Proceedings, Eleventh International Conference on Principles of Knowledge Representation and Reasoning (2008)

How Many Toes Do I Have?
Parthood and Number Restrictions in Description Logics

Lutz Schröder
DFKI-Lab Bremen and
Department of Computer Science, University of Bremen

Dirk Pattinson
Department of Computing
Imperial College London

Abstract

The modelling of parthood relations in description logics via transitive roles often leads to undecidability when combined with number restrictions and role hierarchies. Here, we introduce the description logic $\mathcal{PHQ}$ that explicitly supports reasoning about parthood in the presence of qualified number restrictions. Our main results are completeness and decidability in NEXPTIME. Conceptually, we argue that $\mathcal{PHQ}$ provides a better semantic fit for many applications: more often than not, parthoods occurring e.g. in biomedical ontologies are expected to be tree-like. In such cases, $\mathcal{PHQ}$ supports stronger inferences than standard description logics. Technically this is achieved by explicitly excluding the merging of descendants, which, at the same time, eliminates the prime source of undecidability. We work in the general setting of *coalgebraic modal logic*, a generic semantic framework for not-necessarily-normal modal logics. This added generality allows the re-use of many of our results for other logics of sometimes quite different flavour.

Introduction

Expressive description logics (Baader et al. 2003) are receiving increasing attention as the formalism underlying ontologies e.g. in biomedical applications such GALEN (Rector and Horrocks 1997) and as the foundation of ontology languages for the semantic web such as OWL (Horrocks, Patel-Schneider, and van Harmelen 2003). The goal in the design of these logics is to build a high level of expressivity while at the same time keeping the resulting logics decidable, preferably in low complexity classes. Here, classes up to at least $NEXPTIME$ still count as 'low', as experience has shown that many logics of this and even higher complexity can still be handled effectively by highly optimised reasoners.

Two features of particular interest are *qualified number restrictions*, occurring in statements such as 'humans are animals with two legs', and *transitive roles* that are primarily motivated by the need to model parthood without committing to a particular level of decomposition — briefly, one often wants parts of parts to be also parts of the whole (Sattler 1996). Unfortunately, these features interact badly, and in particular may lead to undecidability in combination with certain role hierarchies (Horrocks, Sattler, and Tobies 1999). Recent results (Kazakov, Sattler, and Zolin 2007) show that decidability can be salvaged for some types of role hierarchies; however, no complexity bound has so far been proved for this case.

Here, we introduce a description logic $\mathcal{PHQ}$ that directly expresses parthood and allows both role hierarchies and unlimited use of number restrictions, while keeping decidability in $NEXPTIME$. We achieve this by requiring parthood relations to be treelike. This restriction not only eliminates the root of the above undecidability results, but also provides a better fit for many applications where transitive roles are used to model parthood. Consequently, $\mathcal{PHQ}$ admits stronger reasoning principles compared to the standard approach. As a simple example, consider the terminological axiom

$$\mathsf{human} \sqsubseteq {=}\,2\,\mathsf{hasPart}.\,(\mathsf{foot} \sqcap {=}\,5\,\mathsf{hasPart}.\,\mathsf{toe}).$$

Unlike standard description logics that use transitivity to model parthood, $\mathcal{PHQ}$ allows concluding that humans have at least 10 toes, since by treelikeness of $\mathsf{hasPart}$, the sets of toes on both feet are disjoint.

Technically, we obtain the complexity bound for $\mathcal{PHQ}$ as an instance of a new generic complexity bound in *coalgebraic modal logic*, a semantic framework that unifies many structurally different (typically non-normal) modal logics; we give a self-contained introduction to the essential concepts needed. For coalgebraic techniques to be applicable to $\mathcal{PHQ}$, we substantially extend previously known generic complexity bounds (Schröder 2007), which were limited to *rank-1* modal logics, defined by axioms that do not nest modal operators; our generic result improves on this in being applicable to modal logics that employ arbitrary frame conditions, which may nest modal operators. Technically, this is achieved by extending the finite model construction of (Schröder 2007) for a given consistent formula to a construction of pre-models that satisfy enough instances of the frame conditions to support a closure process which guarantees full satisfaction of the frame conditions while preserving validity of the given formula.

The more abstract coalgebraic view both allows for a more high-level treatment of the logic $\mathcal{PHQ}$ itself and provides the foundation upon which corresponding results for other modal logics can be established. Coalgebraic modal logic has proved to be an extremely useful level of abstraction which covers a wide variety of modal logics including e.g. probabilistic and graded modal logics, conditional logics (Chellas 1980), or coalition logic (Pauly 2002), while at the same time allowing for a generic complexity theory up to and including tight PSPACE upper bounds (Schröder and Pattinson 2006). We believe that the genericity of our results is of particular relevance in knowledge representation, where new modal logics are being proposed at regular intervals. In passing, we substantiate the general applicability of our techniques with examples in standard modal logic such as $S4$ and in conditional logic.

Transitive Roles and Number Restrictions

We briefly recall some aspects of the syntax and semantics of expressive description logics, largely following (Kazakov, Sattler, and Zolin 2007). The family of description logics $\mathcal{ALCQ}(\mathcal{R})$ is parametrised over an *RBox* $\mathcal{R}$ which consists of a set of *roles* with ordering $\sqsubseteq$, the *role hierarchy*, and a distinguished subset Tr of *transitive roles*. The set of *concepts* C is then given by the grammar

$$C ::= A \mid \neg C \mid C_1 \sqcap C_2 \mid \geq n\, R.\, C,$$

where $R \in \mathcal{R}$, $n \in \mathbb{N}$, and A ranges over over a set CN of *concept names*. Boolean operators $\sqcup$, $\bot$, $\top$ and additional modalities $\leq n\, R$, $= n\, R$, $\exists R$, $\forall R$ are defined by abbreviation in the standard way (e.g. $\exists R.\, C \equiv \geq 1\, R.\, C$). A *TBox* is a finite set of concept inclusions $C \sqsubseteq D$. Formulas $\geq n\, R.\, C$ are called *(qualified) number restrictions*; they correspond (omitting the role name R) to formulas of the form $\Diamond_{n-1} C$ in *graded modal logic* (Fine 1972).

Models for $\mathcal{ALCQ}(\mathcal{R})$ are Kripke models $\mathcal{I}$ with set of states (*domain*) $\Delta^{\mathcal{I}}$, interpretations $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ of the concept names, and with a binary relation $R^{\mathcal{I}}$ for each role R in $\mathcal{R}$, such that $R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$ whenever $R \sqsubseteq S$ in $\mathcal{R}$ and $R^{\mathcal{I}}$ is transitive whenever R is a transitive role. Concepts C are interpreted as subsets $C^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, with the obvious clauses for boolean operators and

$$(\geq n\, R.\, C)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \#\{y \in C^{\mathcal{I}} \mid xR^{\mathcal{I}}y\} \geq n\},$$

where $\#$ denotes cardinality. The model $\mathcal{I}$ satisfies a concept inclusion $C \sqsubseteq D$ if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, and a TBox $\mathcal{T}$ if it satisfies all concept inclusions in $\mathcal{T}$. A concept C is *satisfiable* w.r.t. $\mathcal{T}$ if there exists an $\mathcal{ALCQ}(\mathcal{R})$-model $\mathcal{I}$ satisfying $\mathcal{T}$ such that $C^{\mathcal{I}} \neq \emptyset$.

It is shown in (Horrocks, Sattler, and Tobies 1999; Kazakov, Sattler, and Zolin 2007) that for many RBoxes $\mathcal{R}$, concept satisfiability is undecidable. The proofs are via reduction from tiling, the crucial point being to force relations to form $\mathbb{N} \times \mathbb{N}$ grids. The core argument of (Horrocks, Sattler, and Tobies 1999) is roughly as follows. Suppose the role hierarchy includes $X \sqsubseteq S$ and $Y \sqsubseteq S$, with S transitive. Now let some state satisfy $A \sqcap \exists X.\, (B \sqcap \exists Y.\, D) \sqcap \exists Y.\, (C \sqcap \exists X.\, D)$, where A, B, C, D are mutually exclusive. By imposing $\leq 3\, S.\, \top$, one can then enforce that the model locally has the shape

$$\begin{array}{ccc} C & \xrightarrow{X} & D \\ {\scriptstyle Y}\uparrow & & \uparrow{\scriptstyle Y} \\ A & \xrightarrow[X]{} & B \end{array}$$

thus forming a cell in the grid. To sum up, undecidability is caused by the phenomenon that branches of the model (in this case, the top left and the bottom right paths of the above diagram) can be merged. Due to these problems, typical description logics such as $\mathcal{SHQ}$ disallow number restrictions on transitive roles and their superroles. Kazakov, Sattler, and Zolin (2007) prove that $\mathcal{ALCQ}(\mathcal{R})$ *is* decidable if the transitive roles are totally ordered in each connected component of $\mathcal{R}$, but do not give a complexity bound. We complement the results of *loc. cit.* by showing that the restrictions on role hierarchies can be lifted in the logic $\mathcal{PHQ}$ where transitive roles are replaced by descendancy relations in a treelike structure.

$\mathcal{PHQ}$: A Logic for Parthood

We now discuss the design of a description logic $\mathcal{PHQ}$ for parthood that imposes a treelike structure on parthood relations, thus simultaneously avoiding the undecidability issues discussed above and providing for stronger inferences in typical applications.

As indicated above, the main motivation for the introduction of transitive roles in description logics are parthood relations, where one often does not wish to distinguish between different levels of decomposition. A typical approach is to consider for each type of parthood a role R for immediate parthood and a transitive superrole $R^{\oplus}$ (not necessarily the smallest such), the *transitive orbit*, of R to approximate general parthood across several inheritance steps (Sattler 1996). We follow a similar approach syntactically, but make different choices in the semantics.

