

A Finite Model Construction for Coalgebraic Modal Logic

Lutz Schröder

Department of Computer Science, University of Bremen

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 improves known completeness results for coalgebraic modal logic, which we push further by establishing that every coalgebraic modal logic admits a complete axiomatization of 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.

1 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 [21]. Many concepts of concurrency theory can be cast in the coalgebraic framework; this includes general notions of bisimulation, coinduction, and corecursion, as well as a generic modal logic [10, 19, 12, 15, 17]. The role of this *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 behavioral equivalence of states [17, 23], and is sufficiently intuitive for use in actual software specification languages, including object-oriented specification [20, 12, 14]. 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 [16] and subsequent work [5, 11], a (necessarily 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 set of axioms of rank 1 is weakly complete, provided that the axioms are in a precise sense sufficiently strong; the latter property is referred to as *reflexivity*. Here, we exhibit a finite model construction which relies on reflexivity. We thus reprove the mentioned weak completeness result. Moreover, we show that every coalgebraic modal logic admits a reflexive axiomatization, 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 complexity bound.

The material is organized as follows. Section 2 gives an introduction to coalgebra and coalgebraic modal logic, including a number of examples. In Section 3, we recall the deduction system of coalgebraic modal logic [16, 5, 11] and the above-mentioned notion of reflexivity, and prove that reflexive axiomatizations always exist. We then prove the finite model property in Section 4, from which we obtain the generic decision procedure and the arising complexity bound in Section 5.

2 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 $T : \mathbf{Set} \rightarrow \mathbf{Set}$ 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 an *evolution map* $\xi : X \rightarrow TX$. A *morphism* $(X_1, \xi_1) \rightarrow (X_2, \xi_2)$ of T -coalgebras is a map $f : X_1 \rightarrow X_2$ such that $\xi_2 \circ f = Tf \circ \xi_1$. A T -coalgebra C is called *final* if there exists, for each T -coalgebra A , a unique morphism $A \rightarrow C$.

Intuitively, the evolution map describes the successor states of a state, organized in a data structure given by T . These data encode the observable *behavior* of a system, and morphisms of coalgebras preserve this behavior.

We shall occasionally require a technical condition on the signature functor:

Definition 2. A functor T is called κ -*accessible*, where κ is a regular cardinal, if T preserves κ -directed colimits.

Intuitively, this amounts to a restriction on the branching degree, stating that every state in a T -coalgebra has less than κ successors. Many of the examples given below are ω -accessible; however, the central results presented here hold for arbitrary functors.

We explicitly fix some logical terminology:

Definition 3. Let T be a functor. A *language for T -coalgebras* is a set \mathcal{L} of *formulae*, equipped with a family of *satisfaction* relations \vDash_C (or just \vDash) between states of T -coalgebras $C = (X, \xi)$ and formulae $\phi \in \mathcal{L}$; we define $\llbracket \phi \rrbracket_C$ (or just $\llbracket \phi \rrbracket$) as the set $\{x \in X \mid x \vDash_C \phi\}$. The coalgebra C *satisfies* a formula ϕ if $x \vDash \phi$ for all $x \in X$; in this case, we write $C \vDash \phi$. We say that C *satisfies* a set Φ of formulae, and write $C \vDash \Phi$, if $C \vDash \phi$ for each $\phi \in \Phi$.

A formula $\psi \in \mathcal{L}$ is a *global consequence* of a set $\Phi \subset \mathcal{L}$ if, for every T -coalgebra C , $C \vDash \psi$ whenever $C \vDash \Phi$. In this case, we write $\Phi \vDash_g \psi$. We say that ψ is *valid* if $\emptyset \vDash_g \psi$. Moreover, ψ is a *local consequence* of Φ if, for every

state x in every T -coalgebra, $x \models \psi$ whenever $x \models \Phi$. A set Φ of formulae is *locally satisfiable* if it is satisfied in some state in some T -coalgebra; Φ is *globally satisfiable* if $\Phi \not\models_g \perp$ (i.e. if Φ is satisfied by some non-empty T -coalgebra).

Every local consequence is also a global consequence, and every globally satisfiable formula is locally satisfiable. 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 [17], generalizing previous results [10, 19, 12, 15]. The semantics is based on the following central notion.

Definition 4. A *predicate lifting* for a functor T is a natural transformation

$$\lambda : 2^- \rightarrow 2^T,$$

where 2^- denotes the contravariant powerset functor $\mathbf{Set}^{op} \rightarrow \mathbf{Set}$, with $2^f(A) = f^{-1}[A]$.

