Modal Characterisation Theorems over Special Classes of Frames

Anuj Dawar
University of Cambridge Computer Laboratory
15 J.J. Thomson Avenue
Cambridge, CB3 0FD, U.K.

Martin Otto
Fachbereich Mathematik
Technische Universität Darmstadt
64289 Darmstadt, Germany

Abstract

We investigate model theoretic characterisations of the expressive power of modal logics in terms of bisimulation invariance. The paradigmatic result of this kind is van Benthem's theorem which says that a first-order formula is invariant under bisimulation if, and only if, it is equivalent to a formula of basic modal logic. The present investigation primarily concerns ramifications for specific classes of structures. We study in particular model classes defined through conditions on the underlying frames, with a focus on frame classes that play a major role in modal correspondence theory and often correspond to typical application domains of modal logics. Classical model theoretic arguments do not apply to many of the most interesting classes – for instance, rooted connected frames, well-founded frames, finite rooted connected frames, finite transitive frames, finite equivalence frames – as these are not elementary. Instead we develop and extend the game-based analysis (first-order Ehrenfeucht-Fraïssé versus bisimulation games) over such classes and provide bisimulation preserving model constructions within these classes.

1. Introduction

Characterisation theorems are precise correspondences between semantic conditions imposed on logical formulae, typically in the form of closure under certain morphisms or equivalences, and syntactic restrictions. Such theorems have played an important role in model theory, often under the name of preservation theorems. However, the *preservation* property (namely that formulae of a particular syntactic form are preserved under a given semantic morphism) is usually much less significant than the corresponding *expressive completeness* property that any formula satisfying the semantic invariance condition is equivalent to one of the restricted syntactic form.

In the context of modal logics, commonly used for the specification of behaviours of reactive and concurrent systems, the most useful semantic invariance condition is *bisimulation*. Any modal formula is naturally preserved under bisimulations of Kripke structures. Important characterisation theorems have been established in this context, in particular that a sentence of first-order logic, interpreted on Kripke structures is invariant under bisimulations if, and only if, it is logically equivalent to a modal formula (van Benthem [10]) and that a sentence of monadic second-order logic is invariant under bisimulations if, and only if, it is logically equivalent to a formula of the μ-calculus (Janin and Walukiewicz [6]).

We are interested in the question of when such characterisations can be established for restricted classes of structures. For instance, for a given class $\mathcal{C}$, is it the case that any first-order formula that is bisimulation-invariant *on* $\mathcal{C}$ is equivalent *over* $\mathcal{C}$ to a modal formula. In general, such results are not obtained as consequences of the general characterisation theorem on the class of all structures. Van Benthem's theorem is established by classical methods of first-order model theory, relying heavily on the compactness of first-order logic. As a consequence, one obtains the characterisation theorem also for all classes $\mathcal{C}$ that are elementary. However, establishing the result for non-elementary classes requires different techniques.

The classes of structures which we explore are those that are of particular importance in modal correspondence theory. We are mainly interested in classes defined by conditions on the underlying frames, such as requiring the frames to be connected, symmetric or transitive, often combined with a condition of finiteness. While transitivity and symmetry are elementary conditions, they are non-elementary when we also require the structures to be finite. On the other hand, connectivity is non-elementary whether or not infinite structures are permitted. Rosen [9] established that the modal characterisation theorem holds even in restriction to finite structures. Again, this does not yield similar characterisation theorems for restricted classes of finite structures. Connectedness is a particularly interesting case as it is a natural condition to impose when we are interested in modal properties. As the truth of modal properties (or bisimulation-invariant properties in general) at a node v in a Kripke structure depends only on the nodes that are reachable from v, it is natural to consider *rooted* Kripke structures in which every node is reachable from the root so that we need not worry about first-order definable properties that depend on the existence of nodes that are not visible. However, Rosen's proof relies on the presence of such disjoint components. It is easily seen that the modal characterisation theorem fails when we restrict ourselves to finite connected structures. We show here that a variant can be recovered when we also allow global modalities.

