

Nominals for Everyone*

Lutz Schröder

DFKI Bremen and University of Bremen
Lutz.Schroeder@dfki.de

Dirk Pattinson and Clemens Kupke

Department of Computing
Imperial College London
{dirk,ckupke}@doc.ic.ac.uk

Abstract

It has been recognised that the expressivity of ontology languages benefits from the introduction of non-standard modal operators beyond the usual existential restrictions and the number restrictions already featured by many description logics. Such operators serve to support notions such as uncertainty, defaults, agency, obligation, or evidence, which are hard to capture using only the standard operators, and whose semantics often goes beyond relational structures. We work in a unified theory for logics that combine non-standard modal operators and nominals, a feature of established description logics that provides the necessary means for reasoning about individuals; in particular, the logics of this framework allow for internalisation of ABoxes. We reinforce the general framework by proving decidability in EXPTIME of concept satisfiability over general TBoxes; moreover, we discuss example instantiations in various probabilistic logics with nominals.

1 Introduction

Description logics [Baader *et al.*, 2003] are one of the core areas of research in knowledge representation, and semantically underpin the web ontology language OWL [Horrocks *et al.*, 2003] as well as many stand-alone ontologies. A key feature of many description logics is support for *nominals*, which are explicit names designating individual entities to be used *within concepts*, rather than only in a separate collection of assertions about individuals, the ABox. Nominals allow in particular for a direct combination of knowledge about individuals with terminological knowledge.

Another group of features which is often recognised as desirable, but is *not* currently included in standard description logics, is formed by reasoning paradigms which go beyond the standard relational perspective. The latter is the semantic basis e.g. of existential or universal restrictions $\exists R.C /$

$\forall R.C$ along roles asserting that some or all R -successors, respectively, of an individual satisfy a concept C , and of the more general qualified number restrictions $\geq nR.C / \leq nR.C$ which give explicit numerical bounds on the number of R -successors satisfying C . Features not supported by relational models include e.g. reasoning with uncertainty, default implication, coalitional reasoning, or notions of agency. Some effort has recently been invested in designing probabilistic extensions of description logics such as *SHOIN*, some of them supporting a form of defaults (a good overview is found in [Lukasiewicz, 2008]). A common feature of these approaches is that they typically add a new reasoning principle only at the outermost level, e.g. by generalising concept inclusions in the TBox to conditional probabilities.

Here, we propose a framework that allows for a close integration of a wide variety of reasoning principles in combination with nominals and general TBoxes. We obtain this framework by extending the framework of *coalgebraic hybrid logic* recently introduced by Myers *et al.* [2009] with reasoning support for general TBoxes. This general framework supports a wide variety of reasoning principles embodied as modal operators, thus in particular allowing for nested application of the new operators. The semantics of these operators, which include e.g. probabilistic modal operators and other notions of uncertainty, non-monotonic conditional implication, and operators relating to the power of coalitions, often goes far beyond standard relational semantics, being based e.g. on probabilistic structures, selection function models, or game frames. The common umbrella for all these structures is a coalgebra-based semantics [Pattinson, 2003], which we recall in some detail below.

Technically, we prove that under natural assumptions on the axiomatisation of the reasoning principles used in the logic at hand, in fact the same assumptions as used by Schröder and Pattinson to establish a generic PSPACE upper bound for various purely modal logics [2009], concept satisfiability over general TBoxes is in *EXPTIME*, typically a tight upper bound. We achieve this by first reducing the satisfiability problem to the existence of tableaux, and, in a second step, to the existence of winning strategies in parity games.

We conclude with an extended discussion of how our framework may be fruitfully applied in ontological reasoning. We exploit in particular that coalgebraic semantics is *modular* [Schröder and Pattinson, 2007] and hence allows for flex-

*Work of the first author forms part of the DFG project *Generic Algorithms and Complexity Bounds in Coalgebraic Modal Logic* (SCHR 1118/5-1). Work of the second author partially supported by EPSRC grant EP/F031173/1

ibly tailored combinations of reasoning principles and algorithms. We illustrate this point using different combinations of probabilistic and relational semantics in an ontology of the Tudor dynasty.

2 Nominals in Coalgebraic Logic

We recall the generic framework of *coalgebraic hybrid logic* recently introduced by Myers et al. [2009]. It covers a range of logics that feature modal operators interpreted over a wide variety of system types, *nominals* designating individuals within a system, and *satisfaction operators* that permit to assert properties of individuals at any place within a formula, thus in particular allowing for internalisation of ABoxes.

