Centrum voor Wiskunde en Informatica

REPORT*RAPPORT*

Modal model theory

M. de Rijke

Computer Science/Department of Software Technology

CS-R9517 1995

Report CS-R9517
ISSN 0169-118X

CWI
P.O. Box 94079
1090 GB Amsterdam
The Netherlands

Modal Model Theory

Maarten de Rijke

CWI
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands
mdr@cwi.nl

Abstract

This paper contributes to the model theory of modal logic using bisimulations as the fundamental tool. A uniform presentation is given of modal analogues of well-known definability and preservation results from first-order logic. These results include algebraic characterizations of modal equivalence, and of the modally definable classes of models; the preservation results concern preservation of modal formulas under submodels, unions of chains, and homomorphisms.

AMS Subject Classification (1991): 03B45, 03C95, 03C13
CR Subject Classification (1991): F.4.0, F.4.1, I.2.3, I.2.4.
Keywords & Phrases: Modal logic, abstract model theory, finite structures, expressive power, definability, preservation results, bisimulations.

1. Introduction

The guiding theme of this paper is the 'equation'

$$\frac{\cong^p}{\text{first-order logic}} = \frac{\leftrightarroweq}{\text{modal logic}}.$$

That is: bisimulations are to modal logic, what partial isomorphisms are for first-order logic. We substantiate this claim by establishing key results from first-order logic (and beyond) for modal logic, using bisimulations instead of partial isomorphism.

Specifically, after some background material has been presented in §2, §3 introduces basic bisimulations. In §4 these are linked to basic modal languages, resulting in an analogue of the Keisler-Shelah Theorem from first-order logic, as well as modal analogues of Karp's Theorem and the Scott Isomorphism Theorem from $\mathcal{L}_{\infty\omega}$ and $\mathcal{L}_{\omega_1\omega}$ respectively in §5. Building on those results §6 supplies a series of definability results. Then §7 pushes the idea that bisimulations are a fundamental tool in modal model theory even further by using them to establish modal analogues of three well-known preservation results from first-order logic: Łoś's Theorem, the Chang-Łoś-Suszko Theorem, and Lyndon's Theorem. The final section, §8, is devoted to extensions, questions and suggestions for further work.

2. Basic modal languages

As basic modal formulas live inside fragments of classical languages, we need a few notions from classical logic before we can specify our modal languages. We use τ, τ_1, ... to denote (relational) vocabularies of predicate logical languages. For τ a classical vocabulary, $\mathrm{Str}[\tau]$ denotes the class of τ-structures. For $\mathfrak{A} \in \mathrm{Str}[\tau]$ and P in τ, $P^{\mathfrak{A}}$ denotes the extension of P in $\mathfrak{A}$.

We assume that modal languages have modal operators $\#$ equipped with *patterns* $\delta_{\#}$ describing the semantics of $\#$ by means of a formula in classical logic. The bulk of this paper deals with basic modal languages.

Definition 2.1 For τ a classical vocabulary with unary predicate symbols, the *basic modal language over* τ is the *finitary* modal language $\mathcal{BML}_{\omega\omega}(\tau)$ having proposition letters p_0, p_1, ... corresponding to the unary predicate symbols in τ, and also having n-ary modal operators $\#$ with patterns

$$\delta_{\#} = \lambda x.\, \exists x_1 \ldots \exists x_n \Big(Rxx_1 \ldots x_n \wedge p_1(x_1) \wedge \ldots \wedge p_n(x_n) \Big),$$

for every $(n+1)$-ary relation symbol R in τ. In addition $\mathcal{BML}_{\omega\omega}(\tau)$ has the usual Boolean connectives, and constants $\bot$ and $\top$.

We also need infinitary basic modal languages. Let κ be a regular cardinal. The basic infinitary modal language $\mathcal{BML}_{\kappa\omega}$ has proposition letters, modal operators, connectives and constants as in $\mathcal{BML}_{\omega\omega}(\tau)$, but it also has conjunctions $\bigwedge$ and disjunctions $\bigvee$ over sets of formulas of cardinality less than κ. We write $\mathcal{BML}_{\infty\omega}(\tau) = \bigcup_{\kappa} \mathcal{BML}_{\kappa\omega}(\tau)$, and for λ singular, $\mathcal{BML}_{\lambda\omega}(\tau) = \bigcup_{\kappa<\lambda} \mathcal{BML}_{\kappa\omega}(\tau)$.

Definition 2.2 Define the *rank* of a modal formula, $\mathrm{rank}(\phi)$ as follows:

$$\begin{aligned} \mathrm{rank}(p) &= 0 \\ \mathrm{rank}(\neg\phi) &= \mathrm{rank}(\phi) \\ \mathrm{rank}(\textstyle\bigvee \Phi) &= \sup(\{\, \mathrm{rank}(\phi) : \phi \in \Phi \,\}) \\ \mathrm{rank}(\#(\phi_1, \ldots, \phi_n)) &= 1 + \max\{\, \mathrm{rank}(\phi_i) \mid 1 \leq i \leq n \}. \end{aligned}$$

Basic modal languages are interpreted on τ-structures of the form $\mathfrak{A} = (A, R_1, R_2, \ldots, P_1, P_2, \ldots)$, where P_1, P_2, ... interpret the proposition letters of the modal language. We will let valuations V take care of proposition letters; thus we will write $(W, R_1, R_2, \ldots, V)$, where $V(p_i) = P_i$.

Using the patterns of a modal logic a translation ST can be defined that takes modal formulas to formulas in the classical language in which those patterns live: fix an individual variable x, and put

$$\begin{aligned} ST_x(p) &= Px \\ ST_x(\bot) &= (x \neq x) \\ ST_x(\top) &= (x = x) \\ ST_x(\neg\phi) &= \neg ST_x(\phi) \\ ST_x(\phi \wedge \psi) &= ST_x(\phi) \wedge ST_x(\psi) \\ ST_x(\#(\phi_1, \ldots, \phi_n)) &= \exists x_1 \ldots \exists x_n \Big(Rxx_1 \ldots x_n \wedge ST_{x_1}(\phi_1) \wedge \ldots \wedge ST_{x_n}(\phi_n) \Big), \end{aligned}$$

where the semantics of $\#$ is based on R, and $ST_{x_i}(\phi_i)$ is the standard translation of ϕ_i with x_i as its free variable. Then, for all basic modal formulas ϕ: $(A, R_1, R_2, \ldots, V), a \models \phi$

iff $(A, R_1, R_2, \ldots, V) \models ST_x(\phi)[a]$. This equivalence allows us to freely move back and forth between modal formulas and certain classical formulas. Also, as basic modal formulas are equivalent to their (classical) ST-translations, they inherit important properties of classical logic; for $\mathcal{BML}_{\omega\omega}$-formulas this means that they enjoy the usual compactness and Löwenheim-Skolem properties (when interpreted on models).

For the remainder of this paper we will work with so-called *pointed models* $(\mathfrak{A}, a)$, where $\mathfrak{A}$ is a τ-structure as explained above, and a is a point in $\mathfrak{A}$; a is called the *distinguished point* of $(\mathfrak{A}, a)$. There are several reasons to work with pointed models rather than plan τ-structures. Pointed models simply are the semantic units at which modal formulas are evaluated; and second, most of the results and proofs below can be stated in a simpler form when pointed models are used.

3. Bisimulations

Definition 3.1 For τ a classical vocabulary and $\mathfrak{A}$, $\mathfrak{B} \in \mathrm{Str}[\tau]$, we say that $(\mathfrak{A}, a)$, $(\mathfrak{B}, b)$ are *τ-bisimilar* ($(\mathfrak{A}, a) \leftrightarrows_\tau (\mathfrak{B}, b)$) if there exists a non-empty relation Z between the elements of $\mathfrak{A}$ and $\mathfrak{B}$ (called a *τ-bisimulation*, and written $Z : (\mathfrak{A}, a) \leftrightarrows_\tau (\mathfrak{B}, b)$) such that

1. Z links the distinguished points of $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$: Zab,
2. for all unary predicate symbols P in τ, Za_0b_0 implies $a_0 \in P^{\mathfrak{A}}$ iff $b_0 \in P^{\mathfrak{B}}$,
3. if Za_0b_0, $a_1, \ldots, a_n \in \mathfrak{A}$ and $(a_0, a_1, \ldots, a_n) \in R^{\mathfrak{A}}$, then there are $b_1, \ldots, b_n \in \mathfrak{B}$ such that $(b_0, b_1, \ldots, b_n) \in R^{\mathfrak{B}}$ and Za_ib_i, where $1 \le i \le n$ and R is an $(n+1)$-ary relation symbol in τ (*forth condition*),
4. if Za_0b_0, $b_1, \ldots, b_n \in \mathfrak{B}$ and $(b_0, b_1, \ldots, b_n) \in R^{\mathfrak{B}}$, then there are $a_1, \ldots, a_n \in \mathfrak{A}$ such that $(a_0, a_1, \ldots, a_n) \in R^{\mathfrak{A}}$ and Za_ib_i, where $1 \le i \le n$ and R is an $(n+1)$-ary relation symbol in τ (*back condition*).

It is easily verified that isomorphism implies bisimilarity, and that the relational composition and union of bisimulations is again a bisimulation; moreover, bisimilarity is an equivalence relation on the class of all models.

Many familiar constructions on relational structures arise as special examples of bisimulations. For disjoint τ-structures $(\mathfrak{A}_i, a_i)$ $(i \in I)$ their *disjoint union* is the structure $\mathfrak{A}$ which has the union of the domains of $\mathfrak{A}_i$ as its domain, while $R^{\mathfrak{A}} = \bigcup_i R^{\mathfrak{A}_i}$. For each of the components $\mathfrak{A}_i$ there is a bisimulation $Z : (\mathfrak{A}_i, a_i) \leftrightarrows_\tau (\mathfrak{A}, a_i)$ defined by Zxy iff $x = y$.

$(\mathfrak{A}, a)$ is a *generated submodel* of $(\mathfrak{B}, b)$ whenever (i) $a = b$, (ii) the domain of $\mathfrak{A}$ is a subset of the domain of $\mathfrak{B}$, (iii) $R^{\mathfrak{A}}$ is simply the restriction of $R^{\mathfrak{B}}$ to $\mathfrak{A}$, and (iv) if $a_0 \in \mathfrak{A}$ and $R^{\mathfrak{A}}a_0b_1 \ldots b_n$, then $b_1, \ldots, b_n$ are in $\mathfrak{A}$. If $(\mathfrak{A}, a)$ is a generated submodel of $(\mathfrak{B}, b)$, there is a τ-bisimulation $Z : (\mathfrak{A}, a) \leftrightarrows_\tau (\mathfrak{B}, b)$ defined by Zxy iff $x = y$.

A mapping $f : (\mathfrak{A}, a) \to (\mathfrak{B}, b)$ is a *p-morphism* if (i) $f(a) = b$, (ii) it is a homomorphism for all $R \in \tau$, that is: $R^{\mathfrak{A}}aa_1 \ldots a_n$ implies $R^{\mathfrak{B}}f(a)f(a_1) \ldots f(a_n)$, and (iii) if $R^{\mathfrak{B}}f(a)b_1 \ldots b_n$ then there are $a_1, \ldots, a_n$ such that $R^{\mathfrak{A}}aa_1 \ldots a_n$ and $f(a_i) = b_i$. If $f : (\mathfrak{A}, a) \to (\mathfrak{B}, b)$ is a p-morphism, putting Zxy iff $f(x) = y$ defines a bisimulation $Z : (\mathfrak{A}, a) \leftrightarrows_\tau (\mathfrak{B}, b)$.

