

MODEL THEORY OF MODAL LOGIC

Valentin Goranko and Martin Otto

1	Semantics of modal logic	4
1.1	Modal languages	4
1.2	Kripke frames and structures	4
1.3	The standard translations into first- and second-order logic	5
1.4	Theories and equivalence	6
1.5	Definability	7
1.6	Polyadic modalities	8
2	Bisimulation and basic model constructions	8
2.1	Bisimulation and invariance	9
2.2	Classical truth-preserving constructions	10
2.3	Proving non-definability	14
3	Bisimulation: a closer look	15
3.1	Bisimulation games	15
3.2	Finite bisimulations	16
3.3	Characteristic formulae	17
3.4	Finite model property	19
3.5	Hennessy–Milner theorem	22
3.6	Bisimulation and infinitary modal equivalence	23
3.7	Largest bisimulations as greatest fixed points	25
3.8	Bisimulation quotients and canonical representatives	26
3.9	Robinson consistency, local interpolation, and Beth definability	28
3.10	Bisimulation-safe modal operators	30
4	Modal logic as a fragment of first-order logic	30
4.1	Finite variable fragments of first-order logic	31
4.2	Van Benthem’s characterisation theorem	33
4.3	Ramifications of the characterisation theorem	36
4.4	Guarded fragments of first-order logic	37
5	Variations, extensions, and comparisons of modal logics	40
5.1	Variations through refined notions of bisimulation	41
5.2	Extensions beyond first-order	42
5.3	Model theoretic criteria	47
6	Further model-theoretic constructions	48
6.1	Ultrafilter extensions	49
6.2	Ultraproducts	50
6.3	Modal saturation and bisimulations	52
6.4	Modal definability of properties of Kripke structures	53
7	General frames	55
7.1	General frames as semantic structures in modal logic	56
7.2	Constructions and truth preservation results on general frames	57
7.3	Special types of general frames and persistence of modal formulae	59
8	Modal logic on frames	63

8.1	Modal definability of frame properties	63
8.2	Modal logic versus first-order logic on frames	64
8.3	Modal logic and first-order logic with least fixed points	69
8.4	Modal logic and second-order logic	70
9	Finite model theory of modal logics	71
9.1	Finite versus classical model theory	71
9.2	Capturing bisimulation invariant Ptime	73
9.3	0-1 laws in modal logic	74

INTRODUCTION

Model theory is about semantics; it studies the interplay between a logical language (logic) and the models (structures) for that language. Key issues therefore are *expressiveness* and *definability*. At the basic level these concern the questions which structural properties are expressible and which classes of structures are definable in the logic. These basic questions immediately lead to the study of *model constructions*; to the *analysis of models and of model classes* for given formulae or theories; to notions of *equivalence between structures* with respect to the truth of formulae; and to the study of preservation phenomena.

Modal logics¹ come as members of a loosely knit family and have various links to other logics – classical first- and second-order logic as well as, for instance, temporal and process logics stemming from particular applications. Correspondingly, the key issues mentioned above may also be studied comparatively, both within the family and in relation to other relevant logics. Such a comparative view can support an understanding of the internal coherence of the rich family of modal logics. It also offers a perspective to place modal logics in the wider logical and model theoretic context.

In regard to the coherence of the family of modal logics, it is important to understand in model theoretic terms what it is that makes a logic ‘modal’. For that aim we devote a major part of this chapter to the discussion of *bisimulation*. Many other features of the ‘modal character’ can be understood in terms of bisimulation invariance; this is true most notably of the *local* and *restricted nature of quantification*. Due to these features modal logic enjoys very specific features, and in many respects its model theory can be developed along lines that have no direct counterparts in classical model theory.

In regard to the wider logical context, there is a rich body of classical work in modal model theory that measures modal logic against the backdrop of classical first- and second-order logic into which it can be naturally embedded. But, beside this ‘classical picture’, there are also many links with other logics, partly designed for other purposes or studied with a different perspective from that of classical model theory.

In the classical picture, both first- and second-order logic have their role to play. This is because modal logic actually offers several distinct semantic levels, as will be reviewed in the following section which provides an introduction to the model theoretic semantics of modal logic. So, a modal formula is traditionally viewed in four different ways, subject to two orthogonal dichotomies – *Kripke structures* (also called *Kripke models*) versus *Kripke frames* and *local* versus *global*.

¹In this chapter we use the term *modal logic* (despite the established tradition in the literature on modal logic) in a typical model-theoretic sense, as a (propositional) modal language equipped with suitable relational (Kripke) semantics, rather than proof systems over such languages, determined by a set of axioms and inference rules, such as K, S4, etc. We refer to the latter as ‘axiomatic extensions’.

The fundamental semantic notion in basic modal logic is *truth of a formula at a state in a Kripke structure*; this notion is *local* and of a *first-order* nature. Semantics in Kripke frames is obtained, if instead one looks at all possible propositional valuations over the given frame (in effect an abstraction through implicit universal second-order quantification over all valuations); this semantics, accordingly, is of essentially *second-order* nature. On the other hand, the passage from local to *global* semantics is achieved if one looks at truth in all states (an abstraction through implicit universal first-order quantification over all states).

While all these semantic levels are ultimately based on the local semantics in Kripke structures, the two independent directions of generalisation, and in particular the divide between the (first-order) Kripke structure semantics and the (second-order) frame semantics, give rise to very distinct model theoretic flavours, each with their own tradition in the model theory of modal logic. Still, these two semantics meet through the notion of a *general frame* (closely related to a *modal algebra*).

History. The origins of model theory of modal logic go back to the fundamental papers of Jónsson and Tarski [78, 79], and Kripke [86, 87] laying the foundations of the *relational (Kripke) semantics*, followed by the classical work of Lemmon and Scott [91].

Some of the most influential themes and directions of the classical development of the model theory of modal logic in the 1970/80s have been: the *completeness theory* of modal axiomatic systems with respect to the frame-based semantics of modal logic, and the closely related *correspondence theory* between that semantics and first-order logic [117, 28, 123, 124, 113, 46, 51, 125, 127, 128]; and the *duality theory* between Kripke frames and modal algebras, via general frames [46, 47, 48, 49, 114]. Also at that time, the *theory of bisimulations and bisimulation invariance* emerged in the semantic analysis of modal languages in [125, 128]. For detailed historical and bibliographical notes see [5], and the survey [50] for a recent and comprehensive historical account of the development of modal logic, and in particular its model theory.

Overview. The sections of this chapter are roughly arranged in three parts or main tracks, reflecting the semantic distinctions outlined above.

The first part provides a common basic introduction to some of the key notions, in particular the different levels of semantics in section 1, followed by the concept of bisimulation and bisimulation respecting model constructions in section 2. This more general thread is taken up again in section 6 with some more advanced model constructions, and also in the final section 9 devoted to some ideas in the finite model theory of modal logic.

A second track, comprising sections 3 to 5, is primarily devoted to modal logic as a logic of Kripke structures (first-order semantics): section 3 continues the bisimulation theme; section 4 is specifically devoted to the role of modal logic as a fragment of first-order logic; section 5 illustrates some of the richness of modal logics over Kripke structures in terms of variations and extensions.

The third track is devoted to a study of modal logic as a logic of frames (the second-order semantics). This comprises more advanced constructions such as ultrafilter extensions and ultraproducts in section 6, basic model theory of general frames in section 7, and a survey of classical results on frame definability and relations with second-order logic in section 8.

Most of the other chapters in this handbook supplement this chapter with important

model-theoretic topics and results. In particular, we refer the reader to Chapters 1, 3, 6, 7 and 8.

1 SEMANTICS OF MODAL LOGIC

1.1 Modal languages

A (unary, poly-) *modal similarity type* is a set τ of *modalities* $\alpha \in \tau$. Beside τ , we fix a (countable) set Φ of *propositional variables* or *atomic propositions*. With τ and Φ we associate the *modal language* $\text{ML}(\tau, \Phi)$, in which every $\alpha \in \tau$ labels a modal *diamond operator* $\langle \alpha \rangle$. The formulae of $\text{ML}(\tau, \Phi)$ are recursively defined as follows:

$$\varphi := \perp \mid p \mid (\varphi_1 \rightarrow \varphi_2) \mid \langle \alpha \rangle \varphi,$$

where $p \in \Phi$ and $\alpha \in \tau$, and unnecessary outer parentheses are dropped. The logical constant $\top$ and connectives $\neg, \wedge, \vee, \leftrightarrow$ may be introduced on an equal footing or are regarded as standard abbreviations. The operator $[\alpha]$, defined by $[\alpha]\varphi := \neg\langle \alpha \rangle\neg\varphi$, is the *box operator dual to* $\langle \alpha \rangle$. A formula not containing atomic propositions is called a *constant formula*.

To keep the notation simple, we regard the set Φ as fixed, and will usually not mention it explicitly. So we write $\text{ML}(\tau)$, or also just ML when τ is clear from the context or irrelevant. We use the same notation for the set of all formulae of $\text{ML}(\tau, \Phi)$, and in general identify notationally logical languages with their sets of formulae. In the *mono-modal* case of a modal similarity type consisting of a single unary modality, the only diamond and box are denoted by just $\Diamond$ and $\Box$, respectively.

DEFINITION 1. The *nesting depth* δ of a formula is defined recursively as follows:

$$\begin{aligned} \delta(\perp) &= \delta(p) = 0; \\ \delta(\varphi_1 \rightarrow \varphi_2) &= \max(\delta(\varphi_1), \delta(\varphi_2)); \\ \delta(\langle \alpha \rangle \varphi) &= \delta(\varphi) + 1. \end{aligned}$$

The fragment $\text{ML}_n(\tau)$ comprises all formulae of $\text{ML}(\tau)$ with nesting depth $\leq n$.

1.2 Kripke frames and structures

With the modal similarity type τ we associate a *relational similarity type* consisting of binary relations R_α for $\alpha \in \tau$. For simplicity we also denote this derived relational type by τ .

DEFINITION 2. A (*Kripke*) τ -*frame* is a relational τ -structure $\mathfrak{F} = \langle W, \{R_\alpha\}_{\alpha \in \tau} \rangle$ where $W \neq \emptyset$ and $R_\alpha \subseteq W \times W$ for each $\alpha \in \tau$. The domain W of $\mathfrak{F}$ is denoted by $\text{dom}(\mathfrak{F})$. The relations $(R_\alpha)_{\alpha \in \tau}$ are the *accessibility* or *transition relations* in $\mathfrak{F}$. The elements of W , traditionally called *possible worlds*, will also be referred to, depending on the context, as *states*, *points*, or *nodes*. A *pointed τ -frame* is a pair $(\mathfrak{F}, w)$ where $w \in \text{dom}(\mathfrak{F})$.

We also write $wR_\alpha u$ rather than $R_\alpha wu$ or $(w, u) \in R_\alpha$. Given a τ -frame $\mathfrak{F} = \langle W, \{R_\alpha\}_{\alpha \in \tau} \rangle$, every R_α defines two unary operators, $\langle R_\alpha \rangle$ and its dual $[R_\alpha]$, on $\mathcal{P}(W)$ as follows:

$$\langle R_\alpha \rangle(X) := \{w \in W \mid wR_\alpha u \text{ for some } u \in X\} \quad \text{and} \quad [R_\alpha](X) := \overline{\langle R_\alpha \rangle(\overline{X})}$$

where $\overline{X} := W \setminus X$ denotes the complement of X in W .

DEFINITION 3. A *Kripke structure* (*Kripke model*) over a τ -frame $\mathfrak{F} = \langle W, \{R_\alpha\}_{\alpha \in \tau} \rangle$ is a pair $\mathfrak{M} = \langle \mathfrak{F}, V \rangle$ where $V : \Phi \rightarrow \mathcal{P}(W)$ is a *valuation*, assigning to every atomic proposition p the set of states in W where p is declared true. The set W is the *domain* of $\mathfrak{M}$, denoted $\text{dom}(\mathfrak{M})$. We often specify Kripke structures directly: $\mathfrak{M} = \langle W, \{R_\alpha\}_{\alpha \in \tau}, V \rangle$. A *pointed Kripke structure* is a pair $(\mathfrak{M}, w)$ where $w \in \text{dom}(\mathfrak{M})$.

In any Kripke structure $\mathfrak{M} = \langle \mathfrak{F}, V \rangle$ the valuation V can be extended to a valuation of all formulae, which is again denoted by V . That extension is defined recursively as follows:²

$$\begin{aligned} V(\perp) &:= \emptyset; \\ V(\varphi_1 \rightarrow \varphi_2) &:= \overline{V(\varphi_1)} \cup V(\varphi_2); \\ V(\langle \alpha \rangle \varphi) &:= \langle R_\alpha \rangle (V(\varphi)) \quad (\text{and } V([\alpha] \varphi) = [R_\alpha] (V(\varphi))). \end{aligned}$$

While first-order sentences express properties of a structure as a whole, modal formulae always make implicit reference to a distinguished (current) state in a Kripke structure. So the basic semantic notion in modal logic is *truth of a formula at a state of a Kripke structure*, with derived notions of validity also in Kripke structures and frames.

DEFINITION 4. A τ -formula φ is:

- (i) *true at the state w of the τ -structure $\mathfrak{M} = \langle \mathfrak{F}, V \rangle$* , denoted $\mathfrak{M}, w \models \varphi$, if $w \in V(\varphi)$.
This is the same as saying that φ is *true in the pointed structure* $(\mathfrak{M}, w)$.
A formula that is true at a state of some τ -structure is *satisfiable*.
- (ii) *valid in $\mathfrak{M}$* , denoted $\mathfrak{M} \models \varphi$, if $\mathfrak{M}, w \models \varphi$ for every $w \in \text{dom}(\mathfrak{F})$, i.e., if $V(\varphi) = \text{dom}(\mathfrak{F})$.
- (iii) *(locally) valid at the state w of $\mathfrak{F}$* , denoted $\mathfrak{F}, w \models \varphi$, if $\mathfrak{M}, w \models \varphi$ for every τ -structure $\mathfrak{M}$ over $\mathfrak{F}$.
This is the same as saying that φ is *valid in the pointed frame* $(\mathfrak{F}, w)$.
- (iv) *valid in $\mathfrak{F}$* , denoted $\mathfrak{F} \models \varphi$, if $\mathfrak{F}, w \models \varphi$ for every $w \in \text{dom}(\mathfrak{F})$.
Equivalently: $\mathfrak{M} \models \varphi$ for every τ -structure $\mathfrak{M}$ over $\mathfrak{F}$.
- (v) *valid*, denoted $\models \varphi$, if $\mathfrak{F} \models \varphi$ for every τ -frame $\mathfrak{F}$.

1.3 The standard translations into first- and second-order logic

With the modal language $\text{ML}(\tau, \Phi)$, we associate the following purely relational vocabularies:

- the relational version of τ itself, consisting of R_α for $\alpha \in \tau$, and again denoted by just τ .
- the expansion τ_Φ of the relational vocabulary τ by unary predicates $\{P_0, P_1, \dots\}$ associated with the atomic propositions $p_0, p_1, \dots \in \Phi$.

Correspondingly, $\text{FO}(\tau)$ and $\text{FO}(\tau_\Phi)$ are the first-order languages with vocabularies τ and τ_Φ , respectively. We regard a τ -frame as a τ -structure in the usual sense, and a Kripke structure over a τ -frame as a τ_Φ -structure, with P_i interpreted as $V(p_i)$. We use the same notation for Kripke structures and for the associated first-order structures, as

²In algebraic terms (see Chapter 6), the extended valuation is the unique homomorphism from the free τ -algebra of formulae to the modal algebra associated with the model $\mathfrak{M}$, extending V .

this causes no confusion. Wherever necessary, we will highlight the distinction by writing $\models_{\text{FO}}$ to explicitly appeal to first-order semantics.

Truth and validity of a modal formula in a Kripke structure are *first-order notions* in the following sense. Let $\text{VAR} = \{x_0, x_1, \dots\}$ be the set of first-order variables of $\text{FO}(\tau_\Phi)$. The formulae of $\text{ML}(\tau)$ are translated into $\text{FO}(\tau_\Phi)$ by means of the following *standard translation* [124, 127], parameterised with the variables from VAR :

- $\text{ST}(p_i; x_j) := P_i x_j$ for every $p_i \in \Phi$;
- $\text{ST}(\perp; x_j) := \perp$;
- $\text{ST}(\varphi_1 \rightarrow \varphi_2; x_j) := \text{ST}(\varphi_1; x_j) \rightarrow \text{ST}(\varphi_2; x_j)$;
- $\text{ST}(\langle \alpha \rangle \varphi; x_j) := \exists y (x_j R_\alpha y \wedge \text{ST}(\varphi; y))$, where y is the first variable in $\text{VAR} \setminus \{x_j\}$.

Note that only x_j is free in $\text{ST}(\varphi; x_j)$. Furthermore, for the standard translation it suffices to use only the variables x_0 and x_1 (free or bound) in an alternating fashion. This yields a translation into the *two-variable fragment* FO^2 of first-order logic. Also, the standard translation of any modal formula falls into the *guarded fragment* of first-order logic. These observations are taken up in section 4.

The standard translation is semantically faithful in the following sense.

PROPOSITION 5. *For every pointed Kripke structure $(\mathfrak{M}, w)$ and $\varphi \in \text{ML}(\tau)$,*

$$\mathfrak{M}, w \models \varphi \text{ iff } \mathfrak{M}, w \models_{\text{FO}} \text{ST}(\varphi; x_0).$$

While the semantics and validity for modal formulae over Kripke structures is thus essentially first-order, validity of a modal formula in a frame goes beyond first-order logic. Indeed, paraphrasing the definition in terms of the standard translation, a modal formula φ is valid in a frame iff its standard translation is true in that frame under *every* interpretation of the unary predicates occurring in it.

PROPOSITION 6. *For every pointed Kripke frame $(\mathfrak{F}, w)$ and $\varphi \in \text{ML}(\tau)$ with atomic propositions among $p_0, \dots, p_n$:*

$$\mathfrak{F}, w \models \varphi \text{ iff } \mathfrak{F}, w \models \forall P_0 \dots \forall P_n \text{ST}(\varphi; x_0).$$

Consequently, $\mathfrak{F} \models \varphi$ iff $\mathfrak{F} \models \forall P_0 \dots \forall P_n \forall x_0 \text{ST}(\varphi; x_0)$.

1.4 Theories and equivalence

With every logic $\mathcal{L}$ comes an associated notion of logical equivalence between structures. Two structures of the appropriate type are equivalent with respect to $\mathcal{L}$ if no property expressible in $\mathcal{L}$ distinguishes between them, i.e., if their $\mathcal{L}$ -theories are the same. In this sense, first-order logic gives rise to the notion of elementary equivalence. Correspondingly, *modal equivalence* is indistinguishability in modal logic. Each view of the semantics of modal logic – in terms of (pointed or plain) structures or frames – corresponds to a notion of modal theories and modal equivalence.

DEFINITION 7. The *modal theory* of a pointed Kripke τ -structure $(\mathfrak{M}, w)$ is the set of all formulae of $\text{ML}(\tau)$ satisfied in $(\mathfrak{M}, w)$: $\text{Th}_{\text{ML}}(\mathfrak{M}, w) := \{\varphi \in \text{ML}(\tau) \mid \mathfrak{M}, w \models \varphi\}$.

Correspondingly, the modal theory of $\mathfrak{M}$ is $\text{Th}_{\text{ML}}(\mathfrak{M}) := \{\varphi \in \text{ML}(\tau) \mid \mathfrak{M} \models \varphi\}$. The modal theories of a frame and pointed frame, as well as of classes of (pointed) Kripke structures or frames, are defined likewise.

The basic notion of modal equivalence, corresponding to the notion of truth at a state of a Kripke structure, is an equivalence relation on the class of *pointed Kripke structures* $(\mathfrak{M}, w)$. Natural variants cover the derived notions for plain Kripke structures, and for pointed or plain frames.

DEFINITION 8. For two pointed Kripke τ -structures $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$: $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ are *ML-equivalent*, denoted $(\mathfrak{M}, w) \equiv_{\text{ML}} (\mathfrak{M}', w')$, iff they satisfy exactly the same formulae of ML, i.e., iff $\text{Th}_{\text{ML}}(\mathfrak{M}, w) = \text{Th}_{\text{ML}}(\mathfrak{M}', w')$. Modal equivalence between Kripke structures, frames, and pointed frames are defined likewise.

1.5 Definability

Definability in modal logic means different things corresponding to the different levels of the semantics. We distinguish local versus global definability (truth at a state versus validity throughout a frame/structure), and definability at the level of structures versus frames (truth/validity for a given valuation versus for all valuations).

Given a formula $\varphi \in \text{ML}(\tau)$, the classes of pointed Kripke structures, Kripke structures, pointed frames and frames defined by φ are denoted as $\text{KS}(\varphi)$, $\text{PKS}(\varphi)$, $\text{FR}(\varphi)$, and $\text{PFR}(\varphi)$, respectively:

$$\begin{aligned} \text{PKS}(\varphi) &= \{(\mathfrak{M}, w) \mid \mathfrak{M}, w \models \varphi\} & \text{PFR}(\varphi) &= \left\{ (\mathfrak{F}, w) \mid \begin{array}{l} (\mathfrak{F}, V), w \models \varphi \\ \text{for all valuations } V \end{array} \right\} \\ \text{KS}(\varphi) &= \{\mathfrak{M} \mid \mathfrak{M} \models \varphi \text{ for all } w \in \text{dom}(\mathfrak{M})\} & \text{FR}(\varphi) &= \left\{ \mathfrak{F} \mid \begin{array}{l} (\mathfrak{F}, V), w \models \varphi \text{ for all } w \\ \text{and for all valuations } V \end{array} \right\} \end{aligned}$$

DEFINITION 9. A class $\mathcal{P}$ of pointed Kripke τ -structures is (*modally*) *definable* in the language $\text{ML}(\tau)$ if $\mathcal{P} = \text{PKS}(\varphi)$ for some formula $\varphi \in \text{ML}(\tau)$. Definable classes of Kripke structures, frames, and pointed frames are defined likewise.

EXAMPLE 10. Here are some examples of modally definable classes of Kripke frames and structures.

The class of pointed Kripke structures $(\mathfrak{M}, w)$, where $\mathfrak{M} = \langle W, R, V \rangle$, such that w has at least one successor not satisfying p for which every successor satisfies q , is defined by the formula $\Diamond(\neg p \wedge \Box q)$.

The formula $p \rightarrow \Box p$ defines the class of Kripke structures in which the valuation of p is closed under the accessibility relation.

The class of frames in which every state has a successor is defined by the formula $\Diamond \top$; the same formula defines the class of pointed frames $(\mathfrak{F}, w)$ in which w has a successor.

The formula $\Diamond p \rightarrow \Box p$ defines the class of frames $\mathcal{K}$ in which every state has at most one successor. It is straightforward to show that the formula is valid in every such frame. For the converse: if the formula fails at some state w of a Kripke structure over a frame $\mathfrak{F}$, then p is true at some successor of w . But since $\Box p$ is false at w , there must be another successor of w where p fails. Hence $\mathfrak{F}$ does not satisfy the defining property of $\mathcal{K}$.

Other standard examples of modally definable classes of frames include the classes of: reflexive frames, defined by $\Box p \rightarrow p$; transitive frames, defined by $\Box p \rightarrow \Box \Box p$; symmetric frames, defined by $\Diamond \Box p \rightarrow p$, etc. For more examples see [117, 75, 127, 128].

Proposition 5 implies that the definability of classes of (pointed) Kripke structures by modal formulae is a special case of first-order definability. Consequently, modal logic shares many basic model-theoretic results with first-order logic, such as compactness and Löwenheim–Skolem theorems (see [12, 68]). We will discuss the model theoretic aspects of modal logic as a fragment of first-order logic on Kripke structures in section 4.

On the other hand, Proposition 6 indicates that modal definability of (pointed) frames is a form of Π_1^1 -definability, and the model-theoretic consequences of that fact will be discussed in section 8. In particular, we will see that it is indeed essentially second-order.

1.6 Polyadic modalities

In polyadic modal logics one considers modalities α of arbitrary arities $r(\alpha) \in \mathbb{N}$, which give rise to formulae $\langle \alpha \rangle(\varphi_1, \dots, \varphi_n)$ if $n = r(\alpha)$. The interpretation of an n -ary modal operator α is given in terms of $(n + 1)$ -ary relations R_α in corresponding frames, and an n -ary operator on subsets of these frames, in such a way that the semantics is faithfully captured in the standard translation $\text{ST}(\langle \alpha \rangle(\varphi_1, \dots, \varphi_n; x_j)) := \exists y_1 \dots \exists y_n (x_j R_\alpha y_1 \dots y_n \wedge \bigwedge_{i=1}^n \text{ST}(\varphi_i; y_i))$, where $y_1 \dots y_n$ are the first n variables in $\text{VAR} \setminus \{x_j\}$ (and $x_j R_\alpha y_1 \dots y_n$ is just a notational variant for $R_\alpha x_j y_1 \dots y_n$).

Polyadic modalities were first studied from an algebraic perspective, as normal and additive operators in Boolean algebras, by Jónsson and Tarski [78, 79]. All the essential model theoretic features of modal logic can be generalised to this more liberal setting, albeit with some care and sometimes unavoidable notational complications. In [41] Goguadze et al define and develop systematically an interpretation of polyadic languages into monadic ones, and simulations of polyadic by monadic logics, which transfer a number of important properties, such as frame completeness, finite model property, canonicity and first-order definability. On the other hand, so called *purely modal polyadic languages* are defined in [54], where all logical connectives except negation are treated as binary modalities, and modalities can be composed. Thus, all polyadic modal formulae are built from (composite) boxes and diamonds applied to literals, making their syntactic structure much simpler.

Throughout this chapter we will only treat monadic modalities explicitly.

2 BISIMULATION AND BASIC MODEL CONSTRUCTIONS

A major concern in model theory is the analysis of logical equivalence of structures in comparison with other natural notions of structural equivalence, in particular equivalences of a more combinatorial or algebraic nature. Bisimulation equivalences as studied below prove to be the algebraic/combinatorial counterparts to modal equivalence.

For first-order logic this combinatorial approach leads to the well-known characterisation of elementary equivalence via *Ehrenfeucht–Fraïssé games* (see [68, 108, 26, 25]). Variations of the basic Ehrenfeucht–Fraïssé idea apply to many other logics including modal logic. Modal equivalence can thus be put into the general Ehrenfeucht–Fraïssé framework. We shall sketch this connection in section 4. The very natural game associated with modal equivalence has, however, also been invented and studied independently and in its own right, with the notions of *zig-zag relation* (van Benthem) and *bisimulation equivalence* (Hennessy, Milner, Park). We therefore put an autonomous, modal treatment

before the discussion of relationships with the general framework of Ehrenfeucht–Fraïssé and pebble games.

2.1 Bisimulation and invariance


While the notion of logical equivalence is static, it can often be characterised in more dynamic, game-theoretic terms. The concept of bisimulation equivalence, which is closely related to corresponding games, is one of the most productive ideas in the model theory of modal logics, temporal logics, logics for concurrency, etc. Just as it has multiple roots in these various branches of logic, many variants have been employed to capture specific notions of “behavioural equivalence” between all kinds of transition systems that are interesting in their own right for various application areas – and not necessarily with any ‘logic’ in mind.

DEFINITION 11. Let $\mathfrak{M} = \langle W, \{R_\alpha\}_{\alpha \in \tau}, V \rangle$ and $\mathfrak{M}' = \langle W', \{R'_\alpha\}_{\alpha \in \tau}, V' \rangle$ be two Kripke τ -structures. A *bisimulation* between $\mathfrak{M}$ and $\mathfrak{M}'$ is a non-empty relation $\rho \subseteq W \times W'$ satisfying the following conditions for any $w\rho w'$:

Atom equivalence: w and w' satisfy the same atomic propositions, hereafter denoted by $w \simeq w'$.

Forth: For any $\alpha \in \tau$, if $wR_\alpha u$ for some $u \in W$, then there is some $u' \in W'$ such that $w'R'_\alpha u'$ and $u\rho u'$. (Any α -transition at w in $\mathfrak{M}$ can be matched at w' in $\mathfrak{M}'$.)

Back: Similarly, in the opposite direction: for any $\alpha \in \tau$, and $w'R'_\alpha u'$ there is some $u \in W$ such that $wR_\alpha u$ and $u\rho u'$. (Any α -transition at w' in $\mathfrak{M}'$ can be matched at w in $\mathfrak{M}$.)



That ρ is a bisimulation between $\mathfrak{M}$ and $\mathfrak{M}'$ is denoted as $\rho: \mathfrak{M} \rightleftharpoons \mathfrak{M}'$. If, moreover, ρ is such that every element in $\mathfrak{M}$ is linked to some element of $\mathfrak{M}'$ and vice versa, we say that ρ is a *global bisimulation* and that $\mathfrak{M}$ and $\mathfrak{M}'$ are *globally bisimilar*.

DEFINITION 12. Two pointed Kripke structures $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ are *bisimilar* or *bisimulation equivalent*, denoted $(\mathfrak{M}, w) \rightleftharpoons (\mathfrak{M}', w')$, if there is a bisimulation ρ between $\mathfrak{M}$ and $\mathfrak{M}'$ such that $w\rho w'$.

Bisimulations between (pointed) frames can be defined likewise, by omitting *atom equivalence*. Thus, a relation ρ is a bisimulation between two frames $\mathfrak{F}$ and $\mathfrak{F}'$, iff it is a bisimulation between the respective Kripke structures $\langle \mathfrak{F}, V_\perp \rangle$ and $\langle \mathfrak{F}', V'_\perp \rangle$ where the valuations $V_\perp$ and $V'_\perp$ render every atomic proposition false at every state of the respective frame.

DEFINITION 13. Let $\mathcal{C}$ be a class of structures appropriate for the logical language $\mathcal{L}$ (e.g., pointed Kripke structures for ML). Let $\approx$ be an equivalence relation on $\mathcal{C}$. Then $\mathcal{L}$ is *preserved under $\approx$* over $\mathcal{C}$, or $\mathcal{L}$ is *$\approx$ -invariant* over $\mathcal{C}$, iff for any $\mathfrak{A} \approx \mathfrak{A}'$ in $\mathcal{C}$ and

any $\varphi \in \mathcal{L}$: $\mathfrak{A} \models \varphi \Leftrightarrow \mathfrak{A}' \models \varphi$, i.e., $\mathfrak{A}$ and $\mathfrak{A}'$ are $\mathcal{L}$ -equivalent. In other words: $\approx$ is a refinement of $\equiv_{\mathcal{L}}$, or $\approx \subseteq \equiv_{\mathcal{L}}$.

Invariance phenomena give insights into the semantics of the logic involved, and also often provide key tools for the model theoretic study of the logic (e.g., model constructions guided by $\approx$ equivalence). The relationship between modal logics and bisimulation equivalences provides an excellent example of such a fruitful companionship.

It would be straightforward to prove the following by induction on the structure of modal formulae, straight from the definition of bisimulations. However, this will also fall out as a corollary of the more instructive analysis of the associated bisimulation games. We therefore meanwhile only state the fact.

THEOREM 14 (bisimulation invariance).

$\text{ML}(\tau)$ is bisimulation invariant: if $(\mathfrak{M}, w) \rightleftharpoons (\mathfrak{M}', w')$ then $(\mathfrak{M}, w) \equiv_{\text{ML}} (\mathfrak{M}', w')$.

Consequently, for every constant formula $\theta \in \text{ML}(\tau)$ and pointed τ -frames $(\mathfrak{F}, w)$ and $(\mathfrak{F}', w')$: if $(\mathfrak{F}, w) \rightleftharpoons (\mathfrak{F}', w')$, then $(\mathfrak{F}, w) \models \theta$ iff $(\mathfrak{F}', w') \models \theta$.

2.2 Classical truth-preserving constructions

Bisimulations induced by maps from one frame to another have classically been studied as *bounded morphisms* or *p-morphisms*. We state the corresponding back-and-forth conditions, which are slightly simpler in the case of such a functional relationship, and treat some particularly important special cases. The use of *generated and rooted substructures*, *bounded morphic images*, *tree unfoldings* and *disjoint unions* in connection with classical model constructions for modal logic is based on truth preservation for modal formulae. These constructions were introduced for basic modal logic [117, 7] before the notion of bisimulation was developed and its importance for modal logic realised. Via *duality theory*, which connects the relational semantics for modal logic with an algebraic semantics, bounded morphisms, generated subframes and disjoint unions correspond respectively to the fundamental universal algebraic notions of subalgebras, homomorphic images, and direct products. For details see Chapter 6 of this handbook, as well as [78, 47, 48, 114] and [5, Ch. 5].

The preservation results encountered in these special cases of a passage to bisimilar structures highlight to various degrees one of the key characteristic features of the semantics of modal logic: its *explicit locality* and *restricted nature of quantification*. Unlike first-order logic, whose global quantification over the entire universe makes truth generally dependent on the entire structure, the truth of a modal formula in a Kripke structure is evaluated relative to a ‘current’ state and admits access to the rest of the structure only along the edges of the accessibility relations.

Passage from a given structures to a bisimilar tree structure, obtained via a simple bounded morphism, shows for instance that any satisfiable formula of basic modal logic is satisfied at the root of a tree structure (tree model property, see Corollary 24; this can be further strengthened to a finite tree model property, see Lemma 35). Conversely, preservation results can be used to show that certain properties are not modally definable. We shall see some classical examples of this in section 2.3.

Bounded morphisms

DEFINITION 15. Let $\mathfrak{M} = \langle W, \{R_\alpha\}_{\alpha \in \tau}, V \rangle$ and $\mathfrak{M}' = \langle W', \{R'_\alpha\}_{\alpha \in \tau}, V' \rangle$ be Kripke structures. A function $\rho: W \rightarrow W'$ is a *bounded morphism* from $\mathfrak{M}$ to $\mathfrak{M}'$ if its graph is a bisimulation between $\mathfrak{M}$ and $\mathfrak{M}'$. We denote a bounded morphism as in $\rho: \mathfrak{M} \xrightarrow{\text{b}} \mathfrak{M}'$.

Bounded morphisms between frames are similarly defined.

If ρ is onto, then $\mathfrak{M}'$ is a *bounded morphic image* of $\mathfrak{M}$ (and similarly for frames).

Thus, for each $u \in W$, a bounded morphism ρ uniquely singles out a bisimilar state $\rho(u)$ in W' . The bisimulation conditions for a bounded morphism between two Kripke structures correspondingly become:

Atom equivalence: $w \simeq \rho(w)$ for every $w \in W$.

Forth: For any $w \in W$ and $\alpha \in \tau$, if $wR_\alpha u$ for some $u \in W$, then $\rho(w)R'_\alpha \rho(u)$.

Back: For any $w \in W$ and $\alpha \in \tau$, if $\rho(w)R'_\alpha u'$ for some $u' \in W'$ then $u' = \rho(u)$ for some $u \in W$ such that $wR_\alpha u$.

Bisimulation invariance yields the following preservation results.

COROLLARY 16. *Bounded morphisms preserve truth and validity of modal formulae. More specifically, if $\rho: \mathfrak{M} \xrightarrow{\text{b}} \mathfrak{M}'$ is a bounded morphism and $\varphi \in \text{ML}(\tau)$, then:*

1. *for all $u \in \text{dom}(\mathfrak{F})$: $\mathfrak{M}, u \models \varphi$ iff $\mathfrak{M}', \rho(u) \models \varphi$.*
2. *If ρ is onto, then $\mathfrak{M} \models \varphi$ iff $\mathfrak{M}' \models \varphi$, i.e., $\text{Th}_{\text{ML}}(\mathfrak{M}) = \text{Th}_{\text{ML}}(\mathfrak{M}')$.*
3. *If $\mathfrak{F}, u \models \varphi$ then $\mathfrak{F}', \rho(u) \models \varphi$.*
4. *If ρ is onto, then $\mathfrak{F} \models \varphi$ implies $\mathfrak{F}' \models \varphi$.*

For the latter two claims one just has to note that each model $\mathfrak{M}' = \langle \mathfrak{F}', V' \rangle$ over the frame $\mathfrak{F}'$ can be pulled back to give a model $\mathfrak{M} = \langle \mathfrak{F}, V \rangle$ over the frame $\mathfrak{F}$ via $V(p) := \rho^{-1}[V'(p)] = \{w \in \text{dom}(\mathfrak{F}) \mid \rho(w) \in V'(p)\}$. This turns ρ into a bounded morphism from $\mathfrak{M}$ to $\mathfrak{M}'$. Note, however, that not every model over $\mathfrak{F}$ is obtained in this manner.

We turn to several basic model constructions involving bounded morphisms: generated substructures, rooted substructures, tree unfoldings and disjoint unions.

Generated and rooted substructures

If $R \subseteq W^2$ is any binary relation over W , and $W' \subseteq W$, we write $R \upharpoonright W'$ for the restriction of R to W' , $R \upharpoonright W' = R \cap (W' \times W')$. Similarly for a valuation V on W , $V \upharpoonright W'$ stands for its restriction to W' .

DEFINITION 17. Let $\mathfrak{F} = \langle W, \{R_\alpha\}_{\alpha \in \tau} \rangle$ be a frame, or $\mathfrak{M} = \langle \mathfrak{F}, V \rangle$ a Kripke structure over $\mathfrak{F}$, respectively, and $W' \subseteq W$.

- (i) The induced *subframe of $\mathfrak{F}$ over W'* is the frame $\mathfrak{F}' := \mathfrak{F} \upharpoonright W' = \langle W', \{R_\alpha \upharpoonright W'\}_{\alpha \in \tau} \rangle$. The subframe relationship is denoted $\mathfrak{F}' \leq \mathfrak{F}$.
- (ii) $\mathfrak{F}' = \mathfrak{F} \upharpoonright W'$ is a *generated subframe of $\mathfrak{F}$* , denoted $\mathfrak{F}' \leq \mathfrak{F}$, if W' is closed under all accessibility relations in the sense that $wR_\alpha u$ for $w \in W'$ implies $u \in W'$.

- (iii) The induced *substructure of $\mathfrak{M}$ over W'* is the Kripke structure $\mathfrak{M}' = \mathfrak{M} \upharpoonright W' = \langle \mathfrak{F} \upharpoonright W', V \upharpoonright W' \rangle$, denoted $\mathfrak{M}' \leq \mathfrak{M}$. If $\mathfrak{F} \upharpoonright W' \leq \mathfrak{F}$, then $\mathfrak{M}'$ is a *generated substructure of $\mathfrak{M}$* , denoted $\mathfrak{M}' \trianglelefteq \mathfrak{M}$.

Obviously, for $\mathfrak{M}' \trianglelefteq \mathfrak{M}$ the inclusion map $\rho: W' \rightarrow W$ is a bounded morphism. By bisimulation invariance, we therefore have the following.