The framework is parametric in both syntax and semantics. The syntax of a given logic is determined by a (*modal*) *similarity type* Λ consisting of modal operators with associated arities, *which we fix 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 operators $\rightarrow, \leftrightarrow, \vee, \top, \perp$. The set of nominals occurring in a formula ϕ is denoted by $N(\phi)$. A formula of the form $@_i \phi$ is called an *@-formula*. For $\Sigma \subseteq \mathcal{F}(\Lambda)$, we put $N(\Sigma) = \bigcup_{\phi \in \Sigma} N(\phi)$ and $@\Sigma = \{\phi \in \Sigma \mid \phi \text{ @-formula}\}$. Semantically, nominals i denote individual points in a model, and an @-formula $@_i \phi$ stipulates that ϕ holds at i .

The parametrisation of the semantics is essentially the standard coalgebraic semantics of modal logics [Pattinson, 2003]. In particular, the type of systems underlying the semantics is determined by the choice of an endofunctor $T : \mathbf{Set} \rightarrow \mathbf{Set}$ on the category of sets, to be thought of informally as a parametrised datatype (formally, T maps sets X to sets TX and maps $X \rightarrow Y$ to maps $TX \rightarrow TY$, compatibly with identities and composition). Then, T -coalgebras play the roles of *frames*. A T -coalgebra is a pair (C, γ) where C is a set of *states* (or *individuals*) and $\gamma : C \rightarrow TC$ is the *transition function*. When γ is clear from the context, we identify a T -coalgebra (C, γ) with its state space C .

Example 2.1. 1. The (covariant) powerset functor \mathcal{P} maps a set X to its powerset $\mathcal{P}(X)$; its coalgebras $C \rightarrow \mathcal{P}(C)$ are in bijection with Kripke frames $(C, R \subseteq C \times C)$.

2. The multiset functor \mathcal{B} maps a set X to the set of *multisets* over X , i.e. maps $X \rightarrow \mathbb{N} \cup \{\infty\}$ assigning multiplicities to elements of X . Its coalgebras are *multigraphs*, a variant of Kripke frames where edges are annotated with positive integer multiplicities [D’Agostino and Visser, 2002].

3. The distribution functor \mathcal{D} maps a set X to the set of finitely supported probability distributions on X ; its coalgebras are Markov chains, also variously referred to as probabilistic type spaces [Heifetz and Mongin, 2001] or probabilistic transition systems.

4. Coalgebras for the functor \mathcal{CF} taking a set X to the set $\mathcal{P}(X) \rightarrow \mathcal{P}(X)$ of *selection functions* over X are pre-

cisely conditional frames [Chellas, 1980], also called selection function models.

5. Coalgebras for the functor \mathcal{G}_n taking a set X to the set $\{(S_1, \dots, S_n, f) \mid S_1, \dots, S_n \text{ nonempty sets (of strategies), } f : (\prod S_i) \rightarrow X\}$ of *n -player strategic games* over X are Pauly’s game frames [2002].

The interpretation of an n -ary modal operator $\heartsuit \in \Lambda$ is given by an *n -ary predicate lifting* $\llbracket \heartsuit \rrbracket$, i.e. a family of maps $\llbracket \heartsuit \rrbracket_X : \mathcal{P}(X)^n \rightarrow \mathcal{P}(TX)$, indexed over all sets X , such that

$$\llbracket \heartsuit \rrbracket_X(h^{-1}[A_1], \dots, h^{-1}[A_n]) = (Th)^{-1}[\llbracket \heartsuit \rrbracket_Y(A_1, \dots, A_n)]$$

for all $h : X \rightarrow Y$, $A_1, \dots, A_n \in \mathcal{P}Y$.

The semantics induced by these parameters, *which we fix throughout*, is a satisfaction relation \models between states $c \in C$ in (*hybrid*) T -models $M = (C, \gamma, \pi)$ and formulas $\phi \in \mathcal{F}(\Lambda)$. Here, M consists of a T -coalgebra (C, γ) and a *hybrid valuation* π , i.e. a map $P \cup N \rightarrow \mathcal{P}(C)$ that assigns singleton sets to all nominals $i \in N$, where we often identify the singleton set $\pi(i)$ with its unique element. Satisfaction is inductively defined by the obvious clauses for the propositional part, and by

$$\begin{aligned} c, M \models x \text{ iff } c \in \pi(x) \quad c, M \models @_i \phi \text{ iff } \pi(i), M \models \phi \\ c, M \models \heartsuit(\phi_1, \dots, \phi_n) \text{ iff } \gamma(c) \in \llbracket \heartsuit \rrbracket_C(\llbracket \phi_1 \rrbracket_M, \dots, \llbracket \phi_n \rrbracket_M) \end{aligned}$$

where $x \in N \cup P$, $i \in N$, $\heartsuit \in \Lambda$ n -ary, and $\llbracket \phi \rrbracket_M = \{c \in C \mid c, M \models \phi\}$. The focus of the present work is on reasoning over so-called general TBoxes: Given a set $\Gamma \subseteq \mathcal{F}(\Lambda)$ of global assumptions, the *TBox*, we say that M is a Γ -model if $c, M \models \phi$ for all $c \in C$ and all $\phi \in \Gamma$. A formula ϕ (a set Φ of formulas) is Γ -satisfiable if there exists a state satisfying ϕ (all formulas in Φ) in some Γ -model. Note that thanks to the satisfaction operator, an ABox, i.e. a set of assertions about individuals, may be encoded either in the formula ϕ itself or in the TBox Γ .