Just like partial isomorphisms in Abstract Model Theory, bisimulations too are naturally built up by means of approximations. Let $\mathfrak{A}$, $\mathfrak{B} \in \mathrm{Str}[\tau]$. We define a notion of *τ-bisimilarity up to n* by requiring that there exists a sequence of binary relations $Z_0, \ldots, Z_n$ between $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$ such that

1. $Z_n \subseteq \cdots \subseteq Z_0$ and $Z_n ab$,
2. for each i, if $Z_i xy$ then x and y agree on all unary predicates,
3. for $i + 1 \leq n$ the back-and-forth properties are satisfied relative to the indices:
   (a) if $Z_{i+1}xy$ and $R^{\mathfrak{A}}xx_1 \ldots x_m$, then for some $y_1, \ldots, y_m$ in $\mathfrak{B}$: $R^{\mathfrak{B}}yy_1 \ldots y_m$, and for all $j = 1, \ldots, m$: $Z_i x_j y_j$,
   (b) if $Z_{i+1}xy$ and $R^{\mathfrak{B}}yy_1 \ldots y_m$, then for some $x_1, \ldots, x_m$ in $\mathfrak{A}$: $R^{\mathfrak{A}}xx_1 \ldots x_m$, and for all $j = 1, \ldots, m$: $Z_i x_j y_j$.

Clearly, every τ-bisimulation gives rise to a relations of bisimularity up to n, for every n. But conversely, if there exist exist $Z_0, \ldots, Z_n, \ldots$ satisfying the above back-and-forth conditions, then $Z = \bigcap_i Z_i$ need not define a τ-bisimulation — see Example 5.4 below for a counterexample.

If, for some n, there is a bisimulation up to n between $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$, we write $(\mathfrak{A}, a) \leftrightarroweq^n_\tau (\mathfrak{B}, b)$, and say that $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$ are *τ-bisimilar up to n*.

We need two concepts for measuring certain aspects of models: in-degree and depth.

Definition 3.2 (In-degree) Let $(\mathfrak{A}, a)$ be a model, and $c \in \mathfrak{A}$. A *path from a through c* is any sequence of sequences $\vec{x}^1, \vec{x}^2, \ldots, \vec{x}^n$ such that (i) $x^1_0 = a$, (ii) for each i $(1 \leq i \leq n)$ there exists an R^i in the similarity type of $\mathfrak{A}$ such that $R^i\vec{x}^i$ holds in $\mathfrak{A}$, (iii) for each i $(1 < i \leq n)$ x^i_0 is in an argument in $\vec{x}^{i-1}$, that is: $x^i_0 \in \{x^{i-1}_1, x^{i-1}_2 \ldots\}$, and (iv) c is an argument in the final tuple $\vec{x}^n$, that is $c \in \{x^n_1, x^n_2 \ldots\}$.

The *in-degree of c in* $(\mathfrak{A}, a)$ is the number of paths from a through c.

The second notion we need measures the distance from a given element in a model to its distinguished point.

Definition 3.3 (Depth) Let $(\mathfrak{A}, a)$ be a $\tau-$structure; the *τ-hulls H^n_τ* around a are defined as follows

- $H^0_\tau(\mathfrak{A}, a) = \{a\}$,
- $H^{n+1}_\tau(\mathfrak{A}, a) = H^n_\tau(\mathfrak{A}, a) \cup \ \{b \text{ in } \mathfrak{A} \mid$ for some $R \in \tau$, $u \in H^n_\tau(\mathfrak{A}, a)$ and $v_1, \ldots, v_n$ in $\mathfrak{A}$: b is one of the v_i and $R^{\mathfrak{A}}uv_1 \ldots v_n\}$.

So, the n-hull H^κ_τ around a contains all elements in $\mathfrak{A}$ that can be reached from a in at most n relational steps.

For c in $(\mathfrak{A}, a)$, the *depth of c in* $(\mathfrak{A}, a)$ is the smallest n such that $c \in H^n_\tau(\mathfrak{A}, a)$, if such n exists. Otherwise the depth if c is ∞.

For $n \in \omega$, the model $(\mathfrak{A} \upharpoonright n, a)$ is the *restriction of* $(\mathfrak{A}, a)$ *to points of depth* n; it is defined as the submodel of $(\mathfrak{A}, a)$ whose domain is $H^n(\mathfrak{A}, a)$.

Proposition 3.4 *Let $(\mathfrak{A}, a)$, $(\mathfrak{B}, b)$ be two models such that every element has in-degree at most 1, and depth at most n. The following are equivalent:*

1. $(\mathfrak{A}, a) \leftrightarroweq^n_\tau (\mathfrak{B}, b)$,
2. $(\mathfrak{A}, a) \leftrightarroweq_\tau (\mathfrak{B}, b)$.

Proof. We only prove the implication from 1 to 2. Let $Z_0 \supseteq \cdots \supseteq Z_n$ be given. Define $Z \subseteq A \times B$ by

$$Zxy \text{ iff } \mathrm{depth}(x) = \mathrm{depth}(y) \text{ and } (x, y) \in Z_{n-\mathrm{depth}(x)}.$$

To see that this Z satisfies the forth condition, assume Zxy and $R^{\mathfrak{A}}xx_1 \ldots x_m$. Then $\mathrm{depth}(x) = k \leq n-1$, so it follows that there are $y_1, \ldots, y_m$ in $\mathfrak{B}$ with $R^{\mathfrak{B}}yy_1 \ldots y_m$ and $Z_{k-1}x_iy_i$ $(1 \leq i \leq m)$. Now, to conclude the proof, it suffices to show that Zx_iy_i $(1 \leq i \leq m)$, and to this end it suffices to show that $\mathrm{depth}(x_i) = \mathrm{depth}(y_i)$, but this follows from $\mathrm{depth}(x) = \mathrm{depth}(y)$ together with the fact that all states in $\mathfrak{A}$, $\mathfrak{B}$ have in-degree at most 1. ⊣

4. Forcing properties of models

Below we will want to get models that have nice properties, such as a low in-degree for each of its elements, or finite depth for each of its elements. To obtain such models the following comes in handy.

Fix a vocabulary τ. A property P of models is $\leftrightarroweq^b_\tau$*-enforceable*, or simply *enforceable*, iff for every $(\mathfrak{A}, a) \in \mathrm{Str}[\tau]$, there is a $(\mathfrak{B}, b) \in \mathrm{Str}[\tau]$ with $(\mathfrak{A}, a) \leftrightarroweq^b_\tau (\mathfrak{B}, b)$ and $(\mathfrak{B}, b)$ has P.

Proposition 4.1 *The property "every element has finite depth" is enforceable.*

Proof. Let $\mathfrak{A} \in \mathrm{Str}[\tau]$, and let $(\mathfrak{B}, a)$ be the submodel of $\mathfrak{A}$ that is generated by $H^\omega_\tau(\mathfrak{A}, a)$. In $(\mathfrak{B}, a)$ every element has finite depth. Moreover, $(\mathfrak{A}, a) \leftrightarroweq^b_\tau (\mathfrak{B}, a)$, as $(\mathfrak{B}, a)$ is a generated submodel of $(\mathfrak{A}, a)$. ⊣

Proposition 4.2 *Let $(\mathfrak{A}, a)$ a model, $(\mathfrak{B}, b)$ a generated submodel of $\mathfrak{A}$. The property "$(\mathfrak{A}, a)$ contains at least n copies of $\mathfrak{B}$" is enforceable* $(n \geq 1)$.

Proof. Let $\mathfrak{A} \in \mathrm{Str}[\tau]$. Let $\mathfrak{B}$ be the generated submodel we want to copy. (We can assume that $\mathfrak{B}$ is a proper submodel of $\mathfrak{A}$, otherwise $\mathfrak{A}$ with a copy of $\mathfrak{B}$ added to it is simply the disjoint union of two copies of $\mathfrak{A}$.) It suffices to show that we can enforce the property of containing one extra copy of $\mathfrak{B}$.

Let $\mathfrak{B}'$ denote a disjoint copy of $\mathfrak{B}$. Add $\mathfrak{B}'$ to $\mathfrak{A}$ by linking elements in $\mathfrak{B}'$ to all and only the elements in $\mathfrak{A} \setminus \mathfrak{B}$ to which the corresponding original elements in $\mathfrak{B}$ are linked. Let $(\mathfrak{C}, a)$ be $(\mathfrak{A}, a)$ denote the result, and let Z denote the identity relation on $\mathfrak{A}$; so $Z : (\mathfrak{A}, a) \leftrightarroweq^b_\tau (\mathfrak{A}, a)$. Extend Z to a bisimulation $Z' : (\mathfrak{A}, a) \leftrightarroweq^b_\tau (\mathfrak{C}, a)$ by linking elements in $\mathfrak{B}'$ to the corresponding elements in the original $\mathfrak{B}$. ⊣

Proposition 4.3 below generalizes the *unraveling construction* from standard modal logic over a vocabulary with a single binary relation symbol R [18] to arbitrary vocabularies; this generalization will be used frequently below.

Proposition 4.3 *The property "every element has in-degree at most 1" is enforceable.*

Proof. We may assume that $(\mathfrak{A}, a)$ is generated by a. Expand τ to a vocabulary τ^+ that has constants for all elements in $\mathfrak{A}$. Define a *path conjunction* to be a first-order formula that is a conjunction of closed atomic formulas (over τ^+) taken from the smallest set X such that (i) $a = a$ is in X; (ii) $Rac_1 \ldots c_n$ is in X for any R and $c_1, \ldots, c_n$ such that $(\mathfrak{A}, a) \models Rac_1 \ldots c_n$; and (iii) if $\alpha \wedge Rcc_1 \ldots c_n$ is in X and for some S and i, $(\mathfrak{A}, c_i) \models Sc_i d_1 \ldots d_m$, then $\alpha \wedge Rcc_1 \ldots c_n \wedge Sc_i d_1 \ldots d_m$ is in X. A path conjunction $\alpha \equiv \alpha' \wedge Sdd_1 \ldots d_m$ is admissible for a constant c in $\tau^+ \setminus \tau$ if c is one of the d_i occurring in the last conjunct of α.

Define a model $\mathfrak{B}$ whose domain contains, for every constant c in $\tau^+ \setminus \tau$, a copy c_α, for every α that is admissible for c. Define $R^{\mathfrak{B}} cc_1 \ldots c_n$ to hold if each of the $c_1, \ldots, c_n$ is labeled with the same path conjunction $\alpha \equiv \alpha' \wedge Rcc_1 \ldots c_n$. And define a valuation V' on $\mathfrak{B}$ by putting $c_\alpha \in V(p)$ iff $c \in V(p)$.