PROPOSITION 18. *For all Kripke structures $\mathfrak{M}' \trianglelefteq \mathfrak{M}$ and for every formula φ of $ML(\tau)$:*

(i) *for every $u \in \text{dom}(\mathfrak{M}')$: $\mathfrak{M}, u \models \varphi$ iff $\mathfrak{M}', u \models \varphi$.*

(ii) *$\mathfrak{M} \models \varphi$ implies $\mathfrak{M}' \models \varphi$.*

Likewise, for frames $\mathfrak{F}' \trianglelefteq \mathfrak{F}$ and $u \in \text{dom}(\mathfrak{F}')$: $\mathfrak{F}, u \models \varphi$ iff $\mathfrak{F}', u \models \varphi$, and $\mathfrak{F} \models \varphi$ implies $\mathfrak{F}' \models \varphi$.

The latter claim holds since every Kripke structure over $\mathfrak{F}'$ is induced by a Kripke structure on $\mathfrak{F}$.

A particularly important case of generated subframes deals with the set of all states reachable from a fixed state. A *path* in a frame $\mathfrak{F} = \langle W, \{R_\alpha\}_{\alpha \in \tau} \rangle$ is a sequence $\vec{w} = (w_0, \alpha_1, w_1, \dots, \alpha_k, w_k)$, where $w_{i-1} R_{\alpha_i} w_i$ for $i = 1, \dots, k$ (this path is *rooted* at w_0 and has *length* k). A path of length $k = 0$, $\vec{w} = (w_0)$, is identified with its root w_0 . For $u \in W$, we denote the set of all paths rooted at u by $\vec{W}[u]$. For every path $\vec{w}$ as above we define the ‘terminal state’ function $f(\vec{w}) = w_k$ where k is the length of $\vec{w}$. Then

$$W[u] := \{f(\vec{w}) \mid \vec{w} \in \vec{W}[u]\}$$

is the set of all states in $\mathfrak{F}$ reachable from u (including u itself).

DEFINITION 19. Let $\mathfrak{F} = \langle W, \{R_\alpha\}_{\alpha \in \tau} \rangle$ be a frame, $\mathfrak{M} = \langle \mathfrak{F}, V \rangle$ a Kripke structure over $\mathfrak{F}$, and $u \in W$.

(i) The *subframe of $\mathfrak{F}$ rooted at u* is the frame $\mathfrak{F}[u] = \mathfrak{F} \upharpoonright W[u]$.

(ii) The *substructure of $\mathfrak{M}$ rooted at u* is the Kripke structure $\mathfrak{M}[u] = \mathfrak{M} \upharpoonright W[u]$.

(iii) $\mathfrak{F}$ (respectively $\mathfrak{M}$) is *rooted at u* if $W[u] = W$.

Clearly, for any $u \in W$: $\mathfrak{F}[u] \trianglelefteq \mathfrak{F}$ and $\mathfrak{M}[u] \trianglelefteq \mathfrak{M}$, respectively. Therefore, we obtain the following.

COROLLARY 20. *For every Kripke structure $\mathfrak{M}' \trianglelefteq \mathfrak{M}$ and formula φ of $ML(\tau)$:*

(i) *for all $u \in W$: $\mathfrak{M}, u \models \varphi$ iff $\mathfrak{M}[u], u \models \varphi$.*

(ii) *$\mathfrak{M} \models \varphi$ implies $\mathfrak{M}[u] \models \varphi$.*

(iii) *Likewise for (pointed) frames.*

Thus, any satisfiable formula is satisfiable at the root of a rooted Kripke structure.

Tree unfoldings

An important model construction based on a canonical bounded morphism is the *unfolding* or *tree unravelling* of a Kripke structure $\mathfrak{M} = \langle W, \{R_\alpha\}_{\alpha \in \tau}, V \rangle$ from some $u \in W$. This construction was introduced in [113], where the tree model property (cf. Corollary 24 below) was proved, too.

Recall the map $f: \vec{W}[u] \rightarrow W$, which maps the path $\vec{w} = (u = w_0, \alpha_1, w_1, \dots, \alpha_k, w_k)$ to its terminal state $f(\vec{w}) = w_k$. The unfolding of $\vec{\mathfrak{M}}[u]$ of $\mathfrak{M}$ at u is based on the set $\vec{W}[u]$ of all paths rooted at u , with the natural definition of accessibility relations and a valuation that turns f into a bounded morphism.

DEFINITION 21. The *unfolding* (or, *unravelling*) of $\mathfrak{M} = \langle W, \{R_\alpha\}_{\alpha \in \tau}, V \rangle$ from some $u \in W$ is the rooted Kripke structure $\vec{\mathfrak{M}}[u] := \langle \vec{W}[u], \{\vec{R}_\alpha\}_{\alpha \in \tau}, \vec{V} \rangle$ with root $u = (u)$, where

$$\begin{aligned} \vec{R}_\alpha &:= \{(\vec{w}, (\vec{w}, \alpha, w')) \mid \vec{w} \in \vec{W}[u], f(\vec{w})R_\alpha w'\}, \\ \vec{V}(p) &:= f^{-1}[V(p)]. \end{aligned}$$

Indeed, $\vec{\mathfrak{M}}[u]$ is a tree structure with root $u = (u)$ in the sense of the following definition.

DEFINITION 22. A pointed frame $\mathfrak{F} = \langle W, \{R_\alpha\}_{\alpha \in \tau} \rangle$ with distinguished state $u \in W$ is a *tree* with root u if $\mathfrak{F}$ is rooted at u and every state $w \in W$ is reachable from u by a unique path. Accordingly, every Kripke structure over $(\mathfrak{F}, u)$ is a tree structure.

OBSERVATION 23. For every pointed Kripke structure $(\mathfrak{M}, u)$ the terminal state map $f: \vec{W}[u] \rightarrow W[u]$ is a bounded morphism of the unfolding $\vec{\mathfrak{M}}[u]$ onto $\mathfrak{M}[u]$.

As $\vec{\mathfrak{M}}[u]$ is a tree structure with root $u = (u)$, we obtain the following.

COROLLARY 24 (tree-model property). *Every satisfiable modal formula is satisfiable at the root of a tree.*

Also compare Lemma 35 below, where this is strengthened to a *finite tree* model property.

Disjoint unions

Disjoint unions are well known for relational structures: the component structures are put side by side without any relational links between the components. Assuming that the given family of Kripke structures or frames is based on universes that are pairwise disjoint, we may just take the set-theoretic union of the universes, accessibility relations, and valuations, respectively. If the given frames are not disjoint, they first need to be replaced by isomorphic copies over universes that are pairwise disjoint.

To be specific, define the *disjoint union* of an arbitrary family $\{W^i\}_{i \in I}$ of (not necessarily disjoint) sets as $\uplus_{i \in I} W^i := \bigcup_{i \in I} (W^i \times \{i\})$. With this formalisation, we have the natural injection or embedding $\varepsilon_j: W^j \rightarrow \bigcup_{i \in I} (W^i \times \{i\})$ of each component set into the disjoint union, which maps $w \in W^j$ to $(w, j) \in \bigcup_{i \in I} (W^i \times \{i\})$.

DEFINITION 25. Consider a family of τ -frames $\{\mathfrak{F}^i = \langle W^i, \{R_\alpha^i\}_{\alpha \in \tau} \rangle\}_{i \in I}$ and a family of Kripke structures $\{\mathfrak{M}^i = \langle \mathfrak{F}^i, V^i \rangle\}_{i \in I}$ over these.

- (i) The *disjoint union* of $\{\mathfrak{F}^i\}_{i \in I}$ is the frame $\uplus_{i \in I} \mathfrak{F}^i = \langle \uplus_{i \in I} W^i, \{R_\alpha\}_{\alpha \in \tau} \rangle$, where $(w_0, i_0)R_\alpha(w_1, i_1)$ iff $i_0 = i_1 = i$ and $w_0R_\alpha^i w_1$.
- (ii) The *disjoint union* of $\{\mathfrak{M}^i\}_{i \in I}$ is the Kripke τ -structure $\uplus_{i \in I} \mathfrak{M}^i = \langle \uplus_{i \in I} \mathfrak{F}^i, V \rangle$ where $V(p) = \uplus_{i \in I} V^i(p)$.

It is immediate that the natural injection $\varepsilon_j: W^j \rightarrow \uplus_{i \in I} W^i$ isomorphically embeds $\mathfrak{M}^j$ into $\uplus_{i \in I} \mathfrak{M}^i$ and is indeed a bounded morphism with image $\varepsilon_j[\mathfrak{M}^j] \simeq \mathfrak{M}^j$ and

$\epsilon_j[\mathfrak{M}^j] \trianglelefteq \biguplus_{i \in I} \mathfrak{M}^i$. We therefore obtain the following, by bisimulation invariance and based on previous observations.

PROPOSITION 26. *Given a family of τ -frames $\{\mathfrak{F}^i = \langle W^i, \{R_\alpha^i\}_{\alpha \in \tau} \rangle\}_{i \in I}$, a family of Kripke structures $\{\mathfrak{M}^i = \langle \mathfrak{F}^i, V^i \rangle\}_{i \in I}$ over these frames, and $\varphi \in \text{ML}(\tau)$:*

1. *For every $j \in I$ and $w \in \text{dom}(\mathfrak{M}^j)$: $\mathfrak{M}^j, w \models \varphi$ iff $\biguplus_{i \in I} \mathfrak{M}^i, (w, j) \models \varphi$.*
2. *For every $j \in I$ and $w \in \text{dom}(\mathfrak{F}^j)$: $\mathfrak{F}^j, w \models \varphi$ iff $\biguplus_{i \in I} \mathfrak{F}^i, (w, j) \models \varphi$.*
3. *$\biguplus_{i \in I} \mathfrak{M}^i \models \varphi$ iff $\mathfrak{M}^i \models \varphi$ for every $i \in I$.*
4. *$\biguplus_{i \in I} \mathfrak{F}^i \models \varphi$ iff $\mathfrak{F}^i \models \varphi$ for every $i \in I$.*

The following structural observation [127] links some of the ideas explored in this section.

PROPOSITION 27. *Any Kripke structure is the bounded morphic image of a disjoint union of rooted Kripke structures, and indeed of tree structures.*

For a proof of the proposition, consider the families of the $\{\mathfrak{M}[u]\}_{u \in W}$ or $\{\vec{\mathfrak{M}}[u]\}_{u \in W}$. The desired bounded morphisms (from the disjoint unions of these families back onto $\mathfrak{M}$) are the unions of the projection and inclusion or terminal state maps defined on the components of these disjoint unions.

2.3 Proving non-definability

To show that a given property of structures is definable in a given logic, it suffices simply to find a defining formula. Showing that a property is *not* definable, however, is not so straightforward, and often requires elaborate arguments. A standard method for establishing non-definability of a property $\mathcal{P}$ (i.e., of the class of structures satisfying that property) in a logic $\mathcal{L}$ is to show that $\mathcal{P}$ is not closed under some construction preserving truth (validity) of all formulae of $\mathcal{L}$. Now that we have at hand constructions that preserve truth and validity of modal formulae, we can use them to show that various properties of frames and structures are not modally definable. Compare Definition 9 for the relevant notions of definability.

At the level of pointed Kripke structures, modal formulae capture only properties that are local in the sense that whether or not $\mathfrak{M}, w \models \varphi$ only depends on $(\mathfrak{M}[w], w)$. In other words, modal formulae are incapable of expressing any property of $(\mathfrak{M}, w)$ that involves points beyond $\mathfrak{M}[w]$. For instance, there is no $\varphi \in \text{ML}$ such that $\mathfrak{M}, w \models \varphi$ iff $\mathfrak{M} \models p$. Indeed, one can always add to $\mathfrak{M}$ an extra point (as a disjoint union), not reachable from w , where p is false. The resulting pointed structure $(\mathfrak{M}', w)$ is bisimilar to $(\mathfrak{M}, w)$, whence φ would have to be equally true or false at w in both.

Likewise, at the level of Kripke structures, there is no $\varphi \in \text{ML}$ such that $\mathfrak{M} \models \varphi$ iff the accessibility relation of the underlying frame $\mathfrak{F}$ is reflexive. This follows for instance from the fact that the unfolding of any frame is irreflexive. If $\mathfrak{M}$ is reflexive, then so is the generated substructure $\mathfrak{M}[u]$, which however is also a bounded morphic image of the irreflexive $\vec{\mathfrak{M}}[u]$. Reflexivity, however, is well-known to be definable in terms of frame validity by the formula $\Box p \rightarrow p$. In other words, the class of reflexive frames is definable by the second-order sentence $\forall P \forall x (\forall y (Rxy \rightarrow Py) \rightarrow Px)$. Intuitively, in terms of truth

in Kripke structures, modal formulae can make very little reference to the underlying frame.

We turn to properties of frames and non-definability in terms of frame validity, which is maybe the most interesting facet of modal expressiveness. One can show, for instance, that (unlike reflexivity) *irreflexivity* is not a modally definable frame property. This property is not preserved under surjective bounded morphisms, while surjective bounded morphisms preserve frame validity. One may consider unfoldings as above, or also the (irreflexive) frame $\mathfrak{F} = \langle \{w_1, w_2\}, \{(w_1, w_2), (w_2, w_1)\} \rangle$ and its bounded morphic image $\mathfrak{F}' = \langle \{w\}, \{(w, w)\} \rangle$, which is reflexive. Hence any φ valid in the former would also be valid in the latter.

Similarly, the class of non-reflexive frames (i.e., ones having at least one irreflexive point) is not definable in terms of validity of modal formulae, because it is not closed under passage to generated subframes. Likewise, the classes of finite frames, connected frames, or of frames with a universal accessibility relation ($R = W^2$), are not definable in terms of frame validity of modal formulae, as they are not closed under disjoint unions.

For another interesting example, consider the property of a frame to be a reflexive partial ordering. It is not modally definable, because anti-symmetry is not preserved under surjective bounded morphisms. Indeed, $\langle \mathbb{Z}, \leq \rangle$ is antisymmetric, but the mapping of it onto the symmetric frame $\mathfrak{F}$ above, sending all odd numbers to w_1 and all even ones to w_2 , is a surjective bounded morphism (and remains so, even when we add an inverse or past modality, as in basic temporal logic).

However, the preservation results we have discussed so far are insufficient to capture frame non-definability. A witness is the following more subtle example: the property of *continuity*, or *Dedekind completeness* is not modally definable in modal logic, but to see that using a non-preservation argument is not easy. Ultimately, this follows from the fact that $\langle \mathbb{R}, \leq \rangle$ (which is continuous) and $\langle \mathbb{Q}, \leq \rangle$ (which is not) have the same *modal theory* (i.e., the same formulae of basic modal logic are valid in these frames); see [42].³

Finally, note that preservation under generated subframes, surjective bounded morphisms and disjoint unions is *not* sufficient to guarantee modal definability in terms of frame validity, even for first-order definable properties. For instance, the class of frames defined by the first-order sentence $\forall x \exists y (xRy \wedge yRy)$ (see [51, 128, 74]) is not modally definable, despite being closed under these three constructions. We will come back to this example in section 6.1.

For more examples of modal non-definability see [5, Section 3.3] and [128] where syntactic characterisations of the first-order properties preserved by each one of the three constructions mentioned above have been obtained.

3 BISIMULATION: A CLOSER LOOK

3.1 Bisimulation games

Bisimulation relations may be understood as descriptions of (non-deterministic) winning strategies for one player in corresponding model comparison games. We illustrate the

³On the other hand, continuity is definable in temporal logic by the formula $\Box([P]p \rightarrow \langle F \rangle [P]p) \rightarrow ([P]p \rightarrow [F]p)$, where F and P are respectively the *future* and *past* modality, and $\Box\varphi = [P]\varphi \wedge \varphi \wedge [F]\varphi$ is the *always* modality. See [42].

concept in the case of bisimulations for basic modal logic – or Kripke structures with a single binary transition relation R – writing $\Diamond$ and $\Box$ for the associated modalities. All considerations admit canonical ramifications to the more general poly-modal setting (as well as to polyadic modalities).

Let $\mathfrak{M} = (W, R, V)$ and $\mathfrak{M}' = (W', R', V')$ be Kripke structures of this basic type. The *bisimulation game* over $\mathfrak{M}$ and $\mathfrak{M}'$ is played by two players **I** and **II** with one pebble in $\mathfrak{M}$ and one in $\mathfrak{M}'$ to mark a single ‘current’ state in each structure. A configuration in the game consists of a current placement of the two pebbles and is described by the pair of pointed Kripke structures $(\mathfrak{M}, w; \mathfrak{M}', w')$, with distinguished w and w' for the current states (pebble positions).

A single *round* in the game is played as follows. The first player, **I**, or *challenger*, selects one of the two pebbles and moves it forward along an edge in the respective structure to a successor state. The second player, **II**, or *defender*, has to respond by similarly moving forward the pebble in the opposite structure.

During the game, **II** loses when no such response is possible or if the resulting new configuration fails to have the two pebbles in *atom equivalent* states (i.e., the new positions are distinguished by at least one atomic proposition, cf. Definition 29 (ii) below). **I** loses during the game if no further round can be played because both pebbles are in states without successors. An infinite run of the game, which continues through an infinite sequence of rounds played according to the above rules, is won by **II**.

We say that **II** has a *winning strategy* in the bisimulation game starting from configuration $(\mathfrak{M}, w; \mathfrak{M}', w')$, if she has responses to any challenges from the first player that guarantee her to win the game (either because **I** gets stuck, or because she can respond with good moves indefinitely).

Intuitively, we think of **I** as challenging the claim of bisimilarity in the current configuration, while **II** defends that bisimilarity claim. This is borne out by the following.

PROPOSITION 28. *Player **II** has a winning strategy in the bisimulation game starting from the initial configuration $(\mathfrak{M}, w; \mathfrak{M}', w')$ if, and only if, $(\mathfrak{M}, w) \rightleftharpoons (\mathfrak{M}', w')$.*

Indeed, an actual bisimulation $\rho: (\mathfrak{M}, w) \rightleftharpoons (\mathfrak{M}', w')$ is a non-deterministic winning strategy for **II**: she merely needs to select her responses so that the currently pebbled states remain linked by ρ . The atom equivalence condition on ρ guarantees that atom equivalence between pebbled states is maintained; the *forth* condition guarantees a matching response to challenges played by **I** in $\mathfrak{M}$; the *back* condition similarly guarantees a matching response to challenges played in $\mathfrak{M}'$.

Conversely, the set of pairs (u, u') in all configurations $(\mathfrak{M}, u; \mathfrak{M}', u')$ from which **II** has a winning strategy, if non-empty, is a bisimulation.

3.2 Finite bisimulations

The games view of a bisimulation suggest that we look at finite approximations corresponding to the existence of winning strategies for a fixed finite number of rounds. These approximations also hold the key to the connection between bisimulation equivalence and modal equivalence. Natural approximations to $\equiv_{\text{ML}}$ are induced by the stratification of ML with respect to the nesting depth of modal formulae (cf. Definition 1), as follows.

DEFINITION 29. For two pointed Kripke τ -structures $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$:

- (i) For $n \geq 0$, $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ are ML_n -equivalent, denoted $(\mathfrak{M}, w) \equiv_{\text{ML}}^n (\mathfrak{M}', w')$, iff they satisfy exactly the same formulae of ML_n .
- (ii) At the level of $\equiv_{\text{ML}}^0$, we write $w \simeq w'$ instead of $(\mathfrak{M}, w) \equiv_{\text{ML}}^0 (\mathfrak{M}', w')$ and say that w and w' are *atom equivalent* (or, isomorphic when viewed as isolated states with atomic propositions according to V, V').

The n -round bisimulation game is played like the (unbounded) bisimulation game but terminates after n rounds (or beforehand if either player loses during one of these rounds). Now **II** also wins if the n -th round is completed without violating atom equivalence. The notion of a winning strategy is correspondingly adapted.

DEFINITION 30. Let $n \geq 0$. Two pointed Kripke structures $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ are

- (i) n -bisimilar, or n -bisimulation equivalent, denoted $(\mathfrak{M}, w) \rightleftharpoons_n (\mathfrak{M}', w')$, if **II** has a winning strategy in the n -round bisimulation game starting from $(\mathfrak{M}, w; \mathfrak{M}', w')$.
- (ii) *finitely bisimilar*, $(\mathfrak{M}, w) \rightleftharpoons_\omega (\mathfrak{M}', w')$, if $(\mathfrak{M}, w) \rightleftharpoons_n (\mathfrak{M}', w')$ for all $n \in \mathbb{N}$.

Note that 0-bisimulation equivalence is atom equivalence or modal equivalence $\equiv_{\text{ML}}^0$, indistinguishability at the propositional level.

Clearly n -bisimulation equivalence implies m -bisimulation equivalence for any $m \leq n$; (full) bisimulation equivalence implies finite bisimulation equivalence; and finite bisimulation equivalence implies n -bisimulation equivalence for any n . We shall return to the interesting relationship between finite and full bisimulation equivalence below, in connection with the Hennessy–Milner Theorem; see section 38.

A first connection between $\rightleftharpoons_n$ and n -equivalence is made in the following.

LEMMA 31. $(\mathfrak{M}, w) \rightleftharpoons_n (\mathfrak{M}', w') \Rightarrow (\mathfrak{M}, w) \equiv_{\text{ML}}^n (\mathfrak{M}', w')$.

Indeed, if $\mathfrak{M}, w \models \varphi$ and $\mathfrak{M}', w' \models \neg\varphi$ for some $\varphi \in \text{ML}_n$, then **I** has a winning strategy in the n -round game from $(\mathfrak{M}, w; \mathfrak{M}', w')$. This is shown by induction on the nesting depth n of the distinguishing formula φ . At level $n = 0$, a distinction in ML_0 means atomic inequivalence – corresponding to a configuration in which **II** has lost.

In the induction step, assume that $(\mathfrak{M}, w)$ is distinguished from $(\mathfrak{M}', w')$ by a formula $\varphi \in \text{ML}_{n+1}$. Propositional connectives in φ can be unravelled so that without loss of generality φ is of the form $\Diamond\psi$ for some $\psi \in \text{ML}_n$. Suppose then that for instance $\mathfrak{M}', w' \models \neg\varphi$, while $\mathfrak{M}, w \models \varphi$. Let in that case **I** move the pebble in $\mathfrak{M}$ from w to some u , where $\mathfrak{M}, u \models \psi$. As $\mathfrak{M}', w' \models \neg\Diamond\psi$, any available response for **II** can only lead to a configuration $(\mathfrak{M}, u; \mathfrak{M}', u')$ in which $(\mathfrak{M}, u)$ and $(\mathfrak{M}', u')$ are distinguished by $\psi \in \text{ML}_n$. Therefore, by the inductive hypothesis, **I** has a winning strategy for the remaining n rounds of the game.

3.3 Characteristic formulae

For the converse to the previous lemma, or for capturing the bounded bisimulation game in terms of modal logic, it is essential that the underlying vocabulary is finite: both τ and Φ need to be finite. We again stick with the case of a single binary accessibility relation R , but that restriction is purely for expository simplicity.

The crucial step in the transition from the bisimulation game to modal logic is the formalisation, as a formula $\chi_{[\mathfrak{M}, w]}^n \in \text{ML}_n$, of

“**II** has a winning strategy in the n -round game from $(\mathfrak{M}, w; \mathfrak{M}', w')$ ”

as a property of $(\mathfrak{M}', w')$, for fixed reference structure $(\mathfrak{M}, w)$ and depth n . In fact $\chi_{[\mathfrak{M}, w]}^n$ may be constructed by induction on n , simultaneously for all $(\mathfrak{M}, w)$. Along with the induction one observes that $\rightleftharpoons_n$ has finite index, and that, correspondingly, we generate only finitely many non-equivalent formulae $\chi_{[\mathfrak{M}, w]}^n$ at level n (for finite τ and Φ !).

For $n = 0$, $\chi_{[\mathfrak{M}, w]}^0$ is purely propositional and consists of the conjunction of all $p \in \Phi$ that are true in w and all $\neg p$ for those that are false at w . This fixes the atomic equivalence type, as it should.

Inductively, let

$$\chi_{[\mathfrak{M}, w]}^{n+1} := \chi_{[\mathfrak{M}, w]}^0 \wedge \underbrace{\bigwedge_{(w, u) \in R} \Diamond \chi_{[\mathfrak{M}, u]}^n}_{\text{forth}} \wedge \underbrace{\Box \bigvee_{(w, u) \in R} \chi_{[\mathfrak{M}, u]}^n}_{\text{back}}.$$

Even for infinitely branching $\mathfrak{M}$, the conjunctions and disjunctions in this formula remain finite up to logical equivalence as there are only finitely many formulae of the respective kind.

Clearly $\mathfrak{M}, w \models \chi_{[\mathfrak{M}, w]}^{n+1}$. But for arbitrary $(\mathfrak{M}', w')$, $\mathfrak{M}', w' \models \chi_{[\mathfrak{M}, w]}^{n+1}$ indeed guarantees **II** a winning strategy in the $(n+1)$ -round game from $(\mathfrak{M}, w; \mathfrak{M}', w')$. The conjunct $\chi_{[\mathfrak{M}, w]}^0$ guarantees that the game is not lost already. The *back*-and-*forth* attributions in the two main conjuncts suggest how these are used to guarantee suitable responses, in the first round, to challenges from **I** played in either $\mathfrak{M}$ (*forth*) or $\mathfrak{M}'$ (*back*), respectively.

The *forth* part says that for all moves from w to some u in $\mathfrak{M}$, $\mathfrak{M}', w' \models \Diamond \chi_{[\mathfrak{M}, u]}^n$, and any R' -successor u' of w' such that $\mathfrak{M}', u' \models \chi_{[\mathfrak{M}, u]}^n$ provides a response for **II** that will allow her to succeed through another n rounds.

Similarly the *back* part says that for all moves from w' to some u' in $\mathfrak{M}$, $\mathfrak{M}', u' \models \chi_{[\mathfrak{M}, u]}^n$ for some R -successor u of w in $\mathfrak{M}$ – a response that is good for another n rounds for **II**.

That failure of $\mathfrak{M}', w'$ to satisfy $\chi_{[\mathfrak{M}, w]}^n$ affords **I** a win within n rounds follows from Lemma 31. Together, these observations yield the following tight connection between the bisimulation game and modal equivalence.

THEOREM 32. *Let $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ be pointed Kripke structures of the same finite type with finitely many atomic propositions. Then the following are equivalent:*

- (i) $(\mathfrak{M}, w) \rightleftharpoons_n (\mathfrak{M}', w')$.
- (ii) **II** has a winning strategy in the n -round game from $(\mathfrak{M}, w; \mathfrak{M}', w')$.
- (iii) $\mathfrak{M}', w' \models \chi_{[\mathfrak{M}, w]}^n$.
- (iv) $(\mathfrak{M}, w) \equiv_{\text{ML}}^n (\mathfrak{M}', w')$.

As corollaries we obtain a corresponding characterisation of full modal equivalence, and a normal form for ML formulae.

COROLLARY 33. *Over Kripke structures of finite type and with finitely many atomic propositions, finite bisimulation equivalence $\rightleftharpoons_\omega$ coincides with modal equivalence.*

COROLLARY 34. *Any formula $\varphi \in \text{ML}_n$ is logically equivalent to the disjunction $\bigvee_{\mathfrak{M}, w \models \varphi} \chi_{[\mathfrak{M}, w]}^n$, which is in fact finite as there are only finitely many such χ^n up to logical equivalence (in the vocabulary of φ).*

Similarly, for finite vocabularies, any class $\mathcal{C}$ of pointed Kripke structures that is closed under n -bisimulation is definable in ML_n by the disjunction $\bigvee_{(\mathfrak{M}, w) \in \mathcal{C}} \chi_{[\mathfrak{M}, w]}^n$.

Bisimulation-invariance of ML, Theorem 14, also becomes a simple corollary of the analysis of the game. Indeed, $\varphi \in \text{ML}_n$ is even invariant under n -bisimulation equivalence $\rightleftharpoons_n$, which of course implies invariance under (full) bisimulation.

3.4 Finite model property

A logic $\mathcal{L}$ has the *finite model property* (FMP) iff every satisfiable formula of $\mathcal{L}$ is satisfiable in a finite model, i.e., if satisfiability and satisfiability in finite structures coincide for $\mathcal{L}$.

For specific modal logics (e.g., normal extensions of basic modal logic) the implicit restriction to a prescribed class of admissible frames corresponds to a relativisation of the above criterion to the respective classes of (infinite or finite) admissible models. So the finite model property for S5 say, states that any formula of ML that is satisfiable over some equivalence frame is also satisfiable over some finite equivalence frame.

The finite model property is a characteristic feature of many modal logics. For any modal logic with a recursive axiomatisation (such that it is recursively enumerable for validity) whose class of admissible finite frames is also recursively enumerable, the finite model property provides a standard method for proving decidability.⁴ Here we briefly discuss the general *filtration method* for establishing the finite model property for modal logics. For basic modal logic we illustrate in section 3.4 that it even has a *finite tree model property*: every satisfiable formula has a finite tree model.⁵

Filtration

Filtration is the most widely used method for proving the finite model property in modal logics, particularly those determined by classes of frames with specific properties of the accessibility relation. This method is originally due to McKinsey who first applied an algebraic version of it in modal logic. Filtration was introduced in its present form by Lemmon and Scott [91] and further developed and applied by Segerberg [117]. Gabbay [31] introduced a different version, called *selective filtration*. Later, Fisher and Ladner [29] proved the finite model property of propositional dynamic logic PDL using filtration.

Given a formula φ of a modal logic L and a Kripke structure $\mathfrak{M}$ (of type appropriate for L) satisfying φ , we want to produce a *finite* Kripke structure $\tilde{\mathfrak{M}}$ (of appropriate type) satisfying φ . The method of filtration provides a transformation from models $\mathfrak{M}$ to finite models $\tilde{\mathfrak{M}}$ in a uniform manner with respect to φ and $\mathfrak{M}$. Before outlining the construction let us note that the satisfiability of a modal formula φ in a Kripke structure only depends on the truth of the (finitely many) subformulae of φ across that structure. Therefore, two states in a Kripke structure that satisfy the same subformulae of φ are *indistinguishable* from the viewpoint of φ . Sometimes it is necessary to extend the set of subformulae of φ to a wider but still finite set of formulae, called the *closure* of φ and denoted by $cl(\varphi)$. Thus, $cl(\varphi)$ partitions the model into *finitely many equivalence classes*

⁴However, not every modal logic with the finite model property is decidable; see for instance $K \times K \times K$ in [35].

⁵In general, the finite model property and the tree model property are independent. While the tree model property can account for the decidability of the basic systems of modal logic and many of its extensions and variations (see section 5), it does not apply to axiomatic extensions which impose specific restrictions on the frames that are incompatible with tree-like structures like symmetry or confluence of the accessibility relation.

of states, all states in each class satisfying the same subset of $cl(\varphi)$. The underlying idea of the filtration method is to collapse the infinite model to its finite quotient with respect to the equivalence relation generated by that partition, in a way that preserves the truth of all formulae in $cl(\varphi)$, and hence of φ itself.

The equivalence relation itself can be thought of as coarse-grained approximation to bisimulation equivalence that is specific to the given formula φ . It is meant to preserve φ but needs to do so at a coarser level than bisimulation to be of finite index. Note that n -bisimulation $\rightleftharpoons_n$ can also serve as a finite index approximation but, because of its graded nature, does not lend itself to taking quotients in the desired global manner. This is because $\mathfrak{M}, u \rightleftharpoons_n \mathfrak{M}, u'$ (i.e., that u and u' are of the same n -bisimulation type) does not imply that the same n -bisimulation types are accessible from u and u' .

Here are the formal details. Take any Kripke structure $\mathfrak{M} = \langle W, R, V \rangle$ and a set of formulae Γ , which is assumed to be closed under subformulae, single negation (i.e., if $\varphi \in \Gamma$ is not a negation itself, then $\neg\varphi \in \Gamma$) and under $\Box/\Diamond$ duality. Define an equivalence relation $\sim_\Gamma$ on W as follows:

$$u \sim_\Gamma w \text{ iff for every } \psi \in \Gamma : \mathfrak{M}, u \models \psi \Leftrightarrow \mathfrak{M}, w \models \psi.$$

Let $[w]_\Gamma$ be the equivalence class of w with respect to $\sim_\Gamma$ and $W_\Gamma = \{[w]_\Gamma \mid w \in W\}$. Note that if Γ is finite, then W_Γ is finite, too. Further, the valuation V is collapsed to a valuation V_Γ in W_Γ for all $p \in \Gamma$ canonically: $V_\Gamma(p) = \{[w]_\Gamma \mid w \in V(p)\}$; for all other variables q , V_Γ is defined arbitrarily, e.g., $V_\Gamma(q) = \emptyset$.

Now, we say that a Kripke structure $\mathfrak{M} = \langle W_\Gamma, \tilde{R}, V_\Gamma \rangle$ is a *filtration of M with respect to Γ* if for every $\psi \in \Gamma$ and $w \in W$: $\mathfrak{M}, w \models \psi$ iff $\mathfrak{M}, [w]_\Gamma \models \psi$. With a slight abuse of terminology, we also say that $\tilde{R}$ is a filtration of R with respect to Γ .

There are two simple conditions on the relation $\tilde{R}$ which guarantee that it is a filtration of R with respect to Γ . They give lower and upper bounds for that relation, respectively:

MIN. For every $u, w \in W$, if uRw then $[u]_\Gamma \tilde{R} [w]_\Gamma$.

MAX. For every $[u]_\Gamma, [w]_\Gamma \in W_\Gamma$, if $[u]_\Gamma \tilde{R} [w]_\Gamma$ then for every $\Box\psi \in \Gamma$:
if $\mathfrak{M}, u \models \Box\psi$ then $\mathfrak{M}, w \models \psi$.⁶

By induction on ψ one can prove that for every $\tilde{R}$ satisfying these conditions, the structure $\mathfrak{M} = \langle W_\Gamma, \tilde{R}, V_\Gamma \rangle$ is indeed a filtration of $\mathfrak{M}$ with respect to Γ , and this claim is known as the *filtration lemma*. Often, the conditions *MIN* and *MAX* are adopted as *the definition* of a filtration of R , and the filtration lemma then claims that they imply that $\mathfrak{M}_\Gamma$ has the desired property.

Does every Kripke structure have a filtration with respect to any set of formulae Γ ? Yes: converting the implication to equivalence in each of the conditions *MIN* and *MAX* defines a relation that satisfies the other condition, too, and hence renders a filtration:

- the *minimal filtration* $\mathfrak{M}_\Gamma^{\min} = \langle W_\Gamma, R_\Gamma^{\min}, V_\Gamma \rangle$, $R_\Gamma^{\min} = \{([u]_\Gamma, [w]_\Gamma) \mid (u, w) \in R\}$;
- the *maximal filtration* $\mathfrak{M}_\Gamma^{\max} = \langle W_\Gamma, R_\Gamma^{\max}, V_\Gamma \rangle$, where $[u]_\Gamma R_\Gamma^{\max} [w]_\Gamma$ holds iff for every $\Box\psi \in \Gamma$, $\mathfrak{M}, u \models \Box\psi$ implies $\mathfrak{M}, w \models \psi$.

Clearly, every relation $\tilde{R}$ such that $R_\Gamma^{\min} \subseteq \tilde{R} \subseteq R_\Gamma^{\max}$ is a filtration, too.

Now, given a formula φ and a pointed Kripke structure $(\mathfrak{M}, u)$ such that $\mathfrak{M}, u \models \varphi$, applying the filtration construction to $\Gamma = cl(\varphi)$ produces a finite pointed Kripke

structure $(\tilde{\mathfrak{M}}, [u]_\Gamma)$ that satisfies φ , whence basic modal logic (modal logic K) has the finite model property.

This method can be refined to establish the finite model property for axiomatic extensions L of K , too, by adjusting the definition of $\tilde{R}$ so as to preserve the desired properties of the original structure $\mathfrak{M}$, such as transitivity, linearity etc., or to impose such properties on the resulting structure $\tilde{\mathfrak{M}}$ and thus to eventually guarantee that it is a model appropriate for the desired modal logic L . Examples of filtrations for a number of important modal and temporal logics, such as T , $K4$, $S4$, and the logics of various linear orderings can be found in [91, 117, 42]. A more general result, extending a theorem of Lewis [92], is [122, Theorem 2.6.8] stating that every modal logic axiomatised by a finite set of shallow formulae (see Section 8.2) admits filtration.

Finite tree model property

Before returning to the relationship between bisimulation and finite bisimulation in the next section, we apply the preservation result of Lemma 31 to an alternative, simple proof of the finite model property for basic modal logic, by establishing a stronger *finite tree model property*.

LEMMA 35 (finite tree model property). *For every $n \in \mathbb{N}$, every pointed Kripke structure of finite relational type is n -bisimilar to a finite tree structure. Consequently, any satisfiable formula of ML is satisfied at the root of a finite tree.*

Proof. According to section 2.2 the unfolding $\mathfrak{M}[u]$ of $(\mathfrak{M}, u)$ provides a bisimilar tree structure. As we only need n -bisimulation equivalence, we may cut off $\mathfrak{M}[u]$ at depth n from its root u , to obtain a tree structure $(\mathfrak{M}[u] \upharpoonright U^n(u), u) \rightleftharpoons_n (\mathfrak{M}, u)$ whose depth is bounded by n , where $U^n(u)$ stands for the set of nodes at distance up to n from u . This tree structure may still be infinite, due to infinite branching. In that case, however, we may prune successors at every node to retain at most one representative of each $\rightleftharpoons_n$ equivalence class. As $\rightleftharpoons_n$ has finite index (for finite vocabulary; see section 3.3), the resulting tree structure is finite. ■

Finite branching, as well as a finite bound on the number of bisimulation types, are obvious for finite $\mathfrak{M}$, but a finite pointed Kripke structure $(\mathfrak{M}, u)$ in which a directed cycle is reachable from u cannot be bisimilar to a finite tree structure. Locally, however, this can be achieved in partial unfoldings.

LEMMA 36. *Let $n \in \mathbb{N}$. Every finite pointed Kripke structure $(\mathfrak{M}, u)$ is bisimilar to a finite pointed structure $(\hat{\mathfrak{M}}, \hat{u})$ whose restriction to depth n from the distinguished node $\hat{u}$ is a tree structure.*

Proof. Let $(\mathfrak{M}[u] \upharpoonright U^n(u), u)$ be as in the proof of the last lemma (now finite). For each leaf node of this structure, take a new disjoint isomorphic copy of $\mathfrak{M}$ itself and identify the leaf node with its bisimilar partner node in that copy of $\mathfrak{M}$. The resulting structure is finite, bisimilar to $(\mathfrak{M}, u)$ and tree-like up to distance n from the distinguished node. ■

Remarks. Results of this type can be carried much further. For instance, a more involved construction yields finite two-way bisimilar companions which are acyclic in restriction to the n -neighbourhood of any node [106, 107]. Such locally acyclic finite *bisimilar covers*

are available also in restriction to various other non-elementary classes of frames, e.g., within the classes of all (finite) rooted frames or finite equivalence frames [17].