Example 2.2. We recall a few basic examples that use the functors from Example 2.1.

1. The hybrid version of the modal logic K , *hybrid K* for short, has a single unary modal operator \Box , interpreted over the powerset functor \mathcal{P} by $\llbracket \Box \rrbracket_X(A) = \{B \in \mathcal{P}(X) \mid B \subseteq A\}$. This coalgebraic definition of satisfaction translates to the usual semantics of the box operator along the bijection between \mathcal{P} -coalgebras and Kripke frames, inducing the standard semantics of hybrid logic [Areces and ten Cate, 2007]. The description logic \mathcal{ALCCO} is a notational variant of a sublogic of multi-agent hybrid K (captured coalgebraically using multiple copies of the powerset functor).

2. *Graded hybrid logic* has modal operators \diamond_k ‘in more than k successors, it holds that’. It is interpreted over the multiset functor \mathcal{B} 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 [D’Agostino and Visser, 2002]. One can encode the description logic \mathcal{ALCCOQ} (which features *qualified number restrictions* $\geq nR$ and has a relational semantics) into multi-agent graded hybrid logic with multigraph semantics by adding formulas $\neg \diamond_1 i$ for all occurring nominals i to the TBox.

3. *Probabilistic hybrid logic*, the hybrid extension of probabilistic modal logic [Larsen and Skou, 1991; Heifetz and Mongin, 2001], 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 distribution functor \mathcal{D} by putting $\llbracket L_p \rrbracket_X(A) = \{P \in \mathcal{D}(X) \mid PA \geq p\}$.

4. *Hybrid CK*, the hybrid extension of the basic conditional logic *CK*, has a single binary modal operator \Rightarrow , written in infix notation and read e.g. as a non-monotonic default implication. Hybrid *CK* is interpreted over the functor \mathcal{CF} by putting $\llbracket \Rightarrow \rrbracket_X(A, B) = \{f : \mathcal{P}(X) \rightarrow \mathcal{P}(X) \mid f(A) \subseteq B\}$. Other conditional logics with additional axioms, e.g. cautious monotony, are captured similarly. As a simple example, the fact that the national football championship is typically won by team i (an observation that fits a number of countries) is expressed in hybrid conditional logics by the formula $\text{champion} \Rightarrow i$.

5. *Hybrid coalition logic*, the hybrid version of Pauly’s coalition logic [2002], has modal operators $[C]$ ‘the coalition $C \subseteq \{1, \dots, n\}$ of agents may force ...’. These are interpreted by suitable predicate liftings for the functor \mathcal{G}_n [Schröder and Pattinson, 2009]. Given a \mathcal{G}_n -coalgebra (C, γ) , C is the set of states in a strategic game, and nominals therefore encode individual positions.

Our generic complexity result will be based on axiomatisations in a certain format; we require the following notation.

Definition 2.3. The set of boolean combinations over a set V is denoted $\text{Prop}(V)$. A *clause* over a set V is a disjunction of *literals* over V , i.e. elements of $V \cup \{\neg v \mid v \in V\}$. The set of clauses over V is denoted $\text{Cl}(V)$. A *conjunctive normal form (CNF)* of $\phi \in \text{Prop}(V)$ is a subset of $\text{Cl}(V)$ whose disjunction is propositionally equivalent to ϕ . For $\Phi \subseteq \text{Prop}(V)$, $\psi \in \text{Prop}(V)$, we write $\Phi \vdash_{PL} \psi$ (‘ Φ propositionally entails ψ ’) if there exist $\phi_1, \dots, \phi_n \in \Phi$ such that $\phi_1 \wedge \dots \wedge \phi_n \rightarrow \psi$ is a propositional tautology. A valuation $\tau : V \rightarrow \mathcal{P}(X)$ for some set X induces in the obvious way an interpretation $\llbracket \phi \rrbracket_\tau \subseteq X$. Moreover, we put $\Lambda(X) = \{\heartsuit(x_1, \dots, x_n) \mid \heartsuit \in \Lambda \text{ } n\text{-ary}, x_1, \dots, x_n \in X\}$. Given τ as above, one obtains for $\phi \in \text{Prop}(\Lambda(\text{Prop}(V)))$ a *one-step semantics* $\llbracket \phi \rrbracket \subseteq TX$ extending the assignment $\llbracket \heartsuit(\phi_1, \dots, \phi_n) \rrbracket_\tau = \llbracket \heartsuit \rrbracket_X(\llbracket \phi_1 \rrbracket_\tau, \dots, \llbracket \phi_n \rrbracket_\tau)$.

