Completeness of Flat Coalgebraic Fixpoint Logics

LUTZ SCHRÖDER, Friedrich-Alexander-Universität Erlangen-Nürnberg
YDE VENEMA, ILLC, Universiteit van Amsterdam

Modal fixpoint logics traditionally play a central role in computer science, in particular in artificial intelligence and concurrency. The μ-calculus and its relatives are among the most expressive logics of this type. However, popular fixpoint logics tend to trade expressivity for simplicity and readability and in fact often live within the single variable fragment of the μ-calculus. The family of such *flat* fixpoint logics includes, e.g., Linear Temporal Logic (LTL), Computation Tree Logic (CTL), and the logic of common knowledge. Extending this notion to the generic semantic framework of *coalgebraic logic* enables covering a wide range of logics beyond the standard μ-calculus including, e.g., flat fragments of the graded μ-calculus and the alternating-time μ-calculus (such as alternating-time temporal logic), as well as probabilistic and monotone fixpoint logics. We give a generic proof of completeness of the Kozen-Park axiomatization for such flat coalgebraic fixpoint logics.

CCS Concepts: • **Theory of computation → Modal and temporal logics**; *Proof theory*; *Description logics*;

Additional Key Words and Phrases: Completeness, Kozen/Park axioms, branching-time temporal logics, coalgebraic logic, alternating-time temporal logic, graded μ-calculus, algebraic semantics

ACM Reference format:
Lutz Schröder and Yde Venema. 2018. Completeness of Flat Coalgebraic Fixpoint Logics. *ACM Trans. Comput. Logic* 19, 1, Article 4 (January 2018), 34 pages.
https://doi.org/10.1145/3157055

4

1 INTRODUCTION

Many of the most well-known logics in program verification, concurrency, and other areas of computer science and artificial intelligence can be cast as modal fixpoint logics, that is, embedded into some variant of the μ-calculus. Typical examples are Propositional Dynamic Logic (PDL) (Pratt 1976) where, say, the formula $\langle a^* \rangle p$ ("p can be reached by finite iteration of a") can be expressed as the least fixpoint

$$\mu X.\, p \vee \langle a \rangle X;$$

Computation Tree Logic (CTL) (Emerson and Clarke 1982), whose formula AFp ("p eventually holds on all paths") is just the fixpoint

$$\mu X.\, p \vee \Box X;$$

Work by the first author forms part of the DFG project *Generic Algorithmic Methods in Modal and Hybrid Logics (GenMod3, SCHR 1118/5-3).*

Authors' addresses: L. Schröder, Chair of Theoretical Computer Science, Friedrich-Alexander-Universität Erlangen-Nürnberg, Martensstr. 3, 91058 Erlangen, Germany; email: lutz.schroeder@fau.de; Y. Venema, Institute for Logic, Language and Computation, University of Amsterdam, P.O. Box 94242, 1090 GE Amsterdam, The Netherlands; email: y.venema@uva.nl.

https://doi.org/10.1145/3157055

and the common knowledge operator C of epistemic logic (Lewis 1969), where Cp ("it is common knowledge that p") can be expressed as the fixpoint

$$\nu X.\, (p \wedge \bigwedge_{i=1}^{n} K_i X),$$

with n the number of agents and K_i read as "agent i knows that." A common feature of these examples is that they trade off expressivity for simplicity of expression in comparison to the full μ-calculus.

One of the reasons why the full μ-calculus is hard for both end users and logicians is that it requires keeping track of bound variables. Indeed, we note that the simpler logics listed above (in the case of PDL, the $*$-nesting-free fragment) live in the single-variable fragment of the μ-calculus (a subfragment of the alternation-free fragment (Emerson and Lei 1986)), which is precisely what enables one to abandon variables altogether in favour of variable-free fixpoint operators such as AF or C. We refer to logics that embed into a single-variable μ-calculus as *flat fixpoint logics* (Santocanale and Venema 2010).

Here, we study flat fixpoint logics in the more general setting of *coalgebraic logic*. Coalgebra has emerged as the right level of generality for a unified treatment of a wide range of modalities with seemingly disparate semantics beyond the realm of pure relational structures. Examples include monotone modalities (Chellas 1980), probabilistic modalities (Larsen and Skou 1991), graded modalities (Fine 1972; D'Agostino and Visser 2002), coalitional/alternating-time modalities (Alur et al. 2002; Pauly 2002), and various non-monotonic conditionals (Friedman and Halpern 1994; Olivetti et al. 2007). The semantics of coalgebraic logic is parametrized over the choice of an endofunctor on the category of sets, whose coalgebras play the role of frames. Besides standard Kripke frames, the notion of coalgebra encompasses, e.g., Markov chains, weighted automata, multigraphs, neighbourhood frames, selection function frames (Chellas 1980), and concurrent game structures (Alur et al. 2002). Generic completeness results in coalgebraic logic are parametrized over sets of rules or axioms that satisfy a local form of completeness called *one-step completeness*. That is, they require completeness of a restricted logic without fixpoints and nesting of modalities that is interpreted over the successor structure of a single (implicit) state rather than full-blown frames or models (e.g., in the relational base case, over a single subset of the base set thought of as a local view on a set of successors), a condition that is typically quite easy to establish. In fact, suitable one-step complete axiomatizations for many examples can already be found in the literature on coalgebraic logic (Pattinson 2003; Schröder 2007; Schröder and Pattinson 2009; Kupke and Pattinson 2010).

In our *flat coalgebraic fixpoint logics* one thus can express statements such as "the coalition C of agents can maintain p forever," "the present state is the root of a binary tree all whose leaves satisfy p," or "p is commonly believed with reasonable certainty." In particular, we cover flat fragments of the graded μ-calculus (Kupferman et al. 2002) and the alternating-time μ-calculus (AMC) (Alur et al. 2002); one such flat fragment is alternating-time temporal logic (ATL).

Our main result on flat coalgebraic fixpoint logics is completeness of the natural axiomatization that makes the fixpoint definitions explicit, generalizing the well-known Kozen-Park axiomatization. The axiomatization is parametric both w.r.t. the underlying (coalgebraic) system type and the choice of flat fragment, under mild restrictions on the form of fixpoint operators. This result generalizes results by Santocanale and Venema (2010) to the level of coalgebraic logic and relies on the notion of *$\mathcal{O}$-adjointness* (Santocanale 2008) to prove that fixpoints in the Lindenbaum algebra are *constructive*, i.e., approximable in ω steps. The crucial ingredient here are *one-step complete* rule sets (Schröder 2007). These enable generalizations of both the key *rigidity lemma* and the $\mathcal{O}$-adjointness theorem of Santocanale and Venema (2010), the latter to the effect that *all uniform-depth modal operators are $\mathcal{O}$-adjoint.*

Our completeness result follows a long tradition of non-trivial completeness proofs for fixpoint logics, e.g., PDL (Kozen and Parikh 1981; Segerberg 1982), CTL (Emerson and Halpern 1985), Linear Temporal Logic (LTL) (Gabbay et al. 1980; Lichtenstein and Pnueli 2000), the aconjunctive μ-calculus (Kozen 1983), and the full μ-calculus (Walukiewicz 2000). Note that all these results are independent, as completeness is not in general inherited by sublogics and in fact employs quite different methods. Instantiating our generic results to concrete logics yields new results in nearly all cases that go beyond the classical relational μ-calculus, noting that neither Kupferman et al. (2002) nor Cîrstea et al. (2011) cover axiomatizations. In particular, we obtain for the first time a completeness result for graded fixpoint logics, i.e., fragments of the graded μ-calculus, and we generalize the completeness of ATL (Goranko and van Drimmelen 2006) to arbitrary flat fragments of the AMC.

Further Related Work. The present article is an extended and revised version of a previous conference paper (Schröder and Venema 2010). The technical treatment differs from the conference version in that we weaken the assumptions of the generic completeness theorem to require only a one-step complete rule set instead of a one-step cutfree complete one; moreover, we opt for a propositional basis with unrestricted negation, converting to negation normal form only for purposes of the model construction in Section 5. Flat coalgebraic fixpoint logics are fragments of coalgebraic μ-calculi and as such known to be decidable in ExpTime under reasonable assumptions (Cîrstea et al. 2011). A tableau-based global caching algorithm for flat fixpoint logics has recently been developed by Hausmann and one of the authors (Schröder) (2015), so we omit discussion of the (less practical) tableau algorithm given in the conference version.

Organization. We recall the requisite background in coalgebraic logic and introduce the syntax and semantics of flat coalgebraic fixpoint logics in Section 2. We proceed to discuss the generic axiomatization in Section 3. In Section 4, we prove the central $\mathcal{O}$-adjointness theorem and then present the ensuing model construction in Section 5. Section 6 concludes.

2 FLAT COALGEBRAIC FIXPOINT LOGICS

We briefly recall the generic framework of coalgebraic modal logic (Pattinson 2004; Schröder 2007) and define its extension with flat fixpoint operators, a fragment of the coalgebraic μ-calculus (Cîrstea et al. 2011). We present the calculus in a form that includes negation and therefore need to pay attention to positive and negative occurrences of variables and subformulas; to avoid excessive repetition, we fix these notions for all notions of formula that include negation and possibly a notion of (propositional) variable.

Definition 2.1. In any logic with negation $\neg$, an occurrence of a subformula in a formula is *positive* if it is in the scope of an even number of negations $\neg$ and otherwise *negative*. In logics featuring a notion of (propositional) *variable*, a formula ϕ is *positive* (*negative*) in a variable v if all occurrences of v in ϕ are positive (negative).

For the rest of the article, we fix a countably infinite set V of *parameter* or *propositional variables*, typically called p or $p_1, p_2, \ldots$, and additionally a single distinguished *recursion variable* $x \notin V$. The parameter variables will serve as placeholders for formulas in fixpoint schemes and in proof rules. We use the word *variable* to refer to both the parameter variables and the recursion variable.

The first parameter of the syntax of a *flat coalgebraic fixpoint logic* is a *(modal) similarity type* Λ, i.e., an at most countable set of modal operators with associated finite arities. The set of *modal fixpoint schemes* γ, δ is given by the grammar

$$\gamma, \delta ::= \bot \mid x \mid p \mid \neg\gamma \mid \gamma \wedge \delta \mid \heartsuit(\gamma_1, \ldots, \gamma_n),$$

where $\heartsuit \in \Lambda$ is n-ary and $p \in V$ (and x is the fixed recursion variable); we require additionally that fixpoint schemes are positive in all variables (this is essential in case of the recursion variable x to ensure existence of fixpoints; for parameter variables, it is a mere technical convenience, as negative occurrences of a parameter variable can be replaced with positive occurrences of a fresh variable, with negation moved into the parameter formula). Further Boolean operations $\top$, $\vee$, $\to$, $\leftrightarrow$ are defined as usual. Moreover, we abbreviate

$$\overline{\heartsuit}(\phi_1, \ldots, \phi_n) = \neg\heartsuit(\neg\phi_1, \ldots, \neg\phi_n)$$

for $\heartsuit \in \Lambda$ and refer to $\overline{\heartsuit}$ as the *dual* of $\heartsuit$. We intend variables as place holders for arguments and parameters of formulas defining fixpoint operators; as such, they serve only technical purposes and will not form part of the actual fixpoint language defined below (other than as part of modal fixpoint schemes indexing flat fixpoint operators). In particular, variables should not be confused with propositional *atoms*, which can appear also in actual formulas. Propositional atoms are incorporated into the modal similarity type Λ as nullary operators when needed (Example 2.3(1)); this approach not only simplifies the technical presentation, but it also enhances generality in that it allows covering logics that do not have propositional atoms, such as Hennessy-Milner logic.

The second syntactic parameter is a set Γ of modal fixpoint schemes γ determining the choice of fixpoint operators. We require that all $\gamma \in \Gamma$ are *guarded*, i.e., that all occurrences of the recursion variable x are under the scope of at least one modal operator; this is not an essential restriction as every μ-calculus formula is provably equivalent to a guarded formula (Walukiewicz 2000). We denote substituted formulas $\gamma[\phi_1/p_1; \ldots; \phi_n/p_n; \psi/x]$ as $\gamma(\phi_1, \ldots, \phi_n, \psi)$. The set $\mathcal{F}_\sharp(\Lambda, \Gamma)$ or just $\mathcal{F}_\sharp$ of *(fixpoint) formulas* ϕ, ψ is then defined by the grammar

$$\phi, \psi ::= \bot \mid \neg\phi \mid \phi \wedge \psi \mid \heartsuit(\phi_1, \ldots, \phi_n) \mid \sharp_\gamma(\phi_1, \ldots, \phi_n),$$

where $\heartsuit \in \Lambda$ is n-ary and $\gamma \in \Gamma$. The intended meaning of the connective $\sharp_\gamma$ is to take least fixpoints of the parametrized operation interpreting γ:

$$\sharp_\gamma(\phi_1, \ldots, \phi_n) = \mu x.\gamma(\phi_1, \ldots, \phi_n, x)$$

(we will give a formal semantics later in this section). The name *flat* for the fixpoint operators $\sharp_\gamma$ relates to the fact that modal fixpoint schemes γ do not contain fixpoint operators. Note, however, that nesting of flat fixpoint operators is unrestricted, i.e., the ϕ_i can be arbitrary fixpoint formulas in $\sharp_\gamma(\phi_1, \ldots, \phi_n)$. We introduce greatest fixpoint operators as duals of least-fixpoint operators: For a modal fixpoint scheme γ, we denote by $\overline{\gamma}$ its *dual*, i.e., the modal fixpoint scheme $\neg\gamma\sigma$, where $\sigma(v) = \neg v$ for all variables v. We then define the greatest fixpoint operator $\flat_{\overline{\gamma}}$ by

$$\flat_{\overline{\gamma}}(\phi_1, \ldots, \phi_n) = \neg\sharp_\gamma(\neg\phi_1, \ldots, \neg\phi_n)$$

so that $\flat_{\overline{\gamma}}(\phi_1, \ldots, \phi_n) = \nu x.\, \overline{\gamma}(\phi_1, \ldots, \phi_n, x)$ (again, pending the introduction of the formal semantics).

A standard example of a flat fixpoint logic is CTL, whose operators AU, EG can be coded as

$$A[\phi\, U\, \psi] = \sharp_{(p_2 \vee (p_1 \wedge \Box x))}(\phi, \psi) \quad \text{and} \quad EG\,\phi = \flat_{p \wedge \Diamond x}\phi.$$

Note here that $p \wedge \Diamond x$ is equivalent to the dual $\neg(\neg p \vee \Box\neg x)$ of $p \vee \Box x$.

Syntactically, $\sharp_\gamma$ is regarded as a primitive operator that just happens to have a long name; in particular, occurrences of variables in γ do not count as occurrences in formulas $\sharp_\gamma\phi$. For the sake of readability, we restrict the further technical development to unary modalities $\heartsuit \in \Lambda$ and unary fixpoint operators, i.e., we assume that every $\gamma \in \Gamma$ has only one parameter variable, *denoted by p throughout*; the extension to higher arities is a mere notational issue, and in fact we continue to use higher-arity modalities and fixpoint operators in the examples.

We have not included variables in the definition of fixpoint formulas. A *(fixpoint) formula with variables* is an expression of the form $\gamma\sigma$, where γ is a modal fixpoint scheme and σ is a substitution of some of the variables in γ with fixpoint formulas (i.e., variables never appear under fixpoint operators). In the following, the term *formula* will refer to fixpoint formulas without variables unless variables are explicitly mentioned. We sometimes indicate the occurrence of a variable v in a formula ψ by writing $\psi(v)$, and then write $\psi(\rho)$ for the formula obtained by substituting a formula ρ for v. For a modal fixpoint scheme $\gamma(p, x)$, we denote the function taking a formula ψ to $\gamma(\phi, \psi)$ by $\gamma(\phi)$ and by $\gamma(\phi)^k$ its k-fold iteration.

The logic is further parametrized *semantically* over the underlying class of systems and the interpretation of the modal operators. The former is determined by the choice of a functor

$$T : \mathsf{Set} \to \mathsf{Set},$$

i.e., an operation T that maps sets X to sets TX and functions $f : X \to Y$ to functions $Tf : TX \to TY$, preserving identities and composition. The role of models is then played by T*-coalgebras*, i.e., pairs (X, ξ) where X is a set of *states* and

$$\xi : X \to TX$$

is the *structure map*. Thinking of TX informally as a parametrized datatype over X, we regard ξ as associating with each state x a structured collection $\xi(x)$ of successor states and observations. For example, for $TX = \mathcal{P}X \times \mathcal{P}U$, with U a fixed set of propositional atoms and $\mathcal{P}$ denoting the covariant powerset functor (with $\mathcal{P}f(A) = f[A]$ for $f : X \to Y$ and $A \in \mathcal{P}(X)$), we obtain that T-coalgebras are Kripke models, as they associate with each state a set of successor states and a set of valid propositional atoms. Our main interest here is in examples beyond Kripke semantics, see Example 2.3.

Given T, the interpretation of the modalities is determined by associating with each $\heartsuit \in \Lambda$ a predicate lifting $[\![\heartsuit]\!]$ for T. Here, a *predicate lifting* (for T) is a family of maps $\lambda_X : \mathcal{P}X \to \mathcal{P}TX$, where X ranges over all sets, satisfying the *naturality* condition

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

for all $f : X \to Y$, $A \in \mathcal{P}Y$. In other words, a predicate lifting is a natural transformation

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

where $Q : \mathsf{Set}^{op} \to \mathsf{Set}$ denotes the contravariant powerset functor, given by QX being the powerset of X and $Qf(A) = f^{-1}[A]$ for $f : X \to Y$ and $A \subseteq Y$. The idea is that a predicate lifting λ_X converts a predicate on the set X of states into a predicate on the set TX of structured collections over X. The basic example, for $TX = \mathcal{P}(X) \times \mathcal{P}(U)$ as above, is

$$[\![\Box]\!]_X(A) = \{(B, Q) \in \mathcal{P}(X) \times \mathcal{P}(U) \mid B \subseteq A\},$$

which induces precisely the usual semantics of the box when composed with taking preimages under the structure map, as in the clause for the semantics of modalities given below. Given a predicate lifting $[\![\heartsuit]\!]$, we also have a predicate lifting $[\![\overline{\heartsuit}]\!]$ for the dual modality $\overline{\heartsuit}$, given by

$$[\![\overline{\heartsuit}]\!]_X(A) = X - [\![\heartsuit]\!]_X(X - A).$$

For example, for $\Box$ as above and $\Diamond = \overline{\Box}$, we have $[\![\Diamond]\!]_X(A) = \{(B, Q) \in \mathcal{P}(X) \times \mathcal{P}(U) \mid B \cap A \neq \emptyset\}$.

