A Finite Model Construction For Coalgebraic Modal Logic

Lutz Schröder

DFKI-Lab Bremen and
Department of Computer Science, University of Bremen, Germany

Abstract

In recent years, a tight connection has emerged between modal logic on the one hand and coalgebras, understood as generic transition systems, on the other hand. Here, we prove that (finitary) coalgebraic modal logic has the finite model property. This fact not only reproves known completeness results for coalgebraic modal logic, which we push further by establishing that every coalgebraic modal logic admits a complete axiomatisation in rank 1; it also enables us to establish a generic decidability result and a first complexity bound. Examples covered by these general results include, besides standard Hennessy-Milner logic, graded modal logic and probabilistic modal logic.

Key words: modal logic, coalgebra, weak completeness, finite models, decidability

Introduction

Coalgebra has recently had increasing success as a generic theory of reactive systems, providing a unifying perspective on a wide variety of system types [24]. Many concepts of concurrency theory can be cast in the coalgebraic framework; these include general notions of bisimulation, coinduction, and corecursion, as well as generic modal logics. The latter include the seminal *coalgebraic logic* [16] and the more recent *coalgebraic modal logic* [12,22,14,19,21]; here, we focus on coalgebraic modal logic, as it stays close to the traditional syntax and semantics of modal logic. The role of coalgebraic modal logic is twofold: on the one hand, one obtains a suitable generic reactive specification language, which respects encapsulation of the state space, i.e. relates well to behavioural equivalence of states [21,26], and is sufficiently intuitive for use in actual software specification languages, including object-oriented specification [23,14,17]. On the other hand, coalgebraic modal logics frequently correspond to known modal logics such as graded modal logic or probabilistic modal logic, and thus provide these logics with a coalgebraic semantics.

In [20] and subsequent work [5,13], a weak completeness result for coalgebraic modal logic has been established stating that a deductive system consisting of propositional entailment, a congruence rule, and a given axiomatisation in rank 1 (i.e. with nesting depth of modal operators uniformly equal to 1) is weakly complete, provided that the axiomatisation is *one-step complete*. (Strong completeness cannot be expected in general, as many coalgebraic modal logics fail to be compact.) Here, we exhibit a finite model construction which relies on one-step completeness. We thus reprove the mentioned weak completeness result. Moreover, we show that every coalgebraic modal logic has a one-step complete axiomatisation in rank 1, which then implies that coalgebraic modal logic has the *finite model property*, i.e. every satisfiable formula is satisfiable in a finite model. We further exploit the finite model construction to obtain a generic decision procedure which reduces the satisfiability problem for a coalgebraic modal logic to the much simpler *one-step satisfiability* problem. This yields not only decidability of a large number of modal logics, including the above-mentioned graded and probabilistic modal logics, but also, under mild conditions, a first upper complexity bound.

The material is organised as follows. Section 1 gives an introduction to coalgebra and coalgebraic modal logic, including a number of examples. In Section 2, we recall the deduction system of coalgebraic modal logic [20,5,13] and the above-mentioned notion of one-step completeness, and prove that one-step complete axiomatisations always exist. We then prove the finite model property in Section 3, from which we obtain the generic decision procedure and the arising upper complexity bound in Section 4. This work is an extended version of [27].

1 Coalgebraic Modal Logic

We briefly recapitulate the basics of the coalgebraic modelling of reactive systems and of the specification of such systems by means of coalgebraic modal logic.

Definition 1 Let $T : \mathbf{Set} \to \mathbf{Set}$ be a functor (in this work, all functors will implicitly be set functors), referred to as the *signature functor*. A T-*coalgebra* $A = (X, \xi)$ consists of a set X of *states* and a *transition map* $\xi : X \to TX$.

Intuitively, the transition map describes the successor states and observations of a state, organised in a data structure given by T.

We explicitly fix some logical terminology:

Definition 2 Let T be a functor. A *language for T-coalgebras* is a set $\mathcal{L}$ of *formulas*, equipped with a family of *satisfaction* relations $\models_C$ (or just $\models$) between states of T-coalgebras $C = (X, \xi)$ and formulas $\phi \in \mathcal{L}$; we define $[\![\phi]\!]_C$ (or just $[\![\phi]\!]$) as the set $\{x \in X \mid x \models_C \phi\}$. For $\Phi \subseteq \mathcal{L}$, we write $x \models \Phi$ if $x \models \phi$ for all $\phi \in \Phi$. We say that ψ is a *local consequence* of Φ if, for every state x in every T-coalgebra, $x \models \psi$ whenever $x \models \Phi$. We say that ψ is *valid* if $\emptyset \models \psi$. A set Φ of formulas is *(locally) satisfiable* if it is satisfied in some state in some T-coalgebra.

Note that a formula ϕ is valid iff $\neg\phi$ is locally unsatisfiable.

As a specification logic for coalgebraically modelled reactive systems, *coalgebraic modal logic* in the form considered here has been introduced in [21], generalising previous frameworks [12,22,14,19]. The semantics is based on the following central notion.

Definition 3 A *predicate lifting* for a functor T is a natural transformation

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

where $\mathcal{Q}$ denotes the contravariant powerset functor $\mathbf{Set}^{op} \to \mathbf{Set}$, with $\mathcal{Q}f(A) = f^{-1}[A]$. Explicitly, a predicate lifting assigns to each $A \subseteq X$ a set $\lambda_X(A) \subseteq TX$ such that

$$Tf^{-1}[\lambda_Y(A)] = \lambda_X(f^{-1}[A])$$

for all maps $f : X \to Y$.

In the terminology introduced above, a (finitary) coalgebraic modal logic is a language $\mathcal{L}(\Lambda)$ for T-coalgebras, determined by a set Λ of predicate liftings for T. Formulas $\phi, \psi \in \mathcal{L}(\Lambda)$ are defined by the grammar

$$\phi ::= \bot \mid \phi \wedge \psi \mid \neg\phi \mid [\lambda]\,\phi,$$

where λ ranges over Λ. Disjunction $\phi \vee \psi$, truth $\top$, implication $\phi \to \psi$, and biimplication $\phi \leftrightarrow \psi$ are then defined as usual. In the definition of satisfaction, the clauses for the boolean operators $\bot$, $\wedge$, and $\neg$ are as expected; the clause for the modal operator $[\lambda]$ is

$$x \models_{(X,\xi)} [\lambda]\,\phi \iff \xi(x) \in \lambda_X [\![\phi]\!]_{(X,\xi)}.$$

The *size* $|\phi|$ of a formula ϕ is the number of subformulas of ϕ (not to be confused with the *representation size* of ϕ introduced below).

Remark 4 A more general form of coalgebraic modal logic interprets polyadic modal operators by polyadic predicate liftings [26]. The results of this paper extend straightforwardly to polyadic modal logic, essentially by just replacing single formulas and sets by indexed families where appropriate; we restrict the exposition to the unary case purely in the interest of readability.

Example 5 [21,5,26]

(1) Let $\mathcal{P}$ be the covariant powerset functor. Then $\mathcal{P}$-coalgebras are graphs, thought of as transition systems or indeed Kripke frames. We define a predicate lifting $\lambda^\forall$ by

$$\lambda^\forall_X(A) = \{B \subseteq X \mid B \subseteq A\}.$$

We thus obtain the standard box modality $\Box = [\lambda^\forall]$. This setup is easily adapted to transition systems with branching degree limited by a regular cardinal κ, described as coalgebras for the functor $\mathcal{P}_\kappa$ defined by $\mathcal{P}_\kappa(X) = \{A \subseteq X \mid |A| < \kappa\}$.