Using these notions, we can now define the crucial prerequisites for the generic reasoning algorithm.

Definition 2.4. A *(one-step) rule* $R = \phi/\psi$ over a set V of propositional variables consists of a premise $\phi \in \text{Prop}(V)$ and a conclusion $\psi \in \text{Cl}(\Lambda(V))$. The rule R is *one-step sound* if whenever $\llbracket \phi \rrbracket_\tau = X$ for a valuation $\tau : V \rightarrow \mathcal{P}(X)$, then $\llbracket \psi \rrbracket_\tau = TX$. A set \mathcal{R} of one-step rules is *strictly one-step complete* if whenever $\llbracket \chi \rrbracket_\tau = TX$ for some $\tau : V \rightarrow \mathcal{P}(X)$ and some $\chi \in \text{Cl}(\Lambda(V))$, then there exist a rule $\phi/\psi \in \mathcal{R}$ and a V -substitution σ such that $\psi\sigma \vdash_{PL} \chi$ and $X, \tau \models \phi\sigma$. Here, \mathcal{R}_C denotes the extension of \mathcal{R} with *congruence rules* $a_1 \leftrightarrow b_1; \dots; a_n \leftrightarrow b_n / \heartsuit(a_1, \dots, a_n) \leftrightarrow \heartsuit(b_1, \dots, b_n)$ for $\heartsuit \in \Lambda$ n -ary.

Strict one-step completeness essentially amounts to absorption of cut by the rule system. Strictly one-step complete rule sets for the logics of Example 2.2 are given in [Schröder

and Pattinson, 2009; Pattinson and Schröder, 2008]. E.g. for hybrid *K*, the set of rules $a_1 \wedge \dots \wedge a_n \rightarrow b / \Box a_1 \wedge \dots \wedge \Box a_n \rightarrow \Box b$ is strictly one-step complete. The axiomatisation of graded and probabilistic logics is more complicated, but still tractable in a sense recalled below. In the following, we assume given a strictly one-step complete set \mathcal{R} .

3 Generic Complexity Bounds

We proceed to develop a decision procedure for global consequence in coalgebraic hybrid logic, i.e. for Γ -satisfiability of formulas given a TBox Γ , by means of a translation of the satisfiability problem into the problem of finding a winning strategy in a parity game. The latter will be played on a game board built from a Γ -closed set Σ of formulas.

Definition 3.1. Let $\Sigma \subseteq \mathcal{F}(\Lambda)$. A (Σ) -*Hintikka set* is a subset of Σ which is maximally consistent w.r.t. propositional reasoning. We say that Σ is closed if Σ is closed under subformulas, negation, and $@_t$ with $t \in \mathbf{N}(\Sigma)$, where we identify $\neg\neg\phi$ with ϕ , $@_t\neg\phi$ with $\neg@_t\phi$, and $@_s@_t\phi$ with $@_t\phi$. We say that Σ is Γ -closed if $\Gamma \subseteq \Sigma$ and Σ is closed. The Γ -closure of a set Δ is the smallest Γ -closed set containing Δ .

Let ϕ be a formula, to be checked for Γ -satisfiability. As ϕ is Γ -satisfiable iff $@_t\phi$ is Γ -satisfiable for a fresh nominal t , we can assume that ϕ is an $@$ -formula. We form the Γ -closure Σ of $\{\phi\}$ (which is of polynomial size in Γ, ϕ). Note that ϕ is Γ -satisfiable iff there exists a Γ -satisfiable $@\Sigma$ -Hintikka set K such that $\phi \in K$; as going through all such K yields an exponential factor and we are aiming for *EXPTIME* decidability, we can focus on deciding Γ -satisfiability of $@\Sigma$ -Hintikka sets. We can then apply the technique of *@-elimination* from [Myers et al., 2009]:

Definition 3.2. 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.

One shows easily that a model satisfies K iff it satisfies the $@$ -eliminated set $\{@_i\rho[K] \mid @_i\rho \in K\}$. Thus, we assume in the following w.l.o.g. that K is *@-eliminated* and hence that the Σ -Hintikka sets $K_i = \{\rho \mid @_i\rho \in K\}$ ($i \in \mathbf{N}(\Sigma)$) are $@$ -free; intuitively, we have thus reduced to checking Γ -satisfiability of an ABox K . Note that the K_i need not be pairwise distinct. If one of the K_i does not contain Γ , then K is immediately rejected as Γ -unsatisfiable.

