

Coalgebraic Hybrid Logic

Rob Myers¹, Dirk Pattinson^{*1}, and Lutz Schröder^{**2}

¹ Department of Computing, Imperial College London

² DFKI Bremen and Department of Computer Science, Universität Bremen

Abstract. We introduce a generic framework for hybrid logics, i.e. modal logics additionally featuring *nominals* and *satisfaction operators*, thus providing the necessary facilities for reasoning about individual states in a model. This framework, coalgebraic hybrid logic, works at the same level of generality as coalgebraic modal logic, and in particular subsumes, besides normal hybrid logics such as hybrid K , a wide variety of logics with non-normal modal operators such as probabilistic, graded, or coalitional modalities and non-monotonic conditionals. We prove a generic finite model property and an ensuing weak completeness result, and we give a semantic criterion for decidability in PSPACE. Moreover, we present a fully internalised PSPACE tableau calculus. These generic results are easily instantiated to particular hybrid logics and thus yield a wide range of new results, including e.g. decidability in PSPACE of probabilistic and graded hybrid logics.

Introduction

The ability to represent and reason about individuals is a core feature of many formalisms in the field of logic-based knowledge representation. Individuals may be persons, parts of the human body, or even positions in a strategic game. Reasoning about individuals is a prominent feature in description logics [3], and is supported by a number of reasoning engines, including Fact, Racer and Pellet [34,19,32]. Both description logics and the associated reasoning tools are based on *relational models*, usually Kripke structures with a fixed number of relations. While this is adequate for a large number of applications, description logics can neither formulate nor reason about statements of a non-relational nature, such as assertions involving quantitative uncertainty (‘The likelihood that John is a son of Mary is greater than 23 %’) or non-monotonic conditionals (‘John normally goes to work on Mondays’ — unless e.g. he decides to call in sick).

Features of this kind are usually catered for by specific logics, such as probabilistic modal logic [20] or conditional logic [14], neither of which admits a semantics in terms of Kripke structures. On the other hand, these logics cannot be used off the shelf in most applications in knowledge representation, as they lack facilities to represent and reason about individuals. Of course, one may opt to study extensions of these logics on a case-by-case basis. However, we can do better: both probabilistic modal logic and conditional logic, as well as many others including coalition logic [26], graded modal

* Partially supported by EPSRC grant EP/F031173/1

** Work performed as part of the DFG project *Generic Algorithms and Complexity Bounds in Coalgebraic Modal Logic* (SCHR 1118/5-1)

logic [18], and logics for counterfactual reasoning [22], as well as most description logics naturally fit under the umbrella of *coalgebraic modal logic* [24]. On the semantical side, this is achieved by replacing the notion of *model* of a particular logic by *coalgebras* for an endofunctor T on sets. As we illustrate by means of examples, the semantics of particular logics then arises by instantiating the framework with a concrete endofunctor T . This opens the road for a uniform treatment of a large class of modal logics. In particular, reasoning about individuals in coalgebraic logic is tantamount to extending a large class of logics simultaneously.

This is the starting point of this paper. We introduce *coalgebraic hybrid logic*, that is, an extension of coalgebraic logics with both nominals (that denote individual states) and satisfaction operators (which are used to formulate assertions concerning specific individuals) in the same way that hybrid logic [2] extends the modal logic K . Our main results are concerned with completeness and complexity of coalgebraic hybrid logics. We do not treat the downarrow binder \downarrow , as hybrid logics involving \downarrow are known to be undecidable, even in the special case of relational (Kripke) semantics [1]. As we are working in a general framework, these statements are formulated in terms of abstract *coherence conditions* that relate the syntax (modal operators and proof rules) and the semantics, which is abstractly given in terms of coalgebras. Conceptually, we show that the *same* coherence conditions that give rise to completeness and complexity results in the non-hybrid case also guarantee completeness and (the same) complexity bounds for the hybrid extension. In more detail, we prove completeness of hybrid coalgebraic logics both with respect to a Hilbert-style proof calculus in Section 2 and a cut-free sequent system 3. We exploit the latter to show that the satisfiability problem is decidable in polynomial space by means of backward proof search in Section 4 before we give a purely semantical account of complexity in Section 5. As the coherence conditions that guarantee completeness and complexity are already known to hold for a large class of logics, instantiations of the general framework give rise to a number of new results concerning particular logics:

- Both *hybrid conditional logic* and *hybrid probabilistic logic* are complete and decidable in PSPACE
- Moreover, *graded hybrid logic* is decidable in polynomial space with numbers coded in binary, and these results immediately carry over to a setting with role hierarchies, giving a new (tight) PSPACE upper bound for the description logic \mathcal{ALCHQ} over the empty TBox.
- The semantic analysis yields previously unknown PSPACE-upper bounds for *Presburger hybrid logic* [16] and a version of probabilistic modal logic featuring linear inequalities [17].

In a nutshell, the addition of nominals and satisfaction operators greatly increases the expressive power of (coalgebraic) modal logics and is still decidable in PSPACE, i.e. the same complexity class as the modal logic K . In particular, the ability to cater for a large number of logical features, such as quantitative uncertainty and non-monotonic conditionals in the coalgebraic framework offers a new (coalgebraic) perspective on description logics, with the vision of being able to represent a large class of concepts of information-theoretic relevance in a single framework.

1 Syntax and Semantics of Coalgebraic Hybrid Logic

To make our treatment parametric in the concrete syntax of any given modal logic, we fix a modal similarity type Λ consisting of modal operators with associated arities throughout. For given countably infinite and disjoint sets P of propositional variables and N of nominals, the set $\mathcal{F}(\Lambda)$ of *hybrid Λ -formulas* is given by the grammar

$$\mathcal{F}(\Lambda) \ni \phi, \psi ::= p \mid i \mid \phi \wedge \psi \mid \neg\phi \mid \heartsuit(\phi_1, \dots, \phi_n) \mid @_i\phi$$

where $p \in P$, $i \in N$ and $\heartsuit \in \Lambda$ is an n -ary modal operator. We use the standard definitions for the other propositional connectives $\rightarrow, \leftrightarrow, \vee$. The set of nominals occurring in a formula ϕ is denoted by $N(\phi)$, and the nesting depth of modal operators (excluding satisfaction operators) by $\text{rank}(\phi)$. A formula of the form $@_i\phi$ is called an *@-formula*. Semantically, nominals i denote individual states in a model, and an @-formula $@_i\phi$ stipulates that ϕ holds at i .