Thus, we assume given a *role hierarchy*, i.e. a set $\mathcal{R}$ of *simple* roles with ordering $\sqsubseteq$. Then, a *role* is either a simple role or a *descendant role* R^D, where $R \in \mathcal{R}$. We put $R^D \sqsubseteq S^D$ whenever $R \sqsubseteq S$. The syntax for concepts is then exactly as above, where we explicitly allow number restrictions on descendant roles.

Models for $\mathcal{PHQ}$ are Kripke models $\mathcal{I}$ with data $\Delta^{\mathcal{I}}$, $A^{\mathcal{I}}$, $R^{\mathcal{I}}$ (for all roles R) as above, such that $R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$ whenever $R \sqsubseteq S$ for roles R, S. We refer to $R^{\mathcal{I}}$-successors as R-*children* and to $(R^D)^{\mathcal{I}}$)-successors as R-*descendants*. We require that for each $R \in \mathcal{R}$,

$$(R^D)^{\mathcal{I}} = R^{\mathcal{I}} \cup R^{\mathcal{I}}; (R^D)^{\mathcal{I}}, \qquad (1)$$

where ';' denotes diagrammatic composition of relations, and that $R^{\mathcal{I}}$-siblings never have common $(R^D)^{\mathcal{I}}$-successors, i.e. that $xR^{\mathcal{I}}y(R^D)^{\mathcal{I}}w$ and $xR^{\mathcal{I}}z(R^D)^{\mathcal{I}}w$ imply $y = z$. These conditions imply that $(R^{\mathcal{I}})^+ \subseteq (R^D)^{\mathcal{I}}$, and consequently that $R^{\mathcal{I}}$ is tree-like in the sense that $R^{\mathcal{I}}$-siblings never have common $(R^{\mathcal{I}})^+$-successors ($R^{\mathcal{I}}$ may have loops; however, different $R^{\mathcal{I}}$-loops cannot intersect). As shown below, this form of treelikeness suffices to guarantee decidability and indeed low complexity. Note that

TBoxes can be internalised in $\mathcal{PHQ}$-concepts using the descendancy role of a largest role in the hierarchy; we therefore do not discuss TBox reasoning explicitly below.

Descendant roles R^D replace transitive orbits $R^\oplus$ as approximations of the transitive hull R^+ (which, although technically in the same complexity class as the transitive orbit, is computationally harder in practice due to the presence of induction axioms (Horrocks and Gough 1997)). The latter may be defined as a superrole of R satisfying either (1) or transitivity, and an appropriate induction principle. In using either $R^\oplus$ or R^D instead of R^+, one gives up induction, i.e. one may have 'infinitely distant' descendants. In the case of $R^\oplus$, one additionally loses (1) (retaining only the inclusion $\supseteq$), while in using R^D, one gives up transitivity but retains (1), including unlike in the case of $R^\oplus$ the property that every R-descendant is either an R-child or an R-descendant of an R-child. We believe that the latter property may be more relevant in applications than transitivity, as it allows for upper bounds on the number of descendants. It should be noted here that (1) implies that descendants of children, grandchildren etc. are again descendants; the only type of inference that is lost by giving up full transitivity is a pathological one, namely that descendants of infinitely distant descendants are again descendants.

In comparison to the use of transitive superroles to model parthood relations, the fact that siblings in $\mathcal{PHQ}$-models never have common descendants allows for stronger inferences and, in many application contexts, a more natural modelling of parthood relations. Consider e.g. a concept A of the form

$$\geq n R. \geq k R^D. C.$$

Every $\mathcal{PHQ}$-model will satisfy the concept subsumption

$$A \sqsubseteq \geq nk\, R^D. C. \quad (2)$$

Consequently, taking up the example from the introduction, $\mathcal{PHQ}$ indeed allows deducing that every mammal with two feet, each of which has five toes, has at least ten toes in total. Contrastingly, when we replace R^D by $R^\oplus$ in A, i.e. regard A as an $\mathcal{ALCQ}(\mathcal{R})$-formula, where $\mathcal{R}$ consists of $R \sqsubseteq R^\oplus$ with $R^\oplus$ transitive, then (2) with $n = 2$, $k = 1$ fails e.g. at the root node of an $\mathcal{ALCQ}(\mathcal{R})$-model of the shape

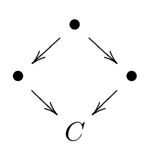

.

Apart from the fact that modelling with $\mathcal{PHQ}$ is often more natural, $\mathcal{PHQ}$ also enjoys better decidability properties compared to the use of transitive roles for modelling parthood. We show below that $\mathcal{PHQ}$ is decidable in $NEXPTIME$ for every RBox, including e.g. $\mathcal{R} = \{R_1 \sqsubseteq S, R_2 \sqsubseteq S\}$. In $\mathcal{ALCQ}$, this would be modelled by an RBox $\mathcal{R}^\oplus$ that adds for each role $R \in \mathcal{R}$ a transitive superrole $R^\oplus$, with $R_1^\oplus \sqsubseteq S^\oplus, R_2^\oplus \sqsubseteq S^\oplus$. By the results of (Kazakov, Sattler, and Zolin 2007), the arising description logic $\mathcal{ALCQ}(\mathcal{R}^\oplus)$ is undecidable.

Remark 1. We do not, of course, claim that treelikeness is a universal feature of parthood in general — e.g. it fails for some corner cases in anatomy (such as the interventricular septum, which is treated as a layer of both the right and the left ventricle in GALEN), and will typically fail for economic parthood relations such as company ownership. However, in many application areas, e.g. engineering, treelikeness is a reasonable assumption, and its improved computational properties would seem to justify forcing treelikeness (by attributing shared components to one parent alone) in cases such as anatomy where failures of treelikeness are the exception rather than the rule. Some anatomical parthood relations such as hierarchical subdivision are standardly regarded as treelike (Bittner and Donnelly 2007).

Remark 2. It is possible (but not always useful due to the mentioned undecidability issues) to encode $\mathcal{PHQ}$ into $\mathcal{ALCQ}(\mathcal{R})$, i.e. to emulate treelikeness in $\mathcal{ALCQ}(\mathcal{R})$ by inserting additional disjointness predicates, e.g. by explicitly distinguishing left and right toes. However, this becomes impractical as soon as either the width or the depth of tree structures increases — e.g. replace, mutatis mutandis, the human by a dog, a spider, or a centipede in the running example — and indeed may formally lead to doubly exponential blowup (with one degree of exponentiality due nesting of number restrictions and the second due to binary coding of numbers). Moreover, natural disjointness predicates are not always available, e.g. in the case of motor valves in a car.

Remark 3. An obvious variant $\mathcal{PHQ}_+$ of $\mathcal{PHQ}$ features exact transitive closure R^+ instead of loosely determined descendancy R^D. Using essentially the same encoding into the monadic second order logic of trees as in (Fattorosi-Barnaba and de Caro 1985), one easily proves that $\mathcal{PHQ}_+$ is decidable. The complexity of $\mathcal{PHQ}_+$, however, remains unclear. Existing automata-theoretic methods such as encodings into graded alternating tree automata, which have been used to prove e.g. an $EXPTIME$ upper bound for the graded μ-calculus (Kupferman, Sattler, and Vardi 2002), do not seem to be immediately applicable here: a graded tree automaton checks satisfaction of number restrictions on immediate successors by sending appropriate numbers of copies of itself to children of the current node. This method of counting does not seem to be suitable for counting descendants across several layers of decomposition, which instead appears to require maintaining actual counters in the state of the automaton. The latter, however, would lead to an exponential blowup of the state space of the automaton (recall that numbers are coded in binary) and hence at best yield a $2EXPTIME$ upper bound for $\mathcal{PHQ}_+$, as the emptyness problem of graded tree automata is $EXPTIME$-complete w.r.t. the size of the state space. It should also be noted that unlike in the case without number restrictions, $\mathcal{PHQ}_+$ cannot just be coded into the graded μ-calculus: while e.g. $\exists R^+. \phi$ is equivalent to the formula $\mu X. \exists R. (\phi \vee X)$ of the (standard) μ-calculus, no obvious analogue holds for $\geq n\, R^+. \phi$. Indeed, the axiomatisation of $\mathcal{PHQ}$ given below suggests that a fixed point definition of $\geq n\, R^+. \phi$ needs to mention formulas of the form $\geq k\, R^+. \phi$ for all $k \leq n$ and hence requires stronger expressive means than provided by the graded μ-calculus (besides being of exponential size).

We do conjecture that $\mathcal{PHQ}_+$ is of the same complexity as $\mathcal{PHQ}$, i.e. at most $NEXPTIME$, but as indicated above we believe that $\mathcal{PHQ}$ is nevertheless practically more tractable than $\mathcal{PHQ}_+$ due to the absence of induction. While $\mathcal{PHQ}_+$ is, of course, preferable from a modeling perspective, note that the semantic gap between $\mathcal{PHQ}$ and $\mathcal{PHQ}_+$ is rather smaller than the one between transitive orbits and transitive closure: while the transitive orbit can introduce unreachable descendants in a rather arbitrary manner as long as transitivity is respected, $\mathcal{PHQ}$ differs from $\mathcal{PHQ}_+$ only in allowing infinitely distant descendants. E.g. making for a moment the assumption that the whole is always larger than its parts, we could use the relation 'larger than' as a transitive orbit of parthood; contrastingly, 'larger than' generally cannot be used as a descendancy relation for parthood in $\mathcal{PHQ}$, as this would require that whenever some object A is larger than an object B, then either B is a part of A or A has a part which is larger than B.

Coalgebraic Modal Logic

For the purpose of the technical development, we view $\mathcal{PHQ}$ as a modal logic, equipped with a *coalgebraic semantics*; below, we give a self-contained account of the essential concepts involved. This affords a more high-level view of the logic $\mathcal{PHQ}$ itself and moreover paves the way to corresponding results for other modal logics. Indeed, as we shall demonstrate, we obtain completeness and decidability also for the modal logic $S4$ and the conditional logic $CK+CMi$ as a by-product of the coalgebraic treatment.