A predicate lifting λ induces a *transpose* $\lambda^b : T \rightarrow 2^{(2^-)}$, given by $\lambda_X^b(t) = \{A \subset X \mid t \in \lambda_X(A)\}$. A set Λ of predicate liftings for T is called *separating* if for each set X , the source of maps $(\lambda_X^b : T \rightarrow 2^{(2^-)})_{\lambda \in \Lambda}$ is jointly injective.

In the terminology introduced above, a coalgebraic modal logic is a language $\mathcal{L}^\kappa(\Lambda)$ for T -coalgebras, determined by a set Λ of predicate liftings for T and a regular cardinal κ which serves as a bound for conjunctions. Since we are aiming at finite model results here, we restrict the exposition to the finite case $\kappa = \omega$, and we write $\mathcal{L}(\Lambda)$ in place of $\mathcal{L}^\omega(\Lambda)$. Formulae $\phi, \psi \in \mathcal{L}(\Lambda)$ are defined by the grammar

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

where λ ranges over Λ . Disjunctions $\phi \vee \psi$, truth \top , and other boolean operations are then defined as usual. In the definition of satisfaction, the clauses for boolean operators 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 subformulae of ϕ .

Remark 5. It is shown in [17, 23] that if Λ is separating and T is ω -accessible, then $\mathcal{L}(\Lambda)$ is *adequate* and *expressive*: states x and y in T -coalgebras A and B , respectively, satisfy the same $\mathcal{L}(\Lambda)$ -formulae iff they are *behaviorally equivalent* in the sense that there exists a coalgebra C and morphisms $f : A \rightarrow C, g : B \rightarrow C$ such that $f(x) = g(y)$. This is one reason why modal logic is regarded as a suitable means of expression for coalgebraic specification — it automatically ensures encapsulation of the state space, allowing judgements precisely about the observable behavior of states.

Remark 6. By the results of [23], obtaining an expressive logic for a given ω -accessible functor may require the use of polyadic modal operators obtained from

polyadic predicate liftings. The results of this paper extend straightforwardly to polyadic modal logic; we restrict the exposition to the unary case purely in the interest of readability.

Example 7. [17, 5, 23]

1. Let \mathcal{P} be the covariant powerset functor. Then \mathcal{P} -coalgebras are graphs, thought of as transition systems or indeed Kripke frames. A separating set of predicate liftings is formed by the single predicate lifting λ^\forall defined by

$$\lambda_X^\forall(A) = \{B \subset X \mid B \subset A\}.$$

This lifting gives rise to the standard box modality $\Box = [\lambda^\forall]$. All this is easily adapted to transition systems with branching degree limited by a regular cardinal κ , described as coalgebras for the κ -accessible functor \mathcal{P}_κ defined by $\mathcal{P}_\kappa(X) = \{A \subset X \mid |A| < \kappa\}$.

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. Separation is then ensured by adding predicate liftings λ^a , $a \in V$, defined by

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

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.

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 $B : X \rightarrow \mathbb{N}$ 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 $f : X \rightarrow Y$, $\mathcal{B}_\mathbb{N}(f)(\sum n_i x_i) = \sum n_i f(x_i)$. Coalgebras for $\mathcal{B}_\mathbb{N}$ are directed graphs with \mathbb{N} -weighted edges, often referred to as *multigraphs* [6].

A separating set of predicate liftings λ^k , $k \in \mathbb{N}$, is defined by

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

The arising modal operators are exactly the modalities \Diamond_k of *graded modal logic* (cf. e.g. [6]), i.e. $x \vDash \Diamond_k \phi$ iff ϕ holds for more than k successor states of x , taking into account multiplicities. 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 $\Box(p \wedge q) \rightarrow \Box p$, and *normal* if it satisfies $\Box(p \rightarrow q) \rightarrow \Box p \rightarrow \Box q$).

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 finite maps $X \rightarrow \mathbb{Z}$, called *generalized multisets* (this set is also familiar as the free abelian group over X).

A separating set of predicate liftings λ^k , $k \in \mathbb{Z}$, with induced modal operators \Diamond_k is defined analogously as for multisets. Note that \Box_k fails to be monotone even for $k = 0$. One may imagine an interpretation of $\mathcal{B}_\mathbb{Z}$ -coalgebras

as transition systems that allow some form of trading — formulae may be violated in ‘positive’ successor states, as long as this is counterbalanced by violations in ‘negative’ successor states.

