

Flat Coalgebraic Fixed Point Logics

Lutz Schröder¹ and Yde Venema²

¹ DFKI Bremen and Department of Computer Science, Universität Bremen

² ILLC, Universiteit van Amsterdam

Abstract. Fixed point logics are widely used in computer science, in particular in artificial intelligence and concurrency. The most expressive logics of this type are the μ -calculus and its relatives. However, popular fixed point 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* fixed point logics includes, e.g., CTL, the $*$ -nesting-free fragment of PDL, and the logic of common knowledge. Here, we extend this notion to the generic semantic framework of *coalgebraic logic*, thus 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 ATL), as well as probabilistic and monotone fixed point logics. Our main results are completeness of the Kozen-Park axiomatization and a timed-out tableaux method that matches EXPTIME upper bounds inherited from the coalgebraic μ -calculus but avoids using automata.

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 fixed point logics, that is, embedded into some variant of the μ -calculus. Typical examples are PDL [25] where, say, the formula $\langle a^* \rangle p$ (p can be reached by finite iteration of a) can be expressed as the least fixed point $\mu X. p \vee \langle a \rangle X$; CTL [7], whose formula AFp (p eventually holds on all paths) is just the fixed point $\mu X. p \vee \square X$; and the common knowledge operator C of epistemic logic [19], where Cp (p is common knowledge that p) can be expressed as the fixed point $\nu X. \bigwedge_{i=1}^n K_i(p \wedge X)$ 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 in comparison to the full μ -calculus.

One of the reasons why the full μ -calculus is both hard to read and algorithmically problematic in practice is that one has to keep 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 [10]), which is precisely what enables one to abandon variables altogether in favour of variable-free fixed point operators such as AF or C . We refer to logics that embed into a single-variable μ -calculus as *flat fixed point logics* [27].

Here, we study flat fixed point logics in the more general setting of *coalgebraic logic*. Coalgebra has recently emerged as the right framework for a unified treatment of a wide range of modal logics with seemingly disparate semantics beyond the realm of pure relational structures. Examples include monotone modal logic, probabilistic modal

logics [17], graded modal logic [11,5], and coalition logic [23]. This level of generality is achieved by parametrizing the semantics over a type functor 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, and concurrent game structures. The theory of coalgebraic modal logic has evolved quite rapidly, and presently includes, e.g., generic upper bounds PSPACE for satisfiability in next-step logics [29], and EXPTIME for satisfiability under global assumptions in hybrid next-step logics [31].

In our *flat coalgebraic fixed point logics* one can express operators 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 the single-variable fragments of the graded μ -calculus [16] and the alternating-time μ -calculus (AMC) [1], including alternating-time temporal logic (ATL). Flat coalgebraic fixed point logics are fragments of coalgebraic μ -calculi, and as such known to be decidable in EXPTIME under reasonable assumptions [4]. However, the decision procedure for the coalgebraic μ -calculus is, like the one for the standard μ -calculus [9], based on automata and as such has exponential average-case run time, while tableaux methods as suggested, e.g., by Emerson and Halpern for CTL [8] and by Kozen for the aconjunctive fragment of the μ -calculus [14] are expected to offer the possibility of feasible average-case behaviour.

Our main results on flat coalgebraic fixed point logics, parametric both w.r.t. the coalgebraic branching type and the choice of flat fragment, are

- completeness of the natural axiomatization that makes the fixed point definitions explicit, generalizing the well-known Kozen-Park axiomatization; and
- a construction of *timed-out tableaux* similar in spirit to Kozen’s tableaux for the aconjunctive μ -calculus,

both under mild restrictions on the form of fixed point operators. The completeness result generalizes results of [27] to the level of coalgebraic logic, and relies on the notion of \mathcal{O} -adjointness [26] to prove that fixed points *in the Lindenbaum algebra* are *constructive*, i.e. approximable in ω steps. The crucial ingredient here are the *one-step cutfree complete* rule sets of [29,22]. These enable significant generalizations of both the key *rigidity lemma* and the \mathcal{O} -adjointness theorem of [27], the latter to the effect that *all uniform-depth modal operators are \mathcal{O} -adjoint*. The novel tableaux construction is instrumental in the completeness proof, and at the same time confirms the known EXPTIME upper bound, avoiding however the use of automata and thus raising hopes for efficient implementation.

Our completeness result follows a long tradition of non-trivial completeness proofs, e.g. for PDL [15,32], CTL [8], the aconjunctive μ -calculus [14], and the full μ -calculus [33]. Note that all these results are independent, as completeness is not in general inherited by sublogics, and in fact employ 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 [16] nor [4] cover axiomatizations. In particular, we obtain for the first time a completeness result and a tableau procedure for graded fixed point logics, i.e. fragments of the graded μ -calculus, and we generalize the completeness of ATL [12] to arbitrary flat fragments of AMC.

2 Flat Coalgebraic Fixed Point Logics

We briefly recall the generic framework of coalgebraic modal logic [21,28], and define its extension with flat fixed point operators, a fragment of the coalgebraic μ -calculus [4].

The first parameter of the syntax is a (*modal*) *similarity type* Λ , i.e. a set of modal operators with associated finite arities. We work with formulas in negation normal form throughout, and therefore assume that every modal operator $\heartsuit \in \Lambda$ comes with a *dual* operator $\overline{\heartsuit} \in \Lambda$ of the same arity, where $\overline{\overline{\heartsuit}} = \heartsuit$. This determines the set $\mathcal{F}(\Lambda)$, or just \mathcal{F} , of *modal formulas* γ, δ by the grammar

$$\gamma, \delta ::= \perp \mid \top \mid v \mid \neg v \mid \gamma \wedge \delta \mid \gamma \vee \delta \mid \heartsuit(\gamma_1, \dots, \gamma_n)$$

where $\heartsuit \in \Lambda$ is n -ary and $v \in V$ for a fixed countably infinite set V of *variables*. Negation \neg , admitted in the above grammar only for variables v , then becomes a derived operation on all formulas in the standard way; e.g., $\neg \heartsuit(\gamma_1, \dots, \gamma_n) = \overline{\heartsuit}(\neg \gamma_1, \dots, \neg \gamma_n)$, and $\neg \neg v = v$. Further derived operations $\rightarrow, \leftrightarrow$ are defined as usual. Moreover, we define the *dual* $\overline{\gamma}$ of γ as $\overline{\gamma} \equiv \neg \gamma \sigma$ where the substitution σ is given by $\sigma(v) = \neg v$ for all $v \in V$. We intend variables as place holders for arguments and parameters of formulas defining fixed point operators; as such, they serve only technical purposes and will not form part of the actual fixed point language defined below. Instead, propositional atoms are incorporated into the modal similarity type Λ as nullary operators when needed.

