Journal of Artificial Intelligence Research 48 (2013) 347-414 Submitted 05/13; published 11/13

Beth Definability in Expressive Description Logics

Balder ten Cate BTENCATE@UCSC.EDU
UC Santa Cruz

Enrico Franconi FRANCONI@INF.UNIBZ.IT
Free University of Bozen-Bolzano

İnanç Seylan SEYLAN@INFORMATIK.UNI-BREMEN.DE
University of Bremen

Abstract

The *Beth definability property*, a well-known property from classical logic, is investigated in the context of description logics: if a general $\mathscr{L}$-TBox implicitly defines an $\mathscr{L}$-concept in terms of a given signature, where $\mathscr{L}$ is a description logic, then does there always exist over this signature an explicit definition in $\mathscr{L}$ for the concept? This property has been studied before and used to optimize reasoning in description logics. In this paper a complete classification of Beth definability is provided for extensions of the basic description logic $\mathcal{ALC}$ with transitive roles, inverse roles, role hierarchies, and/or functionality restrictions, both on arbitrary and on finite structures. Moreover, we present a tableau-based algorithm which computes explicit definitions of at most double exponential size. This algorithm is optimal because it is also shown that the smallest explicit definition of an implicitly defined concept may be double exponentially long in the size of the input TBox. Finally, if explicit definitions are allowed to be expressed in first-order logic, then we show how to compute them in single exponential time.

1. Introduction

We address the Beth definability property (Beth, 1953) in the context of description logics (DLs). The Beth definability property relates two notions of definability in a logic $\mathscr{L}$, implicit definability and explicit definability. Implicit definability is a semantic notion: it asks whether the interpretation of a given $\mathscr{L}$-formula φ is fully determined by the universe of discourse and the interpretation of some given predicates Σ in all models of a theory $\mathcal{T}$. Explicit definability on the other hand is more syntactic: it asks whether there is some $\mathscr{L}$-formula ψ over the set of predicates Σ that is equivalent to φ under $\mathcal{T}$. Clearly, explicit definability implies implicit definability. If the converse holds as well, then the logic $\mathscr{L}$ is said to have the Beth definability property. Logics having this property are considered to be well-balanced in terms of their syntax and semantics since it connects the model-theoretic notion of implicit definability to explicit definability.

The Beth definability property can be naturally formulated for DLs by slightly changing the terminology in the paragraph above: formulas become *concepts*, theories become *TBoxes*, and Σ consists of unary and binary predicates (respectively called *concept names* and *role names*).

Example 1.1. *Consider the following $\mathcal{ALC}$-TBox $\mathcal{T}$.*

$$\begin{array}{lcl} \mathsf{Parent} & \equiv & \exists\mathsf{hasChild}.\top \\ \mathsf{Parent} & \equiv & \mathsf{Father} \sqcup \mathsf{Mother} \\ \mathsf{Father} & \sqsubseteq & \mathsf{Man} \\ \mathsf{Mother} & \sqsubseteq & \mathsf{Woman} \\ \mathsf{Man} & \sqsubseteq & \neg\mathsf{Woman} \end{array}$$

The concept name Mother *is implicitly definable from* $\Sigma = \{\mathsf{hasChild}, \mathsf{Woman}\}$ *under* $\mathcal{T}$. *Precisely what we mean by this will be clear once we present Definition 4.1; intuitively, we mean that the instances of* Mother *in a model* $\mathcal{I}$ *of* $\mathcal{T}$ *can be exactly determined once we know the domain of* $\mathcal{I}$ *and the instances of* Σ *in* $\mathcal{I}$. *In fact, we can spell this implicit definition out as the* $\mathcal{ALC}$*-concept* $\mathsf{Woman} \sqcap \exists\mathsf{hasChild}.\top$. *This concept is an explicit definition of* Mother *from* Σ *under* $\mathcal{T}$ *because* $\mathcal{T} \models \mathsf{Mother} \equiv \mathsf{Woman} \sqcap \exists\mathsf{hasChild}.\top$ *(cf. Definition 4.5).*

Beth definability in DLs has found applications in optimizing reasoning. The first application is related to extracting an equivalent acyclic $\mathscr{L}$-terminology from a general TBox in $\mathscr{L}$ (Baader & Nutt, 2003; ten Cate, Conradie, Marx, & Venema, 2006). An acyclic terminology consists only of acyclic definitions for concept names and they are of particular interest because reasoning with them is 'easier' than with general TBoxes. For example, satisfiability of an $\mathcal{ALC}$-terminology is a PSpace-complete problem whereas the same problem for general $\mathcal{ALC}$-TBoxes is ExpTime-complete (Donini, 2003). The second application is related to an ontology-based data access setting, which assumes the existence of a database instance (also referred to as 'DBox' in this context) and a TBox that may speak about more predicates than the database instance (Seylan, Franconi, & de Bruijn, 2009). In this setting, the user may ask concept queries over the signature of the TBox; and the idea is to find an equivalent rewriting of the original query in terms of the predicates that appear in the DBox. If such a rewriting exists, then determining the certain answers of the query can be reduced to query answering in relational databases, which is known to be in AC^0 in data complexity in contrast to the general coNP-completeness of concept querying in $\mathcal{ALC}$ with DBoxes (Seylan et al., 2009).

Both of these applications involve computing explicit definitions on the basis of implicit definitions. Here, the problem is that this may not always be possible for some DLs, i.e., some DLs may lack the Beth definability property.

Example 1.2. *In this example, we model a scenario about cars, their owners, and the relationships between the owners and their cars. Consider the following* $\mathcal{ALCH}$*-TBox* $\mathcal{T}$ *consisting of the concept inclusion axioms*

$$\begin{array}{lcl} \mathsf{SportsCar} & \sqsubseteq & \mathsf{Car} \\ \mathsf{FuelEfficientCar} & \sqsubseteq & \mathsf{Car} \\ \mathsf{SportsCar} & \sqsubseteq & \neg\mathsf{FuelEfficientCar} \\ \neg\exists\mathsf{proudOwner}.\mathsf{Car} & \sqsubseteq & (\forall\mathsf{loves}.\mathsf{SportsCar} \sqcap \neg\exists\mathsf{owns}.\mathsf{SportsCar})) \sqcup \\ & & (\forall\mathsf{loves}.\mathsf{FuelEfficientCar} \sqcap \neg\exists\mathsf{owns}.\mathsf{FuelEfficientCar})) \end{array}$$

and the role inclusion axioms

$$\begin{array}{lcl} \mathsf{proudOwner} & \sqsubseteq & \mathsf{owns} \\ \mathsf{proudOwner} & \sqsubseteq & \mathsf{loves} \end{array}$$

The concept $\exists\mathsf{proudOwner}.\mathsf{Car}$ *is implicitly definable from* $\Sigma = \{\mathsf{owns}, \mathsf{loves}\}$ *under* $\mathcal{T}$*, in the sense that the instances of* $\exists\mathsf{proudOwner}.\mathsf{Car}$ *in a model* $\mathcal{I}$ *of* $\mathcal{T}$ *can be exactly determined once we know the domain of* $\mathcal{I}$ *and the instances of the roles in* Σ*. Indeed, an individual is a proud owner of a car if and only if the individual owns something that he/she loves. The fact that the left-to-right direction of this equivalence holds in every model of* $\mathcal{T}$ *follows immediately from the role inclusion axioms, and similarly, the fact that the (contrapositive of the) right-to-left direction holds in models of* $\mathcal{T}$ *follows immediately from the other TBox axioms. This implicit definition can be made explicit using the role conjunction operator as the concept* $\exists(\mathsf{owns} \sqcap \mathsf{loves}).\top$*. However, it can be shown that no* $\mathcal{ALCH}$*-concept is an explicit definition of* $\exists\mathsf{proudOwner}.\mathsf{Car}$ *from* Σ *under* $\mathcal{T}$*. We will not formally prove this here, but see the proof of Theorem 4.18 in Section 4.2 for a similar example. In particular, this shows that* $\mathcal{ALCH}$ *lacks the Beth definability property.*

A natural research agenda in this case is to identify DLs that have the Beth definability property. Since this property is useful for computing explicit definitions on the basis of implicit definitions, a vital question then is the complexity of this task, both in terms of the time needed to compute the explicit definitions, and in terms of the size of the explicit definitions obtained. This question was first studied by ten Cate et al. (2006) for a weaker Beth definability property, which considers only concept names in the signature. In this paper we are interested in the more general Beth definability property that takes into account role names in the signature. We believe that this is more natural for DLs because in a DL knowledge base, role names are considered to be a part of the signature. We present a worst-case optimal algorithm for constructing explicit definitions.

Since the work of Craig (1957), it has been customary to establish Beth definability via an interpolation theorem; and our work is no exception. In particular, we obtain our positive results on Beth definability through a worst-case optimal algorithm for constructing interpolants in the description logics that we consider.

Our contributions in this paper are as follows.

- We obtain a complete classification of the Beth definability property for extensions of $\mathcal{ALC}$ with transitive roles, inverse roles, role hierarchies, and/or functionality restrictions, both on arbitrary structures (BP) and on finite structures (BPF). These results are summarized in Table 1. Note that the finite model property (FMP) of all sub-logics of $\mathcal{SHOQ}$ is shown by Lutz, Areces, Horrocks, and Sattler (2005); FMP of all sub-logics of $\mathcal{SHIO}_+$ by Duc and Lamolle (2010); and the failure of FMP in $\mathcal{ALCFI}$ and all its extensions is well-known (cf. Calvanese & Giacomo, 2003).

- We present a constructive algorithm based on an interpolating tableau calculus to compute explicit definitions in $\mathcal{ALC}$ and all of its considered extensions having the Beth definability property. This algorithm runs in double exponential time and computes in the worst case an explicit definition of double exponential size if the concept is implicitly definable. In this respect, the algorithm is optimal because we also show that the smallest explicit definition of an implicitly defined concept may be double exponentially long in the size of the input TBox for each of these DLs.

- We consider the case where explicit definitions are allowed to be expressed in first-order logic. This is particularly relevant for the use case for computing certain answers

$\mathcal{S}$	$\mathcal{H}$	$\mathcal{I}$	$\mathcal{F}$	FMP	BP	BPF
				+	+	+
			•	+	+	+
		•		+	+	+
		•	•	-	+	-
	•			+	-	-
	•		•	+	-	-
	•	•		+	-	-
	•	•	•	-	-	-
•				+	+	+
•			•	+	+	+
•		•		+	+	+
•		•	•	-	+	-
•	•			+	-	-
•	•		•	+	-	-
•	•	•		+	-	-
•	•	•	•	-	-	-

Table 1: BP and BPF from $\mathcal{ALC}$ to $\mathcal{SHIF}$

of a query given a DBox and a TBox. We present an algorithm that computes a first-order explicit definition of an implicitly defined concept in single exponential time for all DLs with BP or BPF.

1.1 Related Work

The Beth definability property, in the general sense, has been first shown to hold for first-order logic (Beth, 1953). Beth definability comes in different flavors and the one we are interested in is related more to *projective Beth definability.* Here, *projective* refers to the ability to specify the set of predicates Σ. The projective version is known be stronger than Beth's original formulation (cf. Hoogland, 2001) and first shown to hold for first-order logic by Craig (1957). Since the seminal works of Beth and Craig, Beth definability has been studied for many other logics.

Lang and Marquis (2008), also motivated from AI, study the propositional variant. The modal and temporal variants have been extensively studied (cf. Gabbay & Maksimova, 2005). The k-variable fragment of first-order logic, for $k \geq 2$, is known to lack the Beth definability property, whereas the Guarded and Packed Fragments satisfy a non-projective version of the Beth property (cf. Hoogland, 2001). The guarded-negation fragment was recently shown to have the Beth definability property as well (Bárány, Benedikt, & ten Cate, 2013).

Beth definability has practical applications in relational databases for query rewriting using exact views (Nash, Segoufin, & Vianu, 2010; Afrati, 2011; Marx, 2007; Pasaila, 2011; Bárány et al., 2013). Here, the idea is to decide if the answers to a given query can be inferred from the content of a collection of views (that is, whether the theory consisting of the view definitions implicitly defines the query in terms of the view predicates), and, if this is indeed the case, to rewrite the query into a query over the schema consisting of the view predicates (that is, an explicit definition of the query in terms of the view predicates). View-based query rewriting naturally arises in various settings, including query optimization, querying under access restrictions, data integration, and privacy analysis.

Beth definability has also been studied in the DL literature. Similarly to the relational database case, it finds applications in computing explicit definitions on the basis of implicit definitions (Baader & Nutt, 2003; ten Cate et al., 2006; Seylan et al., 2009; Seylan, Franconi, & de Bruijn, 2010). Some of these papers also present results on the size of explicit definitions that can be obtained for implicitly defined concepts. Ten Cate et al. establish a single exponential lower bound and a triple exponential upper bound for $\mathcal{ALC}$. It is not hard to see that the lower bound proof by ten Cate et al. carries to the Beth definability property we consider. A matching single exponential upper bound on the size of explicit definitions was claimed to be established by Seylan et al. (2010) in Theorem 1; however, this theorem is wrong since a crucial step for its proof, namely Lemma 1, is erroneous. In this paper, we improve the single exponential lower bound of ten Cate et al. to double exponential and correct the single exponential upper bound of Seylan et al. to double exponential, thus obtaining tight complexity bounds. These bounds in DLs are in sharp contrast to first-order logic since there is no recursive bound on the minimal number of quantifier alternations in explicit definitions in first-order logic (Friedman, 1976). BP has been first shown to hold for $\mathcal{ALC}$ by Seylan et al. (2009) and it is stronger than the variant studied by ten Cate et al.. Specifically, we show that all DLs we consider that support role hierarchies actually lack BP, whereas they satisfy the variant of BP studied by ten Cate et al.. In this respect, Theorem 10 by Seylan et al. (2010) claiming that these DLs have BP is erroneous. The mistake in the proof is that Theorem 9, which presents a reduction from the concept satisfiability problem w.r.t. TBoxes in $\mathcal{SHI}$ to the same problem in $\mathcal{ALC}$, can not actually be used for computing $\mathcal{SHI}$-interpolants.

Since the work of Craig (1957), it has been customary to establish Beth definability via an interpolation lemma; and our work is no exception. An interpolation lemma is usually established by a model-theoretic or a proof-theoretic argument (Hoogland, 2001). The advantage of the latter over the former is that it yields a procedure to construct the interpolant. Several interpolation properties formulated for general TBoxes were studied in the $\mathcal{ALC}$- (ten Cate et al., 2006; Ghilardi, Lutz, & Wolter, 2006; Konev, Lutz, Walther, & Wolter, 2009a; Seylan et al., 2009; Konev, Lutz, Ponomaryov, & Wolter, 2010; Lutz & Wolter, 2011) and $\mathcal{EL}$-family of DLs (Konev, Walther, & Wolter, 2009b; Lutz, Piro, & Wolter, 2010; Nikitina & Rudolph, 2012; Lutz, Seylan, & Wolter, 2012a). A notable variant is the *uniform interpolation* property. A uniform interpolant of a given $\mathscr{L}$-TBox $\mathcal{T}$ and a set of predicates Σ is another $\mathscr{L}$-TBox $\mathcal{T}'$ such that $\mathcal{T}'$ uses only predicates from Σ and the logical consequences of $\mathcal{T}$ and $\mathcal{T}'$ formulated over Σ coincide. In this paper, we do not consider uniform interpolation because it is not the right interpolation property for establishing tight bounds on the size of explicit definitions. This is witnessed by the following observations. Deciding the existence of a uniform interpolant for a given $\mathcal{ALC}$-TBox and a set of predicates is known to be 2-ExpTime-complete (Lutz & Wolter, 2011), whereas the same problem formulated for the interpolation property we study here is in ExpTime. In the simpler DL $\mathcal{EL}$, uniform interpolants are also more 'expensive' than the non-uniform ones. In particular, deciding the existence of uniform interpolants in $\mathcal{EL}$ is ExpTime-complete (Lutz et al., 2012a); and Nikitina and Rudolph (2012) establish triple exponential tight bounds on the size of uniform interpolants. On the other hand, deciding the existence of interpolants, as we consider in this paper but for the description logic $\mathcal{EL}$, is in PTime because this problem can be reduced to concept subsumption w.r.t. a TBox in $\mathcal{EL}$ by Lemma 3 of Lutz, Seylan, and Wolter (2012b).

Most of the results in this paper were announced by ten Cate, Franconi, and Seylan (2011) in an extended abstract. The current paper extends this work by full proofs of the claimed results and the new material in Section 4.4.

1.2 Outline

We start by introducing in Section 2 the DLs for which we study BP and some reasoning problems that are relevant for us in this paper. We also fix in this section our first-order notation and the standard translation of DLs to first-order logic. We will be using first-order logic extensively in Section 3.3. The hammer with which we nail all the positive results, i.e., +, to the columns BP and BPF in Table 1, is a worst-case optimal algorithm for constructing interpolants. Section 3 is dedicated to this interpolation result. Finally, all our results about BP are presented in Section 4. Since the interpolation results are used to prove BP, Section 3 naturally comes before Section 4, but the reader who is less interested in the interpolation results may prefer to skip Section 3 initially.

2. Preliminaries

In this section, we introduce the description logics that we will study. They are frequently used logics in the expressive $\mathcal{ALC}$-family of description logics.

2.1 Description Logics

Let N_C and N_R be countably infinite and mutually disjoint sets of *concept names* and *role names*, respectively. For reasons that will become clear in a moment, we also assume a countably infinite subset of N_R, denoted by N_{R^+}, where $N_R \setminus N_{R^+}$ is also countably infinite. The role names in N_{R^+} are, intuitively, designated as being *transitive*, and are allowed to be used only in description logics with transitive roles. An element of N_C or N_R is also called a *predicate*, and a set $\Sigma \subseteq N_C \cup N_R$ of concept and role names is called a *signature*.

To ease the exposition, we first introduce the description logic $\mathcal{ALCFI}$, and we then define the other description logics that we study. The concept language of $\mathcal{ALCFI}$ is defined as follows:

$$
\begin{array}{lll}
\text{Concepts:} & C, D & ::= \top \mid A \mid \neg C \mid C \sqcap D \mid \exists R.C \mid \leq 1 R \\
\text{Roles:} & R & ::= P \mid P^-
\end{array}
$$

where $A \in N_C$ and $P \in N_R \setminus N_{R^+}$. The concept constructors $\bot$, $\sqcup$, $\forall R.C$, and $\geq 2R$ are defined as abbreviations in the usual way. Also, by a slight abuse of notation, we will sometimes write $(P^-)^-$, for $P \in N_R$, in which case it refers to the role name P itself. An $\mathcal{ALCFI}$*-TBox* $\mathcal{T}$ is a finite set of *concept inclusion axioms* (CIAs) $C \sqsubseteq D$, where C and D are $\mathcal{ALCFI}$-concepts.

The semantics of $\mathcal{ALCFI}$-concepts and roles is given in terms of *interpretations*. An interpretation is a pair $\mathcal{I} = \langle \Delta^\mathcal{I}, \cdot^\mathcal{I} \rangle$ where $\Delta^\mathcal{I}$ is a non-empty set called the *domain* of

$$
\begin{array}{rcl}
\top^{\mathcal{I}} & = & \Delta^{\mathcal{I}}, \\
(\neg C)^{\mathcal{I}} & = & \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}, \\
(C \sqcap D)^{\mathcal{I}} & = & C^{\mathcal{I}} \cap D^{\mathcal{I}}, \\
(\exists R.C)^{\mathcal{I}} & = & \{s \in \Delta^{\mathcal{I}} \mid \text{there exists } t \in \Delta^{\mathcal{I}} \text{ such that } \langle s,t\rangle \in R^{\mathcal{I}} \text{ and } t \in C^{\mathcal{I}}\}, \\
(\leq 1R)^{\mathcal{I}} & = & \{s \in \Delta^{\mathcal{I}} \mid \text{for all } t,u \in \Delta^{\mathcal{I}},\ \text{if } \langle s,t\rangle \in R^{\mathcal{I}} \text{ and } \langle s,u\rangle \in R^{\mathcal{I}} \text{ then } t = u\}, \\
(P^{-})^{\mathcal{I}} & = & \{\langle s,t\rangle \mid \langle t,s\rangle \in P^{\mathcal{I}}\}.
\end{array}
$$

Table 2: Semantics of complex $\mathcal{ALCFI}$-concepts and roles

$\mathcal{I}$, and $\cdot^{\mathcal{I}}$ is a function that maps each concept name $A \in N_C$ to a subset $A^{\mathcal{I}}$ of $\Delta^{\mathcal{I}}$ and each role name $P \in N_R$ to a binary relation $P^{\mathcal{I}}$ on $\Delta^{\mathcal{I}}$. In anticipation of the discussion of description logics with transitive roles below, we require also that for each $P \in N_{R^+}$, the relation $P^{\mathcal{I}}$ is transitive. The map $\cdot^{\mathcal{I}}$ is extended to complex concepts and roles by means of the inductive definitions provided in Table 2.

An interpretation $\mathcal{I}$ satisfies (or, is a model of) a CIA $C \sqsubseteq D$ if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, and $\mathcal{I}$ satisfies (or, is a model of) a TBox $\mathcal{T}$ if it satisfies every CIA in $\mathcal{T}$. We use the notation $\mathcal{I} \models C \sqsubseteq D$ and $\mathcal{I} \models \mathcal{T}$ to express that $\mathcal{I}$ satisfies $C \sqsubseteq D$, respectively, that $\mathcal{I}$ satisfies $\mathcal{T}$.

The description logic $\mathcal{ALCFI}$ that we defined above is a member of a larger family of description logics. The "basic" description logic $\mathcal{ALC}$ is defined as $\mathcal{ALCFI}$ without *inverse roles* (i.e., without roles of the form P^{-}) and without *functionality restrictions* (i.e., without concepts of the form $\leq 1R$). For $X \subseteq \{\mathcal{S}, \mathcal{H}, \mathcal{I}, \mathcal{F}\}$, the description logic $\mathcal{ALC}X$ extends $\mathcal{ALC}$ with

1. Functionality restrictions (as in $\mathcal{ALCFI}$) if $\mathcal{F} \in X$,

2. Inverse roles (as in $\mathcal{ALCFI}$) if $\mathcal{I} \in X$,

3. Transitive roles if $\mathcal{S} \in X$. By this, we mean that the role names in N_{R^+} are allowed to be used.

4. Role hierarchies if $\mathcal{H} \in X$. By this, we mean that a TBox may contain *role inclusion axioms* (*RIAs*) of the form $R \sqsubseteq S$, where R and S are roles, which are satisfied in an interpretation $\mathcal{I}$ if $R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$.

For DLs that include both transitive roles ($\mathcal{S}$) and functionality restrictions ($\mathcal{F}$), a further syntactic restriction is imposed: whenever $\leq 1R$ occurs in a concept, then R is required to be a *simple role* with respect to the TBox at hand (Horrocks, Sattler, & Tobies, 2000). A simple role is, intuitively, a role does not have a transitive subrole. The formal definition of simplicity is as follows: let us write $R \sqsubseteq_{\mathcal{T}} S$ if either $R = S$ or there are roles $R_1, \ldots, R_n$ such that $R_1 = R$, $R_n = S$, and for all $1 \leq i < n$, $\mathcal{T}$ contains either the RIA $R_i \sqsubseteq R_{i+1}$ or the RIA $R_i^{-} \sqsubseteq R_{i+1}^{-}$. We say R is simple with respect to $\mathcal{T}$ if there does not exist a role S such that $S \sqsubseteq_{\mathcal{T}} R$ and such that S is of the form P or P^{-} with $P \in N_{R^+}$. The motivation for this standard syntactic restriction is that, without it, basic decision problems such as satisfiability and concept subsumption with respect to a TBox (defined below) quickly become undecidable (Horrocks et al., 2000).

For $X \subseteq \{\mathcal{S}, \mathcal{H}, \mathcal{I}, \mathcal{F}\}$ with $\mathcal{S} \in X$, it is customary to omit the prefix $\mathcal{ALC}$ in the notation $\mathcal{ALC}X$. In particular, the description logic $\mathcal{ALCSHIF}$ (which is the most expressive description logic we consider in this paper) is referred to simply as $\mathcal{SHIF}$. $\mathcal{SHIF}$ is also the theoretical basis of the Web Ontology Language OWL-Lite (Horrocks, Patel-Schneider, & van Harmelen, 2003), which makes it an important DL from a practical viewpoint.

For an $\mathscr{L}$-concept C, the set $\mathsf{sub}(C)$ consists of C and all its subconcepts. For a concept C and a TBox $\mathcal{T}$, $\mathsf{rol}(C, \mathcal{T})$ denotes the set of roles occurring in C or $\mathcal{T}$; and $\mathsf{sig}(C, \mathcal{T})$ denotes the set of concept names and role names occurring in C or $\mathcal{T}$, i.e., the signature of C and $\mathcal{T}$. We use $\mathsf{sig}(C)$ as an abbreviation for $\mathsf{sig}(C, \emptyset)$. The *size of an $\mathscr{L}$-concept C* ($\mathscr{L}$-role R), written $|C|$ (resp. $|R|$), is the number of occurrences of symbols needed to write C (resp. R). The *size of an $\mathscr{L}$-TBox $\mathcal{T}$*, written $|\mathcal{T}|$, is defined analogously. Later on, in Section 3.3 we will also consider other, more succinct, ways of representing concepts.

There are alternative ways to represent functionality restrictions and transitive roles in the DL literature. For example, functionality and transitivity axioms of the form $\mathsf{funct}\, R$ or $\mathsf{Trans}(R)$ are sometimes treated as axioms in the TBox. Although such syntactic differences can be considered minor as far as the standard reasoning tasks (cf. Section 2.2) are concerned, interpolation results are sensitive to changes in the language. For example, opting for TBox axioms of the form $\mathsf{funct}\, R$ instead of freely allowing $\leq 1R$ as a construct in the concept language would change the expressive power of languages we consider. In Section 3.1, we show that $\mathcal{ALCF}$ has the interpolation property. The interested reader is invited to check if our proof can be adapted to the case where we allow functionality restrictions only as TBox axioms.

2.2 Decision Problems

A concept C is *satisfiable* with respect to TBox $\mathcal{T}$ if there exists a model $\mathcal{I}$ of $\mathcal{T}$ such that $C^{\mathcal{I}} \neq \emptyset$. A CIA $C \sqsubseteq D$ *follows from* a TBox $\mathcal{T}$ (denoted by $\mathcal{T} \models C \sqsubseteq D$), if every model of $\mathcal{T}$ is a model of $C \sqsubseteq D$. We write $\mathcal{T} \models C \equiv D$ if both $\mathcal{T} \models C \sqsubseteq D$ and $\mathcal{T} \models D \sqsubseteq C$ hold true.

The following decision problems will be relevant for us:

- *Concept satisfiability with respect to a TBox*:
  Given C and $\mathcal{T}$, to determine if C is satisfiable w.r.t. $\mathcal{T}$.

- *Concept subsumption with respect to a TBox*:
  Given $C \sqsubseteq D$ and $\mathcal{T}$, to determine if $\mathcal{T} \models C \sqsubseteq D$.

Both problems are parametrized by a description logic $\mathscr{L}$, in which the input concept(s) and TBox are specified. The two problems are reducible to each (or, more accurately, to each others complement) for all the logics we consider, due to the fact that their concept languages are closed under negation. In fact, both problems are ExpTime-complete for each of the description logics that we consider (Tobies, 2001).

The same decision problems can also be considered over the restricted class of *finite interpretations*, i.e., interpretations whose domain is a finite set. We will refer to these variants of the above decision problems as *finite concept satisfiability* and *finite concept subsumption*. Thus, finite concept satisfiability with respect to a TBox is the problem of deciding whether a given concept has a non-empty denotation in some finite model of a given TBox. It is known (Lutz, Sattler, & Tendera, 2005) that finite concept satisfiability and finite concept subsumption are also ExpTime-complete for all the description logics we consider here.

When the finite concept satisfiability problem coincides with the unrestricted satisfiability problem, then we say that the description logic in question has the *finite model property*.

Definition 2.1 (Finite model property)**.** *A DL $\mathscr{L}$ is said to have the* finite model property *(FMP) if for every $\mathscr{L}$-concept C and every $\mathscr{L}$-TBox $\mathcal{T}$, if C is satisfiable w.r.t. $\mathcal{T}$, then there is some finite interpretation $\mathcal{I}$ such that $\mathcal{I}$ is a model of $\mathcal{T}$ and $C^{\mathcal{I}} \neq \emptyset$.*

It is well-known that $\mathcal{ALCFI}$ and its extensions lack the finite model property (Calvanese & Giacomo, 2003).

2.3 First-Order Translation

It is well-known from the correspondence theory of modal/description logics that description logic concepts can, in general, be translated into first-order logic formulae with one free variable (Sattler, Calvanese, & Molitor, 2003). In this translation, each concept name A is viewed as a unary predicate symbol and each role name R is viewed as a binary predicate symbol of our first-order language. An interpretation, then, corresponds to a first-order structure.

We assume that the reader is familiar with basic notation and terminology for first-order logic. In particular, we will use the notation $\mathcal{I}, \alpha \models \varphi$ to express that the first-order formula φ is satisfied in the structure $\mathcal{I}$ under the first-order variable assignment α. Sometimes, it will be convenient to use a different notation to express the same thing: if $\varphi(x_1, \dots, x_n)$ is a first-order formula whose free variables are $x_1, \dots, x_n$, and if $a_1, \dots, a_n$ are elements of the domain of a structure $\mathcal{I}$, we will write $\mathcal{I} \models \varphi\ [a_1, \dots, a_n]$ to express that φ is satisfied in $\mathcal{I}$ under the variable assignment that sends each variable x_i to the corresponding element a_i. Note that this notation implicitly assumes an order on the free variables of φ, which will always be clear from the context.

Definition 2.2. *The mapping π_x from $\mathcal{SHIF}$-concepts to first-order formulae is defined as follows:*

$$
\begin{aligned}
\pi_x(\top) &= \top, \\
\pi_x(A) &= A(x), \\
\pi_x(\neg C) &= \neg \pi_x(C), \\
\pi_x(C \sqcap D) &= \pi_x(C) \wedge \pi_x(D), \\
\pi_x(\exists P.C) &= \exists y[P(x,y) \wedge \pi_y(C)], \\
\pi_x(\exists P^-.C) &= \exists y[P(y,x) \wedge \pi_y(C)], \\
\pi_x(\leq 1 P) &= \forall z_1 \forall z_2 [P(x,z_1) \wedge P(x,z_2) \rightarrow z_1 = z_2], \\
\pi_x(\leq 1 P^-) &= \forall z_1 \forall z_2 [P(z_1,x) \wedge P(z_2,x) \rightarrow z_1 = z_2],
\end{aligned}
$$

where π_y is obtained from the above definition by replacing all occurrences of x by y and vice versa. For a $\mathcal{SHIF}$-TBox $\mathcal{T}$, $\pi(\mathcal{T})$ is defined as $\bigwedge_{\varphi \in \mathcal{T}} \pi(\varphi)$, where

$$\begin{array}{rcl} \pi(C \sqsubseteq D) & = & \forall x[\pi_x(C) \to \pi_x(D)] \\ \pi(R \sqsubseteq S) & = & \forall x \forall y[\pi_{xy}(R) \to \pi_{xy}(S)] \end{array}$$

where, for $P \in N_R$, $\pi_{xy}(P) = P(x, y)$ and $\pi_{xy}(P^-) = P(y, x)$.

The translation above is model-preserving, i.e., for all $\mathcal{SHIF}$-concepts C, interpretations $\mathcal{I}$, and first-order assignments α for $\mathcal{I}$, we have $\alpha(x) \in C^{\mathcal{I}}$ iff $\mathcal{I}, \alpha \models \pi_x(C)$; and similarly for CIAs, RIAs, and TBoxes.

3. Constructive Interpolation with Tableaux

This section provides a constructive proof of an interpolation property in the DLs we are interested in. This property will be the essential part of the proof of BP in these DLs (cf. Definition 4.7). Resorting to interpolation to show the Beth definability property in a logic has been a standard technique since the seminal work of Craig (1957). We start by defining this interpolation property.

Definition 3.1 (Interpolation property). *A DL $\mathscr{L}$ is said to have the* interpolation property *if and only if for all $\mathscr{L}$-concepts C_1, C_2 and all $\mathscr{L}$-TBoxes $\mathcal{T}_1$, $\mathcal{T}_2$, if $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2$, then there is some $\mathscr{L}$-concept I such that*

- $\mathsf{sig}(I) \subseteq \mathsf{sig}(C_1, \mathcal{T}_1) \cap \mathsf{sig}(C_2, \mathcal{T}_2)$,
- $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq I$, *and*
- $\mathcal{T}_1 \cup \mathcal{T}_2 \models I \sqsubseteq C_2$.

Such a concept is called an interpolant *of C_1 and C_2 under $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$.*

The interpolation property we consider is defined specifically to prove BP. Normally, the Craig interpolation property for first-order logic is stated as follows: for all first-order formulae φ and ψ, if $\varphi \models \psi$, then there exists a first-order formula ϑ such that $\mathsf{sig}(\vartheta) \subseteq \mathsf{sig}(\varphi) \cap \mathsf{sig}(\psi)$, $\varphi \models \vartheta$, and $\vartheta \models \psi$. We can however relate the interpolation property we consider to first-order Craig interpolation using the standard translation of Definition 2.2. Given $\mathscr{L}$-concepts C_1, C_2 and $\mathscr{L}$-TBoxes $\mathcal{T}_1$, $\mathcal{T}_2$, we have by the standard translation the following equivalences:

$$\begin{array}{rcl} \mathcal{T}_1 \cup \mathcal{T}_2 & \models & C_1 \sqsubseteq C_2 \\ \pi(\mathcal{T}_1) \wedge \pi(\mathcal{T}_2) & \models & \pi_x(C_1) \to \pi_x(C_2) \\ \pi(\mathcal{T}_1) \wedge \pi_x(C_1) & \models & \pi(\mathcal{T}_2) \to \pi_x(C_2) \end{array}$$

Thus, by setting $\varphi = \pi(\mathcal{T}_1) \wedge \pi_x(C_1)$ and $\psi = \pi(\mathcal{T}_2) \to \pi_x(C_2)$, we know by Craig's Interpolation Theorem for first-order logic that we always have a first-order interpolant ϑ for φ and ψ, if $\varphi \models \psi$ (Craig, 1957). However, we do not know in general whether such an interpolant can be expressed as an $\mathscr{L}$-concept. Because of this reason we will work in the DL setting instead of full first-order. Our proofs are *constructive* in the sense that we present effective procedures for computing the interpolants. This also allows us to establish upper bounds on the size of interpolants.

This section is organized as follows. In Section 3.1, we show directly that the interpolation property holds for $\mathcal{ALC}$ and $\mathcal{ALCF}$ using a worst-case optimal tableau (plural: tableaux) algorithm in the style of Goré and Nguyen (2007). Then in Section 3.2, we show that the interpolation property also holds in the extensions of $\mathcal{ALC}$ and $\mathcal{ALCF}$ with transitive and inverse roles. Instead of establishing these results directly using tableaux, we make use of some satisfiability and signature preserving reductions to $\mathcal{ALC}$ and $\mathcal{ALCF}$. Our main result says that the interpolants in these logics can be computed in double exponential time. In Section 3.3, we study what happens when interpolants are allowed to be expressed in full first-order logic and show that first-order interpolants can be computed in single exponential time.

3.1 A Direct Algorithm for Computing Interpolants in $\mathcal{ALCF}$

In this section, we assume that $\mathcal{ALCF}$-concepts are defined recursively as in Section 2.1 using also $\bot$, $\sqcup$, $\forall R.C$, and $\geq 2R$ as primitives, i.e., we assume that, e.g., $\geq 2R$ is a constructor of our concept language and not an abbreviation for $\neg(\leq 1R)$ anymore. Moreover, we assume that all concepts are in *negation normal form* (NNF), i.e., the negation occurs only in front of concept names. It is well-known that every $\mathcal{ALCF}$-concept can easily be transformed to an equivalent one in NNF by pushing the negation inwards using the dualities between concept constructors (Tobies, 2001), e.g., $\forall R.C$ and $\neg\exists R.\neg C$. The NNF of the complement of a concept C is written as $\dot\neg C$. Another assumption we make is that $\mathcal{ALCF}$-TBoxes consist only of axioms of the form $\top \sqsubseteq C$. These assumptions make our tableau notation more compatible with the standard tableau notation for DLs. More precisely, we want to have a separate rule for each concept constructor in the language (Horrocks et al., 2000). The main result we present in this section, namely Theorem 3.10, can easily be shown to hold in the case where we do not make these assumptions.

Definition 3.2. *Let C be an $\mathcal{ALCF}$-concept and let $\mathcal{T}$ be an $\mathcal{ALCF}$-TBox. The* concept closure $\mathsf{cl}(C,\mathcal{T})$ *of C and $\mathcal{T}$ is the smallest set of concepts satisfying the following conditions:*

- $C \in \mathsf{cl}(C,\mathcal{T})$*;*
- *if* $\top \sqsubseteq D \in \mathcal{T}$*, then* $D \in \mathsf{cl}(C,\mathcal{T})$*;*
- *if* $D \in \mathsf{cl}(C,\mathcal{T})$ *and* $E \in \mathsf{sub}(D)$*, then* $E \in \mathsf{cl}(C,\mathcal{T})$*;*
- *if* $\exists R.D \in \mathsf{cl}(C,\mathcal{T})$ *then* $\forall R.D \in \mathsf{cl}(C,\mathcal{T})$*.*

For the rest of this section, fix two $\mathcal{ALCF}$-concepts C_0, D_0 and two $\mathcal{ALCF}$-TBoxes $\mathcal{T}_\mathbf{l}$, $\mathcal{T}_\mathbf{r}$. We will denote the union $\mathcal{T}_\mathbf{l} \cup \mathcal{T}_\mathbf{r}$ by $\mathcal{T}$. $\mathbf{l}$ stands for *left* and $\mathbf{r}$ for *right* and it is a naming scheme adopted from Fitting (1996). It will allow us to identify from which TBox ($\mathcal{T}_\mathbf{l}$ or $\mathcal{T}_\mathbf{r}$) or concept (C_0 or D_0) an inference is made. A *biased concept* is an expression of the form C^λ, where C is an $\mathcal{ALCF}$-concept and $\lambda \in \{\mathbf{l},\mathbf{r}\}$ is a *bias.* Two relevant biased concept closures cll and clr are defined as follows.

$$\mathsf{cll} = \{C^\mathbf{l} \mid C \in \mathsf{cl}(C_0,\mathcal{T}_\mathbf{l})\} \text{ and } \mathsf{clr} = \{C^\mathbf{r} \mid C \in \mathsf{cl}(\dot\neg D_0,\mathcal{T}_\mathbf{r})\}.$$

We use the Greek letters λ, κ to denote a bias.

