

STRONG COMPLETENESS OF COALGEBRAIC MODAL LOGICS

LUTZ SCHRÖDER¹ AND DIRK PATTINSON²

¹ DFKI Bremen and Department of Computer Science, Universität Bremen
E-mail address: Lutz.Schroeder@dfki.de

² Department of Computing, Imperial College London
E-mail address: dirk@doc.ic.ac.uk

ABSTRACT. Canonical models are of central importance in modal logic, in particular as they witness strong completeness and hence compactness. While the canonical model construction is well understood for Kripke semantics, non-normal modal logics often present subtle difficulties – up to the point that canonical models may fail to exist, as is the case e.g. in most probabilistic logics. Here, we present a generic canonical model construction in the semantic framework of coalgebraic modal logic, which pinpoints coherence conditions between syntax and semantics of modal logics that guarantee strong completeness. We apply this method to reconstruct canonical model theorems that are either known or folklore, and moreover instantiate our method to obtain new strong completeness results. In particular, we prove strong completeness of graded modal logic with finite multiplicities, and of the modal logic of exact probabilities.

In modal logic, completeness proofs come in two flavours: *weak* completeness, i.e. derivability of all universally valid formulas, is often proved using *finite model* constructions, and *strong* completeness, which additionally allows for a possibly infinite set of assumptions. The latter entails recursive enumerability of the set of consequences of a recursively enumerable set of assumptions, and is usually established using (infinite) *canonical models*. The appeal of the first method is that it typically entails decidability. The second method yields a stronger result and has some advantages of its own. First, it applies in some cases where finite models fail to exist, which often means that the logic at hand is undecidable. In such cases, a completeness proof via canonical models will at least salvage recursive enumerability. Second, it allows for schematic axiomatisations, e.g. pertaining to the infinite evolution of a system or to observational equivalence, i.e. statements to the effect that certain states cannot be distinguished by any formula.

In the realm of Kripke semantics, canonical models exist for a large variety of logics and are well understood, see e.g. [2]. But there is more to modal logic than Kripke semantics, and indeed the natural semantic structures used to interpret a large class of modal logics go beyond pure relations. This includes e.g. the selection function semantics of conditional logics [4], the semantics

1998 ACM Subject Classification: F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic — modal logic; I.2.4 [Artificial Intelligence]: Knowledge Representation Formalisms and Methods — modal logic, representation languages.

Key words and phrases: Logic in computer science, semantics, deduction, modal logic, coalgebra.

Work of the first author performed as 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.



of probabilistic logics in terms of probability distributions, and the game frame semantics of coalition logic [16]. To date, there is very little research that provides systematic criteria, or at least a methodology, for establishing strong completeness for logics not amenable to Kripke semantics. This is made worse as the question of strong completeness crucially depends on the chosen semantic domain, which as illustrated above may differ widely. It is precisely this variety in semantics that makes it hard to employ the strong-completeness-via-canonicity approach, as in many cases there is no readily available notion of canonical model. The present work improves on this situation by providing a widely applicable generic canonical model construction. More precisely, we establish the existence of quasi-canonical models, that is, models based on the set of maximally consistent sets of formulas that satisfy the truth lemma, as there may be no unique, or canonical, such model in our more general case. In order to cover the large span of semantic structures, we avoid a commitment to a particular class of models, and instead work within the framework of coalgebraic modal logic [15] which precisely provides us with a semantic umbrella for all of the examples above. This is achieved by using coalgebras for an endofunctor T as the semantic domain for modal languages. As we illustrate in examples, the semantics of particular logics is then obtained by particular choices of T . Coalgebraic modal logic serves in particular as a general semantic framework for non-normal modal logics. As such, it improves on neighbourhood semantics in that it retains the full semantic structure of the original models (neighbourhood semantics offers only very little actual semantic structure, and in fact may be regarded as constructed from syntactic material [18]).

In this setting, our criterion can be formulated as a set of coherence conditions that relate the syntactic component of a logic to its coalgebraic semantics, together with a purely semantic condition stating that the endofunctor T that defines the semantics needs to preserve inverse limits weakly, and thus allows for a passage from the finite to the infinite. We are initially concerned with the existence of quasi-canonical models relative to the class of *all* T -coalgebras, that is, with logics that are axiomatisable by formulas of modal depth uniformly equal to one [17]. As in the classical theory, the corresponding result for logics with extra frame conditions requires that the logic is canonical, i.e. the frame that underlies a quasi-canonical model satisfies the frame conditions, which holds in most cases, but for the time being needs to be established individually for each logic.

Our new criterion is then used to obtain both previously known and novel strong completeness results. In addition to positive results, we dissect a number of logics for which strong completeness fails and show which assumption of our criterion is violated. In particular, this provides a handle on adjusting either the syntax or the semantics of the logic at hand to achieve strong completeness. For example, we demonstrate that the failure of strong completeness for probabilistic modal logic (witnessed e.g. by the set of formulas assigning probability $\geq 1 - 1/n$ to an event for all n but excluding probability 1) disappears in the logic of exact probabilities. Moreover, we show that graded modal logic, and more generally any description logic [1] with qualified number restrictions, role hierarchies, and reflexive, transitive, and symmetric roles, is strongly complete over the multigraph model of [5], which admits infinite multiplicities. While strong completeness fails for the naive restriction of this model to multigraphs allowing only finite multiplicities, we show how to salvage strong completeness using additive (finite-)integer-valued measures. Finally, we prove strong completeness of several conditional logics w.r.t. conditional frames (also known as selection function models); for at least one of these logics, strong completeness was previously unknown.

1. Preliminaries and Notation

Our treatment of strong completeness is parametric in both the syntax and the semantics of a wide range of modal logics. On the syntactic side, we fix a *modal similarity type* Λ consisting of modal

operators with associated arities. Given a similarity type Λ and a countable set P of atomic propositions, the set $\mathcal{F}(\Lambda)$ of Λ -formulas is inductively defined by the grammar

$$\mathcal{F}(\Lambda) \ni \phi, \psi ::= p \mid \perp \mid \neg\phi \mid \phi \wedge \psi \mid L(\phi_1, \dots, \phi_n)$$

where $p \in P$ and $L \in \Lambda$ is n -ary; further boolean operators ($\vee, \rightarrow, \leftrightarrow, \top$) are defined as usual. Given any set X (e.g. of formulas, atomic propositions, or sets (!)), we write $\text{Prop}(X)$ for the set of propositional formulas over X and $\Lambda(X) = \{L(x_1, \dots, x_n) \mid L \in \Lambda \text{ is } n\text{-ary}, x_1, \dots, x_n \in X\}$ for the set of formulas arising by applying exactly one operator to elements of X . We instantiate our results to a variety of settings later with the following similarity types:

Examples 1.1. 1. The similarity type Λ_K of standard modal logic consists of a single unary operator \Box .