In the presentation below, we emphasise the distinct methodologies used to prove the characterisation theorems. One set of techniques is based on locality properties of first-order logic and is used in Section 3 to establish characterisation results for classes of rooted connected frames (both finite and infinite) and classes of frames based on equivalence relations. Another set of techniques, based on decomposition methods is used in Section 4 to establish the modal characterisation theorem for finite tree-like and transitive frames. Both of these general classes of techniques could be useful in establishing similar results for other classes of structures.

2. Background and Methodology

2.1. Bisimulation and modal logics

Bisimulation. Bisimilarity is the fundamental notion of equivalence for the model theory of modal logics. When we regard Kripke structures as transition systems for the description of processes, bisimulation equivalence is the natural equivalence relation that captures behavioural equivalence. Bisimulation equivalence $\sim$ has an intuitive definition in terms of back-and-forth systems or games, which also shows the close analogy with Ehrenfeucht-Fraïssé games. In fact bisimulation equivalence $\sim$ and its finite approximations $\sim^\ell$ are precisely the modal variants of partial isomorphy and ℓ-partial isomorphy of the classical Ehrenfeucht-Fraïssé or pebble games. I.e., they describe strategies in infinite or ℓ-round model theoretic games whose rules reflect the semantics of modal quantification.

We consider Kripke structures (transition systems) over finite relational vocabularies with one or more binary accessibility relations (corresponding to labelled transitions) and finitely many unary predicates (corresponding to basic propositions). We mostly use letters E, E_i, R, R_i, etc. for the binary relations and P, P_i, etc. for the unary predicates. A Kripke structure of type $R_1, \ldots, R_m, P_1, \ldots, P_n$ is a relational structure $\mathfrak{A} = (A, R_1^{\mathfrak{A}}, \ldots, R_m^{\mathfrak{A}}, P_1^{\mathfrak{A}}, \ldots, P_m^{\mathfrak{A}})$, with $R_i^{\mathfrak{A}} \subseteq A \times A$ and $P_j^{\mathfrak{A}} \subseteq A$. The underlying *frame* is the $R_1, \ldots, R_m$ reduct $(A, R_1^{\mathfrak{A}}, \ldots, R_m^{\mathfrak{A}})$. By a *pointed Kripke structure*, $\mathfrak{A}, a$, we mean a structure $\mathfrak{A}$ with a distinguished element $a \in A$.

A *back-and-forth system* between two Kripke structures $\mathfrak{A} = (A, (R_i^{\mathfrak{A}}), (P_j^{\mathfrak{A}}))$ and $\mathfrak{B} = (B, (R_i^{\mathfrak{B}}), (P_j^{\mathfrak{B}}))$ of the same type is given by either a single non-empty subset $Z \subseteq A \times B$ or by a family (Z_k) of non-empty subsets $Z_k \subseteq A \times B$, whose elements are pairs $(a, b) \in A \times B$ that respect the P_j ($a \in P_j^{\mathfrak{A}}$ iff $b \in P_j^{\mathfrak{B}}$) and satisfy certain *back-and-forth conditions* which are characteristic of various kinds of bisimulation. In the following R can be one of the R_i or also a derived binary relation, e.g. the inverse R_i^{-1} of some R_i. For $Z, Z' \subseteq A \times B$, we say that Z' satisfies the *back-and-forth conditions* with respect to R for Z if

(*forth*) for any $(a, b) \in Z$ and any $a' \in A$ such that $(a, a') \in R^{\mathfrak{A}}$, there is some $b' \in B$ such that $(b, b') \in R^{\mathfrak{B}}$ and $(a', b') \in Z'$.

(*back*) for any $(a, b) \in Z$ and any $b' \in B$ such that $(b, b') \in R^{\mathfrak{B}}$, there is some $a' \in B$ such that $(a, a') \in R^{\mathfrak{A}}$ and $(a', b') \in Z'$.

If these conditions are satisfied for $Z' = Z$, we say that Z itself satisfies the back-and-forth conditions.