Our tableau rules will be producing subsets of $\mathsf{cll} \cup \mathsf{clr}$ in a systematic way. To this aim, we make use of the metaphor of a burden and relief. Intuitively, a subset Φ of $\mathsf{cll} \cup \mathsf{clr}$ has a burden if the satisfiability of Φ depends on the satisfiability of one or more subsets of $\mathsf{cll} \cup \mathsf{clr}$ that we call the reliefs of Φ.

Definition 3.3. *Let $\Phi \subseteq \mathsf{cll} \cup \mathsf{clr}$. Then*

- $(C_1 \sqcap C_2)^\lambda$ *is an* $\sqcap$-burden *of Φ iff $(C_1 \sqcap C_2)^\lambda \in \Phi$ and $\{(C_1)^\lambda, (C_2)^\lambda\} \not\subseteq \Phi$;*
- $(C_1 \sqcup C_2)^\lambda$ *is an* $\sqcup$-burden *of Φ iff $(C_1 \sqcup C_2)^\lambda \in \Phi$ and $\{(C_1)^\lambda, (C_2)^\lambda\} \cap \Phi = \emptyset$;*
- $(\leq 1R)^\lambda$ *is an* ≤ 1-burden *of Φ iff $(\leq 1R)^\lambda \in \Phi$ and $\{(\forall R.C)^\kappa \mid (\exists R.C)^\kappa \in \Phi\} \not\subseteq \Phi$;*
- $(\exists R.C)^\lambda$ *is an* $\exists$-burden *of Φ iff $(\exists R.C)^\lambda$ is in Φ;*
- $(\geq 2R)^\lambda$ *is an* ≥ 2-burden *of Φ iff $(\geq 2R)^\lambda$ is in Φ.*

A burden *of Φ is any type of burden from above.*

Definition 3.4. *Let $\Phi \subseteq \mathsf{cll} \cup \mathsf{clr}$, C^λ be a burden of Φ, and $S = \{D^{\mathbf{l}} \mid \top \sqsubseteq D \in \mathcal{T}_{\mathbf{l}}\} \cup \{D^{\mathbf{r}} \mid \top \sqsubseteq D \in \mathcal{T}_{\mathbf{r}}\}$. Then $\Psi \subseteq \mathsf{cll} \cup \mathsf{clr}$ is called a C^λ-*relief *of Φ if*

- *$C = (C_1 \sqcap C_2)^\lambda$ and $\Psi = \{(C_1)^\lambda, (C_2)^\lambda\} \cup \Phi$;*
- *$C = (C_1 \sqcup C_2)^\lambda$ and either $\Psi = \Phi \cup \{(C_1)^\lambda\}$ or $\Psi = \Phi \cup \{(C_2)^\lambda\}$;*
- *$C = (\leq 1R)^\lambda$ and $\Psi = \Phi \cup \{(\forall R.C)^\kappa \mid (\exists R.C)^\kappa \in \Phi\}$;*
- *$C = (\exists R.C)^\lambda$ and $\Psi = \{C^\lambda\} \cup \{D^\kappa \mid (\forall R.D)^\kappa \in \Phi\} \cup S$;*
- *$C = (\geq 2R)^\lambda$ and $\Psi = \{D^\kappa \mid (\forall R.D)^\kappa \in \Phi\} \cup S$.*

A *biased $\langle C_0 \sqsubseteq D_0, \mathcal{T}\rangle$-tableau* ($\langle C_0 \sqsubseteq D_0, \mathcal{T}\rangle$-tableau for short) is a vertex-labeled directed graph $\langle \mathcal{V}, \mathcal{E}\rangle$ with the labeling $\mathsf{content} : \mathcal{V} \rightarrow 2^{\mathsf{cll} \cup \mathsf{clr}}$. Intuitively, for all edges $\langle g, g'\rangle$ constructed by our algorithm, $g'.\mathsf{content}$ will correspond to some C^λ-relief of $g.\mathsf{content}$. Note that a tableau is neither required to be a tree nor a directed acyclic graph (DAG) because cycles may occur in general. We say that a node g in a tableau contains a *clash* if and only if either one of the following holds.

- $\bot^\lambda \in g.\mathsf{content}$,
- $\{A^\lambda, (\neg A)^\kappa\} \subseteq g.\mathsf{content}$,
- $\{(\leq 1R)^\lambda, (\geq 2R)^\kappa\} \subseteq g.\mathsf{content}$.

The *tableau expansion rules* given in Figure 1 expand a tableau by making use of the semantics of concepts. A rule is said to be *applicable* to a node g if and only if its condition is satisfied in g, no rule was applied to g before, and g does not contain a clash. In order to guarantee a finite expansion, we use *proxies* in the following way. Whenever a rule creates a new node g' from g, before attaching the edge $\langle g, g'\rangle$ to $\mathcal{E}$, the tableau is searched for a

The $\mathsf{R}_\sqcap$ rule	
Condition:	$(C_1 \sqcap C_2)^\lambda$ is an $\sqcap$-burden of $g.\mathsf{content}$.
Action:	$\mathcal{E} \leftarrow \mathcal{E} \cup \{\langle g, g' \rangle\}$ and $g'.\mathsf{content} \leftarrow \Phi$, where Φ is the $(C_1 \sqcap C_2)^\lambda$-relief of $g.\mathsf{content}$.
The $\mathsf{R}_\sqcup$ rule	
Condition:	$(C_1 \sqcup C_2)^\lambda$ is an $\sqcup$-burden of $g.\mathsf{content}$.
Action:	$\mathcal{E} \leftarrow \mathcal{E} \cup \{\langle g, g_1 \rangle, \langle g, g_2 \rangle\}$, $g_1.\mathsf{content} \leftarrow \Phi_1$, and $g_2.\mathsf{content} \leftarrow \Phi_2$, where Φ_1, Φ_2 are $(C_1 \sqcup C_2)^\lambda$-reliefs of $g.\mathsf{content}$.
The $\mathsf{R}_{\leq 1}$ rule	
Condition:	$(\leq 1R)^\lambda$ is an ≤ 1-burden of $g.\mathsf{content}$.
Action:	$\mathcal{E} \leftarrow \mathcal{E} \cup \{\langle g, g' \rangle\}$ and $g'.\mathsf{content} \leftarrow \Phi$, where Φ is the $(\leq 1R)^\lambda$-relief of $g.\mathsf{content}$.
The $\mathsf{R}_\exists$ rule	
Condition:	$\Phi = \{(C_1)^{\lambda_1}, \ldots, (C_n)^{\lambda_n}\}$ such that $C^\lambda \in \Phi$ iff C^λ is an $\exists$- or ≥ 2-burden of $g.\mathsf{content}$.
Action:	$\mathcal{E} \leftarrow \mathcal{E} \cup \{\langle g, g_i \rangle \mid 1 \leq i \leq n\}$ and for $1 \leq i \leq n$, $g_i.\mathsf{content} \leftarrow \Phi_i$, where Φ_i is the $(C_i)^{\lambda_i}$-relief of $g.\mathsf{content}$.

Figure 1: Tableau expansion rules for $\mathcal{ALCF}$

node $g'' \in \mathcal{V}$ such that $g'.\mathsf{content} = g''.\mathsf{content}$. If such a g'' is found, then the edge $\langle g, g'' \rangle$ is added to $\mathcal{E}$ and g' is discarded.

We are interested in deciding $\mathcal{T} \models C_0 \sqsubseteq D_0$. The tableau algorithm consists of two phases. The *first phase* starts with the *initial* $\langle C_0 \sqsubseteq D_0, \mathcal{T}\rangle$*-tableau* $\mathbf{T} = \langle \{g_0\}, \emptyset \rangle$, where $g_0.\mathsf{content} = \{(C_0)^{\mathbf{l}}, (\dot\neg D_0)^{\mathbf{r}}\} \cup \{E^{\mathbf{l}} \mid \top \sqsubseteq E \in \mathcal{T}_{\mathbf{l}}\} \cup \{E^{\mathbf{r}} \mid \top \sqsubseteq E \in \mathcal{T}_{\mathbf{r}}\}$. $\mathbf{T}$ is then expanded by repeatedly applying the tableau expansion rules in such a way that if more than one rule is applicable to a node at the same time, then the first applicable rule in the list $[\mathsf{R}_\sqcap, \mathsf{R}_\sqcup, \mathsf{R}_{\leq 1}, \mathsf{R}_\exists]$ is chosen. The first phase continues as long as some rule is applicable to $\mathbf{T}$. A $\langle C_0 \sqsubseteq D_0, \mathcal{T}\rangle$-tableau is called *complete* if and only if it is the output of the first phase of the tableau algorithm.

Lemma 3.5. *The first phase of the tableau algorithm terminates in time* $2^{O(n)}$*, where* $n = |\mathsf{cll} \cup \mathsf{clr}|$*. Moreover for the complete* $\langle C_0 \sqsubseteq D_0, \mathcal{T}\rangle$*-tableau* $\mathbf{T} = \langle \mathcal{V}, \mathcal{E} \rangle$ *it produces, we have* $|\mathcal{V}| \leq 2^n$ *and* $|\mathcal{E}| \in 2^{O(n)}$.

Proof. By definition, the first phase continues as long as some rule is applicable to some node in the tableau. Then by the definition of applicability, we have that at most one rule is applied to a node in the tableau.

Let $n = |\mathsf{cll} \cup \mathsf{clr}|$. By the definition of a proxy, we have $|\mathcal{V}| \leq 2^n$ since there are 2^n distinct subsets of $\mathsf{cll} \cup \mathsf{clr}$. Combining this with the fact that there is at most one rule application per node, we obtain 2^n as a bound on the number of rule applications. Now, it is easy to see that each rule executes in time polynomial in n, i.e., the execution time of each rule is bounded by n^k, where k is a constant. Then we have that the whole running time of the first phase is $2^n \cdot n^k$. That is,

$$\begin{aligned} 2^n \cdot n^k &= 2^{n+\log n^k} \\ &= 2^{n+k\cdot\log n} \\ &\in 2^{O(n)}. \end{aligned}$$

It only remains to show the bound on $|\mathcal{E}|$. By the definition of tableau rules, the out-degree of a node cannot exceed n. Therefore, $|\mathcal{E}| \leq n \cdot 2^n$, i.e., $|\mathcal{E}| \in 2^{O(n)}$. $\square$

Let $\mathbf{T}$ be the complete $\langle C_0 \sqsubseteq D_0, \mathcal{T}\rangle$-tableau obtained from the first phase of the algorithm. The purpose of the *second phase* of the tableau algorithm, i.e., Algorithm 1, is to construct the following functions:

1. $\mathsf{status} : \mathcal{V} \to \{\mathtt{sat}, \mathtt{unsat}\}$ is a total function,
2. int is a partial function from $\mathcal{V}$ to $\mathcal{ALCF}$-concepts.

For a $g \in \mathcal{V}$, the values that are assigned to g by these functions are denoted by $g.\mathsf{status}$ and $\mathsf{int}(g)$. Intuitively, the status of a node g denotes if $\bigsqcap_{C^\lambda \in g.\mathsf{content}} C$ is satisfiable or not w.r.t. $\mathcal{T}$; and $\mathsf{int}(g)$, if defined, is an interpolant of $g.\mathsf{content}$ in the following sense.

Definition 3.6. *Let* $\Phi \subseteq \mathsf{cll} \cup \mathsf{clr}$. *A concept* I *is called an* interpolant of Φ *if and only if*

- $\mathcal{T} \models \bigsqcap_{C^{\mathbf{l}} \in \Phi} C \sqsubseteq I$ *and* $\mathcal{T} \models I \sqsubseteq \bigsqcup_{C^{\mathbf{r}} \in \Phi} \neg C$
- $\mathsf{sig}(I) \subseteq \mathsf{sig}(\bigsqcap_{C^{\mathbf{l}} \in \Phi} C) \cap \mathsf{sig}(\bigsqcup_{C^{\mathbf{r}} \in \Phi} \neg C)$,

By the definition of Algorithm 1, it will be that for all $g \in \mathcal{V}$, $\mathsf{int}(g)$ is defined if, and only if, $g.\mathsf{status} = \mathtt{unsat}$. In order to compute $\mathsf{int}(g)$ for a node $g \in \mathcal{V}$ with $g.\mathsf{status} = \mathtt{unsat}$, Algorithm 1 uses the interpolant calculation rules that are presented in Figures 2, 3, 4. The rules in Figure 2 compute $\mathsf{int}(g)$ based solely on $g.\mathsf{content}$; ones in Figure 3 take into account $g.\mathsf{content}$ and for some successor g' of g, the values $g'.\mathsf{content}$ and $\mathsf{int}(g')$; and finally, ones in Figure 4 take into account $g.\mathsf{content}$ and for every successor g' of g, the values $g'.\mathsf{content}$ and $\mathsf{int}(g')$. We invite the reader to verify that, indeed, whenever Algorithm 1 assigns $\mathtt{unsat}$ to $g.\mathsf{status}$, for a node g of the tableau, then there is an interpolant calculation rule that can be applied to compute $\mathsf{int}(g)$. Furthermore, each interpolant calculation rule is easily seen to be sound. For example, the interpolant calculation rule $\mathsf{C}_\sqcap$ in Figure 3 is sound because, if for a successor g' of g, $g'.\mathsf{content}$ is the $(C_1 \sqcap C_2)^\lambda$-relief of $g.\mathsf{content}$ and $\mathsf{int}(g')$ is an interpolant of $g'.\mathsf{content}$ (in the sense of Definition 3.6) then it is necessarily also an interpolant of $g.\mathsf{content}$.

Let $\mathbf{T} = \langle \mathcal{V}, \mathcal{E}\rangle$ be a complete $\langle C_0 \sqsubseteq D_0, \mathcal{T}\rangle$-tableau which is an output of the second phase. $\mathbf{T}$ is said to be *open* if and only if $g_0.\mathsf{status} = \mathtt{sat}$; and it is said to be *closed* if and only if $g_0.\mathsf{status} = \mathtt{unsat}$. If $\mathbf{T}$ is determined to be open after the second phase, then the tableau algorithm returns "$\mathcal{T} \not\models C_0 \sqsubseteq D_0$", otherwise it returns "$\mathcal{T} \models C_0 \sqsubseteq D_0$".

The next three results establish some important properties of our tableau algorithm and we use them to prove Theorem 3.10. The proofs of these results require the introduction of standard but substantial amount of notation from the DL and modal logic literature. In order to present Theorem 3.10 more clearly, we defer these proofs to Appendix C.

Algorithm 1 Second phase of the tableau algorithm

Propagate:
do

- $\mathsf{done} \leftarrow \mathsf{true}$.
- For every $g \in \mathcal{V}$ with $g.\mathsf{status} \neq \mathtt{unsat}$:
  - if g contains a clash, then
    1. $g.\mathsf{status} \leftarrow \mathtt{unsat}$,
    2. apply one of $\{\mathsf{C}^{\mathbf{l}}_{\bot}, \mathsf{C}^{\mathbf{r}}_{\bot}, \mathsf{C}^{\mathbf{ll}}_{\neg}, \mathsf{C}^{\mathbf{rr}}_{\neg}, \mathsf{C}^{\mathbf{lr}}_{\neg}, \mathsf{C}^{\mathbf{rl}}_{\neg}\}$, one whose condition is satisfied, to calculate $\mathsf{int}(g)$,
    3. $\mathsf{done} \leftarrow \mathsf{false}$.
  - if $\exists g' \in \mathcal{V}$ with $\langle g, g' \rangle \in \mathcal{E}$, $g'.\mathsf{status} = \mathtt{unsat}$, and $g'.\mathsf{content}$ is some $(C_1 \sqcap C_2)^{\lambda}$-, $(\leq 1R)^{\lambda}$-, $(\exists R.C)^{\lambda}$, or $(\geq 2R)^{\lambda}$-relief of $g.\mathsf{content}$, then
    1. $g.\mathsf{status} \leftarrow \mathtt{unsat}$,
    2. apply one of $\{\mathsf{C}_{\sqcap}, \mathsf{C}^{\mathbf{l}\not{R}}_{\leq 1}, \mathsf{C}^{\mathbf{r}\not{R}}_{\leq 1}, \mathsf{C}^{\mathbf{l}R}_{\leq 1}, \mathsf{C}^{\mathbf{r}R}_{\leq 1}, \mathsf{C}^{\mathbf{l}\not{R}}_{\exists}, \mathsf{C}^{\mathbf{r}\not{R}}_{\exists}, \mathsf{C}^{\mathbf{l}R}_{\exists}, \mathsf{C}^{\mathbf{r}R}_{\exists}\}$, one whose condition is satisfied, to calculate $\mathsf{int}(g)$,
    3. $\mathsf{done} \leftarrow \mathsf{false}$.
  - if $\exists g_1, g_2 \in \mathcal{V}$ with $g_1 \neq g_2$, $\langle g, g_1 \rangle, \langle g, g_2 \rangle \in \mathcal{E}$, $g_i.\mathsf{status} = \mathtt{unsat}$ for each $i \in \{1, 2\}$, $g_i.\mathsf{content}$ is a $(C_1 \sqcup C_2)^{\lambda}$-relief of $g.\mathsf{content}$ for each $i \in \{1, 2\}$, then
    1. $g.\mathsf{status} \leftarrow \mathtt{unsat}$,
    2. apply one of $\{\mathsf{C}^{\mathbf{l}}_{\sqcup}, \mathsf{C}^{\mathbf{r}}_{\sqcup}\}$, one whose condition is satisfied, to calculate $\mathsf{int}(g)$,
    3. $\mathsf{done} \leftarrow \mathsf{false}$.

while $\mathsf{done} = \mathsf{false}$.
Assign:
For every $g \in \mathcal{V}$ with $g.\mathsf{status} \neq \mathtt{unsat}$, $g.\mathsf{status} \leftarrow \mathtt{sat}$.

The $\mathsf{C}^{\mathbf{l}}_{\perp}$ rule	
Condition:	$\perp^{\mathbf{l}} \in g.\mathsf{content}$.
Action:	$\mathsf{int}(g) \leftarrow \perp$
The $\mathsf{C}^{\mathbf{r}}_{\perp}$ rule	
Condition:	$\perp^{\mathbf{r}} \in g.\mathsf{content}$.
Action:	$\mathsf{int}(g) \leftarrow \top$
The $\mathsf{C}^{\mathbf{ll}}_{\neg}$ rule	
Condition:	$\{C^{\mathbf{l}}, (\dot\neg C)^{\mathbf{l}}\} \subseteq g.\mathsf{content}$, for a C of the form A or $\leq 1R$.
Action:	$\mathsf{int}(g) \leftarrow \perp$
The $\mathsf{C}^{\mathbf{rr}}_{\neg}$ rule	
Condition:	$\{C^{\mathbf{r}}, (\dot\neg C)^{\mathbf{r}}\} \subseteq g.\mathsf{content}$, for a C of the form A or $\leq 1R$.
Action:	$\mathsf{int}(g) \leftarrow \top$
The $\mathsf{C}^{\mathbf{lr}}_{\neg}$ rule	
Condition:	$\{C^{\mathbf{l}}, (\dot\neg C)^{\mathbf{r}}\} \subseteq g.\mathsf{content}$, for a C of the form A or $\leq 1R$.
Action:	$\mathsf{int}(g) \leftarrow C$
The $\mathsf{C}^{\mathbf{rl}}_{\neg}$ rule	
Condition:	$\{C^{\mathbf{r}}, (\dot\neg C)^{\mathbf{l}}\} \subseteq g.\mathsf{content}$, for a C of the form A or $\leq 1R$.
Action:	$\mathsf{int}(g) \leftarrow \dot\neg C$

Figure 2: Interpolant calculation rules for $\mathcal{ALCF}$ ($\mathsf{content}$ dependent rules)

Lemma 3.7. *Let $\mathbf{T} = \langle \mathcal{V}, \mathcal{E} \rangle$ be the output of the second phase. For all $g \in \mathcal{V}$, if $g.\mathsf{status} =$ **unsat**, then*

1. *$g.\mathsf{content}$ is unsatisfiable w.r.t. $\mathcal{T}$;*
2. *$\mathsf{int}(g)$ is defined and it is an interpolant of $g.\mathsf{content}$; and*
3. *$|\mathsf{int}(g)| \in O(2^{2^n})$, where $n = |\mathsf{cll} \cup \mathsf{clr}|$.*

The next lemma establishes a double exponential upper bound for the runtime of Algorithm 1. This is a consequence of interpolant calculation and our double exponential upper bound on the size of these interpolants (cf. Lemma 3.7).

Lemma 3.8. *The second phase of the tableau algorithm, i.e., Algorithm 1, runs in time $O(2^{2^n})$, where $n = |\mathsf{cll} \cup \mathsf{clr}|$.*

The next proposition establishes the soundness and the completeness of our algorithm for concept subsumption w.r.t. TBoxes in $\mathcal{ALCF}$.

Proposition 3.9. $\mathbf{T}$ *is a closed $\langle C_0 \sqsubseteq D_0, \mathcal{T} \rangle$-tableau if and only if $\mathcal{T} \models C_0 \sqsubseteq D_0$.*

The tableau algorithm we presented in this section with the two phases is actually an algorithm to compute interpolants of at most double exponential size in $\mathcal{ALCF}$. This upper bound is optimal because the results we establish in Section 4 imply that smallest interpolants can be of double exponential size.

Theorem 3.10. *For all $\mathcal{ALCF}$-concepts C, D and all $\mathcal{ALCF}$-TBoxes $\mathcal{T}_1, \mathcal{T}_2$ if $\mathcal{T}_1 \cup \mathcal{T}_2 \models C \sqsubseteq D$ then there exists an interpolant of C and D under $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$ that can be computed in time double exponential in $|\mathcal{T}_1| + |\mathcal{T}_2| + |C| + |D|$.*

The $\mathsf{C}_{\sqcap}$ rule	
Condition:	$g'.\mathsf{content}$ is the $(C_1 \sqcap C_2)^{\lambda}$-relief of $g.\mathsf{content}$.
Action:	$\mathsf{int}(g) \leftarrow \mathsf{int}(g')$.
The $\mathsf{C}^{\mathbf{l}\not R}_{\leq 1}$ rule	
Condition:	$g'.\mathsf{content}$ is the $(\leq 1R)^{\mathbf{l}}$-relief of $g.\mathsf{content}$ and there is no biased concept of the form $(\exists R.C)^{\mathbf{r}} \in g.\mathsf{content}$.
Action:	$\mathsf{int}(g) \leftarrow \mathsf{int}(g')$.
The $\mathsf{C}^{\mathbf{r}\not R}_{\leq 1}$ rule	
Condition:	$g'.\mathsf{content}$ is the $(\leq 1R)^{\mathbf{r}}$-relief of $g.\mathsf{content}$ and there is no biased concept of the form $(\exists R.C)^{\mathbf{l}} \in g.\mathsf{content}$.
Action:	$\mathsf{int}(g) \leftarrow \mathsf{int}(g')$.
The $\mathsf{C}^{\mathbf{l}R}_{\leq 1}$ rule	
Condition:	$g'.\mathsf{content}$ is the $(\leq 1R)^{\mathbf{l}}$-relief of $g.\mathsf{content}$ and there is some biased concept of the form $(\exists R.C)^{\mathbf{r}} \in g.\mathsf{content}$.
Action:	$\mathsf{int}(g) \leftarrow \mathsf{int}(g') \sqcap \leq 1R$.
The $\mathsf{C}^{\mathbf{r}R}_{\leq 1}$ rule	
Condition:	$g'.\mathsf{content}$ is the $(\leq 1R)^{\mathbf{r}}$-relief of $g.\mathsf{content}$ and there is some biased concept of the form $(\exists R.C)^{\mathbf{l}} \in g.\mathsf{content}$.
Action:	$\mathsf{int}(g) \leftarrow \mathsf{int}(g') \sqcup \geq 2R$.
The $\mathsf{C}^{\mathbf{l}\not R}_{\exists}$ rule	
Condition:	$g'.\mathsf{content}$ is the $(\exists R.C)^{\mathbf{l}}$- or $(\geq 2R)^{\mathbf{l}}$-relief of $g.\mathsf{content}$, there is no biased concept of the form $(\forall R.D)^{\mathbf{r}} \in g.\mathsf{content}$.
Action:	$\mathsf{int}(g) \leftarrow \bot$.
The $\mathsf{C}^{\mathbf{r}\not R}_{\exists}$ rule	
Condition:	$g'.\mathsf{content}$ is the $(\exists R.C)^{\mathbf{r}}$- or $(\geq 2R)^{\mathbf{r}}$-relief of $g.\mathsf{content}$, there is no biased concept of the form $(\forall R.D)^{\mathbf{l}} \in g.\mathsf{content}$.
Action:	$\mathsf{int}(g) \leftarrow \top$.
The $\mathsf{C}^{\mathbf{l}R}_{\exists}$ rule	
Condition:	$g'.\mathsf{content}$ is the $(\exists R.C)^{\mathbf{l}}$- or $(\geq 2R)^{\mathbf{l}}$-relief of $g.\mathsf{content}$, there is some biased concept of the form $(\forall R.D)^{\mathbf{r}} \in g.\mathsf{content}$.
Action:	$\mathsf{int}(g) \leftarrow \exists R.\mathsf{int}(g')$.
The $\mathsf{C}^{\mathbf{r}R}_{\exists}$ rule	
Condition:	$g'.\mathsf{content}$ is the $(\exists R.C)^{\mathbf{r}}$- or $(\geq 2R)^{\mathbf{r}}$-relief of $g.\mathsf{content}$, there is some biased concept of the form $(\forall R.D)^{\mathbf{l}} \in g.\mathsf{content}$.
Action:	$\mathsf{int}(g) \leftarrow \forall R.\mathsf{int}(g')$.

Figure 3: Interpolant calculation rules for $\mathcal{ALCF}$ (single successor dependent rules)

Proof. Suppose C, D are $\mathcal{ALCF}$-concepts and $\mathcal{T}_1$, $\mathcal{T}_2$, and $\mathcal{T}$ are $\mathcal{ALCF}$-TBoxes such that $\mathcal{T}_1 \cup \mathcal{T}_2 = \mathcal{T}$ and $\mathcal{T} \models C \sqsubseteq D$. Then by Proposition 3.9, there is a closed $\langle C \sqsubseteq D, \mathcal{T}\rangle$-tableau $\mathbf{T} = \langle \mathcal{V}, \mathcal{E}\rangle$. This means $g_0.\mathsf{status} = \mathtt{unsat}$, and thus by Lemma 3.7, there is some $\mathcal{ALCF}$-concept I such that $\mathsf{int}(g_0) = I$ and I is an interpolant of $g_0.\mathsf{content}$. Let $X = \bigsqcap_{\top \sqsubseteq E \in \mathcal{T}_1} E$ and $Y = \bigsqcup_{\top \sqsubseteq E \in \mathcal{T}_2} \neg E$. Since I is an interpolant of $g_0.\mathsf{content}$, we have $\mathcal{T} \models C \sqcap X \sqsubseteq I$, $\mathcal{T} \models I \sqsubseteq D \sqcup Y$, and $\mathsf{sig}(I) \subseteq \mathsf{sig}(C \sqcap X) \cap \mathsf{sig}(D \sqcup Y)$. Then by the

The $\mathsf{C}^{\mathbf{l}}_{\sqcup}$ rule	
Condition:	$g_1.\mathsf{content}, g_2.\mathsf{content}$ are $(C_1 \sqcup C_2)^{\mathbf{l}}$-reliefs of $g.\mathsf{content}$.
Action:	$\mathsf{int}(g) \leftarrow \mathsf{int}(g_1) \sqcup \mathsf{int}(g_2)$.
The $\mathsf{C}^{\mathbf{r}}_{\sqcup}$ rule	
Condition:	$g_1.\mathsf{content}, g_2.\mathsf{content}$ are $(C_1 \sqcup C_2)^{\mathbf{r}}$-reliefs of $g.\mathsf{content}$.
Action:	$\mathsf{int}(g) \leftarrow \mathsf{int}(g_1) \sqcap \mathsf{int}(g_2)$.

Figure 4: Interpolant calculation rules for $\mathcal{ALCF}$ (multiple successor dependent rules)

fact that $\mathcal{T} \models X \equiv \top$ and $\mathcal{T} \models Y \equiv \bot$, we obtain $\mathcal{T} \models C \sqsubseteq I$ and $\mathcal{T} \models I \sqsubseteq D$; and by $\mathsf{sig}(I) \subseteq \mathsf{sig}(C \sqcap X) \cap \mathsf{sig}(D \sqcup Y)$, we obtain $\mathsf{sig}(I) \subseteq \mathsf{sig}(C, \mathcal{T}_1) \cap \mathsf{sig}(D, \mathcal{T}_2)$. Hence I is an interpolant of C and D under $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$. Finally by Lemma 3.8, I can be computed in time double exponential in $|\mathcal{T}_1| + |\mathcal{T}_2| + |C| + |D|$. □

We end this section with a discussion of the techniques we used. The tableau algorithm we defined is based on a tableau algorithm by Goré and Nguyen (2007). Here we extended this algorithm for $\mathcal{ALCF}$ and added more machinery to compute interpolants. In general interpolation follows as a corollary to a cut-free sequent or tableau calculus[1] for a logic (e.g., see Rautenberg, 1983; Fitting, 1996; Kracht, 2007); but such a corollary does not give upper bounds on the size and computation time of interpolants unless the calculus is combined with a decision procedure. In this section, our goal was to obtain tight upper bounds on the size and computation time of interpolants in $\mathcal{ALCF}$. More traditional tableau algorithms for DLs, e.g., the one by Horrocks et al. (2000), can also be used to establish similar results (Seylan et al., 2009). Here the crucial idea is that the tableau algorithm should provide an explicit representation of the tableau rule applications so that an interpolant can be calculated by induction on the rule applications. We chose a non-traditional DL tableau algorithm for our purposes because it is based on a non-labeled[2] tableau calculus and such calculi are actually more commonly used for proving interpolation results in modal logics (e.g., Rautenberg, 1983).

3.2 Extending Interpolation to Transitive and Inverse Roles

In this section, we extend Theorem 3.10 to more logics in order to obtain our main interpolation result Theorem 3.22. To this aim, we present various polynomial reductions from reasoning in one DL to another. The purpose of these reductions is to eliminate some constructors in the language. The technique we use for these reductions is well-known in the DL literature and it is called the *axiom schema instantiation technique* (Calvanese, Giacomo, & Rosati, 1998; Calvanese, Giacomo, Lenzerini, & Nardi, 2001). Similar techniques also appear in modal logic (Kracht, 2007). The idea behind this technique can be summarized as follows.

DLs are syntactic variants of modal logics. It is well-known that an axiom schema that is valid in a modal logic corresponds to a certain condition on the accessibility relation in the frames of that logic (Blackburn, de Rijke, & Venema, 2001). For example the axiom schema

1. A *tableau calculus* is defined as a set of tableau rules.
2. A non-labeled tableau calculus provides no explicit representation of individuals in the interpretations.

$4 : \Box\varphi \rightarrow \Box\Box\varphi$ defines the class of transitive frames. The axiom schema instantiation technique is based on instantiating an axiom schema a finite number of times for each concept in cl or a relevant concept closure, and adding these instances to the TBox to obtain an equi-satisfiable TBox. The resulting TBox will then be free of the constructor in the language for which we instantiated the axiom schema.

We note that the input in these reductions is normally a concept and a TBox; but for interpolation, we are given a pair of concepts C_1, C_2 and a pair of TBoxes $\mathcal{T}_1$, $\mathcal{T}_2$. Therefore, we require from these reductions that they do not mix the signature of $\mathsf{sig}(C_1, \mathcal{T}_1)$ and $\mathsf{sig}(C_2, \mathcal{T}_2)$ in an 'uncontrolled' way. What exactly we mean by this will be clear in Lemma 3.14 and Lemma 3.19. Naturally, this calls for extra notation.

Definition 3.11. *An injective function $\zeta : X \rightarrow N_R$, where X is a finite subset of $N_R \cup \{P^- \mid P \in N_R\}$, is called a* role renaming *if for all $P \in N_R$, we have $\{P, P^-\} \not\subseteq X$. A role renaming ζ is called* safe for a signature Σ *if $\mathsf{range}(\zeta) \cap \Sigma = \emptyset$.*

Given an $\mathscr{L}$-concept C and a role renaming ζ, $Z_\zeta(C)$ is the concept obtained from C by replacing every occurrence of every $R \in \mathsf{dom}(\zeta)$ by $\zeta(R)$.

Intuitively, we use role renamings, as the name suggests, to rename roles in concepts. We need to make sure that the renaming operation is well-defined and thus, we avoid mappings where a role and its inverse are in the domain of the mapping. Safeness of the mapping w.r.t. a signature is a property that we desire in the following reductions. We start with transitive roles and thus, instantiate the axiom schema $\Box\varphi \rightarrow \Box\Box\varphi$.

Definition 3.12. *Let C_0 be a $\mathcal{SIF}$-concept, $\mathcal{T}$ be a $\mathcal{SIF}$-TBox, and ζ be a safe role renaming for $\mathsf{sig}(C_0, \mathcal{T})$ with $\mathsf{dom}(\zeta) = \mathsf{sig}(C_0, \mathcal{T}) \cap N_{R^+}$ and $\mathsf{range}(\zeta) \cap N_{R^+} = \emptyset$. Then $\tau_\mathcal{S}(C_0, \mathcal{T}, \zeta)$ is defined as the $\mathcal{ALCFI}$-TBox $\tau^1_\mathcal{S}(C_0, \mathcal{T}, \zeta) \cup \tau^2_\mathcal{S}(C_0, \mathcal{T}, \zeta)$, where $\tau^1_\mathcal{S}(C_0, \mathcal{T}, \zeta) = \{\top \sqsubseteq Z_\zeta(C) \mid \top \sqsubseteq C \in \mathcal{T}\}$ and*

$$\tau^2_\mathcal{S}(C_0, \mathcal{T}, \zeta) \quad = \quad \{Z_\zeta(\forall R.C) \sqsubseteq Z_\zeta(\forall R.\forall R.C) \mid \forall R.C \in \mathsf{cl}(C_0, \mathcal{T}) \text{ and } \{R, R^-\} \cap N_{R^+} \neq \emptyset\}$$

Note that in the definition above, the signature of the resulting $\mathcal{ALCFI}$-TBox will not be equal to the signature of the original $\mathcal{SIF}$-TBox $\mathcal{T}$ if C_0 or $\mathcal{T}$ contains transitive roles. Introducing these new non-transitive role names is necessary because we are not allowed to use symbols from N_{R^+} in logics without transitive roles (cf. Section 2.1). Although the formulation of the following proposition is slightly different from the one of Lemma 6.23 by Tobies (2001), the proof idea is the same.

Proposition 3.13. *A $\mathcal{SIF}$-concept C_0 is satisfiable w.r.t. a $\mathcal{SIF}$-TBox $\mathcal{T}$ if and only if the $\mathcal{ALCFI}$-concept $Z_\zeta(C_0)$ is satisfiable w.r.t. the $\mathcal{ALCFI}$-TBox $\tau_\mathcal{S}(C_0, \mathcal{T}, \zeta)$, where ζ is a safe role renaming for $\mathsf{sig}(C_0, \mathcal{T})$ with $\mathsf{dom}(\zeta) = \mathsf{sig}(C_0, \mathcal{T}) \cap N_{R^+}$ and $\mathsf{range}(\zeta) \cap N_{R^+} = \emptyset$.*

The reduction (for concept satisfiability w.r.t. TBoxes) in Definition 3.12 satisfies the following property that will be essential for extending our interpolation results to logics with transitive roles. In this respect, it also resembles the *splitting reduction functions* of Kracht (2007).

Lemma 3.14. *Let $\mathcal{T}_1, \mathcal{T}_2$ be $\mathcal{SIF}$-TBoxes and let C_1, C_2 be $\mathcal{SIF}$-concepts. Then*

$$\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2 \text{ iff } \tau_\mathcal{S}(C_1, \mathcal{T}_1, \zeta) \cup \tau_\mathcal{S}(\dot\neg C_2, \mathcal{T}_2, \zeta) \models Z_\zeta(C_1) \sqsubseteq Z_\zeta(C_2)$$

where ζ is a safe role renaming for $\mathsf{sig}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1 \cup \mathcal{T}_2)$ *with* $\mathsf{dom}(\zeta) = \mathsf{sig}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1 \cup \mathcal{T}_2) \cap N_{R^+}$ *and* $\mathsf{range}(\zeta) \cap N_{R^+} = \emptyset$.

Proof. Let ζ be a safe role renaming for $\mathsf{sig}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1 \cup \mathcal{T}_2)$ as specified in the lemma. We will use the following claims for the proof.

Claim 3.15. $\tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1 \cup \mathcal{T}_2, \zeta) = \tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_2, \zeta)$.

Proof of claim. ($\Rightarrow$) Suppose $C \sqsubseteq D \in \tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1 \cup \mathcal{T}_2, \zeta)$. Then either $C \sqsubseteq D \in \tau^1_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1 \cup \mathcal{T}_2, \zeta)$ or $C \sqsubseteq D \in \tau^2_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1 \cup \mathcal{T}_2, \zeta)$. If the former holds, then we immediately obtain the desired result; thus, suppose the latter holds. Then $C \sqsubseteq D$ is of the form $Z_\zeta(\forall R.C) \sqsubseteq Z_\zeta(\forall R.\forall R.C)$, $\forall R.C \in \mathsf{cl}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1 \cup \mathcal{T}_2)$, and $\{R, R^-\} \cap N_{R^+} \neq \emptyset$. By Definition 3.2 and $\forall R.C \in \mathsf{cl}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1 \cup \mathcal{T}_2)$, we obtain $\forall R.C$ is in $\mathsf{cl}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1)$ or $\mathsf{cl}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_2)$. Then by Definition 3.12 and the fact that either $R \in N_{R^+}$ or $R^- \in N_{R^+}$, we have that $Z_\zeta(\forall R.C) \sqsubseteq Z_\zeta(\forall R.\forall R.C) \in \tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_2, \zeta)$, which is what we wanted to show.

($\Leftarrow$) It is rather easy to see that this direction of the claim holds. ⊣

Claim 3.16. $\tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_2, \zeta) = \tau_{\mathcal{S}}(C_1, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{S}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta)$.

Proof of claim. ($\Leftarrow$) Suppose $C \sqsubseteq D \in \tau_{\mathcal{S}}(C_1, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{S}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta)$. The desired result follows immediately if $C \sqsubseteq D = \top \sqsubseteq Z_\zeta(C')$, for some $\top \sqsubseteq C' \in \mathcal{T}_1 \cup \mathcal{T}_2$. Otherwise, we have by Definition 3.12 that $C \sqsubseteq D$ is of the form $Z_\zeta(\forall R.C) \sqsubseteq Z_\zeta(\forall R.\forall R.C)$, $\forall R.C \in \mathsf{cl}(C_1, \mathcal{T}_1) \cup \mathsf{cl}(\dot{\neg} C_2, \mathcal{T}_2)$ and either $R \in N_{R^+}$ or $R^- \in N_{R^+}$. Then by $\forall R.C \in \mathsf{cl}(C_1, \mathcal{T}_1) \cup \mathsf{cl}(\dot{\neg} C_2, \mathcal{T}_2)$ and $\mathsf{cl}(C_1, \mathcal{T}_1) \cup \mathsf{cl}(\dot{\neg} C_2, \mathcal{T}_2) \subseteq \mathsf{cl}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1) \cup \mathsf{cl}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_2)$, we obtain $\forall R.C \in \mathsf{cl}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1) \cup \mathsf{cl}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_2)$. Then by Definition 3.12 and the fact that either $R \in N_{R^+}$ or $R^- \in N_{R^+}$, we have $Z_\zeta(\forall R.C) \sqsubseteq Z_\zeta(\forall R.\forall R.C) \in \tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_2, \zeta)$, which is what we wanted to show.