2. Conditional logic [4] is defined over the similarity type $\Lambda_{\text{CL}} = \{\Rightarrow\}$ where the binary operator \Rightarrow is read as a non-monotonic conditional (default, relevant etc.), usually written in infix notation.

3. Graded modal operators [8] appear in expressive description logics [1] in the guise of so-called qualified number restrictions; although we discuss only modal aspects, we use mostly description logic notation and terminology below. The operators of graded modal logic (GML) are $\Lambda_{\text{GML}} = \{(\geq k) \mid k \in \mathbb{N}\}$ with $(\geq k)$ unary. We write $\geq k. \phi$ instead of $(\geq k)\phi$. A formula $\geq k. \phi$ is read as ‘at least k successor states satisfy ϕ ’, and we abbreviate $\Box\phi = \neg \geq 1. \neg\phi$.

4. The similarity type Λ_{PML} of probabilistic modal logic (PML) [14] contains the unary modal operators L_p for $p \in \mathbb{Q} \cap [0, 1]$, read as ‘with probability at least p, \dots ’.

We split axiomatisations of modal logics into two parts: the first group of axioms is responsible for axiomatising the logic w.r.t. the class of *all* (coalgebraic) models, whereas the second consists of frame conditions that impose additional conditions on models. As the class of all coalgebraic models, introduced below, can always be axiomatised by formulas of *rank 1*, i.e. containing exactly one level of modal operators [17] (and conversely, every collection of such axioms admits a complete coalgebraic semantics [18]), we restrict the axioms in the first group accordingly. More formally:

Definition 1.2. A (modal) logic is a triple $\mathcal{L} = (\Lambda, \mathcal{A}, \Theta)$ where Λ is a similarity type, $\mathcal{A} \subseteq \text{Prop}(\Lambda(\text{Prop}(P)))$ is a set of *rank-1 axioms*, and $\Theta \subseteq \mathcal{F}(\Lambda)$ is a set of *frame conditions*. We say that \mathcal{L} is a *rank-1 logic* if $\Theta = \emptyset$. If $\phi \in \mathcal{F}(\Lambda)$, we write $\vdash_{\mathcal{L}} \phi$ if ϕ can be derived from $\mathcal{A} \cup \Theta$ with the help of propositional reasoning, uniform substitution, and the congruence rule: from $\phi_1 \leftrightarrow \psi_1, \dots, \phi_n \leftrightarrow \psi_n$ infer $L(\phi_1, \dots, \phi_n) \leftrightarrow L(\psi_1, \dots, \psi_n)$ whenever $L \in \Lambda$ is n -ary. For a set $\Phi \subseteq \mathcal{F}(\Lambda)$ of assumptions, we write $\Phi \vdash_{\mathcal{L}} \phi$ if $\vdash_{\mathcal{L}} \phi_1 \wedge \dots \wedge \phi_n \rightarrow \phi$ for (finitely many) $\phi_1, \dots, \phi_n \in \Phi$. A set Φ is *\mathcal{L} -inconsistent* if $\Phi \vdash_{\mathcal{L}} \perp$, and otherwise *\mathcal{L} -consistent*.

Examples 1.3. 1. The modal logic K comes about as the rank-1 logic $(\Lambda_K, \mathcal{A}_K, \emptyset)$ where $\mathcal{A}_K = \{\Box\top, \Box(p \rightarrow q) \rightarrow (\Box p \rightarrow \Box q)\}$. The logics $K4, S4, KB, \dots$ arise as $(\Lambda_K, \mathcal{A}_K, \Theta)$ where Θ contains the additional axioms that define the respective logic [2], e.g. $\Theta = \{\Box p \rightarrow \Box\Box p\}$ in the case of $K4$.

2. For conditional logic, we take the similarity type Λ_{CL} together with rank-1 axioms $r \Rightarrow \top, r \Rightarrow (p \rightarrow q) \rightarrow ((r \Rightarrow p) \rightarrow (r \Rightarrow q))$ stating that the binary conditional is normal in its second argument. Typical additional rank-1 axioms are

(ID)	$a \Rightarrow a$	(identity)
(DIS)	$(a \Rightarrow c) \wedge (b \Rightarrow c) \rightarrow ((a \vee b) \Rightarrow c)$	(disjunction)
(CM)	$(a \Rightarrow c) \wedge (a \Rightarrow b) \rightarrow ((a \wedge b) \Rightarrow c)$	(cautious monotony)

which together form the so-called *System C*, a modal version of the well-known KLM (Krauss/Lehmann/Magidor) axioms of default reasoning due to Burgess [3].

3. The axiomatisation of GML given in [8] consists of the rank-1 axioms

$$\begin{aligned} & \Box(p \rightarrow q) \rightarrow (\Box p \rightarrow \Box q) \\ & \geq k. p \rightarrow \geq l. p \text{ for } l < k \\ & \geq k. p \leftrightarrow \bigvee_{i=0, \dots, k} \geq i. (p \wedge q) \wedge \geq (k-i). (p \wedge \neg q) \\ & \Box(p \rightarrow q) \rightarrow (\geq k. p \rightarrow \geq k. q) \end{aligned}$$

Frame conditions of interest include e.g. reflexivity ($p \rightarrow \geq 1. p$), symmetry ($p \rightarrow \Box \geq 1. p$), and transitivity ($\geq 1. \geq n. p \rightarrow \geq n. p$).

To keep our results parametric also in the semantics of modal logic, we work in the framework of *coalgebraic modal logic* in order to achieve a uniform and coherent presentation. In this framework, the particular shape of models is encapsulated by an endofunctor $T : \mathbf{Set} \rightarrow \mathbf{Set}$, the *signature functor* (recall that such a functor maps every set X to a set TX , and every map $f : X \rightarrow Y$ to a map $Tf : TX \rightarrow TY$ in such a way that composition and identities are preserved), which may be thought of as a parametrised data type. We fix the data Λ, \mathcal{L}, T etc. throughout the generic part of the development. The role of models is then played by T -coalgebras:

Definition 1.4. A T -coalgebra is a pair $\mathbb{C} = (C, \gamma)$ where C is a set (the *state space* of \mathbb{C}) and $\gamma : C \rightarrow TC$ is a function, the transition structure of \mathbb{C} .

We think of TC as a type of successors, polymorphic in C . The transition structure γ associates a structured collection of successors $\gamma(c)$ to each state $x \in C$. The following choices of signature functors give rise to the semantics of the modal logics discussed in Expl. 1.3.