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* [8]) 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 require a more elaborate measure theoretic treatment [27].)

A separating set for T is obtained by combining the propositional symbols (cf. Example 2) with the 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 $\langle p \rangle = [\lambda^p]$ of *probabilistic modal logic (PML)* [13, 8], where $\langle p \rangle \phi$ reads ‘ ϕ holds in the next step with probability at least p ’. Note that $[p]$, defined as $\neg \langle p \rangle \neg$, is monotone, but not normal.

6. For a field k , the *linear space functor* $k \cdot _$ takes a set X to the free k -vector space $k \cdot X$, i.e. the set of formal k -linear combinations, over X . A coalgebra for $k \cdot _$ is a *linear automaton* [3, 26] (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 $k = \mathbb{R}$, a separating set of predicate liftings can be constructed in the same way as for D_ω , giving rise to modal operators $\langle p \rangle$ for $p \in \mathbb{Q}$. Here, $\langle p \rangle \phi$ 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 inputs from an alphabet I , i.e. by passing from T to one of the functors S and R given by $SX = I \rightarrow TX$ and $RX = T(I \times X)$, respectively. When I 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).

An expressive set of modal operators is then obtained by indexing modal operators over $a \in I$ in the form $[_]_a$. In the case $T = \mathcal{P}_\omega$, this leads to the usual operators \square of Hennessy-Milner Logic [9]. In the probabilistic case, the meaning of $\langle p \rangle_a \phi$ in reactive probabilistic automata is that on input a , ϕ holds in the next step with probability at least p , and in generative probabilistic automata that with probability at least p , the input is a and ϕ holds in the next step.

Remark 8. Graded modal logic is more standardly interpreted over Kripke models by just counting successor states (as e.g. in [25]), rather than in multigraphs as in the above example and e.g. in [6]. One can regard Kripke models as multigraph models by just regarding sets as multisets where all elements have

multiplicity 1; conversely, one can unroll a state in a multigraph model into a tree-like Kripke model by making copies of elements according to their multiplicity. Both constructions preserve satisfaction of graded modal formulae, so that the two semantics induce the same local consequence relations.

3 Proof Systems for Coalgebraic Modal Logic

We now discuss completeness of derivation for coalgebraic modal logic, partly following [16, 5, 11]. Since an otherwise unstructured 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 axiomatization of a coalgebraic modal logic for T it is enough to consider modal axioms of rank 1. The approach taken in [16, 5, 11] is based on this expectation; we shall prove below that it is indeed formally the case that axioms of rank 1 are sufficient. This fact will be crucial for our finite model result to be proved in Section 4.

To begin, we note that both the global and the local consequence relation (Definition 3) of a coalgebraic modal logic in general fail to be compact:

Example 9. In the case of Hennessy-Milner logic over finitely branching systems [9] with two inputs a, b , the set

$$\Phi = \{ \diamond (\Box^{n+1} \perp \wedge \diamond^n \top) \mid n \in \mathbb{N} \},$$

where \Box^n stands for n consecutive boxes, is locally (and hence globally) 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 globally (and hence 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 formulae over $\mathcal{B}_{\mathbb{N}}$. Non-compactness of PML is observed in [8]; in this case, non-compactness does *not* have to do with bounded branching.

Thus, in general neither the local nor the global consequence relations of a coalgebraic modal logic admit a finitary complete proof system. Instead, one is lead to study weak completeness, where a proof system is called *weakly complete* if it proves all valid formulae. This notion is equivalent to completeness in the sense used in [5, 11, 16], where only local consequence with singleton sets of premises is considered (ψ is a local consequence of $\{\phi\}$ iff $\phi \rightarrow \psi$ 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 [16, 5]:

Definition 10. Let V be a set. We denote the set of propositional formulae over V by $\text{Prop}(V)$. Here, we regard \neg and \wedge as the basic connectives, with all other connectives defined in the standard way. A *literal* over V is either an

element of V or the negation of such an element. A (*conjunctive*) *clause* is a finite, possibly empty, disjunction (conjunction) of literals. Moreover, we denote by $\text{Up}(V)$ the set $\{[\lambda]a \mid \lambda \in \Lambda, a \in V\}$.