Finally, define a relation Z between $\mathfrak{A}$ and $\mathfrak{B}$ by putting Zxy iff $y = x_\alpha$ for some path conjunction α. Then $Z : (\mathfrak{A}, a) \leftrightarrows^b_\tau (\mathfrak{B}, a_{a=a})$. ⊣

Following the above proof of Proposition 4.3, let us call a first-order formula $\alpha(x, y)$ a *path formula of length n over τ* if it is of the form

$$\alpha(x,y) = \exists \vec{x}^1 \cdots \vec{x}^n \left(R^1 x \vec{x}^1 \wedge \left(\bigvee_i R^2 x_i^1 \vec{x}^2 \right) \wedge \ldots \wedge \left(\bigvee_i R^n x_i^{n-1} \vec{x}^n \right) \right),$$

where each of the R^i is a relation symbol in τ. Two path formulas are called *different* if either they have different lengths or they involve different relation symbols.

For future purposes it is useful to observe that in a τ-model $(\mathfrak{A}, a)$ every element has in-degree at most 1 iff the model satisfies the following collection of first-order sentences:

$$\{\forall y \, \neg(\alpha(a, y) \wedge \beta(a, y)) \mid \alpha,\ \beta \text{ are different path formulas over } \tau\}.$$

Definition 4.4 A τ-structure $(\mathfrak{A}, a)$ is called *smooth* if all elements in $(\mathfrak{A}, a)$ have finite depth and in-degree at most 1, and for all R and every R-tuple $(x, x_1, \ldots, x_n)$ in $\mathfrak{A}$, we have that all x_i have the same finite depth.

Proposition 4.5 *Smoothness is enforceable.*

Proof. By the the proofs of Propositions 4.1 and 4.3. ⊣

5. Modal equivalence and bisimulations

In this section we determine the relationship between bisimilarity and modal equivalence. For $\mathcal{BML}(\tau)$ a basic modal language over τ, let $(\mathfrak{A}, a) \equiv_{\mathcal{BML}(\tau)} (\mathfrak{B}, b)$ denote that $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$ satisfy the same $\mathcal{BML}(\tau)$-formulas.

Proposition 5.1 *Let τ be a classical vocabulary, and let $\mathcal{ML}(\tau)$ be a basic modal language over τ. Then $\leftrightarrows^b_\tau \subseteq \equiv_{\mathcal{BML}(\tau)}$.*

A similar relation holds between finite approximations of bisimulations and restricted fragments of modal languages. We need the following notation. We write $(\mathfrak{A}, a) \equiv^n_{\mathcal{BML}(\tau)} (\mathfrak{B}, b)$ for $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$ verify the same $\mathcal{BML}(\tau)$-formulas of rank at most n.

Proposition 5.2 *Let τ be a classical vocabulary, and let $\mathcal{BML}(\tau)$ be any basic modal language over τ. Then $\leftrightarrows_\tau^{b,n} \subseteq \equiv^n_{\mathcal{BML}(\tau)}$.*

Proposition 5.3 *Let $(\mathfrak{A}, a)$, $(\mathfrak{B}, b)$ be two finite models such that every element has in-degree at most 1, and depth at most n. The following are equivalent:*

1. $(\mathfrak{A}, a) \equiv^n_{\mathcal{BML}(\tau)} (\mathfrak{B}, b)$,
2. $(\mathfrak{A}, a) \leftrightarrows_\tau^{b,n} (\mathfrak{B}, b)$,
3. $(\mathfrak{A}, a) \equiv_{\mathcal{BML}(\tau)} (\mathfrak{B}, b)$,
4. $(\mathfrak{A}, a) \leftrightarrows_\tau^{b} (\mathfrak{B}, b)$.

Proof. The implication $4 \Rightarrow 2$ is Proposition 3.4. The implication $2 \Rightarrow 1$ is immediate, and the implication $3 \Rightarrow 4$ may be proved by an argument similar to the one in Theorem 5.8. To complete the proof we need to show that 1 implies 3. It suffices to observe that on models of depth $\leq n$, every basic modal formula is equivalent to a formula of rank $\leq n$. $\dashv$

Example 5.4 (Hennessy and Milner [12]) The converse of the inclusion in Proposition 5.1 does not hold: as is well-known from the general literature on bisimulations, there are $\mathcal{BML}$-equivalent models that are not bisimilar.

Let τ be a vocabulary with just a single binary relation symbol R. Define models $\mathfrak{A}$ and $\mathfrak{B}$ as in Figure 1 below, where arrows denote R-transitions: Then $(\mathfrak{A}, a) \equiv_{\mathcal{BML}(\tau)} (\mathfrak{B}, b)$, but

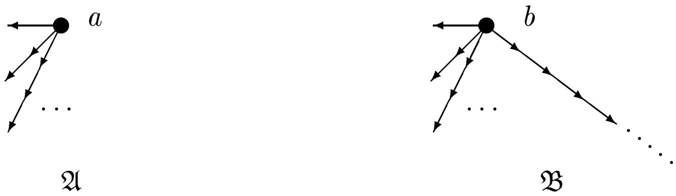

Figure 1: Equivalent but not bisimilar.

$(\mathfrak{A}, a) \not\leftrightarrows_\tau^{b} (\mathfrak{B}, b)$. The first claim is obvious; to see that the second is true, observe that any candidate bisimulation Z has to link points on the infinite branch of $\mathfrak{B}$ to points of $\mathfrak{A}$ having only finitely many successors. This violates the back-and-forth conditions.

To determine the exact relation between $\leftrightarrows_\tau^{b}$ and $\equiv_{\mathcal{BML}(\tau)}$ we need the following.

Definition 5.5 A model $\mathfrak{A} \in \mathrm{Str}[\tau]$ is said to be *ω-saturated* if for every finite subset Y of $\mathfrak{A}$, every type $\Gamma(x)$ of $\mathcal{L}_{\omega\omega}[\tau^+]$, where $\tau^+ = \tau \cup \{ c_a \mid a \in Y \}$, that is consistent with $\mathrm{Th}_{\mathcal{L}_{\omega\omega}}((\mathfrak{A}, a)_{a \in Y})$ is realized in $(\mathfrak{A}, a)_{a \in Y}$. By a routine argument the restriction to types in a single free variable may be lifted to finitely many.

Recall that an ultrafilter is *countably incomplete* if it is not closed under arbitrary intersection.

Lemma 5.6 (Keisler [14]) *Let τ be countable, $\mathfrak{A} \in \mathrm{Str}[\tau]$, and let U be a countably incomplete ultrafilter over an index set I. The ultrapower $\prod_U \mathfrak{A}$ is ω-saturated.*

Theorem 5.7 (Bisimulation Theorem) *Let $\mathfrak{A}$, $\mathfrak{B} \in \mathrm{Str}[\tau]$. $(\mathfrak{A}, a) \equiv_{\mathcal{BML}(\tau)} (\mathfrak{B}, b)$ iff $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$ have τ-bisimilar ultrapowers.*

Proof. The direction from right to left is obvious. For the converse, assume $(\mathfrak{A}, a) \equiv_{\mathcal{BML}(\tau)} (\mathfrak{B}, b)$. We construct elementary extensions $\mathfrak{A}' \succcurlyeq \mathfrak{A}$ and $\mathfrak{B}' \succcurlyeq \mathfrak{B}$, and a bisimulation between $\mathfrak{A}'$ and $\mathfrak{B}'$ that relates a and b.

First, let $\tau^+ = \tau \cup \{ c \}$, and expand $\mathfrak{A}$ and $\mathfrak{B}$ to τ^+-structures $\mathfrak{A}^+$ and $\mathfrak{B}^+$ by interpreting c as a in $\mathfrak{A}^+$, and as b in $\mathfrak{B}^+$. Let I be an infinite index set; by Chang and Keisler [7, Proposition 4.3.5] there is a countably incomplete ultrafilter U over I. By Lemma 5.6 the ultrapowers $\prod_U(\mathfrak{A}, a) =: (\mathfrak{A}', a')$ and $\prod_U(\mathfrak{B}, b) =: (\mathfrak{B}', b')$ are ω-saturated. Observe that both a' in $\mathfrak{A}'$ and b' in $\mathfrak{B}'$ realize the set of $\mathcal{BML}(\tau)$-formulas realized by a in $\mathfrak{A}$.

Define a relation Z on the universes of $\mathfrak{A}'$ and $\mathfrak{B}'$ by putting

$$Zxy \text{ iff for all } \mathcal{BML}(\tau)\text{-formulas } \phi\text{: } (\mathfrak{A}', x) \models \phi \text{ iff } (\mathfrak{B}', y) \models \phi.$$

We verify that Z is a τ-bisimulation. First, as a and b verify the same $\mathcal{BML}(\tau)$-formulas, Z must be non-empty. The condition on unary predicates is trivially met. To check the *forth* condition, assume Za_0b_0, $a_1, \ldots, a_n \in \mathfrak{A}'$, and $Ra_0a_1 \ldots a_n$ in $\mathfrak{A}'$. Define

$$\Psi_i(x_i) := \{ ST(\phi)(x_i) \mid \phi \in \mathcal{BML}(\tau),\ \mathfrak{A}', a_i \models \phi \} \ (1 \leq i \leq n).$$

Then $\bigcup_i \Psi_i(x_i) \cup \{ R\underline{b}_0 x_1 \ldots x_n \}$ is finitely satisfiable in $(\mathfrak{B}', b', b_0)$. To see this, assume $\Phi_i(x_i) \subseteq \Psi_i(x_i)$ is finite. Then

$$(\mathfrak{A}', a', a_0) \models \{ Rx_0x_1 \ldots x_n \} \cup \textstyle\bigcup_i \Phi_i(x_i)[a_1 \ldots a_n].$$

As Za_0b_0 and $\exists x_1 \ldots \exists x_n (Rx_0x_1 \ldots x_n \wedge \bigwedge \Phi_1(x_1) \wedge \ldots \wedge \bigwedge \Phi_n(x_n))$ is really a modal formula, it follows that for some $b_1, \ldots, b_n$ in $(\mathfrak{B}', b')$

$$(\mathfrak{B}', b', b_0) \models \{ Rx_0x_1 \ldots x_n \} \cup \textstyle\bigcup_i \Phi_i(x_i)[b_1 \ldots b_n],$$

Hence, by saturation, $(\mathfrak{B}', b', b_0) \models \bigcup_i \Psi_i(x_i) \cup \{ R\underline{b}_0 x_1 \ldots x_n \}[b_1 \ldots b_n]$ for some $b_1, \ldots, b_n$ in $\mathfrak{B}'$. But then we have Za_ib_i and $Rb_0b_1 \ldots b_n$ $(1 \leq i \leq n)$, as required. The *back* condition is checked similarly.

As $\mathfrak{A}'$ and $\mathfrak{B}'$ are reducts to the original vocabulary τ of the ultrapowers $\prod_U(\mathfrak{A}, a)$ and $\prod_U(\mathfrak{B}, b)$, respectively, this shows that $\mathfrak{A}$ and $\mathfrak{B}$ have τ-bisimilar ultrapowers. ⊣

The Bisimulation Theorem should be compared to a weak version of the Keisler-Shelah Theorem in first-order logic: two first-order models are *elementary equivalent* iff they have *partially isomorphic* ultrapowers (Van Benthem and Doets [8]); the original strong version of the result replaces 'partially isomorphic' with 'isomorphic' (Chang and Keisler [7, Theorem 6.1.15]).

Now that we know that finitary modal equivalence between two models means bisimilarity 'somewhere else,' the obvious next question is: for which modal language $\mathcal{L}$ *does* $\equiv_{\mathcal{L}}$ coincide with $\leftrightarrows^b_\tau$?