As we work with fixpoints, we insist that *all modal operators are monotone*, i.e., $[\![\heartsuit]\!]_X : \mathcal{P}(X) \to \mathcal{P}(TX)$ is monotone w.r.t. set inclusion for each $\heartsuit \in \Lambda$. The semantics of a formula ϕ with recursion variable x (no other variables will ever be evaluated in unsubstituted form) is a subset $[\![\phi]\!]_{(X,\xi)}(B) \subseteq X$, depending on a T-coalgebra (X, ξ) and a set $B \subseteq X$ serving as the interpretation of x. The semantics of formulas ϕ without variables (in particular of $\sharp$-formulas) will not depend on B and hence will be denoted just by $[\![\phi]\!]_{(X,\xi)}$; given $s \in X$, we write $s \models_{(X,\xi)} \phi$ for $s \in [\![\phi]\!]_{(X,\xi)}$. The set $[\![\phi]\!]_{(X,\xi)}(B)$ is defined by recursion over ϕ:

$$\begin{aligned}
[\![x]\!]_{(X,\xi)}(B) &= B \\
[\![\neg\phi]\!]_{(X,\xi)}(B) &= X \setminus [\![\phi]\!]_{(X,\xi)}(B) \\
[\![\phi \wedge \psi]\!]_{(X,\xi)}(B) &= [\![\phi]\!]_{(X,\xi)}(B) \cap [\![\psi]\!]_{(X,\xi)}(B) \\
[\![\heartsuit\phi]\!]_{(X,\xi)}(B) &= \xi^{-1}[\![\heartsuit]\!]_X([\![\phi]\!]_{(X,\xi)}(B)) \\
[\![\sharp_\gamma\phi]\!]_{(X,\xi)} &= \bigcap\{B \subseteq X \mid [\![\gamma(\phi)]\!]_{(X,\xi)}(B) \subseteq B\}.
\end{aligned}$$

The clause for $\sharp_\gamma\phi$ just says that $[\![\sharp_\gamma\phi]\!]_{(X,\xi)}$ is the least fixpoint of the map $[\![\gamma(\phi)]\!]_{(X,\xi)} : \mathcal{P}(X) \to \mathcal{P}(X)$, which is monotone, because all modalities are monotone and modal fixpoint schemes are positive in the recursion variable. *We fix the data T, Λ, Γ, and so on, throughout.* A formula $\phi \in \mathcal{F}_\sharp$ is *valid* if $[\![\phi]\!]_{(X,\xi)} = X$ for every T-coalgebra (X,ξ), and *satisfiable* if $\neg\phi$ is not valid, i.e., if $[\![\phi]\!]_{(X,\xi)} \neq \emptyset$ for some T-coalgebra (X,ξ).

We recall that given T-coalgebras (X,ξ) and (Y,ζ), a *T-coalgebra morphism $f : (X,\xi) \to (Y,\zeta)$* is a map $f : X \to Y$ making the diagram

$$\begin{tikzcd}
X \arrow[r, "f"] \arrow[d, "\xi"'] & Y \arrow[d, "\zeta"] \\
TX \arrow[r, "Tf"'] & TY
\end{tikzcd}$$

commute. Flat fixpoint formulas are invariant under coalgebra morphisms:

Lemma 2.2. *Let $f : (X,\xi) \to (Y,\zeta)$ be a T-coalgebra morphism, let $B \subseteq Y$, and let ϕ be a flat fixpoint formula. Then*

$$[\![\phi]\!]_{(X,\xi)}(f^{-1}[B]) = f^{-1}[[\![\phi]\!]_{(Y,\zeta)}(B)]. \tag{1}$$

(The lemma holds more generally for the full coalgebraic μ-calculus, with essentially the same proof; we refrain from stating it in more generality here to avoid introducing additional notation.)

Proof. Induction over ϕ; the Boolean cases are trivial, and the modal cases are by naturality of predicate liftings (cf. (Pattinson 2004)). We are left with the fixpoint case; we work with μ-calculus notation, i.e., our remaining case is of the form $\mu x.\,\phi$, a closed formula, because there is only one recursion variable. That is, we are to show that

$$[\![\mu x.\,\phi]\!]_{(X,\xi)} = f^{-1}[[\![\mu x.\,\phi]\!]_{(Y,\zeta)}]. \tag{2}$$

It is well known that we can approximate least fixpoints of monotone functions from below using ordinal-indexed chains. Specifically, $[\![\mu x.\,\phi]\!]_{(X,\xi)}$ is the union of the sets $[\![\phi]\!]^\alpha_{(X,\xi)}(\emptyset)$ indexed over ordinals α, defined by $[\![\phi]\!]^0_{(X,\xi)}(\emptyset) = \emptyset$, by $[\![\phi]\!]^{\alpha+1}_{(X,\xi)}(\emptyset) = [\![\phi]\!]_{(X,\xi)}([\![\phi]\!]^\alpha_{(X,\xi)}(\emptyset))$ in the successor step and by $[\![\phi]\!]^\alpha_{(X,\xi)}(\emptyset) = \bigcup_{\beta<\alpha}[\![\phi]\!]^\beta_{(X,\xi)}(\emptyset)$ in the limit step; analogously, $[\![\mu x.\,\phi]\!]_{(Y,\zeta)}$ is approximated from below by sets $[\![\phi]\!]^\alpha_{(Y,\zeta)}(\emptyset)$. Since taking preimages under f commutes with $[\![\phi]\!]$ by the inductive hypothesis, and generally commutes with (infinite) unions and preserves $\emptyset$, an easy transfinite induction shows that

$$[\![\phi]\!]^\alpha_{(X,\xi)}(\emptyset) = f^{-1}[[\![\phi]\!]^\alpha_{(Y,\zeta)}(\emptyset)]$$

for all α. The inductive claim follows by forming the union over all α on both sides, again using commutation of preimages with unions. □

Example 2.3 (Coalgebraic logics). We discuss select examples covered by the coalgebraic approach, starting with a more detailed exposition of the basic example of Kripke semantics and then moving on beyond. More examples are found, e.g., in Schröder and Pattinson (2009) and Pattinson and Schröder (2010). For the sake of readability, we elide the treatment of propositional atoms in all examples except Kripke semantics, as the technicalities are the same in all cases.

(1) *Kripke semantics:* In mild generalization of the basic example discussed above, fixpoint extensions of the modal logics K_m have modal operators $\Box_i$ for $i = 1, \ldots, m$, interpreted over the functor T given on sets by $TX = (\mathcal{P}X)^m \times \mathcal{P}U$ using the predicate liftings $[\![\Box_i]\!]_X(A) = \{(B_1, \ldots, B_m, P) \in (\mathcal{P}X)^m \times \mathcal{P}U \mid B_i \subseteq A\}$. Coalgebras for T are in 1-1 correspondence with m-relation Kripke models, and $[\![\Box_i]\!]$ captures the usual semantics of the box operators. Atomic propositions are modelled as nullary modalities $a \in \Lambda$, interpreted by nullary predicate liftings

$$[\![a]\!]_X = \{(B_1, \ldots, B_m, P) \in (\mathcal{P}X)^m \times \mathcal{P}U \mid a \in P\}.$$

CTL, $*$-nesting-free PDL, and the logic of common knowledge all are flat fixpoint logics in this setting. The CTL operators have been exemplified above; the operator $[a^*]$ of PDL is the greatest fixpoint operator

$$[a^*] = \flat_{p \wedge [a]x};$$

and the common knowledge operator C is the greatest fixpoint operator $\flat_\gamma$ for

$$\gamma = p \wedge \bigwedge_{i=1}^{m} \Box_i x.$$

(Of course, reading $\Box_i$ as a knowledge operator "agent i knows that" would standardly be combined with imposing additional frame axioms, e.g., $S5$, an aspect that we do not pursue here. Note, however, that the above definition of C effectively incorporates the actual truth of commonly known facts, $Cp \to p$, and positive introspection, $Cp \to CCp$.) Although our main result (Theorem 5.16) does not support completely arbitrary fixpoint operators, it does (like Santocanale and Venema (2010)) cover operators that go beyond CTL. For example, the operator *even*(ϕ) "on every path, every even state satisfies ϕ," which is not expressible in CTL (Emerson 1990; Wolper 1983), is the greatest fixpoint operator

$$\flat_{p \wedge \Box\Box x},$$

to which our main result does apply.

Strictly speaking, one should note that, so far, we have not covered CTL in the standard sense, according to which models are assumed to be *serial*, i.e., every state is required to have at least one successor. We model this requirement coalgebraically by replacing T as above with the functor $T'X = \mathcal{P}^*X \times \mathcal{P}U$, where $\mathcal{P}^*$ is the *non-empty powerset functor*, i.e., $\mathcal{P}^*X$ is the set of non-empty subsets of X. (Since CTL standardly uses only one next-step modality AX, we omit the exponent m.) We distinguish the arising fixpoint logics by the adjective *serial* and refer to absence of the seriality requirement by the adjective *non-serial*.

Further restricting the semantics to require that every state has *exactly* one successor in fact produces the semantics of LTL (without past) in terms of infinite sequences of valuations (where $\Box$ is typically denoted $\bigcirc$). Coalgebraically, this is reflected in using the functor $T''X = X \times \mathcal{P}U$; we refer to T''-coalgebras as *deterministic* Kripke models. (For purposes of model checking, LTL is more generally interpreted over labelled transition systems, but the more restrictive infinite-sequence semantics is equivalent for purposes of satisfiability.)

(2) *Graded fixpoint logics* are sublogics of the *graded μ-calculus* (Kupferman et al. 2002). They have modal operators $\Diamond_k$ "in more than k successors," with duals $\Box_k = \overline{\Diamond_k}$ "in all but k successors." We interpret them over the functor $\mathcal{B}$ that takes a set X to the set

$$\mathcal{B}X = X \to \mathbb{N} \cup \{\infty\}$$

of multisets over X (with possibly infinite multiplicities) by

$$[\![\Diamond_k]\!]_X(A) = \{\mu \in \mathcal{B}X \mid \mu(A) > k\},$$

where we use $\mu \in \mathcal{B}X$ as an $\mathbb{N} \cup \{\infty\}$-valued measure, i.e., write $\mu(A) = \sum_{x \in A} \mu(x)$. This captures the semantics of graded modalities over *multigraphs* (D'Agostino and Visser 2002), which is equivalent to the more customary Kripke semantics (Fine 1972) w.r.t. satisfiability of fixpoint formulas (Lemma 2.4). In description logic, graded operators are called *qualified number restrictions* (Baader et al. 2003). The example mentioned by Kupferman et al. (2002), a graded fixpoint formula expressing that the current state is the root of a finite binary tree all whose leaves satisfy p, can be expressed by the $\sharp$-operator for

$$p \vee \Diamond_1 x.$$

Similarly, the $\sharp$-operator for

$$p \vee \Box_k x$$

expresses that p holds somewhere on every infinite $k + 1$-ary tree starting at the current state. To add an example where the recursion variable x appears under more than one modality, the $\sharp$-operator for

$$p \vee \Diamond_1 \Diamond_1 x$$

expresses that the current state is the root of a finite binary tree all whose leaves are at even distance from the root and satisfy p.

(3) *Probabilistic fixpoint logics*, i.e., fixpoint extensions of probabilistic modal logic (Larsen and Skou 1991; Fagin and Halpern 1994; Heifetz and Mongin 2001), have modal operators L_r "in the next step, it holds with probability at least r that" for $r \in [0, 1] \cap \mathbb{Q}$. They are interpreted over the functor $\mathcal{D}$ that maps a set X to the set of discrete probability distributions on X by putting

$$[\![L_r]\!]_X(A) = \{P \in \mathcal{D}X \mid PA \geq r\}.$$

Recall here that a probability distribution P on X is *discrete* if it is defined on the σ-algebra of all subsets of X, which implies that P is determined by its values on singleton sets, i.e., $P(A) = \sum_{x \in A} P(\{x\})$ for all $A \subseteq X$ (and then necessarily $P(\{x\}) > 0$ for at most countably many x). Coalgebras for $\mathcal{D}$ are Markov chains. Flat probabilistic fixpoint logics in this sense are fragments of the *probabilistic μ-calculus* in the sense introduced by Cirstea et al. (2011) (to be distinguished from the $[0, 1]$-valued logic of the same name (Morgan and McIver 1997; Huth and Kwiatkowska 1997)) as an instance of the coalgebraic μ-calculus. The probabilistic μ-calculus was subsequently reinvented by Liu et al. (2015) (Liu et al. prove that model checking the probabilistic μ-calculus is in UP∩coUP and satisfiability is in 2ExpTime while Cirstea et al. already show that satisfiability is ExpTime-complete) and taken up again by Larsen et al. (2016). More generally, one can admit linear inequalities involving probabilities, as, e.g., in work by Fagin and Halpern (1994), as long as one pays attention to monotonicity.

In a view of probabilistic logic as a logic of reactive systems, we can use the $\flat$-operator

$$AG_r = \flat_{p \wedge L_r x}$$

to express formulas like $AG_r\, \neg\mathsf{fail}$, stating that the system will, at any point during its runtime, fail with probability at most $1 - r$, a sensible specification for systems that may sometimes fail but should not fail excessively often (as announced, we silently include propositional atoms such as fail in the syntax).

Alternatively, we may interpret the operators L_p epistemically. We extend the logic to multiple agents in the same way as for K in Item 1, obtaining a logic with probabilistic operators L^i_r read "agent i believes with confidence r that." We then have an uncertain variant C_r of the common knowledge operator, namely the $\flat$-operator for

$$\bigwedge_{i=1}^{m} L_r^i(p \wedge x).$$

Thus, $C_r\phi$ is read "everyone believes with confidence r that ϕ holds and that everyone believes with confidence r that ϕ holds and so on"; in short "ϕ is commonly believed with confidence r."

A variant is the operator $C_{r,q}$ that, when applied to a formula ϕ, separates belief in ϕ from beliefs about other agents: $C_{r,q}$ is the $\flat$-operator for

$$\bigwedge_{i=1}^{n}(L_r^i p \wedge L_q^i x),$$

and thus $C_{r,q}\phi$ states that everyone believes ϕ with confidence r and believes with confidence q that all agents believe the same, and so on. (Note that $C_{r,r}$ is not the same as C_r!)

Systems and logics with non-negative real-valued weights (not required to sum up to 1) are treated coalgebraically in a similar manner (Kupke and Pattinson 2010).

(4) *Conditional fixpoint logics* have a single binary modal operator $\Rightarrow$, written in infix notation. The intended reading of $a \Rightarrow b$ is "if a, then *normally* b." Conditional logics come with a wide variety of axiomatizations and semantics (see, e.g., Pattinson and Schröder (2010) for an overview). For example, the minimal conditional logic CK is interpreted over the functor Cf that maps a set X to the set $\mathcal{P}X \to \mathcal{P}X$ (more precisely $\mathcal{Q}X \to \mathcal{P}X$, where $\mathcal{Q}$ is contravariant powerset), whose coalgebras are selection function models (Chellas 1980), by putting

$$[\![\Rightarrow]\!]_X(A, B) = \{f \in CfX \mid f(A) \subseteq B\}.$$

(Thus, $\Rightarrow$ is monotone in the second but neither monotone nor antitone (order reversing) in the first argument. This makes $\Rightarrow$ different from material implication, which is precisely the point of inventing it. Technically, this is fine as long as we form fixpoints only over the second argument, i.e., let the recursion variable x appear in $\gamma \in \Gamma$ only to the right of $\Rightarrow$ as in the example below.)

We can combine coalgebraic logics freely using results of Schröder and Pattinson (2011). For example, combining conditional logic with multi-agent K in its description logic incarnation $\mathcal{ALC}$ (Baader et al. 2003), we can define an abstract concept of animal taxa with two-gender descendancy as $\flat_\gamma(\mathsf{FirstOfItsKind}, \mathsf{Male}, \mathsf{Female})$ with γ given as

$$(\neg p_1) \Rightarrow ((\exists\, \mathsf{hasParent}.\, p_2 \sqcap x) \sqcap (\exists\, \mathsf{hasParent}.\, p_3 \sqcap x)).$$

A taxon is a fixpoint of $\gamma(\mathsf{FirstOfItsKind}, \mathsf{Male}, \mathsf{Female})$ if all individuals that do not belong to some assumed first ancestor generation of the species *normally* have two parents of the same species, and the greatest fixpoint subsumes all animals belonging to such taxa. The use of the conditional $\Rightarrow$ instead of standard material implication $\to$ takes into account that these days, a given exceptional sheep might, e.g., be a clone and thus not have parents in the strict sense.

(5) *The alternating-time μ-calculus (AMC)* (Alur et al. 2002) has modal operators $\langle\langle C\rangle\rangle\bigcirc$ read "coalition C of agents can enforce ... in one step," where a *coalition* C is a subset $C \subseteq N$ of a fixed set $N = \{1, \ldots, n\}$ of agents; we shall also write $[C]$ in place of $\langle\langle C\rangle\rangle\bigcirc$ as in coalition logic (Pauly 2002). The semantics of the coalitional modalities is defined over *concurrent game structures* (or *game frames*) and can be captured coalgebraically (Schröder and Pattinson 2009): We define a functor G by

$$\mathrm{G}X = \{(f, (k_i)_{i\in N}) \mid f : (\textstyle\prod_{i\in N})\, [k_i]\} \to X,$$

where $k_i \in \mathbb{N}$ and $[k_i] = \{1, \ldots, k_i\}$ and by $\mathrm{G}g(f, (k_i)_{i\in N}) = (g \circ f, (k_i)_{i\in N})$ for $g : X \to Y$. This captures a form of concurrent game where each agent $i \in N$ chooses a move $j_i \in [k_i]$ and the joint choice determines an *outcome* $f(j_1, \ldots, j_n) \in X$. (The semantics given by Pauly (2002) differs slightly in that the agents can have unrestricted sets of available moves rather than only finite ones.) Coalgebras (X, ξ) for G are *concurrent game structures* (Alur et al. 2002); they associate to each state $x \in X$ a concurrent game $\xi(x) \in \mathrm{G}X$ whose outcomes are states and thus allow for plays with multiple successive moves. The semantics of the modalities $[C]$ is given by the liftings

$$[\![[C]]\!]_X(A) = \{(f, (k_i)_{i\in N}) \in \mathrm{G}X \mid \exists (j_i \in [k_i])_{i\in C}.\, \forall (j_i \in [k_i])_{i\in N\setminus C}.\, f((j_i)_{i\in N}) \in A\}.$$

That is, a state x in a concurrent game structure satisfies $[C]\phi$ if the agents in C have a joint choice of moves such that regardless of the choice of moves by the other agents, the outcome satisfies ϕ.

One of the flat fragments of the AMC is ATL (Alur et al. 2002). For example, the ATL-operator $\langle\!\langle C\rangle\!\rangle p_1 U p_2$, read "coalition C can eventually force p_2 and meanwhile maintain p_1," is the $\sharp$-operator for

$$p_2 \vee (p_1 \wedge [C]x).$$

As in the case of CTL, flat fixpoints in the AMC go considerably beyond ATL; e.g., the $\flat$-operator for $p \wedge [\emptyset][\emptyset]x$ ("p holds in all even states along any path") is not even in the more expressive logic ATL* (Alur et al. 2002; Dam 1994). A similar flat operator, the $\flat$-operator for $[C](p \wedge [D](q \wedge x))$, expresses that coalitions C and D can forever play ping-pong between p and q.

