Conservative Rewritability of Description Logic TBoxes

Boris Konev[1], **Carsten Lutz**[2], **Frank Wolter**[1] and **Michael Zakharyaschev**[3]
[1]Univ. of Liverpool, UK [2]Univ. of Bremen, Germany [3]Birkbeck, Univ. of London, UK
{konev,wolter}@liverpool.ac.uk clu@informatik.uni-bremen.de michael@dcs.bbk.ac.uk

Abstract

We investigate the problem of conservative rewritability of a TBox $\mathcal{T}$ in a description logic (DL) $\mathcal{L}$ into a TBox $\mathcal{T}'$ in a weaker DL $\mathcal{L}'$. We focus on model-conservative rewritability ($\mathcal{T}'$ entails $\mathcal{T}$ and all models of $\mathcal{T}$ are expandable to models of $\mathcal{T}'$), subsumption-conservative rewritability ($\mathcal{T}'$ entails $\mathcal{T}$ and all subsumptions in the signature of $\mathcal{T}$ entailed by $\mathcal{T}'$ are entailed by $\mathcal{T}$), and standard DLs between $\mathcal{ALC}$ and $\mathcal{ALCQI}$. We give model-theoretic characterizations of conservative rewritability via bisimulations, inverse p-morphisms and generated subinterpretations, and use them to obtain a few rewriting algorithms and complexity results for deciding rewritability.

1 Introduction

Over the past 30 years, a multitude of description logics (DLs) have been designed, investigated, and used in practice as ontology languages. The introduction of new DLs has been driven by (*i*) the need for additional expressive power (e.g., transitive roles in the 1990s), and (*ii*) applications that require efficient reasoning of a novel type (e.g., ontology-based data access in the 2000s). While the resulting flexibility in choosing DLs has had the positive effect of making DLs available for a large number of domains and applications, it has also led to the development of ontologies with language constructors that are not really required to represent their knowledge. A 'not required' constructor can mean different things here, ranging from the high-level 'this domain can be represented in an adequate way in a weaker DL' to the very concrete 'this ontology is logically equivalent to an ontology in a weaker DL'. In this paper, we take the latter understanding as a starting point. Equivalent rewritability of a DL ontology (TBox) to a weaker language has been investigated by Lutz, Piro, and Wolter [2011] who established model-theoretic characterizations in terms of (various types of) global bisimilations and applied them to the problem of deciding equivalent rewritability. However, equivalent rewritability seems to be an unnecessarily strong condition for multiple applications where *fresh* symbols can be used in rewritings.

Therefore, in this paper, we propose a more flexible notion of *conservative rewritability* that allows the use of fresh symbols in a rewriting $\mathcal{T}'$ of a given TBox $\mathcal{T}$. We demand that $\mathcal{T}'$ entails $\mathcal{T}$. On the other hand, to avoid uncontrolled additional consequences of $\mathcal{T}'$, we also require that (*i*) it does not entail any new subsumptions in the signature of $\mathcal{T}$, or even that (*ii*) every model of $\mathcal{T}$ can be expanded to a model of $\mathcal{T}'$. The latter type of conservative rewriting is known as *model-conservative extension* [Konev et al., 2013], and we call a TBox $\mathcal{T}$ model-conservatively $\mathcal{L}$-rewritable if there is a model-conservative extension of $\mathcal{T}$ in the DL $\mathcal{L}$. The former type is known as a *subsumption* or *deductive conservative extension* [Ghilardi, Lutz, and Wolter, 2006] and, given a DL $\mathcal{L}$, an $\mathcal{L}$ TBox $\mathcal{T}$ and a weaker DL $\mathcal{L}'$, we call $\mathcal{T}$ subsumption-conservatively $\mathcal{L}'$-rewritable if there is a TBox $\mathcal{T}'$ in $\mathcal{L}'$ such that $\mathcal{T}'$ entails the same $\mathcal{L}$-subsumptions in the signature of $\mathcal{T}$ as $\mathcal{T}$. Model-conservative rewritability is a more robust notion as it is language-independent and not only leaves unchanged the entailed subsumptions of the original TBox but also, for example, certain answers in case the ontologies are used to access data.

The main aim of this paper is to show that, for many important DLs, model- and subsumption-conservative rewritabilities can be characterized in terms of natural model-theoretic preservation conditions. In fact, the role played by global bisimilations for equivalent rewritability is now played by generated subinterpretations and p-morphisms (or bounded morphisms), that is, functional bisimilations introduced in modal logic as basic truth-preserving operations on Kripke frames and models [Goranko and Otto, 2006]. We also observe that, in some cases, these characterizations give rise to rewriting algorithms and complexity bounds for deciding conservative rewritability. The latter results are in sharp contrast to the fact that it is typically undecidable whether a given TBox is a model-conservative rewriting of another TBox [Lutz and Wolter, 2010; Konev et al., 2013]. We focus on standard DLs between $\mathcal{ALC}$ and $\mathcal{ALCQI}$, but also briefly consider rewritings into the lightweight DL *DL-Lite*$_{horn}$.

Our model-theoretic characterizations are summarized in Table 1, where the criteria for equivalent rewritability are taken from [Lutz, Piro, and Wolter, 2011]. Thus, for example, model-conservative $\mathcal{ALCI}$-to-$\mathcal{ALC}$ rewritability coincides with subsumption-conservative $\mathcal{ALCI}$-to-$\mathcal{ALC}$-rewritability, and both are characterized by preservation under generated subinterpretations or, equivalently, inverse p-morphisms. In contrast, model-conservative $\mathcal{ALCQ}$-to-$\mathcal{ALC}$ rewritabil-

Rewritability	Equivalent	Model Conservative	Subsumption Conservative
$\mathcal{ALCI}$-to-$\mathcal{ALC}$	global bisimulations	generated subinterpretations/p-morphisms^{-1}	
$\mathcal{ALCQI}$-to-$\mathcal{ALCQ}$	global counting bisimulations	counting p-morphisms^{-1}	
$\mathcal{ALCQ}$-to-$\mathcal{ALC}$	global bisimulations		p-morphisms^{-1}
$\mathcal{ALCQI}$-to-$\mathcal{ALCI}$	global i-bisimulations	?	i-p-morphisms^{-1}
$\mathcal{ALCQI}$-to-$\mathcal{ALC}$	global bisimulations	?	p-morphisms^{-1}
$\mathcal{ALCI}$-to-*DL-Lite*$_{horn}$	products and succ-simulations		

Table 1: Model-theoretic characterizations of rewritability.

ity coincides with equivalent $\mathcal{ALCQ}$-to-$\mathcal{ALC}$ rewritability, but not with subsumption-conservative $\mathcal{ALCQ}$-to-$\mathcal{ALC}$ rewritability. The situation is yet again different for $\mathcal{ALCI}$-to-*DL-Lite*$_{horn}$ rewritability, in which case all three notions coincide. The question marks indicate two cases where characterizations are unknown.

An in-depth exploration of the applicability of our model-theoretic characterizations is beyond the scope of this paper. We only mention in passing three cases that come naturally along with the preservation criteria. Thus, we show that the preservation conditions for $\mathcal{ALCI}$-to-$\mathcal{ALC}$ rewritability are decidable in ExpTime and give an algorithm constructing polynomial-size rewritings, while those for model-conservative and subsumption-conservative $\mathcal{ALCQ}$-to-$\mathcal{ALC}$ rewritabilities give rise to 2ExpTime decision algorithms.