Theorem 5.8 *The relations $\leftrightarrows^b_\tau$ and $\equiv_{\mathcal{BML}_{\infty\omega}(\tau)}$ coincide.*

Proof. The inclusion $\leftrightarroweq_\tau^b \subseteq \equiv_{\mathcal{BML}_{\infty\omega}(\tau)}$ is immediate by an inductive argument. For the converse, we adopt an argument due to Hennessy and Milner [12]. We show that the relation Z defined by Zab whenever a and b satisfy the same $\mathcal{BML}_{\infty\omega}(\tau)$-formulas is a τ-bisimulation. Assume it is not. If a_0 and b_0 disagree on some proposition letter, then they can't have the same $\mathcal{BML}_{\infty\omega}(\tau)$-theory. Hence, for some R and $a_1, \ldots, a_n$ we have $Ra_0a_1 \ldots a_n$, while for all $b_1, \ldots, b_n$ in $\mathfrak{B}$ $Rb_0b_1 \ldots b_n$ implies that for some i a_i and b_i disagree on some formula in $\mathcal{BML}_{\infty\omega}(\tau)$. Let $X = \{ (b_1, \ldots, b_n) : Rb_0b_1 \ldots b_n \}$. Clearly $X \neq \emptyset$, and for every $(b_1, \ldots, b_n) \in X$ there is an i such that for some ϕ_i $a_i \models \phi_i$ and $b_i \not\models \phi_i$ $(1 \leq i \leq n)$. Put $\Phi_i := \bigwedge \phi_i$ (letting the empty conjunction denote $\top$). Then, for $\#_R$ the modal operator whose semantics is based on R, we have $a_0 \models \#_R(\Phi_1, \ldots, \Phi_n)$, but $b_0 \not\models \#_R(\Phi_1, \ldots, \Phi_n)$, contradicting Za_0b_0. $\dashv$

For countable structures a sharper form of Theorem 5.8 is possible: Van Benthem and Bergstra [4] show that, for vocabularies τ not containing symbols of arity > 2, countable structures are characterized up to bisimilarity by a single $\mathcal{BML}_{\omega_1\omega}(\tau)$-formula; the generalization to arbitrary vocabularies is due to Holger Sturm (personal communication). The reader should compare this result with Scott's Isomorphism Theorem saying that countable structures are characterized up to isomorphism by a single $\mathcal{L}_{\omega_1\omega}$-sentence (Scott [19]).

Theorem 5.9 (Van Benthem & Bergstra [4]) *Let τ be a countable vocabulary. For every countable structure $\mathfrak{A} \in \mathrm{Str}[\tau]$ there is a formula ϕ in $\mathcal{BML}_{\omega_1\omega}(\tau)$ such that for all a in $\mathfrak{A}$, all countable $\mathfrak{B}$ and all b in $\mathfrak{B}$, we have $(\mathfrak{A}, a) \leftrightarroweq_\tau^b (\mathfrak{B}, b)$ iff $(\mathfrak{B}, b) \models \phi$.*

We conclude the section with two brief comments on related work. First, De Rijke [17] proves a Lindström Theorem for basic modal logic: basic modal logic is the strongest logic whose formulas are invariant for bisimulations and preserved under ultrapowers over ω. Second, Goldblatt [11] and Hollenberg [13] describe so-called Hennessy-Milner classes of models; these are classes K on which modal equivalence and bisimilarity coincide; by the proof of the Bisimulation Theorem 5.7 the ω-saturated models form such a class.

6. Definability

As stated before, the standard translation ST embeds our basic modal languages into fragments of classical languages. Combined with known definability results and techniques for the classical background languages, this fact allows for easy proofs of definability results for basic modal languages. The general strategy here is to 'bisimulate' results and proofs from classical logic, for instance by replacing $\cong$, $\cong^p$ and $\preccurlyeq$ with $\leftrightarroweq$. As a corollary we find that basic τ-bisimulations cut out precisely the basic modal fragment of first-order logic.

We need some further definitions. A class of (pointed) models K is called an *$\mathcal{L}$-elementary class* (or: K is EC in $\mathcal{L}$) if $\mathsf{K} = \{ (\mathfrak{A}, a) \mid (\mathfrak{A}, a) \models \phi$, for some $\mathcal{L}$-formula $\phi \}$. We write K is EC_Δ in $\mathcal{L}$ if it is the intersection of classes that are EC in $\mathcal{L}$. For K a class of models $\overline{\mathsf{K}}$ denotes the complement of K, $\mathbf{Pr}(\mathsf{K})$ denotes the class of ultraproducts of models in K, $\mathbf{Po}(\mathsf{K})$ denotes the class of ultrapowers of models in K, and $\mathbf{B}_b(\mathsf{K})$ is the class of all models that are basically bisimilar to a model in K.

Proposition 6.1 *Let I be an index set, U an ultrafilter over I.*

1. *If for all* i, $(\mathfrak{A}_i, a_i) \leftrightarrows^b_\tau (\mathfrak{B}_i, b_i)$, *then* $\prod_U(\mathfrak{A}_i, a_i) \leftrightarrows^b_\tau \prod_U(\mathfrak{B}_i, b_i)$,
2. *If* $(\mathfrak{A}, a) \leftrightarrows^b_\tau (\mathfrak{B}, b)$, *then* $\prod_U(\mathfrak{A}, a) \leftrightarrows^b_\tau \prod_U(\mathfrak{B}, b)$.

Proof. 1. Assume that $Z_i : (\mathfrak{A}_i, a_i) \leftrightarrows^b_\tau (\mathfrak{B}_i, b_i)$. For x in $\prod_U(\mathfrak{A}_i, a_i)$ and y in $\prod_U(\mathfrak{B}_i, b_i)$ define Zxy iff $\{ i \in I \mid Z_i x(i) y(i) \} \in U$. Then Z defines a basic bisimulation $\prod_U(\mathfrak{A}_i, a_i) \leftrightarrows^b_\tau \prod_U(\mathfrak{B}_i, b_i)$ linking the distinguished points a and b of $\prod_U(\mathfrak{A}_i, a_i)$ and $\prod_U(\mathfrak{B}_i, b_i)$, respectively, where for all i in I, $a(i) = a_i$, $b(i) = b_i$.

2. This is immediate from item 1. (Alternatively, the diagonal map $d : a \to f_a$, where f_a is the constant map with value a, induces a bisimulation $(\mathfrak{A}, a) \leftrightarrows^b_\tau \prod_U(\mathfrak{A}, a)$. Likewise, one has $(\mathfrak{B}, b) \leftrightarrows^b_\tau \prod_U(\mathfrak{B}, b)$, hence $(\mathfrak{A}, a) \leftrightarrows^b_\tau (\mathfrak{B}, b)$ yields $\prod_U(\mathfrak{A}, a) \leftrightarrows^b_\tau \prod_U(\mathfrak{B}, b)$.) ⊣

Corollary 6.2 *Let* K *be a class of* τ*-models.*

1. $\mathbf{PrB}_b(\mathsf{K}) \subseteq \mathbf{B}_b\mathbf{Pr}(\mathsf{K})$, *hence* K *is closed under basic bisimulations and ultraproducts iff* $\mathsf{K} = \mathbf{B}_b\mathbf{Pr}(\mathsf{K})$,
2. $\mathbf{PoB}_b(\mathsf{K}) \subseteq \mathbf{B}_b\mathbf{Po}(\mathsf{K})$, *hence* K *is closed under basic bisimulations and ultrapowers iff* $\mathsf{K} = \mathbf{B}_b\mathbf{Po}(\mathsf{K})$.

Proof. 1. Assume $(\mathfrak{A}, a) \in \mathbf{PrB}_b(\mathsf{K})$. Then there are an index set I, models $(\mathfrak{A}_i, a_i)$ and $(\mathfrak{B}_i, b_i)$ $(i \in I)$ such that $(\mathfrak{B}_i, b_i) \in \mathsf{K}$, $(\mathfrak{A}_i, a_i) \leftrightarrows^b_\tau (\mathfrak{B}_i, b_i)$, and $(\mathfrak{A}, a) = \prod_U(\mathfrak{A}_i, a_i)$, for some ultrafilter U over I. Trivially, $\prod_U(\mathfrak{B}_i, b_i) \in \mathbf{Pr}(\mathsf{K})$. By Proposition 6.1, item 1, $(\mathfrak{A}, a) = \prod_U(\mathfrak{A}_i, a_i) \leftrightarrows^b_\tau \prod_U(\mathfrak{B}_i, b_i)$. Hence, $(\mathfrak{A}, a) \in \mathbf{B}_b\mathbf{Pr}(\mathsf{K})$. As a consequence, if $\mathbf{B}_b\mathbf{Pr}(\mathsf{K}) = \mathsf{K}$, then, as both $\mathbf{B}_b$ and $\mathbf{Pr}$ are idempotent, applying $\mathbf{B}_b$ or $\mathbf{Pr}$ does not take us outside K; this is clear for $\mathbf{B}_b$, and for $\mathbf{Pr}$ we have $\mathbf{PrB}_b\mathbf{Pr}(\mathsf{K}) \subseteq \mathbf{B}_b\mathbf{PrPr}(\mathsf{K}) \subseteq \mathbf{B}_b\mathbf{Pr}(\mathsf{K}) \subseteq \mathsf{K}$.

2. The proof is similar to the proof of item 1; use Proposition 6.1, item 2. ⊣

Theorem 6.3 (Definability Theorem) *Let* $\mathcal{L}$ *denote* $\mathcal{BML}(\tau)$, *and let* K *be a class of* τ*-models. Then*

1. K *is* EC_Δ *in* $\mathcal{L}$ *iff* $\mathsf{K} = \mathbf{B}_b\mathbf{Pr}(\mathsf{K})$ *and* $\overline{\mathsf{K}} = \mathbf{B}_b\mathbf{Po}(\overline{\mathsf{K}})$,
2. K *is* EC *in* $\mathcal{L}$ *iff* $\mathsf{K} = \mathbf{B}_b\mathbf{Pr}(\mathsf{K})$ *and* $\overline{\mathsf{K}} = \mathbf{B}_b\mathbf{Pr}(\overline{\mathsf{K}})$.

Proof. 1. The *only if* direction is easy. For the converse, assume K is closed under ultraproducts and basic bisimulations, while $\overline{\mathsf{K}}$ is closed under ultrapowers. Let

$$T = \mathrm{Th}_{\mathcal{L}}(\mathsf{K}) = \{ \phi \mid (\mathfrak{A}, a) \models \phi, \text{ for all } (\mathfrak{A}, a) \in \mathsf{K} \}.$$

Then $\mathsf{K} \models T$. Let $(\mathfrak{B}, b) \models T$. Let $\Sigma = \mathrm{Th}_{\mathcal{L}}(\mathfrak{B}, b)$, and define $I = \{ \sigma \subseteq \Sigma : |\sigma| < \omega \}$. For each $i = \{ \sigma_1, \ldots, \sigma_n \} \in I$ there is a model $(\mathfrak{A}_i, a_i)$ of i. By standard model-theoretic arguments there exists an ultraproduct $\prod_U(\mathfrak{A}_i, a_i)$ which is a model of Σ. As $\mathbf{Pr}(\mathsf{K}) \subseteq \mathsf{K}$, $\prod_U(\mathfrak{A}_i, a_i) \in \mathsf{K}$. But, if $(\mathfrak{A}, a) \models \Sigma$, then $(\mathfrak{A}, a) \equiv_{\mathcal{L}} (\mathfrak{B}, b)$, so $\prod_U(\mathfrak{A}_i, a_i) \equiv_{\mathcal{L}} (\mathfrak{B}, b)$. By the Bisimulation Theorem there is an ultrafilter U' such that $\prod_{U'}(\prod_U(\mathfrak{A}_i, a_i)) \leftrightarrows^b_\tau \prod_{U'}(\mathfrak{B}, b)$. Hence, the latter is in K, and, by the closure condition on $\overline{\mathsf{K}}$, this implies $(\mathfrak{B}, b) \in \mathsf{K}$. Therefore, K is the class of all models of T, and so K is EC_Δ in $\mathcal{L}$.