($\Rightarrow$) Suppose $C \sqsubseteq D \in \tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_2, \zeta)$. The desired result follows immediately if $C \sqsubseteq D = \top \sqsubseteq Z_\zeta(C')$, for some $\top \sqsubseteq C' \in \mathcal{T}_1 \cup \mathcal{T}_2$. Otherwise, we have by Definition 3.12 that $C \sqsubseteq D$ is of the form $Z_\zeta(\forall R.C) \sqsubseteq Z_\zeta(\forall R.\forall R.C)$, $\forall R.C \in \mathsf{cl}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1) \cup \mathsf{cl}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_2)$, and either $R \in N_{R^+}$ or $R^- \in N_{R^+}$. Then by $\forall R.C \in \mathsf{cl}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1) \cup \mathsf{cl}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_2)$, the fact that $C_1 \sqcap \dot{\neg} C_2 \neq \forall R.C$, and Definition 3.2, we obtain $\forall R.C \in \mathsf{cl}(C_1, \mathcal{T}_1) \cup \mathsf{cl}(\dot{\neg} C_2, \mathcal{T}_2)$. Then by Definition 3.12 and the fact that either $R \in N_{R^+}$ or $R^- \in N_{R^+}$, we have $Z_\zeta(\forall R.C) \sqsubseteq Z_\zeta(\forall R.\forall R.C) \in \tau_{\mathcal{S}}(C_1, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{S}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta)$, which is what we wanted to show. ⊣

Now the lemma can be shown in the following way.

- $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2$, iff
- $C_1 \sqcap \dot{\neg} C_2$ is unsatisfiable w.r.t. $\mathcal{T}_1 \cup \mathcal{T}_2$, iff
- $Z_\zeta(C_1 \sqcap \dot{\neg} C_2)$ is unsatisfiable w.r.t. $\tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1 \cup \mathcal{T}_2, \zeta)$ (Proposition 3.13), iff
- $Z_\zeta(C_1 \sqcap \dot{\neg} C_2)$ is unsatisfiable w.r.t. $\tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{S}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_2, \zeta)$ (first claim), iff

- $Z_\zeta(C_1 \sqcap \dot{\neg} C_2)$ is unsatisfiable w.r.t. $\tau_{\mathcal{S}}(C_1, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{S}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta)$ (second claim), iff
- $\tau_{\mathcal{S}}(C_1, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{S}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta) \models Z_\zeta(C_1) \sqsubseteq Z_\zeta(C_2)$.

□

We need a similar reduction to eliminate inverse roles. De Giacomo (1996) presents a method to reduce converse-**PDL** satisfiability to **PDL** satisfiability using the axiom schema instantiation technique. Since DLs are notational variants of **PDL**s, this technique can easily be adapted to DLs as done by Calvanese et al. (1998, 2001). The idea is to instantiate the converse-**PDL** axiom schemas $\varphi \to [\alpha]\langle\alpha^-\rangle\varphi$ and $\varphi \to [\alpha^-]\langle\alpha\rangle\varphi$.

Definition 3.17. *Let C_0 be an $\mathcal{ALCFI}$-concept, let $\mathcal{T}$ be an $\mathcal{ALCFI}$-TBox, and ζ be a safe role renaming for $\mathsf{sig}(C_0, \mathcal{T})$ with $\mathsf{dom}(\zeta)$ consisting of all inverse roles appearing in C_0 or $\mathcal{T}$, and $\mathsf{range}(\zeta) \cap N_{R^+} = \emptyset$. Then $\tau_{\mathcal{I}}(C_0, \mathcal{T}, \zeta)$ is defined as the $\mathcal{ALCF}$-TBox $\tau^1_{\mathcal{I}}(C_0, \mathcal{T}, \zeta) \cup \tau^2_{\mathcal{I}}(C_0, \mathcal{T}, \zeta)$, where $\tau^1_{\mathcal{I}}(C_0, \mathcal{T}, \zeta) = \{\top \sqsubseteq Z_\zeta(C) \mid \top \sqsubseteq C \in \mathcal{T}\}$ and*

$$\tau^2_{\mathcal{I}}(C_0, \mathcal{T}, \zeta) \quad = \quad \{Z_\zeta(\dot{\neg} C) \sqsubseteq Z_\zeta(\forall R^-.\exists R.\dot{\neg} C) \mid \forall R.C \in \mathsf{cl}(C_0, \mathcal{T})\}$$

Note that in the definition above, the signature of the resulting $\mathcal{ALCF}$-TBox will not be equal to the signature of the original $\mathcal{ALCFI}$-TBox $\mathcal{T}$ if C_0 or $\mathcal{T}$ contains inverse roles. Proposition 3.18 establishes the correctness of this reduction for concept satisfiability w.r.t. TBoxes. A full proof of this proposition is given by Seylan (2012).

Proposition 3.18. *An $\mathcal{ALCFI}$-concept C_0 is satisfiable w.r.t. an $\mathcal{ALCFI}$-TBox $\mathcal{T}$ if and only if the $\mathcal{ALCF}$-concept $Z_\zeta(C_0)$ is satisfiable w.r.t. the $\mathcal{ALCF}$-TBox $\tau_{\mathcal{I}}(C_0, \mathcal{T}, \zeta)$, where ζ is a safe role renaming for $\mathsf{sig}(C_0, \mathcal{T})$ with $\mathsf{dom}(\zeta)$ consisting of all inverse roles appearing in C_0 or $\mathcal{T}$ and $\mathsf{range}(\zeta) \cap N_{R^+} = \emptyset$.*

The following property of this reduction will be useful in our interpolation results.

Lemma 3.19. *Let $\mathcal{T}_1, \mathcal{T}_2$ be $\mathcal{ALCFI}$-TBoxes and let C_1, C_2 be $\mathcal{ALCFI}$-concepts. Then*

$$\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2 \text{ iff } \tau_{\mathcal{I}}(C_1, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{I}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta) \models Z_\zeta(C_1) \sqsubseteq Z_\zeta(C_2)$$

where ζ is a safe role renaming for $\mathsf{sig}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1 \cup \mathcal{T}_2)$ with $\mathsf{dom}(\zeta)$ consisting of all inverse roles appearing in $C_1 \sqcap \dot{\neg} C_2$ or $\mathcal{T}_1 \cup \mathcal{T}_2$ and $\mathsf{range}(\zeta) \cap N_{R^+} = \emptyset$.

Proof. The following claims can be shown analogously to Claim 3.15 and Claim 3.16, respectively.

Claim 3.20. $\tau_{\mathcal{I}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1 \cup \mathcal{T}_2, \zeta) = \tau_{\mathcal{I}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{I}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_2, \zeta)$.

Claim 3.21. $\tau_{\mathcal{I}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{I}}(C_1 \sqcap \dot{\neg} C_2, \mathcal{T}_2, \zeta) = \tau_{\mathcal{I}}(C_1, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{I}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta)$.

Then the argument is the same as the last step in the proof of Lemma 3.14. □

Theorem 3.22. *Let $\mathscr{L}$ be $\mathcal{ALC}$ or any of its extensions with constructors from $\{\mathcal{S}, \mathcal{I}, \mathcal{F}\}$. For all $\mathscr{L}$-concepts C_1, C_2 and all $\mathscr{L}$-TBoxes $\mathcal{T}_1, \mathcal{T}_2$, if $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2$, then there exists an interpolant of C_1 and C_2 under $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$ that can be computed in time double exponential in $|\mathcal{T}_1| + |\mathcal{T}_2| + |C_1| + |C_2|$.*

Proof. Theorem 3.10 already covers the case for $\mathscr{L} = \mathcal{ALCF}$.

For $\mathscr{L} = \mathcal{ALC}$. The tableau algorithm for $\mathcal{ALCF}$ (with which we proved Theorem 3.10) can be used without modification to decide concept satisfiability w.r.t. a TBox in $\mathcal{ALC}$. In other words, given $\mathcal{ALC}$-concepts C_1, C_2 and an $\mathcal{ALC}$-TBox $\mathcal{T} = \mathcal{T}_1 \cup \mathcal{T}_2$, we can check if $\mathcal{T} \models C_1 \sqsubseteq C_2$ using the same algorithm. Observe that during the execution of the algorithm, $\mathsf{R}_{\leq 1}$ will never be applied and there will be no clashes involving a concept of the form $\leq 1R$. If the algorithm constructs a closed $\langle C_1 \sqsubseteq C_2, \mathcal{T}\rangle$-tableau, then the interpolant calculation algorithm will calculate an interpolant in $\mathcal{ALCF}$. Since $\mathsf{R}_{\leq 1}$ was never applied in the first phase and there is no clash involving a concept of the form $\leq 1R$ in the resulting tableau, the interpolant calculation rules producing concepts of the form $\leq 1R$ or $\geq 2R$, namely $\mathsf{C}^{\mathbf{l}R}_{\leq 1}$, $\mathsf{C}^{\mathbf{r}R}_{\leq 1}$, and the ones in Figure 2, will never be applied in the second phase. Hence the resulting interpolant is actually an $\mathcal{ALC}$-concept. That there is always an interpolant if $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2$ and the double exponential upper bound on its computation time can be shown as in Theorem 3.10.

For $\mathscr{L} \in \{\mathcal{ALCI}, \mathcal{ALCFI}\}$. Let C_1, C_2 be $\mathscr{L}$-concepts and $\mathcal{T}_1, \mathcal{T}_2$ be $\mathscr{L}$-TBoxes such that $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2$. Let ζ be a role renaming as specified in Lemma 3.19: such a role renaming always exists. Then by Lemma 3.19, $\tau_{\mathcal{I}}(C_1, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{I}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta) \models Z_\zeta(C_1) \sqsubseteq Z_\zeta(C_2)$, where C_1, C_2 are $\mathcal{ALC}$-concepts ($\mathcal{ALCF}$-concepts) and $\tau_{\mathcal{I}}(C_1, \mathcal{T}_1, \zeta)$, $\tau_{\mathcal{I}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta)$ are $\mathcal{ALC}$-TBoxes (respectively $\mathcal{ALCF}$-TBoxes). We compute an interpolant I of $Z_\zeta(C_1)$ and $Z_\zeta(C_2)$ under $\langle \tau_{\mathcal{I}}(C_1, \mathcal{T}_1, \zeta), \tau_{\mathcal{I}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta)\rangle$ in time that is at most double exponential in the size of the input. We have

1. $\mathsf{sig}(I) \subseteq \mathsf{sig}(Z_\zeta(C_1), \tau_{\mathcal{I}}(C_1, \mathcal{T}_1, \zeta)) \cap \mathsf{sig}(Z_\zeta(C_2), \tau_{\mathcal{I}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta))$,
2. $\tau_{\mathcal{I}}(C_1, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{I}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta) \models Z_\zeta(C_1) \sqsubseteq I$,
3. $\tau_{\mathcal{I}}(C_1, \mathcal{T}_1, \zeta) \cup \tau_{\mathcal{I}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta) \models I \sqsubseteq Z_\zeta(C_2)$.

Let ζ_1 be the restriction of ζ to $\mathsf{rol}(C_1, \mathcal{T}_1)$ and ζ_2 be the restriction of ζ to $\mathsf{rol}(\dot{\neg} C_2, \mathcal{T}_2)$; and set $\Sigma_1 = \mathsf{range}(\zeta_1)$ and $\Sigma_2 = \mathsf{range}(\zeta_2)$. Intuitively, Σ_1 and Σ_2 are exactly the sets of new role names we introduced in $\tau_{\mathcal{I}}(C_1, \mathcal{T}_1, \zeta)$ and $\tau_{\mathcal{I}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta)$, respectively. It is easy to see that $\mathsf{sig}(Z_\zeta(C_1), \tau_{\mathcal{I}}(C_1, \mathcal{T}_1, \zeta)) \subseteq \mathsf{sig}(C_1, \mathcal{T}_1) \cup \Sigma_1$, and $\mathsf{sig}(Z_\zeta(C_2), \tau_{\mathcal{I}}(\dot{\neg} C_2, \mathcal{T}_2, \zeta)) \subseteq \mathsf{sig}(C_2, \mathcal{T}_2) \cup \Sigma_2$. Then by item 1 above, we have

$$\mathsf{sig}(I) \subseteq (\mathsf{sig}(C_1, \mathcal{T}_1) \cup \Sigma_1) \cap (\mathsf{sig}(C_2, \mathcal{T}_2) \cup \Sigma_2)$$

By a simple distributivity argument, we obtain

$$\begin{aligned} \mathsf{sig}(I) \quad \subseteq \quad & (\mathsf{sig}(C_1, \mathcal{T}_1) \cap \mathsf{sig}(C_2, \mathcal{T}_2)) \cup (\Sigma_1 \cap \Sigma_2) \cup \\ & (\mathsf{sig}(C_1, \mathcal{T}_1) \cap \Sigma_2) \cup (\mathsf{sig}(C_2, \mathcal{T}_2) \cap \Sigma_1) \end{aligned}$$

Since $\mathsf{sig}(C_1, \mathcal{T}_1) \cap \Sigma_2 = \emptyset$ and $\mathsf{sig}(C_2, \mathcal{T}_2) \cap \Sigma_1 = \emptyset$,

$$\mathsf{sig}(I) \subseteq (\mathsf{sig}(C_1, \mathcal{T}_1) \cap \mathsf{sig}(C_2, \mathcal{T}_2)) \cup (\Sigma_1 \cap \Sigma_2)$$

Now let D be the $\mathscr{L}$-concept that is obtained from I by replacing all occurrences of each role name $P \in \Sigma_1 \cap \Sigma_2$ by the only role R^- such that $\zeta(R^-) = P$. Since ζ is injective, this is well defined. Moreover, we have $Z_\zeta(D) = I$.

We claim that for every $P \in \Sigma_1 \cap \Sigma_2$, the role name R with $\zeta(R^-) = P$ is in $\mathsf{sig}(C_1, \mathcal{T}_1) \cap \mathsf{sig}(C_2, \mathcal{T}_2)$. Suppose $P \in \Sigma_1 \cap \Sigma_2$. Then $P \in \mathsf{range}(\zeta_1) \cap \mathsf{range}(\zeta_2)$. Since ζ_1 and ζ_2 are defined as restrictions of ζ to $\mathsf{rol}(C_1, \mathcal{T}_1)$ and $\mathsf{rol}(C_2, \mathcal{T}_2)$, respectively, there is some $R^- \in \mathsf{rol}(C_1, \mathcal{T}_1) \cap \mathsf{rol}(C_2, \mathcal{T}_2)$ such that $\zeta(R^-) = P$. But then $R \in \mathsf{sig}(C_1, \mathcal{T}_1) \cap \mathsf{sig}(C_2, \mathcal{T}_2)$.

Now by the claim we have just shown, $\mathsf{sig}(I) \subseteq (\mathsf{sig}(C_1, \mathcal{T}_1) \cap \mathsf{sig}(C_2, \mathcal{T}_2)) \cup (\Sigma_1 \cap \Sigma_2)$, and the construction of D, we have $\mathsf{sig}(D) \subseteq \mathsf{sig}(C_1, \mathcal{T}_1) \cap \mathsf{sig}(C_2, \mathcal{T}_2)$. Moreover, by $Z_\zeta(D) = I$, items 2 and 3 above, and Lemma 3.19, we obtain $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq D$ and $\mathcal{T}_1 \cup \mathcal{T}_2 \models D \sqsubseteq C_2$. Hence D is an interpolant of C_1 and C_2 under $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$. It is easy to see that the time required to compute D is as stated in the theorem.

For $\mathscr{L} \in \{\mathcal{S}, \mathcal{SI}, \mathcal{SF}, \mathcal{SIF}\}$. In what follows, let $\mathscr{L}'$ be $\mathscr{L}$ without the transitive role constructor, e.g., if $\mathscr{L} = \mathcal{SIF}$, then $\mathscr{L}' = \mathcal{ALCFI}$. We know by now that $\mathscr{L}'$ satisfies what is stated in the theorem. Suppose that C_1, C_2 are $\mathscr{L}$-concepts and $\mathcal{T}_1, \mathcal{T}_2$ are $\mathscr{L}$-TBoxes such that $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2$. The proof proceeds analogously to the inverse role case, except of course we use Lemma 3.14.

To conclude, we have shown for each logic $\mathscr{L}$ stated in the theorem a constructive way to compute an interpolant, if one exists, in time double exponential in the size of the input. Hence the theorem follows. ☐

3.3 Shorter First-Order Interpolants

We will now show that our interpolation algorithm can be adapted to compute first-order interpolants in single exponential time. The proof will proceed along the following lines. First we will show that the double exponential size of the interpolants is only due to the repeated occurrence of subformulas and that our algorithm yields single exponential size interpolants using a succinct (DAG-shaped as opposed to tree-shaped) concept representation. Next we apply an idea implicit in the work of Avigad (2003), namely that succinctly represented first-order formulas can be transformed in polynomial time into equivalent ordinary tree-shaped first-order formulas over structures with at least two elements. This allows us to compute single exponential first-order interpolants over structures with at least two elements. After that, we show that single exponential interpolants over structures with one element can be constructed by a reduction to propositional logic. By combining the interpolants obtained via these two methods, we finally obtain the desired single exponential first-order interpolant over arbitrary structures.

Step 1: Singly-exponential interpolants via succinct representation We start by defining the notions that will allow us to represent DAG-shaped concepts.

Definition 3.23. *Fix a description logic $\mathscr{L}$. An axiom of the form $A \equiv C$, where $A \in N_C$ and C is an $\mathscr{L}$-concept, is called a* concept definition axiom *in $\mathscr{L}$ (or, an $\mathscr{L}$-CDA). Let Σ be a signature. An* acyclic terminology over Σ *in $\mathscr{L}$ is a set of $\mathscr{L}$-CDAs*

$$\mathcal{T} = \{A_1 \equiv C_1, \ldots, A_n \equiv C_n\}$$

where $\{A_1, \ldots, A_n\} \cap \Sigma = \emptyset$ and $\mathsf{sig}(C_i) \subseteq \Sigma \cup \{A_1, \ldots, A_{i-1}\}$ for $i \in \{1, \ldots, n\}$.

A succinct-$\mathscr{L}$-concept over Σ *is a pair $\langle A, \mathcal{T} \rangle$, where $\mathcal{T}$ is an acyclic terminology over Σ in $\mathscr{L}$ and A is a concept name belonging to $\mathsf{sig}(\mathcal{T}) \setminus \Sigma$. The* unfolding *of a succinct-$\mathscr{L}$-concept $\langle A, \mathcal{T} \rangle$ is the $\mathscr{L}$-concept over Σ that is obtained from A by repeatedly "applying"* *the CDAs in $\mathcal{T}$, i.e., replacing occurrences of their left-hand side by their right-hand side, until no more CDA can be applied.*

Note that acyclic terminologies are well-known in the DL literature (Baader & Nutt, 2003).

Example 3.24. *Let $\mathcal{T}$ consist of the following.*

$$\begin{array}{rcl}\mathsf{Woman} & \equiv & \mathsf{Person} \sqcap \mathsf{Female}\\ \mathsf{Man} & \equiv & \mathsf{Person} \sqcap \mathsf{Male}\\ \mathsf{Human} & \equiv & \mathsf{Woman} \sqcup \mathsf{Man}\end{array}$$

Then $\mathcal{T}$ is an acyclic terminology over $\{\mathsf{Person}, \mathsf{Female}, \mathsf{Male}\}$. The unfolding of the succinct-concept $\langle \mathsf{Human}, \mathcal{T}\rangle$ is

$$(\mathsf{Person} \sqcap \mathsf{Female}) \sqcup (\mathsf{Person} \sqcap \mathsf{Male}).$$

The unfolding of a succinct-concept is in general exponentially longer.

Proposition 3.25. *Let $\mathscr{L}$ be any description logic. For each succinct-$\mathscr{L}$-concept $\langle A, \mathcal{T}\rangle$ with unfolding C, $|C| \in 2^{|\mathcal{T}|^{O(1)}}$.*

Theorem 3.26. *Let $\mathscr{L}$ be $\mathcal{ALC}$ or any of its extensions with constructors from $\{\mathcal{S}, \mathcal{I}, \mathcal{F}\}$. For all $\mathscr{L}$-concepts C_1, C_2 and all $\mathscr{L}$-TBoxes $\mathcal{T}_1, \mathcal{T}_2$, if $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2$, then there exists a succinct-$\mathscr{L}$-concept $\langle A, \mathcal{T}\rangle$ over $\mathsf{sig}(C_1, \mathcal{T}_1) \cap \mathsf{sig}(C_2, \mathcal{T}_2)$ such that*

- *the unfolding of $\langle A, \mathcal{T}\rangle$ is an interpolant of C_1 and C_2 under $\langle \mathcal{T}_1, \mathcal{T}_2\rangle$, and*
- *$\langle A, \mathcal{T}\rangle$ can be computed in time single exponential in $|\mathcal{T}_1| + |\mathcal{T}_2| + |C_1| + |C_2|$.*

Proof. Let $\mathscr{L}$ be one of the DLs mentioned in the theorem, let $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2$, where $\mathcal{T}_1, \mathcal{T}_2$ are $\mathscr{L}$-TBoxes and C_1, C_2 are $\mathscr{L}$-concepts, and let $m = |\mathcal{T}_1| + |\mathcal{T}_2| + |C_1| + |C_2|$. As in the proof of Theorem 3.22, we first reduce $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2$ to $\mathcal{T}_1' \cup \mathcal{T}_2' \models D_1 \sqsubseteq D_2$, where $\mathcal{T}_1', \mathcal{T}_2'$ are $\mathcal{ALC}$-TBoxes ($\mathcal{ALCF}$-TBoxes) and D_1, D_2 are $\mathcal{ALC}$-concepts (resp. $\mathcal{ALCF}$-concepts).

We show that the interpolant calculation step in Algorithm 1 for $\mathcal{ALCF}$ (and thus $\mathcal{ALC}$, see Figures 2, 3, 4) can be modified to compute a succinct-concept of single exponential size as an interpolant, instead of a concept.

We associate to every node g in the tableau a distinct fresh concept name X_g. The new algorithm still uses the same interpolation calculation rules but instead of directly assigning an interpolant to every node g with $g.\mathsf{status} = \mathtt{unsat}$, we construct an acyclic terminology $\mathcal{T}'$ over $\mathsf{sig}(D_1, \mathcal{T}_1') \cap \mathsf{sig}(D_2, \mathcal{T}_2')$, where the acyclic terminology makes use of the new concept names X_g, and such that the unfolding of the succinct-concept $\langle X_g, \mathcal{T}'\rangle$ is an interpolant for $g.\mathsf{content}$ whenever $g.\mathsf{status} = \mathtt{unsat}$. The set $\mathcal{T}'$ is initialized as an empty set, and throughout the computation of the algorithm, $\mathcal{T}'$ is extended in the natural way. For instance, suppose $\mathsf{C}^{\mathbf{l}}_{\sqcup}$ is applied to g. Then $\mathsf{C}^{\mathbf{l}}_{\sqcup}$ adds to $\mathcal{T}'$ the CDA $X_g \equiv X_{g_1} \sqcup X_{g_2}$, where g_1 and g_2 are the successors of the node g in the tableau. Another example is a clash rule. Suppose $\mathsf{C}^{\mathbf{lr}}_{\neg}$ is applied to g for some $\{C^{\mathbf{l}}, (\dot{\neg} C)^{\mathbf{r}}\} \subseteq g.\mathsf{content}$. Then $\mathsf{C}^{\mathbf{lr}}_{\neg}$ adds to $\mathcal{T}'$ the CDA $X_g \equiv C$. By Lemma 3.5, it follows that $|\mathcal{T}'| \leq 2^{O(m)}$; and by the definition of Algorithm 1, it follows that $\mathcal{T}'$ is an acyclic terminology over $\mathsf{sig}(D_1, \mathcal{T}_1') \cap \mathsf{sig}(D_2, \mathcal{T}_2')$. Moreover, by $\mathcal{T}_1' \cup \mathcal{T}_2' \models D_1 \sqsubseteq D_2$, there is some $X_{g_0} \equiv C \in \mathcal{T}'$. Then $\langle X_{g_0}, \mathcal{T}' \rangle$ is a succinct-concept over $\mathsf{sig}(D_1, \mathcal{T}_1') \cap \mathsf{sig}(D_2, \mathcal{T}_2')$ and its unfolding can easily be shown to be an interpolant of D_1 and D_2 under $\mathcal{T}_1' \cup \mathcal{T}_2'$.

In a way similar to the proof of Theorem 3.22, i.e., by replacing back the newly introduced role names for inverse and transitive roles in $\mathcal{T}'$ with the originals, we obtain a new terminology $\mathcal{T}''$. Then the unfolding of $\langle X_{g_0}, \mathcal{T}'' \rangle$ is guaranteed to be an interpolant of C_1 and C_2 under $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$. Moreover, $\langle X_{g_0}, \mathcal{T}'' \rangle$ is of size single exponential in m. $\square$

For the rest of the section, our purpose is to obtain an equivalent first-order formula from a given succinct-concept in polynomial time. We will make use of the standard translation (see Definition 2.2). In the following, we will not distinguish between DL interpretations and first-order structures (we choose the unary and binary predicates of our first-order language to be the symbols in N_C and N_R, respectively).

Step 2: Singly-exponential FO interpolants for interpretations with two elements For a first-order formula $\varphi(x)$ and an interpretation $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$ with $s \in \Delta^{\mathcal{I}}$, we write $\mathcal{I}, s \models \varphi(x)$ if and only if there is some first-order assignment α such that $\alpha(x) = s$ and $\mathcal{I}, \alpha \models \varphi(x)$. By $\models_{\geq 2}$, we denote the restriction of the relation $\models$ that only considers interpretations $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$ where $|\Delta^{\mathcal{I}}| \geq 2$. Similarly, by $\models_{=1}$, we denote the restriction of the relation $\models$ that only considers interpretations $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$ where $|\Delta^{\mathcal{I}}| = 1$.

The proof of the following theorem is inspired by a result of Avigad (2003), which states that, over structures with at least two elements, one can efficiently eliminate acyclic definitions from proofs. Theorem 3.27 can be viewed as an adaptation of this result to the first-order translation of succinct-concepts in description logic.

Theorem 3.27. *Given a succinct-$\mathcal{SHIF}$-concept $\langle B, \mathcal{T} \rangle$ over a signature Σ, we can construct in polynomial time a first-order formula $\psi(x)$ over Σ, such that $\models_{\geq 2} \psi(x) \leftrightarrow \pi_x(C)$, where C is the unfolding of $\langle B, \mathcal{T} \rangle$.*

Our proof of Theorem 3.27 will be based on a lemma that we state next. For expository reasons, it is more convenient to state the lemma in terms of structures with constant symbols. These constant symbols are not needed for Theorem 3.27. They are only used to make the statement and proof of the following lemma more readable.

Lemma 3.28. *Given an acyclic terminology $\mathcal{T} = \{A_1 \equiv C_1, \ldots, A_n \equiv C_n\}$ in $\mathcal{SHIF}$, we can construct in polynomial time a first-order formula $\varphi_{\mathcal{T}}(x, y_1, \ldots, y_n, z)$ with additional constant symbols $\mathbf{0}$ and $\mathbf{1}$, such that, for all interpretations $\mathcal{I}$ satisfying $\mathbf{0}^{\mathcal{I}} \neq \mathbf{1}^{\mathcal{I}}$, and for all elements $a, \vec{b}, c \in \Delta^{\mathcal{I}}$ (where $\vec{b} = b_1, \ldots, b_n$),*

$$\mathcal{I} \models \varphi_{\mathcal{T}}\ [a, \vec{b}, c] \textit{ if and only if } \vec{b} = \underline{k} \textit{ for some } k \in \{1, \ldots, n\}, \textit{ and } c = \begin{cases} \mathbf{1}^{\mathcal{I}} & \textit{if } a \in C_k^{\mathcal{I}} \\ \mathbf{0}^{\mathcal{I}} & \textit{otherwise} \end{cases}$$

where $\underline{k} = \underbrace{\mathbf{0}^{\mathcal{I}} \cdots \mathbf{0}^{\mathcal{I}}}_{k-1 \textit{ times}} \mathbf{1}^{\mathcal{I}} \underbrace{\mathbf{0}^{\mathcal{I}} \cdots \mathbf{0}^{\mathcal{I}}}_{n-k \textit{ times}}$ *and C_k is the unfolding of the succinct-concept $\langle A_k, \mathcal{T} \rangle$.*

Proof. We define $\varphi_{\mathcal{T}}$ by induction on the number n of CDAs in $\mathcal{T}$. If $n = 1$, then we can simply define $\varphi_{\mathcal{T}}(x, y, z)$ as

$$\varphi_{\mathcal{T}}(x,y,z) \quad = \quad (y = \mathbf{1}) \wedge ((\pi_x(C_1) \wedge z = \mathbf{1}) \vee (\neg\pi_x(C_1) \wedge z = \mathbf{0}))$$

Now, let $n > 1$ and let $\mathcal{T}'$ be obtained from $\mathcal{T}$ by removing the last CDA. In other words, let $\mathcal{T} = \mathcal{T}' \cup \{A_n \equiv C_n\}$. By induction hypothesis, there is a formula $\varphi_{\mathcal{T}'}(u, \vec{v}, w)$ satisfying the required conditions w.r.t. $\mathcal{T}'$ (where $\vec{v} = v_1, \ldots, v_{n-1}$). We can distinguish the following cases:

1. C_n is an atomic concept or functionality restriction over the signature Σ. In this case, we can define $\varphi_{\mathcal{T}}$ as follows, where $\vec{y} = y_1 \ldots y_n$ and $\vec{v} = v_1 \ldots v_{n-1}$.

$$\begin{aligned}\varphi_{\mathcal{T}}(x,\vec{y},z) \quad = \quad & \exists u, \vec{v}, w(\varphi_{\mathcal{T}'}(u,\vec{v},w) \wedge x = u \wedge \vec{y} = \vec{v}\mathbf{0} \wedge z = w) \vee \\ & (\vec{y} = \mathbf{0}\cdots\mathbf{01} \wedge ((\pi_x(C_n) \wedge z = \mathbf{1}) \vee (\neg\pi_x(C_n) \wedge z = \mathbf{0}))))\end{aligned}$$

Here, $\vec{y} = \vec{v}\mathbf{0}$ is a shorthand for the formula $\bigwedge_{i<n} y_i = v_i \wedge y_n = \mathbf{0}$, and, similarly, $\vec{y} = \mathbf{0}\cdots\mathbf{01}$ is shorthand for the formula $\bigwedge_{i<n} y_i = \mathbf{0} \wedge y_n = \mathbf{1}$.

2. C_n is of the form $\neg A_i$ with $i < n$. In this case, we can define $\varphi_{\mathcal{T}}$ as follows:

$$\begin{aligned}\varphi_{\mathcal{T}}(x,\vec{y},z) \quad = \quad & \exists u, \vec{v}, w\big(\varphi_{\mathcal{T}'}(u,\vec{v},w) \wedge \\ & ((x = u \wedge \vec{y} = \vec{v}\mathbf{0} \wedge z = w) \vee \\ & \;(\vec{y} = \mathbf{0}\cdots\mathbf{01} \wedge u = x \wedge \vec{v} = \underline{i} \wedge ((w = \mathbf{1} \wedge z = \mathbf{0}) \vee (w = \mathbf{0} \wedge z = \mathbf{1}))))\big)\end{aligned}$$

Here, the same notation conventions apply as in the previous item. In addition, $\vec{v} = \underline{i}$ is used as a shorthand for the formula $v_i = \mathbf{1} \wedge \bigwedge_{j \neq i} v_j = \mathbf{0}$. The notations will also be used in the following items.

3. C_n is of the form $A_i \sqcap A_j$ with $i, j < n$. As a first attempt, define $\varphi_{\mathcal{T}}(x, \vec{y}, z)$ as follows:

$$\begin{aligned}\varphi_{\mathcal{T}}(x,\vec{y},z) \quad = \quad & \exists u, \vec{v}, w\ \exists u', \vec{v}', w'\big(\varphi_{\mathcal{T}'}(u,\vec{v},w) \wedge \varphi_{\mathcal{T}'}(u',\vec{v}',w') \wedge \\ & ((x = u \wedge \vec{y} = \vec{v}\mathbf{0} \wedge z = w) \vee \\ & \;(\vec{y} = \mathbf{0}\cdots\mathbf{01} \wedge u = u' = x \wedge \vec{v} = \underline{i} \wedge \vec{v}' = \underline{j} \wedge \\ & \;\;(w = w' = z = \mathbf{1} \vee ((w = \mathbf{0} \vee w' = \mathbf{0}) \wedge z = \mathbf{0}))))\big)\end{aligned}$$

This works, except for the fact that $\varphi_{\mathcal{T}'}$ occurs twice in the formula, which may result in an exponential blowup. We solve this problem by replacing the conjunction $\varphi_{\mathcal{T}'}(u, \vec{v}, w) \wedge \varphi_{\mathcal{T}'}(u', \vec{v}', w')$ by

$$\forall u'', \vec{v}'', w''((u'' = u \wedge \vec{v}'' = v \wedge w'' = w) \vee (u'' = u' \wedge \vec{v}'' = v' \wedge w'' = w') \rightarrow \varphi_{\mathcal{T}'}(u'', \vec{v}'', w''))$$

4. C_n is of the form $\exists P.A_i$ with $i < n$. This is the most difficult case. The following formula expresses the required property:

$$\begin{aligned}\varphi_{\mathcal{T}}(x,\vec{y},z) \quad = \quad & \exists u, \vec{v}, w\ \big(\varphi_{\mathcal{T}'}(u,\vec{v},w) \wedge \\ & ((x = u \wedge \vec{y} = \vec{v}\mathbf{0} \wedge z = w) \vee \\ & (\vec{y} = \mathbf{0}\cdots\mathbf{01} \wedge z = \mathbf{1} \wedge Pxu \wedge \vec{v} = \underline{i} \wedge w = \mathbf{1}) \vee \\ & (\vec{y} = \mathbf{0}\cdots\mathbf{01} \wedge z = \mathbf{0} \wedge \\ & \quad \forall u', \vec{v}', w'(\varphi_{\mathcal{T}'}(u',\vec{v}',w') \wedge Pxu' \wedge \vec{v}' = \underline{i} \rightarrow w' = \mathbf{0}))))\big)\end{aligned}$$

However, as before, this formula still has the problem that it contains two copies of $\varphi_{\mathcal{T}'}$. We fix this in two steps. First, we bring the universal quantifiers to the front, and transform the above formula into the following equivalent formula:

$$\begin{aligned}&\exists u, \vec{v}, w\ \forall u', \vec{v}', w' \big(\varphi_{\mathcal{T}'}(u, \vec{v}, w) \wedge \varphi_{\mathcal{T}'}(u', \vec{v}', w') \ \rightarrow \\ &((x = u \wedge \vec{y} = \vec{v}\mathbf{0} \wedge z = w) \vee \\ &\ (\vec{y} = \mathbf{0}\cdots\mathbf{01} \wedge z = \mathbf{1} \wedge Pxu \wedge \vec{v} = \underline{i} \wedge w = \mathbf{1}) \vee \\ &\ (\vec{y} = \mathbf{0}\cdots\mathbf{01} \wedge z = \mathbf{0} \wedge (Pxu' \wedge \vec{v}' = \underline{i} \rightarrow w' = \mathbf{0}))))\big)\end{aligned}$$

Finally, as before, we replace the conjunction $\varphi_{\mathcal{T}'}(u, \vec{v}, w) \wedge \varphi_{\mathcal{T}'}(u', \vec{v}'', w')$ by

$$\forall u'', \vec{v}'', w''((u'' = u \wedge \vec{v}'' = v \wedge w'' = w) \vee (u'' = u' \wedge \vec{v}'' = v' \wedge w'' = w') \rightarrow \varphi_{\mathcal{T}'}(u'', \vec{v}'', w''))$$

5. C_n is of the form $\exists P^-.A_i$ with $i < n$. This case is handled like the previous one.

Note that, in general, C_n could be a complex concept in which various A_i with $i < n$ occur. However, such complex CDAs can always be decomposed into multiple simpler CDAs of the above kinds, at the cost of a polynomial increase in the size of the terminology.

It is clear from the construction that the formula $\varphi_{\mathcal{T}}$ obtained as above satisfies the conditions stated in the lemma. That $\varphi_{\mathcal{T}}$ is obtained from $\mathcal{T}$ in polynomial-time follows from the fact that, in the above inductive definition of $\varphi_{\mathcal{T}}$, the previously constructed formula $\varphi_{\mathcal{T}'}$ occurs only once. $\square$

We are now ready for the proof of Theorem 3.27.

Proof of Theorem 3.27. Let a succinct-concept $\langle A_i, \mathcal{T}\rangle$ be given, where $\mathcal{T} = \{A_1 \equiv C_1, \ldots, A_n \equiv C_n\}$. Let $\varphi(x) = \varphi_{\mathcal{T}}(x, \underline{i}, \mathbf{1})$ and let $\psi(x) = \exists u, v(u \neq v \wedge \varphi'(x))$, where $\varphi'(x)$ is obtained from $\varphi(x)$ by replacing $\mathbf{0}$ and $\mathbf{1}$ by u and v, respectively. Then we have that, for every interpretation $\mathcal{I}$ with a domain of at least two elements, and for every $a \in \Delta^{\mathcal{I}}$, the following conditions are all equivalent:

1. $\mathcal{I}, a \models \psi(x)$
2. $\mathcal{I}', a \models \varphi(x)$, for some interpretation $\mathcal{I}'$ that extends $\mathcal{I}$ by mapping the constant symbols $\mathbf{0}$ and $\mathbf{1}$ to distinct elements of $\Delta^{\mathcal{I}}$.
3. $\mathcal{I}', a \models C$, where C is the unfolding of $\langle A_i, \mathcal{T}\rangle$.
4. $\mathcal{I}, a \models C$, where C is the unfolding of $\langle A_i, \mathcal{T}\rangle$.