Related work. Conservative rewritings of TBoxes are ubiquitous in DL research. For example, transformations of TBoxes into normal forms are often model-conservative [Baader, Brandt, and Lutz, 2005; Kazakov, 2009]. We note, however, that some well known DL rewritings introducing fresh symbols that are used as a pre-processing step in reasoning [Ding, Haarslev, and Wu, 2007; Carral et al., 2014b; 2014a] or to prove complexity results for reasoning [De Giacomo, 1995] are not conservative rewritings but only satisfiability preserving. There has been significant work on rewritings of ontology-mediated queries (OMQs), which preserve their certain answers, into datalog or OMQs in weaker DLs [Kaminski and Cuenca Grau, 2013; Bienvenu et al., 2014]. It seems that, from a technical viewpoint, rewritability of OMQs is not related to TBox conservative rewritability. Baader [1996] considers the expressive power of DLs and corresponding notions of rewritability based on a variant of model-conservative extension and discusses the relationship to subsumption-conservative extensions. Another closely related problem is TBox approximation. In this case, rather than aiming at a conservative rewriting, the aim is to compute a TBox in a weaker DL that approximates the consequences of the original TBox [Ren, Pan, and Zhao, 2010; Console et al., 2014].

Detailed proofs can be found in [Konev et al., 2016].

2 Conservative Rewritability

In DLs, *concepts* and *roles* are defined inductively starting from countably infinite sets $\mathsf{N_C}$ of *concept names* and $\mathsf{N_R}$ of *role names* and using a set of constructors. The constructors available in $\mathcal{ALCQI}$ are shown in the table below, where the formation of inverse roles is the only role constructor and the remaining four are concept constructors. The third column defines the *extensions* of roles and concepts with these constructors in an interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\cdot^{\mathcal{I}}$ maps each concept name A to a subset $A^{\mathcal{I}}$ of the domain $\Delta^{\mathcal{I}}$ of $\mathcal{I}$, and each role name r to $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. In the table, r stands for a role (i.e., a role name or its inverse), A, B for concept names, and C, D for (possibly compound) concepts; $r^{\mathcal{I}}(d) = \{d' \mid (d, d') \in r^{\mathcal{I}}\}$ and $|\Delta|$ is the cardinality of a set Δ. As usual, we define $\top$, $\bot$, $\sqcup$, $\rightarrow$ and $\leftrightarrow$ as standard Boolean abbreviations, $\exists r.C$ (*existential restriction*) as an abbreviation for $(\geqslant 1\ r\ C)$, and $\forall r.C$ (*universal restriction*) for $(\leqslant 0\ r\ \neg C)$. In the sublanguage $\mathcal{ALCQ}$ of $\mathcal{ALCQI}$, inverse roles are disallowed; in $\mathcal{ALCI}$, at-least and at-most restrictions are limited to $\exists r.C$ and $\forall r.C$; and $\mathcal{ALC}$ is the common part of $\mathcal{ALCQ}$ and $\mathcal{ALCI}$.

constructor	syntax	semantics
inverse role	r^-	$(r^{\mathcal{I}})^{-1} = \{(d, e) \mid (e, d) \in r^{\mathcal{I}}\}$
negation	$\neg C$	$\Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$
conjunction	$C \sqcap D$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
at-least restriction	$(\geqslant n\ r\ C)$	$\{d \in \Delta^{\mathcal{I}} \mid \lvert r^{\mathcal{I}}(d) \cap C^{\mathcal{I}} \rvert \geq n\}$
at-most restriction	$(\leqslant n\ r\ C)$	$\{d \in \Delta^{\mathcal{I}} \mid \lvert r^{\mathcal{I}}(d) \cap C^{\mathcal{I}} \rvert \leq n\}$

An $\mathcal{L}$ *TBox*, $\mathcal{T}$, for a DL $\mathcal{L}$ is a finite set of *concept inclusions* (CI) of the form $C \sqsubseteq D$, where C and D are $\mathcal{L}$-concepts. We write $\mathcal{I} \models C \sqsubseteq D$ if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ and $\mathcal{I} \models \mathcal{T}$ if this holds for all CIs in $\mathcal{T}$, in which case $\mathcal{I}$ is said to be a *model* of $\mathcal{T}$. $\mathcal{T}$ is *consistent* if it has a model. By a *signature*, Σ, we mean any set of concept and role names. The *signature* $\mathsf{sig}(\mathcal{T})$ *of a TBox* $\mathcal{T}$ is the set of concept and role names occurring in $\mathcal{T}$. By $\mathsf{sub}(\mathcal{T})$ we denote the closure under single negation of the set of subconcepts of concepts in $\mathcal{T}$.

Now we define three notions of TBox rewritability for DLs $\mathcal{L}$ and $\mathcal{L}'$, where $\mathcal{L}$ is typically more expressive than $\mathcal{L}'$.

Definition 1 An $\mathcal{L}'$ TBox $\mathcal{T}'$ is an *equivalent* $\mathcal{L}'$*-rewriting* of an $\mathcal{L}$ TBox $\mathcal{T}$ if $\mathcal{T}$ and $\mathcal{T}'$ have the same models. $\mathcal{T}$ is *equivalently* $\mathcal{L}'$*-rewritable* if it has an equivalent $\mathcal{L}'$-rewriting.

Equivalent $\mathcal{L}$-to-$\mathcal{L}'$ rewritability was studied by Lutz, Piro, and Wolter [2011], who gave the semantic characterizations in the first column of Table 1. For example, if $\mathcal{L}$ is $\mathcal{ALCQI}$, $\mathcal{ALCI}$ or $\mathcal{ALCQ}$ and $\mathcal{L}'$ is $\mathcal{ALC}$, then an $\mathcal{L}$ TBox is equivalently $\mathcal{L}'$-rewritable iff its class of models is preserved under *global bisimulations*, which are defined as follows. Given a signature Σ, a Σ*-bisimulation* between interpretations $\mathcal{I}_1$ and $\mathcal{I}_2$ is a relation $S \subseteq \Delta^{\mathcal{I}_1} \times \Delta^{\mathcal{I}_2}$ that satisfies conditions

[Atom], [Forth] and [Back] in the table below, for $r, A \in \Sigma$. In [Back] and elsewhere, 'dual' refers to swapping the rôles of $\mathcal{I}_1, d_1, d_1'$ and $\mathcal{I}_2, d_2, d_2'$. The relation S is a *global* Σ-bisimulation between $\mathcal{I}_1$ and $\mathcal{I}_2$ if $\Delta^{\mathcal{I}_1}$ is the domain of S and $\Delta^{\mathcal{I}_2}$ its range. $\mathcal{I}_1$ and $\mathcal{I}_2$ are *globally* Σ*-bisimilar* if there is a global Σ-bisimulation between them. If $\Sigma = \mathsf{N_C} \cup \mathsf{N_R}$, we omit Σ and say simply '(global) bisimulation.' A TBox $\mathcal{T}$ is *preserved under global bisimulations* if any interpretation that is globally bisimilar to a model of $\mathcal{T}$ is a model of $\mathcal{T}$.

[Atom]	for all $(d_1, d_2) \in S,\ d_1 \in A^{\mathcal{I}_1}$ iff $d_2 \in A^{\mathcal{I}_2}$
[Forth]	if $(d_1, d_2) \in S$ and $d_1' \in r^{\mathcal{I}_1}(d_1)$, $r \in \mathsf{N_R}$, then there is a $d_2' \in r^{\mathcal{I}_2}(d_2)$ with $(d_1', d_2') \in S$.
[Back]	dual of [Forth]
[QForth]	if $(d_1, d_2) \in S$ and $D_1 \subseteq r^{\mathcal{I}_1}(d_1)$ is finite, $r \in \mathsf{N_R}$, then there is a $D_2 \subseteq r^{\mathcal{I}_2}(d_2)$ such that S contains a bijection between D_1 and D_2.
[QBack]	dual of [QForth]

Example 1 The $\mathcal{ALCI}$ TBox $\{\exists r^-.B \sqsubseteq A\}$ can be equivalently rewritten to the $\mathcal{ALC}$ TBox $\{B \sqsubseteq \forall r.A\}$. However, the $\mathcal{ALCI}$ TBox $\mathcal{T} = \{\exists r^-.B \sqcap \exists s^-.B \sqsubseteq A\}$ is not equivalently $\mathcal{ALC}$-rewritable. Indeed, the interpretation $\mathcal{I}$ in the picture below is a model of $\mathcal{T}$ and globally bisimilar to the interpretation $\mathcal{J}$, which is not a model of $\mathcal{T}$.