Examples 1.5. 1. Coalgebras for the covariant powerset functor \mathcal{P} defined on sets X by $\mathcal{P}(X) = \{A \mid A \subseteq X\}$ and on maps f by $\mathcal{P}(f)(A) = f[A]$ are Kripke frames, as relations $R \subseteq W \times W$ on a set W of worlds are in bijection with functions of type $W \rightarrow \mathcal{P}(W)$. Restricting the powerset functor to *finite* subsets, i.e. putting $\mathcal{P}_\omega(X) = \{A \subseteq X \mid A \text{ finite}\}$, one obtains the class of image finite Kripke frames as \mathcal{P}_ω -coalgebras.

2. The semantics of conditional logic is captured coalgebraically by the endofunctor \mathcal{S} that maps a set X to the set $(\mathcal{P}(X) \rightarrow \mathcal{P}(X))$ of selection functions over X (the action of \mathcal{S} on functions $f : X \rightarrow Y$ is given by $\mathcal{S}(f)(s)(B) = f[s(f^{-1}[B])]$). The ensuing \mathcal{S} -coalgebras are precisely the conditional frames of [4].

3. The (*infinite*) *multiset functor* \mathcal{B}_∞ maps a set X to the set $\mathcal{B}_\infty X$ of multisets over X , i.e. functions of type $X \rightarrow \mathbb{N} \cup \{\infty\}$. Accordingly, \mathcal{B}_∞ -coalgebras are *multigraphs* (graphs with edges annotated by multiplicities). Multigraphs provide an alternative semantics for GML which is in many respects more natural than the original Kripke semantics [5], as also confirmed by new results below.

4. Finally, if $\text{supp}(\mu) = \{x \in X \mid \mu(x) \neq 0\}$ is the support of a function $\mu : X \rightarrow [0, 1]$ and $\mathcal{D}(X) = \{\mu : X \rightarrow [0, 1] \mid \text{supp}(\mu) \text{ finite}, \sum_{x \in X} \mu(x) = 1\}$ is the set of finitely supported probability distributions on X , then \mathcal{D} -coalgebras are probabilistic transition systems, the semantic domain of PML.

The link between coalgebras and modal languages is provided by predicate liftings [15], which are used to interpret modal operators. Essentially, predicate liftings convert predicates on the state space X into predicates on the set TX of structured collections of states:

Definition 1.6. [15] An n -ary *predicate lifting* ($n \in \mathbb{N}$) for T is a family of maps $\lambda_X : \mathcal{P}X^n \rightarrow \mathcal{P}TX$, 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 $f : X \rightarrow Y, A_1, \dots, A_n \in \mathcal{P}Y$. (For the categorically minded, λ is a natural transformation $\mathcal{Q}^n \rightarrow \mathcal{Q} \circ T^{op}$, where \mathcal{Q} denotes contravariant powerset.) A *structure* for a similarity type Λ over an endofunctor T is the assignment of an n -ary predicate lifting $\llbracket L \rrbracket$ to every n -ary modal operator $L \in \Lambda$.

Given a valuation $V : P \rightarrow \mathcal{P}(C)$ of the propositional variables and a T -coalgebra (C, γ) , a structure for Λ allows us to define a satisfaction relation $\models_{(C, \gamma, V)}$ between states of C and formulas $\phi \in \mathcal{F}(\Lambda)$ by stipulating that $c \models_{(C, \gamma, V)} p$ iff $c \in V(p)$ and

$$c \models_{(C, \gamma, V)} L(\phi_1, \dots, \phi_n) \text{ iff } \gamma(c) \in \llbracket L \rrbracket_C(\llbracket \phi_1 \rrbracket, \dots, \llbracket \phi_n \rrbracket),$$

where $\llbracket \phi \rrbracket = \{c \in C \mid c \models_{(C, \gamma, V)} \phi\}$. An \mathcal{L} -*model* is now a *model*, i.e. a triple (C, γ, V) as above, such that $c \models_{(C, \gamma, V)} \psi$ for all $c \in C$ and all substitution instances ψ of $\mathcal{A} \cup \Theta$. An \mathcal{L} -*frame* is a T -coalgebra (C, γ) such that (C, γ, V) is an \mathcal{L} -model for all valuations V . The reader is invited to check that the following predicate liftings induce the standard semantics for the modal languages introduced in Expl. 1.1.

Examples 1.7. 1. A structure for Λ_K over the covariant powerset functor \mathcal{P} is given by $\llbracket \Box \rrbracket_X(A) = \{Y \in \mathcal{P}(X) \mid Y \subseteq A\}$. The frame classes defined by the frame conditions mentioned in Expl. 1.3.1 are well-known; e.g. a Kripke frame (X, R) is a $K4$ -frame iff R is transitive.

2. Putting $\llbracket \Rightarrow \rrbracket_X(A, B) = \{f \in \mathcal{S}(X) \mid f(A) \subseteq B\}$ reconstructs the semantics of conditional logic in a coalgebraic setting.

3. A structure for GML over \mathcal{B}_∞ is given by $\llbracket (\geq k) \rrbracket_X(A) = \{f : X \rightarrow \mathbb{N} \cup \{\infty\} \mid \sum_{x \in A} f(x) \geq k\}$. The frame conditions mentioned in Expl. 1.3.3 correspond to conditions on multigraphs that can be read off directly from the logical axioms. E.g. a multigraph satisfies the transitivity axiom $\geq 1. \geq n. p \rightarrow \geq n. p$ iff whenever x has non-zero transition multiplicity to y and y has transition multiplicity at least n to z , then x has transition multiplicity at least n to z .

4. The structure over \mathcal{D} that captures PML coalgebraically is given by the the predicate lifting $\llbracket L_p \rrbracket_X(A) = \{\mu \in \mathcal{D}(X) \mid \sum_{x \in A} \mu(x) \geq p\}$ for $p \in [0, 1] \cap \mathbb{Q}$.

From now on, fix a modal logic $\mathcal{L} = (\Lambda, \mathcal{A}, \Theta)$ and a structure for Λ over a functor T . We say that \mathcal{L} is *strongly complete* for some class of models if every \mathcal{L} -consistent set of formulas is satisfiable in some state of some model in that class. Restricting to *finite* sets Φ defines the notion of *weak completeness*; many coalgebraic modal logics are only weakly complete [17].