The second syntactic parameter of a *flat coalgebraic fixed point logic* is a set Γ of modal formulas γ , where we distinguish a single fixed *argument* variable x and regard all other variables p_1, \dots, p_n in γ as *parameters*; we require that γ is *monotone* in all variables, i.e. does not contain $\neg x$ (an essential condition for the existence of fixed points) or $\neg p_i$ (a mere technical convenience, and not an actual restriction as one can always negate the actual parameter instead of the parameter variable). We require moreover that all $\gamma \in \Gamma$ are *guarded*, i.e. that all occurrences of the argument variable x are under the scope of at least one modal operator; as shown in [33], this is not an essential restriction as every μ -calculus formula is provably equivalent to a guarded formula. We denote substituted formulas $\gamma[\varphi_1/p_1; \dots; \varphi_n/p_n; \psi/x]$ as $\gamma(\varphi_1, \dots, \varphi_n, \psi)$. The set $\mathcal{F}_{\#}(\Lambda, \Gamma)$ or just $\mathcal{F}_{\#}$ of (*fixed point*) *formulas* φ, ψ is then defined by the grammar

$$\varphi, \psi ::= \perp \mid \top \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \heartsuit(\varphi_1, \dots, \varphi_n) \mid \#_{\gamma}(\varphi_1, \dots, \varphi_n) \mid b_{\overline{\gamma}}(\varphi_1, \dots, \varphi_n)$$

where $\heartsuit \in \Lambda$ is n -ary and $\gamma \in \Gamma$. The operator $\#_{\gamma}$ represents the least fixed point

$$\#_{\gamma}(\varphi_1, \dots, \varphi_n) = \mu x. \gamma(\varphi_1, \dots, \varphi_n, x),$$

while $b_{\overline{\gamma}}(\varphi_1, \dots, \varphi_n)$ represents the greatest fixed point $\nu x. \overline{\gamma}(\varphi_1, \dots, \varphi_n, x)$. The name *flat* for the fixed point operators $\#_{\gamma}, b_{\overline{\gamma}}$ relates to the fact that we require the formula γ to belong to the basic (fixed point free) modal language. Note that nesting of fixed point operators is unrestricted, e.g. φ can be an arbitrary fixed point formula in $\#_{\gamma}\varphi$. Syntactically, $\#_{\gamma}$ is an atomic operator, and occurrences of variables in γ do not count as occurrences in formulas $\#_{\gamma}\phi$. For the sake of readability, we restrict the further technical development (but not the examples) to unary modalities \heartsuit and unary fixed point 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.

Negation extends to fixed points by $\neg\sharp_\gamma\varphi = \flat_{\bar{\gamma}}(\neg\varphi)$ and $\neg\flat_{\bar{\gamma}}\varphi = \sharp_\gamma(\neg\varphi)$. Note that unlike in the case of modal formulas, we have not included variables in the definition of fixed point formulas. A *(fixed point) formula with variables* is an expression of the form $\gamma\sigma$, where γ is a modal formula and σ is a substitution of some of the variables in γ with fixed point formulas (i.e. variables never appear under fixed point operators). In the following, the term *formula* will refer to fixed point formulas without variables unless variables are explicitly mentioned. For $\gamma \in \Gamma$, we denote the function taking a formula ψ to $\gamma(\varphi, \psi)$ by $\gamma(\varphi)$, and by $\gamma(\varphi)^k$ its k -fold iteration. We assume a *reasonable size measure on Λ* and hence on formulas and sets of formulas [30], in particular that numbers (e.g. in graded or probabilistic operators) are *coded in binary*.

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 *type functor* $T : \text{Set} \rightarrow \text{Set}$, i.e. an operation T that maps sets X to sets TX and functions $f : X \rightarrow Y$ to functions $Tf : TX \rightarrow TY$, preserving identities and composition, and the latter by the choice of a predicate lifting $\llbracket \heartsuit \rrbracket$ for each $\heartsuit \in \Lambda$. Here, a *predicate lifting* (for T) is a family of maps $\lambda_X : \mathcal{P}X \rightarrow \mathcal{P}TX$, where X ranges over all sets, satisfying the *naturality* condition $\lambda_X(f^{-1}[A]) = (Tf)^{-1}[\lambda_Y(A)]$ for all $f : X \rightarrow Y$, $A \in \mathcal{P}Y$. As we work with fixed points, we insist that *all modal operators are monotone*, i.e. $\llbracket \heartsuit \rrbracket : \mathcal{P}(X) \rightarrow \mathcal{P}(TX)$ is monotone w.r.t. set inclusion for each $\heartsuit \in \Lambda$. Moreover, the assignment of predicate liftings must respect duality of operators: for $\heartsuit \in \Lambda$, $\llbracket \heartsuit \rrbracket_X(A) = TX - \llbracket \heartsuit \rrbracket_X(X - A)$. Given these data, the role of models is played by *T -coalgebras*, i.e. pairs (X, ξ) where X is a set of *states* and $\xi : X \rightarrow TX$ is the transition function; thinking of TX informally as a parametrised datatype over X , we regard ξ as associating with each state x a structured collection $\xi(x)$ of successor states and observations. E.g. for $TX = \mathcal{P}X \times \mathcal{P}(U)$, given a set U of propositional atoms, 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 1.

As indicated above, the choice of predicate liftings determine the interpretation of modal operators. The semantics of a formula φ with argument variable x (no other variables will ever be evaluated in unsubstituted form) is a subset $\llbracket \varphi \rrbracket_{(X, \xi)}(B) \subseteq X$, given a T -coalgebra (X, ξ) and a set $B \subseteq X$. The semantics of formulas φ without variables (in particular of \sharp - or \flat -formulas) does not depend on B and hence will be denoted just by $\llbracket \varphi \rrbracket_{(X, \xi)}$. One has obvious clauses for Boolean operators, $\llbracket x \rrbracket_{(X, \xi)}(B) = B$, and

$$\begin{aligned} \llbracket \heartsuit \varphi \rrbracket_{(X, \xi)}(B) &= \xi^{-1} \llbracket \heartsuit \rrbracket_X(\llbracket \varphi \rrbracket_{(X, \xi)}(B)) \\ \llbracket \sharp_\gamma \varphi \rrbracket_{(X, \xi)} &= \bigcap \{B \subseteq X \mid \llbracket \gamma(\varphi) \rrbracket_{(X, \xi)}(B) \subseteq B\} \\ \llbracket \flat_{\bar{\gamma}} \varphi \rrbracket_{(X, \xi)} &= \bigcup \{B \subseteq X \mid B \subseteq \llbracket \bar{\gamma}(\varphi) \rrbracket_{(X, \xi)}(B)\}. \end{aligned}$$

The clause for $\sharp_\gamma\varphi$ just says that $\llbracket \sharp_\gamma\varphi \rrbracket_{(X, \xi)}$ is the least fixed point of the monotone map $\llbracket \gamma(\varphi) \rrbracket_{(X, \xi)} : \mathcal{P}(X) \rightarrow \mathcal{P}(X)$, and similarly $\llbracket \flat_{\bar{\gamma}}\varphi \rrbracket_{(X, \xi)}$ is the greatest fixed point of $\llbracket \bar{\gamma}(\varphi) \rrbracket_{(X, \xi)}$. We fix the data T, Λ, Γ etc. throughout.