In this approach, the type of structure underlying the semantics is embodied in the choice of a *(set) functor*, i.e. an operation T that maps sets X to sets TX and functions $f : X \to Y$ to functions $Tf : TX \to TY$, preserving identities and composition. The role of models is then played by T-*coalgebras*, i.e. pairs (X, ξ) where X is a set of *states* and $\xi : X \to TX$ is the transition function; thinking of TX informally as a parametrised datatype over X, we regard ξ as associating to each state x a structured collection $\xi(x)$ of successor states and observations. Choosing e.g. $TX = \mathcal{P}(X)$ we obtain that T-coalgebras are Kripke frames (since they associate to each state a set of successor states), whereas $TX = \mathcal{B}(X)$, where $\mathcal{B}(X)$ is the set of multisets over X (explained in more detail further below), gives the multigraph models of (D'Agostino and Visser 2002). In this setting, establishing decidability or completeness of particular logics then amounts to checking coherence conditions for the semantics (defined by the functor T) and the logical axiomatisation.

To define the syntax, we fix a *(modal) similarity type* Λ, i.e. a set of modal operators with associated arities, and a set V of propositional variables that play the role of concept names. For the logic $\mathcal{PHQ}$, the similarity type will consist of the operators $\geq nR$ for roles R and nonnegative integers $n \in \mathbb{N}$. The set $\mathcal{F}(\Lambda, V)$ of formulas over Λ and V is then given by the grammar

$$\phi ::= a \mid \bot \mid \neg\phi \mid \phi \wedge \psi \mid L(\phi_1, \dots, \phi_n),$$

where a ranges over V and L over Λ; n is the arity of L. A formula over $V = \emptyset$ is *closed*; the set of closed formulas is denoted by $\mathcal{F}(\Lambda)$. For the sake of readability, we will treat the satisfaction problem explicitly only for closed formulas, admitting variables only in the frame axioms defining a modal logic. However, our results extend straightforwardly to formulas with variables. For any set Z, we let $\mathsf{Prop}(Z)$ denote the set of propositional formulas over Z.

The interpretation of modal operators requires a means of converting predicates on the state space X into predicates on the set TX of structured collections of states:

Definition 4. (Pattinson 2003) An *n-ary predicate lifting* ($n \in \mathbb{N}$) for T is a family of maps $\lambda_X : \mathcal{P}(X)^n \to \mathcal{P}(TX)$, where X ranges over all sets, satisfying the *naturality* condition

$$\lambda_X(f^{-1}[A_1], \dots, f^{-1}[A_n]) = (Tf)^{-1}[\lambda_Y(A_1, \dots, A_n)]$$

for all $f : X \to Y$, $A_1, \dots, A_n \in \mathcal{P}(Y)$. (For the categorically-minded reader, a predicate lifting is a natural transformation $\lambda : \mathcal{Q}^n \to \mathcal{Q} \circ T^{op}$, where $\mathcal{Q}$ is contravariant powerset.)

The interpretation of modal operators over coalgebraic models as defined below is determined by the choice of an n-ary predicate lifting $[\![L]\!]$ for every n-ary modality L.

Given a set V of variables, a T-*model* $M = ((X, \xi), \pi)$ is a T-coalgebra (X, ξ) equipped with an assignment $\pi(x) \in \mathcal{P}(V)$ to each state x. A satisfaction relation $\models_M$ between states in M and formulas is defined recursively with the obvious clauses for variables and boolean operators; the clause for an n-ary modal operator $L \in \Lambda$ is

$$x \models_M L(\phi_1, \dots, \phi_n) \text{ iff } \xi(x) \in [\![L]\!]([\![\phi_1]\!]_M, \dots, [\![\phi_n]\!]_M),$$

where $[\![\phi]\!]_M = \{y \mid y \models_M \phi\}$. Closed formulas can be evaluated over T-coalgebras, without an assignment π.

The generic deductive system for coalgebraic modal logics is parametrised by a set of *one-step axioms*. These axioms are of a restricted form: they are of *rank 1*, i.e. they are propositional combinations of formulas $L(\phi_1, \dots, \phi_n)$, where $L \in \Lambda$ is n-ary and the ϕ_i are propositional formulas over a set V of propositional variables. A typical example is the K-axiom $\Box(a \to b) \to \Box a \to \Box b$. Deduction may then employ propositional reasoning, replacement of equivalents (i.e. from $\phi_1 \leftrightarrow \psi_1, \dots, \phi_n \leftrightarrow \psi_n$ infer $L(\phi_1, \dots, \phi_n) \leftrightarrow L(\psi_1, \dots, \psi_n)$), and introduction of substitution instances of the axioms. Generally, we denote propositional entailment of formulas ψ from sets Φ of formulas by $\Phi \vdash_{PL} \psi$.

One has notions of *one-step soundness* and *one-step completeness* for sets of one-step axioms which entail soundness and completeness w.r.t. the class of *all* coalgebras for a given functor, and are typically verified quite easily in concrete examples (Pattinson 2003). Concrete one-step complete axiomatisations for various modal logics, including graded modal logic (whose operators are slight variants of the number restrictions appearing in $\mathcal{PHQ}$), are listed e.g. in (Schröder and Pattinson 2006). Our approach here is to assume a sound and (one-step) complete *background axiomatisation* in rank-1 and then study completeness and decidability for logics interpreted over classes of T-coalgebras defined by arbitrary additional axioms that we call *frame* *conditions*. We write $\mathcal{L}$ for the background axioms together with the additional frame conditions and sometimes refer to formulas of $\mathcal{F}(\Lambda)$ as $\mathcal{L}$-formulas. We write $\mathcal{L} \vdash \phi$ for an $\mathcal{L}$-formula ϕ if ϕ can be derived using the deduction rules described above and additionally substitution instances of the frame conditions. We say that a set Φ of $\mathcal{L}$-formulas is *$\mathcal{L}$-consistent* if there do not exist formulas $\phi_1, \dots, \phi_n \in \Phi$ such that $\mathcal{L} \vdash \neg(\phi_1 \wedge \dots \wedge \phi_n)$. An *$\mathcal{L}$-frame* is a T-coalgebra C such that all states in all models (C, π) satisfy the frame conditions of $\mathcal{L}$. We are interested in the *satisfiability problem* of $\mathcal{L}$, i.e. in deciding for a given (closed) formula ϕ whether there exist an $\mathcal{L}$-frame C and a state x in C such that $x \models_C \phi$. The generality of our approach is demonstrated by the following examples (and many other ones listed e.g. in (Schröder 2007; Schröder and Pattinson 2006)); the detailed treatment of $\mathcal{PHQ}$ is given further below.

Example 5. 1. *Modal logic $S4$:* The Kripke semantics of normal modal logics is captured coalgebraically by the *powerset functor* $\mathcal{P}$, which maps each set X to its powerset $\mathcal{P}(X)$, and each map $f : X \to Y$ to the map $\mathcal{P}(f) : \mathcal{P}(X) \to \mathcal{P}(Y)$ which takes direct images under f. Here, a $\mathcal{P}$-coalgebra $C = (X, \xi : X \to \mathcal{P}X)$ corresponds to the Kripke frame $(X, R \subseteq X \times X)$ defined by xRy iff $y \in \xi(x)$. The similarity type of simple normal modal logics (K, T, $S4$ etc.) consists of a single unary modal operator $\Box$, whose interpretation over $\mathcal{P}$-coalgebras is given as the predicate lifting defined by

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

The naturality condition for this lifting translates into the set-theoretic fact that $A \subseteq f^{-1}[B]$ iff $f[A] \subseteq B$. With ξ, R as above, $\xi(x) \in [\![\Box]\!]_X([\![\phi]\!])$ iff $y \models \phi$ for all y such that xRy; i.e. the coalgebraic semantics yields the standard interpretaion of $\Box$. A one-step complete background axiomatisation is given by $\Box\top$ and the K-axiom $\Box(a \to b) \to \Box a \to \Box b$. The logic $S4$ is then determined by the frame conditions $\Box a \to a$ and $\Box a \to \Box\Box a$, and a Krikpe frame (X, R) is an $S4$-frame iff R is transitive and reflexive.

2. *Conditional logic $CK + CMi$:* The similarity type of conditional logic consists of a single binary infix modal operator $\Rightarrow$, read as a non-monotonic (e.g. default) conditional. The semantics of the conditional logic CK is defined over *(standard) conditional frames* (Chellas 1980), captured coalgebraically by the functor Cf defined by $Cf(X) = (\mathcal{P}(X) \to \mathcal{P}(X))$, with $A \to B$ denoting the set of functions from A to B. Conditional frames, i.e. Cf-coalgebras $(X, \xi : X \to Cf(X))$, may be identified with structures $(X, (R_A)_{A \subseteq X})$, where the R_A are binary relations on X. The semantics of $\Rightarrow$ is given by

$$[\![\Rightarrow]\!]_X(A, B) = \{f : \mathcal{P}(X) \to \mathcal{P}(X) \mid f(A) \subseteq B\}$$

— i.e. $x \models \phi \Rightarrow \psi$ iff $xR_{[\![\phi]\!]}y$ implies $y \models \psi$. A complete (background) axiomatisation of the conditional logic of the class of all conditional frames, the logic CK, consists essentially of commutation with conjunction in the second argument, and only replacement of equivalents in the first argument (Chellas 1980). A frame condition of interest is e.g. the axiom

$$(CMi) \quad ((a \wedge b) \Rightarrow c) \to (a \Rightarrow (a \wedge b \Rightarrow c))$$

that allows material consequences of assumptions to be introduced as further assumptions (dually, the well-known cautious monotony axiom $(a \Rightarrow b) \wedge (a \Rightarrow c) \to (a \wedge b \to c)$ allows strengthening the assumption by its default consequences). A conditional frame $(X, (R_A)_{A \subseteq X})$ satisfies CMi iff for all $B \subseteq A \subseteq X$,

$$R_A ; R_B \subseteq R_B.$$

Finite Models and Decidability

We now exhibit a generic model construction that establishes the finite model property for many modal logics with frame conditions in a uniform way. We proceed in two steps: firstly, we construct finite T-coalgebras by a filtration technique, where attention is restricted to a finite closed set Σ of formulas, following (Schröder 2007); however, while in the earlier construction such models would only satisfy instances of the frame conditions which are in $\mathsf{Prop}(\Sigma)$, we note that one can additionally force satisfaction of instances containing a finite number of distinguished *spikes*, i.e. subformulas that protrude above Σ by exactly one modal operator. For many logics, in particular for $\mathcal{PHQ}$ as we shall see later, these *spike instances* are sufficient to allow completing such models to actual $\mathcal{L}$-frames while preserving satisfaction of Σ-formulas. For the sake of readability, we will restrict the further development to unary modal operators; the extension to polyadic operators is a purely notational exercise.