Definition 1.8. Let X be a set. If $\psi \in \mathcal{F}(\Lambda)$ and $\tau : P \rightarrow \mathcal{P}(X)$ is a valuation, we write $\psi\tau$ for the result of substituting $\tau(p)$ for p in ψ , with propositional subformulas evaluated according to the boolean algebra structure of $\mathcal{P}(X)$. (Hence, $\psi\tau$ is a formula over the set $\mathcal{P}(X)$ of atoms.) A formula $\phi \in \text{Prop}(\Lambda(\mathcal{P}(X)))$ is *one-step \mathcal{L} -derivable*, denoted $\vdash_{\mathcal{L}}^1 \phi$, if ϕ is propositionally entailed by the set $\{\psi\tau \mid \tau : P \rightarrow \mathcal{P}(X), \psi \in \mathcal{A}\}$. A set $\Phi \subseteq \text{Prop}(\Lambda(\mathcal{P}(X)))$ is *one-step \mathcal{L} -consistent* if there do not exist formulas $\phi_1, \dots, \phi_n \in \Phi$ such that $\vdash_{\mathcal{L}}^1 \neg(\phi_1 \wedge \dots \wedge \phi_n)$. Dually, the *one-step semantics* $\llbracket \phi \rrbracket_X^1 \subseteq TX$ of a formula $\phi \in \text{Prop}(\Lambda(\mathcal{P}(X)))$ is defined inductively by $\llbracket L(A_1, \dots, A_n) \rrbracket_X^1 = \llbracket L \rrbracket_X(A_1, \dots, A_n)$ for $A_1, \dots, A_n \subseteq X$. A set $\Phi \subseteq \text{Prop}(\Lambda(\mathcal{P}(X)))$ is *one-step satisfiable* if $\bigcap_{\phi \in \Phi} \llbracket \phi \rrbracket_X^1 \neq \emptyset$. We say that \mathcal{L} (or Λ) is *separating* if $t \in TX$ is uniquely determined by the set $\{\phi \in \Lambda(\mathcal{P}(X)) \mid t \in \llbracket \phi \rrbracket_X^1\}$. We call \mathcal{L} (or \mathcal{A}) *one-step sound* if every one-step derivable formula $\phi \in \text{Prop}(\Lambda(\mathcal{P}(X)))$ is one-step valid, i.e. $\llbracket \phi \rrbracket_X^1 = X$.

Henceforth, we assume that \mathcal{L} is one-step sound, so that every T -coalgebra satisfies the rank-1 axioms; in the absence of frame conditions ($\Theta = \emptyset$), this means in particular that every T -coalgebra

is an \mathcal{L} -frame. The above notions of one-step satisfiability and one-step consistency are the main concepts employed in the proof of strong completeness in the following section.

Given a structure for Λ over T , every set \mathcal{B} of rank-1 axioms over Λ defines a subfunctor $T_{\mathcal{B}}$ of \mathcal{B} with $T_{\mathcal{B}}(X) = \bigcap \{ \llbracket \phi \tau \rrbracket_X^1 \mid \phi \in \mathcal{B}, \tau : P \rightarrow \mathcal{P}(X) \} \subseteq TX$. This functor induces a structure for which \mathcal{B} is one-step sound.

Example 1.9. The additional rank-1 axioms of Expl. 1.3.2 induce subfunctors $\mathcal{S}_{\mathcal{B}}$ of the functor \mathcal{S} of Expl. 1.5.2. E.g. we have

$$\mathcal{S}_{\{ID\}}X = \{f \in \mathcal{S}(X) \mid \forall A \subseteq X. f(A) \subseteq A\}$$

$$\mathcal{S}_{\{ID,DIS\}}X = \{f \in \mathcal{S}(X) \mid \forall A, B \subseteq X. f(A) \subseteq A \wedge f(A \cup B) \subseteq f(A) \cup f(B)\}$$

$$\mathcal{S}_{\{ID,DIS,CM\}}X = \{f \in \mathcal{S}(X) \mid \forall A, B \subseteq X. f(A) \subseteq A \wedge (f(B) \subseteq A \Rightarrow f(A) \cap B \subseteq f(B))\}$$

(it is an amusing exercise to verify the last claim).

2. Strong Completeness Via Quasi-Canonical Models

We wish to establish strong completeness of \mathcal{L} by defining a suitable T -coalgebra structure ζ on the set S of maximally \mathcal{L} -consistent subsets of $\mathcal{F}(\Lambda)$, equipped with the standard valuation $V(p) = \{\Gamma \in S \mid p \in \Gamma\}$. The crucial property required is that ζ be *coherent*, i.e.

$$\zeta(\Gamma) \in \llbracket L \rrbracket(\hat{\phi}_1, \dots, \hat{\phi}_n) \iff L(\phi_1, \dots, \phi_n) \in \Gamma,$$

where $\hat{\phi} = \{\Delta \in S \mid \phi \in \Delta\}$, for $L \in \Lambda$ n -ary, $\Gamma \in S$, and $\phi_1, \dots, \phi_n \in \mathcal{F}(\Lambda)$, as this allows proving, by a simple induction over the structure of formulas,

Lemma 2.1 (Truth lemma). *If ζ is coherent, then for all formulas ϕ , $\Gamma \models_{(S,\zeta,V)} \phi$ iff $\phi \in \Gamma$.*

We define a *quasi-canonical model* to be a model (S, ζ, V) with ζ coherent; the term quasi-canonical serves to emphasise that the coherence condition does not determine the transition structure ζ uniquely. By the truth lemma, quasi-canonical models for \mathcal{L} are \mathcal{L} -models, i.e. satisfy all substitution instances of the frame conditions. The first question is now under which circumstances quasi-canonical models exist; we proceed to establish a widely applicable criterion. This criterion has two main aspects: a *local* form of strong completeness involving only finite sets, and a preservation condition on the functor enabling passage from finite sets to certain infinite sets. We begin with the latter part:

Definition 2.2. A *surjective ω -cochain (of finite sets)* is a sequence $(X_n)_{n \in \mathbb{N}}$ of (finite) sets equipped with surjective functions $p_n : X_{n+1} \rightarrow X_n$ called *projections*. The *inverse limit* $\varprojlim X_n$ of (X_n) is the set $\{(x_i) \in \prod_{i \in \mathbb{N}} X_i \mid \forall n. p_n(x_{n+1}) = x_n\}$ of *coherent families* (x_i) . The *limit projections* are the maps $\pi_i((x_n)_{n \in \mathbb{N}}) = x_i$, $i \in \mathbb{N}$; note that the π_i are surjective, i.e. every $x \in X_i$ can be extended to a coherent family. Since all set functors preserve surjections, (TX_n) is a surjective ω -cochain with projections Tp_n . The functor T *weakly preserves inverse limits of surjective ω -cochains of finite sets* if for every surjective ω -cochain (X_n) of finite sets, the canonical map $T(\varprojlim X_n) \rightarrow \varprojlim TX_n$ is surjective, i.e. every coherent family (t_n) in $\prod TX_n$ is *induced* by a (not necessarily unique) $t \in T(\varprojlim X_n)$ in the sense that $T\pi_n(t) = t_n$ for all n .