The equivalence of 1 and 2 is immediate from the construction of ψ. The equivalence of 2 and 3 follows from Lemma 3.28. The equivalence of 3 and 4 is immediate, since $\mathbf{0}$ and $\mathbf{1}$ do not occur in C. This concludes the proof. $\square$

Definition 3.29. *Let C, D be $\mathscr{L}$-concepts and let $\mathcal{T}_1, \mathcal{T}_2$ be $\mathscr{L}$-TBoxes such that $\mathcal{T}_1 \cup \mathcal{T}_2 \models C \sqsubseteq D$. A first-order formula $\varphi(x)$ is called a* FO interpolant *of C and D under $\langle \mathcal{T}_1, \mathcal{T}_2\rangle$ if the following conditions hold:*

- $\mathsf{sig}(\varphi(x)) \subseteq \mathsf{sig}(C, \mathcal{T}_1) \cap \mathsf{sig}(D, \mathcal{T}_2)$,

- $\pi(\mathcal{T}_1) \cup \pi(\mathcal{T}_2) \models \forall x.\pi_x(C) \to \varphi(x)$, *and*
- $\pi(\mathcal{T}_1) \cup \pi(\mathcal{T}_2) \models \forall x.\varphi(x) \to \pi_x(D)$.

FO $\models_{\geq 2}$*-interpolant and FO* $\models_{=1}$*-interpolant are defined in the same way as above, except that we replace all occurrences of* $\models$ *by* $\models_{\geq 2}$ *and* $\models_{=1}$*, respectively.*

Proposition 3.30. *Let* $\mathscr{L}$ *be* $\mathcal{ALC}$ *or any of its extensions with constructors from* $\{\mathcal{S}, \mathcal{I}, \mathcal{F}\}$*. For all* $\mathscr{L}$*-concepts* C_1, C_2 *and all* $\mathscr{L}$*-TBoxes* $\mathcal{T}_1, \mathcal{T}_2$*, if* $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2$*, then there exists a FO* $\models_{\geq 2}$*-interpolant of* C_1 *and* C_2 *under* $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$ *that can be computed in time single exponential in* $|\mathcal{T}_1| + |\mathcal{T}_2| + |C_1| + |C_2|$*.*

Proof. Suppose $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2$. By Theorem 3.26, there is some succinct-concept $\langle A, \mathcal{T} \rangle$ over $\mathsf{sig}(C_1, \mathcal{T}_1) \cap \mathsf{sig}(C_2, \mathcal{T}_2)$ such that the unfolding I of $\langle A, \mathcal{T} \rangle$ is an interpolant of C_1 and C_2 under $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$, and $\langle A, \mathcal{T} \rangle$ can be computed in time single exponential in $|\mathcal{T}_1|+|\mathcal{T}_2|+|C_1|+|C_2|$. Then by Theorem 3.27, there is some first-order formula $\varphi(x)$ that can be constructed in time polynomial in $|\mathcal{T}|$ (hence single exponential in $|\mathcal{T}_1|+|\mathcal{T}_2|+|C_1|+|C_2|$) such that

- $\mathsf{sig}(\varphi(x)) \subseteq \mathsf{sig}(I)$,
- $\models_{\geq 2} \varphi(x) \leftrightarrow \pi_x(I)$.

It follows that $\varphi(x)$ is a FO $\models_{\geq 2}$-interpolant of C_1 and C_2 under $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$ whose size is single exponential in $|\mathcal{T}_1| + |\mathcal{T}_2| + |C_1| + |C_2|$. ◻

Step 3: Singly-exponential FO interpolants for interpretations with one element We still have to obtain interpolants over structures with only one element. We will show how to do this in Proposition 3.35. The essential idea is that interpolants over structures with singleton domains are not much different from propositional interpolants. First, we give a reduction from concept subsumption w.r.t. TBoxes over interpretations with singleton domains to entailment in propositional logic.

Definition 3.31. *Let* C *be a* $\mathcal{SHIF}$*-concept. Then the mapping* $\tau_{\mathsf{PL}}(C)$ *is defined inductively as follows.*

$$\begin{array}{rcl}
\tau_{\mathsf{PL}}(\top) & = & \top, \\
\tau_{\mathsf{PL}}(A) & = & A, \\
\tau_{\mathsf{PL}}(\neg C) & = & \neg \tau_{\mathsf{PL}}(C), \\
\tau_{\mathsf{PL}}(C \sqcap D) & = & \tau_{\mathsf{PL}}(C) \sqcap \tau_{\mathsf{PL}}(D), \\
\tau_{\mathsf{PL}}(\exists R.C) & = & A_R \sqcap \tau_{\mathsf{PL}}(C), \\
\tau_{\mathsf{PL}}(\leq 1 R) & = & \top,
\end{array}$$

where $A_P = A_{P^-}$ *is a fresh concept name for every role name* $P \in N_R$*. For a* $\mathcal{SHIF}$*-TBox* $\mathcal{T}$*, we define*

$$\tau_{\mathsf{PL}}(\mathcal{T}) = \{\tau_{\mathsf{PL}}(C) \sqsubseteq \tau_{\mathsf{PL}}(D) \mid C \sqsubseteq D \in \mathcal{T}\} \cup \{A_R \sqsubseteq A_S \mid R \sqsubseteq S \in \mathcal{T}\}.$$

Here, the concept name A_P, intuitively, expresses the non-emptiness of the role P. Note that all transitive roles are ignored in the above translation, as their semantics is trivially satisfied in interpretations whose domain is a singleton set. For a $\mathcal{SHIF}$-concept C, $\tau_{\mathsf{PL}}(C)$ is an $\mathcal{ALC}$-concept without role constructors. We view $\tau_{\mathsf{PL}}(C)$ as a propositional formula (where the concept names are the propositions, and we identify $\sqcap$ and $\sqcup$ with the propositional connectives $\wedge$ and $\vee$, respectively). Similarly, for a $\mathcal{SHIF}$-TBox $\mathcal{T}$, $\tau_{\mathsf{PL}}(\mathcal{T})$ is a set of $\mathcal{ALC}$ CIAs without role constructors, which we view as a set of propositional formulae.

Proposition 3.32. *Let $\mathcal{T}$ be a $\mathcal{SHIF}$-TBox and let C, D be $\mathcal{SHIF}$-concepts. Then*

$$\mathcal{T} \models_{=1} C \sqsubseteq D \quad \textit{if and only if} \quad \tau_{\mathsf{PL}}(\mathcal{T}) \models \tau_{\mathsf{PL}}(C) \sqsubseteq \tau_{\mathsf{PL}}(D).$$

Proof. ($\Rightarrow$) Let $\mathcal{T} \models_{=1} C \sqsubseteq D$, and suppose $\mathcal{I} \models \tau_{\mathsf{PL}}(\mathcal{T})$, and $s \in \tau_{\mathsf{PL}}(C)^{\mathcal{I}}$. We need to show that $s \in \tau_{\mathsf{PL}}(D)^{\mathcal{I}}$. Let $\mathcal{J}$ be obtained by restricting the domain of $\mathcal{I}$ to the element s and "reading off" the interpretation of each role name P from the concept name A_P. Formally,

- $\Delta^{\mathcal{J}} = \{s\}$;
- for all $A \in N_C$, $s \in A^{\mathcal{J}}$ iff $s \in A^{\mathcal{I}}$;
- for all $P \in N_R$, $P^{\mathcal{J}} = \{\langle s, s\rangle\}$ if $s \in A_P^{\mathcal{I}}$, and $P^{\mathcal{J}} = \emptyset$, otherwise.

By the definition above, it trivially follows for every $P \in N_{R^+}$ that $P^{\mathcal{J}}$ is transitive. Moreover, for every role R, we have

$$\langle s, s\rangle \in R^{\mathcal{J}} \text{ if and only if } s \in A_R^{\mathcal{I}}. \tag{1}$$

To see this, suppose first $\langle s, s\rangle \in R^{\mathcal{J}}$. If $R = P$ for some $P \in N_R$, then $s \in A_P^{\mathcal{I}}$, i.e., $s \in A_R^{\mathcal{I}}$; and if $R = P^-$ for some $P \in N_R$, then again $s \in A_P^{\mathcal{I}}$ and by the fact that $A_P = A_{P^-}$ (see Definition 3.31), we obtain $s \in A_R^{\mathcal{I}}$. Hence $s \in A_R^{\mathcal{I}}$. For the other direction, suppose $s \in A_R^{\mathcal{I}}$. If $R = P$ for some $P \in N_R$, then $\langle s, s\rangle \in P^{\mathcal{J}}$, i.e., $\langle s, s\rangle \in R^{\mathcal{J}}$; and if $R = P^-$ for some $P \in N_R$, then by the fact that $A_P = A_{P^-}$, we have $\langle s, s\rangle \in P^{\mathcal{J}}$ and thus, $\langle s, s\rangle \in R^{\mathcal{J}}$. Hence (1) follows.

Claim 3.33. *For every $\mathcal{SHIF}$-concept C', we have $s \in \tau_{\mathsf{PL}}(C')^{\mathcal{I}}$ if and only if $s \in (C')^{\mathcal{J}}$.*

Proof of claim. The proof is by induction on the structure of C'. The base case, where $C' = A$ or $C' = \top$, is trivial, and the boolean cases follow immediately by the inductive hypothesis. For $C' = \exists R.D'$, we have the following equivalences:

- $s \in \tau_{\mathsf{PL}}(C')^{\mathcal{I}}$;
- $s \in A_R^{\mathcal{I}}$ and $s \in \tau_{\mathsf{PL}}(D')^{\mathcal{I}}$ (by semantics);
- $\langle s, s\rangle \in R^{\mathcal{J}}$ and $s \in (D')^{\mathcal{J}}$ (by (1) and the inductive hypothesis);
- $s \in (C')^{\mathcal{J}}$ (by semantics).

Finally, for $C' = \leq 1R$, since $\tau_{\mathsf{PL}}(C') = \top$ and $s \in \Delta^{\mathcal{I}}$, we have $s \in \tau_{\mathsf{PL}}(C')^{\mathcal{I}}$. Moreover, by the definition of $\mathcal{J}$, we have $s \in (C')^{\mathcal{J}}$. But then $s \in \tau_{\mathsf{PL}}(C')^{\mathcal{I}}$ iff $s \in (C')^{\mathcal{J}}$, which is what we wanted to show. $\dashv$

We now show that $\mathcal{J} \models \mathcal{T}$, i.e., $\mathcal{J}$ satisfies every CIA and RIA in $\mathcal{T}$. That $\mathcal{J}$ satisfies every CIA in $\mathcal{T}$ is a direct consequence of the previous claim; so we proceed with the case for RIAs. Let $R \sqsubseteq S \in \mathcal{T}$ and $\langle s, t\rangle \in R^{\mathcal{J}}$. By the definition of $\mathcal{J}$, we have $s = t$. Hence w.l.o.g. suppose that $\langle s, s\rangle \in R^{\mathcal{J}}$. Then by (1), $s \in A_R^{\mathcal{I}}$. Since $A_R \sqsubseteq A_S \in \tau_{\mathsf{PL}}(\mathcal{T})$ and $\mathcal{I} \models \tau_{\mathsf{PL}}(\mathcal{T})$, we then have $s \in A_S^{\mathcal{I}}$. By (1) again, this implies $\langle s, s\rangle \in S^{\mathcal{J}}$. Hence $\mathcal{J}$ satisfies $R \sqsubseteq S$.

Now we proceed towards our goal $s \in \tau_{\mathsf{PL}}(D)^{\mathcal{I}}$ as follows. By $\mathcal{I} \models \tau_{\mathsf{PL}}(\mathcal{T})$, $s \in \tau_{\mathsf{PL}}(C)^{\mathcal{I}}$, and the previous claim, we obtain $s \in C^{\mathcal{J}}$. Since $\mathcal{J} \models \mathcal{T}$, it follows by $\mathcal{T} \models C \sqsubseteq D$ that $s \in D^{\mathcal{J}}$. Then using the previous claim, we conclude that $s \in \tau_{\mathsf{PL}}(D)^{\mathcal{I}}$.

($\Leftarrow$) Let $\tau_{\mathsf{PL}}(\mathcal{T}) \models \tau_{\mathsf{PL}}(C) \sqsubseteq \tau_{\mathsf{PL}}(D)$, and suppose $\mathcal{I} \models \mathcal{T}$, and $s \in C^{\mathcal{I}}$, where $\Delta^{\mathcal{I}} = \{s\}$. We need to show that $s \in D^{\mathcal{I}}$. Define the interpretation $\mathcal{J}$ as follows:

- $\Delta^{\mathcal{J}} = \{s\}$;
- for all $A \in N_C$, $A^{\mathcal{J}} = A^{\mathcal{I}}$
- for all $P \in N_R$, $(A_P)^{\mathcal{J}} = \{s\}$ if $P^{\mathcal{I}} = \{\langle s, s\rangle\}$, and $(A_P)^{\mathcal{J}} = \emptyset$ otherwise

We first show for every role R that

$$s \in A_R^{\mathcal{J}} \text{ if and only if } \langle s, s\rangle \in R^{\mathcal{I}}. \tag{2}$$

For left-to-right, suppose $s \in A_R^{\mathcal{J}}$. If $R = P$ for some $P \in N_R$, then $\langle s, s\rangle \in P^{\mathcal{I}}$, i.e., $\langle s, s\rangle \in R^{\mathcal{I}}$; and if $R = P^-$ for some $P \in N_R$, then by the fact that $A_P = A_{P^-}$, we have $s \in A_P^{\mathcal{J}}$, which implies $\langle s, s\rangle \in R^{\mathcal{I}}$. Hence $\langle s, s\rangle \in R^{\mathcal{I}}$. For the other direction, suppose $\langle s, s\rangle \in R^{\mathcal{I}}$. If $R = P$ for some $P \in N_R$, then $s \in A_P^{\mathcal{J}}$, i.e., $s \in A_R^{\mathcal{J}}$; and if $R = P^-$ for some $P \in N_R$, then $\langle s, s\rangle \in P^{\mathcal{I}}$, which implies by $A_P = A_{P^-}$ that $s \in A_R^{\mathcal{J}}$. Hence (2) follows.

Claim 3.34. *For every $\mathcal{SHIF}$-concept C', we have $s \in (C')^{\mathcal{I}}$ if and only if $s \in \tau_{\mathsf{PL}}(C')^{\mathcal{J}}$.*

Proof of claim. The proof is by induction on the structure of C'. The base case, where $C' = A$ or $C' = \top$, is trivial, and the boolean cases follow immediately by the inductive hypothesis. For $C' = \exists R.D'$, we have the following equivalences

- $s \in (C')^{\mathcal{I}}$;
- $\langle s, s\rangle \in R^{\mathcal{I}}$ and $s \in (D')^{\mathcal{I}}$ (by semantics and $\Delta^{\mathcal{I}} = \{s\}$);
- $s \in A_R^{\mathcal{J}}$ and $s \in \tau_{\mathsf{PL}}(D')^{\mathcal{J}}$ (by (2) and the inductive hypothesis).
- $s \in \tau_{\mathsf{PL}}(C')^{\mathcal{J}}$ (by semantics).

Finally, for $C' = \leq 1R$, since $\Delta^{\mathcal{I}} = \{s\}$, we have $\mathcal{I} \models \top \sqsubseteq \leq 1R$, and thus, $s \in (C')^{\mathcal{I}}$. Moreover, by $\tau_{\mathsf{PL}}(C') = \top$, we have $s \in \tau_{\mathsf{PL}}(C')^{\mathcal{J}}$. But then $s \in (C')^{\mathcal{I}}$ iff $s \in \tau_{\mathsf{PL}}(C')^{\mathcal{J}}$. $\dashv$

We now show that $\mathcal{J} \models \tau_{\mathsf{PL}}(\mathcal{T})$. By definition, every CIA in $\tau_{\mathsf{PL}}(\mathcal{T})$ is of the form (i) $\tau_{\mathsf{PL}}(C') \sqsubseteq \tau_{\mathsf{PL}}(D')$, where $C' \sqsubseteq D' \in \mathcal{T}$; or of the form (ii) $A_R \sqsubseteq A_S$, where $R \sqsubseteq S \in \mathcal{T}$.

That $\mathcal{J}$ satisfies CIAs of the form (i) is a direct consequence of the previous claim; so we focus on CIAs of the form (ii). Let $A_R \sqsubseteq A_S \in \tau_{\mathsf{PL}}(\mathcal{T})$ and $s \in A_R^{\mathcal{J}}$. Then by (2), $\langle s, s \rangle \in R^{\mathcal{I}}$. Since $\mathcal{I} \models \mathcal{T}$ and $R \sqsubseteq S \in \mathcal{T}$, we then have $\langle s, s \rangle \in S^{\mathcal{I}}$. Then by (2) again, $s \in A_S^{\mathcal{J}}$. Hence $\mathcal{J}$ satisfies $A_R \sqsubseteq A_S$.

Now we proceed towards our goal $s \in D^{\mathcal{I}}$ as follows. By $s \in C^{\mathcal{I}}$ and the previous claim, we have $s \in \tau_{\mathsf{PL}}(C)^{\mathcal{J}}$. Then by $\mathcal{J} \models \tau_{\mathsf{PL}}(\mathcal{T})$ and $\tau_{\mathsf{PL}}(\mathcal{T}) \models \tau_{\mathsf{PL}}(C) \sqsubseteq \tau_{\mathsf{PL}}(D)$, we obtain $s \in \tau_{\mathsf{PL}}(D)^{\mathcal{J}}$. Using the previous claim again, we conclude that $s \in D^{\mathcal{I}}$. □

Proposition 3.35. *Let $\mathscr{L}$ be $\mathcal{ALC}$ or any of its extensions with constructors from $\{\mathcal{S}, \mathcal{H}, \mathcal{I}, \mathcal{F}\}$. For all $\mathscr{L}$-concepts C_1, C_2 and all $\mathscr{L}$-TBoxes $\mathcal{T}_1, \mathcal{T}_2$, if $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2$, then there exists a FO $\models_{=1}$-interpolant of C_1 and C_2 under $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$ that can be computed in time single exponential in $|\mathcal{T}_1| + |\mathcal{T}_2| + |C_1| + |C_2|$.*

Proof. Let $\mathscr{L}$ be one of the DLs mentioned in the theorem and let C_1, C_2 be $\mathscr{L}$-concepts and let $\mathcal{T}_1, \mathcal{T}_2$ be $\mathscr{L}$-TBoxes such that $\mathcal{T}_1 \cup \mathcal{T}_2 \models C_1 \sqsubseteq C_2$. Then it immediately follows that $\mathcal{T}_1 \cup \mathcal{T}_2 \models_{=1} C_1 \sqsubseteq C_2$. By Proposition 3.32, $\mathcal{T}_1 \cup \mathcal{T}_2 \models_{=1} C_1 \sqsubseteq C_2$ implies $\tau_{\mathsf{PL}}(\mathcal{T}_1) \cup \tau_{\mathsf{PL}}(\mathcal{T}_2) \models \tau_{\mathsf{PL}}(C_1) \sqsubseteq \tau_{\mathsf{PL}}(C_2)$. Now by Theorem 3.10, there is some interpolant I of $\tau_{\mathsf{PL}}(C_1)$ and $\tau_{\mathsf{PL}}(C_2)$ under $\langle \tau_{\mathsf{PL}}(\mathcal{T}_1), \tau_{\mathsf{PL}}(\mathcal{T}_2) \rangle$ that can be computed in time double exponential in $|\mathcal{T}_1| + |\mathcal{T}_2| + |C_1| + |C_2|$. However, in this case we are only dealing with propositional formulae and the tableau algorithm can easily be modified to construct a tree-shaped proof instead of a general graph-shaped one by eliminating the use of proxies. In fact, we have just described a standard tableau algorithm for propositional logic. It is well-known that each node in the tree has a polynomial out-degree in the size of the input and the height of the tree is polynominal in the size of the input. By inspecting the proof of Theorem 3.10, one can easily see that in this case I can be computed in time single exponential in $|\mathcal{T}_1| + |\mathcal{T}_2| + |C_1| + |C_2|$. Finally let D be the concept obtained from I by replacing each occurrence of a concept name A_R by $\exists R.\top$. We have that $\pi_x(D)$ is a FO $\models_{=1}$-interpolant of C_1 and C_2 under $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$. It is easy to see that the time required to compute $\pi_x(D)$ is as stated in the proposition. □

Step 4: Putting it all together The result that we were after now follows, by putting the FO $\models_{=1}$-interpolants and the FO $\models_{\geq 2}$-interpolants together:

Theorem 3.36. *Let $\mathscr{L}$ be $\mathcal{ALC}$ or any of its extensions with constructors from $\{\mathcal{S}, \mathcal{I}, \mathcal{F}\}$. For all $\mathscr{L}$-concepts C, D and $\mathscr{L}$-TBoxes $\mathcal{T}_1$, $\mathcal{T}_2$, if $\mathcal{T}_1 \cup \mathcal{T}_2 \models C \sqsubseteq D$, then there exists an FO interpolant $\varphi(x)$ of C and D under $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$ and $\varphi(x)$ can be computed in time single exponential in $|\mathcal{T}_1| + |\mathcal{T}_2| + |C| + |D|$.*

Proof. Let $\mathscr{L}$ be one of the DLs mentioned in the theorem, let C, D be $\mathscr{L}$-concepts, and let $\mathcal{T}_1$, $\mathcal{T}_2$ be $\mathscr{L}$-TBoxes such that $\mathcal{T}_1 \cup \mathcal{T}_2 \models C \sqsubseteq D$. By Proposition 3.35, there is some FO $\models_{=1}$-interpolant $\psi(x)$ of C and D under $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$ that can be computed in time single exponential in $|\mathcal{T}_1| + |\mathcal{T}_2| + |C| + |D|$; and by Proposition 3.30, there is some FO $\models_{\geq 2}$-interpolant $\varphi(x)$ of C and D under $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$ that can be computed in time single exponential in $|\mathcal{T}_1| + |\mathcal{T}_2| + |C| + |D|$. Let

$$\vartheta(x) = (\exists y \exists z (y \neq z) \rightarrow \varphi(x)) \wedge (\forall y \forall z (y = z) \rightarrow \psi(x)).$$

Claim 3.37. $\pi(\mathcal{T}_1) \cup \pi(\mathcal{T}_2) \models \forall x. \pi_x(C) \rightarrow \vartheta(x)$ *and* $\pi(\mathcal{T}_1) \cup \pi(\mathcal{T}_2) \models \forall x. \vartheta(x) \rightarrow \pi_x(D)$.

Proof of claim. We prove the first part. The proof of the second part is analogous.

Let $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$ be a model of $\mathcal{T}_1 \cup \mathcal{T}_2$, i.e., of $\pi(\mathcal{T}_1) \cup \pi(\mathcal{T}_2)$, and α be a first-order variable assignment with $\mathcal{I}, \alpha \models \pi_x(C)$. We need to show that $\mathcal{I}, \alpha \models \vartheta(x)$. To this aim, we show $\mathcal{I}, \alpha \models (\exists y \exists z (y \neq z) \rightarrow \varphi(x))$ and $\mathcal{I}, \alpha \models (\forall y \forall z (y = z) \rightarrow \psi(x))$.

First suppose that $\mathcal{I}, \alpha \models \exists y \exists z (y \neq z)$. We are done if we prove that $\mathcal{I}, \alpha \models \varphi(x)$. $\mathcal{I}, \alpha \models \exists y \exists z (y \neq z)$ implies $|\Delta^{\mathcal{I}}| \geq 2$. Then by $\mathcal{I}, \alpha \models \pi_x(C)$ and $\pi(\mathcal{T}_1) \cup \pi(\mathcal{T}_2) \models_{\geq 2} \forall x.\pi_x(C) \rightarrow \varphi(x)$, we obtain $\mathcal{I}, \alpha \models \varphi(x)$, and we are done.

Now suppose that $\mathcal{I}, \alpha \models \forall y \forall z (y = z)$. We are done if we prove that $\mathcal{I}, \alpha \models \psi(x)$. $\mathcal{I}, \alpha \models \forall y \forall z (y = z)$ implies $|\Delta^{\mathcal{I}}| = 1$. Then by $\mathcal{I}, \alpha \models \pi_x(C)$ and $\pi(\mathcal{T}_1) \cup \pi(\mathcal{T}_2) \models_{=1} \forall x.\pi_x(C) \rightarrow \psi(x)$, we obtain $\mathcal{I}, \alpha \models \psi(x)$, and we are done.

Thus, both of the conjuncts of $\vartheta(x)$ are satisfied by $\mathcal{I}, \alpha$. But then $\mathcal{I}, \alpha \models \vartheta(x)$. $\dashv$

By assumption we have that $\mathsf{sig}(\varphi(x)), \mathsf{sig}(\psi(x)) \subseteq \mathsf{sig}(C, \mathcal{T}_1) \cap \mathsf{sig}(D, \mathcal{T}_2)$. Since the formulas $\exists y \exists z (y \neq z)$ and $\forall y \forall z (y = z)$ do not introduce new predicates, we have that $\mathsf{sig}(\vartheta(x)) \subseteq \mathsf{sig}(C, \mathcal{T}_1) \cap \mathsf{sig}(D, \mathcal{T}_2)$. Therefore, $\vartheta(x)$ is a FO interpolant for C and D under $\langle \mathcal{T}_1, \mathcal{T}_2 \rangle$. Moreover, since both of its conjuncts can be computed in single exponential time, so can $\vartheta(x)$. Hence the theorem follows. $\square$

4. Results on Beth Definability

In this section, we present the main technical contributions of the paper. We first introduce the notions of *implicit* and *explicit definability* for concepts and define the (projective) *Beth definability property*, which are in fact the primary notions of interest in this paper. In what follows, $\mathscr{L}$ denotes any of the description logics $\mathcal{ALC}X$ with $X \subseteq \{\mathcal{S}, \mathcal{H}, \mathcal{I}, \mathcal{F}\}$.

Definition 4.1 (Implicit definability)**.** *Let C be an $\mathscr{L}$-concept, $\mathcal{T}$ an $\mathscr{L}$-TBox, and $\Sigma \subseteq \mathsf{sig}(C, \mathcal{T})$. C is* implicitly definable *from Σ under $\mathcal{T}$ if, for every two models $\mathcal{I}$ and $\mathcal{J}$ of $\mathcal{T}$ satisfying $\Delta^{\mathcal{I}} = \Delta^{\mathcal{J}}$ and, for all $P \in \Sigma$, $P^{\mathcal{I}} = P^{\mathcal{J}}$, it holds that $C^{\mathcal{I}} = C^{\mathcal{J}}$.*

In other words, given a TBox, a concept C is implicitly definable if the set of all its instances depends only on the extension of the predicates in Σ and the domain of discourse. *Deciding implicit definability in $\mathscr{L}$* means, given an $\mathscr{L}$-concept C, $\mathscr{L}$-TBox $\mathcal{T}$, and a set of predicates $\Sigma \subseteq \mathsf{sig}(C, \mathcal{T})$, to check whether C is implicitly definable from Σ under $\mathcal{T}$. For every predicate $P \in \mathsf{sig}(C, \mathcal{T}) \setminus \Sigma$, introduce a new predicate P' which is not in $\mathsf{sig}(C, \mathcal{T})$. Now let $\widetilde{C}$ (respectively, $\widetilde{\mathcal{T}}$) be the concept (respectively, TBox) obtained by replacing every occurrence of a predicate $P \notin \Sigma$ in C (respectively, in $\mathcal{T}$) by P'. Lemma 4.2, whose proof is a routine adaptation of an analogous result for first-order logic (Boolos, Burgess, & Jeffrey, 2007), provides a characterization of implicit definability in terms of entailment. This well-known characterization is often used as a definition of implicit definability (Hoogland & Marx, 2002; Conradie, 2002).

Lemma 4.2. *Let C be an $\mathscr{L}$-concept, $\mathcal{T}$ be an $\mathscr{L}$-TBox, and $\Sigma \subseteq \mathsf{sig}(C, \mathcal{T})$. Then C is implicitly definable from Σ under $\mathcal{T}$ if and only if $\mathcal{T} \cup \widetilde{\mathcal{T}} \models C \equiv \widetilde{C}$.*

In particular, Lemma 4.2 reduces implicit definability in $\mathscr{L}$ to the concept subsumption problem in $\mathscr{L}$ w.r.t. TBoxes. It is also possible to reduce the concept subsumption problem in $\mathscr{L}$ w.r.t. TBoxes to the problem of deciding implicit definability in $\mathscr{L}$.

Lemma 4.3. *Let $C \sqsubseteq D$ be an $\mathscr{L}$-CIA, $\mathcal{T}$ be an $\mathscr{L}$-TBox, $\Sigma = \mathsf{sig}(C \sqcap D, \mathcal{T})$, and $A_0 \in N_C \setminus \Sigma$. Then $\mathcal{T} \models C \sqsubseteq D$ if and only if A_0 is implicitly definable from Σ under $\mathcal{T} \cup \{A_0 \sqsubseteq C \sqcap \neg D\}$.*

Proof. ($\Rightarrow$) Suppose $\mathcal{T} \models C \sqsubseteq D$. Let $\mathcal{I}$ and $\mathcal{J}$ be models of $\mathcal{T} \cup \{A_0 \sqsubseteq C \sqcap \neg D\}$ such that $\Delta^{\mathcal{I}} = \Delta^{\mathcal{J}}$ and for all $P \in \Sigma$, we have $P^{\mathcal{I}} = P^{\mathcal{J}}$. Obviously, $\mathcal{I}$ and $\mathcal{J}$ are also models of $\mathcal{T}$. Then by $\mathcal{T} \models C \sqsubseteq D$, we have that $(C \sqcap \neg D)^{\mathcal{I}} = (C \sqcap \neg D)^{\mathcal{J}} = \emptyset$. But then $A_0^{\mathcal{I}} = A_0^{\mathcal{J}} = \emptyset$. Hence, A_0 is implicitly definable from Σ under $\mathcal{T} \cup \{A_0 \sqsubseteq C \sqcap \neg D\}$.

($\Leftarrow$) We show the contrapositive, i.e., if $\mathcal{T} \not\models C \sqsubseteq D$, then A_0 is not implicitly definable from Σ under $\mathcal{T} \cup \{A_0 \sqsubseteq C \sqcap \neg D\}$. Suppose $\mathcal{T} \not\models C \sqsubseteq D$. Then there is some model $\mathcal{I}$ of $\mathcal{T}$ and some $s \in \Delta^{\mathcal{I}}$ such that $s \in (C \sqcap \neg D)^{\mathcal{I}}$. Let $\mathcal{I}_1 = \langle \Delta^{\mathcal{I}_1}, \cdot^{\mathcal{I}_1} \rangle$ and $\mathcal{I}_2 = \langle \Delta^{\mathcal{I}_2}, \cdot^{\mathcal{I}_2} \rangle$ be such that

- $\Delta^{\mathcal{I}_1} = \Delta^{\mathcal{I}_2} = \Delta^{\mathcal{I}}$;
- $A^{\mathcal{I}_1} = A^{\mathcal{I}_2} = A^{\mathcal{I}}$, for all $A \in (N_C \setminus A_0)$;
- $R^{\mathcal{I}_1} = R^{\mathcal{I}_2} = R^{\mathcal{I}}$, for all $R \in N_R$;
- $A_0^{\mathcal{I}_1} = \{s\}$ and $A_0^{\mathcal{I}_2} = \emptyset$.

It is easy to see that $\mathcal{I}_1$ and $\mathcal{I}_2$ are models of $\mathcal{T} \cup \{A_0 \sqsubseteq C \sqcap \neg D\}$. Also observe that $\mathcal{I}_1$ and $\mathcal{I}_2$ are two interpretations with the same domain and they agree on what they assign to predicates in Σ. But $A_0^{\mathcal{I}_1} \neq A_0^{\mathcal{I}_2}$. Hence A_0 is not implicitly definable from Σ under $\mathcal{T} \cup \{A_0 \sqsubseteq C \sqcap \neg D\}$. ☐

Using Lemma 4.2 (for the upper bound) and Lemma 4.3 (for the lower bound), the following theorem follows immediately, since the concept subsumption problem w.r.t. TBoxes is ExpTime-complete for the description logics in question (Tobies, 2001).

Theorem 4.4. *In $\mathcal{ALC}$ and any of its extensions with constructors from $\{\mathcal{S}, \mathcal{H}, \mathcal{I}, \mathcal{F}\}$, implicit definability is* ExpTime*-complete.*

Explicit definability is the syntactic counterpart of implicit definability. Given a concept C, signature Σ, and TBox $\mathcal{T}$, it asks for the existence of a concept D formulated over Σ such that the C and D denote the same set in every model of $\mathcal{T}$.

Definition 4.5 (Explicit definability)**.** *Let C be an $\mathscr{L}$-concept, $\mathcal{T}$ a $\mathscr{L}$-TBox, and $\Sigma \subseteq \mathsf{sig}(C, \mathcal{T})$. We say that C is* explicitly definable *from Σ under $\mathcal{T}$ if there is some $\mathscr{L}$-concept D such that $\mathcal{T} \models C \equiv D$ and $\mathsf{sig}(D) \subseteq \Sigma$. Such a concept D is called an* explicit definition *of C from Σ under $\mathcal{T}$.*

Proposition 4.6. *Let C be an $\mathscr{L}$-concept, $\mathcal{T}$ an $\mathscr{L}$-TBox, and $\Sigma \subseteq \mathsf{sig}(C, \mathcal{T})$. If C is explicitly definable from Σ under $\mathcal{T}$, then C is implicitly definable from Σ under $\mathcal{T}$.*

Proof. Suppose C is explicitly definable from Σ under $\mathcal{T}$. Then there is some concept D such that $\mathcal{T} \models C \equiv D$. This implies by the definition of $\widetilde{\mathcal{T}}$ and $\widetilde{C}$, and $\mathsf{sig}(D) \subseteq \Sigma$ that $\widetilde{\mathcal{T}} \models \widetilde{C} \equiv D$. Then we have $\mathcal{T} \cup \widetilde{\mathcal{T}} \models C \equiv D$ and $\mathcal{T} \cup \widetilde{\mathcal{T}} \models \widetilde{C} \equiv D$ by the monotonicity of $\models$. These yield $\mathcal{T} \cup \widetilde{\mathcal{T}} \models C \equiv \widetilde{C}$. Then by Lemma 4.2, C is implicitly definable from Σ under $\mathcal{T}$. ☐

Definition 4.7 (Beth definability property)**.** *$\mathscr{L}$ has the* Beth definability property *(BP) if for all $\mathscr{L}$-concepts C, all $\mathscr{L}$-TBoxes $\mathcal{T}$, and all signatures $\Sigma \subseteq \mathsf{sig}(C, \mathcal{T})$, if C is implicitly definable from Σ under $\mathcal{T}$, then C is explicitly definable from Σ under $\mathcal{T}$.*

Observe that, in the above definition, Σ restricts the concept names and the role names that are allowed to appear in the explicit definition. We can obtain a weaker version of the Beth definability property by restricting only the concept names occurring in the explicit definition. This is called the *concept-name Beth definability property* (CBP). In other words, the CBP refers to the existence of explicit definitions over signatures of the form $\Sigma \cup N_R$.

As we will explain later, we also have reasons to be interested in whether description logics satisfy the Beth definability property over the restricted class of finite interpretations. It is known that the Beth definability property, when restricted to finite structures, fails for first-order logic (see e.g., Hoogland, 2001), in spite of the fact that it holds in the unrestricted case. We will specifically investigate Beth definability for description logics restricted to finite interpretations. We call this the *Beth definability property in the finite* (BPF). Formally, BPF is defined in the same way as BP, except that we replace, in the definition, all occurrences of the word 'interpretation' or 'model' by 'finite interpretation' or 'finite model', and we replace the symbol $\models$ by $\models_f$, where $\models_f$ considers only finite interpretations. In addition, we will speak about f-implicit definability and f-explicit definability. It follows from Lemma 4.2 that, if $\mathscr{L}$ has FMP, then BP are BPF are equivalent for $\mathscr{L}$. Hence it only makes sense to specifically study BPF for logics without FMP.

4.1 Bounds on the Size of Explicit Definitions

We start by a positive result on BP which is a direct application of the interpolation theorem, i.e., Theorem 3.22.

Theorem 4.8 (BP)**.** *Let $\mathscr{L}$ be $\mathcal{ALC}$ or any of its extensions with constructors from $\{\mathcal{S}, \mathcal{I}, \mathcal{F}\}$. Then for all $\mathscr{L}$-concepts C, all $\mathscr{L}$-TBoxes $\mathcal{T}$, and all signatures $\Sigma \subseteq \mathsf{sig}(C, \mathcal{T})$, if C is implicitly definable from Σ under $\mathcal{T}$, then C is explicitly definable from Σ under $\mathcal{T}$, and the explicit definition of C can be computed in time double exponential in $|\mathcal{T}| + |C|$.*

Proof. Let $\mathscr{L}$ be one of the DLs stated in the theorem, C be an $\mathscr{L}$-concept, $\mathcal{T}$ be an $\mathscr{L}$-TBox, and $\Sigma \subseteq \mathsf{sig}(C, \mathcal{T})$ such that C is implicitly definable from Σ under $\mathcal{T}$. By Lemma 4.2, we have that $\mathcal{T} \cup \widetilde{\mathcal{T}} \models C \equiv \widetilde{C}$ (where $\widetilde{\mathcal{T}}$ and $\widetilde{C}$ are obtained from $\mathcal{T}$ and C, respectively, by replacing all occurrences of predicates $P \notin \Sigma$ by fresh predicates P' that are not in $\mathsf{sig}(C, \mathcal{T})$). Now by Theorem 3.22, there is an interpolant I of C and $\widetilde{C}$ under $\langle \mathcal{T}, \widetilde{\mathcal{T}} \rangle$ that can be computed in time double exponential in $|\mathcal{T}| + |\widetilde{\mathcal{T}}| + |C| + |\widetilde{C}|$. Since it is an interpolant, $\mathsf{sig}(I) \subseteq \mathsf{sig}(C, \mathcal{T}) \cap \mathsf{sig}(\widetilde{C}, \widetilde{\mathcal{T}}) = \Sigma$, and both (a) $\mathcal{T} \cup \widetilde{\mathcal{T}} \models C \sqsubseteq I$ and (b) $\mathcal{T} \cup \widetilde{\mathcal{T}} \models I \sqsubseteq \widetilde{C}$. By (b) and $\mathcal{T} \cup \widetilde{\mathcal{T}} \models \widetilde{C} \sqsubseteq C$, we have $\mathcal{T} \cup \widetilde{\mathcal{T}} \models I \sqsubseteq C$, from which $\mathcal{T} \cup \widetilde{\mathcal{T}} \models C \equiv I$ follows by (a). From the structure of $\widetilde{\mathcal{T}}$, it now straightforwardly follows that $\mathcal{T} \models C \equiv I$.

As for the time needed to compute I, observe that $|\mathcal{T}| + |\widetilde{\mathcal{T}}| + |C| + |\widetilde{C}| = 2 \cdot (|\mathcal{T}| + |C|)$. Hence I can be computed in time double exponential in $|\mathcal{T}| + |C|$. □

The proof of Theorem 4.8 uses Theorem 3.22. Similarly, if we use Theorem 3.36 instead, we can show that first-order explicit definitions of implicitly defined concepts can be computed in single exponential time. Note that Theorem 4.8 also establishes a double exponential upper bound on the size explicit definitions in the considered logics. This upper bound is optimal because we show in Theorem 4.11 below that explicit definitions in $\mathscr{L}$ may need to be double exponentially big.