(6) *Monotone fixpoint logics* have a modal operator $\Box$, interpreted over the *monotone neighbourhood functor* defined by

$$\mathcal{M}X = \{\mathfrak{A} \in \mathcal{P}(\mathcal{P}X) \mid \mathfrak{A} \text{ upwards closed}\}$$

by means of the predicate lifting

$$[\![\Box]\!]_X(A) = \{\mathfrak{A} \in \mathcal{M}X \mid A \in \mathfrak{A}\}.$$

(The functor $\mathcal{M}$ acts on maps $f : X \to Y$ by $\mathcal{M}f(\mathfrak{A}) = \{B \in \mathcal{P}Y \mid f^{-1}[B] \in \mathfrak{A}\}$ and hence is a subfunctor of the double contravariant powerset functor.) Often, the axioms $\Box\top$ and $\Diamond\top$ are imposed where $\Diamond = \overline{\Box}$ denotes the dual of $\Box$. This amounts to using the subfunctor $\mathcal{M}_s$ of $\mathcal{M}$ given by

$$\mathcal{M}_s X = \{\mathfrak{A} \in \mathcal{M}X \mid \emptyset \notin \mathfrak{A} \ni X\},$$

the *serial monotone neighbourhood functor*, whose coalgebras are *serial monotone neighbourhood frames*. In particular, these form the semantic setting of concurrent PDL (Peleg 1987) and game logic (Parikh 1985), where operators are indexed over atomic programs or games, respectively; this is modelled coalgebraically in the same way as multi-modal K (Item 1). The $*$-nesting-free fragments of concurrent PDL and game logic are flat fixpoint logics. For example, game logic has operators $\langle\gamma\rangle$, indexed over composite games γ and read "Angel has a strategy to enforce … in game γ". Games are formed from atomic games using the usual constructs for regular expressions as in PDL, and, additionally, the dualizing operator $(-)^d$ that swaps the roles of the players in a game. For example, one has a demonic iteration operator $(-)^\times$ defined by $\gamma^\times = ((\gamma^d)^*)^d$. The formula $\langle\gamma^\times\rangle\phi$ thus says that Angel has a strategy to enforce ϕ in the game where γ is played repeatedly, with Demon choosing the number of rounds. If γ is star free (i.e., contains neither $(-)^*$ nor $(-)^\times$), then we can phrase $\langle\gamma^\times\rangle$ as the $\flat$-operator for $p \wedge \langle\gamma\rangle x$.

It remains to prove the mentioned equivalence of the Kripke semantics and the multigraph semantics of the graded μ-calculus (Example 2.3(2)), generalizing the equivalence for the fixpoint-free case (Schröder 2007). Since the Kripke semantics of graded modalities fails to be coalgebraic (as it violates the naturality condition for predicate liftings), this equivalence is needed to obtain completeness of flat fragments of the graded μ-calculus from our generic completeness result.

Lemma 2.4. *A formula in the flat graded μ-calculus is satisfiable over (finite) Kripke frames iff it is satisfiable over (finite) multigraphs.*

Proof. For the sake of simplicity, we continue to elide propositional atoms. Since every Kripke frame can be regarded as a multigraph, "only if" is clear. To show "if", let ϕ be a flat graded fixpoint formula, and let x_0 be a state in a $\mathcal{B}$-coalgebra (X, ξ) such that $x_0 \models_{(X,\xi)} \phi$. Let k_0 be maximal such that $\Diamond_{k_0}$ occurs in ϕ. Observe that ϕ remains satisfied if we replace ξ with ξ', where $\xi'(x)(y) = \min(\xi(x)(y), k_0 + 1)$ (formally, this is proved by induction on ϕ), so we can assume that all $\xi(x)(y)$ are finite (in fact, at most $k_0 + 1$). Now construct a Kripke model $(\bar{X}, R)$ by making sufficiently many copies of states, as in Schröder (2007, Remark 6): Take as states in $\bar{X}$ all pairs $(y, j) \in X \times \mathbb{N}$ such that $\xi(x)(y) > j$ for some x, and in this case put $(x, i)R(y, j)$ for all i such that $(x, i) \in \bar{X}$. Note that $\bar{X}$ is finite if X is finite. Like for any Kripke frame, we can equivalently regard $(\bar{X}, R)$ as a multigraph $(\bar{X}, \bar{\xi})$, where $\bar{\xi}(x, i)(y, j) = 1$ if $(x, i)R(y, j)$, and $\bar{\xi}(x, i)(y, j) = 0$ otherwise. Let $\pi : \bar{X} \to X$ denote the projection that maps (x, i) to x. By Lemma 2.2, it suffices to show that

$$\pi : (\bar{X}, \bar{\xi}) \to (X, \xi)$$

is a $\mathcal{B}$-coalgebra morphism, i.e., that

$$\mathcal{B}\pi(\bar{\xi}(x, i)) = \xi(x).$$

Indeed, the multiplicity of $y \in X$ in the multiset on the left-hand side is the cardinality of the set $\{(y, j) \in \bar{X} \mid (x, i)R(y, j)\}$, which equals $\xi(x)(y)$ by construction of R. □

3 THE GENERIC AXIOMATIZATION

The generic semantic and syntactic framework of the previous section comes with a generic, parametrized *deduction system*, whose completeness will be our main result. We begin with the fixed part of the deduction system. We include full *propositional reasoning*, i.e., introduction of substituted propositional tautologies and modus ponens. Fixpoints are governed by the obvious *Kozen-Park axiomatization*: We have the *unfolding* axiom

$$\gamma(\phi, \sharp_\gamma \phi) \to \sharp_\gamma \phi$$

and the *fixpoint induction* rule

$$\frac{\gamma(\phi, \chi) \to \chi}{\sharp_\gamma \phi \to \chi},$$

for all formulas ϕ, χ; together, these axioms capture the fact that $\sharp_\gamma \phi$ is the least prefixpoint of $\gamma(\phi)$.

Lemma 3.1. *The monotonicity and congruence rules*

$$\frac{\phi \to \psi}{\sharp_\gamma \phi \to \sharp_\gamma \psi} \qquad \frac{\phi \leftrightarrow \psi}{\sharp_\gamma \phi \leftrightarrow \sharp_\gamma \psi}$$

are derivable.

Proof. The congruence rule is derivable from the monotonicity rule. To derive the latter, assume $\phi \to \psi$. Since γ is positive in the parameter variable, we can then derive $\gamma(\phi, \sharp_\gamma \psi) \to \gamma(\psi, \sharp_\gamma \psi)$. By unfolding, we derive $\gamma(\phi, \sharp_\gamma \psi) \to \sharp_\gamma \psi$ and then $\sharp_\gamma \phi \to \sharp_\gamma \psi$ by fixpoint induction. □

The variable part of the proof system is the axiomatization of the modal operators, which turns out to be completely orthogonal to the fixpoint axiomatization. In fact, we can just reuse complete rule sets for the purely modal part of the logic (Schröder 2007; Schröder and Pattinson 2009). First some notation.

Definition 3.2. We denote the set of propositional formulas over a set Z by $\mathrm{Prop}(Z)$ and the set $\{\heartsuit a \mid \heartsuit \in \Lambda, a \in Z\}$ by $\Lambda(Z)$. A *literal* over Z is either an element of Z or the negation of such an element, i.e., has the form ϵz, where $z \in Z$ and $\epsilon \in \{\cdot, \neg\}$ is either nothing or negation. A (disjunctive) *clause* is a finite (possibly empty) disjunction of literals; a *conjunctive clause* is a finite conjunction of literals. A clause (disjunctive or conjunctive) is *contracted* if it contains every literal at most once. For $\phi, \psi \in \mathsf{Prop}(Z)$, we say that ϕ *propositionally entails* ψ, and write $\phi \vdash_{PL} \psi$, if $\phi \to \psi$ is a propositional tautology. Similarly, $\Phi \subseteq \mathsf{Prop}(Z)$ propositionally entails ψ ($\Phi \vdash_{PL} \psi$) if there exist $\phi_1, \ldots, \phi_n \in \Phi$ such that $\phi_1 \wedge \cdots \wedge \phi_n \vdash_{PL} \psi$. We write 2 for the set $\{\bot, \top\}$ of truth values. For $\phi \in \mathsf{Prop}(Z)$, we denote the evaluation of ϕ in the Boolean algebra $\mathcal{P}X$ under a valuation $\tau : Z \to \mathcal{P}X$ by $[\![\phi]\!]_{X,\tau}$ and write $X, \tau \models \phi$ if $[\![\phi]\!]_{X,\tau} = X$. For $\psi \in \mathsf{Prop}(\Lambda(\mathsf{Prop}(Z)))$, the interpretation

$$[\![\psi]\!]_{TX,\tau} \subseteq TX$$

of ψ in the Boolean algebra $\mathcal{P}(TX)$ under τ is the inductive extension of the assignment

$$[\![\heartsuit(z)]\!]_{TX,\tau} = [\![\heartsuit]\!]_X \tau(z).$$

We write $TX, \tau \models \psi$ if $[\![\psi]\!]_{TX,\tau} = TX$. A propositional formula over $\Lambda(V)$ is *clean* if it mentions every variable at most once.

We can now give the formal definition of the modal rule format. We continue to use the set V of variables as placeholders for formulas; when variables appear in rules, we prefer to call them *propositional* variables.

Definition 3.3. A *one-step rule* $R = \phi/\chi$ consists of a *premise* $\phi \in \mathsf{Prop}(V)$ and a *conclusion* χ that is a clean (disjunctive) clause over $\Lambda(V)$, where every variable in ϕ appears also in χ. We say that R is *monotone* (a notion similar to one introduced by Cirstea et al. (2011) for rules phrased without negation) if whenever χ is positive (negative) in a variable $a \in V$, then ϕ is positive (negative) in a (note that because χ is clean, we have for every variable a that χ is either positive or negative in a). The rule R is *one-step sound* if whenever $X, \tau \models \phi$ for a valuation $\tau : V \to \mathcal{P}X$, then $TX, \tau \models \chi$. A set $\mathcal{R}$ of one-step rules is *one-step complete* if, whenever $TX, \tau \models \psi$ for a set X, a clean clause ψ over $\Lambda(V)$, and a $\mathcal{P}X$-valuation τ, then ψ is *provable over* X, τ, i.e., propositionally entailed by clauses $\chi\sigma$, where $\phi/\chi \in \mathcal{R}$ and σ is a $\mathsf{Prop}(V)$-substitution such that $X, \tau \models \phi\sigma$. Moreover, $\mathcal{R}$ is *one-step cutfree complete* if, whenever $TX, \tau \models \psi$ for X, τ, ψ as above, then ψ is *cutfree provable over* X, τ, i.e., $\chi \vdash_{PL} \psi$ for some $\phi/\chi \in \mathcal{R}$ such that $X, \tau \models \phi$.

Remark 3.4. In the terminology of Schröder and Pattinson (2009), one-step cutfree complete rule sets correspond to one-step complete rule sets that are closed under contraction, resolution, and injective renamings of the propositional variables. Notice in particular that χ as in the definition of cutfree provability is clean, being the conclusion of a one-step rule, and hence contracted.

As the last parameter of the framework, we

> *fix from now on a one-step complete set $\mathcal{R}$ of one-step sound one-step rules, and denote the arising logic by $\mathcal{L}_\sharp$.*

Rules $\phi/\psi \in \mathcal{R}$ are applied in substituted form, i.e., for every substitution σ, we may conclude $\psi\sigma$ from $\phi\sigma$. In summary, the proof system consists of propositional reasoning, the unfolding axiom, the fixpoint induction rule, and the rules in $\mathcal{R}$. It is easy to see that this system is sound; explicitly:

THEOREM 3.5 (SOUNDNESS). *Every provable fixpoint formula is valid.*

We will show that the system is also complete. Note that that the system without the fixpoint rules (i.e., comprising only propositional reasoning and a one-step complete set of modal rules) is known to be complete for the fixpoint-free modal language (Schröder 2007). As usual, we write $\vdash \phi$ if a formula ϕ is provable. We say that ϕ is *consistent* if $\neg\phi$ is not provable.

We proceed to record some facts on the relationship between the two notions of one-step completeness.

Definition 3.6. A one-step rule ϕ/χ is *$\mathcal{R}$-derivable* if χ is propositionally entailed by conclusions $\psi\sigma$ of rules $\rho/\psi \in \mathcal{R}$, where σ is a $\mathsf{Prop}(V)$-substitution and $\phi \vdash_{PL} \rho\sigma$.

Of course, $\mathcal{R}$-derivable rules are sound. The converse holds thanks to one-step completeness:

Lemma 3.7. *All one-step sound one-step rules are $\mathcal{R}$-derivable.*

Proof. Let ϕ/χ be one-step sound. Let $V_0 \subseteq V$ be the set of propositional variables that occur in χ, and put

$$X = \{\kappa : V_0 \to 2 \mid \kappa \models \phi\} \qquad \tau(a) = \{\kappa \in X \mid \kappa(a) = \top\}.$$

Then $X, \kappa \models \phi$, so that $TX, \kappa \models \chi$ by one-step soundness. By one-step completeness of $\mathcal{R}$, χ is propositionally entailed by clauses $\psi\sigma$, where $\rho/\psi \in \mathcal{R}$ and σ is a $\mathsf{Prop}(V)$-substitution such that $X, \tau \models \rho\sigma$. It suffices to show that $\phi \vdash_{PL} \rho\sigma$, which, however, is clear by construction of X, τ. □

Lemma 3.8. *The set of $\mathcal{R}$-derivable monotone one-step rules is one-step cutfree complete.*

In other words, *a clause over $\Lambda(V)$ is provable over X, τ iff it is cutfree provable over X, τ using an $\mathcal{R}$-derivable monotone one-step rule.*

Proof (Lemma 3.8). By Lemma 3.7, this follows once we show that the set of one-step sound monotone one-step rules is one-step cutfree complete. This is proved by Cirstea et al. (2011, Proposition 4.7) for a formally even more restrictive class of rules (also called monotone[1]). □

Definition 3.9. A conjunctive clause ρ over $\Lambda(V)$ is *one-step τ-consistent* for $\tau : V \to \mathcal{P}X$ if $\neg\rho$ is not provable over X, τ. Moreover, ρ is *one-step τ-satisfiable* if $[\![\rho]\!]_{TX,\tau} \neq \emptyset$. We extend the notion of one-step τ-consistency to infinite sets Φ of literals over $\Lambda(V)$: We say that Φ is *one-step τ-consistent* if, for all $\rho_1, \ldots, \rho_n \in \Phi$, the conjunctive clause $\rho_1 \wedge \cdots \wedge \rho_n$ is one-step τ-consistent.

Note that by the above, ρ is one-step τ-consistent iff $\neg\rho$ is not cutfree provable over X, τ using an $\mathcal{R}$-derivable monotone one-step rule. One-step completeness can be rephrased as saying that *every one-step τ-consistent conjunctive clause over $\Lambda(V)$ is one-step τ-satisfiable.*

Example 3.10 (One-step rules). One-step complete rule systems have been exhibited for all logics of Example 2.3 and many more (Pattinson 2003; Cîrstea and Pattinson 2007; Schröder and Pattinson 2009; Pattinson and Schröder 2010; Kupke and Pattinson 2010). In some cases (Cîrstea and Pattinson 2007; Pattinson 2003), axiomatizations have been phrased in terms of one-step *axioms*, i.e., formulas in $\phi \in \mathsf{Prop}(\Lambda(\mathsf{Prop}(V)))$ that can be introduced in substituted form $\phi\sigma$ with σ a substitution of propositional variables by formulas; it has been shown that this format is interconvertible with one-step rules (Schröder 2007). We recall some examples in more detail, converting to one-step rules where necessary.

(1) *Kripke semantics:* The standard axiomatization of the modal logic K_1, with $\Box_1$ written as just $\Box$, can be phrased in terms of one-step rules for necessitation, monotonicity, and normality as

$$\frac{a}{\Box a} \quad \frac{a \to b}{\Box a \to \Box b} \quad \frac{a \wedge b \to c}{\Box a \wedge \Box b \to \Box c},$$

[1]Specifically, a rule is monotone in the sense of Cirstea at al. if its premise is a positive formula and its conclusion is a disjunction of atoms of the form $\heartsuit a$ or $\overline{\heartsuit} a$. All such rules are clearly monotone in our sense. In fact, we also have a converse: Monotone rules in our sense can be transformed into the more restrictive format by taking negation normal forms and substituting away negated variables.

where here and in the following, we write clauses $\neg a_1 \vee \cdots \vee \neg a_n \vee b_1 \vee \ldots b_m$ as implications $a_1 \wedge \cdots \wedge a_n \rightarrow b_1 \vee \cdots \vee b_m$. Given this axiomatization, a one-step rule ϕ/ψ is $\mathcal{R}$-derivable iff ψ contains disjuncts $\neg\Box a_1, \ldots, \neg\Box a_n, \Box b$ (with $n \geq 0$) such that $\phi \vdash_{PL} a_1 \wedge \cdots \wedge a_n \rightarrow b$ (all such rules are clearly derivable, and the converse is easily seen semantically (Pattinson and Schröder 2010, Example 4.6)).

A more restricted semantics over serial Kripke models as for serial CTL (Example 2.3.1) is reflected in the additional rule

$$\frac{\neg a}{\neg\Box a}.$$

Restricting additionally to deterministic models as used for LTL (and switching from $\Box$ to $\bigcirc$ as customary) logically corresponds to adding the rule

$$\frac{a \vee b}{\bigcirc a \vee \bigcirc b}$$

(which, given the other rules, is mutually interderivable with the better-known axiom $\neg \bigcirc a \rightarrow \bigcirc\neg a$).

(2) *Graded fixpoint logics:* Rephrasing a known complete axiomatization of graded modal logic (De Caro 1988), we obtain the rules

$$(RG1)\ \frac{a \rightarrow b}{\Diamond_{n+1} a \rightarrow \Diamond_n b} \qquad (A1)\ \frac{c \rightarrow a \vee b}{\Diamond_{n_1+n_2} c \rightarrow \Diamond_{n_1} a \vee \Diamond_{n_2} b}$$

$$(A2)\ \frac{\begin{array}{c} a \vee b \rightarrow c \\ a \wedge b \rightarrow d \end{array}}{\Diamond_{n_1} a \wedge \Diamond_{n_2} b \rightarrow \Diamond_{n_1+n_2+1} c \vee \Diamond_0 d} \qquad (RN)\ \frac{\neg a}{\neg\Diamond_0 a}.$$

These rules are clearly one-step sound. They have previously been shown to be one-step complete (Schröder and Pattinson 2009) by reference to the previous completeness proof for graded modal logic; we give a simple stand-alone proof in Lemma 3.12.