Example 2.3. Let A be a finite alphabet; then the sets A^n , $n \in \mathbb{N}$, form a surjective ω -cochain of finite sets with projections $p_n : A^{n+1} \rightarrow A^n$, $(a_1, \dots, a_{n+1}) \mapsto (a_1, \dots, a_n)$. The inverse limit $\varprojlim A^n$ is the set A^ω of infinite sequences over A . The covariant powerset functor \mathcal{P} preserves this inverse limit weakly: given a coherent family of subsets $B_n \subseteq A^n$, i.e. $p_n[B_{n+1}] = B_n$ for all n ,

we define the set $B \subseteq A^\omega$ as the set of all infinite sequences $(a_n)_{n \geq 1}$ such that $(a_1, \dots, a_n) \in B_n$ for all n ; it is easy to check that indeed B induces the B_n , i.e. $\pi_n[B] = B_n$. However, B is by no means uniquely determined by this property: Observe that B as just defined is a safety property. The intersection of B with any liveness property C , e.g. the set C of all infinite sequences containing infinitely many occurrences of a fixed letter in A , will also satisfy $\pi_n[B \cap C] = B_n$ for all n .

The second part of our criterion is an infinitary version of a local completeness property called one-step completeness, which has been used previously in *weak* completeness proofs [15, 17].

Definition 2.4. We say that \mathcal{L} is *strongly one-step complete over finite sets* if for finite X , every one-step consistent subset Φ of $\text{Prop}(\Lambda(\mathcal{P}(X)))$ is one-step satisfiable.

The difference with plain one-step completeness is that Φ above may be infinite. Consequently, strong and plain one-step completeness coincide in case the modal similarity type Λ is finite, since in this case, $\text{Prop}(\Lambda(\mathcal{P}(X)))$ is, for finite X , finite up to propositional equivalence. The announced strong completeness criterion is now the following.

Theorem 2.5. *If \mathcal{L} is strongly one-step complete over finite sets and separating, Λ is countable, and T weakly preserves inverse limits of surjective ω -cochains of finite sets, then \mathcal{L} has a quasi-canonical model.*

Proof sketch. The most natural argument is via the dual adjunction between sets and boolean algebras that associates to a set the boolean algebra of its subsets, and to a boolean algebra the set of its ultrafilters. For economy of presentation, we outline a direct proof instead: we prove that

- (*) every maximally one-step consistent $\Phi \subseteq \text{Prop}(\Lambda(\mathfrak{A}))$ is one-step satisfiable, where $\mathfrak{A} = \{\hat{\phi} \mid \phi \in \mathcal{F}(\Lambda)\} \subseteq \mathcal{P}(S)$.

The existence of the required coherent coalgebra structure ζ on S follows immediately, since the coherence requirement for $\zeta(\Gamma)$, $\Gamma \in S$, amounts to one-step satisfaction of a maximally one-step consistent subset of $\text{Prop}(\Lambda(\mathfrak{A}))$.

To prove (*), let $\Lambda = \{L_n \mid n \in \mathbb{N}\}$, let $P = \{p_n \mid n \in \mathbb{N}\}$, let \mathcal{F}_n denote the set of Λ -formulas of modal nesting depth at most n that employ only modal operators from $\Lambda_n = \{L_0, \dots, L_n\}$ and only the atomic propositions p_0, \dots, p_n , and let S_n be the set of maximally consistent subsets of \mathcal{F}_n . Then S is (isomorphic to) the inverse limit $\varprojlim S_n$, where the projections $S_{n+1} \rightarrow S_n$ and the limit projections $S \rightarrow S_n$ are just intersection with \mathcal{F}_n . As the sets S_n are finite, we obtain by strong one-step completeness $t_n \in TS_n$ such that $t_n \models_{S_n}^1 \Phi \cap \text{Prop}(\Lambda(\mathfrak{A}_n))$, where $\mathfrak{A}_n = \{\hat{\phi} \cap S_n \mid \phi \in \mathcal{F}_n\}$. By separation, $(t_n)_{n \in \mathbb{N}}$ is coherent, and hence is induced by some $t \in TS$ by weak preservation of inverse limits; then, $t \models_S^1 \Phi$. ■

Together with the Lindenbaum Lemma we obtain strong completeness as a corollary.

Corollary 2.6. *Under the conditions of Thm. 2.5, \mathcal{L} is strongly complete for \mathcal{L} -models.*

Both Thm. 2.5 and Cor. 2.6 do apply to the case that \mathcal{L} has frame conditions. When \mathcal{L} is of rank 1 (i.e. $\Theta = \emptyset$), Cor. 2.6 implies that \mathcal{L} is strongly complete for (models based on) \mathcal{L} -frames. In the presence of frame conditions, the underlying frame of an \mathcal{L} -model need not be an \mathcal{L} -frame, so that the question arises whether \mathcal{L} is also strongly complete for \mathcal{L} -frames. In applications, positive answers to this question, usually referred to as the canonicity problem, typically rely on a judicious choice of quasi-canonical model to ensure that the latter is an \mathcal{L} -frame, often the largest quasi-canonical model under some ordering on TS . Detailed examples are given in Sec. 3.

Remark 2.7. It is shown in [13] that T admits a strongly complete modal logic if T weakly preserves (arbitrary) inverse limits and preserves finite sets. The essential contribution of the above

result is to remove the latter restriction, which fails in important examples. Moreover, the observation that we need only consider *surjective* ω -cochains is relevant in some applications, see below.

Remark 2.8. A last point that needs clearing up is whether strong completeness of coalgebraic modal logics can be established by some more general method than quasi-canonical models of the quite specific shape used here. The answer is negative, at least in the case of rank-1 logics \mathcal{L} : it has been shown in [12] that every such \mathcal{L} admits models which consist of the maximally *satisfiable* sets of formulas and obey the truth lemma. Under strong completeness, such models are quasi-canonical.

This seems to contradict the fact that some canonical model constructions in the literature, notably the canonical Kripke models for graded modal logics [8, 6], employ state spaces which have multiple copies of maximally consistent sets. The above argument indicates that such logics fail to be coalgebraic, and indeed this is the case for GML with Kripke semantics. As mentioned above, GML has an alternative coalgebraic semantics over multigraphs, and we show below that this semantics does admit quasi-canonical models in our sense.

3. Examples

We now show how the generic results of the previous section can be applied to obtain canonical models and associated strong completeness and compactness theorems for a large variety of structurally different modal logics. We have included some negative examples where canonical models necessarily fail to exist due to non-compactness, and we analyse which conditions of Thm. 2.5 fail in each case. We emphasise that in the positive examples, the verification of said conditions is entirely stereotypical. Weak preservation of inverse limits of surjective ω -cochains usually holds without the finiteness assumption, which is therefore typically omitted.

