The craft of model making: PSPACE bounds for non-iterative modal logics

Lutz Schröder
DFKI-Lab Bremen
and Department of Computer Science, Universität Bremen

Dirk Pattinson
Department of Computing
Imperial College London

January 22, 2008

Abstract

The methods used to establish *PSPACE*-bounds for modal logics can roughly be grouped into two classes: syntax driven methods establish that exhaustive proof search can be performed in polynomial space whereas semantic approaches directly construct shallow models. In this paper, we follow the latter approach and establish generic *PSPACE*-bounds for a large and heterogeneous class of modal logics in a coalgebraic framework. In particular, no complete axiomatisation of the logic under scrutiny is needed. This does not only complement our earlier, syntactic, approach conceptually, but also covers a wide variety of new examples which are difficult to harness by purely syntactic means. Apart from re-proving known complexity bounds for a large variety of structurally different logics, we apply our method to obtain previously unknown *PSPACE*-bounds for Elgesem's logic of agency and for graded modal logic over reflexive frames.

1 Introduction

Special purpose modal logics often combine expressivity and decidability, usually in a low complexity class. In the absence of fixed point operators, these logics are frequently decidable in PSPACE, i.e. not dramatically worse than propositional logic. While lower PSPACE bounds for modal logics can typically be obtained directly from seminal results of Ladner [20] by embedding a PSPACE-hard logic such as K or KD, upper bounds are often non-trivial to establish. In particular PSPACE upper bounds for non-normal logics have recently received much attention:

- A PSPACE upper bound for graded modal logic [11] is obtained using a constraint set algorithm in [30]. This corrects a previously published incorrect algorithm and refutes a previous EXPTIME hardness conjecture.
- More recently, a *PSPACE* upper bound for Presburger modal logic (which contains graded modal logic and majority logic [23]) has been established using a Ladner-type algorithm [6] .
- Using a variant of a shallow neighbourhood frame construction from [31], a PSPACE upper bound for coalition logic is established in [26].

• PSPACE upper bounds for CK and related conditional logics [3] are obtained in [22] by a detailed proof-theoretic analysis of a labelled sequent calculus.

The methods used to obtain these results can be broadly grouped into two classes. Syntactic approaches presuppose a complete tableaux or Gentzen system and establish that proof search can be performed in polynomial space. Semantics-driven approaches, on the other hand, directly construct shallow tree models. Both approaches are intimately connected in the case of normal modal logics interpreted over Kripke frames: counter models can usually be derived directly from search trees [16]. It should be noted that this method is not immediately applicable in the non-normal case, where the structure of models often goes far beyond mere graphs.

Using coalgebraic techniques, we have previously shown [29] that the syntactic approach uniformly generalises to a large class of modal logics: starting from a *one-step complete* axiomatisation, we have applied *resolution closure* to obtain complete tableaux systems. Generic PSPACE-bounds follow if the ensuing rule set is *PSPACE-tractable*. Here, we present a different, semantic, set of methods to establish uniform PSPACE bounds by directly constructing shallow models for logics subject to the *one-step polysize model property*, or a variant of the latter. In particular, no axiomatisation of the logic itself is needed.

Apart from the fact that both methods use substantially different techniques, they apply to different classes of examples. While it is e.g. relatively easy to obtain a resolution closed rule set for coalition logic [26], proving the one-step polysize model property for (the coalgebraic semantics of) coalition logic is a non-trivial task. On the other hand, small one-step models are comparatively easy to construct for complex modal logics such as probabilistic modal logic [8] or Presburger modal logic [6] that are not straightforwardly amenable to the syntactic approach via resolution closure, either because no axiomatisation has been given or because the complexity of the axiomatisation makes the resolution closure hard to harness.

Moreover, the present semantic approach to PSPACE-bounds takes a significant step to overcome an important barrier in the coalgebraic treatment of modal logics. Existing decidability and completeness results [25, 5, 28, 29] are limited to *rank-1 logics*, given by axioms whose modal nesting depth is uniformly equal to one. While this already encompasses a large class of examples (including all logics mentioned so far), the semantic model construction in the present paper applies to *non-iterative logics* [21], i.e. logics axiomatised without nested modalities (rank-1 logics additionally exclude top-level propositional variables). Despite the seemingly minute difference between the two classes of logics, this generalisation is not only technically non-trivial but also substantially extends the scope of the coalgebraic method. Besides the modal logic T, the class of non-iterative logics includes e.g. all conditional logics covered in [22] (of which only 4 are rank-1), in particular $CK+MP$ [3], as well as Elgesem's logic of agency [7, 12] and the graded version Tn of T [11].

As in [29], we work in the framework of *coalgebraic modal logic* [25] to obtain results that are parametric in the underlying semantics of particular logics. While normal modal logics are usually interpreted over Kripke frames, non-normal logics see a large variety of different semantics, e.g. probabilistic systems [8], frames with ordered branching [6], game frames [26], or conditional frames [3]. The coalgebraic treatment allows us to encapsulate the semantics in the choice of a *signature functor*, whose coalgebras then play the role of models, leading to results that are unformly applicable to a large class of different logics.

Since the class of *all* coalgebras for a given signature functor can always be completely axiomatised in rank 1 [28], in analogy to the fact that the K-axioms are complete for the class of *all* Kripke frames, the standard coalgebraic approach is not directly applicable to non-iterative logics. To overcome this limitation, we introduce the new concept of interpreting modal logics over coalgebras for *copointed functors*, i.e. functors T equipped with a natural transformation of type $T \to Id$.

In this setting, our main technical tool is to cut back model constructions from modal logics to the level of *one-step logics* which semantically do not involve state transitions, and then amalgamate the corresponding *one-step models* into shallow models for the full modal logic, which ideally can be traversed in polynomial space. For this approach to work, the logic at hand needs to support a small model property for its one-step fragment, the *one-step polysize model property (OSPMP)*. Our first main theorem shows that the OSPMP guarantees decidability in polynomial space. Crucially, the OSPMP is much easier to establish than a shallow model property for the logic itself. To reprove e.g. Ladner's PSPACE upper bound for K, one just observes that to construct a set that intersects n given sets, one needs at most n elements. For the conditional logics CK, $CK+ID$, and $CK+MP$, the OSPMP is similarly easy to check. For other logics, in particular various logics of quantitative uncertainty, the OSPMP can be obtained by sharpening known off-the-shelf results. As a new result, we establish the OSPMP for Elgesem's logic of agency to obtain a previously unknown PSPACE upper bound.

As a by-product of our construction, we obtain NP-bounds for the bounded rank fragments of all logics with the OSPMP, generalizing the corresponding result for the logics K and T from [13] to a large variety of structurally different (non-iterative) logics.

While the OSPMP is usually easy to establish, a weaker property, the *one-step pointwise polysize model property (OSPPMP)*, can be used in cases where the OSPMP fails, provided that the signature functor supports a notion of pointwise smallness for overall exponential-sized one-step models. This allows traversing exponentially branching shallow models in polynomial space by dealing with the successor structures of single states in a pointwise fashion. Our second main result, which yields PSPACE upper bounds for logics with the OSPPMP, is applied to reprove the known PSPACE bound for Presburger modal logic [6] and to derive a new PSPACE bound for Presburger modal logic over reflexive frames, and hence for Tn [11] (which was so far only known to be decidable [10]). The latter result extends straightforwardly to a description logic with role hierarchies, qualified number restrictions, and reflexive roles.

2 Coalgebraic Modal Logic

We recall the coalgebraic interpretation of modal logic and extend it to non-iterative logics using copointed functors.

A *modal signature* Λ is a set of modal operators with associated finite arity. The signature Λ determines two languages: firstly, the *one-step logic* of Λ, whose formulas $\psi, \dots$ (the *one-step formulas*) over a set V of propositional variables are defined by the grammar

$$\psi ::= \bot \mid \psi_1 \wedge \psi_2 \mid \neg\psi \mid L(\phi_1, \dots, \phi_n),$$

where $L \in \Lambda$ is n-ary and the ϕ_i are propositional formulas over V; and secondly, the *modal logic* of Λ, whose set $\mathcal{F}(\Lambda)$ of Λ-*formulas* $\psi, \dots$ is defined by the grammar

$$\psi ::= \bot \mid \psi_1 \wedge \psi_2 \mid \neg\psi \mid L(\psi_1, \dots, \psi_n).$$

Thus, the modal logic of Λ is distinguished from the one-step logic in that it admits nested modalities. The boolean operations $\vee$, $\rightarrow$, $\leftrightarrow$, $\top$ are defined as usual. The *rank* $\mathrm{rank}(\phi)$ of $\phi \in \mathcal{F}(\Lambda)$ is the maximal nesting depth of modalities in ϕ (note however that the notion of rank-1 logic [28, 29] is stricter than suggested by this definition, as it excludes top-level propositional variables in axioms; the latter are allowed only in non-iterative logics). We denote by $\mathcal{F}_n(\Lambda)$ the set of formulas of rank at most n; we refer to the languages $\mathcal{F}_n(\Lambda)$ as *bounded-rank fragments*.