To reflect parametricity in the particular underlying logic also semantically, we equip hybrid logics with a *coalgebraic semantics* extending the standard coalgebraic semantics of modal logics [24]: we fix throughout a Λ -structure consisting of an endofunctor $T : \text{Set} \rightarrow \text{Set}$ on the category of sets, together with an assignment of an n -ary *predicate lifting* $\llbracket \heartsuit \rrbracket$ to every n -ary modal operator $\heartsuit \in \Lambda$, i.e. a set-indexed family of mappings $(\llbracket \heartsuit \rrbracket_X : \mathcal{P}(X)^n \rightarrow \mathcal{P}(TX))_{X \in \text{Set}}$ that satisfies

$$\llbracket \heartsuit \rrbracket_X \circ (f^{-1})^n = (Tf)^{-1} \circ \llbracket \heartsuit \rrbracket_Y$$

for all $f : X \rightarrow Y$. In categorical terms, $\llbracket \heartsuit \rrbracket$ is a natural transformation $\mathcal{Q}^n \rightarrow \mathcal{Q} \circ T^{op}$ where $\mathcal{Q} : \text{Set}^{op} \rightarrow \text{Set}$ is the contravariant powerset functor.

In this setting, T -coalgebras play the roles of *frames*. A T -coalgebra is a pair (C, γ) where C is a set of *states* and $\gamma : C \rightarrow TC$ is the *transition function*. If clear from the context, we identify a T -coalgebra (C, γ) with its state space C . A (*hybrid*) T -model (C, γ, π) consists of a T -coalgebra (C, γ) together with a *hybrid valuation* π , i.e. a map $P \cup N \rightarrow \mathcal{P}(C)$ that assigns singleton sets to all nominals $i \in N$. We often identify the singleton set $\pi(i)$ with its unique element.

The semantics of $\mathcal{F}(\Lambda)$ is a satisfaction relation \models between states $c \in C$ in hybrid T -models $M = (C, \gamma, \pi)$ and formulas $\phi \in \mathcal{F}(\Lambda)$, inductively defined as follows. For $x \in N \cup P$ and $i \in N$, put

$$c, M \models x \text{ iff } c \in \pi(x) \quad c, M \models @_i\phi \text{ iff } \pi(i), M \models \phi.$$

Modal operators are interpreted using their associated predicate liftings, that is,

$$c, M \models \heartsuit(\phi_1, \dots, \phi_n) \iff \gamma(c) \in \llbracket \heartsuit \rrbracket_C(\llbracket \phi_1 \rrbracket_M, \dots, \llbracket \phi_n \rrbracket_M)$$

where $\heartsuit \in \Lambda$ n -ary and $\llbracket \phi \rrbracket_M = \{c \in C \mid M, c \models \phi\}$ denotes the truth-set of ϕ relative to M . A formula ϕ is *satisfiable* if it is satisfied in some state in some model, and *valid*, written $\models \phi$, if it is satisfied in all states in all models.

The distinguishing feature of the coalgebraic approach to hybrid and modal logics is the parametricity in both the logical language and the notion of frame: concrete instantiations of the general framework, in other words a choice of modal operators Λ and a Λ -structure T , capture the semantics of a wide range of modal logics, as witnessed by the following examples.

Example 1. 1. The hybrid version of the modal logic K , *hybrid K* for short, has a single unary modal operator \Box , interpreted over the structure given by the powerset functor \mathcal{P} that takes a set X to its powerset $\mathcal{P}(X)$ and $\llbracket \Box \rrbracket_X(A) = \{B \in \mathcal{P}(X) \mid B \subseteq A\}$. It is clear that \mathcal{P} -coalgebras $(C, \gamma : C \rightarrow \mathcal{P}(C))$ are in 1-1 correspondence with Kripke frames, and that the coalgebraic definition of satisfaction specialises to the usual semantics of the box operator.

2. *Graded hybrid logic* has modal operators \Diamond_k ‘in more than k successors, it holds that’. It is interpreted over the functor \mathcal{B} that takes a set X to the set $\mathcal{B}(X)$ of multisets over X , i.e. maps $B : X \rightarrow \mathbb{N} \cup \{\infty\}$, by $\llbracket \Diamond_k \rrbracket_X(A) = \{B \in \mathcal{B}(X) \mid \sum_{x \in A} B(x)\}$. This captures the semantics of graded modalities over *multigraphs* [15], which are precisely the \mathcal{B} -coalgebras. Unlike in the modal case [27], the multigraph semantics does not engender the same notion of satisfiability as the more standard Kripke semantics of graded modalities, as the latter validates all formulas $\neg \Diamond_1 i$, $i \in \mathbb{N}$. One can however polynomially encode Kripke semantics into multigraph semantics: a graded hybrid formula ϕ , w.l.o.g. an @-formula, is satisfiable over Kripke frames iff $\phi \wedge \bigwedge_{j \in \mathbb{N}(\phi)} @_j \bigwedge_{n < \text{rank}(\phi)} \Box_0^n \bigwedge_{i \in \mathbb{N}(\phi)} \neg \Diamond_1 i$, where \Box_k is defined as $\neg \Diamond_k \neg$ and \Box_0^n denotes n boxes, is satisfiable over multigraphs. Thus, our completeness and complexity results for multigraph semantics derived below transfer to Kripke semantics. In particular they apply to many description logics, which commonly feature both nominals and graded modal operators (*qualified number restrictions*).

3. *Probabilistic hybrid logic*, the hybrid extension of probabilistic modal logic [21], has modal operators L_p ‘in the next step, it holds with probability at least p that’, for $p \in [0, 1] \cap \mathbb{Q}$. It is interpreted over the functor D_ω that maps a set X to the set of finitely-supported probability distributions on X by putting $\llbracket L_p \rrbracket_X(A) = \{P \in D_\omega(X) \mid PA \geq p\}$. Coalgebras for D_ω are just Markov chains. A simple valid formula of hybrid probabilistic logic is $L_p i \wedge L_q j \rightarrow L_q i \vee L_1(\neg(i \wedge j))$, $i, j \in \mathbb{N}$.

4. *Hybrid CK*, the hybrid extension of the basic conditional logic CK , has a single binary modal operator \Rightarrow , written in infix notation. Hybrid CK is interpreted over the functor Cf that maps a set X to the set of $\mathcal{P}(X) \rightarrow \mathcal{P}(X)$, whose coalgebras are selection function models [14], by putting $\llbracket \Rightarrow \rrbracket_X(A, B) = \{f : \mathcal{P}(X) \rightarrow \mathcal{P}(X) \mid f(A) \subseteq B\}$.