If the elements of V are, or have an interpretation as, subsets of a given set X , then $\phi \in \text{Prop}(V)$ can be interpreted as a subset $\llbracket \phi \rrbracket_X$ of X ; we say that ϕ *holds in* X and write $\models_X \phi$ if $\llbracket \phi \rrbracket_X = X$, and we say that ϕ is *satisfiable in* X if $\llbracket \phi \rrbracket_X \neq \emptyset$. Similarly, if $a \in V$ is interpreted as a subset A of X , then we interpret $[\lambda]a \in \text{Up}(V)$ as the subset $\llbracket [\lambda]a \rrbracket = \lambda_X(A)$ of TX . (This can of course be iterated, leading to interpretations $\llbracket \phi \rrbracket \subset TX$ of $\phi \in \text{Up}(\text{Prop}(V))$ etc.)

In case the elements of V are formulae in $\mathcal{L}(\Lambda)$, we also regard propositional formulae over V as formulae in $\mathcal{L}(\Lambda)$. We sometimes explicitly designate V as consisting of *propositional variables*; propositional variables retain their status across further applications of Up and Prop (e.g. if V is a set of propositional variables, then V and not $\text{Prop}(V)$ is the set of propositional variables for $\text{Up}(\text{Prop}(V))$). Given a set L , an *L-substitution* is a substitution σ of the propositional variables by elements of L . Then, $\phi\sigma$ is called an *instance* of ϕ over L . If $L \subset \mathcal{P}(X)$ for some X , then we also refer to σ as an *L-valuation* or a *P(X)-valuation*.

The format we impose on axioms is essentially equivalent to the formal notion of *axiom* used in [16, 11]:

Definition 11. A *rank-1-clause* over a set V of propositional variables is a clause over $\text{Up}(\text{Prop}(V))$. Such a clause is *valid* if all its instances over $\mathcal{L}(\Lambda)$ are valid.

Proposition 12. *Let ϕ be a rank-1-clause. If $\phi\sigma$ holds in TX for every set X and every $\mathcal{P}(X)$ -valuation σ , then ϕ is valid. The converse holds if T is ω -accessible, Λ is separating, and the final T -coalgebra is infinite.*

(The condition on ϕ in the above proposition has been called *admissibility* in [16], where also the first implication is proved.)

A given set Ax of valid rank-1-clauses, called *axioms*, induces a proof system for $\mathcal{L}(\Lambda)$ as follows [11].

Definition 13. The set of formulae *derivable* from Ax is the smallest set containing all instances of axioms over $\mathcal{L}(\Lambda)$ and closed under propositional entailment and the *congruence rule*

$$\frac{\phi \leftrightarrow \psi}{[\lambda]\phi \leftrightarrow [\lambda]\psi}.$$

It is easy to see that this proof system is sound. The completeness results in [16, 5, 11] require the presence of ‘enough’ axioms in the following sense.

Definition 14. The set Ax is *reflexive* if, for each set X and each $\mathfrak{A} \subset \mathcal{P}(X)$, every clause ϕ over $\text{Up}(\mathfrak{A})$ that holds in TX is *derivable*, i.e. propositionally entailed by instances of axioms over \mathfrak{A} and by formulae of the form $[\lambda]\phi \leftrightarrow [\lambda]\psi$, where $\phi, \psi \in \text{Prop}(\mathfrak{A})$ and $\phi \leftrightarrow \psi$ holds in X .

Of course, we can restrict ourselves to finite \mathfrak{A} . The definition originally used in [16] to establish weak completeness is stronger than the notion above in that derivations of clauses over L are restricted to use only the subset relation on \mathfrak{A} rather than propositional formulae that hold in X . By the results of [5], the weaker definition above, which is essentially equivalent to one given in [11], suffices to establish weak completeness. Examples of reflexive axiomatizations are given in [16, 5].

The weak completeness theorem, stating that reflexive sets induce weakly complete proof systems [16, 5], will appear as a corollary to our finite model result in Section 4. We now proceed to establish that every functor indeed admits a reflexive set, i.e. that the set of all valid rank-1-clauses is reflexive.

As a preparation, we note that rank-1-clauses are equivalent to rules of a restricted format.

Definition 15. An (*extended*) *one-step rule* R over a set V of propositional variables is a rule ϕ/ψ , where $\phi \in \text{Prop}(V)$, and ψ is a clause over $\text{Up}(V)$ (over $\text{Up}(\text{Prop}(V))$). The rule R is *valid* if, whenever $\phi\sigma$ is valid for an $\mathcal{L}(A)$ -substitution σ , then $\psi\sigma$ is valid. As part of a proof system, R allows deriving $\psi\sigma$ from $\phi\sigma$ for each $\mathcal{L}(A)$ -substitution σ .