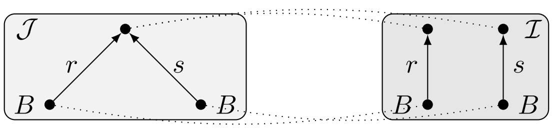

Equivalent $\mathcal{ALCQI}$-to-$\mathcal{ALCQ}$ rewritability is characterized by *counting bisimulations* defined but replacing [Forth] and [Back] in the definition of bisimulations with [QForth] and [QBack]. For equivalent $\mathcal{ALCQI}$-to-$\mathcal{ALCI}$ rewritability, we need *i-bisimulations*, that is, bisimulations for which [Forth] and [Back] hold for inverse roles as well.

We now introduce two subtler notions of TBox rewritability, which allow the use of fresh concept and role names in rewritings. For an interpretation $\mathcal{I}$ and a signature Σ, the Σ*-reduct of* $\mathcal{I}$ is the interpretation $\mathcal{I}_{|\Sigma}$ coinciding with $\mathcal{I}$ on Σ and having $X^{\mathcal{I}_{|\Sigma}} = \emptyset$ for $X \notin \Sigma$. We say that interpretations $\mathcal{I}$ and $\mathcal{J}$ *coincide on* Σ and write $\mathcal{I} =_\Sigma \mathcal{J}$ if the Σ-reducts of $\mathcal{I}$ and $\mathcal{J}$ coincide. A TBox $\mathcal{T}'$ is called a *model-conservative* (or m-conservative) *extension* of $\mathcal{T}$ if $\mathcal{T}' \models \mathcal{T}$ and, for every $\mathcal{I} \models \mathcal{T}$, there is $\mathcal{I}' \models \mathcal{T}'$ such that $\mathcal{I} =_{\mathsf{sig}(\mathcal{T})} \mathcal{I}'$.

Definition 2 An $\mathcal{L}'$ TBox $\mathcal{T}'$ is called an m*-conservative* $\mathcal{L}'$*-rewriting* of an $\mathcal{L}$ TBox $\mathcal{T}$ if $\mathcal{T}'$ is an m-conservative extension of $\mathcal{T}$. An $\mathcal{L}$ TBox $\mathcal{T}$ is m*-conservatively* $\mathcal{L}'$*-rewritable* if it has an m-conservative $\mathcal{L}'$-rewriting.

Any equivalent $\mathcal{L}'$-rewriting of a TBox $\mathcal{T}$ is also an m-conservative $\mathcal{L}'$-rewriting of $\mathcal{T}$, but not the other way round:

Example 2 The $\mathcal{ALCI}$ TBox $\{\exists r^-.B \sqcap \exists s^-.B \sqsubseteq A\}$ from Example 1 is m-conservatively $\mathcal{ALC}$-rewritable to

$$\{B \sqsubseteq \forall r.B_{\exists r^-.B}, B \sqsubseteq \forall s.B_{\exists s^-.B}, B_{\exists r^-.B} \sqcap B_{\exists s^-.B} \sqsubseteq A\},$$

where $B_{\exists r^-.B}$, $B_{\exists s^-.B}$ are *fresh* concept names.

A TBox $\mathcal{T}'$ is a *subsumption-conservative* (s-conservative) *extension* of an $\mathcal{L}$ TBox $\mathcal{T}$ if $\mathcal{T}' \models \mathcal{T}$ and $\mathcal{T}' \models C \sqsubseteq D$ implies $\mathcal{T} \models C \sqsubseteq D$, for any $\mathcal{L}$-CI $C \sqsubseteq D$ given in $\mathsf{sig}(\mathcal{T})$.

Definition 3 An $\mathcal{L}'$ TBox $\mathcal{T}'$ is an s*-conservative* $\mathcal{L}'$*-rewriting* of an $\mathcal{L}$ TBox $\mathcal{T}$ if $\mathcal{T}'$ is an s-conservative extension of $\mathcal{T}$. An $\mathcal{L}$ TBox $\mathcal{T}$ is s*-conservatively* $\mathcal{L}'$*-rewritable* if it has an s-conservative $\mathcal{L}'$-rewriting.

Note that it makes sense to speak about an s-conservative $\mathcal{L}'$-rewriting of a TBox $\mathcal{T}$ only if the language of $\mathcal{T}$ is understood. For example, the $\mathcal{ALC}$ TBox $\{\top \sqsubseteq \exists r.A \sqcap \exists r.\neg A\}$ is an s-conservative rewriting of $\mathcal{T} = \{\top \sqsubseteq \exists r.\top\}$ when $\mathcal{T}$ is regarded as an $\mathcal{ALC}$ TBox, but not as an $\mathcal{ALCQ}$ TBox.

Every m-conservatively $\mathcal{L}'$-rewritable TBox $\mathcal{T}$ is s-conservatively $\mathcal{L}'$-rewritable, but not the converse:

Example 3 The $\mathcal{ALCQ}$ TBox $\mathcal{T} = \{A \sqsubseteq\ \geq 2\, r.B\}$ is s-conservatively $\mathcal{ALC}$-rewritable to

$$\mathcal{T}' = \{A \sqsubseteq \exists r.B_1,\ A \sqsubseteq \exists r.B_2,\ B_1 \sqsubseteq \neg B_2,\ B_1 \sqcup B_2 \sqsubseteq B\},$$

where B_1 and B_2 are *fresh* concept names. To show this, note first that $\mathcal{T}' \models \mathcal{T}$. Second, recall that $\mathcal{ALCQ}$ TBoxes are complete for *ditree interpretations*, that is, interpretations $\mathcal{I}$ such that $r^{\mathcal{I}} \cap s^{\mathcal{I}} = \emptyset$ for $r \neq s$ and the directed graph with nodes $\Delta^{\mathcal{I}}$ and edges $(d, d') \in \bigcup_{r \in \mathsf{N_R}} r^{\mathcal{I}}$ is a directed tree. Thus, if $\mathcal{T} \not\models C \sqsubseteq D$, for an $\mathcal{ALCQ}$-CI $C \sqsubseteq D$ in $\mathsf{sig}(\mathcal{T})$, then there is a ditree model $\mathcal{I}$ of $\mathcal{T}$ with $\mathcal{I} \not\models C \sqsubseteq D$. Clearly, there exists a model $\mathcal{J}$ of $\mathcal{T}'$ with $\mathcal{J} =_{\mathsf{sig}(\mathcal{T})} \mathcal{I}$. But then $\mathcal{J} \not\models C \sqsubseteq D$, and so $\mathcal{T}' \not\models C \sqsubseteq D$, as required.

However, $\mathcal{T}'$ is not an m-conservative rewriting of $\mathcal{T}$ because (in contrast to ditree models of $\mathcal{T}$) the model $\mathcal{I}$ of $\mathcal{T}$ shown below is not the $\mathsf{sig}(\mathcal{T})$-reduct of any model of $\mathcal{T}'$.

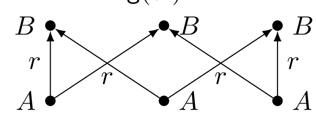

It is not difficult to generalize this argument to prove that there is *no* m-conservative $\mathcal{ALC}$-rewriting of $\mathcal{T}$.