(2) It is straightforward to extend a given coalgebraic modal logic for T with a set V of propositional symbols. This amounts to considering the functor $T \times \mathcal{P}(V)$, where $\mathcal{P}(V)$ stands for the corresponding constant functor. We then have predicate liftings λ^a, $a \in V$, defined by

$$\lambda^a_X(A) = \{(t, B) \in TX \times \mathcal{P}(V) \mid a \in B\}.$$

Since λ^a is independent of its argument, the induced modal 'operator' can be written as just the propositional symbol a, with the expected meaning. (Of course, propositional symbols are more naturally modelled as nullary predicate liftings in a framework with predicate liftings of arbitrary finite arities as discussed in Remark 4.)

(3) The *finite multiset* (or *bag*) functor $\mathcal{B}_\mathbb{N}$ is given as follows. The set $\mathcal{B}_\mathbb{N}(X)$ consists of the maps $B : X \to \mathbb{N}$ with finite support; we say that B contains $x \in X$ with *multiplicity* $B(x)$. We write multisets additively, denoting by $\sum n_i x_i$ the multiset that contains x with multiplicity $\sum_{x_j = x} n_j$. For $f : X \to Y$, $\mathcal{B}_\mathbb{N}(f)(\sum n_i x_i) = \sum n_i f(x_i)$. Coalgebras for $\mathcal{B}_\mathbb{N}$ are directed graphs with $\mathbb{N}$-weighted edges, often referred to as *multigraphs* [6].

One has predicate liftings λ^k, $k \in \mathbb{N}$, defined by

$$\lambda^k_X(A) = \{\textstyle\sum n_i x_i \in \mathcal{B}_\mathbb{N} X \mid \sum_{x_i \in A} n_i > k\}.$$

The arising modal operators are exactly the modalities $\Diamond_k$ of *graded modal logic* (cf. e.g. [6]), i.e. $x \models \Diamond_k \phi$ iff ϕ holds for more than k successor states of x, taking into account multiplicities. (The semantics of graded modal logic has originally been defined over Kripke frames; this is however equivalent for purposes of satisfiability, cf. Remark 6 below.) Note that $\Box_k$, defined as $\neg\Diamond_k\neg$, is monotone, but fails to be normal unless $k = 0$. (Recall that a modal operator $\Box$ is called *monotone* if it satisfies $\Box(p \wedge q) \to \Box p$, and *normal* if it satisfies $\Box(p \to q) \to \Box p \to \Box q$.)

(4) A similar functor, denoted $\mathcal{B}_{\mathbb{Z}}$, is given by a slight modification of the multiset functor where we allow elements to have also *negative* multiplicities, i.e. $\mathcal{B}_{\mathbb{Z}}X$ consists of finitely supported maps $X \to \mathbb{Z}$, called *generalised multisets* (this set is also familiar as the free abelian group over X). Coalgebras for $\mathcal{B}_{\mathbb{Z}}$ appear in the literature as *integer weighted automata* [9].

   One has predicate liftings λ^k, $k \in \mathbb{Z}$, with induced modal operators $\Diamond_k$ defined analogously as for multisets. We refer to the arising logic as *generalised graded modal logic*. Note that in this case, $\Diamond_k$ fails to be monotone, even for $k = 0$.

(5) The *finite distribution functor* D_ω maps a set X to the set of probability distributions on X with finite support. Coalgebras for the functor $T = D_\omega \times \mathcal{P}(V)$, where $V \neq \emptyset$ is a set of propositional symbols, are probabilistic transition systems (also called *probabilistic type spaces* [10]) with finite branching degree. (The example is easily extended to countable branching by considering instead the functor D_{ω_1} of probability distributions with countable support, while higher branching degrees as admitted by the general definition of probabilistic type space require a more elaborate measure theoretic treatment [36].)

   Besides the propositional symbols (cf. Example 2), we have predicate liftings λ^p, $p \in [0,1] \cap \mathbb{Q}$, defined by

$$\lambda^p(A) = \{P \in D_\omega X \mid PA \geq p\}.$$

   These induce the modal operators $L_p = [\lambda^p]$ of *probabilistic modal logic* [15,10], where $L_p\,\phi$ reads 'ϕ holds in the next step with probability at least p'.

(6) For a field k, the *linear space functor* $k \cdot \_$ takes a set X to the free k-vector space over X, i.e. the set $k \cdot X$ of formal k-linear combinations over X. A coalgebra for $k \cdot \_$ is a *linear automaton* [3,34] (where one would in general also assume linear output in a vector space V, corresponding to the functor $(k \cdot \_) \times V$). In the case $k = \mathbb{R}$, predicate liftings λ^p, $p \in \mathbb{Q}$, may be defined analogously as for D_ω, giving rise to modal operators L_p for $p \in \mathbb{Q}$. Here, $L_p\phi$ holds if the sum of the coefficients of successor states satisfying ϕ is at least p.

(7) The above examples may be extended by adding labels from an alphabet L, i.e. by passing from T to one of the functors S and R given by $SX = L \to TX$ and $RX = T(L \times X)$, respectively. When L is finite, these functors are isomorphic for $T \in \{\mathcal{P}_\omega, \mathcal{B}_{\mathbb{N}}, \mathcal{B}_{\mathbb{Z}}\}$ but not for $T = D_\omega$. In the latter case, S-coalgebras are reactive probabilistic automata, and R-coalgebras are generative probabilistic automata [1] (more precisely, one allows for terminal states by additionally introducing the constant functor 1 as a summand).

   A natural set of modal operators is then obtained by additionally indexing modal operators over $a \in L$. In the case $T = \mathcal{P}_\omega$, this leads to the usual operators $[a]$ of Hennessy-Milner Logic [11]. In the probabilistic case, the meaning of $L_p^a\,\phi$ in reactive probabilistic automata is that ϕ holds with probability at least p after an a-transition takes place, and in generative probabilistic automata that with probability at least p, the next step is an a-transition leading into a state satisfying ϕ.

Remark 6 Graded modal logic is more standardly interpreted over Kripke frames by just counting successor states (as e.g. in [33]), rather than in multigraphs as in the above example and e.g. in [6]. However, the two semantics induce the same local consequence relations; this is seen as follows. On the one hand, one can regard Kripke frames as multigraphs by just regarding sets as multisets where all elements have multiplicity 1, and this identification is compatible with the satisfaction of graded modal formulas. Conversely, one turn a multigraph into a semantically equivalent Kripke frame by making copies of elements according to their multiplicity; explicitly: Let X be a multigraph. Construct a Kripke frame $\bar{X}$ with transition relation R by taking as states all pairs $(y, j) \in X \times \mathbb{N}$ for which there exists x such that y is a successor of x with multiplicity $n > j$ in X, and in this case put $(x, i)R(y, j)$ for all i such that (x, i) is a state in $\bar{X}$. By induction over graded modal formulas ϕ, one shows easily that $x \models \phi$ in X iff $(x, i) \models \phi$ in $\bar{X}$. Note that $\bar{X}$ is finite if X is finite.

2 Proof Systems For Coalgebraic Modal Logic

We now discuss completeness of derivation for coalgebraic modal logic, partly following [20,5,13]. Since a signature functor T contains information only about the one-step evolution of the system (as opposed to a comonad, which may contain information also about further steps), it is natural to expect that for the axiomatisation of a coalgebraic modal logic for T it is enough to consider modal axioms of rank 1. The approach taken in [20,5,13] is based on this expectation; we shall prove below that it is indeed formally the case that axioms, or altenatively rules, of rank 1 are sufficient. This fact will be crucial for our finite model result to be proved in Section 3.

