Universal Simulands and Subsumption Checking in Lightweight Coalgebraic Logics

Daniel Gorín and Lutz Schröder

Department of Computer Science, FAU Erlangen-Nürnberg

Abstract. While reasoning in logics extending a Boolean propositional base is necessarily at least coNP-hard, some families of specifically designed lightweight logics allow for tractable, i.e. polynomial-time reasoning, and hence may be expected to scale to large reasoning problems. One example of this type is the $\mathcal{EL}$ family of description logics; in this case, efficient reasoning may be based on simulation checking between suitable small models. In the current work, we lift this principle to the level of generality of coalgebraic logic. We thus not only identify tractable fragments of non-relational logics whose semantics features, e.g., neighbourhoods or integer weights, but we also obtain new insights in the standard relational setting, e.g. on polynomial-time reasoning with global assumptions in modal logics featuring only box and conjunction.

1 Introduction

One of the most successful applications of modal logics in computer science is in *knowledge representation*, where *description logics* (DLs) are prevalent [4]; these are essentially notational variants of relational modal logics. One of the most basic reasoning tasks is then to decide whether a formula (or *concept*) ϕ is *subsumed* by a formula ψ, notation $\phi \sqsubseteq \psi$, i.e. whether ϕ is a particular case of ψ, possibly taking into account a so-called *terminology* or *TBox*, i.e. a set of global assumptions.

Figure 1 displays a very simple TBox (in German DL notation). There, a 'person' is said to be something whose 'offspring' are 'persons' as well (a property that is also true of any infertile entity), while a 'parent' is defined precisely as a 'person' with 'offspring'. It is not hard to verify that this TBox implies

$$\mathsf{Happy} \sqcap \mathsf{GrandParent} \sqsubseteq \exists \mathsf{hasOffspring}.\, \exists \mathsf{hasOffSpring}.\, (\mathsf{Happy} \sqcap \mathsf{Person}).$$

Many DLs extend the basic language $\mathcal{ALC}$, where concepts can be built using conjunctions ($\sqcap$), disjunction ($\sqcup$), negation ($\neg$), and relational modalities ($\exists\mathsf{R}/\forall\mathsf{R}$). Modulo syntax, the concept language of $\mathcal{ALC}$ is just the basic multi-modal logic $\mathbf{K}_m$. The computational complexity of subsumption in $\mathcal{ALC}$ is PSPACE-complete without a TBox, and EXPTIME-complete over unrestricted TBoxes, respectively. The term *lightweight description logics* designates families of DLs with tractable reasoning. Their concept language is necessarily non-Boolean, and typically lacks disjunction and negation. In particular, languages from the $\mathcal{EL}$ family lack also universal restriction ($\forall\mathsf{R}$), but enjoy a polynomial-time subsumption checking problem [2, 3]. Their expressiveness however suffices to accommodate, e.g., large-scale medical ontologies such as SNOMED CT[1].

[1] http://www.ihtsdo.org/snomed- ct/

$$\begin{array}{ll} \mathsf{Person} \sqsubseteq \forall \mathsf{hasOffspring}.\mathsf{Person} & \mathsf{Parent} \equiv \mathsf{Person} \sqcap \exists \mathsf{hasOffspring}.\top \\ \mathsf{Happy} \sqsubseteq \forall \mathsf{hasOffspring}.\mathsf{Happy} & \mathsf{GrandParent} \equiv \mathsf{Parent} \sqcap \exists \mathsf{hasOffspring}.\mathsf{Parent} \end{array}$$

Fig. 1. A very simple terminology of parenthood in $\mathcal{ALC}$.

Maybe surprisingly, $\mathcal{FL}_0$, the counterpart of $\mathcal{EL}$ with universal instead of existential restriction, has a subsumption problem that is tractable only over the empty TBox [1, 5, 12]. This shows that there is more to lightweight DLs than just dropping disjunctions. Here, we aim to develop the conceptual tools to identify lightweight modal formalisms beyond the relational realm, covering also semantic concepts such as gradedness (*qualified number restrictions* in DL parlance) and monotone neighbourhoods (and, in principle, any other semantics satisfying the criteria we develop, ATL-style game-based semantics being one candidate). We achieve this by working in the setting of *coalgebraic logic* [14], where the notions of model and modal operators are abstracted to deal with relational and non-relational features in a uniform fashion. Coalgebraic logic is, of course, agnostic about the interpretation of its instance logics, which may equally well be understood as logics of labelled transition systems and concurrency.

Our main new contribution is the identification of criteria (necessarily highly restrictive) that allow replacing formulas by models in the sense that satisfaction of the formula is equivalent to simulation of the model, following an approach for $\mathcal{EL}$ [2] and building on a recently developed notion of coalgebraic simulation [10]. Calling such models *universal simulands*, one then reduces subsumption to similarity of universal simulands. We believe that universal simulands are an important model-theoretic concept, and in particular expect that they will prove to be instrumental not only for subsumption checking but also for more general reasoning problems prominently including *least common subsumers* [5]. Tractability of reasoning depends on smallness of universal simulands; we develop criteria for this that not only reproduce known results on $\mathcal{EL}$ and $\mathcal{FL}_0$ but also yield new tractability results for conjunctive monomodal K with only boxes as well as fragments of monotone and graded modal logic. This will allow for TBoxes that, varying Figure 1, say things like 'happy parents have at most 2 unhappy children' (graded modal logic) or 'happy parents win arguments about disco visits' (monotone modal logic, read as a logic of games in the spirit of [13]).

2 Preliminaries: Coalgebraic Logic

We begin with a brief introduction to the basic concepts and terminology of coalgebraic logics. The generality of coalgebraic modal logics stems from the parametricity of its syntax and semantics. The language depends on a *similarity type* Λ, i.e. a set of *modal operators* with finite arities ≥ 0. (Atomic) propositions are just modalities of arity 0. To simplify notation, we will pretend that all operators are unary; however, all results generalize straightforwardly to higher arities.

Definition 1. The set $L(\Lambda)$ of Λ-*formulas* is given by the grammar

$$\phi, \psi ::= \top \mid \neg\phi \mid \phi \wedge \psi \mid \heartsuit\phi \qquad (\heartsuit \in \Lambda).$$

(with $\vee$ derived as usual). By $rank(\phi)$ we denote the maximal nesting depth of modal operators $\heartsuit \in \Lambda$ in ϕ. We are interested in two types of fragments of $L(\Lambda)$: a formula is *positive* if it is generated by the grammar

$$\phi, \psi ::= \top \mid \bot \mid \phi \wedge \psi \mid \phi \vee \psi \mid \heartsuit\phi \qquad (\heartsuit \in \Lambda);$$

moreover, a positive formula is *conjunctive* if it does not mention $\bot$ and disjunction. Generally, given an operator $\heartsuit \in \Lambda$ we refer to the operator $\bar{\heartsuit}$ with $\bar{\heartsuit}\phi$ interpreted like $\neg\heartsuit\neg\phi$ as the *dual* of $\heartsuit$. We do *not* assume that Λ is closed under duals, and inclusion or non-inclusion of mutually dual operators in Λ usually makes a big difference for the lightweight logics we are considering here.

The semantics is parametrized by associating a Λ-structure $\langle T, \{[\![\heartsuit]\!]\}_{\heartsuit\in\Lambda}\rangle$ to a similarity type Λ. Here, T is an endofunctor on the category Set of sets, and each $[\![\heartsuit]\!]$ is a *predicate lifting*, that is, a natural transformation $[\![\heartsuit]\!] : \mathcal{Q} \dot{\to} \mathcal{Q} \circ T^{op}$, where $\mathcal{Q} : \mathsf{Set}^{op} \to \mathsf{Set}$ is the contravariant powerset functor (that is, $\mathcal{Q}X = 2^X$ for every set X, and for a map f, Qf takes preimages under f; thus, naturality of $[\![\heartsuit]\!]$ means that $[\![\heartsuit]\!]_X(f^{-1}[A]) = Tf^{-1}[[\![\heartsuit]\!]_Y(A)]$ for $f : X \to Y$). Note that from $[\![\heartsuit]\!]$ we obtain a dual predicate lifting correctly interpreting $\bar{\heartsuit}$ by $[\![\bar{\heartsuit}]\!]_X(A) = TX - [\![\heartsuit]\!]_X(X - A)$.

Assumption 2. We can assume w.l.o.g. that T preserves injective maps [6]. For convenience, we will in fact assume that T preserves subsets.

Abusing notation, we shall identify a similarity type Λ with its associated Λ-structure, and refer to both as Λ, with the underlying functor denoted by T throughout.