Thanks to the congruence rule, one-step rules can replace extended one-step rules (just introduce premises abbreviating propositional formulae as propositional variables). In particular, every rank-1-clause can be replaced by a one-step rule. Conversely, we have

Proposition 16. *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.*

The proof needs the following fact from propositional logic.

Lemma 17. *Let $\phi \in \text{Prop}(V)$ be satisfiable. Then there exists a $\text{Prop}(V)$ -substitution σ such that*

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

are tautologies.

Proof (Proposition 16). We can assume that the the premise ϕ of $R = \phi/\psi$ is satisfiable. Thus, fix σ as in the above lemma for ϕ . Then R and the rank-1-clause $\psi\sigma$ are mutually interderivable as claimed. \square

We are now ready to prove the announced axiomatizability result:

Theorem 18. *The set of all valid rank-1-clauses is reflexive.*

Proof. (Sketch) Let X be a set, let $\mathfrak{A} \subset \mathcal{P}(X)$ be finite, and let the clause ψ over $\text{Up}(\mathfrak{A})$ hold in TX . Let the formula ϕ be the ‘propositional theory’ of \mathfrak{A} , i.e. the (finite) conjunction of all clauses over \mathfrak{A} that hold in X . Then one can

show that the one-step rule $R \equiv \phi/\psi$ over \mathfrak{A} , abused as a set of propositional variables, is valid. By Proposition 16, it follows that R is derivable from the set of all valid rank-1-clauses; combined with the fact that ϕ holds in X , this yields a derivation of ψ over TX in the sense of Definition 14. \square

4 The Finite Model Construction

The non-compactness of coalgebraic modal logic (cf. Example 9) 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 consistent subsets of closed sublanguages.

We recall a few basic definitions:

Definition 19. Given a set Ax of axioms, a finite set $\{\phi_1, \dots, \phi_n\}$ of formulae is called *consistent* (w.r.t Ax) if $\neg(\phi_1 \wedge \dots \wedge \phi_n)$ is not derivable according to Definition 13. A set Σ of formulae is called *closed* if it is closed under subformulae and under *normalized 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 $\perp \notin A$ and, for $\phi, \psi \in \Sigma$, $\phi, \psi \in A$ iff $\phi \wedge \psi \in A$, 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 consistent subsets of Σ .

Thus, a Σ -atom is just a consistent Σ -Hintikka set.

Lemma 20 (Lindenbaum Lemma). *Every consistent subset of Σ is contained in a Σ -atom.*

Given a closed set Σ , the carrier of the model to be constructed will be the set S of Σ -atoms; the main problem is then to define the coalgebra structure on S . The following lemma is crucial for this purpose.

Lemma 21. *Let V be a set of propositional variables, let $\phi \in \text{Prop}(V)$, and let σ be a Σ -substitution. Then $\phi\sigma$ is derivable iff $\phi\tau$ holds in the set S of Σ -atoms, where τ is the $\mathcal{P}(S)$ -valuation given by $\tau(a) = \{A \in S \mid \sigma(a) \in A\}$.*

Expecting 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

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

for all $A \in S$ and all formulae $[\lambda]\phi$ in Σ . This is where reflexivity comes in:

Lemma 22 (Existence Lemma). *If Ax is reflexive and Σ is finite, then $\xi(A)$ satisfying (*) exists for each $A \in S$.*

Proof. Assume that $\xi(A)$ does not exist. Let V be the set $\{a_\phi \mid \phi \in \Sigma\}$ of propositional variables. Let ψ be the clause over $\text{Up}(V)$ 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 $\psi\tau$ holds in TS by assumption. By

reflexivity, $\psi\tau$ is derivable in the sense of Definition 14 from the propositional formulae of the form $\chi\tau$ that hold in S . This derivation can be copied to obtain a derivation of $\psi\sigma$ from those $\chi\sigma$ for which $\chi\tau$ holds in S , where σ is the Σ -substitution taking a_ϕ to ϕ . These $\chi\sigma$ are derivable by Lemma 21. Thus, $\psi\sigma$ is derivable, in contradiction to the consistency of A . \square

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 5:

Lemma 23 (Truth Lemma). *Let Σ be a closed set, let S be a set of Σ -Hintikka sets, and let $\xi : S \rightarrow TS$ satisfy condition (*) above. Then for all $\phi \in \Sigma$ and all $A \in \Sigma$,*

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

Proof. Straightforward induction over ϕ . \square

This is all we need in order to establish

Theorem 24. *Let Ax be reflexive. Then every formula ϕ that is consistent w.r.t. Ax is locally satisfiable in a finite T -coalgebra of size at most $2^{|\phi|}$.*

The proof is just an application of the lemmas above to the (finite) smallest closed set $\Sigma(\phi)$ containing a given consistent formula ϕ .