We treat one-step logics as a technical tool in the study of modal logics. However, one-step logics also appear as logics of independent interest in the literature [9, 14, 15]. One of the central ideas of coalgebraic modal logic is that properties of the full modal logic, such as soundness, completeness, and decidability, can be reduced to properties of the much simpler one-step logic. This is also the spirit of the present work, whose core is a construction of polynomially branching shallow models for the modal logic assuming a small model property for the one-step logic.

The semantics of both the one-step logic and the modal logic of Λ are parametrized coalgebraically by the choice of a set functor. The standard setup of coalgebraic modal logic using all coalgebras for a plain set functor covers only *rank-1 logics*, i.e. logics axiomatised one-step formulas [28] (a typical example is the K-axiom $\Box(a \rightarrow b) \rightarrow \Box a \rightarrow \Box b$). Here, we improve on this by considering the class of coalgebras for a given *copointed* set functor, which enables us to cover the more general class of *non-iterative logics*, axiomatised by arbitrary formulas without nested modalities (such as the T-axiom $\Box a \rightarrow a$). We follow a purely semantic approach and hence do not formally consider axiomatisations in the present work (where we do mention axioms, this is for solely explanatory purposes). However, the extended scope of the new framework and its relation to non-iterative modal logics (which can be made precise in the same way as for plain functors and rank-1 logics [28]) will become clear in the examples.

In general, a copointed functor (T, ϵ) consists of a functor $T : \mathsf{Set} \rightarrow \mathsf{Set}$, where Set is the category of sets, and a natural transformation $\epsilon : T \rightarrow Id$. For our present purposes, a slightly restricted notion is more convenient:

Definition 2.1. A *(restricted) copointed functor* S with *signature functor* $S_0 : \mathsf{Set} \rightarrow \mathsf{Set}$ is a subfunctor of $S_0 \times Id$ (where $(S_0 \times Id)X = S_0X \times X$). We say that S is *trivially copointed* if $S = S_0 \times Id$. An S-*coalgebra* $A = (X, \xi)$ consists of a set X of *states* and a *transition* function $\xi : X \rightarrow S_0X$ such that $(\xi(x), x) \in SX$ for all x.

Remark 2.2. The modal logic $\mathcal{F}(\Lambda)$ does not explicitly include propositional variables. These may be regarded as nullary modal operators in Λ; their semantics is then defined over coalgebras for $S_0 \times \mathcal{P}(V)$, where V is the set of variables (cf. also e.g. [28]). We omit discussion of propositional variables in the examples, even in cases like the modal logic of probability that become trivial in the absence of variables; our treatment extends straightforwardly to the case with variables in the manner just indicated.

We view coalgebras as generalized transition systems: the transition function maps a state to a structured set of successors and observations, with the structure prescribed by the signature functor. Thus, the latter encapsulates the branching type of the underlying transition systems. Copointed functors additionally impose local frame conditions that relate a state to the collection of its successors.

Assumption 2.3. We assume w.l.o.g. that S_0 preserves injective maps [2], and even $S_0X \subseteq S_0Y$ in case $X \subseteq Y$, and that S is *non-trivial*, i.e. $SX = \emptyset \Rightarrow X = \emptyset$.

Generalising earlier work (e.g. [17, 19]), *coalgebraic modal logic* abstractly captures the interpretation of modal operators as polyadic predicate liftings [25, 27],

Definition 2.4. An *n-ary predicate lifting* ($n \in \mathbb{N}$) for S_0 is a natural transformation

$$\lambda : \mathcal{Q}^n \to \mathcal{Q} \circ S_0^{op},$$

where $\mathcal{Q}$ denotes the contravariant powerset functor $\mathsf{Set}^{op} \to \mathsf{Set}$ (i.e. $\mathcal{Q}X$ is the powerset $\mathcal{P}(X)$, and $\mathcal{Q}f(A) = f^{-1}[A]$), and $\mathcal{Q}^n$ is defined by $\mathcal{Q}^n X = (\mathcal{Q}X)^n$.

A coalgebraic semantics for Λ is formally defined as a Λ-*structure* $\mathcal{M}$ (*over* S) consisting of a copointed functor S with signature functor S_0 and an assignment of an n-ary predicate lifting $[\![L]\!]$ for S_0 to every n-ary modal operator $L \in \Lambda$. When S is trivially copointed, we will also call $\mathcal{M}$ a *simple Λ-structure (over S_0). We fix the notation Λ, $\mathcal{M}$, S, S_0 throughout the paper.* The semantics of the modal language $\mathcal{F}(\Lambda)$ is then given in terms of a satisfaction relation $\models_C$ between states x of S-coalgebras $C = (X, \xi)$ and $\mathcal{F}(\Lambda)$-formulas over V. The relation $\models_C$ is defined inductively, with the usual clauses for boolean operators. The clause for an n-ary modal operator L is

$$x \models_C L(\phi_1, \dots, \phi_n) \;\Leftrightarrow\; \xi(x) \in [\![L]\!]([\![\phi_1]\!]_C, \dots, [\![\phi_n]\!]_C)$$

where $[\![\phi]\!]_C = \{x \in X \mid x \models_C \phi\}$. We drop the subscripts C when clear from the context. Our main interest is in the (local) *satisfiability problem* over $\mathcal{M}$:

Definition 2.5. An $\mathcal{F}(\Lambda)$-formula ϕ is *satisfiable* if there exist an S-coalgebra C and a state x in C such that $x \models_C \phi$. Dually, ϕ is *valid* if $x \models_C \phi$ for all C, x.

Contrastingly, the semantics of the one-step logic is given in terms of satisfaction relations $\models^1_{X,\tau}$ between elements $t \in S_0X$ and one-step formulas over V, where X is a set and τ is a $\mathcal{P}(X)$-*valuation* for V, i.e. a map $\tau : V \to \mathcal{P}(X)$. The valuation τ canonically induces an interpretation of $[\![\phi]\!]^0\tau \subseteq X$ of propositional formulas ϕ over V. We write $X, \tau \models^0 \phi$ if $[\![\phi]\!]^0\tau = X$. The relation $\models^1_{X,\tau}$ is then defined by the usual clauses for boolean operators, and

$$t \models^1_{X,\tau} L(\phi_1, \dots, \phi_n) \;\Leftrightarrow\; t \in [\![L]\!]([\![\phi_1]\!]^0\tau, \dots, [\![\phi_n]\!]^0\tau).$$

Note in particular that the semantics of the one-step logic does not involve a notion of state transition.

Definition 2.6. A *one-step model* (X, τ, t, x) over V consists of a set X, a $\mathcal{P}(X)$-valuation τ for V, $t \in S_0X$, and $x \in X$ such that $(t, x) \in SX$. The latter condition is vacuous if S is trivially copointed, in which case we omit the mention of x. For a one-step formula ψ over V, (X, τ, t, x) is a *one-step model of* ψ if $t \models^1_{X,\tau} \psi$.

We recall some basic notation:

Definition 2.7. We denote the set of propositional formulas over a set Z, generated by the basic connectives $\neg$ and $\wedge$, by $\mathsf{Prop}(Z)$. We use variables ϵ etc. to denote either nothing or $\neg$. Thus, a *literal* over Z is a formula of the form ϵa, with $a \in Z$. A *(conjunctive) clause* is a finite, possibly empty, disjunction (conjunction) of literals. We denote by $\Lambda(Z)$ the set $\{L(a_1, \dots, a_n) \mid L \in \Lambda \; n\text{-ary}, a_1, \dots, a_n \in Z\}$.

In the above notation, the set of one-step formulas over V is $\mathsf{Prop}(\Lambda(\mathsf{Prop}(V)))$. The one-step logic may alternatively be presented in terms of pairs of formulas separating out the lower propositional layer:

Definition 2.8. A *one-step pair* (ϕ, ψ) over V consists of formulas $\psi \in \mathsf{Prop}(\Lambda(V))$ and $\phi \in \mathsf{Prop}(V)$. A one-step model (X, τ, t, x) is a *one-step model of* (ϕ, ψ) if $X, \tau \models^0 \phi$ and $t \models^1_{X,\tau} \psi$.

In analogy to the equivalence between axioms and one-step rules described in [28], one-step pairs and one-step formulas may replace each other for purposes of satisfiability:

Lemma 2.9. *For every one-step pair (ϕ, ψ) over V with ϕ satisfiable, there exists a $\mathsf{Prop}(V)$-substitution σ such that the one-step formula $\psi\sigma$ is equivalent to (ϕ, ψ) in the sense that if $t \models^1_{X,\tau} \psi\sigma$ then $t \models^1_{X,\sigma\tau} (\phi, \psi)$, and if $t \models^1_{X,\tau} (\phi, \psi)$ then $t \models^1_{X,\tau} \psi\sigma$ (and $\sigma\tau = \tau$). Here, $\sigma\tau$ denotes the $\mathcal{P}(X)$-valuation taking a to $[\![\sigma(a)]\!]^0\tau$.*

Conversely, we have, for $\psi \in \mathsf{Prop}(\Lambda(\mathsf{Prop}(V)))$, an equivalent one-step pair (ϕ, ψ_1) over $V \cup W$, where ψ decomposes as $\psi \equiv \psi_1\sigma$, with $\psi_1 \in \mathsf{Prop}(\Lambda(W))$, σ a $\mathsf{Prop}(V)$-substitution, and $V \cap W = \emptyset$, and where ϕ is the conjunction of the formulas $a \leftrightarrow \sigma(a)$, $a \in W$. Here, restricting valuations to V induces a bijection between one-step models of (ϕ, ψ_1) and one-step models of ψ.