Example 1. 1. *Kripke semantics:* Fixed point extensions of the modal logic K have a single modal operator \Box , interpreted over the powerset functor \mathcal{P} (which takes a set

X to its powerset $\mathcal{P}(X)$) by the predicate lifting $\llbracket \Box \rrbracket_X(A) = \{B \in \mathcal{P}(X) \mid B \subseteq A\}$. \mathcal{P} -coalgebras $(X, \xi : X \rightarrow \mathcal{P}(X))$ are in 1-1 correspondence with Kripke frames, and $\llbracket \Box \rrbracket$ captures the usual semantics of the box operator. Multi-agent extensions are interpreted over $TX = \mathcal{P}(A \times X)$ where A is the set of agents. CTL, *-nesting-free PDL, and the logic of common knowledge are flat fixed point logics in this setting; e.g. AU and EU are the \sharp -operators for $p_2 \vee (p_1 \wedge \Box x)$ and for $p_2 \vee (p_1 \wedge \Diamond x)$, respectively.

2. *Graded fixed point logics* are sublogics of the *graded μ -calculus* [16]. They have modal operators \Diamond_k ‘in more than k successors’, with duals \Box_k ‘in all but k successors’, interpreted over the functor \mathcal{B} that takes a set X to the set $\mathcal{B}(X) = X \rightarrow \omega + 1$ of multisets over X by $\llbracket \Diamond_k \rrbracket_X(A) = \{B \in \mathcal{B}(X) \mid \sum_{x \in A} B(x) > k\}$. This captures the semantics of graded modalities over *multigraphs* [5], which is equivalent to the more customary Kripke semantics [11] w.r.t. satisfiability of fixed point formulas. In description logic, graded operators are called *qualified number restrictions* [2]. The example mentioned in [16], a graded fixed point 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.

3. *Probabilistic fixed point logics*, i.e. fixed point extensions of probabilistic modal logic [17], have modal operators L_p ‘in the next step, it holds with probability at least p that’, for $p \in [0, 1] \cap \mathbb{Q}$. They are interpreted over the functor \mathcal{D} that maps a set X to the set of discrete probability distributions on X by putting $\llbracket L_p \rrbracket_X(A) = \{P \in \mathcal{D}(X) \mid PA \geq p\}$. Coalgebras for \mathcal{D} are Markov chains. We can use the \flat -operator AG_p for $p \wedge L_p x$ to express formulas like $AG_p \text{–fail}$, stating that the system will, at any point during its run time, fail with probability at most $1 - p$; a sensible specification for systems that may sometimes fail but should not fail excessively often. In an epistemic reading of probabilities, flat probabilistic fixed point logics support, e.g., a common belief operator ‘it is commonly believed with confidence p that’.

4. *The alternating-time μ -calculus (AMC)* [1] has modal operators $\langle\langle A \rangle\rangle \bigcirc$ read ‘coalition A has a joint strategy to enforce ... in one step’, where a *coalition* is a subset of a fixed set N of agents (in coalition logic [23], these operators are denoted $[A]$). Their semantics is defined over *concurrent game structures* (or *game frames*), and can be captured coalgebraically [29]. One of the flat fragments of AMC is Alternating-Time Temporal Logic (ATL) [1]. E.g., the ATL-operator $\langle\langle A \rangle\rangle p_1 \mathcal{U} p_2$, read ‘coalition A can eventually force p_2 and meanwhile maintain p_1 ’, is the \sharp -operator for $p_2 \vee (p_1 \wedge \langle\langle A \rangle\rangle \bigcirc x)$. Flat fixed points in AMC go considerably beyond ATL; e.g. the \flat -operator for $p \wedge \langle\langle \rangle\rangle \bigcirc \langle\langle \rangle\rangle \bigcirc x$ (‘ p holds in all even states along any path’) is not even in ATL^* [1,6]. A similar flat operator, the \flat -operator for $\langle\langle A \rangle\rangle \bigcirc (p \wedge \langle\langle B \rangle\rangle \bigcirc (q \wedge x))$, expresses that coalitions A and B can forever play ping-pong between p and q .

5. *Monotone fixed point 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 $\llbracket \Box \rrbracket_X(A) = \{\mathfrak{A} \in \mathcal{M}(X) \mid A \in \mathfrak{A}\}$. In multi-modal versions of this, boxes and their semantics are indexed, e.g. over agents, programs, or games. This is the semantic setting of logics such as concurrent PDL [24] and Parikh’s game logic [20], the flat fragments of which are the *-nesting-free fragments. E.g., using $\langle \gamma \rangle$ to denote the game logic operator ‘Angel has a strategy

to enforce ... in game γ' , the operator $\langle \gamma^\times \rangle$ for a $*$ -free game γ , where \times denotes demonic iteration (Demon chooses the number of rounds), is the \flat -operator for $p \wedge \langle \gamma \rangle x$.

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 one of our main results. 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. Fixed points are governed by the obvious generalization of the *Kozen-Park axiomatization*: we have the *unfolding* axiom

$$\sharp_\gamma \varphi \leftrightarrow \gamma(\varphi, \sharp_\gamma \varphi)$$

and the *fixed-point induction* rule

$$\gamma(\varphi, \chi) \rightarrow \chi / \sharp_\gamma \varphi \rightarrow \chi,$$

for all formulas φ, χ . (Here α / β denotes the rule ‘from α infer β ’).

The variable part is now the axiomatization of the modal operators, which turns out to be completely orthogonal to the fixed point axiomatization. In fact, we can just re-use complete rule sets for the purely modal logic as developed in [29]. First some notation.

Definition 2. We denote the set of positive propositional formulas (formed using only \wedge and \vee) over a set Z by $\text{Pos}(Z)$, and the set $\{\heartsuit a \mid \heartsuit \in \Lambda, a \in Z\}$ by $\Lambda(Z)$. We say that a conjunction (disjunction) is *contracted* if no conjunct (disjunct, respectively) occurs twice in it. For $\varphi, \psi \in \text{Pos}(Z)$, we say that φ *propositionally entails* ψ and write $\varphi \vdash_{PL} \psi$ if $\varphi \rightarrow \psi$ is a propositional tautology. Similarly, $\Phi \subseteq \text{Pos}(Z)$ propositionally entails ψ ($\Phi \vdash_{PL} \psi$) if there exist $\varphi_1, \dots, \varphi_n \in \Phi$ such that $\varphi_1 \wedge \dots \wedge \varphi_n \vdash_{PL} \psi$. For $\varphi \in \text{Pos}(Z)$, we denote the evaluation of φ in the Boolean algebra $\mathcal{P}(X)$ under a valuation $\tau : Z \rightarrow \mathcal{P}(X)$ by $\llbracket \varphi \rrbracket_{X, \tau}$, and write $X, \tau \models \varphi$ if $\llbracket \varphi \rrbracket_{X, \tau} = X$. For $\psi \in \text{Pos}(\Lambda(Z))$, the interpretation $\llbracket \psi \rrbracket_{TX, \tau}$ of ψ in the Boolean algebra $\mathcal{P}(TX)$ under τ is the inductive extension of the assignment $\llbracket \heartsuit(z) \rrbracket_{TX, \tau} = \llbracket \heartsuit \rrbracket_{X, \tau}(z)$. We write $TX, \tau \models \psi$ if $\llbracket \psi \rrbracket_{TX, \tau} = TX$.