To begin, we note that the local consequence relation (Definition 2) of a coalgebraic modal logic in general fails to be compact:

Example 7 In the case of Hennessy-Milner logic over finitely branching systems [11] with two labels a, b, the set

$$\Phi = \{\langle a\rangle([b]^{n+1}\bot \wedge \langle b\rangle^{n}\top) \mid n \in \mathbb{N}\},$$

where $[b]^n$ and $\langle b\rangle^n$ stand for n consecutive boxes or diamonds, respectively, is locally unsatisfiable, since it requires, for each n, the existence of an a-successor from which exactly n b-steps are possible. However, every finite subset of Φ is locally satisfiable. For an example of the same kind, but of bounded rank, consider the set

$$\{\Diamond_k \top \mid k \in \mathbb{N}\}$$

of graded modal formulas over $\mathcal{B}_{\mathbb{N}}$, which requires that the size of the multiset of successors exceeds every $k \in \mathbb{N}$. Non-compactness of probabilistic modal logic is observed in [10]; in this case, non-compactness does *not* have to do with bounded branching.

Remark 8 The same examples as above show that also the *global* consequence relations of the mentioned logics fail to be compact. A formula ψ is a *global consequence* of a set Φ of formulas if for every T-coalgebra C, $C \models \psi$ whenever $C \models \Phi$, where we write $C \models \psi$ if $x \models \psi$ for all states x of C.

Thus, the local consequence relation of a coalgebraic modal logic in general fails to admit a finitary strongly complete proof system, where we call a proof system with induced entailment relation $\vdash$ *strongly complete* w.r.t. a consequence relation $\models$ if, for every set Φ of formulas and every formula ψ, $\Phi \vdash \psi$ whenever $\Phi \models \psi$. Instead, one is lead to study weak completeness, where a proof system is called *weakly complete* if it proves all valid formulas (i.e. $\emptyset \vdash \psi$ whenever $\emptyset \models \psi$). This notion is equivalent to completeness in the sense used in [5,13,20], where only local consequence with singleton sets of premises is considered (ψ is a local consequence of $\{\phi\}$ iff $\phi \to \psi$ is valid).

For the remainder of the paper, we assume given a functor T and a set Λ of predicate liftings for T. We recall a few basic notions from propositional logic, as well as notation for coalgebraic modal logic introduced in [20,5]:

Definition 9 We denote the set of propositional formulas over a set V, generated by the basic connectives $\neg$ and $\wedge$, by $\mathsf{Prop}(V)$. We use variables ϵ etc. to denote either nothing or $\neg$. Thus, a *literal* over V is a formula of the form ϵa, with $a \in V$. A *(conjunctive) clause* is a finite, possibly empty, disjunction (conjunction) of literals. Moreover, we denote by $\mathsf{Up}(V)$ the set $\{[\lambda]a \mid \lambda \in \Lambda, a \in V\}$. If $V \subseteq \mathcal{L}(\Lambda)$, we also regard formulas over V as formulas in $\mathcal{L}(\Lambda)$, i.e. we assume $\mathsf{Prop}(\mathcal{L}(\Lambda)) \subseteq \mathcal{L}(\Lambda)$ and $\mathsf{Up}(\mathcal{L}(\Lambda)) \subseteq \mathcal{L}(\Lambda)$.

We sometimes explicitly designate V as consisting of *propositional variables*; these retain their status across further applications of Up and Prop (e.g. V is also the set of propositional variables for $\mathsf{Up}(\mathsf{Prop}(V))$). An *$L$-substitution* is a substitution σ of the propositional variables by elements of a set L; for a formula ϕ over V, we call $\phi\sigma$ an *L-instance* of ϕ. If $L = \mathcal{P}(X)$ for some X, then we also refer to σ as a *$\mathcal{P}(X)$-valuation.* If $L = \{\top, \bot\}$, then we refer to σ as a *valuation*, and denote by $\phi\sigma$ the truth value of ϕ under σ.

Given a set X and a $\mathcal{P}(X)$-valuation τ, we define interpretations $[\![\phi]\!]\tau \subseteq X$

and $[\![\psi]\!]\tau \subseteq TX$ for $\phi \in \mathsf{Prop}(V)$ and $\psi \in \mathsf{Prop}(\mathsf{Up}(\mathsf{Prop}(V)))$, respectively, by the usual clauses for boolean operators and by $[\![[\lambda]\phi]\!]\tau = \lambda[\![\phi]\!]\tau$. We write $X, \tau \models \phi$ if $[\![\phi]\!]\tau = X$, and $TX, \tau \models \psi$ if $[\![\psi]\!]\tau = TX$. We say that ψ is *(one-step) satisfiable over* τ if $[\![\psi]\!]\tau \neq \emptyset$.

If $\Phi \subseteq \mathsf{Prop}(V)$, $\psi \in \mathsf{Prop}(V)$, $\phi_1 \wedge \ldots \wedge \phi_n \to \psi$ is a propositional tautology for some $\phi_1, \ldots, \phi_n \in \Phi$, and σ is an L-substitution, then we say that $\Phi\sigma$ (i.e. the set $\{\phi\sigma \mid \phi \in \Phi\}$) *propositionally entails* $\psi\sigma$, and in proof systems, we refer to steps deriving $\psi\sigma$ from $\Phi\sigma$ as *propositional reasoning over* L.

Moreover, we fix notions of rank-1 axioms and one-step rules, i.e. rules with a purely propositional premise and a conclusion which is a modal formula of rank 1:

Definition 10 A *one-step rule* R over a set V of propositional variables is a rule ϕ/ψ, where $\phi \in \mathsf{Prop}(V)$ and ψ is a clause over $\mathsf{Up}(V)$. An *extended one-step rule* has the form ϕ/ψ, where $\phi \in \mathsf{Prop}(V)$ and ψ is a clause over $\mathsf{Up}(\mathsf{Prop}(V))$. We will refer to extended one-step rules just as *rules* when this is unlikely to cause confusion. A *rank*-1 *clause* is an extended one-step rule with empty premise, i.e. a clause over $\mathsf{Up}(\mathsf{Prop}(V))$. A rule ϕ/ψ is *sound* if, whenever $\phi\sigma$ is valid for an $\mathcal{L}(\Lambda)$-substitution σ, then $\psi\sigma$ is valid, and *one-step sound* if, whenever $Y, \tau \models \phi$ for a set Y and a $\mathcal{P}(Y)$-valuation τ, then $TY, \tau \models \psi$.

In particular, a rank-1 clause ψ is sound if all its $\mathcal{L}(\Lambda)$-instances are valid formulas, and one-step sound if $TX, \tau \models \psi$ for every set X and every $\mathcal{P}(X)$-valuation τ. Note that rank-1 clauses are essentially equivalent to the formal notion of *axiom* used in [20,13]. One-step soundness of rank-1 clauses has been called *admissibility* in [20]. The term 'axiomatisation in rank 1' formally refers to axiomatisations by rank-1 clauses or, equivalently, (extended) one-step rules (cf. Proposition 15 below).

Proposition 11 *Every one-step sound rule is sound.*

PROOF. Straightforward. □

Remark 12 The converse of the previous proposition holds under mild additional assumptions (namely when T is ω-accessible, the set Λ of predicate liftings is separating in the sense of [21], and the final T-coalgebra is infinite); this is a straightforward generalisation of [26], Theorem 30.

Proof systems for coalgebraic modal logic are parametrised over a given *axiomatisation*, i.e. a set $\mathcal{R}$ of one-step sound rules, similarly as in [13]; rank-1 clauses in $\mathcal{R}$ are referred to as *axioms*.