A *model* for $L(\Lambda)$ is just a T-*coalgebra* $C = (X, \xi)$, i.e. a set X (of *states*) and a *transition function* $\xi : X \to TX$. A *pointed* T-*coalgebra* is just a pair (C, r), where r, a state of C, is called the *point* or *root*. Given $x \in X$, satisfaction of $L(\Lambda)$-formulas ϕ at x ($x \models_C \phi$) is defined by the expected clauses for Boolean operators, and

$$x \models_C \heartsuit\phi \iff \xi(x) \models \heartsuit[\![\phi]\!]_C$$

where $[\![\phi]\!]_C = \{x \in X \mid x \models_C \phi\}$ is the *extension* of ϕ in C, and for $t \in TX$ and $A \subseteq X$, $t \models \heartsuit A$ is a more suggestive notation for $t \in [\![\heartsuit]\!]_X A$.

When restricting to positive formulas we can no longer encode all reasoning tasks as validity or satisfiability; rather, we consider as the core reasoning task *local consequence* or, in description logic terms, *subsumption*: For formulas ϕ and ψ, we say that ψ *subsumes* ϕ, and write $\phi \sqsubseteq \psi$, if $[\![\phi]\!]_C \subseteq [\![\psi]\!]_C$ in all T-coalgebras C.

Example 3. Coalgebras for the (covariant) finite powerset functor $\mathcal{P}_\omega$ are finitely branching directed graphs. For $\Lambda = \{\Box, \Diamond\}$ one has predicate liftings

$$[\![\Box]\!]_X(A) := \{B \mid B \subseteq A\}$$
$$[\![\Diamond]\!]_X(A) := \{B \mid B \cap A \neq \emptyset\}.$$

To obtain the full basic modal logic K one additionally needs to enrich the coalgebra structure with an interpretation for propositions. So let V be a set of propositions (nullary modal operators), and let C_V be the constant functor that maps every set X to 2^V. For each $p \in V$, the (nullary) predicate lifting $[\![p]\!]_X := \{\pi \in 2^V \mid p \in \pi\}$ describes structures satisfying p. The Kripke functor K is then defined as $K = C_V \times \mathcal{P}$, and the similarity type $\Lambda = V \cup \{\Diamond, \Box\}$ is interpreted using the corresponding predicate liftings on the appropriate projections. We largely forget about propositions until Section 5.

Example 4. The language of *graded (modal) logic* has the similarity type $\Lambda = \{\Diamond_k \mid k \in \mathbb{N}\}$ (with $\Diamond_k$ read 'in more than k successors') and is interpreted over the multiset functor $\mathcal{B}_\infty$, i.e., $\mathcal{B}_\infty X = X \to \mathbb{N} \cup \{\infty\}$. We regard $b \in \mathcal{B}_\infty X$ as an $\mathbb{N} \cup \{\infty\}$-valued measure on X, and correspondingly write $b(A) = \sum_{x \in A} b(x)$ for a subset $A \subseteq X$. Coalgebras for $\mathcal{B}_\infty$ are *multigraphs*, i.e. directed graphs whose edges are annotated with multiplicities from $\mathbb{N} \cup \{\infty\}$. Interpretation of the modal operators is by way of the following family of predicate liftings, for each $k \in \mathbb{N}$:

$$[\![\Diamond_k]\!]_X(A) := \{b \in \mathcal{B}_\infty X \mid b(A) > k\}.$$

Example 5. Consider the subfunctor $\mathcal{M}$ of the neighbourhood functor $\mathcal{Q} \circ \mathcal{Q}$ given by $\mathcal{M}X = \{\mathfrak{N} \in \mathcal{Q}\mathcal{Q}X \mid \mathfrak{N} \text{ is upwards closed under } \subseteq\}$. Over this functor one obtains the monotone neighbourhood semantics of modal logic with $\Lambda = \{\Box\}$ using the predicate lifting $[\![\Box]\!]_X(A) := \{S \in \mathcal{M}X \mid A \in S\}$. The dual of $\Box$ is written $\Diamond$.

Generally, a modal operator $\heartsuit$ is *monotone* if $A \subseteq B \subseteq X$ implies $[\![\heartsuit]\!]_X A \subseteq [\![\heartsuit]\!]_X B$. All examples above are monotone. In general, coalgebraic logic does support non-monotone logics. However, since we are exploring a method of subsumption checking via simulations, we need to restrict to monotone logics (basically, simulations preserve but do not reflect satisfaction of formulas, so that inductive proofs need monotonicity).

Assumption 6. In the following, we assume all modal operators to be monotone.

Coalgebraic logic adopts the local perspective of modal languages (cf. Slogan 2 on [7, p.ix]). In fact, many phenomena such as *derivability*, *satisfiability* or, as we shall see in the following sections, *similarity* can be studied in the simpler setting of *one-step models* (roughly, the result of forgetting the structure of a pointed coalgebra everywhere except at the root), and results then extrapolate to the general case (e.g. [15]). With one-step models come *one-step formulas*, i.e. shallow modal formulas where propositional variables are introduced as placeholders for the 'missing' recursive structure.

Definition 7 (One-step models and formulas). Let V be a set of propositional variables (not fixed, and typically finite); a *one-step model over V* is a tuple (X, τ, t) where X is a set (possibly empty), $\tau : V \to \mathcal{P}X$ interprets propositional variables, and $t \in TX$. The dual representation of τ is $\breve{\tau} : X \to \mathcal{P}V$, i.e. $\breve{\tau}(x) = \{p \mid x \in \tau(p)\}$. A *(simple) one-step ($\Lambda$-)formula* is a Boolean combination of atoms $\heartsuit p$, where $\heartsuit \in \Lambda$, $p \in V$ (here, we use, and immediately omit, the word *simple* to differentiate from the more general case where arguments of modal operators can be Boolean combinations of propositional variables). The satisfaction relation $t \models_\tau \phi$ is given by the usual Boolean clauses plus $t \models_\tau \heartsuit p \iff t \in [\![\heartsuit]\!]_X \tau(p)$. We say that a one-step formula is *positive* if it mentions only atoms, $\vee$, and $\wedge$, and *conjunctive* if it mentions only atoms and $\wedge$.

The transfer of results between the one-step and the full logic is often by way of what we shall term *collages* (pasting pointed coalgebras into a one-step model to form a new coalgebra) and *décollages* (tearing away most of the structure of a pointed coalgebra to obtain a one-step model); e.g. this technique has been used in the construction of shallow models for coalgebraic modal [15] and hybrid [11] logics. Explicitly:

Definition 8. A pointed coalgebra (C, r), with $C = (Y, \xi)$, is a *collage* over a one-step model (X, τ, t) if there is a family of coalgebras $C_x = (Y_x, \xi_x)$ with $x \in Y_x$ for all $x \in X$ such that Y is the disjoint union of $\{r\}$ and the Y_x, and

$$\xi(y) := \begin{cases} T(i_X)(t) & \text{if } y = r \\ T(i_{Y_x})(\xi_x(y)) & \text{otherwise, for the } x \text{ such that } y \in Y_x \end{cases}$$

where, for every $Y_0 \subseteq Y$, $i_{Y_0} : Y_0 \hookrightarrow Y$ denotes the inclusion function.

In a nutshell, a collage over (X, τ, t) is obtained by replacing every $x \in X$ by a pointed coalgebra (C_x, x). The following is immediate by construction:

Lemma 9 (Collage lemma). *For a collage (C, r) over (X, τ, t), with $C = (Y, \xi)$,*

1. *$x \models_C \phi \iff x \models_{C_x} \phi$, for all $x \in X$, and*
2. *$t \in \heartsuit_X(A \cap X) \iff \xi(r) \in \heartsuit_Y A$, for all $A \subseteq Y$ and $\heartsuit \in \Lambda$.*

One typically needs collages based on interpretations of propositional variables as modal formulas. Here, we will be interested in *preserving* the interpretation of the satisfied atoms; more precisely:

Definition 10. A collage (C, r) over a one-step model (X, τ, t) over V *(positively) fulfills* a substitution $\rho : V \to L(\Lambda)$ if for all $x \in X$, $x \models_{C_x} \rho(p)$ iff (if) $x \in \tau(p)$.

Corollary 11. *If (C, r) is a collage over (X, τ, t) (positively) fulfilling ρ, then*

1. *$x \in \tau(p)$ iff (implies) $x \models_C \rho(p)$, and*
2. *$t \models_\tau \heartsuit p$ iff (implies) $r \models_C \heartsuit\rho(p)$.*

The converse process is as follows.