2. Again, the *only if* direction is easy. Assume $\mathsf{K}, \overline{\mathsf{K}}$ satisfy the stated conditions. Then both are closed under ultrapowers, hence, by item 1, there are sets of $\mathcal{L}$-formulas T_1, T_2 witnessing

that K is EC_Δ in $\mathcal{L}$, and that $\overline{\mathsf{K}}$ is EC_Δ in $\mathcal{L}$, respectively. Obviously, $T_1 \cup T_2 \models \bot$, so by compactness for some $\phi_1, \ldots, \phi_n \in T_1$, $\psi_1, \ldots, \psi_m \in T_2$, we have $\bigwedge_i \phi_i \models \bigvee_j \neg\psi_j$. Then K is the class of all models of $\bigwedge_i \phi_i$. ⊣

The definability results for first-order logic that correspond to Theorem 6.3 say that a class of models K is EC_Δ in first-order logic iff $\mathsf{K} = \mathbf{IPr}(\mathsf{K})$ and $\overline{\mathsf{K}} = \mathbf{IPo}(\overline{\mathsf{K}})$, and similarly for EC classes in first-order logic.

Corollary 6.4 (Separation Theorems) *Let $\mathcal{L}$ denote $\mathcal{BML}(\tau)$. Let K, L be classes of τ-models such that $\mathsf{K} \cap \mathsf{L} = \emptyset$.*

1. *If $\mathbf{B}_b\mathbf{Pr}(\mathsf{K}) = \mathsf{K}$, $\mathbf{B}_b\mathbf{Po}(\mathsf{L}) = \mathsf{L}$, then there exists a class M that is EC_Δ in $\mathcal{L}$ with $\mathsf{K} \subseteq \mathsf{M}$ and $\mathsf{L} \cap \mathsf{M} = \emptyset$,*

2. *If $\mathbf{B}_b\mathbf{Pr}(\mathsf{K}) = \mathsf{K}$, $\mathbf{B}_b\mathbf{Pr}(\mathsf{L}) = \mathsf{L}$, then there exists a class M that is EC in $\mathcal{L}$ with $\mathsf{K} \subseteq \mathsf{M}$ and $\mathsf{L} \cap \mathsf{M} = \emptyset$.*

Proof. 1. Let K' be the class of all τ-models $(\mathfrak{A}, a)$ such that for some $(\mathfrak{B}, b) \in \mathsf{K}$, $(\mathfrak{A}, a) \equiv_{\mathcal{L}} (\mathfrak{B}, b)$. Define L' similarly. Then $\mathsf{K} \subseteq \mathsf{K}'$, $\mathsf{L} \subseteq \mathsf{L}'$ and K' and L' are both closed under $\equiv_{\mathcal{L}}$.

Our first claim is that $\mathsf{K}' \cap \mathsf{L}' = \emptyset$. For suppose $(\mathfrak{A}, a) \in \mathsf{K}' \cap \mathsf{L}'$; then there exist $(\mathfrak{B}, b) \in \mathsf{K}$, $(\mathfrak{C}, c) \in \mathsf{L}$ such that $(\mathfrak{B}, b) \equiv_{\mathcal{L}} (\mathfrak{A}, a) \equiv_{\mathcal{L}} (\mathfrak{C}, c)$. By the Bisimulation Theorem $(\mathfrak{B}, b)$ and $(\mathfrak{C}, c)$ have basically τ-bisimilar ultrapowers $\prod_U(\mathfrak{B}, b)$ and $\prod_U(\mathfrak{C}, c)$. As K, L are closed under $\mathbf{B}_b$ and $\mathbf{Po}$, this implies $\prod_U(\mathfrak{B}, b) \in \mathsf{K} \cap \mathsf{L}$, contradicting $\mathsf{K} \cap \mathsf{L} = \emptyset$.

Let $T = \mathrm{Th}_{\mathcal{L}}(\mathsf{K}')$. Then K' is the class of models of T. As $\mathsf{K} \subseteq \mathsf{K}'$ and $\mathsf{K}' \cap \mathsf{L} = \emptyset$, we are done.

2. This may be proved analogously to 1. Use the assumption that $\mathbf{B}_b\mathbf{Pr}(\mathsf{L}) = \mathsf{L}$ to conclude that L' is EC_Δ in $\mathcal{L}$, and then apply a compactness argument as in the proof of Theorem 6.3, part 2. ⊣

The Separation Theorems are 'bisimilar' to corresponding results in first-order logic. Observe that the Craig Interpolation Theorem is a special case of 6.4:

Theorem 6.5 *If K, L are EC in $\mathcal{BML}(\tau')$ for some $\tau' \supseteq \tau$, and $\mathsf{K} \cap \mathsf{L} = \emptyset$, then there is a class M that is EC in $\mathcal{BML}(\tau)$ with $\mathsf{K} \subseteq \mathsf{M}$ and $\mathsf{M} \cap \mathsf{L} = \emptyset$.*

The Definability Theorem 6.3 is difficult to apply in practice, as ultrapowers are rather abstract objects. The following Fraïssé type result supplies a more manageable criterion for EC classes.

Theorem 6.6 *Let τ be a finite vocabulary, and let K be a class of τ-models. Then K is EC in $\mathcal{BML}(\tau)$ iff, for some $n \in \mathbb{N}$, K is closed under basic τ-bisimulations up to n.*

Proof. The *only if* direction is clear. If K is closed under basic τ-bisimulations up to n, let $(\mathfrak{A}, a) \in \mathsf{K}$, and define $\phi_{(\mathfrak{A},a)}$ to be the conjunction of all $\mathcal{BML}$-formulas of rank at most n that are true at a. (Observe that over a finite vocabulary there are only finitely many basic modal formulas of any given rank). Modulo equivalence there are only finitely many such formulas $\phi_{(\mathfrak{A},a)}$ for $(\mathfrak{A}, a) \in \mathsf{K}$; let Φ be their disjunction. Then Φ defines K. For let $(\mathfrak{B}, b) \models \Phi$; then $(\mathfrak{B}, b) \equiv^n_{\mathcal{BML}} (\mathfrak{A}, a)$ for some $(\mathfrak{A}, a) \in \mathsf{K}$. By a routine induction, $(\mathfrak{B}, b) \leftrightarrows^{b,n}_{\tau} (\mathfrak{A}, a)$; hence $(\mathfrak{B}, b) \in \mathsf{K}$. ⊣

To conclude our list of results on definability we give a theorem that characterizes the modal fragment of first-order logic. For the standard modal language $\mathcal{ML}(\Diamond)$ a semantic description of the corresponding first-order fragment in terms of bisimulations was first given by Van Benthem [2, Theorem 1.9].

We need a definition. Let $\alpha(x)$ be a first-order formula over τ; $\alpha(x)$ is called *invariant for basic τ-bisimulations* if for all $(\mathfrak{A}, a)$, $(\mathfrak{B}, b) \in \mathrm{Str}[\tau]$, all basic τ-bisimulations $Z : (\mathfrak{A}, a) \leftrightarroweq^b_\tau (\mathfrak{B}, b)$, and all $x \in \mathfrak{A}$, $y \in \mathfrak{B}$ we have that Zxy implies $\mathfrak{A} \models \alpha[x]$ iff $\mathfrak{B} \models \alpha[y]$.

Theorem 6.7 (Fragment Theorem) *Let $\alpha(x)$ be a first-order formula over τ. The following are equivalent.*

1. *$\alpha(x)$ is equivalent to (the ST-translation of) a modal formula in $\mathcal{BML}(\tau)$.*
2. *$\alpha(x)$ is invariant under basic τ-bisimulations.*
3. *for some $n \in \mathbb{N}$, α is invariant under basic τ-bisimulations up to n.*

Proof. We only prove the implication $2 \Rightarrow 1$. Let K be the class of models of $\alpha(x)$. Then K and $\overline{\mathsf{K}}$ (being defined by $\neg\alpha(x)$) are closed under ultraproducts. As α is invariant under $\leftrightarroweq^b_\tau$, it follows that $\mathsf{K} = \mathbf{B}_b\mathbf{Pr}(\mathsf{K})$ and $\overline{\mathsf{K}} = \mathbf{B}_b\mathbf{Pr}(\overline{\mathsf{K}})$. By Theorem 6.3 K must be EC in $\mathcal{BML}(\tau)$. This means that α is equivalent to (the translation of) some modal formula ϕ. ⊣

7. Preservation

Preservation results formed the backbone of model theory for first-order logic until the early sixties. More recently there has been a renewed interest in preservation results with the growing importance of restricted fragments and restricted model classes. The best known examples of preservation results in first-order logic include

- Łoś's Theorem: A first-order formula is preserved under submodels iff it is equivalent to a universal first-order formula (Chang and Keisler [7, Theorem 3.2.2]).
- The Chang-Łoś-Suszko Theorem: A first-order formula is preserved under unions of chains iff it is equivalent to a 'universal-existential' first-order formula (Chang and Keisler [7, Theorem 3.2.3]).
- Lyndon's Theorem: A first-order formula is preserved under homomorphisms iff it is equivalent to a positive first-order formula (Chang and Keisler [7, Theorem 3.2.4]).

To further substantiate our main claim that bisimulations form the basic tools for the model theory of modal logic, we will prove modal versions of each of the above preservation results.

Submodels

Definition 7.1 A formula in $\mathcal{BML}(\tau)$ is *existential* if it has been built using (negated) proposition letters, $\vee$, $\wedge$, $\bot$, $\top$ and modal operators $\#$ only. A formula in $\mathcal{BML}(\tau)$ is *universal* if it has been built using (negated) proposition letters, $\vee$, $\wedge$, $\bot$, $\top$ and duals $\overline{\#}$ of modal operators $\#$ in $\mathcal{BML}(\tau)$ only.

Definition 7.2 Let $(\mathfrak{A}, a)$, $(\mathfrak{B}, b)$ be two models for the same vocabulary. $(\mathfrak{A}, a)$ is a *submodel* of $(\mathfrak{B}, b)$ if $a = b$, and for every R, $R^{\mathfrak{A}}$ is the restriction of $R^{\mathfrak{B}}$ to the (appropriate) domain(s) of $\mathfrak{A}$. A basic modal formula is *preserved under submodels* if $(\mathfrak{B}, b) \models \phi$ implies $(\mathfrak{A}, a) \models \phi$ whenever $(\mathfrak{A}, a)$ is a submodel of $(\mathfrak{B}, b)$.