Definition 13 Let $\mathcal{R}_C$ denote the set of rules obtained by extending $\mathcal{R}$ with the *congruence rule*

$$\frac{a \leftrightarrow b}{[\lambda]a \rightarrow [\lambda]b}$$

(this rule is also known under the name *replacement of equivalents*). The set of formulas *derivable* under $\mathcal{R}$ is the smallest set closed under propositional entailment (cf. Definition 10) and the application of rules ϕ/ψ in $\mathcal{R}_C$, i.e. inference of $\psi\sigma$ from $\phi\sigma$ for a $\mathcal{L}(\Lambda)$-substitution σ.

It is easy to see that this proof system is sound. The completeness results in [20,5,13] require the presence of 'enough' axioms in the following sense.

Definition 14 The set $\mathcal{R}$ is *one-step complete* if, whenever $TX, \tau \models \chi$ for a set X, a clause χ over $\mathsf{Up}(V)$, and a $\mathcal{P}(X)$-valuation τ, then χ is *one-step derivable*; i.e. χ is propositionally entailed by clauses $\psi\sigma$, where ϕ/ψ is in $\mathcal{R}_C$ and σ is a $\mathsf{Prop}(V)$-substitution such that $X, \tau \models \phi\sigma$.

Of course, we can restrict ourselves to finite V. Examples of one-step complete axiomatisations are given in [20,5], where also a weak completeness theorem is proved, stating that one-step complete sets induce weakly complete proof systems; this theorem will appear as a corollary to our finite model result in Section 3.

In fact, the proof systems in [20,5,13] employ only axioms rather than rules. However, we can always interchange one-step rules, extended one-step rules, and axioms: to begin, one-step rules can replace extended one-step rules thanks to the congruence rule (just introduce premises abbreviating propositional formulas as propositional variables). In particular, every axiom can be replaced by a one-step rule. Conversely, we can replace one-step rules by axioms:

Proposition 15 *For each one-step rule R over V, there exists a rank-1 clause χ over V such that χ and R are derivable from each other by propositional reasoning and the congruence rule. Explicitly: for $R = \phi/\psi$, ψ can be derived from ϕ and χ using the congruence rule and propositional reasoning over V and over $\mathsf{Up}(\mathsf{Prop}(V))$ (cf. Definition 10), and conversely, χ can be derived using R and propositional reasoning over V.*

The proof needs the following fact from propositional logic.

Lemma 16 *Let $\phi \in \mathsf{Prop}(V)$ be satisfiable. Then there exists a $\mathsf{Prop}(V)$-substitution σ such that*

$$\phi \rightarrow (a \leftrightarrow \sigma(a)) \quad \text{(for each } a \in V\text{) and}$$

$$\phi\sigma$$

are tautologies.

PROOF. Let κ be a valuation such that $\phi\kappa = \top$, and put

$$\sigma(a) = \begin{cases} a \wedge \phi, & \text{if } \kappa(a) = \bot; \\ \phi \to a & \text{otherwise.} \end{cases}$$

It is then clear that $\phi \to (a \leftrightarrow \sigma(a))$ is a tautology. It remains to be shown that $\phi\sigma\tau = \top$ for each valuation τ. If $\phi\tau = \top$, then $\phi\sigma\tau$ is, by the preceding observation, equivalent to $\phi\tau$ and hence true. If $\phi\tau = \bot$, then $\sigma(a)\tau = \kappa(a)$ and hence $\phi\sigma\tau = \phi\kappa = \top$. □

PROOF. [Proposition 15] We can assume that the premise ϕ of $R = \phi/\psi$ is satisfiable. Thus, fix σ as in Lemma 16 for ϕ. Then R and the rank-1 clause $\psi\sigma$ are mutually interderivable: to derive R from $\psi\sigma$, assume ϕ. Then $a \leftrightarrow \sigma(a)$ for all a by propositional reasoning over V, and hence we can derive ψ from $\psi\sigma$ by the congruence rule and propositional reasoning over $\mathsf{Up}(\mathsf{Prop}(V))$. Conversely, we can derive $\psi\sigma$ from the tautology $\phi\sigma$ using R. □

Remark 17 Lemma 16 and Proposition 15 are of independent interest; it is not clear to the authour whether these results are previously known in the literature. As a simple example application, consider the monotonicity rule

$$\frac{a \to b}{\Box a \to \Box b}.$$

By the proofs of Lemma 16 and Proposition 15, we obtain three axioms which may replace this rule (assuming the congruence rule) from the three satisfying valuations for $a \to b$, namely

$$\begin{aligned} &\Box a \to \Box(a \vee b), \\ &\Box(a \wedge b) \to \Box(a \vee b), \text{ and} \\ &\Box(a \wedge b) \to \Box b. \end{aligned}$$

We now proceed to establish that every coalgebraic modal logic indeed admits a one-step complete axiomatisation.

Theorem 18 *The set of all one-step sound one-step rules is one-step complete.*

PROOF. Let $TX, \tau \models \psi$ for a set X, a clause ψ over $\mathsf{Up}(V)$, with V assumed to be finite, and a $\mathcal{P}(X)$-valuation τ. Let the formula ϕ be the 'propositional theory' of τ, i.e. the (finite) conjunction of all clauses χ over V such that $X, \tau \models \chi$. We will show that the one-step rule $R \equiv \phi/\psi$ over V is one-step sound; it then follows that ψ is one-step derivable by applying R to $X, \tau \models \phi$.

Thus, let Y be a set, and let σ be a $\mathcal{P}(Y)$-valuation such that $Y, \sigma \models \phi$. We have to show $TY, \sigma \models \psi$. For each $y \in Y$, there exists $x \in X$ such that

$$x \in \tau(a) \iff y \in \sigma(a)$$

for all $a \in V$: the non-existence of x would amount to a clause χ over V such that $X, \tau \models \chi$ but $Y, \sigma \not\models \chi$ by virtue of the existence of y, in contradiction to $Y, \sigma \models \phi$.

Thus, we have $f : Y \to X$ such that

$$\sigma(a) = f^{-1}[\tau(a)] \quad \text{for all } a \in V.$$

By naturality of the $\lambda \in \Lambda$, and since preimages preserve intersections and complements, we now obtain

$$[\![\psi]\!]\sigma = Tf^{-1}[[\![\psi]\!]\tau].$$

Since $TX, \tau \models \psi$, we conclude $TY, \sigma \models \psi$ as required. □

Since by Proposition 15, rank-1 clauses and one-step rules are interchangeable, we obtain

Corollary 19 *The set of all one-step sound rank-*1 *clauses is one-step complete.*

Remark 20 Note that the proof of Theorem 18 establishes one-step derivability using only a single V-instance rather than several $\mathsf{Prop}(V)$-instances of the given rules, and thus actually shows that the set of all one-step sound one-step rules is *strictly* one-step complete in the sense of [31].

Remark 21 Corollary 19 may also be viewed as a limitative result saying that coalgebraic modal logic covers only modal logics that are axiomatisable by rank-1 clauses (in fact, exactly the modal logics axiomatisable by rank-1 clauses [32]). However, modal logics axiomatised outside rank 1, such as the classical logic $S4$, may be modelled by restricting the semantics to a suitable subclass of the class of all coalgebras for the given signature functor. Research aimed at generic model-theoretic and algorithmic results for such logics is under way.

One can turn the above results into one-step completeness criteria stating e.g. that a set of rank-1 clauses is one-step complete if its closure under substitution, congruence, and propositional entailment contains all one-step sound rank-1 clauses. However, in concrete examples one-step completeness proofs using such criteria are apparently more or less identical to direct proofs. For purposes of our finite model result, we are more interested in the fact that every coalgebraic modal logic has a one-step complete axiomatisation, if not necessarily a recursive one. Note however the following facts.