In our examples so far, we used fresh concept names but no fresh role names. This is no accident: for the DLs considered in this paper, fresh role names in conservative rewritings are not required. Say that a DL $\mathcal{L}$ *reflects disjoint unions* if, for any $\mathcal{L}$ TBox $\mathcal{T}$, whenever the *disjoint* union $\biguplus_{i \in I} \mathcal{I}_i$ of interpretations $\mathcal{I}_i$ is a model of $\mathcal{T}$, then each $\mathcal{I}_i$, $i \in I$, is a model of $\mathcal{T}$. All of our DLs are known to reflect disjoint unions.

Theorem 1 *Let $\mathcal{L}$ be a DL reflecting disjoint unions, $\mathcal{T}$ an $\mathcal{L}$ TBox, and let $\mathcal{L}' \in \{\mathcal{ALCQ}, \mathcal{ALCI}, \mathcal{ALC}\}$. If $\mathcal{T}$ is* m*-conservatively* (*or* s*-conservatively*) *$\mathcal{L}'$-rewritable, then $\mathcal{T}$ has a* m*-conservative* (*or, respectively,* s*-conservative*) *$\mathcal{L}'$-rewriting not using role names outside* $\mathsf{sig}(\mathcal{T})$.

Proof. To illustrate the idea, consider an m-conservative $\mathcal{ALC}$-rewriting $\mathcal{T}'$ of $\mathcal{T}$. For any $C \in \mathsf{sub}(\mathcal{T}')$ of the form $\exists r.C'$ or $\forall r.C'$ with $r \notin \mathsf{sig}(\mathcal{T})$, take a fresh concept name B_C and denote by $D^\sharp$ the result of replacing all top-most occurrences of such C in $D \in \mathsf{sub}(\mathcal{T}')$ by B_C. The required m-conservative $\mathcal{ALC}$-rewriting $\mathcal{T}^\dagger$ is given by the inclusions $\bigsqcap_{C \in \boldsymbol{t}} C^\sharp \sqsubseteq \bot$, where $\boldsymbol{t}$ ranges over maximal subsets of $\mathsf{sub}(\mathcal{T})$ such that $\bigsqcap_{C \in \boldsymbol{t}} C$ is not satisfiable with respect to $\mathcal{T}'$. Indeed, for any $\mathcal{I} \models \mathcal{T}$, there is $\mathcal{J} \models \mathcal{T}^\dagger$ with $\mathcal{J} =_{\mathsf{sig}(\mathcal{T})} \mathcal{I}$. To show $\mathcal{T}^\dagger \models \mathcal{T}$, suppose $\mathcal{I} \models \mathcal{T}^\dagger$ and $\mathcal{I} \not\models \mathcal{T}$, with $r^{\mathcal{I}} = \emptyset$ for $r \notin \mathsf{sig}(\mathcal{T})$. By the definition of $\mathcal{T}^\dagger$, for every $d \in \Delta^{\mathcal{I}}$, there is a ditree model $\mathcal{I}_d$ of $\mathcal{T}'$ with root d (and no other shared elements) such that $d \in (C^\sharp)^\mathcal{I}$ iff $d \in C^{\mathcal{I}_d}$, for $C \in \mathsf{sub}(\mathcal{T}')$. We remove all $(d, d') \in r^{\mathcal{I}_d}$ with $r \in \mathsf{sig}(\mathcal{T})$ from $\mathcal{I}_d$, $d \in \Delta^\mathcal{I}$, and take the union $\mathcal{J}$ of the resulting interpretations with $\mathcal{I}$. Then $\mathcal{J} \models \mathcal{T}'$ but $\mathcal{J} \not\models \mathcal{T}$ (because $\mathcal{T}$ reflects disjoint unions and $\mathcal{J}_{|\mathsf{sig}(\mathcal{T})}$ is the disjoint union of the $\mathsf{sig}(\mathcal{T})$-reduct of $\mathcal{I}$ and the $\mathsf{sig}(\mathcal{T})$-reducts of $\mathcal{I}_d$ with d removed), which is a contradiction. ❏

Note that the size of $\mathcal{T}^\dagger$ is exponential in $|\mathcal{T}|$. It is an interesting open problem whether a polynomial rewriting exists. To see why reflection of disjoint unions is essential, consider the $\mathcal{ALCU}$ TBox $\mathcal{T} = \{\top \sqsubseteq \exists u.A\}$ with the *universal role* u, which is a logical symbol and not part of the signature of $\mathcal{T}$ [Krötzsch, Simančík, and Horrocks, 2012]. Then $\{\top \sqsubseteq \exists r.A\}$ is an m-conservative $\mathcal{ALC}$-rewriting of $\mathcal{T}$ but no such rewriting without fresh role names exists.

3 Rewriting Inverse Roles

In this section, we investigate conservative TBox rewritability from DLs with inverse roles to the corresponding DLs without them. First, we give a natural characterization of m- and s-conservative $\mathcal{ALC}$-rewritability of $\mathcal{ALCI}$-TBoxes in terms of generated subinterpretations. Motivated by the observation that preservation under generated subinterpretations does *not* characterize conservative $\mathcal{ALCQI}$-to-$\mathcal{ALCQ}$ rewritability, we then give an alternative characterization of conservative $\mathcal{ALCI}$-to-$\mathcal{ALC}$ rewritability in terms of p-morphisms. In contrast to generated subinterpretations, p-morphisms can be lifted to $\mathcal{ALCQI}$, and we show that m- and s-conservative $\mathcal{ALCQI}$-to-$\mathcal{ALCQ}$ rewritability is characterized in terms of counting p-morphisms.

An interpretation $\mathcal{I}$ is a *subinterpretation* of $\mathcal{J}$ if $\Delta^\mathcal{I} \subseteq \Delta^\mathcal{J}$, $A^\mathcal{I} = A^\mathcal{J} \cap \Delta^\mathcal{I}$, and $r^\mathcal{I} = r^\mathcal{J} \cap (\Delta^\mathcal{I} \times \Delta^\mathcal{I})$ for all A and r. $\mathcal{I}$ is a *generated subinterpretation* of $\mathcal{J}$ if, in addition, $d \in \Delta^\mathcal{I}$ and $(d, d') \in r^\mathcal{J}$ imply $d' \in \Delta^\mathcal{I}$. A TBox $\mathcal{T}$ is *preserved under generated subinterpretations* if every generated subinterpretation of a model of $\mathcal{T}$ is also a model of $\mathcal{T}$. As well known, all $\mathcal{ALC}$ TBoxes enjoy this property.

Suppose we want to construct an m-conservative $\mathcal{ALC}$-rewriting of an $\mathcal{ALCI}$ TBox $\mathcal{T}$. Without loss of generality, we can assume that $\mathcal{T}$ uses the concept constructors $\neg$, $\sqcap$ and $\exists$ only. For any role name r in $\mathcal{T}$, take a fresh role name $\bar{r}$. Then, for any $\exists r.C$ in $\mathsf{sub}(\mathcal{T})$, where r is a role (a role name or its inverse), take a fresh concept name $B_{\exists r.C}$. Denote by $D^\sharp$ the $\mathcal{ALC}$-concept obtained from any $D \in \mathsf{sub}(\mathcal{T})$ by replacing every top-most occurrence of a subconcept of the form $\exists r.C$ in D with $B_{\exists r.C}$. Now, let $\mathcal{T}^\dagger$ be an $\mathcal{ALC}$ TBox containing $C^\sharp \sqsubseteq D^\sharp$, for $C \sqsubseteq D \in \mathcal{T}$, and for $r \in \mathsf{N_R}$,

$$C^\sharp \sqsubseteq \forall \bar{r}.B_{\exists r.C}, \quad B_{\exists r.C} \equiv \exists r.C^\sharp, \quad \text{for } \exists r.C \in \mathsf{sub}(\mathcal{T}),$$

$$C^\sharp \sqsubseteq \forall r.B_{\exists r^-.C}, \; B_{\exists r^-.C} \equiv \exists \bar{r}.C^\sharp, \; \text{for } \exists r^-.C \in \mathsf{sub}(\mathcal{T}).$$

Clearly, $\mathcal{T}^\dagger$ can be constructed in polynomial time in $|\mathcal{T}|$.