We can now give the formal definition of the modal rule format, where due to monotonicity of the modal operators we can restrict to monotone rules following [4]. To understand the following, note that every rule of the form φ/χ , which says that if φ is provable then χ is provable, comes with a dual *tableau rule* $\bar{\chi}/\bar{\varphi}$ saying that if $\bar{\chi}$ is consistent then $\bar{\varphi}$ is consistent.

Definition 3. A (*monotone one-step*) rule $R = \varphi/\chi$ consists of a *premise* $\varphi \in \text{Pos}(V)$ and a *conclusion* χ which is a disjunction over $\Lambda(V)$ (recall that V is the set of variables), where every variable appears at most once in φ and every variable in φ appears also in χ . The rule R is *one-step sound* if whenever $X, \tau \models \varphi$ for a valuation $\tau : V \rightarrow \mathcal{P}(X)$, then $TX, \tau \models \chi$. Given a set \mathcal{R} of one-step rules, we say that a conjunction ψ over $\Lambda(V)$ is *one-step cut-free τ -consistent* for a set X and $\tau : V \rightarrow \mathcal{P}(X)$

if whenever $\varphi/\chi \in \mathcal{R}$ and $\sigma : V \rightarrow V$ is a renaming such that $\chi\sigma$ is contracted and $\psi \vdash_{PL} \bar{\chi}\sigma$ (note that propositional entailment between conjunctions is just reverse containment), then $\llbracket \bar{\varphi}\sigma \rrbracket_{X,\tau} \neq \emptyset$. We say that \mathcal{R} is *one-step cutfree complete* if $\llbracket \psi \rrbracket_{TX,\tau} \neq \emptyset$ whenever ψ is one-step cut-free τ -consistent. A set $\Psi \subseteq \Lambda(V)$ is *one-step cut-free τ -consistent* if for all $\psi_1, \dots, \psi_n \in \Psi$, $\psi_1 \wedge \dots \wedge \psi_n$ is one-step cut-free τ -consistent.

(In the terminology of [29], one-step cutfree complete rule sets correspond to one-step complete rule sets which are closed under contraction and resolution.) As the last parameter of the framework, we *fix from now on a one-step cutfree complete set \mathcal{R} of one-step sound monotone one-step rules, and denote the arising logic by \mathcal{L}_{\sharp}* . Rules $\varphi/\psi \in \mathcal{R}$ are applied in substituted form, i.e. for every substitution σ , we may conclude $\psi\sigma$ from $\varphi\sigma$. It is easy to see that the arising parametrized deduction system is sound. As usual, we write $\vdash \varphi$ if φ is provable, and $\varphi \vdash \psi$ if $\vdash \varphi \rightarrow \psi$. We say that φ is *consistent* if $\neg\varphi$ is not provable. It has been shown that one-step cutfree complete rule sets engender complete cut-free sequent systems for the purely modal logic, and suitable rule systems have been exhibited for all logics of Example 1 and many more [29,22]. E.g., a one-step cutfree complete set of monotone one-step rules for K is

$$\frac{\bigvee_{i=1}^n a_i \vee b}{\bigvee_{i=1}^n \diamond a_i \vee \square b} \quad (n \geq 0).$$

As a more complex example, we recall the one-step cutfree complete rule schema for graded operators [29], reformulated to fit the monotone rule format:

$$\frac{\sum_{i=1}^n -r_i(\neg a_i) + \sum_{j=1}^m s_j b_j \geq 0}{\bigvee_{i=1}^n \square_{k_i} a_i \vee \bigvee_{j=1}^m \diamond_{l_j} b_j},$$

where $n + m \geq 1$ and $r_1, \dots, r_n, s_1, \dots, s_m > 0$, subject to the side condition $\sum_{i=1}^n r_i(k_i + 1) \geq 1 + \sum_{j=1}^m s_j l_j$. Here, the premise represents a linear inequality between the characteristic functions of the a_i and the b_j , i.e. count s_j when b_j holds and $-r_i$ when a_i does not hold; this is easily seen to be expressible by a positive propositional formula (cf. [29]).

4 Constructive Fixed Points

Our next aim is to prove that the Lindenbaum algebra of \mathcal{L}_{\sharp} is *constructive*, i.e. its fixed points can be iteratively approximated in ω steps. In terms of consistency of formulas, this means that whenever a formula of the form $\sharp_{\gamma}\varphi \wedge \psi$ is consistent, then already $\gamma^i(\varphi)(\perp) \wedge \psi$ is consistent for some $i < \omega$; this fact plays a pivotal role in our tableau model construction. We begin by introducing the requisite algebraic tools.

We define a Λ -*modal algebra* A as a Boolean algebra extended with a monotone operation $\heartsuit^A : A \rightarrow A$ for each $\heartsuit \in \Lambda$. In such an algebra, every modal formula $\varphi(v_1, \dots, v_n)$ is naturally interpreted as an operation $\varphi^A : A^n \rightarrow A$. Now we say that A *validates* a rule $R = \varphi/\psi$ if $\psi^A(a_1, \dots, a_n) = \top$ whenever $\varphi^A(a_1, \dots, a_n) = \top$. A \sharp -*algebra* is a Λ -algebra A that is endowed with operations \sharp_{γ}^A and $b_{\overline{\gamma}}^A$ for each $\gamma \in \Gamma$

such that for each $a \in A$, $\#_\gamma^A(a)$ is the least fixed point of the map $\gamma^A(a, -) : A \rightarrow A$ and $\flat_\gamma^A(a)$ is the greatest fixed point of $\bar{\gamma}^A(a, -)$ (in particular, these fixed points *exist* in a $\#$ -algebra). An $\mathcal{L}_\#$ -algebra is a $\#$ -algebra A that validates every rule R of our fixed set \mathcal{R} of one-step rules. In the tradition of algebraic logic, the class of these algebras provides an algebraic encoding of the proof system.