Proof. The second claim is clear. The first claim is proved as follows. As in [28], let κ be a satisfying truth valuation for ϕ and put $\sigma(a) = a \wedge \phi$ if $\kappa(a) = \bot$, and $\sigma(a) = \phi \to a$ otherwise. Then $\phi\sigma$ and the formulas $\phi \to (a \leftrightarrow \sigma(a))$, for $a \in V$, are tautologies [28]. Both directions of the first claim now follow straightforwardly. $\Box$

The coalgebraic approach subsumes many interesting modal logics, including e.g. graded and probabilistic modal logics and coalition logic [29]. Below, we present the most basic examples, the modal logics K and T, as well as various conditional logics and logics of quantitative uncertainty. The treatment of Elgesem's modal logic of agency is deferred to Sect. 4.

Example 2.10. 1. *Modal logic K:* Let $\Lambda = \{\Box\}$, with $\Box$ a unary modal operator. We define a simple Λ-structure over the covariant powerset functor $\mathcal{P}$ (i.e. $\mathcal{P}X$ is powerset, and $\mathcal{P}f(A) = f[A]$) by putting $[\![\Box]\!]_X(A) = \{B \in \mathcal{P}X \mid B \subseteq A\}$. Naturality of $[\![\Box]\!]$ is just the equivalence $f[B] \subseteq A \iff B \subseteq f^{-1}[A]$.

$\mathcal{P}$-coalgebras are Kripke frames, and $\mathcal{P}$-models are Kripke models. The modal logic of Λ is precisely the modal logic K, equipped with its standard Kripke semantics. Contrastingly, a one-step formula over V is a propositional combination of atoms of the form $\Box\phi$, where $\phi \in \mathsf{Prop}(V)$. For $A \in \mathcal{P}X$, we have $A \models^1_{X,\tau} \Box\phi$ iff $A \subseteq [\![\phi]\!]^0\tau$. One easily checks that the one-step logic is NP-complete, while the modal logic K is PSPACE-complete [20].

2. *Modal logic T:* The logic T has the same syntax as K. Its coalgebraic semantics is a structure over the copointed functor R with signature functor $\mathcal{P}$, given by

$$RX = \{(A, x) \in \mathcal{P}X \times X \mid x \in A\}.$$

Thus, R-coalgebras are reflexive Kripke frames. The interpretation of $\Box$ is defined as for K. (Axiomatically, T is determined by the non-iterative axiom $\Box a \to a$.)

3. *Conditional logic CK:* The signature of conditional logic has a single binary modal operator $\Rightarrow$, written in infix notation. Formulas $\phi \Rightarrow \psi$ are read as non-monotonic conditionals. The semantics of the conditional logic *CK* [3] is given by a simple structure over the functor *Cf* given by $Cf(X) = (\mathcal{Q}X \to \mathcal{P}X)$, with $\to$ denoting function space and $\mathcal{Q}$ contravariant powerset, cf. Definition 2.4. *Cf*-coalgebras are *conditional frames* [3]. The operator $\Rightarrow$ is interpreted over *Cf* by

$$[\![\Rightarrow]\!]_X(A, B) = \{f : \mathcal{Q}X \to \mathcal{P}X \mid f(A) \subseteq B\}.$$

4. *Conditional logic CK+ID:* The conditional logic *CK+ID* [3] extends *CK* with the rank-1 axiom $a \Rightarrow a$, referred to as *ID*. The semantics of *CK+ID* is modelled by restricting the structure for *CK* to the subfunctor Cf_{ID} of *Cf* defined by

$$Cf_{ID}(X) = \{f \in Cf(X) \mid \forall A \in \mathcal{Q}X.\, f(A) \subseteq A\}.$$

5. *Conditional Logic CK+MP:* The logic *CK+MP* [3] extends *CK* with the non-iterative axiom

$$(MP) \quad (a \Rightarrow b) \to (a \to b).$$

(This axiom is undesirable in default logics, but reasonable in relevance logics.) Semantically, this amounts to passing from the functor *Cf* to the copointed functor Cf_{MP} with signature functor *Cf*, defined by

$$Cf_{MP}(X) = \{(f, x) \in Cf(X) \times X \mid \forall A \subseteq X.\, x \in A \Rightarrow x \in f(A)\}.$$

6. *Modal logics of quantitative uncertainty:* The modal signature of *likelihood* has n-ary modal operators $\sum_{i=1}^{n} a_i l(\_) \geq b$ for $a_1, \dots, a_1, b \in \mathbb{Q}$. The terms $l(\phi)$ are called *likelihoods*. The interpretation of likelihoods varies. E.g. the semantics of the *modal logic of probability* [8] is modelled coalgebraically by a structure over the (finite) distribution functor D_ω, where $D_\omega X$ is the set of finitely supported probability distributions on X, and $D_\omega f$ acts as image measure formation. Coalgebras for D_ω are probabilistic transition systems (i.e. Markov chains). Likelihoods are interpreted as probabilities; i.e.

$$[\![\textstyle\sum_{i=1}^{n} a_i \cdot l(\_) \geq b]\!]_X(A_1, \dots, A_n) = \{P \in D_\omega X \mid \textstyle\sum_{i=1}^{n} a_i P(A_i) \geq b\}.$$

Alternatively, likelihoods may be interpreted as *upper probabilities* [14], i.e. the functor D_ω is replaced by $\mathcal{P} \circ D_\omega$, and in the above definition, $P(A_i)$ is replaced by $\mathfrak{P}^*(A_i)$, where for $\mathfrak{P} \in \mathcal{P}D_\omega X$, $\mathfrak{P}^*(A) = \sup\,\{PA \mid P \in \mathfrak{P}\}$. This setting describes situations where agents are unsure about the actual probability distribution. Further alternative notions of likelihood include Dempster-Shafer belief functions and Dubois-Prade possibility measures [15]. An extension of the modal signature of likelihood is the modal signature of *expectation* [15], where instead of likelihoods one more generally considers expectations $e(\sum_{j=1}^{n_k} b_{ij}\phi_{ij})$. Here, linear combinations of formulas represent *gambles*, i.e. real-valued outcome functions, where the payoff of ϕ is the characteristic function of ϕ. The exact definition of expectation depends on the underlying notion of likelihood.

One-step logics of quantitative uncertainty are often considered to be of independent interest. E.g. the one-step logic of probability, i.e. a logic without nesting of likelihoods that talks only about a single probability distribution, is introduced independently [9] and only later extended to a full modal logic [8]. In fact, logics of expectation [15] and the logic of upper probability [14] so far appear in the literature only as one-step logics; the corresponding modal logics are of interest as natural variations of the modal logic of probability.

Convention 2.11. We assume that Λ is equipped with a size measure, thus inducing a size measure for $\mathcal{F}(\Lambda)$. For one-step formulas ϕ over V, we assume w.l.o.g. that $|V| \leq \text{size}(\phi)$. For finite X, we assume given a representation of elements $(t,x) \in SX$ as strings of size size(t,x) over some finite alphabet. We do not require that all elements of SX are representable, nor that all strings denote elements of SX. When S is trivially copointed, we represent only $t \in S_0(X)$. We require that inclusions $SX \subseteq SY$ induced according to Assumption 2.3 by inclusions $X \subseteq Y$ into a finite base set Y preserve representable elements and increase their size by at most $\log |Y|$.

We make these issues explicit for the above examples:

Example 2.12. 1. *Modal logics K and T:* For X finite, elements of $\mathcal{P}X$ are represented as lists of elements of X.

2. *Conditional logics CK and CK+ID:* For X finite, elements of $\mathcal{Q}X \to \mathcal{P}X$ are represented as partial maps $f_0 : \mathcal{Q}X \rightharpoonup \mathcal{P}X$; such an f_0 represents the total map f that extends f_0 by $f(A) = \emptyset$ when $f_0(A)$ is undefined. (The use of partial maps avoids exponential blowup.)

3. *Conditional logic $CK + MP$:* For X finite, a pair (f_0, x) consisting of a partial map $f_0 : \mathcal{Q}X \rightharpoonup \mathcal{P}X$ and $x \in X$ represents the pair (f, x), where $f : \mathcal{Q}X \to \mathcal{P}X$ extends f_0 by $f(A) = A \cap \{x\}$ in case $f_0(A)$ is undefined.

4. *Modal logics of quantitative uncertainty:* Suitable compact representations are described in [9, 14, 15].