(3) *Probabilistic fixpoint logics:* We can reuse the one-step complete rule set for probabilistic modal logic (Heifetz and Mongin 2001; Cîrstea and Pattinson 2007). For the extended language with linear inequalities on probabilities, one has the one-step cutfree complete rule set given by Kupke and Pattinson (2010), noting that one-step cutfree complete rule sets can be restricted to any subset of the modal operators (Schröder and Pattinson 2009), in particular to monotone linear inequalities.

(4) *Conditional fixpoint logics:* One-step complete rule sets are known for various flavours of conditional logic (Pattinson and Schröder 2010; Schröder et al. 2010).

(5) *Alternating-time μ-calculus:* The following one-step complete set of rules is implicit in Pauly (2002) (see also Schröder and Pattinson (2009)):

$$\frac{\neg a}{\neg[C]a} \qquad \frac{a}{[C]a} \qquad \frac{a \vee b}{[0]a \vee [N]b} \qquad \frac{a \wedge b \rightarrow c}{[C]a \wedge [D]b \rightarrow [C \cup D]c},$$

where C and D are disjoint in the last rule. In words: No coalition can enforce the logically impossible; every coalition can enforce logical tautologies; either a is unavoidable or $\neg a$ can be enforced by all agents in collaboration; and disjoint coalitions can combine their abilities.

(6) *Monotone fixpoint logics:* When we interpret $\Box$ over the monotone neighbourhood functor, we have only the rule

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

Seriality is captured by the additional rules

$$\frac{a}{\Box a} \qquad \frac{\neg a}{\neg \Box a}.$$

Remark 3.11. We point out that in all examples with finite modal similarity type Λ, the rule sets given above are finite, so that our completeness result will establish finite axiomatizability; this holds in particular for alternating-time logics. (Note also that one-step rules can be converted into axioms (Schröder 2007)). When Λ is infinite, we cannot reasonably expect a finite axiomatization. The rules for graded modalities are *locally finite* in the sense that every modality is mentioned only in finitely many axioms. The rules for probabilistic logic (Heifetz and Mongin 2001; Cîrstea and Pattinson 2007) do not have this property, and it seems unlikely that a locally finite axiomatization is possible in this case.

As announced we provide a stand-alone proof of one-step completeness of the rules for graded modal logic:

LEMMA 3.12. *The rules* $(RG1)$*,* $(A1)$*,* $(A2)$*, and* (RN) *(Example 3.10.2) are one-step complete for graded modal logic.*

PROOF. Let $\tau : V \to \mathcal{P}X$. We prove one-step completeness by showing that every one-step τ-consistent conjunctive clause ϕ is one-step τ-satisfiable (Definition 3.9). We can assume that X is finite (see (Schröder 2007, Proposition 23), (Myers et al. 2009, Lemma 30)), and then that $\tau : V \to \mathcal{P}X$ is surjective. By a standard argument, there exists a maximal one-step τ-consistent set Φ (Definition 3.9) of literals over $\Lambda(V)$ such that $\Phi \vdash_{PL} \phi$. (Note that the argument is entirely generic up to this point.) We construct a multiset $\mu \in \mathcal{B}X$ that satisfies Φ over τ, i.e., $\mu \in \bigcap_{\phi \in \Phi} [\![\phi]\!]_{TX,\tau}$. Specifically, we define μ as an $\mathbb{N} \cup \{\infty\}$-valued measure on X by

$$\mu(A) = \max\{k + 1 \mid \Diamond_k a \in \Phi, \tau(a) = A\}$$

for $A \subseteq X$, where by convention $\max \emptyset = 0$. To see that μ is really a measure, note first that $\mu(\emptyset) = 0$ by (RN), applied to a such that $\tau(a) = \emptyset$. Moreover, μ is additive, i.e., $\mu(A \cup B) = \mu(A) + \mu(B)$ for disjoint A, B. (Since X is finite, there is no issue about μ being σ-additive.) Here, $\geq$ is by $(A2)$ and (RN), applied to a, b, c, d such that $\tau(a) = A$, $\tau(b) = B$, $\tau(d) = \emptyset$, and $\tau(c) = A \cup B$. Similarly, $\leq$ is by $(A1)$ applied to a, b, c such that $\tau(a) = A$, $\tau(b) = B$, and $\tau(c) = A \cup B$. It remains to show that $\mu \in [\![\Diamond_k a]\!]_{TX,\tau}$ iff $\Diamond_k a \in \Phi$. "If" is immediate by the definition of μ, and "only if" is by $(RG1)$. $\Box$

Completeness Proofs: Coalgebraic Logics and Flat Fixpoint Logics

Our completeness proof brings together ideas from earlier work of the authors and others on, respectively, coalgebraic modal logic and flat modal fixpoint logics. To provide some orientation in our proof, we very briefly summarize the salient points of the mentioned two lines of research.

Coalgebraic Modal Logic: Completeness via One-Step Completeness. The simplest case where one-step completeness of a set of one-step rules has been used to establish completeness of a full-blown modal logic is the basic next-step logic, given by omitting all fixpoint operators from $\mathcal{F}_\sharp$ (Schröder 2007). Here are some of the main points of this model construction:

- As usual, we start with a consistent formula ϕ and need to show that ϕ is satisfiable; in fact, we show that ϕ has a finite model.
- To this end, we form the (finite) set Σ of all subformulas of ϕ and their negations, with a view to building a (finite) model whose set X of states consists of the Σ-*atoms*, i.e., the maximally consistent subsets of Σ.

— Of course, we intend to construct the model in such a way that the *truth lemma* holds: Every Σ-atom should satisfy, as a state of the model, precisely the Σ-formulas that it contains. Thus, every formula $\rho \in \Sigma$ comes with a putative extension $\hat{\rho}$, consisting of all Σ-atoms that contain ρ.

— The truth lemma is proved by induction, and the only case in the induction that refers to the yet-to-be-defined coalgebra structure ξ on the set of Σ-atoms is the one for the modalities. The property that ξ needs to satisfy to make this step work is *coherence*:

$$\xi(A) \in [\![\heartsuit]\!]_X(\hat{\rho}) \quad \text{iff} \quad \heartsuit\rho \in A \tag{3}$$

for every Σ-atom A and every $\heartsuit\rho \in \Sigma$. Indeed, in the induction for the truth lemma, the inductive hypothesis implies that $\hat{\rho}$ is not just the putative but the actual extension of ρ, and then the left-hand side of (3) is precisely the definition of A satisfying $\heartsuit\rho$.

— The key step completing the proof is thus to establish the *existence lemma*, which states that a coherent coalgebra structure on the set of atoms does exist. This is where one-step completeness comes in: If for some Σ-atom A there does *not* exist any choice of $\xi(A)$ satisfying the coherence requirement, then this means that $TX, \tau \models \psi$ where ψ is a (disjunctive) clause over $\Lambda(V)$ containing $\neg\heartsuit a_\rho$ whenever $\heartsuit\rho \in A$ and $\heartsuit a_\rho$ whenever $\neg\heartsuit\rho \in A$, and $\tau(a_\rho) = \hat{\rho}$. It follows by one-step completeness that ψ is one-step provable over X, τ, which readily leads to a contradiction.

Flat Modal Fixpoint Logic: Algebraic Completeness via Constructive Fixpoints. The completeness result by Santocanale and the second author (Santocanale and Venema 2010) deals with systems that are simple in a complementary direction, in the sense that they do allow flat fixpoint operators, while their coalgebraic language is that of basic modal logic interpreted on standard Kripke models.

The proof method here follows the tradition of algebraic logic, encoding the derivation system into the so-called *Lindenbaum* algebra, a special algebraic structure that is built from the equivalence classes of formulas modulo provable equivalence. Where this Lindenbaum algebra corresponds to the deductive side of the logic, the semantic side is provided by assigning to each Kripke structure a so-called complex algebra, whose carrier is just the powerset of the set of states. Completeness of the logic is then proved via a representation theorem stating that the Lindenbaum algebra can be embedded into one of these complex algebras.

This approach is standard in the context of basic modal logic (Blackburn et al. 2001); what makes it work in the setting of flat modal fixpoint logics as well is that for many of these logics, the Lindenbaum algebra has some special properties. Most significantly, the goal of Santocanale and Venema's proof is to ensure that every least $\sharp$-fixpoint in the Lindenbaum algebra is *constructive*, that is, reached after at most ω many steps in the standard ordinal approximation of fixpoints. The central tool for proving this constructivity is the notion of a *finitary O-adjoint*, introduced by Santocanale (Santocanale 2008). Further on we provide a precise definition and an example of this notion; for the moment we just note that O-adjointness is a generalisation of ordinary adjointness. The point for the completeness proof is that finitary O-adjoints provide constructive fixpoints, as a fairly straightforward proof will reveal.

Summarizing, this approach yields an automatic completeness result for flat modal fixpoint languages that restricts to fixpoint operators $\sharp_\gamma$ for which the formula γ induces a finitary O-adjoint on the Lindenbaum algebra. Finally, Santocanale and Venema provide some sufficient syntactic conditions on the fixpoint schemes that ensure finitary O-adjointness and, thus, completeness.

Roughly speaking, our completeness proof for flat coalgebraic fixpoint logics combines the algebraic approach for flat modal fixpoint logics with the model construction of coalgebraic logic. The salient point about this coalgebraic model construction for the case with fixpoints, which will be given in Section 5, is that plain atoms no longer suffice; we will additionally annotate least-fixpoint formulas with suitable *time-outs* to differentiate between multiple occurrences of the same atom (see Example 5.13). The results of the next section show, in a nutshell, that such time-outs, which determine how often a least fixpoint can be unfolded, can indeed be introduced consistently. The key observation here is that we can generalize the concept of O-adjointness to the coalgebraic setting, and, under some syntactic conditions on the fixpoint scheme, guarantee the constructiveness of the Lindenbaum algebra of the logic.

4 CONSTRUCTIVE FIXPOINTS

We now launch into the details of the completeness proof for the system introduced in the previous section. Where we give examples of technical definitions along the way, these will, for readability, mostly be cast in the basic, relational instance of coalgebraic logic (Example 2.3.(1)). This should not distract from the fact that all results apply to coalgebraic logic in full generality as illustrated in Example 2.3.

Our first aim is to prove that fixpoints in the Lindenbaum algebra of $\mathcal{L}_\sharp$ are *constructive*, i.e., can be iteratively approximated in ω steps. In terms of consistency of formulas, this means that whenever a formula of the form $\sharp_\gamma \phi \wedge \psi$ is consistent, then already $\gamma^i(\phi)(\bot) \wedge \psi$ is consistent for some $i < \omega$; this fact will play a pivotal role in the model construction in Section 5. We begin by introducing the requisite algebraic tools.

Definition 4.1. A Λ-*algebra* is a Boolean algebra A with a monotone operation $\heartsuit^A : A \to A$ for each $\heartsuit \in \Lambda$. In a Λ-algebra A, a modal fixpoint scheme $\gamma(p,x)$ is interpreted as an operation $\gamma^A : A^2 \to A$ in the evident way. A $\sharp$-*algebra* is a Λ-algebra A that is endowed with operations $\sharp_\gamma^A$ for each $\gamma(p,x) \in \Gamma$ such that for each $a \in A$, $\sharp_\gamma^A(a)$ is the least fixpoint of the map $\gamma^A(a,-) : A \to A$ (in particular, these fixpoints *exist* in a $\sharp$-algebra). In a $\sharp$-algebra A, every fixpoint formula with variables $\phi(v_1,\dots,v_n)$ is interpreted as a function $\phi^A : A^n \to A$. We say that A *validates* a rule $R = \phi/\psi$ if $\psi^A(a_1,\dots,a_n) = \top$ whenever $\phi^A(a_1,\dots,a_n) = \top$ for $a_1,\dots,a_n \in A$. An $\mathcal{L}_\sharp$-*algebra* is a $\sharp$-algebra A that validates every rule in our fixed set $\mathcal{R}$ of one-step rules.

As mentioned earlier, in the tradition of algebraic logic, $\mathcal{L}_\sharp$-algebras provide an algebraic encoding of our proof system. More specifically, we will be interested in the *Lindenbaum algebra* $A(\mathcal{L}_\sharp)$ of our logic. As usual, this algebra (almost always infinite) is defined as the quotient of the set $\mathcal{F}_\sharp$ of formulas under the congruence relation $\equiv$ of provable equivalence ($\phi \equiv \psi$ iff $\vdash \phi \leftrightarrow \psi$), equipped with the algebra structure that just interprets every connective as itself. (The congruence property of $\equiv$ follows from the fact that both the modalities and the fixpoint operators come with congruence rules, the former by one-step completeness and the latter by Lemma 3.1.) An easy induction shows that every formula ϕ is interpreted as the element $\phi^{A(\mathcal{L}_\sharp)} = [\phi]$ of this algebra, where $[\phi]$ denotes the equivalence class of ϕ under $\equiv$; when there is no danger of confusion, we may write ϕ in place of $[\phi]$. More generally, when ϕ is a formula with n variables, then $\phi^{A(\mathcal{L}_\sharp)}([\rho_1],\dots,[\rho_n]) = [\phi(\rho_1,\dots,\rho_n)]$. The Kozen-Park axiomatization ensures that $A(\mathcal{L}_\sharp)$ actually is an $\mathcal{L}_\sharp$-algebra:

Lemma 4.2. *The Lindenbaum algebra is the initial $\mathcal{L}_\sharp$-algebra.*

Proof. The unfolding axiom makes $\sharp_\gamma \phi$ a prefixpoint of $\gamma^{A(\mathcal{L}_\sharp)}(\phi,-)$ in $A(\mathcal{L}_\sharp)$, and the fixpoint induction rule ensures that $\sharp_\gamma \phi$ is the least such, because every element of $A(\mathcal{L}_\sharp)$ is the denotation of a formula. This shows that $A(\mathcal{L}_\sharp)$ is a $\sharp$-algebra. To show that $A(\mathcal{L}_\sharp)$ is an $\mathcal{L}_\sharp$-algebra, it remains to establish that $A(\mathcal{L}_\sharp)$ validates every one-step rule ϕ/ψ in $\mathcal{R}$. So let $[\rho_1],\dots,[\rho_n] \in A(\mathcal{L}_\sharp)$ such that $\phi^{A(\mathcal{L}_\sharp)}([\rho_1],\dots,[\rho_n]) = \top$. Then the formula $\phi(\rho_1,\dots,\rho_n)$ is valid; since ϕ/ψ is sound, it follows that $\psi(\rho_1,\dots,\rho_n)$ is valid, i.e., $\psi^{A(\mathcal{L}_\sharp)}([\rho_1],\dots,[\rho_n]) = \top$, as required.

Initiality of $A(\mathcal{L}_\sharp)$ is then straightforward: If B is an $\mathcal{L}_\sharp$-algebra, then B validates the unfolding axiom and the fixpoint induction rule so that we have a well-defined map $f : A(\mathcal{L}_\sharp) \to B$ given by $f([\phi]) = \phi^B$. By construction of the algebra structure on $A(\mathcal{L}_\sharp)$, f is homomorphic w.r.t. all algebraic operators (Boolean, modal, and fixpoint). Uniqueness of f is shown by induction over formulas in the usual way, using the fact that the fixpoint operators are explicitly included in the algebra structure. □

In these terms, our target property is phrased as follows.

Definition 4.3. We say that $\gamma \in \Gamma$ is *constructive* if for all ϕ,

$$\sharp_\gamma \phi = \bigvee_{i<\omega} \gamma(\phi)^i(\bot)$$

in the Lindenbaum algebra $A(\mathcal{L}_\sharp)$, i.e., if $\vdash \sharp_\gamma \phi \to \psi$ whenever $\vdash \gamma(\phi)^i(\bot) \to \psi$ for all $i < \omega$.

We explicitly state the dual formulation of this property:

Lemma 4.4. *Let γ be constructive. If $\sharp_\gamma \phi \wedge \psi$ is consistent, then $\gamma(\phi)^i(\bot) \wedge \psi$ is consistent for some $i < \omega$.*

(In algebraic terms, the above lemma says that if γ is constructive and $\sharp_\gamma \phi \wedge \psi > \bot$ in $A(\mathcal{L}_\sharp)$, then $\gamma(\phi)^i(\bot) \wedge \psi > \bot$ for some $i < \omega$.)

Example 4.5 (Constructivity). Constructivity of the fixpoint

$$\sharp_{p \vee \Diamond x} \phi$$

defining the CTL formula $EF\phi$ means that $\sharp_{p\vee\Diamond x}\phi$ is the supremum of the infinite ascending chain

$$\begin{array}{l} \bot \\ \phi \\ \phi \vee \Diamond\phi \\ \phi \vee \Diamond\phi \vee \Diamond\Diamond\phi \\ \dots \end{array}$$

in the Lindenbaum algebra (noting that algebraically, $\rho \vee \Diamond\bot = \rho$).

The central tool for proving constructivity, introduced by Santocanale (2008) and featuring prominently in subsequent work by Santocanale and Venema (2010), is the notion of a finitary O-adjoint. We recall that a map f between partially ordered sets (e.g., between Boolean algebras) is a *left adjoint* if there exists a map g in the opposite direction such that $f(x) \le y$ iff $x \le g(y)$. This requirement is suitably weakened in the definition of O-adjointness:

Definition 4.6. We say that a formula $\psi(x)$ is an *O-adjoint* if for all $\phi \in \mathcal{F}_\sharp$, there exists a finite set $G_\psi(\phi)$ of formulas such that for all $\rho \in \mathcal{F}_\sharp$,

$$\vdash \psi(\rho) \to \phi \text{ iff } \vdash \rho \to \chi \text{ for some } \chi \in G_\psi(\phi), \tag{4}$$

i.e., $\psi(\rho) \le \phi$ in $A(\mathcal{L}_\sharp)$ iff $\rho \le \chi$ for some $\chi \in G_\psi(\phi)$. Moreover, ψ is a *finitary O-adjoint* if G_ψ can be chosen such that for every ϕ, the closure of ϕ under G_ψ, i.e., the least set $\mathcal{A}$ with $\phi \in \mathcal{A}$ and $\chi \in \mathcal{A} \Rightarrow G_\psi(\chi) \subseteq \mathcal{A}$, is finite. We say that a modal fixpoint scheme $\gamma \in \Gamma$ is a *(finitary) O-adjoint* if $\gamma(\phi)$ is a (finitary) O-adjoint for all $\phi \in \mathcal{F}_\sharp$.