2 A Generic Complete Hilbert System

We proceed to present a Hilbert-system for coalgebraic hybrid logics, and prove its soundness and its weak completeness over finite models, provided that the logic at hand satisfies certain coherence conditions between the axiomatisation and the semantics — in fact the *same* conditions as in the modal case, which are easily verified *local* properties that can be verified without reference to T -models and are already known to hold for a large variety of logics [24,27]. We recall notation from earlier work:

Notation 2. For a set or multiset Σ of formulas and a set O of operators, we write $O\Sigma$ or $O(\Sigma)$ for the set or multiset of formulas arising by prefixing elements of Σ with exactly one operator from O ; e.g. $\Lambda(\Sigma) = \{\heartsuit(\phi_1, \dots, \phi_n) \mid \heartsuit \in \Lambda \text{ } n\text{-ary}, \phi_1, \dots, \phi_n \in \Sigma\}$ and $@\Sigma := \{@_i \mid i \in \mathbb{N}\}(\Sigma) = \{@_i \phi \mid i \in \mathbb{N}, \phi \in \Sigma\}$. Moreover, $\text{Prop}(\Sigma)$

denotes the set of propositional combinations of Σ -formulas. For a propositional formula $\phi \in \text{Prop}(\mathbf{P} \cup \mathbf{N})$, we write $\kappa \models \phi$ if $\kappa : \mathbf{P} \cup \mathbf{N} \rightarrow 2 = \{\perp, \top\}$ is a satisfying valuation for ϕ , and $X, \tau \models \phi$ if ϕ evaluates to \top in the boolean algebra $\mathcal{P}(X)$ under a hybrid valuation $\tau : \mathbf{P} \cup \mathbf{N} \rightarrow \mathcal{P}(X)$. For $\psi \in \text{Prop}(\Lambda(\mathbf{P} \cup \mathbf{N}))$, 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 $\llbracket \heartsuit(p_1, \dots, p_n) \rrbracket_{TX, \tau} = \llbracket \heartsuit \rrbracket_X(\tau(p_1), \dots, \tau(p_n))$. We write $TX, \tau \models \psi$ if $\llbracket \psi \rrbracket_{TX, \tau} = TX$, and $t \models_{TX, \tau} \psi$ if $t \in \llbracket \psi \rrbracket_{TX, \tau}$.

Proof systems for coalgebraic logics are most conveniently described in terms of one-step rules, for which the announced coherence conditions take the following form.

Definition 3. A *one-step rule* over Λ is an $n + 1$ -tuple $r = (\Gamma_1, \dots, \Gamma_n, \Gamma_0)$, written

$$\frac{\Gamma_1 \dots \Gamma_n}{\Gamma_0},$$

where $\Gamma_0 \subseteq \Lambda(\mathbf{P}) \cup \neg\Lambda(\mathbf{P})$ and $\Gamma_1, \dots, \Gamma_n \subseteq \mathbf{P} \cup \neg\mathbf{P}$. The rule $\Gamma_1 \dots \Gamma_n / \Gamma_0$ is *one-step sound* if $TX, \tau \models \bigvee \Gamma_0$ whenever $X, \tau \models \bigvee \Gamma_i$ for all $i = 1, \dots, n$. A set \mathcal{R} of one-step rules is *one-step (cut-free) complete* if whenever $TX, \tau \models \bigvee \Delta$ for $\Delta \subseteq \Lambda(\mathbf{P}) \cup \neg\Lambda(\mathbf{P})$ then $\bigvee \Delta$ is propositionally entailed by formulas $\bigvee \Gamma_0 \sigma$ ($\Delta \supseteq \Gamma_0$ is a super-sequent of a single sequent Γ_0) where $\Gamma_1 \dots \Gamma_n / \Gamma_0$ is in \mathcal{R} and $\sigma : \mathbf{P} \rightarrow \text{Prop}(\mathbf{P})$ is a substitution such that $X, \tau \models \bigvee \Gamma_i \sigma$ for all i .

Example 4. One-step cut-free complete rule sets, which can just be inherited from the corresponding modal systems, for graded and probabilistic logics, conditional logics, and many others are found in [28,25]. We recall that the one-step cut-free complete rule set for (hybrid) K consists of the rules

$$\frac{\neg a_1, \dots, \neg a_n, b}{\neg \Box a_1, \dots, \neg \Box a_n, \Box b} \quad (n \geq 0)$$

A set \mathcal{R} of one-step rules now gives rise to a Hilbert-system HR by adjoining the congruence rule (replacement of equivalents under modal operators), propositional tautologies, modus ponens, uniform substitution, standard axioms for satisfaction operators $@_i$ stating that $@_i$ commutes with all propositional connectives ($\neg @_i \perp$, $\neg @_i \phi \leftrightarrow @_i \neg \phi$, $@_i(\phi \wedge \psi) \leftrightarrow (@_i \phi \wedge @_i \psi)$) and that $i \sim j \equiv @_i j$ defines an equivalence relation on nominals ($@_i i$, $@_i j \leftrightarrow @_j i$, $@_i k \wedge @_j p \rightarrow @_i p$), the $@$ -generalisation rule $a / @_i a$, and the axioms

$$\begin{aligned} \text{(in)} \quad & i \wedge \phi \rightarrow @_i \phi \\ \text{(mob)} \quad & @_i p \rightarrow (\heartsuit(q_1, \dots, q_n) \leftrightarrow \heartsuit(@_i p \wedge q_1, \dots, @_i p \wedge q_n)) \end{aligned}$$

called $@$ -introduction (in) and *make-or-break* (mob), respectively. We write $\text{HR} \vdash \phi$ if ϕ is derivable in this system. The (mob) axiom captures the fact that the truth set of an $@$ -formula is either empty or the whole model; in the case of hybrid K , it is deductively equivalent to the standard back axiom $@_i \phi \rightarrow \Box @_i \phi$.

From $@$ -introduction one easily derives $@$ -elimination $i \wedge @_i a \rightarrow a$. Moreover, the (mob) axiom readily generalises to any $@$ -formula in place of $@_i p$. Hence we can derive

a relativised congruence rule

$$(@cong) \frac{\phi \rightarrow ((a_1 \leftrightarrow b_1) \wedge \dots \wedge (a_n \leftrightarrow b_n))}{\phi \rightarrow (\heartsuit(a_1, \dots, a_n) \leftrightarrow \heartsuit(b_1, \dots, b_n))} (\phi \text{ @-formula})$$

It is clear that HR is sound if all rules in \mathcal{R} are one-step sound. For the remainder of the section, we assume that \mathcal{R} is *one-step complete* and proceed to prove weak completeness of HR over finite models by extending the finite model construction of [27], i.e. by constructing a model for a consistent formula ϕ whose states are maximally consistent subsets of a suitable finite closure of ϕ .

In the following, let Σ be a finite and *closed* set of formulas, i.e. closed under subformulas (where we count i as a subformula of $@_i\rho$), negation, and $@_i$ for $i \in \mathbb{N}(\Sigma)$, where we identify $\neg\neg\phi$ with ϕ , $\neg@_i\phi$ with $@_i\neg\phi$, and $@_i@_j\phi$ with $@_j\phi$. Denote by $@\Sigma$ the set of @-formulas in Σ . Fix a maximally HR -consistent subset K of $@\Sigma$. Let S denote the set of *atoms*, i.e. maximally HR -consistent subsets of Σ , and let S_K be the set of *K-atoms*, i.e. atoms containing K .