3 Polynomially branching shallow models

We now turn to the announced construction of polynomially branching shallow models for modal logics whose one-step logic has a small model property; this construction leads to a PSPACE decision procedure.

Definition 3.1. We say that $\mathcal{M}$ has the *one-step polysize model property (OSPMP)* if there exist polynomials p and q such that, whenever a one-step pair (ϕ, ψ) over V has a one-step model (X, τ, t, x), then it has a one-step model (Y, κ, s, y) such that $|Y| \leq p(|\psi|)$, (s, y) is representable with size$(s, y) \leq q(|\psi|)$, and $y \in \kappa(a)$ iff $x \in \tau(a)$ for all $a \in V$.

In analogy to the transition between rules and axioms described in [28], one-step pairs are interchangeable with one-step formulas. In particular, we have

Proposition 3.2. *The Λ-structure $\mathcal{M}$ has the OSPMP iff there exist polynomials p, q such that, whenever a one-step formula ψ over V has a one-step model (X, τ, t, x), then it has a one-step model (Y, κ, s, y) such that $y \in \kappa(a)$ iff $x \in \tau(a)$ for all $a \in V$, $|Y| \leq p(|\psi_1|)$, and (s, y) is representable with* size$(s, y) \leq q(|\psi_1|)$*, where $\psi \equiv \psi_1\sigma$ with $\psi_1 \in \mathsf{Prop}(\Lambda(W))$ and σ a $\mathsf{Prop}(V)$-substitution.*

Proof. Only if: Let (X, τ, t, x) be a one-step model of a one-step formula ψ over V. By Lemma 2.9, ψ is equivalent to a one-step pair of the form (ϕ, ψ_1), with ψ_1 as in the statement. By the OSPMP, (ϕ, ψ_1) has a one-step model (Y, κ, s, y) such that $|Y| \leq p(|\psi_1|)$, size$(s, y) \leq q(|\psi_1|)$, and $y \in \kappa(a)$ iff $x \in \tau(a)$ for all $a \in V$; by Lemma 2.9, this model gives rise to a one-step model of ψ with the components Y, s, y unchanged.

If: Let (X, τ, t, x) be a one-step model of a one-step pair (ϕ, ψ) over V. By Lemma 2.9, (ϕ, ψ) is equivalent to a one-step formula of the form $\psi\sigma$, where σ is a $\mathsf{Prop}(V)$-substitution. By assumption, $\psi\sigma$ has a one-step model (Y, κ, s, y) such that $|Y| \leq p(|\psi|)$, $\mathrm{size}(s, y) \leq q(|\psi|)$, and $y \in \kappa(a)$ iff $x \in \tau(a)$ for all $a \in V$. By Lemma 2.9, this model gives rise to a one-step model of (ϕ, ψ) with the components Y, s, y unchanged. □

Both formulations of the OSPMP easily reduce to the case that ψ is a conjunctive clause.

Remark 3.3. It is shown in [28] that the one-step logic always has an exponential-size model property: a one-step formula ψ over V has a one-step model iff it has a one-step model with carrier set $\mathcal{P}(V)$.

We are now ready to prove the shallow model theorem.

Definition 3.4. A *supporting Kripke frame* of an S-coalgebra (X, ξ) is a Kripke frame (X, R) such that for each $x \in X$,

$$\xi(x) \in S_0\{y \mid xRy\} \subseteq S_0X$$

(equivalently $(\xi(x), x) \in S\{y \mid xRy\}$). A state $x \in X$ is a *loop* if xRx.

Theorem 3.5 (Shallow model property)**.** *If $\mathcal{M}$ has the OSPMP, then $\mathcal{F}(\Lambda)$ has the* polynomially branching shallow model property*: There exist polynomials p, q such that every satisfiable $\mathcal{F}(\Lambda)$-formula ψ is satisfiable at the root of an S-coalgebra (X, ξ) which has a supporting Kripke frame (X, R) such that removing all loops from (X, R) yields a tree of depth at most* $\mathrm{rank}(\psi)$ *and branching degree at most $p(|\psi|)$, and $(\xi(x), x) \in S\{y \mid xRy\}$ is representable with* $\mathrm{size}(\xi(x), x) \leq q(|\psi|)$.

Definition 3.6. For $x \in X$ and a $\mathcal{P}(X)$-valuation τ, we put $\mathsf{Th}_\tau(x) \equiv \bigwedge_{x \in \tau(a)} a \wedge \bigwedge_{x \notin \tau(a)} \neg a$.

Proof of Theorem 3.5. Induction over the rank of ψ. If $\mathrm{rank}(\psi) = 0$, then ψ evaluates to $\top$ and hence is satisfied in a singleton S-coalgebra (X, ξ), which exists by Assumption 2.3.

Now let $\mathrm{rank}(\psi) = n+1$. Let z_0 be a state in an S-coalgebra (Z, ζ) such that $z_0 \models_{(Z,\zeta)} \psi$. Let $MSub(\psi)$ denote the set of subformulas of ψ occuring in ψ within the scope of a modal operator, let V be the set of variables a_ρ, indexed over $\rho \in MSub(\psi)$, and let σ denote the substitution taking a_ρ to ρ for all ρ. Let $\bar{\psi}$ be the conjunction of all literals $\epsilon L(a_{\rho_1}, \dots, a_{\rho_n})$ such that $L(\rho_1, \dots, \rho_n)$ is a subformula of ψ and $z_0 \models_{(Z,\zeta)} \epsilon L(\rho_1, \dots, \rho_n)$. (Recall that ϵ denotes either nothing or negation.) Moreover, let ϕ denote the propositional theory of σ, i.e. the conjunction of all clauses χ over V such that $\chi\sigma$ is $\mathcal{L}$-valid.

Then $(Z, \kappa, \zeta(z_0), z_0)$ is a one-step model of $(\phi, \bar{\psi})$, where $\kappa(a) = [\![\sigma(a)]\!]_{(Z,\zeta)}$. By the OSPMP, it follows that $(\phi, \bar{\psi})$ has a one-step model (Y, τ, t, x_0) of polynomial size in $|\bar{\psi}|$ such that for all $\rho \in MSub(\psi_1)$, $x_0 \in \tau(a_\rho)$ iff $z_0 \in \kappa(\rho)$, which in turn is equivalent to $z_0 \models_{(Z,\zeta)} \rho$.

From this model, we now construct a shallow model (X, ξ) for ψ. To begin, note that $\mathsf{Th}_\tau(y)\sigma$ is $\mathcal{L}$-satisfiable for every $y \in Y$. For suppose not; then $\neg\mathsf{Th}_\tau(y)\sigma$ is $\mathcal{L}$-valid, hence $\neg\mathsf{Th}_\tau(y)$ is a conjunct of ϕ. Thus, $Y, \tau \models^0 \neg\mathsf{Th}_\tau(y)$, in contradiction to the fact that $y \in [\![\mathsf{Th}_\tau(y)]\!]\tau$ by construction. By induction, we thus have, for every $y \in Y$, a shallow model (X_y, ξ_y) of $\mathsf{Th}_\tau(y)\sigma$, where we may assume $y \in X_y$ and $y \models_{(X_y,\xi_y)} \mathsf{Th}_\tau(y)\sigma$, with depth at most $\mathrm{rank}(\mathsf{Th}_\tau(y)\sigma) = n$. We take (X, ξ) as the disjoint union of the (X_y, ξ_y) over $y \in Y - \{x_0\}$, extended by the state x_0, for which we put $\xi(x_0) = t \in S_0Y \subseteq S_0X$.

We have to verify that $x_0 \models_{(X,\xi)} \psi$. We will prove the stronger statement $x_0 \models_{(X,\xi)} \bar{\psi}\sigma$, i.e.

$$t \models^1_{X,\theta} \bar{\psi}, \tag{1}$$

where $\theta(a_\rho) = [\![\rho]\!]_{(X,\xi)}$ for $\rho \in MSub(\psi)$.

By induction over χ and naturality of predicate liftings, $y \models_{(X,\xi)} \chi$ iff $y \models_{(X_y,\xi_y)} \chi$ for $y \in Y - \{x_0\}$ and for every formula χ. In particular, $y \models_{(X,\xi)} \mathsf{Th}_\tau(y)\sigma$ for all $y \in Y - \{x_0\}$, i.e.

$$y \models_{(X,\xi)} \rho \iff y \in \tau(a_\rho) \tag{2}$$

for all $\rho \in MSub(\psi)$. We prove by induction over $\rho \in MSub(\psi)$ that

$$x_0 \models_{(X,\xi)} \rho \iff x_0 \in \tau(a_\rho), \tag{3}$$

which in connection with (2) yields

$$[\![\rho]\!]_{(X,\xi)} \cap Y = \tau(a_\rho). \tag{4}$$

The steps for boolean operations are straightforward. For $L(\rho_1, \dots, \rho_n) \in MSub(\psi)$, we have

$$\begin{aligned} & x_0 \models_{(X,\xi)} L(\rho_1, \dots, \rho_n) \\ \iff & t \in [\![L]\!]_Y([\![\rho_i]\!]_{(X,\xi)} \cap Y)_{i=1,\dots,n} = [\![L]\!]_Y(\tau(a_{\rho_1}), \dots, \tau(a_{\rho_n})) \\ \iff & t \models^1_{(Y,\tau)} L(a_{\rho_1}, \dots, a_{\rho_n}), \end{aligned}$$

using naturality of $[\![L]\!]$ in the first step and the inductive hypothesis in the shape of (4) in the subsequent equality. Since $t \models_{(Y,\tau)} \bar{\psi}$, the last statement is equivalent to $z_0 \models_{(Z,\zeta)} L(\rho_1, \dots, \rho_n)$. By the definition of κ, this is equivalent to $z_0 \in \kappa(a_{L(\rho_1,\dots,\rho_n)})$, which in turn is equivalent to $x_0 \in \tau(a_{L(\rho_1,\dots,\rho_n)})$ by construction of (Y, τ, t, x_0).

