ with $d \in K^{\mathcal{I}}$, $d_1, d_2 \in K'^{\mathcal{I}}$, and $d_1 \neq d_2$. We distinguish the following cases:

- (d, d_1) was added by an application of (c2), (d, d_2) was not. By the former, there is $\lambda = s \overset{1}{\leftrightarrow}_r s'$ such that $d \in I_s \setminus R_{\lambda}^1$, $d_1 \in I_{s'} \setminus R_{\lambda}^2$, and $(d, d_1) \in \pi_{\lambda}$. By Lemma 7, $d \in I_s \setminus R_{\lambda}^1$ yields $\text{tp}_{\mathcal{I}}(d) = s$. Moreover, $d_1 \in I_{s'} \setminus R_{\lambda}^2$ implies $d_1 \in I_{s'} \setminus R_{\lambda^-}^1$, and thus another application of Lemma 7 yields $\text{tp}_{\mathcal{I}}(d_1) = s'$.
Since (d, d_2) was not added by step (c2), (i2) gives the following subcases:
 - $\text{tp}_{\mathcal{I}}(d) \rightarrow_r \text{tp}_{\mathcal{I}}(d_2)$. Thus $\mathcal{T}_f \models \text{tp}_{\mathcal{I}}(d) \sqsubseteq \exists r. \text{tp}_{\mathcal{I}}(d_2)$ and $\text{tp}_{\mathcal{I}}(d_2)$ is maximal with this property. Since $\text{tp}_{\mathcal{I}}(d) = s$ and by λ , $\mathcal{T}_f \models \text{tp}_{\mathcal{I}}(d) \sqsubseteq \exists r. \text{tp}_{\mathcal{I}}(d_2)$.
Using the facts that $\mathcal{T}_f \models K \sqsubseteq (\leq 1 r K')$, $K \subseteq \text{tp}_{\mathcal{I}}(d) = s$, $K' \subseteq \text{tp}_{\mathcal{I}}(d_2)$,

and $K' \subseteq \text{tp}_{\mathcal{I}}(d_1) = s'$, an easy semantic argument shows that $\mathcal{T}_f \models \text{tp}_{\mathcal{I}}(d) \sqsubseteq \exists r.(\text{tp}_{\mathcal{I}}(d_2) \cup s')$. The maximality of $\text{tp}_{\mathcal{I}}(d_2)$ thus yields $s' \subseteq \text{tp}_{\mathcal{I}}(d_2)$. Thus, $d \in (\exists r.s')^{\mathcal{I}}$ was true before step (c2) was applied, which is a contradiction to $d \notin R_{\lambda}^1$.

- $\text{tp}_{\mathcal{I}}(d_2) \rightarrow_{r^-} \text{tp}_{\mathcal{I}}(d)$. Then $\mathcal{T}_f \models \text{tp}_{\mathcal{I}}(d_2) \sqsubseteq \exists r^- .s$. By λ , we have $\mathcal{T}_f \models s \sqsubseteq \exists r.s'$. Since $K \subseteq s$, $K \subseteq \text{tp}_{\mathcal{I}}(d_2)$, $K \subseteq \text{tp}_{\mathcal{I}}(d_1) = s'$, and $\mathcal{T}_f \models K \sqsubseteq (\leq 1 \ r \ K')$, a simple semantic argument shows that $s' \subseteq \text{tp}_{\mathcal{I}}(d_2)$. This again means that $d \in (\exists r.s')^{\mathcal{I}}$ was true before step (c2) was applied, in contradiction to $d \notin R_{\lambda}^1$.
- both (d, d_1) and (d, d_2) were added by an application of (c2). Then there are λ_1 and λ_2 , such that, for $i \in \{1, 2\}$, we have

$$\lambda_i = s_i \stackrel{1}{\leftrightarrow}_r s'_i, \quad (d, d_i) \in \pi_{\lambda_i}, \quad d \in I_{s_i} \setminus R_{\lambda_i}^1, \quad d_i \in I_{s'_i} \setminus R_{\lambda_i}^2.$$

Applying Lemma 7 to λ_i and $d \in I_{s_i} \setminus R_{\lambda_i}^1$ yields $s_i = \text{tp}_{\mathcal{I}}(d)$, for $i \in \{1, 2\}$. Consequently, $s_1 = s_2$. We next show $s'_1 = s'_2$, thus $\lambda_1 = \lambda_2$.