An essential tool in the proof of Theorem 4.11, will be the *path-set construction* that was previously used by Lutz (2006) to characterize the succinctness of public announcement logic compared to epistemic logic. The path-set construction has also been used by Ghilardi et al. (2006) to establish a lower bound on the size of concepts 'witnessing' that a TBox is not a conservative extension of another TBox.

Definition 4.9. *If C is an $\mathcal{ALC}$-concept, then the* path-set P_C *of C is defined by structural induction as follows, where ε denotes the empty sequence and $\cdot$ denotes concatenation of finite sequences:*

- $P_\top = P_A = \{\varepsilon\}$, *for* $A \in N_C$*;*
- $P_{\neg C} = P_C$*;*
- $P_{C \sqcap D} = P_C \cup P_D$*;*
- $P_{\exists R.C} = \{\varepsilon\} \cup \{R \cdot p \mid p \in P_C\}$.

Intuitively, P_C describes the nestings of role constructors in C. We will use P_C as a tool for establishing lower bounds on the size of concepts.

Lemma 4.10. *For every $\mathcal{ALC}$-concept C, we have $|C| \geq |P_C|$.*

Proof. The proof is by induction on the structure of C.

If C is an atomic concept of the form $\top$ or A (with $A \in N_C$), then, by definition, $|C| = 1$ and $|P_C| = 1$ since $P_C = \{\varepsilon\}$. Hence $|C| \geq |P_C|$.

Next, let $C = \neg D$. By the inductive hypothesis, we have $|D| \geq |P_D|$. Then by $|P_{\neg D}| = |P_D|$, we obtain $|D| \geq |P_{\neg D}|$. Finally, by the fact that $|\neg D| = |D| + 1$, we obtain $|\neg D| \geq |P_{\neg D}|$. Hence $|C| \geq |P_C|$.

Next, let $C = C_1 \sqcap C_2$. By the inductive hypothesis, we have $|C_1| \geq |P_{C_1}|$ and $|C_2| \geq |P_{C_2}|$. This implies $|C_1|+|C_2| \geq |P_{C_1}|+|P_{C_2}|$. Then by the fact that $|C_1 \sqcap C_2| = |C_1|+|C_2|+1$, we obtain $|C_1 \sqcap C_2| \geq |P_{C_1}|+|P_{C_2}|$. Finally, by $|P_{C_1}|+|P_{C_2}| \geq |P_{C_1 \sqcap C_2}|$, we have $|C_1 \sqcap C_2| \geq |P_{C_1 \sqcap C_2}|$. Hence $|C| \geq |P_C|$.

Finally, let $C = \exists R.D$. By the inductive hypothesis, we have $|D| \geq |P_D|$. This implies $|D| + 2 \geq |P_D| + 2$. Then by the fact that $|\exists R.D| = |D| + 2$, we obtain $|\exists R.D| \geq |P_D| + 2$. Finally, by $|P_D| + 1 = |P_{\exists R.D}|$, we have $|\exists R.D| \geq |P_{\exists R.D}|$. Hence $|C| \geq |P_C|$. ◻

Theorem 4.11 (Explicit definition lower bound)**.** *Let $\Sigma = \{R, S\} \subseteq N_R$. Then for every $n \in \mathbb{N}$, there is an $\mathcal{ALC}$-concept C_n and an $\mathcal{ALC}$-TBox $\mathcal{T}_n$ such that $\Sigma \subseteq \mathsf{sig}(C_n, \mathcal{T}_n)$, $|\mathcal{T}_n|$ and $|C_n|$ are polynomial in n, C_n is implicitly definable from Σ under $\mathcal{T}_n$, and the smallest explicit definition of C_n from Σ under $\mathcal{T}_n$ is double exponentially long in n.*

Proof. Fix an $n \in \mathbb{N}$. Let $A_1, \ldots, A_n$ be pairwise distinct concept names. We use these concept names and their negations to represent in binary format a number in $\{0, \ldots, 2^n - 1\}$. More precisely, $\neg A_n \sqcap \ldots \sqcap \neg A_1$ *represents* 0, $\neg A_n \sqcap \neg A_{n-1} \ldots \sqcap A_1$ *represents* 1, and so on. Obviously, this implies that the least significant bit is at position 1. For every $i \in \{0, \ldots, 2^n - 1\}$, we denote the concept that represents i by C_i. Note that in each C_i, either A_j or $\neg A_j$ is a *conjunct of* C_i, for all $j \in \{1, \ldots, n\}$.

For $k \in \{1, \ldots, n\}$,

- let $X_k = \neg A_1 \sqcap \ldots \sqcap \neg A_{k-1} \sqcap A_k$ and
- let $Y_k = A_1 \sqcap \ldots \sqcap A_{k-1} \sqcap \neg A_k$.

Note that X_k and Y_k are not concept names and we will use them only to abbreviate our CIAs. We define $\mathcal{T}_n$ as the $\mathcal{ALC}$-TBox consisting of the following CIAs.

- $\neg A_n \sqcap \ldots \sqcap \neg A_1 \sqsubseteq \forall R.\bot \sqcap \forall S.\bot$
- $A_1 \sqcup \ldots \sqcup A_n \sqsubseteq \exists R.\top \sqcup \exists S.\top$
- For every $k \in \{1, \ldots, n\}$ and $\sigma \in \Sigma$,

$$\begin{array}{rcl} X_k & \sqsubseteq & \forall \sigma.Y_k \sqcap \\ & & \displaystyle\bigsqcap_{k < l \leq n} ((A_l \sqcap \forall \sigma.A_l) \sqcup (\neg A_l \sqcap \forall \sigma.\neg A_l)) \end{array}$$

Intuitively, the last item above allows us to decrease the counter value by one by flipping the respective bits. Note that $|\mathcal{T}_n|$ is polynomial in n and that $\mathcal{T}_n$ is satisfiable. In fact, we present models of $\mathcal{T}_n$ in Claim 4.14. If $\mathcal{I}$ is a model of $\mathcal{T}_n$, we have for every $s \in \Delta^{\mathcal{I}}$ and every $i \in \{1, \ldots, n\}$, either $s \in A_i^{\mathcal{I}}$ or $s \in (\neg A_i)^{\mathcal{I}}$ by the virtue of $\mathcal{I}$ being an interpretation. Therefore, for every $s \in \Delta^{\mathcal{I}}$ there is exactly one $i \in \{0, \ldots, 2^n - 1\}$ such that $s \in C_i^{\mathcal{I}}$.

Claim 4.12. *Let* $i \in \{1, \ldots, 2^n - 1\}$. *Then*

1. $\mathcal{T}_n \models C_i \sqsubseteq \forall R.C_{i-1} \sqcap \forall S.C_{i-1}$
2. $\mathcal{T}_n \models C_i \equiv \exists R.C_{i-1} \sqcup \exists S.C_{i-1}$

Proof of claim. For 1, suppose $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$ is a model of $\mathcal{T}_n$, $s \in \Delta^{\mathcal{I}}$ with $s \in C_i^{\mathcal{I}}$, and $\sigma \in \{R, S\} = \Sigma$. It suffices to show that $s \in (\forall \sigma.C_{i-1})^{\mathcal{I}}$. If there is no $t \in \Delta^{\mathcal{I}}$ such that $\langle s, t \rangle \in \sigma^{\mathcal{I}}$ then we are done immediately; therefore, suppose $\langle s, t \rangle \in \sigma^{\mathcal{I}}$. We need to show $t \in C_{i-1}^{\mathcal{I}}$.

We have that $C_i = B_n \sqcap \ldots \sqcap B_1$, where $B_j = A_j$ or $B_j = \neg A_j$, for each $j \in \{1, \ldots, n\}$. Denote by $\overline{B}_j$ the concept $\neg A_j$ if $B_j = A_j$, or else the concept A_j if $B_j = \neg A_j$. Since $s \in C_i^{\mathcal{I}}$, there is exactly one $k \in \{1, \ldots, n\}$ such that $s \in X_k^{\mathcal{I}}$. Then by the CIA

$$\begin{array}{rcl} X_k & \sqsubseteq & \forall \sigma.Y_k \sqcap \\ & & \displaystyle\bigsqcap_{k < l \leq n} ((A_l \sqcap \forall \sigma.A_l) \sqcup (\neg A_l \sqcap \forall \sigma.\neg A_l)) \end{array}$$

in $\mathcal{T}_n$, we have that $t \in (B_n \sqcap \ldots \sqcap B_{k+1} \sqcap \overline{B}_k \sqcap \ldots \sqcap \overline{B}_1)^{\mathcal{I}}$. It is not hard to see that

$$C_{i-1} = B_n \sqcap \ldots \sqcap B_{k+1} \sqcap \overline{B}_k \sqcap \ldots \sqcap \overline{B}_1.$$

Hence we conclude that $t \in C_{i-1}^{\mathcal{I}}$, which is what we wanted to show.

For 2. ($\Rightarrow$) Suppose $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$ is a model of $\mathcal{T}_n$ and $s \in \Delta^{\mathcal{I}}$ with $s \in C_i^{\mathcal{I}}$. Since $i \neq 0$, by $A_1 \sqcup \ldots \sqcup A_n \sqsubseteq \exists R.\top \sqcup \exists S.\top \in \mathcal{T}_n$, $C_i \in (\exists R.\top)^{\mathcal{I}}$ or $C_i \in (\exists S.\top)^{\mathcal{I}}$. That is, there is some $t \in \Delta^{\mathcal{I}}$ such that either $\langle s, t \rangle \in R^{\mathcal{I}}$ or $\langle s, t \rangle \in S^{\mathcal{I}}$. In both cases, $t \in C_{i-1}^{\mathcal{I}}$ by 1. Hence $s \in (\exists R.C_{i-1} \sqcup \exists S.C_{i-1})^{\mathcal{I}}$.

($\Leftarrow$) Suppose $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$ is a model of $\mathcal{T}_n$ and $s \in \Delta^{\mathcal{I}}$ with $s \in (\exists R.C_{i-1} \sqcup \exists S.C_{i-1})^{\mathcal{I}}$. This means there is some $t \in \Delta^{\mathcal{I}}$ such that $t \in C_{i-1}^{\mathcal{I}}$ and either $\langle s, t \rangle \in R^{\mathcal{I}}$ or $\langle s, t \rangle \in S^{\mathcal{I}}$. We proceed towards a contradiction so further suppose that $s \notin C_i^{\mathcal{I}}$, i.e., $s \in (\neg C_i)^{\mathcal{I}}$. Then by the definition of an interpretation, $s \in C_j^{\mathcal{I}}$, where $j \neq i$ and $j \in \{0, \ldots, 2^n - 1\}$. If $j = 0$, then by $\neg A_n \sqcap \ldots \sqcap \neg A_1 \sqsubseteq \forall R.\bot \sqcap \forall S.\bot \in \mathcal{T}_n$, we immediately get a contradiction. If $j \neq 0$, then by 1, $t \in C_{j-1}^{\mathcal{I}}$. Since $i \neq j$, we have $(i-1) \neq (j-1)$. Thus, the binary representation of $i-1$ and $j-1$ must differ in at least one bit. This implies by $t \in C_{j-1}^{\mathcal{I}}$ and $t \in C_{i-1}^{\mathcal{I}}$ that there is some $k \in \{1, \ldots, n\}$ such that $t \in A_k^{\mathcal{I}}$ and $t \notin A_k^{\mathcal{I}}$. Hence a contradiction. $\dashv$

Now define concepts $D_0 \ldots D_{2^n-1}$ inductively as follows.

$$\begin{array}{rcl} D_0 & = & \forall R.\bot \sqcap \forall S.\bot \\ D_i & = & \exists R.D_{i-1} \sqcup \exists S.D_{i-1} \end{array}$$

Intuitively, D_i has the shape of a binary tree (due to role names R, S) and the height of the tree is $O(i)$. This implies $|C_{2^n-1}|$ is double exponential in n.

Claim 4.13. *For every $i \in \{0, \ldots, 2^n - 1\}$, we have $\mathcal{T}_n \models C_i \equiv D_i$.*

Proof of claim. The proof is by induction on i. The base case is when $i = 0$. Then by the axioms in $\mathcal{T}_n$, it trivially follows that $\mathcal{T}_n \models \neg A_1 \sqcap \ldots \neg A_n \equiv \forall R.\bot \sqcap \forall S.\bot$. In other words, $\mathcal{T}_n \models C_0 \equiv D_0$. Hence the claim holds in the base case.

For the inductive step, suppose $i > 0$. By the previous claim, $\mathcal{T}_n \models C_i \equiv \exists R.C_{i-1} \sqcup \exists S.C_{i-1}$; and by the inductive hypothesis, $\mathcal{T}_n \models C_{i-1} \equiv D_{i-1}$. But then $\mathcal{T}_n \models C_i \equiv \exists R.D_{i-1} \sqcup \exists S.D_{i-1}$ which is what we wanted to show. $\dashv$

By the previous claim, we have that for all $i \in \{0, \ldots, 2^n - 1\}$, D_i is an explicit definition of C_i from $\Sigma = \{R, S\}$ under $\mathcal{T}_n$. Then by Proposition 4.6, C_i is implicitly definable from Σ under $\mathcal{T}_n$. In the rest of the proof, we show that each explicit definition of C_i from Σ under $\mathcal{T}_n$ is at least double exponentially long. To this aim, we introduce interpretations that are based on some elements of Σ^*, where Σ^* denotes the set of all strings over the symbols in Σ. More precisely, for every $p \in \Sigma^*$ with $0 \leq |p| \leq 2^n - 1$, we define the interpretation $\mathcal{I}_p$ as follows.

- $\Delta^{\mathcal{I}_p} = \{p' \in \Sigma^* \mid p' \text{ is a prefix of } p\}$;
- for all $A \in N_C$,
  - if $A = A_j$ for some $j \in \{1, \ldots, n\}$, then

$$A^{\mathcal{I}_p} = \{p' \in \Delta^{\mathcal{I}_p} \mid A \text{ is a conjunct of } C_{|p|-|p'|}\},$$

  - if $A \neq A_j$ for all $j \in \{1, \ldots, n\}$, then $A^{\mathcal{I}_p} = \emptyset$;

- for all $T \in N_R$,
  - $T^{\mathcal{I}_p} = \{\langle p_1, p_2 \rangle \in \Delta^{\mathcal{I}_p} \times \Delta^{\mathcal{I}_p} \mid p_2 = p_1 \cdot T\}$, if $T \in \Sigma$,
  - $T^{\mathcal{I}_p} = \emptyset$, if $T \in N_R \setminus \Sigma$.

The following claim is easy to show.

Claim 4.14. *For every $p \in \Sigma^*$ with $0 \leq |p| \leq 2^n - 1$, we have*

- *$\mathcal{I}_p \models \mathcal{T}_n$, and*
- *$\varepsilon \in C_{|p|}^{\mathcal{I}_p}$.*

Denote for every $i \in \{0, \ldots, 2^n - 1\}$, the set of all $p \in \Sigma^*$ such that $|p| = i$ by Σ^i.

Claim 4.15. *Let $i \in \{0, \ldots, 2^n - 1\}$ and let C be an $\mathcal{ALC}$-concept such that $\mathsf{sig}(C) \subseteq \Sigma = \{R, S\}$ and $\mathcal{T}_n \models C_i \equiv C$. Then $\Sigma^i \subseteq P_C$.*

Proof of claim. Suppose first $i = 0$. Then $\Sigma^i = \{\varepsilon\}$. Moreover, by definition we have $\varepsilon \in P_C$. Hence $\Sigma^i \subseteq P_C$, which is what we wanted to show.

Now suppose $i > 0$. We proceed towards a contradiction. Suppose that there is some $p_a \in \Sigma^i \setminus P_C$. Let p_b be the prefix of p_a with $|p_b| = i - 1$. Since $i > 0$, p_b is well-defined. We claim that for all $s \in \Delta^{\mathcal{I}_{p_b}} \subseteq \Delta^{\mathcal{I}_{p_a}}$ and $D \in \mathsf{sub}(C)$ such that $\{s \cdot p \mid p \in P_D\} \subseteq P_C$,

$$s \in D^{\mathcal{I}_{p_b}} \text{ if and only if } s \in D^{\mathcal{I}_{p_a}}. \tag{3}$$

The proof is by induction on the structure of D. Since the base and the boolean cases are trivial, we only treat the case $D = \exists \sigma.E$, where $\sigma \in \Sigma$.

- ($\Rightarrow$). Suppose $s \in D^{\mathcal{I}_{p_b}}$. Then there is a $t \in \Delta^{\mathcal{I}_{p_b}}$ such that $\langle s, t \rangle \in \sigma^{\mathcal{I}_{p_b}}$ and $t \in E^{\mathcal{I}_{p_b}}$. Since t is in $\Delta^{\mathcal{I}_{p_a}}$ as well, the former yields $\langle s, t \rangle \in \sigma^{\mathcal{I}_{p_a}}$. It thus remains to show that $t \in E^{\mathcal{I}_{p_a}}$. By definition, $t = s \cdot \sigma$ and $P_D = \{\varepsilon\} \cup \{\sigma \cdot p \mid p \in P_E\}$. Thus, our assumption $\{s \cdot p \mid p \in P_D\} \subseteq P_C$ yields $\{s\} \cup \{t \cdot p \mid p \in P_E\} \subseteq P_C$. This implies $\{t \cdot p \mid p \in P_E\} \subseteq P_C$. Then by the induction hypothesis and $t \in E^{\mathcal{I}_{p_b}}$, we obtain $t \in E^{\mathcal{I}_{p_a}}$, which is what we wanted to show for this direction of the proof.

- ($\Leftarrow$). Suppose $s \in D^{\mathcal{I}_{p_a}}$. Then there is a $t \in \Delta^{\mathcal{I}_{p_a}}$ such that $\langle s, t \rangle \in \sigma^{\mathcal{I}_{p_a}}$ and $t \in E^{\mathcal{I}_{p_a}}$. By definition, $t = s \cdot \sigma$ and $P_D = \{\varepsilon\} \cup \{\sigma \cdot p \mid p \in P_E\}$. Thus, our assumption $\{s \cdot p \mid p \in P_D\} \subseteq P_C$ yields $\{s\} \cup \{t \cdot p \mid p \in P_E\} \subseteq P_C$. This implies

$$\{t \cdot p \mid p \in P_E\} \subseteq P_C. \tag{4}$$

  By (4) and $\varepsilon \in P_E$, we obtain $t \in P_C$. Since $p_a \notin P_C$, this means $t \neq p_a$. But then $t \in \Delta^{\mathcal{I}_{p_b}}$ and $\langle s, t \rangle \in \sigma^{\mathcal{I}_{p_b}}$. By (4), $t \in E^{\mathcal{I}_{p_a}}$, and the induction hypothesis, we have $t \in E^{\mathcal{I}_{p_b}}$. Hence $s \in (\exists \sigma.E)^{\mathcal{I}_{p_b}}$.

Thus, we have shown that (3) holds. Now we arrive at a contradiction as follows. By Claim 4.14, we have $\varepsilon \in C_i^{\mathcal{I}_{p_a}}$ and $\varepsilon \in C_{i-1}^{\mathcal{I}_{p_b}}$, since $|p_a| = i$ and $|p_b| = i - 1$, respectively. $\varepsilon \in C_{i-1}^{\mathcal{I}_{p_b}}$ implies by the definition of C_i that $\varepsilon \notin C_i^{\mathcal{I}_{p_b}}$. Then by $\mathcal{T}_n \models C_i \equiv C$ and Claim 4.14, we obtain $\varepsilon \in C^{\mathcal{I}_{p_a}}$ and $\varepsilon \notin C^{\mathcal{I}_{p_b}}$. But this contradicts with an immediate consequence of (3), namely $\varepsilon \in C^{\mathcal{I}_{p_b}}$ iff $\varepsilon \in C^{\mathcal{I}_{p_a}}$. Hence a contradiction. Thus, we conclude that $\Sigma^i \subseteq P_C$ for $i > 0$. $\dashv$

To show the theorem, we argue as follows. Suppose that C is an $\mathcal{ALC}$-concept such that $\mathcal{T}_n \models C_{2^n-1} \equiv C$ and $\mathsf{sig}(C) \subseteq \Sigma$. Then by the previous claim $\Sigma^{2^n-1} \subseteq P_C$. By its definition, $|\Sigma^{2^n-1}| = 2^{2^n-1}$ and thus, $2^{2^n-1} \leq |P_C|$. Then by Lemma 4.10, $2^{2^n-1} \leq |C|$. Hence the theorem follows. $\square$

Remark 4.16. *With the* role disjunction *constructor, which is not present in $\mathcal{ALC}$, C_n would admit a single exponentially long explicit definition from Σ under $\mathcal{T}_n$ in Theorem 4.11.*

Remark 4.17. *The lower bound argument in Theorem 4.11 works for CBP as well by just setting $\Sigma = \emptyset$.*

Combined with Theorem 4.8, Theorem 4.11 implies that implicit definitions using general TBoxes are exactly double exponentially more succinct than explicit definitions using acyclic terminologies. This closes the open problem of ten Cate et al. (2006) about the size of explicit definitions. Moreover, the same theorems establish an exact bound on the size of equivalent rewritings of concept queries as considered by Seylan et al. (2009). Theorem 4.11 also shows that Theorem 1 by Seylan et al. (2010), which claims a single exponential upper bound on the size of explicit definitions in $\mathcal{ALC}$, is wrong. The source of the problem in the proof of Theorem 1 is Lemma 1, which claims a single exponential upper bound on the size of interpolants in $\mathcal{ALC}$.

4.2 Failure of Beth Definability in the Presence of Role Hierarchies

We now show that BP fails in the description logics that we consider that include role hierarchies ($\mathcal{H}$). This shows that BP is indeed a stronger property than CBP because the same logics have CBP (ten Cate et al., 2006).

Theorem 4.18. *Let $\mathscr{L}$ be $\mathcal{ALCH}$ or any of its extensions with constructors from $\{\mathcal{S}, \mathcal{I}, \mathcal{F}\}$. Then $\mathscr{L}$ does not have BP.*

Proof. Let $\Sigma = \{R_1, R_2\}$ and consider the $\mathcal{ALCH}$-TBox $\mathcal{T}$ that consists of

$$\begin{array}{rcl} S & \sqsubseteq & R_1 \\ S & \sqsubseteq & R_2 \\ \exists R_1.A \sqcap \forall S.\bot & \sqsubseteq & \forall R_2.\neg A \\ \exists R_1.\neg A \sqcap \forall S.\bot & \sqsubseteq & \forall R_2.A \end{array}$$

It is easy to see that $\mathcal{T}$ is satisfiable. In fact, we will present two models of $\mathcal{T}$ below.

Claim 4.19. *$\exists S.\top$ is implicitly definable from Σ under $\mathcal{T}$.*

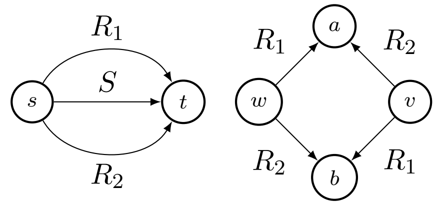

Figure 5: Interpretations $\mathcal{I}$ and $\mathcal{J}$ that are used for disproving BP for $\mathcal{ALCH}$

Proof of claim. Define $X_\mathcal{I} = \{s \in \Delta^\mathcal{I} \mid \exists t \in \Delta^\mathcal{I}.\langle s,t\rangle \in R_1^\mathcal{I} \cap R_2^\mathcal{I}\}$. We will show that, whenever $\mathcal{I} \models \mathcal{T}$, then $(\exists S.\top)^\mathcal{I} = X_\mathcal{I}$. This establishes the claim, since $X_\mathcal{I}$ depends only on $R_1^\mathcal{I}$ and $R_2^\mathcal{I}$.

First, we show $(\exists S.\top)^\mathcal{I} \subseteq X_\mathcal{I}$. Suppose that $s \in (\exists S.\top)^\mathcal{I}$. Then there is some $t \in \Delta^\mathcal{I}$ such that $\langle s,t\rangle \in S^\mathcal{I}$. By the RIAs in $\mathcal{T}$, we then have that $\langle s,t\rangle \in R_1^\mathcal{I} \cap R_2^\mathcal{I}$. Hence $s \in X_\mathcal{I}$.

Next, we show $X_\mathcal{I} \subseteq (\exists S.\top)^\mathcal{I}$. For contradiction, suppose that $s \in X_\mathcal{I}$ and $s \notin (\exists S.\top)^\mathcal{I}$, i.e., $s \in (\forall S.\bot)^\mathcal{I}$. Then there is some $t \in \Delta^\mathcal{I}$ such that $\langle s,t\rangle \in R_1^\mathcal{I} \cap R_2^\mathcal{I}$ and $\langle s,t\rangle \notin S^\mathcal{I}$. By the definition of an interpretation, either (i) $t \in A^\mathcal{I}$ or (ii) $t \in (\neg A)^\mathcal{I}$. If (i), then by $\exists R_1.A \sqcap \forall S.\bot \sqsubseteq \forall R_2.\neg A \in \mathcal{T}$, we have $t \in (\neg A)^\mathcal{I}$, which is a contradiction. If (ii), then by $\exists R_1.\neg A \sqcap \forall S.\bot \sqsubseteq \forall R_2.A \in \mathcal{T}$, we have $t \in A^\mathcal{I}$, which is a contradiction. Hence $X_\mathcal{I} \subseteq (\exists S.\top)^\mathcal{I}$. $\dashv$

Let $\mathcal{I} = \langle \Delta^\mathcal{I}, \cdot^\mathcal{I}\rangle$ be the interpretation where

- $\Delta^\mathcal{I} = \{s,t\}$,
- $R_1^\mathcal{I} = R_2^\mathcal{I} = S^\mathcal{I} = \{\langle s,t\rangle\}$;
- $R^\mathcal{I} = \emptyset$, for all $R \in N_R \setminus (\Sigma \cup \{S\})$;
- $B^\mathcal{I} = \emptyset$, for all $B \in N_C$.

Let $\mathcal{J} = \langle \Delta^\mathcal{J}, \cdot^\mathcal{J}\rangle$ be the interpretation where

- $\Delta^\mathcal{J} = \{w,v,a,b\}$,
- $R_1^\mathcal{J} = \{\langle w,a\rangle, \langle v,b\rangle\}$, $R_2^\mathcal{J} = \{\langle w,b\rangle, \langle v,a\rangle\}$;
- $R^\mathcal{J} = \emptyset$, for all $R \in N_R \setminus \Sigma$;
- $A^\mathcal{J} = \{a\}$;
- $B^\mathcal{J} = \emptyset$, for all $B \in (N_C \setminus \{A\})$.

The interpretations $\mathcal{I}$ and $\mathcal{J}$ are depicted in Figure 5. It is not hard to see that $\mathcal{I}$ and $\mathcal{J}$ are models of $\mathcal{T}$. Furthermore, the two structures are indistinguishable by concepts in the signature Σ, in the following sense:

Claim 4.20. *For all $\mathcal{SHIF}$-concepts C with $\mathsf{sig}(C) \subseteq \Sigma = \{R_1, R_2\}$, we have*

1. *$s \in C^\mathcal{I}$ if and only if $w \in C^\mathcal{J}$;*

2. $s \in C^{\mathcal{I}}$ *if and only if* $v \in C^{\mathcal{J}}$*;*

3. $t \in C^{\mathcal{I}}$ *if and only if* $a \in C^{\mathcal{J}}$*;*

4. $t \in C^{\mathcal{I}}$ *if and only if* $b \in C^{\mathcal{J}}$

The proof of this claim is straightforward, by simultaneous induction on the structure of the concept C (alternatively, bisimulations can be used to establish the same result).

Since $s \in (\exists S.\top)^{\mathcal{I}}$ and $w \notin (\exists S.\top)^{\mathcal{J}}$, it follows that there is no $\mathcal{SHIF}$-concept C such that $\mathsf{sig}(C) \subseteq \Sigma$ and $\mathcal{T} \models \exists S.\top \equiv C$. In summary, we have that the $\mathcal{ALCH}$-concept $\exists S.\top$ is implicitly definable from Σ under the $\mathcal{ALCH}$-TBox $\mathcal{T}$, but $\exists S.\top$ is not explicitly definable from Σ under $\mathcal{T}$ in $\mathcal{SHIF}$. We can conclude that BP fails for every description logic that includes $\mathcal{ALCH}$ and that is included in $\mathcal{SHIF}$. $\square$

Theorem 4.18 shows that Theorem 10 by Seylan et al. (2010), which claims that $\mathcal{ALCH}$ and its extensions with $\mathcal{S}$ and/or $\mathcal{I}$ have BP, is incorrect. The mistake in the proof is that Theorem 9, which presents a reduction from the concept satisfiability problem w.r.t. TBoxes in $\mathcal{SHI}$ to the same problem in $\mathcal{ALC}$, can not actually be used for computing $\mathcal{SHI}$-interpolants.

4.3 Failure of Beth Definability in the Finite

We now consider BPF (the analogue of Beth Definability over finite structures). Before we start, we explain our motivations. Seylan et al. (2009) consider an ontology-based data access setting, where traditional ABoxes are replaced by DBoxes. Syntactically, DBoxes are defined in the same way as ABoxes, but their semantics is different: while an ABox is merely assumed to express true facts, a DBox is assumed to list all true facts for some specified subset of the signature (known as the set of *data predicates*). Thus, for example, $\mathcal{D} = \{A(a), R(a, b)\}$ is a DBox for data predicates A and R, and, by the definition of the semantics of DBoxes, we have that, in every model $\mathcal{I}$ of $\mathcal{D}$, $A^{\mathcal{I}} = \{a^{\mathcal{I}}\}$ and $R^{\mathcal{I}} = \{\langle a^{\mathcal{I}}, b^{\mathcal{I}}\rangle\}$. In this setting, the TBox may contain other predicates than the data predicates and the authors use BP to determine whether a concept query over the signature of the TBox can be rewritten to an equivalent first-order query over the data predicates. When this is possible, computing the certain answers of the original query can then be reduced to computing the answers of the rewriting over the DBox, viewed as a database. In the setting we have described here, and for DLs without FMP, it is more natural to consider BPF than BP. The reason is that, in every interpretation of a DBox, the data predicates are, by definition, finite relations. In fact, the appropriate analogue of BP in this setting is one that is restricted to interpretations in which the data predicates are finite and the rest of the signature is unrestricted. This variant of BP can be viewed as a common generalization of BP and BFP. We do not study it here, but the negative results that we will present below for BFP apply to it as well.

Theorem 4.21 below establishes that BPF fails in $\mathscr{L}$, where $\mathscr{L}$ is any DL (among the ones we consider) lacking FMP. More precisely, we show that there is an $\mathscr{L}$-TBox $\mathcal{T}$, $\mathscr{L}$-concept C, and signature Σ such that C is f-implicitly definable from Σ under $\mathcal{T}$, and that there is no f-explicit definition in $\mathscr{L}$, i.e., there is no $\mathscr{L}$-concept D such that $\mathsf{sig}(D) \subseteq \Sigma$ and $\mathcal{T} \models_f C \equiv D$. Intuitively, the reason for the failure of BPF in these logics will be that they can not express the transitive closure of a role (see also the discussion below after the proof of Theorem 4.21).

Theorem 4.21. *Let $\mathscr{L}$ be $\mathcal{ALCFI}$ or any of its extensions with constructors from $\{\mathcal{S}, \mathcal{H}\}$. Then $\mathscr{L}$ does not have BPF.*

Proof. We will, in fact, prove something stronger: we will construct an implicit definition for which there is no corresponding explicit definition even in full first-order logic.

Let A, B, X be concept names and let R be a role name. Suppose $\Sigma = \{R, A\}$. Consider the $\mathcal{ALCFI}$-TBox $\mathcal{T}$ that consists of the following.

$$\begin{aligned}
\top &\sqsubseteq \ \leq 1R \ \sqcap \ \leq 1R^- \\
B &\sqsubseteq \exists R.B \\
A &\sqsubseteq X \\
\exists R.(A \sqcap \neg B) &\sqsubseteq \neg X \\
\exists R.\neg X &\sqsubseteq \neg X
\end{aligned}$$

We will show that some concept is f-implicitly definable from Σ under $\mathcal{T}$ but it is not f-explicitly definable from Σ under $\mathcal{T}$. The concept in question is $A \sqcap B$. Note that this concept is finitely satisfiable w.r.t. $\mathcal{T}$, i.e., there is some finite model $\mathcal{I}$ of $\mathcal{T}$ such that $(A \sqcap B)^\mathcal{I} \neq \emptyset$. In fact, we provide such a model $\mathcal{I}_n$ below.

For an interpretation $\mathcal{I} = \langle \Delta^\mathcal{I}, \cdot^\mathcal{I} \rangle$. A sequence $s_0, \ldots, s_n$ of elements of $\Delta^\mathcal{I}$ is called a *finite R-path* if $n > 0$ and $\langle s_i, s_{i+1} \rangle \in R^\mathcal{I}$ for all $i < n$. An *infinite R-path* is defined analogously. An R-path such that the start and the end nodes are the same is called an *R-cycle*. Now we will show two claims that will be useful for the proof of the theorem.

Claim 4.22. *Let $\mathcal{I}$ be a finite model of $\mathcal{T}$. If $s \in B^\mathcal{I}$, then $\langle s, s \rangle \in (R^\mathcal{I})^+$, where $(R^\mathcal{I})^+$ is the transitive closure of $R^\mathcal{I}$.*

Proof of claim. Suppose that $s \in B^\mathcal{I}$. Then the axiom $B \sqsubseteq \exists R.B \in \mathcal{T}$ implies the existence of the following infinite R-path:

$$p = s_0, s_1, \ldots$$

where $s_0 = s$ and for all $i \geq 0$, we have $s_i \in B^\mathcal{I}$.

Since $\mathcal{I}$ is finite, there is some $0 \leq n < m$ such that $s_n = s_m$. If $n = 0$, we immediately have that $\langle s, s \rangle \in (R^\mathcal{I})^+$. Otherwise, we claim that for all pairs $\langle s_i, s_j \rangle$ in the sequence $\langle s_n, s_m \rangle, \langle s_{n-1}, s_{m-1} \rangle, \ldots, \langle s_0, s_{m-n} \rangle$, we have $s_i = s_j$. The base case follows immediately from $s_n = s_m$. For the inductive step, we have by the inductive hypothesis that $s_{n-i} = s_{m-i} = t$, for some $t \in \Delta^\mathcal{I}$. Then by the definition of p, $\langle s_{n-i-1}, t \rangle \in R^\mathcal{I}$ and $\langle s_{m-i-1}, t \rangle \in R^\mathcal{I}$, which imply by the axiom $\top \sqsubseteq \leq 1R^- \in \mathcal{T}$ and $\mathcal{I} \models \mathcal{T}$ that $s_{n-i-1} = s_{m-i-1}$. Hence we have that $s = s_0 = s_{m-n}$. But then $\langle s, s \rangle \in (R^\mathcal{I})^+$, which is what we wanted to show. $\dashv$

Claim 4.23. *$A \sqcap B$ is f-implicitly definable from Σ under $\mathcal{T}$.*

Proof of claim. For all interpretations $\mathcal{I}$, define

$$Y_\mathcal{I} = \{s \in \Delta^\mathcal{I} \mid \langle s, s \rangle \in (R^\mathcal{I})^+ \wedge s \in A^\mathcal{I}\}.$$

We will show that, for all finite models $\mathcal{I}$ of $\mathcal{T}$, $(A \sqcap B)^{\mathcal{I}} = Y_{\mathcal{I}}$. This implies the claim, since $Y_{\mathcal{I}}$ depends only on $R^{\mathcal{I}}$ and $A^{\mathcal{I}}$.

($\Rightarrow$) Suppose $s \in (A \sqcap B)^{\mathcal{I}}$. By Claim 4.22, we know that $\langle s, s\rangle \in (R^{\mathcal{I}})^{+}$. But then $s \in Y_{\mathcal{I}}$.

($\Leftarrow$) Suppose $s \in Y_{\mathcal{I}}$. Then $s \in A^{\mathcal{I}}$ and $\langle s, s\rangle \in (R^{\mathcal{I}})^{+}$. Since $s \in A^{\mathcal{I}}$, it is enough to show that $s \in B^{\mathcal{I}}$. By the definition of an interpretation, either $s \in B^{\mathcal{I}}$ or $s \in (\neg B)^{\mathcal{I}}$. If $s \in (\neg B)^{\mathcal{I}}$, then by $\langle s, s\rangle \in (R^{\mathcal{I}})^{+}$, $s \in (A \sqcap \neg B)^{\mathcal{I}}$, and the axioms $\exists R.(A \sqcap \neg B) \sqsubseteq \neg X \in \mathcal{T}$ and $\exists R.\neg X \sqsubseteq \neg X \in \mathcal{T}$, we have $s \in (\neg X)^{\mathcal{I}}$ which would be a contradiction by $s \in A^{\mathcal{I}}$ and $A \sqsubseteq X \in \mathcal{T}$. So it must be that $s \in B^{\mathcal{I}}$. $\dashv$

The rest of the proof is all about showing that there is no f-explicit definition of $A \sqcap B$ from Σ under $\mathcal{T}$. To this aim, we start by defining an interpretation $\mathcal{I}_n$, parameterized by a natural number $n > 0$.

- $\Delta^{\mathcal{I}_n} = \{s_0, \ldots, s_{2n+1}\} \cup \{t_0, \ldots, t_{2n+1}\}$
- $R^{\mathcal{I}_n} = \{\langle s_i, s_{i+1}\rangle \mid 0 \leq i \leq 2n\} \cup \{\langle t_i, t_{i+1}\rangle \mid 0 \leq i \leq 2n\} \cup \{\langle t_{2n+1}, t_0, \}\rangle$
- $A^{\mathcal{I}_n} = \{s_n, t_n\}$
- $B^{\mathcal{I}_n} = X^{\mathcal{I}_n} = \{t_i \mid 0 \leq i \leq 2n+1\}$

Claim 4.24. *For every first-order formula $\varphi(x)$ there is an $n > 0$ such that $\mathcal{I}_n \models \varphi\ [s_n]$ if and only if $\mathcal{I}_n \models \varphi\ [t_n]$.*

Proof. We apply the Gaifman locality theorem (cf. Libkin, 2004). In the present setting, where we only have unary and binary relations, and we are concerned with formulas in a single free variable, the Gaifman locality theorem is particularly easy to state. Given an interpretation $\mathcal{I}$ and elements $a, b \in \Delta^{\mathcal{I}}$, we say that a and b have *distance at most* n relative to a signature Σ, if there is a sequence $s_0, \ldots, s_m$ with $0 \leq m \leq n$ such that $s_0 = a$, $s_m = b$, and for all $0 \leq i < m$, pair $\langle s_i, s_{i+1}\rangle$ belongs to $P^{\mathcal{I}} \cup (P^{\mathcal{I}})^{-}$ for some binary relation (i.e., role name) $P \in \Sigma$. For any interpretation $\mathcal{I}$, element $a \in \Delta^{\mathcal{I}}$, and natural number $n \geq 0$, we denote by $\mathcal{I} \restriction_{a,n}$ the interpretation whose domain consists of the elements from $\Delta^{\mathcal{I}}$ that have distance at most n from a, and whose relations are the ones from $\mathcal{I}$ restricted to this subset of the domain. The Gaifman locality theorem can then be stated as follows: for every first-order formula $\varphi(x)$, there is a natural number $n > 0$ such that, for all structures $\mathcal{I}$ and elements $a, b \in \Delta^{\mathcal{I}}$, if $\mathcal{I} \restriction_{a,n}$ is isomorphic to $\mathcal{I} \restriction_{b,n}$, via an isomorphism that maps a to b, then $\mathcal{I} \models \varphi\ [a]$ if and only if $\mathcal{I} \models \varphi\ [b]$.