Definition 12. Given a pointed coalgebra (C, r) with $C = (X, \xi)$ and a substitution $\rho : V \to L(\Lambda)$, we say that (X, τ, t) is the *décollage of* (C, r) *by* ρ if $t = \xi(r)$ and $\tau(p) = [\![\rho(p)]\!]_C$.

Lemma 13 (Décollage lemma). *If (X, τ, t) is a décollage of (C, r) by $\rho : V \to L(\Lambda)$ then for all one-step formulas ϕ over V we have $t \models_\tau \phi \iff r \models_C \phi\rho$.*

3 Simulations

We now proceed to introduce our notion of modal simulation. Given a binary relation $S \subseteq X \times Y$ and $A \subseteq X$, we denote by $S[A]$ the relational image $S[A] = \{y \mid \exists x \in A.\, xSy\}$, and by S^- the relational inverse of S.

Definition 14 (**Λ-Simulation).** Let $C = (X, \xi)$ and $D = (Y, \zeta)$ be T-coalgebras. A *Λ-simulation $S : C \to D$ (of C by D)* is a relation $S \subseteq X \times Y$ such that whenever xSy then for all $\heartsuit \in \Lambda$ and all $A \subseteq X$

$$\xi(x) \models \heartsuit A \text{ implies } \zeta(y) \models \heartsuit S[A]. \tag{1}$$

If xSy for a Λ-simulation S, then we say that (C, x) *and* (D, y) *are Λ-similar*.

The crucial properties of Λ-simulations that we need here are sufficient stability under standard constructions and preservation of positive formulas:

Proposition 15. *Λ-simulations are stable under relational composition; moreover, (graphs of) identities are Λ-simulations.*

Proposition 16. *Let $S : C \to D$ be a Λ-simulation of coalgebras $C = (X, \xi)$ and $D = (Y, \zeta)$, and let ϕ be a positive Λ-formula. Then xSy and $x \models_C \phi$ imply $y \models_D \phi$.*

The effect of dualizing modal operators is to turn around the notion of simulation:

Proposition 17. *Let $\bar{\Lambda} = \{\bar{\heartsuit} \mid \heartsuit \in \Lambda\}$. Then a relation S between T-coalgebras is a $\bar{\Lambda}$-simulation iff S^- is a Λ-simulation.*

Example 18. (See [10] for details)

1. Over Kripke frames, when $\Lambda = \{\Diamond\}$, then a Λ-simulation $S : C \to D$ is just a simulation $C \to D$ in the usual sense.
2. By Proposition 17, when $\Lambda = \{\Box\}$, then a Λ-simulation $S : C \to D$ is just a simulation $D \to C$ in the usual sense.
3. Consequently, a $\{\Box, \Diamond\}$-simulation is a bisimulation in the usual sense.
4. For graded modal logic, with $\Lambda = \{\Diamond_k \mid k \in \mathbb{N}\}$, a relation $S \subseteq X \times Y$ between $\mathcal{D}$-coalgebras (X, ξ) and (Y, ζ) is a Λ-simulation iff for all xSy and all $A \subseteq X$,

$$\zeta(y)(S[A]) \geq \xi(x)(A) \tag{2}$$

(keep in mind that we view $\xi(x) \in \mathcal{B}_\infty(X)$, $\zeta(y) \in \mathcal{B}_\infty(Y)$ as discrete $\mathbb{N} \cup \{\infty\}$-valued measures). Consequently, for $\bar{\Lambda} = \{\Box_k \mid k \in \mathbb{N}\}$, S is a $\bar{\Lambda}$-simulation iff $\xi(x)(S^-[B]) \geq \zeta(y)(B)$ for all xSy and all $B \subseteq Y$.
5. For monotone neighbourhood logic, with $\Lambda = \{\Box\}$, we have that a relation $S \subseteq X \times Y$ between $\mathcal{M}$-coalgebras (X, ξ) and (Y, ζ) is a Λ-simulation iff for xSy, $A \in \xi(x)$ implies $S[A] \in \zeta(y)$.

As our method of subsumption checking will be based on checking similarity, we need to analyse the complexity of the latter, which formally is the following problem:

Definition 19. The *Λ-similarity problem* is to decide whether two finite pointed coalgebras (C, x) and (D, y) are Λ-similar.

(Here and in the following, we assume a suitable representation format for elements of TX determining the input sizes for decision problems such as this one; in the examples, exact choices will play only a minor role and hence will mostly be glossed over.)

We follow the usual paradigm of reducing this problem to one formulated on the one-step level, specifically:

Definition 20. A relation $S \subseteq X \times Y$ is a *one-step Λ-simulation* between $t \in TX$ and $s \in TY$ if for every $\heartsuit \in \Lambda$ and every $A \subseteq X$,

$$t \models \heartsuit A \quad \text{implies} \quad s \models \heartsuit S[A].$$

The *one-step Λ-simulation problem* is to decide, given $S \subseteq X \times Y$ with X, Y finite, $t \in TX$, and $s \in SX$, whether S is a one-step Λ-simulation between t and s.

Given this notion, we can evidently reformulate the definition of S being a Λ-simulation between $C = (X, \xi)$ and $D = (Y, \zeta)$ as saying that whenever xSy then S is a one-step Λ-simulation between $\xi(x)$ and $\zeta(y)$. In other words, Λ-simulations are post-fixpoints of the (monotone) functional F taking $S \subseteq X \times Y$ to $\{(x, y) \in S \mid S \text{ is a one-step } \Lambda\text{-simulation from } \xi(x) \text{ to } \zeta(y)\}$. When C and D are finite, we can therefore compute the greatest simulation between C and D by iterating F, starting with $S = X \times Y$ as the initial value. We thus have the desired criterion:

Proposition 21. *If the one-step Λ-simulation problem is in P, then the Λ-similarity problem is in P.*

Example 22. 1. For $\Lambda \subseteq \{\Box, \Diamond\}$, interpreted over Kripke frames, one-step Λ-simulation is easily seen to be in P using suitable characterizations (e.g. S is a one-step $\Diamond$-simulation from $t \in \mathcal{P}(X)$ to $s \in \mathcal{P}(Y)$ iff for every $x \in t$ there exists $y \in s$ such that xSy). We thus regain the well-known result that standard similarity is in P.

2. For graded logic (with $\Lambda = \{\Diamond_k \mid k \in \mathbb{N}\}$), the situation is less opportune, as to check one-step Λ-simulation there seems to be no universally applicable improvement over the naive idea of checking Equation (2) for all sets A. In fact, the one-step Λ-similarity problem for this case looks suspiciously similar to (the complement of) the subset sum problem, and we conjecture that it is coNP-hard, hence coNP-complete, being clearly in coNP. (We will see later that whether or not hardness really holds in this case is not all that relevant for our current purposes.)

3. For monotone modal logic, with $\Lambda = \{\Box\}$, we represent elements of TX, i.e. upwards closed systems of subsets of X, by subsets $\mathfrak{A} \subseteq \mathcal{P}(X)$, to be understood as their upwards closure (this makes the representation more concise, so our decision problems become harder). Given two such representations $\mathfrak{A} \subseteq \mathcal{P}(X)$ and $\mathfrak{B} \subseteq \mathcal{P}(Y)$, a relation $S \subseteq X \times Y$ is a one-step Λ-simulation iff for all $A \in \mathfrak{A}$, there exists $B \in \mathfrak{B}$ such that $B \subseteq S[A]$, which can clearly be checked in polynomial time. Therefore, Λ-similarity for monotone modal logic is in P.

4 Universal simulands

We now introduce the second of our core technical notions, that of *universal simulands*. It is well-known that in general, models of a given modal formula may look very different, and will certainly not necessarily be related by simulations. *Lightweight* logics in the sense considered here are basically those where one does have smallest models under the simulation preorder. Formally:

Definition 23. Let ϕ be a positive Λ-formula. We say that a pointed coalgebra (C_ϕ, x_ϕ) is a *universal simuland for* ϕ if for any pointed coalgebra (D, y), $y \models_D \phi$ iff there is a Λ–simulation $S : C_\phi \to D$ with $x_\phi S y$.

Remark 24. Since identities are Λ-simulations, a universal simuland for ϕ must actually satisfy ϕ. Thus, by Proposition 16, (C_ϕ, x_ϕ) is a universal simuland for ϕ iff (i) $x_\phi \models \phi$ and (ii) whenever (D, y) is a pointed coalgebra such that $y \models_D \phi$, then (C_ϕ, x_ϕ) and (D, y) are Λ-similar.