As announced, the above result implies weak completeness [16, 5]; explicitly:

Corollary 25 (Weak completeness). *The proof system induced by a reflexive set of axioms is weakly complete.*

Combining Theorems 18 and 24, we obtain independently of deduction:

Corollary 26 (Finite model property). *Every locally satisfiable formula ϕ is locally satisfiable in a finite T -coalgebra of size at most $2^{|\phi|}$.*

Remark 27. The finite model property does not generalize to the case where global axioms are present, i.e. it may be the case that ϕ is locally satisfiable in some model globally satisfying a formula ψ , but not in any finite such model. Examples are found already in standard modal logic [18].

5 Decidability

Unlike in the classical case, the finite model construction of the preceding section does not immediately imply decidability, even though it gives a computable bound on the size of the model, since 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 [16].) 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 formulae is r.e.. We then obtain decidability provided that the set of valid formulae is also r.e., which by the weak completeness theorem and Theorem 18 is the case if the set of all valid rank-1-clauses is r.e.

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 the set H of $\Sigma(\phi)$ -Hintikka sets. We are then faced with the following decision problem:

Definition 28. The *one-step satisfiability* problem is to decide, given a finite set X and a *conjunctive* clause ϕ over $\mathbf{Up}(\mathcal{P}(X))$, whether ϕ is satisfiable in TX .

Remark 29. For purposes of determining the input size for the above problem, we assume that subsets of X can be represented in $|X|$ bits. Moreover, we assume that Λ is countable, with a reasonable size measure for the representation of modal operators supposed as given (e.g., size $\log k$ for the graded modal operator $[k]$, and size $\log n + \log m$ for the probabilistic modal operator $[n/m]$, with $n, m \in \mathbb{N}$ relatively prime). The same applies to the size of a formula ϕ as input for the satisfiability problem, which is thus larger than the size $|\phi|$ as defined in Section 2.

A decision procedure for one-step satisfiability leads to the following algorithm for satisfiability of ϕ .

Algorithm 1. 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

$$\bigwedge_{[\lambda]\psi \in A} [\lambda]\{B \in S \mid \psi \in B\} \wedge \bigwedge_{\neg[\lambda]\psi \in A} \neg[\lambda]\{B \in S \mid \psi \in B\}$$

is satisfiable in TS . 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’.

The correctness of the algorithm is guaranteed by the Truth Lemma. Thus,

Theorem 30. *If one-step satisfiability is decidable, then satisfiability of $\mathcal{L}(\Lambda)$ -formulae is decidable.*

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

Algorithm 2. Nondeterministically choose $S \subset H$; 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. Since acceptance sets of non-deterministic algorithms are r.e., we thus have

Theorem 31. *If one-step satisfiability is semi-decidable, then satisfiability of $\mathcal{L}(\Lambda)$ -formulae is semi-decidable.*

(Note that semi-decidability of one-step satisfiability is weaker than the above-mentioned condition that T takes finite sets to r.e. sets. E.g., the one-step satisfiability problem will turn out to be decidable for Λ as in Example 7.5, although $D_\omega(X)$ is uncountable for $|X| \geq 2$.)

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

Theorem 32. *If the one-step satisfiability problem is in NP, then satisfiability of $\mathcal{L}(\Lambda)$ -formulae is in NEXPTIME.*

Remark 33. 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 $S \circ T$: here, one has to do with conjunctive clauses over $\text{Up}(S(TX))$, where the application of T may produce an exponential blowup or lead to infinite sets.

Besides the examples whose decidability is already captured by the finite model result of [16], i.e. functors preserving finite sets, such as \mathcal{P} , our results cover the following cases.