Lemma 4.7 (Santocanale 2008). *Every finitary O-adjoint is constructive.*

Example 4.8 ($\mathcal{O}$-adjointness). Our proof of $\mathcal{O}$-adjointness (Theorem 4.22) will feature an explicit construction of $G_\psi(\phi)$ as in Definition 4.6, which we will illustrate after the proof (Example 4.23). For now, we look at an ad hoc example for specific ψ, ϕ in probabilistic modal logic (Example 2.3.(3)), anticipating, for simplicity, the completeness theorem (Theorem 5.16): Consider

$$\psi = L_{1/2}x \qquad \phi = \neg(L_{5/6}a \wedge L_{5/6}b \wedge L_{5/6}c),$$

where a, b, c are propositional atoms. Then we can take

$$G_\psi(\phi) = \{\neg(a \wedge b), \neg(b \wedge c), \neg(a \wedge c)\}.$$

To show that this really does the job, we argue semantically—and this suffices by our soundness and (anticipated) completeness results. For the "if" direction in (4), assume w.l.o.g. that $\models \rho \to \neg(a \wedge b)$, let (X, ξ) be a $\mathcal{D}$-coalgebra (i.e., a Markov chain), and let $y \in X$ such that $y \models_{(X,\xi)} L_{1/2}\rho$. Assume for a contradiction that $y \models_{(X,\xi)} L_{5/6}a \wedge L_{5/6}b$. Then $y \models_{(X,\xi)} L_{2/3}(a \wedge b)$ but by our other assumptions, $y \models_{(X,\xi)} L_{1/2}\neg(a \wedge b)$, contradiction.

We prove the "only if" direction in Equation (4) by contraposition. So assume that all of $\rho \wedge a \wedge b$, $\rho \wedge b \wedge c$, and $\rho \wedge a \wedge c$ are satisfiable. We will show that $L_{1/2}\rho \wedge L_{5/6}a \wedge L_{5/6}b \wedge L_{5/6}a$ is satisfiable as well. To start with, the formula $a \wedge b \wedge c$ is, of course, satisfiable. Forming the disjoint union of coalgebras in which the mentioned satisfiable formulas are indeed satisfied in some state and adding a new root state, we can thus construct a $\mathcal{D}$-coalgebra (X, ξ) containing a root state y_0 and states y_1, y_2, y_3, y_4 such that $y_1 \models_{(X,\xi)} \rho \wedge a \wedge b$, $y_2 \models_{(X,\xi)} \rho \wedge b \wedge c$, $y_3 \models_{(X,\xi)} \rho \wedge a \wedge c$, and $y_4 \models_{(X,\xi)} a \wedge b \wedge c$, with $\xi(y_0)(y_1) = \xi(y_0)(y_2) = \xi(y_0)(y_3) = 1/6$ and $\xi(y_0)(y_4) = 1/2$:

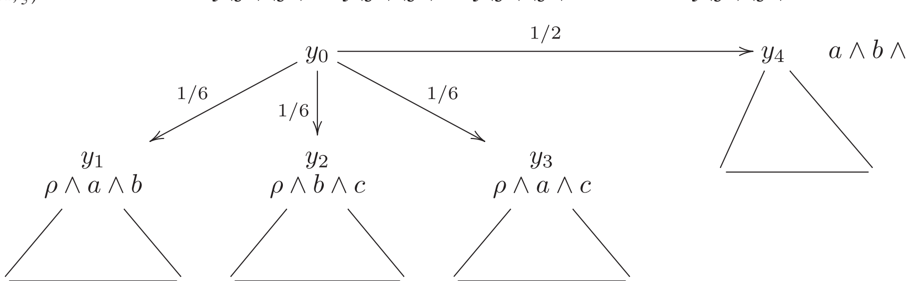

(the triangles are meant to indicate the coalgebras whose disjoint union we formed). Then $y_0 \models_{(X,\xi)} L_{1/2}\rho \wedge L_{5/6}a \wedge L_{5/6}b \wedge L_{5/6}c$.

Definition 4.9 (Top-level decomposition). Let ψ be a formula (possibly with variables) such that ψ contains some modal operator and every occurrence of a variable in ψ is in the scope of some modal operator. Let ψ' be the equivalent formula that arises by unfolding all top-level occurrences of $\sharp$ once (where an occurrence is *top-level* if it is not in the scope of a modal or fixpoint operator). By guardedness of fixpoint operators, ψ' is of the form $\psi_0\sigma$, where $\psi_0 \in \mathsf{Prop}(\Lambda(V))$ is clean and σ is a substitution. We refer to the equivalence $\psi \equiv \psi_0\sigma$ (or, more precisely, the pair (ψ_0, σ)) as the *top-level decomposition* of ψ.

Example 4.10 (Top-level decomposition). The top-level decomposition $\chi \equiv \chi_0\sigma$ of a formula of the shape

$$\chi = \Box\phi \vee \sharp_{p \wedge \Diamond x}\heartsuit\psi$$

(we give no semantics for the modalities as the definition is purely syntactic) is given by

$$\begin{aligned}\chi_0 &= \Box a \vee (\heartsuit b \wedge \Diamond c)\\ \sigma(a) &= \phi\\ \sigma(b) &= \psi\\ \sigma(c) &= \sharp_{p\wedge\Diamond x}\heartsuit\psi.\end{aligned}$$

The first step in the proof of O-adjointness for a large class of operators is a generalization of Santocanale's rigidity lemma (2008):

Lemma 4.11 (Rigidity). *Let ψ be a clause over $\Lambda(\mathcal{F}_\sharp)$, with top-level decomposition $\psi \equiv \psi_0\sigma_0$. Then ψ is provable iff there exists an $\mathcal{R}$-derivable monotone one-step rule ϕ/χ such that $\phi\sigma_0$ is provable and $\chi \vdash_{PL} \psi_0$.*

The proof relies on the *one-point extension* of an algebra (so called because it mimics the addition of a new root point in a coalgebraic model on the algebraic side), in generalization of a similar construction by Santocanale and Venema (2010):

Let A be a countable $\mathcal{L}_\sharp$-algebra, let $\mathcal{S}(A)$ be the set of ultrafilters of A, fix a surjective map $\sigma : V \to A$, and let a clean conjunctive clause ρ over $\Lambda(V)$ be one-step θ-consistent for $\theta : V \to \mathcal{P}(\mathcal{S}(A))$ given by $\theta = j \circ \sigma$ (applicative composition), where $j : A \to \mathcal{P}(\mathcal{S}(A))$ is the usual canonical map $j(a) = \{u \in \mathcal{S}(A) \mid a \in u\}$. We construct the one-point extension A^ρ, an $\mathcal{L}_\sharp$-algebra emulating the addition of a new point whose successor structure is described by ρ, as follows. To begin, we can find a maximally one-step θ-consistent set Φ (Definition 3.9) of literals over $\Lambda(V)$ such that $\Phi \vdash_{PL} \rho$. As we emulate adding a single point, the carrier of A^ρ is $A \times 2$, where 2 is the Boolean algebra $\{\bot, \top\}$; we thus have projection maps $\pi_1 : A^\rho \to A$ and $\pi_2 : A^\rho \to 2$. We make A^ρ into a Λ-algebra by putting

$$\heartsuit^{A^\rho}(a, d) = (\heartsuit^A(a), \heartsuit^\rho(a)), \tag{5}$$

where $\heartsuit^\rho : A \to 2$ is defined by $\heartsuit^\rho(a) = \top$ iff $\heartsuit a \in \Phi\sigma$. (Thus, $\heartsuit^{A^\rho}(a, d)$ is independent of d, in agreement with the semantic fact that the interpretation of modal operators depends only on the successor structure of the current state, not on the current state itself.)

Remark 4.12. We are interested only in one-point extensions A^ρ where A is countably infinite. Nevertheless, it may be helpful to look at the case where A is the complex algebra $\mathcal{P}(X)$ of a finite T-coalgebra (X, ξ) (Pattinson and Schröder 2008); the operation interpreting a modality $\heartsuit \in \Lambda$ is then defined by $\heartsuit(B) = \xi^{-1}[\![\heartsuit]\!]_X(B)$ for $B \in \mathcal{P}(X)$. In this case, we can identify $\mathcal{S}(A)$ with X and take j to be identity (as for finite X, every ultrafilter in $\mathcal{P}(X)$ is fixed), so $\theta = \sigma$. By one-step completeness, ρ is one-step θ-satisfiable (Definition 3.9), so we can pick a $t \in [\![\rho]\!]_{TX,\theta} \subseteq TX$ and then take Φ to be the set $\{\heartsuit a \mid t \in [\![\heartsuit a]\!]_{TX,\theta}\}$. (In fact, this way of obtaining Φ would work in the general case.) Now A^ρ can be identified with the complex algebra $\mathcal{P}(X \cup \{x_0\})$ of a coalgebra $(X \cup \{x_0\}, \bar{\xi})$, where $x_0 \notin X$ and $\bar{\xi}$ extends ξ to $X \cup \{x_0\}$ by $\bar{\xi}(x_0) = Ti(t)$, with i denoting the inclusion $X \hookrightarrow X \cup \{x_0\}$: We can decompose $B \in \mathcal{P}(X \cup \{x_0\})$ into a pair $(B \cap X, d)$ consisting of the subset $B \cap X \in \mathcal{P}(X)$ and an element $d \in 2$ indicating whether $x_0 \in B$, and $x_0 \in \heartsuit(B)$ iff $t \in [\![\heartsuit]\!]_X(B \cap X)$ iff $\heartsuit(B \cap X) \in \Phi\sigma$, so the effect of the algebraic operation $\heartsuit$ on d is exactly as in Equation (5). In other words, in the restricted case at hand, the construction of A^ρ really does correspond exactly to extending (X, ξ) with a new root point whose behaviour is specified by ρ.

We proceed with the general case and define a valuation $\hat{\sigma} : V \to A^\rho$ by

$$\hat{\sigma}(v) = \begin{cases}(\sigma(v), \bot) & \text{if } v \text{ occurs in a positive literal in } \rho\\ (\sigma(v), \top) & \text{otherwise.}\end{cases}$$

We then have

$$\rho\hat{\sigma} > \bot \text{ in } A^\rho : \tag{6}$$

by definition of A^ρ, the second component of $\rho\hat{\sigma}$ is $\top$ because $\Phi\sigma \vdash_{PL} \rho\sigma$.

LEMMA 4.13. *The algebra A^ρ is an $\mathcal{L}_\sharp$-algebra.*

PROOF. The proof that A^ρ is a $\sharp$-algebra is a simple application of Bekic's theorem, as in Santocanale and Venema (2010): Since A^ρ has carrier $A \times 2$, the fixpoint definition of $(\psi, c) = \sharp_\gamma(\phi, b)$ in A^ρ can be seen as a mutually recursive definition of two variables ψ, c. By guardedness of fixpoints and the interpretation of modalities in A^ρ, the definition of the first variable ψ does not mention the second variable c, so we can calculate the solution for ψ separately, using the fact that A is a $\sharp$-algebra, and then replace ψ with this solution $\sharp_\gamma(\phi)$ in the recursive definition of the second variable c. We end up with c being defined as the least fixpoint of a monotone function on 2, which exists because 2 is a complete lattice.

It remains to prove that A^ρ validates the one-step rules in $\mathcal{R}$. The first component of A^ρ behaves just like A, so that we have to verify the rules only on the second component, 2. That is, whenever we have a one-step rule $\chi/\psi \in \mathcal{R}$ and a valuation $\tau : V \to A^\rho$ such that $\chi\tau = \top$ in A^ρ, we have to prove that $\pi_2\psi\tau = \top$. Since $\psi\tau$ depends only on $\pi_1\tau$, we thus have to prove that whenever we have $\tau' : V \to A$ such that $\chi\tau' = \top$ in A, the interpretation of $\psi\tau'$ in 2 is $\top$, where the interpretation of $\psi\tau'$ is determined by means of the $\heartsuit^\rho$ and the Boolean algebra structure of 2; that is, we have to show that $\Phi\sigma \vdash_{PL} \psi\tau'$. Since σ is surjective, we have $\tau' = \sigma \circ \hat{\tau}$ (applicative composition) for some renaming $\hat{\tau} : V \to V$; our goal then reduces to showing that $\Phi \vdash_{PL} \psi\hat{\tau}$. We proceed by contradiction: Assume $\Phi \nvdash_{PL} \psi\hat{\tau}$. Then by maximality of Φ, $\Phi \vdash_{PL} \neg\psi\hat{\tau}$. Since Φ is one-step θ-consistent, this implies that $[\![\neg\chi\hat{\tau}]\!]_{\mathcal{S}(A),\theta} \neq \emptyset$. But since $\theta \circ \hat{\tau} = j \circ \sigma \circ \hat{\tau} = j \circ \tau'$, $[\![\neg\chi\hat{\tau}]\!]_{\mathcal{S}(A),\theta}$ is the image of $\neg\chi\tau' \in A$ under the Boolean homomorphism $j : A \to \mathcal{P}(\mathcal{S}(A))$. This implies $\neg\chi\tau' \neq \bot$, in contradiction to $\chi\tau' = \top$. □

In consequence of the fact that $A(\mathcal{L}_\sharp)$ is the *initial* $\mathcal{L}_\sharp$-algebra, we thus have

LEMMA 4.14. *Let $\sigma : V \to A(\mathcal{L}_\sharp)$ be surjective, and let ρ be a conjunctive clause over $\Lambda(V)$. If ρ is one-step θ-consistent for $\theta(v) = \{u \in \mathcal{S}(A(\mathcal{L}_\sharp)) \mid \sigma(v) \in u\}$, then $\rho\sigma$ is consistent, i.e., $\rho\sigma > \bot$ in $A(\mathcal{L}_\sharp)$.*

PROOF. Let f be the unique $\mathcal{L}_\sharp$-algebra homomorphism $A(\mathcal{L}_\sharp) \to A(\mathcal{L}_\sharp)^\rho$, and take $\hat{\sigma} : V \to A(\mathcal{L}_\sharp)^\rho$ as above. Note that the first projection $\pi_1 : A(\mathcal{L}_\sharp)^\rho \to A(\mathcal{L}_\sharp)$ is a homomorphism of $\mathcal{L}_\sharp$-algebras by construction of $A(\mathcal{L}_\sharp)^\rho$. By the uniqueness part of initiality, it follows that $\pi_1 \circ f = id$. For $a \in V$, we therefore have $\pi_1(f(\sigma(a)) = \sigma(a) = \pi_1(\hat{\sigma}(a))$; moreover, if a occurs in a positive literal in ρ, then $\pi_2(f(\sigma(a))) \geq \bot = \pi_2(\hat{\sigma}(a))$, and if a occurs in a negative literal in ρ (hence does not occur in a positive literal, since ρ is clean), then $\pi_2(f(\sigma(a))) \leq \top = \pi_2(\hat{\sigma}(a))$. Since the order on $A(\mathcal{L}_\sharp)^\rho$ is componentwise, this implies that $f \circ \sigma(a) \geq \hat{\sigma}(a)$ if a occurs in a positive literal in ρ, and $f \circ \sigma(a) \leq \hat{\sigma}(a)$ if a occurs in a negative literal in ρ. Since f is homomorphic and the modalities are monotone, it follows by (6) that $f(\rho\sigma) = \rho(f \circ \sigma) \geq \rho\hat{\sigma} > \bot$. Since f preserves $\bot$, this implies $\rho\sigma > \bot$. □

From Lemma 4.14, one easily proves rigidity (Lemma 4.11) using the fact that every consistent formula is contained in some ultrafilter of $A(\mathcal{L}_\sharp)$:

PROOF (LEMMA 4.11). "If" is clear. For "only if," we prove the dual statement: Let ψ be a conjunctive clause over $\Lambda(\mathcal{F}_\sharp)$, with top-level decomposition $\psi \equiv \psi_0\sigma_0$, such that $\neg\phi\sigma_0$ is consistent for all $\mathcal{R}$-derivable monotone one-step rules ϕ/χ such that $\chi \vdash_{PL} \neg\psi_0$; we show that ψ is consistent. W.l.o.g. σ_0 is a surjection $V \to \mathcal{F}_\sharp$, which we prolong to a surjection $\bar{\sigma}_0 : V \to A(\mathcal{L}_\sharp)$. By

Lemma 4.14, it suffices to prove that ψ_0 (a conjunctive clause over $\Lambda(V)$) is one-step θ-consistent for $\theta(v) = \{u \in \mathcal{S}(A(\mathcal{L}_\sharp)) \mid \bar{\sigma}_0(v) \in u\}$. Thus, let ϕ/χ be an $\mathcal{R}$-derivable rule such that $\chi \vdash_{PL} \neg\psi_0$. We have to show that $\mathcal{S}(A(\mathcal{L}_\sharp)), \theta \not\models \phi$. By assumption and because $\chi \vdash_{PL} \neg\psi_0$, we have that $\neg\phi\sigma_0$ is consistent, hence $\neg\phi\bar{\sigma}_0 > \bot$ in $A(\mathcal{L}_\sharp)$, so that there exists $u \in \mathcal{S}(A(\mathcal{L}_\sharp))$ with $\neg\phi\bar{\sigma}_0 \in u$. Now one shows by induction over $\rho \in \mathsf{Prop}(V)$ that $u \in [\![\rho]\!]_{\mathcal{S}(A(\mathcal{L}_\sharp)),\theta}$ iff $\rho\bar{\sigma}_0 \in u$: The cases for Boolean connectives are by the ultrafilter property of u, and the base case $\rho = a \in V$ is by definition of θ. In particular, we have $u \in [\![\neg\phi]\!]_{\mathcal{S}(A(\mathcal{L}_\sharp)),\theta}$, showing that $\mathcal{S}(A(\mathcal{L}_\sharp)), \theta \not\models \phi$ as required. □

In a nutshell, rigidity enables us to prove $\mathcal{O}$-adjointness of all (monotone) modal operators and, even more generally, all modal fixpoint schemes in which the recursion variable x occurs at uniform depth (such as $\Box\Diamond x \wedge \Diamond\Box x$). Formally:

Definition 4.15. A formula ψ with variables is *uniform of depth k* if every occurrence of the recursion variable x in ψ is in the scope of exactly k modal operators (including the case that x does not occur in ψ; recall, moreover, that variables never occur under fixpoint operators). Moreover, ψ is *uniform* if ψ is uniform of depth k for some k; the minimal such k is the *depth of uniformity* of ψ (in fact, k is always unique except when x does not appear in ψ, in which case the depth of uniformity is 0). Formally, uniform formulas are defined inductively by following inductive clauses:

— x is uniform of depth 0;
— formulas not containing x are uniform of depth k, for any k;
— any Boolean combination of uniform formulas of depth k is uniform of depth k; and
— if ψ is uniform of depth k, then $\heartsuit\psi$ is uniform of depth $k+1$, for $\heartsuit \in \Lambda$.

Finitariness of $\mathcal{O}$-adjoints will use the standard Fischer-Ladner closure:

Definition 4.16. A set Σ of formulas is *Fischer-Ladner closed* if Σ is closed under subformulas and negation (where the negation of a negated formula $\neg\phi$ is taken to be ϕ), and whenever $\sharp_\gamma\phi \in \Sigma$, then $\gamma(\phi, \sharp_\gamma\phi) \in \Sigma$. We denote the *Fischer-Ladner closure* of a formula ψ, i.e., the smallest Fischer-Ladner closed set containing ψ, by $FL(\psi)$.

