Subsumption Checking in Conjunctive Coalgebraic Fixpoint Logics

Daniel Gorín Lutz Schröder

Friedrich-Alexander-Universität Erlangen-Nürnberg

Abstract

While reasoning in a logic extending a complete Boolean basis is coNP-hard, restricting to conjunctive fragments of modal languages sometimes allows for tractable reasoning even in the presence of greatest fixpoints. One such example is the $\mathcal{EL}$ family of description logics; here, efficient reasoning is based on satisfaction checking in suitable small models that characterize formulas in terms of simulations. It is well-known, though, that not every conjunctive modal language has a tractable reasoning problem. Natural questions are then how common such tractable fragments are and how to identify them. In this work we provide sufficient conditions for tractability in a general way by considering unlabeled tableau rules for a given modal logic. We work in the framework of coalgebraic modal logics as unifying semantic setting. Apart from recovering known results for description logics such as $\mathcal{EL}$ and $\mathcal{FL}_0$, we obtain new ones for conjunctive fragments of relational and non-relational modal logics with greatest fixpoints. Most notably we find tractable fragments of game logic and the alternating-time μ-calculus.

Keywords: Materializers, convexity, tractable reasoning, fixpoints.

1 Introduction

The complexity of reasoning in logics extending a complete Boolean basis is at least coNP. For modal logics, it is typically even harder: already the basic (multi-)modal logic K_m is PSPACE-complete [21] and if fixed points are added to the mix, the complexity typically goes up to at least EXPTIME which, e.g., is the complexity of PDL and the μ-calculus [15]. Practical reasoning in these logics requires highly optimized heuristic strategies and will ultimately have only a limited degree of scalability.

This motivated the study of fragments in which core reasoning problems become tractable, i.e., decidable in polynomial time. Such fragments typically exclude negation and disjunction. Perhaps the best-known example is the $\mathcal{EL}$ family of *lightweight description logics*, where also universal restrictions (i.e. $\Box$-modalities) are excluded. In the absence of negation, satisfiability is no longer the central reasoning problem, being in fact often trivial (e.g., in $\mathcal{EL}$ every formula is satisfiable). Instead, one focuses on the entailment problem, alternatively called *subsumption checking* in the DL community. Indeed $\mathcal{EL}$

turns out to have a polynomial-time subsumption problem [4,6], even when extended with greatest fixed points [22]. Despite the limited syntax, $\mathcal{EL}$ can in practice accommodate large ontologies such as Snomed ct.

Rather surprisingly, the subsumption problem of $\mathcal{FL}_0$, the counterpart of $\mathcal{EL}$ with universal instead of existential restriction, becomes intractable when greatest fixed points (or even just non-recursive global definitions, i.e. acyclic TBoxes) are added to the language [5,8,26]. This shows that there is more to lightweight logics than just dropping disjunctions. Here, we aim to develop conceptual tools to identify lightweight modal formalisms beyond the purely relational realm. For uniformity, we work in the setting of *coalgebraic modal logic* [28], where the notions of model and modal operators are suitably abstracted. We then state and prove a general version of each result and just obtain the featured instances as corollaries.

Tractability of relational lightweight logics exploits the existence of what are called *materializations* of formulas (which moreover need to be computable and small). (Alternatively, tractability can be shown by proof-theoretic methods [19].) A materialization for ϕ is a model that satisfies *only* the formulas that ϕ entails; thus, subsumption can be reduced to model-checking in a materialization [22]. Moreover, while there seems to be a strong connection between tractability of subsumption for a given fragment and *convexity* of its formulas, meaning that they imply at least one of the disjuncts of every disjunction they entail, the precise nature of such a connection is still partly unclear (see, e.g., [20,23]).

For coalgebraic logics, we show that a stronger (infinitary) version of convexity of their conjunctive fragments is actually equivalent to the existence of materializations (for the relational description logic $\mathcal{ALCFI}$, a more fine-grained connection at the level of TBoxes has been established by Lutz and Wolter [23]). As in the relational case, our materializations moreover have an even stronger property — they can be taken as complete replacements of the materialized formulas, in the sense that the satisfaction relation for the former corresponds to a similarity relation for the latter in a sense we developed recently [16]. However, the mere existence of materializations is not enough for tractability; one requires additional conditions that ensure that materializations for a given conjunctive fragment of a modal logic $\mathcal{L}$ can be obtained in polynomial time. For this, we develop a simple syntactic criterion based on the set of (unlabeled) tableau rules for $\mathcal{L}$, and show how to compute small materializations from them even in the presence of greatest fixpoints. With this result we can show tractability of several conjunctive modal (fixpoint) logics, including fragments of game logic [27] and the alternating-time μ-calculus [2].

Proofs are sketched or omitted; more technical details can be found in the arXiv preprint 1401.6359.

2 Preliminaries

We first present the various concrete logics that will serve as case studies throughout the paper and then briefly introduce the basic concepts of coalgebraic logic that are used in the generic development. For each concrete logic we also consider its reasoning principles, in the form of unlabeled tableau-rules $\Gamma_0/\Gamma_1 \mid \cdots \mid \Gamma_n$, where the Γ_i are sets of formulas. A set of rules $\mathcal{R}$ is meant to be used in the usual way: to show satisfiability of a set Γ (interpreted conjunctively), one needs to show the satisfiability of at least one conclusion of every rule in $\mathcal{R}$ applicable to Γ. All tableau systems are understood to extend a set of propositional rules.

Basic modal logic. We assume the reader to be familiar with the syntax and semantics of the basic modal logic K interpreted over Kripke models. We shall also consider its restriction KD to *serial* models, where every node has at least one successor, making $\Diamond\top$ valid. The set of rules $\mathcal{R}_K = \{K_n : n \geq 0\}$ (Fig. 1) induces a complete tableau system for K. For KD one needs to add to $\mathcal{R}_K$ the rules D_n for $n \geq 0$.

Monotone neighbourhood logic. The minimal monotone logic M uses the same language as K but is interpreted over monotone neighbourhoods, i.e., neighbourhood models where the set of neighbourhoods of each point is upwards closed w.r.t. set inclusion [13]. We read $\Box\phi$ as 'there is a neighbourhood where ϕ holds'. It is well known that this logic can be encoded in K, replacing $\Box$ with $\Diamond\Box$ and $\Diamond$ with $\Box\Diamond$ (e.g. [27]). A complete tableau system for M is obtained simply by taking rule K_1.

Here too, we will be interested in the *serial* case which corresponds to the case where $\Box\top$ and $\Diamond\top$ are taken as axioms; we shall denote the resulting logic by M_s. Serial monotone neighbourhood frames underlie the semantics of game logic [27], discussed in more detail in Section 6. Seriality means that each state has some neighbourhood and the empty set is never a neighbourhood. Notice that in the mentioned encoding of monotone modal logic into normal modal logic, serial monotone neighbourhood frames correspond exactly to serial Kripke frames. It is easy to see that the set of rules $\mathcal{R}_{M_s} = \{K_1, K_0, D_1\}$ is a complete tableau system for M_s (notice that K_0 and D_1 are just the instances of K_1 for $\Box\top$ and $\Diamond\top$, respectively).

Coalition logic and alternating-time logics. Coalition logic [29] is essentially the next-step fragment of the alternating-time μ-calculus AMC [2], discussed in Section 6. A *coalition* is a subset of a fixed set $N = \{1, \ldots, n\}$ of agents and one has a modal operator $[C]$ for each coalition C. Intuitively, we read formula $[C]\phi$ as 'coalition C has a joint strategy to enforce that ϕ shall hold in the next state'. Formally, the semantics is over *game frames*, where for each state x we have a function f_x with domain $S_1 \times \cdots \times S_n$, each S_q being a finite set of actions available to agent $q \in N$ in state x. Intuitively, the choice of an action by each agent determines a successor state as specified by the *outcome function* f_x. One then defines the semantics of $[C]$ by putting $x \models [C]\phi$ iff there exists a joint choice $(s_q)_{q\in C}$ of actions for the agents in C such that for each joint choice $(s_q)_{q\in N-C}$ for the agents outside C, $f((s_q)_{q\in N}) \models \phi$. Note that each choice of N defines a different logic CL_N (in the sense that extending CL_{N_0} to CL_{N_1} for $N_0 \subsetneq N_1$ does *not* preserve subsumption), since the seman-

$$K_n \ \frac{\Box a_1, \ldots, \Box a_n, \Diamond b}{a_1, \ldots, a_n, b} \qquad\qquad D_n \ \frac{\Box a_1, \ldots, \Box a_n}{a_1, \ldots, a_n}$$

$$C_{nm} \ \frac{[C_1]a_1, \ldots, [C_n]a_n, \langle D\rangle b, \langle N\rangle c_1, \ldots, \langle N\rangle c_m}{a_1, \ldots, a_n, b, c_1, \ldots, c_m} \ \dagger\ddagger \qquad C'_n \ \frac{[C_1]a_1, \ldots, [C_n]a_n}{a_1, \ldots, a_n} \ \dagger$$

Fig. 1. Tableau rules, with side conditions: (†) $i \neq j \Rightarrow C_i \cap C_j = \emptyset$, and (‡) $C_i \subseteq D$.

tics of $[C]$ depends on how many agents there are outside C. For a fixed N, the set of rules $\mathcal{R}_{\mathrm{CL}_N} = \{C_{ij}, C'_k : i, j \geq 0, k > 0\}$ (Fig. 1) yields a complete tableau system for CL [14,32].