The most obvious example where universal simulands fail to exist are disjunctive formulas; e.g. even for a purely propositional formula such as $a \vee b$ we will clearly not be able to find a universal simuland. Formally, we have

Proposition 25. *If $\phi \vee \psi$ has a universal simuland, then $\phi \sqsubseteq \psi$ or $\psi \sqsubseteq \phi$.*

It is thus no surprise that known examples of lightweight logics such as $\mathcal{EL}$ and $\mathcal{FL}_0$ exclude disjunction; also here, we will henceforth restrict attention to conjunctive formulas. However, if we move beyond the purely relational realm, even conjunctive formulas may fail to have canonical models:

Example 26. In graded logic, with $\Lambda = \{\Diamond_k \mid k \in \mathbb{N}\} \cup \{\Box_k \mid k \in \mathbb{N}\}$, consider the formula $\Diamond_0 a \wedge \Diamond_0 b \wedge \Diamond_0 c \wedge \Box_2 \bot$, specifying that one has successors satisfying a, b, and c, respectively, but only two successors altogether. Any model of this formula will need to include a successor satisfying either $a \wedge b$ or $b \wedge c$ or $a \wedge c$, but of course none of these choices will yield a universal simuland.

Worse, with the wrong choice of Λ, even $\top$ may fail to have a universal simuland: with $\Lambda = \{\Box, \Diamond\}$, interpreted over Kripke frames, a universal simuland for $\top$ would need to be *bisimilar* in the usual sense to any other state (see Example 18), which is impossible because states with successors fail to be bisimilar to deadlocked states.

Universal simulands will thus need to rely on a judicious choice of the modal signature Λ. We next complete the discussion of how universal simulands can be used to decide subsumption (without TBoxes for the time being; TBoxes will be discussed in Section 5). We will then proceed to analyse how universal simulands can be constructed from their one-step counterparts, and discuss examples.

Theorem 27. *If conjunctive formulas ϕ and ψ have universal simulands (C_ϕ, x_ϕ) and (C_ψ, x_ψ), then $\phi \sqsubseteq \psi$ iff there is a Λ-simulation $S : C_\psi \to C_\phi$ such that $x_\psi S x_\phi$.*

This property is the basis for tractable reasoning, if all goes well:

Corollary 28. *If every conjunctive Λ-formula has a polynomial-size universal simuland and the one-step Λ-simulation problem is in P, then subsumption between conjunctive Λ-formulas is in P.*

Universal simulands can be obtained from similar structures at the one-step level:

Definition 29 (One-step universal simulands). A one-step model (X, τ, t) is a *one-step universal simuland* for a one-step formula ϕ over V if, for every one-step model (Y, ϑ, s), $s \models_\vartheta \phi$ iff there exists a one-step Λ-simulation S between t and s such that $S[\tau(a)] \subseteq \vartheta(a)$ for all $a \in V$.

Somewhat analogously to Remark 24, we have

Lemma 30. *A one-step model (X, τ, t) is a one-step universal simuland for a one-step formula ϕ over V iff $t \models_\tau \phi$ and for every one-step model (Y, ϑ, s) such that $s \models_\vartheta \phi$, the relation $xSy \iff \breve{\tau}(x) \subseteq \breve{\vartheta}(y)$ is a one-step Λ-simulation between t and s.*

Remark 31. Universal simulands need not be unique if they exist. However, one can assume w.l.o.g. that in a universal simuland (X, τ, t), every $x \in X$ is uniquely determined by $\breve{\tau}(x)$ (quotient (X, τ, t) by the equivalence relation induced by $\breve{\tau}(y)$).

The main technical result of this section is then the following.

Definition 32. The T-structure Λ *admits (one-step) universal simulands* if every conjunctive (one-step) formula has a universal simuland.

Theorem 33. *If Λ admits one-step universal simulands, then Λ admits universal simulands.*

Proof ((sketch)). Induction on $rank(\phi)$. We have $\phi = \bigwedge_{i \in I} \heartsuit_i \chi_i$ for a finite (possibly empty) set I. Take $V_\phi = \{a_{\chi_i} \mid i \in I\}$ and decompose ϕ as $\phi = \tilde{\phi}\rho$ into a one-step formula $\tilde{\phi} = \bigwedge_{i \in I} \heartsuit_i a_{\chi_i}$ and a substitution $\rho(a_{\chi_i}) = \chi_i$. Let (X, τ, t) be a one-step universal simuland for $\tilde{\phi}$. By induction, we have, for each $x \in X$, a universal simuland (C_x, x) with $C_x = (Y_x, \xi_x)$ for $\bigwedge_{p \in \breve{\tau}(x)} \rho(p)$ with root x; w.l.o.g. the Y_x are pairwise disjoint. Pick a fresh x_ϕ, and let (C_ϕ, x_ϕ) be the resulting collage over (X, τ, t), with $C_\phi = (Y, \xi)$. Then (C_ϕ, x_ϕ) is a universal simuland for ϕ. $\square$

We note in passing that the theorem has, under mild additional assumptions, a converse:

Definition 34. A T-structure Λ is called non-trivial if there are infinitely many *independent* formulas $I = \{\chi_1, \chi_2, \ldots\}$, independent in the sense that for every finite subset H of I and every $\chi_i \notin H$, $H \not\models \chi_i$.

This assumption is obviously rather harmless; in particular, is satisfied whenever Λ contains infinitely many propositional atoms. Observe that any finite conjunction of independent formulas is by definition satisfiable.

Theorem 35. *If Λ is non-trivial and admits universal simulands, then Λ admits one-step universal simulands.*

For tractability, universal simulands are not enough, they also need to be small. We reiterate, however, that we believe that universal simulands are of independent interest, even when they are exponentially large; we do therefore later pay attention also to examples where universal simulands exist but fail to be small.

Definition 36. The T-structure Λ *admits small universal simulands* if every conjunctive Λ-formula ϕ has a polynomial-size universal simuland. Moreover, Λ *admits linear one-step universal simulands* if every conjunctive one-step Λ-formula over V has a polynomial-size one-step universal simuland (X, τ, t) such that $|\tau(a)| \leq 1$ for all $a \in V$ (then w.l.o.g. $|X| \leq |V|$). Here, sizes of formulas rely on a suitable representation for modal operators; in particular, we assume that numbers are coded in binary.

The point of *linear* one-step universal simulands is that in the construction of universal simulands, we never need to copy subformulas, so we avoid exponential blowup. The following criterion is, then, rather immediate.

Theorem 37. *If Λ admits linear one-step universal simulands, then Λ admits small universal simulands.*

Corollary 38. *If Λ admits linear one-step universal simulands and one-step Λ-simulation is in P, then subsumption between conjunctive Λ-formulas is in P.*

We now proceed to see some examples. It is not a surprise that universal simulands turn out to be a rare phenomenon requiring heavy restriction of the logic, and all the more so for small universal simulands. We do however recover the known examples, and come across some that are, to our current best knowledge, new. In the examples below, the term *subsumption checking* refers to reasoning over the empty TBox; TBox reasoning is discussed in the next section.

Example 39. Over Kripke frames, we have the following situations depending on Λ.

1. $\Lambda = \{\Diamond\}$ admits linear one-step universal simulands: for $\bigwedge_{i\in I} \Diamond a_i$, the one-step model (I, τ, I) with $\tau(a_i) = \{i\}$ is a linear one-step universal simuland. This extends to the multimodal case, of which $\mathcal{EL}$ [5] is a syntactic variant (up to the presence of propositional atoms, treated as in Example 3). We thus recover the known result that subsumption checking in $\mathcal{EL}$ is in P.

2. $\Lambda = \{\Box\}$ admits linear one-step universal simulands: for $\bigwedge_{i\in I} \Box a_i$, the singleton one-step model $(\{*\}, \tau, \{*\})$ with $\tau(a_i) = \{*\}$ is a linear one-step universal simuland. (This may seem surprising, but recall that $\Box$-simulations are converse simulations.) This extends straightforwardly to the multimodal case (for k modalities, one needs k states in a one-step universal simuland). The DL $\mathcal{FL}_0$ (see, e.g., [1]) is a syntactic variant of this, so we recover the known result that subsumption checking in $\mathcal{FL}_0$ is in P.

3. $\Lambda = \{\Box, \Diamond\}$ does not admit universal simulands, see Example 26.