By (4) and naturality of predicate liftings, our remaining goal (1) reduces to $t \models^1_{Y,\tau} \bar{\psi}$, which holds by construction.

Finally, we have to establish that the overall branching degree of the model is polynomial in $|\psi|$. The model is recursively constructed from polynomial-size one-step models for pairs whose second components are conjunctive clauses over atoms $L(a_{\rho_1}, \dots, a_{\rho_n})$, where $L(\rho_1, \dots, \rho_n)$ is a subformula of ψ. Such conjunctive clauses are of at most quadratic size in $|\psi|$ (even $O(|\psi| \log |\psi|)$ if subformulas of ψ are represented by pointers into ψ); this proves the claim. □

Remark 3.7. While it is to be expected that the construction of polynomially branching models depends on a condition like the OSPMP, it does not seem to be the case that the precise formulation of this condition is implicit in the literature (not even for the trivially copointed case). Note in particular that the polynomial bound depends only on the second component of a one-step pair. This is crucial, as the first component of the one-step pair constructed in the above proof may be of exponential size. When we say in the introduction that the OSPMP can be obtained from off-the-shelf results (e.g. [9, 14, 15]), we refer to polynomial-size model theorems in which the polynomial bound depends, in the notation of Proposition 3.2, on $|\psi|$, which may be exponentially larger than $|\psi_1|$; typically, only an inspection of the given proofs shows that the bound can be sharpened to be polynomial in $|\psi_1|$.

The proof of Theorem 3.5 leads to the following nondeterministic decision procedure.

Algorithm 3.8. (Decide satisfiability of an $\mathcal{F}(\Lambda)$-formula ψ) Let $\mathcal{M}$ have the OSPMP, and let p, q be polynomial bounds as in Definition 3.1.

1. If $\mathrm{rank}(\psi) = 0$, terminate successfully if ψ evaluates to $\top$, else unsuccessfully. Otherwise:
2. Take V and σ as in the proof of Theorem 3.5, and guess a conjunctive clause $\bar{\psi}$ over $\Lambda(V)$ containing for each subformula $L(\rho_1, \dots, \rho_n)$ of ψ either $L(a_{\rho_1}, \dots, a_{\rho_n})$ or $\neg L(a_{\rho_1}, \dots, a_{\rho_n})$ such that $\bar{\psi}\sigma$ propositionally entails ψ.
3. Guess a $\mathcal{P}(Y)$-valuation τ for V and $(t, x) \in SY$ with $\mathrm{size}(t, x) \leq q(|\bar{\psi}|)$, where $Y = \{1, \dots, p(|\bar{\psi}|)\}$, such that $t \models^1_{Y,\tau} \bar{\psi}$.
4. For each $y \in Y$, check recursively that $\mathsf{Th}_\tau(y)\sigma$ is satisfiable.

Since the rank decreases with each recursive call, the above algorithm can be implemented in polynomial space, provided that Step 3 can be performed in polynomial space.

Definition 3.9. The *one-step model checking problem* is to check, given a string s, a finite set X, $A_1, \dots, A_n \subseteq X$, and $L \in \Lambda$ n-ary, whether s represents some $(t, x) \in SX$ and whether $t \in [\![L]\!]_X(A_1, \dots, A_n)$.

This property and the above algorithm lead to a PSPACE bound for the modal logic. Moreover, for bounded-rank fragments, the polynomially branching shallow model property becomes a polynomial size model property, thus leading to an NP upper bound:

Corollary 3.10. *Let $\mathcal{M}$ have the OSPMP.*

1. If one-step model checking is in PSPACE, then the satisfiability problem of $\mathcal{F}(\Lambda)$ is in PSPACE.

2. If one-step model checking is in P, then the satisfiability problem of $\mathcal{F}_n(\Lambda)$ is in NP for every $n \in \mathbb{N}$.

Proof. 1: By Algorithm 3.8.

2: Let p and q be polynomial bounds on the branching degree of supporting Kripke frames and on the size of successor structures $\xi(x)$ as guaranteed by Theorem 3.5. Let $n \in \mathbb{N}$. Then by Theorem 3.5, every satisfiable formula $\psi \in \mathcal{F}_n(\Lambda)$ is satisfiable in a model (X, ξ) such that $|X| \leq \sum_{i=0}^{n} p(|\psi|)^i =: N$ and $\mathrm{size}(\xi(x)) \leq \log(N) q(|\psi|)$ for all $x \in X$, where the second inequality relies also on Convention 2.11. Thus, the entire representation size of the model (X, ξ) is bounded by $M := N \log(N) q(|\psi|)$, which is polynomial in $|\psi|$. Thus, the following non-deterministic algorithm decides satisfiability of ψ in polynomial time:

1. Guess a model (X, ξ) of size at most M
2. Check that (X, ξ) is an S-coalgebra.
3. Check that $[\![\psi]\!]_{(X,\xi)} \neq \emptyset$.

The second step can be performed in polynomial time because one-step model checking is in P. The third step can be performed in polynomial time by recursively computing extensions $[\![\phi]\!]_{(X,\xi)}$, again because one-step model checking is in P. □

This generalises results for the modal logics K and T established in [13].

Example 3.11. 1. *Modal logics K and T:* One-step model checking for K and T amounts to checking a subset inclusion and, in the case of T, additionally an elementhood; this is clearly in P. To verify the OSPMP for K, let (X, τ, A) be a one-step model of a one-step pair (ϕ, ψ) over V; w.l.o.g. ψ is a conjunctive clause over atoms $\Box a$, where $a \in V$. For $\neg\Box a$ in ψ, there exists $x_a \in A$ such that $x_a \notin \tau(a)$. Taking Y to be the set of these x_a, we obtain a polynomial-size one-step model (Y, τ_Y, Y) of (ϕ, ψ), where $\tau_Y(a) = \tau(a) \cap Y$ for all a. The construction for T is the same, except that the point x of the original one-step model (X, τ, A, x) is retained in the carrier set Y, and becomes the point of the small model. By Corollary 3.10, this reproves Ladner's PSPACE upper bounds for K and T [20], as well as Halpern's NP upper bounds for bounded-rank fragments [13].

2. *Conditional logic:* It is easy to see that one-step model checking for CK, $CK+ID$, and $CK+MP$ is in P. (In particular, deciding whether a given string represents an element of $Cf_{ID}(X)$ just amounts to checking subset inclusions. Moreover, deciding whether $(f, x) \in Cf_{MP}(X)$, i.e. whether $x \in A$ implies $x \in f(A)$, can be done in polynomial time thanks to the choice of default value for f; cf. Example 2.12.2.)

To prove that CK has the OSPMP, let (X, τ, f) be a one-step model of a one-step pair (ϕ, ψ), where w.l.o.g. ψ is a conjunctive clause $\bigwedge_{i=1}^{n} \epsilon_i(a_i \Rightarrow b_i)$. If $\tau(a_i) \neq \tau(a_j)$, fix an element y_{ij} in the symmetric difference of $\tau(a_i)$ and $\tau(a_j)$. Moreover, if ϵ_i is negation, fix $z_i \in f(\tau(a_i)) \setminus \tau(b_i)$. Let Y be the set of all y_{ij} and all z_i. Let τ_Y be the $\mathcal{P}(Y)$-valuation defined by $\tau_Y(v) = \tau(v) \cap Y$, and let $f_Y \in Cf(Y)$ be represented by the partial map taking $\tau_Y(a_i)$ to $f(\tau(a_i)) \cap Y$ for all i (this is well-defined by construction of Y). Then (Y, τ_Y, f_Y) is a one-step model of (ϕ, ψ). The cardinality of Y is quadradic in ψ, and the representation size of f_Y is polynomial.

Thanks to the choice of default value, this construction of polynomial-size one-step models works also for $CK+ID$. The construction for $CK+MP$ is almost identical, except that the point x of (X, τ, f, x) is retained in the small one-step model (Y, τ_Y, f_Y, x); here, $(f_Y, x) \in Cf_{MP}(Y)$ due to the different choice of default value.