More specifically, we will be interested in the *Lindenbaum algebra* $A(\mathcal{L}_\#)$ of our logic. As usual, this algebra is defined as the quotient of the formula/term algebra (or absolutely free algebra) under the congruence relation \equiv of provable equivalence ($\varphi \equiv \psi$ iff $\varphi \leftrightarrow \psi$ is derivable). Observe that in a natural way, every sentence φ is interpreted as the element $\varphi^{A(\mathcal{L}_\#)} = [\varphi]$ of this algebra; we will mostly write φ rather than $[\varphi]$. *The Kozen-Park axiomatization ensures that $A(\mathcal{L}_\#)$ actually is an $\mathcal{L}_\#$ -algebra*, and then of course, the *initial $\mathcal{L}_\#$ -algebra*. In these terms, our target property is phrased as follows.

Definition 4. We say that $\gamma \in \Gamma$ is *constructive* if

$$\#_\gamma \varphi = \bigvee_{i < \omega} \gamma(\varphi)^i(\perp)$$

in the Lindenbaum algebra $A(\mathcal{L}_\#)$, i.e. if $\#_\gamma \varphi \vdash \psi$ whenever $\gamma(\varphi)^i(\perp) \vdash \psi$ for all $i < \omega$. If all $\gamma \in \Gamma$ are constructive, then $A(\mathcal{L}_\#)$ is *constructive*.

We explicitly state the dual formulation of this property:

Lemma 5. *Let γ be constructive. If $\#_\gamma \varphi \wedge \psi$ is consistent, then $\gamma(\varphi)^i(\perp) \wedge \psi$ is consistent for some $i < \omega$.*

The central tool for proving constructivity, introduced in [26] and featuring prominently in [27], is the notion of a finitary \mathcal{O} -adjoint:

Definition 6. We say that γ is an \mathcal{O} -adjoint if for all $\varphi, \psi \in \mathcal{F}_\#$, there exists a finite set $G_{\gamma(\varphi)}(\psi)$ of formulas such that for all $\rho \in \mathcal{F}_\#$,

$$\gamma(\varphi, \rho) \vdash \psi \text{ iff } \rho \vdash \chi \text{ for some } \chi \in G_{\gamma(\varphi)}(\psi),$$

i.e. $\gamma(\varphi, \rho) \leq \psi$ in $A(\mathcal{L}_\#)$ iff $\rho \leq \chi$ for some $\chi \in G_{\gamma(\varphi)}(\psi)$. Moreover, γ is a *finitary \mathcal{O} -adjoint* if $G_{\gamma(\varphi)}$ can be chosen such that for every ψ , the closure of ψ under $G_{\gamma(\varphi)}$, i.e. the smallest set \mathcal{A} with $\psi \in \mathcal{A}$ and $\chi \in \mathcal{A} \Rightarrow G_{\gamma(\varphi)}(\chi) \subseteq \mathcal{A}$, is finite.

Lemma 7. [26] *Every finitary \mathcal{O} -adjoint is constructive.*

The first step in the proof of \mathcal{O} -adjointness for a large class of operators is a generalization of the rigidity lemma of [26]:

Lemma 8 (Rigidity). *Let ψ be a disjunction over $\Lambda(A(\mathcal{L}_\#))$. Then ψ is provable iff there exists a one-step rule φ/χ and a substitution σ such that $\varphi\sigma$ is provable, $\chi\sigma$ is contracted, and $\chi\sigma \vdash_{PL} \psi$.*

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 in [27]:

Let A be a countable $\mathcal{L}_\#$ -algebra, let $\mathcal{S}(A)$ be the set of ultrafilters of A , fix a surjective map $\sigma : V \rightarrow A$, and let a conjunction ρ over $\Lambda(V)$ be one-step θ -consistent for $\theta : V \rightarrow \mathcal{P}(\mathcal{S}(A))$ given by $\theta(v) = \{u \in \mathcal{S}(A) \mid \sigma(v) \in u\}$. We construct the one-point extension A^ρ , an $\mathcal{L}_\#$ -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 $\Phi \subseteq \Lambda(V)$ such that $\Phi \vdash_{PL} \rho$. As we emulate adding a single point, the carrier of A^ρ is $A \times 2$. We make A^ρ into a Λ -modal algebra by putting

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

where $\heartsuit^\rho : A \rightarrow 2$ is defined by $\heartsuit^\rho(a) = \top$ iff $\heartsuit a \in \Phi$. (Thus, $\heartsuit^{A^\rho}(a, d)$ is independent of d , in agreement with the intuition that the interpretation of modal operators depends only on the successor structure of the current state, not on the state itself.) In particular, this implies that $\rho\sigma > \perp$ in A^ρ .

Lemma 9. *The algebra A^ρ is an $\mathcal{L}_\#$ -algebra.*

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

Lemma 10. *Let $\sigma : V \rightarrow A(\mathcal{L}_\#)$ be surjective. If a conjunction ρ over $\Lambda(\mathcal{F}_\#)$ is one-step θ -consistent for $\theta(v) = \{u \in \mathcal{S}(A(\mathcal{L}_\#)) \mid \sigma(v) \in u\}$, then ρ is consistent, i.e. $\rho > \perp$ in $A(\mathcal{L}_\#)$.*

From Lemma 10, one easily proves Lemma 8 using the fact that every consistent formula is contained in some ultrafilter of $A(\mathcal{L}_\#)$.

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

Definition 11. A formula φ with variables is *uniform of depth k* if every occurrence of the fixed argument variable x in φ is in the scope of exactly k modal operators (including the case that x does not occur in φ ; recall moreover that variables never occur under fixed point operators). Moreover, φ is *uniform* if φ is uniform of depth k for some k ; the minimal such k is the *depth of uniformity* of φ .

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

Definition 12. A set Σ of formulas is *Fischer-Ladner closed* if Σ is closed under subformulas and negation, and whenever $\star_\gamma\varphi \in \Sigma$, then $\gamma(\varphi, \star_\gamma\varphi) \in \Sigma$ for $\star \in \{\#, b\}$. We denote the Fischer-Ladner closure of a formula φ by $FL(\varphi)$.

Lemma 13. [14] *The set $FL(\varphi)$ is finite and of polynomial size in φ .*

The further development revolves largely around *admissible rules*, i.e. rules φ/ψ where φ and ψ are formulas with variables v_1, \dots, v_n such that $A(\mathcal{L}_\#)$ validates φ/ψ , i.e. whenever $\vdash \varphi(\rho_1, \dots, \rho_n)$ for formulas ρ_1, \dots, ρ_n then $\vdash \psi(\rho_1, \dots, \rho_n)$.

Lemma 14. *Let ψ be uniform, and put*

$$G = \{\varphi \in \text{Pos}(FL(\psi)) \mid \varphi/\psi \text{ admissible, } \varphi \text{ uniform of depth } 0\}.$$

Then we have that for all ρ , $\psi(\rho)$ is provable iff $\varphi(\rho)$ is provable for some $\varphi \in G$.