Example 3.1 (Strong completeness of Kripke semantics for K). Recall from Expl. 1.5.1 that Kripke frames are coalgebras for the powerset functor $TX = \mathcal{P}(X)$. Strong completeness of K with respect to Kripke semantics is, of course, well known. We briefly illustrate how this can be derived from our coalgebraic treatment. To see that K is strongly one-step complete over finite sets X , let $\Phi \subseteq \text{Prop}(\Lambda_K(\mathcal{P}(X)))$ be maximally one-step consistent. It is easy to check that $\{x \in X \mid \diamond\{x\} \in \Phi\}$ satisfies Φ . To prove that the powerset functor weakly preserves inverse limits, let (X_n) be an ω -cochain, and let $(A_n \in \mathcal{P}(X_n))$ be a coherent family. Then (A_n) is itself a cochain, and the set $A = \varprojlim A_n \subseteq \varprojlim X_n$ induces (A_n) (w.r.t. the subset ordering on $\mathcal{P}(X)$). Separation is clear. By Thm. 2.5, there exists a quasi-canonical Kripke model for all normal modal logics. In particular, the standard canonical model [4] is quasi-canonical; it witnesses strong completeness (w.r.t. frames) of all canonical logics such as $K4$, $S4$, $S5$.

Example 3.2 (Failure of strong completeness of K over finitely branching models). As seen in Expl. 1.5.1, finitely branching Kripke frames are coalgebras for the finite powerset functor \mathcal{P}_ω . It is clear that quasi-canonical models fail to exist in this case, as compactness fails over finitely branching frames: one can easily construct formulas ϕ_n that force a state to have at least n different successors. The obstacle to the application of Thm. 2.5 is that the finite powerset functor fails to preserve inverse limits weakly, as the inverse limit of an ω -cochain of finite sets may fail to be finite.

Example 3.3 (Conditional logic). Recall from Expl. 1.5.2 that the conditional logic CK is interpreted over the functor $\mathcal{S}(X) = \mathcal{P}(X) \rightarrow \mathcal{P}(X)$. To prove strong one-step completeness over finite sets X , let $\Phi \subseteq \text{Prop}(\Lambda_{\text{CL}}(\mathcal{P}(X)))$ be maximally one-step consistent. Define $f : \mathcal{P}(X) \rightarrow \mathcal{P}(X)$ by $f(A) = \bigcap \{B \subseteq X \mid A \Rightarrow B \in \Phi\}$; it is mechanical to check that $f \models^1 \Phi$. To see that \mathcal{S} weakly preserves inverse limits, let (X_n) be a surjective ω -cochain, let $X = \varprojlim X_n$, and let $(f_n \in \mathcal{S}(X_n))$

be coherent. Define $f : \mathcal{P}(X) \rightarrow \mathcal{P}(X)$ by letting $(x_n) \in f(A)$ for a coherent family $(x_n) \in X$ iff whenever $A = \pi_n^{-1}[B]$ for some n and some $B \subseteq X_n$, then $x_n \in f_n(B)$. Using surjectivity of the projections of (X_n) , it is straightforward to prove that f induces (f_n) . Finally, separation is clear. By Thm. 2.5, it follows that the conditional logic CK has a quasi-canonical model, and hence that CK is strongly complete for conditional frames. In the case of the additional rank-1 axioms mentioned in Expl. 1.3.2 and the corresponding subfunctors of \mathcal{S} described in Expl. 1.9, the situation is as follows.

Identity: The functor $\mathcal{S}_{\{ID\}}$ weakly preserves inverse limits of surjective ω -cochains. In the notation above, put $(x_n) \in f(A)$ iff the condition above holds and $(x_n) \in A$.

Identity and disjunction: The functor $\mathcal{S}_{\{ID,DIS\}}$ weakly preserves inverse limits of surjective ω -cochains: put $(x_n) \in f(A)$ iff $(x_n) \in A$ and whenever $(x_n) \in \pi_m^{-1}B \subseteq A$, then $x_m \in f_m(B)$.

System C: It is open whether the the functor $\mathcal{S}_{\{ID,DIS,CM\}}$ weakly preserves inverse limits of surjective ω -cochains, and whether System C is strongly complete over conditional frames.

Indeed it appears to be an open problem to find *any* semantics for which System C is strongly complete, other than the generalised neighbourhood semantics as described e.g. in [18], which is strongly complete for very general reasons but provides little in the way of actual semantic information. The classical preference semantics according to Lewis is only known to be weakly complete [3]. Friedman and Halpern [9] do silently prove strong completeness of System C w.r.t. plausibility measures; however, on close inspection the latter turn out to be essentially equivalent to the above-mentioned generalised neighbourhood semantics. Moreover, Segerberg [19] proves strong completeness for a whole range of conditional logics over *general* conditional frames, where, in analogy to corresponding terminology for Kripke frames, a general conditional frame is equipped with a distinguished set of *admissible propositions* limiting both the range of valuations and the domain of selection functions. In contrast, our method yields full conditional frames in which the frame conditions hold for *any* valuation of the propositional variables. While in the case of CK and its extension by ID alone, these models differ from Segerberg's only in that they insert default values for the selection function on non-admissible propositions, the canonical model for the extension of CK by $\{ID, DIS\}$ has non-trivial structure on non-admissible propositions, and we believe that our strong completeness result for this logic is genuinely new.

Example 3.4 (Strong completeness of GML over multigraphs). Recall from Expl. 1.5.3 that graded modal logic (GML) has a coalgebraic semantics in terms of the multiset functor \mathcal{B}_∞ . To prove strong one-step completeness over finite sets X , let $\Phi \subseteq \text{Prop}(\Lambda_{GML}(\mathcal{P}(X)))$ be maximally one-step consistent. We define $B \in \mathcal{B}_\infty(X)$ by $B(A) \geq n \iff \geq n. A \in \Phi$; it is easy to check that B is well-defined and additive. To prove weak preservation of inverse limits, let (X_n) be an ω -cochain, let $X = \varprojlim X_n$, and let $(B_n \in \mathcal{B}_\infty(X_n))$ be coherent. Then define $B \in \mathcal{B}_\infty(X)$ pointwise by

$$B((x_n)) = \min_{n \in \mathbb{N}} B_n(x_n),$$

noting that the sequence $(B_n(x_n))$ is decreasing by coherence. A straightforward computation shows that B induces (B_n) . Separation is clear.