In the tableau system for global entailment, possible non-termination arises both from the presence of global assumptions, which may propagate indefinitely, as well as from the presence of nominals, which may force loops. The game-theoretic approach that we apply below allows us to deal with infinite paths in tableaux, and eliminates the need to consider blocking conditions. We introduce a notion of tableau graph that captures all possible tableaux, i.e. all possible rule applications at every node, within a single object:

Definition 3.3. If H is a Σ -Hintikka set, $\chi/\psi \in \mathcal{R}$, and σ is a substitution such that $\psi\sigma \in \text{Prop}(\Sigma)$ and $H \vdash_{PL} \neg\psi\sigma$, then $\neg\chi\sigma$ is a *demand* of H . A Γ -tableau graph for K is a graph

whose set of nodes consists of Σ -Hintikka sets and includes the Σ -Hintikka sets K_i , such that

1. for every demand ρ of a node H , there exists an edge $H \rightarrow G$ such that $G \vdash_{PL} \rho$
2. whenever $H \vdash_{PL} i$ for some node H and some $i \in N(\Sigma)$, then $H = K_i$,
3. $H \supseteq \Gamma$ for every node H .

Theorem 3.4. *The set K is Γ -satisfiable iff there exists a Γ -tableau graph for K .*

Sketch. ‘Only if’ is by straightforward extraction of a tableau graph from a Γ -model for K . ‘If’ is by construction of a so-called coherent coalgebra structure ξ on the set of nodes in a tableau graph such that the graph becomes a *supporting Kripke frame*, i.e. for every node H , $\xi(H) \in TY$ where Y is the set of successor nodes of H . Here, ξ is called *coherent* if for all $\heartsuit(\rho_1, \dots, \rho_n) \in \Sigma$ and all nodes H ,

$$\xi(H) \in \llbracket \heartsuit \rrbracket(\hat{\rho}_1, \dots, \hat{\rho}_n) \text{ iff } \heartsuit(\rho_1, \dots, \rho_n) \in H$$

where $\hat{\rho}$ is the set of successor nodes G of H such that $\rho \in G$. Existence of a coherent structure ξ is proved by means of strict one-step completeness, analogously as in [Schröder and Pattinson, 2009] but avoiding induction over the depth of nodes. Coherence then allows the inductive proof of a truth lemma, which entails that the model constructed satisfies both Γ and K . \square

As the nodes of the tableau graph are subsets of Σ , we obtain a small model property for hybrid coalgebraic logic relative to an arbitrary background theory.

Corollary 3.5. *Every Γ -satisfiable formula ϕ is satisfiable in a Γ -model of exponential size in Γ and ϕ .*

Having reduced the satisfiability problem to existence of tableau graphs, we now show that the latter can be further reduced to existence of winning strategies in certain parity games, as follows. The game is played by two players, Abelard (\forall) and Eloise (\exists); \exists tries to prove that K is Γ -satisfiable, while \forall tries to prove the opposite. A move by \forall consists in the choice of rule to be applied, giving rise to a demand, while a move by \exists consists in the choice of a Hintikka set that satisfies the demand. Formally:

Definition 3.6 (Tableau Game). The Γ -tableau game for K is a graph game

$$S = (B_\exists, B_\forall, E)$$

where

- B_\forall , the set of positions owned by \forall , consists of all Σ -Hintikka sets containing Γ (including the K_i) and containing $i \in N(\Sigma)$ only in case $H = K_i$, and an additional initial position *init*.
- B_\exists , the set of positions owned by \exists , consists of pairs (R, σ) , where $R = \chi/\psi$ is a rule in \mathcal{R} and σ is a substitution such that $\psi\sigma \in \text{Prop}(\Sigma)$
- E is the set of permissible moves, where \forall may move from a Σ -Hintikka-set H to a pair $(\chi/\psi, \sigma)$ such that $H \vdash_{PL} \neg\psi\sigma$, and \exists may move from $(\chi/\psi, \sigma)$ to a Σ -Hintikka set H such that $H \vdash_{PL} \neg\chi\sigma$. Additionally, \forall may move to any of the K_i from *init*.

The set of all positions on the game board is $B = B_\exists \cup B_\forall$. (Note that B_\forall is a priori infinite if there are infinitely many rules; we will introduce additional assumptions later that all-reducing to a finite board.)

A *full play* in the tableau game is a finite or infinite sequence of moves (b_0, b_1, b_2, \dots) such that $b_0 = \text{init}$, $(b_i, b_{i+1}) \in E$ for all $i \geq 0$, and – in case the sequence is finite – the last position has no permissible moves. A finite full play is lost by the player who owns the last position (and hence cannot move), and infinite full plays are won by \exists .