Now, let $\varphi(x)$ be any first-order formula, and let $n > 0$ be the natural number given by the Gaifman locality theorem. Consider the instance $\mathcal{I}_n$ that we constructed earlier. It is immediately clear from the construction of $\mathcal{I}_n$ that $\mathcal{I}_n \restriction_{s_n,n}$ is isomorphic to $\mathcal{I}_n \restriction_{t_n,n}$, via an isomorphism that maps s_n to t_n. Therefore, $\mathcal{I}_n \models \varphi\ [s_n]$ if and only if $\mathcal{I}_n \models \varphi\ [t_n]$. $\square$

Claim 4.25. *There is no $\mathcal{SHIF}$-concept C such that $\mathsf{sig}(C) \subseteq \Sigma$ and $\mathcal{T} \models_f (A \sqcap B) \equiv C$.*

Proof of claim. We proceed towards a contradiction. Suppose C is an $\mathcal{ALCFI}$-concept such that $\mathsf{sig}(C) \subseteq \Sigma = \{R, A\}$ and $\mathcal{T} \models_f A \sqcap B \equiv C$. Let $\varphi(x) = \pi_x(C)$. Since $s_n \in (A \sqcap B)^{\mathcal{I}_n}$, $\mathcal{T} \models_f A \sqcap B \equiv C$, and the fact that $\mathcal{I}_n$ is a finite model of $\mathcal{T}$, we have $\mathcal{I}_n \not\models \varphi\ [s_n]$; and by the same reasoning, we have $\mathcal{I}_n \models \varphi\ [t_n]$. But by the previous claim, $\mathcal{I}_n \models \varphi\ [s_n]$ if and only if $\mathcal{I}_n \models \varphi\ [t_n]$, which is a contradiction. $\dashv$

In summary, we have that the $\mathcal{ALCFI}$-concept $A \sqcap B$ is f-implicitly definable from Σ under the $\mathcal{ALCFI}$-TBox $\mathcal{T}$, but $A \sqcap B$ is not f-explicitly definable from Σ under $\mathcal{T}$ even in $\mathcal{SHIF}$ (or in first-order logic, for that matter). It follows that, if $\mathscr{L}$ is any proper extension of $\mathcal{ALCFI}$ with constructors from $\{\mathcal{S}, \mathcal{H}\}$, then $\mathscr{L}$ does not have BPF. $\square$

We point out that the specific counterexample to BPF described in the proof of Theorem 4.21 actually admits an explicit definition if one were to allow the use of transitive closure. Specifically, it can be shown that $A \sqcap (\exists R.\top) \sqcap (\neg\exists R^+.\neg\exists R.\top)$ is an explicit definition, where R^+ denotes the transitive closure of the role R.

4.4 The Transitive Closure Operator

The proof of Theorem 4.21 suggests that the failure of BPF in the considered logics may be caused by the fact that they can not express transitive closure. This raises the question whether one can regain BPF by adding the transitive closure constructor to $\mathcal{ALCFI}$. In this section, we show that $\mathcal{ALCFI}$ extended with the transitive closure constructor still lacks BPF.

In the following, we denote by $\mathscr{L}_+$ the language obtained from $\mathscr{L}$ by additionally allowing R^+ as a role for every role R in $\mathscr{L}$. This allows us to include such roles in the inductive definition of concepts. However, if $\mathscr{L}$ includes functionality restrictions, then, as usual, we forbid the use of transitive closure inside these functionality restrictions. In other words, in concepts of the form $\leq 1R$, R is not allowed to make use of the transitive closure constructor.

The semantics of the transitive closure construct is as expected, namely, $(R^+)^\mathcal{I}$ is the relation

$$\{\langle s,t\rangle \mid \text{ there are } s_1,\dots,s_n\ (n>1) \text{ with } s_1=s,\ s_n=t, \text{ and } \langle s_i,s_{i+1}\rangle \in R^\mathcal{I} \text{ for } 1\leq i<n\}.$$

Theorem 4.26. *$\mathcal{ALCFI}_+$ does not have BPF.*

Proof. Consider the following $\mathcal{ALCFI}_+$-TBox $\mathcal{T}$.

$$\begin{array}{rcl@{\qquad\qquad}rcl}
\top & \sqsubseteq & \leq 1R^- & A & \sqsubseteq & \neg B \\
\top & \sqsubseteq & \exists R.\top & \exists R.B & \sqsubseteq & \neg B \\
\top & \sqsubseteq & \exists R^+.A & \exists R.\neg B & \sqsubseteq & B
\end{array}$$

It is easy to see that $\mathcal{T}$ is finitely satisfiable, i.e., $\mathcal{T}$ has a finite model. In fact, we provide finite models $\mathcal{I}_n$, for $n > 0$, of $\mathcal{T}$ below. We first show that the concept name B is f-implicitly definable from $\Sigma = \{R, A\}$ under $\mathcal{T}$ and then show that there is no f-explicit definition of this concept from Σ in the language. For an interpretation $\mathcal{I}$, an R-path and R-cycle in $\mathcal{I}$ are defined as in the proof of Theorem 4.21.

Claim 4.27. *Let $\mathcal{I}$ be a finite model of $\mathcal{T}$. Then for all $s \in \Delta^\mathcal{I}$, we have $\langle s,s\rangle \in (R^+)^\mathcal{I}$.*

Proof of claim. Identical to the proof of Claim 4.22 in the proof of Theorem 4.21. $\dashv$

Claim 4.28. *We have*

(a) $\mathcal{T} \models_f \top \sqsubseteq \exists R^{-}.\top$*, and*

(b) $\mathcal{T} \models_f \top \sqsubseteq \leq 1R$*.*

Proof of claim. Part (a) follows immediately from the previous claim.

To prove part (b), suppose, for the sake of a contradiction that $\mathcal{T} \not\models_f \top \sqsubseteq \leq 1R$. Then there is some finite model $\mathcal{I}$ of $\mathcal{T}$ and $s, t, u \in \Delta^{\mathcal{I}}$ such that $t \neq u$ and $\langle s, t\rangle, \langle s, u\rangle \in R^{\mathcal{I}}$. Since $\top \sqsubseteq \leq 1R^{-} \in \mathcal{T}$ and $\mathcal{I} \models \mathcal{T}$ and by part (a) of our claim, we have that $(R^{-})^{\mathcal{I}}$ is the graph of a total function on $\Delta^{\mathcal{I}}$. Since $t \neq u$ and $\langle s, t\rangle, \langle s, u\rangle \in R^{\mathcal{I}}$, we also know that the cardinality of the image Y of this function must be strictly smaller than the cardinality of the domain, i.e., $|Y| < |\Delta^{\mathcal{I}}|$. This implies that $Y \subsetneq \Delta^{\mathcal{I}}$. But this contradicts with $\top \sqsubseteq \exists R.\top \in \mathcal{T}$ and $\mathcal{I} \models \mathcal{T}$. Hence we conclude that $\mathcal{T} \models_f \top \sqsubseteq \leq 1R$. ⊣

For $s, t \in \Delta^{\mathcal{I}}$, we write $\mathsf{odd}(s, t)$ if there is an R-path of odd length from s to t, that is, an R-path $s_0, \ldots, s_n$ such that $s = s_0$, $t = s_n$, and n is odd. Note that in an R-path like $s_0, \ldots, s_n$, we always have $n > 0$.

Claim 4.29. *For all finite models* $\mathcal{I}$ *of* $\mathcal{T}$*, we have*

$$B^{\mathcal{I}} = \{s \in \Delta^{\mathcal{I}} \mid \exists t \in \Delta^{\mathcal{I}}.\mathsf{odd}(s, t) \wedge t \in A^{\mathcal{I}}\}.$$

Proof of claim. ($\Rightarrow$) Suppose $s \in B^{\mathcal{I}}$. By the first claim, there is some R-cycle $p = s_0, \ldots, s_n$, where $s_0 = s_n = s$ and $n > 0$. Since $\top \sqsubseteq \exists R^{+}.A \in \mathcal{T}$ and $\mathcal{I} \models \mathcal{T}$, there is some $t \in \Delta^{\mathcal{I}}$ such that $\langle s, t\rangle \in (R^{+})^{\mathcal{I}}$ and $t \in A^{\mathcal{I}}$. We claim that $t = s_i$, for some $i \in \{1, \ldots, n-1\}$. To show this, we proceed towards a contradiction.

Suppose that our claim does not hold. Then there is some R-path $t_0, \ldots, t_m$, where $t_0 = s$ and $t_m = t$. Obviously, this path is different from the R-cycle p since t does not occur in p. Now by using $\mathcal{T} \models_f \top \sqsubseteq \leq 1R$ from the previous claim, we can show that every individual t_i actually appears in p, which contradicts with the fact that t does not appear in p. Hence we conclude that there is some $i \in \{1, \ldots, n-1\}$ such that $s_i = t$.

We will now show that i is odd. By $A \sqsubseteq \neg B \in \mathcal{T}$, $\mathcal{I} \models \mathcal{T}$, and $t \in A^{\mathcal{I}}$, we have $t \notin B^{\mathcal{I}}$. Then by using $s \in B^{\mathcal{I}}$ and the axioms $\exists R.B \sqsubseteq \neg B \in \mathcal{T}$, $\exists R.\neg B \sqsubseteq B \in \mathcal{T}$, one can easily show by induction that i is odd. This implies $\mathsf{odd}(s, t)$.

Hence there is some $t \in \Delta^{\mathcal{I}}$ such that $\mathsf{odd}(s, t)$ and $t \in A^{\mathcal{I}}$, which is what we wanted to show.

($\Leftarrow$) Suppose $s \in \Delta^{\mathcal{I}}$ such that there is some $t \in \Delta^{\mathcal{I}}$ with $\mathsf{odd}(s, t)$ and $t \in A^{\mathcal{I}}$. This implies that there is some R-path $s_0, \ldots, s_n$ such that $s_0 = s$, $s_n = t$, and n is odd. By $t \in A^{\mathcal{I}}$, $A \sqsubseteq \neg B \in \mathcal{T}$, and $\mathcal{I} \models \mathcal{T}$, we have $t \notin B^{\mathcal{I}}$. Then by using the fact that n is odd, the axioms $\exists R.B \sqsubseteq \neg B \in \mathcal{T}$ and $\exists R.\neg B \sqsubseteq B \in \mathcal{T}$, one can easily show by induction that $s \in B^{\mathcal{I}}$. ⊣

Claim 4.29 implies that B is f-implicitly definable from Σ under $\mathcal{T}$. The rest of the proof shows that there is no f-explicit definition of B from Σ under $\mathcal{T}$. For each $n \geq 0$, let $\mathcal{I}_n$ be the following interpretation:

- $\Delta^{\mathcal{I}_n} = \{s_0, \ldots, s_{2n+3}\}$

- $R^{\mathcal{I}_n} = \{\langle s_i, s_{i+1}\rangle \mid 0 \le i < 2n+3\} \cup \{\langle s_{2n+3}, s_0\rangle\}$
- $A^{\mathcal{I}_n} = \{s_{n+2}\}$,
- $B^{\mathcal{I}_n} = \{s_i \mid 0 \le i \le 2n+3 \text{ and } \mathsf{odd}(s_i, s_{n+2})\}$.

Intuitively, $\mathcal{I}_n$ is an R-cycle of (even) length $2n+4$, some of whose elements satisfy the concept name A and/or B. Observe that $\mathcal{I}_n$, for $n \ge 0$, is a model of $\mathcal{T}$. Define the function $d : \Delta^{\mathcal{I}_n} \to \mathbb{N}$ as follows.

$$d(s_i) = \begin{cases} i-(n+2) & \text{if } i \ge n+2 \\ (n+2)-i & \text{if } i < n+2 \end{cases}$$

In other words, $d(s)$ is the distance between s and s_{n+2}.

For each $\mathcal{ALCFI}_+$-concept C, we will denote by $\mathsf{md}(C)$ the *modal depth* of C, that is, the maximal nesting depth of role constructors in C. Formally,

- $\mathsf{md}(A) = \mathsf{md}(\top) = \mathsf{md}(\le 1R) = 0$
- $\mathsf{md}(\neg C) = \mathsf{md}(C)$
- $\mathsf{md}(C \sqcap D) = \max\{\mathsf{md}(C), \mathsf{md}(D)\}$
- $\mathsf{md}(\exists R.C) = \mathsf{md}(\exists R^+.C) = \mathsf{md}(C)+1$

where $A \in C_A$ and R is of the form P or P^- with $P \in N_R$.

Claim 4.30. *For all $i \in \{0,\dots,n\}$ and all $s, s' \in \Delta^{\mathcal{I}_n} \setminus \{s \in \Delta^{\mathcal{I}_n} \mid d(s) \le i\}$, we have*

$$s \in C^{\mathcal{I}_n} \textit{ iff } s' \in C^{\mathcal{I}_n}$$

for all $\mathcal{ALCFI}_+$-concepts with $\mathsf{md}(C) \le i$ and $\mathsf{sig}(C) \subseteq \Sigma$.

Proof of claim. Let $i \in \{0,\dots,n\}$, $s, s' \in \Delta^{\mathcal{I}_n} \setminus \{s \in \Delta^{\mathcal{I}_n} \mid d(s) \le i\}$, and C be an $\mathcal{ALCFI}_+$-concept with $\mathsf{md}(C) \le i$ and $\mathsf{sig}(C) \subseteq \Sigma$. The proof is by induction on i.

For $i = 0$. Since $\mathsf{md}(C) = 0$ and $\mathsf{sig}(C) \subseteq \Sigma$, C obeys the following grammar:

$$C ::= \top \mid A \mid \le 1S \mid \neg C \mid C \sqcap C$$

where $S = R$ or $S = R^-$ (recall that we forbid the use of transitive closure inside functionality restrictions).

By induction on the structure of C, we show that $s \in C^{\mathcal{I}_n}$ iff $s' \in C^{\mathcal{I}_n}$.

- $C = \top$. By the definition of an interpretation, we have $s \in \top^{\mathcal{I}_n}$ and $s' \in \top^{\mathcal{I}_n}$. Hence $s \in \top^{\mathcal{I}_n}$ iff $s' \in \top^{\mathcal{I}_n}$.
- $C = A$ (recall that A is the only concept name in Σ). By assumption, we have $s \ne s_{n+2}$ and $s' \ne s_{n+2}$. Then by the definition of $\mathcal{I}_n$, we obtain $s \notin A^{\mathcal{I}_n}$ and $s' \notin A^{\mathcal{I}_n}$. Hence $s \in A^{\mathcal{I}_n}$ iff $s' \in A^{\mathcal{I}_n}$.
- $C = \le 1S$. By the definition of $\mathcal{I}_n$, we have for all $t \in \Delta^{\mathcal{I}_n}$ that $|S^{\mathcal{I}_n}(t)| = 1$. Then in particular, $|S^{\mathcal{I}_n}(s)| = |S^{\mathcal{I}_n}(s')| = 1$. Hence $s \in (\le 1S)^{\mathcal{I}_n}$ iff $s' \in (\le 1S)^{\mathcal{I}_n}$.

- $C = \neg D$. Follows easily by the inductive hypothesis for C.
- $C = C_1 \sqcap C_2$. Follows easily by the inductive hypothesis for C.

Hence we conclude that $s \in C^{\mathcal{I}_n}$ iff $s' \in C^{\mathcal{I}_n}$, for $i = 0$.

Next, consider the case that $i > 0$, and let $\mathsf{md}(C) \leq i$ and $\mathsf{sig}(C) \subseteq \Sigma$. Then we have that C obeys the following grammar:

$$C ::= \exists S.E \mid \exists S^+.E \mid \neg C \mid C \sqcap C$$

where $\mathsf{md}(E) \leq i - 1$, and $S = R$ or $S = R^-$. By induction on the structure of C, we show that $s \in C^{\mathcal{I}_n}$ iff $s' \in C^{\mathcal{I}_n}$.

- $C = \exists S.E$ with $\mathsf{md}(E) \leq i - 1$, and $S = R$ or $S = R^-$. By the definition of $\mathcal{I}_n$, we have that there is exactly one $t \in \Delta^{\mathcal{I}_n}$ with $\langle s, t\rangle \in S^{\mathcal{I}_n}$ and exactly one $t' \in \Delta^{\mathcal{I}_n}$ with $\langle s', t'\rangle \in S^{\mathcal{I}_n}$. Moreover, $t, t' \in \Delta^{\mathcal{I}_n} \setminus \{s \in \Delta^{\mathcal{I}_n} \mid d(s) \leq i - 1\}$. We have that the following are equivalent:
  - $s \in (\exists S.E)^{\mathcal{I}_n}$
  - $t \in E^{\mathcal{I}_n}$ (since t is the only individual with $\langle s, t\rangle \in S^{\mathcal{I}_n}$)
  - $t' \in E^{\mathcal{I}_n}$ (by the inductive hypothesis for i)
  - $s' \in (\exists S.E)^{\mathcal{I}_n}$ (since t is the only individual with $\langle s', t'\rangle \in S^{\mathcal{I}_n}$).
- $D = \exists S^+.E$ with $\mathsf{md}(E) = i-1$, and $S = R$ or $S = R^-$. Suppose first $s \in (\exists S^+.E)^{\mathcal{I}_n}$. Then there is some $t \in \Delta^{\mathcal{I}}$ such that $\langle s, t\rangle \in (S^{\mathcal{I}_n})^+$ and $t \in E^{\mathcal{I}_n}$. We distinguish the following cases:
  - $t \neq s'$. Then by the definition of $\mathcal{I}_n$, we immediately obtain $\langle s', t\rangle \in (S^{\mathcal{I}_n})^+$; and by $t \in E^{\mathcal{I}_n}$ this implies $s' \in (\exists S^+.E)^{\mathcal{I}_n}$.
  - $t = s'$. By the definition of $\mathcal{I}_n$, we have $\langle s', s\rangle \in (S^{\mathcal{I}_n})^+$. Moreover, $s, s' \in \Delta^{\mathcal{I}_n} \setminus \{s \in \Delta^{\mathcal{I}_n} \mid d(s) \leq i - 1\}$. Then by the inductive hypothesis on i and $s' \in E^{\mathcal{I}_n}$, we have $s \in E^{\mathcal{I}_n}$, which implies by $\langle s', s\rangle \in (S^{\mathcal{I}_n})^+$ that $s' \in (\exists S^+.E)^{\mathcal{I}_n}$.

  Hence $s' \in (\exists S^+.E)^{\mathcal{I}_n}$ in both cases, which is what we wanted to show. The direction from right to left can be shown analogously.
- The other cases can be shown easily by the inductive hypothesis on C.

Hence the claim follows. $\dashv$

Claim 4.31. *There is no $\mathcal{ALCFI}_+$-concept C such that $\mathsf{sig}(C) \subseteq \{A, R\}$ and $\mathcal{T} \models_f B \equiv C$.*

Proof of claim. We proceed towards a contradiction so suppose the existence of such a concept C. By definition, $\mathsf{md}(C) = n$, for some $n \geq 0$; and $s_0, s_1 \in \Delta^{\mathcal{I}_n} \setminus \{s \in \Delta^{\mathcal{I}_n} \mid d(s) \leq n\}$. Then by the previous claim, we have $s_0 \in C^{\mathcal{I}_n}$ iff $s_1 \in C^{\mathcal{I}_n}$. Then by the fact that $\mathcal{I}_n$ is a finite model of $\mathcal{T}$ and $\mathcal{T} \models_f B \equiv C$, we have $s_0 \in B^{\mathcal{I}_n}$ iff $s_1 \in B^{\mathcal{I}_n}$. This implies by the definition of $\mathcal{I}_n$ and Claim 4.29 that $\mathsf{odd}(s_0, s_{n+2})$ iff $\mathsf{odd}(s_1, s_{n+2})$, which is a contradiction. Hence we conclude that there exists no $\mathcal{ALCFI}_+$-concept C such that $\mathsf{sig}(C) \subseteq \{A, R\}$ and $\mathcal{T} \models B \equiv C$. $\dashv$

Now the proof of the theorem is as follows. By Claim 4.29, B is f-implicitly definable from $\Sigma = \{A, R\}$ under $\mathcal{T}$. But by Claim 4.31, B is not f-explicitly definable from Σ under $\mathcal{T}$. Hence $\mathcal{ALCFI}_+$ does not have BPF. $\square$

5. Concluding Remarks

In this paper, we studied BP in expressive DLs with commonly used concept constructors. All of these constructors appear in the Web Ontology Language OWL-Lite (Horrocks et al., 2003). OWL-Lite is now superseded by OWL 2, which supports some other important constructors such as *nominals*, denoted by $\mathcal{O}$ in the language, and *qualified number restrictions*, denoted by $\mathcal{Q}$ in the language. There are already some results available regarding BP in logics having $\mathcal{Q}$ or $\mathcal{O}$.

$\mathcal{Q}$ is a generalization of $\mathcal{F}$ and ten Cate et al. (2006) show via a model-theoretic argument that CBP holds in $\mathcal{ALCQ}$. We believe that BP can also be shown to hold for $\mathcal{ALCQ}$ and $\mathcal{ALCQI}$ using a model-theoretic argument; although such an argument gives no upper bound on the size of explicit definitions. Extending our upper bound results on the size of explicit definitions to these logics appears to be more difficult because of the unavailability of a natural and optimal tableau algorithm for these logics.

In logics with $\mathcal{O}$, besides the concept and role names, we assume a set $N_I = \{i, j, \dots\}$ of nominals. Syntactically, nominals are treated as atomic concepts but semantically each nominal is interpreted as a singleton set. The presence of nominals gives rise to two different Beth definability properties. In the first one, we are allowed to restrict the nominals appearing in implicit/explicit definitions by making them part of the signature Σ; in the second one, definitions are allowed to use any nominal from N_I. Obviously, the first one is a stronger property. Ten Cate et al. (2006) show that even the second property fails in $\mathcal{ALCO}$. They also observe that extending $\mathcal{ALCO}$ with concepts of the form $@_iC$ is enough to regain CBP. Intuitively, $@_iC$ says that the point satisfying the nominal i also satisfies the concept C.

In a similar way, one can try to identify an extension of $\mathcal{ALCH}$ that has BP. In the proof of Theorem 4.18, our argument for the failure of BP in the considered logics was that they can not express role conjunction. It remains open if $\mathcal{ALCH}$ extended with the role conjunction constructor has BP. Another interesting open question is to identify a minimal extension of $\mathcal{ALCFI}$ having BPF.

By Theorem 3.36, we know how to compute first-order explicit definitions of single exponential size, given that a concept is implicitly defined under a TBox. We leave as another open problem the existence of a matching lower bound, i.e., is there a family of TBoxes implicitly defining a concept such that smallest explicit definitions in first-order logic are single exponentially big?

Acknowledgments

We are grateful to Carsten Lutz and Maarten Marx for helpful discussions on the topic. A substantial part of the research was carried out during an extended visit of İnanç Seylan to

UC Santa Cruz in 2010, and we thank Phokion Kolaitis for his hospitality. We also thank the anonymous reviewers for their extensive comments.

Balder ten Cate was supported by the NSF grants IIS-0905276 and IIS-1217869.

Appendix A. Quasimodels

Decision procedures based on semantic tableau do not construct a model of the given formula/concept, but a finite representation of a model from which the model can be unfolded. In this paper, we will use the term 'quasimodel' to denote such a finite representation following Andréka, Németi, and van Benthem (1998). Various other names have been used in the literature, including Hintikka structures (Schwendimann, 1998), model graph (Goré, 1999), and even tableau (Horrocks & Sattler, 2007). Modulo some differences, the building blocks of these structures are sets of finite concepts each of which is a subset of a relevant *concept closure*. We will be using the definition of concept closure $\mathsf{cl}(C, \mathcal{T})$ given in Section 3.1.

Remark A.1. *For the rest of the appendix, we assume that $\mathcal{ALCF}$-concepts are defined recursively as in Section 2.1 using also $\bot$, $\sqcup$, $\forall R.C$, and $\geq 2R$ as primitives; all concepts are in NNF; and $\mathcal{ALCF}$-TBoxes consist only of axioms of the form $\top \sqsubseteq C$. For a discussion of these assumptions, we refer the reader to the beginning of Section 3.1.*

Not every subset of the concept closure is suitable to take part in a quasimodel. Depending on the logic at hand, these sets satisfy some basic consistency requirements. Following, e.g., Lutz et al. (2005), we will use the term '*type*' to denote these sets satisfying these requirements. Note, however, that the non-membership of a concept in a type does not imply the membership of the negation of the concept in the type. In this respect, our types are similar to Hintikka sets, which are also called downward-saturated sets (cf. Fitting, 1996).

Definition A.2. *Let C_0 be an $\mathcal{ALCF}$-concept and let $\mathcal{T}$ be an $\mathcal{ALCF}$-TBox. A $\tau \subseteq \mathsf{cl}(C_0, \mathcal{T})$ is called an $\langle C_0, \mathcal{T}\rangle$-*type *for $\mathcal{ALCF}$ if and only if for all $A, C, C_1, C_2, \exists R.C, \geq 2R, \leq 1R \in \mathsf{cl}(C_0, \mathcal{T})$,*

($\mathsf{P}_\bot$) *$\bot \notin \tau$;*

($\mathsf{P}_\neg$) *$\{A, \neg A\} \not\subseteq \tau$;*

($\mathsf{P}_\sqcap$) *if $C_1 \sqcap C_2 \in \tau$, then $C_1 \in \tau$ and $C_2 \in \tau$;*

($\mathsf{P}_\sqcup$) *if $C_1 \sqcup C_2 \in \tau$, then $C_1 \in \tau$ or $C_2 \in \tau$;*

($\mathsf{P}_\sqsubseteq$) *if $\top \sqsubseteq C \in \mathcal{T}$, then $C \in \tau$;*

($\mathsf{P}_\bowtie$) *$\{\leq 1R, \geq 2R\} \not\subseteq \tau$;*

($\mathsf{P}_{\leq 1}$) *if $\{\leq 1R, \exists R.C\} \subseteq \tau$, then $\forall R.C \in \tau$.*

When a type belongs to a quasimodel, it may force some other type to also belong to the quasimodel, for instance to witness an existential statement. In fact, a quasimodel is a collection of types coherent with each other in this sense.

Definition A.3. *Let C_0 be an $\mathcal{ALCF}$-concept, $\mathcal{T}$ an $\mathcal{ALCF}$-TBox and τ, υ two $\langle C_0, \mathcal{T}\rangle$-types for $\mathcal{ALCF}$.*

- *We write $\tau \xrightarrow{\exists R.C} \upsilon$ if $\exists R.C \in \tau$ and $\{C\} \cup \{C' \mid \forall R.C' \in \tau\} \subseteq \upsilon$.*
- *We write $\tau \xrightarrow{\geq 2R} \upsilon$ if $\geq 2R \in \tau$ and $\{C' \mid \forall R.C' \in \tau\} \subseteq \upsilon$.*
- *A set Q of $\langle C_0, \mathcal{T}\rangle$-types for $\mathcal{ALCF}$ is a $\langle C_0, \mathcal{T}\rangle$-*quasimodel for $\mathcal{ALCF}$ *if it satisfies:*
  (a) there is some $\tau_0 \in Q$ such that $C_0 \in \tau_0$;
  (b) for every $\tau \in Q$ and $\exists R.C \in \tau$, there is a type $\upsilon \in Q$ such that $\tau \xrightarrow{\exists R.C} \upsilon$; and
  (c) for every $\tau \in Q$ and $\geq 2R \in \tau$, there is a type $\upsilon \in Q$ such that $\tau \xrightarrow{\geq 2R} \upsilon$.

The following theorem will be useful in soundness and completeness proofs of the tableau and interpolation algorithms. Its proof is inspired by Marx and Venema (2007).

Theorem A.4. *An $\mathcal{ALCF}$-concept C_0 is satisfiable w.r.t. an $\mathcal{ALCF}$-TBox $\mathcal{T}$ if and only if there is some $\langle C_0, \mathcal{T}\rangle$-quasimodel for $\mathcal{ALCF}$.*

Proof. ($\Rightarrow$) Given a model $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}}\rangle$ of $\mathcal{T}$ with $C_0^{\mathcal{I}} \neq \emptyset$, we carve out for all $s \in \Delta^{\mathcal{I}}$, a set of concepts $\mathcal{L}(s) \subseteq \mathsf{cl}(C_0, \mathcal{T})$ as follows.

$$\mathcal{L}(s) = \{C \in \mathsf{cl}(C_0, \mathcal{T}) \mid s \in C^{\mathcal{I}}\}.$$

Now let $Q = \{\mathcal{L}(s) \mid s \in \Delta^{\mathcal{I}}\}$.

Claim A.5. *Each $\tau \in Q$ is a $\langle C_0, \mathcal{T}\rangle$-type for $\mathcal{ALCF}$.*

Proof of claim. Suppose $\tau \in Q$. Then $\tau = \mathcal{L}(s)$ for some $s \in \Delta^{\mathcal{I}}$. We verify the conditions in Definition A.2.

- By definition, $\bot^{\mathcal{I}} = \emptyset$ and thus $s \notin \bot^{\mathcal{I}}$ and thus $\bot \notin \mathcal{L}(s)$. Hence ($\mathsf{P}_{\bot}$) is satisfied.
- By the virtue of $\mathcal{I}$ being an interpretation, it is not the case that $s \in A^{\mathcal{I}}$ and $s \in (\neg A)^{\mathcal{I}}$. Hence ($\mathsf{P}_{\neg}$) is satisfied.
- If $C_1 \sqcap C_2 \in \mathcal{L}(s)$, then $s \in (C_1 \sqcap C_2)^{\mathcal{I}}$. Since $\mathcal{I}$ is an interpretation, $s \in C_1^{\mathcal{I}}$ and $s \in C_2^{\mathcal{I}}$. But then $C_1, C_2 \in \mathcal{L}(s)$. Hence ($\mathsf{P}_{\sqcap}$) is satisfied.
- If $C_1 \sqcup C_2 \in \mathcal{L}(s)$, then $s \in (C_1 \sqcup C_2)^{\mathcal{I}}$. Since $\mathcal{I}$ is an interpretation, $s \in C_1^{\mathcal{I}}$ or $s \in C_2^{\mathcal{I}}$. But then $C_1 \in \mathcal{L}(s)$ or $C_2 \in \mathcal{L}(s)$. Hence ($\mathsf{P}_{\sqcup}$) is satisfied.
- If $\top \sqsubseteq C \in \mathcal{T}$, then $\Delta^{\mathcal{I}} \subseteq C^{\mathcal{I}}$ and thus $s \in C^{\mathcal{I}}$. But then $C \in \mathcal{L}(s)$. Hence ($\mathsf{P}_{\sqsubseteq}$) is satisfied.
- Suppose for a contradiction that ($\mathsf{P}_{\bowtie}$) does not hold. Then $s \in (\leq 1R)^{\mathcal{I}}$ and $s \in (\geq 2R)^{\mathcal{I}}$. But this is a contradiction. Hence ($\mathsf{P}_{\bowtie}$) is satisfied.
- Suppose $\{\leq 1R, \exists R.C\} \subseteq \tau$. By assumption $s \in (\leq 1R)^{\mathcal{I}}$ and $s \in (\exists R.C)^{\mathcal{I}}$. Then it follows that there is exactly one $t \in \Delta^{\mathcal{I}}$ such that $\langle s, t\rangle \in R^{\mathcal{I}}$ and $t \in C^{\mathcal{I}}$. But then $s \in (\forall R.C)^{\mathcal{I}}$. Hence ($\mathsf{P}_{\leq 1}$) is satisfied.

Since we have shown that all the conditions in Definition A.2 are satisfied, we conclude that τ is a $\langle C_0, \mathcal{T}\rangle$-type for $\mathcal{ALCCF}$. $\dashv$

We claim that Q is a $\langle C_0, \mathcal{T}\rangle$-quasimodel. By Claim A.5, if $\tau \in Q$, then τ is a $\langle C_0, \mathcal{T}\rangle$-type for $\mathcal{ALCCF}$. Thus it remains to show that condition (a), (b), and (c) from Definition A.3 are satisfied.

For (a), since $C_0^{\mathcal{I}} \neq \emptyset$, there is some $s_0 \in \Delta^{\mathcal{I}}$ such that $s_0 \in C_0^{\mathcal{I}}$ and by the construction of Q, $\mathcal{L}(s_0)$ is in Q. Hence, condition (a) is satisfied.

For condition (b), suppose $\exists R.C \in \mathcal{L}(s)$ for some $s \in \Delta^{\mathcal{I}}$. This means $s \in (\exists R.C)^{\mathcal{I}}$, i.e., there is some individual t such that $\langle s, t\rangle \in R^{\mathcal{I}}$ and $t \in C^{\mathcal{I}}$. Then by the construction of Q, we have $C \in \mathcal{L}(t)$. Now let $\forall R.D \in \mathcal{L}(s)$. Then by the construction of Q, we have $s \in (\forall R.D)^{\mathcal{I}}$. This implies by $\langle s, t\rangle \in R^{\mathcal{I}}$ that $t \in D^{\mathcal{I}}$. By the construction of Q again, we obtain $D \in \mathcal{L}(t)$. Hence, $\mathcal{L}(s) \xRightarrow{\exists R.C} \mathcal{L}(t)$; and we conclude that (b) is satisfied.

The proof for (c) is analogous.

($\Leftarrow$) Suppose that Q is a $\langle C_0, \mathcal{T}\rangle$-quasimodel for $\mathcal{ALCCF}$. The idea of the proof is to construct an interpretation $\mathcal{I}$ inductively using Q and then show that $\mathcal{I} \models \mathcal{T}$ and $C_0^{\mathcal{I}} \neq \emptyset$. For this construction, we need to introduce some notation first.

Let $\mathcal{I}$ be an interpretation and let $\mathcal{L} : \Delta^{\mathcal{I}} \to Q$. A pair $\langle s, C\rangle$ with $s \in \Delta^{\mathcal{I}}$ and $C \in \mathsf{cl}(C_0, \mathcal{T})$ is called a *defect of $\mathcal{I}$ w.r.t. $\mathcal{L}$*, if and only if,

- $\exists R.C \in \mathcal{L}(s)$ and there is no $t \in \Delta^{\mathcal{I}}$ such that $\langle s, t\rangle \in R^{\mathcal{I}}$ and $C \in \mathcal{L}(t)$, or
- $\geq 2R \in \mathcal{L}(s)$ and $|\{t \in \Delta^{\mathcal{I}} \mid \langle s, t\rangle \in R^{\mathcal{I}}\}| < 2$.

Fix a map $f : \mathsf{cl}(C_0, \mathcal{T}) \to \mathbb{N}$, and let $\preceq$ be any linear order on the Cartesian product $\mathbb{N} \times \mathbb{N}$ of order type ω (recall that a countably infinite linear order is said to have order type ω if for each element in the order there are only finitely many elements that are less than it; it is well known that there are linear orders on $\mathbb{N} \times \mathbb{N}$ of order type ω).

We are now ready to define by induction the interpretations $\mathcal{I}_i = \langle \Delta^{\mathcal{I}_i}, \cdot^{\mathcal{I}_i}\rangle$ with $\Delta^{\mathcal{I}_i} \subseteq \mathbb{N}$ and mappings $\mathcal{L}_i : \Delta^{\mathcal{I}_i} \to Q$, for $i \in \mathbb{N}$.

Base case. By condition (a) from Definition A.3, there is some type $\tau_0 \in Q$ with $C_0 \in \tau_0$. Define the interpretation $\mathcal{I}_0$ as follows.

- $\Delta^{\mathcal{I}_0} = \{s\}$, for some $s \in \mathbb{N}$;
- for all $A \in N_C$,
  - if $A \in \tau_0$, then $A^{\mathcal{I}_0} = \{s\}$,
  - if $A \notin \tau_0$, then $A^{\mathcal{I}_0} = \emptyset$;
- for all $R \in N_R$, $R^{\mathcal{I}_0} = \emptyset$.

Set $\mathcal{L}_0 = \{s \mapsto \tau_0\}$.

Inductive step. If there is no defect of $\mathcal{I}_i$ w.r.t. $\mathcal{L}_i$, then set $\mathcal{I}_{i+1} = \mathcal{I}_i$ and $\mathcal{L}_{i+1} = \mathcal{L}_i$; otherwise, let $\langle s, C\rangle$ be the least defect of $\mathcal{I}_i$ w.r.t. $\mathcal{L}_i$, i.e., for every defect $\langle t, D\rangle$ of $\mathcal{I}_i$ w.r.t. $\mathcal{L}_i$, we have $\langle s, f(C)\rangle \preceq \langle t, f(D)\rangle$ (using the fact that $\prec$ has order type ω). By $\mathcal{L}_i(s) \in Q$ and conditions (b) and (c) from Definition A.3, there is some $\tau \in Q$ such that $\mathcal{L}_i(s) \xRightarrow{C} \tau$. If $C = \exists R.D$, then let $t \in \mathbb{N} \setminus \Delta^{\mathcal{I}_i}$ and define

- $\Delta^{\mathcal{I}_{i+1}} = \Delta^{\mathcal{I}_i} \cup \{t\}$,
- for all $A \in N_C$,
  - if $A \in \tau$, then $A^{\mathcal{I}_{i+1}} = A^{\mathcal{I}_i} \cup \{t\}$,
  - if $A \notin \tau$, then $A^{\mathcal{I}_{i+1}} = A^{\mathcal{I}_i}$;
- for all $S \in N_R$,
  - if $S = R$, then $S^{\mathcal{I}_{i+1}} = \{\langle s, t\rangle\} \cup S^{\mathcal{I}_i}$,
  - if $S \neq R$, then $S^{\mathcal{I}_{i+1}} = S^{\mathcal{I}_i}$.

Also set $\mathcal{L}_{i+1} = \mathcal{L}_i \cup \{t \mapsto \tau\}$. If $C = \geq 2R$, then let $t_1, t_2 \in \mathbb{N} \setminus \Delta^{\mathcal{I}_i}$ with $t_1 \neq t_2$ and define

- $\Delta^{\mathcal{I}_{i+1}} = \Delta^{\mathcal{I}_i} \cup \{t_1, t_2\}$,
- for all $A \in N_C$,
  - if $A \in \tau$, then $A^{\mathcal{I}_{i+1}} = A^{\mathcal{I}_i} \cup \{t_1, t_2\}$,
  - if $A \notin \tau$, then $A^{\mathcal{I}_{i+1}} = A^{\mathcal{I}_i}$;
- for all $S \in N_R$,
  - if $S = R$, then $S^{\mathcal{I}_{i+1}} = \{\langle s, t_1\rangle, \langle s, t_2\rangle\} \cup S^{\mathcal{I}_i}$,
  - if $S \neq R$, then $S^{\mathcal{I}_{i+1}} = S^{\mathcal{I}_i}$.