Theorem 2 *The following conditions are equivalent for any $\mathcal{ALCI}$ TBox $\mathcal{T}$:*

(1) *$\mathcal{T}$ is m-conservatively $\mathcal{ALC}$-rewritable;*

(2) *$\mathcal{T}$ is s-conservatively $\mathcal{ALC}$-rewritable;*

(3) *$\mathcal{T}$ is preserved under generated subinterpretations;*

(4) *$\mathcal{T}^\dagger$ is an m-conservative $\mathcal{ALC}$-rewriting of $\mathcal{T}$.*

Proof. We only briefly discuss the proof of (3) $\Rightarrow$ (4) here. Assume (3). Clearly, for every model $\mathcal{I}$ of $\mathcal{T}$, there is a model $\mathcal{J}$ of $\mathcal{T}^\dagger$ with $\mathcal{J} =_{\mathsf{sig}(\mathcal{T})} \mathcal{I}$. It remains to show that $\mathcal{T}^\dagger \models \mathcal{T}$. Suppose $\mathcal{I} \models \mathcal{T}^\dagger$. The extension $\mathcal{I}_1$ of $\mathcal{I}$ in which the interpretation of every $\bar{r}$ is extended by the inverse of $r^\mathcal{I}$ is also a model of $\mathcal{T}^\dagger$. Let $\mathcal{I}_2$ be $\mathcal{I}_1$ with every $d \in \Delta^{\mathcal{I}_1}$ renamed to d'. Take the disjoint union $\mathcal{J}$ of $\mathcal{I}_1$ and $\mathcal{I}_2$, and replace each $(d, e) \in \bar{r}^\mathcal{J}$ such that $d, e \in \Delta^{\mathcal{I}_1}$ and $(e, d) \notin r^\mathcal{J}$ with $(e', d) \in r^\mathcal{J}$ and $(d, e') \in \bar{r}^\mathcal{J}$, and add $(e', d') \in r^\mathcal{J}$ for any $(d', e') \in \bar{r}^\mathcal{J}$ with $d', e' \in \Delta^{\mathcal{I}_2}$ and $(e', d') \notin r^\mathcal{J}$. Then $\mathcal{J} \models \mathcal{T}$, with the $\mathsf{sig}(\mathcal{T})$-reduct of $\mathcal{I}$ being a generated subinterpretation of the $\mathsf{sig}(\mathcal{T})$-reduct of $\mathcal{J}$. Thus $\mathcal{I} \models \mathcal{T}$. ❏

It is open whether a polynomial-size rewriting *without* additional role names exists. The proof above shows that to decide whether $\mathcal{T}$ is m-conservatively $\mathcal{ALC}$-rewritable, it is enough to check whether $\mathcal{T}^\dagger \models \mathcal{T}$, which can be done in ExpTime [Baader et al., 2003]. A matching ExpTime lower bound is obtained by reduction of $\mathcal{ALCI}$ TBox satisfiability.

Corollary 1 *Deciding m-conservative $\mathcal{ALCI}$-to-$\mathcal{ALC}$ rewritability is ExpTime-complete.*

The next example shows that preservation under generated subinterpretations does not guarantee conservative $\mathcal{ALCQI}$-to-$\mathcal{ALCQ}$ rewritability.

Example 4 Any subinterpretation of a model of the $\mathcal{ALCQI}$ TBox $\mathcal{T} = \{A \sqsubseteq (\leq 1\; r^-.\top)\}$ is also a model of $\mathcal{T}$, and so $\mathcal{T}$ is preserved under generated subinterpretations. We prove below that $\mathcal{T}$ is not m-conservatively $\mathcal{ALCQ}$ rewritable.

The reason why $\mathcal{T}$ cannot be conservatively rewritten into an $\mathcal{ALCQ}$ TBox is that, without inverse roles, one cannot restrict the number of r-predecessors. To capture this intuition, we introduce a functional version of (counting) bisimulations.

Definition 4 A (*counting*) *Σ-p-morphism* from $\mathcal{I}_1$ to $\mathcal{I}_2$ is any global (counting) Σ-bisimulation S between $\mathcal{I}_1$ and $\mathcal{I}_2$ such that S is a function. If $\Sigma = \mathsf{N_C} \cup \mathsf{N_R}$, we refer to S as a (counting) p-morphism. A TBox $\mathcal{T}$ is *preserved under inverse* (*counting*) *p-morphisms* if $\mathcal{I} \models \mathcal{T}$ whenever there is a (counting) p-morphism from $\mathcal{I}$ to a model of $\mathcal{T}$.

A fundamental property of p-morphisms is established by

Lemma 1 *Suppose $\mathcal{T}$ is an $\mathcal{ALC}$ (or $\mathcal{ALCQ}$) TBox, Σ contains all role names in $\mathsf{sig}(\mathcal{T})$, and there is a* (*counting*) *Σ-p-morphism f from an interpretation $\mathcal{I}$ to some model $\mathcal{I}'$ of $\mathcal{T}$. Then there is a model $\mathcal{J}$ of $\mathcal{T}$ such that $\mathcal{J} =_\Sigma \mathcal{I}$.*

Proof. We define $\mathcal{J}$ in the same way as $\mathcal{I}$ except that we set $A^\mathcal{J} = f^{-1}(A^{\mathcal{I}'})$ for $A \in \mathsf{sig}(\mathcal{T}) \setminus \Sigma$. Then f is a (counting) $\mathsf{sig}(\mathcal{T})$-bisimulation from $\mathcal{J}$ to $\mathcal{I}'$, and so $\mathcal{J} \models \mathcal{T}$. ❏

It follows that if an $\mathcal{ALCI}$ (or $\mathcal{ALCQI}$) TBox $\mathcal{T}$ is m-conservatively $\mathcal{ALC}$- (or $\mathcal{ALCQ}$-) rewritable, then $\mathcal{T}$ is preserved under inverse (counting) p-morphisms. Indeed, let $f : \mathcal{I}_1 \to \mathcal{I}_2$ be a p-morphism and $\mathcal{T}'$ an m-conservative $\mathcal{ALC}$-rewriting of $\mathcal{T}$. By Theorem 1, we may assume that the role names in $\mathsf{sig}(\mathcal{T}')$ belong to $\mathsf{sig}(\mathcal{T})$. By Lemma 1, there is a model $\mathcal{J}_1$ of $\mathcal{T}'$ with $\mathcal{J}_1 =_{\mathsf{sig}(\mathcal{T})} \mathcal{I}_1$, from which $\mathcal{I}_1 \models \mathcal{T}$.

Example 5 The map $f\colon \mathcal{I}_1 \to \mathcal{I}_2$ below is a counting p-morphism. Since $\mathcal{I}_2$ is a model of $\mathcal{T}$ from Example 4 but $\mathcal{I}_1$ is not, $\mathcal{T}$ is not m-conservatively $\mathcal{ALCQ}$-rewritable.