We include only the basic definitions of coalgebraic logic, which is more comprehensively presented elsewhere [28,31,34]. The generality of the framework stems from the parametricity of its syntax and semantics. The language depends on a *similarity type* Λ, which may include atomic propositions, seen as modalities of arity 0. To simplify notation, we pretend that all modal operators are unary. The grammar for the set $L(\Lambda)$ of *positive* Λ*-formulas* is

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

The set of *conjunctive* Λ-formulas is obtained by dropping the clauses for $\bot$ and $\vee$ from the grammar above. When $\mathcal{L}$ is a logic, we refer to the restriction of $\mathcal{L}$ to conjunctive formulas as *conjunctive* $\mathcal{L}$.

Given a modality $\heartsuit \in \Lambda$ we use $\bar{\heartsuit}$ to denote the *dual* of $\heartsuit$, with $\bar{\heartsuit}\phi$ interpreted as $\neg\heartsuit\neg\phi$ (under the usual meaning of $\neg$); we also use $\bar{\Lambda} := \{\bar{\heartsuit} : \heartsuit \in \Lambda\}$. We do *not* assume that Λ is closed under duals, as inclusion or non-inclusion of dual operators in Λ usually makes a big difference for the existence and size of materializations (Section 4).

The semantics is parametrized, first, in terms of an endofunctor T on the category Set of sets and maps, which determines the class of models. For a fixed T, a *model* is then just a T*-coalgebra* $C = (X, \xi)$, consisting of a set X (of *states*) and a *transition function* $\xi : X \to TX$. A *pointed model* is a pair (C, r), where r is a state of C, called the *point* or *root*. The intuition here is that $\xi(x)$ is the *local view* of the model standing on a state x; e.g., in a Kripke model, $\xi(w)$ would consist of the set of immediate successors of world w, plus the set of propositions that hold at w; thus, the class of all Kripke models arises as the class of all T-coalgebras for the functor $TX = \mathcal{P}(X) \times \mathcal{P}(\mathsf{Prop})$. As usual, we assume w.l.o.g. that T is non-trivial, i.e. $TX = \emptyset \implies X = \emptyset$ (otherwise, $TX = \emptyset$ for all X) and preserves subsets, i.e. $TX \subseteq TY$ whenever $X \subseteq Y$. (This is w.l.o.g. as we can assume that T preserves injective maps, possibly after changing $T\emptyset$ in a way that does not affect the class of coalgebras [9].)

The second parameter of the semantics is the interpretation of the modal operators, which relies on associating to each $\heartsuit \in \Lambda$ a *predicate lifting* $[\![\heartsuit]\!]$, i.e. 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. In particular, naturality of $[\![\heartsuit]\!]$ means that $[\![\heartsuit]\!]_X(f^{-1}[A]) = (Tf)^{-1}[[\![\heartsuit]\!]_Y(A)]$ for any map $f : X \to Y$.

Intuitively, a predicate lifting $[\![\heartsuit]\!]$ tells us what the local view of a state in X should be for it to satisfy a formula $\heartsuit\phi$ where ϕ has extension $A \subseteq X$; explicitly, the local view $\xi(x)$ of x should be an element of the set $[\![\heartsuit]\!]_X(A)$. E.g., one interprets $\Box$ on the Kripke functor T above using the predicate lifting

$$[\![\Diamond]\!]_X(A) := \{(S, V) : S \subseteq A, V \in \mathcal{P}(\mathsf{Prop})\}.$$

Formally, the notion of *satisfaction* of Λ-formulas ϕ at states x of C (denoted $x \models_C \phi$) is then defined by the expected clauses for Boolean operators, plus:

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

where $[\![\phi]\!]_C = \{x \in X : x \models_C \phi\}$ 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 $t \in [\![\heartsuit]\!]_X(A)$. From $[\![\heartsuit]\!]$ we obtain the predicate lifting interpreting $\bar{\heartsuit}$ by $[\![\bar{\heartsuit}]\!]_X(A) = TX - [\![\heartsuit]\!]_X(X - A)$.

On positive formulas, the core reasoning task is *subsumption checking*: 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.

Abusing notation, we identify a similarity type Λ with this semantic structure $\langle T, [\![\heartsuit]\!]_{\heartsuit\in\Lambda}\rangle$ used to interpret it, and refer to both as Λ. We shall use T for the underlying functor throughout.

Example 2.1 All logics discussed above are coalgebraic; see, e.g., [34,16]. As an additional example, *graded (modal) logic*, which we call G, has the similarity type $\Lambda = \{\Diamond_k : k \in \mathbb{N}\}$, with $\Diamond_k\phi$ read 'ϕ holds 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 any subset $A \subseteq X$ (then, for a map f, $\mathcal{B}_\infty f$ acts by taking image measures, i.e. $\mathcal{B}_\infty f(\mu)(y) = \mu(f^{-1}[\{y\}]$.) Coalgebras for $\mathcal{B}_\infty$ are *multigraphs*, i.e. directed graphs whose edges are annotated with multiplicities from $\mathbb{N} \cup \{\infty\}$. Each $\Diamond_k$ is interpreted by the predicate lifting

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

A multigraph (X, ξ) is essentially a more concise representation of a Kripke frame, with $\xi(x)(y) = n$ standing for n distinct successors of x, all of them isomorphic copies of y. Thus, $[\![\Diamond_k]\!]$ clearly captures the informal reading of $\Diamond_k$.

This framework is modular [34], and in particular supports fusion of modal logics by taking products of functors. For instance, the functor inducing Kripke models with m relations, supporting the interpretation of m relational modalities, can be seen as arising from the product $TX = \prod_{i=1}^{m} \mathcal{P}(X) \times 2^{\mathsf{Prop}}$ of m copies of the covariant powerset functor $\mathcal{P}$, and a copy of the constant functor

2 given by $2X = 2 = \{0, 1\}$ for each proposition symbol in Prop (the associated predicate liftings are derived in the obvious way).

Although coalgebraic logic supports non-monotone modalities, we assume operators to be *monotone* ($A \subseteq B \subseteq X \Rightarrow [\![\heartsuit]\!]_X A \subseteq [\![\heartsuit]\!]_X B$): to characterize formulas by simulations, we need monotonicity in inductive proofs, since simulations preserve but do not reflect satisfaction of formulas. Crucially, all monotone coalgebraic logics admit complete sets of tableau rules consisting (besides the standard propositional rules) of rules of the form $\Gamma_0/\Gamma_1 \mid \cdots \mid \Gamma_n$ where Γ_0 contains only formulas $\heartsuit a$, with $\heartsuit \in \Lambda \cup \bar{\Lambda}$, and $\Gamma_1, \ldots, \Gamma_n$ contain only variables, as in Fig. 1 [14]; we fix such a rule set $\mathcal{R}$ throughout.

In coalgebraic logic one exploits locality and reduces logical phenomena such as derivability or satisfiability from the full logic to the simpler setting of *one-step models*, which are, roughly, the result of forgetting the structure of a pointed model everywhere except at the root; see, e.g., [31]. With one-step models come *one-step formulas*, i.e. shallow modal formulas where propositional variables are introduced as placeholders for complex argument formulas under modal operators.

Definition 2.2 (One-step logic) Let V be a set of propositional variables (not fixed, and typically finite); a *one-step model over* V is just 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. $\breve{\tau}(x) = \{p : x \in \tau(p)\}$. A *conjunctive one-step* Λ*-formula* is a finite conjunction of atoms $\heartsuit p$, where $\heartsuit \in \Lambda$, $p \in V$. The satisfaction relation is given by $(X, \tau, t) \models_\tau \bigwedge_{i \in I} \heartsuit_i p_i$ iff $t \models \heartsuit_i \tau(p_i)$ for all i. Similarly, a *positive one-step* Λ*-formula* is an element of $\mathsf{Pos}(\Lambda(\mathsf{Pos}(V)))$, where $\Lambda(W) = \{\heartsuit w : \heartsuit \in \Lambda, w \in W\}$ and Pos denotes positive propositional combinations (using $\top$, $\bot$, $\vee$, $\wedge$), with the expected semantics. We write $\sqsubseteq_1$ for the subsumption relation in the one-step logic: $\phi \sqsubseteq_1 \psi$ if $(X, \tau, t) \models \psi$ whenever $(X, \tau, t) \models \phi$.