By the above and Thm. 2.5, all extensions of GML have quasi-canonical multigraph models. While the technical core of the construction is implicit in the work of Fine [8] and de Caro [6], these authors were yet unaware of multigraph semantics, and hence our result that *GML is strongly complete over multigraphs* has not been obtained previously.

The standard frame conditions for reflexivity, symmetry, and transitivity (Expls. 1.5.3 and 1.7. 3) and arbitrary combinations thereof are easily seen to be satisfied in the quasi-canonical model constructed above. We point out that this contrasts with Kripke semantics in the case of the

graded version of $S4$, i.e. GML extended with the reflexivity and transitivity axioms of Expl. 1.5.3: as shown in [7], the complete axiomatisation of graded modal logic over transitive reflexive Kripke frames includes two rather strange combinatorial artefacts, which by the above disappear in the multigraph semantics. The reason for the divergence (which we regard as an argument in favour of multigraph semantics) is that, while in many cases multigraph models are easily transformed into equivalent Kripke models by just making copies of states, no such translation exists in the transitive reflexive case (transitivity alone is unproblematic).

Observe moreover that the above extends straightforwardly to decription logics $\mathcal{ALCQ}(\mathcal{R})$ with qualified number restrictions and a role hierarchy \mathcal{R} where roles may be distinguished as, in any combination, transitive, reflexive, or symmetric. As shown in [10, 11], $\mathcal{ALCQ}(\mathcal{R})$ is undecidable for many \mathcal{R} , even when only transitive roles are considered. For undecidable logics, completeness is in some sense the ‘next best thing’, as it guarantees if not recursiveness then at least recursive enumerability of all valid formulas, and hence enables automatic reasoning. Essentially, our results show that the natural axiomatisation of $\mathcal{ALCQ}(\mathcal{R})$ with transitive, symmetric and reflexive roles is strongly complete over multigraphs, a result which fails for the standard Kripke semantics.

Example 3.5 (Failure of strong completeness of image-finite GML). Similarly to the case of image-finite Kripke frames, one can model an image-finite version of graded modal logic coalgebraically by exchanging the functor \mathcal{B}_∞ for the *finite multiset functor* \mathcal{B} , where $\mathcal{B}(X)$ consists of all maps $X \rightarrow \mathbb{N}$ with finite support. Of course, the resulting logic is non-compact and hence fails to admit a canonical model. This is witnessed not only by the same family of formulas as in the case of image-finite Kripke semantics, which targets finiteness of the number of different successors, but also by the set of formulas $\{\geq n. a \mid n \in \mathbb{N}\}$, which targets finiteness of multiplicities. Analysing the conditions of Thm. 2.5, we detect two violations: not only does weak preservation of inverse limits fail, but there is also no way to find an axiomatisation which is strongly one-step complete over finite sets (again, consider sets $\{\geq n. \{x\} \mid n \in \mathbb{N}\}$).

Strong completeness of image-finite GML can be recovered by slight adjustments to the syntax and semantics. We formulate a more general approach, as follows.

Example 3.6 (Strong completeness of the logic of additive measures). We fix an at most countable commutative monoid M (e.g. $M = \mathbb{N}$). We think of the elements of M as describing the measure of a set of elements. To ensure compactness, we have to allow some sets to have undefined measure. That is, we work with coalgebras for the endofunctor T_M defined by

$$T_M(X) = \{(\mathfrak{A}, \mu) \mid \mathfrak{A} \subseteq \mathcal{P}(X) \text{ closed under disjoint unions, } \mu : \mathfrak{A} \rightarrow M \text{ additive}\}$$

The modal logic of additive M -valued measures is given by the similarity type $\Lambda_M = \{E_m \mid m \in M\}$ where $E_m \phi$ expresses that ϕ has measure m , i.e.

$$\llbracket E_m \rrbracket_X B = \{(\mathfrak{A}, \mu) \in T_M(X) \mid B \in \mathfrak{A}, \mu(B) = m\}.$$

Λ_M is clearly separating. The logic is axiomatised by the following two axioms:

$$E_m a \rightarrow \neg E_n a \quad (n \neq m) \quad \text{and} \quad E_m(a \wedge b) \wedge E_n(a \wedge \neg b) \rightarrow E_{m+n} a.$$

These axioms are strongly one-step complete over finite sets X : if $\Phi \subseteq \text{Prop}(\Lambda_M(\mathcal{P}(X)))$ is maximally one-step consistent, then $(\mathfrak{A}, \mu) \models^1 \Phi$ where $A \in \mathfrak{A}$ iff $E_m A \in \Phi$ for some necessarily unique m , in which case $\mu(A) = m$. Moreover, T_M weakly preserves inverse limits $X = \varprojlim X_n$, with finite X_n : a coherent family $((\mathfrak{A}_n, \mu_n) \in T_M(X_n))$ is induced by $(\mathfrak{A}, \mu) \in T_M(X)$, where $\mathfrak{A} = \{\pi_n^{-1}[B] \mid n \in \mathbb{N}, B \in \mathfrak{A}_n\}$ and $\mu(\pi_n^{-1}[B]) = \mu_n(B)$ is easily seen to be well-defined and additive. Theorem 2.5 now guarantees existence of quasi-canonical models. A simple example is $M = \mathbb{Z}/2\mathbb{Z}$, which induces a logic of even and odd.

For the case $M = \mathbb{N}$, we obtain a variant of graded modal logic with finite multiplicities, where we code $\geq k.\phi$ as $\neg \bigvee_{0 \leq i < k} E_i \phi$. However, it may still be the case that a state has a family of successor sets of unbounded measure, so that undefinedness of the measure of the entire state space just hides an occurrence of infinity. This defect is repaired by insisting that the measure of the whole state space is finite at the expense of disallowing the modal operator E_0 in the language, as follows.

Example 3.7 (Strong completeness of finitely branching GML^-). To force the entire state space to have finite measure, we additionally introduce a *measurability* operator E , interpreted by $\llbracket E \rrbracket B = \{(\mathfrak{A}, \mu) \mid B \in \mathfrak{A}\}$, and impose obvious axioms guaranteeing that measures on X are defined on boolean subalgebras of $\mathcal{P}(X)$, in particular $E\top$ (i.e. $\mu(X)$ is finite), and $E_n a \rightarrow E a$. In order to achieve compactness, we now leave a bolt hole on the syntactical side and exclude the operator E_0 . In other words, the syntax of GML^- is given by the similarity type $\Lambda_{\text{GML}}^- = \{E\} \cup \{E_n \mid n > 0\}$, and we interpret GML^- over coalgebras for the functor \mathcal{B}_M defined by

$$\mathcal{B}_M(X) = \{(\mathfrak{A}, \mu) \mid \mathfrak{A} \text{ boolean subalgebra of } \mathcal{P}(X), \mu : \mathfrak{A} \rightarrow \mathbb{N} \text{ additive}\}.$$

Separation is clear. The axiomatisation of GML^- is given by the axiomatisation of the modal logic of additive measures, the above-mentioned axioms on E , and the additional axiom

$$E_n a \wedge E b \rightarrow E_n(a \wedge b) \vee E_n(a \wedge \neg b) \vee \bigvee_{0 < k < n} (E_k(a \wedge b) \wedge E_{n-k}(a \wedge \neg b))$$

which compensates for the absence of E_0 . Strong one-step completeness over finite sets and weak preservation of inverse limits is shown analogously as in Expl. 3.6, so that we obtain a *strongly complete finitely branching graded modal logic* GML^- . The tradeoff is that the operator $\geq k.\phi$ is no longer expressible as $\neg \bigvee_{0 \leq i < k} E_i \phi$ in GML^- which only allows to formulate the implication $\geq 1.\phi \rightarrow \geq n.\phi$.