We thus obtain that CK, $CK+ID$, and $CK+MP$ are in PSPACE (hence PSPACE-complete, as these logics contain K and — in the case of $CK+MP$— T, respectively, as sublogics). This has previously been proved using a detailed analysis of a labelled sequent calculus [22] (the method of [22] yields an explicit polynomial bound on space usage which is not matched by the generic algorithm). The NP upper bound for bounded-rank fragments of CK, $CK+ID$, and $CK+MP$ arising from Corollary 3.10.2 is, to our knowledge, new.

3. *Modal logics of quantitative uncertainty:* Polynomial size model properties for one-step logics have been proved for the logic of probability [9], the logic of upper probability [14], and various logics of expectation [15]. As indicated in Remark 3.7, the polynomial bounds are stated in the cited work as depending on the size of the entire one-step formula ψ; however, inspection of the given proofs shows that the polynomial bound in fact depends only on the number of likelihoods or expectations in ψ, respectively, and on the representation size of the largest coefficient. By Proposition 3.2, it follows that the respective logics have the OSPMP. Suitable complexity estimates for one-step model checking are also found in the cited work.

By the above results, it follows that the respective modal logics of quantitative uncertainty are in PSPACE (hence PSPACE-complete, as one can embed KD by mapping $\Diamond$ to $l(\_) > 0$), and in NP when the modal nesting depth is bounded. For the modal logic of probability, a proof of the PSPACE upper bound is sketched in [8]. The PSPACE upper bounds for the remaining cases (e.g. the modal logic of upper probability and the various modal logics of expectation) seem to be new, if only for the reason that only the one-step versions of these logics appear in the literature. Similarly, all NP upper bounds for bounded-rank fragments are, to our knowledge, new. Moreover, the upper bounds extend easily to modal logics of uncertainty with non-iterative axioms, e.g. an axiom $a \to l(a) \geq p$ which states that the present state remains stationary with likelihood at least p.

4 Extended Example: Elgesem's modal logic of agency

There have been numerous approaches to capturing the notion of agents bringing about certain states of affairs, one of the most recent ones being Elgesem's modal logic of agency ([7] and references therein, [12]). Modal logics of agency play a role e.g. in planning and task assignment in multi-agent systems (cf. e.g. [4, 18]).

Elgesem defines a logic with two modalities E and C (in general indexed over agents; all results below easily generalise to the multi-agent case), read 'the agent brings about' and 'the agent is capable of realising', respectively. The semantics is given by a class of conditional frames $(X, f : X \to \mathcal{Q}X \to \mathcal{P}X)$ (Example 2.10.3), called *selection function models* in this context. The clauses for the modal operators are

$$\begin{aligned} x \models E\phi \quad &\text{iff} \quad x \in f(w)(\llbracket\phi\rrbracket) \quad \text{and} \\ x \models C\phi \quad &\text{iff} \quad f(w)(\llbracket\phi\rrbracket) \neq \emptyset. \end{aligned}$$

The relevant class of selection function models (X, f) is defined by the conditions

(E1)	$f(x)(X) = \emptyset$
(E2)	$f(x)(A) \cap f(x)(B) \subseteq f(x)(A \cap B)$
(E3)	$f(x)(A) \subseteq A$.

It is shown in [7, 12] that the logic of agency is completely axiomatised by $\neg C\top$, $\neg C\bot$, $Ea \wedge Eb \to E(a \wedge b)$, $Ea \to a$, and $Ea \to Ca$. Notably, the agent is incapable of realising what is logically necessary ($\neg C\top$), i.e. the notion of realising a state of affairs entails actual attributability (this axiom is weaker than previous formulations using *avoidability*; cf. the baby food example in [7]). Monotonicity is not imposed. The axiom $\neg C\bot$ is due to [12].

Most of the information in selection function models (motivated by philosophical considerations in [7]) is irrelevant for the semantics of E and C: one only needs to know whether $f(x)(A)$ is non-empty, and whether it contains x. Moreover, the selection function semantics fails to be coalgebraic, as the naturality condition fails for the (generalised) predicate lifting implicit in the clause for E. Both problems are easily remedied by moving to the following coalgebraic semantics: put $3 = \{\bot, *, \top\}$ (to represent the cases $f(x)(A) = \emptyset$, $x \notin f(x)(A) \neq \emptyset$, and $x \in f(x)(A)$, respectively), and take as signature functor the 3-*valued neighborhood functor* N_3 given by $N_3(X) = \mathcal{Q}(X) \to 3$ (with $\mathcal{Q}(X)$ denoting contravariant powerset). We define the copointed functor $\mathcal{A}$ as the subfunctor of $N_3 \times Id$ such that $(f, x) \in \mathcal{A}(X)$ iff for all $A, B \subseteq X$,

(E1′)	$f(X) = \bot$
(E2′)	$f(A) \wedge f(B) \leq f(A \cap B)$
(E3a′)	$f(\emptyset) = \bot$
(E3b′)	$f(A) = \top \implies x \in A$

where $\wedge$ and $\leq$ refer to the ordering $\bot < * < \top$. We define a structure over $\mathcal{A}$ for the modal logic of agency by

$$[\![E]\!]_X A = \{f : \mathcal{Q} \to 3 \mid f(A) = \top\}$$
$$[\![C]\!]_X A = \{f : \mathcal{Q} \to 3 \mid f(A) \neq \bot\}.$$

Proposition 4.1. *A formula of the modal logic of agency is satisfiable in a selection function model iff it is satisfiable over $\mathcal{A}$.*

Proof. 'Only if:' Given a selection function model (X, f), define an N_3-coalgebra $(X, \tilde{f})$ by

$$\tilde{f}(x)(A) = \begin{cases} \top & \text{if } x \in f(x)(A) \\ * & \text{if } x \notin f(x)(A) \neq \emptyset \\ \bot & \text{if } f(x)(A) = \emptyset. \end{cases}$$

It is clear that $(X, \tilde{f})$ is an $\mathcal{A}$-coalgebra and that $x \in X$ satisfies the same formulas in $(X, \tilde{f})$ as in (X, f).

'If': Let (X, f) be an $\mathcal{A}$-coalgebra. We can assume that $|[\![\phi]\!]_{(X,f)}| \neq 1$ for all formulas ϕ (otherwise, form the coproduct of (X, f) with itself, so that each state has a twin satisfying the same formulas). We define a selection function model $(X, \bar{f})$ by

$$\bar{f}(x)(A) = \begin{cases} A & \text{if } f(x)(A) = \top \\ A - \{x\} & \text{if } f(x)(A) = * \\ \emptyset & \text{if } f(x)(A) = \bot. \end{cases}$$

It is clear that $(X, \bar{f})$ satisfies E1–E3. One shows by induction over the formula structure that $x \in X$ satsifies the same formulas in $(X, \bar{f})$ as in (X, f), with the only non-trivial point being that in the step for the modal operator C, one has to note that, by the above assumption, $[\![\phi]\!]_{(X,f)} - \{x\} \neq \emptyset$ whenever $f(x)([\![\phi]\!]_{(X,f)}) = *$. ☐

To avoid exponential explosion, we represent elements of $N_3(X)$, for X finite, using partial maps $f_0 : \mathcal{Q}(X) \rightharpoonup 3$. To enforce (E2′), we let such an f_0 represent the map $f : \mathcal{Q}(X) \to 3$ that maps $B \subseteq X$ to the maximum of $\bigwedge_{i=1}^{n} f_0(A_i)$, taken over all sets $A_1, \dots, A_n \subseteq X$ such that $\bigcap A_i = B$ and $f_0(A_i)$ is defined for all i; when no such sets exist, the maximum is understood to be $\bot$.

Lemma 4.2. *Let f_0 and f be as above.*

1. *Whenever $f_0(A)$ is defined, then $f_0(A) \leq f(A)$.*
2. *Let $b \in 3$. Then $f(A) \geq b$ iff $\bigcap\{B \subseteq X \mid A \subseteq B, f_0(B) \geq b \text{ defined}\} = A$.*
3. *The pair (f, x) satisfies (E1′) iff $f_0(X)$ is either undefined or equals $\bot$.*
4. *The pair (f, x) satisfies (E2′).*
5. *The pair (f, x) satisfies (E3a′) iff $\bigcap\{A \subseteq X \mid f_0(A) > \bot \text{ defined}\} \neq \emptyset$.*
6. *The pair (f, x) satisfies (E3b′) iff whenever $f_0(A) = \top$ is defined, then $x \in A$.*

Proof. 1.: Trivial.
2.: 'If' is trivial. 'Only if': by assumption, $A = \bigcap_{i=1}^{n} B_i$ for some B_i such that $f_0(B_i) \geq b$ is defined for all i; the claim follows immediately.