The transfer of results between the one-step and the full logic is done by way *collages*, i.e., 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 (see e.g. the construction of shallow models in [31,25]). Explicitly:

Definition 2.3 Given $t \in TX$, a family of pairwise disjoint pointed coalgebras $(C_x, x) = ((Y_x, \xi_x), x)$ for all $x \in X$, and a fresh root state r, the *collage* of these *collage data* is the pointed coalgebra (C, r), with $C = (Y, \xi)$, where Y is the (disjoint) union of $\{r\}$ and the Y_x, and

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

As indicated earlier, we assume that T preserves subsets, so, e.g., $TX \subseteq TY$.

In a nutshell, the collage is obtained from a root state r with successor structure $t \in TX$ by replacing every $x \in X$ with a pointed coalgebra (C_x, x). The following is immediate by construction:

Lemma 2.4 (Collage lemma) *For a collage (C, r) with collage data as in Definition 2.3, and all $x \in X$, $A \subseteq Y$ and $\heartsuit \in \Lambda$,*

(i) $x \models_C \phi \iff x \models_{C_x} \phi$, *and*

(ii) $t \in \heartsuit_X(A \cap X) \iff \xi(r) \in \heartsuit_Y A$.

Proof. 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:

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

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 2.5 Given collage data as in Definition 2.3, a valuation $\tau : V \to \mathcal{P}(X)$ *(positively) matches* a substitution $\rho : V \to L(\Lambda)$ if for all $x \in X$, $x \models_{C_x} \rho(p)$ iff (if) $x \in \tau(p)$.

Lemma 2.6 *Let* $\tau : V \to \mathcal{P}(X)$ (positively) match $\rho : V \to L(\Lambda)$. *Then*

(i) $x \in \tau(p)$ *iff (implies)* $x \models_C \rho(p)$, *and*

(ii) $t \models_\tau \heartsuit p$ *iff (implies)* $r \models_C \heartsuit\rho(p)$.

The converse process is as follows.

Definition 2.7 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 2.8 (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* $(X, \tau, t) \models \phi \iff r \models_C \phi\rho$.

3 Coalgebraic Simulations

We recall the notion of coalgebraic modal simulation from [16]. Given a binary relation $S \subseteq X \times Y$, we denote by S^- its relational inverse. Moreover, for $A \subseteq X$, the relational image of S over A is given by $S[A] := \{y : \exists x \in A.\, xSy\}$.

Definition 3.1 (Λ-Simulation) Let $C = (X, \xi)$ and $D = (Y, \zeta)$ be two given T-coalgebras. A Λ-*simulation* $S : C \to D$ (*of* C *by* D) is a relation $S \subseteq X \times Y$ such that xSy and $\xi(x) \models \heartsuit A$ imply $\zeta(y) \models \heartsuit S[A]$, for all $\heartsuit \in \Lambda$ and $A \subseteq X$. When xSy for a Λ-simulation S, we say that (D, y) Λ-*simulates* (C, x).

The properties of Λ-simulations that we need here are the following (cf. [16]):

Lemma 3.2 *Λ-simulations are stable under relational composition; moreover, (graphs of) identities are Λ-simulations.*

Lemma 3.3 *Let $S : C \to D$ be a Λ-simulation and 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 3.4 *Let $\bar{\Lambda} := \{\bar{\heartsuit} : \heartsuit \in \Lambda\}$. A relation S between T-coalgebras is a $\bar{\Lambda}$-simulation iff S^- is a Λ-simulation.*

Example 3.5 (See [16] for details.)

(i) Over Kripke frames and for $\Lambda = \{\Diamond\}$, a Λ-simulation $S : C \to D$ is just a simulation $C \to D$ in the usual sense. By Proposition 3.4, for $\Lambda = \{\Box\}$, a Λ-simulation $S : C \to D$ is then a simulation $D \to C$ in the usual sense. Consequently, a $\{\Box, \Diamond\}$-simulation is just a standard bisimulation.

(ii) A $\{p\}$-simulation for a proposition p is just a relation that preserves p.

(iii) Over monotone neighbourhoods with $\Lambda = \{\Box\}$, $S \subseteq X \times Y$ is a Λ-simulation between $\mathcal{M}$-coalgebras (X, ξ) and (Y, ζ) iff xSy and $A \in \xi(x)$ imply $S[A] \in \zeta(y)$.

4 Weakly Simulation-Initial Models

In general, modal formulas need not have smallest models under the simulation preorder. In some cases, however, such smallest models do exist. Formally, we define this property as follows.

Definition 4.1 (wsi models) Let ϕ be a positive Λ-formula. A pointed model (C_ϕ, x_ϕ) is called *weakly simulation-initial (wsi) for ϕ* if for any other (D, y), $y \models_D \phi$ iff (D, y) Λ-simulates (C_ϕ, x_ϕ).

In the relational setting, the term *sim-initial* has been used for an analogous notion [23]. Initiality in this sense is rather weak, though, since the witnessing simulations are not necessarily unique.

Remark 4.2 Since identities are Λ-simulations, a wsi model for ϕ satisfies ϕ. Thus, by Lemma 3.3, (C_ϕ, x_ϕ) is wsi for ϕ iff (i) $x_\phi \models \phi$, and (ii) whenever (D, y) is such that $y \models_D \phi$, then (D, y) Λ-simulates (C_ϕ, x_ϕ).

Definition 4.3 ([22]) A pointed coalgebra (C, x) is a *materialization* of ϕ if for all positive Λ-formulas ψ, $x \models_C \psi$ iff $\phi \sqsubseteq \psi$. In this case, ϕ is *materializable.*

Of course, this definition implies that a materialization of ϕ is a model of ϕ. By Lemma 3.3, the following is immediate:

Lemma 4.4 *Every wsi model is a materialization.* □

Thus, subsumption reduces to model checking in wsi models when they exist.

Definition 4.5 (Convexity) [6] A satisfiable Λ-formula ϕ is *(strongly) convex* if whenever $\phi \sqsubseteq \bigvee_{i \in I} \psi_i$ for some (possibly infinite) index set I and positive

Λ-formulas ψ_i (with the expected semantics of $\bigvee$), then already $\phi \sqsubseteq \psi_i$ for some $i \in I$.

Lemma 4.6 *If ϕ is materializable then ϕ is strongly convex.* $\square$

Remark 4.7 Convexity is generally felt to be necessary for tractability; see, e.g., [20,23] (where it is considered w.r.t. *finite* disjunctions). It is not only an important structural property but also provides a good handle for showing that certain formulas are *not* materializable. E.g. a formula that is itself a disjunction can have a materialization only when it is equivalent to one of its disjuncts. It is thus no surprise that tractable logics such as $\mathcal{EL}$ and TBox-free $\mathcal{FL}_0$ exclude disjunction; also here, we will henceforth *restrict attention to conjunctive formulas.*

But even conjunctive formulas may fail to be materializable. E.g., in G with $\Lambda = \{\Diamond_k : k \in \mathbb{N}\}$ we have $\Diamond_1 a \wedge \Diamond_1 b \sqsubseteq \Diamond_2(a \vee b) \vee \Diamond_1(a \wedge b)$ but the left hand side is not subsumed by any of the disjuncts of the right hand side, so convexity fails (cf. [6]). Similarly, conjunctive $\{\Box_1\}$-formulas may fail to be convex, as witnessed by

$$\Box_1(a \wedge b) \wedge \Box_1(b \wedge c) \wedge \Box_1(c \wedge d) \wedge \Box_1(d \wedge a) \sqsubseteq \\ \Box_1(a \wedge b \wedge c) \vee \Box_1(b \wedge c \wedge d) \vee \Box_1(c \wedge d \wedge a) \vee \Box_1(d \wedge a \wedge b)\,.$$

Worse, with the wrong choice of Λ, even $\top$ may fail to be materializable: in K with $\Lambda = \{\Box, \Diamond\}$ we have $\top \sqsubseteq \Box\Diamond\top \vee \Diamond\top$ but $\top \not\sqsubseteq \Box\Diamond\top$ and $\top \not\sqsubseteq \Box\Diamond\top$. Similarly, in M with $\Lambda = \{\Box, \Diamond\}$, one has that $\top \sqsubseteq \Box\top \vee \Diamond\top$ and yet $\top \not\sqsubseteq \Box\top$ and $\top \not\sqsubseteq \Diamond\top$.