For uniformity, we use s to denote s_1 and s_2 . From λ_i , we obtain $\mathcal{T}_f \models s \sqsubseteq \exists r.s_i$ and s_i is maximal with this property, for $i \in \{1, 2\}$. Note that $d_i \in I_{s'_i} \setminus R_{\lambda_i}^2$ implies $d_i \in I_{s'_i} \setminus R_{\lambda_i}^1$. Applications of Lemma 7 to λ_i^- and $d_i \in I_{s'_i} \setminus R_{\lambda_i}^1$ yield $\text{tp}_{\mathcal{I}}(d_i) = s_i$. Using the facts that $\mathcal{T}_f \models s \sqsubseteq \exists r.s_i$ for $i \in \{1, 2\}$, $K \subseteq \text{tp}_{\mathcal{I}}(d) = s$, $K' \subseteq \text{tp}_{\mathcal{I}}(d_i) = s_i$ for $i \in \{1, 2\}$, and $\mathcal{T}_f \models K \sqsubseteq (\leq 1 \ r \ K')$, an easy semantic argument shows that $\mathcal{T}_f \models s \sqsubseteq \exists r.(s_1 \cup s_2)$. The maximality of s_1 and s_2 thus implies $s_1 = s_2$ as desired.

Hence, $\lambda_1 = \lambda_2$ and $(d, d_1), (d, d_2) \in \pi_{\lambda_1}$. Since π_{λ_1} is a bijection, we obtain $d_1 = d_2$, a contradiction.

Application of (c3) and (c4) preserves all invariants. It is obvious that the invariants (i1) and (i2) are preserved with each single application of (c3) or (c4). To show that the same is true for (i3), we can use the same proof as for (c1) because the assumptions of (c3) and (c4) differ from that of (c1) in weakening $\text{tp}_{\mathcal{I}}(d) \Rightarrow_r^1 t$ to $\text{tp}_{\mathcal{I}}(d) \rightarrow_r^1 t$ and $\text{tp}_{\mathcal{I}}(d) \rightarrow_r t$, respectively, which is sufficient for that proof.

4.3 Termination of Model Construction

We show that the constructed interpretation \mathcal{I} is indeed finite.

Proposition 8. $\Delta^{\mathcal{I}}$ is finite.

Proof. To analyze the termination of the construction of \mathcal{I} , we associate a certain directed tree $T = (V, E)$ with the model \mathcal{I} that makes explicit the way in which \mathcal{I} was constructed. Note that only the completion steps (c1) to (c3) introduce new domain elements and that (c1) and (c3) introduce a single new element with each application whereas (c2) introduces a whole (finite) set of fresh elements. Also note that each application of a completion step is triggered by a single domain element d for which some existential restriction is not yet satisfied. Now, the tree T is defined as follows:

- V consists of all subsets of $\Delta^{\mathcal{I}}$ that were introduced together by a single application of one of the completion steps (c1) to (c3); additionally, the set of all elements in the initial interpretation \mathcal{I} is a node in V (in fact, the root node);
- the edge (v, v') is included in E if the elements in v' were introduced by an application of a completion step to an element d of v . We call this element the *trigger* of v' and denote it with $d_{v'}$.

To show that $\Delta^{\mathcal{I}}$ is finite, it clearly suffices to show that V is finite. The outdegree of T is finite since every rule application introduces only finitely many elements. By König's Lemma, it thus remains to show that T is of finite depth. We first note that an easy analysis of the completion steps (c1) to (c3) reveals the following property:

- (*) if $(v_1, v_2), (v_2, v_3) \in E$, then there are $d_0, \dots, d_k \in v_2$ and roles r_0, \dots, r_{k-1} s.t.
- $d_0 = d_{v_2} \in v_1, d_1, \dots, d_k \in v_2$, and $d_k = d_{v_3} \in v_2$;
 - $\text{tp}_{\mathcal{I}}(d_i) \xrightarrow{1}_{r_i} \text{tp}_{\mathcal{I}}(d_{i+1})$ for all $i < k$.

Now assume towards a contradiction that the depth of T is larger than $2|\text{TP}(\mathcal{T}_{\bar{f}})| + 1$ and choose a concrete path $v_1 \cdots v_n$ with v_1 the root of T and $n > 2|\text{TP}(\mathcal{T}_{\bar{f}})| + 1$. This path gives rise to a corresponding sequence of triggers d_{v_1}, \dots, d_{v_n} . Since the length of this sequence exceeds $2|\text{TP}(\mathcal{T}_{\bar{f}})|$, there must be i, j with $1 \leq i < j \leq n$ and such that $\text{tp}_{\mathcal{I}}(d_{v_i}) = \text{tp}_{\mathcal{I}}(d_{v_j})$ and $j > i + 1$. By applying (*) multiple times, we obtain a sequence of domain elements d_0, \dots, d_k and roles r_0, \dots, r_{k-1} such that

1. $d_0 = d_{v_i} \in v_{i-1}, d_1 \in v_i$, and $d_p = d_{v_j} \in v_{j-1}$;
2. $\text{tp}_{\mathcal{I}}(d_\ell) \xrightarrow{1}_{r_\ell} \text{tp}_{\mathcal{I}}(d_{\ell+1})$ for $\ell < k$.
3. d_0, \dots, d_k contains all elements d_{v_i}, \dots, d_{v_j} .