Remark 3.7. We note that the tableau game is a parity game where we assign priority 0 to all positions of the game so that – by the parity condition – \exists wins all infinite games [Mazala, 2001].

Definition 3.8. A *history-free* strategy for \exists is a function $f : B_\exists \rightarrow B$ such that $(b, f(b)) \in E$ for all $b \in B_\exists$. We say that f is a winning strategy for \exists if \exists wins all full plays that conform with f in the obvious sense.

Lemma 3.9. *Eloise has a history-free winning strategy in the Γ -tableau game for K iff there exists a Γ -tableau graph for K .*

Proof Sketch. ‘If’ is clear. ‘Only if’: construct the Γ -tableau graph starting from the initial set of nodes $\{K_i \mid i \in N(\Sigma)\}$ and successively introducing additional nodes and edges according to the strategy of \exists for every possible move of \forall , i.e. for all arising demands. \square

We now show that the existence of a winning strategy for \exists in the tableau game can be decided in exponential time, subject to a mild condition on the rule sets that is satisfied in all our examples. We require the modal tableau rules to be *tractable* in a similar sense as in [Schröder and Pattinson, 2009]; the main condition here is that one may restrict to rule sets with at most polynomial-size codes (regarding the remaining conditions, we can be slightly more generous in the context of *EXPTIME* bounds relevant here).

Definition 3.10. The set \mathcal{R} of modal rules is *EXPTIME-tractable* if there exists a coding of the rules such that, up to propositional equivalence, all demands of a Hintikka set can be generated by rules with codes of polynomially bounded size, and such that validity of codes, matching of rule codes for $\chi/\psi \in \mathcal{R}$ to Hintikka sets H (in the sense of finding σ such that $H \vdash_{PL} \neg\psi\sigma$), and membership of clauses in a CNF of a rule premise are all decidable in *EXPTIME*.

Lemma 3.11. *If \mathcal{R} is EXPTIME-tractable, then it can be decided in EXPTIME whether \exists has a winning strategy in the Γ -tableau game for K .*

Proof Sketch. Given that the rule set is tractable, we may replace the positions B_\exists owned by \exists by codes of polynomial size in Γ, K . This leads to a game board whose size n is at most exponential in Γ, K . As we have a parity game with only one priority, it takes at most $\mathcal{O}(n^4) * k$ steps to determine whether \exists has a winning strategy [Klauck, 2001], where k is such that one can decide in at most k steps whether $(b, b') \in E$. Tractability of the rule set guarantees that k is at most exponential, so that we obtain overall complexity *EXPTIME*. \square

Corollary 3.12. *If \mathcal{R} is EXPTIME-tractable, then Γ -satisfiability of formulas ϕ over general TBoxes Γ is decidable in EXPTIME.*

The above corollary yields decidability in EXPTIME of reasoning over general TBoxes for all logics mentioned in Example 2.2. In particular, this reproves the known tight upper bound for hybrid K (which follows from an EXPTIME upper bound for the graded μ -calculus [Arecas and ten Cate, 2007]), as well as for the description logic \mathcal{ALCOQ} (and, with minor modifications, \mathcal{ALCHOQ}) [Tobies, 2000], however by embedding the latter into a more expressive logic that internalises ABoxes by means of satisfaction operators. The use of games in the context of TBox reasoning appears to be new. We emphasise moreover that the algorithm essentially analyses a tableau; already in the case without nominals, the ad-hoc analysis of tableaux for TBox reasoning in the basic description logic \mathcal{ALC} has proved to be quite complex [Donini and Massacci, 2000]. The (tight) upper bounds for TBox reasoning in probabilistic hybrid logic, conditional hybrid logic, and hybrid coalition logic appear to be new. We discuss some examples of this type in more detail below.

4 Two Views on Probabilistic Successors

We discuss two applications that highlight the generality of our results. In both examples, we combine classical relational successors and uncertainty, but in two different ways. Both are phrased in terms of descendancy, where states in a model represent persons.

4.1 Probabilistic Successors

We imagine a situation where we only have probabilistic knowledge about the offspring of a certain person. Suppose for instance that the probability that c (Catherine Carey) is a child of h (Henry VIII) is known to be at least 0.8, and similarly we know that h' (Henry Carey) is a child of h with likelihood 0.6. To model this situation, we consider a structure of type $C \rightarrow \mathcal{DP}(C)$ where $\mathcal{P}(X)$ is the powerset of a set X and $\mathcal{D}(X) = \{\mu : X \rightarrow [0, 1] \mid \text{supp}(\mu) \text{ is finite, } \sum_{x \in X} \mu(x) = 1\}$ is the set of finitely supported probability distributions over X as in Example 2.2. In other words, given a person $x \in C$, an application of the structure map yields a probability distribution over sets (!) of persons. If this distribution assigns probability p to a set $C' \subseteq C$, we interpret this as the fact that the probability that C' are (precisely) the children of x equals p . (Note that this model applies primarily when x is male.)