Example 3.8 (Failure of strong completeness for PML over finitely supported probability distributions). Like image-finite graded modal logic, probabilistic modal logic as introduced in Expl. 1.5.4 fails to be compact, and violates the conditions of Thm. 2.5 on two counts, namely weak preservation of inverse limits and strong one-step completeness over finite sets. The first issue is related to image-finiteness, while the second is rooted in the structure of the real numbers: e.g. the set $\{L_{1/2-1/n} a \mid n \in \mathbb{N}\} \cup \{\neg L_{1/2} a\}$ is finitely satisfiable but not satisfiable.

Example 3.9 (Strong completeness of the logic of exact probabilities). In order to remove the above-mentioned failure of compactness, we consider the fragment of probabilistic modal logic containing only operators E_p stating that a given event has probability exactly p . (This is, of course, less expressive than the operators L_p but still allows reasonable statements such as that rolling a six on a die happens with probability $1/6$.) Moreover, we require probabilities to be rational and allow probabilities to be undefined, thus following the additive measures approach as outlined above, where we consider a subfunctor of $T_{\mathbb{Q}}$ defined by the requirement that the whole set has measure 1. However, we are able to impose stronger conditions on the domain $\mathfrak{A} \subseteq \mathcal{P}(X)$ of a probability measure P on X : we require that $X \in \mathfrak{A}$ and that $A, B \in \mathfrak{A}$, $B \subseteq A$ imply $A - B \in \mathfrak{A}$, which is reflected in the additional axioms $E_1 \top$ and $E_p a \wedge E_q(a \wedge b) \rightarrow E_{p-q}(a \wedge \neg b)$. It is natural that we cannot force closure under intersection, as there is in general no way to infer the exact probability of $A \cap B$ from the probabilities of A and B . Along the same lines as above, we now obtain quasi-canonical models, and hence strong completeness and compactness, of the arising modal logic of exact probabilities.

4. Conclusion

We have laid out a systematic method of proving existence of canonical models in a generic semantic framework encompassing a wide range of structurally different modal logics. We have shown how this method turns the construction of canonical models into an entirely mechanical exercise where applicable, and points the way to obtaining compact fragments of non-compact logics. As example applications, we have reproved a number of known strong completeness result and established several new results of this kind; specifically, the latter includes strong completeness of the following logics.

- The modal logic of exact probabilities, with operators E_p ‘with probability exactly p ’.
- Graded modal logic over transitive reflexive multigraphs, i.e. the natural graded version of $S4$, and more generally description logic with role hierarchies including transitive, reflexive, and symmetric roles and qualified number restrictions also on non-simple (e.g. transitive) roles.
- The conditional logic $CK + \{ID, DIS\}$, i.e. with the standard axioms of identity and disjunction, interpreted over conditional frames.

A number of interesting open problems remain, e.g. to find further strongly complete variants of probabilistic modal logic or to establish strong completeness of the full set of standard axioms of default logic, Burgess’ System C [3], over the corresponding class of conditional frames.

References

- [1] F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, eds. *The Description Logic Handbook*. Cambridge University Press, 2003.
- [2] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.
- [3] J. Burgess. Quick completeness proofs for some logics of conditionals. *Notre Dame J. Formal Logic*, 22:76–84, 1981.
- [4] B. Chellas. *Modal Logic*. Cambridge University Press, 1980.
- [5] G. D’Agostino and A. Visser. Finality regained: A coalgebraic study of Scott-sets and multisets. *Arch. Math. Logic*, 41:267–298, 2002.
- [6] F. De Caro. Graded modalities II. *Stud. Log.*, 47:1–10, 1988.
- [7] M. Fattorosi-Barnaba and C. Cerrato. Graded modalities III. *Stud. Log.*, 47:99–110, 1988.
- [8] K. Fine. In so many possible worlds. *Notre Dame J. Formal Logic*, 13:516–520, 1972.
- [9] N. Friedman and J. Y. Halpern. Plausibility measures and default reasoning. *J. ACM*, 48(4):648–685, 2001.
- [10] I. Horrocks, U. Sattler, and S. Tobies. Practical reasoning for expressive description logics. In *Logic for Programming and Automated Reasoning, LPAR 99*, vol. 1705 of *LNCS*, pp. 161–180. Springer, 1999.
- [11] Y. Kazakov, U. Sattler, and E. Zolin. How many legs do I have? Non-simple roles in number restrictions revisited. In *Logic for Programming, Artificial Intelligence, and Reasoning, LPAR 2007*, vol. 4790 of *LNCS*, pp. 303–317. Springer, 2007.
- [12] A. Kurz and D. Pattinson. Coalgebraic modal logic of finite rank. *Math. Struct. Comput. Sci.*, 15:453–473, 2005.
- [13] A. Kurz and J. Rosický. Strongly complete logics for coalgebras. Draft, July 2006.
- [14] K. Larsen and A. Skou. Bisimulation through probabilistic testing. *Inform. Comput.*, 94:1–28, 1991.
- [15] D. Pattinson. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoret. Comput. Sci.*, 309:177–193, 2003.
- [16] M. Pauly. A modal logic for coalitional power in games. *J. Logic Comput.*, 12:149–166, 2002.
- [17] L. Schröder. A finite model construction for coalgebraic modal logic. *J. Log. Algebr. Prog.*, 73:97–110, 2007.
- [18] L. Schröder and D. Pattinson. Rank-1 modal logics are coalgebraic. In *Theoretical Aspects of Computer Science, STACS 07*, vol. 4393 of *LNCS*, pp. 573–585. Springer, 2007. Full version to appear in *J. Log. Comput.*
- [19] K. Segerberg. Notes on conditional logic. *Stud. Log.*, 48:157–168, 1989.