Definition 6. A set Σ of $\mathcal{L}$-formulas is *closed* if Σ is closed under subformulas and negation, identifying $\neg\neg\phi$ and ϕ. A subset $H \subseteq \Sigma$ that is maximally consistent w.r.t. propositional reasoning is called a *Σ-Hintikka set*, and a *Σ-atom* is an $\mathcal{L}$-consistent Σ-Hintikka set. A *Σ-spike* is a formula $L(\psi_1, \dots, \psi_n)$, where the ψ_i are in $\mathsf{Prop}(\Sigma)$, and a *Σ-spike set* is a finite closed superset $\overline{\Sigma} \supseteq \Sigma$ that arises as the closure of $\Sigma \cup \Theta$, where Θ is a set of Σ-spikes. The set $\overline{\Sigma}$ gives rise to the set $\mathcal{L}(\overline{\Sigma}) = \{\phi \in \mathsf{Prop}(\overline{\Sigma}) \mid \vdash_{\mathcal{L}} \phi\}$ of *spike instances* of the frame conditions.

The proof of the finite model property hinges on two facts: given a finite closed set Σ, we can (a) construct a coalgebra structure on the set of Σ-atoms that additionally satisfies the spike instances, and (b) the resulting model can be completed to an $\mathcal{L}$-frame. The first part is discharged using an extension of the known filtration construction for rank-1 logics (Schröder 2007):

Definition 7. For a set S of Σ-Hintikka sets, a *sieve* on S is a map ν assigning to each $A \in S$ a $\bar{\Sigma}$-Hintikka set $\nu(A) \supseteq A$. We say that ν is *$\mathcal{L}$-consistent* if $\nu(A)$ is $\mathcal{L}$-consistent for all $A \in S$. We call a coalgebra structure ξ on S *ν-coherent* if

$$L\phi \in \nu(A) \quad \text{iff} \quad \xi(A) \in [\![L]\!]\phi{\uparrow} S \tag{3}$$

whenever $L\phi \in \bar{\Sigma}$, where $\phi{\uparrow} S = \{B \mid B \vdash_{PL} \phi\}$.

Lemma 8. *An $\mathcal{L}$-consistent sieve on S exists iff S consists of $\mathcal{L}$-atoms.*

Lemma 9 (Lindenbaum)**.** *If $\phi \in \Sigma$ is $\mathcal{L}$-consistent, then there exists an atom A such that $\phi \in A$.*

Lemma 10 (Existence lemma)**.** *Let ν be a consistent Γ-sieve on the set S of Σ-atoms. Then there exists a ν-coherent coalgebra structure on S.*

The proof of the existence lemma relies crucially on one-step completeness of the background axiomatisation.

Lemma 11 (Truth lemma)**.** *Let ν be a sieve on a set S of Σ-Hintikka sets, and let ξ be a ν-coherent coalgebra structure on S. Then*

$$A \models_{(S,\xi)} \phi \text{ iff } \phi \in \nu(A) \qquad \text{for all } A \in S, \phi \in \bar{\Sigma}.$$

Corollary 12. *If ν is an $\mathcal{L}$-consistent sieve on S and ξ is a ν-coherent coalgebra structure on S, then $(S, \xi) \models \mathcal{L}(\Sigma)$.*

In other words, the model whose existence is guaranteed by the existence lemma already satisfies the spike instances of the frame conditions. The second step of the model construction completes this model so that all instances of the frame conditions are satisfied. This construction succeeds if the logic under consideration is closure stable:

Definition 13. A coalgebra C is Σ-*filtered* if different states in C can be distinguished by Σ-formulas. Coalgebras $A = (X, \xi_A)$ and $C = (X, \xi_C)$ are Σ-*equivalent* if for all $x \in X$ and all $\phi \in \Sigma$, $x \models_A \phi$ iff $x \models_C \phi$. A set $\mathcal{S}$ of formulas is Σ-*sufficient* if, whenever $C \models \mathcal{S}$ for a Σ-filtered coalgebra $C = (X, \xi)$, then there exists an $\mathcal{L}$-frame on X which is Σ-equivalent to C. We say that $\mathcal{L}$ is *closure stable* if it can be equipped with assignments of

- a finite closed set $\Sigma(\phi)$ containing ϕ
- a spike set $\overline{\Sigma}(\phi)$ for $\Sigma(\phi)$, and
- a $\Sigma(\phi)$-sufficient set $\mathcal{S}(\phi) \subseteq \mathcal{L}(\overline{\Sigma}(\phi))$

to every $\mathcal{L}$-formula ϕ.

For purposes of completeness proofs, one may always take $\mathcal{S}(\phi) = \mathcal{L}(\overline{\Sigma}(\phi))$. Complexity considerations often require more economic choices. Closure stable logics enjoy the finite model property, and hence completeness:

Theorem 14 (Weak completeness and finite model property)**.** *If $\mathcal{L}$ is closure stable, then $\mathcal{L}$ is weakly complete for finite $\mathcal{L}$-frames. Indeed, every $\mathcal{L}$-consistent formula ϕ is satisfiable in an $\mathcal{L}$-frame with at most $2^{|\Sigma(\phi)|/2}$ states.*

Proof. By the Lindenbaum lemma, ϕ is contained in a $\Sigma(\phi)$-atom A_ϕ, and by Lemma 8, there exists a sieve ν on the set S of atoms. The existence lemma thus guarantees the existence of a T-coalgebra C on S which by Corolloary 12 satisfies $\mathcal{S}(\phi)$ and in which $A_\phi \models_C \phi$ by the truth lemma. By closure stability, this induces an $\mathcal{L}$-frame with the same properties. ◻

Proving that a logic $\mathcal{L}$ is closure-stable crucially requires the completion of Σ-filtered coalgebras to $\mathcal{L}$-frames. This can often be done using inductive arguments; in the examples, the situation is as follows:

Example 15. In the examples below, $\Sigma(\phi)$ is simply the closure of $\{\phi\}$, and provisionally we let $\overline{\Sigma}(\phi)$ consist of $\mathsf{Prop}(\Sigma(\phi))$ and all applications of modal operators to formulas from $\mathsf{Prop}(\Sigma(\phi))$, and put $\mathcal{S}(\phi) = \mathcal{L}(\overline{\Sigma}(\phi))$. Smaller spike sets and sufficient sets can be read from the proofs.

1. The closure process that completes a Kripke frame (X, R) satisfying the spike instances of the $S4$-axioms to an $S4$-frame is transitive reflexive closure; denote the resulting $S4$-frame by (X, R^*). We have to show that (X, R) and (X, R^*) are $\Sigma(\phi)$-equivalent. The proof is by structural induction over formulas, with trivial boolean steps. In the induction step for formulas $\Box\psi \in \Sigma(\phi)$, it is clear that $x \models_{(X,R)} \Box\psi$ whenever $x \models_{(X,R^*)} \Box\psi$. Conversely, we have to prove that, whenever $x \models_{(X,R)} \Box\psi$ and we have a chain $x = x_0 R x_1 R \dots R x_n$, then $x_n \models_{(X,R)} \psi$ (and hence $x_n \models_{(X,R^*)} \psi$ by induction). We prove the stronger claim that $x_n \models_{(X,R)} \Box\psi$ ($x_n \models_{(X,R)} \psi$ follows using the spike instance $\Box\psi \to \psi$) by induction over n. The base case is trivial. If $n > 0$ and $x_{n-1} \models_{(X,R)} \Box\psi$, then $x_{n-1} \models_{(X,R)} \Box\Box\psi$ by satisfaction of the spike instance $\Box\psi \to \Box\Box\psi$, and hence $x_n \models_{(X,R)} \Box\psi$. These arguments essentially appear already in (Halpern and Moses 1992). Note that this example is mainly educational, as one can modify the definition of the transition relation on the set of atoms to obtain a transitive reflexive relation directly (namely, put ARB for atoms A, B iff $\Box\psi \in A$ implies $\Box\psi \in B$ for $\Box\psi \in \Sigma(\phi)$), avoiding a closure process; similar shortcuts do not, however, seem to be available in the other examples, including $\mathcal{PHQ}$.

2. *Conditional logic $CK + CMi$:* The logic $CK + CMi$ (Example 5.2) is closure stable. The closure process completing a $\Sigma(\phi)$-filtered conditional frame $C = (X, (R_A)_{A \subseteq X})$ satisfying the spike instances of CMi to a CMi-frame $\overline{C} = (X, (\overline{R}_A)$ is given by

$$\overline{R}_B = \bigcup_{B \subseteq A_1, \dots, A_n \subseteq X} R_{A_1}; \dots; R_{A_n}; R_B. \qquad (4)$$

The inductive proof that this process preserves satisfaction of $\Sigma(\phi)$-formulas requires spike instances

$$(\rho \Rightarrow \psi) \to (\rho \lor \chi_i \Rightarrow (\rho \Rightarrow \psi))$$

(note $\rho \land (\rho \lor \chi_i) \leftrightarrow \rho$) where the $\chi_i \in \mathsf{Prop}(\Sigma(\phi))$ are chosen such that $A_i = [\![\chi_i]\!]_C$ for given A_i as in (4); such formulas exist by $\Sigma(\phi)$-filteredness.