This situation can be syntactically described using modal operators of the form $L_p \diamond$, where $L_p \diamond \phi$ reads ‘the probability that there exists a successor that satisfies ϕ is at least p ’, together with the companion modality $L_p \square$ expressing the same statement relative to all successors.

Assume we know that the probability that every king has at least one illegitimate child is at least 0.8. This is expressed using the global assumption

$$\text{king} \rightarrow L_{0.8} \diamond \text{illegitimate},$$

while the above assertions about Catherine and Henry Carey take the form

$$@_h L_{0.8} \diamond c \quad \text{and} \quad @_h L_{0.6} \diamond h'.$$

Moreover, we know that Henry is a king, and m (Mary) is certainly a child of Henry, and either c or h' is illegitimate, whereas e is legitimate, which we express by

$$\begin{aligned} & @_h \text{king} \quad @_h (1)m \quad @_e \neg \text{illegitimate} \\ & @_c \text{illegitimate} \vee @_h \text{illegitimate}. \end{aligned}$$

Now consider the concept

$$\text{king} \wedge L_{0.9} \square \neg \text{illegitimate}$$

asserting that all children of a king are legitimate with probability at least 0.9. This concept is satisfiable, but h is not an instance of c_2 (as e is a legitimate child of h).

Similarly, the ABox

$$@_h L_1 \square (\text{illegitimate} \rightarrow (c \vee h')) \quad @_c \neg \text{illegitimate}$$

formalising that c is legitimate and c and h' are the only possible illegitimate children of h is satisfiable (and the global assumption forces that h' be a child of h with likelihood at least 0.8) but it becomes unsatisfiable if we stipulate for example that h' is a child of h with probability at most 0.7. The proof rules that govern this situation are a straightforward combination of the rules discussed in [Schröder and Pattinson, 2009], which immediately yields tractability of the ensuing combined rule set. As a consequence, we have that global consequence for the logic of probabilistic successors is decidable in EXPTIME.

4.2 Probabilistic Identities

Now let us suppose that someone internal to Henry’s court has observed that none of the children of c' (Catherine of Aragon) had really died, but they were rather removed from court, and we are only left with probabilistic knowledge concerning their identities.

To model this situation, we need to consider a different combination of relational successors and probability distributions. Our knowledge base is modelled by structures of the form $C \rightarrow \mathcal{P}(\mathcal{D}(C))$ where \mathcal{P} and \mathcal{D} are as above. The main difference is now that, from each state of the model, we can observe a set of relational successors (corresponding to the person’s offspring), but each successor carries a probability distribution over the model that expresses uncertainty concerning the successor’s identity.

Syntactically, this leads to modal operators of the form $\diamond L_p$ asserting that there exists a (relational) successor that satisfies a given formula with probability at least p . However, we use a slightly richer set of modal operators, where we can interpose propositional connectives between the relational and probabilistic operators. Formally, this leads to a two-sorted logical language where one sort describes relational successors and the second models quantitative uncertainty; tractability of such combinations is established in [Schröder and Pattinson, 2007], so that we obtain decidability in EXPTIME of global consequence in this situation (formally, the logic of the previous section arises as a similar combination, the second sort corresponding to possible worlds).

Suppose that c' had one female child, whose identity is known to be m (Queen Mary) with certainty, and she also

had a male child, believed to be a with likelihood at least 0.2 and b with likelihood at least 0.8. We take it for granted that a queen's offspring is always legitimate child, leading to the global assumption

$$\text{queen} \rightarrow \Box \neg L_1 \text{illegitimate}$$

where we have made use of the ability to apply propositional connectives to subformulas of the same (here: probabilistic) type. This assumption expresses that at least one of the possible candidates (with non-zero probability) for any given child of a queen, namely the actual child, must be legitimate.

In addition, we have the ABox

$$\begin{aligned} @_{c'} \text{queen} \quad @_{c'} \diamond (L_{0.2} a \wedge L_{0.8} b) \quad @_{c'} \diamond L_1 m \\ @_m \text{female} \quad @_a \neg \text{female} \quad @_b \neg \text{female} \end{aligned}$$

that formalises our assumptions concerning c' 's offspring. We may now ask whether it is possible that both a and b are illegitimate, i.e.

$$@_a \text{illegitimate} \wedge @_b \text{illegitimate}.$$

This formula is not satisfiable as it would violate the global assumption. In contrast, the statement that a queen has at least one child who will be queen with likelihood at least 0.7, i.e. the formula

$$\text{queen} \rightarrow \diamond L_{0.7} \text{queen}$$

is consistent with our (hypothetical) knowledge – we may e.g. consider models satisfying $@_m \text{queen}$.