The existence of wsi models thus depends strongly on the chosen T-structure Λ, as well as on slight variations in the semantics (e.g. w.r.t. seriality). We now proceed to show that one can limit the study of the phenomenon to the level of the much simpler one-step logic (Section 2). As suggested by Remark 4.2, we define in this case:

Definition 4.8 (one-step wsi models) A one-step model (X, τ, t) is *weakly simulation-initial (wsi)* for a conjunctive one-step formula ϕ over V if (i) $t \models_\tau \phi$, and (ii) for every (Y, ϑ, s), $A \subseteq X$ and $\heartsuit \in \Lambda$, $t \in \heartsuit_X A$ implies $s \in \heartsuit_Y S[A]$, where $xSy \iff \breve{\tau}(x) \subseteq \breve{\vartheta}(y)$.

Remark 4.9 One-step wsi models are never unique. However, one can assume w.l.o.g. that if (X, τ, t) is wsi, then every $x \in X$ is uniquely determined by $\breve{\tau}(x)$ (quotient (X, τ, t) by the equivalence relation induced by $\breve{\tau}$), and hence that X is of at most exponential size on the number of variables.

Definition 4.10 We say that Λ *admits (one-step) wsi models* if every conjunctive (one-step) formula has a (one-step) wsi model.

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

Theorem 4.11 *Λ admits wsi models whenever it admits one-step wsi models.*

Proof (Sketch) Induction on ϕ. We have $\phi = \bigwedge_{i \in I} \heartsuit_i \chi_i$ for a finite (possibly empty) set I. Take $V_\phi := \{a_{\chi_i} : i \in I\}$ and decompose ϕ as $\phi = \phi^* \rho$ with $\phi^* := \bigwedge_{i \in I} \heartsuit_i a_{\chi_i}$ a one-step formula and $\rho(a_{\chi_i}) := \chi_i$ a substitution. Let (X, τ, t) be a wsi for ϕ^*. By IH, there is, for each $x \in X$, a wsi model (C_x, x) with $C_x = (Y_x, \xi_x)$ for $\bigwedge_{p \in \check{\tau}(x)} \rho(p)$ with root x; the Y_x can be assumed pairwise disjoint. Pick a fresh x_ϕ, and obtain (C_ϕ, x_ϕ) by taking $\xi(x_\phi) = t$ and attaching C_x at each $x \in X$ (cf. [31]). One easily shows that (C_ϕ, x_ϕ) is wsi for ϕ. □

We now analyze under which conditions one-step wsi models exist. To begin, note that at the one-step level, wsi models coincide with materializations (recall that $\sqsubseteq_1$ is the one-step subsumption relation of Def. 2.2):

Definition 4.12 A one-step model (X, τ, t) is a *one-step materialization* of a conjunctive one-step Λ-formula ϕ over V if for every literal $\heartsuit\rho$ with $\heartsuit \in \Lambda$ and $\rho \in \mathsf{Pos}(V)$, $t \models_\tau \heartsuit\rho$ iff $\phi \sqsubseteq_1 \heartsuit\rho$. In this case, ϕ is *materializable*.

Again, this implies that a one-step materialization of ϕ is a model of ϕ.

Lemma 4.13 *A one-step model is wsi for a conjunctive one-step Λ-formula ϕ iff it is a materialization of ϕ.*

Moreover, existence of materializations is equivalent to convexity:

Definition 4.14 A satisfiable one-step Λ-formula ϕ over V is *strongly convex* if whenever $\phi \sqsubseteq_1 \bigvee_{i \in I} \psi_i$ for positive one-step Λ-formulas ψ_i over V and a (possibly infinite) index set I, then already $\phi \sqsubseteq_1 \psi_i$ for some $i \in I$.

Remark 4.15 In case Λ is finite, strong convexity of one-step formulas is the same as convexity (the notion obtained by restricting I to be finite in Definition 4.14), as then there are, up to equivalence, only finitely many positive one-step Λ-formulas over V.

Lemma 4.16 *A one-step Λ-formula is materializable iff it is strongly convex.*

Remark 4.17 Summing up, at the one-step level the notions of *being materializable*, *having a wsi model* and *being strongly convex* coincide. For the full logic, we have already noted that wsi models are materializations and materializable formulas are strongly convex. We leave the equivalence of these notions for individual formulas, i.e. to show that every strongly convex formula has a wsi model, to future research (for some relational logics, this equivalence is known [1,11]). Under mild additional assumptions, it does follow at the current stage that the equivalence holds between the respective properties of the logic as a whole: assume for simplicity that Λ contains infinitely many proposition symbols (actually, it suffices that the logic is *non-trivial*, i.e. contains infinitely many propositionally independent formulas). If all conjunctive Λ-formulas are strongly convex, then this holds (emulating propositional variables by proposition symbols from Λ) also for conjunctive one-step Λ-formulas. By the above, it follows that Λ admits one-step wsi models, and hence admits wsi models.

Next, we show how to read off convexity from the structure of the tableau rules for Λ. At the same time, we obtain a description of the structure of one-step materializations.

Definition 4.18 We call a tableau rule *definite* if it has exactly one conclusion, i.e. is of the form Γ/Δ with $\Gamma \subseteq (\Lambda \cup \bar{\Lambda})(V)$ and $\Delta \subseteq V$. A set $\mathcal{R}$ of definite one-step rules *preserves* Λ*-convexity* if whenever a rule R over V in $\mathcal{R}$ can be written in the form $\Gamma_1, \Gamma_2/\Delta_1, \Delta_2 \in \mathcal{R}$ with $\Gamma_1 \subseteq \Lambda(V_1)$, $\Gamma_2 \subseteq \bar{\Lambda}(V_2)$, $\Delta_1 \subseteq V_1$, $\Delta_2 \subseteq V_2$, with V_1, V_2 a disjoint decomposition of V (we call this a Λ*-splitting* of R), then for each $\bar{\heartsuit}a \in \Gamma_2$, the rule $\Gamma_1, \bar{\heartsuit}a/\Delta_1, a$ is also in $\mathcal{R}$.

The next theorem will show that preservation of Λ-convexity is sufficient for convexity of conjunctive Λ-formulas. It is fairly clear that, in cases where all rules are definite, necessity also holds for a sufficiently carefully formulated weakening of preservation of Λ-convexity (e.g. in the above notation, it clearly suffices to have $\Gamma_1, \bar{\heartsuit}a/\Delta_1, a$ derivable from $\mathcal{R}$ in the obvious sense); we refrain from exploring details.

Remark 4.19 In case Λ is closed under duals, the rule set $\mathcal{R}$ preserves Λ-convexity iff whenever Γ/Δ is a rule over V in $\mathcal{R}$ and $\emptyset \neq V_0 \subseteq V$, then $(\Gamma \cap \Lambda(V_0))/(\Delta \cap V_0)$ is in $\mathcal{R}$ – that is, iff $\mathcal{R}$ is stable under deleting variables.

Theorem 4.20 *Let Λ be finite (for brevity; in fact it suffices to assume a more sophisticated form of completeness [33]). If $\mathcal{R}$ preserves Λ-convexity, then Λ admits wsi models. Moreover, a one-step materialization for a conjunctive one-step Λ-formula $\phi = \bigwedge_{i \in I} \heartsuit_i a_i$ (read also as the set $\{\heartsuit_i a_i : i \in I\}$) is then obtained as follows. First put $W = \{a_i : i \in I\}$, and define (X, τ) to consist of*

- *a state x with $\breve{\tau}(x) = \Delta\sigma$, for each rule Γ/Δ over V in $\mathcal{R}$ and each renaming $\sigma : V \to W$ with $\Gamma\sigma \subseteq \phi$;*
- *a state x with $\breve{\tau}(x) = \Delta_1\sigma$, for each rule $\Gamma, \bar{\heartsuit}b/\Delta_1, \Delta_2$ over $V \uplus \{b\}$ in $\mathcal{R}$ with $\Delta_2 \subseteq \{b\}$ and each renaming $\sigma : V \to W$ with $\Gamma\sigma \subseteq \phi$.*

(In both cases, we can restrict to rules and renamings for which $\Gamma\sigma$ becomes maximal.) Then there exists $t \in TX$ such that (X, τ, t) is a materialization of ϕ.

Remark 4.21 The rule sets in all examples are built in such a way that σ can be restricted to be injective in the construction of Theorem 4.20 [32]; however, it is easy to see that in such cases, this restriction does not actually affect the result of the construction.