Lemma 4.17 (Kozen 1983). *The set $FL(\psi)$ is finite.*

The further development revolves largely around derivable rules:

Definition 4.18. A *rule* $R = \phi/\psi$ consists of a *premise* ϕ and a *conclusion* ψ, both being formulas with variables. The rule R is *derivable* if ψ can be derived from the assumption ϕ using the rules of the system (propositional reasoning, unfolding, fixpoint induction, modal rules).

Lemma 4.19. *If ψ is positive in x and $\vdash \chi \to \rho$, then $\vdash \psi(\chi) \to \psi(\rho)$.*

Proof. Induction over ψ, simultaneously with a corresponding statement on formulas that are negative in x. The Boolean cases are straightforward, and the cases for fixpoint operators are trivial, because x never appears under fixpoint operators. The case for a modal operator $\heartsuit \in \Lambda$ is discharged by the fact that by monotonicity of $\heartsuit$, the monotonicity rule $a \to b/\heartsuit a \to \heartsuit b$ is one-step sound and, therefore, by Lemma 3.7, $\mathcal{R}$-derivable. □

Lemma 4.20. *Let $\psi = \psi(x)$ be a uniform formula, and put*

$$G = \{\phi \in \mathsf{Prop}(FL(\psi)) \mid \phi/\psi \textit{ derivable}, \phi \textit{ uniform of depth } 0\}.$$

Then, given a formula ρ, $\psi(\rho)$ is provable iff $\phi(\rho)$ is provable for some $\phi(x) \in G$.

Proof. *"If"*: This is essentially just the claim that derivability is stable under substitution, proved by induction over derivations.

"Only if": We proceed by induction over the depth of uniformity, with trivial base case. Thus, let ψ be uniform of depth $k > 0$, and let $\psi(\rho)$ be provable. By unfolding (guarded) top-level fixpoints (i.e., those not in the scope of a modality) and then applying propositional reasoning to transform into CNF, we reduce to the case that ψ is a clause over $\Lambda(FL(\psi))$; note here that these transformations remain within $\mathsf{Prop}(FL(\psi))$. Let $\psi \equiv \psi_0\sigma_0$ be the top-level decomposition of ψ; then the top-level decomposition of $\psi(\rho)$ has the form $\psi_0\sigma_0^\rho$ where $\sigma_0^\rho(a) = \sigma_0(a)(\rho)$. By rigidity (Lemma 4.11), we have an $\mathcal{R}$-derivable monotone one-step rule χ/ψ' such that $\chi\sigma_0^\rho$ is provable and $\psi' \vdash_{PL} \psi_0$. Then the rule $\chi\sigma_0/\psi$ is derivable, $\chi\sigma_0 \in \mathsf{Prop}(FL(\psi))$ is uniform of depth $k-1$, and $(\chi\sigma_0)(\rho)$ is provable (being equal to $\chi\sigma_0^\rho$). By the inductive assumption, applied to $\chi\sigma_0$, there exists $\phi(x) \in \mathsf{Prop}(FL(\chi\sigma_0)) \subseteq \mathsf{Prop}(FL(\psi))$ such that $\phi/\chi\sigma_0$ is derivable, ϕ is uniform of degree 0, and $\phi(\rho)$ is provable. Since $\phi/\chi\sigma_0$ and $\chi\sigma_0/\psi$ are derivable, so is ϕ/ψ, hence $\phi \in G$, which proves the claim. □

We recall a trick from propositional logic:

Lemma 4.21. *Let ψ be a formula with variables containing only top-level occurrences (Definition 4.9) of the variable y. Then there is a formula a such that*

$$\psi[\top/y] \vee \psi[\bot/y] \vdash_{PL} \psi[a/y].$$

Proof. By standard Boolean expansion. Specifically, $a = \psi[\top/y]$ does the job. This is seen by case distinction over $\psi[\top/y]$: First, assume $\psi[\top/y]$; then $\psi[\top/y]$ is equivalent to $\top$, so $\psi[\top/y]$ propositionally entails $\psi[\psi[\top/y]/y]$, because y has only top-level occurrences in ψ. Second, assume $\neg\psi[\top/y]$. From $\psi[\top/y] \vee \psi[\bot/y]$, we then conclude $\psi[\bot/y]$. Since by assumption, $\psi[\top/y]$ is equivalent to $\bot$, this, again, propositionally entails $\psi[\psi[\top/y]/y]$. □

We are now set to prove the main result of this section:

Theorem 4.22 (Finitary $\mathcal{O}$-adjointness). *If the formula ψ with recursion variable x is positive and uniform in x, then the operation $\psi^{A(\mathcal{L}_\sharp)} : A(\mathcal{L}_\sharp) \to A(\mathcal{L}_\sharp)$ induced by ψ is a finitary $\mathcal{O}$-adjoint.*

Proof. For readability, we phrase the arguments using formulas in $\mathcal{F}_\sharp$ rather than elements of $A(\mathcal{L}_\sharp)$. Let $\phi \in \mathcal{F}_\sharp$. We have to construct a set $G_\psi(\phi)$ of formulas such that for all $\rho \in \mathcal{F}_\sharp$, $\vdash \psi(\rho) \to \phi$ iff $\vdash \rho \to \chi$ for some $\chi \in G_\psi(\phi)$; moreover we require $G_\psi(\phi)$ to be finite up to provable equivalence.

The formula $\psi' :\equiv \psi \to \phi$ is uniform (as ϕ does not contain x). Let $G \subseteq \mathsf{Prop}(FL(\psi'))$ be as in Lemma 4.20, applied to ψ'; notice that $\mathsf{Prop}(FL(\psi'))$ is finite up to provable equivalence. Then put

$$G_\psi(\phi) = \{\chi(\top) \mid \chi(x) \in G, \vdash \chi(\top) \vee \chi(\bot)\}.$$

Now let $\vdash \rho \to \chi(\top)$ for some $\chi(x) \in G$ such that $\vdash \chi(\top) \vee \chi(\bot)$. To show that $\vdash \psi(\rho) \to \phi$, it suffices by construction of G and positivity of ψ to prove that $\vdash \chi(a) \wedge (\rho \to a)$ for some formula a: Then it follows that $\vdash (\psi(a) \to \phi) \wedge (\rho \to a)$, and hence by Lemma 4.19, that $\vdash \psi(\rho) \to \phi$. Since $\chi \wedge (\rho \to x)$ is uniform of depth 0, existence of such an a follows by Lemma 4.21 once we show that

$$\vdash (\chi(\top) \wedge (\rho \to \top)) \vee (\chi(\bot) \wedge (\rho \to \bot));$$

equivalently,

$$\vdash \chi(\top) \vee (\neg\rho \wedge \chi(\bot)).$$

By distributing disjunction over conjunction, this last formula is equivalent to

$$(\rho \to \chi(\top)) \wedge (\chi(\bot) \vee \chi(\top))$$

and hence provable by assumption.

Conversely, let $\vdash \psi(\rho) \to \phi$. By Lemma 4.20, there exists $\chi \in G$ such that $\chi(\rho)$ is provable. Since χ is uniform of depth 0, $\chi(\rho)$ is propositionally equivalent to its Boolean expansion $(\rho \to \chi(\top)) \wedge (\neg\rho \to \chi(\bot))$ and hence propositionally entails $\chi(\bot) \vee \chi(\top)$, which is therefore provable. That is, we have $\chi(\top) \in G_\psi(\phi)$ and $\vdash \rho \to \chi(\top)$, as required.

This proves that $\psi^{A(\mathcal{L}_\sharp)}$ is O-adjoint. From the above description of G_ψ, one sees immediately that $\psi^{A(\mathcal{L}_\sharp)}$ is in fact a finitary O-adjoint, as $\mathrm{Prop}(FL(\psi'))$ is closed under G_ψ and finite up to provable equivalence. □

Example 4.23 (O-adjointness). We illustrate the construction of $G_\psi(\phi)$ in the proof of Theorem 4.22 for the relational base case (Example 2.3(1)), for simplicity without propositional atoms, so $\Lambda = \{\Box\}$. We can generally restrict attention to the case that ϕ is a clause over $\Lambda(\mathcal{F}_\sharp)$, i.e., ϕ has the form

$$\phi = \bigwedge_{i\in I} \Box\phi_i \to \bigvee_{j\in J} \Box\chi_j,$$

where I and J are finite index sets. We reuse the notation in the proof of Theorem 4.22 but instead of G allow ourselves to use any set $G_0 \subseteq G$ such that $G = \{\chi \mid \exists \chi' \in G_0.\ \models \chi \to \chi'\}$. We start off with

$$\psi = \Box x.$$

By the description of derivable rules given in Example 3.10(1), we can take

$$G_0 = \{x \to (\textstyle\bigwedge_{i\in I} \phi_i \to \chi_j) \mid j \in J\},$$

and hence (replacing, as announced, G with G_0 in the construction of $G_\psi(\phi)$)

$$G_{\Box x}(\phi) = \{\chi(\top) \mid \chi \in G_0, \vdash \chi(\top) \vee \chi(\bot)\} = \{\textstyle\bigwedge_{i\in I} \phi_i \to \chi_j \mid j \in J\}.$$

As a second example, take

$$\psi = \Diamond x = \neg\Box\neg x.$$

Again by the description of derivable rules, we can now take

$$G_0 = \{\textstyle\bigwedge_{i\in I} \phi_i \to \neg x\} \cup \{\textstyle\bigwedge_{i\in I} \phi_i \to \chi_j \mid j \in J\}. \tag{7}$$

Of course, if any formula $\bigwedge_{i\in I} \phi_i \to \chi_j$ is derivable, then ϕ is derivable, so we can just take $G_\psi(\phi) = \{\top\}$. Otherwise, we can omit the second component of the union in Equation (7) for purposes of constructing $G_\psi(\phi)$ and obtain

$$G_{\Diamond x}(\phi) = \{\neg \textstyle\bigwedge_{i\in I} \phi_i\}.$$

In both cases, $G_{\Diamond x}(\phi)$ is a singleton, so $\Diamond$ is actually left adjoint.

Using uniform formulas as a base, we can now exploit some known closure properties of finitary O-adjoints (Santocanale 2008).

Definition 4.24. The set of *permissible* modal fixpoint schemes is the closure of the set of uniform modal fixpoint schemes under disjunction, conjunction with modal fixpoint schemes not containing the recursion variable x, and substitution for x, the latter in the sense that if $\gamma(x)$ and δ are permissible, then $\gamma(\delta)$ is permissible.

Example 4.25 (Permissibility). The modal fixpoint scheme

$$\Box(\Diamond x \vee (p \wedge \Diamond\Diamond x)) \wedge \Diamond(\Diamond x \vee (p \wedge \Diamond\Diamond x))$$

is permissible, as it arises by substituting the disjunction $\Diamond x \vee (p \wedge \Diamond\Diamond x)$ of the (trivially) uniform modal fixpoint schemes $\Diamond x$ and $p \wedge \Diamond\Diamond x$ into the uniform modal fixpoint scheme $\Box x \wedge \Diamond x$.

Corollary 4.26. *If $\gamma \in \Gamma$ is permissible, then γ is a finitary O-adjoint and hence constructive.*

Proof. The set of finitary $\mathcal{O}$-adjoints is closed under joins, meets with constants, and composition (Santocanale 2008). □

From now on, *we require that every* $\gamma \in \Gamma$ *is permissible* and hence constructive; a flat coalgebraic fixpoint logic satisfying this requirement will be called *permissible*. All fixpoint operators mentioned in Example 2.3 are based on permissible fixpoint schemes (in fact, on uniform ones).

Remark 4.27. The sufficient criterion for $\mathcal{O}$-adjointness given by Santocanale and Venema (2010) is that modal fixpoint schemes be *harmless*, which in the single-modality case means that modal fixpoint schemes γ, δ are generated by the grammar

$$\gamma, \delta ::= \top \mid x \mid \gamma \vee \delta \mid \chi \wedge \gamma \mid \textstyle\bigwedge_{i=1}^{k} \Diamond\gamma_i \mid \Box\gamma,$$

where χ is a modal formula not mentioning the recursion variable x (but possibly mentioning parameter variables). This notion is incomparable to permissibility in the sense of Definition 4.24; e.g., $\Diamond x \wedge \Diamond\Diamond x$ is harmless but not permissible, and $\Diamond\Box x \wedge \Box\Diamond x$ is permissible (in fact, uniform) but not harmless. We leave it as an open problem to find a sufficient criterion for $\mathcal{O}$-adjointness that subsumes both permissibility and harmlessness. The fixpoint schemes generating the CTL operators are both harmless and permissible.

5 THE MODEL CONSTRUCTION

We proceed to describe a model construction that uses sets of *timed-out formulas* as states; a timed-out formula has some of the fixpoints that appear in it annotated with time-outs indicating how often they need to be unfolded. Our time-outs are related to Kozen's μ-counters (Kozen 1983) but, as indicated, are integrated into formulas rather than maintained independently in a tableau construction. The use of time-outs is justified by constructivity of fixpoint operators as proved in the previous section.

Since only some of the fixpoint subformulas contained in a state will be annotated with finite time-outs, only one implication of the truth lemma will hold (every state satisfies the timed-out formulas it contains but not conversely—the model we construct will be finite, so every fixpoint will be satisfied with some time-out, which the state may fail to specify). Consequently, in the inductive proof of the truth lemma, the step for negation would fail. We therefore work with formulas in negation normal form, defined in detail as follows. A modal fixpoint scheme is in *negation normal form (NNF)* if it can be generated by the grammar

$$\gamma, \delta ::= \bot \mid \top \mid x \mid p \mid \gamma \wedge \delta \mid \gamma \vee \delta \mid \heartsuit\gamma \mid \overline{\heartsuit}\gamma \qquad (p \in V, \heartsuit \in \Lambda)$$

(recall that $\overline{\heartsuit}$ abbreviates $\neg\heartsuit\neg$). We can clearly transform every modal fixpoint scheme into a provably equivalent one in NNF (recall that modal fixpoint schemes are positive in all variables) and therefore assume from now on that Γ consists of modal fixpoint schemes in NNF (this does not substantially affect the syntax as the modal fixpoint schemes just serve as indices of fixpoint operators). A formula ϕ is in *NNF* if it can be generated by the grammar

$$\phi, \psi ::= \bot \mid \top \mid \phi \wedge \psi \mid \phi \vee \psi \mid \heartsuit\phi \mid \overline{\heartsuit}\phi \mid \sharp_\gamma\phi \mid \flat_\gamma\phi.$$

The dual $\overline{\gamma}$ of a modal fixpoint scheme γ as defined in Section 2 is clearly equivalent to the modal fixpoint scheme obtained from γ by swapping $\wedge$ with $\vee$, $\top$ with $\bot$, and $\heartsuit$ with $\overline{\heartsuit}$, and we regard $\overline{\gamma}$ as being syntactically defined in this way from now on; e.g., if $\Diamond = \overline{\Box}$, then the dual of $\gamma = p \vee \Box x$ is $\overline{\gamma} = p \wedge \Diamond x$. To show that we can transform every formula into NNF, it suffices as usual to show that we can implement negation on NNFs; this is by the standard procedure of pushing negation inside through the other connectives using provable equivalences. In particular, the equivalence

$$\neg\sharp_\gamma\phi \leftrightarrow \flat_{\overline{\gamma}}\neg\phi$$

is trivially provable, as the right-hand side is just an abbreviation for $\neg\sharp_\gamma\neg\neg\phi$. Summing up, we have

Lemma 5.1. *Every formula is provably equivalent to a formula in NNF.*

It remains to adapt the notion of Fischer-Ladner closure; as no confusion is likely and the changes are rather inessential, we continue to use the same term:

Definition 5.2. A set Σ of formulas in NNF is *Fischer-Ladner closed* if Σ is closed under subformulas and fixpoint unfolding, i.e., whenever $\star_\gamma\phi \in \Sigma$ for $\star \in \{\sharp, \flat\}$, then $\gamma(\phi, \star_\gamma\phi) \in \Sigma$.

The analogue of Lemma 4.17 remains true for the modified definition; so from now on we fix a finite Fischer-Ladner closed set Σ of formulas in NNF. We proceed to introduce the announced notion of time-out:

Definition 5.3. The set of *timed-out formulas* ϕ, ψ is generated by the grammar

$$\phi, \psi ::= \bot \mid \top \mid \phi \wedge \psi \mid \phi \vee \psi \mid \heartsuit\phi \mid \overline{\heartsuit}\phi \mid \sharp^\kappa_\gamma\rho \mid \flat_{\overline{\gamma}}\rho \qquad (\kappa \in \omega + 1, \rho \in \mathcal{L}_\sharp),$$

where $\gamma \in \Gamma, \heartsuit \in \Lambda$, subject to the restriction that ϕ is a timed-out formula only in case ϕ has at most one subformula of the form $\sharp^\kappa_\gamma\chi$ with $\kappa < \omega$ (which, however, may occur any number of times), and for this $\sharp^\kappa_\gamma\chi$,

— $\sharp^\omega_\gamma\chi$ is not a subformula of ϕ; and
— whenever $\sharp^\omega_\delta\rho$ is a subformula of ϕ, then $\sharp_\delta\rho$ is a subformula of χ.

In this case, we define the *time-out* $\tau(\phi)$ of ϕ to be κ, and $\tau(\phi) = \omega$ otherwise (i.e., if ϕ does not contain any subformula of the form $\sharp^\kappa_\gamma\chi$ with $\kappa < \omega$). The time-out gives the number of steps left until satisfaction of the eventuality $\sharp_\gamma\chi$, with time-out ω signifying an unspecified number of steps (note that time-outs are never associated with $\flat$-formulas).

We define two translations $(-)^s$ and $(-)^t$ of timed-out formulas into $\mathcal{F}_\sharp$, given by commutation with Boolean and modal operators, $(\flat_{\overline{\gamma}}\rho)^s = (\flat_{\overline{\gamma}}\rho)^t = \flat_{\overline{\gamma}}\rho$, and

$$(\sharp^\omega_\gamma\rho)^s = \sharp_\gamma\rho \qquad (\sharp^i_\gamma\rho)^s = \gamma(\rho)^i(\bot) \quad (i < \omega) \qquad (\sharp^\kappa_\gamma\rho)^t = \sharp_\gamma\rho.$$

Thus, s unfolds fixpoints as prescribed by their time-outs, and t just removes time-outs. In the converse direction, we define a trivial translation $(-)^\omega$ that converts a formula $\phi \in \mathcal{F}_\sharp$ into a timed-out formula by adding ω as a time-out at the outermost level; that is, $(-)^\omega$ is given by commutation with Boolean and modal operators, $(\flat_{\overline{\gamma}}\rho)^\omega = \flat_{\overline{\gamma}}\rho$, and $(\sharp_\gamma\rho)^\omega = \sharp^\omega_\gamma\rho$. All translations extend to sets A of formulas, e.g., $A^s = \{\phi^s \mid \phi \in A\}$. For timed-out formulas ϕ, ψ, we put $\phi \preceq \psi$ iff $\phi^t = \psi^t$ and $\tau(\phi) \leq \tau(\psi)$. That is, $\phi \preceq \psi$ iff ϕ is the same as ψ up to possible decrease of the time-out. Given a set Σ of formulas, a timed-out formula ϕ is a *timed-out Σ-formula* if $\phi^t \in \Sigma$.