Since $\text{tp}_{\mathcal{I}}(d_{v_i}) = \text{tp}_{\mathcal{I}}(d_{v_j})$ and by Point 2, $\text{tp}_{\mathcal{I}}(d_0), r_0, \dots, r_{k-1}, \text{tp}_{\mathcal{I}}(d_k)$ is a finmod cycle in $\mathcal{T}_{\bar{f}}$. Since all finmod cycles in $\mathcal{T}_{\bar{f}}$ have been reversed, we have

$$\text{tp}_{\mathcal{I}}(d_0) \xrightarrow{1\leftarrow 1}_{r_0} \text{tp}_{\mathcal{I}}(d_1) \xrightarrow{1\leftarrow 1}_{r_1} \cdots \xrightarrow{1\leftarrow 1}_{r_{n-1}} \text{tp}_{\mathcal{I}}(d_n). \quad (\dagger)$$

In the appendix, we prove the following claim:

Claim. If step (c1) or step (c3) is triggered by $d \in \Delta^{\mathcal{I}}$ and generates a new element $e \in \Delta^{\mathcal{I}}$, then there is no role r such that $\text{tp}_{\mathcal{I}}(d) \xrightarrow{1\leftarrow 1}_r \text{tp}_{\mathcal{I}}(e)$.

Since $d_{v_i} \in v_{i-1}$ and $d_1 \in v_i$, d_1 was generated by the application of a completion step triggered by d_0 . By (\dagger) and the claim, this completion step must be (c2). By definition of (c2) and (\dagger) , all elements d_1, \dots, d_k have been introduced by exactly this application of (c2). This leads to a contradiction: we have $d_1 \in v_i$ and $d_k \in v_{j-1}$, and since $j > i + 1, v_i \neq v_{j-1}$. Consequently, d_1 and d_k were introduced by different applications of completion steps. \square

4.4 Correctness of Model Construction

To complete the proof of the “if” direction of Theorem 4, it remains to show the following.

Proposition 9. \mathcal{I} is a model of \mathcal{T}_f .

Proof. We show that for every axiom $K \sqsubseteq C \in \mathcal{T}_f$, we have that $\mathcal{I} \models K \sqsubseteq C$. We distinguish the following cases:

- $C = A$. Let $d \in K^{\mathcal{I}}$. Then $K \subseteq \text{tp}_{\mathcal{I}}(d)$ and by (i1) $\text{tp}_{\mathcal{I}}(d) \in \text{TP}(\mathcal{T}_f)$. Since $\mathcal{T}_f \models K \sqsubseteq A$, this yields $A \in \text{tp}_{\mathcal{I}}(d)$ and thus $d \in A^{\mathcal{I}}$.
- $C = \perp$. Follows from (i1). Indeed since for every $d \in \Delta^{\mathcal{I}}$, $\text{tp}(d) \in \text{TP}(\mathcal{T}_f)$, $K^{\mathcal{I}} = \emptyset$.
- $C = \exists r.K'$. Let $d \in K^{\mathcal{I}}$. Then we have that $K \subseteq \text{tp}_{\mathcal{I}}(d)$. Since $\mathcal{T}_f \models K \sqsubseteq \exists r.K'$, we have that $\text{tp}_{\mathcal{I}}(d) \rightarrow_r t'$ for some t' with $K' \subseteq t'$. Thus there is some d' with $K' \subseteq \text{tp}_{\mathcal{I}}(d')$ such that $(d, d') \in r^{\mathcal{I}}$ was added by an application of (c1), (c2), (c3) or (c4). In fact, if no such d' is added by (c1) to (c3), then the edge $(d, d') \in r^{\mathcal{I}}$ will clearly be added by (c4).
- $C = \forall r.K'$. Let $d \in K^{\mathcal{I}}$ and $(d, d') \in r^{\mathcal{I}}$. We have that $K \subseteq \text{tp}_{\mathcal{I}}(d)$. Further, by (i2), we can distinguish the following cases:
 - $\text{tp}_{\mathcal{I}}(d) \rightarrow_r \text{tp}_{\mathcal{I}}(d')$. Then $\mathcal{T}_f \models \text{tp}_{\mathcal{I}}(d) \sqsubseteq \exists r.\text{tp}(d')$ and $\text{tp}(d')$ is maximal with this property. Since $\mathcal{T}_f \models K \sqsubseteq \forall r.K'$, we have that $\mathcal{T}_f \models \text{tp}_{\mathcal{I}}(d) \sqsubseteq \exists r.\text{tp}(d') \cup K'$, and the maximality of $\text{tp}(d')$ yields $K' \subseteq \text{tp}(d')$, and thus $d' \in K'^{\mathcal{I}}$.
 - $\text{tp}_{\mathcal{I}}(d') \rightarrow_{r^-} \text{tp}_{\mathcal{I}}(d)$. Then $\mathcal{T}_f \models \text{tp}_{\mathcal{I}}(d') \sqsubseteq \exists r^-. \text{tp}_{\mathcal{I}}(d)$. Together with $\mathcal{T}_f \models K \sqsubseteq \forall r.K'$, we obtain $\mathcal{T}_f \models \text{tp}_{\mathcal{I}}(d') \sqsubseteq K'$. Since $\text{tp}_{\mathcal{I}}(d') \in \text{TP}(\mathcal{T}_f)$ by (i1), we obtain $K' \subseteq \text{tp}_{\mathcal{I}}(d')$ and thus $d' \in K'^{\mathcal{I}}$.
- $C = (\leq 1 r K)'$. Follows from (i3). □