Example 4.22 Over the proposition functor 2, $\Lambda = \{p\}$ and $\Lambda = \{\bar{p}\}$ (but not, of course, $\Lambda = \{p, \overline{p}\}$) are easily seen to admit wsi models; e.g. $(\emptyset, \emptyset, 1)$ is wsi for p. This is our only positive example not matching Theorem 4.20: the one-step rule $p, \bar{p}/\bot$ fails to be definite, having no conclusion.

Example 4.23 Over Kripke frames, we have the following.

(i) $\Lambda = \{\Diamond\}$ admits wsi models: a $\{\Diamond\}$-splitting $\Gamma_1, \Gamma_2/\Delta_1, \Delta_2$ of (K_n) in Fig. 1 is of the form $\Gamma_1 = \Diamond b$, $\Gamma_2 = \Box a_1, \dots, \Box a_n$, and for each j we have a rule $\Diamond b, \Box a_j/b, a_j$ in $\mathcal{R}_K$, as required. The one-step wsi model for $\bigwedge_{i \in I} \Diamond a_i$ according to Theorem 4.20 is (I, τ, I) with $\tau(a_i) = \{i\}$. An example is depicted in Fig. 2(a). This extends to the multimodal case (see Remark 4.26), essentially, to $\mathcal{EL}$.

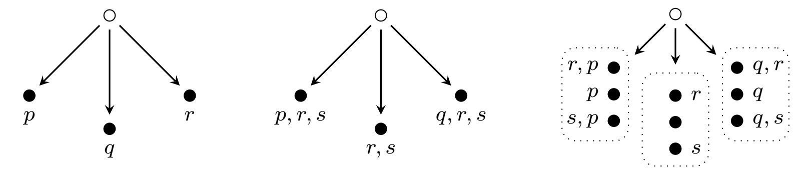

(a) K: $\Diamond p \wedge \Diamond q \wedge \Diamond r$ (b) KD: $\Diamond p \wedge \Diamond q \wedge \Box r \wedge \Box s$ (c) M_s: $\Box p \wedge \Box q \wedge \Diamond r \wedge \Diamond s$

Fig. 2. One-step wsi models for the indicated formulas. The white node is the implicit root, the black ones its domain. For M_s, minimal neighbourhoods are depicted (dotted boxes), not their supersets.

(ii) $\Lambda = \{\Box\}$ admits wsi models: any $\{\Box\}$-splitting $\Gamma_1, \Gamma_2/\Delta_1, \Delta_2$ of K_n already has Γ_2 of the form $\Diamond b$. Restricting to maximal rule matches, the one-step wsi model for $\bigwedge_{i \in I} \Box a_i$ according to (the second clause of) Theorem 4.20 is $(\{*\}, \tau, \{*\})$ with $\tau(a_i) = \{*\}$ for all i. This extends straightforwardly to the multi-modal case (Remark 4.26), of which $\mathcal{FL}_0$ is a syntactic variant.

(iii) In K, $\Lambda = \{\Box, \Diamond\}$ fails to be convex (Remark 4.7). Note that $\mathcal{R}_K$ fails to preserve convexity: deleting b from K_n yields $D_n \notin \mathcal{R}_K$ (Fig. 1, Remark 4.19). In KD, however, $\{\Box, \Diamond\}$ does admit wsi models, as the rules K_n and D_n together are stable under deleting occurrences of variables. Restricting to maximal matches, the one-step wsi model for $\bigwedge_{i \in I} \Box a_i \wedge \bigwedge_{j \in J} \Diamond b_j$ is $(J \cup \{*\}, \tau, J \cup \{*\})$ (with $* \notin J$) given by $\tau(a_i) = J \cup \{*\}$ and $\tau(b_j) = \{j\}$ (see Fig. 2(b)).

Example 4.24 Over monotone neighbourhoods, the situation is analogous as over Kripke frames, due to the similarity of the rule sets: both $\Lambda = \{\Box\}$ and $\Lambda = \{\Diamond\}$ admit wsi models in M, but not $\Lambda = \{\Box, \Diamond\}$ (M validates $\Box\top \vee \Diamond\top$ but none of the disjuncts, so no $\{\Box, \Diamond\}$-formula is convex). In M_s, $\{\Box, \Diamond\}$ does admit wsi models, though, for essentially the same reasons as in KD. The one-step wsi model from Theorem 4.20 for $\bigwedge_{i \in I} \Box a_i \wedge \bigwedge_{j \in J} \Diamond b_j$ (I, J disjoint) is $(X, \tau, \mathfrak{N})$:

$$X := \{K \subseteq I \cup J : |K \cap I| \leq 1, |K \cap J| \leq 1\} \qquad \tau(a_i) := \{K \in X : i \in K\}$$
$$\mathfrak{N} := \uparrow(\{\tau(a_i) : i \in I\} \cup \{\{K \in X : K \subseteq J\}\}) \quad \tau(b_j) := \{K \in X : j \in K\}$$

where $\uparrow$ is closure under taking supersets. Fig. 2(c) depicts the construction.

Example 4.25 In coalition logic, $\Lambda = \{[C], \langle C \rangle : C \subseteq N\}$ admits wsi models: its rules are stable under deleting occurrences of variables by Remark 4.19.

Remark 4.26 When one models fusion of modal logics by taking products of functors as noted in Sec. 2 (see [34]) this is reflected in the construction of one-step wsi models by just taking disjoint unions of the domains and pairing the transition structures (prolonged into the disjoint union). For instance, in the multimodal logic $\Lambda = \{\Box_1, \ldots, \Box_n\}$ over Kripke frames, one-step wsi models for one-step formulas $\bigwedge_{i=1}^{n} \bigwedge_{j=1}^{m_i} \Box_i a_{ij}$ are formed by taking the disjoint union of the one-step wsi models for the formulas $\bigwedge_{j=1}^{m_i} \Box_i a_{ij}$ as described in Example 4.23, and thus have n states, with the i-th state satisfying the propositional variables $a_{i1}, \ldots, a_{im_i}$. By Example 4.22, adding atomic propositions does not enlarge the carriers of wsi models at all.

For tractability, studied in the next section, we need wsi models to be small. However, existence of wsi models is of independent interest, even in those cases in which they may be exponentially large. For instance, from Example 4.23, we can already conclude that *conjunctive KD is convex.*

5 Tractability

Assume from now on that Λ admits one-step wsi models. Lemma 4.4 then allows us to reduce subsumption to satisfaction in such models. (This is also the principle underlying state-of-the-art consequence-based reasoning procedures, which for $\mathcal{EL}$ go back to [6].) In the previous sections, we have refrained from giving explicit descriptions of t when (X, τ, t) is wsi, and in fact it is not necessary to actually know t. Instead, we opt for a different representation of wsi models: in the recursive construction of a wsi model (C_ϕ, x_ϕ) for a conjunctive Λ-formula ϕ (see proof sketch of Theorem 4.11), we have calculated a one-step formula ϕ^* and used a one-step wsi model (X, τ, t) for it. For algorithmic purposes, we now drop t but store ϕ^*, X, and τ; we call the arising object an *abstract wsi model for* ϕ. We face then the following problem:

Definition 5.1 The *conjunctive one-step consequence problem* of Λ is to decide, given a conjunctive one-step Λ-formula ψ over V, $\heartsuit \in \Lambda$, and $\rho \in \mathsf{Pos}(V)$, if $\psi \sqsubseteq_1 \heartsuit\rho$.

If the conjunctive one-step consequence problem for Λ is in P, then we can check in time polynomial in the size of an abstract wsi model (C_ϕ, x_ϕ) for ϕ whether $x_\phi \models_{C_\phi} \psi$ for a positive Λ-formula ψ, e.g. by calculating extensions of subformulas of ψ bottom up. Now in the positive examples of the previous section, deciding whether, in the notation of the above definition, $\psi \sqsubseteq_1 \heartsuit\rho$ can be done using the respective rule sets to check whether $\psi \wedge \bar{\heartsuit}\neg\rho$ is satisfiable, which in turn will lead to checking satisfiability of a propositional formula of the form $\chi \wedge \neg\rho$ where χ is a conjunction over V, a trivial task given that ρ is positive. Thus, the conjunctive one-step consequence problem of Λ is in P provided that we can polynomially bound the number of rule matches to a given conjunction over $\Lambda(W)$, which is easily seen for all relevant examples.

Polynomial-time computability (entailing polynomially bounded size) of abstract wsi models will then imply tractability of subsumption. In some cases, tractability will hold only if we bound certain parameters. To avoid overformalization, we will call any set of conjunctive Λ-formulas a *conjunctive* Λ*-fragment* and apply notions defined so far w.r.t. the set of all conjunctive formulas, such as *admitting wsi models*, also to fragments. Note that sometimes restricting to a fragment will also restrict the relevant set of one-step formulas.