It is also interesting to note that finite and bisimilar $(\mathfrak{M}, u)$ and $(\mathfrak{M}', u')$ admit finite bisimilar companions $(\hat{\mathfrak{M}}, \hat{u})$ and $(\hat{\mathfrak{M}}', \hat{u}')$, respectively, whose restrictions to depth n from their distinguished nodes $\hat{u}$ and $\hat{u}'$ are even isomorphic tree structures. For this, we take $(\hat{\mathfrak{M}}, \hat{u})$ and $(\hat{\mathfrak{M}}', \hat{u}')$ as from the proof above and modify them by merely attaching extra isomorphic copies of substructures at nodes in the tree parts so as to achieve equal multiplicities for all bisimulation types at each node in the tree parts. It then follows that the tree parts are isomorphic.

3.5 Hennessy–Milner theorem

For the relationship between finite and full bisimulation equivalence, we note that finite bisimulation equivalence can be strictly weaker in structures with infinite branching. A typical example of tree structures with $(\mathfrak{M}, w) \rightleftharpoons_\omega (\mathfrak{M}', w')$ but $(\mathfrak{M}, w) \not\equiv (\mathfrak{M}', w')$ is the following.

EXAMPLE 37. Let $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ be tree Kripke structures with trivial valuations, rooted at w and w' , respectively. Let the roots have countably many distinct successors u_i , $i \geq 1$ in $\mathfrak{M}$ and u'_i , $i \geq 0$ in $\mathfrak{M}'$. For $i \geq 1$, we let each of u_i and u'_i be the starting point of a simple finite path of length i . We let the extra node u'_0 in $\mathfrak{M}'$ be the root of a simple infinite path. Then $(\mathfrak{M}, w) \not\equiv (\mathfrak{M}', w')$: let **I** move in $\mathfrak{M}'$ from w' to u'_0 ; the second player must move to one of the u_i for $i \geq 1$ in $\mathfrak{M}$; let then **I** lead the play in $\mathfrak{M}'$ along the infinite path: **II** gets stuck and loses in round $i+2$ when the end of the length i path from u_i has been reached. On the other hand, $(\mathfrak{M}, w) \rightleftharpoons_n (\mathfrak{M}', w')$ for every $n \in \mathbb{N}$, since any two paths of lengths greater than or equal to n look exactly the same in an n -round game.

However, infinite branching is essential to this phenomenon, as the following shows.

THEOREM 38 (Hennessy–Milner theorem). *Let $\mathfrak{M}$ and $\mathfrak{M}'$ both be finitely branching, i.e., every state in either structure has only finitely many immediate successors.*

Then $(\mathfrak{M}, w) \rightleftharpoons_\omega (\mathfrak{M}', w')$ implies $(\mathfrak{M}, w) \equiv (\mathfrak{M}', w')$. Consequently, over finitely branching Kripke structures, modal equivalence coincides with bisimulation equivalence.

Proof. The argument is best given via the games. We claim that **II** can maintain $(\mathfrak{M}, w) \rightleftharpoons_\omega (\mathfrak{M}', w')$ indefinitely – which gives her a winning strategy for the infinite game. For instance, let **I** play in $\mathfrak{M}$ and move the pebble from w to u . Suppose that for all responses u' available to **II** in $\mathfrak{M}'$, $(\mathfrak{M}, u) \not\equiv_\omega (\mathfrak{M}', u')$. As there are only finitely many choices for u' due to finite branching, we can find a sufficiently large $n \in \mathbb{N}$ such that $(\mathfrak{M}, u) \not\equiv_n (\mathfrak{M}', u')$ for all u' with $(w', u') \in R'$. But this would imply $(\mathfrak{M}, w) \not\rightleftharpoons_{n+1} (\mathfrak{M}', w')$, contradicting the assumption $(\mathfrak{M}, w) \rightleftharpoons_\omega (\mathfrak{M}', w')$. ■

Unlike the Hennessy–Milner theorem, which is rather specific for bisimulation, the following observation rests on arguments from classical model theory, to do with saturation properties, and highlights a more general principle that applies to any finitary versus unbounded game equivalences of the Ehrenfeucht–Fraïssé variety; see for instance [108]. *Saturation* properties refer to the realisation of *types*. We think of a type as the formalisation of the properties of an element, through a set of formulae using constants for parameters from a given structure.

With a first-order language L and a set of parameters $A \subseteq W$ of the universe W of some structure $\mathfrak{M}$, associate the expansion of L_A of L with a constant name for each element of A ; the corresponding expansion of $\mathfrak{M}$ is denoted $\mathfrak{M}_A$.

DEFINITION 39. An *element type with parameters in A* (in the first-order language L_A) is a set Σ of L_A -formulae in a single free element variable x . The type Σ is a *type of $\mathfrak{M}_A$* if it is (finitely) *consistent* with the theory of $\mathfrak{M}_A$ in the sense that $\mathfrak{M}_A \models \exists x \bigwedge \Sigma_0$ for every finite $\Sigma_0 \subseteq \Sigma$. The type Σ is *realised* in $\mathfrak{M}_A$ if $\mathfrak{M}_A, w \models \Sigma$ for some element w .

A structure $\mathfrak{M}$ is ω -saturated if for every finite subset A every type of $\mathfrak{M}_A$ is realised in $\mathfrak{M}_A$.

Interesting properties are often expressible by types rather than by an individual formula. For instance, an R -successor of $w \in \mathfrak{M}$ from which there are arbitrarily long R -paths is described by the type $\Sigma_w := \{Rwx\} \cup \{\text{ST}(\Diamond^n \top; x) \mid n \in \mathbb{N}\}$ with parameter w from $\mathfrak{M}$. Note that this is a type of $\mathfrak{M}$ iff there are arbitrarily long R -paths from $w \in \mathfrak{M}$. This does not imply that $\mathfrak{M}$ itself has a realisation of the type – a successor u of w that simultaneously satisfies all the requirements in Σ_w . By compactness, however, every structure $\mathfrak{M}$ has an ω -saturated elementary extension, cf. [12]. Let $\mathfrak{M}^*$ be such an elementary extension of $\mathfrak{M}$. For $w \in \mathfrak{M}$, Σ_w is a type of $\mathfrak{M}^*$ if it is a type of $\mathfrak{M}$. If Σ_w for $w \in \mathfrak{M}^*$ is a type of $\mathfrak{M}^*$, then there also is some R -successor u of w in $\mathfrak{M}^*$ such that $\mathfrak{M}^*, u \models \Diamond^n \top$ for all n ; hence Σ_u will also be a type of $\mathfrak{M}^*$ and repeating the argument inductively we find that $\mathfrak{M}^*$ has an infinite path from w . In ω -saturated models, therefore, any element from which there are arbitrarily long paths, will also have an infinite path. Similar reasoning extends to provide responses for **II** in the infinite game over $\mathfrak{M}^*$ to meet any challenge from **I**, provided she has responses that are good for n rounds, for each n . In other words, playing over ω -saturated structures, **II** has a winning strategy in the infinite game whenever she has, for every n , a winning strategy for the n -round game. The proof is analogous to that given for Proposition 87 below; in the terminology to be introduced there, the class of ω -saturated structures has the *Hennessy–Milner property*.

REMARK 40. As shown in section 6.3, $\rightleftharpoons_\omega$ coincides with $\rightleftharpoons$ in restriction to ω -saturated structures.

In the bisimulation context weaker forms of saturation suffice, and in that sense the Hennessy–Milner theorem may be regarded as a special case. See section 6.3 for more on (modal) saturation.

3.6 Bisimulation and infinitary modal equivalence

Since, over infinitely branching structures, equivalence with respect to ordinary modal logic only reaches up to the level of finite bisimulation equivalence, $\rightleftharpoons_\omega$, the question of the actual logical counterpart to full bisimulation equivalence arises. The situation is entirely similar to that in classical first-order logic, where it is clarified by Karp’s theorem [82] (see also [68]). While the classical Ehrenfeucht–Fraïssé theorem associates finitary game equivalence (the back-and-forth notion of *finite isomorphism* between structures) with elementary equivalence, full infinitary game equivalence (the back-and-forth notion of *partial isomorphism* between structures) corresponds to equivalence with respect to the *infinitary logic* $L_{\infty\omega}$ whose syntax allows for disjunctions and conjunctions over arbitrary sets of formulae, [68, 108]. In order to extend modal logic ML to its infinitary variant

ML_∞ , we put the following additional clause for formula formation. If Ψ is any set of formulae of ML_∞ , then $\bigwedge \Psi$ and $\bigvee \Psi$ are formulae of ML_∞ . These formulae have an ordinal-valued nesting depth, based on the usual rules for the finitary constructors of ML (see Definition 1) together with the extra stipulation that the nesting depth of an infinitary conjunction or disjunction is the supremum of the nesting depths of the constituent formulae. The semantics of the infinite conjunctions and disjunctions is the natural one; with, e.g., $(M, w) \models \bigvee \Psi$ iff $(M, w) \models \psi$ for some $\psi \in \Psi$.

Completely analogous to the treatment of the finitary game in relation to finitary modal logic, we then get the following. (As a side effect of the availability of infinitary conjunctions and disjunctions, we need not restrict the underlying vocabularies to be finite.) Comparison with the classical version of Karp's theorem highlights the observation that bisimulation is for modal model theory what partial isomorphism is for classical model theory.

THEOREM 41 (Karp's theorem for modal logic). *Let $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ be Kripke structures of the same type. Then the following are equivalent:*

- (i) $(\mathfrak{M}, w) \rightleftharpoons (\mathfrak{M}', w')$.
- (ii) **II** has a winning strategy in the infinite bisimulation game from $(\mathfrak{M}, w; \mathfrak{M}', w')$.
- (iii) $(\mathfrak{M}, w) \equiv_{ML_\infty} (\mathfrak{M}', w')$.

Proof. (i) $\Leftrightarrow$ (ii) is obvious. For (ii) $\Rightarrow$ (iii) compare Lemma 31: similar to there, if $(\mathfrak{M}, w)$ is distinguished from $(\mathfrak{M}', w')$ by a formula of nesting depth α , then one can find a move for **I** which will force a successor configuration in which the positions are distinguished at some nesting depth $\beta < \alpha$. By well-foundedness this gives **I** a winning strategy. For (iii) $\Rightarrow$ (ii) one observes that **II** can maintain ML_∞ equivalence indefinitely. ■

Remark. Characteristic formulae $\chi_{[\mathfrak{M}, u]}^\alpha$ with an ordinal parameter α for their nesting depth, can still be defined inductively in a canonical way. (The analogous infinitary formulae for the infinite first-order Ehrenfeucht–Fraïssé game are known as Scott formulae, see for instance [68].) For the infinite game over infinitely branching $\mathfrak{M}$, a position in which players may have infinitely many non-equivalent choices for a next move, is adequately described by an infinite conjunction $\bigwedge \Diamond \varphi_i$ in conjunction with $\Box \bigvee \varphi_i$, where each φ_i describes the bisimulation type of one potential successor in the game over $\mathfrak{M}$, at a nesting depth level that typically needs to be an infinite ordinal. A sufficiently high nesting depth that can be used uniformly across a given $\mathfrak{M}[u]$ is the least ordinal α such that any two states in $\mathfrak{M}[u]$ that are equivalent at nesting depth α in ML_∞ are equivalent at nesting depth $\alpha + 1$. For this α , equivalence at nesting depth α implies equivalence at any nesting depth, i.e., full ML_∞ equivalence, and hence bisimilarity. (The minimal such α is the *closure ordinal* of the co-inductive definition of the bisimulation relation over $\mathfrak{M}[u]$, also compare section 3.7.)

That a given α has this property for $\mathfrak{M}[u]$ is itself expressible in ML_∞ . The defining property of α is equivalent to the assertion that, for all $v \in \mathfrak{M}[u]$, $\mathfrak{M}[u] \models \chi_{[\mathfrak{M}, v]}^\alpha \rightarrow \chi_{[\mathfrak{M}, v]}^{\alpha+1}$. Let ψ^α be the conjunction of all formulae

$$\Box^n \bigwedge_{v \in \mathfrak{M}[u]} (\chi_{[\mathfrak{M}, v]}^\alpha \rightarrow \chi_{[\mathfrak{M}, v]}^{\alpha+1}).$$

Then $\mathfrak{M}, u \models \psi^\alpha$ iff within $\mathfrak{M}[u]$, the ML_∞ type at nesting depth $\alpha+1$ is fully determined by the type at nesting depth α . The conjunction $\chi_{[\mathfrak{M}, u]} := \psi^\alpha \wedge \chi_{[\mathfrak{M}, u]}^\alpha$ for suitable α , characterises $(\mathfrak{M}, u)$ up to bisimulation, thus providing a canonical characteristic formula in ML_∞ .

Unlike the definability assertion of Corollary 34, however, bisimulation closure of a class $\mathcal{C}$ of Kripke structures on its own does not guarantee definability in ML_∞ . In the example of observation 42 below, the relevant disjunction of characteristic formulae would be class-sized, and hence not in ML_∞ . However, definability of $\mathcal{C}$ in ML_∞ does follow, for instance, if $\mathcal{C}$ comprises only set-many different bisimulation types (which is in particular the case for the setting of finite model theory, or in restriction to any other class of bounded cardinality). This is sufficient to ensure that $\mathcal{C}$ is definable by a disjunction over characteristic formulae analogous to Corollary 34.

OBSERVATION 42. Well-foundedness, or the class of all pointed Kripke structures $(\mathfrak{M}, u)$ in which there is no infinite path from u , is not definable in infinitary modal logic ML_∞ .

This class is definable by the modal μ -calculus L_μ formula $\mu X. \Box X$ (see section 5.2) and hence in monadic second-order logic MSO .⁷ On the other hand, well-foundedness is not even definable in infinitary first-order logic $L_{\infty\omega}$, [95]. We sketch a direct proof of non-definability in ML_∞ .

For an ordinal α consider the Kripke structure $\mathfrak{M}_\alpha = \langle \{\beta \mid \beta \leq \alpha\}, R \rangle$ with $R = \{(\beta, \beta') \mid \beta' < \beta < \alpha\}$ the inverse of the order relation on these ordinals, and its modification $\mathfrak{M}'_\alpha$ with R replaced by $R' = R \cup \{(\alpha, \alpha)\}$. We show by induction on the ordinal γ that $(\mathfrak{M}'_\alpha, \alpha) \equiv_\gamma (\mathfrak{M}_\alpha, \beta)$ (equivalence in ML_∞ up to nesting depth γ) for all $\alpha \geq \beta \geq \gamma$. It follows that no formula of ML_∞ can separate the well-founded $(\mathfrak{M}_\alpha, \alpha)$ from the non-wellfounded $(\mathfrak{M}'_\alpha, \alpha)$ for all α .

The claim is obvious for $\gamma = 0$; also the limit steps are trivial. For the successor step, from γ to $\gamma + 1$, consider $\alpha \geq \beta \geq \gamma + 1$; it suffices to show that then even $\mathfrak{M}'_\alpha, \alpha \models \Diamond \psi \Leftrightarrow \mathfrak{M}_\alpha, \beta \models \Diamond \psi$ for ψ of nesting depth γ . The only non-trivial instance of this assertion is when $\mathfrak{M}'_\alpha, \alpha \models \Diamond \psi$ because $\mathfrak{M}'_\alpha, \alpha \models \psi$. But then $\mathfrak{M}_\alpha, \gamma \models \psi$ by the inductive hypothesis. It follows that $\mathfrak{M}_\alpha, \beta \models \Diamond \psi$ as $\beta \geq \gamma + 1$ implies that $(\beta, \gamma) \in R$.

3.7 Largest bisimulations as greatest fixed points

The union of all bisimulation relations between two given Kripke structures is again a bisimulation relation, and hence a maximal bisimulation in the sense of set inclusion. Such *largest bisimulations* can also be defined co-inductively, and be understood as the greatest fixed-point of suitable monotone operators. Again, and purely for expository purposes, we sketch this approach in the simple case of a single accessibility relation R .

Let $X \subseteq W \times W'$, and let $w \in W$ and $w' \in W'$ be atom equivalent ($w \simeq w'$). Let us say that the pair (w, w') has the *back-and-forth property w.r.t. X* iff player **II** has a single round strategy to lead the bisimulation game from $(\mathfrak{M}, w; \mathfrak{M}', w')$ to a configuration $(\mathfrak{M}, u; \mathfrak{M}', u')$ such that $(u, u') \in X$. (Note that the *back-and-forth* conditions for a bisimulation relation say that each of its pairs has the *back-and-forth* property w.r.t. the relation itself.)

⁷Note that this is definability in the sense of (local) Kripke structure semantics, albeit in a logic which is itself of a second-order nature, and should not be confused with the modal definability of the class of transitive well-founded frames.

Consider the following operator F on subsets $X \subseteq W \times W'$:

$$F(X) := \{(w, w') \in X \mid (w, w') \text{ has the } \textit{back-and-forth} \text{ property w.r.t. } X\}.$$

The operator F is monotone in the sense that $X \subseteq Y \Rightarrow F(X) \subseteq F(Y)$. It therefore has a unique greatest fixed point in restriction to any subset of $W \times W'$. We are interested in the greatest fixed point of F that respects atom equivalence, and therefore consider the restriction F_0 of F to $X_0 := \{(u, u') \in W \times W' \mid u \simeq u'\}$. Let $\rho := \text{gfp}(F_0) \subseteq X_0$ be this greatest fixed point. Being a fixed point of F within X_0 , ρ respects atom equivalence; being a fixed point of F , ρ has the *back-and-forth* property. So ρ is a bisimulation. As any bisimulation between $\mathfrak{M}$ and $\mathfrak{M}'$ must also be a fixed-point of F_0 , ρ is the largest such.

REMARK 43. The stages of the evaluation of $\text{gfp}(F_0)$ produce a monotone decreasing ordinal-indexed sequence of subsets $X_\alpha \subseteq W \times W'$ according to

$$\begin{aligned} X_0 &= \{(u, u') \in W \times W' \mid u \simeq u'\} \\ X_{\alpha+1} &= F_0(X_\alpha) \quad (\text{successor stage}) \\ X_\lambda &= \bigcap_{\alpha < \lambda} X_\alpha \quad (\text{limit stage}) \end{aligned}$$

which is eventually constant with value $\text{gfp}(F_0)$. The least ordinal α such that $X_{\alpha+1} = X_\alpha$ is called the *closure ordinal* of this greatest fixed point evaluation over $\mathfrak{M}$ and $\mathfrak{M}'$. This closure ordinal is bounded by the number of bisimulation types realised in $\mathfrak{M}$ and $\mathfrak{M}'$. For cardinality reasons it is in particular strictly less than the successor cardinal of $|W| + |W'|$.

Over finite Kripke structures in particular, the limit $\text{gfp}(F_0)$ is reached within a number of iterations bounded by $|W| + |W'|$, whence the largest bisimulation is polynomial time computable.

One verifies by induction that, for $n \in \mathbb{N}$, X_n is the subset

$$X_n = \{(u, u') \in W \times W' \mid (\mathfrak{M}, u) \rightleftharpoons_n (\mathfrak{M}', u')\}$$

and correspondingly that

$$X_\omega = \{(u, u') \in W \times W' \mid (\mathfrak{M}, u) \rightleftharpoons_\omega (\mathfrak{M}', u')\}.$$

Closure within $m := |W| + |W'|$ steps for finite Kripke structures, implies that, in restriction to $\mathfrak{M}$ and $\mathfrak{M}'$, m -bisimulation equivalence $\rightleftharpoons_m$ and hence equivalence in ML_m coincide with full bisimulation equivalence $\rightleftharpoons$ and equivalence in ML . This quantitative analysis provides a direct proof of the Hennessy–Milner theorem (with additional a priori bounds) in the special case of finite (rather than just finitely branching) Kripke structures.

3.8 Bisimulation quotients and canonical representatives

Bisimulation quotients provide canonical minimal bisimilar companions, in which every bisimulation type is realised only once. They thus form succinct representations of the overall bisimulation type of a structure $\mathfrak{M}$. There is an analogy with filtrations (compare section 3.4), but here the quotient is taken with respect to the largest bisimulation within the given structure, rather than with respect to some coarser equivalence induced by some set of modal formulae. Passage to bisimulation quotients is often desirable for complexity

reasons, for instance for model checking of bisimulation invariant properties. Bisimulation quotients of finite structures are polynomial time computable, as the largest bisimulation is polynomial time computable over finite structures as a greatest fixed point.

For a Kripke structure $\mathfrak{M} = \langle W, \{R_\alpha\}_{\alpha \in \tau}, V \rangle$, consider the largest bisimulation within $\mathfrak{M}$ itself, $\rho^{\mathfrak{M}} = \{(u, u') \mid (\mathfrak{M}, u) \rightleftharpoons (\mathfrak{M}, u')\} \subseteq W^2$, as an equivalence relation on W . Let us write $[u]_\rho$ for the equivalence class of $u \in W$. Note that $\rho^{\mathfrak{M}}$ is a congruence w.r.t. the valuation V (by atom equivalence). Therefore V induces a natural quotient valuation $V/\rho^{\mathfrak{M}}$ on the quotient $W/\rho^{\mathfrak{M}}$. While $\rho^{\mathfrak{M}}$ is not in general a congruence w.r.t. the R_α , clearly $(w, u) \in R_\alpha$ implies that for any $w' \in [w]_\rho$ there is $u' \in [u]_\rho$ such that $(w', u') \in R_\alpha$ (by the back and forth conditions). A natural quotient interpretation for the R_α over $W/\rho^{\mathfrak{M}}$ therefore is

$$R_\alpha/\rho^{\mathfrak{M}} := \{([w]_\rho, [u]_\rho) \in (W/\rho^{\mathfrak{M}})^2 \mid (w, u) \in R_\alpha^{\mathfrak{M}}\}.$$

DEFINITION 44. The *bisimulation quotient* $\mathfrak{M}/\rho^{\mathfrak{M}}$ is the Kripke structure with universe $W/\rho^{\mathfrak{M}} = \{[u]_\rho \mid u \in W\}$, accessibility relations $R_\alpha/\rho^{\mathfrak{M}}$ and valuation $V/\rho^{\mathfrak{M}}$.

LEMMA 45. The canonical projection $\pi: W \rightarrow W/\rho^{\mathfrak{M}}$ from $\mathfrak{M}$ onto its bisimulation quotient $\mathfrak{M}/\rho^{\mathfrak{M}}$ is a surjective bounded morphism.

$\mathfrak{M}/\rho^{\mathfrak{M}}$ is minimal among all globally bisimilar companion structures of $\mathfrak{M}$, as any other such must also have at least one representative of each bisimulation type realised in $\mathfrak{M}$. Moreover, any global bisimulation between two such quotient structures is uniquely determined by bisimulation types and is necessarily an isomorphism. The analogue for ordinary (rather than global) bisimulation equivalence of pointed Kripke structures $(\mathfrak{M}, u)$ needs to be based on quotients $\mathfrak{M}[u]/\rho^{\mathfrak{M}}$ taken after restriction to the generated substructure rooted at u . The bisimulation quotient associated with a (pointed) Kripke structure thus provides a *canonical representative* of its bisimulation type, ‘canonical’ in the sense of being uniquely determined up to isomorphism.

COROLLARY 46. Kripke structures $\mathfrak{M}$ and $\mathfrak{M}'$ are globally bisimilar iff their bisimulation quotients are isomorphic. Pointed Kripke structures $(\mathfrak{M}, u)$ and $(\mathfrak{M}', u')$ are bisimilar iff the bisimulation quotients $(\mathfrak{M}[u]/\rho^{\mathfrak{M}}, [u]_\rho^{\mathfrak{M}})$ and $(\mathfrak{M}'[u']/\rho^{\mathfrak{M}'}, [u']_\rho^{\mathfrak{M}'})$ are isomorphic.

For other kinds of canonical representatives of the bisimulation type of a pointed Kripke structure we may look to trees. Via tree unfoldings any pointed Kripke structure is bisimilar to a tree structure. In order to associate a companion tree structure which is uniquely determined up to isomorphism, though, one needs to impose conditions on the multiplicities among bisimilar siblings in the tree. For countably branching Kripke structures, for instance, in which every state has at most countably many immediate successors, ω -branching tree unfoldings $\vec{\mathfrak{M}}^\omega[u]$ may be used. These are defined in complete analogy with ordinary tree unfoldings, cf. Definition 21, but based on the set of all ω -labelled paths rooted at u . An ω -labelled path in $\mathfrak{M}$ is a sequence $\vec{w} = (w_0, \alpha_1, m_1, w_1, \dots, \alpha_k, m_k, w_k)$, where $\vec{w} = (w_0, \alpha_1, w_1, \dots, \alpha_k, w_k)$ is a path in $\mathfrak{M}$ in the usual sense, and with labels $m_i \in \mathbb{N}$. Two ω -labelled paths $\vec{w}, \vec{w}'$ are linked by an R_α -edge in $\vec{\mathfrak{M}}^\omega[u]$ if $\vec{w}'$ is an α -extension of $\vec{w}$: $\vec{w}' = (\vec{w}, \alpha, m, w')$ for some $m \in \mathbb{N}$. Through the ω -labelling, the multiplicity of each bisimulation type in each successor set w.r.t. R_α is countably infinite. It is then easy to see that any two bisimilar ω -branching tree unfoldings of countably branching Kripke structures are isomorphic.

This observation may be extended in a straightforward manner to κ -tree unfoldings $\vec{\mathfrak{M}}^\kappa[u]$ based on κ -labelled paths, for any infinite cardinal κ .

COROLLARY 47. *For any infinite cardinal κ , and pointed Kripke structures $(\mathfrak{M}, u)$ and $(\mathfrak{M}', u')$ whose branching degree is bounded by κ : $(\mathfrak{M}, u) \rightleftharpoons (\mathfrak{M}', u')$ if, and only if, $(\vec{\mathfrak{M}}^\kappa[u], u) \simeq (\vec{\mathfrak{M}}'^\kappa[u'], u')$.*

3.9 Robinson consistency, local interpolation, and Beth definability

We illustrate the usefulness of the canonicity property expressed in Corollary 47 with a proof of the following analogue of the Robinson joint consistency property [12] for poly-modal logic.

PROPOSITION 48 (Robinson consistency). *For $i = 1, 2$ let $\tau^{(i)}$ be modal similarity types; $\Phi^{(i)}$ sets of atomic propositions; and $\Gamma^{(i)} \subseteq \text{ML}[\tau^{(i)}, \Phi^{(i)}]$. If $\Gamma^{(1)} \cap \Gamma^{(2)}$ is a complete modal theory (in the local sense), and if both $\Gamma^{(1)}$ and $\Gamma^{(2)}$ are consistent, then $\Gamma = \Gamma^{(1)} \cup \Gamma^{(2)}$ is also consistent.*

Proof. Let $\tau^{(0)} := \tau^{(1)} \cap \tau^{(2)}$, $\Phi^{(0)} := \Phi^{(1)} \cap \Phi^{(2)}$, $\Gamma^{(0)} := \Gamma^{(1)} \cap \Gamma^{(2)}$.

Let $\mathfrak{M}, u \models \Gamma^{(1)}$ and $\mathfrak{N}, v \models \Gamma^{(2)}$. Without loss of generality assume that both structures are ω -saturated, which implies that also their $(\tau^{(0)}, \Phi^{(0)})$ reducts $\mathfrak{M}^{(0)}$ and $\mathfrak{N}^{(0)}$ are ω -saturated. Then $(\mathfrak{M}^{(0)}, u) \equiv_{\text{ML}} (\mathfrak{N}^{(0)}, v)$, as both satisfy the complete theory $\Gamma^{(0)}$. By the Hennessy–Milner property for ω -saturated structures: $(\mathfrak{M}^{(0)}, u) \rightleftharpoons (\mathfrak{N}^{(0)}, v)$ (cf. Remark 40).

Let $\kappa \geq |\mathfrak{M}|, |\mathfrak{N}|$ and consider the κ -tree unfoldings $\hat{\mathfrak{M}} := \vec{\mathfrak{M}}^\kappa[u]$ and $\hat{\mathfrak{N}} := \vec{\mathfrak{N}}^\kappa[v]$. The generated $(\tau^{(0)}, \Phi^{(0)})$ -subtrees of $(\hat{\mathfrak{M}}, u)$ and $(\hat{\mathfrak{N}}, v)$ are themselves κ -tree unfoldings of corresponding generated $(\tau^{(0)}, \Phi^{(0)})$ -substructures of $\mathfrak{M}$ and $\mathfrak{N}$. So they are isomorphic as $(\tau^{(0)}, \Phi^{(0)})$ -trees. We may assume that $(\hat{\mathfrak{M}}, u)$ and $(\hat{\mathfrak{N}}, v)$ intersect precisely in these isomorphic subtrees. Let then $\mathfrak{K} := \hat{\mathfrak{M}} \cup \hat{\mathfrak{N}}$ be their union (note that $u = v$). The component structures $(\hat{\mathfrak{M}}, u)$ and $(\hat{\mathfrak{N}}, v)$ are the generated $(\tau^{(i)}, \Phi^{(i)})$ -subtrees for $i = 1$ and $i = 2$, respectively. By bisimulation invariance, $\mathfrak{K}, u \models \Gamma^{(i)}$ for $i = 1, 2$. Therefore Γ is satisfiable. $\blacksquare$

Consistency properties can usually be directly related to interpolation [12]. Here we obtain the *local interpolation* theorem for poly-modal logic as a corollary. For modal similarity types and sets of atomic propositions as above: let $\models \varphi \rightarrow \psi$ be a valid (local) consequence, $\varphi \in \text{ML}[\tau^{(1)}, \Phi^{(1)}]$, $\psi \in \text{ML}[\tau^{(2)}, \Phi^{(2)}]$. We want to show that there is an interpolant $\chi \in \text{ML}[\tau^{(0)}, \Phi^{(0)}]$ (i.e., in the common language):

$$\models (\varphi \rightarrow \chi) \wedge (\chi \rightarrow \psi).$$

Assume there was no interpolant. One can then find a complete theory $\Gamma^{(0)}$ in the common vocabulary for which both $\Gamma^{(1)} := \Gamma^{(0)} \cup \{\varphi\}$ and $\Gamma^{(2)} := \Gamma^{(0)} \cup \{\neg\psi\}$ are consistent (see below). With the consistency property established above, however, this would show that $\varphi \wedge \neg\psi$ is satisfiable, invalidating the implication $\varphi \rightarrow \psi$. Assuming without loss of generality that the common language $\text{ML}[\tau^{(0)}, \Phi^{(0)}]$ is countable, one generates $\Gamma^{(0)}$ inductively as a union of an increasing chain of finite sets $\Gamma_n^{(0)}$. The sets $\Gamma_n^{(0)}$ are inductively augmented towards completion, by adding one formula or its negation at a time, guided by the condition that there be no interpolant χ with $\Gamma_n^{(0)} \models (\varphi \rightarrow \chi) \wedge (\chi \rightarrow \psi)$.

COROLLARY 49. *Poly-modal logic satisfies the interpolation theorem for local consequence.*

The interpolation property can be relativised to particular modal logics or classes of frames. We mention one general result of this kind. Following [122], a subframe $\mathfrak{F}$ of the direct product $\prod_{i \in I} \mathfrak{F}_i$ (see section 6.2) is said to be a *bisimulation product* of the family of frames $\{\mathfrak{F}_i\}_{i \in I}$ if the canonical projection $\pi_i : \mathfrak{F} \rightarrow \mathfrak{F}_i$ is a surjective bounded morphism, for each $i \in I$. The following has been established in [122, Thm 2.5.3].

PROPOSITION 50. *Let $\mathcal{K}$ be an elementary class of frames closed under generated subframes and bisimulation products. Then modal logic over $\mathcal{K}$ has interpolation.*

The interpolation property is intimately related to the *Beth definability property* which links *implicit* with *explicit* definability.

Consider a fixed (poly-)modal language $\text{ML}[\tau, \Phi]$. For any list of propositional variables $\mathbf{q}$ from Φ , we denote by $\text{ML}[\mathbf{q}]$ the sublanguage of $\text{ML}[\tau, \Phi]$ restricted to the propositional variables listed in $\mathbf{q}$. Let $p \in \Phi$ be a propositional variable not in $\mathbf{q}$, and $\Gamma = \Gamma(p, \mathbf{q}) \subseteq \text{ML}[p, \mathbf{q}]$ a modal theory. Intuitively, Γ defines p implicitly if it uniquely determines the valuation of p relative to the rest. Formally, let p' be a propositional variable not occurring in $\Gamma(p, \mathbf{q})$ and $\Gamma' = \Gamma(p', \mathbf{q})$ the result of substituting p' for p throughout Γ . Γ *defines p implicitly* if the following is valid (in the sense of local consequence):

$$\Gamma \cup \Gamma' \models p \leftrightarrow p'.$$

On the other hand, p is said to be *explicitly definable* relative to Γ if for some $\varphi(\mathbf{q}) \in \text{ML}[\mathbf{q}]$ (thus, not containing p):

$$\Gamma \models p \leftrightarrow \varphi(\mathbf{q}).$$

Such φ is then called an *explicit definition* of p relative to Γ .

Semantically, implicit definability means that in any pointed Kripke structure $(\mathfrak{M}, u)$ for $\text{ML}(\mathbf{q})$, there is at most one valuation for p such that the resulting expansion $(\mathfrak{M}^p, u)$ satisfies $\Gamma(p, \mathbf{q})$. Thus, in order to show that Γ does *not* define p implicitly it suffices to find two models of $\Gamma(p, \mathbf{q})$ that differ in the valuation of p but are otherwise identical. This is the idea of *Padoa's method* for disproving definability in classical logic.

As for explicit definability in modal logic, Conradie [13] has shown that it can be characterised semantically as follows. p is explicitly definable relative to $\Gamma(p, \mathbf{q})$ iff for every two pointed Kripke structures $(\mathfrak{M}_1, u_1)$ and $(\mathfrak{M}_2, u_2)$ satisfying $\Gamma(p, \mathbf{q})$: $(\mathfrak{M}_1, u_1) \equiv_{\text{ML}[\mathbf{q}]} (\mathfrak{M}_2, u_2) \Rightarrow (\mathfrak{M}_1, u_1) \equiv_{\text{ML}[p, \mathbf{q}]} (\mathfrak{M}_2, u_2)$. Here $\equiv_{\text{ML}[\mathbf{q}]}$ denotes equivalence in $\text{ML}[\mathbf{q}]$. In fact, $\equiv_{\text{ML}[\mathbf{q}]}$ may be replaced for this condition by the corresponding bisimulation relation $\rightleftarrows_{\text{ML}[\mathbf{q}]}$.

Clearly, explicit definability entails implicit definability. *Beth's definability theorem* (proved in the early 1950s for first-order logic) states the converse: *implicit definability entails explicit definability*. A standard proof technique is by reduction to interpolation.

Let $\Gamma(p, \mathbf{q}) \cup \Gamma(p', \mathbf{q}) \models p \leftrightarrow p'$. By compactness, $\gamma(p, \mathbf{q}) \wedge \gamma(p', \mathbf{q}) \models p \leftrightarrow p'$ for some formula γ from Γ (we may assume Γ to be closed under $\wedge$). This implies the (local) validity of

$$\models (\gamma(p, \mathbf{q}) \wedge p) \rightarrow (\gamma(p', \mathbf{q}) \rightarrow p').$$

Local interpolation yields an interpolant $\varphi \in \text{ML}[\mathbf{q}]$ in the common language and thus not containing p or p' , such that both $\models (\gamma(p, \mathbf{q}) \wedge p) \rightarrow \varphi$ and $\models \varphi \rightarrow (\gamma(p', \mathbf{q}) \rightarrow p')$.

Together these two establish that φ explicitly defines p relative to γ and hence relative to Γ . We have thus obtained the following.

COROLLARY 51. *Modal logic satisfies Beth's definability theorem for local consequence.*

The notions of implicit and explicit definability, and the Beth property (just like the interpolation property) admit *global* versions, with respect to the global consequence relation. Global versions easily follow from the local versions formulated above.

For more on interpolation, Beth definability, and algebraic equivalents to them on various extensions of modal logic, see Chapter 8 of this handbook as well as [98, 10, 16] for uniform interpolation in the modal μ -calculus, [73] for a general study of interpolation and its links with amalgamation, [122] for results on interpolation in extended modal languages, and [36] for a comprehensive exposition of the state of the art on interpolation and definability.

3.10 Bisimulation-safe modal operators

It is easy to find examples of bisimilar pointed Kripke structures $(\mathfrak{M}, w) \rightleftharpoons (\mathfrak{M}', w')$, over the modal similarity type with a single modality associated with an accessibility relation R say, such that the corresponding expansions with new accessibility relations interpreted by the converse relations $R^{-1} := \{(u, v) \mid (v, u) \in R\}$ are not bisimilar. On the other hand, $\rho: (\mathfrak{M}, w) \rightleftharpoons (\mathfrak{M}', w')$ for pointed (poly-modal) Kripke structures $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ implies that the same ρ also is a bisimulation for the expansions by accessibility relations generated from the R_α by the constructors provided in propositional dynamic logic PDL: union, composition, star, as well as test (compare Lemma 70). Thus, the question arises: which operations on relations are 'safe for bisimulations', i.e., preserve bisimulations which hold for their arguments? This question was raised and analysed by van Benthem. In particular, he answered that question completely for the case of *first-order definable* operations on binary relations (see [130, Section 5.3], also [5, Section 2.7]). The operation $\sim$ of *domain-complementation* is defined as an operation on binary relations according to $\sim R := \{(x, x) \mid \neg \exists z R x z\}$.

THEOREM 52. *A first-order definable operation $O(R_1, \dots, R_n)$ on binary relations is safe for bisimulation iff it can be constructed from $R_1, \dots, R_n$ using atomic tests $p?$, unions, compositions and the operation of domain-complementation.*

This characterisation was extended in [131] to operations definable in infinitary languages, by allowing infinite unions, too. Since the star operation or iteration, $*$, is definable as an infinite union of compositions, this accounts for the bisimulation safety of PDL as stated above. The notion of bisimulation-safety and the results above were further extended by Hollenberg [71].

4 MODAL LOGIC AS A FRAGMENT OF FIRST-ORDER LOGIC

The embedding of modal logics into a fragment of first-order logic via the standard translation makes results and techniques for that fragment directly available to the analysis of the modal logic. In this section we discuss further aspects of the relationship between modal and first-order logic.

4.1 Finite variable fragments of first-order logic

DEFINITION 53. Over a purely relational vocabulary and for $k \geq 1$ let k -variable first-order logic be the syntactic fragment $\text{FO}^k \subseteq \text{FO}$ consisting of those FO formulae that only use k distinct variable symbols, say $x_0, \dots, x_{k-1}$, free or bound.