4. We define a modality $\boxtimes$ by $\boxtimes\phi := \Box\phi \wedge \Diamond\top$, corresponding to the allsome operator occasionally considered in DLs [9]. Then $\Lambda = \{\boxtimes, \Diamond\}$ does admit one-step unversal simulands: for $\bigwedge_{i\in I} \boxtimes a_i \wedge \bigwedge_{j\in J} \Diamond b_j$, the one-step model $(J \cup \{*\}, \tau, J \cup \{*\})$ (with $* \notin J$) given by $\tau(a_i) = J \cup \{*\}$ and $\tau(b_j) = \{j\}$ is a one-step universal simuland. These one-step universal simulands are never linear unless we interdict $\Diamond$. We thus have that $\Lambda = \{\boxtimes, \Diamond\}$ *admits universal simulands* but do not obtain tractability. Nevertheless, we believe that the fact that conjunctive formulas using $\boxtimes$ and $\Diamond$ have universal simulands is of interest, certainly as a pleasant model-theoretic property, and potentially also computationally by identifying global restrictions on formulas that prevent exponential blow-up in universal simulands. (One obvious if maybe a bit high-handed restriction is to bound the number of modal depths at which diamonds are allowed to occur in formulas.)

Example 40. Over monotone neighbourhood frames, we have:

1. $\Lambda = \{\Box\}$ admits linear one-step universal simulands — for $\bigwedge_{i\in I} \Box a_i$, the one-step model $(I, \tau, \mathfrak{N})$ with $\tau(a_i) = \{i\}$ and $\mathfrak{N} = \{A \subseteq 2^I \mid A \neq \emptyset\}$ (the upwards

closure of the polynomial-sized system $\{\{i\} \mid i \in I\}$) is a one-step universal simuland. Comparison with Example 39 thus reveals that this logic is essentially the same as $\mathcal{EL}$. This equivalence, however, breaks down under slight variations, such as the following.

2. Take $\Lambda = \{\boxtimes\}$, again abbreviating $\boxtimes a = \Box a \wedge \Diamond\top$. For $\bigwedge_{i\in I} \boxtimes a_i$ with $I \neq \emptyset$ (the case $I = \emptyset$ is simpler), the one-step model $(I \cup \{*\}, \tau, \mathfrak{N})$ with $\tau(a_i) = \{i\}$ and $\mathfrak{N} = \{A \subseteq 2^{I\cup\{*\}} \mid A \neq \emptyset\}$ is a one-step universal simuland. We obtain that *subsumption between conjunctive* $\{\boxtimes\}$*-formulas in monotone modal logic is in* P, to our knowledge a new result, which does not seem to correspond to a variation of $\mathcal{EL}$. The modalitiy $\boxplus$ with $\boxplus a = \Box a \wedge \Diamond a$ can be treated similarly (with $\mathfrak{N} = \{A \subseteq 2^{I\cup\{*\}} \mid A \neq \emptyset \wedge * \in A\}$). The functor $\mathcal{M}$ of Example 5 has more predicate liftings than $\mathcal{P}$, and we expect that our approach will work for further variants.

Example 41. For graded logic, we have the following.

1. $\Lambda \subseteq \{\Diamond_k \mid k \in \mathbb{N}\}$ does not admit universal simulands unless $|\Lambda| \leq 1$, a boring case that is not really different from $\{\Diamond\}$ over Kripke frames. The reason for this is roughly that for a conjunct $\Diamond_k a$ in a conjunctive one-step Λ-formula one will need to have a multiplicity-$(k+1)$ element satisfying exactly a. But then for n such conjuncts $\Diamond_{k_i} a_i$, the model will satisfy $\Diamond_{-1+\sum(k_i+1)}\{a_1, \dots, a_n\}$, which will not necessarily hold in other models of $\bigwedge \Diamond_{k_i} a_i$; and this is visible to the logic unless $|\Lambda| \leq 1$.
2. $\Lambda = \{\Box_k \mid k \in \mathbb{N}\}$: For $\bigwedge_{i\in I} \Box_{k_i} a_i$, the one-step model (X, τ, b) with

$$
\begin{aligned}
X &= \{k_i \mid i \in I\} \cup \{\infty\} \\
\tau(a_i) &= \{n \in X \mid k_i < n\} \\
b(n) &= n - \max\{k \in X \mid k < n\}
\end{aligned}
$$

(where $\max \emptyset = 0$) is a one-step universal simuland, so that *conjunctive diamond-free formulas in graded modal logic have universal simulands.* One-step universal simulands for $\bigwedge_{i\in I} \Box_{k_i} a_i$ are linear iff $|\{k_i \mid i \in I\}| \leq 1$. Thus, $\{\Box_k\}$ *admits small universal simulands*. As states in universal simulands then effectively have only two successors, the problem with computing simulation discussed in Example 22 disappears, so that *subsumption checking between conjunctive graded formulas mentioning only one graded box is in* P, to our knowledge a new result. (Note that $\Lambda = \{\Box_k\}$ is not quite as boring as $\Lambda = \{\Diamond_k\}$; we can read $\Box_k\phi$ as 'almost necessarily ϕ', and universal simulands for $\{\Box_k\}$ look different from the ones for $\{\Box\}$.)

5 Terminologies

When coalgebraic logics are applied in knowledge representation, i.e. are seen as a form of generalized description logic, one will typically wish to reason with global assumptions; in description logic (DL) terms, this amounts to reasoning over a *terminology* or *TBox*, which specifies properties to be satisfied *everywhere* in a model. Formally, a *general TBox* is a finite set $\mathcal{T}$ of *equivalence axioms* of the form $\phi \equiv \psi$, where $\phi, \psi \in L(\Lambda)$. A model C *satisfies* $\mathcal{T}$ ($C \models \mathcal{T}$) if $[\![\phi]\!]_C = [\![\psi]\!]_C$ for every axiom $\phi \equiv \psi \in \mathcal{T}$. *Inclusion axioms* of the form $\phi \sqsubseteq \psi$ (as showcased in Fig. 1) can then be encoded as $\phi \wedge a \equiv \psi$ with a being a *fresh* proposition. Reasoning relative to $\mathcal{T}$ amounts to restricting attention to models satisfying $\mathcal{T}$.

We assume from now on that Λ is a disjoint union $\Lambda = \Lambda_p \cup \Delta$, where Δ consists of finitely many propositions (note that Λ_p may contain propositions as well), interpreted coalgebraically as indicated in Example 3. That is, we assume that T decomposes as $T = T_p \times 2^{\Delta}$, with Λ_p interpreted over the first projection and Δ over the second projection. We call Λ_p the set of *primitive* operators while the propositions in Δ are called *defined propositions*. Correspondingly, we will refer to T_p-coalgebras as *primitive*, and we will say that a T-coalgebra D is *based* on a primitive coalgebra C if $C = (X, \xi)$ arises from $D = (X, \zeta)$ by composing with the first projection $TX \to T_pX$; that is, D extends C by interpretations of the defined propositions. Moreover, instead of dealing with *general TBoxes*, we will restrict our attention to *classical* TBoxes that contain only definitions of the form $a \equiv \phi$ where $a \in \Delta$ and $\phi \in L(\Lambda)$, under the proviso that each a occurs on the left-hand-side of exactly one definition. There are three established interpretations of such TBoxes using least fixpoints, greatest fixpoints, and loose (*descriptive*) semantics, respectively [12]; here, we focus on greatest fixpoint semantics.

Formally, we define a partial order $\leq_C$ on coalgebras based on a common primitive coalgebra C by setting $D \leq_C E \iff [\![a]\!]_D \subseteq [\![a]\!]_E$ for all $a \in \Delta$. Clearly, this defines a complete lattice, so $\leq_C$-monotone functions have greatest fixpoints. Now let $f_{\mathcal{T},C}$ be the function mapping a coalgebra D based on C to the coalgebra E based on C such that $[\![a]\!]_E = [\![\phi]\!]_D$ for all $a \equiv \phi \in \mathcal{T}$ (this is well-defined because $\mathcal{T}$ is classical). Since all the operators in Λ_p are monotone, $f_{\mathcal{T},C}$ is $\leq_C$-monotone. Moreover, $D \models \mathcal{T}$ iff D is a fixpoint of $f_{\mathcal{T},C}$.

Definition 42. We say that a coalgebra D based on a primitive coalgebra C is a *gfp interpretation of* $\mathcal{T}$ if D is the greatest fixpoint of $f_{\mathcal{T},C}$. For $a, b \in \Delta$, we say that *a is subsumed by b in $\mathcal{T}$ under gfp semantics*, notation $a \sqsubseteq^g_{\mathcal{T}} b$, if $[\![a]\!]_D \subseteq [\![b]\!]_D$ for every gfp interpretation D of $\mathcal{T}$.