Definition 5.2 A conjunctive Λ-fragment $\mathcal{L}$ *admits polynomial wsi models* if every $\mathcal{L}$-formula has a polynomial-time computable abstract wsi model.

Lemma 5.3 *If $\mathcal{L}$ admits polynomial wsi models and the conjunctive one-step consequence problem of Λ is in P, then subsumption $\phi \sqsubseteq \psi$ between $\mathcal{L}$-formulas ϕ and positive Λ-formulas ψ is in P.*

We identify tractability criteria at the one-step level:

Definition 5.4 We say that one-step wsi models (X, τ, t) of one-step formulas $\phi = \bigwedge \heartsuit_i a_i$ are *linear* if $|\tau(a_i)| \leq 1$ for all i, *k-bounded* if $|\breve{\tau}(x)| \leq k$ for all $x \in X$, and *polynomial* if $|X|$ is polynomially bounded in the size of ϕ.

In words, linearity means that every propositional variable is satisfied in at most one state, while k-boundedness means that each state satisfies at most k propositional variables.

Proposition 5.5 a) *If a conjunctive Λ-fragment $\mathcal{L}$ admits linear or k-bounded one-step wsi models, then $\mathcal{L}$ admits polynomial wsi models.* b) *If Λ admits polynomial one-step wsi models, then conjunctive Λ-fragments defined by bounding the modal depth admit polynomial wsi models.*

(The complexity of bounded-depth fragments of modal logics over a complete Boolean basis has been studied, e.g., in [17].)

Proof (Sketch) Linearity implies that a wsi model for ϕ has at most as many states as ϕ has subformulas. On the other hand, k-boundedness ensures that wsi models, constructed as trees in the proof of Theorem 4.11, can be collapsed into polynomial-sized dags by identifying states realizing the same target formula; by k-boundedness, at most $|\phi|^k$ target formulas will arise in the construction. $\Box$

Example 5.6 (i) One-step wsi models for $\{\Diamond\}$ and for $\{\Box\}$ over Kripke frames (Example 4.23) are linear; those for $\{\Diamond\}$ are in addition 1-bounded. By Remark 4.26, this extends straightforwardly to the case with multiple modalities and atomic propositions. We thus recover the known results that subsumption checking in conjunctive multimodal K with only diamonds ($\mathcal{EL}$) or only boxes ($\mathcal{FL}_0$) is in P. As an aside, the conjunctive fragment of the co-contravariant modal logic of [1], which is essentially positive Hennessy-Milner logic with only diamonds for some actions and only boxes for the others, can be seen as a fusion of a logic of boxes with a logic of diamonds, and thus also has linear one-step wsi models, i.e. has a polynomial-time subsumption problem.

(ii) One-step wsi models for $\{\Box, \Diamond\}$ over serial Kripke frames, i.e. for conjunctive KD, are polynomial, so that *subsumption in bounded-depth fragments of conjunctive KD is in P* (with unboundedly many atomic propositions). This may be seen as a companion result to the (easily proved) coNP upper bound for bounded-depth fragments of full K [17].

Alternatively, if one restricts conjunctive KD formulas to use at most k boxes at each modal depth, then one-step wsi models for them become $k + 1$-bounded, so that this restriction also ensures tractable reasoning. Again, this extends easily to the multimodal case with unboundedly many atomic propositions. Since one has a straightforward embedding of $\mathcal{EL}$ into multimodal KD (using a fresh propositional atom e marking 'existing' states to simulate arbitrary Kripke frames with serial ones), this result can be seen as generalizing the tractability of $\mathcal{EL}$ (which is just the case $k = 0$).

It is worth observing that the more specific problem of *satisfiability* but over unrestricted conjunctive KD extended with atomic negation (called *poor man's logic*) is known to be in P [18].

(iii) In M_s (Example 4.24), wsi models for $\Lambda = \{\Box, \Diamond\}$ are 2-bounded, so that *conjunctive* M_s *is tractable*. Similarly, wsi models for the structure $\Lambda = \{[C], \langle C\rangle : C \subseteq N\} - \{[\emptyset], \langle N\rangle\}$ in coalition logic / alternating-time logic are n-bounded, where n is the (fixed!) total number of agents (since n is also the maximal number of disjoint non-empty coalitions). Thus, for each finite set N of agents, *conjunctive coalition logic over* N *without* $[\emptyset]$ *and* $\langle N\rangle$ *is tractable.*

6 Greatest Fixpoints

We now proceed to extend the base logic with a fixpoint operator. This will allow us to cover global definitions (e.g.. classical terminological boxes, in DL parlance) and fragments of game logic and the alternating-time μ-calculus. We can only expect to get wsi models for formulas with *greatest* fixpoints, which are similar in flavour to infinite conjunctions, while least fixed points are disjunctive (e.g., $\nu x.(p \wedge \Diamond x)$ can be seen as the infinitary formula $p \wedge \Diamond(p \wedge \Diamond(p \wedge \dots)))$ which characterizes an infinite path of nodes satisfying p).

Following [22], we will actually allow for mutually recursive auxiliary definitions, as in the vectorial μ-calculus [3]. The resulting logic can be shown to be no more expressive than the one with only single-variable ν, but to admit exponentially more succinct definitions [22]. Syntactically, the grammar of *positive* Λ-ν-*formulas* extends that of positive Λ-formulas with fixpoint variables from a set Δ and, for $\alpha \in \{\nu, \mu\}$, formulas $\alpha(y; y_1, \dots y_n).(\phi, \phi_1 \dots \phi_n)$, where $y, y_1, \dots y_n \in \Delta$ must be distinct and $\phi, \phi_1, \dots \phi_n$ are positive Λ-ν-formulas. A formula $\nu(y; y_1, \dots y_n).(\phi; \phi_1, \dots, \phi_n)$ defines $y, y_1, \dots, y_n$ as a simultaneous greatest fixpoint, and then returns y; similarly for μ with least fixpoints. A *sentence* is a formula where every fixpoint variable is bound by a ν or μ. *Conjunctive* fixpoint Λ-formulas extend conjunctive Λ-formulas with ν only.

We define the semantics of this language over a T-coalgebra $C = (X, \zeta)$ and a valuation $\mathcal{V} : \Delta \to \mathcal{P}(X)$; by $[\![\phi]\!]_{C,\mathcal{V}}$ we denote the extension of ϕ in C assuming that the fixpoints variables are interpreted using $\mathcal{V}$. The propositional and modal cases are defined like before (with $[\![x]\!]_{C,\mathcal{V}} = \mathcal{V}(x)$); moreover, $[\![\nu(y_0; y_1, \dots y_n).(\phi_0; \phi_1, \dots \phi_n)]\!]_{C,\mathcal{V}}$ is the first projection of the greatest fixed point of the map taking $(A_0, \dots, A_n)$ to $([\![\phi_i]\!]_{C,\mathcal{V}[y_0 \mapsto A_0 \dots y_n \mapsto A_n]})_{i=1,\dots,n}$. The semantics of μ is dual. For a sentence ϕ, the initial $\mathcal{V}$ is irrelevant, so we may write just $[\![\phi]\!]_C$. Preservation of positive formulas by simulations extends to fixpoint formulas:

Lemma 6.1 *Let S be a Λ-simulation of a coalgebra $C = (X, \xi)$ by a coalgebra*

D, and let $\mathcal{V} : \Delta \to \mathcal{P}(X)$ be a valuation. Then for every positive Λ-ν-formula ϕ, $S[[\![\phi]\!]_{C,\mathcal{V}}] \subseteq [\![\phi]\!]_{D,S[\mathcal{V}]}$, where $S[\mathcal{V}]$ denotes the valuation taking x to $S[\mathcal{V}(x)]$.

Extending the definition of *wsi models* literally to positive Λ-ν-formulas, we thus obtain a generalization of Lemma 4.4, i.e. a wsi model for a fixpoint formula ϕ is a *materialization*, so that subsumption of ϕ by positive Λ-ν-formulas reduces to satisfaction in the wsi model.