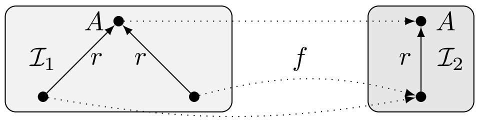

Note that if a TBox $\mathcal{T}$ reflects disjoint unions and is preserved under inverse p-morphisms, then it is preserved under generated subinterpretations. Indeed, let $\mathcal{I}$ be a generated subinterpretation of $\mathcal{J} \models \mathcal{T}$. Take the disjoint union $\mathcal{I}' = (\mathcal{I} \times \{0\}) \cup (\mathcal{J} \times \{1\})$ of $\mathcal{I}$ and $\mathcal{J}$. The map $f\colon \mathcal{I}' \to \mathcal{J}$ defined by setting $f(d,i) = d$ for $i = 0,1$ is a p-morphism. Then $\mathcal{I}' \models \mathcal{T}$, and so $\mathcal{I} \models \mathcal{T}$. Thus, we obtain:

Theorem 3 *An $\mathcal{ALCI}$ TBox is* m*-conservatively (and* s*-conservatively) $\mathcal{ALC}$-rewritable iff it is preserved under inverse p-morphisms.*

Counting p-morphisms characterize both m- and s-conservative $\mathcal{ALCQ}$-rewritabilities:

Theorem 4 *The following conditions are equivalent for any $\mathcal{ALCQI}$ TBox $\mathcal{T}$:*

(1) *$\mathcal{T}$ is* m*-conservatively $\mathcal{ALCQ}$-rewritable;*

(2) *$\mathcal{T}$ is* s*-conservatively $\mathcal{ALCQ}$-rewritable;*

(3) *$\mathcal{T}$ is preserved under inverse counting p-morphisms.*

Proof. We sketch the proof of (3) $\Rightarrow$ (1) where, unlike Theorem 2, we construct an *infinite* rewriting $\mathcal{T}'$ from which a finite one is obtained by compactness. $\mathcal{T}'$ is defined by brute force: given $\mathcal{T}$, it includes all $C^\sharp \sqsubseteq D^\sharp$ with $\mathcal{T} \models C \sqsubseteq D$, where C, D are $\mathcal{ALCQI}$ concepts over $\mathsf{sig}(\mathcal{T})$ and $C^\sharp, D^\sharp$ are the results of replacing uniformly any top-most qualified number restriction with inverse role by a fresh concept name. The crucial step now is to prove that $\mathcal{T}' \models \mathcal{T}$ if $\mathcal{T}$ is preserved under inverse counting morphisms. Suppose this is not so. Take an ω-saturated model $\mathcal{I}$ of $\mathcal{T}'$ that is not a model of $\mathcal{T}$ [Chang and Keisler, 1990, p. 100]. By unraveling $\mathcal{I}$ into a tree-shaped interpretation and using preservation under inverse counting p-morphisms, we construct a new $\mathcal{I}'$ with $\mathcal{I}' \models \mathcal{T}'$ and $\mathcal{I}' \not\models \mathcal{T}$, in which no node has more than one r-predecessor (r a role name) satisfying the same $\mathcal{ALCQ}$-concepts; cf. Example 5. Now we construct a model $\mathcal{J}$ of $\mathcal{T}$ containing $\mathcal{I}'$ as a generated subinterpretation, contrary to $\mathcal{T}$ being preserved under inverse counting p-morphisms. ❏

The decidability of rewritability and the size of rewritings in Theorem 4 remain open.

4 Rewriting Number Restrictions

Now we consider TBox rewritability from DLs with qualified number restrictions to the corresponding DLs without them. We first characterize s-conservative $\mathcal{ALCQ}$-to-$\mathcal{ALC}$ rewritability and $\mathcal{ALCQI}$-to-$\mathcal{ALCI}$ rewritability in terms of p-morphisms and, respectively, i-p-morphisms. We then generalize Example 3 and show that m-conservative $\mathcal{ALCQ}$-to-$\mathcal{ALC}$ rewritability coincides with equivalent $\mathcal{ALC}$-rewritability by characterizing it in terms of preservation under global bisimulations. Finally, we show that this is not the case for m-conservative $\mathcal{ALCQI}$-to-$\mathcal{ALCI}$ rewritability.

The next lemma shows that s-conservative $\mathcal{ALCQ}$-to-$\mathcal{ALC}$ rewritability can be regarded as a principled approximation of m-conservative rewritability (cf. Example 3).

Lemma 2 *An $\mathcal{ALC}$ TBox $\mathcal{T}'$ is an* s*-conservative rewriting of an $\mathcal{ALCQ}$ TBox $\mathcal{T}$ iff $\mathcal{T}'$ is an* m*-conservative rewriting of $\mathcal{T}$ over ditree interpretations of finite outdegree.*

Suppose we need an s-conservative $\mathcal{ALC}$-rewriting of an $\mathcal{ALCQ}$-TBox $\mathcal{T}$. As before, we assume that $\mathcal{T}$ is built using $\neg$, $\sqcap$ and $(\geqslant n\ r\ C)$ only. Take fresh concept names $B_D, B_1^D, \dots, B_n^D$, for $D = (\geqslant n\ r\ C) \in \mathsf{sub}(\mathcal{T})$, and let Σ be $\mathsf{sig}(\mathcal{T})$ together with the fresh concept names. For $C \in \mathsf{sub}(\mathcal{T})$, let $C^\sharp$ be the $\mathcal{ALC}$-concept obtained from C by replacing all top-most occurrences of $D = (\geqslant n\ r\ D')$ in C with B_D. Let $\mathcal{T}^\dagger$ be the *infinite* TBox containing $C^\sharp \sqsubseteq D^\sharp$, for $C \sqsubseteq D \in \mathcal{T}$, and for $D = (\geqslant n\ r\ C) \in \mathsf{sub}(\mathcal{T})$,

– $B_i^D \sqsubseteq \neg B_j^D$ for $i \neq j$,

– $B_D \sqsubseteq \exists r.(C^\sharp \sqcap B_1^D) \sqcap \dots \sqcap \exists r.(C^\sharp \sqcap B_n^D)$,

– $\bigsqcap_{1 \leq i \leq n} (\exists r.(C^\sharp \sqcap C_i^\sharp \sqcap \bigsqcap_{j \neq i} \neg C_j^\sharp)) \sqsubseteq B_D$, for any $\mathcal{ALC}$-concepts C_i with $\mathsf{sig}(C_i) \subseteq \Sigma$.

Theorem 5 *The following conditions are equivalent for any $\mathcal{ALCQ}$ TBox $\mathcal{T}$:*

(1) *$\mathcal{T}$ is* s*-conservatively $\mathcal{ALC}$-rewritable;*

(2) *$\mathcal{T}^\dagger$ is an* s*-conservative (infinite) $\mathcal{ALC}$-rewriting of $\mathcal{T}$;*

(3) *$\mathcal{T}$ is preserved under inverse p-morphisms.*

Proof. We sketch (3) $\Rightarrow$ (2). The interesting step is to prove that $\mathcal{T}^\dagger \models \mathcal{T}$. Suppose this is not the case. We find an ω-saturated model $\mathcal{I}$ of $\mathcal{T}^\dagger$ such that $\mathcal{I} \not\models \mathcal{T}$. Let $\mathcal{J}$ be the quotient $\mathcal{I}/_\sim$, where $d \sim d'$ if (d,d') is contained in the largest Σ-bisimulation on $\mathcal{I}$. The map that sends each $d \in \Delta^\mathcal{I}$ to its equivalence class $d/_\sim$ in $\mathcal{J}$ is a Σ-p-morphism, and by carefully analysing $\mathcal{T}^\dagger$ one can show that $\mathcal{J} \models \mathcal{T}$. By (2), $\mathcal{I} \models \mathcal{T}$, which is a contradiction. ❏

Although we do not know how to decide preservation under inverse counting p-morphisms from Theorem 4, preservation under inverse p-morphisms of $\mathcal{ALCQ}$ TBoxes can be decided in 2ExpTime (with numbers coded in unary). The algorithm uses a type elimination argument similar to the one employed for deciding equivalent $\mathcal{ALC}$-rewritability of $\mathcal{ALCI}$ TBoxes [Lutz, Piro, and Wolter, 2011]. So we have:

Theorem 6 *The problem of* s*-conservative $\mathcal{ALC}$-rewritability of $\mathcal{ALCQ}$ TBoxes is decidable in* 2ExpTime.