Since we can always add fresh definitions to a TBox, the restriction to elements of Δ in the definition of $\sqsubseteq^g_{\mathcal{T}}$ is without loss of generality. For examples in $\mathcal{EL}$ see, e.g., [2].

We will now study universal simulands for classical TBoxes, and establish tractability of subsumption checking if universal simulands are small. As is common in lightweight DLs, we will rely on a normal form of TBoxes that can be obtained in polynomial time [2].

Definition 43. A TBox $\mathcal{T}$ is *normalized* if $a \equiv \phi \in \mathcal{T}$ implies that ϕ is a conjunctive one-step Λ_p-formula over $V = \Delta$ (i.e. $\phi = \bigwedge_{i \in I} \heartsuit_i a_i$ with $\heartsuit_i \in \Lambda_p$ and $a_i \in \Delta$).

Proposition 44. *Every classical TBox $\mathcal{T}$ can be translated in polynomial time into a normalized $\mathcal{T}'$ (possibly containing fresh defined propositions) such that subsumption in $\mathcal{T}$ and in $\mathcal{T}'$ with respect to gfp semantics coincide.*

For the rest of this section, we *assume TBoxes to be normalized*. Now, if Λ_p admits one-step universal simulands, we can obtain, from every normalized TBox $\mathcal{T}$, a T_p-coalgebra $C_{\mathcal{T}}$ that will be a universal simuland for $\mathcal{T}$ in a sense to be made precise shortly.

We assume a fixed choice of a one-step universal simuland $(X_\phi, \tau_\phi, t_\phi)$ for each conjunctive one-step Λ_p-formula ϕ over $V = \Delta$, which we then call *the* one-step universal simuland of ϕ (recall that Δ is the set of defined propositions in $\mathcal{T}$). We assume w.l.o.g. that $X_\phi \subseteq \mathcal{P}(\Delta)$ and $\tau_\phi(a) = \{A \in X_\phi \mid x \in A\}$ (Remark 31). We then construct the carrier $X_\mathcal{T}$ of $C_\mathcal{T}$ as a subset of $\mathcal{P}(\Delta)$. For $A \subseteq \Delta$, we let ϕ_A denote the conjunction of all right-hand sides of equivalences $a = \phi$ in $\mathcal{T}$ with $a \in A$ (i.e. ϕ_A is a conjunctive one-step Λ_p-formula over Δ, the expansion of $\bigwedge A$ according to $\mathcal{T}$). Then, $X_\mathcal{T}$ is the smallest subset of Δ such that

$$X_{\phi_A} \subseteq X_\mathcal{T} \quad \text{for each } A \in X_\mathcal{T}.$$

We define a T_p-coalgebra structure $\xi_\mathcal{T}$ on $X_\mathcal{T}$ by

$$\xi_\mathcal{T}(A) = T(i_A) t_{\phi_A}$$

where i_A is the inclusion $X_{\phi_A} \hookrightarrow X_\mathcal{T}$. We then have a universality property analogous to the one established for $\mathcal{EL}$ in [2]:

Theorem 45. *If Λ admits one-step universal simulands, then for every normalized TBox $\mathcal{T}$, $C_\mathcal{T}$ as constructed above is a* universal simuland*, i.e. for any gfp interpretation D of $\mathcal{T}$ based on a T_p-coalgebra C, any state x in D, and any $a \in \Delta$, $x \models_D a$ iff there is a Λ_p-simulation $S : C_\mathcal{T} \to D$ such that $\{a\} S x$. Consequently, for $a, b \in \Delta$, $a \sqsubseteq^g_\mathcal{T} b$ iff there exists a Λ_p-simulation $S : C_\mathcal{T} \to C_\mathcal{T}$ such that $\{b\} S \{a\}$.*

Thus, those of the logics listed as having one-step universal simulands in Examples 39–41 have universal simulands for normalized TBoxes under gfp semantics, among them the conjunctive fragments of multimodal K with only boxes, only diamonds, and $\{\boxtimes, \Diamond\}$, respectively, as well as conjunctive graded modal logic with only boxes. This does not yet imply tractability, even when Λ_p admits linear one-step universal simulands, since the closure process defining $X_\mathcal{T}$ may still produce an exponentially large set. One very simple (and limitative) criterion for smallness of $X_\mathcal{T}$ is the following.

Theorem 46. *If every one-step Λ_p-formula ϕ over $V = \Delta$ has a one-step universal simuland (X, τ, t) such that $|\breve{\tau}(x)| \leq 1$ for all $x \in X$, then every normalized TBox $\mathcal{T}$ has a universal simuland $C_\mathcal{T} = (X_\mathcal{T}, \xi_\mathcal{T})$ such that $|X_\mathcal{T}| \leq |\Delta| + 1$. If additionally one-step Λ-similarity is in P, then subsumption checking over classical TBoxes in the conjunctive fragment of $L(\Lambda)$ is in P.*

Example 47. By the above criterion and the description of one-step universal simulands in Examples 39–41, we regain the known result that subsumption checking over classical TBoxes with gfp semantics in $\mathcal{EL}$ is in P [2]. Moreover, we obtain the (to our knowledge, new) result that *subsumption checking over classical TBoxes with gfp semantics in the conjunctive fragment of the monotone modal logic of* $\boxtimes$ (with $\boxtimes a = \Box a \wedge \Diamond\top$) *is in P*, similarly for the modality $\boxplus a = \Box a \wedge \Diamond a$.

For the remaining cases with linear one-step universal simulands, we need suitable restrictions on $\mathcal{T}$ and on one-step universal simulands to guarantee smallness of $C_\mathcal{T}$.

One strong condition on one-step universal simulands (X, τ, t) is to require that there is at most one $x \in X$ such that $\breve{\tau}(x) \neq \emptyset$. By the description of one-step universal simulands in Examples 39 and 41, this includes two of our remaining examples, the conjunctive fragment of K with only boxes and the conjunctive fragment of graded modal logic with only one graded box $\Box_k$. It *excludes* the multimodal version of the former, i.e. $\mathcal{FL}_0$, and in fact reasoning over even the most restrictive (i.e. acyclic, see below) TBoxes in $\mathcal{FL}_0$ is known to be at least $coNP$-hard [12].

Additionally, we need to restrict $\mathcal{T}$. We define a relation U ('uses') on Δ by aRb iff b occurs in the (unique) ϕ such that $a = \phi \in \mathcal{T}$. Following standard terminology, $\mathcal{T}$ is *acyclic* if U is acyclic. Moreover, call $\mathcal{T}$ *reflexive* if U is reflexive; roughly, this means that the propositions defined by $\mathcal{T}$ as greatest fixpoints must hold 'without gaps', i.e. are always inherited down to direct successor states (rather than, say, only to grandchildren). Under the above restriction on one-step universal simulands, $C_{\mathcal{T}}$ is easily seen to be polynomial-sized if $\mathcal{T}$ is either acyclic or reflexive; summing up: *Subsumption checking in gfp semantics over acyclic or reflexive TBoxes in the conjunctive fragments of K with only boxes or graded modal logic with only one graded box is in P*, both to our knowledge new results. The condition 'acyclic or reflexive' can in fact be relaxed to require only that all $a \in \Delta$ that are on a U-loop are on a U-loop of bounded length.

6 Conclusions

We have developed the basic concepts of *lightweight coalgebraic logics*. These are conjunctive fragments of coalgebraic modal logics characterized by allowing for *universal simulands* associated to formulas in such a way that satisfaction of a formula is equivalent to simulation of its universal simuland. Although additional restrictions are needed to obtain smallness of universal simulands and hence tractable reasoning, we have the impression that universal simulands are of independent interest as a model-theoretic phenomenon, and may prove to be computationally useful even in cases where smallness does not hold in general. We have established that various logics of interest admit universal simulators, including the diamond-free conjunctive fragment of graded modal logic and the conjunctive fragment of multimodal K (i.e. of $\mathcal{ALC}$).

By an additional analysis of the size of universal simulands, we can establish that checking subsumption (i.e. local consequence) is in P under suitable additional restrictions, occasionally even over so-called classical TBoxes with gfp semantics. We thus recover known results for the relational logics $\mathcal{EL}$ and (without TBoxes) $\mathcal{FL}_0$, and moreover obtain several new tractability results, specifically for

- the conjunctive fragment of the monotone modal logic of $\boxtimes$ (where $\boxtimes a = \Box a \wedge \Diamond\top$ — a logic of reasonable parents that do not win arguments advocating absurd propositions) with classical TBoxes under gfp semantics; similarly for $\boxplus$ defined by $\boxplus a = \Box a \wedge \Diamond a$ in place of $\boxtimes$.
- the diamond-free conjunctive fragment of K over acyclic or *reflexive* TBoxes; and
- the conjunctive fragment of graded modal logic with only one graded box over acyclic or reflexive TBoxes.
- the conjunctive diamond-free fragment of K over reflexive TBoxes, which may be thought of as a logic of safety properties in transition systems (essentially the fragment of the μ-calculus with only $\Box$, $\wedge$, and ν).