Example 5.4 (Timed-out formulas). A formula of the shape $\phi = \sharp^3_\gamma\sharp_\delta a \wedge \Box\sharp^3_\gamma\sharp_\delta a \wedge \sharp^\omega_\delta a$ is a timed-out formula (with time-out $\tau(\phi) = 3$): $\sharp^3_\gamma\sharp_\delta a$ is the only subformula of ϕ with a finite time-out, and the only other $\sharp$-formula appearing in ϕ outside the scope of a fixpoint, $\sharp^\omega_\delta a$, appears (without time-out) as a subformula in $\sharp^3_\gamma\sharp_\delta a$.

Note that unlike formulas in NNF, timed-out formulas are not closed under negation, because time-outs can only appear on least fixpoints. The point of the definition of timed-out formulas is that every standard formula ϕ has at most one candidate subformula at which one can insert a time-out, namely the greatest element under the subformula ordering among the subformulas of ϕ that are $\sharp$-formulas and appear in ϕ outside the scope of other fixpoints, if such a greatest element exists. This enables the simple definition of $\preceq$, which by the preceding discussion has the following property.

Lemma 5.5. *For every formula ϕ, the preimage of ϕ under the translation t is non-empty (specifically, contains ϕ^ω) and well-ordered by $\preceq$.*

At the same time, timed-out formulas are stable under unfolding:

Lemma 5.6. *For every fixpoint formula $\rho \in \mathcal{F}_\sharp$ and every $\kappa \in \omega + 1$, $\gamma(\rho^\omega, \sharp^\kappa_\gamma \rho)$ is a timed-out formula.*

Proof. As ρ is a (standard) formula, it is clear that $\gamma(\rho^\omega, \sharp^\kappa_\gamma \rho)$ cannot contain formulas of the form $\sharp^\lambda_\delta \rho$ with $\lambda < \omega$ other than $\sharp^\kappa_\gamma \rho$. By well-foundedness of the subformula relation, $\gamma(\rho^\omega, \sharp^\kappa_\gamma \rho)$ cannot contain $\sharp^\omega_\gamma \rho$. Finally, since γ itself does not contain any fixpoint operators, the only way $\sharp$-subformulas can arise in $\gamma(\rho^\omega, \sharp^\kappa_\gamma \rho)$ is as subformulas of ρ. □

We have the expected relationship between the translations regarding provable entailment:

Lemma 5.7.

(1) *For every timed-out formula ϕ, $\vdash \phi^s \to \phi^t$.*
(2) *If $\phi \preceq \psi$ for timed-out formulas ϕ, ψ, then $\vdash \phi^s \to \psi^s$.*

Proof. By iterated application of the unfolding axiom and monotonicity of modal fixpoint schemes. □

States of the model will be sets of formulas satisfying a timed-out version of the usual expandedness requirement.

Definition 5.8 (Atom, timed-out atom). As usual, a Σ-*atom* is a maximally consistent subset of Σ. A *timed-out* Σ-*atom* is a maximal set A of timed-out Σ-formulas such that A^s is consistent.

Lemma 5.9. *Every timed-out Σ-atom is upwards closed under $\preceq$.*

To prove finiteness of the model we construct, we use the fact that finite product orderings $(\omega + 1)^k$ are well-quasi-orders and in particular have only finite anti-chains (Laver 1976).

Lemma 5.10 (Timed-out Lindenbaum lemma).

(1) *The set of timed-out Σ-atoms is finite.*
(2) *For every set A_0 of timed-out Σ-formulas such that A_0^s is consistent, there exists a timed-out Σ-atom A such that $A_0 \subseteq A$.*
(3) *For every consistent subset $C \subseteq \Sigma$, there exists a timed-out Σ-atom A such that $C \subseteq A^t$.*

Proof. Claim (3) follows from (2) for $A_0 = C^\omega$. We prove (1) and (2) in one go. The set of Σ-atoms is finite, so for every set A_0 as in the statement there exists a Σ-atom C such that $A_0^t \subseteq C$. Let $\mathfrak{A}$ be the set of sets A of timed-out Σ-formulas such that A^s is consistent, $A^t = C$, and A is upwards closed under $\preceq$; then the timed-out Σ-atoms whose t-image is C are the maximal elements of $\mathfrak{A}$. Now every element $A \in \mathfrak{A}$ is uniquely determined by the C-tuple $\iota(A)$ of minimal time-outs it induces (explicitly, for $\phi \in C$, the ϕ-component of $\iota(A)$ is $\min\{\tau(\phi') \mid \phi' \in A, (\phi')^t = \phi\}$), and for $A, B \in \mathfrak{A}$, $A \subseteq B$ iff $\iota(B) \leq \iota(A)$ in the componentwise ordering. Thus, $\mathfrak{A}$ is order-isomorphic to a subset of the finite power $(\omega + 1)^C$ of the well-ordering $\omega + 1$. It follows from the theory of *well-quasi-orders* that $(\omega + 1)^C$ is a well-quasi-order, which means that for every subset F of $(\omega + 1)^C$, every element of F is above one of finitely many minimal elements of F (Laver 1976); applying this to $\mathfrak{A}$ proves (1) and (2) (for (1), recall additionally that the set of Σ-atoms is finite). □

As usual, the proof of the truth lemma will depend on a set of Hintikka-like properties:

Lemma 5.11. *If A is a timed-out Σ-atom, then*

(1) if $\phi \wedge \psi \in A$, then $\phi, \psi \in A$;
(2) if $\phi \vee \psi \in A$, then $\phi \in A$ or $\psi \in A$;
(3) $\bot \notin A$;
(4) if $\sharp_\gamma^\omega \phi \in A$, then $\sharp_\gamma^\kappa \phi \in A$ for some $\kappa < \omega$;
(5) for $\sharp_\gamma \phi \in \Sigma$, $\sharp_\gamma^\kappa \phi \in A$ iff $\gamma(\phi, \sharp_\gamma^{\kappa-1}\phi) \in A$;
(6) for $\sharp_\gamma \phi \in \Sigma$, $\flat_\gamma \phi \in A$ iff $\gamma(\phi, \flat_\gamma \phi) \in A$.

Proof. 1: If $\phi \wedge \psi \in A$, then $(A \cup \{\phi\})^s$ and $(A \cup \{\psi\})^s$ are consistent, so $\phi, \psi \in A$ by maximality.
2: If $\phi \vee \psi \in A$, then either $(A \cup \{\phi\})^s$ or $(A \cup \{\psi\})^s$ is consistent, and hence one of ϕ and ψ is in A by maximality.
3: Clear.
4: If $\sharp_\gamma^\omega \phi \in A$, then by Lemma 4.4 and by finiteness of A, there is $\kappa < \omega$ such that $(A \cup \{\sharp_\gamma^\kappa \phi\})^s$ is consistent; hence $\sharp_\gamma^\kappa \phi \in A$ by maximality.
5: Note that both $\sharp_\gamma^\kappa \phi$ and $\gamma(\phi, \sharp_\gamma^{\kappa-1}\phi)$ are timed-out Σ-formulas (the latter by Fischer-Ladner closedness of Σ and Lemma 5.6), and their s-translations are syntactically equal. Therefore if, e.g., $\sharp_\gamma^\kappa \phi \in A$, then $(A \cup \{\gamma(\phi, \sharp_\gamma^{\kappa-1}\phi)\})^s$ is consistent, so $\gamma(\phi, \sharp_\gamma^{\kappa-1}\phi) \in A$ by maximality. The converse implication is similar.
6: Similar to 5. □

We denote by $\mathcal{T}_\Sigma$ the (by Lemma 5.10, finite) set of timed-out Σ-atoms and proceed to construct a model with carrier $\mathcal{T}_\Sigma$. We need to construct a coalgebra rather than just a relational structure; this coalgebra should adequately implement the formulas contained in the states. This requirement is encapsulated in the notion of *coherence*:

Definition 5.12. A coalgebra structure ξ on $\mathcal{T}_\Sigma$ is *coherent* if for every $A \in \mathcal{T}_\Sigma$ and every timed-out Σ-formula $\spadesuit\rho$ where $\spadesuit$ is either a modal operator $\heartsuit \in \Lambda$ or a dual modal operator $\overline{\heartsuit}$,

$$\xi(A) \in [\![\spadesuit]\!]_{\mathcal{T}_\Sigma}(\hat{\rho}) \text{ whenever } \spadesuit\rho \in A,$$

where $\hat{\rho} = \{B \in \mathcal{T}_\Sigma \mid \rho \in B\}$.

Example 5.13 (Timed-out atoms, coherence). We exemplify the intended construction of models from timed-out atoms in the basic relational setting, specifically in serial CTL (Example 2.3(1)), continuing to write $\Box$ in place of AX for brevity. We slightly adapt an example formula from work on global caching (Hausmann and Schröder 2015): take

$$\psi = AG(F_1 \wedge F_2 \wedge \phi),$$

where

$$F_1 = AF\, a \qquad F_2 = AF\, b$$

$$\phi = (a \vee b \vee c) \wedge (\neg a \vee \neg b) \wedge (\neg b \vee \neg c) \wedge (\neg a \vee \neg c) \wedge$$
$$(a \to \Box c) \wedge (b \to \Box c) \wedge (c \to \Box(a \vee b)),$$

and let $\Sigma = FL(\psi)$. In words, $AG\,\phi$ ensures that every reachable state satisfies exactly one of a, b, c, and that c alternates with $a \vee b$ along every path. Of course, every state reachable from a state satisfying ψ must itself satisfy ψ. There are only three possible Σ-*atoms* A_a^0, A_b^0, A_c^0 containing ψ. They all contain $\Box\psi, F_1, F_2, \phi, \Box F_1, \Box F_2$, and are then determined by

$$A_a^0 \ni a, \neg b, \neg c, \Box c$$
$$A_b^0 \ni \neg a, b, \neg c, \Box c$$
$$A_c^0 \ni \neg a, \neg b, c, \Box(a \vee b).$$

This shows that constructing a model with the Σ-atoms as states will not work: Since both A^0_a and A^0_b must be reachable from A^0_c, the only candidate has the shape

$$\begin{tikzcd} A^0_a \arrow[r, shift left] & A^0_c \arrow[l, shift left] \arrow[r, shift left] & A^0_b \arrow[l, shift left] \end{tikzcd},$$

which, however, fails to satisfy $AF\,a$ and $AF\,b$ at A^0_c.

Contrastingly, there are four *timed-out* Σ-atoms A_a, A_b, A^1_c, A^2_c containing ψ. Their images under t are $A^t_a = A^0_a$, $A^t_b = A^0_b$, and $(A^1_c)^t = (A^2_c)^t = A^0_c$. They are then determined by the minimal time-outs on F_1 and F_2 they contain, which we just write on the operator $AF = \sharp_{p\vee\Box x}$:

$$\begin{aligned} A_a &\ni AF^0 a, AF^2 b \\ A_b &\ni AF^2 a, AF^0 b \\ A^1_c &\ni AF^1 a, AF^3 b \\ A^2_c &\ni AF^3 a, AF^1 b. \end{aligned}$$

Note that, e.g., $\Box F_1$ and $\Box F_2$ also receive time-outs (unlike ψ, as in this formula, the least fixpoints F_1, F_2 are in the scope of a greatest fixpoint). For example we have $A^1_c \ni \Box AF^0 a, \Box AF^2 b$. Observe that while ψ does allow paths visiting a (alternating with c) several times before reaching b, such paths will not be realized in the model we are constructing, because the associated sets of timed-out formulas are non-maximal, i.e., fail to contain minimal time-outs.

From these timed-out Σ-atoms, we build the coherent model

$$\begin{tikzcd} & A^1_c \arrow[dl] & \\ A_a \arrow[dr] & & A_b \arrow[ul] \\ & A^2_c \arrow[ur] & \end{tikzcd}$$

For example, since $\Box(a \vee b) \in A^1_c$, coherence at A^1_c requires that $\xi(A^1_c) \in [\![\Box]\!](\widehat{a \vee b})$, i.e., $\xi(A^1_c) \subseteq \widehat{a \vee b}$, where ξ is the structure map of the model, understood as a $\mathcal{P}^*$-coalgebra. This holds because $\xi(A^1_c) = \{A_a\}$ and $\widehat{a \vee b}$ consists of the timed-out atoms containing $a \vee b$, i.e., $\widehat{a \vee b} = \{A_a, A_b\}$.

In general, existence of a coherent coalgebra structure relies on one-step completeness:

Lemma 5.14 (Existence lemma). *There exists a coherent coalgebra structure on $\mathcal{T}_\Sigma$.*

Proof. The proof follows the same pattern as the one for the fixpoint-free case (Schröder 2007) (see also the discussion in Section 3). We can construct the coalgebra structure ξ pointwise. So let $A \in \mathcal{T}_\Sigma$; in the notation of Definition 5.12, we have to show that there exists $t \in T\mathcal{T}_\Sigma$ such that $t \in [\![\spadesuit]\!]_{\mathcal{T}_\Sigma}(\hat{\phi})$ whenever $\spadesuit\phi \in A$. With a view to deriving a contradiction, assume the contrary. Then $T\mathcal{T}_\Sigma, \tau \models \psi$, where

$$\begin{aligned} \psi &= \bigvee\nolimits_{\heartsuit\rho\in A} \neg\heartsuit a_{\heartsuit\rho} \vee \bigvee\nolimits_{\overline{\heartsuit}\rho\in A} \heartsuit b_{\overline{\heartsuit}\rho} \\ \tau(a_{\heartsuit\rho}) &= \hat{\rho} \\ \tau(b_{\overline{\heartsuit}\rho}) &= \mathcal{T}_\Sigma - \hat{\rho} \end{aligned}$$

for pairwise distinct propositional variables $a_{\heartsuit\rho}, b_{\overline{\heartsuit}\rho}$. By one-step completeness, it follows that ψ is provable over $\mathcal{T}_\Sigma, \tau$, so by Lemma 3.8 there is an $\mathcal{R}$-derivable monotone one-step rule $\phi/\chi \in \mathcal{R}$ such that $\chi \vdash_{PL} \psi$ and $\mathcal{T}_\Sigma, \tau \models \phi$. Now let θ be the substitution defined by $\theta(a_{\heartsuit\rho}) = \rho^s$ and $\theta(b_{\overline{\heartsuit}\rho}) = \neg(\rho^s)$. Then A^s propositionally entails $\neg\psi\theta$. Since A^s is consistent, we are done once we show that $\psi\theta$ is provable. To this end, it suffices to show that $\phi\theta$ is provable. Assume the contrary, i.e., $\neg\phi\theta$ is consistent. Then there is a conjunctive clause ϕ_0 in the disjunctive normal form of $\neg\phi$ such that $\phi_0\theta$ is consistent. Since ϕ/χ is monotone and $\chi \vdash_{PL} \psi$ just means that χ is contained in ψ, ϕ is negative in the $a_{\heartsuit\rho}$ and positive in the $b_{\overline{\heartsuit}\rho}$; thus, ϕ_0 is positive in the $a_{\heartsuit\rho}$ and negative in the $b_{\overline{\heartsuit}\rho}$. Therefore, $\phi_0\theta$ is, after removing double negations, of the form $\bigwedge A^s$ for a set A of timed-out Σ-formulas. By the timed-out Lindenbaum lemma (Lemma 5.10), A is contained in a timed-out Σ-atom $B \in \mathcal{T}_\Sigma$. Then $B \in [\![\phi_0]\!]\tau$, so $B \notin [\![\phi]\!]\tau$, in contradiction to $\mathcal{T}_\Sigma, \tau \models \phi$. □

Next, we establish that a coherent coalgebra does what we expect:

Lemma 5.15 (Truth lemma). *Let ξ be a coherent coalgebra structure on $\mathcal{T}_\Sigma$. If $A \in \mathcal{T}_\Sigma$ and $\phi \in A$, then $A \models_{(\mathcal{T}_\Sigma,\xi)} \phi^s$.*

Proof. Induction over timed-out Σ-formulas ϕ using the lexicographic product of the subterm ordering on ϕ^t with $\preceq$ as the induction measure. The case $\phi = \top$ is trivial. The steps for $\bot$, $\wedge$, and $\vee$ are taken care of by Lemma 5.11. The step for modal operators $\heartsuit$ or $\overline{\heartsuit}$ is by coherence and monotonicity.

Next, we discharge the case $\phi = \sharp^\kappa_\gamma\psi$. By Lemma 5.11, we can assume $\kappa < \omega$, and have $\gamma(\psi, \sharp^{\kappa-1}_\gamma\psi) \in A$. We prove by a further induction on modal fixpoint schemes δ that $\delta(\psi, \sharp^{\kappa-1}_\gamma\psi) \in A$ implies that $A \models_{(\mathcal{T}_\Sigma,\xi)} (\delta(\psi, \sharp^{\kappa-1}_\gamma\psi))^s$. The case for the parameter variable is discharged by the inductive hypothesis applied to ψ, as ψ^t is a proper subterm of ϕ^t, while the case for the recursion variable is discharged by the inductive hypothesis applied to $\sharp^{\kappa-1}_\gamma\psi$ (which is strictly below $\sharp^\kappa_\gamma\psi$ w.r.t. $\preceq$ since $\kappa < \omega$). The cases for Boolean operations and modal operators are as in the outer induction. This finishes the inner induction, so that $A \models_{(\mathcal{T}_\Sigma,\xi)} \gamma(\psi, \sharp^{\kappa-1}_\gamma\psi)$ and hence $A \models_{(\mathcal{T}_\Sigma,\xi)} \sharp^\kappa_\gamma\psi$.