We now turn to satisfiability algorithms. In the following, fix a closure stable logic $\mathcal{L}$ with assignments $\Sigma(\phi)$, $\overline{\Sigma}(\phi)$ and $\mathcal{S}(\phi) \subseteq \mathcal{L}(\overline{\Sigma}(\phi))$ as in Definition 13, and moreover upper approximations $H(\phi)$ and $\overline{H}(\phi)$ of the sets of $\Sigma(\phi)$-atoms and $\overline{\Sigma}(\phi)$-atoms, respectively (i.e. sets of Hintikka sets containing all atoms, such as the set of all Hintikka sets). For purposes of the algorithmic development, we refine the notion of *sieve* to mean a map $\nu : H(\phi) \to \overline{H}(\phi)$ such that $A \subseteq \nu(A)$ for all A. The complexity of the satisfiability problem now depends on the complexity of two subproblems, the *one-step satisfiability problem* and the *local conformance problem*:

Definition 16. The *local conformance* problem of $\mathcal{L}$ is to decide, given $B \in \overline{H}(\phi)$, whether $B \vdash_{PL} \mathcal{S}(\phi)$ (recall that $\vdash_{PL}$ denotes propositional entailment). The *one-step* *satisfiability* problem (Schröder 2007) of $\mathcal{L}$ is to decide, given a conjunctive clause $\phi = \bigwedge_{i\in I} L_i a_i \wedge \bigwedge_{j\in J} \neg L_j b_j$ where the a_i and b_j are variables from a set V, a finite set X, and a valuation $\tau : V \to \mathcal{P}(X)$, whether $[\![\phi]\!]\tau = \bigcap_{i\in I}[\![L_i]\!]_X(\tau(a_i)) \cap \bigcap_{j\in J}(TX \setminus [\![L_j]\!]_X(\tau(b_j))) \neq \emptyset$.

We now present two algorithms to decide satisfiability of $\phi \in \mathcal{F}(\Lambda)$. Both algorithms exploit the fact that a satisfying model for a formula ϕ will necessarily be based on a set of *consistent* Hintikka sets over $\Sigma(\phi)$. The task of removing all inconsistent Hintikka sets from the given initial over-approximation $H(\phi)$ is performed by the following elimination operator, for the description of which we fix a propositional variable a_ψ for each $\psi \in \mathsf{Prop}(\Sigma(\phi))$. For $S \subseteq H(\phi)$, define the $\mathcal{P}(S)$-valuation τ_S by $\tau_S(a_\psi) = \{B \in S \mid B \vdash_{PL} \phi\}$. Given a sieve ν, the *elimination operator* $\mathcal{E}_\nu$ on $\mathcal{P}(H(\phi))$ is given by

$$\mathcal{E}_\nu(S) = \{A \in S \mid [\![\bigwedge_{L\psi\in\nu(A)} La_\psi \wedge \bigwedge_{\neg L\psi\in\nu(A)} \neg La_\psi]\!]\tau_S \neq \emptyset\}.$$

The first algorithm is deterministic and relies on the fact that the elimination operator is monotone and every fixed point of $\mathcal{E}_\nu$ contains $\Sigma(\phi)$-atoms only.

Algorithm 17 (Decide satisfiability of $\phi \in \mathcal{L}(\Lambda)$ deterministically)**.** For every sieve ν on $H(\phi)$, perform the following steps. If all ν have been checked unsuccessfully, output 'unsatisfiable'.

1. Compute the greatest fixed point S of $\mathcal{E}_\nu$ by iterated application of $\mathcal{E}_\nu$, starting with $\{A \in H(\phi) \mid \nu(A) \vdash_{PL} \mathcal{S}(\phi)\}$.
2. Check whether $\phi \in A$ for some $A \in S$; if yes, output 'satisfiable', otherwise continue.

The non-deterministic counterpart of the above algorithm replaces the systematic computation of the greatest fixed point of $\mathcal{E}_\nu$ by non-deterministically guessing a suitable subset $S \subseteq H(\phi)$ and subsequently checking whether this guess is a fixed point of $\mathcal{E}_\nu$.

Algorithm 18 (Decide satisfiability of $\phi \in \mathcal{L}(\Lambda)$ non-deterministically)**.** Guess $S \subseteq H(\phi)$ and a sieve ν such that $\mathcal{E}_\nu(S) = S$, $\phi \in A$ for some $A \in S$, and $\nu(A) \vdash_{PL} \mathcal{S}(\phi)$ for all $A \in S$.

These algorithms establish decidability and the following complexity bounds:

Theorem 19 (Decidability, complexity)**.** *1. If $\Sigma(\phi)$ and $\overline{\Sigma}(\phi)$ are computable, and membership in $H(\phi)$ and in $\overline{H}(\phi)$, local conformance, and one-step satisfiability are semi-decidable, then $\mathcal{L}$ is decidable.*

2. If $H(\phi)$ and $\overline{\Sigma}(\phi)$ are of at most exponential size in ϕ and computable in exponential time, membership in $\overline{H}(\phi)$ is in NP, local conformance is in $NEXPTIME$, and one-step satisfiability is in NP, then $\mathcal{L}$ is in $NEXPTIME$.

3. If $\Sigma(\phi)$ and $\overline{\Sigma}(\phi)$ are of polynomial size in ϕ and computable in exponential time, membership in $H(\phi)$ and in $\overline{H}(\phi)$ and local conformance are in $EXPTIME$, and one-step satisfiability is in P, then $\mathcal{L}$ is in $EXPTIME$.

Proof. 1.: By Algorithm 18, the semi-decidability conditions imply semi-decidability of the satisfiability problem, which is hence decidable, as semi-decidability of unsatisfiability follows already by completeness.

2.: By a time analysis of Algorithm 18.

3.: By a time analysis of Algorithm 17. □

Example 20. Using a completion process similar to the one described in Example 15, the above theorem establishes the (known) tight $EXPTIME$ upper bound for description logics with transitive roles and role hierarchies (but without number restrictions) as long as transitive roles are upwards closed in the role hierarchy. Moreover, we immediately obtain decidability of $CK+CMi$ as this logic is closure stable.

The $NEXPTIME$-upper bound established in Theorem 19 is pertinent to the logic $\mathcal{PHQ}$, which we discuss next.

Decidability of $\mathcal{PHQ}$

We now apply the above generic results to $\mathcal{PHQ}$, using the customary notation and terminology of modal logic whenever this is more convenient. In particular we identify the notions of formula and concept.

Coalgebraic semantics We note that the Kripke semantics of $\mathcal{PHQ}$ (which itself fails to be coalgebraic, essentially because the family of maps $\lambda_X : \mathcal{P}(X) \to \mathcal{P}(\mathcal{P}(X))$ taking $A \subseteq X$ to $\{B \in \mathcal{P}(X) \mid |A \cap B| \geq n\}$ violates the naturality condition from Definition 4) is equivalent to a coalgebraic semantics, defined as follows. We work with *multigraphs* (D'Agostino and Visser 2002), a variant of Kripke frames where edges are annotated with positive integer *multiplicities*. Multigraphs may be regarded as coalgebras for (multiple copies of) the *bag functor* $\mathcal{B}$ which takes a set X to the set $\mathcal{B}(X)$ of all bags (or *multisets*) over X; recall that bags B are distinguished from sets as containing elements x with a given multiplicity $B(x) \in \mathbb{N}\cup\{\infty\}$. The semantics of $\mathcal{PHQ}$ over multigraphs is as expected, where the interpretation of number restrictions takes into account multiplicities. Standard complete axiomatisations of graded modal logic, e.g. (Fine 1972), in combination with obvious axioms reflecting the role hierarchy, immediately yield a one-step complete background axiomatisation of $\mathcal{PHQ}$.

Explicitly, a multigraph model $\mathcal{I}$ with domain $\Delta^\mathcal{I}$ assigns to each role R a multiplicity function $R^\mathcal{I} : \Delta^\mathcal{I} \times \Delta^\mathcal{I} \to \mathbb{N} \cup \{\infty\}$, which we extend in the obvious way (i.e. additively) to $\Delta^\mathcal{I} \times \mathcal{P}(\Delta^\mathcal{I})$. The interpretation of a qualified number restriction $\geq n\, R$ over $\mathcal{I}$ is then given by

$$x \models \geq n\, R.\, \phi \iff R^\mathcal{I}(x, [\![\phi]\!]) \geq n.$$

We write $xR^\mathcal{I}y$ if $R^\mathcal{I}(x, y) > 0$. We say that $R^\mathcal{I}$ is *finitely branching* if $R^\mathcal{I}(x, \Delta^\mathcal{I}) < \infty$ for all $x \in \Delta^\mathcal{I}$.