Also set $\mathcal{L}_{i+1} = \mathcal{L}_i \cup \{t_1 \mapsto \tau, t_2 \mapsto \tau\}$. This finishes our inductive construction. Now define the interpretation $\mathcal{I}$ as follows:

- $\Delta^{\mathcal{I}} = \bigcup_{i \geq 0} \Delta^{\mathcal{I}_i}$,
- for all $P \in N_C \cup N_R$, $P^{\mathcal{I}} = \bigcup_{i \geq 0} P^{\mathcal{I}_i}$.

Also set $\mathcal{L} = \bigcup_{i \geq 0} \mathcal{L}_i$. Observe that $\mathcal{L}$ is a total mapping from $\Delta^{\mathcal{I}}$ to Q.

Claim A.6. *For all concepts $C \in \mathsf{cl}(C_0, \mathcal{T})$ and all $s \in \Delta^{\mathcal{I}}$, if $C \in \mathcal{L}(s)$ then $s \in C^{\mathcal{I}}$.*

Proof of claim. Let s and C be as stated in the claim. Suppose $C \in \mathcal{L}(s)$. Then by the definition of $\mathcal{L}$, there is some $i \in \mathbb{N}$ such that $C \in \mathcal{L}_i(s)$; let i be the smallest natural number satisfying $C \in \mathcal{L}_i(s)$, i.e., $\mathcal{I}_i$ is the interpretation that we introduced s. By induction on the structure of C, we show that $s \in C^{\mathcal{I}}$. Since for all $\tau \in Q$, we have $\bot \notin \tau$ by $(\mathsf{P}_\bot)$, it follows that $C \neq \bot$. Hence, we consider the remaining cases for C.

- $C = \top$. We have by assumption that $s \in \Delta^{\mathcal{I}}$, i.e., $s \in \top^{\mathcal{I}}$.
- $C = A$, for some $A \in N_C$. Then by the definition of $\mathcal{I}_i$ and $A \in \mathcal{L}_i(s)$, it immediately follows that $s \in A^{\mathcal{I}_i}$. This implies by the definition of $\mathcal{I}$ that $s \in A^{\mathcal{I}}$.
- $C = \neg A$, for some $A \in N_C$. Since $\mathcal{L}_i(s) \in Q$, $\mathcal{L}_i(s)$ satisfies $(\mathsf{P}_\neg)$. Then by $\neg A \in \mathcal{L}_i(s)$, we have $A \notin \mathcal{L}_i(s)$. One can now easily show by induction that for all $k \geq i$, we have $s \notin A^{\mathcal{I}_k}$. This implies by our assumption about i that for all $k \in \mathbb{N}$, $s \notin A^{\mathcal{I}_k}$. Then by the definition of $\mathcal{I}$, we obtain $s \notin A^{\mathcal{I}}$, i.e., $s \in (\neg A)^{\mathcal{I}}$.

- $C = C_1 \sqcap C_2$. Follows easily by the inductive hypothesis and $(\mathsf{P}_\sqcap)$.
- $C = C_1 \sqcup C_2$. Follows easily by the inductive hypothesis and $(\mathsf{P}_\sqcup)$.
- $C = \forall R.D$. Let $t \in \Delta^\mathcal{I}$ such that $\langle s, t\rangle \in R^\mathcal{I}$. We need to show that $t \in D^\mathcal{I}$. By $\langle s, t\rangle \in R^\mathcal{I}$, there is some $k \in \mathbb{N}$ such that $\langle s, t\rangle \in R^{\mathcal{I}_k}$; w.l.o.g. assume that $\mathcal{I}_k$ is the interpretation that we introduced t. It follows that there is some $E \in \mathsf{cl}(C_0, \mathcal{T})$ such that $\langle s, E\rangle$ is a defect of $\mathcal{I}_{k-1}$ w.r.t. $\mathcal{L}_{k-1}$ and $\mathcal{L}_k(s) \overset{E}{\Rightarrow} \mathcal{L}_k(t)$. This implies $D \in \mathcal{L}_k(t)$. Then by the definition of $\mathcal{L}$, we obtain $D \in \mathcal{L}(t)$. By the inductive hypothesis, this implies $t \in D^\mathcal{I}$. Hence, $s \in (\forall R.D)^\mathcal{I}$.
- $C = \exists R.D$. By our assumption about i, we have that $\langle s, C\rangle$ is a defect of $\mathcal{I}_i$ w.r.t. $\mathcal{L}_i$. By the definition of $\preceq$, there are finitely many pairs $\langle t, E\rangle$ with $t \in \mathbb{N}$ and $E \in \mathsf{cl}(C_0, \mathcal{T})$ such that $\langle t, f(E)\rangle \preceq \langle s, f(C)\rangle$. This implies that there is some $k > i$ such that we 'fix' the defect $\langle s, C\rangle$ at step k. Then there is some $t \in \Delta^{\mathcal{I}_k}$ such that $\langle s, t\rangle \in R^{\mathcal{I}_k}$ and $D \in \mathcal{L}_k(t)$. By the definition of $\mathcal{I}$, we then have $\langle s, t\rangle \in R^\mathcal{I}$ and $D \in \mathcal{L}(t)$. By the inductive hypothesis, the latter implies $t \in D^\mathcal{I}$. Hence, $s \in (\exists R.D)^\mathcal{I}$.
- $C = \leq 1R$. Suppose for a contradiction that there are $t_1, t_2 \in \Delta^\mathcal{I}$ such that $t_1 \neq t_2$ and $\langle s, t_1\rangle, \langle s, t_2\rangle \in R^\mathcal{I}$. Then there are $k_1, k_2 \in \mathbb{N}$ such that $\langle s, t_i\rangle \in R^{\mathcal{I}_{k_i}}$ and $\mathcal{I}_{k_i}$ is the interpretation that we introduced t_i, for each $i \in \{1, 2\}$. By our construction, this implies that there are concepts $C_1, C_2 \in \mathsf{cl}(C_0, \mathcal{T})$ such that $\langle s, C_i\rangle$ is a defect of $\mathcal{I}_{k_i-1}$ w.r.t. $\mathcal{L}_{k_i-1}$ and $\mathcal{L}_{k_i}(s) \overset{C_i}{\Longrightarrow} \mathcal{L}_{k_i}(t_i)$, for each $i \in \{1, 2\}$. It follows that $C_i \neq \geq 2R$, for each $i \in \{1, 2\}$; otherwise, we would obtain a contradiction by $(\mathsf{P}_\bowtie)$. Thus, $C_1 = \exists R.D_1$ and $C_2 = \exists R.D_2$. Then by the definition of $\mathsf{cl}(C_0, \mathcal{T})$, we have $\forall R.D_1, \forall R.D_2 \in \mathsf{cl}(C_0, \mathcal{T})$; and by $(\mathsf{P}_{\leq 1})$, this implies $\forall R.D_1, \forall R.D_2 \in \mathcal{L}_{k_1}(s) = \mathcal{L}_{k_2}(s)$. Suppose w.l.o.g. that $k_1 < k_2$. Then $D_2 \in \mathcal{L}_{k_1}(t_1)$. But this contradicts with the fact that $\langle s, \exists R.D_2\rangle$ is a defect of $\mathcal{I}_{k_2-1}$ w.r.t $\mathcal{L}_{k_2-1}$.
- $C = \geq 2R$. This case can be shown similarly to the case $C = \exists R.D$.

Since we considered all the possible cases, we conclude that the claim holds. $\dashv$

Using Claim A.6, this direction of the Theorem can now be shown easily as follows. By the base case of our inductive construction, there is some $s \in \Delta^{\mathcal{I}_0}$ such that $C_0 \in \mathcal{L}_0(s)$. This implies $C_0 \in \mathcal{L}(s)$ and then by Claim A.6, we obtain $s \in C_0^\mathcal{I}$. Moreover, by Claim A.6 and $(\mathsf{P}_\sqsubseteq)$, we have $\mathcal{I} \models \mathcal{T}$. Hence C_0 is satisfiable w.r.t. $\mathcal{T}$. $\square$

Appendix B. Useful Lemmas for Tableau Correctness and Interpolation

For all $\Phi \subseteq \mathsf{cll} \cup \mathsf{clr}$, we define

$$\Phi(\mathbf{l}) = \{C \mid C^\mathbf{l} \in \Phi \cap \mathsf{cll}\} \text{ and } \Phi(\mathbf{r}) = \{C \mid C^\mathbf{r} \in \Phi \cap \mathsf{clr}\}.$$

$\Phi(\lambda)$ is a shorthand for $\Phi(\mathbf{l}) \cup \Phi(\mathbf{r})$. In the following the *signature of a set of $\mathcal{ALCF}$-concepts S* will be of concern. We define $\mathsf{sig}(S) = \bigcup_{C \in S} \mathsf{sig}(C)$. Let τ be a finite set of $\mathcal{ALCF}$-concepts and $\mathcal{T}$ be an $\mathcal{ALCF}$-TBox. We say that τ *is satisfiable w.r.t.* $\mathcal{T}$ if and only if $\bigsqcap_{D \in \tau} D$ is satisfiable w.r.t. $\mathcal{T}$. Moreover a $\Phi \subseteq \mathsf{cll} \cup \mathsf{clr}$ *is satisfiable w.r.t.* $\mathcal{T}$ if and only if $\Phi(\lambda)$ is satisfiable w.r.t. $\mathcal{T}$.

Lemma B.1. *Let $\Phi \subseteq \mathsf{cll} \cup \mathsf{clr}$ be satisfiable w.r.t. $\mathcal{T}$. We have*

- *if χ is an $\star$-burden of Φ for $\star \in \{\sqcap, \leq 1, \exists, \geq 2\}$ and Ψ is the χ-relief of Φ, then Ψ is satisfiable w.r.t. $\mathcal{T}$;*
- *if χ is an $\sqcup$-burden of Φ, then there is some χ-relief Ψ of Φ such that Ψ is satisfiable w.r.t. $\mathcal{T}$.*

Proof. Suppose that Φ is as stated in the Theorem, i.e., it is satisfiable w.r.t. $\mathcal{T}$. This means $\Phi(\lambda)$ is satisfiable w.r.t. $\mathcal{T}$. By Theorem A.4 we then have that there is some $\langle C, \mathcal{T}\rangle$-quasimodel Q for $\mathcal{ALCF}$, where $C = \bigsqcap_{D \in \Phi(\lambda)} D$. This means there is some $\tau \in Q$ such that $\Phi(\lambda) \subseteq \tau$. We will also use the term $\langle \tau, \mathcal{T}\rangle$*-quasimodel* for Q.

Assume that $(C_1 \sqcap C_2)^\lambda$ is an $\sqcap$-burden of Φ. Then the $(C_1 \sqcap C_2)^\lambda$-relief of Φ is $\Psi = \Phi \cup \{(C_1)^\lambda, (C_2)^\lambda\}$. By $(\mathsf{P}_\sqcap)$, $\{C_1, C_2\} \subseteq \tau$ and thus $\Psi(\lambda) \subseteq \tau$. Hence $\Psi(\lambda)$ is satisfiable w.r.t. $\mathcal{T}$.

Assume that $(\leq 1R)^\lambda$ is an ≤ 1-burden of Φ. Then the $(\leq 1R)^\lambda$-relief of Φ is $\Psi = \Phi \cup \{(\forall R.C)^\kappa \mid (\exists R.C)^\kappa \in \Phi\}$. If $(\exists R.C)^\kappa \in \Phi$, then $\exists R.C \in \tau$ and by $(\mathsf{P}_{\leq 1})$, $\forall R.C \in \tau$. Hence $\Psi(\lambda) \subseteq \tau$ and Ψ is satisfiable w.r.t. $\mathcal{T}$.

Assume that $(\exists R.C)^\lambda$ is an $\exists$-burden of Φ. $\Phi(\lambda) \subseteq \tau$ so by condition (b) of Definition A.3, there is some $\upsilon \in Q$ such that $\upsilon \supseteq \{C\} \cup \{D \mid \forall R.D \in \Phi(\lambda)\}$; and by $(\mathsf{P}_\sqsubseteq)$, $\{E \mid \top \sqsubseteq E \in \mathcal{T}\} \subseteq \upsilon$. Let Ψ be the $(\exists R.C)^\lambda$-relief of Φ. Then we have $\Psi(\lambda) \subseteq \upsilon$. Hence Ψ is satisfiable w.r.t. $\mathcal{T}$.

Assume that $(\geq 2R)^\lambda$ is an ≥ 2-burden of Φ. $\Phi(\lambda) \subseteq \tau$ so by condition (b) of Definition A.3, there is some $\upsilon \in Q$ such that $\upsilon \supseteq \{D \mid \forall R.D \in \Phi(\lambda)\}$; and by $(\mathsf{P}_\sqsubseteq)$, $\{E \mid \top \sqsubseteq E \in \mathcal{T}\} \subseteq \upsilon$. Let Ψ be the $\geq 2R$-relief of Φ. Then we have $\Psi(\lambda) \subseteq \upsilon$. Hence Ψ is satisfiable w.r.t. $\mathcal{T}$.

Assume that $(C_1 \sqcup C_2)^\lambda$ is an $\sqcup$-burden Φ. Then for some $(C_1 \sqcup C_2)^\lambda$-relief Ψ of Φ, we have $\Psi(\lambda) \subseteq \tau$ by $(\mathsf{P}_\sqcup)$. Hence, there is some $(C_1 \sqcup C_2)^\lambda$-relief of Φ that is satisfiable w.r.t. $\mathcal{T}$. ☐

Proposition B.2. *Let $\mathcal{T}$ be an $\mathcal{ALCF}$-TBox and $C_0, C_1, \ldots, C_n, D$ be $\mathcal{ALCF}$-concepts.*

1. *If $\mathcal{T} \models C_0 \sqcap C_1 \sqcap \ldots \sqcap C_n \sqsubseteq D$, then*
$$\mathcal{T} \models \exists R.C_0 \sqcap \forall R.C_1 \sqcap \ldots \sqcap \forall R.C_n \sqsubseteq \exists R.D.$$
2. *If $\mathcal{T} \models D \sqsubseteq C_1 \sqcup \ldots \sqcup C_n$, then*
$$\mathcal{T} \models \exists R.D \sqsubseteq \exists R.C_1 \sqcup \ldots \sqcup \exists R.C_n.$$
3. *If $\mathcal{T} \models C_1 \sqcap \ldots \sqcap C_n \sqsubseteq D$, then*
$$\mathcal{T} \models \geq 2R \sqcap \forall R.C_1 \sqcap \ldots \sqcap \forall R.C_n \sqsubseteq \exists R.D.$$

Proof. For 1, we proceed towards a contradiction. Suppose $\mathcal{T} \models C_0 \sqcap C_1 \sqcap \ldots \sqcap C_n \sqsubseteq D$ and $\mathcal{T} \not\models \exists R.C_0 \sqcap \forall R.C_1 \sqcap \ldots \sqcap \forall R.C_n \sqsubseteq \exists R.D$. Then there is some model $\mathcal{I}$ of $\mathcal{T}$ such that $\mathcal{I} \models C_0 \sqcap C_1 \sqcap \ldots \sqcap C_n \sqsubseteq D$ and $\mathcal{I} \not\models \exists R.C_0 \sqcap \forall R.C_1 \sqcap \ldots \sqcap \forall R.C_n \sqsubseteq \exists R.D$. By the latter there is some $s \in \Delta^{\mathcal{I}}$ such that $s \in (\exists R.C_0 \sqcap \forall R.C_1 \sqcap \ldots \sqcap \forall R.C_n)^{\mathcal{I}}$ and $s \notin (\exists R.D)^{\mathcal{I}}$. That is, there is some $t \in \Delta^{\mathcal{I}}$ such that $\langle s, t\rangle \in R^{\mathcal{I}}$, $t \in (C_0 \sqcap C_1 \sqcap \ldots \sqcap C_n)^{\mathcal{I}}$, and $t \in (\neg D)^{\mathcal{I}}$. But this contradicts with $\mathcal{I} \models C_0 \sqcap C_1 \sqcap \ldots \sqcap C_n \sqsubseteq D$.

For 2, we proceed towards a contradiction. Suppose $\mathcal{T} \models D \sqsubseteq C_1 \sqcup \ldots \sqcup C_n$ and $\mathcal{T} \not\models \exists R.D \sqsubseteq \exists R.C_1 \sqcup \ldots \sqcup \exists R.C_n$. Then there is some model $\mathcal{I}$ of $\mathcal{T}$ such that $\mathcal{I} \models D \sqsubseteq C_1 \sqcup \ldots \sqcup C_n$ and $\mathcal{I} \not\models \exists R.D \sqsubseteq \exists R.C_1 \sqcup \ldots \sqcup \exists R.C_n$. By the latter there is some $s \in \Delta^{\mathcal{I}}$ such that $s \in (\exists R.D)^{\mathcal{I}}$ and $s \notin (\exists R.C_1 \sqcup \ldots \sqcup \exists R.C_n)^{\mathcal{I}}$. That is there is some $t \in \Delta^{\mathcal{I}}$ such that $\langle s, t\rangle \in R^{\mathcal{I}}$, $t \in D^{\mathcal{I}}$, and $t \in (\neg C_1 \sqcap \ldots \sqcap \neg C_n)^{\mathcal{I}}$. But this contradicts with $\mathcal{I} \models D \sqsubseteq C_1 \sqcup \ldots \sqcup C_n$. □

Proposition B.3. *Let C be an $\mathcal{ALCF}$-concept and R be a role name. Then*

$$\models \leq 1R \sqcap \exists R.C \sqcap \forall R.C \equiv \leq 1R \sqcap \exists R.C.$$

Proof. That $\models \leq 1R \sqcap \exists R.C \sqcap \forall R.C \sqsubseteq \leq 1R \sqcap \exists R.C$ is trivial. For the other direction, suppose for a contradiction that $\not\models \leq 1R \sqcap \exists R.C \sqsubseteq \leq 1R \sqcap \exists R.C \sqcap \forall R.C$. This means there is some interpretation $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}}\rangle$ such that $\mathcal{I} \not\models \leq 1R \sqcap \exists R.C \sqsubseteq \leq 1R \sqcap \exists R.C \sqcap \forall R.C$. Thus there is some $s \in \Delta^{\mathcal{I}}$ such that $s \in (\leq 1R)^{\mathcal{I}}$, $s \in (\exists R.C)^{\mathcal{I}}$, and $s \in (\exists R.\neg C)^{\mathcal{I}}$. By the last two there are $t_1, t_2 \in \Delta^{\mathcal{I}}$ such that $\langle s, t_1\rangle, \langle s, t_2\rangle \in R^{\mathcal{I}}$, $t_1 \in C^{\mathcal{I}}$, and $t_2 \in (\neg C)^{\mathcal{I}}$. But by $s \in (\leq 1R)^{\mathcal{I}}$, $t_1 = t_2$ which is a contradiction. □

Lemma B.4. *Let $\Phi \subseteq \mathsf{cll} \cup \mathsf{clr}$. We have*

1. *if $\perp^{\mathbf{l}} \in \Phi$, then $\perp$ is an interpolant of Φ;*
2. *if $\perp^{\mathbf{r}} \in \Phi$, then $\top$ is an interpolant of Φ;*
3. *for a concept C of the form A or $\leq 1R$,*
   (a) *if $\{C^{\mathbf{l}}, (\dot\neg C)^{\mathbf{l}}\} \subseteq \Phi$, then $\perp$ is an interpolant of Φ;*
   (b) *if $\{C^{\mathbf{r}}, (\dot\neg C)^{\mathbf{r}}\} \subseteq \Phi$, then $\top$ is an interpolant of Φ;*
   (c) *if $\{C^{\mathbf{l}}, (\dot\neg C)^{\mathbf{r}}\} \subseteq \Phi$, then C is an interpolant of Φ;*
   (d) *if $\{C^{\mathbf{r}}, (\dot\neg C)^{\mathbf{l}}\} \subseteq \Phi$, then $\dot\neg C$ is an interpolant of Φ;*
4. *if Ψ is the $(C_1 \sqcap C_2)^{\lambda}$-relief of Φ and I is an interpolant of Ψ, then I is an interpolant of Φ;*
5. *if Ψ_1 and Ψ_2 are $(C_1 \sqcup C_2)^{\mathbf{l}}$-reliefs of Φ, and I_1, I_2 are interpolants of Ψ_1, Ψ_2 respectively, then $I_1 \sqcup I_2$ is an interpolant of Φ;*
6. *if Ψ_1 and Ψ_2 are $(C_1 \sqcup C_2)^{\mathbf{r}}$-reliefs of Φ, and I_1, I_2 are interpolants of Ψ_1, Ψ_2 respectively, then $I_1 \sqcap I_2$ is an interpolant of Φ;*
7. *if Ψ is the $(\leq 1R)^{\mathbf{l}}$-relief of Φ, there is no biased concept of the form $(\exists R.C)^{\mathbf{r}} \in \Phi$, and I is an interpolant of Ψ, then I is an interpolant of Φ;*
8. *if Ψ is the $(\leq 1R)^{\mathbf{r}}$-relief of Φ, there is no biased concept of the form $(\exists R.C)^{\mathbf{l}} \in \Phi$, and I is an interpolant of Ψ, then I is an interpolant of Φ;*

9. *if Ψ is the $(\leq 1R)^{\mathbf{l}}$-relief of Φ, there is some biased concept of the form $(\exists R.C)^{\mathbf{r}} \in \Phi$, and I is an interpolant of Ψ, then $I \sqcap \leq 1R$ is an interpolant of Φ;*

10. *if Ψ is the $(\leq 1R)^{\mathbf{r}}$-relief of Φ, there is some biased concept of the form $(\exists R.C)^{\mathbf{l}} \in \Phi$, and I is an interpolant of Ψ, then $I \sqcup \geq 2R$ is an interpolant of Φ;*

11. *if Ψ is the $(\exists R.C)^{\mathbf{l}}$-relief of Φ, I is an interpolant of Ψ, and there is no biased concept of the form $(\forall R.D)^{\mathbf{r}} \in \Phi$, then $\perp$ is an interpolant of Φ;*

12. *if Ψ is the $(\exists R.C)^{\mathbf{r}}$-relief of Φ, I is an interpolant of Ψ, and there is no biased concept of the form $(\forall R.D)^{\mathbf{l}} \in \Phi$, then $\top$ is an interpolant of Φ;*

13. *if Ψ is the $(\exists R.C)^{\mathbf{l}}$-relief of Φ, I is an interpolant of Ψ, and there is some biased concept of the form $(\forall R.D)^{\mathbf{r}} \in \Phi$, then $\exists R.I$ is an interpolant of Φ;*

14. *if Ψ is the $(\exists R.C)^{\mathbf{r}}$-relief of Φ, I is an interpolant of Ψ, and there is some biased concept of the form $(\forall R.D)^{\mathbf{l}} \in \Phi$, then $\forall R.I$ is an interpolant of Φ;*

15. *if Ψ is the $(\geq 2R)^{\mathbf{l}}$-relief of Φ, I is an interpolant of Ψ, and there is no biased concept of the form $(\forall R.D)^{\mathbf{r}} \in \Phi$, then $\perp$ is an interpolant of Φ;*

16. *if Ψ is the $(\geq 2R)^{\mathbf{r}}$-relief of Φ, I is an interpolant of Ψ, and there is no biased concept of the form $(\forall R.D)^{\mathbf{l}} \in \Phi$, then $\top$ is an interpolant of Φ;*

17. *if Ψ is the $(\geq 2R)^{\mathbf{l}}$-relief of Φ, I is an interpolant of Ψ, and there is some biased concept of the form $(\forall R.D)^{\mathbf{r}} \in \Phi$, then $\exists R.I$ is an interpolant of Φ;*

18. *if Ψ is the $(\geq 2R)^{\mathbf{r}}$-relief of Φ, I is an interpolant of Ψ, and there is some biased concept of the form $(\forall R.D)^{\mathbf{l}} \in \Phi$, then $\forall R.I$ is an interpolant of Φ.*

Proof. For 1. Suppose $\Phi(\mathbf{l}) = \{X_1, \ldots, X_n\} \cup \{\perp\}$ and $\Phi(\mathbf{r}) = \{Y_1, \ldots, Y_m\}$. But $\mathcal{T} \models \perp \sqcap X_1 \sqcap \ldots, X_n \sqsubseteq \perp$ and $\mathcal{T} \models \perp \sqsubseteq \neg Y_1 \sqcup \ldots \sqcup \neg Y_m$ hold trivially. Since $\perp$ is a logical constant, $\emptyset = \mathsf{sig}(\perp) \subseteq \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r}))$. Hence 1 is satisfied.

For 2. Suppose $\Phi(\mathbf{r}) = \{Y_1, \ldots, Y_m\} \cup \{\perp\}$ and $\Phi(\mathbf{l}) = \{X_1, \ldots, X_n\}$. But $\mathcal{T} \models X_1 \sqcap \ldots, X_n \sqsubseteq \top$ and $\mathcal{T} \models \top \sqsubseteq \neg Y_1 \sqcup \ldots \sqcup \neg Y_m \sqcup \top$ hold trivially. Since $\top$ is a logical constant, $\emptyset = \mathsf{sig}(\top) \subseteq \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r}))$. Hence 2 is satisfied.

For 3a. Suppose $\Phi(\mathbf{l}) = \{X_1, \ldots, X_n\} \cup \{C, \dot{\neg} C\}$ and $\Phi(\mathbf{r}) = \{Y_1, \ldots, Y_m\}$. But $\mathcal{T} \models X_1 \sqcap \ldots, X_n \sqcap C \sqcap \dot{\neg} C \sqsubseteq \perp$ and $\mathcal{T} \models \perp \sqsubseteq \neg Y_1 \sqcup \ldots \sqcup \neg Y_m$ hold trivially. Since $\perp$ is a logical constant, $\emptyset = \mathsf{sig}(\perp) \subseteq \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r}))$. Hence 3a is satisfied.

The argument for 3b is analogous to the previous case.

For 3c. Suppose $\Phi(\mathbf{l}) = \{X_1, \ldots, X_n\} \cup \{C\}$ and $\Phi(\mathbf{r}) = \{Y_1, \ldots, Y_m\} \cup \{\dot{\neg} C\}$. But $\mathcal{T} \models X_1 \sqcap \ldots \sqcap X_n \sqcap C \sqsubseteq C$, $\mathcal{T} \models C \sqsubseteq \neg Y_1 \sqcup \ldots \sqcup \neg Y_m \sqcup \neg(\dot{\neg} C)$, and $\mathsf{sig}(C) \subseteq \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r}))$ hold trivially. Hence 3c is satisfied.

The argument for 3d is analogous to the previous case.

For 4. Suppose Ψ is a $(C_1 \sqcap C_2)^{\mathbf{l}}$-relief of Φ, I is an interpolant of Ψ, $\Phi(\mathbf{l}) = \{X_1, \ldots, X_n\} \cup \{C_1 \sqcap C_2\}$, and $\Phi(\mathbf{r}) = \{Y_1, \ldots, Y_m\}$. By assumption, $\mathcal{T} \models X_1 \sqcap \ldots \sqcap X_n \sqcap (C_1 \sqcap C_2) \sqcap C_1 \sqcap C_2 \sqsubseteq I$, i.e., $\mathcal{T} \models X_1 \sqcap \ldots \sqcap X_n \sqcap (C_1 \sqcap C_2) \sqsubseteq I$ and $\mathcal{T} \models I \sqsubseteq \neg Y_1 \sqcup \ldots \sqcup \neg Y_m$. By assumption again, $\mathsf{sig}(\Psi(\mathbf{l})) = \mathsf{sig}(\Phi(\mathbf{l}))$ and $\mathsf{sig}(\Psi(\mathbf{r})) = \mathsf{sig}(\Phi(\mathbf{r}))$, and thus $\mathsf{sig}(I) \subseteq \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r}))$.

Therefore I is an interpolant of Φ. The case for when Ψ is a $(C_1 \sqcap C_2)^{\mathbf{r}}$-relief of Φ and I is an interpolant of Ψ can be shown analogously. Hence 4 is satisfied.

For 5. Suppose Ψ_1 and Ψ_2 are $(C_1 \sqcup C_2)^{\mathbf{l}}$-reliefs of Φ, I_1, I_2 are interpolants of Ψ_1, Ψ_2 respectively, $\Phi(\mathbf{l}) = \{X_1, \ldots, X_n\} \cup \{C_1 \sqcup C_2\}$, and $\Phi(\mathbf{r}) = \{Y_1, \ldots, Y_m\}$. By assumption, $\mathcal{T} \models X_1 \sqcap \ldots \sqcap X_n \sqcap (C_1 \sqcup C_2) \sqcap C_1 \sqsubseteq I_1$ and $\mathcal{T} \models X_1 \sqcap \ldots \sqcap X_n \sqcap (C_1 \sqcup C_2) \sqcap C_2 \sqsubseteq I_2$. Then we have the following.

$$\begin{aligned}
\mathcal{T} \models I_1 \sqcup I_2 &\sqsupseteq (X_1 \sqcap \ldots \sqcap X_n \sqcap (C_1 \sqcup C_2) \sqcap C_1) \sqcup \\
&\quad\; (X_1 \sqcap \ldots \sqcap X_n \sqcap (C_1 \sqcup C_2) \sqcap C_2) \\
\mathcal{T} \models I_1 \sqcup I_2 &\sqsupseteq (X_1 \sqcap \ldots \sqcap X_n \sqcap (C_1 \sqcup C_2)) \sqcap (C_1 \sqcup C_2) \\
\mathcal{T} \models I_1 \sqcup I_2 &\sqsupseteq X_1 \sqcap \ldots \sqcap X_n \sqcap (C_1 \sqcup C_2)
\end{aligned}$$

For the other half, by assumption $\mathcal{T} \models I_1 \sqsubseteq \neg Y_1 \sqcup \ldots \sqcup \neg Y_m$ and $\mathcal{T} \models I_2 \sqsubseteq \neg Y_1 \sqcup \ldots \sqcup \neg Y_m$. But then $\mathcal{T} \models I_1 \sqcup I_2 \sqsubseteq \neg Y_1 \sqcup \ldots \sqcup \neg Y_m$. Clearly, $\mathsf{sig}(I_1 \sqcup I_2) \subseteq \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r}))$. Hence 5 is satisfied.

The argument for 6 is analogous to the previous case.

For 7. Suppose

- Ψ is a $(\leq 1R)^{\mathbf{l}}$-relief of Φ,
- I is an interpolant of Ψ,
- $\Phi(\mathbf{l}) = \{X_1, \ldots, X_n\} \cup \{\leq 1R\} \cup \{\exists R.C_1, \ldots, \exists R.C_k\}$, where $\{\exists R.C_1, \ldots, \exists R.C_k\} = \{\exists R.C \in \Phi(\mathbf{l})\}$,
- $\Phi(\mathbf{r}) = \{Y_1, \ldots, Y_m\}$,
- there is no biased concept of the form $(\exists R.C)^{\mathbf{r}} \in \Phi$.

Let $E = \forall R.C_1 \sqcap \ldots \sqcap \forall R.C_k$. By assumption,

$$\mathcal{T} \models X_1 \sqcap \ldots \sqcap X_n \sqcap \leq 1R \sqcap \exists R.C_1 \sqcap \ldots \sqcap \exists R.C_k \sqcap E \sqsubseteq I.$$

Then by Proposition B.3

$$\mathcal{T} \models X_1 \sqcap \ldots \sqcap X_n \sqcap \leq 1R \sqcap \exists R.C_1 \sqcap \ldots \sqcap \exists R.C_k \sqsubseteq I$$

which is what we wanted to show. For the other half, since there is no biased concept of the form $(\exists R.C)^{\mathbf{r}} \in \Phi$, we have $\Phi(\mathbf{r}) = \Psi(\mathbf{r})$. But then

$$\mathcal{T} \models I \sqsubseteq \neg Y_1 \sqcup \ldots \sqcup \neg Y_m$$

which is what we wanted to show. By assumption $\mathsf{sig}(\Psi(\mathbf{l})) = \mathsf{sig}(\Phi(\mathbf{l}))$ and $\mathsf{sig}(\Psi(\mathbf{r})) = \mathsf{sig}(\Phi(\mathbf{r}))$, and thus $\mathsf{sig}(I) \subseteq \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r}))$. Therefore I is an interpolant of Φ. Hence 7 is satisfied.

8 can be shown analogously to the previous case.

For 9. Suppose

- Ψ is a $(\leq 1R)^{\mathbf{l}}$-relief of Φ,

- I is an interpolant of Ψ,
- $\Phi(\mathbf{l}) = \{X_1, \dots, X_n\} \cup \{\leq 1R\} \cup \{\exists R.C_1, \dots, \exists R.C_k\}$, where $\{\exists R.C_1, \dots, \exists R.C_k\} = \{\exists R.C \in \Phi(\mathbf{l})\}$,
- $\Phi(\mathbf{r}) = \{Y_1, \dots, Y_m\} \cup \{\exists R.D_1, \dots, \exists R.D_l\}$, where $\{\exists R.D_1, \dots, \exists R.D_l\} = \{\exists R.C \in \Phi(\mathbf{r})\}$,
- there is some biased concept of the form $(\exists R.C)^{\mathbf{r}} \in \Phi$.

Let $E = \forall R.C_1 \sqcap \dots \sqcap \forall R.C_k$. By assumption,

$$\mathcal{T} \models X_1 \sqcap \dots \sqcap X_n \sqcap \leq 1R \sqcap \exists R.C_1 \sqcap \dots \sqcap \exists R.C_k \sqcap E \sqsubseteq I.$$

Also, we trivially have the following.

$$\mathcal{T} \models X_1 \sqcap \dots \sqcap X_n \sqcap \leq 1R \sqcap \exists R.C_1 \sqcap \dots \sqcap \exists R.C_k \sqcap E \sqsubseteq \leq 1R.$$

Combining these two, we get

$$\mathcal{T} \models X_1 \sqcap \dots \sqcap X_n \sqcap \leq 1R \sqcap \exists R.C_1 \sqcap \dots \sqcap \exists R.C_k \sqcap E \sqsubseteq I \sqcap \leq 1R$$

which is what we wanted to show.

For the other half, let $F = \exists R.\neg C_1 \sqcup \dots \sqcup \exists R.\neg C_l$. By the assumption about I,

$$\mathcal{T} \models I \sqsubseteq \neg Y_1 \sqcup \dots \sqcup \neg Y_m \sqcup \forall R.\neg C_1 \sqcup \dots \sqcup \forall R.\neg C_l \sqcup F.$$

From this, we trivially get

$$\mathcal{T} \models I \sqsubseteq \neg Y_1 \sqcup \dots \sqcup \neg Y_m \sqcup \forall R.\neg C_1 \sqcup \dots \sqcup \forall R.\neg C_l \sqcup F \sqcup \geq 2R.$$

Then by Proposition B.3,

$$\mathcal{T} \models I \sqsubseteq \neg Y_1 \sqcup \dots \sqcup \neg Y_m \sqcup \forall R.\neg C_1 \sqcup \dots \sqcup \forall R.\neg C_l \sqcup \geq 2R,$$

which implies

$$\mathcal{T} \models I \sqcap \leq 1R \sqsubseteq \neg Y_1 \sqcup \dots \sqcup \neg Y_m \sqcup \forall R.\neg C_1 \sqcup \dots \sqcup \forall R.\neg C_l,$$

and this is what we wanted to show. By assumption $\mathsf{sig}(\Psi(\mathbf{l})) = \mathsf{sig}(\Phi(\mathbf{l}))$ and $\mathsf{sig}(\Psi(\mathbf{r})) = \mathsf{sig}(\Phi(\mathbf{r}))$, and thus $\mathsf{sig}(I) \subseteq \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r}))$. Moreover, since there is some biased concept of the form $(\exists R.C)^{\mathbf{r}} \in \Phi$, $R \in \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r}))$. In conclusion, $\mathsf{sig}(I \sqcap \leq 1R) \in \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r}))$. Hence 9 is satisfied.

10 can be shown analogously to the previous case.

For 11. Suppose the following:

- Ψ is the $(\exists R.C)^{\mathbf{l}}$-relief of Φ,
- I is an interpolant of Ψ,
- $E = \bigsqcap_{\top \sqsubseteq C \in \mathcal{T}_{\mathbf{l}}} C$, $F = \bigsqcup_{\top \sqsubseteq C \in \mathcal{T}_{\mathbf{r}}} \neg C$,

- $\Phi(\mathbf{l}) = \{X_1, \ldots, X_n\} \cup \{\exists R.C\} \cup \{\forall R.D_1, \ldots, \forall R.D_k\}$, where $\{\forall R.D_1, \ldots, \forall R.D_k\} = \{\forall R.D \in \Phi(\mathbf{l})\}$,
- there is no biased concept of the form $(\forall R.D)^{\mathbf{r}} \in \Phi$.

By the last assumption, $\Psi(\mathbf{r}) = \{C \mid \top \sqsubseteq C \in \mathcal{T}_{\mathbf{r}}\}$. By assumption, $\mathcal{T} \models I \sqsubseteq F$. Since $\mathcal{T} \models F \sqsubseteq \bot$, we have that $\mathcal{T} \models I \equiv \bot$ and thus $\mathcal{T} \models \exists R.I \equiv \bot$. By assumption again $\mathcal{T} \models C \sqcap D_1 \sqcap \ldots \sqcap D_k \sqcap E \sqsubseteq I$. Since $\mathcal{T} \models \top \sqsubseteq E$, we have that $\mathcal{T} \models C \sqcap D_1 \sqcap \ldots \sqcap D_k \sqsubseteq I$. By Proposition B.2

$$\mathcal{T} \models \exists R.C \sqcap \forall R.D_1 \sqcap \ldots \sqcap \forall R.D_k \sqsubseteq \exists R.I.$$

However by $\mathcal{T} \models \exists R.I \equiv \bot$ this means

$$\mathcal{T} \models \exists R.C \sqcap \forall R.D_1 \sqcap \ldots \sqcap \forall R.D_k \sqsubseteq \bot.$$

Hence,

$$\mathcal{T} \models X_1 \sqcap \ldots \sqcap X_n \sqcap \exists R.C \sqcap \forall R.D_1 \sqcap \ldots \sqcap \forall R.D_k \sqsubseteq \bot$$

which is what we wanted to show. For the other half, let $\Phi(\mathbf{r}) = \{Y_1, \ldots, Y_m\}$. Since $\mathcal{T} \models I \equiv \bot$, we have trivially

$$\mathcal{T} \models \bot \sqsubseteq \neg Y_1 \sqcup \ldots \sqcup \neg Y_m$$

As the final step, we need to show that

$$\mathsf{sig}(\bot) \subseteq \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r})).$$

But this follows easily since $\bot$ is a logical constant. Hence 11 is satisfied.