3.: 'Only if' is immediate by 1., and 'if' holds because $X = \bigcap A_i$ implies $A_i = X$ for all i.
4.: By construction.
5.: Immediate by 2.
6.: 'Only if' holds by 1., and 'if' holds because $f(B) = \top$ implies that $B = \bigcap A_i$ for sets A_i such that $f_0(A_i) = \top$ for all i. □

By Lemma 4.2, it is immediate that one-step model checking is in P. To prove the OSPMP, let $(X, \tau, f : \mathcal{Q}(X) \to 3, x)$ be a one-step model of a one-step pair (ϕ, ψ) over V. By Remark 3.3, we can assume that X is finite. Let the set $Y \subseteq X$ consist of

- the element x;
- an element $y_{ab} \in \tau(a) \setminus \tau(b)$ for each pair $(a, b) \in V^2$ such that $\tau(a) \not\subseteq \tau(b)$;
- an element $z_a \in \bigcap\{\tau(b) \mid b \in V, \tau(a) \subseteq \tau(b),\, f(\tau(b)) > f(\tau(a))\} \setminus \tau(a)$ for each $a \in V$ (z_a exists by (E2′)); and
- an element $w_0 \in \bigcap\{\tau(b) \mid f(\tau(b)) > \bot\}$ (w_0 exists by (E2′) and (E3a′)).

Put $\tau_Y(a) = \tau(a) \cap Y$ for $a \in V$, and let f_Y be represented by the partial map f_0 taking $\tau_Y(a)$ to $f(\tau(a))$ (f_0 is well-defined by construction of Y). Then Y and (f_Y, x) are of polynomial size in ψ, and $(Y, \tau_Y) \models \phi$. By Lemma 4.2, (f_Y, x) is in $\mathcal{A}(X)$, with the criterion for (E3a′) satisfied due to $w_0 \in Y$. By Lemma 4.2.2, the $z_a \in Y$ ensure that $f_Y(\tau_Y(a)) = f(\tau(a))$ for all $a \in V$, so that $f_Y \models^1_{(Y,\tau_Y)} \psi$.

By Corollary 3.10, we obtain that *the modal logic of agency is in PSPACE*, and that bounding the modal nesting depth brings the complexity down to NP. Both results (and even decidability) seem to be new. In the light of the previous observation that the agglomeration axiom $Ea \wedge Eb \to E(a \wedge b)$ tends to cause PSPACE-hardness [31], we conjecture that the PSPACE upper bound is tight.

5 Exponential Branching

In cases where the OSPMP fails, it may still be possible to obtain a PSPACE upper bound by traversing an exponentially branching shallow model (by Remark 3.3, branching is never worse than exponential). The crucial prerequisite is that exponential-size one-step models can be traversed pointwise, accumulating during the traversal a polynomial amount of information that suffices for one-step model checking. This requires additional assumptions on the signature functor S_0:

Definition 5.1. We say that S_0 is *pointwise bounded* w.r.t. a set C if for all sets X, there exists an injection $S_0X \hookrightarrow (X \to C)$ (i.e. $|S_0X| \leq |X \to C|$). We then identify S_0X with a subset of $X \to C$.

(Note that the above does *not* require that $\lambda X.\, X \to C$ is functorial.) Recall from [27] that the signature functor S_0 admits a *separating* set of unary predicate liftings (separation is a necessary condition for the generalised Hennessy-Milner property) iff the family of maps $S_0 f : S_0X \to S_0 2$, indexed over all maps $f : X \to 2 = \{\bot, \top\}$, is jointly injective for each set X. Typically, functors S_0 satisfying this condition satisfy the stronger requirement that already the family of maps $(S_0 \mathbb{1}_{\{x\}} : S_0X \to S_0 2)_{x \in X}$ is jointly injective, where $\mathbb{1}_A$ denotes the characteristic function of $A \subseteq X$, so that S_0 is pointwise bounded w.r.t. $S_0 2$; often, even a quotient of $S_0 2$ will suffice. Of the signature functors mentioned in Example 2.10, $\mathcal{P}$ and D_ω are pointwise bounded (w.r.t. 2 and $[0, 1]$, respectively), while *Cf* and $\mathcal{P} \circ D_\omega$ fail to be so.

Further examples of pointwise bounded functors include the game frame functor appearing in the semantics of coalition logic [29] and the multiset functor introduced below.

Assume from now on that S_0 is pointwise bounded w.r.t. C, with a given representation of elements of C (Convention 2.11 is no longer needed). For $t : X \to C$, we define

$$\mathrm{maxsize}(t) = \max_{x \in X} \mathrm{size}(t(x)),$$

and put $\mathrm{maxsize}(t, x) = \mathrm{maxsize}(t)$ for $x \in X$.

Definition 5.2. We say that $\mathcal{M}$ has the *one-step pointwise polysize model property (OSPPMP)* if there exists a polynomial p such that, whenever a one-step pair (ϕ, ψ) over V has a one-step model (X, τ, t, x), then it has a one-step model (Y, κ, s, y) such that $|Y| \leq 2^{|V|}$, $\mathrm{maxsize}(s) \leq p(|\psi|)$, and $y \in \kappa(a)$ iff $x \in \tau(a)$ for all $a \in V$; such a model is called *pointwise polysize.*

By Remark 3.3, the actual content of the OSPPMP is the polynomial bound on $\mathrm{maxsize}(s)$. The OSPPMP holds for all pointwise bounded functors mentioned so far, trivially so in cases where C is finite. We have a variant of Theorem 3.5, proved entirely analogously, which states that *under the OSPPMP, every satisfiable formula ψ is satisfied in a shallow model (X, ξ) with branching degree at most $2^{|\psi|}$ and $\mathrm{maxsize}(\xi(x))$ polynomially bounded in $|\psi|$*. For the ensuing algorithmic treatment, we need a refined notion of one-step model checking:

Definition 5.3. The *pointwise one-step model checking problem* is to check, given a map $t : X \to C$, $x \in X$, a $\mathcal{P}(X)$-valuation τ for V, $Y \subseteq X$, and a conjunctive clause ψ over $\Lambda(V)$, whether $(t, x) \in SY \subseteq (X \to C) \times X$ and $t \models^1_{Y,\tau_Y} \psi$, where $\tau_Y(a) = \tau(a) \cap Y$ for $a \in V$. We say that this problem is *PSPACE-tractable* if it is decidable on a non-deterministic Turing machine with input tape that uses space polynomial in $\mathrm{maxsize}(t)$ and accesses each input symbol at most once.

Theorem 5.4. *If $\mathcal{M}$ has the OSPPMP and pointwise one-step model checking is PSPACE-tractable, then the satisfiability problem of $\mathcal{F}(\Lambda)$ is in PSPACE.*

Proof. Let M be a decision procedure for pointwise one-step model checking as required in the definition of PSPACE-tractability (Defn. 5.3). Let p be a polynomial witnessing the OSPPMP as in Definition 5.2. Then the following non-deterministic algorithm decides satisfiability of $\mathcal{F}(\Lambda)$-formulas:

Algorithm 5.5. 1. If $\mathrm{rank}(\psi) = 0$, terminate successfully if ψ evaluates to $\top$, else unsuccessfully. Otherwise:

2. Take V and σ as in the proof of Theorem 3.5, and guess a conjunctive clause $\bar{\psi}$ over $\Lambda(V)$ containing for each subformula $L(\rho_1, \dots, \rho_n)$ of ψ either $L(a_{\rho_1}, \dots, a_{\rho_n})$ or $\neg L(a_{\rho_1}, \dots, a_{\rho_n})$ such that $\bar{\psi}\sigma$ propositionally entails ψ.

3. Call M with arguments X, τ, Y, t as in Definition 5.3 to check that $t \in TY$ and $t \models^1_{(Y,\tau_Y)} \bar{\psi}$, with τ_Y as in Definition 5.3. Here, $X = 2^V$, $\tau(a) = \{B \in X \mid a \in B\}$,

$$Y = \{B \in X \mid \bigwedge_{a_\rho \in B} \rho \wedge \bigwedge_{a_\rho \notin B} \neg\rho \text{ satisfiable}\}$$

is calculated recursively, and $t \in (X \to C)$ with $\mathrm{maxsize}(t) \leq p(|\phi|)$ is guessed.

It remains to see that the above algorithm can be implemented in polynomial space although the input to M in Step 3 is of overall exponential size. This is achieved by replacing read operations on the input tape in M by calls to a procedure passed by the caller, which produces the k-th input symbol on demand, and then calling the modified pointwise model checker M' with such a procedure instead of the full argument. By the assumption that M accesses each symbol on the input tape at most once, there is no need to keep the symbols representing the guessed value t in memory after they have been passed to M'. Therefore, only polynomial space overhead is generated by the input to M' (the input procedure depends on ϕ and hence has representation size $O(|\phi|)$); the space usage of M' itself is polynomial in $\mathrm{maxsize}(t)$ and therefore in $|\phi|$. □

Example 5.6. Theorem 5.4 applies e.g. to the modal logics K and T, as well as to probabilistic modal logic; however, as all these logics in fact enjoy the OSPMP, the method of Sect. 3 is preferable in these cases. A more interesting application is given by graded modal logic [11], or more generally Presburger modal logic [6].