In coalgebraic terms, such multigraph models are coalgebras for the functor $\mathcal{B}^{\bar{\mathcal{R}}}$ taking a set X to the set $\mathcal{B}(X)^{\bar{\mathcal{R}}}$, where $\bar{\mathcal{R}}$ denotes the set of all roles (simple or descendant). The interpretation of qualified number restrictions is induced by the predicate liftings defined by

$$[\![\geq n\, R]\!]_X(A) = \{(B_R)_{R\in\bar{\mathcal{R}}} \in \mathcal{B}(X)^{\bar{\mathcal{R}}} \mid \textstyle\sum_{x\in A} B_R(x) \geq n\}.$$

To capture treelikeness, we interpret $\mathcal{PHQ}$ over *descendancy multigraphs*, defined by the *descendancy equation*

$$(R^D)^{\mathcal{I}}(x,z) = R^{\mathcal{I}}(x,z) + \sum_{xR^{\mathcal{I}}y} R^{\mathcal{I}}(x,y)(R^D)^{\mathcal{I}}(y,z)$$

for all simple roles R. The crucial point that makes the proof of closure stability work for $\mathcal{PHQ}$ where it fails for arbitrary transitive relations is that the descendancy relation generalises to subsets, i.e. we have

Lemma 21. *For any subset A of the domain $\Delta^{\mathcal{I}}$ of a descendancy multigraph,*

$$(R^D)^{\mathcal{I}}(x,A) = R^{\mathcal{I}}(x,A) + \sum_{xR^{\mathcal{I}}y} R^{\mathcal{I}}(x,y)(R^D)^{\mathcal{I}}(y,A).$$

Contrastingly, mere transitivity of a role R translates into the condition $R^{\mathcal{I}}(x,z) \geq \sup_{xR^{\mathcal{I}}y} R^{\mathcal{I}}(y,z)$ on multigraphs, which without the treelikeness assumption does not generalise to sets $A \subseteq \Delta^{\mathcal{I}}$ in place of z – in a nutshell, sums do not commute with suprema, but sums do commute with sums.

Satisfiability of $\mathcal{PHQ}$ concepts over descendancy multigraphs is equivalent to satisfiability over Kripke frames as defined above. This is formally captured by the following theorem, whose (not entirely trivial) proof we omit:

Theorem 22. *A $\mathcal{PHQ}$ concept is satisfiable in a Kripke frame with $R^{\mathcal{I}}$ finitely branching for all simple roles R iff it is satisfiable in a corresponding descendancy multigraph.*

Remark 23. The finite model result proved below implies that the finite branching condition in Theorem 22 can be removed.

Axiomatisation The class of descendancy multigraphs is defined by the following frame conditions: one has *lower bounds* (on numbers of descendants)

$$\bigwedge_{i=1}^{m} \geq n_i R.\,(C_i \wedge \geq k_i R^D.\,b) \rightarrow \geq \big(\textstyle\sum_{i=1}^{q} n_i + \sum_{i=1}^{m} n_i k_i\big) R^D.\,b,$$

where $1 \leq q \leq m$, $0 \leq n_i, k_i$, $C_i = a_i \wedge \bigwedge_{i \neq j} \neg a_j \wedge b$ for $1 \leq i \leq q$, $C_i = a_i \wedge \bigwedge_{i \neq j} \neg a_j$ for $q+1 \leq i \leq n$, and b and the a_i are variables; and *upper bounds*

$$\geq r R^D.\,b \rightarrow \geq 1 R.\,\geq r R^D.\,b \vee \bigvee_{(n_i,k_i),q} \bigwedge_{i=1}^{m} \geq n_i R.\,(C_i \wedge = k_i R^D.\,b),$$

where $r \geq 1$, $C_i = b$ for $i \leq q$, $C_i = \neg b$ otherwise, and the big disjunction ranges over all families $(n_i,k_i)_{i=1,\dots,m}$ and all $1 \leq q \leq m$ such that $n_i, k_i \geq 1$, $k_i < r$ and $n_i \leq r$ for all i, $\sum_{i=1}^{q} n_i + \sum_{i=1}^{m} n_i k_i \geq r$, and $k_i \neq k_j$ for $i \neq j$ and $i,j \leq q$ or $i,j > q$. These axioms capture natural intuitions about counting descendants in trees: the right hand side of the lower bound counts R-children ($\sum_{i=1}^{q} n_i$) and R-descendants of R-children ($\sum_{i=1}^{m} n_i k_i$) satisfying b; and the right hand side of the upper bound enumerates all possible ways in which r or more R-descendants satisfying b may arise as either R-children or R-descendants of R-children.

Note that the lower bound axiom implies

$$\bigwedge_{i=1}^{m} \geq n_i R.\,(\phi_i \wedge \geq k_i R^D.\,b) \rightarrow \geq r R^D.\,b \qquad (5)$$

whenever $\sum_{i \leq q} n_i + \sum n_i k_i \geq r$, where $1 \leq q \leq m$, $0 \leq n_i, k_i$, and the ϕ_i are are mutually exclusive propositional formulas such that ϕ_i entails b for $i \leq q$. Instances of this formula will be used in the proof of closure stability.

Closure stability In the finite model construction, we exploit that theories of states are closed under decreasing numbers in number restrictions $\geq nR$. This is reflected in the following concept.

Definition 24. A set A of $\mathcal{PHQ}$-formulas is called *downward closed* if for every role R, whenever $\geq nR.\,\psi \in A$ and $n \geq m$ then $\geq mR.\,\psi \in A$.

Given a formula ϕ, let $\Sigma(\phi)$ be the smallest closed set containing ϕ which is moreover downward closed (note that $\Sigma(\phi)$ is of exponential size). Spikes may protrude above $\Sigma(\phi)$ by those modal operators $\geq nR$ that occur in $\Sigma(\phi)$. Formally, let Γ_ϕ denote the subsignature containing all modal operators $\geq rS$ occurring in $\Sigma(\phi)$. We shall prove

Theorem 25. *$\mathcal{PHQ}$ is closure stable w.r.t. the choice of spike set $\Gamma_\phi(\mathsf{Prop}(\Sigma(\phi))) \cup \mathsf{Prop}(\Sigma(\phi))$ for each formula ϕ.*

A smaller spike set $\overline{\Sigma}(\phi)$ and a small sufficient set $\mathcal{S}(\phi)$ will be read from the proofs later.

The definition of the closure process which transforms a $\Sigma(\phi)$-filtered multigraph $\mathcal{I}$ satisfying the spike instances into a descendancy multigraph $\mathcal{J}$ by adding missing descendants uses a case distinction which parallels the types of possible solutions of the descendancy equation: There are two types of exceptional finite solutions, the zero case and the case of infinitely deferred descendancy along so-called stable cycles, explained in more detail below. In the remaining (standard) case, the descendancy equation is just taken as a recursive definition of $(R^D)^{\mathcal{I}}$. We present details of the development for a trivial role hierarchy $\mathcal{R} = \{R\}$, so that we have only two roles R and R^D, and thereafter comment on the extension to full role hierarchies.

Thus let $\mathcal{I}$ be a $\Sigma(\phi)$-filtered descendancy multigraph that satisfies all spike instances of $\mathcal{PHQ}$ for ϕ. Put $\Sigma(\phi)/R^D = \{\psi \mid \geq 1R^D.\,\psi \in \Sigma(\phi)\}$; for $\psi \in \Sigma(\phi)/R^D$, let $k(\psi) = \max\{n \mid \geq nR^D.\,\psi \in \Sigma(\phi)\}$. We proceed to construct a $\Sigma(\phi)$-equivalent $\mathcal{PHQ}$-frame $\mathcal{J}$ on $\Delta^{\mathcal{I}}$; in preparation for this, we make the solution types of the descendancy equation outlined above precise.

Definition 26. An *R-path* is a sequence $w = x_1 \dots x_n$ of states $x_i \in \Delta^{\mathcal{I}}$ such that $x_i R^{\mathcal{I}} x_{i+1}$ for $i = 1, \dots, n-1$. We understand *cyclic R-paths* $x_1 \dots x_n x_1$, $n \geq 1$, as indexed modulo n.

Lemma 27. *If $x_1 \dots x_n$, $n \geq 1$, is an R-path, $\psi \in \Sigma(\phi)/R^D$, and $x_n \models_{\mathcal{I}} \geq mR^D.\,\psi$, $m \leq k(\psi)$, then $x_1 \models_{\mathcal{I}} \geq mR^D.\,\psi$.*

Proof. The formula $\geq 1R.\,\geq mR^D.\,\psi \rightarrow \geq mR^D.\,\psi$ is a spike instance of the lower bound. □

Definition 28. For $A \subseteq \Delta^{\mathcal{I}}$, the *zero set of* A is the set

$$Z_A = \{x \in \Delta^{\mathcal{I}} \mid \forall y.\, x(R^{\mathcal{I}})^* y \implies R^{\mathcal{I}}(y, A) = (R^D)^{\mathcal{I}}(y, A) = 0\},$$

where $*$ denotes transitive reflexive closure. A cyclic R-path $x_1 \dots x_n x_1$ is A-*stable* if for $i = 1, \dots, n$, $R^{\mathcal{I}}(x_i, x_{i+1}) = 1$, $x_i \notin A$, and $y \in Z_A$ whenever $x_i R^{\mathcal{I}} y$ and $y \neq x_{i+1}$. The set S_A is the set of all states in $\Delta^{\mathcal{I}}$ that lie on some A-stable cyclic R-path. Put $U_A := \Delta^{\mathcal{I}} - (S_A \cup Z_A)$; thus, U_A denotes the 'regular' case w.r.t. A. For $x \in \Delta^{\mathcal{I}}$, we write $Z_x = Z_{\{x\}}$, $S_x = S_{\{x\}}$, and $U_x = U_{\{x\}}$.

By Lemma 27, $x \in Z_A$ iff $(R^D)^{\mathcal{I}}(x, A) = 0$ when $A = [\![\psi]\!]$, $\psi \in \Sigma(\phi)/R^D$.

Lemma 29. *Let $A \subseteq \Delta^{\mathcal{I}}$. If $x \in S_A - Z_A$, then x lies on a unique A-stable cyclic R-path.*

Lemma 30. *If $w = x_1 \dots x_n x_1$, $n \geq 0$, is a cyclic R-path, then for all $\psi \in \Sigma(\phi)/R^D$, either w is $[\![\psi]\!]$-stable, or for all i, $x_i \models_{\mathcal{I}} \geq k(\psi) R^D.\, \psi$, or for all i, $x_i \models_{\mathcal{I}} \leq 0 R^D.\, \psi$.*

Proof. Let $\psi \in \Sigma(\phi)/R^D$, put $A = [\![\psi]\!]$, and assume that none of the alternatives holds. By Lemma 27, failure of the last two alternatives implies that we have $1 \leq m < k(\psi)$ such that $(R^D)^{\mathcal{I}}(x_i, A) = m$ for all i.