Thus, given an $\mathcal{ALCQ}$ TBox $\mathcal{T}$, one can first decide s-conservative $\mathcal{ALC}$-rewritability and then, in case of a positive answer, effectively construct a rewriting by going through the finite subsets of $\mathcal{T}^\dagger$ in a systematic way until a finite $\mathcal{T}' \subseteq \mathcal{T}^\dagger$ with $\mathcal{T}' \models \mathcal{T}$ is found, which must exist by compactness.

Our analysis of s-conservative $\mathcal{ALC}$-rewritability of $\mathcal{ALCQ}$ TBoxes can be lifted to s-conservative $\mathcal{ALCI}$-rewritability of $\mathcal{ALCQI}$ TBoxes by replacing (*i*) ditree interpretations with *tree interpretations* (in which $r^\mathcal{I} \cap s^\mathcal{I} = \emptyset$ for all *roles* $r \neq s$, and the undirected graph with nodes $\Delta^\mathcal{I}$ and edges $\{d,d'\}$ for $(d,d') \in \bigcup_{r \in \mathsf{N_R}} r^\mathcal{I}$ is a tree); (*ii*) p-morphisms with i-p-morphisms (functional i-bisimulations); and (*iii*) using fresh concept names B_D for qualified number restrictions D with inverse roles as well. These modifications give the required generalizations of Lemma 2 and Theorem 5. However, decidability of s-conservative $\mathcal{ALCI}$-rewritability of $\mathcal{ALCQI}$ TBoxes remains open.

As to m-conservative $\mathcal{ALCQ}$-to-$\mathcal{ALC}$ rewritability, Example 3 shows that the straightforward s-conservative $\mathcal{ALC}$-rewriting $\mathcal{T}'$ of $\mathcal{T} = \{A \sqsubseteq\ \geq 2\, r.B\}$ is not an m-conservative rewriting because there is a *non*-tree interpretation $\mathcal{I}$ for which no $\mathcal{J} \models \mathcal{T}'$ with $\mathcal{J} =_{\mathsf{sig}(\mathcal{T})} \mathcal{I}$ exists. A generalization of this argument shows that only $\mathcal{ALCQ}$ TBoxes that are preserved under global bisimulations are m-conservatively $\mathcal{ALC}$-rewritable. Thus, we obain:

Theorem 7 *An $\mathcal{ALCQ}$ TBox is* m*-conservatively $\mathcal{ALC}$-rewritable iff it is equivalently $\mathcal{ALC}$-rewritable.*

Using type elimination, one can prove that deciding preservation of $\mathcal{ALCQ}$ TBoxes under global bisimulations is in 2EXPTIME. Thus, m-conservative $\mathcal{ALC}$-rewritability of $\mathcal{ALCQ}$ TBoxes is decidable in 2EXPTIME.

Surprisingly, the situation is different for m-conservative $\mathcal{ALCQI}$-to-$\mathcal{ALCI}$ rewritability, where one would also expect that only equivalently $\mathcal{ALCI}$-rewritable TBoxes (those that are preserved under global i-bisimulations) are m-conservatively $\mathcal{ALCI}$-rewritable. However, the following example shows that this is not the case:

Example 6 The TBox $\mathcal{T} = \{\exists r.\top \sqsubseteq \exists r.(\geq\ 2r^-.\top)\}$ in $\mathcal{ALCQI}$ has the m-conservative $\mathcal{ALCI}$-rewriting $\mathcal{T}' = \{\exists r.\top \sqsubseteq \exists r.(\exists r^-.B \sqcap \exists r^-.\neg B)\}$. No equivalent $\mathcal{ALCI}$-rewriting of $\mathcal{T}$ exists because it is not preserved under global i-bisimulations. The proof that, for every $\mathcal{I} \models \mathcal{T}$, there is $\mathcal{J} \models \mathcal{T}'$ with $\mathcal{J} =_{\mathsf{sig}(\mathcal{T})} \mathcal{I}$ relies on the observation that non-tree shaped counterexamples such as the one in Example 3 do not exist because of the interaction between $\mathcal{T}$'s constraints for r-successors and r^--successors.

We do not have any conjecture as to a natural semantic characterization of m-conservative $\mathcal{ALCQI}$-to-$\mathcal{ALCI}$ rewritability. In fact, Theorem 7 and Example 6 together suggest that such a characterization does not exist.

5 $\mathcal{ALCQI}$-to-$\mathcal{ALC}$ Rewritability

At first sight, $\mathcal{ALCQI}$-to-$\mathcal{ALC}$ rewritability easily reduces to the two-step $\mathcal{ALCQI}$-to-$\mathcal{ALCQ}$-to-$\mathcal{ALC}$ rewritability. Note, however, that the first step introduces fresh concept names that are not regarded as auxiliary in the second step. In fact, to smoothly compose the two steps, a more general notion of rewritability with a distinguished set of symbols in the input TBox is needed. Call a TBox $\mathcal{T}$ m*-conservatively $\mathcal{L}$ rewritable relative to a signature* $\Sigma \subseteq \mathsf{sig}(\mathcal{T})$ if there exists an $\mathcal{L}$-TBox $\mathcal{T}'$ such that $\{\mathcal{I}_{|\Sigma} \mid \mathcal{I} \models \mathcal{T}\} = \{\mathcal{I}_{|\Sigma} \mid \mathcal{I} \models \mathcal{T}'\}$. Investigating this notion is beyond the scope of this paper. We only mention one unexpected result, which can be proved by reduction of the undecidable problem whether an $\mathcal{ALC}$ TBox is an m-conservative rewriting of the empty TBox [Konev et al., 2013] (cf. Corollary 1):

Theorem 8 *The problem of* m*-conservative $\mathcal{ALCI}$-to-$\mathcal{ALC}$ rewritability relative to a signature Σ is undecidable.*

As $\mathcal{ALC}$-rewritable $\mathcal{ALCQI}$ TBoxes are not preserved under global bisimulations (see Example 2), we cannot simply put together the corresponding characterizations from the previous two sections in order to characterize m-conservative $\mathcal{ALC}$-rewritability of $\mathcal{ALCQI}$ TBoxes. Nevertheless, by applying the s-conservative $\mathcal{ALCQ}$ rewriting above to the rewriting in the proof of Theorem 4, we obtain an s-conservative $\mathcal{ALC}$-rewriting of an input $\mathcal{ALCQI}$ TBox $\mathcal{T}$ iff $\mathcal{T}$ is preserved under inverse p-morphisms iff such a rewriting exists at all.

Theorem 9 *An $\mathcal{ALCQI}$ TBox is* s*-conservatively $\mathcal{ALC}$-rewritable iff it is preserved under inverse p-morphisms.*

For m-conservative rewritability, we have:

Theorem 10 *If an $\mathcal{ALCQI}$ TBox is preserved under global i-bisimulations and inverse p-morphisms, then it is* m*-conservatively $\mathcal{ALC}$-rewritable.*

Proof. From preservation under global i-bisimulations of $\mathcal{T}$ follows the existence of an equivalent $\mathcal{ALCI}$ TBox $\mathcal{T}'$. Then $\mathcal{T}$ is m-conservatively $\mathcal{ALC}$-rewritable iff $\mathcal{T}'$ is m-conservatively $\mathcal{ALC}$-rewritable iff $\mathcal{T}'$ is preserved under inverse p-morphisms (Theorem 3). ❏

We conjecture that the converse also holds. By Lemma 1, m-conservatively $\mathcal{ALC}$-rewritable $\mathcal{ALCQI}$ TBoxes are preserved under inverse p-morphims. Thus, the conjecture would follow from preservation under global i-bisimulations.