Lemma 5. *Let $\phi \in \text{Prop}(\mathbb{P})$, and let σ be a Σ -substitution. Then $K \rightarrow \phi\sigma$ is derivable iff $S_K, \tau \models \phi$, where τ is the $\mathcal{P}(S_K)$ -valuation $\tau(a) = \{A \in S_K \mid \sigma(a) \in A\}$.*

We define a hybrid $\mathcal{P}(S_K)$ -valuation π in the standard way by $\pi(a) = \{A \mid a \in A\}$ for $a \in \mathbb{P} \cup \mathbb{N}(\Sigma)$, and taking $\pi(i)$ to be an arbitrary singleton set otherwise.

Lemma 6. *The valuation τ is hybrid, i.e. $\tau(i)$ is a singleton for each $i \in \mathbb{N}(\Sigma)$, namely $\tau(i) = \{K_i\}$, where $K_i = \{\phi \in \Sigma \mid @_i\phi \in K\}$.*

Call a coalgebra (S_K, ξ) *coherent* if for all $\heartsuit(\phi_1, \dots, \phi_n) \in \Sigma$, $A \in S_K$,

$$\xi(A) \in \llbracket \heartsuit \rrbracket(\hat{\phi}_1, \dots, \hat{\phi}_n) \iff \heartsuit(\phi_1, \dots, \phi_n) \in A$$

where $\hat{\phi} = \{A \in S_K \mid \phi \in A\}$. Making crucial use of the relativised congruence rule (@-cong), one proves

Lemma 7 (Relativised existence lemma). *There exists a coherent coalgebra (S_K, ξ) .*

It is then straightforward to establish

Lemma 8 (Truth lemma). *If (S_K, ξ) is coherent, then $A \models_{(S_K, \xi, \pi)} \phi$ iff $\phi \in A$ for all $\phi \in \Sigma$, $A \in S_K$.*

Theorem 9 (Weak completeness of HR over finite models). *Every HR -consistent hybrid formula ϕ is satisfiable in a T -model with at most $2^{|\phi|}$ states.*

The above theorem establishes weak completeness of the Hilbert calculus not only for hybrid K , but also for graded and probabilistic hybrid logic and hybrid CK , as well as hybrid versions of many other modal logics treated e.g. in [27,28].

3 Hybrid Sequent Calculi and Cut Elimination

We now introduce a sound and complete sequent calculus for coalgebraic hybrid logics. Completeness of the calculus, initially with the cut-rule, is proved using completeness of the Hilbert system of the previous section by showing that both provability predicates coincide. Subsequently, we prove cut elimination.

The sequent calculus shares one characteristic trait with the tableaux calculus studied in [6]: sequents are composed of @-prefixed formulas. As a consequence, the sequent calculus presented here can be understood as (the dual of) a labelled tableau [35] by reading an @-prefixed formula $@_i\phi$ as a labelled sequent $i : \phi$. As universal validity of $\phi \in \mathcal{F}(A)$ is equivalent to validity of $@_t\phi$ for some $t \in \mathbb{N}$ which doesn't occur in ϕ , this suffices to obtain completeness. The main difference between our calculus and those found in the literature is its modularity: it comes about by extending a standard, cut-free sequent calculus for the underlying logic by means of proof rules for nominals and satisfaction operators. We first consider a sequent calculus with the cut-rule and relegate cut-elimination to the next section. We begin by fixing our notation regarding sequent calculi; in the tradition of the sequent calculus literature, we denote formulas by capital letters A, B, \dots for the next two sections.

Definition 10. If $\Sigma \subseteq @F(A)$ is a set of formulas, a Σ -sequent is a finite multiset of formulas in Σ . We write $\mathcal{S}(\Sigma)$ for the set of Σ -sequents, and \mathcal{S} for the set of $@F(A)$ -sequents. If $\Gamma, \Delta \in \mathcal{S}$, then Γ, Δ denotes their multiset union, and we identify the singleton sequent $\{A\}$ with A for $A \in @F(A)$. We put $\text{rank}(\Gamma) = \max\{\text{rank}(A) \mid A \in \Gamma\}$.

By virtue of the above definition, all elements of a sequent are necessarily @-prefixed formulas. As the modal logics we consider are extensions of classical propositional logic, it is most convenient to use a right-handed (or Gentzen-Schütte) calculus where a sequent intuitively stands for the disjunction of its elements. The rules for propositional reasoning and reasoning about names then take the following form:

$$\begin{array}{ll}
(\text{Ax}) \quad @_t\neg A, @_t A, \Gamma & (\text{Ref}) \quad @_t t, \Gamma \quad (@\top) \quad @_t \top, \Gamma \\
(\neg\neg) \quad \frac{@_t A, \Gamma}{@_t \neg\neg A, \Gamma} & (\wedge) \quad \frac{@_t A, \Gamma \quad @_t B, \Gamma}{@_t (A \wedge B), \Gamma} \\
(\neg\wedge) \quad \frac{@_t \neg A, @_t \neg B, \Gamma}{@_t \neg(A \wedge B), \Gamma} & (\text{At}) \quad \frac{@_t A, \Gamma}{@_s @_t A, \Gamma} \\
(\text{Sd}) \quad \frac{@_s \neg A, \Gamma}{@_t \neg @_s A, \Gamma} & (\text{Eq}) \quad \frac{\Gamma[t := i]}{@_t \neg i, \Gamma}
\end{array}$$

In the above, $s, t \in \mathbb{N}$, $A, B \in @F(A)$, and $\Gamma \in \mathcal{S}(@F(A))$.

In combination with a set \mathcal{R} of one-step sound one-step rules, which we fix throughout, we obtain the following notion of derivability.

Definition 11. The set of \mathcal{R} -derivable sequents is the least set that

- contains all instances of (Ax), (Ref) and ($@\top$)
- is closed under the rules ($\neg\neg$), (\wedge), ($\neg\wedge$), (At), (Sd) and (Eq)
- is closed under the rules

$$(R) \frac{@_n\Gamma_1\sigma, @_t\Gamma_0\sigma, \Delta \quad \dots \quad @_n\Gamma_k\sigma, @_t\Gamma_0\sigma, \Delta}{@_t\Gamma_0\sigma, \Delta} (n \notin \Gamma_0\sigma, \dots, \Gamma_k\sigma, \Delta)$$

where $\Gamma_1, \dots, \Gamma_k/\Gamma_0 \in \mathcal{R}$, σ is a substitution and $\Delta \in \mathcal{S}$. In the side condition, $n \notin \Sigma$ denotes that the nominal n does not occur in the sequent Σ . The instances of the rule schema (R) are called *modal rules* and the remaining rules are referred to as *static rules*.