Example 6.2 DLs are logics for knowledge representation, where terminologies are defined via axioms in *TBoxes* which effectively constrain the classes of models over which one reasons. In particular, one is sometimes interested in so-called *classical TBoxes with greatest fixpoint semantics* [7, Chapter 2]. Here, axioms of a TBox $\mathcal{T}$ are definitions of the form $a \equiv \phi$ with a a proposition symbol that is allowed to occur as a left-hand side of only one definition. Such an a is said to be a *derived* concept of $\mathcal{T}$. Each model C interpreting the non-derived propositions is extended to a unique model $C^{\mathcal{T}}$ which arises as the greatest fixpoint of the function mapping an extension C' of C interpreting also the derived propositions to the extension C'' where for each $a \equiv \phi \in \mathcal{T}$, $[\![a]\!]_{C''} = [\![\phi]\!]_{C'}$. One writes $\mathcal{T} \models \psi \sqsubseteq \chi$ if for each model C, $\psi \sqsubseteq \chi$ holds in $C^{\mathcal{T}}$. It is then clear that *subsumption over* $\mathcal{T}$, i.e. to decide whether $\mathcal{T} \models \psi \sqsubseteq \chi$, reduces to subsumption of fixpoint formulas: assume $\mathcal{T} = \{a_1 \equiv \phi_1, \dots a_n \equiv \phi_n\}$; we have $\mathcal{T} \models \psi \sqsubseteq \chi$ iff $\nu(z; a_1, \dots, a_n).(\psi; \phi_1, \dots, \phi_n) \sqsubseteq \nu(z; a_1, \dots, a_n).(\chi; \phi_1, \dots, \phi_n)$ where z is a fresh variable. Additional details are given by Lutz et al. [22].

Example 6.3 (Game logic) Model-checking a PDL formula $\langle\alpha\rangle\top$ can be seen as finding a winning strategy in a one-player game, where α describes the rules of the game and the model encodes the possible moves of the player on a fixed game board. In Game Logic (GL) [27], this notion is extended to two-player games (of perfect information). Composite games α are built from atomic games using the program constructors of PDL plus a *dualization* operator $(\cdot^d)$, which corresponds to players swapping roles, so that $\langle\alpha^d\rangle\phi \equiv [\alpha]\phi$ (and hence $[\alpha]$ can be omitted from the language). The two-player view disables normality (i.e. one no longer has $\langle\alpha\rangle(\phi \vee \psi) \to \langle\alpha\rangle\psi \vee \langle\alpha\rangle\psi$); hence, models of GL are products of monotone neighbourhood frames S_a, one per atomic game a. Intuitively, a set $A \in S_a(x)$ corresponds to (an upper bound on) positions that could be reached from x when following a fixed strategy for a; allowing for different responses of player II, we see that A need not be a singleton. As a notational infelicity, the predicate lifting interpreting $\langle a\rangle$ in GL (for a atomic) is that of $\Box$ in standard notation for monotone modal logic. Serial models are those where atomic games never get stuck, no matter which player begins. We note that GL has a well-known sublogic, concurrent propositional dynamic logic CPDL [30], which omits dualization $\cdot^d$ but retains $\cap$, the dual of $\cup$.

GL can be embedded into the fixpoint extension ${M_s}^{\nu}_m$ of multi-modal M_s (with duals of atomic propositions), much like PDL can be embedded into the relational μ-calculus. Two fixpoint variables suffice for this [10]. It is not hard to see that using fixpoint variables as a form of let-expressions, one can avoid the exponential blowup present in the original encoding. The *conjunctive* fragment of GL is swiftly defined as the preimage of the conjunctive fragment of ${M_s}^{\nu}_m$ under this embedding.

Example 6.4 (Alternating time) The *alternating-time μ-calculus (AMC)* is essentially the extension of coalition logic with fixpoint operators (its actual notation is slightly different) [2]. The *conjunctive fragment* of the AMC can by defined in the obvious way excluding $\vee$, $\neg$, and μ. In this fragment, we can still express 'always' formulas from alternating-time temporal logic (ATL) such as $\langle\langle C\rangle\rangle\Box\phi$, which is read 'coalition C can maintain ϕ forever', and is equivalent to the fixpoint formula $\nu x.\,(\phi \wedge [C]x)$.

We proceed to show that if Λ admits one-step wsi models, we also obtain wsi models for conjunctive Λ-ν-formulas. We exploit the fact that any such sentence can be put, in polynomial time, in a *shallow* normal form, i.e. without nested occurrences of ν (using Bekič's law [3]) and without nesting of modal operators (using abbreviations for subformulas in analogy to standard TBox normalizations [4]).

Thus, let $\phi = \nu(x_0; x_1, \dots, x_n)(\phi_0; \phi_1, \dots, \phi_n)$ be a shallow sentence. We shall assume, for each conjunctive one-step Λ-formula ψ over $V = \Delta$, a fixed one-step wsi model $(X_\psi, \tau_\psi, t_\psi)$ which we then call *the* one-step wsi model for ψ. We assume w.l.o.g. that $X_\psi \subseteq \mathcal{P}(V(\psi))$, where $V(\psi)$ is the set of variables mentioned in ψ, and $\tau_\psi(x) = \{A \in X_\psi : x \in A\}$ (Remark 4.9). We then construct the carrier X_ϕ of C_ϕ as a subset of $\mathcal{P}(V)$. For $A \subseteq V$, we let ϕ_A denote the conjunctive one-step formula given by $\bigwedge_{x_i \in A} \phi_i$. Then, X_ϕ is the smallest subset of $\mathcal{P}(V)$ containing $r_\phi = \{x_0\}$ such that $X_{\phi_A} \subseteq X_\phi$, for each $A \in X_\phi$. We define a T-coalgebra structure ξ_ϕ on X_ϕ by $\xi_\phi(A) = T(i_A)t_{\phi_A}$, where i_A is the inclusion $X_{\phi_A} \hookrightarrow X_\phi$.

Theorem 6.5 *If Λ admits one-step wsi models, then for every shallow $L^\nu(\Lambda)$-sentence ϕ, (C_ϕ, r_ϕ) as constructed above is a wsi model.*

Proof (Sketch) Let ϕ have the form $\nu(x_0; x_1, \dots x_n).(\phi_0; \phi_1, \dots, \phi_n)$, so $V = \{x_0, \dots x_n\}$. We have to show that (i) $r_\phi \models_{C_\phi} \phi$ and (ii) that if $d \models_D \phi$, then $r_\phi S d$ for some simulation $S : C_\phi \to D$ (Remark 4.2).

(i): By coinduction – taking $\mathcal{V}(x_i) = \{A \in X_\phi : x_i \in A\}$, one shows that $\mathcal{V}(x_i) \subseteq [\![\phi_i]\!]_{C_\phi,\mathcal{V}}$ for all $x_i \in V$. The gfp property of ϕ then implies $\mathcal{V}(x_0) \subseteq [\![\phi]\!]$, and clearly $r_\phi \in \mathcal{V}(x_0)$.

(ii): For $i = 0, \dots n$, let $\phi^{(i)}$ denote the formula obtained by projecting the i-th component of ϕ:

$$\phi^{(i)} = \nu(x_i; x_0 \dots x_{i-1}, x_{i+1} \dots x_n).(\phi_i; \phi_0 \dots \phi_{i-1}, \phi_{i+1} \dots),$$

so in particular $\phi^{(0)} = \phi$. Assume $d \models_D \phi$ for some coalgebra $D = (Y, \zeta)$. Define a relation $S \subseteq X_\phi \times Y$ by

$$ASy \iff y \models_D \bigwedge_{x_i \in A} \phi^{(i)}.$$

Then clearly $r_\phi S d$, for by definition $r_\phi = \{x_0\}$ and $\phi^{(0)} = \phi$. One can show that S is a Λ-simulation. $\Box$

Clearly, all conjunctive logics listed as having one-step wsi models in the examples of Sec. 4 have wsi models when extended with greatest fixpoints, in particular remain convex. Of course, the wsi models constructed above may be exponentially large, even when Λ admits linear one-step wsi models. However, under k-boundedness, elements of $X_\phi \subseteq \mathcal{P}(\Delta)$ have at most k elements, leading to our main criterion for smallness of wsi models under greatest fixpoints:

Theorem 6.6 *If Λ admits k-bounded one-step wsi models for some k, then conjunctive Λ-ν-formulas have polymomial-size wsi models.*

By Theorem 6.6 and the description of one-step wsi models in Section 4, and using abstract wsi models as in Section 5, we regain the known result that subsumption checking over classical TBoxes with gfp semantics in $\mathcal{EL}$ is in P [4], and in fact can extend it to allow a bounded number of universal restrictions, always in conjunction with $\Diamond\top$. As new results, we obtain:

Corollary 6.7 *Subsumption checking for conjunctive Game Logic is in P.*

Corollary 6.8 *Subsumption checking for the conjunctive alternating-time μ-calculus (AMC) without $[\emptyset]$ and $\langle N\rangle$ is in P.*

Remark 6.9 There is one case where we do obtain polynomial-size wsi models without k-boundedness, namely $\Lambda = \{\Box\}$ over Kripke frames – here, one-step wsi models have only one state, so that wsi models for fixpoint formulas are lassos, i.e. chains of states ending in a loop. For smallness, one still needs to impose additional restrictions on shallow fixpoints $\nu(y; y_1, \dots, y_n).(\phi; \phi_1, \dots, \phi_n)$, e.g. that y_i always appears in ϕ_i, or that the fixpoint is acyclic, i.e. not actually recursive. This example does not extend to the multi-modal case since the property of one-step wsi models being singletons is not stable under taking disjoint sums (Remark 4.26). Indeed, the multimodal version is FL, and reasoning over even the most restrictive (i.e. acyclic) TBoxes in $\mathcal{FL}_0$ is known to be coNP-hard [26].