Definition 1. Let $\mathfrak{A}$ and $\mathfrak{B}$ be Kripke structures of type $R_1, \ldots, R_m, P_1, \ldots, P_n$. A *bisimulation* between $\mathfrak{A}$ and $\mathfrak{B}$ is a back-and-forth system $Z \subseteq A \times B$ satisfying the back-and-forth conditions w.r.t. each R_i.

A bisimulation is a *two-way bisimulation* if it also satisfies the back-and-forth conditions w.r.t. the inverses R_i^{-1}. A (forward or two-way) bisimulation Z between $\mathfrak{A}$ and $\mathfrak{B}$ is a *global bisimulation*, if it also satisfies the back-and-forth conditions w.r.t. the universal relations $U^{\mathfrak{A}} = A \times A$ and $U^{\mathfrak{B}} = B \times B$ [i.e., if $\pi_1(Z) = A$ and $\pi_2(Z) = B$].

We write $Z\colon \mathfrak{A} \sim \mathfrak{B}$ to indicate that Z is a bisimulation between $\mathfrak{A}$ and $\mathfrak{B}$; $Z\colon \mathfrak{A} \sim_{-} \mathfrak{B}$ to denote that Z is a two-way bisimulation; $Z\colon \mathfrak{A} \sim_{\forall} \mathfrak{B}$ for a global bisimulation; and $Z\colon \mathfrak{A} \approx \mathfrak{B}$ for a global two-way bisimulation.

The usual conventions as to distinguished nodes in pointed structures apply. For instance, we write $\mathfrak{A}, a \sim \mathfrak{B}, b$ if there is a bisimulation $Z\colon \mathfrak{A} \sim \mathfrak{B}$ with $(a, b) \in Z$.

Corresponding *bisimulation games* are played between two players, I and II, over the given structures $\mathfrak{A}$ and $\mathfrak{B}$, with one pebble marking one node in each structure. The rules of the game are such that in each round player I moves the pebble in one of the structures along an R-edge; player II has to respond by moving the other pebble along an R-edge in the opposite structure. Player II loses either if she cannot move in response to a move of player I or when her response leads to a pebbled pair (a, b) that does not respect the unary P_j; II wins the (infinite) game if she can maintain play indefinitely without losing. A bisimulation Z is then easily seen to correspond to a non-deterministic winning strategy for II in the (forward, two-way or global) bisimulation game (in which respective kinds of moves are admitted).

We turn to the back-and-forth systems corresponding to finite approximations associated with bisimulation games of ℓ rounds, for $\ell \in \mathbb{N}$. Here a winning strategy for player II corresponds to a back-and-forth system $(Z_k)_{0\leqslant k\leqslant \ell}$ where each Z_{k-1} satisfies the back-and-forth condition for Z_k of the respective kind (forward, two-way or global). We correspondingly write, for instance, $\mathfrak{A}, a \sim^\ell \mathfrak{B}, b$ if there is a back-and-forth system $(Z_k)_{0\leqslant k\leqslant \ell}$ for $\mathfrak{A}$ and $\mathfrak{B}$ w.r.t. the ordinary forward back-and-forth conditions, and such that $(a, b) \in Z_\ell$. Variants like $\mathfrak{A}, a \sim^\ell_\forall \mathfrak{B}, b$ or $\mathfrak{A}, a \approx^\ell \mathfrak{B}, b$ are similarly defined, with reference to the stronger back-and-forth conditions that govern those bisimulations.

Modal logics. We regard basic modal logic ML as a fragment of first-order logic FO, with restricted forms of quantification that capture the nature of modalities associated with the R_i. In vocabulary $R_1, \ldots, R_m, P_1, \ldots, P_n$, the formulae of basic modal logic ML are generated from atomic formulae $\varphi(x) = P_j x$ by Boolean connectives and quantification rules (here introduced via their first-order translations)

$$\begin{array}{rcl} ([R]\varphi)(x) & := & \forall y\big(Rxy \rightarrow \varphi(y)\big) \\ (\langle R\rangle\varphi)(x) & := & \exists y\big(Rxy \wedge \varphi(y)\big) \end{array}$$

for $R = R_1, \ldots, R_m$. In the case of one single accessibility relation R, one writes $\Box$ and $\Diamond$ instead of $[R]$ and $\langle R\rangle$.