Proof (sketch). Induction over the depth of uniformity, with trivial base case, using rigidity (Lemma 8) in the inductive step. \square

Theorem 15 (Finitary \mathcal{O} -adjointness). *If the formula ψ with argument variable x is monotone and uniform in x , then the operation $\psi^{A(\mathcal{L}_\#)} : A(\mathcal{L}_\#) \rightarrow A(\mathcal{L}_\#)$ induced by ψ is a finitary \mathcal{O} -adjoint.*

Proof (sketch). For $\varphi \in \mathcal{F}_\#$, we have to construct a set $G_\psi(\varphi)$ of formulas such that for all $\rho \in \mathcal{F}_\#$, $\psi(\rho) \vdash \varphi$ iff $\rho \vdash \chi$ for some $\chi \in G_\psi(\varphi)$. Now $\psi' := \psi \rightarrow \varphi$ is uniform. Let $G \subseteq \text{Pos}(FL(\psi'))$ be as in Lemma 14, applied to ψ' , and let G_0 be a finite set of representatives of G modulo propositional equivalence. Then we can put

$$G_\psi(\varphi) = \{\chi(\top) \mid \chi \in G_0, \vdash \chi(\top) \vee \chi(\perp)\}. \quad \square$$

Using uniform formulas as a base, we can now exploit some known closure properties of finitary \mathcal{O} -adjoints [26].

Definition 16. The set of *admissible* modal formulas is the closure of the set of monotone uniform modal formulas in x under disjunction, conjunction with modal formulas not containing x , and substitution for the argument variable, the latter in the sense that if γ and δ are admissible, then $\gamma(\delta)$ is admissible.

Corollary 17. *If $\gamma \in \Gamma$ is admissible, then γ is a finitary \mathcal{O} -adjoint, and hence constructive.*

From now on, we require that every $\gamma \in \Gamma$ is admissible, and hence $A(\mathcal{L}_\#)$ is constructive. All fixpoint operators mentioned in Example 1 are based on admissible formulas (in fact, on uniform ones).

5 The Tableau Construction

We proceed to describe a construction of *timed-out tableaux* for consistent formulas, which we shall then use as carrier sets for coalgebraic models. (Note that in coalgebraic logics, tableaux, being only relational structures, cannot directly serve as models.) Our time-outs are related to Kozen's μ -counters [14] but are integrated into the formulas appearing in the tableau (rather than maintained independently in the *construction* of the tableau), and in particular govern the way modal successor nodes are generated. The use of time-outs is justified by constructivity of fixed point operators as proved in the previous section. In the following, we fix a finite Fischer-Ladner closed set Σ .

Definition 18. The set of *timed-out formulas* φ, ψ is generated by the grammar

$$\varphi, \psi ::= \perp \mid \top \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \heartsuit \varphi \mid \#_\gamma(\rho)^\kappa \mid \flat_\gamma(\rho) \quad (\kappa \in \omega + 1, \rho \in \mathcal{L}_\#)$$

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 $\#_\gamma(\chi)^\kappa$ with $\kappa < \omega$ (which however may occur any number of times), and for this $\#_\gamma(\chi)^\kappa$, (i) $\#_\gamma(\chi)^\omega$ is not a subformula of φ ; and (ii) whenever $\#_\delta(\rho)^\omega$ is a subformula of φ , then $\#_\delta(\rho)$ is a subformula of χ .

In this case, we define the *time-out* $\tau(\varphi)$ of φ to be κ , and $\tau(\varphi) = \omega$ otherwise (i.e. if φ does not contain any subformula of the form $\sharp_\gamma(\chi)^\kappa$ 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{L}_\sharp , given by commutation with Boolean and modal operators, $(\flat_{\overline{\gamma}}(\rho))^s = (\flat_{\overline{\gamma}}(\rho))^t = \flat_{\overline{\gamma}}(\rho)$, and

$$(\sharp_\gamma(\rho)^\omega)^s = \sharp_\gamma(\rho) \quad (\sharp_\gamma(\rho)^i)^s = \gamma(\rho)^i(\perp) \quad (i < \omega) \quad (\sharp_\gamma(\rho)^\kappa)^t = \sharp_\gamma(\rho).$$

Thus, s unfolds fixed points as prescribed by their time-outs, and t just removes time-outs. Both translations extend to sets of formulas. For timed-out formulas φ, ψ , we put $\varphi \preceq \psi$ iff $\varphi^t = \psi^t$ and $\tau(\varphi) \leq \tau(\psi)$. That is, $\varphi \preceq \psi$ 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 $\varphi^t \in \Sigma$.

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 which are \sharp -formulas, if such a greatest element exists and is not under the scope of a \flat -operator. This enables the simple definition of \preceq , which trivially has the following property.

Lemma 19. *For every formula φ , the preimage of φ under the translation t is linearly ordered by \preceq .*

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

Lemma 20. *If $\sharp_\gamma\varphi^\kappa$ is a timed-out formula, then so is $\gamma(\varphi, \sharp_\gamma\varphi^\kappa)$.*

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

Definition 21. A *timed-out Σ -atom* is a maximal set A of timed-out Σ -formulas such that (i) the translation t is injective on A , and (ii) A^s is consistent. Here, maximality is w.r.t. \sqsubseteq where $A \sqsubseteq B$ iff for all $\varphi \in A$, there exists a (necessarily unique) $\varphi' \in B$ such that $\varphi' \preceq \varphi$; intuitively: B contains A up to possible decrease of time-outs. We write \overline{A} for the closure of A under \preceq (i.e. if $\varphi \in \overline{A}$ and $\varphi \preceq \varphi'$ then $\varphi' \in \overline{A}$).

The following lemma uses the fact that finite product orderings $(\omega+1)^k$ are well-quasi-orders, and in particular have only finite anti-chains [18].

Lemma 22. *The set of timed-out Σ -atoms is finite.*

Lemma 23 (Timed-out Lindenbaum lemma). *For every set A_0 of timed-out Σ -formulas such that A_0^s is consistent and t is injective on A_0 , there exists a timed-out Σ -atom A such that $A_0 \sqsubseteq A$.*

The proof of the truth lemma crucially depends on a set of Hintikka-like properties:

Lemma 24. *If A is a timed-out Σ -atom, then*

1. if $\varphi \wedge \psi \in A$ then $\varphi \in \bar{A}$ and $\psi \in \bar{A}$;
2. if $\varphi \vee \psi \in A$ then $\varphi \in \bar{A}$ or $\psi \in \bar{A}$;
3. $\perp \notin \bar{A}$;
4. if $\#_\gamma \varphi^\kappa \in A$, then $\kappa < \omega$;
5. $\#_\gamma \varphi^\kappa \in A$ iff $\gamma(\varphi, \#_\gamma \varphi^{\kappa-1}) \in A$;
6. $\flat_\gamma \varphi \in A$ iff $\gamma(\varphi, \flat_\gamma \varphi) \in A$.