5 Conclusion

We have extended the algorithmic framework that surrounds coalgebraic hybrid logic [Myers *et al.*, 2009] to deal with global logical consequence. While hybrid constructs allow us to make assertions about individuals, global consequence allows us to algorithmically decide satisfiability of formulas relative to a global background theory. In description logic terms, we internalise the ABox and provide support for *concept satisfiability* and *instance checking* relative to a general TBox. The prime achievement of this work is its generality: suitable instantiations of the general coalgebraic framework yield EXPTIME bounds for a large number of modal and description logics that go far beyond Kripke semantics. In particular, we have established

- a small model property for coalgebraic hybrid logic over general TBoxes (Corollary 3.5), and
- EXPTIME complexity of the global consequence problem in coalgebraic hybrid logic (Corollary 3.12)

for the class of all logics that can be formalised in the coalgebraic framework, which includes various conditional logics, coalition logic, and logics for uncertainty. Further extensions of the framework and the analysis of further reasoning tasks are the subject of ongoing investigation.

References

[Arecas and ten Cate, 2007] Carlos Arecas and Balder ten Cate. Hybrid logics. In P. Blackburn, J. van Benthem, and F. Wolter, editors, *Handbook of Modal Logic*. Elsevier, 2007.

- [Baader *et al.*, 2003] Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook*. Cambridge University Press, 2003.
- [Chellas, 1980] B. Chellas. *Modal Logic*. Cambridge University Press, 1980.
- [D'Agostino and Visser, 2002] G. D'Agostino and A. Visser. Finality regained: A coalgebraic study of Scott-sets and multisets. *Arch. Math. Logic*, 41:267–298, 2002.
- [Donini and Massacci, 2000] Francesco M. Donini and Fabio Massacci. EXPTIME tableaux for ALC. *Artif. Intell.*, 124:87–138, 2000.
- [Heifetz and Mongin, 2001] Aviad Heifetz and Philippe Mongin. Probabilistic logic for type spaces. *Games and Economic Behavior*, 35:31–53, 2001.
- [Horrocks *et al.*, 2003] Ian Horrocks, Peter F. Patel-Schneider, and Frank van Harmelen. From SHIQ and RDF to OWL: the making of a web ontology language. *J. Web Sem.*, 1:7–26, 2003.
- [Klauck, 2001] Hartmut Klauck. Algorithms for parity games. In *Automata, Logics, and Infinite Games*, volume 2500 of LNCS, pages 107–129. Springer, 2001.
- [Larsen and Skou, 1991] K. Larsen and A. Skou. Bisimulation through probabilistic testing. *Inform. Comput.*, 94:1–28, 1991.
- [Lukasiewicz, 2008] Thomas Lukasiewicz. Expressive probabilistic description logics. *Artif. Intell.*, 172:852–883, 2008.
- [Mazala, 2001] René Mazala. Infinite games. In *Automata, Logics, and Infinite Games*, volume 2500 of LNCS, pages 23–42. Springer, 2001.
- [Myers *et al.*, 2009] Rob Myers, Dirk Pattinson, and Lutz Schröder. Coalgebraic hybrid logic. In *Foundations of Software Science and Computation Structures, FOSSACS 2009*, LNCS. Springer, 2009. To appear.
- [Pattinson and Schröder, 2008] Dirk Pattinson and Lutz Schröder. Admissibility of cut in coalgebraic logics. In *Coalgebraic Methods in Computer Science, CMCS 08*, volume 203 of ENTCS, pages 221–241. Elsevier, 2008.
- [Pattinson, 2003] Dirk Pattinson. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoret. Comput. Sci.*, 309:177–193, 2003.
- [Pauly, 2002] Marc Pauly. A modal logic for coalitional power in games. *J. Logic Comput.*, 12:149–166, 2002.
- [Schröder and Pattinson, 2007] Lutz Schröder and Dirk Pattinson. Modular algorithms for heterogeneous modal logics. In *Automata, Languages and Programming, ICALP 07*, volume 4596 of LNCS, pages 459–471. Springer, 2007.
- [Schröder and Pattinson, 2009] Lutz Schröder and Dirk Pattinson. PSPACE bounds for rank-1 modal logics. *ACM Trans. Comput. Log.*, 10, 2009. To appear.
- [Tobies, 2000] Stephan Tobies. The complexity of reasoning with cardinality restrictions and nominals in expressive description logics. *J. Artif. Intell. Res.*, 12:199–217, 2000.