For 12. Suppose the following:

- Ψ is the $(\exists R.C)^{\mathbf{r}}$-relief of Φ,
- I is an interpolant of Ψ,
- $E = \bigsqcap_{\top \sqsubseteq C \in \mathcal{T}_{\mathbf{l}}} C$, $F = \bigsqcup_{\top \sqsubseteq C \in \mathcal{T}_{\mathbf{r}}} \neg C$,
- $\Phi(\mathbf{r}) = \{Y_1, \ldots, Y_m\} \cup \{\exists R.C\} \cup \{\forall R.D_1, \ldots, \forall R.D_k\}$, where $\{\forall R.D_1, \ldots, \forall R.D_k\} = \{\forall R.D \in \Phi(\mathbf{r})\}$,
- there is no biased concept of the form $(\forall R.D)^{\mathbf{l}} \in \Phi$.

By the last assumption, $\Psi(\mathbf{l}) = \{C \mid \top \sqsubseteq C \in \mathcal{T}_{\mathbf{l}}\}$. By assumption, $\mathcal{T} \models E \sqsubseteq I$. Since $\mathcal{T} \models \top \sqsubseteq E$, we have that $\mathcal{T} \models I \equiv \top$ and thus $\mathcal{T} \models \forall R.I \equiv \top$. By assumption again $\mathcal{T} \models I \sqsubseteq \neg C \sqcup \neg D_1 \sqcup \ldots \sqcup \neg D_k \sqcup F$. Since $\mathcal{T} \models F \sqsubseteq \bot$, we have that $\mathcal{T} \models I \sqsubseteq \neg C \sqcup \neg D_1 \sqcup \ldots \sqcup \neg D_k$. By Proposition B.2

$$\mathcal{T} \models \forall R.I \sqsubseteq \forall R.\neg C \sqcup \exists R.\neg D_1 \sqcup \ldots \sqcup \exists R.\neg D_k.$$

However by $\mathcal{T} \models \top \equiv \forall R.I$, this means

$$\mathcal{T} \models \top \sqsubseteq \forall R.\neg C \sqcup \exists R.\neg D_1 \sqcup \ldots \sqcup \exists R.\neg D_k.$$

Hence,

$$\mathcal{T} \models \top \sqsubseteq \neg Y_1 \sqcup \ldots \sqcup \neg Y_m \sqcup \forall R.\neg C \sqcup \exists R.\neg D_1 \sqcup \ldots \sqcup \exists R.\neg D_k$$

which is what we wanted to show. For the other half, let $\Phi(\mathbf{l}) = \{X_1, \ldots, X_n\}$. Since $\mathcal{T} \models I \equiv \top$, we have trivially

$$\mathcal{T} \models X_1 \sqcap \ldots \sqcap X_n \sqsubseteq \top$$

As the final step, we need to show that

$$\mathsf{sig}(\top) \subseteq \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r})).$$

But this follows easily since $\top$ is a logical constant. Hence 12 is satisfied.

For 13. Suppose the following:

- Ψ is the $(\exists R.C)^{\mathbf{l}}$-relief of Φ,
- I is an interpolant of Ψ,
- $E = \bigsqcap_{\top \sqsubseteq C \in \mathcal{T}_{\mathbf{l}}} C$, $F = \bigsqcup_{\top \sqsubseteq C \in \mathcal{T}_{\mathbf{r}}} \neg C$,
- $\Phi(\mathbf{l}) = \{X_1, \ldots, X_n\} \cup \{\exists R.C\} \cup \{\forall R.D_1, \ldots, \forall R.D_k\}$, where $\{\forall R.D_1, \ldots, \forall R.D_k\} = \{\forall R.D \in \Phi(\mathbf{l})\}$,
- $\Phi(\mathbf{r}) = \{Y_1, \ldots, Y_m\} \cup \{\forall R.C_1, \ldots, \forall R.C_l\}$, where $l \geq 1$ and $\{\forall R.C_1, \ldots, \forall R.C_l\} = \{\forall R.C \in \Phi(\mathbf{r})\}$.

By assumption, $\mathcal{T} \models C \sqcap D_1 \sqcap \ldots \sqcap D_k \sqcap E \sqsubseteq I$. Since $\mathcal{T} \models \top \sqsubseteq E$, we have that $\mathcal{T} \models C \sqcap D_1 \sqcap \ldots \sqcap D_k \sqsubseteq I$. By Proposition B.2

$$\mathcal{T} \models \exists R.C \sqcap \forall R.D_1 \sqcap \ldots \sqcap \forall R.D_k \sqsubseteq \exists R.I \tag{5}$$

Now by (5), we have

$$\mathcal{T} \models X_1 \sqcap \ldots \sqcap X_n \sqcap \exists R.C \sqcap \forall R.D_1 \sqcap \ldots \sqcap \forall R.D_k \sqsubseteq \exists R.I$$

which is what we wanted to show. Now we argue for the other half. By assumption, $\mathcal{T} \models I \sqsubseteq \neg C_1 \sqcup \ldots \sqcup \neg C_l \sqcup F$. Since $\mathcal{T} \models F \sqsubseteq \bot$, we have that $\mathcal{T} \models I \sqsubseteq \neg C_1 \sqcup \ldots \sqcup \neg C_l$. By Proposition B.2

$$\mathcal{T} \models \exists R.I \sqsubseteq \neg \forall R.C_1 \sqcup \ldots \sqcup \neg \forall R.C_l \tag{6}$$

Now by (6), we have

$$\mathcal{T} \models \exists R.I \sqsubseteq \neg Y_1 \sqcup \ldots \sqcup \neg Y_m \sqcup \neg \forall R.C_1 \sqcup \ldots \sqcup \neg \forall R.C_l$$

which is what we wanted to show. As the final step, we need to show that

$$\mathsf{sig}(\exists R.I) \subseteq \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r})).$$

But this follows easily since by assumption $\mathsf{sig}(I) \subseteq \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r}))$ and $R \in \mathsf{sig}(\Phi(\mathbf{l})) \cap \mathsf{sig}(\Phi(\mathbf{r}))$, where the latter is a consequence of $l \geq 1$.

The argument for 14 is analogous to the previous case. Moreover 15, 16, 17, 18 can be shown similarly to 11, 12, 13, 14, respectively. $\square$

Appendix C. Tableau Correctness, Termination, and Interpolation

Lemma C.1. *Let $\mathbf{T} = \langle \mathcal{V}, \mathcal{E} \rangle$ be the output of the second phase. Then for every node $g \in \mathcal{V}$:*

1. *g.status is either **sat** or **unsat**.*

2. *If g.status $=$ **unsat**, then either one of the following holds.*

   - *g is a sink node[3] containing a clash;*
   - *there is exactly one successor of g' of g such that for some $(C_1 \sqcap C_2)^\lambda \in g$.content, g'.content is a $(C_1 \sqcap C_2)^\lambda$-relief of g.content and g'.status $=$ **unsat**;*
   - *there is exactly one successor of g' of g such that for some $(\leq 1R)^\lambda \in g$.content, g'.content is a $(\leq 1R)^\lambda$-relief of g.content and g'.status $=$ **unsat**;*
   - *there are exactly n successors $g_1, \dots, g_n$ of g, where n is the cardinality of the set $\{(C_1)^{\lambda_1}, \dots, (C_n)^{\lambda_n}\}$ of all $\exists$- or ≥ 2-burdens of g.content, g_i.content is the $(C_i)^{\lambda_i}$-relief of g.content for $i \in \{1, \dots, n\}$, and there is some $i \in \{1, \dots, n\}$ such that g_i.status $=$ **unsat**; or*
   - *there are exactly two successors g_1, g_2 of g such that for some $(C_1 \sqcup C_2)^\lambda \in g$.content, g_i.content is a $(C_1 \sqcup C_2)^\lambda$-relief of g.content for $i \in \{1, 2\}$, g_1.content $\neq g_2$.content, and g_i.status $=$ **unsat** for $i \in \{1, 2\}$.*

3. *If g.status $=$ **sat**, then either one of the following holds.*

   - *g is a sink node not containing a clash,*
   - *there is exactly one successor of g' of g such that for some $(C_1 \sqcap C_2)^\lambda \in g$.content, g'.content is a $(C_1 \sqcap C_2)^\lambda$-relief of g.content and g'.status $=$ **sat**;*
   - *there is exactly one successor of g' of g such that for some $(\leq 1R)^\lambda \in g$.content, g'.content is a $(\leq 1R)^\lambda$-relief of g.content and g'.status $=$ **sat**;*
   - *there are exactly n successors $g_1, \dots, g_n$ of g, where n is the cardinality of the set $\{(C_1)^{\lambda_1}, \dots, (C_n)^{\lambda_n}\}$ of all $\exists$- or ≥ 2-burdens of g.content, g_i.content is the $(C_i)^{\lambda_i}$-relief of g.content for $i \in \{1, \dots, n\}$, and for all $i \in \{1, \dots, n\}$ we have g_i.status $=$ **sat**; or*
   - *there are exactly two successors g_1, g_2 of g such that for some $(C_1 \sqcup C_2)^\lambda \in g$.content, g_i.content is a $(C_1 \sqcup C_2)^\lambda$-relief of g.content for $i \in \{1, 2\}$, g_1.content $\neq g_2$.content, and there is some $i \in \{1, 2\}$ such that g_i.status $=$ **sat**.*

Proof. 1 clearly follows from the fact that every node that is not assigned the status `unsat` during the Propagate step of Algorithm 1 gets the status `sat` at the end (Assign) of Algorithm 1.

Let $g \in \mathcal{V}$. By the definition of the tableau algorithm, g satisfies exactly one of the following structural conditions:

- g is a sink node;

3. a node with no outgoing edges

- there are exactly two successors g_1, g_2 of g such that for some $(C_1 \sqcup C_2)^\lambda \in g.\mathsf{content}$, $g_i.\mathsf{content}$ is a $(C_1 \sqcup C_2)^\lambda$-relief of $g.\mathsf{content}$ for $i \in \{1, 2\}$;
- there is exactly one successor of g' of g such that for some $(C_1 \sqcap C_2)^\lambda \in g.\mathsf{content}$, $g'.\mathsf{content}$ is a $(C_1 \sqcap C_2)^\lambda$-relief of $g.\mathsf{content}$;
- there is exactly one successor of g' of g such that for some $(\leq 1R)^\lambda \in g.\mathsf{content}$, $g'.\mathsf{content}$ is a $(\leq 1R)^\lambda$-relief of $g.\mathsf{content}$;
- there are exactly n successors $g_1, \ldots, g_n$ of g, where n is the cardinality of the set $\{(C_1)^{\lambda_1}, \ldots, (C_n)^{\lambda_n}\}$ of all $\exists$- or ≥ 2-burdens of $g.\mathsf{content}$, and $g_i.\mathsf{content}$ is the $(C_i)^{\lambda_n}$-relief of $g.\mathsf{content}$ for $i \in \{1, \ldots, n\}$.

Suppose first $g.\mathsf{status} = \mathtt{unsat}$ then g clearly respects 2 because these are the only ways for a node to get status $\mathtt{unsat}$ in Propagate. Suppose now $g.\mathsf{status} = \mathtt{sat}$. Then $g.\mathsf{status}$ is determined in Assign of Algorithm 1. We distinguish between the structural properties of g above.

Suppose g is a sink node. This means that no rule is applied to g. Then by $g.\mathsf{status} = \mathtt{sat}$, we immediately obtain that $g.\mathsf{status}$ does not contain a clash; because if $g.\mathsf{status}$ contains a clash, we would have $g.\mathsf{status} = \mathtt{unsat}$ and this contradicts with the fact that for every node $g \in \mathcal{V}$, the value of $g.\mathsf{status}$ is only calculated once.

Suppose there are exactly two successors g_1, g_2 of g such that for some $(C_1 \sqcup C_2)^\lambda \in g.\mathsf{content}$, $g_i.\mathsf{content}$ is a $(C_1 \sqcup C_2)^\lambda$-relief of $g.\mathsf{content}$ for $i \in \{1, 2\}$. Since $g.\mathsf{status} = \mathtt{sat}$, $g.\mathsf{status}$ was undefined right before Assign. This implies that there is some $i \in \{1, 2\}$ such that $g_i.\mathsf{status}$ was undefined because otherwise $g.\mathsf{status} = \mathtt{unsat}$. Then after Assign, $g_i.\mathsf{status} = \mathtt{sat}$. Hence, g satisfies 3.

Suppose there is exactly one successor of g' of g such that for some $(C_1 \sqcap C_2)^\lambda \in g.\mathsf{content}$, $g'.\mathsf{content}$ is a $(C_1 \sqcap C_2)^\lambda$-relief of $g.\mathsf{content}$. Since $g.\mathsf{status} = \mathtt{sat}$, $g.\mathsf{status}$ was undefined right before Assign. This implies that, $g'.\mathsf{status}$ was undefined before Assign because otherwise $g.\mathsf{status} = \mathtt{unsat}$. Then after Assign, $g'.\mathsf{status} = \mathtt{sat}$. Hence, g satisfies 3.

The remaining cases can be shown analogously. Hence the lemma follows. □

Proof of Lemma 3.7. We start with some observations. Algorithm 1 assigns the status $\mathtt{unsat}$ to nodes in $\mathcal{V}$ during the Propagate phase. These status assignment steps induce a sequence $0, 1, 2.\ldots.$ To each assignment step i, we can associate a set $\mathcal{V}_i$ such that $\mathcal{V}_i$ are all the nodes with status $\mathtt{unsat}$ so far. Observe that from step i to step $i+1$, we extend $\mathcal{V}_i$ by a single node only. By induction on the number of status assignment steps, we first show that for all $g \in \mathcal{V}_i$,

- $g.\mathsf{content}$ is unsatisfiable w.r.t. $\mathcal{T}$;
- there is some $\mathcal{ALCF}$-concept C such that
  - $\mathsf{int}(g) = C$,
  - C is an interpolant of $g.\mathsf{content}$,
  - $|\mathsf{int}(g)| \leq 2^{i+2} - 1$.

As the base case, we have that $\mathcal{V}_0 = \{g\}$ for some sink node $g \in \mathcal{V}$ containing a clash. Obviously, $g.\mathsf{content}$ is unsatisfiable. The interpolant calculation rules of Figure 2 cover all the cases for this clash and thus, some $\mathcal{ALCF}$-concept is assigned to $\mathsf{int}(g)$. By Lemma B.4, $\mathsf{int}(g)$ is an interpolant of $g.\mathsf{content}$. We claim that $|\mathsf{int}(g)| \leq 2$. For $\mathsf{int}(g)$ of the form $\top$, $\bot$, A, or $\neg A$, this is clear; and for $\mathsf{int}(g)$ of the form $\leq 1R$ (or $\geq 2R$), we observe that it can be encoded using one symbol for ≤ 1 (resp. ≥ 2) and one symbol for R. Hence, $|\mathsf{int}(g)| \leq 2 \leq 2^{i+2} - 1$ and the inductive hypothesis holds for the base case.

For the inductive step, let $\mathcal{V}_{i+1} = \mathcal{V}_i \cup \{g\}$. The inductive hypothesis holds for every $g' \in \mathcal{V}_i$, trivially; thus, we only consider the case for g. By Lemma C.1, we have five cases to distinguish:

1. g is a sink node containing a clash. This can be shown analogously to the base case.
2. There is exactly one successor of g' of g such that for some $(C_1 \sqcap C_2)^\lambda \in g.\mathsf{content}$, $g'.\mathsf{content}$ is a $(C_1 \sqcap C_2)^\lambda$-relief of $g.\mathsf{content}$ and $g'.\mathsf{status} = \mathtt{unsat}$. By the inductive hypothesis, $g'.\mathsf{content}$ is unsatisfiable w.r.t. $\mathcal{T}$, $\mathsf{int}(g')$ is an interpolant of g', and $|\mathsf{int}(g')| \leq 2^{i+2} - 1$. Then by (the contrapositive version of) Lemma B.1, $g.\mathsf{content}$ is unsatisfiable w.r.t. $\mathcal{T}$. Moreover, $\mathsf{C}_\sqcap$ was applied to calculate $\mathsf{int}(g)$ and $\mathsf{int}(g) = \mathsf{int}(g')$. By Lemma B.4, $\mathsf{int}(g)$ is an interpolant of $g.\mathsf{content}$. We have by $\mathsf{int}(g) = \mathsf{int}(g')$ and $|\mathsf{int}(g')| \leq 2^{i+2} - 1$ that $|\mathsf{int}(g)| \leq 2^{i+3} - 1$. Hence the inductive hypothesis holds for this case.
3. There are exactly two successors g_1, g_2 of g such that for some $(C_1 \sqcup C_2)^\lambda \in g.\mathsf{content}$, $g_j.\mathsf{content}$ is a $(C_1 \sqcup C_2)^\lambda$-relief of $g.\mathsf{content}$ for $j \in \{1,2\}$, $g_1.\mathsf{content} \neq g_2.\mathsf{content}$, and $g_j.\mathsf{status} = \mathtt{unsat}$ for $j \in \{1,2\}$. By the inductive hypothesis, $g_j.\mathsf{content}$ is unsatisfiable w.r.t. $\mathcal{T}$, $\mathsf{int}(g_j)$ is an interpolant of g_j, and $|\mathsf{int}(g_j)| \leq 2^{i+2} - 1$, for $j \in \{1,2\}$. Then by (the contrapositive version of) Lemma B.1, $g.\mathsf{content}$ is unsatisfiable w.r.t. $\mathcal{T}$. Moreover, depending on λ, either $\mathsf{C}^{\mathsf{l}}_\sqcup$ or $\mathsf{C}^{\mathsf{r}}_\sqcup$ was applied to calculate $\mathsf{int}(g)$. By Lemma B.4, $\mathsf{int}(g)$ is an interpolant of $g.\mathsf{content}$. We have that $|\mathsf{int}(g)| = |\mathsf{int}(g_1)| + |\mathsf{int}(g_2)| + 1$. Then by the inductive hypothesis, we obtain
$$|\mathsf{int}(g)| \leq (2^{i+2} - 1) + (2^{i+2} - 1) + 1 = 2^{i+3} - 1.$$
Thus, the inductive hypothesis holds for this case.
4. The other cases can be shown similarly.

Hence, our claim follows.

Now, we use the claim that we have just shown to prove the lemma. Let $g \in \mathcal{V}$ with $g.\mathsf{status} = \mathtt{unsat}$. By Lemma 3.5, we have $|\mathcal{V}| \leq 2^n$, where $n = |\mathsf{cll} \cup \mathsf{clr}|$. Thus, in the worst case, there are 2^n status assignment steps in Propagate because then $\mathcal{V}_{2^n} = \mathcal{V}$. Since $g.\mathsf{status} = \mathtt{unsat}$, it follows that $g \in \mathcal{V}_{2^n}$. Then by our claim, $g.\mathsf{content}$ is unsatisfiable w.r.t. $\mathcal{T}$, $\mathsf{int}(g)$ is defined and it is an interpolant of $g.\mathsf{content}$, and $|\mathsf{int}(g)| \leq 2^{2^n+2} - 1 = 4 \cdot 2^{2^n} - 1$. But then $\mathsf{int}(g) \in O(2^{2^n})$. Hence the lemma follows. □

Proof of Lemma 3.8. Algorithm 1 consists of two stages: Propagate and Assign.

In Assign, we make 2^n assignments because by Lemma 3.5 the number of nodes in the tableau is bounded by that number. Moreover, each assignment step takes a constant time. So the whole Assign stage takes time $O(2^n)$.

In Propagate, we have a for loop inside a do-while loop. The for loop iterates over 2^n nodes and assigns, if possible, the status $\texttt{unsat}$ to a node by checking in the worst case n direct successors of the node. If the algorithm assigns the status $\texttt{unsat}$ to a node g, then it also assigns a concept to $\mathsf{int}(g)$. By Lemma 3.7, $|\mathsf{int}(g)| \leq O(2^{2^n})$. Thus it spends most of its time calculating $\mathsf{int}(g)$. Suppose the for loop finished its iteration over all nodes in the tableau. Now if during its execution, none of the nodes got the status $\texttt{unsat}$, the do-while loop terminates because $\texttt{done} = \texttt{true}$. In the worst case, a status will be assigned only to one node in each iteration of the do-while loop. Hence the do-while loop iterates at most 2^n times and as we discussed each iteration takes time at most $O(2^{2^n})$ because of interpolation calculation. Since this dominates the runtime of Algorithm 1, the lemma follows. $\square$

Lemma C.2 (Soundness)**.** *If* $\mathbf{T}$ *is a closed* $\langle C_0 \sqsubseteq D_0, \mathcal{T}\rangle$*-tableau, then* $\mathcal{T} \models C_0 \sqsubseteq D_0$*.*

Proof. Let $\mathbf{T}$ be a closed $\langle C_0 \sqsubseteq D_0, \mathcal{T}\rangle$-tableau. Since $\mathbf{T}$ is closed, $g_0.\mathsf{content} = \texttt{unsat}$. By $g_0.\mathsf{content} = \{(C_0)^{\mathbf{l}}, (\dot\neg D_0)^{\mathbf{r}}\} \cup \{E^{\mathbf{l}} \mid \top \sqsubseteq E \in \mathcal{T}_{\mathbf{l}}\} \cup \{E^{\mathbf{r}} \mid \top \sqsubseteq E \in \mathcal{T}_{\mathbf{r}}\}$ and Lemma 3.7, this implies

$$\mathcal{T} \models C_0 \sqcap \bigsqcap_{\top \sqsubseteq E \in \mathcal{T}_{\mathbf{l}}} E \sqsubseteq D_0 \sqcup \bigsqcup_{\top \sqsubseteq E \in \mathcal{T}_{\mathbf{r}}} \neg E$$

But then $\mathcal{T} \models C_0 \sqsubseteq D_0$. $\square$

Definition C.3. *Let* $\mathbf{T} = \langle \mathcal{V}, \mathcal{E}\rangle$ *be the output of the second phase. We say that a node* $g \in \mathcal{V}$ *is* saturated *if and only if*

- $g.\mathsf{status} = \texttt{sat}$ *and* g *is a sink node, or*
- $g.\mathsf{status} = \texttt{sat}$ *and* $\mathsf{R}_\exists$ *was applied to* g*.*

For $g, g' \in \mathcal{V}$ *,* g' *is called a* saturation *of* g *if and only if* g' *is saturated, and there is a path* $g = g_0, g_1, \ldots, g_k = g'$ *with* $k \geq 0$ *in* $\mathbf{T}$ *such that for each* $0 \leq i < k$*, we have* $g_i.\mathsf{status} = \texttt{sat}$ *and the edge* $\langle g_i, g_{i+1}\rangle$ *was created by an application of a rule in* $\{\mathsf{R}_\sqcap, \mathsf{R}_\sqcup, \mathsf{R}_{\leq 1}\}$*.*

Lemma C.4. *Let* $\mathbf{T} = \langle \mathcal{V}, \mathcal{E}\rangle$ *be a complete tableau for* $\langle C_0 \sqsubseteq D_0, \mathcal{T}\rangle$*. Then we have*

1. *If* $g \in \mathcal{V}$ *is saturated, then* $g.\mathsf{content}(\lambda)$ *is a* $\langle C_0 \sqcap \dot\neg D_0, \mathcal{T}\rangle$*-type.*
2. *If* $g \in \mathcal{V}$ *with* $g.\mathsf{status} = \texttt{sat}$*, then there is some saturation* g' *of* g *with* $g'.\mathsf{content} \supseteq g.\mathsf{content}$*.*

Proof. For 1, suppose that g is saturated. We need to show that $g.\mathsf{content}(\lambda)$ satisfies Definition A.2. We start with $g.\mathsf{content}(\lambda) \subseteq \mathsf{cl}(C_0 \sqcap \dot\neg D_0, \mathcal{T})$. Let $C \in g.\mathsf{content}(\lambda)$. Then it follows that $C^\lambda \in g.\mathsf{content}$, for some $\lambda \in \{\mathbf{l}, \mathbf{r}\}$. Since $g.\mathsf{content} \subseteq \mathsf{cll} \cup \mathsf{clr}$, we have that $C^\lambda \in \mathsf{cll} \cup \mathsf{clr}$. This implies $C \in \mathsf{cl}(C_0, \mathcal{T}_{\mathbf{l}}) \cup \mathsf{cl}(\dot\neg D_0, \mathcal{T}_{\mathbf{r}})$. But then $C \in \mathsf{cl}(C_0 \sqcap \dot\neg D_0, \mathcal{T})$, which is what we wanted to show.

Now we show that the properties in Definition A.2 are satisfied. By definition, $g.\mathsf{status} = \texttt{sat}$. g does not contain a clash because otherwise $g.\mathsf{status} = \texttt{unsat}$ which would contradict our assumption. Hence, $(\mathsf{P}_\bot)$, $(\mathsf{P}_\neg)$, and $(\mathsf{P}_\bowtie)$ are satisfied. By definition, g is a sink node or $\mathsf{R}_\exists$ was applied to g. In both cases, we have that none of $\{\mathsf{R}_\sqcap, \mathsf{R}_\sqcup, \mathsf{R}_{\leq 1}\}$ is applicable to g: for the former, this follows from the fact that no rule is applicable to g; and for the latter, this follows from our rule precedence. Hence, ($\mathsf{P}_\sqcap$), ($\mathsf{P}_\sqcup$), ($\mathsf{P}_{\leq 1}$) are satisfied. Finally, we have $\{C \mid \top \sqsubseteq C \in \mathcal{T}\} \subseteq g.\mathsf{content}(\lambda)$ as an easy consequence of the definition of the tableau algorithm. This means ($\mathsf{P}_\sqsubseteq$) is satisfied. Hence, we conclude that 1 holds.

For 2, suppose $g \in \mathcal{V}$ with $g.\mathsf{status} = \mathtt{sat}$. That there is some saturation g' of g with $g'.\mathsf{content} \supseteq g.\mathsf{content}$ follows easily by Lemma C.1. $\square$

Lemma C.5 (Completeness)**.** *If* $\mathbf{T}$ *is an open* $\langle C_0 \sqsubseteq D_0, \mathcal{T}\rangle$*-tableau, then* $\mathcal{T} \not\models C_0 \sqsubseteq D_0$*.*

Proof. Suppose $\mathbf{T} = \langle \mathcal{V}, \mathcal{E}\rangle$ is an open $\langle C_0 \sqsubseteq D_0, \mathcal{T}\rangle$-tableau. Since $\mathbf{T}$ is open, we have $g_0.\mathsf{status} = \mathtt{sat}$. Then by Lemma C.4, there is some saturation $g_\star$ of g_0 such that $g_\star.\mathsf{content} \supseteq g_0.\mathsf{content}$. Since $g_\star$ is saturated, it follows by Lemma C.4 that $g_\star.\mathsf{content}(\lambda)$ is a $\langle C_0 \sqcap \dot\neg D_0, \mathcal{T}\rangle$-type. Let $\tau_0 = \{C_0 \sqcap \dot\neg D_0\} \cup g_\star.\mathsf{content}(\lambda)$. We claim that τ_0 is also a $\langle C_0 \sqcap \dot\neg D_0, \mathcal{T}\rangle$-type. Suppose for a contradiction that it is not. Since $g_\star.\mathsf{content}(\lambda)$ is such a type, it follows that ($\mathsf{P}_\sqcap$) is violated for $C_0 \sqcap \dot\neg D_0 \in \tau_0$, i.e., $\{C_0, \dot\neg D_0\} \not\subseteq \tau_0$. Then by the definition of τ_0, this means $\{C_0, \dot\neg D_0\} \not\subseteq g_\star.\mathsf{content}(\lambda)$. But we know that $\{C_0, \dot\neg D_0\} \subseteq g_0.\mathsf{content}(\lambda)$ and by $g_\star.\mathsf{content} \supseteq g_0.\mathsf{content}$, this implies $\{C_0, \dot\neg D_0\} \subseteq g_\star.\mathsf{content}(\lambda)$, i.e., a contradiction. Hence we conclude that τ_0 is a $\langle C_0 \sqcap \dot\neg D_0, \mathcal{T}\rangle$-type. Define

$$Q = \{\tau_0\} \cup \{g.\mathsf{content}(\lambda) \mid g \in \mathcal{V} \text{ is saturated}\}.$$

We show that Q is a $\langle C_0 \sqcap \dot\neg D_0, \mathcal{T}\rangle$-quasimodel because then $\mathcal{T} \not\models C_0 \sqsubseteq D_0$ follows by Theorem A.4. It is easy to see that Q is a set of $\langle C_0 \sqcap \dot\neg D_0, \mathcal{T}\rangle$-types: we have already shown that τ_0 is such a type; and for $g.\mathsf{content}(\lambda)$ with $g \in \mathcal{V}$ is saturated, this fact follows immediately by Lemma C.4. It remains to show that conditions (a), (b), (c) from Definition A.3 hold.

Condition (a) holds since $\tau_0 \in Q$ and $C_0 \sqcap \dot\neg D_0 \in \tau_0$.

Suppose that $\exists R.C \in \tau$ for some $\tau \in Q$. We distinguish between $\tau = \tau_0$ and $\tau = g.\mathsf{content}(\lambda)$ for some saturated $g \in \mathcal{V}$. We first argue for the latter. Since g is saturated and $\exists R.C \in g.\mathsf{content}(\lambda)$, $\mathsf{R}_\exists$ was applied to g; since g is saturated, we have $g.\mathsf{status} = \mathtt{sat}$. Then by Lemma C.1, there is some successor g' of g in $\mathbf{T}$ such that $\{C\} \cup \{D \mid \forall R.D \in \tau\} \subseteq g'.\mathsf{content}(\lambda)$ and $g'.\mathsf{status} = \mathtt{sat}$. Then by Lemma C.4, there is some saturation g'' of g' such that $g''.\mathsf{content} \supseteq g'.\mathsf{content}$. Since g'' is saturated, we have that $g''.\mathsf{content}(\lambda) \in Q$ and $g''.\mathsf{content}(\lambda)$ is a $\langle C_0 \sqcap \dot\neg D_0, \mathcal{T}\rangle$-type. But then $\tau \xRightarrow{\exists R.C} g''.\mathsf{content}(\lambda)$. The case for $\tau = \tau_0$ follows analogously by using the fact that there is some successor g' of $g_\star$ in $\mathbf{T}$ such that $\{C\} \cup \{D \mid \forall R.D \in \tau\} \subseteq g'.\mathsf{content}(\lambda)$ and $g'.\mathsf{status} = \mathtt{sat}$. Hence condition (b) from Definition A.3 is satisfied.

That condition (c) holds can be shown very similarly to the previous case; we leave it to the reader to verify this. Hence, we conclude that Q is a $\langle C_0 \sqcap \dot\neg D_0, \mathcal{T}\rangle$-quasimodel. $\square$

Proposition 3.9 now follows immediately from Lemma C.2 and Lemma C.5.

References

Afrati, F. N. (2011). Determinacy and query rewriting for conjunctive queries and views. *Theoretical Computer Science*, *412*(11), 1005–1021.

Andréka, H., Németi, I., & van Benthem, J. (1998). Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, *27*, 217–274.

Avigad, J. (2003). Eliminating definitions and skolem functions in first-order logic. *ACM Transactions on Computational Logic*, *4*, 402–415.

Baader, F., & Nutt, W. (2003). Basic description logics. In *The Description Logic Handbook*, pp. 43–95. Cambridge University Press.

Bárány, V., Benedikt, M., & ten Cate, B. (2013). Rewriting guarded negation queries. In *MFCS13*, pp. 98–110.

Beth, E. W. (1953). On Padoa's methods in the theory of definitions. *Indagationes Mathematicae*, *15*, 330–339.

Blackburn, P., de Rijke, M., & Venema, Y. (2001). *Modal logic*. Cambridge University Press.

Boolos, G. S., Burgess, J. P., & Jeffrey, R. C. (2007). *Computability and Logic*. Cambridge University Press.

Calvanese, D., & Giacomo, G. D. (2003). Expressive description logics. In *The Description Logic Handbook*, pp. 178–218. Cambridge University Press.

Calvanese, D., Giacomo, G. D., Lenzerini, M., & Nardi, D. (2001). Reasoning in expressive description logics. In *Handbook of Automated Reasoning*, pp. 1581–1634.

Calvanese, D., Giacomo, G. D., & Rosati, R. (1998). A note on encoding inverse roles and functional restrictions in $\mathcal{ALC}$ knowledge bases. In *Description Logics*, Vol. 11. CEUR-WS.org.

ten Cate, B., Conradie, W., Marx, M., & Venema, Y. (2006). Definitorially complete description logics. In *KR*, pp. 79–89.

ten Cate, B., Franconi, E., & Seylan, İ. (2011). Beth definability in expressive description logics. In *IJCAI*, pp. 1099–1106.

Conradie, W. (2002). Definability and changing perspectives: The beth property for three extensions of modal logic. Master's thesis, University of Amsterdam.

Craig, W. (1957). Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. *The Journal of Symbolic Logic*, *22*(3), 269–285.

De Giacomo, G. (1996). Eliminating "converse" from Converse PDL. *Journal of Logic, Language and Information*, *5*(2), 193–208.

Donini, F. M. (2003). Complexity of reasoning. In *The Description Logic Handbook*, pp. 96–136. Cambridge University Press.

Duc, C. L., & Lamolle, M. (2010). Decidability of description logics with transitive closure of roles in concept and role inclusion axioms. In *Description Logics*, Vol. 573, pp. 372–383. CEUR-WS.org.

Fitting, M. (1996). *First-order logic and automated theorem proving (2nd ed.)*. Springer-Verlag.

Friedman, H. (1976). The complexity of explicit definitions. *Advances in Mathematics*, *20*(1), 18–29.

Gabbay, D. M., & Maksimova, L. (2005). *Interpolation and Definability in Modal Logics (Oxford Logic Guides)*. Clarendon Press.

Ghilardi, S., Lutz, C., & Wolter, F. (2006). Did I damage my ontology? A case for conservative extensions in description logics. In *KR*, pp. 187–197.

Goré, R. (1999). Tableau methods for modal and temporal logics. In *Handbook of Tableau Methods*, pp. 297–396. Kluwer.

Goré, R., & Nguyen, L. A. (2007). Exptime tableaux with global caching for description logics with transitive roles, inverse roles and role hierarchies. In *TABLEAUX*, pp. 133–148.

Hoogland, E. (2001). *Definability and Interpolation: Model-theoretic investigations*. Ph.D. thesis, University of Amsterdam.

Hoogland, E., & Marx, M. (2002). Interpolation and definability in guarded fragments. *Studia Logica*, *70*(3), 373–409.

Horrocks, I., Patel-Schneider, P. F., & van Harmelen, F. (2003). From $\mathcal{SHIQ}$ and RDF to OWL: The making of a web ontology language. *Journal of Web Semantics*, *1*(1), 7–26.

Horrocks, I., & Sattler, U. (2007). A tableau decision procedure for $\mathcal{SHOIQ}$. *Journal of Automated Reasoning*, *39*(3), 249–276.

Horrocks, I., Sattler, U., & Tobies, S. (2000). Practical reasoning for very expressive description logics. *Logic Journal of the IGPL*, *8*(3), 239–264.

Konev, B., Lutz, C., Ponomaryov, D., & Wolter, F. (2010). Decomposing description logic ontologies. In *KR*, pp. 236–246.

Konev, B., Lutz, C., Walther, D., & Wolter, F. (2009a). Formal properties of modularisation. In *Modular Ontologies*, pp. 25–66. Springer.

Konev, B., Walther, D., & Wolter, F. (2009b). Forgetting and uniform interpolation in large-scale description logic terminologies. In *IJCAI*, pp. 830–835.

Kracht, M. (2007). Modal consequence relations. In *Handbook of Modal Logic*, pp. 491–545. Elsevier.

Lang, J., & Marquis, P. (2008). On propositional definability. *Artificial Intelligence*, *172*, 991–1017.

Libkin, L. (2004). *Elements of Finite Model Theory*. Springer.

Lutz, C. (2006). Complexity and succinctness of public announcement logic. In *AAMAS*, pp. 137–143.

Lutz, C., Areces, C., Horrocks, I., & Sattler, U. (2005). Keys, nominals, and concrete domains. *Journal of Artificial Intelligence Research*, *23*, 667–726.

Lutz, C., Piro, R., & Wolter, F. (2010). Enriching $\mathcal{EL}$-concepts with greatest fixpoints. In *ECAI*, pp. 41–46.

Lutz, C., Sattler, U., & Tendera, L. (2005). The complexity of finite model reasoning in description logics. *Information and Computation*, *199*(1-2), 132–171.

Lutz, C., Seylan, İ., & Wolter, F. (2012a). An automata-theoretic approach to uniform interpolation and approximation in the description logic $\mathcal{EL}$. In *KR*.

Lutz, C., Seylan, İ., & Wolter, F. (2012b). Mixing open and closed world assumption in ontology-based data access: Non-uniform data complexity. In *Description Logics*, Vol. 846, pp. 268–278. CEUR-WS.org.

Lutz, C., & Wolter, F. (2011). Foundations for uniform interpolation and forgetting in expressive description logics. In *IJCAI*, pp. 989–995.

Marx, M. (2007). Queries determined by views: pack your views. In *PODS*, pp. 23–30.

Marx, M., & Venema, Y. (2007). Local variations on a loose theme: Modal logic and decidability. In *Finite Model Theory and Its Applications*, pp. 371–426. Springer.

Nash, A., Segoufin, L., & Vianu, V. (2010). Views and queries: Determinacy and rewriting. *ACM Transactions on Database Systems*, *35*(3).

Nikitina, N., & Rudolph, S. (2012). Expexpexplosion: Uniform interpolation in general $\mathcal{EL}$ terminologies. In *ECAI*, pp. 618–623.

Pasaila, D. (2011). Conjunctive queries determinacy and rewriting. In *ICDT*, pp. 220–231.

Rautenberg, W. (1983). Modal tableau calculi and interpolation. *Journal of Philosophical Logic*, *12*(4), 403–423.

Sattler, U., Calvanese, D., & Molitor, R. (2003). Relationships with other formalisms. In *The Description Logic Handbook*, pp. 137–177. Cambridge University Press.

Schwendimann, S. (1998). A new one-pass tableau calculus for PLTL. In *TABLEAUX*, pp. 277–292.

Seylan, İ. (2012). *DBoxes and Beth Definability in Description Logics*. Ph.D. thesis, Free University of Bozen-Bolzano.

Seylan, İ., Franconi, E., & de Bruijn, J. (2009). Effective query rewriting with ontologies over DBoxes. In *IJCAI*, pp. 923–929.

Seylan, İ., Franconi, E., & de Bruijn, J. (2010). Optimal rewritings in definitorially complete description logics. In *Description Logics*, Vol. 573, pp. 125–136. CEUR-WS.org.

Tobies, S. (2001). *Complexity Results and Practical Algorithms for Logics in Knowledge Representation*. Ph.D. thesis, RWTH-Aachen.