- Example 34.**
1. Let Λ be the set of predicate liftings λ_k for the multiset functor of Example 7.3. Then the one-step satisfiability problem amounts to deciding the solvability of systems of linear inequations over the naturals; this problem is in NP [22]. By Theorem 32, we obtain that *the satisfiability problem for graded modal logic is in NEXPTIME*. (In fact, this problem is in PSPACE [25].)
 2. By the same line of reasoning, the satisfiability problem for generalized graded modal logic over coalgebras for the generalized multiset functor (Example 7.4) is in NEXPTIME.
 3. Let Λ be the set of predicate liftings for PML as in Example 7.5. Then the one-step satisfiability problem amounts to the solvability of systems of rational linear inequations over the reals, which is decidable in polynomial time by standard linear programming methods (using Motzkin's transposition theorem to treat also strict inequalities) [22]. By Theorem 30, it follows that *probabilistic modal logic is in NEXPTIME*.
 4. By the same reasoning, the modal logic for linear automata of Example 7.6 is decidable in NEXPTIME.
 5. It is straightforward to extend the above results to include proposition symbols, where not already present, or inputs (cf. Examples 7.2 and 7.7).

Remark 35. A decision algorithm for PML is announced in [7], but not explicitly contained in the full version [8]. The latter proves the finite model property for PML, from which decidability does follow by the argument sketched at the beginning of this section, with some additional work required to reduce to models

with rational probabilities in order to ensure recursive enumerability of the set of finite models. Our algorithm improves this result by giving an upper bound on the complexity, albeit a rather generous one (see below).

Remark 36. It is, of course, desirable to obtain better general complexity bounds; this is the subject of ongoing research. The best general bound we can hope for is *PSPACE*, since the decision problem for K , i.e. the modal logic of \mathcal{P} , is known to be *PSPACE*-complete [2]. The following further results are forthcoming [24]:

- (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 $\mathcal{L}(\Lambda)$ -formulae is in *EXPTIME* if one-step satisfiability is in P .
- (ii) Given a tractable axiomatization of $\mathcal{L}(\Lambda)$, one can show that satisfiability of $\mathcal{L}(\Lambda)$ -formulae is in *PSPACE* by means of a shallow model construction.

(Neither of these results makes Theorem 32 obsolete, since both rely on stronger assumptions.) By (i), one immediately improves the bound for PML as well as for the modal logic of linear automata from *NEXPTIME* (Example 34) to *EXPTIME*. The method of (ii) reproduces the known *PSPACE* bounds for K and for graded modal logic, and very likely leads to novel *PSPACE* bounds for all other logics mentioned in Example 34, including PML. An open problem that remains is whether there is a semantic criterion (not involving deduction) that guarantees a *PSPACE* bound.

We conclude this section with a few remarks on axiomatizability.

Definition 37. The *one-step validity* problem is to decide, given a finite set X and a (disjunctive) clause ϕ over $\text{Up}(\mathcal{P}(X))$, whether ϕ holds in TX .

Of course, one-step validity and one-step satisfiability are, via negation, reducible to each other's complements.

Proposition 38. *Let T be ω -accessible, let Λ be separating, and let the final T -coalgebra be infinite. Then a rank-1-clause ϕ over a set V of propositional variables is valid iff $\phi\sigma$ holds in $T(\mathcal{P}(V))$, where σ is the $\mathcal{P}(\mathcal{P}(V))$ -valuation taking a $a \in V$ to the set*

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

Corollary 39. *Let T be ω -accessible, let Λ be separating, and let the final T -coalgebra be infinite. Then the validity of rank-1-clauses is decidable (semi-decidable) if the one-step validity problem is decidable (semi-decidable).*

We have seen in Example 34 that the one-step validity problem is decidable in many important cases. Thus, Theorem 18 does frequently supply a feasible axiomatization of $\mathcal{L}(\Lambda)$, although one will, of course, in general strive for a more compact, preferably finite axiomatization.

6 Conclusion

Coalgebraic modal logic in general fails to be compact, so that completeness results are necessarily restricted to weak completeness and moreover cannot rely on constructing full canonical models. Above, we have described a finite model construction for coalgebraic modal logic, using the ‘small canonical model’ method known from standard modal logic; weak completeness is a corollary to this result. Here, the notion of reflexive axiom sets, which has appeared as a prerequisite for existing weak completeness results for coalgebraic modal logic [16, 5, 11], has played a crucial role. In particular, we have proved that every coalgebraic modal logic admits a reflexive axiomatization by axioms of rank 1; this not only means that the mentioned completeness results are, in principle, always applicable, but also implies a finite model property which states that all satisfiable formulae can be satisfied in a finite model whose size is exponentially bounded by the size of the formula.