Gabbay [32] first observed that the standard translation, with thrifty re-use of variables as presented in section 1.3, embeds basic modal logic ML into FO^2 , the two-variable fragment of first-order logic. (For polyadic modalities of arities up to m , one similarly gets an embedding into the $(m + 1)$ -variable fragment.)

LEMMA 54. *The standard translation based on $\text{ST}(-; x_0)$ and $\text{ST}(-; x_1)$ embeds ML into FO^2 .*

For instance, for ML with a single unary modality $\Diamond$ associated with the binary accessibility relation R , the standard translation operates with an alternate use of two variables, x_0 and x_1 , as in

$$\text{ST}(\Diamond \Box \Diamond p; x_0) = \exists x_1 (R x_0 x_1 \wedge \forall x_0 (R x_1 x_0 \rightarrow \exists x_1 (R x_0 x_1 \wedge P x_1))).$$

It should be noted that this re-use of variable symbols is at odds for instance with a prenex formalisation in first-order logic. It has several other benefits, however, to be discussed below. And even though the embedding into the guarded fragment of first-order logic which has emerged more recently (see section 4.4 below) may have greater explanatory power for some characteristic features of modal logics, the straightforward embedding into finite variable fragments has also been put to good use.

Consider the embedding of basic modal logic into FO^2 . By results of Scott [116] (valid for FO^2 without equality) and Mortimer [102] (with equality), FO^2 has the finite model property. In fact FO^2 has an exponential bound on small models [59]. Therefore, the finite model property for basic modal logic and decidability for satisfiability may be inferred via the translation into FO^2 . The complexity and small model bounds obtained in this way, however, are not optimal.

The fact that ML embeds into a finite-variable fragment also provides upper bounds on its model checking complexity. Consider the so-called combined complexity of checking whether $\mathfrak{M}, w \models \varphi$, with both the finite structure $\mathfrak{M}$ and the formula φ as input. The standard translation of modal logic into FO is itself linear time computable. While the combined model checking complexity for FO over finite relational structures is complete for Pspace, it becomes Ptime for FO^k . Moreover, even for FO^2 and basic modal logic the problem is Ptime-hard. For FO^2 one also obtains a bound of $\mathcal{O}(|\varphi||\mathfrak{M}|)$, linear in both input components.⁸ FO^2 thus constitutes a natural syntactic fragment of classical first-order logic which matches the finite model property, the decidability and model checking complexity of basic modal logic. These parallels and their limitations are further discussed in [60, 135, 57].

Remark. At the level of FO^3 and higher, which becomes relevant for instance for polyadic modalities, the target logic FO^k fails to have the finite model property and is just as undecidable for satisfiability as full first-order logic, and also does not have linear time

⁸This bound refers to a random access model of computation and a succinct representations of the binary accessibility relations R_α through adjacency lists. The input size for the structure is then linear in the number of states plus the number of accessibility edges.

model checking. For many purposes, including satisfiability and model checking, however, natural reductions from polyadic into unary modal logics are available that still make the special status of the two-variable fragment available for polyadic modal logics. See for instance [41].

The k -variable fragments of FO play an interesting role in finite model theory and for algorithmic issues, primarily because they give rise to natural and algorithmically manageable *pebble games*. The k -pebble game is precisely the variant of the classical (first-order) Ehrenfeucht–Fraïssé game associated with the restriction to k variable symbols. Bisimulation games in their turn may be regarded as restrictions of these k -pebble games.

Just as the ordinary n -round Ehrenfeucht–Fraïssé game captures elementary equivalence up to quantifier rank n [26, 68, 108], and just as the n -round bisimulation game captures modal equivalence in ML_n , so the n -round k -pebble game captures equivalence in FO^k up to quantifier rank n .

In the classical (variable-unconstrained) Ehrenfeucht–Fraïssé game for FO over two relational structures $\mathfrak{A}$ versus $\mathfrak{A}'$, the players, **I** and **II**, mark finite configurations of elements in these structures with matching pebbles. A configuration in the game is specified by two tuples of marked elements $\mathbf{a}$ in $\mathfrak{A}$ and $\mathbf{a}'$ in $\mathfrak{A}'$, denoted $(\mathfrak{A}, \mathbf{a}; \mathfrak{A}', \mathbf{a}')$. In each round, **I** chooses one of the structures, and places another marker on one of the elements of that structure; **II** has to respond by marking an element in the opposite structure. In one round the game thus proceeds from a configuration $(\mathfrak{A}, \mathbf{a}; \mathfrak{A}', \mathbf{a}')$ to some configuration $(\mathfrak{A}, \mathbf{a}, a; \mathfrak{A}', \mathbf{a}', a')$ with newly pebbled elements a and a' . **II** loses as soon as the partial map induced by the correspondence between pebbled elements $f: \mathbf{a} \mapsto \mathbf{a}'$ is not a local isomorphism. The existence of a winning strategy for **II** in the n -round game then precisely captures elementary equivalence up to quantifier rank n .

The variant for FO^k is obtained by changing the rules in such a manner that no more than k elements of each structure are ever pebbled simultaneously; the game is restricted to configurations $\mathfrak{A}, \mathbf{a}; \mathfrak{A}', \mathbf{a}'$ with tuples $\mathbf{a}$ and $\mathbf{a}'$ of lengths up to k . In any round starting from a configuration of full length k , **I** first removes one of the pebbles and then repositions that same pebble in its structure, and **II** has to do likewise with the matching pebble in the opposite structure. This game then captures levels of equivalence in FO^k , [25].

It is an obvious consequence of Lemma 54 that equivalence in FO^2 implies equivalence with respect to basic modal logic with unary modalities. However, this may also be inferred directly at the level of the games. One observes that the relevant bisimulation game can be emulated by the 2-pebble game in the sense that

- any challenge available to player **I** in the modal game is also available in the 2-pebble game.
- any responses for **II** that are good for the 2-pebble game are good in the modal game, too.

A move along an R -edge in the bisimulation game is emulated in the two-pebble game by means of a placement of the second pebble in the target node. The formerly active pebble now only plays the auxiliary role to guarantee that the right kind of edge is used in an admissible manner also in the response by **II**. But, clearly a strategy in the two-pebble game guarantees more than just bisimulation equivalence, illustrating the gap in expressive power between modal logic and the two-variable fragment of first-order logic

into which it can be embedded.

Consider the expressive power of basic (poly-modal) ML over corresponding Kripke structures $\mathfrak{M} = \langle W, \{R_\alpha\}_{\alpha \in \tau}, \{P_i\} \rangle$. Unlike ML, FO^2 formulae generally define binary predicates over Kripke structures. However, the expressive power of FO is also very limited in this respect. As can be inferred from the 2-pebble game, any FO^2 -formula $\varphi(x_0, x_1)$ is logically equivalent to a Boolean combination of quantifier free formulae of FO^2 (atomic formulae, including equality) and FO^2 formulae in a single free variable $\psi(x_i)$, $i = 0, 1$. In other words, the expressive power of FO^2 , too, is essentially governed by its expressive power in terms of unary relations (properties of single elements in Kripke structures, state properties in process logics). Comparing the expressive power of basic modal logic ML with that of FO^2 for defining properties of elements and the discriminating powers of bisimulation versus two-pebble game equivalence, basic modal logic is lacking

- (i) relativised quantification along backward R_α -edges.
- (ii) quantification relativised by (positive or arbitrary) boolean combinations of accessibility relations (including equality).
- (iii) unrelativised, global first-order quantification in one variable.

Corresponding features can be added to basic modal logic, as for instance through extensions via *inverse modalities* (in temporal settings: past modalities) interpreted w.r.t. to the converses $R_\alpha^- = \{(v, u) \mid (u, v) \in R_\alpha\}$; a *global modality* interpreted w.r.t. to the full binary relation $U = W \times W$ over universe W ; or other constructors for derived accessibilities.⁹ An extension of basic modal logic that provides a minimal set of constructs in the above vein so as to precisely capture the expressive power of FO^2 , is provided in [96]. A comparison of the satisfiability problems of these two logics shows that there is no polynomial time translation from FO^2 into its modal counterpart, under suitable complexity assumptions. Furthermore, on certain classes of frames extended modal logics can reach the full expressiveness of first-order logic. The most prominent example is Kamp's result in [80] that the temporal language with Since and Until is expressively complete for all first-order definable connectives on the class of Dedekind complete linear orders. This line of work was further developed by others, including Stavi, Gabbay, Venema, Reynolds. For further details see [32, 34], as well as Chapter 11 of this handbook.

4.2 Van Benthem's characterisation theorem

The fundamental observation that modal logics are embedded into (fragments of) first-order logic via the standard translation immediately calls for the following question. Given an arbitrary first-order formula (in an appropriate vocabulary of Kripke structures), under which conditions is it equivalently expressible in modal logic? In other words, precisely which first-order properties of Kripke structures are expressible in modal logic? Bisimulation invariance is obviously a necessary condition; van Benthem's Theorem says that it is sufficient as well.

Another point of view is also illuminating. Take bisimulation invariance as the fundamentally important semantic notion. It deserves this status for many non-logical reasons, since it is *the* natural notion of process equivalence (thinking of Kripke structures as tran-

⁹An example of that is the union (*or*) on program formulae in PDL, which, however, is reducible to plain ML in this context, since, for instance $[\alpha \cup \beta]\varphi \equiv [\alpha]\varphi \vee [\beta]\varphi$ and $\langle \alpha \cup \beta \rangle \varphi \equiv \langle \alpha \rangle \varphi \vee \langle \beta \rangle \varphi$.

sition systems), game equivalence (transition systems for games), knowledge equivalence (Kripke structures for knowledge representation), et cetera. From the perspective of first-order logic, then, one would want to isolate the bisimulation invariant properties because just these conform with the underlying semantic intuition. For instance, a first-order property of transition systems captures a property of processes if, and only if, it does not distinguish between bisimulation equivalent transition systems.

Bisimulation invariance is not a decidable property of first-order formulae, as can be seen through reduction of the satisfiability problem. For $\varphi \in \text{FO}(\tau_\Phi \setminus \{R\})$, the formula $\hat{\varphi}(x) := \varphi \wedge Rxx$ is bisimulation invariant iff φ is unsatisfiable. The syntactic subset consisting of those first-order formulae that happen to be bisimulation invariant is therefore *not* the syntax of a reasonable logic. But the characterisation theorem says that modal logic precisely fills this gap. ML, or its translation into FO, provides decidable syntax for just the bisimulation invariant first-order properties; ML is the first-order logic for bisimulation respecting properties.

THEOREM 55 (van Benthem). *Let $\varphi(x) \in \text{FO}$ be in a vocabulary of Kripke structures. Then the following are equivalent:*

- (i) φ is bisimulation invariant: $(\mathfrak{M}, w) \rightleftharpoons (\mathfrak{M}', w')$ implies $\mathfrak{M}, w \models \varphi \Leftrightarrow \mathfrak{M}', w' \models \varphi$.
- (ii) $\varphi(x)$ is logically equivalent to a formula $\tilde{\varphi} \in \text{ML}$.

Note that (ii) $\Rightarrow$ (i) is just Theorem 14 again. The crucial point here is *expressive completeness* of ML for *all* bisimulation invariant first-order properties. The core idea for that is to establish the following – which is reminiscent of a compactness property.

LEMMA 56. *If $\varphi(x) \in \text{FO}$ is bisimulation invariant, then it is invariant under n -bisimulation for some $n \in \mathbb{N}$.*

The lemma implies (i) $\Rightarrow$ (ii) in the theorem, as any n -bisimulation invariant property is clearly definable in ML_n . Indeed, by Corollary 34, φ is then equivalent to a disjunction of characteristic formulae for n -bisimulation equivalence classes. For the lemma, we sketch a version of the classical proof and an alternative argument more closely based on the games.

Via classical model theory. Assume to the contrary that φ was not invariant under n -bisimulation for any $n \in \mathbb{N}$, and hence not equivalent to any modal formula. Enumerate all modal formulae of the appropriate type as $(\psi_i)_{i \in \mathbb{N}}$. Successively choose one of ψ_i or $\neg\psi_i$ to obtain a maximally consistent set T of modal formulae consistent with both φ and $\neg\varphi$. By compactness one obtains pointed Kripke structures $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ such that both satisfy T , while $\mathfrak{M}, w \models \varphi$ and $\mathfrak{M}', w' \models \neg\varphi$. As $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ satisfy the same complete modal theory, $(\mathfrak{M}, w) \equiv_{\text{ML}} (\mathfrak{M}', w')$ and therefore $(\mathfrak{M}, w) \rightleftharpoons_\omega (\mathfrak{M}', w')$. Passage to ω -saturated (or modally saturated, see section 6.3) elementary extensions of $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ would then give us structures $(\tilde{\mathfrak{M}}, w) \rightleftharpoons (\tilde{\mathfrak{M}}', w')$ (cf. Remark 40), which are still distinguished by φ , contradicting bisimulation invariance of φ .

Via games. This alternative proof of the crucial step towards the characterisation theorem admits ramifications that persist where the classical argument fails, in particular in finite model theory. In its present form this argument is based on [105, 107] building on ideas from Rosen's finite model theory version of the characterisation theorem [112], as further discussed in sections 4.3 and 9. The n -neighbourhood of an element u in a Kripke structure $\mathfrak{M}$ consists of all elements whose *Gaifman distance* from u is at most n . Here Gaifman distance is graph theoretic distance in the undirected graph induced by

the symmetrised accessibility relation. We write $\mathfrak{M} \upharpoonright U^n(u)$ for the induced substructure on the n -neighbourhood of u in $\mathfrak{M}$.


DEFINITION 57. A formula $\varphi(x)$ is *n -local* if for any two pointed tree Kripke structures $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ that are isomorphic in restriction to the n -neighbourhoods of their distinguished nodes, $\mathfrak{M}, w \models \varphi \Leftrightarrow \mathfrak{M}', w' \models \varphi$.

It is easy to see that, if $\varphi(x)$ is bisimulation invariant and n -local, then it is n -bisimulation invariant. In fact, this is obvious for trees and then extends to arbitrary pointed Kripke structures through their unfoldings into trees (see section 2.2). See [105, 106] for the following.

LEMMA 58. *Let $\varphi(x) \in \text{FO}$ have quantifier rank q . If φ is bisimulation invariant, then it is n -local, and hence invariant under n -bisimulation, for $n = 2^q - 1$.*

The first-order locality argument is in fact a ramification of the much more general *Gaifman locality* property of first-order logic [38], which is a useful tool in classical as well as finite model theory [25]. In the context of bisimulation invariant properties, locality together with the exponential bound may however also be derived from a straightforward and self-contained analysis based on first-order Ehrenfeucht–Fraïssé games. In fact, the lemma holds for any $\varphi(x)$ that is invariant under disjoint unions, which itself is an easy consequence of bisimulation invariance (see section 2.2). Let $\varphi(x) \in \text{FO}$ have quantifier rank q . Consider a pointed Kripke structure $(\mathfrak{M}, w)$ or, because it may be conceptually easier though not necessary for the argument, without loss of generality a pointed tree structure $(\mathfrak{M}, w)$ with root w . Let $\mathfrak{M}' = \mathfrak{M} \upharpoonright U^n(w)$ be the substructure induced on the n -neighbourhood of w . It suffices to show that $\mathfrak{M}, w \models \varphi$ iff $\mathfrak{M}', w \models \varphi$.

Let $\mathfrak{N}$ be the disjoint union of q copies of $\mathfrak{M}$ and $\mathfrak{M}'$ each. Using invariance under disjoint unions, it suffices to show that $\mathfrak{N} \uplus \mathfrak{M}, w \models \varphi$ iff $\mathfrak{N} \uplus \mathfrak{M}', w \models \varphi$.



It is not hard to exhibit a winning strategy for **II** in the ordinary q -round Ehrenfeucht–Fraïssé game on these structures. **II** merely needs to respect, in round m of the game, the critical distance $d_m = 2^{q-m}$: if **I**'s move in round m goes to within distance d_m of an already pebbled element, **II** plays according to a local isomorphism in the d_m -neighbourhoods of previously pebbled elements; if **I**'s move goes to an element further away from *all* previously pebbled elements, **II** responds in a fresh isomorphic copy of type $\mathfrak{M}$ or $\mathfrak{M}'$, correspondingly.

The exponential bound expressed in the lemma is actually optimal. For a bisimulation invariant property expressible in FO_q but not in ML_n for any $n < 2^q - 1$ consider the property that a state in which p holds is reachable on a path of length less than 2^q . It should be noted that the classical proof of van Benthem's theorem provides no corresponding quantitative information.

COROLLARY 59. *For $\varphi(x) \in \text{FO}$ of quantifier rank q , the following are equivalent for $n = 2^q - 1$:*

- (i) *φ is bisimulation invariant.*
- (ii) *φ is invariant under n -bisimulation and equivalently expressible in ML_n .*

The exponential bound on the modal nesting depth is sharp: FO is exponentially more succinct than ML for expressing bisimulation invariant properties.

Some of the underlying ideas of these results are very robust and extend to various ramified settings, some of which are to be discussed in section 5. The classical proof of the characterisation theorem, in particular, carries through for many natural extensions of basic modal logic associated with refined notions of basic bisimulation equivalence; we mention in particular the corresponding characterisation theorem for the guarded fragment of first-order logic [1] (see Theorem 65 here). But also the game based approach extends to a wide range of settings. One of its main strengths is that it goes through in the setting of finite model theory, as explained below. Another variation that comes naturally from the game based proof is its relativisation to arbitrary bisimulation closed classes [105]. The classical proof, on the other hand, clearly relativises to elementary classes of structures.

COROLLARY 60. *Let $\mathcal{C}$ be a class of Kripke structures that is closed under bisimulation. Then $\varphi(x) \in \text{FO}$ is bisimulation invariant in restriction to $\mathcal{C}$ iff it is equivalent to a formula $\tilde{\varphi} \in \text{ML}$ in restriction to $\mathcal{C}$. Similarly for any elementary class $\mathcal{C}$.*

Theorem 55 characterises the elementary properties of pointed Kripke structures which are definable by single modal formulae. In section 6.4 we will obtain more general preservation results, characterising properties and classes of Kripke structures which are definable by finite or infinite sets of modal formulae, by employing constructions and results from classical model theory.

4.3 Ramifications of the characterisation theorem

We sketch a version of the game and locality based proof of van Benthem's characterisation theorem given above, which applies in finite model theory as well as classically. We thus get the finite model theory version due to Rosen [112], even with the same tight exponential bound on succinctness as in Corollary 59.

THEOREM 61. *For any $\varphi(x) \in \text{FO}$ of quantifier rank q , the following are equivalent:*

- (i) *φ is bisimulation invariant over finite Kripke structures.*
- (ii) *φ is equivalent to a formula of ML_n over finite Kripke structures, for $n = 2^q - 1$.*

Proof. We merely need to adapt the proof outlined in section 4.2 in minor ways to avoid passage through infinite structures. For that we may replace bisimilar companion tree structures by the finite, local versions provided by Lemma 36 rather than full unfoldings. For the proof of n -locality of φ (cf. Lemma 58) no modifications are necessary in the game argument, as it applies to arbitrary relational structures exactly as for trees. In fact we only need to use partial tree unfoldings to argue that n -locality and bisimulation invariance together imply n -bisimulation invariance also in restriction to finite structures, as follows.

Let $\varphi(x)$ be bisimulation invariant and n -local over finite structures. Consider finite structures $(\mathfrak{M}, w) \rightleftharpoons_n (\mathfrak{M}', w')$. We need to show that $\mathfrak{M}, w \models \varphi$ iff $\mathfrak{M}', w' \models \varphi$. As φ is bisimulation invariant, we may replace $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ by bisimilar finite companion structures whose restrictions to $U^n(w)$ and $U^n(w')$ are trees, by Lemma 36. As φ is n -local, these structures may further be replaced by their restrictions to the n -neighbourhoods of w and w' , which are n -bisimilar tree structures of depth n , hence bisimilar. So φ is true in w iff it is true in w' . ■

Compare section 9 for further discussion of the finite model theory context; for further ramifications concerning modal logics based on refined notions of bisimulation also compare section 5; for relativisations to other non-elementary classes of frames, see in particular [17].

4.4 Guarded fragments of first-order logic

The standard translation (see section 1.3) immediately suggested finite variable fragments as an appropriate framework for the study of modal logic within first-order logic. In some ways, however, the finite variables feature fails to give satisfactory insights into the model theoretic behaviour of modal logics. The comparatively smooth finite model theory (see section 9) of modal logics and most notably also their decidability properties (considering robustness under extensions [135]; see section 5.2) are not reflected by the finite variable fragments or even FO^2 in particular [61, 60].

Guarded fragments of first-order logic were introduced by Andr  ka, van Benthem and N  meti in [1]. Compared to the finite variable fragments, the *guarded fragment* GF of first-order logic is much closer to the qualitative characteristics of modal logics. It has greater explanatory power as a framework for the study of modal logics within first-order. For instance, GF and some of its further extensions mirror the decidability as well as finite and tree model properties of modal logics. Crucially, there is a natural notion of *guarded bisimulation* at the root of some of these features. On the other hand guarded logics considerably extend the expressive power of standard modal logics, and in particular still encompass many of their important extensions. Guarded logics have thus come to play an important role in the quest for more expressive fragments of first-order logic that share many of the model theoretic and algorithmic properties that make modal logics so useful for various applications. GF and its relatives extend the scope of essentially modal model theory, including algorithmic and finite model theory aspects, in the direction of first-order.

The guarded fragment GF of FO generalises the *relativised nature* of modal quantification. Let $\alpha(\mathbf{x}, \mathbf{y})$ be an atomic first-order formula in variable tuples as displayed, and consider existential and universal quantification over variables $\mathbf{y}$ where the range of quantification is restricted to those $\mathbf{y}$ that satisfy $\alpha(\mathbf{x}, \mathbf{y})$ in relation to $\mathbf{x}$ (α is called the *guard* of the quantification). The following shorthand syntax is useful for this α -relativised quantification:

$$(\exists \mathbf{y}.\alpha)\varphi := \exists \mathbf{y}(\alpha(\mathbf{x}, \mathbf{y}) \wedge \varphi(\mathbf{x}, \mathbf{y})), \text{ and its dual } (\forall \mathbf{y}.\alpha)\varphi := \forall \mathbf{y}(\alpha(\mathbf{x}, \mathbf{y}) \rightarrow \varphi(\mathbf{x}, \mathbf{y})).$$

Modal quantification (or its standard translation into first-order) displays just this kind of relativisation, where the guards are the atoms $R_\alpha xy$ for accessibility relations

R_α .¹⁰ GF admits relativised quantification of this kind, for any atom α , provided that the variables that occur in α comprise *all* the free variables in the formula φ that is being quantified. The standard translation of modal logics (section 1.3) clearly obeys these restrictions.

DEFINITION 62. For an arbitrary relational vocabulary τ , the formulae of $\text{GF}(\tau) \subset \text{FO}(\tau)$, the guarded fragment, are generated from the atomic formulae by closure under boolean connectives and *guarded quantification*; i.e., if $\varphi(\mathbf{x}, \mathbf{y}) \in \text{GF}(\tau)$ and if $\alpha(\mathbf{x}, \mathbf{y})$ is a τ -atom (also allowing equality) such that $\text{free}(\varphi) \subseteq \text{var}(\alpha)$, then $(\forall \mathbf{y}. \alpha)\varphi(\mathbf{x}, \mathbf{y})$ and $(\exists \mathbf{y}. \alpha)\varphi(\mathbf{x}, \mathbf{y})$ are also in $\text{GF}(\tau)$.

The atom α in these last formulae is called the *guard* of the (universal or existential) quantification. The *nesting depth* is declared for formulae of GF similar to the first-order quantifier rank, with the only exception that it increases by just 1 with every guarded quantification (rather than by the number of quantified variables in $\mathbf{y}$). The semantics of GF is just that of first-order logic. It makes sense, however, to look at the crucial restriction with a view to a semantic understanding.

DEFINITION 63. Let $\mathfrak{A}$ be a τ -structure. A subset $s \subseteq A$ is *guarded* if s is a singleton set or if $s = \{a_1, \dots, a_k\}$ for some tuple $(a_1, \dots, a_k) \in R^{\mathfrak{A}}$ for some relation $R \in \tau$. A tuple $\mathbf{a}$ over $\mathfrak{A}$ is *guarded* if its components are elements of some common guarded subset.

Guarded quantification essentially is quantification over guarded tuples. Intuitively, only the elements of guarded subsets are simultaneously visible in the guarded perspective; this intuition is borne out in the concept of guarded bisimulation (see Definition 64 below).

Clearly the standard translation embeds ML into GF, and actually into the two-variable fragment of GF, $\text{GF} \cap \text{FO}^2$, which is strictly between ML and FO^2 in expressive power, comprising some but not all the features that separate ML from FO^2 as discussed at the end of section 4.1 above. GF naturally comprises

- (i) inverse (or past) modalities, as guardedness is non-directional.
- (ii) positive Boolean operations on accessibilities (including equality), as for instance in $[\alpha \cap \beta]\varphi \equiv (\forall y. (R_\alpha xy \wedge R_\beta xy))\varphi(y) \equiv (\forall y. R_\alpha xy)(R_\beta xy \rightarrow \varphi(y))$.
- (iii) a global modality, or universal/existential quantification over a single free variable, as any singleton set is guarded.

Moreover, it should be noted that GF is genuinely polyadic in the sense of representing no restriction on the arities of definable predicates, whereas even polyadic modal logics are still monadic in that sense. But GF indirectly also has a finite variable nature to it. Note that guarded sets are bounded in size by the width (maximal arity) of the available relation symbols. It is not hard to show that any formula in $\text{GF}(\tau)$, for τ of width k , is equivalent to a boolean combination of atomic formulae and formulae that are in $\text{GF} \cap \text{FO}^k$ (up to a possible renaming of variables).

Guarded bisimulations form the backbone of the model theory of GF, playing the same role for GF that ordinary bisimulations play for modal logics. In essence a guarded bisimulation is a back-and-forth equivalence based on local isomorphisms between guarded subsets.

¹⁰This is good also in the polyadic case, where an n -ary modality α associated with an $(n+1)$ -ary relation R_α gives rise to quantification with guard $R_\alpha(x, \mathbf{y})$.

DEFINITION 64. A *guarded bisimulation* between τ -structures $\mathfrak{A}$ and $\mathfrak{B}$ is a non-empty set Z of local (partial) isomorphisms between $\mathfrak{A}$ and $\mathfrak{B}$ such that

- (i) for every $\rho \in Z$, the domain and image of ρ are guarded subsets of $\mathfrak{A}$ and $\mathfrak{B}$, respectively.
- (ii) Z satisfies the following back-and-forth conditions w.r.t. guarded subsets:
forth: for every $\rho \in Z$ with domain s and every guarded subset s' of $\mathfrak{A}$, there is some $\rho' \in Z$ with domain s' such that ρ and ρ' agree on $s \cap s'$.
back: analogously, w.r.t. to the inverse maps ρ^{-1} and for guarded subsets of $\mathfrak{B}$.

Guarded bisimulations preserve the semantics of GF just as bisimulations preserve the semantics of ML. Moreover, bounded guarded bisimulations – best defined in terms of the restriction of corresponding guarded bisimulation games to a fixed finite number of rounds – precisely capture the levels of equivalence w.r.t. guarded formulae of corresponding nesting depth. Finally, GF is semantically characterised as a fragment of FO precisely through guarded bisimulation invariance. This analogue of van Benthem’s characterisation of modal logic is due to Andr  ka, van Benthem and N  meti [1].

THEOREM 65. *For any first-order formula $\varphi \in \text{FO}(\tau)$ the following are equivalent:*

- (i) *φ is invariant under guarded bisimulation.*
- (ii) *φ is logically equivalent to a formula $\tilde{\varphi} \in \text{GF}(\tau)$.*

It is interesting to note that the full analogue of this characterisation theorem in finite model theory is currently still open. For relational vocabularies of width up to two (essentially coloured directed graphs), the analogue is proved in [106].

Perhaps the most important model theoretic consequence of an analysis of GF w.r.t. guarded bisimulations is a corresponding generalisation of the tree model property. For arbitrary relational structures one obtains guarded bisimilar companion structures through a process of *guarded unravelling* or *unfolding*. These relational structures are close to trees in being *tree-decomposable* by means of guarded subsets. Tree decompositions provide a representation of the underlying relational structure by a tree. This notion from graph and hypergraph theory (see for instance [4]) has been fruitfully employed in relational structures also in applications to relational databases [3]. Tree representations based on guarded subsets work with tree structures whose nodes describe all the guarded substructures of the given structure. Guarded unravellings [56, 58] provide tree decompositions by guarded subsets. As the size of guarded subsets in τ -structures is bounded by the width of τ (the maximal arity of relations in τ), one automatically obtains a bound on the tree width. The resulting *generalised tree model property* from [56] is the following.

THEOREM 66. *Any satisfiable formula $\varphi \in \text{GF}(\tau)$ has a model which is tree decomposable in terms of its guarded subsets and consequently of tree width $m - 1$, where m is the width of τ .*

Such a generalised tree model property can be of eminent model theoretic importance, especially with a view to algorithmic questions, because properties of tree decomposed models may be determined in terms of their tree representations. Using classical model theoretic tools for trees, and in particular automata theoretic methods, the generalised tree model property has strong consequences for decidability and complexity issues. For instance, GF and some of its extensions beyond first-order logic that are invariant un-

der guarded bisimulation and hence satisfy the generalised tree model property, can be decided for satisfiability via reductions to the monadic second-order theory of trees (Rabin's theorem). A direct reduction to emptiness problems for suitable tree automata moreover typically yields optimal complexity bounds. Even finite models for formulae of GF can be built from infinite tree-like models, using finite saturation arguments based on Herwig's extension theorem for partial isomorphisms [67], thus providing an elegant proof of the finite model property for GF [56].

THEOREM 67. *Any satisfiable formula of GF has a finite model: GF has the finite model property.*

The *clique guarded fragment* pushes the basic idea of guarded quantification a bit further by relaxing the notion of guarded subsets. A subset s of a relational structure is *clique guarded* if any pair of elements from the set is guarded (the subset forms a clique in the Gaifman graph). In the clique guarded fragment, quantification is restricted to clique guarded rather than guarded subsets. The resulting logic naturally embeds the first-order translation of the **Until** operator of temporal logic:

$$(\varphi \text{ Until } \psi)(x) \equiv \exists y(x \leq y \wedge \psi(y) \wedge \forall z((x \leq z \wedge z < y) \rightarrow \varphi(z))),$$

because the relevant x, y, z triples form cliques w.r.t. comparability under $\leq$. The clique guarded fragment is no longer restricted to finite variables as clique guarded subsets can have any size. (The **Until** operator, which crucially requires three variables, is expressible in terms of clique guarded triples w.r.t. a binary relation.)

Despite the increase in expressiveness, the clique guarded fragment is still decidable for satisfiability [56] and it also satisfies the finite model property [69, 70] (with links between clique guardedness and extension theorems for partial isomorphisms).

5 VARIATIONS, EXTENSIONS, AND COMPARISONS OF MODAL LOGICS

There is a considerable body of work on ramifications of the familiar classical modal logics. At the level of ordinary semantics in (pointed) Kripke structures or transition systems, many variations and extensions have been proposed. These largely aim at preserving some of the key model theoretic features of basic modal logics while adapting or boosting the expressive power – either for the purposes of a systematic investigation or for the modelling of situations that cannot be captured by the standard modal languages. The many application areas of modal logics contribute to interesting ramifications and continue to trigger new developments. We give but a few examples. Variants of basic modal languages for the purposes of description logics, as treated in depth in Chapter 13 of this handbook, naturally use for instance inverse modalities (for inverse roles) or graded modalities (for number constraints). Various constructors for new modalities based on composite accessibility relations (e.g., relational composition or transitive closures) have long been studied in temporal and process logics (see Chapters 11 and 12 among others). More recently similar extensions have been employed in formalisms developed for the navigation and retrieval of information in data formats like XML (see [100]).

While a more comprehensive concept of a generalised modal model theory may lead to further consolidation of the big picture, we can here only attempt to exemplify some simple model theoretic ideas in this direction. For a tentative framework, let us regard the underlying notion of bisimulation invariance as the key feature of a specifically modal

model theory (at the level of Kripke semantics). We may then tentatively explore this theme along two axes: *variations* in the sense of variations of the underlying notion of bisimulation; and *extensions* of expressive power subject to the requirement of invariance w.r.t. the given notion of bisimulation.

For two typical examples of these orthogonal directions consider, on the one hand, the addition of past modalities (backward moves in the bisimulation game), and, on the other hand, the extension by path quantification (as for reachability assertions or unbounded iteration of $\Diamond$).

For this largely informal sketch we limit ourselves to just a few logics that play a prominent role in connection with transition systems and the behaviour of processes. Some of these and many others are treated at much greater depth in other chapters of this handbook, in particular Chapters 11, 12 and 17 of this handbook and several others in Parts 3 and 4. As criteria for the model theoretic character of the logics under consideration, over and above their expressive power, we look in particular at the corresponding bisimulation games and model theoretic characterisation theorems, at the tree model property and the finite model property, and at satisfiability issues, which are particularly relevant in many applications (compare Chapters 3 and 17 of this handbook).

5.1 Variations through refined notions of bisimulation

A refinement of bisimulation equivalence ought to be matched, on the logic side, by a more expressive logic. We thus encounter extensions of basic modal logic to more expressive fragments of first-order logic, like those considered in sections 4.1 and 4.4.

In terms of the bisimulation game (or the back-and-forth conditions) over Kripke structures with binary accessibility relations one can introduce a variety of additional moves, in order to capture the expressiveness of some natural extensions of basic modal logic, for instance:

- unconstrained moves to arbitrary states (global bisimulation). This corresponds to the addition of a universal modality (or $\forall/\exists$ quantification) to basic ML, which also allows for an explicit transition between global and local semantics (see, e.g., [53, 22]).
- backward moves along edges (two-way bisimulation). This corresponds to the addition of past or inverse modalities to basic ML.
- counting moves, in which the number of available responses is controlled (counting or locally bijective bisimulation). This corresponds to the extension of basic ML by graded or counting modalities (see [21]).

(Also compare [88] for bisimulations for a hierarchy of description logic languages).

In terms of further reaching variations that also involve the format of the underlying structures and game positions, we discussed in section 4.4 guarded bisimulations for arbitrary relational structures – corresponding to guarded rather than ordinary modal quantification and guarded fragments of first-order logic as important intermediaries between modal and first-order logics.

As indicated, these variations typically correspond to natural extensions of ML. These correspondences manifest themselves in terms of

- (i) Ehrenfeucht–Fraïssé relationships: equivalence in the extended logic is characterised by the existence of winning strategies for player **II** in the corresponding, refined bisimulation games.
- (ii) characterisation theorems in the style of Theorems 55 or 65 that characterise the respective logic as a fragment of first-order logic, in terms of invariance under the refined notion of bisimulation.

For instance, the *global bisimulation game* gives player **I** the option to switch, for an individual round, to moves in which both players are allowed to move the pebbles to any element of the respective structure rather than just along accessibility edges. This is the Ehrenfeucht–Fraïssé game for the extension $\text{ML}[\forall]$ of basic ML, in which a global modality is available (corresponding to unrestricted universal first-order quantification in the standard translations). Then **II** has a winning strategy for the n -round game on $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ iff $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$ satisfy exactly the same formulae in $\text{ML}[\forall]$ of quantifier rank up to n . Also the classical proof pattern for the characterisation theorem (compare the classical proof argument for Lemma 56) goes through. This uses compactness and ω -saturated or modally saturated extensions and the analogue of Remark 40, which is good also for this refined bisimulation game. So we have obtained the following.

PROPOSITION 68. *A first-order formula $\varphi(x)$ is invariant under global bisimulation iff it is equivalent to a formula of $\text{ML}[\forall]$.*

This proposition may serve as a representative for a whole family of similar characterisation results for many other variants of basic modal logic. In fact, these game techniques are not at all even restricted to the modal setting. Analogous Ehrenfeucht–Fraïssé and characterisation theorems hold for instance also for the finite variable fragments FO^k in relation to k -pebble game equivalence. Interestingly, as far as the characterisation theorems are concerned, the picture becomes more varied when we shift attention to the finite model theory versions (cf. section 9, in particular Theorem 130).

5.2 Extensions beyond first-order

Extensions induced by variations of the underlying notion of bisimulation in the first instance all lead to modal logics of (pointed) Kripke structures that are still fragments of first-order logic. There is the orthogonal direction of extension that adds expressiveness through stronger constructors in the logic while still adhering to invariance under the given notion of bisimulation. These extensions address some of the expressive deficiencies inherent in first-order, in particular its restriction to essentially local properties (in the sense of Gaifman’s locality theorem). Major process logics, aimed at formalising dynamic properties of processes in terms of Kripke structures as transition systems, need to express fundamental properties – like reachability or well-foundedness – that are non-local and hence not expressible in FO.

The process logics discussed below specifically aim for the formalisation of properties of programs or processes, based on the modelling of states and state transitions in Kripke structures as *transition systems*: atomic propositions model atomic *state properties*, and accessibility relations between states model atomic state transformers or *atomic programs*. This setting calls for logics of a fundamentally modal nature – especially since the intended processes are captured by transition systems only up to bisimulation equivalence. Bisimilar transition systems describe exactly the same processes in the sense that

there is a complete correspondence of possible runs at the level of individual transitions and in terms of mutual step-wise simulation (*bi-simulation*).

We fix a finite similarity type with modalities α corresponding to binary predicates R_α (transition relations for atomic programs α) and a set of atomic propositions p corresponding to unary predicates P interpreted as the set of states satisfying p . The framework of basic modal logic ML provides modalities for the atomic programs α for assertions about the possible results of single-step state transformations. Various additional constructors have been proposed for the formalisation of dynamic, non-local properties, involving for instance unbounded iterations of transitions. We illustrate the examples of PDL, CTL* and L_μ . For one simple concrete example of a dynamic, non-local property, we consider the following (at a state):

- (χ) in any possible future state of the system, there will be
a reachable state in that state's future where p holds.

Propositional dynamic logic

Propositional dynamic logic PDL [29] is based on a dual perspective involving both states and transitions as primary objects of its semantics. Correspondingly, PDL distinguishes two kinds of formulae, *state formulae* and *program formulae*. State formulae, like the familiar modal formulae are evaluated at the states of a transition system and thus define unary predicates on the universe; program formulae on the other hand are evaluated on pairs of states and define binary predicates on the state space, i.e., derived transition relations. Here we work with the following definition; for more on PDL see Chapter 12 of this handbook. We use $\varphi, \psi, \dots$ for state formulae, $\eta, \zeta, \dots$ for program formulae.

DEFINITION 69. State and program formulae of PDL are generated by mutual induction.