To prove a basic modal version of Łoś's Theorem we need a technical lemma. The following triple arrow notation will be useful. For Σ a set of $\mathcal{BML}(\tau)$-formulas, $(\mathfrak{A}, a) \Rrightarrow_{\Sigma} (\mathfrak{B}, b)$ abbreviates: for all $\phi \in \Sigma$, $(\mathfrak{A}, a) \models \phi$ implies $(\mathfrak{B}, b) \models \phi$; in particular we will use $\Rrightarrow_{\mathrm{E}}$, where 'E' denotes the set of all existential formulas.

Recall that a model is smooth if if each of its elements has finite depth and in-degree at most 1, and for all R and every R-tuple $(x, x_1, \ldots, x_n)$, we have that all x_i have the same finite depth (see Definition 4.4).

Lemma 7.3 *Let $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$ be τ-structures such that $(\mathfrak{A}, a)$ is smooth, $(\mathfrak{B}, b)$ is ω-saturated, and $(\mathfrak{A}, a) \Rrightarrow_{\mathrm{E}} (\mathfrak{B}, b)$. Then there exists $(\mathfrak{B}', b') \leftrightarrows_{\tau}^{b} (\mathfrak{B}, b)$ such that $(\mathfrak{A}, a)$ is embeddable in $(\mathfrak{B}', b')$.*

In a diagram the Lemma claims:

$$\begin{array}{ccc} (\mathfrak{A}, a) & \Rrightarrow_{\mathrm{E}} & (\mathfrak{B}, b) \\ = \Big| & & \Big| \leftrightarrows_{\tau}^{b} \\ (\mathfrak{A}, a) & \hookrightarrow & (\mathfrak{B}', b'). \end{array}$$

In a somewhat different form, and restricted to the standard modal language, Lemma 7.3 is due to Van Benthem [3].

Proof of Lemma 7.3. We define a 'forth simulation' F between $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$, that is: a relation F that links two points only if they agree on all proposition letters, and that satisfies the *forth* condition:

> if Fvw, $R^{\mathfrak{A}}vv_1 \ldots v_n$, then there are $w_1, \ldots, w_n$ in $\mathfrak{B}$ with
> $R^{\mathfrak{B}}ww_1 \ldots w_n$ and Fv_iw_i $(1 \leq i \leq n)$.

We define a function F from $(\mathfrak{A}, a)$ to $(\mathfrak{B}, b)$ by induction on the depth of elements in $(\mathfrak{A}, a)$. This function will be a forth simulation, and as such it will satisfy $(\mathfrak{A}, x) \Rrightarrow_{\mathrm{E}} (\mathfrak{B}, Fx)$. Put $Fa = b$. Assume that F has been defined for all elements of depth $< n$. Let x in $(\mathfrak{A}, a)$ have depth n. By the smoothness of $(\mathfrak{A}, a)$ there are *unique* elements y of depth $n - 1$, and $x_1, \ldots, x_n$ of depth n such that x is one of $x_1, \ldots, x_n$, and such that for some R we have $R^{\mathfrak{A}}yx_1 \ldots x_n$. We define F for each of $x_1, \ldots, x_n$. Let E_i be the set of existential modal formulas satisfied by x_i. By $(\mathfrak{A}, y) \Rrightarrow_{\mathrm{E}} (\mathfrak{B}, Fy)$ and saturation there are $x_1', \ldots, x_n'$ in $\mathfrak{B}$ with $x_i' \models E_i$ and $R^{\mathfrak{B}}F(y)x_1' \ldots x_n'$ $(1 \leq i \leq n)$. Put $Fx_i = x_i'$ $(1 \leq i \leq n)$.

The next step is to extend F to a full bisimulation between a supermodel $(\mathfrak{B}', b')$ of $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$. Define $(\mathfrak{B}', b')$ (as in Figure 2) to be the disjoint union of $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$ *in which we identify the two distinguished points of* $(\mathfrak{A}, a)$ *and* $(\mathfrak{B}, b)$, and with the following extension of the relations:

> if $x \in (\mathfrak{A}, a)$, $Fx = y$ and $Ryv_1 \ldots v_n$, then $Rxv_1 \ldots v_n$.

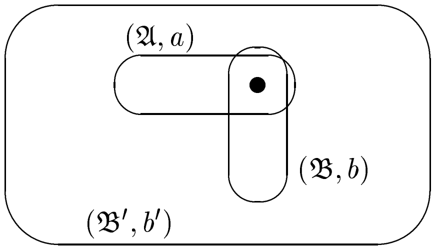

Figure 2: Combining $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$.

Observe that a and b agree on all proposition letters, thus their identification is well-defined. Define a relation Z between the domain of $(\mathfrak{B}', b')$ and the domain of $(\mathfrak{B}, b)$ as follows: for x in $\mathfrak{A}$ we put Zxy whenever $Fx = y$, and for x in $\mathfrak{B}$ we put Zxx. Then $Z : (\mathfrak{B}', b') \leftrightarrows^b_\tau (\mathfrak{B}, b)$:

- Z-related points agree on all proposition letters,
- Assume v in $\mathfrak{B}'$, w in $\mathfrak{B}$ and Zvw. If $R^{\mathfrak{B}'}vv_1 \dots v_k$, then either $v_1, \dots, v_k$ all live in $\mathfrak{A}$, or they all live in $\mathfrak{B}$. In the first case our forth simulation F will find $w_1, \dots, w_k$ with Zv_iw_i $(1 \leq i \leq k)$ and $R^{\mathfrak{B}}ww_1 \dots w_k$. In the second case we have two possibilities: if v in $\mathfrak{B}$, then $v = w$, $R^{\mathfrak{B}}vv_1 \dots v_k$ and Zv_iv_i. The other possibility is that v is not in $\mathfrak{B}$; but then $Fv = w$ and $R^{\mathfrak{B}}wv_1 \dots v_k$, and by construction Zv_iw_i, as required.
- Assume v in $\mathfrak{B}'$, w in $\mathfrak{B}$ and Zvw. Assume also that $R^{\mathfrak{B}}ww_1 \dots w_k$. If v in $\mathfrak{A}$, then by construction $Fv = w$, and $R^{\mathfrak{B}'}vw_1 \dots w_k$ and Zw_iw_i. If v is in $\mathfrak{B}$, then we must have $v = w$, $R^{\mathfrak{B}'}vw_1 \dots w_k$ and Zw_iw_i, and we are done.

Thus $Z : (\mathfrak{B}', b') \leftrightarrows^b_\tau (\mathfrak{B}, b)$. As $(\mathfrak{A}, a)$ lies embedded as a submodel in $(\mathfrak{B}', b')$, this completes the proof. $\dashv$

Theorem 7.4 (Łoś's Theorem) *A basic modal formula is preserved under submodels iff it is equivalent to a universal basic modal formula.*

Proof. Aside from an application of Lemma 7.3 this is a routine argument. First, it is easy to check that if ϕ is equivalent to a universal formula, then it is preserved under submodels.

Second, if ϕ is so preserved, let $\mathrm{CONS}_U(\phi)$ be the set of universal consequences of ϕ. By compactness it suffices to show $\mathrm{CONS}_U(\phi) \models \phi$. So assume $(\mathfrak{A}, a) \models \mathrm{CONS}_U(\phi)$; we may assume that $(\mathfrak{A}, a)$ is smooth. Let E be the set of all existential formulas ψ with $(\mathfrak{A}, a) \models \psi$. Then, by compactness, $E + \phi$ has a model $(\mathfrak{B}, b)$, which may be assumed to be ω-saturated. By Lemma 7.3 $(\mathfrak{B}, b) \models E + \phi$ implies that some supermodel $(\mathfrak{B}', b')$ of $(\mathfrak{A}, a)$ has $(\mathfrak{B}', b') \models \phi$. By preservation under submodels $(\mathfrak{A}, a) \models \phi$. $\dashv$

Unions of chains

Definition 7.5 A formula in $\mathcal{BML}(\tau)$ is *universal existential* if it has been built using existential formulas, $\wedge$, $\vee$, and dual modal operators $\overline{\#}$ only. A formula is *existential universal* if it has been built using universal formulas, $\wedge$, $\vee$, and modal operators $\#$ only.

We write $(\mathfrak{A}, a) \Rrightarrow_{\mathrm{UE}} (\mathfrak{B}, b)$ for: $(\mathfrak{B}, b)$ satisfies all universal existential formulas satisfied by $(\mathfrak{A}, a)$; and similarly for $\Rrightarrow_{\mathrm{EU}}$.

Definition 7.6 A *chain* of τ-structures is a collection $((\mathfrak{A}_i, a_i) : i \in I)$ such that for all i, j, if $i < j$, then $(\mathfrak{A}_i, a_i)$ is a submodel of $(\mathfrak{A}_j, a_j)$. A *bisimilar chain* is a chain $((\mathfrak{A}_i, a_i) : i \in I)$ in which for all $i \leq j \in I$, $(\mathfrak{A}_i, a_i) \mathrel{\underline{\leftrightarrow}}^b_\tau (\mathfrak{A}_j, a_j)$.

The *union of the chain* $((\mathfrak{A}_i, a_i) : i \in I)$ is the model $\mathfrak{A} = \bigcup_{i \in I}(\mathfrak{A}_i, a_i)$ whose universe is the set $\bigcup_{i \in I}$, and whose relations are the unions of the corresponding relations of $(\mathfrak{A}_i, a_i)$: $R^{\mathfrak{A}} = \bigcup R^{\mathfrak{A}_i}$.

Lemma 7.7 *Let $((\mathfrak{A}_i, a_i) : i \in \omega)$ be a bisimilar chain of τ-structures such that fo all $i \in \omega$, $(\mathfrak{A}_i, a_i) \hookrightarrow (\mathfrak{A}_{i+1}, a_{i+1})$. Then, for each j, $(\mathfrak{A}_j, a_j) \mathrel{\underline{\leftrightarrow}}^b_\tau \bigcup_{i \in \omega}(\mathfrak{A}_i, a_i)$.*

Lemma 7.8 *Assume $(\mathfrak{C}, c)$ is a smooth model that lies embedded as a submodel in $(\mathfrak{D}, d)$. Then there exists $(\mathfrak{E}, e) \mathrel{\underline{\leftrightarrow}}^b_\tau (\mathfrak{D}, d)$ such that $(\mathfrak{C}, c)$ lies embedded as a submodel in $(\mathfrak{E}, e)$ and $(\mathfrak{E}, e)$ is smooth.*

Proof. First, take the submodel of $(\mathfrak{D}, d)$ that is generated by d, and then apply the 'unraveling' construction of Proposition 4.3 to the result. As $(\mathfrak{C}, c)$ is smooth neither operation will affect $(\mathfrak{C}, c)$. $\dashv$

Lemma 7.9 *Let $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$ be τ-structures such that $(\mathfrak{A}, a)$ is smooth, $(\mathfrak{B}, b)$ is ω-saturated, and $(\mathfrak{A}, a) \Rrightarrow_{\mathrm{EU}} (\mathfrak{B}, b)$. Then there exists a smooth model $(\mathfrak{B}', b') \mathrel{\underline{\leftrightarrow}}^b_\tau (\mathfrak{B}, b)$ such that $(\mathfrak{A}, a)$ is embeddable in $(\mathfrak{B}', b')$ and $(\mathfrak{A}, a) \Rrightarrow_{\mathrm{U}} (\mathfrak{B}', b')$.*