We write $\mathcal{GR} \vdash \Gamma$ if Γ is \mathcal{R} -derivable and \mathcal{GR} for the above set of rules. If Γ is \mathcal{R} -derivable with additional help of the cut-rule

$$(\text{cut}) \frac{\Gamma, @_tA \quad \Delta, @_t\neg A}{\Gamma, \Delta}$$

then this is denoted by $\mathcal{GRC} \vdash \Gamma$. We write $\mathcal{GR} \vdash_n \Gamma$ if there exists a proof tree with end sequent Γ where (R) has been applied at most n times on any branch.

Proposition 12 (Soundness). $\bigvee \Gamma$ is valid whenever $\mathcal{GRC} \vdash \Gamma$.

The completeness of \mathcal{GRC} is witnessed by the fact that every proof in the (complete) system \mathcal{HR} can be simulated. More precisely:

Theorem 13. Let $\mathcal{HR} \vdash A$. Then for every $t \notin A$, $\mathcal{GRC} \vdash @_tA$.

Together with the completeness theorem for the Hilbert-system, we obtain completeness as an easy corollary.

Theorem 14. Let \mathcal{R} be one-step complete. Then for every $\Gamma \in \mathcal{S}$, $\mathcal{GRC} \vdash \Gamma$ whenever $\models \bigvee \Gamma$.

Our main motivation for introducing the system \mathcal{GR} is to determine the complexity of the satisfiability problem by means of proof search. Hence, completeness of \mathcal{GRC} is insufficient, as the use of the cut-rule leads to an infinite search space. Our next goal is therefore cut-elimination.

Cut elimination is subject to a number of structural properties that are readily established inductively. We begin with the inversion lemma.

Lemma 15 (Inversion Lemma). All static rules of \mathcal{GR} are invertible, i.e. if $\Gamma_1 \dots \Gamma_k/\Gamma_0$ is a static rule of \mathcal{GR} , then $\mathcal{GR} \vdash_k \Gamma_0$ iff $\mathcal{GR} \vdash_k \Gamma_i$ for all $i = 1, \dots, k$ and all $k \geq 0$.

It is easy to see that weakening is also admissible, and moreover does not increase the number of applications of (R) in a proof tree.

Concerning the admissibility of cut and contraction, we adapt the standard double induction method [33] and use an additional outermost induction on the modal depth of the endsequent. The key point here is to observe that an instance of cut or contraction in the principal sequent of an instance of (R) can be eliminated by virtue of one-step cutfree completeness using a different instance of (R).

Theorem 16. *Contraction and cut are admissible in \mathcal{GR} , that is $\mathcal{GR} \vdash_n \Gamma, A$ whenever $\mathcal{GR} \vdash_n \Gamma, A, A$, and $\mathcal{GR} \vdash_n \Gamma, \Delta$ whenever $\mathcal{GR} \vdash_n \Gamma, A$ and $\mathcal{GR} \vdash_n \Delta, \neg A$, for all $\Gamma, \Delta \in \mathcal{S}$ and all $A \in @F(A)$.*

The proof establishes moreover that neither cut elimination nor contraction increase the number of applications of (R) along any branch of the proof tree.

4 Complexity of Proof Search

While cutfreeness is clearly essential to establish complexity bounds by means of proof search, the duplication of the conclusion in the application of the sequent version of a one-step rule remains a further potential source of non-termination. We now proceed to show that the height of a proof tree in the sequent calculus of Section 3 can be bounded polynomially in the size of the endsequent. As a consequence, we obtain a polynomial space bound for proof search, and – dually – for the satisfiability problem.

Definition 17. A *pseudo-subformula* of Γ is of the form B or $\neg B$ for a subformula B of A with $@_t A \in \Gamma$. The set of pseudo-subformulas of Γ is denoted by $\text{PSF}(\Gamma)$. A *core formula* of Γ is a formula $@_t A$ such that $t \in \mathbf{N}(\Gamma)$ and A can be obtained from $B \in \text{PSF}(\Gamma)$ by means of a sequence of renamings that only affect nominals in $\mathbf{N}(\Gamma)$, formally $A = B[i_1 := j_1] \dots [i_n := j_n]$ where $i_1, \dots, i_n, j_1, \dots, j_n \in \mathbf{N}(\Gamma)$ and $\mathbf{N}(\Gamma) = \bigcup \{ \mathbf{N}(A) \mid A \in \Gamma \}$.

Lemma 18. *Let $\mathcal{GR} \vdash \Gamma$. Then the collection of formulas appearing on any branch of a proof tree of Γ contains at most $|\text{PSF}(\Gamma)| \cdot |\mathbf{N}(\Gamma)|^2$ core formulas.*

The previous lemma allows us to argue that we can eliminate branches of the proof tree in case they do not add new core formulas.

Theorem 19. *Let $\mathcal{GR} \vdash \Gamma$. Then Γ has a proof where at most $|\mathbf{N}(\Gamma)|^2 \cdot |\text{PSF}(\Gamma)| \cdot \text{rank}(\Gamma)$ instances of (R) are applied on every branch.*

For the proof of this theorem, we adopt the technique of [8] and attach a label $(t \rightarrow n)$ to an instance of the rule (R) as in Definition 11. Note that the modal depth of the formula decreases when we move from conclusion to premise. We next show that it suffices to consider applications of (R) where either a new core formula is introduced or a new nominal which is used higher up in the branch. In this situation, the collection of labels $(t \rightarrow n)$ forms a forest with at most $|\mathbf{N}(\Gamma)|^2 \cdot |\text{PSF}(\Gamma)|$ leaves by Lemma 18, as every leaf corresponds to an application of (R) that either introduces a new core formula or has empty premise. As the modal depth decreases with each application of (R) , this gives an overall bound of $|\mathbf{N}(\Gamma)|^2 \cdot |\text{PSF}(\Gamma)| \cdot \text{rank}(\Gamma)$ on the number of times (R) was applied.

This bound on the proof depth implies that proof search is in PSPACE, provided the rules are such that the set of possible premises of rule applications can be computed from the conclusions in nondeterministic polynomial time. We adapt the treatment of [28] to the sequent format of one-step rules used here. Recall that a nondeterministic polynomial time multivalued (NPMV) function [9] is a function $f : \Sigma^* \rightarrow \mathcal{P}\Delta^*$, where

Σ and Δ are alphabets, such that there exists a polynomial p such that $|y| \leq p(|x|)$ for all $y \in f(x)$, where $|\cdot|$ denotes size, and the graph $\{(x, y) \mid y \in f(x)\}$ of f is in NP . With a view to implementing proof search on an alternating turing machine [13], this leads to the following definition.