State formulae: the Boolean closure of atomic propositions p , and modal quantification of the form $\langle \eta \rangle \varphi$ and $[\eta] \varphi$ for program formulae η and state formulae φ .

Program formulae: the closure of the atomic program formulae α and of all formulae φ ? (“test” operator on state formulae φ) under union ($\eta \cup \zeta$), composition ($\eta; \zeta$) and star or iteration, (η^*).

The semantics of state formulae is the natural one based on the semantics of the corresponding program formulae that define modalities η in terms of new transition relations R_η . For those, the specific constructors are defined in relational terms: atomic program formulae α refer to the given transition relations R_α ; the union operator is set union: $R_{\eta \cup \zeta} = R_\eta \cup R_\zeta$; composition is relational composition: $R_{\eta; \zeta} = R_\eta \circ R_\zeta = \{(u, w) \mid (u, v) \in R_\eta, (v, w) \in R_\zeta \text{ for some } v\}$; the star operation corresponds to the reflexive transitive closure: $R_{\eta^*} = \bigcup_{n \geq 0} (R_\eta)^n$; finally, the test operator defines a loop relation according to $R_{\varphi?} = \{(u, u) \mid \mathfrak{M}, u \models \varphi\}$.

The PDL state formula $\langle \eta^* \rangle \varphi$, for instance, expresses reachability on an η -path of a state that satisfies φ . Note that this is not expressible in FO, even for atomic η and φ . (χ) of the example above is expressible in PDL using $\eta := \bigcup_{\alpha \in \tau} \alpha$, as $\chi = [\eta^*] \langle \eta^* \rangle p$.

We turn to bisimulation invariance. While the standard notion refers to state formulae, the constructors for PDL program formulae also respect bisimulation equivalence, in the sense of bisimulation safety (see section 3.10).

LEMMA 70. *For Kripke structures $\mathfrak{M}$, let $\mathfrak{M}^*$ denote the expansion with all the accessibility relations defined by PDL program formulae. Then any bisimulation $\rho: \mathfrak{M} \rightleftharpoons \mathfrak{M}'$ is also a bisimulation between these expansions, $\rho: \mathfrak{M}^* \rightleftharpoons \mathfrak{M}'^*$.*

Bisimulation invariance for state formulae is then straightforward. In fact it falls out of the inductive proof of the claim of the lemma, which is best understood in terms of the underlying games. Consider the operations of union, composition and star on accessibility operations. For moves along $R_{\eta \cup \zeta} = R_\eta \cup R_\zeta$, the responses of $\mathbf{II}$ merely need no longer respect η/ζ individually; moves along $R_{\eta;\zeta}$ can be responded to as if they came as individual moves in two consecutive rounds; similarly, a move along an R_{η^*} -edge corresponds to a finite sequence of moves along R_η -edges, which is similarly covered by $\mathbf{II}$'s strategy. If, for some state formula φ , $(u, u') \in \rho$ implies that $\mathfrak{M}, u \models \varphi$ iff $\mathfrak{M}', u' \models \varphi$, then it follows that play according to ρ guarantees that (stationary) $R_{\varphi^?}$ -moves are available in $\mathfrak{M}$ iff they are available in $\mathfrak{M}'$.

COROLLARY 71. *Any state formula of PDL is invariant under bisimulation.*

Computation tree logic

For computation tree logic CTL^* , the emphasis is on branching time temporal behaviour rather than process algebra. It is customary to study CTL^* over transition systems with a single binary transition relation R (corresponding to a single unary modality $\Diamond$) which moreover is required to have no terminal nodes, i.e., we assume $\mathfrak{M} \models \Diamond \top$.

The intuitive idea in CTL^* is to associate the runs from a state u of a transitions system $\mathfrak{M}$ with the tree structure $\vec{\mathfrak{M}}[u]$ (the unfolding or tree unravelling, as defined in section 2.2). The infinite branches of the tree $\vec{\mathfrak{M}}[u]$ are the *computation paths* of $\mathfrak{M}$ at u . Besides state formulae, which define properties of states as usual, CTL^* has *path formulae* that define properties of such computation paths. Here a *path* at u is an infinite R -path rooted at u in the usual graph theoretic sense; we write $\sigma = u_0, u_1, \dots$ for a path at $u = u_0$.

DEFINITION 72. State and path formulae of CTL^* are generated by mutual induction. *State formulae:* Boolean closure of atomic propositions p and formulae $E\gamma$ and $A\gamma$ for path formulae γ (existential and universal path quantification).

Path formulae: Boolean closure of all state formulae φ and formulae $\text{Next } \gamma$ (temporal “next” operator) and $\gamma \text{ Until } \delta$ (temporal until operator) for path formulae γ, δ .

The semantics of atomic propositions (as state formulae) and of the Boolean connectives is the natural one. We just highlight the specific constructors for state and path formulae. The semantics of a state formula φ is given in terms of a state $u \in \mathfrak{M}$, the semantics of path formulae γ, δ in terms of a path $\sigma = u_0, u_1, \dots$ in $\mathfrak{M}$, whose suffixes we denote as in $\sigma^j = u_j, u_{j+1}, \dots$:

$\mathfrak{M}, u \models E\gamma$ iff there is a path σ at u such that $\mathfrak{M}, \sigma \models \gamma$, similarly for the dual A .

$\mathfrak{M}, \sigma \models \varphi$ iff $\mathfrak{M}, u_0 \models \varphi$.

$\mathfrak{M}, \sigma \models \text{Next } \gamma$ iff $\mathfrak{M}, \sigma^1 \models \gamma$.

$\mathfrak{M}, \sigma \models \gamma \text{ Until } \delta$ iff for some $j \geq 0$: $\mathfrak{M}, \sigma^j \models \delta$ and for $0 \leq i < j$, $\mathfrak{M}, \sigma^i \models \gamma$.

Reachability of a state satisfying φ , for instance, becomes expressible as $E(\top \text{ Until } \varphi)$. The formula $\top \text{ Until } \varphi$ is also abbreviated $F\varphi$, “eventually φ ”. Using this abbreviation, our sample property (χ) is expressible as $\chi = \neg EF \neg EF p$.

PROPOSITION 73. *Any state formula of CTL^* is invariant under bisimulation.*

This is a straightforward consequence of the fact that any bisimulation $\rho: \mathfrak{M} \rightleftharpoons \mathfrak{M}'$ preserves paths in the sense that for $(u, u') \in \rho$, every path $\sigma = u_0, u_1, \dots$ at $u_0 = u$ in $\mathfrak{M}$ has a bisimilar companion path $\sigma' = u'_0, u'_1, \dots$ at $u'_0 = u'$ in $\mathfrak{M}'$, which is bisimilar in the sense that $(u_i, u'_i) \in \rho$ for all i .

Interestingly, CTL^* admits a characterisation as the bisimulation invariant fragment of monadic path logic, that fragment of monadic second-order logic (over trees) in which second-order quantifiers range over paths. In the light of Theorem 76 below, this characterisation also clarifies the relationship between CTL^* and the much more expressive modal μ -calculus. The following is due to [101] over arbitrary tree models and to [65] over the binary tree.

THEOREM 74. *State formulae of CTL^* precisely define those state properties that are bisimulation invariant and definable in monadic path logic.*

Modal μ -calculus

The modal μ -calculus L_μ is a particularly natural and powerful extension of basic modal logic, which encompasses both PDL and CTL^* . In many ways it may be regarded as *the* extension of modal logic for the purposes of temporal reasoning about processes and corresponding model checking applications. Its theory is well developed, ranging from more classical model theoretic issues to computational and in particular automata theoretic analysis; see Chapter 12 of this handbook for a thorough treatment. Here, we only very selectively comment on some aspects of L_μ and essentially restrict ourselves to its role as an extension of ML in our bisimulation-oriented perspective on modal model theory.

L_μ is the canonical fixed point extension of basic modal logic. Least (and dually, greatest) fixed points of monotone operators capture natural forms of recursion closely related to inductive (and dually, co-inductive) definitions. In L_μ basic modal logic is augmented by the means to define, as fixed points, the results of recursions based on definable monotone operators.

Consider basic modal logic with free monadic-second order variables $X, Y, \dots$ (treated like monadic predicate letters or variables for propositions). A formula $\psi = \psi(X)$ is *positive* in X if X only appears within the scope of an even number of negations in ψ . Positivity in X ensures that, for each structure $\mathfrak{M}$ that interprets all the remaining variables, the following operation on the power set $\mathcal{P}(W)$ of the universe W of $\mathfrak{M}$ is monotone (in the sense that $X \subseteq X'$ implies $\psi[X] \subseteq \psi[X']$):

$$\begin{aligned} \psi^{\mathfrak{M}}: \mathcal{P}(W) &\longrightarrow \mathcal{P}(W) \\ X &\longmapsto \psi^{\mathfrak{M}}[X] := \{w \in W \mid \mathfrak{M}, X, w \models \psi\}. \end{aligned}$$

This operation therefore has unique $\subseteq$ -minimal and $\subseteq$ -maximal fixed points, the least and greatest fixed points of $\psi(X)$, respectively.

DEFINITION 75. The syntax of L_μ is based on basic modal logic ML with free monadic second-order variables, plus closure under the least and greatest fixed point constructors: if $\psi \in L_\mu$ is positive in X , then $\mu X.\psi$ and $\nu X.\psi$ are also formulae of L_μ (in which X is bound).

The semantics of formulae $\varphi \in L_\mu$ is inductively defined in terms of Kripke structures $\mathfrak{M}$ with interpretations for the free second-order variables; $\mathfrak{M}, u \models \mu X.\psi$ (respectively

$\nu X.\psi$) if u is in the least (respectively greatest) fixed point of the operator associated with ψ over $\mathfrak{M}$.

The least fixed point $\mu X.\psi(X)$ in $\mathfrak{M}$ is also definable as the limit of stages X^α generated by induction over the ordinal α , where $X^0 = \emptyset$, $X^{\alpha+1} = \psi^{\mathfrak{M}}[X^\alpha]$ for successor steps, and $X^\lambda = \bigcup_{\alpha < \lambda} X^\alpha$ for limits λ . By monotonicity, the sequence of the X^α is increasing. Over each $\mathfrak{M}$ it eventually must become constant for cardinality reasons. Then the least fixed point of $\psi^{\mathfrak{M}}$ is $X^\infty = \bigcup_\alpha X^\alpha = X^\gamma$ for the minimal γ such that $X^{\gamma+1} = X^\gamma$. (This γ is the closure ordinal of the fixed point over $\mathfrak{M}$.)

The L_μ formula $\mu X.\psi(X)$ for $\psi(X) = \varphi \vee \Diamond X$, for instance, expresses reachability of a state satisfying φ . The monotone operator $\psi^{\mathfrak{M}}$ maps $X \subseteq W$ to the union of $\varphi^{\mathfrak{M}}$ with $\Diamond(X)$. Stage X^n consists of those states from which a state satisfying φ is reachable on an R -path of length less than n . The least fixed point is reached within ω stages over any $\mathfrak{M}$, with $X^\infty = X^\omega$ being the set of states satisfying $\langle R^* \rangle \varphi$. Similarly, well-foundedness of the converse of R , i.e., non-existence of infinite R -paths from a state, is captured by the least fixed point of the operator defined by the formula $\psi(X) = \Box X$.

Our sample property (χ) is expressible as $\chi = \nu Y.(\Box Y \wedge \mu X.(p \vee \Diamond X))$.

Least and greatest fixed points as provided in L_μ admit straightforward explicit definitions in monadic second-order logic MSO, and L_μ may be regarded as a fragment of MSO via a corresponding translation. The following theorem of Janin and Walukiewicz [77] characterises L_μ as the bisimulation invariant fragment of MSO. This is entirely similar in spirit to Theorem 55 for basic modal logic at the first-order level. Covering a far more expressive setting, its proof is also entirely different and based on a sophisticated use of tree automata that recognise corresponding classes of tree models.

THEOREM 76 (Janin–Walukiewicz). *For any MSO formula $\varphi = \varphi(x)$ the following are equivalent:*

- (i) φ is bisimulation invariant.
- (ii) φ is logically equivalent to a formula of L_μ .

We note that, in a similar modal spirit, fixed point extensions have been explored under variations of the underlying notion of bisimulation. In particular, the so-called full μ -calculus with inverse modalities, as related to two-way bisimulation, is studied in [136]; guarded fixed point logic μGF , [62], is the natural extension of the guarded fragment GF by fixed points. For the latter, an analogue of the above characterisation theorem has also been obtained, with a stronger fragment of second-order logic, guarded second-order logic, in place of MSO, [58].

Infinitary modal logics

We encountered ML_∞ , the extension of basic modal logic ML by conjunctions and disjunctions over arbitrary *sets* of formulae, in section 3.6. Theorem 41 characterises bisimulation equivalence as equivalence in ML_∞ . The restriction to set-size (rather than class-size) disjunctions (or unions) is crucial. Remarkably, L_μ (and CTL^*) cannot be embedded into ML_∞ : the well-foundedness property expressed by $\mu X.\Box X \in L_\mu$, for instance, is not globally definable in ML_∞ (see observation 42). In fact, L_μ (or CTL^*) and ML_∞ are incomparable in expressive power.

On the other hand, the individual stages in the generation of any modal least or greatest fixed point are globally definable in ML_∞ . In the example of $\mu X.\Box X$, the

stages X^α are definable by formulae $\varphi_\alpha \in \text{ML}_\infty$ according to $\varphi_0 = \perp$, $\varphi_{\alpha+1} = \Box\varphi_\alpha$ and $\varphi_\lambda = \bigvee_{\alpha < \lambda} \varphi_\alpha$. The reason that the fixed point X^∞ is not ML_∞ definable is that there is no bound on the closure ordinal of this induction. For many natural (restricted) settings, however, ML_∞ is a maximal bisimulation-invariant logic. For the following compare the remark on characteristic formulae below Theorem 41.

OBSERVATION 77. In restriction to any class of structures that intersects only set-many bisimulation equivalence classes, ML_∞ defines any bisimulation closed state property.

Several extended logics, including PDL as an important fragment of L_μ , also admit direct translations into ML_∞ , though. For PDL this is a consequence of the fact that the closure ordinal of the fixed points needed to capture PDL constructs is uniformly bounded by ω . In fact, PDL therefore embeds into that fragment of ML_∞ in which disjunctions and conjunctions over countable, rather than arbitrary, sets of formulae are admitted, $\text{ML}_{\omega_1} \subset \text{ML}_\infty$. The PDL reachability assertion $\langle \alpha^* \rangle \varphi$, for instance, globally translates into $\bigvee_{n \in \omega} \langle \alpha \rangle^n \varphi$, where $\langle \alpha \rangle^n$ is the n -fold iteration of the diamond operator.

ML_{ω_1} may be studied as a fragment of the corresponding infinitary extension of first-order logic, $L_{\omega_1\omega}$, which itself has a well developed classical model theory [83]. Similar to $L_{\omega_1\omega}$, ML_{ω_1} also admits a complete proof system (including infinitary rules) and even satisfies (Craig and Lyndon type) interpolation theorems. Characterisation, completeness, and preservation theorems for ML_{ω_1} and some of its fragments have been obtained along such lines by Radev [110] and Sturm [119, 120].

5.3 Model theoretic criteria

We briefly discuss three particularly relevant model theoretic properties in the light of some of the variations and extensions mentioned above. These may serve as examples that among others could contribute to a framework for a more comprehensive comparative model theory of modal logics.

Finite model property (FMP). As noted in section 3.4, the basic modal logic itself has the finite model property, as do many of its variations and extensions. The variations of ML discussed in section 5.1 above, by inverse and global modalities, as well as the guarded fragment GF, have the FMP. For the extensions beyond FO the finite model property for L_μ , due to Streett and Emerson [118], implies FMP for all of its sub-logics, like CTL^* and PDL.¹¹ The *full μ -calculus*, L_μ with inverse modalities, on the other hand lacks the FMP [136]. The following counterexample illustrates this. The formula $\nu X. (\langle R \rangle X \wedge \mu Y. [R^{-1}] Y)$ requires an infinite (forward) R -path along which every node is well-founded w.r.t. R (does not admit an infinite backward R -path). This implies that the infinite path cannot fold back onto itself; the formula therefore only admits infinite models.

Tree model property. Recall that a logic has the tree model property if every satisfiable formula is satisfied in a tree model. Basic modal logic has the (finite) tree model property (cf. Lemma 35). In fact any bisimulation invariant logic has the tree model property, based on the existence of bisimilar tree unfoldings (cf. section 2.2). In this sense the tree model property, more than the finite model property, is a hallmark of modal model

¹¹The finite model property of many variations and extensions of modal logic, such as PDL and CTL^* , can be obtained by filtration, see [42]. However, this method does not work for some of the more complex systems such as CTL^* and L_μ , where tableau-like and automata-based methods are applied instead.

theory. Moreover, many important variations, even though no longer invariant under ordinary bisimulations, still retain (variant) tree model properties. This phenomenon carries particularly far in the case of GF (see Theorem 66, which also generalises to any guarded bisimulation invariant logic).

Decidability. Decidability and complexity of the satisfiability problem provides one measure for the comparison of the variations and extensions discussed above. Basic modal logic may be seen to be decidable for a number of distinct reasons, as it were. Firstly, as FO is recursively enumerable for validity, ML is decidable as a fragment of FO that is recursively enumerable for satisfiability due to its finite model property. More specifically, however, the finite (tree) model property for basic modal logic (cf. Lemma 35) may be strengthened by effective bounds on depth and branching degree of the candidate tree models – indeed, a Pspace (or alternating Ptime) procedure for satisfiability can be extracted (cf. Chapter 3 of this handbook). Alternatively, decidability of ML may be attributed to just its tree model property and the fact that its tree models are recognised by tree automata, for which emptiness is decidable (cf. Chapters 3 and 17). In view of the extensions that go beyond FO this second line of reasoning carries much further. Extensions that are ‘modal’ in the sense of being bisimulation invariant share the tree model property. Allowing for the appropriate variations of bisimulation, this approach covers not only L_μ , but even the full μ -calculus [136] or the fixed point extension of the guarded fragment [62], which fail to have the FMP. See [57, 135] in this connection for a discussion of the robustness of decidability of modal logics, with a focus on tree models and the accompanying automata theoretic techniques; also see Chapter 17 of this handbook. A comparison between FO^2 and ML in relation to their extensions by natural constructs (e.g., counting, path quantification, transitive closures, fixed points) has also highlighted the special status of modal logic in regard to decidability of such extensions: even comparatively weak extensions of FO^2 along these lines are highly undecidable [61].

6 FURTHER MODEL-THEORETIC CONSTRUCTIONS

One of the traditional directions of development for model theory of a given logic is to identify a sufficiently rich collection of constructions on models, preserving truth in the logic, so that the fundamental concepts of logical definability and logical equivalence can be characterised in terms of these constructions.

In section 2 we introduced the basic model-theoretic notions of generated substructures, bounded morphisms and disjoint unions of Kripke structures and frames, and established corresponding preservation results. These constructions, however, are not sufficient for a complete description of the modal definability of properties or modal equivalence of structures. In this section we introduce and study two more advanced constructions: *ultrafilter extensions* and *ultraproducts*. The former, stemming from the Jónsson–Tarski representation theorem for Boolean algebras with operators in [78], was introduced in modal logic by Goldblatt [47, 48] and used for model-theoretic characterisations of modal definability in [51, 126, 28]. See also section 8. The latter comes from first-order logic, as the most characteristic construction preserving first-order validity (see [12]). Since modal logic on Kripke structures is a fragment of first-order logic, it is a natural truth-preserving construction here, too, and features in the model-theoretic characterisations of modal definability in Kripke structures in section 6.4. Later in this

section we indicate how ultrafilter extensions and ultraproducts are linked with each other, and how they relate modal equivalence between Kripke structures with bisimulations, through the notion of saturation.

6.1 Ultrafilter extensions

Let $\mathfrak{F} = \langle W, \{R_\alpha\}_{\alpha \in \tau} \rangle$ be a τ -frame and let $\mathbf{U}(W)$ be the set of all ultrafilters over W . For every $w \in W$, $\mathbf{u}[w] = \{X \subseteq W \mid w \in X\}$ is the principal ultrafilter generated by w . Further, for every $X \subseteq W$ we define $\mathbf{u}(X) := \{\mathbf{u} \in \mathbf{U}(W) \mid X \in \mathbf{u}\}$.

For each $\alpha \in \tau$ we define a binary relation $R_\alpha^{\mathbf{u}\epsilon}$ on $\mathbf{U}(W)$ as follows. For $\mathbf{u}, \mathbf{w} \in \mathbf{U}(W)$:

$$\mathbf{u} R_\alpha^{\mathbf{u}\epsilon} \mathbf{w} \text{ iff } \langle R_\alpha \rangle(X) \in \mathbf{u} \text{ for every } X \in \mathbf{w}.$$

In particular, note that for every $\alpha \in \tau$, and $x, y \in W$, $x R_\alpha y$ iff $\mathbf{u}[x] R_\alpha^{\mathbf{u}\epsilon} \mathbf{u}[y]$.

DEFINITION 78. Given a τ -frame $\mathfrak{F} = \langle W, \{R_\alpha\}_{\alpha \in \tau} \rangle$:

1. The *ultrafilter extension* of $\mathfrak{F}$ is the τ -frame $\mathbf{ue}(\mathfrak{F}) := \langle \mathbf{U}(W), \{R_\alpha^{\mathbf{u}\epsilon}\}_{\alpha \in \tau} \rangle$.
2. For every Kripke τ -structure $\mathfrak{M} = \langle \mathfrak{F}, V \rangle$, the *ultrafilter extension* of $\mathfrak{M}$ is the Kripke τ -structure $\mathbf{ue}(\mathfrak{M}) := \langle \mathbf{ue}(\mathfrak{F}), V^{\mathbf{u}\epsilon} \rangle$ where $V^{\mathbf{u}\epsilon}(p) = \mathbf{u}(V(p))$ for each $p \in \Phi$.

Thus, the subframe of $\mathbf{ue}(\mathfrak{F})$ consisting of the principal ultrafilters on $\mathfrak{F}$ is isomorphic to $\mathfrak{F}$ but in general, it is *not* a generated subframe of $\mathbf{ue}(\mathfrak{F})$ (see [5, Example 2.58]). However, every finite frame is isomorphic to its ultrafilter extension. For a proof, see e.g. [5, Proposition 2.59].

Here are two concrete examples of ultrafilter extensions from [129]; also compare [129] for a detailed study of ultrafilter extensions and their use in characterising modal definability in some special classes of frames.

- $\mathbf{ue}(\langle \mathbb{Z}, < \rangle)$, where $\langle \mathbb{Z}, < \rangle$ is the linearly ordered set of integers, comprises an isomorphic copy of $\langle \mathbb{Z}, < \rangle$ represented by the principal ultrafilters, and two infinite clusters of free ultrafilters, one consisting of elements less than all ‘standard’ integers, and the other of elements greater than all ‘standard’ integers. All ultrafilters in each cluster are $<^{\mathbf{u}\epsilon}$ -related.
- $\mathbf{ue}(\langle \mathbb{Q}, < \rangle)$, where $\langle \mathbb{Q}, < \rangle$ is the linearly ordered set of rationals, looks similar. It consists of a copy of the rationals, with infinite clusters on each end, but, since every real number can be approximated from either side by a sequence of rationals, it also has for every *real* number a pair of ‘infinitesimally’ close clusters, one on either side.

LEMMA 79. For every Kripke τ -structure $\mathfrak{M} = \langle \mathfrak{F}, V \rangle$ and any formula φ of $\text{ML}(\tau)$: $V^{\mathbf{u}\epsilon}(\varphi) = \mathbf{u}(V(\varphi))$, i.e., $\mathbf{ue}(\mathfrak{M}), \mathbf{u} \models \varphi$ iff $V(\varphi) \in \mathbf{u}$.

This lemma shows that the notion of ultrafilter extension is *canonical*: a state, being an ultrafilter, contains precisely the valuations of those formulae which are true at that state.

COROLLARY 80. For every Kripke τ -structure $\mathfrak{M} = \langle \mathfrak{F}, V \rangle$, $w \in \text{dom}(\mathfrak{F})$, and any formula φ of $\text{ML}(\tau)$:

1. $\mathfrak{M}, w \models \varphi$ iff $\mathbf{ue}(\mathfrak{M}), u[w] \models \varphi$.
2. If $\mathbf{ue}(\mathfrak{M}) \models \varphi$ then $\mathfrak{M} \models \varphi$.
3. If $\mathbf{ue}(\mathfrak{F}), u[w] \models \varphi$ then $\mathfrak{F}, w \models \varphi$.
4. If $\mathbf{ue}(\mathfrak{F}) \models \varphi$ then $\mathfrak{F} \models \varphi$.

We say that a class of τ -frames $\mathcal{C}$ *reflects ultrafilter extensions* if a τ -frame $\mathfrak{F}$ belongs to $\mathcal{C}$ whenever $\mathbf{ue}(\mathfrak{F}) \in \mathcal{C}$. Thus, $\text{FR}(\Gamma)$ reflects ultrafilter extensions for every set of modal formulae Γ .

That the converses of the latter 3 claims above do not hold can be seen from the following example. The modal formulae preserved in ultrafilter extensions will be characterised in Proposition 114.

EXAMPLE 81. By Proposition 114, the Gödel–Löb formula: $\Box(\Box p \rightarrow p) \rightarrow \Box p$ is not preserved in ultrafilter extensions because it is not canonical (see [75]).

Non-reflection of ultrafilter extensions can be used to prove modal non-definability in frames in cases where the other truth preserving constructions introduced earlier may not work. Going back to the example at the end of section 2.3: the sentence $\delta = \forall x \exists y (xRy \wedge yRy)$ is not captured by frame validity of any ML formula, despite being preserved under generated subframes, surjective bounded morphisms and disjoint unions, because it does not reflect ultrafilter extensions. Indeed, $\langle \mathbb{N}, < \rangle \not\models_{\text{FO}} \delta$ while $\mathbf{ue}(\langle \mathbb{N}, < \rangle) \models_{\text{FO}} \delta$ because every free ultrafilter is a maximal element with respect to the quasi-order $<^{\text{ue}}$ (see [128] or [5, Example 2.58] for details).

6.2 Ultraproducts

The constructions of direct products and ultraproducts of first-order structures can be applied to frames, considered as $\text{FO}(\tau)$ -structures, and to Kripke structures, considered as $\text{FO}(\tau_\Phi)$ -structures.

DEFINITION 82. Let $\{W^i\}_{i \in I}$ be a family of sets indexed by a set I .

1. The *direct product* of $\{W^i\}_{i \in I}$ is the set $\prod_{i \in I} W^i = \{g : I \rightarrow \bigcup_{i \in I} W^i \mid g(i) \in W^i \text{ for all } i \in I\}$.
2. For any ultrafilter $\mathbf{U}$ on I , the *ultraproduct* of $\{W^i\}_{i \in I}$ over $\mathbf{U}$, $\prod_{i \in I}^{\mathbf{U}} W^i$, is the quotient of $\prod_{i \in I} W^i$ w.r.t. the equivalence relation $\sim_{\mathbf{U}}$ defined by $g \sim_{\mathbf{U}} g'$ iff $\{i \in I \mid g(i) = g'(i)\} \in \mathbf{U}$. We write $g^{\mathbf{U}}$ for the $\sim_{\mathbf{U}}$ equivalence class of g .
3. For any family $\{X^i \subseteq W^i\}_{i \in I}$, $\prod_{i \in I}^{\mathbf{U}} X^i = \{g^{\mathbf{U}} \in \prod_{i \in I}^{\mathbf{U}} W^i \mid \{i \in I \mid g(i) \in X^i\} \in \mathbf{U}\}$.

DEFINITION 83. Let $\{\mathfrak{F}^i = \langle W^i, \{R_\alpha^i\}_{\alpha \in \tau} \rangle\}_{i \in I}$ be a family of τ -frames indexed by a set I , and $\{\mathfrak{M}^i = \langle \mathfrak{F}^i, V^i \rangle\}_{i \in I}$ be a family of Kripke τ -structures over these frames.

1. The *direct product* of $\{\mathfrak{F}^i\}_{i \in I}$ is the τ -frame $\prod_{i \in I} \mathfrak{F}^i := \langle \prod_{i \in I} W^i, \{R_\alpha\}_{\alpha \in \tau} \rangle$, where for $\alpha \in \tau$: $g_0 R_\alpha g_1$ iff $g_0(i) R_\alpha^i g_1(i)$ for every $i \in I$.

2. The *direct product* of $\{\mathfrak{M}^i\}_{i \in I}$ is the Kripke τ -structure $\prod_{i \in I} \mathfrak{M}^i := \langle \prod_{i \in I} \mathfrak{F}^i, V \rangle$ where $V(p) := \prod_{i \in I} V^i(p)$ for each $p \in \Phi$.

If, further, $\mathbf{U}$ is an ultrafilter on I :

3. The *ultraproduct* of $\{\mathfrak{F}^i\}_{i \in I}$ over $\mathbf{U}$ is the τ -frame $\prod_{i \in I}^{\mathbf{U}} \mathfrak{F}^i := \langle \prod_{i \in I}^{\mathbf{U}} W^i, \{R_\alpha^{\mathbf{U}}\}_{\alpha \in \tau} \rangle$, where for $\alpha \in \tau$: $g_0^{\mathbf{U}} R_\alpha^{\mathbf{U}} g_1^{\mathbf{U}}$ iff $\{i \in I \mid g_0(i) R_\alpha^i g_1(i)\} \in \mathbf{U}$.
4. The *ultraproduct* of $\{\mathfrak{M}^i\}_{i \in I}$ over $\mathbf{U}$ is the Kripke τ -structure $\prod_{i \in I}^{\mathbf{U}} \mathfrak{M}^i := \langle \prod_{i \in I}^{\mathbf{U}} \mathfrak{F}^i, V^{\mathbf{U}} \rangle$ such that for each $p \in \Phi$, $V^{\mathbf{U}}(p) := \prod_{i \in I}^{\mathbf{U}} V^i(p)$.

If $\mathfrak{F}^i = \mathfrak{F}$ for every $i \in I$, the ultraproduct is called an *ultrapower* of $\mathfrak{F}$, denoted $\prod_I^{\mathbf{U}} \mathfrak{F}$; similarly for Kripke structures, where the ultrapower is denoted $\prod_I^{\mathbf{U}} \mathfrak{M}$.

By the fundamental theorem of Łoś (see, e.g., [12, 68]), every first-order definable property holds in an ultraproduct iff it holds in a ‘large’ (i.e., in the ultrafilter) set of component structures. Moreover, every Σ_1^1 -definable property is preserved by ultraproducts [12, Corollary 4.1.14]. Therefore, validity of modal formulae in (pointed) frames, being a Π_1^1 -definable property in terms of the standard translation, is reflected (i.e., its negation is preserved) by ultraproducts. Using these, we obtain the following preservation results.

PROPOSITION 84. *For every family of Kripke τ -structures $\{\mathfrak{M}^i = \langle \mathfrak{F}^i, V^i \rangle\}_{i \in I}$, ultrafilter $\mathbf{U}$ on I , $g^{\mathbf{U}} \in \prod_{i \in I}^{\mathbf{U}} \mathfrak{F}^i$, and formula φ of $\text{ML}(\tau)$:*

1. $\prod_{i \in I}^{\mathbf{U}} \mathfrak{M}^i, g^{\mathbf{U}} \models \varphi$ iff $\{j \in I \mid \mathfrak{M}^j, g(j) \models \varphi\} \in \mathbf{U}$.
2. $\prod_{i \in I}^{\mathbf{U}} \mathfrak{M}^i \models \varphi$ iff $\{j \in I \mid \mathfrak{M}^j \models \varphi\} \in \mathbf{U}$.
3. If $\prod_{i \in I}^{\mathbf{U}} \mathfrak{F}^i, g^{\mathbf{U}} \models \varphi$ then $\{j \in I \mid \mathfrak{F}^j, g(j) \models \varphi\} \in \mathbf{U}$.
4. If $\prod_{i \in I}^{\mathbf{U}} \mathfrak{F}^i \models \varphi$ then $\{j \in I \mid \mathfrak{F}^j \models \varphi\} \in \mathbf{U}$.

Since, however, not every valuation in an ultraproduct of frames can be obtained as an ultraproduct of valuations in the components, the converse of the latter two claims above does not hold.

The following observation due to Goldblatt [48, 43] blends first-order and modal constructions.

PROPOSITION 85. *For any family $\{\mathfrak{F}^i\}_{i \in I}$ of τ -frames and any ultrafilter $\mathbf{U}$ on I , $\prod_{i \in I}^{\mathbf{U}} \mathfrak{F}^i$ is embeddable as a generated subframe into $\prod^{\mathbf{U}} (\uplus_{i \in I} \mathfrak{F}^i)$.*

The embedding is defined canonically as $g^{\mathbf{U}} \mapsto g^{\mathbf{U},+}$, where $g^{\mathbf{U},+} := (w(i), i)$ for each $i \in I$. Furthermore, as shown in [129], any ultraproduct of frames $\prod_{i \in I}^{\mathbf{U}} \mathfrak{F}^i$ is embeddable as a subframe of $\text{uc}(\uplus_{i \in I} \mathfrak{F}^i)$.

6.3 Modal saturation and bisimulations

A class of (pointed) Kripke structures $\mathcal{C}$ is said to have the *Hennessey–Milner property* if modal equivalence between structures in $\mathcal{C}$ implies (and hence is equivalent to) bisimulation equivalence. For instance, as noted in Theorem 38 the class of all finite structures has the Hennessey–Milner property. Compare Definition 39 for first-order types and ω -saturation. The following weaker notion of saturation is more specific to modal logic.

DEFINITION 86. A Kripke τ -structure $\mathfrak{M} = \langle W, \{R_\alpha\}_{\alpha \in \tau}, V \rangle$ is *modally saturated at a state* $w \in W$ if for every $\alpha \in \tau$ and set of modal formulae Γ , the following saturation condition holds:

if $\mathfrak{M}, w \models \langle \alpha \rangle \bigwedge \Gamma_0$ for all finite $\Gamma_0 \subseteq \Gamma$, then there is some $u \in W$ such that $wR_\alpha u$ and $\mathfrak{M}, u \models \Gamma$.

$\mathfrak{M}$ is *modally saturated* if it is modally saturated at each of its states.

It is clear from Definition 39 that ω -saturated Kripke structures are modally saturated.

PROPOSITION 87. *The class of modally saturated Kripke structures has the Hennessey–Milner property.*

Proof. If $\mathfrak{M}$ and $\mathfrak{M}'$ are modally saturated, then $\rho := \{(w, w') \in W \times W' \mid (\mathfrak{M}, w) \equiv_{\text{ML}} (\mathfrak{M}', w')\}$ is a bisimulation between $\mathfrak{M}$ and $\mathfrak{M}'$. Atom equivalence is obvious. Consider for instance the *forth* condition. Let $(\mathfrak{M}, w) \equiv_{\text{ML}} (\mathfrak{M}', w')$ and let $(w, u) \in R_\alpha$. Put $\Gamma := \text{Th}_{\text{ML}}(\mathfrak{M}, u)$. For finite $\Gamma_0 \subseteq \Gamma$, $\mathfrak{M}, w \models \langle \alpha \rangle \bigwedge \Gamma_0$ and hence also $\mathfrak{M}', w' \models \langle \alpha \rangle \bigwedge \Gamma_0$. By modal saturation of $\mathfrak{M}'$ at w' therefore, there is some u' such that $(w', u') \in R_\alpha$ and $\mathfrak{M}', u' \models \Gamma$. But this means that $(\mathfrak{M}, u) \equiv_{\text{ML}} (\mathfrak{M}', u')$, and u' is as desired for the forth requirement. $\blacksquare$

COROLLARY 88. *The class of ω -saturated Kripke structures has the Hennessey–Milner property.*

It is well-known from classical model theory [12, Corollary 4.3.14] that the ultrapower of any (pointed) Kripke structure w.r.t. a regular ultrafilter is an ω -saturated elementary extension of that structure. Furthermore, two (pointed) Kripke structures are modally equivalent iff any pair of their ω -saturated ultrapowers are modally equivalent, and hence, by Corollary 88, bisimilar. Thus, we obtain the following characterisation of modal equivalence between Kripke structures from [20], as a corollary of the above.

THEOREM 89. *Two (pointed) Kripke structures are modally equivalent iff any pair of their ω -saturated ultrapowers are bisimilar.*

A parallel with first-order logic can be drawn here if we think of bisimulations as the modal analogue of partial isomorphisms between Kripke structures, and note that elementary equivalence on ω -saturated structures coincides with partial isomorphism between them (see [108, 68, 23]). Then Theorem 91 below completes the match. Before getting there, we need the following result, due to van Benthem [126], building on a construction of Fine [28].

THEOREM 90. *For every Kripke τ -structure $\mathfrak{M}$, $\text{uc}(\mathfrak{M})$ is a bounded morphic image of an ω -saturated ultrapower of $\mathfrak{M}$.*

Proof. Let $\mathfrak{M} = \langle \mathfrak{F}, V \rangle$ where $\mathfrak{F} = \langle W, \{R_\alpha\}_{\alpha \in \tau} \rangle$.

The structure $\mathfrak{F}^\times = \langle W, \{R_\alpha\}_{\alpha \in \tau}, \{X \mid X \subseteq W\} \rangle$ has in particular every $V(\varphi)$ as a

distinguished predicate. Take an ω -saturated ultrapower $\mathfrak{F}^* = \prod_I^U \mathfrak{F}^\times$ and for each $f^U \in \prod_I^U W$ define $v(f^U) = \{X \subseteq W \mid f^U \in \prod_I^U X\}$. It is immediate to check that $v(f^U) \in \mathbf{U}(W)$. Considering v as a mapping from $\prod_I^U \mathfrak{M}$ onto $\mathbf{ue}(\mathfrak{M})$ one can show that it is a bounded morphism. The most difficult step (proved in [126] for the case of one unary modality, see also the proof of [5, Proposition 2.61]) is to prove the *back* condition, which uses the saturation of $\mathfrak{F}^\times$. ■

Using this theorem we can now obtain a strengthening of the model-theoretic characterisation of modal equivalence, first proved by Hollenberg [72]. See also [138] and [5, Theorem 2.62].