In a diagram the Lemma claims:

$$\begin{array}{ccc}
(\mathfrak{A}, a) & \Rrightarrow_{\mathrm{EU}} & (\mathfrak{B}, b) \\
= \big| & & \big| \mathrel{\underline{\leftrightarrow}}^b_\tau \\
(\mathfrak{A}, a) & \overset{\hookrightarrow}{\Rrightarrow_{\mathrm{U}}} & (\mathfrak{B}', b').
\end{array}$$

Proof of Lemma 7.9. This is similar to the proof of Lemma 7.3. Define a function F that is a forth simulation from $(\mathfrak{A}, a)$ to $(\mathfrak{B}, b)$ such that $Fa = b$ and $(\mathfrak{A}, x) \Rrightarrow_{\mathrm{EU}} (\mathfrak{B}, Fx)$. Extend F to a full bisimulation between $(\mathfrak{A}, a)$ and a supermodel $(\mathfrak{B}', b') \mathrel{\underline{\leftrightarrow}}^b_\tau (\mathfrak{B}, b)$ of $(\mathfrak{A}, a)$ that has $(\mathfrak{A}, a) \hookrightarrow (\mathfrak{B}', b')$, as in the proof of 7.3. By Lemma 7.8 we may take $(\mathfrak{B}', b')$ to be smooth. To complete the proof we need to show that $(\mathfrak{A}, a) \Rrightarrow_{\mathrm{U}} (\mathfrak{B}', b')$. This is almost trivial: for a universal formula ϕ we have that $(\mathfrak{A}, a) \models \phi$ implies $(\mathfrak{B}, b) \models \phi$, as $(\mathfrak{A}, a) \Rrightarrow_{\mathrm{EU}} (\mathfrak{B}, b)$. Since $(\mathfrak{B}', b') \mathrel{\underline{\leftrightarrow}}^b_\tau (\mathfrak{B}, b)$, this implies $(\mathfrak{B}', b') \models \phi$. $\dashv$

Lemma 7.10 *Assume $(\mathfrak{A}, a)$ is a model of the universal existential consequences of ϕ. Then there exists an ω-saturated model $(\mathfrak{B}, b)$ such that*

1. $(\mathfrak{B}, b) \models \phi$,
2. $(\mathfrak{A}, a)$ *is in embeddable in* $(\mathfrak{B}, b)$*, and*
3. $(\mathfrak{A}, a) \Rrightarrow_{\mathrm{U}} (\mathfrak{B}, b)$.

Proof. Consider the set T of all existential universal modal formulas ψ such that $(\mathfrak{A}, a) \models \psi$. Then, by a simple compactness argument, $T \cup \{\phi\}$ has a model $(\mathfrak{C}, c)$, which we may assume to be ω-saturated. Applying Lemma 7.9, we find that $(\mathfrak{A}, a)$ can be extended to a model $(\mathfrak{B}, b)$ of ϕ that satisfies all the universal formulas satisfied by $(\mathfrak{A}, a)$. Moreover, we can take $(\mathfrak{B}, b)$ to be ω-saturated. $\dashv$

Theorem 7.11 (Chang-Łoś-Suszko Theorem) *A basic modal formula is preserved under unions of chains iff it is equivalent to a universal existential formula.*

Proof. Again, the argument is (bi-)similar to the standard argument proving the result for first-order logic. We only prove the hard direction. Assume ϕ is preserved under unions of chains. Let $\mathrm{CONS}_{UE}(\phi)$ denote the set of universal existential consequences of ϕ. It suffices to prove that $\mathrm{CONS}_{UE}(\phi) \models \phi$. So assume $(\mathfrak{A}_0, a_0) \models \mathrm{CONS}_{UE}(\phi)$; we may of course assume that $(\mathfrak{A}_0, a_0)$ is smooth and ω-saturated. We prove that $(\mathfrak{A}_0, a_0) \models \phi$. To this end we construct a bisimilar chain $((\mathfrak{A}_i, a_i) : i < \omega)$ of smooth, ω-saturated models, ω-saturated extensions $(\mathfrak{B}_i, b_i) \supseteq (\mathfrak{A}_i, a_i)$, and embeddings $g_i : (\mathfrak{B}_i, b_i) \to (\mathfrak{A}_{i+1}, a_{i+1})$ as in the following diagram:

$$\begin{array}{ccccccccc} & & (\mathfrak{B}_0, b_0) & & & & (\mathfrak{B}_1, b_1) & & \\ & \overset{\subseteq}{\nearrow} & & \searrow^{g_0} & & \overset{\subseteq}{\nearrow} & & \searrow^{g_1} & \\ (\mathfrak{A}_0, a_0) & & \leftrightarrows^b_\tau & & (\mathfrak{A}_1, a_1) & & \leftrightarrows^b_\tau & & (\mathfrak{A}_2, a_2) \quad \ldots \end{array} \tag{7.1}$$

We will require that for each $i < \omega$:

$$(\mathfrak{B}_i, b_i) \models \phi \text{ and } (\mathfrak{B}_i, b_i) \Rrightarrow_{\mathrm{E}} (\mathfrak{A}_i, a_i). \tag{7.2}$$

The diagram is constructed as follows. Suppose $(\mathfrak{A}_i, a_i)$ has been defined. As $(\mathfrak{A}_0, a_0) \leftrightarrows^b_\tau (\mathfrak{A}_i, a_i)$ we have $(\mathfrak{A}_i, a_i) \models \mathrm{CONS}_{UE}(\phi)$. By Lemma 7.10 there exists an ω-saturated extension (B_i, b_i) of $(\mathfrak{A}_i, a_i)$ satisfying (7.2). As $(\mathfrak{A}_i, a_i)$ is smooth, we may assume (B_i, b_i) to be smooth and ω-saturated by Lemma 7.7. Applying Lemma 7.3, we find a model $(\mathfrak{A}_{i+1}, a_{i+1})$ bisimilar to $\mathfrak{A}_i, a_i)$ and an embedding $g_i : (\mathfrak{B}_i, b_i) \hookrightarrow (\mathfrak{A}_{i+1}, a_{i+1})$ such that g_i is the identity on $\mathfrak{A}_i$; we may assume that the model $(\mathfrak{A}_{i+1}, a_{i+1})$ is both smooth and ω-saturated.

In the diagram (7.1) we can replace each $(\mathfrak{B}_i, b_i)$ by its image under g_i, and so assume that the maps are inclusions. Then $\bigcup_{i<\omega}(\mathfrak{A}_i, a_i)$ and $\bigcup_{i<\omega}(\mathfrak{B}_i, b_i)$ are the same structure $(\mathfrak{C}, c)$. As ϕ is preserved under unions of chains, $(\mathfrak{B}_i, b_i) \models \phi$ (for all i) implies $(\mathfrak{C}, c) \models \phi$. By Lemma 7.7 $(\mathfrak{A}_0, a_0) \leftrightarrows^b_\tau (\mathfrak{C}, c)$, hence $(\mathfrak{A}_0, a_0) \models \phi$. $\dashv$

Homomorphisms

Definition 7.12 A formula ϕ in $\mathcal{BML}(\tau)$ is *positive* iff it has been built up using only $\bot$, $\top$, proposition letters, $\wedge$, $\vee$, as well as modal operators $\#$ and their duals $\overline{\#}$. A formula ϕ is *negative* iff it has been built up from $\bot$, $\top$, negated proposition letters, $\wedge$, $\vee$, as well as modal operators $\#$ and their duals $\overline{\#}$.

Definition 7.13 Let $(\mathfrak{A}, a)$, $(\mathfrak{B}, b)$ be two τ-structures. A *homomorphism* $f : (\mathfrak{A}, a) \to (\mathfrak{B}, b)$ is a mapping with $f(a) = b$, that preserves all relations and proposition letters. A basic modal formula ϕ is *preserved under surjective homomorphisms* if $(\mathfrak{A}, a) \models \phi$ implies $(\mathfrak{B}, b) \models \phi$ whenever $(\mathfrak{B}, b)$ is a homomorphic image of $(\mathfrak{A}, a)$.

Some more notation: $(\mathfrak{A}, a) \Rrightarrow_{\mathrm{P}} (\mathfrak{B}, b)$ is short for: for all positive formulas ψ, $(\mathfrak{A}, a) \models \psi$ implies $(\mathfrak{B}, b) \models \psi$.

Lemma 7.14 *Let $(\mathfrak{A}, a)$, $(\mathfrak{B}, b)$ be ω-saturated τ-structures with $(\mathfrak{A}, a) \Rrightarrow_{\mathrm{P}} (\mathfrak{B}, b)$, and such that both in $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$ all elements have in-degree at most 1. Then there exist τ-structures $(\mathfrak{A}', a') \,\underline{\leftrightarrow}^b_\tau\, (\mathfrak{A}, a)$ and $(\mathfrak{B}', b') \,\underline{\leftrightarrow}^b_\tau\, (\mathfrak{B}, b)$ with a surjective homomorphism $f : (\mathfrak{A}', a') \to (\mathfrak{B}', b')$.*

In a diagram the Lemma asserts the existence of the following configuration:

$$\begin{array}{ccc} (\mathfrak{A}, a) & \Rrightarrow_{\mathrm{P}} & (\mathfrak{B}, b) \\ \underline{\leftrightarrow}^b_\tau \Big| & & \Big| \underline{\leftrightarrow}^b_\tau \\ (\mathfrak{A}', a') & \overset{f}{\twoheadrightarrow} & (\mathfrak{B}', b'). \end{array}$$

Proof of Lemma 7.14. The strategy of the proof is to move to smooth models where we can inductively define a surjective homomorphism from a model bisimilar to $(\mathfrak{A}, a)$ onto a model bisimilar to $(\mathfrak{B}, b)$. To ensure surjectivity we have to blow up the model bisimilar to $(\mathfrak{A}, a)$.

Let $(\mathfrak{A}'', a)$ be the submodel of $(\mathfrak{A}, a)$ generated by a, and let $(\mathfrak{B}', b)$ be the submodel of $(\mathfrak{B}, b)$ generated by b. Then both $(\mathfrak{A}'', a)$ and $(\mathfrak{B}', b)$ are smooth. By induction on the depth of elements we will add $|\mathfrak{B}'|^+$ many copies of all (tuples of) elements in $(\mathfrak{A}'', a)$. We show how to do this by adding copies of elements of depth 1 in $(\mathfrak{A}'', a)$ to obtain a model $(\mathfrak{A}_1, a) \,\underline{\leftrightarrow}^b_\tau\, (\mathfrak{A}'', a)$.