We have then described a generic decision procedure for satisfiability in coalgebraic modal logic, 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 [16], 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 complexity bound (*NP*) for one-step satisfiability, we have established a first general complexity bound for coalgebraic modal logic (*NEXPTIME*). This result applies to both graded and probabilistic modal logic; while for graded modal logic, a better bound (*PSPACE*) is known, no complexity bound at all has, to our knowledge, so far been given for probabilistic modal logic. Forthcoming work [24] will establish, under additional assumptions, tighter generic bounds which in particular push the bound for PML at least to *EXPTIME* and likely to *PSPACE*.

Acknowledgements

The author wishes to thank Till Mossakowski, Markus Roggenbach, and Horst Reichel for collaboration on CoCASL, Erwin R. Catesbeiana for favoring empty coalgebras, and Dirk Pattinson for useful discussions.

References

- [1] F. Bartels, A. Sokolova, and E. de Vink. A hierarchy of probabilistic system types. In *Coalgebraic Methods in Computer Science*, volume 82 of *ENTCS*. Elsevier, 2003.
- [2] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge, 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, 1980.
- [5] C. Cirstea and D. Pattinson. Modular construction of modal logics. In *Concurrency Theory*, volume 3170 of *LNCS*, 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] A. Heifeitz and P. Mongin. The modal logic of probability. In *Theoretical Aspects of Rationality and Knowledge*, pages 175–186. Morgan Kaufmann, 1998.
- [8] A. Heifeitz and P. Mongin. Probabilistic logic for type spaces. *Games and Economic Behavior*, 35:31–53, 2001.
- [9] M. Hennessy and R. Milner. Algebraic laws for non-determinism and concurrency. *J. ACM*, 32:137–161, 1985.
- [10] B. Jacobs. Towards a duality result in the modal logic of coalgebras. In *Coalgebraic Methods in Computer Science*, volume 33 of *ENTCS*. Elsevier, 2000.
- [11] C. Kupke, A. Kurz, and D. Pattinson. Algebraic semantics for coalgebraic logics. In *Coalgebraic Methods in Computer Science*, volume 106 of *ENTCS*, pages 219–241. Elsevier, 2004.
- [12] A. Kurz. Specifying coalgebras with modal logic. *Theoret. Comput. Sci.*, 260:119–138, 2001.
- [13] K. Larsen and A. Skou. Bisimulation through probabilistic testing. *Inform. Comput.*, 94:1–28, 1991.
- [14] T. Mossakowski, L. Schröder, M. Roggenbach, and H. Reichel. Algebraic-coalgebraic specification in CoCASL. *J. Logic Algebraic Programming*. To appear.
- [15] D. Pattinson. Semantical principles in the modal logic of coalgebras. In *Symposium on Theoretical Aspects of Computer Science*, volume 2010 of *LNCS*, pages 514–526. Springer, 2001.
- [16] D. Pattinson. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoret. Comput. Sci.*, 309:177–193, 2003.
- [17] D. Pattinson. Expressive logics for coalgebras via terminal sequence induction. *Notre Dame J. Formal Logic*, 45:19–33, 2004.
- [18] S. Popkorn. *First Steps in Modal Logic*. Cambridge, 1994.
- [19] M. Röbiger. Coalgebras and modal logic. In *Coalgebraic Methods in Computer Science*, volume 33 of *ENTCS*. Elsevier, 2000.
- [20] J. Rothe, H. Tews, and B. Jacobs. The Coalgebraic Class Specification Language CCSL. *J. Universal Comput. Sci.*, 7:175–193, 2001.
- [21] J. Rutten. Universal coalgebra: A theory of systems. *Theoret. Comput. Sci.*, 249:3–80, 2000.
- [22] A. Schrijver. *Theory of linear and integer programming*. Wiley Interscience, 1986.
- [23] L. Schröder. Expressivity of coalgebraic modal logic: the limits and beyond. In *Foundations of Software Science And Computation Structures*, volume 3441 of *LNCS*, pages 440–454. Springer, 2005.
- [24] L. Schröder and D. Pattinson. *PSPACE* reasoning for coalgebraic modal logic. In preparation.
- [25] S. Tobies. *PSPACE* reasoning for graded modal logics. *J. Logic Computation*, 11:85–106, 2001.
- [26] P. Turakainen. On the minimization of linear space automata. *Ann. Acad. Sci. Fenn. Ser. A I*, 506, 1972. 15 pp.
- [27] I. Viglizzo. Final sequences and final coalgebras for measurable spaces. In *Algebra and Coalgebra in Computer Science*, volume 3629 of *LNCS*, pages 395–407. Springer, 2005.