Definition 20. Let sequents be represented in Σ^* for a finite alphabet Σ . The rule set \mathcal{R} is *tractable* if there are NPMV functions $f : \Sigma^* \rightarrow \mathcal{P}(\Sigma^*)$ and $g : \Sigma^* \rightarrow \mathcal{P}(\Sigma^*)$ such that $\{\{\Gamma_1, \dots, \Gamma_n\} \mid \Gamma_1, \dots, \Gamma_n / \Gamma \in \mathcal{GR}\} = \{g(x) \mid x \in f(\Gamma)\}$ for all $\Gamma \in \mathcal{S}$.

This allows us to formulate the main result of the present section as follows:

Theorem 21. *Let \mathcal{R} be one-step sound, one-step cut-free complete and tractable. Then satisfiability of $A \in \mathcal{F}(\Lambda)$ is in PSPACE.*

Example 22. As all hybrid logics of Example 1 have previously [28,25] been equipped with rule sets satisfying the assumptions of Theorem 21, they are decidable in polynomial space. In the case of hybrid K , this re-proves a known result [2]. The PSPACE-bounds for graded and probabilistic hybrid logic and for hybrid CK are new.

5 Shallow Models and PSPACE Algorithms

Next, we establish a semantics-based criterion for a hybrid logic to be decidable in PSPACE, thus complementing the sequent calculus based approach above by a method that applies also to logics for which no tractable cut-free axiomatisation is known. To this end, we extend the shallow model construction of [30] to the hybrid case. The extension is quite non-trivial for two reasons: shallow models are now forest-shaped rather than tree-shaped; and moreover they are not perfect forests in that they may have loops into the roots. The coalgebraic constructions require the novel concept of *fragments*, i.e. models based on partial coalgebras. We assume that Λ is equipped with a size measure, thus inducing a size measure $|\cdot|$ on $\mathcal{F}(\Lambda)$, with numbers coded in binary. As we now leave the sequent calculus context, we return to designating formulas by small greek letters.

Definition 23. A model *satisfies* a set K of @-formulas if its states satisfy K . A formula ϕ is *K-satisfiable* if it is satisfiable in some model satisfying K , and *K-valid* if it is valid in all models satisfying K .

Let ϕ be satisfied in a model M . Then ϕ is *K-satisfiable*, where K is the set of @-formulas satisfied in M . Moreover, ϕ is satisfiable iff $@_t\phi$ is satisfiable for t fresh. Thus, we assume w.l.o.g. that ϕ is an @-formula. Then K entails ϕ , so that we can forget about ϕ and concentrate on models satisfying K . In the following we fix a finite closed set Σ and a maximally satisfiable set $K \subseteq @\Sigma$, and we put $N = \mathbb{N}(\Sigma)$. For $i \in N$, we put $K_i = \{\rho \mid @_i\rho \in K\}$ as in Section 2.

Definition 24. A hybrid formula is *@-free* if it does not contain occurrences of @. A set of @-formulas is *@-eliminated* if it consists of formulas $@_i\rho$ with ρ @-free. For $\rho \in \Sigma$, $\rho[K]$ denotes the @-free formula obtained by replacing every subformula $@_i\chi$ of ρ not contained in further occurrences of @ by \top if $@_i\chi \in K$, and by \perp otherwise. The *@-eliminated form* $K[K]$ of K is the @-eliminated set $\{@_i\rho[K] \mid @_i\rho \in K\}$.

Lemma 25 (@-Elimination). *A model satisfies K iff it satisfies $K[K]$.*

By @-elimination, we may henceforth assume that K is @-eliminated and hence that the K_i are @-free. We wish to construct a model which satisfies K and which is shallow in the following sense:

Definition 26. A *supporting Kripke frame* of a T -model (C, γ, π) is a Kripke frame (C, R) such that $\gamma(c) \in T\{d \mid cRd\} \subseteq TC$ for all $c \in C$. A T -model of K is *shallow* if it has a supporting Kripke frame which is a forest of depth at most the rank of K up to loops into the roots, and whose roots have names in N .

We now introduce the crucial notion of fragment:

Definition 27. A K -*fragment* is a *partial T -model*, i.e. a triple $F = (C, \gamma, \pi)$ consisting of the same data as a T -model except that $\gamma : C \rightarrow TC$ is a *partial* map, with $K_i \in C$ for all $i \in N$, such that γ is undefined precisely on the K_i and $\pi(i) = \{K_i\}$ for $i \in N$. We define K -*fragmentary satisfaction* $c \models_F^K \rho$ of $\rho \in \Sigma$ in $c \in C$ by

$$K_i \models_F^K \rho \quad \text{iff} \quad \rho \in K_i$$

for $i \in N$, and by the usual recursive clauses for the other states. We put $\llbracket \rho \rrbracket_F^K = \{c \mid c \models_F^K \rho\}$. We say that F is a *shallow K -fragment model* of $\rho \in \Sigma$ if F has a supporting Kripke frame (defined in analogy to Definition 26) which is, up to possible isolated states K_i , a tree with root r of depth at most the rank of ρ such that $r \models_F^K \rho$.

Lemma 28. *Every @-free K -satisfiable $\psi \in \text{Prop}(\Sigma)$ has an exponentially branching shallow K -fragment model.*

The proof is largely analogous to the modal shallow model theorem [30], the crucial point being that one need not yet define the coalgebra structure on the named states K_i .

Definition 29. Let $\kappa : N \rightarrow (V \cup N) \rightarrow 2$ be an N -indexed family of valuations for $V \cup N$, where $V \subseteq P$. A *one-step (κ -)model* (X, τ, t) over V consists of a set X , a hybrid $\mathcal{P}(X)$ -valuation τ for $V \cup N$, and $t \in TX$ (such that $\tau(i) \subseteq \tau(a)$ iff $\kappa(i)(a) = \top$ for all $i \in N, a \in V$). A *one-step N -pair* (η, ψ) over V consists of formulas $\psi \in \text{Prop}(\Lambda(V \cup N))$ and $\eta \in \text{Prop}(V \cup N)$. We say that (X, τ, t) is a *one-step (κ -)model of (η, ψ)* if $X, \tau \models \eta$ and $t \models_{TX, \tau} \psi$. The pair (η, ψ) is *one-step (κ -)satisfiable* if it has a one-step (κ -)model, and *one-step satisfiable over (X, τ)* if it has a one-step model of the form (X, τ, t) .

Lemma 30. *A one-step N -pair (η, ψ) over V is one-step κ -satisfiable for $\kappa : N \rightarrow (V \cup N) \rightarrow 2$ iff it is one-step κ -satisfiable over the set*

$$X = \{\iota : V \cup N \rightarrow 2 \mid \iota \models \eta; \iota(i) = \top \Rightarrow \iota(a) = \kappa(i)(a) \text{ for all } i \in N, a \in V\},$$

equipped with the (hybrid) valuation $\tau(b) = \{\iota \in X \mid \iota(b) = \top\}$ for all $b \in V \cup N$.