Definition 22 The *one-step validity* problem is to decide whether $TX, \tau \models \phi$ for a finite set X, a (disjunctive) clause ϕ over $\mathsf{Up}(V)$, and a $\mathcal{P}(X)$-valuation τ.

Proposition 23 *A rank-*1 *clause* ϕ *over* V *is one-step sound iff* $T(\mathcal{P}(V)), \sigma \models \phi$*, where* σ *is the* $\mathcal{P}(\mathcal{P}(V))$*-valuation given by*

$$\sigma(a) = \{B \in \mathcal{P}(V) \mid a \in B\}.$$

PROOF. The 'only if' direction is trivial. To prove the 'if' direction, let $T(\mathcal{P}(V)), \sigma \models \phi$. We have to show that $TX, \tau \models \phi$ for every $\mathcal{P}(X)$-valuation τ. This follows by naturality of predicate liftings, applied to the map $f : X \to \mathcal{P}(V)$, $x \mapsto \{a \in V \mid x \in \tau(a)\}$. $\square$.

Corollary 24 *One-step soundness of rank-*1 *clauses is decidable (semi-decidable) if the one-step validity problem is decidable (semi-decidable).*

We will see in Example 42 that the dual of one-step validity, *one-step satisfiability* (Definition 34 below), is decidable in many important cases; of course, decidability of one-step validity then follows. Thus, Corollary 19 does frequently supply a feasible axiomatisation, although one will, of course, in general strive for a more compact axiomatisation.

3 The Finite Model Construction

The non-compactness of coalgebraic modal logic (cf. Example 7) means that canonical models, based on the set of all maximally consistent sets w.r.t. a finitary deduction system, do not in general exist. An alternative is to use filtration methods (cf. e.g. [4,2]), in the variant that uses maximally consistent subsets of finite sublanguages. In more detail, we construct a satisfying model for a consistent formula ϕ by taking as the carrier set the set S of all maximally consistent subsets of the closure of ϕ under subformulas and negation; the crucial problem here is the construction of the coalgebra structure on S. We show that there exists a coalgebra structure on S that allows proving the truth lemma, provided that the given axiomatisation is one-step complete. This reproves the weak completeness result of [20,5], and in combination with Theorem 18 implies as a corollary that *coalgebraic modal logic has the finite model property.*

Unlike in [27], we have phrased all results below in terms of axiomatisations by extended one-step rules — this subsumes both axiomatisations by rank-1 clauses, as in the original formulation [27], and axiomatisations by one-step rules.

We recall a few basic definitions:

Definition 25 Given a set $\mathcal{R}$ of rules, a finite set $\{\phi_1, \dots, \phi_n\}$ of formulas is called *$\mathcal{R}$-consistent* if $\neg(\phi_1 \wedge \dots \wedge \phi_n)$ is not derivable under $\mathcal{R}$ (Definition 13). A set Σ of formulas is called *closed* if it is closed under subformulas and under *normalised negation* $\sim$, where $\sim\phi$ is defined to be ψ in case ϕ is of the form $\neg\psi$, and $\neg\phi$ otherwise. A subset A of Σ is called a *Σ-Hintikka set* if $\bot \notin A$ and, for $\phi \wedge \psi \in \Sigma$, $\phi \wedge \psi \in A$ iff $\phi, \psi \in A$, and, for $\neg\phi \in \Sigma$, $\neg\phi \in A$ iff $\phi \notin A$. Moreover, A is called a *Σ-atom* if A is maximal among the $\mathcal{R}$-consistent subsets of Σ.

Thus, a Σ-atom is just an $\mathcal{R}$-consistent Σ-Hintikka set. As usual, one proves

Lemma 26 (Lindenbaum Lemma) *Every $\mathcal{R}$-consistent subset of Σ is contained in a Σ-atom.*

The following lemma is crucial for the construction of a suitable coalgebra structure on the set S of Σ-atoms.

Lemma 27 *Let V be a set of propositional variables, let $\phi \in \mathsf{Prop}(V)$, and let σ be a Σ-substitution. Then $\phi\sigma$ is derivable iff $S, \tau \models \phi$, where S is the set of Σ-atoms and τ is the $\mathcal{P}(S)$-valuation given by $\tau(a) = \{A \in S \mid \sigma(a) \in A\}$.*

PROOF. We prove the equivalent claim that $\phi\sigma$ is $\mathcal{R}$-consistent iff ϕ is satisfiable over τ. Assume w.l.o.g. that ϕ is a conjunctive clause. Let $\Phi \subseteq \Sigma$ be the set containing, for each $a \in V$, the formula $\sigma(a)$ if ϕ contains the literal a and the formula $\sim\sigma(a)$ if ϕ contains the literal $\neg a$. Then $\phi\sigma$ is $\mathcal{R}$-consistent iff Φ is $\mathcal{R}$-consistent iff (by the Lindenbaum Lemma) $\Phi \subseteq A$ for some $A \in S$ iff ϕ is satisfiable over τ, where the last equivalence uses the fact that A is Hintikka. $\square$

In the expectation that the extension of a formula $\phi \in \Sigma$ in the coalgebra (S, ξ) to be constructed will be the set $\{A \in S \mid \phi \in A\}$, we will need to require that

$$\xi(A) \in \lambda_S\{B \in S \mid \phi \in B\} \iff [\lambda]\phi \in A \qquad (*)$$

for all $A \in S$ and all formulas $[\lambda]\phi$ in Σ. This is where one-step completeness comes in:

Lemma 28 (Existence Lemma) *Let S be the set of Σ-atoms. If $\mathcal{R}$ is one-step complete and Σ is finite, then $\xi(A)$ satisfying $(*)$ exists for each $A \in S$.*

PROOF. Assume that $\xi(A)$ does *not* exist. We introduce a propositional variable a_ϕ for each $\phi \in \Sigma$ and put $V = \{a_\phi \mid \phi \in \Sigma\}$. Let ψ be the clause over $\mathsf{Up}(V)$ containing, for each $[\lambda]\phi \in \Sigma$, the literal $\neg[\lambda]a_\phi$ if $[\lambda]\phi \in A$, and the literal $[\lambda]a_\phi$ otherwise. Let τ be the $\mathcal{P}(S)$-valuation taking a_ϕ to $\{B \mid \phi \in B\}$. Then $TS, \tau \models \psi$ by assumption. By one-step completeness, ψ is one-step derivable from propositional formulas χ such that $S, \tau \models \chi$. This derivation can be copied to obtain a derivation of $\psi\sigma$ from formulas $\chi\sigma$ such that $S, \tau \models \chi$, where σ is the Σ-substitution taking a_ϕ to ϕ. These $\chi\sigma$ are derivable by Lemma 27. Thus, $\psi\sigma$ is derivable, in contradiction to the $\mathcal{R}$-consistency of A. □

It remains to prove the truth lemma, which we state in a slightly more general form than needed in this section for reuse in Section 4:

Lemma 29 (Truth Lemma) *Let Σ be a closed set, let S be a set of Σ-Hintikka sets, and let $\xi : S \to TS$ satisfy condition $(*)$ above. Then for all $\phi \in \Sigma$ and all $A \in S$,*

$$A \models_{(S,\xi)} \phi \iff \phi \in A.$$

PROOF. Straightforward induction over ϕ. □

This is all we need in order to establish

Theorem 30 *Let $\mathcal{R}$ be one-step complete. Then every formula ϕ that is $\mathcal{R}$-consistent is locally satisfiable in a finite T-coalgebra of size at most $2^{|\phi|}$.*