We proceed to define the actual tableaux, which relate timed-out atoms in a way that reflects application of dual rules $\bar{\chi}/\bar{\varphi}$ of modal rules $\varphi/\chi \in \mathcal{R}$, while fixed points are in a sense taken care of by the timed-out atoms themselves.

Definition 25. A demand of a Σ -atom A is a formula $\rho \equiv \bar{\varphi}\sigma$, where $\varphi/\chi \in \mathcal{R}$ is a rule with dual rule $\bar{\chi}/\bar{\varphi}$ and σ is a substitution such that $\chi\sigma$ is contracted and $A \vdash_{PL} \bar{\chi}\sigma$. A timed-out Σ -tableau (\mathcal{T}, R, l) consists of a finite graph (\mathcal{T}, R) and a labelling l of the nodes $n \in \mathcal{T}$ with timed-out Σ -atoms $l(n)$ such that for every demand ρ of $l(n)$, there exists nRm such that $l(m) \vdash_{PL} \rho$. The tableau (\mathcal{T}, R, l) is a timed-out Σ -tableau for $\varphi \in \Sigma$ if $(\varphi')^t = \varphi$ for some $\varphi' \in l(n)$, $n \in \mathcal{T}$. A coalgebra structure ξ on \mathcal{T} is coherent if for every n and every $\heartsuit\varphi \in \Sigma$,

$$\xi(n) \in \llbracket \heartsuit \rrbracket n(\varphi) \text{ iff } \heartsuit\varphi \in l(n),$$

where $n(\varphi) = \{m \in \mathcal{T} \mid nRm, \varphi \in l(m)\}$.

The link between timed-out tableaux and coalgebraic models is provided by the following lemma, whose proof relies on one-step cutfree completeness of the rule set.

Lemma 26 (Model existence lemma). For every timed-out Σ -tableau (\mathcal{T}, R, l) , there exists a coherent coalgebra structure on \mathcal{T} .

Lemma 27 (Truth lemma). If (\mathcal{T}, R, l) is a timed-out Σ -tableau and ξ is a coherent coalgebra structure on \mathcal{T} , then $n \in \llbracket \varphi \rrbracket_{(\mathcal{T}, \xi)}$ whenever $\varphi \in l(n)$.

Proof (sketch). Induction over timed-out Σ -formulas φ using the lexicographic product of the subterm ordering on φ^t and \preceq as the induction measure, and with the inductive hypothesis strengthened to apply also to $\varphi \in \overline{l(n)}$. Boolean cases are by Lemma 24; the step for modal operators is by coherence. The case for \flat -operators is by coinduction. For $\varphi = \#_\gamma(\psi)^\kappa$, we have $\kappa < \omega$ and $\gamma(\psi, \#_\gamma(\psi)^{\kappa-1}) \in \overline{l(n)}$ by Lemma 24. Then prove by a further induction over subformulas δ of γ that $n \models_{(\mathcal{T}, \xi)} (\delta(\psi, \#_\gamma(\psi)^{\kappa-1}))^s$ whenever $\delta(\psi, \#_\gamma(\psi)^{\kappa-1}) \in \overline{l(n)}$. Here, the case for the parameter variable x is discharged by the inductive hypothesis applied to $\#_\gamma(\psi)^{\kappa-1}$. \square

The previous two lemmas imply that every formula that has a timed-out tableau is satisfiable. The following lemma provides the link to consistency.

Lemma 28. For any consistent $\varphi \in \Sigma$ there is a finite timed-out Σ -tableau.

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

Theorem 29 (Completeness). If Γ is admissible and \mathcal{R} is one-step cutfree complete, then the $\mathcal{L}_\#$ is complete over finite models.

This result applies to all flat fixed point logics of Example 1, including all admissible flat fragments of AMC and the graded μ -calculus.

6 Complexity

Next we analyse the algorithmic aspects of satisfiability checking. This analysis is independent of the completeness result from Section 5 (except that completeness tells us that satisfiability checking is equivalent to consistency checking) but uses the same model construction. The complexity of the satisfiability problem as such is known: under additional conditions that we shall use below as well, it has been shown that satisfiability in the coalgebraic μ -calculus is in EXPTIME [4] (and therefore typically EXPTIME-complete, with hardness inherited from the standard μ -calculus). However, like known decision procedures for the standard μ -calculus, the algorithm in [4] uses automata-based methods and as such will exhibit exponential average-case behaviour, while a simple tableau method such as the one developed in Section 5 offers the possibility of feasible average-case behaviour using bottom-up construction of tableaux.

What is missing technically from the tableau construction of Section 5 with a view to complexity bounds is a bound on the time-outs. While we are confident that this can be proved directly using the \mathcal{O} -adjointness method (e.g. it is easy to show in this way that in Lemma 5, i can be exponentially bounded in $(\#_{\gamma}\varphi) \wedge \psi$), this is not actually necessary given that it has already been proved in [4] that the coalgebraic μ -calculus has the exponential model property. This implies immediately that time-outs can be exponentially bounded, so that tableaux are at most exponentially large. The key contribution of our tableaux construction here is to make this straightforward idea (which is similar in spirit to, e.g., Kozen's tableaux for the aconjunctive fragment of the μ -calculus [14]) work in a way that handles time-outs economically and consistently.

The size bound on tableaux alone does not yet imply an EXPTIME bound; however, we can obtain such a bound by using the coalgebraic generalization of the global caching method in exactly the same way as done in [13] for coalgebraic modal logic with global assumptions. To this end, we need to assume, as in [13,31], that our set \mathcal{R} of one-step rules is EXPTIME-tractable, i.e. that there exists a coding of the rules such that, up to propositional equivalence, all demands of a conjunction over $\Lambda(\mathcal{F}_{\#})$ can be generated by rules with codes of polynomially bounded size, and such that validity of codes, matching of rule codes for $\varphi/\chi \in \mathcal{R}$ to conjunctions ψ over $\Lambda(\mathcal{F}_{\#})$ (in the sense of finding σ such that $\chi\sigma$ is contracted and $\psi \vdash_{PL} \bar{\chi}\sigma$), and membership of disjunctions in a CNF of a rule premise are all decidable in EXPTIME. Summing up,

if \mathcal{R} is EXPTIME-tractable, then global caching decides existence of tableaux for $\mathcal{L}_{\#}$ in EXPTIME.

Global caching will typically avoid full expansion of tableaux, and provides a handle to achieve feasible average-case performance using suitable heuristics.