7 Conclusions

Representability of formulas by models in the sense that simulation of the model is equivalent to satisfaction of the formula is a highly useful phenomenon in conjunctive fragments of modal fixpoint logics. It implies, for instance, convexity of the formula (and is equivalent to it in the one-step case) and under a polynomial size bound on the model, tractability of reasoning. We have studied the question of existence of such *weakly simulation-initial (wsi) models*, in the framework of coalgebraic logic; in particular, we have proved a reduction of the problem to a local (*one-step*) version. We were able to derive a criterion for tractability from the shape of the tableau rules that enabled us to establish tractability in a number of key examples:

- we have recovered known tractability results for the description logics $\mathcal{EL}$ (over classical TBoxes with gfp semantics) and $\mathcal{FL}_0$ (without a TBox), and shown that reasoning over classical TBoxes with gfp semantics in $\mathcal{EL}$ (equivalently in the fragment of the multi-modal μ-calculus defined by restricting to conjunction, diamonds, and greatest fixed points) remains tractable when we allow a bounded number of universal restrictions (i.e. boxes);

- we established tractability of conjunctive monotone logic with greatest fixed points over serial models, which subsumes corresponding fragments of game logic [27];
- we have shown tractability of the conjunctive fragment (which has greatest but not least fixed points) of the alternating-time μ-calculus AMC [2]; this fragment still includes the game-based versions of EG and AG found in ATL.

Outside the large body of work on $\mathcal{EL}$, there has been only a limited amount of research on wsi models for conjunctive logics. Notable examples are the work on the relationship between relational modal logics and modal transition systems [11,1] where formulas in certain variants of positive Hennessy-Milner logic are shown to have wsi models iff they are convex. (*prime* in the cited works). We exhibited a similar equivalence at the level of conjunctive coalgebraic *logics*; we leave a generalization of the equivalence for individual *formulas* as future work. There is some work on sub-Boolean fragments of temporal logics, which however focuses on satisfiability rather than subsumption (e.g. [24]).

Further points for future investigation include the use of wsi models to calculate so-called *least common subsumers* [8], as well as covering *general* TBoxes (i.e. finite sets of arbitrary inclusion axioms), which is known to remain tractable in the case of $\mathcal{EL}$ [12].

Acknowledgments The authors wish to thank Carsten Lutz for useful discussions, and Erwin R. Catesbeiana for unsolicited remarks regarding the absence of unsatisfiable formulas in $\mathcal{EL}$.

References

[1] Aceto, L., I. Fábregas, D. de Frutos-Escrig, A. Ingólfsdóttir and M. Palomino, *Graphical representation of covariant-contravariant modal formulae*, in: *Expressiveness in Concurrency, EXPRESS 2011*, EPTCS **64**, 2011, pp. 1–15.

[2] Alur, R., T. A. Henzinger and O. Kupferman, *Alternating-time temporal logic*, J. ACM **49** (2002), pp. 672–713.

[3] Arnold, A. and D. Niwiński, "Rudiments of μ-calculus," Elsevier, 2001.

[4] Baader, F., *Terminological cycles in a description logic with existential restrictions*, in: *International Joint Conference on Artificial Intelligence, IJCAI 2003*, pp. 325–330.

[5] Baader, F., *Using automata theory for characterizing the semantics of terminological cycles*, Ann. Math. Artif. Intell. **18** (1996), pp. 175–219.

[6] Baader, F., S. Brandt and C. Lutz, *Pushing the $\mathcal{EL}$ envelope*, in: *International Joint Conference on Artificial Intelligence, IJCAI 2005.*

[7] Baader, F., D. Calvanese, D. L. McGuinness, D. Nardi and P. F. Patel-Schneider, editors, "The Description Logic Handbook," Cambridge University Press, 2007, second edition.

[8] Baader, F., 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.

[9] Barr, M., *Terminal coalgebras in well-founded set theory*, Theoret. Comput. Sci. **114** (1993), pp. 299–315.

[10] Berwanger, D., E. Grädel and G. Lenzi, *The variable hierarchy of the μ-calculus is strict*, Theory Comput. Syst. **40** (2007), pp. 437–466.

[11] Boudol, G. and K. Larsen, *Graphical versus logical specifications*, Theoret. Comput. Sci. **106** (1992).

[12] Brandt, S., *Polynomial time reasoning in a description logic with existential restrictions, GCI axioms, and – what else?*, in: *Eureopean Conference on Artificial Intelligence, ECAI 2004* (2004), pp. 298–302.

[13] Chellas, B., "Modal Logic," Cambridge University Press, 1980.

[14] Cîrstea, C., C. Kupke and D. Pattinson, *EXPTIME tableaux for the coalgebraic μ-calculus*, Log. Methods Comput. Sci. **7** (2011).

[15] Emerson, E. and C. Jutla, *The complexity of tree automata and logics of programs*, SIAM J. Comput. **29** (1999), pp. 132–158.

[16] Gorín, D. and L. Schröder, *Simulations and bisimulations for coalgebraic modal logics*, in: *Algebra and Coalgebra in Computer Science, CALCO 2013*, LNCS **8089** (2013), pp. 253–266.

[17] Halpern, J., *The effect of bounding the number of primitive propositions and the depth of nesting on the complexity of modal logic*, Artif. Intell. **75** (1995), pp. 361–372.

[18] Hemaspaandra, E., *The complexity of poor man's logic*, Journal of Logic and Computation **11** (2001), pp. 609–622.

[19] Hofmann, M., *Proof-theoretic approach to description-logic*, in: *Logic in Computer Science, LICS 2005*, pp. 229–237.

[20] Krisnadhi, A. and C. Lutz, *Data complexity in the $\mathcal{EL}$ family of description logics*, in: *Logic for Programming, Artificial Intelligence, and Reasoning, LPAR 2007*, LNCS **4790** (2007), pp. 333–347.

[21] Ladner, R., *The computational complexity of provability in systems of modal propositional logic*, SIAM J. Comput. **6** (1977).

[22] Lutz, C., R. Piro and F. Wolter, *Enriching $\mathcal{EL}$-concepts with greatest fixpoints*, in: *European Conference on Artificial Intelligence, ECAI 2010*, pp. 41–46.

[23] Lutz, C. and F. Wolter, *Non-uniform data complexity of query answering in description logics*, in: *Principles of Knowledge Representation and Reasoning, KR 2012*.

[24] Meier, A., M. Thomas, H. Vollmer and M. Mundhenk, *The complexity of satisfiability for fragments of CTL and CTL**, Int. J. Found. Comput. Sci. **20** (2009), pp. 901–918.

[25] Myers, R., D. Pattinson and L. Schröder, *Coalgebraic hybrid logic*, in: *Foundations of Software Science and Computational Structures, FOSSACS 2009*, LNCS **5504** (2009), pp. 137–151.

[26] Nebel, B., *Terminological reasoning is inherently intractable*, Artif. Intell. **43** (1990), pp. 235–249.

[27] Parikh, R., *Propositional game logic*, in: *Foundations of Computer Science, FOCS 1983*.

[28] Pattinson, D., *Coalgebraic modal logic: Soundness, completeness and decidability of local consequence*, Theoret. Comput. Sci. **309** (2003), pp. 177–193.

[29] Pauly, M., *A modal logic for coalitional power in games*, J. Log. Comput. **12** (2002), pp. 149–166.

[30] Peleg, D., *Concurrent dynamic logic*, J. ACM **34** (1987), pp. 450–479.

[31] Schröder, L. and D. Pattinson, *Shallow models for non-iterative modal logics*, in: *Advances in Artificial Intelligence, KI 2008*, LNCS **5243** (2008), pp. 324–331.

[32] Schröder, L. and D. Pattinson, *PSPACE bounds for rank-1 modal logics*, ACM Trans. Comput. Log. **10** (2009), pp. 13:1–13:33.

[33] Schröder, L. and D. Pattinson, *Strong completeness of coalgebraic modal logics*, in: *Theoretical Aspects of Computer Science, STACS 2009*, LIPIcs **3** (2009), pp. 673–684.

[34] Schröder, L. and D. Pattinson, *Modular algorithms for heterogeneous modal logics via multi-sorted coalgebra*, Math. Structures Comput. Sci. **21** (2011), pp. 235–266.