THEOREM 91. *For any pointed Kripke structures $(\mathfrak{M}, w)$ and $(\mathfrak{M}', w')$,*

$$(\mathfrak{M}, w) \equiv_{\text{ML}} (\mathfrak{M}', w') \text{ iff } (\mathbf{ue}(\mathfrak{M}), u[w]) \rightleftharpoons (\mathbf{ue}(\mathfrak{M}'), u[w']).$$

Proof. The direction from right to left is immediate from Lemma 79 and bisimulation invariance, Theorem 14. For the converse direction, suppose $(\mathfrak{M}, w) \equiv_{\text{ML}} (\mathfrak{M}', w')$. Then, by Theorem 89, $(\prod_I^U \mathfrak{M}, g_w^U) \rightleftharpoons (\prod_I^U \mathfrak{M}', g_{w'}^U)$ for the ω -saturated ultrapowers defined in the proof above, where $g_w(i) = w$ for each $i \in I$, and likewise for $g_{w'}$. Note that $v(g_w^U) = u[w]$ and $v(g_{w'}^U) = u[w']$. Composing this bisimulation with the surjective bounded morphisms $v : (\prod_I^U \mathfrak{M}, g_w^U) \twoheadrightarrow (\mathbf{ue}(\mathfrak{M}), u[w])$ and $v' : (\prod_I^U \mathfrak{M}', g_{w'}^U) \twoheadrightarrow (\mathbf{ue}(\mathfrak{M}'), u[w'])$, we obtain a bisimulation between the ultrafilter extensions. ■

The following observation is immediate from the definitions.

LEMMA 92. *Bisimulations preserve modal saturation at a state: if $(\mathfrak{M}, w) \rightleftharpoons (\mathfrak{M}', w')$, then $\mathfrak{M}$ is modally saturated at w iff $\mathfrak{M}'$ is modally saturated at w' . Consequently, global bisimulations preserve modal saturation of models.*

From this lemma and Theorem 90, since surjective bounded morphisms are global bisimulations, we obtain the following result from [44], (see also [5, Proposition 2.61])

COROLLARY 93. *The ultrafilter extension of every Kripke structure is modally saturated.*

As Venema argues quite aptly in [138], this result along with Theorem 91 indicates that, for modal logics, ultrafilter extensions can play the role that ultrapowers play in first-order logic for the construction of saturated extensions of structures.

6.4 Modal definability of properties of Kripke structures

Kripke structures serve to give model theoretic semantics to modal logic. Conversely, focusing on Kripke structures in their own right, we regard modal logic as a language for defining classes of Kripke structures. We may ask the natural model-theoretic questions from this angle, like, for instance: what classes/properties of (pointed) Kripke structures are definable by (sets of) modal formulae? A definitive answer to that question was given in the case of elementary properties of pointed Kripke structures defined by single modal formulae, by Theorem 55. Here we address the general question by using classical model-theoretic tools and the constructions introduced earlier in this section.

Since modal formulae express first-order conditions on (pointed) Kripke structures, these are special cases of *first-order definable* (by a single first-order sentence), respectively *elementary* (definable by any set of first-order sentences) classes and properties. Keisler's theorem [12, Theorem 4.1.12] characterising elementary and first-order definable classes is therefore relevant here: a class of first-order structures is elementary iff it is closed under elementary equivalence and ultraproducts; it is first-order definable iff both the class and its complement are elementary. Since modal formulae cover only a fragment of the first-order language $\text{FO}(\tau_\Phi)$, these results give necessary but not sufficient conditions for modal definability of classes of (pointed) Kripke structures. But 'elementary equivalence' for modal logic is modal equivalence. Would that adjustment of Keisler's theorem suffice to guarantee modal definability? The answer is 'yes' in both cases. The following is from [22].

THEOREM 94. *A class $\mathcal{K}$ of (pointed) Kripke structures is definable by a set of modal formulae iff it is closed under modal equivalence and ultraproducts; $\mathcal{K}$ is definable by a single modal formula iff both $\mathcal{K}$ and its complement are definable by a set of modal formulae.*

Proof. These can be proved by adapting the proof of Keisler's theorem. Alternatively, we may invoke a corollary of the Keisler–Shelah theorem (cf. Corollary 6.1.16 and Theorem 6.1.15 in [12]) which states that a class of first-order structures is elementary iff it is closed under isomorphism and ultraproducts while its complement is closed under ultrapowers. The latter condition here follows from closure under modal equivalence. Once $\mathcal{K}$ has been shown to be elementary, a general argument can be applied that works not only for modal formulae but for any other natural fragment Δ of first-order logic (see [12, Lemma 3.2.1]): if $\Delta \subseteq \text{FO}$ is closed under negation and disjunction, then an elementary class is axiomatisable with formulae from Δ iff it is closed under Δ -equivalence.

For definability by a single formula, one may use compactness for ML just as for FO to show that whenever both the given class and its complement are definable by a set of formulae, then the class (and its complement) are definable by a single formula. Alternatively, one may first establish first-order definability of $\mathcal{K}$, and then use Theorem 55 and bisimulation invariance to see that the defining formula must be equivalent to a modal formula. ■

Note that, as an immediate consequence of (the classical proof of) Theorem 55, an elementary class of (pointed) Kripke structures is closed under modal equivalence iff it is closed under bisimulations. Therefore, we can strengthen somewhat the results above, by replacing closure under modal equivalence by bisimulation closure, but at the expense of demanding closure of the complement under ultrapowers. See [22] and [5, Theorems 2.75, 2.76] for the following.

THEOREM 95. *For any class $\mathcal{K}$ of pointed Kripke structures:*

1. *$\mathcal{K}$ is definable by a set of modal formulae iff it is closed under bisimulation and ultraproducts, while its complement is closed under ultrapowers.*
2. *$\mathcal{K}$ is definable by a single modal formula iff it is closed under bisimulation, while both it and its complement are closed under ultrapowers.*

Proof. For the non-trivial part of (i): assuming the closure conditions for $\mathcal{K}$ and its complement, we consider the modal theory $\text{Th}_{\text{ML}}(\mathcal{K})$ and show that it defines $\mathcal{K}$, i.e., every model of it is in $\mathcal{K}$. For details see [22], [5, Theorem 2.75]. Alternatively, we can take a shortcut: by Theorem 89 the closure conditions imply that $\mathcal{K}$ is closed under modal equivalence, and hence Theorem 94 applies.

For the non-trivial part of (ii) we may use (i) and a standard compactness argument as in the proof of Keisler’s theorem (see [22] and [5, Theorems 2.76]), or use Theorem 94 again. ■

Similar results can be obtained for classes of Kripke structures; we leave these to the reader.

Finally, we mention the following results of Venema [138] which characterise modal definability of classes of (pointed) Kripke structures in purely modal terms, i.e., without involving the typical constructions from classical logic. In what follows, a bisimulation $\rho: \mathfrak{M} \rightleftharpoons \mathfrak{M}'$ is *surjective* if every state in $\mathfrak{M}'$ has a bisimilar one in $\mathfrak{M}$; an *ultrafilter union* of a family of pointed Kripke structures $\{\mathfrak{M}^i, w^i\}_{i \in I}$ is a pointed Kripke structure $(\text{ue}(\biguplus_{i \in I} \mathfrak{M}^i), w)$, where w is an ultrafilter containing every co-finite subset of $\{w_i \mid i \in I\}$.

THEOREM 96. *A class of Kripke structures is modally definable iff it is closed under disjoint unions, surjective bisimulations, and ultrafilter extensions, while it reflects ultrafilter extensions.*

A class of pointed Kripke structures is modally definable iff it is closed under bisimulations and ultrafilter unions, and reflects ultrafilter extensions.

To summarise: model theory of modal logic over Kripke structures essentially derives from first-order model theory, with the crucial extra feature of bisimulation invariance. The additional requirement of bisimulation invariance leads us from classical model theory to modal model theory and allows us to develop the analogy between them further.

7 GENERAL FRAMES

Neither of the two kinds of semantic structures we have considered so far, viz. Kripke frames and Kripke structures, provides a completely satisfactory framework for the semantics for modal logic. On the one hand, truth and validity in Kripke structures, with its crucial dependency on given valuations, does not reflect the richer semantics in terms of validity in frames. On the other hand, validity in frames, being an essentially second-order notion, is in general deductively intractable. As a consequence, frame-incomplete modal logics are the rule, rather than the exception (see Chapter 7 of this handbook). It is therefore necessary to look for a new type of semantic structures, ‘hybrids’ between Kripke structures and frames, combining the expressive richness of the frame-based semantics with the flexibility and good deductive behaviour of the one based on Kripke structures.

Such structures, called *general frames*, were introduced in modal logic by Thomason in [124], with precursors in [97] and [28]. General frames are analogues to Henkin’s ‘general models’ for second-order logic, extending first-order structures with a family of ‘admissible sets’, and restricting the second-order quantification to such sets only. Independently, general frames essentially arose from the seminal study by Jónsson and

Tarski [78] of Boolean algebras with operators (see also Chapter 6 of this handbook), since they appear as the ‘concrete’, set-theoretic counterparts of modal algebras, arising in the Jónsson–Tarski representation theorem, and thus providing the link between the algebraic and relational semantics.

In this section we introduce the modal semantics based on general frames, develop the basic model theory of general frames and briefly mention the *duality theory* which relates them to algebras. We then discuss the relevance and use of general frames to the model theory of the frame-based modal semantics, in terms of *persistence* of modal formulae with respect to various important classes of general frames.

7.1 General frames as semantic structures in modal logic

Note that the operators $\langle R \rangle$ and $[R]$ defined in section 1.2 are monotone. Besides, the operators $\langle R \rangle$ are *normal* (preserving falsum) and *additive* (distributive over disjunctions); see Chapter 6 of this handbook. Hence every structure $\langle \mathcal{P}(W); \cap, -, \emptyset, \{\langle R_\alpha \rangle\}_{\alpha \in \tau} \rangle$ is a (complete and atomic) Boolean algebra with operators in the terms of [78] (see also Chapter 6), called a *modal τ -algebra*.

DEFINITION 97. Given a τ -frame $\mathfrak{F} = \langle W, \{R_\alpha\}_{\alpha \in \tau} \rangle$, a *general τ -frame over $\mathfrak{F}$* is a structure $\langle \mathfrak{F}, \mathbb{W} \rangle$ expanding $\mathfrak{F}$ with a τ -algebra of *admissible subsets* of $\mathcal{P}(W)$, closed under boolean operations and the operators $\{\langle R_\alpha \rangle\}_{\alpha \in \tau}$, i.e., $\mathbb{W}$ is a τ -subalgebra of $\langle \mathcal{P}(W); \cap, -, \emptyset, \{\langle R_\alpha \rangle\}_{\alpha \in \tau} \rangle$.

Given a general τ -frame $\mathfrak{G} = \langle \mathfrak{F}, \mathbb{W} \rangle$ we denote $\mathfrak{F}$ by $\mathfrak{G}_\#$ and the τ -algebra $\mathbb{W}$ by $\mathfrak{G}^+$.

EXAMPLE 98. For every Kripke structure $\mathfrak{M} = \langle \mathfrak{F}, V \rangle$, $\langle \mathfrak{F}, \{V(\varphi) \mid \varphi \in \text{ML}(\tau)\} \rangle$ is a general τ -frame over $\mathfrak{F}$, *generated by $\mathfrak{M}$* . In particular, the general τ -frame $\mathfrak{G}_L$ generated by the canonical Kripke structure $\mathfrak{M}_L$ (see Chapter 7 of this handbook) of a normal modal logic L is called the *canonical general frame of L* .

Among the general frames over $\mathfrak{F} = \langle W, \{R_\alpha\}_{\alpha \in \tau} \rangle$ there is a least one, viz. $\mathfrak{F}_{\min} = \langle \mathfrak{F}, \mathbb{W}_{\min} \rangle$ generated from the Kripke structure $\mathfrak{M}_{\min} = \langle \mathfrak{F}, V_{\min} \rangle$ where $V_{\min}(p) = \emptyset$ for every $p \in \Phi$, and a greatest one, viz. the *full general τ -frame* $\mathfrak{F}_{\max} = \langle \mathfrak{F}, \mathcal{P}(W) \rangle$. Clearly, local (as well as global) validity in $\mathfrak{F}$ and $\mathfrak{F}_{\max}$ coincide. So we can safely identify the τ -frame $\mathfrak{F}$ with $\mathfrak{F}_{\max}$. Furthermore, the family of all general frames over a τ -frame $\mathfrak{F} = \langle W, \{R_\alpha\}_{\alpha \in \tau} \rangle$ forms a complete lattice.

DEFINITION 99. Given a general τ -frame $\mathfrak{G} = \langle \mathfrak{F}, \mathbb{W} \rangle$, a *valuation over $\mathfrak{G}$* is any valuation $V : \Phi \rightarrow \mathbb{W}$. A Kripke structure $\langle \mathfrak{F}, V \rangle$ where V is a valuation over $\mathfrak{G}$ is a *Kripke structure over $\mathfrak{G}$* , also denoted by $\langle \mathfrak{G}, V \rangle$ or $\langle \mathfrak{F}, \mathbb{W}, V \rangle$.

It follows by a routine induction that if $\mathfrak{M} = \langle \mathfrak{F}, \mathbb{W}, V \rangle$, then $V(\varphi) \in \mathbb{W}$ for every $\varphi \in \text{ML}(\tau)$.

DEFINITION 100. Given a formula $\varphi \in \text{ML}(\tau)$, a general τ -frame $\mathfrak{G}$, and $w \in W$, we say that φ is *(locally) valid at w in $\mathfrak{G}$* , denoted $\mathfrak{G}, w \models \varphi$, if φ is true at w in every Kripke structure over $\mathfrak{G}$. φ is *valid in $\mathfrak{G}$* , denoted $\mathfrak{G} \models \varphi$, if φ is valid in $\mathfrak{G}$ at every $w \in W$, i.e., φ is valid in every Kripke structure over $\mathfrak{G}$.

Note that local validity of modal formulae in a general τ -frame is preserved under the rule Modus Ponens and under taking uniform substitutions, while validity is also preserved under Necessitation.

All general frames generated from Kripke structures have an at most countable algebra of admissible sets, so not every general frame is of that type. On the other hand, every general frame *can* be generated from a Kripke structure in an extended language with an appropriately large cardinality of the set of atomic propositions. This observation is sufficient to transfer various results and constructions from Kripke structures to general frames.

However, as semantic structures for modal logic, general frames match most closely modal algebras. Indeed, as already noted, every general τ -frame $\mathfrak{G}$ generates a ‘complex τ -algebra’ $\mathfrak{G}^+$. Conversely, every τ -algebra $\mathfrak{A}$ determines a general frame $\mathfrak{A}_+$ based on the *ultrafilter frame* of that algebra (see section 7.2), and is moreover embedded in $(\mathfrak{A}_+)^+$ in a way extending the Stone representation for Boolean algebras. That embedding is the subject of the celebrated *Jónsson–Tarski representation theorem* (see [78], [5, Section 5.3], or Chapter 6 of this handbook). Furthermore, there exists an algebraic-categorical *duality* between general frames and modal algebras, systematically developed by Goldblatt in [47, 48, 43] and later, from a topological perspective by [114] (see also [5, Section 5.4]), discussed in detail in Chapter 6 of this handbook.

7.2 Constructions and truth preservation results on general frames

Bisimulations and special cases

DEFINITION 101. Let $\mathfrak{G} = \langle \mathfrak{F}, \mathbb{W} \rangle$ and $\mathfrak{G}' = \langle \mathfrak{F}', \mathbb{W}' \rangle$ be two general τ -frames. A bisimulation ρ between $\mathfrak{F}$ and $\mathfrak{F}'$ is a *bisimulation between $\mathfrak{G}$ and $\mathfrak{G}'$* if for every valuation V over $\mathfrak{G}$ there is a valuation V' over $\mathfrak{G}'$ such that $\rho: \langle \mathfrak{G}, V \rangle \rightleftharpoons \langle \mathfrak{G}', V' \rangle$, and vice versa.

A bisimulation between pointed general frames is defined likewise.

Note that not every bisimulation between Kripke frames is a bisimulation between them as full general frames, because not every valuation over one of them must have a matching valuation satisfying *atom equivalence*.

COROLLARY 102. If $\rho: (\mathfrak{G}, w) \rightleftharpoons (\mathfrak{G}', w')$ is a bisimulation between pointed general τ -frames $(\mathfrak{G}, w)$ and $(\mathfrak{G}', w')$ then $(\mathfrak{G}, w) \equiv_{\text{ML}} (\mathfrak{G}', w')$. Likewise, if $\rho: \mathfrak{G} \rightleftharpoons \mathfrak{G}'$ then $\mathfrak{G} \equiv_{\text{ML}} \mathfrak{G}'$.

The definitions of generated subframes, bounded morphisms, and disjoint unions can be extended to general frames.

DEFINITION 103. Given a general τ -frame $\mathfrak{G} = \langle \mathfrak{F}, \mathbb{W} \rangle$, a *generated subframe of $\mathfrak{G}$* is any general τ -frame $\mathfrak{G}' = \langle \mathfrak{F}', \mathbb{W}' \rangle$ where $\mathfrak{F}' \leq \mathfrak{F}$ and $\mathbb{W}' = \{X \cap \text{dom}(\mathfrak{F}') \mid X \in \mathbb{W}\}$.

DEFINITION 104. Let $\mathfrak{G} = \langle \mathfrak{F}, \mathbb{W} \rangle$ and $\mathfrak{G}' = \langle \mathfrak{F}', \mathbb{W}' \rangle$ be two general τ -frames and $\rho: \mathfrak{F} \xrightarrow{\Rightarrow} \mathfrak{F}'$ a bounded morphism. Then ρ is a *bounded morphism* from $\mathfrak{G}$ to $\mathfrak{G}'$ if for every $Y \in \mathbb{W}'$, $\rho^{-1}[Y] \in \mathbb{W}$; ρ is a *bounded strong morphism* from $\mathfrak{G}$ to $\mathfrak{G}'$ if it is a bounded morphism from $\mathfrak{G}$ to $\mathfrak{G}'$ and for every $X \in \mathbb{W}$, $\rho[X] \in \mathbb{W}'$ and $X = \rho^{-1}[\rho[X]]$.

DEFINITION 105. The *disjoint union* of the family $\{\mathfrak{G}^i = \langle \mathfrak{F}^i, \mathbb{W}^i \rangle\}_{i \in I}$ of general τ -frames is $\biguplus_{i \in I} \mathfrak{G}^i = \langle \biguplus_{i \in I} \mathfrak{F}^i, \mathbb{W} \rangle$, where $\mathbb{W} = \{\biguplus_{i \in I} X^i \mid X^i \in \mathbb{W}^i \text{ for each } i \in I\}$.

We leave it to the reader to check that generated subframes and disjoint unions of general frames produce general frames indeed, and to see that they, as well as bounded strong morphisms, are particular cases of general frame bisimulations. The associated preservation results are immediate, and are left to the reader, too. As for bounded

morphisms of general frames, in general they are not general frame bisimulations and only preserve validity in the forward direction.

Ultrafilter extensions and ultraproducts

The construction of ultrafilter extensions of frames can be generalised to the Stone representation of modal algebras (see Chapter 6 of this handbook), which in turn are essentially general frames, thus defining ultrafilter extensions of general frames. More precisely, given a general τ -frame $\mathfrak{G} = \langle W, \{R_\alpha\}_{\alpha \in \tau}, \mathbb{W} \rangle$ over a frame $\mathfrak{F}$, let $\mathbf{U}(\mathbb{W})$ be the set of all ultrafilters over the algebra $\mathfrak{G}^+$. For each $\alpha \in \tau$ we define a binary relation $R_\alpha^{\mathbb{W}}$ on $\mathbf{U}(\mathbb{W})$ just like $R_\alpha^{\mathfrak{u}\mathfrak{e}}$ in $\mathfrak{u}\mathfrak{e}(\mathfrak{F})$, i.e., for any $\mathfrak{u}, \mathfrak{w} \in \mathbf{U}(\mathbb{W})$:

$$\mathfrak{u} R_\alpha^{\mathbb{W}} \mathfrak{w} \text{ iff } \langle R_\alpha \rangle(X) \in \mathfrak{u} \text{ for every } X \in \mathfrak{w}.$$

The frame $(\mathfrak{G}^+)_+ = \langle \mathbf{U}(\mathbb{W}), \{R_\alpha^{\mathbb{W}}\}_{\alpha \in \tau} \rangle$ is called the *ultrafilter frame of the τ -algebra $\mathfrak{G}^+$* .

Finally, we put $\mathbb{W}^{\mathfrak{u}\mathfrak{e}} := \{\mathfrak{u}_{\mathbb{W}}(X) \mid X \in \mathbb{W}\}$ where $\mathfrak{u}_{\mathbb{W}}(X) = \{\mathfrak{u} \in \mathbf{U}(\mathbb{W}) \mid X \in \mathfrak{u}\}$. It is routine to check that $\langle R_\alpha^{\mathbb{W}} \rangle(\mathfrak{u}_{\mathbb{W}}(X)) = \mathfrak{u}_{\mathbb{W}}(\langle R_\alpha \rangle(X))$ and hence $\mathbb{W}^{\mathfrak{u}\mathfrak{e}}$ is a modal τ -algebra.

DEFINITION 106. Given a general τ -frame $\mathfrak{G} = \langle W, \{R_\alpha\}_{\alpha \in \tau}, \mathbb{W} \rangle$, the *ultrafilter extension of $\mathfrak{G}$* is the general τ -frame $\mathfrak{u}\mathfrak{e}(\mathfrak{G}) := \langle \mathbf{U}(\mathbb{W}), \{R_\alpha^{\mathbb{W}}\}_{\alpha \in \tau}, \mathbb{W}^{\mathfrak{u}\mathfrak{e}} \rangle$, also known as the *general ultrafilter frame of the τ -algebra $\mathfrak{G}^+$* .

From the basic properties of ultrafilters, and the closure of $\mathbb{W}^{\mathfrak{u}\mathfrak{e}}$ under $\langle R_\alpha^{\mathbb{W}} \rangle$, it follows that $\mathfrak{G}^+ \cong \mathfrak{u}\mathfrak{e}(\mathfrak{G})^+$ for any general τ -frame $\mathfrak{G}$. Note, however, that $\mathfrak{u}\mathfrak{e}(\mathfrak{G}) \cong \mathfrak{G}$ does not hold in general, and in section 7.3 we will characterise the general frames for which this is the case. Still, since validity of modal formulae in $\mathfrak{G}$ and in $\mathfrak{G}^+$ coincide, we obtain the following.

THEOREM 107. For any general τ -frame $\mathfrak{G}$, $\mathfrak{u}\mathfrak{e}(\mathfrak{G}) \equiv_{\text{ML}} \mathfrak{G}$.

DEFINITION 108. Let $\{\mathfrak{G}^i = \langle \mathfrak{F}^i, \mathbb{W}^i \rangle\}_{i \in I}$ be a family of general τ -frames indexed by a set I . For any ultrafilter $\mathbf{U}$ on I , the *ultraproduct of $\{\mathfrak{G}^i\}_{i \in I}$ over $\mathbf{U}$* is the general τ -frame $\prod_{i \in I}^{\mathbf{U}} \mathfrak{G}^i := \langle \prod_{i \in I}^{\mathbf{U}} \mathfrak{F}^i, \mathbb{W}^{\mathbf{U}} \rangle$, where $\mathbb{W}^{\mathbf{U}} = \{\prod_{i \in I}^{\mathbf{U}} X^i \mid X^i \in \mathbb{W}^i \text{ for each } i \in I\}$.

Note that the ultraproduct of a family of Kripke frames regarded as full general frames is *not* a full general frame itself, so it differs from the ultraproduct of frames, as defined earlier. To distinguish these, we call the former *general ultraproduct* of frames. Unlike the latter, every valuation in it is an ultraproduct of respective valuations in the components, whence the following preservation result (see [47, 48, 43]).

PROPOSITION 109. For every family of general τ -frames $\{\mathfrak{G}^i = \langle \mathfrak{F}^i, \mathbb{W}^i \rangle\}_{i \in I}$, ultrafilter $\mathbf{U}$ on I , element $\mathfrak{w}_{\mathbf{U}} \in \prod_{i \in I}^{\mathbf{U}} \mathfrak{F}^i$, and formula φ of $ML(\tau)$:

1. $\prod_{i \in I}^{\mathbf{U}} \mathfrak{G}^i, \mathfrak{w}_{\mathbf{U}} \models \varphi$ iff $\{j \in I \mid \mathfrak{G}^j, \mathfrak{w}(j) \models \varphi\} \in \mathbf{U}$.
2. $\prod_{i \in I}^{\mathbf{U}} \mathfrak{G}^i \models \varphi$ iff $\{j \in I \mid \mathfrak{G}^j \models \varphi\} \in \mathbf{U}$.

7.3 Special types of general frames and persistence of modal formulae

Let $\mathcal{G}$ be the class of all general τ -frames of a fixed modal type τ , and let $\mathcal{C}$ be any subclass of $\mathcal{G}$.

DEFINITION 110. A formula $\varphi \in \text{ML}(\tau)$ is *locally $\mathcal{C}$ -persistent*, if for every general τ -frame $\mathfrak{G} = \langle \mathfrak{F}, \mathbb{W} \rangle \in \mathcal{C}$, and $w \in \text{dom}(\mathfrak{F})$, $\mathfrak{G}, w \models \varphi$ implies $\mathfrak{F}, w \models \varphi$; φ is *$\mathcal{C}$ -persistent*, if for every general τ -frame $\mathfrak{G} = \langle \mathfrak{F}, \mathbb{W} \rangle \in \mathcal{C}$, $\mathfrak{G} \models \varphi$ implies $\mathfrak{F} \models \varphi$.

Clearly, local persistence implies persistence, but the converse does not always hold. While often the practically important notion is the latter, the former is more natural.

A general frame can be thought of as a frame in which a restriction on the valuations is imposed by allowing only those valuations which assign admissible sets to the propositional variables (and hence, to all formulae). Thus, the idea of persistence is that it enables one to conclude (local) validity, i.e., truth under *every* valuation, of a modal formula in a frame, based on its truth under *some* special valuations, viz. the admissible ones. In other words, a formula is $\mathcal{C}$ -persistent if, whenever it is falsified in a Kripke frame $\mathfrak{F}$, it is falsified by *some* admissible valuation in *each* general frame from $\mathcal{C}$ over $\mathfrak{F}$. Thus, persistence gives a measure of the ‘semantic complexity’ of a formula, in terms of its falsifying valuations. Note that a modal formula is locally $\mathcal{G}$ -persistent iff it is semantically equivalent to a constant formula (i.e., a formula without propositional variables). Indeed, every constant formula is $\mathcal{G}$ -persistent. Conversely, if φ is $\mathcal{G}$ -persistent, then for every pointed frame $(\mathfrak{F}, w)$, $\mathfrak{F}, w \models \varphi$ iff $\langle \mathfrak{F}, V_\perp \rangle, w \models \varphi$, where $V_\perp$ assigns $\emptyset$ to every atomic proposition, iff $\mathfrak{F}, w \models \varphi_\perp$ where $\varphi_\perp$ is obtained from φ by replacing all atomic propositions by $\perp$.

We will introduce some important classes of general frames, persistence with respect to which provides sufficient conditions for good expressive or axiomatic behaviour of the formulae.

DEFINITION 111. Let $\mathfrak{G} = \langle W, \{R_\alpha\}_{\alpha \in \tau}, \mathbb{W} \rangle$ be a general τ -frame and $\alpha \in \tau$. The relation R_α is *tight* in $\mathfrak{G}$ if for every $u, w \in W$: $uR_\alpha w$ iff for all $X \in \mathbb{W}$, $w \in X$ implies $u \in \langle R_\alpha \rangle(X)$; equivalently, iff $u \in \bigcap \{ \langle R_\alpha \rangle(X) \mid X \in \mathbb{W} \text{ and } w \in X \}$.

Recall, for the compactness property below, that a family of sets $\mathcal{F}$ has the *finite intersection property* (FIP) if the intersection of every finite sub-family of $\mathcal{F}$ is non-empty.

DEFINITION 112. A general τ -frame $\langle W, \{R_\alpha\}_{\alpha \in \tau}, \mathbb{W} \rangle$ is:

- *differentiated*, if for every $u, u' \in W$, if $u \neq u'$ then there is $X \in \mathbb{W}$ such that $u \in X$ and $u' \notin X$;
- *tight*, if R_α is tight for every $\alpha \in \tau$;
- *discrete*, if $\{u\} \in \mathbb{W}$ for every $u \in W$;
- *elementary*, if every subset of W that is $\text{FO}(\tau)$ -definable with parameters (in the sense of Definition 39) is admissible;
- *compact*, if every family of admissible sets in $\mathfrak{G}$ with FIP has a non-empty intersection;¹²
- *refined*, if it is differentiated and tight;

¹²This is equivalent to the requirement that every ultrafilter over $\mathfrak{G}^+$ consists of all admissible sets containing a fixed state in $\mathfrak{G}$.

– *descriptive*, if it is refined and compact.

Amongst all discrete general frames over a Kripke frame $\mathfrak{F}$, there is a *least* one, viz. $\mathfrak{D}(\mathfrak{F})$, generated from all singletons by closing under the Boolean and modal operators. It contains all finite and co-finite sets in $\mathfrak{F}$. Likewise, amongst all elementary general frames over a Kripke frame $\mathfrak{F}$, there is a *least* one, viz. $\mathfrak{E}(\mathfrak{F})$, in which the admissible sets are precisely the subsets of the domain of $\mathfrak{F}$ that are parametrically first-order definable in $\text{FO}(\tau)$.

Assuming the type τ is fixed, the class of all differentiated (resp. tight, discrete, elementary, refined, descriptive) general τ -frames will be denoted by $\mathcal{DF}$ (resp. $\mathcal{T}, \mathcal{DI}, \mathcal{E}, \mathcal{R}, \mathcal{D}$).

Here are some relationships between these classes.

- Every full general frame is discrete, and therefore, refined (see below). Every finite, but no infinite, discrete general frame is descriptive, for otherwise the intersection of all sets $W \setminus \{w\}$ would have to be non-empty; on the other hand, every finite differentiated frame is full.
- Every discrete frame is refined. Indeed, for tightness note that in every discrete frame $xR_\alpha w$ holds iff $x \in \langle R_\alpha \rangle(\{w\})$. The converse need not hold, e.g., canonical general frames (see Chapter 7 of this handbook) are refined, even descriptive, but not discrete, being infinite.
- Every elementary frame is discrete, while the converse does not hold, as we will see further.

To summarise: $\mathcal{E} \subsetneq \mathcal{DI} \subsetneq \mathcal{R} = \mathcal{DF} \cap \mathcal{T}$; $\mathcal{D} \subsetneq \mathcal{R}$; $\mathcal{D} \not\subseteq \mathcal{DI}$ $\mathcal{DI} \not\subseteq \mathcal{D}$.

Below, we list some remarks on the various notions of persistence and relationships between them. Analogous remarks apply to local persistence.

- First, note that if $\mathcal{C}_1 \subseteq \mathcal{C}_2$, then $\mathcal{C}_2$ -persistence implies $\mathcal{C}_1$ -persistence.
- A formula is $\mathcal{DI}$ -persistent iff it is valid in a frame $\mathfrak{F}$ whenever it is valid in $\mathfrak{D}(\mathfrak{F})$. Likewise, a formula is $\mathcal{E}$ -persistent iff it is valid in a frame $\mathfrak{F}$ whenever it is valid in $\mathfrak{E}(\mathfrak{F})$.
- While every (locally) $\mathcal{R}$ -persistent formula is $\mathcal{DI}$ -persistent, the converse does not hold, a simple witness being e.g., the ‘density’ formula $\Diamond p \rightarrow \Diamond \Diamond p$ (see [5, p.319]).
- Also, not every (even locally) $\mathcal{D}$ -persistent formula is $\mathcal{DI}$ -persistent (and hence, even less $\mathcal{R}$ -persistent), a witness being *Geach’s formula* $\Diamond \Box p \rightarrow \Box \Diamond p$, defining the Church–Rosser confluence property of the accessibility relation (see [5, p.305]). Moreover, not every $\mathcal{D}$ -persistent formula is $\mathcal{E}$ -persistent, as we will see in section 8.2.
- Not every (even locally) $\mathcal{E}$ -persistent is $\mathcal{DI}$ -persistent, again witnessed by Geach’s formula.
- Finally, not every (even locally) $\mathcal{DI}$ -persistent formula is $\mathcal{D}$ -persistent. The formula $\mathbf{vB} = \Box \Diamond \top \rightarrow \Box(\Box(\Box p \rightarrow p) \rightarrow p)$, proposed by van Benthem in [127], is an example. First, note that for every discrete general frame $\mathfrak{G} = \langle \mathfrak{F}, \mathbb{W} \rangle$ and $w \in \text{dom}(\mathfrak{F})$, $\mathfrak{G}, w \models \mathbf{vB}$ implies $\mathfrak{G}, w \models \Box \Diamond \top \rightarrow \Box \perp$; hence $\mathfrak{G}_\#, w \models \Box \Diamond \top \rightarrow \Box \perp$.

Indeed, assuming $\mathfrak{G}, w \models \Box\Diamond\top \wedge \neg\Box\perp$, for any successor u of w the valuation $W \setminus \{u\}$ for p falsifies $\mathbf{vB}$ at w . Furthermore, for every frame $\mathfrak{F}$ and $w \in \text{dom}(\mathfrak{F})$, $\mathfrak{F}, w \models \Box\Diamond\top \rightarrow \Box\perp$ implies $\mathfrak{F}, w \models \mathbf{vB}$. Hence $\mathbf{vB}$ is locally $\mathcal{DI}$ -persistent. On the other hand, $\mathbf{vB}$ is not $\mathcal{D}$ -persistent. Indeed, as shown in [127] (see also [5, p.216]) $\mathbf{vB}$ is valid in a certain general frame $\mathfrak{J}$, the modal logic $\mathbf{KvB}$ of which is incomplete. That is because $\Box\Diamond\top \rightarrow \Box\perp$, not being valid in $\mathfrak{J}$, is not a theorem of $\mathbf{KvB}$ while, as seen above, it is valid in every frame for $\mathbf{KvB}$. Thus, while $\mathbf{vB}$ is valid in the (descriptive) canonical frame of $\mathbf{KvB}$, it fails in the underlying Kripke frame which falsifies $\Box\Diamond\top \rightarrow \Box\perp$.

Consequently, not every locally $\mathcal{E}$ -persistent formula is $\mathcal{D}$ -persistent.

To summarise again, if we denote by $\mathcal{C}^p$ the set of all $\mathcal{C}$ -persistent formulae, we have the following: $\mathcal{DF}^p \cap \mathcal{T}^p = \mathcal{R}^p \subsetneq \mathcal{DI}^p \subsetneq \mathcal{E}^p$; $\mathcal{R}^p \subsetneq \mathcal{D}^p$; $\mathcal{DI}^p \not\subseteq \mathcal{D}^p \not\subseteq \mathcal{E}^p$.

The same relationships hold for local persistence.

Now, we discuss some important results about refined and descriptive frames and the related persistence properties, while elementary frames and elementary persistence will be discussed in section 8.2.

First, note ([124]) that every general frame $\mathfrak{G} = \langle \mathfrak{F}, \mathbb{W} \rangle$ can be ‘refined’ by constructing a refined quotient of it over the set $W^\sim$ of all equivalence classes modulo the equivalence relation $\sim$, defined as $v \sim w$ iff $\forall X \in \mathbb{W} (v \in X \iff w \in X)$, and taking as admissible all sets of the type $X^\sim = \{w^\sim \mid w \in X\}$ for $X \in \mathbb{W}$. It now remains to ‘tighten’ all accessibility relations by closing under the definition of tightness: for every $u^\sim, w^\sim \in W^\sim$, $u^\sim R_\alpha w^\sim$ holds iff for all $X^\sim \in \mathbb{W}^\sim$ and $u' \sim u, w' \sim w$, if $w' \in X$ then $u' \in \langle R_\alpha \rangle(X)$. Note, however, that (see [10, p.263]) while for finite frames this construction produces a bounded morphic image, this is not necessarily the case when applied to infinite general frames.

Descriptive frames typically appear as the *canonical general frames* (see Chapter 7 of this handbook) of every normal modal logic without any special inference rules. Thus, all $\mathcal{D}$ -persistent formulae are valid in the underlying canonical Kripke frames, and hence they axiomatise Kripke complete logics. For that reason the $\mathcal{D}$ -persistent formulae are also called *canonical*.¹³ However, in hybrid logics with nominals (see Chapter 14 of this handbook) or in logics with special additional rules of inference, e.g., the non- ξ rules in [137], $\mathcal{D}$ -persistent formulae *need not* be canonical, because the canonical general frames for such logics are only discrete (for hybrid logics) or refined (in logics with additional ‘context’ rules, see [52]). In such cases, $\mathcal{DI}$ -persistence or $\mathcal{R}$ -persistence is the right notion of canonicity. $\mathcal{DI}$ -persistent formulae have the important property to remain canonical when added as axioms to hybrid logics with nominals, while $\mathcal{R}$ -persistent formulae remain canonical not only in the presence of other axioms, but even if additional rules of inference of the type mentioned above are added to the axiomatic system.

Descriptive frames feature prominently in the duality theory between general frames and modal algebras, as they turn out to be precisely the fixed points of ultrafilter extensions of general frames, which are essentially the Stone representations of modal algebras (see Chapter 6 of this handbook).

¹³Note that across the literature on modal logic the term ‘canonicity’ is used in somewhat different, and not entirely equivalent, senses (see [126, 127]). For instance, Fine defines in [28] canonicity of a set of formulae as validity of every formula of that set in any canonical frame built for a modal language with any cardinality of propositional variables. Since all canonical models generate descriptive frames, the notion of canonicity adopted here following [126] is at least as strong as Fine’s.

PROPOSITION 113. *A general τ -frame $\mathfrak{G}$ is descriptive iff $\mathfrak{G} \cong \mathfrak{ue}(\mathfrak{G})$.*

Indeed, the proof that every ultrafilter extension is descriptive is just a variation of the proof that every canonical general frame is descriptive (see Chapter 7 of this handbook). For the converse, the crucial observation is that, given a descriptive general frame $\mathfrak{G} = \langle \mathfrak{F}, \mathbb{W} \rangle$, for every $w \in \mathbf{F}$, the set $\mathfrak{u}_{\mathbb{W}}[w] = \{X \in \mathbb{W} \mid w \in X\}$ is an ultrafilter in $\mathbb{W}$, and every ultrafilter in $\mathbb{W}$, due to the compactness of $\mathfrak{G}$, is of this type. Thus, the mapping $\lambda w. \mathfrak{u}_{\mathbb{W}}[w]$ is a bijection (since $\mathfrak{G}$ is differentiated) between $\mathfrak{G}$ and $\mathfrak{ue}(\mathfrak{G})$. This bijection is in fact an isomorphism, due to the tightness of $\mathfrak{G}$.