W.l.o.g A-stability fails at $i = 1$. If $R^{\mathcal{I}}(x_1, x_2) > 1$, then by the spike instance

$$\geq 2R.\, \geq m R^D.\, \psi \to \geq \max(2m, k(\psi)) R^D.\, \psi$$

of the lower bound, $x_1 \models_{\mathcal{I}} \geq \max(2m, k(\psi)) R^D.\, \psi$; contradiction, as $m < \max(2m, k(\psi))$. If $x_1 \in A$, then by the spike instance

$$\geq 1R.\, (\psi \wedge \geq m R^D.\, \psi) \to \geq (m+1) R^D.\, \psi$$

of the lower bound, $x_n \models_{\mathcal{I}} \geq (m+1) R^D.\, \psi$, contradiction. Otherwise, we have $y \neq x_2$ such that $x_1 R y$ and $y \notin Z_A$, hence $y \models_{\mathcal{I}} \geq 1 R^D.\, \psi$. As $\mathcal{I}$ is $\Sigma(\phi)$-filtered, we have $\chi \in \Sigma(\phi)$ such that $x_2 \models \chi$, $y \models \neg\chi$. By the spike instance

$$\geq 1R.\, (\chi \wedge \geq 0 R^D.\, \psi) \wedge \geq 1R.\, (\neg\chi \wedge \geq m R^D.\, \psi) \to \geq (m+1) R^D.\, \psi$$

of (5), it follows that $x \models_{\mathcal{I}} \geq (m+1) R^D.\, \psi$, contradiction. □

We now construct a descendancy multigraph $\mathcal{J}$ by completing $(R^D)^{\mathcal{I}}$ to an actual descendancy relation $(R^D)^{\mathcal{J}}$ as follows. Let $z \in \Delta^{\mathcal{I}}$. For $x \in Z_z$, put $(R^D)^{\mathcal{J}}(x, z) = 0$. For every $\{z\}$-stable cyclic R-path w, choose an element $x(w)$ of w and put $(R^D)^{\mathcal{J}}(y, z) = (R^D)^{\mathcal{I}}(x(w), z)$ for every y occurring in w (this is well-defined by Lemma 29 and the fact that every cyclic R-path containing an element of Z_z lies entirely in Z_z). For $x \in U_z$, recursively define

$$(R^D)^{\mathcal{J}}(x, z) = R^{\mathcal{I}}(x, z) + \textstyle\sum_{x R^{\mathcal{I}} y} R^{\mathcal{I}}(x, y) (R^D)^{\mathcal{J}}(y, z), \tag{6}$$

on the understanding that $(R^D)^{\mathcal{J}}(x, z) = \infty$ if the restriction $R^{\mathcal{I}}|_{U_z}$ of $R^{\mathcal{I}}$ to U_z is non-well-founded at x. Otherwise, the recursion defines a finite number.

Lemma 31. *$\mathcal{J}$ is a descendancy multigraph.*

Proof. Straightforward. □

We now have to prove that $\mathcal{J}$ satisfies the same $\Sigma(\phi)$-formulas as $\mathcal{I}$, i.e.

Lemma 32. *$\mathcal{J}$ is $\Sigma(\phi)$-equivalent to $\mathcal{I}$.*

The proof requires two additional lemmas.

Lemma 33. *If $\psi \in \Sigma(\phi)/R^D$ and $x \models_{\mathcal{I}} \geq mR.\, \psi$ for some $m \leq k(\psi)$, then $x_n \models_{\mathcal{I}} \geq m R^D.\, \psi$.*

Proof. The formula $\geq mR.\, (\psi \wedge \geq 0.\psi) \to \geq m R^D.\, \psi$ is a spike instance of the lower bound. □

Lemma 34. *Let $A \subseteq \Delta^{\mathcal{I}}$.*

1. *For $x \in Z_A$, $(R^D)^{\mathcal{J}}(x, A) = 0$.*
2. *For $x \in U_A$, $(R^D)^{\mathcal{J}}(x, A)$ is finite iff $R^{\mathcal{I}}|_{U_A}$ is well-founded at x.* □

The core of the argument is now as follows.

Proof (Lemma 32). Induction over ϕ, with trivial boolean steps. In the step for modal operators, we have to prove for $\geq nS.\psi \in \Sigma(\phi)$, with $S \in \{R, R^D\}$ and $[\![\psi]\!]_{\mathcal{J}} = [\![\psi]\!]_{\mathcal{I}} =: A$, that $x \models_{\mathcal{I}} \geq nS.\psi$ iff $x \models_{\mathcal{J}} \geq nS.\psi$. The case $S = R$ is trivial; let $S = R^D$. If $x \in Z_A$, then $(R^D)^{\mathcal{J}}(x, A) = 0 = (R^D)^{\mathcal{I}}(x, A)$ by Lemma 34, and we are done. If x lies on an A-stable cyclic R-path w, then w is also $\{z\}$-stable for all $z \in Z$, so that $(R^D)^{\mathcal{J}}(x, A) = (R^D)^{\mathcal{I}}(x(w), A)$, i.e. $x \models_{\mathcal{J}} \geq n R^D.\, \psi$ iff $x(w) \models_{\mathcal{I}} \geq n R^D.\, \psi$ iff, by Lemma 27, $x \models_{\mathcal{I}} \geq n R^D.\, \psi$.

For the case $x \in U_A$, we proceed by a case distinction on whether $R^{\mathcal{I}}|_{U_A}$ is well-founded at x. If $R^{\mathcal{I}}|_{U_A}$ is non-well-founded, then $(R^D)^{\mathcal{J}}(x, A) = \infty$ by Lemma 34; in particular, $x \models_{\mathcal{J}} \geq n R^D.\, \psi$. Moreover, one reaches a cyclic path w in U_A from x via $R^{\mathcal{I}}$. By Lemmas 33 and 27, $\geq 1 R^D.\, \psi$ holds on $U_A \subseteq X - Z_A$. By Lemma 30, it follows that $\geq k(\psi) R^D.\, \psi$ holds on w (as w lies outside S_A). By Lemma 27, $x \models_{\mathcal{I}} \geq k(\psi) R^D.\, \psi$ and hence $x \models_{\mathcal{I}} \geq n R^D.\, \psi$.

Finally, if $R^{\mathcal{I}}|_{U_A}$ is well-founded at x, we proceed by well-founded induction on x, where we may now use the inductive hypothesis for all $x R^{\mathcal{I}} y$ (including $y \notin U_A$).

'Only if': By a shallow instance of the upper bound, we have one of the following two cases.

1. $x \models_{\mathcal{I}} \geq 1R.\, \geq n R^D.\, \psi$. Thus we have $x R^{\mathcal{I}} y$ such that $y \models_{\mathcal{I}} \geq n R^D.\, \psi$. By induction, $y \models_{\mathcal{J}} \geq n R^D.\, \psi$, i.e. $(R^D)^{\mathcal{J}}(y, A) \geq n$ and hence by Lemma 21 $(R^D)^{\mathcal{J}}(x, A) \geq n$, i.e. $x \models_{\mathcal{J}} \geq n R^D.\, \psi$.
2. We have $n \leq \sum_{i=1}^{q} n_i + \sum_{i=1}^{m} n_i k_i$, with $n_i, k_i > 0$ for all i, $q \leq m$, and the k_i pairwise distinct within $1 \leq i \leq q$ and within $q < i \leq m$, such that for $1 \leq i \leq m$,

$$x \models_{\mathcal{I}} \geq n_i R.\, \rho_i, \quad \text{with } \rho_i :\equiv \chi_i \wedge = k_i R^D.\, \psi,$$

where $\chi_i = \psi$ if $i \leq q$ and $\chi_i = \neg\psi$ otherwise. Then $y \models_{\mathcal{J}} = k_i R^D.\, \psi$ by induction, i.e. $(R^D)^{\mathcal{J}}(y, A) = k_i$.

Thus

$$\begin{aligned}(R^D)^{\mathcal{J}}(x,A) &\\ &= R^{\mathcal{I}}(x,A) + \textstyle\sum_{xR^{\mathcal{I}}y} R^{\mathcal{I}}(x,y)(R^D)^{\mathcal{J}}(y,A) \\ &\qquad\qquad \text{(Lemma 21)} \\ &\geq R^{\mathcal{I}}(x,A) + \\ &\qquad \textstyle\sum_{i=1}^{m}\sum_{xR^{\mathcal{I}}y\models_{\mathcal{I}}\rho_i} R^{\mathcal{I}}(x,y)(R^D)^{\mathcal{J}}(y,A) \\ &\qquad\qquad (\rho_i \text{ mut. exclusive}) \\ &\geq \textstyle\sum_{i=1}^{q} n_i + \sum_{i=1}^{m} n_i k_i \\ &\geq n,\end{aligned}$$

where in the lower estimate for $R^{\mathcal{I}}(x,A)$ in the penultimate step, we exploit that $\rho_1,\dots,\rho_q$ are mutually exclusive and imply ψ. It follows that $x \models_{\mathcal{J}} \geq nR^D.\,\psi$.