5 Conclusion

We have presented a reduction from finite satisfiability in Horn- \mathcal{SHIQ} to unrestricted satisfiability, extending the technique introduced for $DL\text{-Lite}_{\mathcal{F}}$ in [14]. As discussed in the introduction, we believe that our technique is a more suitable basis for efficient implementation than the techniques for full \mathcal{ALCQI} and \mathcal{SHIQ} based on exponentially large systems of inequalities.

As future work, we plan to develop a consequence based calculus for finite satisfiability in Horn- \mathcal{SHIQ} and to extend the results obtained in this paper from finite satisfiability to answering conjunctive queries over ABoxes, assuming finite models. While we believe that the constructions given in this paper can be easily extended to instance query answering over ABoxes, the treatment of full conjunctive queries requires significant modification of the model construction.

References

1. F. Baader, C. Lutz, and B. Suntisrivaraporn. CEL – a polynomial-time reasoner for life science ontologies. In *Proc. of IJCAR-06*, volume 4130 of *LNCS*, pages 287–291. Springer, 2006.
2. M. Bienvenu, C. Lutz, and F. Wolter. First-order rewritability of atomic queries in Horn description logics. In *Proc. of IJCAI-13*. IJCAI/AAAI, 2013.
3. D. Calvanese. Finite model reasoning in description logics. In *Proc. of KR-96*, pages 292–303. Morgan Kaufmann, 1996.
4. S.S. Cosmadakis, P.C. Kanellakis, and M.Y. Vardi. Polynomial-time implication problems for unary inclusion dependencies. *J. ACM*, 37(1):15–46, 1990.
5. T. Eiter, M. Ortiz, M. Šimkus, T.-K. Tran, and G. Xiao. Towards practical query answering for Horn-*SHIQ*. In *Proc. of DL-12*, volume 846 of *CEUR-WS.org*, 2012.
6. U. Hustadt, B. Motik, and U. Sattler. Reasoning in description logics by a reduction to disjunctive datalog. *J. Autom. Reasoning*, 39(3), 2007.
7. Y. Kazakov. Consequence-driven reasoning for Horn *SHIQ* ontologies. In *Proc. of IJCAI-09*, pages 2040–2045, 2009.
8. Y. Kazakov, M. Krötzsch, and F. Simančík. Unchain my \mathcal{EL} reasoner. In *Proc. of DL-11*, volume 745 of *CEUR-WS.org*, 2011.
9. C. Lutz, U. Sattler, and L. Tendera. The complexity of finite model reasoning in description logics. *Information and Computation*, 199:132–171, 2005.
10. C. Lutz and F. Wolter. Non-uniform data complexity of query answering in description logics. In *Proc. of KR-12*. AAAI Press, 2012.
11. M. Ortiz, S. Rudolph, and M. Šimkus. Query answering in the Horn fragments of the description logics *SHOIQ* and *SRHOIQ*. In *Proc. of IJCAI-11*, pages 1039–1044. IJCAI/AAAI, 2011.
12. L. Pacholski, W. Szwast, and L. Tendera. Complexity results for first-order two-variable logic with counting. *SIAM J. Comput.*, 29(4):1083–1117, 2000.
13. I. Pratt-Hartmann. Complexity of the two-variable fragment with counting quantifiers. *J. of Logic, Language and Information*, 14(3):369–395, 2005.
14. R. Rosati. Finite model reasoning in DL-Lite. In *Proc. of ESWC-08*, volume 5021 of *LNCS*, pages 215–229. Springer, 2008.