We expect that (even without size bounds), universal simulators can be used to establish the existence of least common subsumers. Another core issue for future research is lightweight reasoning over *general* TBoxes (i.e. finite sets of arbitrary inclusion axioms), which is known to remain tractable in the case of $\mathcal{EL}$ [8].

References

1. F. Baader. Using automata theory for characterizing the semantics of terminological cycles. *Ann. Math. Artif. Intell.*, 18:175–219, 1996.
2. F. Baader. Terminological cycles in a description logic with existential restrictions. In *International Joint Conference on Artificial Intelligence, IJCAI 2003*, pp. 325–330. Morgan Kaufmann, 2003.
3. F. Baader, S. Brandt, and C. Lutz. Pushing the $\mathcal{EL}$ envelope. In *International Joint Conference on Artificial Intelligence, IJCAI 2005*. Morgan-Kaufmann, 2005.
4. F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, eds. *The Description Logic Handbook*. Cambridge University Press, 2003.
5. F. Baader, R. Küsters, and R. Molitor. Computing least common subsumers in description logics with existential restrictions. In *International Joint Conference on Artificial Intelligence, IJCAI 1999*, pp. 96–101. Morgan Kaufmann, 1999.
6. M. Barr. Terminal coalgebras in well-founded set theory. *Theoret. Comput. Sci.*, 114:299–315, 1993.
7. P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.
8. S. Brandt. Polynomial time reasoning in a description logic with existential restrictions, GCI axioms, and – what else? In *Eureopean Conference on Artificial Intelligence, ECAI 2004*, pp. 298–302. IOS Press, 2004.
9. R. Dionne, E. Mays, and F. Oles. The equivalence of model-theoretic and structural subsumption in description logics. In *International Joint Conference on Artificial Intelligence, IJCAI 1993*, pp. 710–717. Morgan Kaufmann, 1993.
10. D. Gorín and L. Schröder. Simulations and bisimulations for coalgebraic modal logics. In *Algebra and Coalgebra in Computer Science, CALCO 2013*, vol. 8089 of *LNCS*, pp. 253–266. Springer, 2013.
11. R. Myers, D. Pattinson, and L. Schröder. Coalgebraic hybrid logic. In *Foundations of Software Science and Computational Structures, FOSSACS 2009*, vol. 5504 of *LNCS*, pp. 137–151. Springer, 2009.
12. B. Nebel. Terminological reasoning is inherently intractable. *Artificial Intelligence*, 43:235–249, 1990.
13. R. Parikh. Propositional game logic. In *Foundations of Computer Science, FOCS 1983*, pp. 195–200. IEEE Computer Society, 1983.
14. D. Pattinson. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoret. Comput. Sci.*, 309:177–193, 2003.
15. L. Schröder and D. Pattinson. Shallow models for non-iterative modal logics. In *Advances in Artificial Intelligence, KI 2008*, vol. 5243 of *LNCS*, pp. 324–331. Springer, 2008.

Appendix: Omitted Proof Details

Proof of Lemma 9 (Collage lemma) The second equivalence follows directly from naturality of $\heartsuit$. For the first one, one proceeds by induction on ϕ; the relevant case is the modal one, for which we have:

$$\begin{aligned} x \models_C \heartsuit\psi &\iff T(\hookrightarrow_{Y_x})(\xi_x(x)) \in \heartsuit_Y [\![\psi]\!]_C \\ &\iff \xi_x(x) \in \heartsuit_{Y_x}([\![\psi]\!]_C \cap Y_x) && \text{(naturality)} \\ &\iff \xi_x(x) \in \heartsuit_{Y_x}[\![\psi]\!]_{C_x} && \text{(IH)} \\ &\iff x \models_{C_x} \heartsuit\psi. \end{aligned}$$

□

Proof of Corollary 11 The first case follows directly from the collage lemma. For the second one we prove the positive case:

$$\begin{aligned} t \models_\tau \heartsuit p &\iff t \in \heartsuit_X \tau(p) \\ &\implies t \in \heartsuit_X([\![\rho(p)]\!]_C \cap X) && (\tau(p) \subseteq [\![\rho(p)]\!]_C \text{ by 1 + monotonicity}) \\ &\iff \xi(r) \in \heartsuit_Y [\![\rho(p)]\!]_C && \text{(collage lemma)} \\ &\iff r \models_C \heartsuit\rho(p). \end{aligned}$$

□

Proof of Proposition 16 Let $(C_{\phi\vee\psi}, x_{\phi\vee\psi})$ be a canonical model for $\phi \vee \psi$; in particular, $x_{\phi\vee\psi} \models \phi \vee \psi$, which means that $x \models \phi$ or $x \models \psi$. W.l.o.g. assume the former. Now, take any (D, y) with $y \models \phi\vee\psi$; since $(C_{\phi\vee\psi}, x_{\phi\vee\psi})$ and (D, y) are Λ-similar, we conclude (using Lemma 16) that $y \models_D \phi$. Hence, $\phi \vee \psi \sqsubseteq \phi$ and therefore $\psi \sqsubseteq \phi$. □

Proof of Proposition 25 Induction over the formula structure, with trivial Boolean cases (noting that these do not include negation). For the modal case, we have

$$\begin{aligned} x \models_C \heartsuit\rho &\iff \xi(x) \models \heartsuit[\![\rho]\!]_C \\ &\implies \zeta(y) \models \heartsuit S[[\![\phi]\!]_D] && S\ \Lambda\text{-}\mathfrak{A}\text{-simulation} \\ &\implies \zeta(y) \models \heartsuit[\![\phi]\!]_D && \text{induction, monotonicity} \\ &\iff y \models_\zeta \heartsuit\phi. \end{aligned}$$

□

Proof of Theorem 27 *'Only if'*: We have $x_\phi \models_{C_\phi} \phi$, so by assumption $x_\phi \models_{C_\psi} \psi$, so that we are done by the definition of universal simuland.

'If': Let (D, y) be a pointed coalgebra with $y \models_D \phi$. By the definition of universal simuland, (C_ϕ, x_ϕ) and (D, y) are Λ-similar. By Proposition 15, Λ-similarity is transitive, so that by assumption it follows that (C_ψ, x_ψ) and (D, y) are similar; by Proposition 16, $y \models_D \psi$. □

Proof of Theorem 33 Induction on $rank(\phi)$. We have $\phi = \bigwedge_{i \in I} \heartsuit_i \chi_i$ for a finite (possibly empty) index set I. Take $V_\phi = \{a_{\chi_i} \mid i \in I\}$ and decompose ϕ as $\phi = \tilde{\phi}\rho$ into a one-step formula $\tilde{\phi} = \bigwedge_{i \in I} \heartsuit_i a_{\chi_i}$ and a substitution $\rho(a_{\chi_i}) = \chi_i$. Let (X, τ, t) be a one-step universal simuland for $\tilde{\phi}$. By induction, we have, for each $x \in X$, a universal simuland (C_x, x) with $C_x = (Y_x, \xi_x)$ for $\bigwedge_{p \in \check{\tau}(x)} \rho(p)$ with root x. We assume w.l.o.g. that the Y_x are pairwise disjoint. Pick a fresh x_ϕ, and let (C_ϕ, x_ϕ) be the resulting collage over (X, τ, t), with $C_\phi = (Y, \xi)$.