Definition 31. For $x \in X$ and a hybrid $\mathcal{P}(X)$ -valuation τ for $V \cup N$, where $V \subseteq P$, we put $\text{Th}_\tau(x) \equiv \bigwedge_{x \in \tau(a)} a \wedge \bigwedge_{x \notin \tau(a)} \neg a$, where a ranges over $V \cup N$.

Theorem 32. *The (satisfiable) set K is satisfiable in an exponentially branching shallow T -model.*

In the proof, one constructs a forest-shaped model with roots K_i . One generates suitable exponential-size one-step models according to Lemma 30 that induce the coalgebra structure on the K_i , and then attaches shallow K -fragments obtained from Lemma 28 at the arising successor states, merging identically named states throughout the process.

This construction informs the design of a decision procedure for satisfiability of a hybrid formula ϕ . The algorithm first performs a non-deterministic reduction to satisfiability of $K \subseteq @ \Sigma$ as above and @-elimination, and then reduces satisfiability of K to fragmentary K -satisfiability of certain formulas $\psi \in \text{Prop}(\Sigma)$. The latter is decided by the following recursive procedure.

Algorithm 33 (Check fragmentary K -satisfiability of $\psi \in \text{Prop}(\Sigma)$).

1. Decompose $\psi \equiv \psi_0 \sigma$ with $\psi_0 \in \Lambda(V) \cup \neg \Lambda(V)$ and σ a Σ -substitution for $V \subseteq P$.
2. Recursively compute the *propositional theory* $\eta \in \text{Prop}(V \cup N)$, $N = N(\Sigma)$, of σ as the disjunction of all maximal conjunctive clauses χ over $V \cup N$ (i.e. for each $b \in V \cup N$, χ contains either b or $\neg b$) such that $\chi \sigma$ is fragmentarily K -satisfiable.
3. Check that (η, ψ_0) is one-step κ -satisfiable, where for $i \in N$ and $a \in V$, $\kappa(i)(a) = \top$ iff $@_i \sigma(a) \in K$.

In the implementation of this algorithm (and analogously in the reduction step that precedes it), one cannot keep the whole (potentially exponential-sized) formula η in memory at once. Instead, read access to η by the one-step satisfiability checking procedure queries whether individual conjunctive clauses χ belong to η . Since the recursion depth is bounded by $\text{rank}(\psi)$, such an implementation will run in PSPACE if one-step satisfiability checking can be performed in polynomial space. The exact definition of the relevant decision problem is the following:

Definition 34. The *strict one-step satisfiability* problem is to decide whether a one-step N -pair (η, ψ) over V is one-step κ -satisfiable for a given $\kappa : N \rightarrow (V \cup N) \rightarrow 2$, where the input size is defined to be $|\psi|$, and η is represented as a disjunctive set of maximal conjunctive clauses over $V \cup N$ and stored on an input tape (which does not count towards space consumption). The *lax one-step satisfiability* problem is the same decision problem, but with input size $|(\eta, \psi)|$.

We thus have

Theorem 35. *If strict one-step satisfiability is in PSPACE, then the satisfiability problem of \mathcal{L} is in PSPACE.*

In concrete applications, the above condition may be established either directly or with the help of a local small model property:

Definition 36. The *one-step polysize model property (OSPMP)* holds if there is a polynomial p such that every one-step κ -satisfiable pair (η, ψ) , $\kappa : N \rightarrow (V \cup N) \rightarrow 2$, has a one-step κ -model (X, τ, t) such that $|X| \leq p(|\psi| + |N|)$.

Theorem 37. *Under the OSPMP, strict one-step satisfiability is in PSPACE iff lax one-step satisfiability is in PSPACE.*

Example 38. In all hybrid logics of Example 1 (and many others), strict one-step satisfiability is in PSPACE, which reproves the PSPACE upper bounds of Example 22. This is established in each case in essentially the same way as in the purely modal version as carried out in [30]. In most examples, this involves application of Theorem 37, with the crucial step being the proof of the OSPMP (lax one-step satisfiability is typically even in NP [27]), the only essential change w.r.t. the modal case being that states with names in N are retained in the construction of small submodels of given one-step models. This applies moreover to the extension of probabilistic hybrid logic with linear inequalities [17], which is thus newly established to be in PSPACE.

One notable case that requires a direct proof that strict one-step satisfiability is in PSPACE is graded hybrid logic. Although this, too, is largely analogous to the modal case, we briefly sketch the argument for the sake of illustration. Thus, let (η, ψ) be a one-step N -pair over V in graded hybrid logic, and let $\kappa : N \rightarrow (V \cup N) \rightarrow 2$. One-step satisfiability of ψ over given (X, τ) amounts to solvability of a system of integer linear inequalities [27] whose coefficients occur in ψ . Such a system is solvable iff it has a solution whose components are of size at most $p(|\psi|)$, where p is a polynomial [23]. By Lemma 30, a non-deterministic algorithm which traverses a multiset over η , regarded as a set of maximal conjunctive clauses χ , by successively guessing the multiplicity of each χ and adding up the multiplicity for each $a \in V \cup N$ in order to finally check satisfaction of ψ , decides one-step κ -satisfiability of (η, ψ) ; the algorithm clearly uses polynomial space in $|\psi|$. This establishes the new result that *graded hybrid logic is in PSPACE* with numbers coded in binary.

These arguments extend straightforwardly to the hybrid extension of Presburger modal logic [16], which in generalisation of graded modalities features linear inequalities between satisfaction multiplicities of formulas. It is equally straightforward to add multiple agents, or multiple roles in description logic parlance, and role hierarchies, i.e. inclusion axioms between roles. In summary, we obtain the new result that *Presburger hybrid logic with role hierarchies is in PSPACE*. This logic is substantially stronger than e.g. the description logic \mathcal{ALCHOQ} , which features role hierarchies, nominals, and qualified number restrictions but not linear inequalities or satisfaction operators, so that as a corollary we obtain that *concept satisfiability over the empty TBox in \mathcal{ALCHOQ} is in PSPACE*, a tight upper bound. According to the description logic complexity navigator [36], this bound was previously unknown (a PSPACE upper bound for the sublogic \mathcal{ALCCOQ} , which excludes role hierarchies, is proved for unary coding of numbers, and claimed to extend to binary coding, in [4]). Note that Presburger-type logics and probabilistic logics with linear inequalities [17] are presently not amenable to syntactic complexity analysis, e.g. using the sequent calculus method presented above, as no cut-free axiomatisation is known.

6 Conclusions and Related Work

There is quite a large variety of different proof calculi for hybrid logics: one sees fully internalised [31] and labelled calculi [7], as well as natural deduction systems [10].