Define $\sim$ on the elements of depth 1 in $(\mathfrak{A}'', a)$ by putting $x \sim y$ iff for some R and x_1, $\ldots$, x_n we have that both x and y are among $x_1, \ldots, x_n$ and $R^{\mathfrak{A}''} a x_1 \ldots x_n$. By smoothness this is well defined. For each $\sim$-equivalence class $X = \{ x_1 \ldots x_n \}$ let $\mathfrak{C}_X$ be the submodel of $(\mathfrak{A}'', a)$ that is generated by X. Now, for each $\mathfrak{C}_X$ take $|\mathfrak{B}'|^+$ many disjoint copies of $\mathfrak{C}_X$, and add them to $(\mathfrak{A}'', a)$; for each copy $\mathfrak{C}'_{X'}$ of $\mathfrak{C}_X$ relate the generating points $x'_1, \ldots, x'_n$ to a the way the originals $x_1, \ldots, x_n$ are related to a. Let $(\mathfrak{A}_1, a)$ be the resulting model. Then $(\mathfrak{A}'', a) \,\underline{\leftrightarrow}^b_\tau\, (\mathfrak{A}_1, a)$. Repeat this construction for all depths n to obtain models

$$(\mathfrak{A}, a) \,\underline{\leftrightarrow}^b_\tau\, (\mathfrak{A}'', a) \underset{\underline{\leftrightarrow}^b_\tau}{\overset{\subseteq}{}} (\mathfrak{A}_1, a) \underset{\underline{\leftrightarrow}^b_\tau}{\overset{\subseteq}{}} (\mathfrak{A}_2, a) \cdots$$

Define $(\mathfrak{A}', a) = \bigcup (\mathfrak{A}_i, a)$. Then $(\mathfrak{A}', a) \,\underline{\leftrightarrow}^b_\tau\, (\mathfrak{A}, a)$, and $(\mathfrak{A}', a)$ has at least $|\mathfrak{B}'|^+$ many copies of each of its submodels generated by tuples $x_1, \ldots, x_n$ such that $R^{\mathfrak{A}'} x x_1 \ldots x_n$ for some R and x.

Next we define a function F from $(\mathfrak{A}', a)$ to $(\mathfrak{B}', b)$ by induction on the depth of elements in such a way that $(\mathfrak{A}', x) \Rrightarrow_{\mathrm{P}} (\mathfrak{A}, Fx)$. For each n we first make sure that all elements of depth n in $(\mathfrak{B}', b)$ are in the range of F. After that we give F values to points of depth n in $(\mathfrak{A}', a)$ that are not yet in the domain of F.

Here we go. Put $Fa = b$. Assume that $n > 0$, and that F has been defined for all depths less than n in such a way that all elements of $(\mathfrak{B}', b)$ of depth less than n are already in the range of F. Let y in $(\mathfrak{B}', b)$ have depth n, and choose z of depth $n-1$, $y_1, \ldots, y_n$ of depth n and R such that $R^{\mathfrak{B}'} z y_1 \ldots y_n$ and y is one of the y_i $(1 \leq i \leq n)$. Let N_i be the set of all negative modal formulas satisfied by y_i in $(\mathfrak{B}', b)$. Then $(\mathfrak{B}, y_i) \models N_i$. By assumption there exists x' in $\mathfrak{A}'$ with $Fx' = z$, and $(\mathfrak{A}', x') \Rrightarrow_{\mathrm{P}} (\mathfrak{B}', z)$. Let x in $\mathfrak{A}''$ be such that x' is a copy of x if x' is in $\mathfrak{A}' \setminus \mathfrak{A}''$, and $x = x'$ otherwise. Then $(\mathfrak{A}', x') \leftrightarrows^b_\tau (\mathfrak{A}'', x) \leftrightarrows^b_\tau (\mathfrak{A}, x)$. Hence, $(\mathfrak{A}, x) \Rrightarrow_{\mathrm{P}} (\mathfrak{B}, z)$. By a saturation argument there are $x_1, \ldots, x_n$ in $\mathfrak{A}$ with $R^{\mathfrak{A}} x x_1 \ldots x_n$ and $x_i \models N_i$ $(1 \leq i \leq n)$. Then $x_1, \ldots, x_n$ are in $\mathfrak{A}''$. Now let $x'_1, \ldots, x'_n$ be copies of $x_1, \ldots, x_n$ such that $R^{\mathfrak{A}'} x' x'_1 \ldots x'_n$ and such that $x'_1, \ldots, x'_n$ are not yet in the domain of F (this is possible as we have added $|\mathfrak{B}'|^+$ many copies to $\mathfrak{A}''$), and put $Fx'_i = x_i$ $(1 \leq i \leq n)$.

Once we have included all elements of depth n in $(\mathfrak{B}', b)$ in the range of F, we define what F should do with elements of depth n in $(\mathfrak{A}', a)$ by using a saturation argument as before, but this time using sets P_i of positive modal formulas, rather than sets N_i of negative modal formulas.

Obviously, the function F thus defined is a homomorphism and a surjection. Hence we are done. ⊣

Theorem 7.15 (Lyndon's Theorem) *A basic modal formula is preserved under surjective homomorphisms iff it is equivalent to a positive modal formula.*

Proof. We only prove the hard direction: assume ϕ is preserved under surjective homomorphisms. Let $\mathrm{CONS}_P(\phi)$ be the set of positive formulas ψ with $\phi \models \psi$. It suffices to show that $\mathrm{CONS}_P(\phi) \models \phi$. Assume $(\mathfrak{B}, b) \models \mathrm{CONS}_P(\phi)$. Let N be the set of all negative formulas true at b in $\mathfrak{B}$. Let $(\mathfrak{A}, a) \models N + \phi$. Then $(\mathfrak{A}, a) \Rrightarrow_{\mathrm{P}} (\mathfrak{B}, b)$. We may of course assume that both $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$ are ω-saturated, and that all elements in $(\mathfrak{A}, a)$, $(\mathfrak{B}, b)$ have in-degree at most 1.

By Lemma 7.14 there are $(\mathfrak{A}', a') \leftrightarrows^b_\tau (\mathfrak{A}, a)$ and $(\mathfrak{B}', b') \leftrightarrows^b_\tau (\mathfrak{B}, b)$, as well as a homomorphism $f : (\mathfrak{A}', a') \twoheadrightarrow (\mathfrak{B}', b')$. Now, $(\mathfrak{A}, a) \models \phi$ implies $(\mathfrak{A}', a') \models \phi$; by preservation under surjective homomorphisms this implies $(\mathfrak{B}', b) \models \phi$, which gives $(\mathfrak{B}, b) \models \phi$, as required. ⊣

8. Concluding remarks and historical notes

This paper has developed the model theory of the class of basic modal languages in parallel with the basic model theory of first-order logic, using bisimulations as its key tool. By means of a Bisimulation Theorem, according to which two models are equivalent in basic modal logic iff they have bisimilar ultrapowers, a series of definability and separability results were obtained; in addition, we were able to prove preservation results for universal, universal existential and positive basic modal formulas that by using bisimulations in an essential way.

The paper only covered some rudimentary model theory, and it only did so for basic modal languages and some extensions — a lot remains to be done.

1. First, our Fragment Theorem in §5 only characterizes (finitary) basic modal languages as a fragments of first-order languages. What about characterizations of infinitary basic modal languages as fragments of the corresponding infinitary classical languages?

2. For one well-known first-order preservation result we have not been able yet to obtain a modal counterpart, namely for the result that identifies first-order Horn sentences as the ones that are preserved under reduced products.

3. In a recent manuscript Johan van Benthem characterizes the (first-order) formulas defining operations on relations that preserve bisimilarity. What is the connection

between this 'safety result' and the definability and characterization results obtained here?

4. Both bisimilarity and modal equivalence cut up the universe of all model into equivalence classes. This raises the following question: when does an equivalence relation on the class of all models come from a modal language?

5. Although the above constitute a number of interesting questions, the really big question is: what makes the central equation of this paper work in the first place? How should we understand the 'take a first-order result, and bisimulate it' strategy of this paper? As a connection between two suitable categories? As a kind of duality principle?

We end on a short historical note: modal logic bisimulations have been around since Van Benthem [2]; there they are called p-relations. In the computational tradition bisimulations date back at least to Park [16]. In essence bisimulations are trimmed down versions of the Ehrenfeucht games found in first-order logic [8]. Further references, both on modal and on computational aspects can be found in Van Benthem and Bergstra [4].

Acknowledgment. I want to thank Johan van Benthem and Holger Sturm for valuable input and help in tracking down bugs in earlier versions of the paper.

The author was supported by the Netherlands Organization for Scientific Research (NWO), project NF 102/62-356 'Structural and Semantic Parallels in Natural Languages and Programming Languages'.

REFERENCES

1. J. Barwise and S. Feferman. *Model-Theoretic Logics.* Springer, New York, 1985.
2. J. van Benthem. *Modal Correspondence Theory.* PhD thesis, Mathematisch Instituut & Instituut voor Grondslagenonderzoek, Univ. of Amsterdam, 1976.
3. J. van Benthem. Temporal logic. Technical Report X-91-05, ILLC, Univ. of Amsterdam, 1991.
4. J. van Benthem and J. Bergstra. Logic of transition systems. Technical Report CT-93-03, ILLC, University of Amsterdam, 1993.
5. J. van Benthem, J. van Eijck, and V. Stebletsova. Modal logic, transition systems and processes. *Journal of Logic and Computation*, 4:811–855, 1994.
6. P. Blackburn. Nominal tense logic. *Notre D. Journ. Form. Logic*, 34:56–83, 1993.
7. C. Chang and H. Keisler. *Model Theory.* North-Holland, Amsterdam, 1973. Third edition, 1990.
8. K. Doets and J. van Benthem. Higher-order logic. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic, Volume 1*, pages 275–329. Reidel, Dordrecht, 1983.
9. R. Goldblatt. *Logics of Time and Computation.* CSLI Lecture Notes, vol. 7. CSLI Publications, Stanford, 1987.
10. R. Goldblatt. *Mathematics of Modality.* CSLI Lecture Notes, vol. 43. CSLI Publications, Stanford, 1993.

11. R. Goldblatt. Saturation and the Hennessy-Milner property. In A. Ponse, M. de Rijke, and Y. Venema, editors, *Modal Logic and Process Algebra*, CSLI Lecture Notes, Stanford, 1995. CSLI Publications. To appear.

12. M. Hennessy and R. Milner. Algebraic laws for indeterminism and concurrency. *Journal of the ACM*, 32:137–162, 1985.

13. M. Hollenberg. Hennessy-milner classes and process algebra. In A. Ponse, M. de Rijke, and Y. Venema, editors, *Modal Logic and Process Algebra*, CSLI Lecture Notes, Stanford, 1995. CSLI Publications. To appear.

14. J. Keisler. Ultraproducts and elementary classes. *Inag. Math.*, 23:477–495, 1961.

15. S. Kripke. Semantic analysis of modal logic. *Zeitschrift für mathematische Logik und Grundlagen der Mathematik*, 9:67–96, 1963.

16. D. Park. Concurrency and automata on infinite sequences. In *Proc. 5th GI Conf.*, pages 167–183. Springer, 1981.

17. M. de Rijke. A Lindström theorem for modal logic. In A. Ponse, M. de Rijke, and Y. Venema, editors, *Modal Logic and Process Algebra*, CSLI Lecture Notes, Stanford, 1995. CSLI Publications. To appear. Also appeared as Technical Report CS-R9456, CWI, Amsterdam, 1994.

18. H. Sahlqvist. Completeness and correspondence in the 1st and 2nd order semantics for modal logic. In *Proceedings of the 3rd Scandinavian Logic Symposium*, pages 110–143. North-Holland, Amsterdam, 1975.

19. D. Scott. Logic with denumerably long formulas and finite strings of quantifiers. In J. Addison et al., editor, *The Theory of Models*, pages 329–341. North-Holland, Amsterdam, 1965.