6 Discussion and Future Work

Up to now, our focus has been on rewritability between expressive DLs. However, rewritability to the lightweight DLs from the *DL-Lite* and $\mathcal{EL}$ families is of great interest as well. Table 1 gives our model-theoretic characterization of rewritability from $\mathcal{ALCI}$ to *DL-Lite$_{horn}$* (without role inclusions) [Calvanese et al., 2007; Artale et al., 2009]. The characterization of equivalent rewritability in terms of products and succ-simulations was given by Lutz, Piro, and Wolter [2011]. It is straightforward to prove that it also applies to m- and s-conservative rewritability of $\mathcal{ALCI}$ TBoxes by first showing that Theorem 1 holds for rewritings into *DL-Lite$_{horn}$* as well:

Theorem 11 *For $\mathcal{ALCI}$ TBoxes, equivalent DL-Lite$_{horn}$-rewritability,* m*-conservative DL-Lite$_{horn}$-rewritability, as well as* s*-conservative DL-Lite$_{horn}$-rewritability coincide and are* EXPTIME*-complete.*

Rewritability into *DL-Lite* dialects with role inclusions (where Theorem 1 does not hold) and into $\mathcal{EL}$ appear to be much more challenging and a detailed study remains for future work. More generally, at the moment we only fully understand conservative $\mathcal{ALCI}$-to-$\mathcal{ALC}$-rewritability; in all other cases, it remains to determine the optimal size of rewritings, the complexity of computing them, as well as tight bounds for the complexity of deciding rewritability. Based on the resulting algorithms, it would be of great interest to study conservative rewritability in practice and, in particular, determine the rewritability status of real-world ontologies.

Acknowledgments. This work has been supported by EPSRC UK grants EP/M012646/1 and EP/M012670/1 (iTract), and ERC grant 647289 (CODA).

References

[Artale et al., 2009] Artale, A.; Calvanese, D.; Kontchakov, R.; and Zakharyaschev, M. 2009. The DL-Lite family and relations. *Journal of Artificial Intelligence Research* 36:1–69.

[Baader et al., 2003] Baader, F.; Calvanese, D.; McGuinness, D. L.; Nardi, D.; and Patel-Schneider, P. F., eds. 2003. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press.

[Baader, Brandt, and Lutz, 2005] Baader, F.; Brandt, S.; and Lutz, C. 2005. Pushing the $\mathcal{EL}$ Envelope. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30-August 5, 2005*, 364–369.

[Baader, 1996] Baader, F. 1996. A formal definition for the expressive power of terminological knowledge representation languages. *J. Log. Comput.* 6(1):33–54.

[Bienvenu et al., 2014] Bienvenu, M.; ten Cate, B.; Lutz, C.; and Wolter, F. 2014. Ontology-based data access: A study through disjunctive datalog, CSP, and MMSNP. *ACM Trans. Database Syst.* 39(4):33.

[Calvanese et al., 2007] 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. Autom. Reasoning.* 39(3):385–429.

[Carral et al., 2014a] Carral, D.; Feier, C.; Grau, B. C.; Hitzler, P.; and Horrocks, I. 2014a. *EL*-ifying ontologies. In *Automated Reasoning - 7th International Joint Conference, IJCAR 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 19-22, 2014. Proceedings*, 464–479.

[Carral et al., 2014b] Carral, D.; Feier, C.; Romero, A. A.; Grau, B. C.; Hitzler, P.; and Horrocks, I. 2014b. Is your ontology as hard as you think? Rewriting ontologies into simpler DLs. In *Informal Proceedings of the 27th International Workshop on Description Logics, Vienna, Austria, July 17-20, 2014.*, 128–140.

[Chang and Keisler, 1990] Chang, C., and Keisler, H. 1990. *Model Theory*. North-Holland, Amsterdam.

[Console et al., 2014] Console, M.; Mora, J.; Rosati, R.; Santarelli, V.; and Savo, D. F. 2014. Effective computation of maximal sound approximations of description logic ontologies. In *The Semantic Web - ISWC 2014 - 13th International Semantic Web Conference, Riva del Garda, Italy, October 19-23, 2014. Proceedings, Part II*, 164–179.

[De Giacomo, 1995] De Giacomo, G. 1995. *Decidability of Class-Based Knowledge Representation Formalisms*. Ph.D. Dissertation, Università di Roma.

[Ding, Haarslev, and Wu, 2007] Ding, Y.; Haarslev, V.; and Wu, J. 2007. A new mapping from $\mathcal{ALCI}$ to $\mathcal{ALC}$. In *Proceedings of the 2007 International Workshop on Description Logics, Brixen-Bressanone, Italy, 8-10 June, 2007*.

[Ghilardi, Lutz, and Wolter, 2006] Ghilardi, S.; Lutz, C.; and Wolter, F. 2006. Did I damage my ontology? A case for conservative extensions in description logics. In *Proceedings, Tenth International Conference on Principles of Knowledge Representation and Reasoning, Lake District of the United Kingdom, June 2-5, 2006*, 187–197.

[Goranko and Otto, 2006] Goranko, V., and Otto, M. 2006. *Model Theory of Modal Logic.* Handbook of Modal Logic, 255–325. Elsevier.

[Kaminski and Cuenca Grau, 2013] Kaminski, M., and Cuenca Grau, B. 2013. Sufficient conditions for first-order and datalog rewritability in $\mathcal{ELU}$. In *Proceedings of the 26th International Workshop on Description Logics, Ulm, Germany, July 23 - 26, 2013*, 271–293.

[Kazakov, 2009] Kazakov, Y. 2009. Consequence-driven reasoning for Horn $\mathcal{SHIQ}$ ontologies. In *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11-17, 2009*, 2040–2045.

[Konev et al., 2009] Konev, B.; Lutz, C.; Walther, D.; and Wolter, F. 2009. Formal properties of modularisation. In *Modular Ontologies: Concepts, Theories and Techniques for Knowledge Modularization*. 25–66.

[Konev et al., 2013] Konev, B.; Lutz, C.; Walther, D.; and Wolter, F. 2013. Model-theoretic inseparability and modularity of description logic ontologies. *Artif. Intell.* 203:66–103.

[Konev et al., 2016] Konev, B.; Lutz, C.; Wolter, F.; and Zakharyaschev, M. 2016. Conservative Rewritability of Description Logic TBoxes. Technical report (`cgi.csc.liv.ac.uk/~frank/publ/publ.html`).

[Krötzsch, Simančík, and Horrocks, 2012] Krötzsch, M.; Simančík, F.; and Horrocks, I. 2012. A description logic primer. *CoRR* abs/1201.4089.

[Lutz and Wolter, 2010] Lutz, C., and Wolter, F. 2010. Deciding inseparability and conservative extensions in the description logic $\mathcal{EL}$. *J. Symb. Comput.* 194–228.

[Lutz, Piro, and Wolter, 2011] Lutz, C.; Piro, R.; and Wolter, F. 2011. Description logic TBoxes: Model-theoretic characterizations and rewritability. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, 983–988.

[Ren, Pan, and Zhao, 2010] Ren, Y.; Pan, J. Z.; and Zhao, Y. 2010. Soundness preserving approximation for TBox reasoning. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010*.