Finally, the case $\phi = \flat_\gamma\psi$ is discharged by coinduction. For timed-out Σ-formulas ρ, we put $\mathcal{T}_\Sigma(\rho) = \{A \in \mathcal{T}_\Sigma \mid \rho \in A\}$ (this is like $\hat{\rho}$ in Definition 5.12; we change the notation for typographical reasons as we will need to apply $\mathcal{T}_\Sigma$ to long formulas). As $[\![\flat_\gamma\psi]\!]_{(\mathcal{T}_\Sigma,\xi)}$ is a greatest fixpoint, it suffices to prove that $\mathcal{T}_\Sigma(\flat_\gamma\psi)$ is semantically a postfixpoint of $\gamma(\psi)$, i.e.,

$$\mathcal{T}_\Sigma(\flat_\gamma\psi) \subseteq [\![\gamma(\psi)]\!](\mathcal{T}_\Sigma(\flat_\gamma\psi)). \tag{8}$$

To begin, we prove by induction on modal fixpoint schemes δ that for all timed-out Σ-formulas χ,

$$\mathcal{T}_\Sigma(\delta(\psi,\chi)) \subseteq [\![\delta(\psi)]\!](\mathcal{T}_\Sigma(\chi)). \tag{9}$$

The cases for $\bot$, $\top$, and the recursion variable x are clear. The case for the parameter variable p is discharged by the outer inductive hypothesis applied to ψ. The cases for $\wedge$ and $\vee$ are by Lemma 5.11; e.g., we have

$$\begin{aligned}
&\mathcal{T}_\Sigma((\delta_1 \vee \delta_2)(\psi,\chi)) \\
&\subseteq \mathcal{T}_\Sigma(\delta_1(\psi,\chi)) \cup \mathcal{T}_\Sigma(\delta_2(\psi,\chi)) && \text{(Lemma 5.11)} \\
&\subseteq [\![\delta_2(\psi)]\!](\mathcal{T}_\Sigma(\chi)) \cup [\![\delta_1(\psi)]\!](\mathcal{T}_\Sigma(\chi)) && \text{(induction)} \\
&= [\![(\delta_2 \vee \delta_1)(\psi)]\!](\mathcal{T}_\Sigma(\chi)).
\end{aligned}$$

Finally, the case for modal operators is by coherence: for $\spadesuit$ being either a modal operator $\heartsuit \in \Lambda$ or a dual modal operator $\overline{\heartsuit}$, we have

$$\begin{aligned}
&\mathcal{T}_\Sigma(\spadesuit\delta(\psi,\chi)) \\
&\subseteq \xi^{-1}[[\![\spadesuit]\!]_{\mathcal{T}_\Sigma}(\mathcal{T}_\Sigma(\delta(\psi,\chi)))] && \text{(coherence)} \\
&\subseteq \xi^{-1}[[\![\spadesuit]\!]_{\mathcal{T}_\Sigma}([\![\delta(\psi)]\!]\mathcal{T}_\Sigma(\chi))] && \text{(induction, monotonicity)} \\
&= [\![\spadesuit\delta(\psi)]\!](\mathcal{T}_\Sigma(\chi)).
\end{aligned}$$

by (9), we reduce our goal (8) to

$$\mathcal{T}_\Sigma(\flat_\gamma \phi) \subseteq \mathcal{T}_\Sigma(\gamma(\psi, \flat_\gamma \psi)),$$

which holds by Lemma 5.11. □

In summary, we have proved (weak) completeness of the Kozen-Park axiomatization:

Theorem 5.16 (Weak completeness). *If Γ is permissible and $\mathcal{R}$ is one-step complete, then the logic $\mathcal{L}_\sharp$ is weakly complete over finite models, i.e., every consistent formula has a finite model.*

Proof. Let $\phi \in \mathcal{F}_\sharp$ be consistent; by Lemma 5.1 we can assume that ϕ is in NNF. We have to show that ϕ is satisfiable in a finite T-coalgebra. This follows from the preceding lemmas by the usual pattern: Let Σ be the least Fischer-Ladner closed set of formulas in NNF containing ϕ. Then ϕ is contained in the t-image of some timed-out Σ-atom by the timed-out Lindenbaum lemma (Lemma 5.10), and, hence, by the truth lemma (Lemma 5.15) satisfied in a coherent coalgebra on the set $\mathcal{T}_\Sigma$ of timed-out Σ-atoms, which exists by Lemma 5.14; the set $\mathcal{T}_\Sigma$ is finite by Lemma 5.10. □

Remark 5.17. The full coalgebraic μ-calculus is known to have the finite model property (every satisfiable formula has a finite model) (Cîrstea et al. 2011). Since soundness (Theorem 3.5) holds over unrestricted models, Theorem 5.16 reproves the finite model property for the single-variable fragment of the coalgebraic μ-calculus.

Strong completeness in the sense that every consistent *set* of formulas is satisfiable fails already in the basic relational instance; e.g., the set

$$\{EF\, a\} \cup \{\Box^n \neg a \mid n \geq 0\},$$

where $\Box^n$ stands for n consecutive boxes, is consistent (as all its finite subsets are satisfiable, hence consistent) but unsatisfiable.

We enumerate a few concrete instances of the completeness result:

Example 5.18 (Completeness). We obtain that the Kozen-Park axiomatization (in combination with the modal rules and propositional reasoning) is complete for the following logics.

(1) All permissible flat fragments of the standard relational μ-calculus, interpreted over unrestricted, serial, or deterministic Kripke models. We thus recover the known completeness results for LTL (Gabbay et al. 1980) and serial and non-serial CTL (Emerson and Halpern 1985) but also for logics featuring operators outside CTL as discussed in Example 2.3(1) (note that the modal fixpoint scheme $p \wedge \Box\Box x$ appearing in Wolper's *even* operator is trivially uniform, as there is only a single occurrence of the recursion variable x). Non-serial CTL is already covered by the generic results of Santocanale and Venema (2010).

(2) All permissible flat fragments of the graded μ-calculus, including ones featuring the operators "the current state is the root of a finite binary tree all whose leaves satisfy ... ," "...holds somewhere on every infinite $k + 1$-ary tree starting at the current state," and "the current state is the root of a finite binary tree all whose leaves are at even distance from the root and satisfy ..." discussed in Example 2.3.2.

(3) All permissible flat probabilistic fixpoint logics, including ones featuring linear inequalities on probabilities.

(4) All permissible flat conditional fixpoint logics.

(5) All permissible flat fragments of the monotone μ-calculus and the serial monotone μ-calculus, including the star-nesting-free fragments of CPDL and game logic.

(6) All permissible flat fragments of the alternating-time μ-calculus, including alternating-time temporal logic ATL but also logics going beyond ATL, e.g., ones featuring the operator "…holds in all even states along any path" discussed in Example 2.3.5. In fact, ATL appears to be the only example outside the relational world for which a completeness result of this type was previously known (Goranko and van Drimmelen 2006).

In most of these examples except graded and probabilistic fixpoint logics, we actually obtain finite axiomatizability and locally finite axiomatizability in the case of graded fixpoint logics (Remark 3.11).

6 CONCLUSIONS

We have lifted the completeness theorem for flat modal fixpoint logics (Santocanale and Venema 2010) to the level of generality of coalgebraic logic. Specifically, we have given a Kozen-Park style axiomatization for fixpoint operators, and we have shown this axiomatization to be sound and complete under the conditions that (i) the defining formulas of the fixpoint operators satisfy a mild syntactic restriction, and (ii) the coalgebraic base logic is axiomatized by a one-step complete rule set. This result covers, e.g., probabilistic fixpoint logics and flat fragments of the monotone μ-calculus, the ambient fixpoint logic of Parikh's game logic (Parikh 1985), and concurrent PDL (Peleg 1987). Further instances include completeness of flat fragments of the graded μ-calculus (Kupferman et al. 2002), to our knowledge the first completeness result for any graded fixpoint logic, and completeness of flat fragments of the alternating-time μ-calculus (Alur et al. 2002), with ATL being apparently the only previously known example (Goranko and van Drimmelen 2006). In those examples that have finite modal similarity type, in particular for alternating-time fixpoint logics, the axiomatization we obtain is finite.

A core technical point in the proof was to show that essentially all monotone modal operators (including nested ones like $\Box\Box$, as long as the nesting depth is uniform) are finitary $\mathcal{O}$-adjoints in the sense of Santocanale (2008), and hence induce *constructive* fixpoint operators that can be approximated in ω steps in the Lindenbaum algebra. This has enabled a model construction using explicit time-outs for least fixpoint formulas in the spirit of the completeness proof for the aconjunctive fragment of the μ-calculus (Kozen 1983), which relies on a judicious definition of timed-out formula.

A remaining open problem is to extend the completeness result to larger fragments of the coalgebraic μ-calculus beyond the single variable fragment covered here. In collaboration with Enqvist and Seifan, the second author recently obtained completeness results for the full coalgebraic μ-calculi based on Moss' modality ∇ (Enqvist et al. 2016) (in the case of the coalgebra functor preserving weak pullbacks), and for the full coalgebraic μ-calculi based on a signature Λ of predicate liftings that admit a so-called disjunctive basis (Enqvist et al. 2017).

ACKNOWLEDGMENTS

The authors thank Dirk Pattinson for useful discussions and Erwin R. Catesbeiana for recommendations on inconsistency checking.

REFERENCES

Rajeev Alur, Thomas Henzinger, and Orna Kupferman. 2002. Alternating-time temporal logic. *J. ACM* 49, 5 (2002), 672–713.

Franz Baader, Diego Calvanese, Deborah McGuinness, Daniele Nardi, and Peter Patel-Schneider (Eds.). 2003. *The Description Logic Handbook*. Cambridge University Press.

Patrick Blackburn, Maarten de Rijke, and Yde Venema. 2001. *Modal Logic*. Cambridge University Press.

Brian Chellas. 1980. *Modal Logic*. Cambridge University Press.

Corina Cîrstea, Clemens Kupke, and Dirk Pattinson. 2011. EXPTIME Tableaux for the coalgebraic μ-calculus. *Log. Meth. Comput. Sci.* 7, 3:03 (2011), 1–33.

Corina Cîrstea and Dirk Pattinson. 2007. Modular construction of complete coalgebraic logics. *Theoret. Comput. Sci.* 388, 1–3 (2007), 83–108.

Giovanna D'Agostino and Albert Visser. 2002. Finality regained: A coalgebraic study of Scott-sets and multisets. *Arch. Math. Logic* 41, 3 (2002), 267–298.

Mads Dam. 1994. CTL* and ECTL* as fragments of the modal mu-calculus. *Theoret. Comput. Sci.* 126, 1 (1994), 77–96.

Francesco De Caro. 1988. Graded modalities II. *Stud. Log.* 47, 1 (1988), 1–10.

E. Allen Emerson. 1990. Temporal and modal logic. In *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics*. Elsevier, 995–1072.

E. Allen Emerson and Edmund Clarke. 1982. Using branching time temporal logic to synthesize synchronization skeletons. *Sci. Comput. Program.* 2, 3 (1982), 241–266.

E. Allen Emerson and Joseph Halpern. 1985. Decision procedures and expressiveness in the temporal logic of branching time. *J. Comput. System Sci.* 30, 1 (1985), 1–24.

E. Allen Emerson and Chin-Laung Lei. 1986. Efficient model checking in fragments of the propositional mu-calculus. In *Logic in Computer Science,* Vol. 86. IEEE, 267–278.

Sebastian Enqvist, Fatemeh Seifan, and Yde Venema. 2016. Completeness for coalgebraic fixpoint logic. In *Computer Science Logic*, Jean-Marc Talbotand Laurent Regnier (Eds.), Vol. 62. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 7:1–7:19.

Sebastian Enqvist, Fatemeh Seifan, and Yde Venema. 2017. Completeness for μ-calculi: a coalgebraic approach (unpublished).

Ronald Fagin and Joseph Halpern. 1994. Reasoning about knowledge and probability. *J. ACM* 41, 2 (1994), 340–367.

Kit Fine. 1972. In so many possible worlds. *Notre Dame J. Formal Logic* 13, 4 (1972), 516–520.

Nir Friedman and Joseph Halpern. 1994. On the complexity of conditional logics. In *Proceedings of the Conference on Knowledge Representation and Reasoning (KR'94)*, Jon Doyle, Erik Sandewall, and Pietro Torasso (Eds.). Morgan Kaufmann, 202–213.

Dov Gabbay, Amir Pnueli, Saharon Shelah, and Jonathan Stavi. 1980. On the temporal basis of fairness. In *Proceedings of the Conference on Principles of Programming Languages (POPL'80*), Paul Abrahams, Richard Lipton, and Stephen Bourne (Eds.). ACM Press, 163–173.

Valentin Goranko and Govert van Drimmelen. 2006. Complete axiomatization and decidability of alternating-time temporal logic. *Theoret. Comput. Sci.* 353, 1–3 (2006), 93–117.

Daniel Hausmann and Lutz Schröder. 2015. Global caching for the flat coalgebraic μ-calculus. In *Proceedings of the Conference on Temporal Representation and Reasoning (TIME'15)*, Fabio Grandi, Martin Lange, and Alessio Lomuscio (Eds.). IEEE, 121–130.

Aviad Heifetz and Philippe Mongin. 2001. Probabilistic logic for type spaces. *Games Econ. Behav.* 35, 1–2 (2001), 31–53.

Michael Huth and Marta Kwiatkowska. 1997. Quantitative analysis and model checking. In *Proceedings of the Conference on Logic in Computer Science (LICS'97)*. IEEE Computer Society, 111–122.

Dexter Kozen. 1983. Results on the propositional μ-calculus. *Theoret. Comput. Sci.* 27, 3 (1983), 333–354.

Dexter Kozen and Rohit Parikh. 1981. An elementary proof of the completeness of PDL. *Theoret. Comput. Sci.* 14, 1 (1981), 113–118.

Orna Kupferman, Ulrike Sattler, and Moshe Vardi. 2002. The complexity of the graded μ-calculus. In *Proceedings of the Conference on Automated Deduction (CADE'02)*, Andrei Voronkov (Ed.), Vol. 2392. Springer, 423–437.

Clemens Kupke and Dirk Pattinson. 2010. On modal logics of linear inequalities. In *Proceedings of the Conference on Advances in Modal Logic (AiML'10)*, Lev Beklemishev, Valentin Goranko, and Valentin Shehtman (Eds.). College Publications, 235–255.

Kim Larsen, Radu Mardare, and Bingtian Xue. 2016. Probabilistic μ-calculus: Decidability and complete axiomatization. In *Proceedings of the Foundations of Software Technology and Theoretical Computer Science (FSTTCS'16)*, Akash Lal, S. Akshay, Saket Saurabh, and Sandeep Sen (Eds.), Vol. 65. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 25:1–25:18.

Kim Larsen and Arne Skou. 1991. Bisimulation through probabilistic testing. *Inf. Comput.* 94, 1 (1991), 1–28.

Richard Laver. 1976. Well-quasi-orders and sets of finite sequences. *Math. Proc. Cambr. Philos. Soc.* 79, 1 (1976), 1–10.

David Lewis. 1969. *Convention, A Philosophical Study*. Harvard University Press.

Orna Lichtenstein and Amir Pnueli. 2000. Propositional temporal logics: Decidability and completeness. *Logic J. IGPL* 8, 1 (2000), 55–85.

Wanwei Liu, Lei Song, Ji Wang, and Lijun Zhang. 2015. A simple probabilistic extension of modal mu-calculus. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'15)*, Qiang Yangand Michael Wooldridge (Eds.). AAAI Press, 882–888.

Carroll Morgan and Annabelle McIver. 1997. A probabilistic temporal calculus based on expectations. In *Formal Methods Pacific 1997*. Springer, 4–22.

Robert Myers, Dirk Pattinson, and Lutz Schröder. 2009. Coalgebraic hybrid logic. In *Proceedings of the Foundations of Software Science and Computational Structures (FOSSACS'09)*, Luca de Alfaro (Ed.), Vol. 5504. Springer, 137–151.

Nicola Olivetti, Gian Pozzato, and Camilla Schwind. 2007. A sequent calculus and a theorem prover for standard conditional logics. *ACM Trans. Comput. Log* 8, 4 (2007), 22:1–51.

Rohit Parikh. 1985. The logic of games and its applications. *Ann. Discr. Math.* 24 (1985), 111–140.

Dirk Pattinson. 2003. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoret. Comput. Sci.* 309, 1–3 (2003), 177–193.

Dirk Pattinson. 2004. Expressive logics for coalgebras via terminal sequence induction. *Notre Dame J. Formal Logic* 45, 1 (2004), 19–33.

Dirk Pattinson and Lutz Schröder. 2008. Beyond rank 1: Algebraic semantics and finite models for coalgebraic logics. In *Proceedings of the Foundations of Software Science and Computational Structures (FOSSACS'08)*, Roberto Amadio (Ed.), Vol. 4962. Springer, 66–80.

Dirk Pattinson and Lutz Schröder. 2010. Cut elimination in coalgebraic logics. *Inf. Comput.* 208, 12 (2010), 1447–1468.

Marc Pauly. 2002. A modal logic for coalitional power in games. *J. Log. Comput.* 12, 1 (2002), 149–166.

David Peleg. 1987. Concurrent dynamic logic. *J. ACM* 34, 2 (1987), 450–479.

Vaughan Pratt. 1976. Semantical considerations on floyd-hoare logic. In *Proceedings of the Foundations of Computer Science (FOCS'76)*. IEEE, 109–121.

Luigi Santocanale. 2008. Completions of μ-algebras. *Ann. Pure Appl. Logic* 154, 1 (2008), 27–50.

Luigi Santocanale and Y. de Venema. 2010. Completeness for flat modal fixpoint logics. *Ann. Pure Appl. Logic* 162, 1 (2010), 55–82.

Lutz Schröder. 2007. A finite model construction for coalgebraic modal logic. *J. Log. Algebr. Prog.* 73, 1–2 (2007), 97–110.

Lutz Schröder and Dirk Pattinson. 2009. PSPACE bounds for rank-1 modal logics. *ACM Trans. Comput. Log.* 10, 2 (2009), 13:1–13:33.

Lutz Schröder and Dirk Pattinson. 2011. Modular algorithms for heterogeneous modal logics via multi-sorted coalgebra. *Math. Struct. Comput. Sci.* 21 (2011), 235–266.

Lutz Schröder, Dirk Pattinson, and Daniel Hausmann. 2010. Optimal tableaux for conditional logics with cautious monotonicity. In *Proceedings of the European Conference on Artificial Intelligence (ECAI'10)*, Helder Coelho, Rudi Studer, and Michael Wooldridge (Eds.), Vol. 215. IOS Press, 707–712.

Lutz Schröder and Yde Venema. 2010. Flat coalgebraic fixed point logics. In *Proceedings of the Concurrency Theory (CONCUR'10)*, Paul Gastinand François Laroussinie (Eds.), Vol. 6269. Springer, 524–538.

Krister Segerberg. 1982. A completeness theorem in the modal logic of programs. In *Universal Algebra and Applications*. Banach Centre Publications, Vol. 9. PWN–Polish Scientific Publishers, Warsaw, 31–46.

Igor Walukiewicz. 2000. Completeness of Kozen's axiomatisation of the propositional μ-calculus. *Inf. Comput.* 157, 1–2 (2000), 142–182.

Pierre Wolper. 1983. Temporal logic can be more expressive. *Inf. Contr.* 56, 1–2 (1983), 72–99.

Received June 2016; revised March 2017; accepted October 2017