PROOF. Let $\Sigma(\phi)$ be the smallest closed set containing ϕ, and let S be the set of $\Sigma(\phi)$-atoms. By the Existence Lemma, S can be equipped with a T-coalgebra structure ξ satisfying condition $(*)$. By the Lindenbaum Lemma, there exists $A \in S$ such that $\phi \in A$. By the Truth Lemma, $A \models_{(S,\xi)} \phi$. □

As announced, the above result implies weak completeness [20,5]; explicitly:

Corollary 31 (Weak completeness) *The proof system induced by a one-step complete axiomatisation is weakly complete.*

Combining Theorems 18 and 30, we obtain independently of deduction:

Corollary 32 (Finite model property) *Every locally satisfiable formula ϕ is locally satisfiable in a finite T-coalgebra of size at most $2^{|\phi|}$.*

Remark 33 By Remark 6, the above corollary reproves the fact that graded modal logic has the finite model property when interpreted over Kripke frames [35], although of course not with the same size bound as for the interpretation over multigraphs (e.g. any Kripke frame in which $\phi \equiv \Diamond_k \top$ is satisfiable will have at least k states, while $|\phi| = 2$).

4 Decidability

Unlike in the classical case, the finite model construction of the preceding section does not immediately imply that satisfiability of modal formulas is decidable, even though the construction gives a computable bound on the size of the model. The problem is that there may in general be infinitely many T-coalgebras on a given finite set. (In fact, this is the interesting case; for functors T that preserve finite sets, a finite model construction is given already in [20].) If T takes finite sets to recursively enumerable sets — as is the case e.g. for $\mathcal{B}_\mathbb{N}$, $\mathcal{B}_\mathbb{Z}$, and $\mathbb{Z}[\_]$, but not for D_ω — then the finite model property implies that the set of satisfiable formulas is recursively enumerable. We then obtain decidability provided that the set of valid formulas is also recursively enumerable, which by the weak completeness theorem and Corollary 19 is the case if the set of all one-step sound rank-1 clauses is recursively enumerable. By Corollary 24, the latter condition may be further reduced to semi-decidability of one-step validity.

We can however improve on this by exploiting the details of the finite model construction, as follows. We have no direct access to the set of all $\Sigma(\phi)$-atoms, since this would already require a decision procedure for consistency. We can however easily decide whether a subset of $\Sigma(\phi)$ is Hintikka. We are then faced with the following decision problem:

Definition 34 The *one-step satisfiability* problem is to decide, given a *conjunctive* clause ϕ over $\mathsf{Up}(V)$, a finite set X, and a $\mathcal{P}(X)$-valuation τ, whether ϕ is one-step satisfiable over τ.

Remark 35 For the complexity considerations below, we need to fix a notion of *representation size* of formulas in $\mathcal{L}(\Lambda)$, conjunctive clauses over $\mathsf{Up}(V)$, and $\mathcal{P}(X)$-valuations. We count 1 for every boolean operation and every modal operator, and $1 + \log|V|$ for every propositional variable. Moreover, for a subset of a finite set X, we assume a representation size of $|X| + 1$. Finally, we need to account for the representation of indices of modal operators. The representation size $\mathit{size}(k)$ of an integer k is $\lceil \log_2(|k| + 1) \rceil$, where $\lceil r \rceil = \min\{z \in \mathbb{Z} \mid z \geq r\}$ as usual. The representation size $\mathit{size}(p)$ of a rational number $p = k/l$, with k, l relatively prime, is $1 + \mathit{size}(k) + \mathit{size}(l)$. Note that the representation size thus defined for a formula ϕ is typically larger than the size $|\phi|$ as defined in Section 1.

A decision procedure for one-step satisfiability leads to a decision procedure for $\mathcal{L}(\Lambda)$ which essentially tries to find a set S of Hintikka sets which admits a T-coalgebra structure satisfying Condition $(*)$ of Section 3.

Algorithm 1 (Decide satisfiability of $\phi \in \mathcal{L}(\Lambda)$) Let H denote the set of $\Sigma(\phi)$-Hintikka sets, and introduce a propositional variable a_ψ for each $\psi \in \Sigma$. For all subsets S of H, perform the following steps.

(1) Check whether $\phi \in A$ for some $A \in S$; if not, continue with the next S.
(2) Decide whether for all $A \in S$, the conjunctive clause

$$\bigwedge_{[\lambda]\psi \in A} [\lambda]a_\psi \wedge \bigwedge_{\neg[\lambda]\psi \in A} \neg[\lambda]a_\psi$$

is satisfiable over the $\mathcal{P}(S)$-valuation τ defined by

$$\tau(a_\psi) = \{B \in S \mid \psi \in B\}.$$

If yes, terminate with output 'yes'; otherwise, continue in Step 1 with the next S.

If all S have been checked unsuccessfully, terminate with output 'no'.

Thus, we have the following decidability criterion.

Theorem 36 *If one-step satisfiability is decidable, then satisfiability of $\mathcal{L}(\Lambda)$-formulas is decidable.*

PROOF. We have to prove the correctness of Algorithm 1. It is clear that the algorithm terminates. If the algorithm terminates successfully for a formula ϕ, then it has found a set S of $\Sigma(\phi)$-Hintikka sets satisfying Condition $(*)$ of Section 3 such that $\phi \in A$ for some $A \in S$, and therefore ϕ is satisfiable by the Truth Lemma (Lemma 29). If, conversely, ϕ is satisfiable, then the algorithm will terminate successfully, since by the Lindenbaum Lemma (Lemma 26) and the Existence Lemma (Lemma 28), the set S of $\Sigma(\phi)$-atoms meets the search criteria. □

A non-deterministic variant of Algorithm 1 will also be useful:

Algorithm 2 Nondeterministically choose $S \subseteq H$; then proceed as in Algorithm 1, but fail (i.e. loop infinitely) rather than continue with the next S if one of the checks in Steps 1 or 2 fails.

In this algorithm, we can also employ a *semi*-decision procedure for one-step satisfiability. Correctness of the algorithm is shown in the same way as for Algorithm 1. Since acceptance sets of non-deterministic algorithms are recursively enumerable, we thus have

Theorem 37 *If one-step satisfiability is semi-decidable, then satisfiability of $\mathcal{L}(\Lambda)$-formulas is semi-decidable.*

(Note that semi-decidability of one-step satisfiability is weaker than the above-mentioned condition that T takes finite sets to recursively enumerable sets. E.g., the one-step satisfiability problem will turn out to be decidable for probabilistic modal logic (Example 5.5), although $D_\omega(X)$ is uncountable for $|X| \geq 2$.)

Algorithm 2 yields the not overly tight upper complexity bound to be expected for filtration-based algorithms:

Theorem 38 *If the one-step satisfiability problem is in NP, then satisfiability of $\mathcal{L}(\Lambda)$-formulas is in NEXPTIME.*

PROOF. By a time analysis of Algorithm 2: S can be constructed in exponential time (being a subset of H, whose size is exponential in the size of the formula ϕ), and going through all $A \in S$ gives an exponential factor (since the size of S is exponential). So the algorithm has exponential run time if in Step 2, the check for a given A can be performed in exponential time; since the input for the one-step satisfiability problem is of exponentially bounded size (cf. Remark 35), this is the case if one-step satisfiability is in *NP*. □

Remark 39 In [5], logics for coalgebras are constructed in a modular fashion, following the structure of the signature functor; this raises the question of whether the above decidability and complexity results behave well w.r.t. these constructions. It is easy to see that decision procedures for one-step satisfiability can be combined along products and sums of functors and their logics, while this is not so clear for the case of functor composition $S \circ T$, as the application of T to finite sets may e.g. produce an exponential blowup or lead to infinite sets. This problem may possibly be resolved by moving to multisorted coalgebra [30].