7 Conclusions

We have raised the theory of flat modal fixed point logics [27] to the level of generality of coalgebraic logic. Specifically, we have given a Kozen-Park style axiomatization for fixed point operators, and we have shown this axiomatization to be sound and complete under the conditions that (i) the defining formulas of the fixed point operators satisfy a

mild syntactic criterion, and (ii) the coalgebraic base logic is axiomatized by a one-step cutfree complete rule set. This result is a wide generalization with respect to the case of relational semantics, and covers, e.g., natural fixed point extensions of probabilistic modal logic and monotone modal logic. Most notably, we prove completeness of flat fragments of the graded μ -calculus [16], to our knowledge the first completeness result for any graded fixed point logic, and we generalize completeness of alternating-time temporal logic [12] to flat fragments of the alternating-time μ -calculus [1].

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 [26], and hence induce *constructive* fixed point operators that can be approximated in ω steps. This has enabled a model construction using tableaux with explicit time-outs for least fixed point formulas in the spirit of [14], which relies on a judicious definition of timed-out formula. As a byproduct of this construction, we obtain an optimal (i.e. EXPTIME) tableau calculus which paves the way for efficient implementations of coalgebraic flat fixed point logics, e.g. in the framework of the Coalgebraic Logic Satisfiability Solver CoLoSS [3].

Remaining open problems include the extension of the completeness result to larger fragments of the coalgebraic μ -calculus beyond the single variable fragment covered here, first and foremost the alternation-free fragment, and eventually the full coalgebraic μ -calculus. Similarly, there is the perspective to extend our tableau construction to at least the alternation-free fragment. A further direction for future research includes the development of generic coalgebraic model checking techniques.

References

1. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. ACM*, 49:672–713, 2002.
2. F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, eds. *The Description Logic Handbook*. Cambridge University Press, 2003.
3. G. Calin, R. Myers, D. Pattinson, and L. Schröder. CoLoSS: The Coalgebraic Logic Satisfiability Solver (system description). In *Methods for Modalities (M4M-5, 2007)*, vol. 231 of *ENTCS*, pp. 41–54. Elsevier, 2009.
4. C. Cîrstea, C. Kupke, and D. Pattinson. EXPTIME tableaux for the coalgebraic μ -calculus. In *Computer Science Logic, CSL 09*, vol. 5771 of *LNCS*, pp. 179–193. Springer, 2009.
5. G. D’Agostino and A. Visser. Finality regained: A coalgebraic study of Scott-sets and multisetsets. *Arch. Math. Logic*, 41:267–298, 2002.
6. M. Dam. CTL* and ECTL* as fragments of the modal mu-calculus. *Theoret. Comput. Sci.*, 126:77–96, 1994.
7. E. A. Emerson and E. M. Clarke. Using branching time temporal logic to synthesize synchronization skeletons. *Sci. Comput. Program*, 2(3):241–266, 1982.
8. E. A. Emerson and J. Y. Halpern. Decision procedures and expressiveness in the temporal logic of branching time. *J. Comput. System Sci.*, 30:1–24, 1985.
9. E. A. Emerson and C. S. Jutla. The complexity of tree automata and logics of programs. *SIAM J. Comput.*, 29:132–158, 1999.
10. E. A. Emerson and C.-L. Lei. Efficient model checking in fragments of the propositional mu-calculus. In *Logic in Computer Science, LICS 86*, pp. 267–278. IEEE, 1986.
11. K. Fine. In so many possible worlds. *Notre Dame J. Formal Logic*, 13:516–520, 1972.

12. V. Goranko and G. van Drimmelen. Complete axiomatization and decidability of alternating-time temporal logic. *Theoret. Comput. Sci.*, 353:93–117, 2006.
13. R. Goré, C. Kupke, and D. Pattinson. Optimal tableau algorithms for coalgebraic logics. In *Tools and Algorithms for the Construction and Analysis of Systems, TACAS 10*, vol. 6015 of *LNCS*, pp. 114–128. Springer, 2010.
14. D. Kozen. Results on the propositional μ -calculus. *Theoret. Comput. Sci.*, 27:333–354, 1983.
15. D. Kozen and R. Parikh. An elementary proof of the completeness of PDL. *Theoret. Comput. Sci.*, 14:113–118, 1981.
16. O. Kupferman, U. Sattler, and M. Y. Vardi. The complexity of the graded μ -calculus. In *Automated Deduction, CADE 02*, vol. 2392 of *LNCS*, pp. 423–437. Springer, 2002.
17. K. Larsen and A. Skou. Bisimulation through probabilistic testing. *Inf. Comput.*, 94:1–28, 1991.
18. R. Laver. Well-quasi-orders and sets of finite sequences. *Math. Proc. Cambridge Philos. Soc.*, 79:1–10, 1976.
19. D. Lewis. *Convention, A Philosophical Study*. Harvard University Press, 1969.
20. R. Parikh. The logic of games and its applications. *Annals of Discrete Mathematics*, 24:111–140, 1985.
21. D. Pattinson. Expressive logics for coalgebras via terminal sequence induction. *Notre Dame J. Formal Logic*, 45:19–33, 2004.
22. D. Pattinson and L. Schröder. Cut elimination in coalgebraic logics. *Inf. Comput.* To appear.
23. M. Pauly. A modal logic for coalitional power in games. *J. Log. Comput.*, 12:149–166, 2002.
24. D. Peleg. Concurrent dynamic logic. *J. ACM*, 34:450–479, 1987.
25. V. R. Pratt. Semantical considerations on Floyd-Hoare logic. In *Foundations of Computer Science, FOCS 76*, pp. 109–121. IEEE, 1976.
26. L. Santocanale. Completions of μ -algebras. *Ann. Pure Appl. Logic*, 154:27–50, 2008.
27. L. Santocanale and Y. Venema. Completeness for flat modal fixpoint logics. *Annals Pure Appl. Log.* To appear. Preprint under <http://arxiv.org/abs/0812.2390>.
28. L. Schröder. A finite model construction for coalgebraic modal logic. *J. Log. Algebr. Prog.*, 73:97–110, 2007.
29. L. Schröder and D. Pattinson. PSPACE bounds for rank-1 modal logics. *ACM Trans. Comput. Log.*, 10:13:1–13:33, 2009.
30. L. Schröder and D. Pattinson. Strong completeness of coalgebraic modal logics. In *Theoretical Aspects of Computer Science, STACS 09*, pp. 673–684. Schloss Dagstuhl – Leibniz-Zentrum für Informatik; Dagstuhl, Germany, 2009.
31. L. Schröder, D. Pattinson, and C. Kupke. Nominals for everyone. In *International Joint Conferences on Artificial Intelligence, IJCAI 09*, pp. 917–922. AAAI Press, 2009.
32. K. Segerberg. A completeness theorem in the modal logic of programs. In *Universal Algebra and Applications*, vol. 9 of *Banach Centre Publications*, pp. 31–46. PWN, 1982.
33. I. Walukiewicz. Completeness of Kozen’s axiomatisation of the propositional μ -calculus. *Inf. Comput.*, 157:142–182, 2000.