In its single-agent version, Presburger modal logic has n-ary modal operators $\sum_{i=1}^{n} a_i\#(\_) \sim b$, where b and the a_i are integers and $\sim \in \{<, >, =\} \cup \{\equiv_k \mid k \in \mathbb{N}\}$. A coalgebraic semantics for this logic, equivalent for purposes of satisfiability to the ordered tree semantics given in [6], is defined over the finite multiset functor $\mathcal{B}$, which maps a set X to the set of maps $B : X \to \mathbb{N}$ with finite support, the intuition being that B is a multiset containing $x \in X$ with multiplicity $B(x)$. For $A \subseteq X$, put $B(A) = \sum_{x \in A} B(x)$. $\mathcal{B}$-coalgebras are graphs with $\mathbb{N}$-weighted edges. The above modalities are interpreted by

$$[\![\textstyle\sum_{i=1}^{n} a_i\#(\_) \sim b]\!]_X(A_1, \dots, A_n) = \{B \in \mathcal{B}(X) \mid \textstyle\sum_{i=1}^{n} a_i B(A_i) \sim b\},$$

with $>, <, =$ interpreted as expected, and $\equiv_k$ as equality modulo k. This logic extends graded modal logic, whose operators $\Diamond_k$ now become $\#(\_) > k$.

Of course, $\mathcal{B}$ is pointwise bounded w.r.t. $\mathbb{N}$. It follows easily from estimates on solution sizes of integer linear equalities [24] that Presburger modal logic has the OSPPMP [6]. Moreover, given a conjunctive clause ψ over $\Lambda(V)$, a $\mathcal{P}(X)$-valuation τ, and $B \in \mathcal{B}(X)$, one can check whether $B \models^1_{X,\tau} \psi$ by traversing X and computing the $B(\tau(a))$ by successive summation; it is thus easy to see that pointwise one-step model checking is PSPACE-tractable. It follows that *Presburger modal logic is in PSPACE*. While this is proved already in [6], using essentially the same type of algorithm[1], our method extends straightforwardly to extensions of Presburger modal logic by certain frame conditions such as reflexivity (modelled by the copointed functor $SX = \{(B, x) \in \mathcal{B}X \times X \mid B(x) > 0\}$) or e.g. the condition that at least half of all transitions from a given state are loops (modelled by the copointed functor $SX = \{(B, x) \in \mathcal{B}X \times X \mid B(x) \geq B(X - \{x\})\}$). In particular, this implies that *graded modal logic over reflexive frames* (i.e. the logic Tn of [11]) *is in PSPACE*, to our knowledge a new result. The logic Tn can be seen as a description logic with qualified number restrictions on a single reflexive role. Our arguments extend straightforwardly to show that a description logic with role hierarchies, reflexive roles, and qualified number restrictions has concept satisfiability over the empty T-box in PSPACE. As reflexivity of a role R is expressed by the role inclusion $id(\top) \subseteq R$, where $id(\top)$ denotes the identity role, this logic is a fragment of $\mathcal{ALCHQ}(id)$ [1].

[1] The claim that a (rank-1) logic further extended by regularity constraints is still in PSPACE is retracted in the full version of [6] as being based on possibly erroneous third-party results.

6 Conclusion

We have formulated two local semantic conditions that guarantee PSPACE upper bounds for the satisfiability problem of modal logics in a coalgebraic framework: the OSPMP (one-step polysize model property) and its pointwise variant, the OSPPMP, which is weaker but relies on additional assumptions on the coalgebraic semantics. Both conditions allow a direct construction of shallow models and their traversal in polynomial space. This complements earlier work [29] where syntactic criteria have been used — in particular, both the OSPMP and the OSPPMP can be applied even when no complete axiomatisation of the logic at hand is known. Several instantiations of our results to logics studied in the literature witness both their generality and their usefulness: Apart from re-proving known PSPACE upper bounds for the normal modal logics K and T as well as for the conditional logics CK, $CK+ID$, and $CK+MP$, we have

- given a systematic account of tight PSPACE upper bounds in modal logics of quantitative uncertainty that establishes new complexity bounds in some cases;
- obtained a new PSPACE upper bound for Elgesem's modal logic of agency and for graded (even Presburger [6]) modal logic over reflexive frames [11], and more generally for an extension of the description logic $\mathcal{ALCHQ}$ with reflexive roles [1];
- established (to our knowledge: new) tight NP upper bounds for bounded-rank fragments of the conditional logics CK, $CK+ID$, and $CK+MP$.

Ongoing work focusses on the extension of our results to *iterative* modal logics, defined by frame conditions of higher rank, which however — in particular outside the realm of Kripke semantics — exhibit a tendency towards higher complexity or even undecidability (indeed, it seems to be the case that all known iterative PSPACE-complete modal logics are normal).

References

[1] F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook*. Cambridge University Press, 2003.

[2] M. Barr. Terminal coalgebras in well-founded set theory. *Theoret. Comput. Sci.*, 114:299–315, 1993.

[3] B. Chellas. *Modal Logic*. Cambridge Univ. Press, 1980.

[4] L. Cholvy, C. Garion, and C. Saurel. Ability in a multi-agent context: A model in the situation calculus. In *Computational Logic in Multi-Agent Systems, CLIMA 2005*, volume 3900 of *LNCS*, pages 23–36. Springer, 2006.

[5] C. Cîrstea and D. Pattinson. Modular construction of modal logics. In *Concurrency Theory*, volume 3170 of *LNCS*, pages 258–275. Springer, 2004.

[6] S. Demri and D. Lugiez. Presburger modal logic is only PSPACE-complete. In *Automated Reasoning, IJCAR 06*, volume 4130 of *LNAI*, pages 541–556. Springer, 2006.

[7] D. Elgesem. The modal logic of agency. *Nordic J. Philos. Logic*, 2:1–46, 1997.

[8] R. Fagin and J. Y. Halpern. Reasoning about knowledge and probability. *J. ACM*, 41:340–367, 1994.

[9] R. Fagin, J. Y. Halpern, and N. Megiddo. A logic for reasoning about probabilities. *Inform. Comput.*, 87:78–128, 1990.

[10] M. Fattorosi-Barnaba and F. De Caro. Graded modalities I. *Stud. Log.*, 44:197–221, 1985.

[11] K. Fine. In so many possible worlds. *Notre Dame J. Formal Logic*, 13:516–520, 1972.

[12] G. Governatori and A. Rotolo. On the axiomatisation of Elgesem's logic of agency and ability. *J. Philos. Logic*, 34:403–431, 2005.

[13] J. Halpern. The effect of bounding the number of primitive propositions and the depth of nesting on the complexity of modal logic. *Artificial Intelligence*, 75:361–372, 1995.

[14] J. Halpern and R. Pucella. A logic for reasoning about upper probabilities. *J. Artificial Intelligence Res.*, 17:57–81, 2002.

[15] J. Halpern and R. Pucella. Reasoning about expectation. In *Uncertainty in Artificial Intelligence, UAI 02*, pages 207–215. Morgan Kaufman, 2002.

[16] J. Y. Halpern and Y. O. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54:319–379, 1992.

[17] B. Jacobs. Towards a duality result in the modal logic of coalgebras. In *Coalgebraic Methods in Computer Science, CMCS 2000*, volume 33 of *ENTCS*. Elsevier, 2000.

[18] A. Jones and X. Parent. Conventional signalling acts and conversation. In *Advances in Agent Communication*, volume 2922 of *LNAI*, pages 1–17. Springer, 2004.

[19] A. Kurz. Specifying coalgebras with modal logic. *Theoret. Comput. Sci.*, 260:119–138, 2001.

[20] R. Ladner. The computational complexity of provability in systems of modal propositional logic. *SIAM J. Comput.*, 6:467–480, 1977.

[21] D. Lewis. Intensional logics without iterative axioms. *J. Philos. Logic*, 3:457–466, 1975.

[22] N. Olivetti, G. L. Pozzato, and C. Schwind. A sequent calculus and a theorem prover for standard conditional logics. *ACM Trans. Comput. Logic.* To appear.

[23] E. Pacuit and S. Salame. Majority logic. In *Principles of Knowledge Representation and Reasoning, KR 04*, pages 598–605. AAAI Press, 2004.

[24] C. H. Papadimitriou. On the complexity of integer programming. *J. ACM*, 28:765–768, 1981.

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

[26] M. Pauly. A modal logic for coalitional power in games. *J. Logic Comput.*, 12:149–166, 2002.

[27] L. Schröder. Expressivity of coalgebraic modal logic: the limits and beyond. *Theoret. Comput. Sci.* In press.

[28] L. Schröder. A finite model construction for coalgebraic modal logic. *J. Log. Algebr. Prog.*, 73:97–110, 2007.

[29] L. Schröder and D. Pattinson. PSPACE reasoning for rank-1 modal logics. In *Logic in Computer Science, LICS 06*, pages 231–240. IEEE, 2006. Extended version to appear in *ACM Trans. Comput. Log.*

[30] S. Tobies. PSPACE reasoning for graded modal logics. *J. Logic Comput.*, 11:85–106, 2001.

[31] M. Vardi. On the complexity of epistemic reasoning. In *Logic in Computer Science, LICS 89*, pages 243–251. IEEE, 1989.