Besides the examples whose decidability is already captured by the finite model result of [20], i.e. functors preserving finite sets, such as $\mathcal{P}$, our results cover also more complex logics. The treatment of probabilistic modal logic requires the following fact from linear programming:

Definition 40 A *mixed rational linear inequation system* is a system of inequations

$$Ax < b \quad \text{and} \quad Bx \leq c, \tag{$*$}$$

where A and B are rational matrices and b and c are rational vectors. Following [25], we define the *size* of a vector $a = (a_1, \dots, a_n)$ as $size(a) = n + \sum_i size(a_n)$, and the size of an $n \times m$ matrix $A = (a_{ij})$ as $size(A) = nm + \sum_{i,j} a_{ij}$. The size of a mixed rational linear inequation system as above is $1 + size(A) + size(B) + size(b) + size(c)$.

Theorem 41 *Solvability of mixed rational linear inequation systems is in P.*

PROOF. By Motzkin's transposition theorem ([25], Corollary 7.1k), the system $(*)$ has a solution iff for all row vectors $y, z \geq 0$, the following conditions hold:

(i) if $yA + zB = 0$, then $yb + zc \geq 0$, and
(ii) if $yA + zB = 0$ and $y \neq 0$, then $yb + zc > 0$.

By Farkas' Lemma ([25], Corollary 7.1e), (i) is equivalent to solvability of the system $Ax \leq b$, $Bx \leq c$; solvability of rational linear inequation systems is in P [25].

Concerning (ii), we recall that the following *linear programming problem* is computable in polynomial time [25].

> Given a rational matrix A, a rational column vector b, and a rational row vector c, decide whether $\max\{cx \mid Ax \leq b\}$ is infeasible (i.e. $Ax \leq b$ is unsolvable), finite, or infinite. If it is finite, find an optimal solution.

Under (i), (ii) is equivalent to

> if $yA + zB = 0$ and $yb + zc = 0$, then $y = 0$.

This in turn can be reformulated as the following instance of the linear programming problem:

> $\max\{dy \mid y, z \geq 0, yA + zB = 0, yb + zc = 0\} = 0$, where d is a row vector consisting of 1s.

Thus, also (ii) can be decided in polynomial time once (i) is established. □

Example 42 (1) Let Λ be the set of predicate liftings λ_k for the multiset functor of Example 5.3. For a $\mathcal{P}(X)$-valuation τ with X finite, satisfaction of a positive literal $\Diamond_k a$ by $\sum_{x \in X} n_x x \in \mathcal{B}_{\mathbb{N}}(X)$ then amounts to the inequation

$$\sum_{x \in \tau(a)} n_x \geq k + 1,$$

and satisfaction of the corresponding negated literal is

$$\sum_{x \in \tau(a)} n_x \leq k.$$

Thus, one-step satisfiability amounts to solvability of a system of integer linear inequations over the naturals, which is decidable in *NP* [25]. By Theorem 38, we obtain that graded modal logic is in *NEXPTIME*. (In fact, graded modal logic is in *PSPACE* [33].)

(2) By the same line of reasoning, the satisfiability problem for generalised graded modal logic over coalgebras for the generalised multiset functor (Example 5.4) is in *NEXPTIME*.

(3) Let Λ be the set of predicate liftings for probabilistic modal logic as in Example 5.5. Analogously to the case of graded modal logic, the one-step satisfiability problem reduces to the solvability of mixed rational linear inequation systems over the reals (with strict inequalities arising from negated literals $\neg L_p a$), which by Theorem 41 is in *P*. By Theorem 36, it follows that probabilistic modal logic is in *NEXPTIME*. (By results of [8], probabilistic modal logic is in fact in *PSPACE*.)

(4) By the same reasoning, the modal logic for linear automata of Example 5.6 is decidable in *NEXPTIME*.

(5) It is straightforward to extend the above results to include proposition symbols, where not already present, or labels (cf. Examples 5.2 and 5.7).

The upper bounds in (2) and (4) are, to our knowledge, new, if unsurprising in the light of (1) and (3). They can however be improved to *PSPACE* using more advanced methods (cf. Remark 43).

Remark 43 Better general upper complexity bounds are the subject of ongoing research. Most coalgebraic modal logics are at least *PSPACE*-hard as they embed either K or KD, which are known to be *PSPACE*-complete [2]. The following further results have been obtained so far:

(i) One can show by means of elimination of Hintikka sets (in the same manner as in known algorithms for PDL [2]) that satisfiability of $\mathcal{L}(\Lambda)$-formulas is in *EXPTIME* if one-step satisfiability is in *P* [29].

(ii) Given a tractable axiomatisation of $\mathcal{L}(\Lambda)$, one can show that satisfiability of $\mathcal{L}(\Lambda)$-formulas is in *PSPACE* by means of a shallow model construction [31].

(iii) A semantic criterion for a coalgebraic modal logic to be decidable in *PSPACE* is given in [28]; the relevant conditions here are a small model property and tractability of model checking at the one-step level.

(iv) It is shown in [30] that the method of (ii) is compositional w.r.t. composition of functors $\mathbf{Set}^n \to \mathbf{Set}$.

(None of these results makes Theorem 38 obsolete, since they rely on stronger assumptions.) By (i), one immediately improves the upper bound for probabilistic modal logic as well as for the modal logic of linear automata from *NEXPTIME* (Example 42) to *EXPTIME*. Moreover, the method of (ii) reproduces the known *PSPACE* upper bounds for K, graded modal logic, and probabilistic modal logic (and, in the conference presentation of [31], established a presumably novel *PSPACE* upper bound for majority logic [18] simultaneously with [7]); this requires the construction of a suitable tractable axiomatisation in each case [31]. Adaptation of these upper bounds to generalised graded modal logic and the logic of linear automata (Examples 5.4 and 5.6) is straightforward. The method of (iii) applies in particular to various logics of uncertainty, for which proofs of the relevant criteria are available as off-the-shelf results. Finally, the compositionality result (iv) leads to large numbers of example applications to logics for systems with mixed branching, e.g. with both non-deterministic and probabilistic branching.

It remains an open problem to extend these results to logics axiomatised outside rank 1 such as the classical logics KT and $S4$ (preliminary results have been obtained for logics axiomatised by formulas that combine purely propositional formulas and rank-1 formulas, such as the T-axiom $\Box a \to a$). Moreover, one might conjecture that one may in fact just replace *NEXPTIME* with *PSPACE* in Theorem 38; the present *PSPACE* criteria make stronger assumptions than decidability of one-step satisfiability in *NP*.

Remark 44 While it is easy to construct contrived examples with computationally hard one-step satisfiability problems, we do not know of a natural example of a coalgebraic modal logic that fails to satisfy the criterion of Theorem 38, or indeed the criteria for tighter upper complexity bounds mentioned in Remark 43.

5 Conclusion

We have established that coalgebraic modal logic has the finite model property, and we have described an ensuing generic algorithm for deciding satisfiability, assuming a decision procedure for the rather simpler *one-step satisfiability* problem. We have thus proved decidability for a wide range of modal logics, including graded and probabilistic modal logic. This goes significantly beyond the decidability result of [20], which applies only to signature functors that preserve finite sets, such as the powerset functor (whose coalgebras are standard Kripke frames). Moreover, assuming a mild upper complexity bound (*NP*) for one-step satisfiability, we have established a first general upper complexity bound for coalgebraic modal logic (*NEXPTIME*). This result applies to both graded and probabilistic modal logic. Better upper bounds (*PSPACE*) for these logics are known in the literature; in fact, subsequent work on the complexity of coalgebraic modal logic [31,30,28] establishes, under additional assumptions, a tighter generic upper bound which in particular reproduces the *PSPACE* upper bounds for graded and probabilistic modal logic.