Apart from the fact that the results of this paper are applicable to a much larger variety of logics, the construction of the sequent calculus introduced in Section 3 is canonical, in the sense that the hybrid sequent rules correspond to a system of cut-free sequent rules for the non-hybrid system, which raises hopes that the same rules might also serve in even more expressive logical systems. The analysis of the sequent calculus is complemented by a semantic analysis which confirms the syntactic PSPACE bounds and yields new ones where the syntactic approach has not reached (yet); the precise relationship between the semantics-based algorithms and proof search remains to be explored in detail. The hybrid proof search algorithms will be integrated into the generic reasoner CoLoSS [11]. Our complexity results allow inferring identical complexity bounds for a coalgebraic modal logic and its hybrid companion from the same set of conditions. In the more restrictive realm of relational semantics, one also has so-called *transfer* results that in particular allow inferring the complexity of a hybrid logic from that of its modal companion [5,12]. Coalgebraic transfer results are the subject of future investigation.

Acknowledgments The authors wish to thank Erwin R. Catesbeiana for detailed explanations of his views on consistent formulas.

References

1. C. Areces, P. Blackburn, and M. Marx. A road-map on complexity for hybrid logics. In J. Flum and M. Rodríguez-Artalejo, eds., *Proc. 8th Annual Conference on Computer Science Logic*, number 1683 in LNCS, pp. 307–321, 1999.
2. C. Areces and B. ten Cate. Hybrid logics. In P. Blackburn, F. Wolter, and J. van Benthem, eds., *Handbook of Modal Logics*. Elsevier, 2006.
3. F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, eds. *The Description Logic Handbook*. Cambridge University Press, 2003.
4. F. Baader, M. Milicic, C. Lutz, U. Sattler, and F. Wolter. Integrating description logics and action formalisms for reasoning about web services. LTCS-Report LTCS-05-02, Dresden University of Technology, 2005. Available under <http://lat.inf.tu-dresden.de/research/reports.html>.
5. N. Bezhanishvili and B. ten Cate. Transfer results for hybrid logic. Part 1: the case without satisfaction operators. *J. Logic Comput.*, 16:177–197, 2006.
6. P. Blackburn and M. Marx. Tableaux for quantified hybrid logic. In *Automated Reasoning with Analytic Tableaux and Related Method*, vol. 2381 of *LNAI*, pp. 38–52. Springer, 2002.
7. T. Bolander and P. Blackburn. Termination for hybrid tableaux. *J. Logic Comput.*, 17:517–554, 2007.
8. T. Bolander and T. Braüner. Tableau-based decision procedures for hybrid logic. *J. Log. Comput.*, 16(6):737–763, 2006.
9. R. Book, T. Long, and A. Selman. Quantitative relativizations of complexity classes. *SIAM J. Computing*, 13:461–487, 1984.
10. T. Braüner. Natural deduction for first-order hybrid logic. *J. Logic, Language and Information*, 14:173–198, 2005.
11. G. Calin, R. Myers, D. Pattinson, and L. Schröder. CoLoSS: The coalgebraic logic satisfiability solver (system description). In *Methods for Modalities, M4M-5*, ENTCS. Elsevier, 2008. To appear.
12. B. ten Cate. *Model theory for extended modal languages*. PhD thesis, University of Amsterdam, 2005. ILLC Dissertation Series DS-2005-01.

13. A. Chandra and L. Stockmeyer. Alternation. *J. ACM*, 28:114–133, 1981.
14. B. Chellas. *Modal Logic*. Cambridge University Press, 1980.
15. G. D’Agostino and A. Visser. Finality regained: A coalgebraic study of Scott-sets and multisets. *Arch. Math. Logic*, 41:267–298, 2002.
16. S. Demri and D. Lugiez. Presburger modal logic is only PSPACE-complete. In *Automated Reasoning, IJCAR 06*, vol. 4130 of *LNAI*, pp. 541–556. Springer, 2006.
17. R. Fagin and J. Y. Halpern. Reasoning about knowledge and probability. *J. ACM*, 41:340–367, 1994.
18. K. Fine. In so many possible worlds. *Notre Dame J. Formal Logic*, 13:516–520, 1972.
19. V. Haarslev, R. Möller, R. van der Straeten, and M. Wessel. Extended query facilities for racer and an application to software-engineering problems. In *Description Logics, DL 2004*, pp. 148–157, 2004.
20. A. Heifetz and P. Mongin. Probabilistic logic for type spaces. *Games and Economic Behavior*, 35:31–53, 2001.
21. K. Larsen and A. Skou. Bisimulation through probabilistic testing. *Inform. Comput.*, 94:1–28, 1991.
22. E. Mares. *Relevant Logic: A Philosophical Interpretation*. Cambridge University Press, 2004.
23. C. Papadimitriou. On the complexity of integer programming. *J. ACM*, 28:765–768, 1981.
24. D. Pattinson. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoret. Comput. Sci.*, 309:177–193, 2003.
25. D. Pattinson and L. Schröder. Admissibility of cut in coalgebraic logics. In *Coalgebraic Methods in Computer Science, CMCS 08*, vol. 203 of *ENTCS*, pp. 221–241. Elsevier, 2008.
26. M. Pauly. A modal logic for coalitional power in games. *J. Logic Comput.*, 12:149–166, 2002.
27. L. Schröder. A finite model construction for coalgebraic modal logic. *J. Log. Algebr. Prog.*, 73:97–110, 2007.
28. L. Schröder and D. Pattinson. PSPACE reasoning for rank-1 modal logics. In *Logic in Computer Science, LICS 06*, pp. 231–240. IEEE, 2006. Full version to appear in *ACM TOCL*.
29. L. Schröder and D. Pattinson. Cut elimination in coalgebraic logics. Technical report, Imperial College London, 2008.
30. L. Schröder and D. Pattinson. Shallow models for non-iterative modal logics. In *Advances in Artificial Intelligence, KI 2008*, vol. 5243 of *LNAI*, pp. 324–331. Springer, 2008.
31. J. Seligman. Internalization: The case of hybrid logics. *J. Logic Comput.*, 11:671–689, 2001.
32. E. Sirin, B. Parsia, B. C. Grau, A. Kalyanpur, and Y. Katz. Pellet: A practical OWL-DL reasoner. *J. Web Semantics*, 2006.
33. A. Troelstra and H. Schwichtenberg. *Basic Proof Theory*. Cambridge University Press, 1996.
34. D. Tsarkov and I. Horrocks. FaCT++ description logic reasoner: System description. In *Automated Reasoning, IJCAR 2006*, vol. 4130 of *LNAI*, pp. 292–297. Springer, 2006.
35. M. Tzakova. Tableau calculi for hybrid logics. In N. Murray, ed., *Proc. Automated Reasoning with Analytic Tableaux and Related Methods (TABLEAUX)*, vol. 1617 of *LNCS*, pp. 278–292. Springer, 1999.
36. E. Zolin. The description logic complexity navigator. <http://www.cs.man.ac.uk/~ezolin/dl>, 2007.