Consequently, by Theorem 107, every general frame is modally equivalent to a descriptive frame. Therefore, every $\mathcal{D}$ -persistent formula φ preserves its validity from a frame $\mathfrak{F}$ to the ultrafilter extension of the full general frame $\mathfrak{F}_{\max}$, which is based on $\mathfrak{ue}(\mathfrak{F})$. Since $\mathfrak{ue}(\mathfrak{F}_{\max})$ is descriptive, by $\mathcal{D}$ -persistence, φ preserves validity from $\mathfrak{F}$ to $\mathfrak{ue}(\mathfrak{F})$. Conversely, if φ preserves validity in ultrafilter extensions of frames, then it is $\mathcal{D}$ -persistent by Theorem 115. Thus, we obtain:

PROPOSITION 114. *A modal formula is (locally) $\mathcal{D}$ -persistent iff its validity is (locally) preserved in ultrafilter extensions of frames.*

Every general τ -frame $\mathfrak{G} = \langle W, \{R_\alpha\}_{\alpha \in \tau}, \mathbb{W} \rangle$ determines a *topological space* $T(\mathfrak{G})$ with a base of clopen sets $\mathbb{W}$, and a set of closed sets denoted by $\mathbf{C}(\mathbb{W})$. For a detailed study of this topology, its properties and applications in modal logic see [114]. Hereafter, a *closed* set in the general τ -frame $\mathfrak{G}$ will mean a subset of the domain closed with respect to the topology $T(\mathfrak{G})$, i.e., an intersection of a family of admissible sets.

A number of important properties of general frames can be phrased in terms of their topology. For instance, in every discrete frame $\mathfrak{G}$ the topology $T(\mathfrak{G})$ is discrete. Indeed, every non-empty set is a union of its singleton subsets, which are open in $T(\mathfrak{G})$; hence every subset of $\mathfrak{G}$ is open. Also, differentiatedness of a general frame is equivalent to T_2 -separability (Hausdorffness) of its topology, while compactness, as defined above, is equivalent to the standard topological notion of compactness. Thus, for any compact and differentiated τ -frame $\mathfrak{G}$, $T(\mathfrak{G})$ is a compact Hausdorff space.

Finally, it is instructive to explore which constructions on general frames preserve each of the classes discussed above. For instance, differentiatedness, tightness, and discreteness are preserved in generated subframes and disjoint unions, while compactness is not. Conversely, bounded morphisms preserve compactness, but not discreteness, differentiatedness and tightness. Besides, discreteness, differentiatedness, and tightness (and hence, refinedness), being properties definable in a suitable first-order language for states and admissible sets, and membership between them, are preserved in ultraproducts, while descriptiveness is preserved in finite disjoint unions, but never in infinite ones, nor necessarily in ultraproducts [48, 43].

How does persistence determine the expressiveness of a formula? We will discuss this issue in section 8.2 in connection with first-order definability of modal formulae.

Before closing this section, let us highlight again the role of general frames in the modal theory of modal logic:

- general frames provide a natural link between the first-order semantics on Kripke structures and the second-order semantics on frames, and are thus analogous to Henkin's general models for second-order logic.

- general frames are essentially equivalent to modal algebras, via the duality theory outlined in Chapter 6 of this handbook, and thus provide algebraic semantics for modal logic.
- the notion of persistence of (the truth/validity of) modal formulae with respect to natural classes of general frames is instrumental in characterising their model-theoretic behaviour.

8 MODAL LOGIC ON FRAMES

So far we have mainly studied modal logic as a fragment of first-order logic over Kripke structures. In this section we discuss modal logic as a logic of *frames*, and thus as a fragment of universal monadic second-order logic MSO.

This fragment, while generally not very expressive and missing many simple first-order properties, nevertheless penetrates deeply into MSO. Perhaps its most interesting features are the recursive axiomatisability of validity and its finite model property, together implying decidability – a rare phenomenon in second-order logic when considered over arbitrary structures rather than special ones.

In this section we present some classical results characterising modally definable classes of frames, and discuss how persistence of modal formulae with respect to various classes of general frames can be used to determine their model-theoretic properties.

8.1 Modal definability of frame properties

Here we address the question which classes of frames are definable by modal formulae. A classical result from [51] answers this question in a traditional model-theoretic fashion, albeit using a somewhat ad-hoc construction, called *SA-construction* (‘state-of-affairs construction’). Algebraically, it corresponds to taking a subalgebra of a homomorphic image, thus allowing a ‘translation’ of Birkhoff’s theorem in terms of frame constructions, and so characterising equational classes of algebras as those closed under subalgebras, homomorphic images and direct products (see Chapter 6 of this handbook). Theorem 117 gives a more natural characterisation of the modally definable *elementary* classes. Here is another definability-by-preservation result, due to van Benthem (see [126, Theorem 3.5], [127, Theorem 16.5], [129]).

THEOREM 115. *A class of frames $\mathcal{K}$ is modally definable by a set of $\mathcal{D}$ -persistent formulae iff it is closed under generated subframes, bounded morphisms, disjoint unions and ultrafilter extensions, and reflects ultrafilter extensions.*

Proof. We already know from sections 2 and 6, and Proposition 114 that every $\mathcal{D}$ -persistent formula satisfies all preservation conditions of the theorem, whence the easier direction. Conversely, let $\mathcal{K}$ satisfy the preservation conditions. We show that $\mathcal{K} = \text{FR}(\text{Th}_{\text{ML}}(\mathcal{K}))$. Let $\mathfrak{F} \models \text{Th}_{\text{ML}}(\mathcal{K})$. Recall that $\mathfrak{F}_{\max}$ denotes the full general frame over the frame $\mathfrak{F}$. Using the duality theory between general frames and modal algebras, and Birkhoff’s theorem, one can show that $\text{ue}(\mathfrak{F}_{\max})$ is isomorphic to a generated general subframe of a bounded morphic image of $\text{ue}(\mathfrak{G}_{\max})$ where $\mathfrak{G}$ is a disjoint union of frames from $\mathcal{K}$. Now, $\mathfrak{G} \in \mathcal{K}$; hence $\text{ue}(\mathfrak{G}) \in \mathcal{K}$. So, tracing the underlying frames and using the closure conditions, we eventually find that $\text{ue}(\mathfrak{F}) \in \mathcal{K}$, whence $\mathfrak{F} \in \mathcal{K}$. ■

We note that checking the conditions of the theorem above, even in the case when the class of frames is first-order definable, may be a practically very difficult task. A testimony for that is the fact that preservation of first-order formulae under ultrafilter extensions is Π_1^1 -hard [122, Thm 2.3.17].

In the rest of this section we compare the expressiveness of modal logic over frames with first-order logic and some of its extensions within monadic second-order logic.

8.2 Modal logic versus first-order logic on frames

We have already seen that modal languages are generally incomparable with first-order languages in terms of definability of frame properties. Indeed, while simple elementary properties, such as irreflexivity, escape the basic modal language, it can capture non-elementary properties such as the one defining the class of all transitive frames in which there are no infinite chains of successors. By a simple compactness argument, this class is not elementary, while it is well-known to be defined by the *Gödel–Löb formula* GL (see e.g. [75]). This example also shows that the compactness theorem with respect to frame validity fails in modal logic. The downward Löwenheim–Skolem–Tarski theorem fails here, too. E.g., McKinsey’s formula $\Box \Diamond p \rightarrow \Diamond \Box p$ (see [127], or [5, p.133]) is valid in a certain uncountable frame, but not in any countable elementary subframe of it. Another important example of a non-elementary modal formula (in the extended setting with the star operation for transitive closures) is Segerberg’s induction axiom [117] IND : $[\alpha^*](p \rightarrow [\alpha]p) \rightarrow (p \rightarrow [\alpha^*]p)$.

The model-theoretic interplay

We compare modal formulae (respectively, modally definable properties of frames) and first-order formulae (respectively, properties definable in $\text{FO}(\tau)$) from two perspectives:

- Which modally definable frame properties are first-order definable?
- Which first-order properties of frames are modally definable?

As already mentioned, there are two natural notions of first-order definability: by means of single sentences and by means of theories (possibly infinite sets of sentences). Regarding modally definable classes, however, these turn out to be equivalent. Indeed, if the class of frames $\text{FR}(\varphi)$ is the class of models of an infinite set of $\text{FO}(\tau)$ -formulae Γ , then $\Gamma \models \varphi$ with respect to frame validity, which is a Π_1^1 -property. The compactness theorem of first-order logic applies here, and $\Gamma_0 \models \varphi$ for some finite $\Gamma_0 \subseteq \Gamma$. Hence $\text{FR}(\varphi)$ is defined by the conjunction over Γ_0 . We can therefore refer to modally definable classes which are first-order definable, and to modal formulae defining such classes, as *elementary* without risk of confusion.

On the other hand, it seems to be still unknown whether there is any FO-sentence equivalent to an infinite set of basic modal formulae but not to a single formula.¹⁴

The validity preservation results from sections 2 and 6 imply that every modally definable class of frames $\text{FR}(\varphi)$ is closed under generated subframes, bounded morphic images (in particular, isomorphic copies), and disjoint unions, while it reflects ultrafilter

¹⁴There are known cases, however, where a first-order definable property is infinitely, but not finitely, axiomatisable in some extended modal languages. See, e.g., [53].

extensions and ultraproducts. If, moreover, the formula φ is elementary, then $\text{FR}(\varphi)$ is closed under ultraproducts, too. Conversely, if $\text{FR}(\varphi)$ is closed under ultraproducts then, by the Keisler–Shelah theorem, $\text{FR}(\varphi)$ is elementary. Moreover, by Proposition 85, closure of $\text{FR}(\varphi)$ under ultrapowers suffices, and therefore, closure under elementary equivalence in $\text{FO}(\tau)$ suffices, too. The latter, in turn, characterises $\Sigma\Delta$ -elementary classes, i.e., unions of elementary classes. Thus, we have the following model-theoretic characterisation of the elementary modal formulae (see [48, 43, 127]).

THEOREM 116. *For any modal formula φ the following are equivalent:*

- (i) φ is elementary.
- (ii) $\text{FR}(\varphi)$ is closed under ultraproducts.
- (iii) $\text{FR}(\varphi)$ is closed under ultrapowers.
- (iv) $\text{FR}(\varphi)$ is closed under elementary equivalence, i.e., $\Sigma\Delta$ -elementary.

The result above correspondingly characterises elementary classes of frames that are known to be modally definable. This raises the natural question how to characterise, in model theoretic terms, modal definability of an elementary class of frames. Again, a classical result from [51] answers that question. Here is a somewhat strengthened version (see [5, Theorem 5.54]).

THEOREM 117 (Goldblatt–Thomason). *If a class of frames $\mathcal{K}$ is closed under ultrapowers (in particular, if $\mathcal{K}$ is elementary) then $\mathcal{K}$ is modally definable iff it is closed under generated subframes, bounded morphisms, and disjoint unions, and reflects ultrafilter extensions.*

Proof. One direction is a direct application of the preservation results from sections 2 and 6. For the other direction note that, by Theorem 90 reduced to underlying frames, $\mathcal{K}$ is closed under ultrafilter extensions, too. Thus, Theorem 115 applies, so $\mathcal{K}$ is modally definable, moreover by a set of $\mathcal{D}$ -persistent formulae. ■

We end with an important related result, originally due to Fine [28], later strengthened and proved by van Benthem [127, Theorem 16.7] as a corollary to Theorem 115.¹⁵

We call a modal formula φ *complete* if the modal logic axiomatised by φ is complete for the class of frames defined by φ .

THEOREM 118 (Fine–van Benthem).

Every complete and elementary modal formula φ is $\mathcal{D}$ -persistent.

Proof. $\text{FR}(\varphi)$ satisfies all closure conditions of Theorem 115, so $\text{FR}(\varphi) = \text{FR}(\Gamma)$ for some set of $\mathcal{D}$ -persistent formulae Γ . The modal logic $\mathbf{K}_\tau + \Gamma$, axiomatised with the set of axioms Γ , is canonical and therefore complete. Hence $\mathbf{K}_\tau + \Gamma \vdash \varphi$. By compactness of modal derivations, $\mathbf{K}_\tau + \Gamma_0 \vdash \varphi$ for some finite subset Γ_0 of Γ . By completeness of φ , all formulae from Γ_0 are theorems of $\mathbf{K}_\tau + \varphi$. Hence φ is axiomatically equivalent, and therefore frame-equivalent, too, to the conjunction of Γ_0 , which is itself a $\mathcal{D}$ -persistent formula. ■

It is known ([28], see also section 8.2) that the converse to the above theorem does not hold, viz. not every $\mathcal{D}$ -persistent formula is elementary. An example is $\Diamond\Box(p \vee q) \rightarrow$

¹⁵For a stronger algebraic version of this theorem see [49], [5, Theorem 5.56], or Chapter 6 of this handbook.

$\Diamond(\Box p \vee \Box q)$ (see [28]). Nor is every elementary modal formula $\mathcal{D}$ -persistent, as there are incomplete elementary modal formulae (e.g., van Benthem's formula $\mathbf{vB}$ discussed in Section 7.3, see [128, p.72], also in [5, p.216]).

It had been a longstanding open problem, posed by Fine, whether every modal logic axiomatised by $\mathcal{D}$ -persistent formulae is complete with respect to *some* elementary class. This question has recently been answered negatively in [45].

Persistence and first-order definability

Some persistence properties of modal formulae imply that they are elementary. Perhaps the first interesting result in that vein is due to Lachlan [89] who proved that every $\mathcal{R}$ -persistent formula is elementary. A strengthening of Lachlan's result, using the argument in Goldblatt's proof of it in [48], is that every (locally) $\mathcal{DT}$ -persistent formula is (locally) elementary. First, note that local non-validity of a modal formula, being a Σ_1^1 -property, is preserved by ultraproducts [12, Corollary 4.1.14]. By the Keisler–Shelah theorem it suffices to show that local validity of locally $\mathcal{DT}$ -persistent formulae is preserved under ultraproducts. This follows from the fact that local validity of modal formulae is locally preserved in ultraproducts of general frames (Proposition 109), and that any ultraproduct of full general frames is a discrete general frame.

Let us now turn to $\mathcal{E}$ -persistent formulae. They were first studied by van Benthem in [127] in connection with the *substitution method* which can be used to establish the first-order definability of Sahlqvist formulae (see section 8.2). The idea of the substitution method is to identify finitely many ‘characteristic’ first-order definable valuations of the variables occurring in a given formula, such that the formula is (locally) valid in every frame in which it is (locally) valid for those characteristic valuations. For all Sahlqvist formulae, just one such valuation, the *minimal* one amongst all those satisfying the antecedent of the formula, suffices. Van Benthem provided an alternative characterisation of locally and globally $\mathcal{E}$ -persistent formulae, which implies that they are locally elementary.

Given a $\text{FO}(\tau_\Phi)$ -formula $\beta(\mathbf{x})$ with unary predicates $P_1, \dots, P_n$, assuming that the variables $\mathbf{x}$ do not occur bound in β and the variables $z_1, \dots, z_k, y$ do not occur in β at all, we define a *universally parameterised* $\text{FO}(\tau)$ -*substitution instance* of β to be any $\text{FO}(\tau)$ -formula $\forall z_1 \dots \forall z_k \beta[\sigma_1/P_1, \dots, \sigma_n/P_n]$ obtained from β by selecting $\text{FO}(\tau)$ -formulae $\sigma_i = \sigma_i(\mathbf{x}, z_1, \dots, z_k, y)$ for $i = 1, \dots, n$, uniformly substituting $\sigma_i[x/y]$ for every occurrence of $P_i x$, and then universally quantifying over $z_1, \dots, z_k$. Let $\Theta(\beta)$ be the set of all universally parameterised $\text{FO}(\tau)$ -substitution instances of β .

DEFINITION 119. A modal formula $\varphi = \varphi(p_1, \dots, p_n)$ is a *van Benthem formula* if $\Theta(\text{ST}(\varphi; x_0)) \models \forall P_1 \dots \forall P_n \text{ST}(\varphi; x_0)$. We let VB denote the class of van Benthem formulae (defined slightly differently in [127] as the class M_{sub}^1).

THEOREM 120. A modal formula is locally $\mathcal{E}$ -persistent iff it is a van Benthem formula.

Proof. Recall that $\mathfrak{E}(\mathfrak{F})$ is the minimal elementary general frame over the Kripke frame $\mathfrak{F}$. Let $\varphi(p_1, \dots, p_n) \in \text{VB}$ and suppose $\mathfrak{E}(\mathfrak{F}), w \models \varphi$ for some frame $\mathfrak{F}$. Take any universally parametrised $\text{FO}(\tau)$ -substitution instance $\forall z_1 \dots \forall z_k \text{ST}(\varphi)[\sigma_1/P_1, \dots, \sigma_n/P_n]$. Let $w_1, \dots, w_k \in \text{dom}(\mathfrak{F})$ and $X_i := \{u \in \text{dom}(\mathfrak{F}) \mid \mathfrak{F} \models \sigma_i(w, w_1, \dots, w_k, u)\}$ for $i = 1, \dots, n$. Since $X_1, \dots, X_n$ are admissible in $\mathfrak{E}(\mathfrak{F})$, $(\mathfrak{F}; X_1, \dots, X_n; w) \models \text{ST}(\varphi)(P_1, \dots, P_n; x_0)$.

Therefore, $\mathfrak{F}, w \models \forall z_1 \dots \forall z_k \text{ST}(\varphi)[\sigma_1/P_1, \dots, \sigma_n/P_n]$. Since $\varphi \in \text{VB}$, that implies $\mathfrak{F}, w \models \varphi$.

Conversely, let φ be locally $\mathcal{E}$ -persistent and suppose $\mathfrak{F}, w \models \Theta(\text{ST}(\varphi; x_0))$. Then, reversing the argument above, we find that $\mathfrak{E}(\mathfrak{F}), w \models \varphi$, and therefore $\mathfrak{F}, w \models \varphi$ by local $\mathcal{E}$ -persistence of φ . ■

We can now strengthen the earlier persistence-implies-elementary results.

THEOREM 121. *Every (locally) $\mathcal{E}$ -persistent formula is (locally) elementary.*

Proof. Clearly, for every modal formula φ , $\forall P_1 \dots \forall P_n \text{ST}(\varphi; x_0) \models \Theta(\text{ST}(\varphi; x_0))$. By compactness, every van Benthem formula is a logical consequence of a finite subset of $\Theta(\text{ST}(\varphi; x_0))$, and hence is equivalent to the conjunction over that set. ■

Consequently, not every $\mathcal{D}$ -persistent formula is $\mathcal{E}$ -persistent. Neither is every (locally) elementary modal formula (locally) $\mathcal{E}$ -persistent. An example (see [127]) is the formula $\text{Mk4} = (\Box p \rightarrow \Box \Box p) \wedge (\Box \Diamond p \rightarrow \Diamond \Box p)$, which is elementary and valid in the general frame $\langle \mathbb{N}, <, \mathbb{W} \rangle$ where $\mathbb{W}$ is the set of all finite and co-finite subsets of $\mathbb{N}$, while it fails in $\langle \mathbb{N}, < \rangle$. Since $\mathbb{W}$ contains precisely all parametrically first-order definable sets in $\langle \mathbb{N}, < \rangle$, it is $\mathfrak{E}(\langle \mathbb{N}, < \rangle)$, so Mk4 is not $\mathcal{E}$ -persistent. Similarly, $\text{Mk4}' = (\Box p \rightarrow \Box \Box p) \wedge \Box(\Box p \rightarrow \Box \Box p) \wedge (\Box \Diamond p \rightarrow \Diamond \Box p)$ is locally elementary,¹⁶ but not locally $\mathcal{E}$ -persistent.

Sahlqvist formulae and inductive formulae

The model-theoretic results discussed above, however elegant, are usually not easy to apply, and are of no use to find the actual first-order formula corresponding to the modal formula. It is therefore natural to look for simpler and *effective* sufficient conditions for first-order definability of modal formulae. There can be no completely satisfactory outcome of that search, because that property is not decidable [11], and (at least) in a modal language with more than one modality, not even analytical [122, Thm 2.6.5]. Still, several increasingly general results to that aim were obtained during the 1970's, culminating with the celebrated *Sahlqvist theorem*, which not only identifies a large syntactic class of elementary modal formulae (see a simple definition of that class below), but also proves their canonicity. A variety of expositions of Sahlqvist's theorem can be found in several sources, e.g. [113, 115, 5, 84, 10], Chapters 6 and 7 of this handbook. Here we outline a generalisation of the class of Sahlqvist formulae in monadic poly-modal languages, sharing the same virtues as the original class, viz. the *inductive formulae* introduced and studied for arbitrary polyadic languages in [54].

We fix a modal language $\text{ML}(\tau)$.

DEFINITION 122. Let $\#$ be a symbol not belonging to $\text{ML}(\tau)$. Then a *box-form* of $\#$ in $\text{ML}(\tau)$ is defined recursively as follows:

- (i) $\#$ is a box-form of $\#$;
- (ii) If $\mathbf{B}(\#)$ is a box-form of $\#$ and $\Box$ is a box-modality in $\text{ML}(\tau)$ then $\Box \mathbf{B}(\#)$ is a box-form of $\#$;
- (iii) If $\mathbf{B}(\#)$ is a box-form of $\#$ and A is a positive τ -formula then $A \rightarrow \mathbf{B}(\#)$ is a box-form of $\#$.

¹⁶The fact that Mk4 and $\text{Mk4}'$ are elementary is far from trivial, as the proof requires a form of the Axiom of Choice and cannot be formalised in ZF.

Thus, box-forms of $\#$ are, up to semantic equivalence, of the type $\Box_1(A_1 \rightarrow \Box_2(A_2 \rightarrow \dots \Box_n(A_n \rightarrow \#) \dots))$, where $\Box_1, \dots, \Box_n$ are box-modalities and $A_1, \dots, A_n$ are positive formulae in $\text{ML}(\tau)$.

DEFINITION 123. Given a propositional variable p , a *box-formula* of p is the result $\mathbf{B}(p)$ of substitution of p for $\#$ in any box-form $\mathbf{B}(\#)$. The last occurrence of the variable p is the *head* of $\mathbf{B}(p)$ and every other occurrence of a variable in $\mathbf{B}(p)$ is *inessential* there.

DEFINITION 124. A (*monadic*) *regular formula* is any modal formula built from positive formulae and negations of box-formulae by applying conjunctions, disjunctions, and boxes.

DEFINITION 125. The *dependency digraph* of a set $\mathcal{B} = \{\mathbf{B}_1(p_1), \dots, \mathbf{B}_n(p_n)\}$ of box-formulae is the digraph $G = \langle V, E \rangle$ where $V = \{p_1, \dots, p_n\}$ is the set of heads in $\mathcal{B}$, and $p_i E p_j$ iff p_i occurs as an inessential variable in a box-formula from $\mathcal{B}$ with a head p_j . A digraph is called *acyclic* if it does not contain oriented cycles.

DEFINITION 126. An *inductive formula* is a regular formula with an acyclic dependency digraph of the set of all box-formulae occurring as subformulae in it.

We note that Sahlqvist formulae, up to semantic equivalence, are precisely those regular formulae in which the box-formulae are just *boxed atoms*, i.e., propositional variables prefixed by possibly empty strings of boxes. Thus, all Sahlqvist formulae fall into a simple particular case of inductive formulae, where the dependency digraph has no arcs at all.

The following extension of Sahlqvist's theorem was established in [54].

THEOREM 127. *Each inductive formula is locally elementary and locally $\mathcal{D}$ -persistent. Moreover, its local first-order equivalent can be computed effectively.*

The inductive formulae are van Benthem formulae which, just like Sahlqvist formulae, have first-order definable minimal valuations, but they can only be computed inductively, in steps following the arcs of the dependency digraph, from sources to sinks.

Sahlqvist formulae satisfy a certain persistence property which can be extracted from the syntactic shape of the first-order formulae defining their minimal valuations. In the basic modal language these valuations are either the empty set, or the whole domain, or are finite unions of sets of the type $R^n(y)$ (recall that R^n is the n -fold composition of R with itself). Following [54], let us call a general frame *ample* if it contains all such sets as admissible, and the modal formulae locally persistent with respect to all ample general frames, *locally $\mathcal{A}$ -persistent*. Thus, all Sahlqvist formulae in $\text{ML}(\Diamond)$ are locally $\mathcal{A}$ -persistent, and this property enables us to show that a given formula is not (even semantically equivalent to) a Sahlqvist formula.

EXAMPLE 128. As proved in [54], the formula $D = p \wedge \Box(\Diamond p \rightarrow \Box q) \rightarrow \Diamond \Box \Box q$ is not $\mathcal{A}$ -persistent, and hence not equivalent to any Sahlqvist formula in $\text{ML}(\Diamond)$. However, it is an inductive formula, whose dependency digraph over the set of heads $\{p, q\}$ has only one edge, from p to q . It has a local first-order correspondent $\text{FO}(D) = \exists y(Rxy \wedge \forall z(R^2yz \rightarrow \exists u(Rxu \wedge Rux \wedge Ruz)))$, which is not equivalent to a Kracht formula (i.e., a first-order equivalent to a Sahlqvist formula, see [84]).

The class of inductive formulae does not exhaust the potential of the method of substitutions, (in particular, minimal valuations), since, being syntactically defined (like Sahlqvist formulae), it is not closed even under tautological equivalence.

A more general and robust *algorithmic* approach to identifying elementary and $\mathcal{D}$ -persistent modal formulae (covering all inductive formulae) is outlined in [14]. The algorithm presented there is based on a modal version of Ackermann's lemma (which essentially formalises the idea of minimal valuations) and, when successful, computes effectively a first-order equivalent of the input modal formula and at the same time establishes its $\mathcal{D}$ -persistence.

Shallow formulae and $\mathcal{R}$ -persistence

The property of $\mathcal{R}$ -persistence is much stronger than $\mathcal{D}$ -persistence. Perhaps the largest syntactic class of $\mathcal{R}$ -persistent formulae identified so far is the class of *shallow formulae* [122, Thm 2.4.7]: those in which every occurrence of a propositional variable is in the scope of at most one modal operator. Note that syntactically shallow formulae are not subsumed by the class of Sahlqvist formulae, nor even by the class of inductive formulae.

8.3 Modal logic and first-order logic with least fixed points

With every first-order language $\text{FO}(\tau)$ we associate its extension $\text{LFP}(\tau)$ with *least fixed point operators*. For background on LFP see e.g. [25] or [2]. LFP is a rather expressive proper extension of FO which however still shares nice properties with FO, e.g., the downward Löwenheim–Skolem theorem [30] and the 0-1 law (see [64]).

Which modal formulae are (locally) definable in $\text{LFP}(\tau)$? Which $\text{LFP}(\tau)$ -formulae are modally definable on frames? No explicit model-theoretic criteria seem to be known as yet and these questions are most likely undecidable.

A number of well-known non-elementary modal formulae, such as the Gödel–Löb formula GL and Segerberg's induction axiom IND have local equivalents in $\text{LFP}(\tau)$ while, for instance, the McKinsey formula is outside that class. Indeed, take van Benthem's uncountable frame from [127] in which that formula is valid. Flum's argument from [30], proving the downward Löwenheim–Skolem–Tarski theorem for LFP, produces a countable elementary subframe of it which must satisfy that formula, too, which is not possible, as shown in [127].

Still, a large, effectively defined class of $\text{LFP}(\tau)$ -expressible modal formulae can be identified by noting that the idea of using minimal valuations to eliminate the universal second-order quantifiers in the standard translation of frame validity of modal formulae goes beyond first-order logic. Indeed, the same idea works perfectly for all (polyadic) regular formulae, defined for monadic languages in section 8.2. In cases where the dependency graph has loops and cycles, the minimal valuations are recursively defined and eventually expressed in $\text{LFP}(\tau)$. In particular, this applies to Gödel–Löb and Segerberg formulae, being regular formulae. The following was shown in [54].

THEOREM 129. *Every regular formula has a local correspondent in $\text{LFP}(\tau)$, which can be obtained effectively.*

We illustrate the idea of computing $\text{LFP}(\tau)$ -equivalents of regular formulae with GL.

$$\text{ST}(\text{GL}) = \forall x_1(x_0 R x_1 \rightarrow (\forall x_0(x_1 R x_0 \rightarrow P x_0) \rightarrow P x_1)) \rightarrow \forall x_1(x_0 R x_1 \rightarrow P x_1),$$

which can be rewritten as $\forall x_1(x_0 R x_1 \rightarrow (R[x_1] \subseteq P) \rightarrow P x_1) \rightarrow R[x_0] \subseteq P$ (where $R[x] := \{y \mid x R y\}$). The antecedent can be expressed as

$$\Phi(P) \subseteq P, \text{ where } \Phi(P) = \{x_1 \mid x_0 R x_1 \wedge R[x_1] \subseteq P\}.$$

Note that, since $\Phi(P)$ is positive in P , and hence monotone, there is a $\subseteq$ -minimal valuation for P satisfying $\Phi(P) \subseteq P$, viz. $V_m(p) = \mu X.\Phi(X)$. Then, the local equivalent of GL in $\text{LFP}(\tau)$ is obtained by substituting that minimal valuation in the consequent: $\text{LFP}(\tau)(GL; x_0) = \forall x_1(x_0 R x_1 \rightarrow \mu X.\Phi(X)(x_1))$. By unfolding, based on the Knaster–Tarski theorem, that equivalent is:

$$\begin{aligned} & \forall x_1(x_0 R x_1 \rightarrow \\ & \exists n \geq 0 \forall y_1 \dots \forall y_n (x_1 R y_1 \rightarrow x_0 R y_1 \wedge (\dots (y_{n-1} R y_n \rightarrow x_0 R y_n \wedge R[y_n] = \emptyset) \dots)), \end{aligned}$$

i.e., ‘local’ transitivity and non-existence of infinite R -chains starting at x_0 .

While Theorem 129 may be regarded as an extension of the definability part of the Sahlqvist theorem, it cannot match the canonicity part of it. Not only are there regular formulae which are not $\mathcal{D}$ -persistent (e.g., GL and IND) but there are even ones which are not complete, such as $\Box(\Box p \leftrightarrow p) \rightarrow \Box p$ from [6], which can be easily pre-processed into a semantically equivalent regular formula. It is weaker than GL but has the same class of frames, and is therefore incomplete. On the other hand, it is a plausible conjecture that every modal formula with a minimal valuation expressible in $\text{LFP}(\tau)$ is semantically equivalent to a regular formula.

In order to apply the method of minimal valuations, one has to identify, en route, those $\text{FO}(\tau)$ -formulae γ for which there is a minimal interpretation for each occurring unary predicate P . In recent work van Benthem [132] has obtained syntactic and model theoretic characterisations of these formulae, involving predicates of arbitrary arity (see Chapter 1 of this handbook).

Finally, we note that an algorithm for computing $\text{LFP}(\tau)$ -equivalents of classical modal formulae, based on Ackermann’s method for second-order quantifier elimination, and in particular covering the example above, has been developed in [103].

8.4 Modal logic and second-order logic

The standard translation embeds $\text{ML}(\tau)$, with respect to frame validity, into the monadic Π_1^1 -extension of the first-order language $\text{FO}(\tau_\Phi)$. We already know that the embedding is proper. Still, a natural question arises whether the preservation conditions of Theorem 117 are sufficient to guarantee modal definability of monadic Π_1^1 -formulae, as well. As van Benthem has noted in [127, p.53], this is not the case in the basic modal language, witnessed by the property ‘non-existence of infinite R -chains’ (i.e., well-foundedness of R^{-1}), which satisfies all those preservation conditions and moreover is bisimulation invariant. Still, that property of frames is defined in the extension of the basic modal language with the universal modality $[U]$, by the formula $[U](\Box p \rightarrow p) \rightarrow p$ (see [53]). (Contrast this with observation 42, that as a property of Kripke structures, it is not definable even in ML_∞ .) Thus, one may ask if the natural preservation conditions characterising modal definability of elementary properties (closure under generated subframes, bounded morphisms, and disjoint unions, and reflection of ultrafilter extensions) do not apply also to a wider class (if not the whole of Π_1^1), but for a suitably extended modal language? Surely, some of the results characterising modal definability of properties of Kripke structures would still be useful and relevant here: if the first-order matrix of a Π_1^1 -formula, where all second order quantifiers are in a prefix, meets the conditions for having a modal correspondent on Kripke structures, then the whole formula is frame-definable by the same modal correspondent. It is not currently known if this observation can be turned into a general criterion for modal definability of monadic second-order formulae.

Modal logic penetrates quite deep into monadic second-order logic $\text{MSO}(R)$ (with full quantification over unary predicate variables, over the vocabulary with the single binary relation R). As proved by Thomason [124], logical consequence in terms of frame validity of the latter can be reduced to the former in the following sense. There exists an effective translation t of $\text{MSO}(R)$ into ML, and a special modal formula δ such that for every set Σ of $\text{MSO}(R)$ -sentences and any $\text{MSO}(R)$ -sentence φ : $\Sigma \models_2 \varphi$ iff $\{\delta\} \cup t[\Sigma] \models_{\text{FR}} t(\varphi)$. Here $\models_2$ denotes second-order semantic consequence which, as a consequence from Tarski's non-definability theorem, is not arithmetically definable, and $\Gamma \models_{\text{FR}} \psi$ means that the modal formula ψ is valid in every frame where all modal formulae from Γ are valid. Consequently, $\models_{\text{FR}}$ is not recursively axiomatisable, unlike validity in modal logic.

Furthermore, as noted in [127, p.23], full second-order logic, and even the theory of finite types, can be reduced to $\text{MSO}(R)$, too.

For more on the relations between modal logic and second-order logic, see [127], [24], and Chapter 10 of this handbook. Also, [122, 121, Chapter 12] considers the extension of modal logic with propositional quantifiers, which goes much farther into second-order logic.

9 FINITE MODEL THEORY OF MODAL LOGICS

9.1 *Finite versus classical model theory*

When only finite structures are admitted, the model theoretic basis changes dramatically. For instance, unless the logic under consideration has the finite model property, satisfiability does not imply finite satisfiability, and hence a semantic consequence $\varphi \models \psi$ may be true in the sense of finite models without being classically valid. Crucial tools of classical model theory, most notably the completeness and compactness theorems for FO, fail in restriction to just finite models. From a modelling point of view, on the other hand, the restriction to just finite models is often natural. In applications, in which the intended models ought to be finite, reasoning on the basis also of infinite models may be inadequate and give misleading results. Applications in computer science like specification and verification, or also database theory, for instance, often call for the restriction to finite models, and have had a significant impact on the development of finite model theory.

The methodological shift encountered is highlighted by the failure of classical theorems and tools, most notably of the compactness theorem but also most other key theorems from classical model theory in its wake, see [25]. Certainly results from classical model theory cannot be expected to go through automatically; often they fail, and some still obtain, albeit with new proofs. Modal model theory, in particular, has a number of examples of the latter kind, and sometimes the new proofs shed new light also on the classical version. For some concrete examples, close to (classical) modal model theory, which illustrate the interesting relationship with finite model theory, consider the following.

Interpolation for ML goes through via the finite model property (FMP), treated in section 3.4. If $\models \varphi \rightarrow \psi$ is valid in finite structures, it must also be valid generally, as a counterexample $\mathfrak{M}, w \models \varphi \wedge \neg\psi$ would also yield a counterexample in the sense of finite model theory, by FMP. Clearly a classical interpolant χ with $\models (\varphi \rightarrow \chi) \wedge (\chi \rightarrow \psi)$, is an interpolant also in the sense of finite model theory.

The modal characterisation theorem. Note how both sides of the equivalence expressed in Theorem 55 change their meaning when interpreted in the sense of finite model theory: both bisimulation invariance and logical equivalence only refer to finite structures. In particular, bisimulation invariance in finite structures does not imply bisimulation invariance over all structures. Trivial examples are provided by formulae without finite models that happen not to be bisimulation invariant for infinite models. Also, while Ehrenfeucht–Fraïssé techniques remain valid, compactness does not and the classical proof with its necessary detour through infinite models is no longer available. As discussed in section 4.3, however, the theorem itself persists in the form of Theorem 61 as a theorem of finite model theory due to Rosen [112]. Interestingly the new proofs in [112, 105] are valid classically as well as in finite model theory and have lead to additional insights into the classical result. In contrast, the failure of the corresponding characterisation theorem for FO^2 in finite model theory shows that the finite model property does not guarantee a smooth passage to finite model theory. While an FO sentence that is (classically) invariant under 2-pebble game equivalence is logically equivalent to a sentence in FO^2 , this characterisation breaks down for finite model theory. The FO sentence saying that a binary relation is a linear ordering, which is 2-pebble invariant only in restriction to finite structures, is not expressible in FO^2 even over finite structures.

Similarly, Rosen [112] has a proof of the finite model theory version of the modal existential preservation theorem: $\varphi \in \text{ML}$ is preserved under extensions (holds inside the whole Kripke structure if it holds in a substructure) iff it is equivalent to an existential modal formula (built from positive and negated atoms by means of only $\wedge$, $\vee$ and $\Diamond$ – disallowing $\Box$ or nesting of $\neg$ and $\Diamond$). The corresponding preservation theorem for first-order logic is known to become invalid in restriction to just finite structures.

Modal logic stands out in comparison with first-order logic or the FO^k in having a comparatively smooth finite model theory that preserves a number of classical theorems, as is the case for the above examples.

The variations of basic modal logic mentioned in section 5.1 have partly also been investigated with respect to their finite model theory, with several results that suggest a similarly smooth behaviour. Their characterisations as fragments of FO, in terms of invariance under correspondingly refined notions of bisimulation, have been studied in finite model theory in [106] with further ramifications w.r.t. other restricted classes of finite frames in [17]. Just as is the case with van Benthem–Rosen characterisation, Theorems 55 and 61 surprisingly many of these characterisations go through in restriction to finite Kripke structures just as classically, albeit with rather specific new proofs. The following may serve as a typical representative for several related results from [106, 17]. Also compare Proposition 68; this should be contrasted with the failure of, for instance, the corresponding characterisation of FO^2 in finite model theory.

THEOREM 130. *For any $\varphi(x) \in \text{FO}$, the following are equivalent:*

- (i) *φ is invariant under global bisimulation over finite Kripke structures.*
- (ii) *φ is equivalent to a formula of $\text{ML}[\forall]$ over finite Kripke structures.*

Similarly are equivalent:

- (i) *φ is bisimulation invariant over finite, rooted Kripke structures.*
- (ii) *φ is equivalent to a formula of $\text{ML}[\forall]$ over finite, rooted Kripke structures.*

Related open problems concern the status in finite model theory of Theorem 65, for the

guarded fragment GF in arbitrary relational similarity types, and particularly strikingly of Theorem 76, for the modal μ -calculus.

But finite model theory also deals with new questions, which only arise in the context of finite structures. We devote the rest of this section to two sketches dealing with two very specific issues of this kind: one from descriptive complexity (section 9.2), the other one 0-1 laws (section 9.3). Descriptive complexity deals with the relationship between the algorithmic complexity and the logical definability of properties of finite structures; here finite structures feature as input to algorithmic problems and logic becomes a measure of complexity. In 0-1 laws, and more generally asymptotic probability, one deals with the statistics of logically defined properties over the collection of all size n structures in the limit as n goes to infinity; here finite structures form the sample space for probabilistic analysis. Compare [25, 93] for general background on these topics in finite model theory.