The results of this work and [32] imply that a modal logic can be regarded as a coalgebraic modal logic iff it can be axiomatised by rank-1 formulas, i.e. propositional combinations of atoms of the form $\Box\phi$ (or more generally $\Box(\phi_1, \dots, \phi_n)$ if polyadic modal operators are considered), where $\Box$ is a modal operator and ϕ is a propositional formula; this format excludes logics such as $S4$ or KT. It is an important direction of future research to extend the existing algorithmic results to logics axiomatised outside rank 1, which may be modelled coalgebraically by imposing suitable restrictions on the coalgebras appearing as semantic domains. A further open problem is to obtain generic complete axiomatisations and decision procedures for coalgebraic modal logics with iteration.

Acknowledgements

The authour wishes to thank Till Mossakowski, Markus Roggenbach, and Horst Reichel for collaboration on CoCasl, Erwin R. Catesbeiana for advice on empty coalgebras, Dirk Pattinson for useful discussions, and the anonymous referees for their helpful suggestions for improvement of this work.

References

[1] F. Bartels, A. Sokolova, and E. de Vink. A hierarchy of probabilistic system types. In H. P. Gumm, editor, *Coalgebraic Methods in Computer Science, CMCS 03*, volume 82 of *Electron. Notes Theoret. Comput. Sci.* Elsevier, 2003.

[2] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.

[3] J. W. Carlyle and A. Paz. Realizations by stochastic finite automata. *J. Comput. System Sci.*, 5:26–40, 1971.

[4] B. Chellas. *Modal Logic*. Cambridge University Press, 1980.

[5] C. Cîrstea and D. Pattinson. Modular construction of modal logics. In L. Rutkowski, J. Siekmann, R. Tadeusiewicz, and L. A. Zadeh, editors, *Concurrency Theory, CONCUR 04*, volume 3170 of *Lect. Notes Comput. Sci.*, pages 258–275. Springer, 2004.

[6] G. D'Agostino and A. Visser. Finality regained: A coalgebraic study of Scott-sets and multisets. *Arch. Math. Logic*, 41:267–298, 2002.

[7] S. Demri and D. Lugiez. Presburger modal logic is only PSPACE-complete. In U. Furbach and N. Shankar, editors, *Automated Reasoning, IJCAR 06*, volume 4130 of *Lecture Notes in Artificial Intelligence*, pages 541–556. Springer, 2006.

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

[9] V. Halava and T. Harju. Undecidability in integer weighted finite automata. *Fundam. Inform.*, 38:189–200, 1999.

[10] A. Heifetz and P. Mongin. Probabilistic logic for type spaces. *Games and Economic Behavior*, 35:31–53, 2001.

[11] M. Hennessy and R. Milner. Algebraic laws for non-determinism and concurrency. *J. ACM*, 32:137–161, 1985.

[12] B. Jacobs. Towards a duality result in the modal logic of coalgebras. In H. Reichel, editor, *Coalgebraic Methods in Computer Science, CMCS 00*, volume 33 of *Electron. Notes Theoret. Comput. Sci.* Elsevier, 2000.

[13] C. Kupke, A. Kurz, and D. Pattinson. Algebraic semantics for coalgebraic logics. In J. Adámek and S. Milius, editors, *Coalgebraic Methods in Computer Science, CMCS 04*, volume 106 of *Electron. Notes Theoret. Comput. Sci.*, pages 219–241. Elsevier, 2004.

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

[15] K. Larsen and A. Skou. Bisimulation through probabilistic testing. *Inform. Comput.*, 94:1–28, 1991.

[16] L. Moss. Coalgebraic logic. *Ann. Pure Appl. Logic*, 96:277–317, 1999.

[17] T. Mossakowski, L. Schröder, M. Roggenbach, and H. Reichel. Algebraic-co-algebraic specification in CoCasl. *J. Logic Algebraic Programming*, 67:146–197, 2006.

[18] E. Pacuit and S. Salame. Majority logic. In D. Dubois, C. A. Welty, and M.-A. Williams, editors, *Principles of Knowledge Representation and Reasoning, KR 04*, pages 598–604. AAAI Press, Menlo Park, California, 2004.

[19] D. Pattinson. Semantical principles in the modal logic of coalgebras. In A. Ferreira and H. Reichel, editors, *Symposium on Theoretical Aspects of Computer Science, STACS 01*, volume 2010 of *Lect. Notes Comput. Sci.*, pages 514–526. Springer, 2001.

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

[21] D. Pattinson. Expressive logics for coalgebras via terminal sequence induction. *Notre Dame J. Formal Logic*, 45:19–33, 2004.

[22] M. Rößiger. Coalgebras and modal logic. In H. Reichel, editor, *Coalgebraic Methods in Computer Science, CMCS 00*, volume 33 of *Electron. Notes Theoret. Comput. Sci.* Elsevier, 2000.

[23] J. Rothe, H. Tews, and B. Jacobs. The Coalgebraic Class Specification Language CCSL. *J. Universal Comput. Sci.*, 7:175–193, 2001.

[24] J. Rutten. Universal coalgebra: A theory of systems. *Theoret. Comput. Sci.*, 249:3–80, 2000.

[25] A. Schrijver. *Theory of linear and integer programming.* Wiley Interscience, 1986.

[26] L. Schröder. Expressivity of coalgebraic modal logic: the limits and beyond. *Theoret. Comput. Sci.* To appear. Earlier version in V. Sassone, editor, *Foundations of Software Science And Computation Structures, FOSSACS 05*, volume 3441 of Lect. Notes Comput. Sci., Springer, 2005, 440–454.

[27] L. Schröder. A finite model construction for coalgebraic modal logic. In L. Aceto and A. Ingólfsdóttir, editors, *Foundations Of Software Science And Computation Structures, FOSSACS 06*, volume 3921 of *Lect. Notes Comput. Sci.*, pages 157–171. Springer, 2006.

[28] L. Schröder. A semantic *PSPACE*-criterion for rank 1 modal logic. Technical report, Universität Bremen, 2006.

[29] L. Schröder and D. Pattinson. Work in progress.

[30] L. Schröder and D. Pattinson. Compositional algorithms for heterogeneous modal logics. Technical report, Universität Bremen, 2006.

[31] L. Schröder and D. Pattinson. PSPACE reasoning for rank-1 modal logics. In *Logic in Computer Science, LICS 06*, pages 231–240. IEEE, 2006. Presentation slides available under www.informatik.uni-bremen.de/∼lschrode/slides/rank1pspace.pdf.

[32] L. Schröder and D. Pattinson. Rank 1 modal logics are coalgebraic. Technical report, Universität Bremen, 2006.

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

[34] P. Turakainen. On the minimization of linear space automata. *Ann. Acad. Sci. Fenn. Ser. A I*, 506, 1972. 15 pp.

[35] W. van der Hoek. *Modalities for Reasoning about Knowledge and Quantities.* PhD thesis, Vrije Universiteit van Amsterdam, 1992.

[36] I. Viglizzo. Final sequences and final coalgebras for measurable spaces. In J. L. Fiadeiro, N. Harman, M. Roggenbach, and J. Rutten, editors, *Algebra and Coalgebra in Computer Science, CALCO 05*, volume 3629 of *Lect. Notes Comput. Sci.*, pages 395–407. Springer, 2005.