'If': Let $x \models_{\mathcal{J}} \geq nR^D.\,\psi$, i.e. $(R^D)^{\mathcal{J}}(x,A) \geq n$. For $xR^{\mathcal{I}}y$, put $n_y = \min(k(\psi), R^{\mathcal{I}}(x,y))$ and $k_y = \min(k(\psi), (R^D)^{\mathcal{J}}(y,A))$, and let $\chi_y \in \mathsf{Prop}(\Sigma(\phi))$ be the $\Sigma(\phi)$-theory of y, i.e. the conjunction of all $\Sigma(\phi)$-formulas satisfied by y. By $\Sigma(\phi)$-filteredness, $[\![\chi_y]\!]_{\mathcal{I}} = \{y\}$. Then $y \models_{\mathcal{J}} \geq k_y R^D.\,\psi$, and hence $y \models_{\mathcal{I}} \geq k_y R^D.\,\psi$ by induction. Thus,

$$x \models_{\mathcal{I}} \geq n_y R.\,(\chi_y \wedge \geq k_y R^D.\,\psi)$$

for all $xR^{\mathcal{I}}y$. By a spike instance of (5), it follows that $x \models \geq r.\psi$, where

$$r = \min(k(\psi), \textstyle\sum_{xR^{\mathcal{I}}y, y\models_{\mathcal{I}}\psi} n_y + \sum_{xR^{\mathcal{I}}y} n_y k_y).$$

Since $n \leq k(\psi)$, we are done if $r = k(\psi)$, $k_y = k(\psi)$ for some $xR^{\mathcal{I}}y$, or $n_y = k(\psi)$ for some $xR^{\mathcal{I}}y$ such that $k_y \neq 0$ or $y \models_{\mathcal{I}} \psi$. Moreover, if $k_y = 0$ and $y \not\models_{\mathcal{I}} \psi$, then the value of n_y is immaterial. Thus w.l.o.g. $r = \sum_{xR^{\mathcal{I}}y, y\models_{\mathcal{I}}\psi} n_y + \sum_{xR^{\mathcal{I}}y} n_y k_y$ and $n_y = R^{\mathcal{I}}(x,y)$, $k_y = (R^D)^{\mathcal{J}}(y,A)$ for all $xR^{\mathcal{I}}y$. Hence

$$\begin{aligned} r &= R^{\mathcal{I}}(x,A) + \textstyle\sum_{xR^{\mathcal{I}}y} R^{\mathcal{I}}(x,y)(R^D)^{\mathcal{J}}(y,A) \\ &= (R^D)^{\mathcal{J}}(x,A) \qquad \text{(Lemma 21)} \\ &\geq n,\end{aligned}$$

so that $x \models \geq nR^D.\,\psi$ as required. □

Extending these results to non-trivial role hierarchies essentially requires noting that zero sets and stable cycles are sufficiently well-behaved under the role ordering. This completes the proof of Theorem 25. By Theorem 14, it follows that the above axiomatisation, together with the background axiomatisation for number restrictions and the role hierarchy, is complete, and that $\mathcal{PHQ}$ has the finite model property.

As announced above, we may now read the precise spike instances required from the above proofs and apply Theorem 19 to obtain

Theorem 35. *Satisfiability of $\mathcal{PHQ}$ concepts, with numbers coded in binary, is in $NEXPTIME$.*

Proof. We collect the required spike instances from the proofs of Lemmas 27–34. The crucial instances w.r.t. complexity (where 'crucial' just means 'large') are the instances of the upper bound and (5) appearing in the proof of Lemma 32; recall that the latter are of the form

$$\bigwedge_y \geq n_y R.\,(\chi_y \wedge \geq k_y R^D.\,C) \to \geq rR^D.\,C, \qquad (7)$$

where y ranges over the $R^{\mathcal{I}}$-successors of x and χ_y is the conjunction of all $\Sigma(\phi)$-formulas satisfied by y (and in particular actually subsumes $\geq k_y R^D.\,C$, which in the proof of Lemma 32 is assumed to be satisfied by y). The upper bound instances are easily seen to contribute at most exponentially to the size of $\overline{\Sigma}(\phi)$. While the χ_y in (7), being subsets of the exponential-sized set $\Sigma(\phi)$, might be suspected to be of doubly exponential size, they are in fact downward closed (in slight abuse of terminology) and as such can be represented in exponential size. Thus, we obtain an exponential-sized spike set $\overline{\Sigma}(\phi)$ which is clearly computable in exponential time.

We then take the upper approximation $H(\phi)$ of the set of $\Sigma(\phi)$-atoms to be the set of all downward closed $\Sigma(\phi)$-Hintikka sets, which is of exponential size. Similarly, we take $\overline{H}(\phi)$ to be the set of all downward closed $\overline{\Sigma}(\phi)$-Hintikka sets. It is clear that $H(\phi)$ is computable in exponential time and that membership in $\overline{H}(\phi)$ is in NP. Concerning local conformance, note that in (7), the χ_y are pairwise distinct by $\Sigma(\phi)$-filteredness. Moreover, given $B \in \overline{H}(\phi)$, one can calculate a single relevant maximal n_y for each y because B is downward closed. Thus, there is essentially only a single exponentially large formula of the above type (7) to check for B. Checking the further required spike instances is easier, so that local conformance is in $EXPTIME$. Finally, one-step satisfiability is essentially linear programming (Schröder 2007) and hence in NP. The claim now follows by Theorem 19.2. □

The upper bound of Theorem 35 applies also to reasoning with general TBoxes, since these can be internalised as already remarked. We do not know whether the bound is tight; however, similarly expressive logics including many of the decidable instances of $\mathcal{ALCQ}(\mathcal{R})$ are at least $NEXPTIME$-hard (Pratt-Hartmann 2008).

Conclusion

We have presented a new way of dealing with parthood across several levels of decomposition in description logics, based on requiring parthood to be treelike. The arising logic $\mathcal{PHQ}$ provides a better semantic fit for many applications, and in particular allows for better lower estimates on numbers of parts than the standard approach of using arbitrary transitive relations. Moreover, $\mathcal{PHQ}$ avoids the well-known undecidability problems arising with the latter approach. In fact we have shown, using new generic results in coalgebraic modal logic, that satisfiability of $\mathcal{PHQ}$ concepts is always decidable in $NEXPTIME$. We expect that $\mathcal{PHQ}$ can be handled efficiently using a tableau calculus accompanied by suitable heuristics, the design of which is the subject of future investigation. Moreover, extensions of the expressive power of $\mathcal{PHQ}$ are of interest, in particular the addition of inverse roles.

We have achieved these results using novel methods in coalgebraic modal logic, which thereby apply to a wide range of other logics, including, beyond the examples covered above, e.g. a version of coalition logic with multi-step strategies. From the more general perspective, this work constitutes part of an effort to push the applicability of coalgebraic modal logic beyond its original domain, the rank-1 logics; while some results in this direction have been obtained previously (Pattinson and Schröder 2008), the present results constitute the first generic complexity results for logics outside rank 1. Generic results for higher-rank logics of lower complexity, such as PSPACE, are the subject of ongoing research.

Acknowledgements

The work of the first author forms part of the DFG project *Generic Algorithms and Complexity Bounds in Coalgebraic Modal Logic* (SCHR 1118/5-1). The work of the second author was partially supported by EPSRC grant EP/F031173/1. The authors wish to thank Joseph Halpern for helpful discussions. Erwin R. Catesbeiana has volunteered some remarks on consistency.

References

Baader, F.; Calvanese, D.; McGuinness, D. L.; Nardi, D.; and Patel-Schneider, P. F., eds. 2003. *The Description Logic Handbook*. Cambridge University Press.

Bittner, T., and Donnelly, M. 2007. Logical properties of foundational relations in bio-ontologies. *Artificial Intelligence in Medicine* 39:197–216.

Chellas, B. 1980. *Modal Logic*. Cambridge.

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

Fattorosi-Barnaba, M., and de Caro, F. 1985. Graded modalities. I. *Studia Logica* 44:197–221.

Fine, K. 1972. In so many possible worlds. *Notre Dame J. Formal Logic* 13:516–520.

Halpern, J. Y., and Moses, Y. O. 1992. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence* 54:319–379.

Horrocks, I., and Gough, G. 1997. Description logics with transitive roles. In *Description Logics*, volume 410 of *URA-CNRS*.

Horrocks, I.; Patel-Schneider, P. F.; and van Harmelen, F. 2003. From SHIQ and RDF to OWL: the making of a web ontology language. *J. Web Sem.* 1:7–26.

Horrocks, I.; Sattler, U.; and Tobies, S. 1999. Practical reasoning for expressive description logics. In *Logic for Programming and Automated Reasoning, LPAR 99*, volume 1705 of *LNCS*, 161–180. Springer.

Kazakov, Y.; Sattler, U.; and Zolin, E. 2007. How many legs do I have? Non-simple roles in number restrictions revisited. In *Logic for Programming, Artificial Intelligence, and Reasoning, LPAR 2007*, volume 4790 of *LNCS*, 303–317. Springer.

Kupferman, O.; Sattler, U.; and Vardi, M. Y. 2002. The complexity of the graded μ-calculus. In *Automated Deduction, CADE 02*, volume 2392 of *LNCS*, 423–437. Springer.

Pattinson, D., and Schröder, L. 2008. Beyond rank 1: Algebraic semantics and finite models for coalgebraic logics. In *Foundations of Software Science and Computation Structures, FOSSACS 2008*, volume 4962 of *LNCS*, 66–80. Springer.

Pattinson, D. 2003. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoret. Comput. Sci.* 309:177–193.

Pauly, M. 2002. A modal logic for coalitional power in games. *J. Logic Comput.* 12:149–166.

Pratt-Hartmann, I. 2008. On the computational complexity of the numerically definite syllogistic and related logics. *Bull. Symb. Log.* 14:1–28.

Rector, A., and Horrocks, I. 1997. Experience building a large, re-usable medical ontology using a description logic with transitivity and concept inclusions. In *AAAI Spring Symposium, AAAI 1997*. AAAI Press.

Sattler, U. 1996. A concept language extended with different kinds of transitive roles. In *Künstliche Intelligenz, KI 96*, volume 1137 of *LNAI*. Springer.

Schröder, L., and Pattinson, D. 2006. PSPACE reasoning for rank-1 modal logics. In *Logic in Computer Science, LICS 06*, 231–240. IEEE.

Schröder, L. 2007. A finite model construction for coalgebraic modal logic. *J. Log. Algebr. Prog.* 73:97–110. FOSSACS 06 special issue.