9.2 Capturing bisimulation invariant Ptime

Descriptive complexity aims for the description and analysis of computational complexity by means of logics. A key example is the long open problem of a *logic for Ptime*. One seeks a logic (with effective syntax) whose formulae define precisely those classes of finite relational structures, for which membership can be decided in polynomial time.¹⁷ By a well-known result of Immerman [76] and Vardi [134], the least fixed point extension of first-order logic, LFP, is the solution for classes of finite, linearly ordered relational structures. The problem remains open to date for not necessarily ordered structures. Interestingly, the corresponding problem for bisimulation closed classes of finite Kripke structures does admit a natural solution [104] (cf. [94] for another, related capturing result).

Consider the framework of basic modal logic with a single modality associated with the binary relation R and with finitely many atomic propositions p_i . Let $\mathcal{Q}$ be a class of finite pointed Kripke structures (i.e., a property of finite pointed Kripke structures) of that type. $\mathcal{Q}$ corresponds to a *bisimulation invariant* property if it is closed under bisimulation in the sense that for any two $(\mathfrak{M}, u) \rightleftharpoons (\mathfrak{M}', u')$: $(\mathfrak{M}, u) \in \mathcal{Q}$ iff $(\mathfrak{M}', u') \in \mathcal{Q}$. Recall the bisimulation quotients $\mathfrak{M}[u]/\rho^{\mathfrak{M}}$ of pointed Kripke structures $(\mathfrak{M}, u)$ as discussed in section 3.8. Bisimulation closure of $\mathcal{Q}$ implies that

$$\mathcal{Q} = \{(\mathfrak{M}, u) \mid (\mathfrak{M}[u]/\rho^{\mathfrak{M}}, [u]_{\rho^{\mathfrak{M}}}) \in \mathcal{Q}\}.$$

Membership in $\mathcal{Q}$ can therefore be determined via passage to canonical quotient representations, and in terms of the intersection of $\mathcal{Q}$ with the class $\mathcal{C}$ of all canonical quotient representations. Note that $\mathcal{C}$ consists of all finite rooted Kripke structures of the appropriate type in which each bisimulation type is realised exactly once (in other words, with identity as the largest bisimulation). As largest bisimulations and bisimulation quotients are polynomial time computable, it follows that $\mathcal{Q}$ is in Ptime if, and only if, $\mathcal{Q} \cap \mathcal{C}$ is. The following special property of $\mathcal{C}$ opens up a reduction to the case of linearly ordered structures, which then leads to the desired capturing result. By a canonical linear ordering of a structure we mean an ordering that is determined by the isomorphism type of that structure.

¹⁷One also has to require an effective link from syntax to Ptime algorithms for its evaluation, in order to avoid pathological solutions.

LEMMA 131. *There is a polynomial time algorithm which for every $(\mathfrak{M}, u) \in \mathcal{C}$ computes a canonical linear ordering of the domain.*

In fact, a linear ordering w.r.t. bisimulation type can be generated in an inductive refinement procedure which, in its n -th stage, produces a linear ordering of the $\rightleftharpoons_n$ -classes within any given finite Kripke structure. This is based on a lexicographic lift of the ordering on $\rightleftharpoons_n$ -classes to an ordering of the $\rightleftharpoons_{n+1}$ -classes, similar to the colour refinement technique in graph theory. Over any finite Kripke structure the common refinement of this process is a linear ordering of $\rightleftharpoons$ -classes; for structures in $\mathcal{C}$ one obtains an actual linear ordering, as each $\rightleftharpoons$ -class is inhabited by a single state.

Moreover, a representation of this linearly ordered version of the quotient structure $\mathfrak{M}[u]/\rho^{\mathfrak{M}}$ is uniformly LFP-definable over the given structures $(\mathfrak{M}, u)$ themselves. This means that in LFP over the $(\mathfrak{M}, u)$ one can also uniformly define any LFP definable property of their linearly ordered quotients $\mathfrak{M}[u]/\rho^{\mathfrak{M}}$. By the Immerman–Vardi result this includes all Ptime properties of these quotient structures, since they are linearly ordered. Together these observations yield an abstract capturing result: an effective syntactic normal form for the definition of precisely those bisimulation invariant properties that are in Ptime. As shown in [104] one can further isolate a natural extension of the modal μ -calculus, a multi-dimensional μ -calculus L_μ^ω , with the property that a class $\mathcal{Q}$ of finite pointed Kripke structures is bisimulation closed and in Ptime if, and only if, $\mathcal{Q}$ is the class of finite models of a formula $\varphi \in L_\mu^\omega$. The logic L_μ^ω is the natural bisimulation-safe least fixed-point extension of basic modal logic over the n -th cartesian power of a Kripke structure (intuitively: n -dimensional ML), for arbitrary $n \in \mathbb{N}$.

PROPOSITION 132. *Let $\mathcal{Q}$ be a class of finite pointed Kripke structures of fixed finite type. Then the following are equivalent:*

- (i) $\mathcal{Q}$ is bisimulation closed and in Ptime.
- (ii) $\mathcal{Q}$ is definable by a formula of the multi-dimensional μ -calculus L_μ^ω .

9.3 0-1 laws in modal logic

Another of the major specific topics in finite model theory is the asymptotic behaviour of the probability for a given property $\mathcal{P}$ to be true in a randomly chosen structure of size n (taken up to isomorphism), in a suitably defined probabilistic space. If that probability has a limit as n increases without bound, that limit is called the *(unlabelled) asymptotic probability of $\mathcal{P}$* .

A fundamental result in this area is the *0-1 law* for first-order logic, stating that the asymptotic probability for every first-order definable property of relational structures exists and equals either 0 or 1, i.e., every such property is either *almost surely true* or *almost surely false*. This result was first proved in [40] (using ‘almost sure’ quantifier elimination), later established independently by Fagin [27] who moreover obtained a purely logical characterisation of the set of first-order sentences that are almost surely true, as the first-order theory of the so-called *countable random structure*. Prior to Fagin’s discovery, Gaifman had studied in [37] infinite random structures as probabilistic models for arbitrary relational first-order languages and had proved that the first-order theory of such structures is axiomatised by an infinite set of *extension axioms*: sentences that require every n -tuple to be extendible to an $(n + 1)$ -tuple in every possible (i.e., consistent) way. Furthermore, he showed that the first-order theory of all extension

axioms is complete and ω -categorical.¹⁸ Thus, Fagin established the following *transfer theorem*, which immediately implies the 0-1 law: a first-order property of relational structures is almost surely true iff it is true in the (unique, up to isomorphism) countable random structure. Grandjean [63] proved that the complexity of checking if a given first-order formula is almost surely true is decidable in Pspace, in sharp contrast to Trachtenbrot's theorem that validity of first-order formulae on *all* finite structures is not even recursively axiomatisable.

The transfer theorem was subsequently extended and the 0-1 law proved for several extensions of first-order logic: for first-order logic with fixed point operators by Blass, Gurevich and Kozen, later subsumed by the 0-1 law for infinitary logic with finitely many variables $L_{\infty\omega}^{\omega}$, proved by Kolaitis and Vardi; for some prefix-defined fragments of monadic second-order logic, again by Kolaitis and Vardi, who also established curious parallel between decidability and 0-1 laws for such fragments. On the other hand, the 0-1 law fails in monadic second-order logic, even in its Σ_1^1 -fragment. For references and further details on these results, see, e.g., [64, 25, 93, 55].

In the framework of modal logic, there are two natural notions of (asymptotic) probability 'in the finite': with respect to Kripke structures and with respect to frames. The 0-1 law with respect to Kripke structures follows directly from Fagin's theorem. Moreover, Halpern and Kapron [66] showed that the modal formulae almost surely valid in finite Kripke structures are precisely the theorems of the non-normal *Carnap's logic* [8]. As for almost sure frame validity, a complete axiomatisation of the modal logic ML^r of the countable random frame has been obtained in [55], where it has also been proved that ML^r has the finite model property and is decidable. It is also shown there that not all modal formulae that are almost surely frame-valid are in ML^r , thus refuting the transfer theorem for frame validity in modal logic. Perhaps the simplest such formula, which fails in the countable random frame, is $\neg\Box\Box(p \leftrightarrow \neg\Diamond p)$, proven later in [90] to be almost surely true. Note that no such formula is frame-definable in fixed point logic LFP, or even in $L_{\infty\omega}^{\omega}$, because the transfer theorem does hold for these.

The failure of the transfer theorem for frame validity in modal logic cast a serious doubt on the truth of the 0-1 law there (claimed in [66]) which was soon justified by le Bars [90] who proved that the formula $\neg p \wedge q \wedge \Box\Box((p \vee q) \rightarrow \neg\Diamond(p \vee q)) \rightarrow \Diamond\Box\neg p$ has no asymptotic probability, by using involved combinatorial-probabilistic methods. Thus, basic modal logic provides the smallest currently known natural fragment of monadic Π_1^1 (resp. Σ_1^1), in a vocabulary with just a single binary relation, where the 0-1 law fails.

As noted in [55] the modal formulae which are almost surely frame-valid form a normal modal logic ML^{as} , which contains ML^r . It is a currently open problem whether ML^{as} is decidable, and its complete axiomatisation has not been established yet. However, a conjecture raised in [55] claims that all axioms that have to be added to ML^r in order to axiomatise ML^{as} are of a uniform, semantic nature, namely: there is an infinite collection $\mathcal{F}$ of special finite frames, and each $\mathfrak{F} \in \mathcal{F}$ determines an axiom $\varphi_{\mathfrak{F}}$ valid in 'almost every' finite frame¹⁹ iff that frame cannot be mapped by a bounded morphism onto $\mathfrak{F}$. For instance, the formula $\neg\Box\Box(p \leftrightarrow \neg\Diamond p)$ corresponds to the frame $\langle\{a, b\}, \{\langle a, b \rangle, \langle b, a \rangle, \langle b, b \rangle\}\rangle$.

¹⁸The probabilistic aspect of this result is rather curious: it means that, assuming uniform distribution, any randomly constructed countable relational structure is isomorphic with probability 1 to the countable random structure! In the case of graphs, that structure was previously known as the *Radó graph*.

¹⁹More precisely, in every finite frame in which each state is reachable from any other state by a path of length ≤ 2 .

CONCLUDING REMARKS

In summary, the semantics of modal logic has (at least) two emblematic features which have a crucial impact on its model theory and which we have attempted to reflect in the composition of this chapter.

Modal logic is local. Truth of a formula is evaluated at a current state (possible world); this localisation is preserved (and carried) along the edges of the accessibility relations by the restricted, relativised quantification corresponding to the modal operators. This feature is reflected by the notion of *bisimulation* between states and between Kripke structures, respectively. The notion of bisimulation invariance plays a key role in characterising what is modally definable, as captured in the van Benthem–Rosen theorem (Theorems 55 and 61 here). Moreover, bisimulation (and its game characterisation) plays a role in modal model theory analogous to that of partial isomorphism (and its Ehrenfeucht–Fraïssé characterisation) in classical model theory. From yet another perspective, the characteristic power of preservation under bisimulations in modal logic can be compared to the characteristic power of preservation under ultraproducts in first-order logic. Quite naturally, therefore, bisimulation emerges as the central and unifying truth-preserving model-theoretic construction in modal logic, and all other basic constructions on which the classical model theory of modal logic builds (generated substructures, bounded morphisms, disjoint unions) are definable in terms of it or at least closely related to it. By systematically developing the bisimulation-based approach to modal model theory in this chapter, we hope to have given a modern treatment on this classical theme. Furthermore, the central role of bisimulations and bisimulation invariance properties is so robustly preserved, *mutatis mutandis*, in the rich and diverse variety of extensions of basic modal logic, that it can be adopted as a benchmark of what constitutes a modal language.

Modal logic is multi-layered. On Kripke structures the modal language is a bounded variable, guarded fragment of first-order logic, while on Kripke frames, due to universal quantification over valuations, it becomes a fragment of universal monadic second-order logic. Each of these semantic layers leads to its own model-theoretic agenda and development, but the two interact closely through various model-theoretic constructions and preservation results presented here, and blend together in the notion of general frames, dually re-incarnated as modal algebras. General frames emerge as a third, intermediate semantic layer of modal logic, casting a bridge between the other two. In particular, by means of a hierarchy of persistency properties, general frames provide a yardstick to measure the ‘expressive complexity’ of modal formulae, and determine their model-theoretic behaviour. This chapter presents the basics of the modal model theory in each of these three layers and illustrates the use of the main tools and results arising in each one of them.

While trying to give a comprehensive account of the main issues and results of both classical and modern model theory of modal logic, we have not covered a number of important and relevant topics and research developments, either for lack of space or because they are adequately treated in other chapters of this handbook. A certainly incomplete list of the more conspicuous omissions (in no particular order) includes:

- model theory of extended modal languages: see [18] and [122] for a recent treatise;
- model theory of combined modal logics: see Chapter 15 of this handbook and [35];
- Lindström-type theorems for modal logic: see [19, 133];
- reductions of polyadic to monadic modal languages and their model theoretic implications, including transfer of properties: see [85, 41], and Chapter 8 of this handbook;
- Kracht’s internal definability theory [84];
- Zakharyashev’s canonical formulae, providing a uniform characterisation of normal modal logics extending K4: see [10, 9], and Chapter 7 of this handbook;
- model-building techniques such as mosaics and networks used for more advanced completeness and decidability proofs: see, e.g., [99] and [5, Ch. 6.4 and 7.4].
- model completions in modal logic [39];
- bisimulation quantifiers and their use for proving uniform interpolation of various modal logics by Visser [139], Ghilardi and Zawadowski [39] (where bisimulation quantifiers are related to model completions), and of the modal μ -calculus by D’Agostino and Hollenberg [15].

It is natural to conclude a handbook chapter by attempting to identify main general trends of the current and future development of the topic under consideration.

To begin with, let us recall and revisit van Benthem’s three ‘pillars of wisdom’ supporting the classical edifice of modal logic: the *Definability (Correspondence), Completeness, and Duality theories* [128]. Each of these has played a crucial role in the development of modal model theory, and will continue to play such a role, with an accordingly modernised and updated agenda.

In particular, analysing the expressive power of modal languages with respect to each of its semantic layers remains one of the main directions of research in modal logic, of growing importance and complexity, due to the active expansion and diversification of modal logic. Accordingly, the classical correspondence theory between modal and first-order logic, much of which has been reflected in the chapter, is gradually ramifying into a hierarchy of correspondence theories, aiming at mapping the variety of modal logics into the hierarchy of classical logical languages centered around first-order logic. An example is the currently emerging correspondence theory between modal logic and LFP.

Establishing completeness results of modal deductive systems designed to capture an intended semantics also remains one of the core areas of modal logic (as of logic in general) which requires increasingly sophisticated and powerful techniques to match the more and more complex modal languages and their semantics. The involved completeness proofs for the modal μ -calculus (see Chapter 12 of this handbook) and CTL* (see [111]), and the still open completeness problem for Parikh’s (full) Game Logic (see Chapter 20 of this handbook) are cases in point.

Likewise for decidability and complexity, where model-theoretic tools and techniques, such as the model-building techniques mentioned above as well as game-theoretic methods, are gaining increasing recognition and variety of applications.

New directions and problem areas in modal model theory itself, or using model-theoretic methods, are emerging, too. Many of them, such as finite model theory and descriptive complexity, finite and infinite state model checking, arise from actual or potential applications of modal logic to computer science and related fields and follow recent trends in classical model theory. Let us note, however, that while the present day model theory of modal logic is still using mainly results and techniques from the classical era of first-order model theory, the enormous development and sophistication of that field over the past decades is yet to make its full impact on modal model theory.

In closing, being aware that we cannot possibly offer a definitive treatment of such a rich and dynamic subject as the model theory of modal logic, we hope to have whetted readers' appetites and their desire to explore it further and to add to it new discoveries of their own.

Acknowledgements. We are indebted to Maarten de Rijke for his careful reading, critical remarks, corrections and numerous helpful suggestions for improvement. We also thank Achim Blumensath, Willem Conradie, Ian Hodkinson, Mark Reynolds, Balder ten Cate and Johan van Benthem for valuable comments or corrections. Most of the work of Valentin Goranko on this chapter was done while he worked at the Rand Afrikaans University (now, the University of Johannesburg), and he would like to acknowledge the financial support provided by that university and by the National Research Foundation of South Africa. Martin Otto was partially supported by a British EPSRC grant, originally held at the University of Wales Swansea, during a major period of his work on this chapter.

BIBLIOGRAPHY

- [1] H. Andreka, J. van Benthem, and I. Nemeti. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27:217–274, 1998.
- [2] A. Arnold and D. Niwinski. *Rudiments of μ -Calculus*. North-Holland, 2001.
- [3] C. Beeri, R. Fagin, D. Maier, and M. Yannakakis. On the desirability of acyclic database schemes. *Journal of the ACM*, 30:479–513, 1983.
- [4] C. Berge. *Graphs and Hypergraphs*. North-Holland, 1973.
- [5] P. Blackburn, M. de Rijke, and V. Venema. *Modal Logic*. Cambridge Univ. Press, 2001.
- [6] G. Boolos and G. Sambin. An incomplete system of modal logic. *Journal of Philosophical Logic*, 14:351–358, 1985.
- [7] R.A. Bull and K. Segerberg. Basic modal logic. In Gabbay and Guenther [33], pages 1–88.
- [8] R. Carnap. *Logical Foundations of Probability*. Univ. of Chicago Press, Chicago, 1950.
- [9] A. Chagrov, F. Wolter, and M. Zakharyashev. Advanced modal logic. In *Handbook of Philosophical Logic*, vol 3, pages 83–266. Kluwer Academic Publishers, second edition, 2001.
- [10] A. Chagrov and M. Zakharyashev. *Modal Logic*. Oxford University Press, Oxford, 1997.
- [11] L.A. Chagrova. An undecidable problem in correspondence theory. *Journal of Symbolic Logic*, 56:1261–1272, 1991.
- [12] C.C. Chang and H.J. Keisler. *Model Theory*. North-Holland, Amsterdam, 1973.
- [13] W. Conradie. Definability and changing perspectives. The Beth property for three extensions of modal logic. Technical Report MoL-2002-03, ILLC Publications, Master of Logic Series, 2002.
- [14] W. Conradie, V. Goranko, and D. Vakarelov. Elementary canonical formulae: a survey on syntactic, algorithmic, and model-theoretic aspects. In *M. Reynolds et al (eds.), Advances in Modal Logic*, volume 5, Kings College Publ., London, pages 17–51, 2005.
- [15] G. D'Agostino and M. Hollenberg. Logical questions concerning the μ -calculus: interpolation, Lyndon and Łoś-Tarski. *Journal of Symbolic Logic*, 65:310–332, 2000.
- [16] Giovanna D'Agostino. *Modal Logic and Non-Well-Founded Set Theory: Translation, Bisimulation, Interpolation*. PhD thesis, University of Amsterdam, 1998. ILLC Dissertation Series DS-1998-04.

- [17] A. Dawar and M. Otto. Modal characterisation theorems over special classes of frames. In *Proc. of LICS'2005, IEEE Press*, pages 21–30, 2005.
- [18] M. de Rijke. *Extending Modal Logic*. PhD thesis, ILLC, University of Amsterdam, 1993.
- [19] M. de Rijke. A Lindström theorem for modal logic. In Ponse et al. [109], pages 217–230.
- [20] M. de Rijke. Modal model theory. Technical Report CS-R9517, CWI, Amsterdam, 1995.
- [21] M. de Rijke. A note on graded modal logic. *Studia Logica*, 64(2):271–283, 2000.
- [22] M. de Rijke and H. Sturm. Global definability in basic modal logic. In H. Wansing, editor, *Essays on Non-Classical Logic*, pages 111–135. World Scientific Publishers, 2001.
- [23] K. Doets. *Basic Model Theory*. CSLI Publications, Stanford, California, 1996.
- [24] K. Doets and J. van Benthem. Higher-order logic. In D.M. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic, vol. 1*, pages 275–329. Reidel, Dordrecht, 1983.
- [25] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Perspectives in Mathematical Logic. Springer, Berlin, 1995.
- [26] H.-D. Ebbinghaus, J. Flum, and W. Thomas. *Mathematical Logic*. Springer, Berlin, 1984.
- [27] R. Fagin. Probabilities on finite models. *Journal of Symbolic Logic*, 41:50–58, 1976.
- [28] Kit Fine. Some connections between elementary and modal logic. In Kanger [81], pages 15–31.
- [29] M. Fischer and R. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 18:194–211, 1979.
- [30] Jörg Flum. On the (infinite) model theory of fixed-point logics. Preprint, Inst. für Math. Logik und Grundl. der Math., Albert-Ludwigs-Universität Freiburg, December 1995.
- [31] D. Gabbay. Selective filtration in modal logic. *Theoria*, 30:323–330, 1970.
- [32] D. Gabbay. *Investigations in Modal and Tense Logics with Applications to Problems in Philosophy and Linguistics*. Reidel, Dordrecht, 1976.
- [33] D.M. Gabbay and F. Guenther, editors. *Handbook of Philosophical Logic, vol. 2*. Reidel, Dordrecht, 1984.
- [34] D.M. Gabbay, I. Hodkinson, and M. Reynolds. *Temporal Logic: Mathematical Foundations and Computational Aspects*. Oxford University Press, 1994.
- [35] Dov Gabbay, Agi Kurucz, Frank Wolter, and Michael Zakharyashev. *Many-dimensional modal logics: theory and applications*. Studies in Logic, 148. Elsevier Science, 2003.
- [36] Dov Gabbay and Larisa Maksimova. *Interpolation and Definability. Modal and Intuitionistic Logic*. Oxford Logic Guides, 46. Oxford Univ. Press, 2005.
- [37] H. Gaifman. Concerning measures in first-order calculi. *Israel J. of Math.*, 2:1–18, 1964.
- [38] H. Gaifman. On local and nonlocal properties. In J. Stern, editor, *Logic Colloquium '81*, pages 105–135. North Holland, 1982.
- [39] S. Ghilardi and M. Zawadowski. *Sheaves, Games, and Model Completions*. Trends in Logic, 14. Kluwer, 2002.
- [40] Y.V. Glebskii, D.I. Kogan, M.I. Liogonki, and V.A. Talanov. Range and degree of realizability of formulas in the restricted predicate calculus. *Cybernetics*, 5:142–154, 1969.
- [41] G. Goguadze, C. Piazza, and Y. Venema. Simulating polyadic modal logics by monadic ones. *Journal of Symbolic Logic*, 68:419–462, 2003.
- [42] R. Goldblatt. *Logics of Time and Computation*. Number 7 in CSLI Lecture Notes. CSLI, Stanford, 2nd edition, 1992.
- [43] R. Goldblatt. *Mathematics of Modality*, volume 43 of *CSLI Lecture Notes*. Center for the Study of Language and Information, Stanford University, 1993.
- [44] R. Goldblatt. Saturation and the Hennessy-Milner property. In Ponse et al. [109], pages 107–129.
- [45] R. Goldblatt, I. Hodkinson, and Y. Venema. Erdős graphs resolve Fine's canonicity problem. *Bulletin of Symbolic Logic*, 10:186–208, 2004.
- [46] R.I. Goldblatt. First-order definability in modal logic. *Journal of Symbolic Logic*, 40:35–40, 1975.
- [47] R.I. Goldblatt. Metamathematics of modal logic I. *Reports on Mathematical Logic*, 6:41–78, 1976.
- [48] R.I. Goldblatt. Metamathematics of modal logic II. *Reports on Mathematical Logic*, 7:21–52, 1976.
- [49] R.I. Goldblatt. Varieties of complex algebras. *Annals of Pure and Applied Logic*, 38:173–241, 1989.
- [50] R.I. Goldblatt. Mathematical modal logic: A view of its evolution. *Journal of Applied Logic*, 1:309–392, 2003.
- [51] R.I. Goldblatt and S.K. Thomason. Axiomatic classes in propositional modal logic. In J. Crossley, editor, *Algebra and Logic*, pages 163–173. Springer, 1974.
- [52] V. Goranko. Axiomatizations with context rules of inference in modal logic. *Studia Logica*, 61:179–197, 1998.
- [53] V. Goranko and S. Passy. Using the universal modality: Gains and questions. *Journal of Logic and Computation*, 2:5–30, 1992.
- [54] V. Goranko and D. Vakarelov. Elementary canonical formulae. Extending Sahlqvist's theorem. *Annals of Pure and Applied Logic*, 2005. to appear.
- [55] Valentin Goranko and Bruce Kapron. The modal logic of the countable random frame. *Archive for Mathematical Logic*, 42:221–243, 2003.

- [56] E. Grädel. On the restraining power of guards. *Journal of Symbolic Logic*, 64:1719–1742, 1999.
- [57] E. Grädel. Why are modal logics so robustly decidable? *Bulletin of the EATCS*, 68:90–103, 1999.
- [58] E. Grädel, C. Hirsch, and M. Otto. Back and forth between guarded and modal logics. *ACM Transactions on Computational Logics*, 3(3):418 – 463, 2002.
- [59] E. Grädel, P. Kolaitis, and M. Vardi. On the Decision Problem for Two-Variable First-Order Logic. *Bulletin of Symbolic Logic*, 3:53–69, 1997.
- [60] E. Grädel and M. Otto. On Logics with Two Variables. *Theoretical Computer Science*, 224:73–113, 1999.
- [61] E. Grädel, M. Otto, and E. Rosen. Undecidability Results on Two-Variable Logics. *Archive for Mathematical Logic*, 38:213–354, 1999.
- [62] E. Grädel and I. Walukiewicz. Guarded Fixed Point Logic. In *Proceedings of 14th IEEE Symposium on Logic in Computer Science LICS '99, Trento*, pages 45–54, 1999.
- [63] E. Grandjean. Complexity of the first-order theory of almost all structures. *Information and Computation*, 52:180–204, 1983.
- [64] Y. Gurevich. Zero-one laws: The logic in computer science column. *Bulletin of the European Association for Theoretical Computer Science*, 46:90–106, 1992.
- [65] T. Hafer and W. Thomas. Computation tree logic CTL* and path quantifiers in the monadic theory of the binary tree. In *Proceedings of ICALP 1987*, pages 269–279, 1987.
- [66] J. Halpern and B. Kapron. Zero-one laws for modal logic. *Annals of Pure and Applied Logic*, 69:157–193, 1994.
- [67] B. Herwig and D. Lascar. Extending partial isomorphisms and the profinite topology on free groups. *Transactions of the AMS*, 352:1985–2021, 2000.
- [68] W. Hodges. *Model Theory*. Cambridge University Press, Cambridge, 1993.
- [69] I. Hodkinson. Loosely guarded fragment of first-order logic has the finite model property. *Studia Logica*, 70:205–240, 2002.
- [70] I. Hodkinson and M. Otto. Finite conformal hypergraph covers and Gaifman cliques in finite structures. *The Bulletin of Symbolic Logic*, 9:387–405, 2003.
- [71] M. Hollenberg. Safety for bisimulation in general modal logic. In *Proceedings 10th Amsterdam Colloquium*, pages 359–376, 1996.
- [72] M.J. Hollenberg. Hennessy-Milner classes and process algebra. In Ponse et al. [109], pages 187–216.
- [73] E. Hoogland. *Definability and Interpolation: Model-theoretic investigations*. PhD thesis, University of Amsterdam, 2001. ILLC Dissertation Series DS-2001-05.
- [74] G. Hughes. Every world can see a reflexive world. *Studia Logica*, 49:175–181, 1990.
- [75] G. Hughes and M. Cresswell. *A New Introduction to Modal Logic*. Routledge, London, 1968.
- [76] N. Immerman. Languages that capture complexity classes. *SIAM Journal of Computing*, 16:760–778, 1987.
- [77] D. Janin and I. Walukiewicz. On the expressive completeness of the propositional μ -calculus w.r.t. monadic second-order logic. In *Proceedings CONCUR '96*, pages 263–277, 1996.
- [78] B. Jónsson and A. Tarski. Boolean algebras with operators, Part I. *American Journal of Mathematics*, 73:891–939, 1952.
- [79] J. Jónsson and A. Tarski. Boolean algebras with operators, Part II. *American Journal of Mathematics*, 74:127–162, 1952.
- [80] H. Kamp. *Tense Logic and the Theory of Linear Order*. PhD thesis, University of California, Los Angeles, 1968.
- [81] S. Kanger, editor. *Proceedings of the Third Scandinavian Logic Symposium. Uppsala 1973.*, Amsterdam, 1975. North-Holland.
- [82] C.R. Karp. *Languages with Expressions of Infinite Length*. North-Holland, Amsterdam, 1964.
- [83] H.J. Keisler. *Model Theory for Infinitary Logic*. North-Holland, Amsterdam, 1971.
- [84] M. Kracht. *Tools and Techniques in Modal Logic*, volume 142 of *Studies in Logic*. Elsevier, Amsterdam, 1999.
- [85] Marcus Kracht and Frank Wolter. Simulation and transfer results in modal logic: a survey. *Studia Logica*, 59:149–177, 1997.
- [86] Saul Kripke. A completeness theorem in modal logic. *Journal of Symbolic Logic*, 24:1–14, 1959.
- [87] Saul Kripke. Semantical analysis of modal logic I. Normal modal propositional calculi. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 9:67–96, 1963.
- [88] N. Kurtonina and M. de Rijke. Expressiveness of concept expressions in first-order description logics. *Artificial Intelligence*, 107(2):303–333, 1999.
- [89] A.H. Lachlan. A note on Thomason's refined structures for tense logics. *Theoria*, 40:117–120, 1970.
- [90] J.-M. Le Bars. Zero-one law fails for frame satisfiability in propositional modal logic. In *Proceedings of LICS'2002*, pages 225–234. IEEE Computer Society Press, 2002.
- [91] E.J. Lemmon and D.S. Scott. *An Introduction to Modal Logic*. Blackwell, Oxford, 1977.
- [92] D. K. Lewis. Intensional logics without iterative axioms. *Journal of Philosophical Logic*, 3:457–466, 1974.

- [93] L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
- [94] A. Lisitsa and V. Sazonov. Linear ordering on graphs, anti-founded sets and polynomial time computability. *Theoretical Computer Science*, 224:173–213, 1999.
- [95] E.G.K. Lopez-Escobar. On defining well-orderings. *Fundamenta Mathematicae*, 59:299–300, 1966.
- [96] C. Lutz, U. Sattler, and F. Wolter. Modal logic and the two-variable fragment. In *Proceedings of 15th International Workshop on Computer Science Logic*, volume 2142 of *LNCS*, pages 247–261, 2001.
- [97] D.C. Makinson. A generalization of the concept of relational model. *Theoria*, pages 331–335, 1970.
- [98] L. Maksimova. Definability and interpolation in classical modal logics. *Contemporary Mathematics*, 131(3):583–599, 1992.
- [99] M. Marx. Complexity of modal logics of relations. Technical Report ML-97-02, Institute for Logic, Language and Computation, University of Amsterdam, May 1997.
- [100] M. Marx. Xpath and modal logics of finite DAG's. In *Automated Reasoning with Analytic Tableaux and Related Methods, LNCS 2796*, pages 150–164, 2003.
- [101] F. Moller and A. Rabinovich. On the expressive power of CTL*. In *Proceedings of the 14th IEEE Symposium on Logic in Computer Science*, pages 360–369, 1999.
- [102] M. Mortimer. On languages with two variables. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 21:135–140, 1975.
- [103] A. Nonnenhant and A. Szalas. A fixpoint approach to second-order quantifier elimination with applications to correspondence theory. In E. Orłowska, editor, *Logic at Work, Essays dedicated to the memory of Helena Rasiouwa*, pages 89–108. Springer Physica-Verlag, 1998.
- [104] M. Otto. Bisimulation-invariant Ptime and higher-dimensional mu-calculus. *Theoretical Computer Science*, 224:237–265, 1999.
- [105] M. Otto. Elementary proof of the van Benthem–Rosen characterisation theorem. Technical Report 2342, Fachbereich Mathematik, Technische Universität Darmstadt, 2004.
- [106] M. Otto. Modal and guarded characterisation theorems over finite transition systems. *Annals of Pure and Applied Logic*, 130:173–205, 2004.
- [107] Martin Otto. Bisimulation invariance and finite models. to appear in: *Lecture Notes in Logic, Logic Colloquium 02*, 2005.
- [108] B. Poizat. *A Course in Model Theory*. Universitext. Springer, 2000.
- [109] A. Ponse, M. de Rijke, and Y. Venema, editors. *Modal Logic and Process Algebra. A Bisimulation Perspective.*, volume 53 of *CSLI Lecture Notes*, Stanford University, 1995. CSLI Publications.
- [110] S. Radev. Infinitary propositional normal modal logic. *Studia Logica*, 46:291–309, 1987.
- [111] M. Reynolds. An axiomatization of full computation tree logic. *Journal of Symbolic Logic*, 66:1011–1057, 2001.
- [112] E. Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6:427–439, 1997.
- [113] H. Sahlqvist. Completeness and correspondence in the first and second order semantics for modal logic. In Kanger [81], pages 110–143.
- [114] G. Sambin and V. Vaccaro. Topology and duality in modal logic. *Annals of Pure and Applied Logic*, 37:249–296, 1988.
- [115] G. Sambin and V. Vaccaro. A topological proof of Sahlqvist's theorem. *Journal of Symbolic Logic*, 54:992–999, 1989.
- [116] D. Scott. A decision method for validity of sentences in two variables. *Journal of Symbolic Logic*, 27:477, 1962.
- [117] K. Segerberg. *An Essay in Classical Modal Logic*. Filosofiska Studier 13. University of Uppsala, 1971.
- [118] R. Streett and E. Emerson. An automata theoretic decision procedure for the propositional mu-calculus. *Information and Computation*, 81:249–264, 1989.
- [119] H. Sturm. *Modale Fragmente von $\mathcal{L}_{\omega\omega}$ und $\mathcal{L}_{\omega_1\omega}$* . PhD thesis, CIS, University of Munich, 1997.
- [120] H. Sturm. Interpolation and preservation in ML_{ω_1} . *Notre Dame Journal of Formal Logic*, 39:190–211, 1998.
- [121] B. ten Cate. Expressivity of second-order propositional modal logic. *Journal of Philosophical Logic*, 2005. to appear.
- [122] B. ten Cate. *Model theory for extended modal languages*. PhD thesis, University of Amsterdam, 2005. ILLC Dissertation Series DS-2005-01.
- [123] S.K. Thomason. An incompleteness theorem in modal logic. *Theoria*, 40:150–158, 1974.
- [124] S.K. Thomason. Reduction of second-order logic to modal logic. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 21:107–114, 1975.
- [125] J. van Benthem. *Modal Correspondence Theory*. PhD thesis, Mathematisch Instituut & Instituut voor Grondslagenonderzoek, University of Amsterdam, 1976.
- [126] J. van Benthem. Canonical modal logics and ultrafilter extensions. *Journal of Symbolic Logic*, 44:1–8, 1980.
- [127] J. van Benthem. *Modal Logic and Classical Logic*. Bibliopolis, Naples, 1983.
- [128] J. van Benthem. Correspondence theory. In Gabbay and Guenther [33], pages 167–247.

- [129] J. van Benthem. Notes on modal definability. *Notre Dame Journal of Formal Logic*, 39:20–39, 1989.
- [130] J. van Benthem. *Exploring Logical Dynamics*. Studies in Logic, Language and Information. CSLI Publications, Stanford, 1996.
- [131] J. van Benthem. Modality, bisimulation and interpolation in infinitary logic. Technical Report ML-97-06, Institute for Logic, Language and Computation, University of Amsterdam, 1997.
- [132] J. van Benthem. Minimal predicates, fixed points, and definability. *Journal of Symbolic Logic*, 70:696–712, 2005.
- [133] J. van Benthem. A new Lindström theorem for modal logic. preprint, 2005.
- [134] M. Vardi. The complexity of relational query languages. In *Proceedings of 14th ACM Symposium on Theory of Computing*, pages 137–146, 1982.
- [135] M. Vardi. Why is modal logic so robustly decidable? In N. Immerman and P. Kolaitis, editor, *Descriptive Complexity and Finite Models*, volume 31 of *DIMACS Series in Discrete Mathematics*, pages 149–184. AMS, 1997.
- [136] M. Vardi. Reasoning about the past with two-way automata. In *Proceedings of 25th Symposium on Automata, Language and Programming*, volume 1443 of *LNCS*, pages 628–641, 1998.
- [137] Y. Venema. Derivation rules as anti-axioms in modal logic. *Journal of Symbolic Logic*, 58:1003–1034, 1993.
- [138] Y. Venema. Modal definability, purely modal, 1999. in: J Gerbrandy et al. (editors), *Essays Dedicated to Johan van Benthem on the Occasion of his 50th Birthday*.
- [139] A. Visser. Bisimulations, model descriptions and propositional quantifiers. Technical Report 161, Department of Philosophy, Utrecht University, 1996.

INDEX

- 0-1 laws, 74
- back-and-forth, 18, 23, 25
- Beth property, 29
- bisimulation, 9
- bisimulation game, 16
- bisimulation invariance, 10, 34, 36, 72, 73
- bisimulation invariant Ptime, 73
- bisimulation quotient, 27
- bisimulation safety, 30
- bounded morphism, 11
- characteristic formulae, 17
- computation tree logic, 44
- constant formula, 4
- direct product, 50
- disjoint union, 13
- filtration, 19
- Fine–van Benthem theorem, 65
- finite bisimulation, 17
- finite model property, 19, 47, 71
- finite variable fragments, 31
- frame, 4
- full general frame, 56
- general frames, 56, 58, 59
- generated subframe/substructure, 11, 12
- Goldblatt–Thomason theorem, 65
- guarded bisimulation, 39
- guarded fragment, 37, 40
- Hennessy–Milner property, 23, 52
- Hennessy–Milner theorem, 22
- inductive formula, 68
- infinitary modal equivalence, 24
- infinitary modal logic, 46
- interpolation, 28
- inverse modalities, 33
- Karp’s theorem, 24
- Kripke frame, 4
- Kripke structure, 5, 56
- largest bisimulation, 25
- modal μ -calculus, 45
- modal characterisation theorem, 34, 36, 72
- modal definability, 7, 53, 63
- modal equivalence, 7, 16, 24
- modal theory, 6
- partial isomorphism, 23
- pebble games, 32
- persistent formula, 59
- pointed frame, 4
- pointed Kripke structure, 5
- propositional dynamic logic, 43
- regular formula, 68
- rooted subframe/substructure, 12
- Sahlqvist formula, 67, 68
- saturation, 23, 52
- shallow formula, 69
- standard translation, 6
- tree model property, 13, 21, 39
- ultrafilter extension, 49
- ultrafilter frame, 57
- ultraproduct, 51, 58
- unfolding/unravelling, 12
- van Benthem formula, 66
- van Benthem’s theorem, 34