We claim that (C_ϕ, x_ϕ) is a universal simuland for ϕ. By construction, (C_ϕ, x_ϕ) positively fulfills ρ so that $x_\phi \models_C \phi$ by Corollary 11. It remains to show that given a coalgebra $D = (Z, \zeta)$ and $z_0 \in Z$ such that $z_0 \models_D \phi$, x_ϕ and z_0 are Λ-similar. Define a one-step model $(Z, \vartheta, \zeta(z_0))$ by putting $\vartheta(a) = [\![\rho(a)]\!]_D$. Then $\zeta(z_0) \models_\vartheta \tilde{\phi}$, so we have a one-step Λ-simulation S between t and $\zeta(z_0)$ such that $S[\tau(a)] \subseteq \vartheta(a)$ for all a. This implies that whenever xSz then $z \models_D \bigwedge_{a \in V_\phi} \rho(a)$, so that there exists a Λ-simulation S_{xz} between (C_x, x) and (D, z). Then $R = \{(x_\phi, z_0)\} \cup \bigcup_{xSz} S_{xz}$ is a Λ-simulation between (C_ϕ, x_ϕ) and (D, z_0). $\square$

Proof of Theorem 35 Assume a satisfiable one-step formula $\phi = \bigwedge_{i \in I} \heartsuit_i p_i$ is given. We identify each p_i with one of the infinitely many independent formulas that exist since Λ is non-trivial and thus set $\rho(p_i) := \chi_i$. The first thing to note is that $\phi\rho$ is satisfiable; this follows directly from the Corollary 11 and the fact that every (finite) conjunction of independent formulas is satisfiable. Hence, let $(C_{\phi\rho}, x_{\phi\rho})$ be a canonical model for $\phi\rho$, with $C_{\phi\rho} = (X, \zeta)$, and let (X, τ, t) be its décollage following ρ which, by the décollage lemma satisfies ϕ. In order to see that it is canonical we need to verify that it is one-step simulated by every other model for ϕ.

So let (Y, σ, s) be such that $\sigma, s \models \phi$. As before, we build a collage (D, r) over (Y, σ, s), with $D = (Z, \xi)$, that positively fulfills ρ but we insist that for each $y \in Y$, the rooted coalgebra (D_y, y) used in the collage, with $D_y = (Z_y, \xi_y)$, need be canonical for $\bigwedge_{p \in \check{\sigma}(y)} \rho(p)$. The canonicity condition guarantees that for $q \notin \check{\sigma}(y)$, $y \not\models_{C_y} \rho(q)$ (otherwise we would have $\bigwedge_{p \in \check{\sigma}(y)} \rho(p) \models \rho(q)$, which would violate the independence assumption) and, by the collage lemma, $y \not\models_C \rho(q)$. We thus conclude that, for $y \in Y$, $y \in \sigma(p) \iff y \models_C \rho(p)$. Moreover, by Corollary 11, $r \models_C \phi\rho$, so by canonicity of $(C_{\phi\rho}, x_{\phi\rho})$, there is a simulation $S : C_{\phi\rho} \to D$ with $x_{\phi\rho} S r$.

Now, for $A \subseteq X$ and $\heartsuit \in \Lambda$, we have:

$$
\begin{aligned}
t \in \heartsuit_X A \;&\Longrightarrow \zeta(x_{\phi\rho}) \in \heartsuit_X A \\
&\Longrightarrow \xi(r) \in \heartsuit_{Y'} S[A] && \text{(canonicity)} \\
&\Longleftrightarrow T(\hookrightarrow_Y)(s) \in \heartsuit_{Y'} S[A] && \text{(collage)} \\
&\Longleftrightarrow s \in \heartsuit_Y (S[A] \cap Y) && \text{(naturality)} \\
&\Longrightarrow s \in \heartsuit_Y T_{\tau,\sigma}[A] && \text{(monotony + } S[A] \cap Y \subseteq T_{\tau,\sigma}\text{)}
\end{aligned}
$$

We only need to verify that $S \cap (X \times Y) \subseteq T_{\tau,\sigma}$, which was used on the last step. For the sake of contradiction, assume that for $x \in X$ and $y \in Y$, we have xSy but it is not the case that $xT_{\tau,\sigma}y$. This means that there is a p such that $x \in \tau(p)$ but $y \notin \sigma(p)$.

By construction (as shown above), the latter implies that $y \not\models_C \rho(p)$; but we also have $x \models_\zeta \rho(p)$ and xSy and hence a contradiction. $\square$

Proof of Proposition 44 The corresponding translation for $\mathcal{EL}$ [2] transfers straightforwardly to the general coalgebraic case, as its correctness depends only on replacement of equivalents under modalities, which is a valid proof principle in coalgebraic logic.

$\square$

Proof of Theorem 45 We prove the first claim; the second is an easy corollary.

'If': We first show that when we extend $C_\mathcal{T}$ to a T-coalgebra $D_\mathcal{T}$ by putting $A \models_{D_\mathcal{T}} a \iff a \in A$ (recall that $X_\mathcal{T}$ consists of subsets of Δ) then

$$D_\mathcal{T} \text{ is a post-fixpoint of } f_{\mathcal{T},C^0_\mathcal{T}}. \tag{3}$$

i.e. for all $a = \phi \in \mathcal{T}$,

$$[\![a]\!]_{D_\mathcal{T}} \subseteq [\![\phi]\!]_{D_\mathcal{T}}.$$

So let $A \models_{D_\mathcal{T}} a$, i.e. $a \in A$. We have to show $A \models_{D_\mathcal{T}} \phi$, i.e. $\xi_\mathcal{T}(A) \models_\tau \phi$ where $\tau(a) = [\![a]\!]_{D_\mathcal{T}} = \{B \in X_\mathcal{T} \mid a \in B\}$ (recall here that ϕ is a one-step Λ_p-formula over Δ). Now $\xi_\mathcal{T}(A) = T(i_A)t_{\phi_A}$, so by naturality of predicate liftings our goal is equivalent to $t_{\phi_A} \models_{\tau_{\phi_A}} \phi$, as $\tau_{\phi_A}(a) = \tau(a) \cap X_A$. The latter follows from the fact that $t_{\phi_A} \models_{\tau_{\phi_A}} \phi_A$ by construction, since ϕ is a conjunct of ϕ_A.

Now we proceed by coinduction: We define a T-coalgebra D^S based on C by putting for $a \in \Delta$

$$x \models_{D^S} a \iff (ASx \text{ for some } A \in X_\mathcal{T} \text{ such that } a \in A).$$

Since D is the gfp of $f_{\mathcal{T},C}$, we are done once we show that D^S is a post-fixpoint of $f_{\mathcal{T},C}$; that is, we have to show that for $a \in \Delta$ and the unique $a = \phi \in \mathcal{T}$, we have

$$[\![a]\!]_{D^S} \subseteq [\![\phi]\!]_{D^S}.$$

So let $x \models_{D^S} a$, i.e. we have $A \in X_\mathcal{T}$ such that $a \in A$ and ASx. We have to show that $x \models_{D^S} \phi$. Since by construction, $S : D_\mathcal{T} \to D^S$ is a Λ-simulation, it suffices to show that $A \models_{D_\mathcal{T}} \phi$; this however is immediate from (6).

'Only if': Let D be a gfp interpretation of $\mathcal{T}$ based on $C = (Y, \zeta)$, and let $y \models_D a$. Define a relation $S \subseteq X_\mathcal{T} \times Y$ by

$$ASx \iff x \models_D \bigwedge A.$$

Then clearly $\{a\}Sx$; we claim that S is a Λ_p-simulation. So let ASx, i.e. $x \models_D \bigwedge A$, and let $\xi_\mathcal{T}(A) \models \heartsuit\mathfrak{B}$ for some $\heartsuit \in \Lambda$, $\mathfrak{B} \subseteq X_\mathcal{T}$. We then have to show $\zeta(x) \models \heartsuit S[\mathfrak{B}]$. Since $\xi_\mathcal{T}(A) = Ti_A(t_{\phi_A})$, we have $t_{\phi_A} \models \heartsuit(\mathfrak{B} \cap X_{\phi_A}$ by naturality of predicate liftings. Since D is a (gfp) interpretation of $\mathcal{T}$, we have $x \models_D \phi_A$ and hence $\zeta(x) \models_\vartheta \phi_A$ where $\vartheta(a) = [\![a]\!]_D$ for all $a \in \Delta$. But $(X_{\phi_A}, \tau_{\phi_A}, t_{\phi_A})$ is a universal simuland for ϕ_A, so it follows that $\zeta(x) \models \heartsuit R[\mathfrak{B} \cap X_{\phi_A}]$ where $BRy \iff B \subseteq \breve{\vartheta}(y)$ (recall that $\tau_{\phi_A}(a) = \{B \in X_{\phi_A} \mid a \in B\}$). Now $R[\mathfrak{B} \cap X_{\phi_A}] \subseteq S[\mathfrak{B}]$: when $B \in \mathfrak{B} \cap X_{\phi_A}]$ and BRy, then $y \models_\vartheta \bigwedge B$, hence $y \models_D \bigwedge B$ so that BSy. By monotonicity, we obtain $\zeta(x) \models \heartsuit S[\mathfrak{B}]$ as required. $\square$