Extensions of basic modal logic ML are obtained by allowing modalities corresponding to further (derived) accessibility relations, in particular w.r.t. the inverses of the given accessibility relation and w.r.t. the universal accessibility relation. These accessibility relations give rise to backward and global modalities, respectively.

Definition 2. We let ML denote basic modal logic with (forward) modalities for every given accessibility relation; $\mathrm{ML}^\forall$ stands for the extension of ML by the global modalities (unrestricted $\forall$ and $\exists$); ML^- stands for the extension with backward modalities for the inverses of the given accessibility relations; $\mathrm{ML}^{-\forall}$ is the combined extension by inverse modalities and the global modalities.

Modal *nesting depth* of a formula, which may be defined in the usual inductive manner, precisely corresponds to FO quantifier rank in the first-order translations. For $\ell \in N$ and a logic $\mathcal{L}$ like $\mathrm{ML}, \mathrm{ML}^\forall, \mathrm{ML}^{-\forall}$ we denote by $\mathcal{L}^\ell$ the fragment of $\mathcal{L}$ consisting of formulae of nesting depth up to ℓ. Equivalences $\equiv_\mathcal{L}$ and $\equiv_{\mathcal{L}^\ell}$ stand for $\mathcal{L}$-equivalence and $\mathcal{L}^\ell$-equivalence between structures. For instance, $\mathfrak{A}, a \equiv_\mathcal{L} \mathfrak{B}, b$ if for all $\varphi \in \mathcal{L}$, $\mathfrak{A}, a \models \varphi \Leftrightarrow \mathfrak{B}, b \models \varphi$.

For FO equivalence we write $\mathfrak{A}, a \equiv \mathfrak{B}, b$; and similarly $\mathfrak{A}, a \equiv_q \mathfrak{B}, b$ for FO equivalence up to quantifier rank q.

Ehrenfeucht-Fraïssé analysis of bisimulation. The following is a straightforward consequence of the Ehrenfeucht-Fraïssé style analysis of the back-and-forth games for various modal logics. For background on Ehrenfeucht-Fraïssé techniques see for instance [3] or [2]. For accounts specifically in connection with the modal case see for instance [1] and also [8] or [5].

Observation 3. *The following equivalences link bisimulation games to equivalence in modal logics. For any pair of pointed Kripke structures of the same finite type, $\mathfrak{A}, a$ and $\mathfrak{B}, b$, and for any $\ell \in \mathbb{N}$:*

$$\begin{array}{l} \mathfrak{A}, a \sim^\ell \mathfrak{B}, b \text{ iff } \mathfrak{A}, a \equiv_{\mathcal{L}^\ell} \mathfrak{B}, b \text{ for } \mathcal{L} = \mathrm{ML}, \\ \mathfrak{A}, a \sim^\ell_\forall \mathfrak{B}, b \text{ iff } \mathfrak{A}, a \equiv_{\mathcal{L}^\ell} \mathfrak{B}, b \text{ for } \mathcal{L} = \mathrm{ML}^\forall, \\ \mathfrak{A}, a \sim^\ell_- \mathfrak{B}, b \text{ iff } \mathfrak{A}, a \equiv_{\mathcal{L}^\ell} \mathfrak{B}, b \text{ for } \mathcal{L} = \mathrm{ML}^-, \\ \mathfrak{A}, a \approx^\ell \mathfrak{B}, b \text{ iff } \mathfrak{A}, a \equiv_{\mathcal{L}^\ell} \mathfrak{B}, b \text{ for } \mathcal{L} = \mathrm{ML}^{-\forall}. \end{array}$$

Each one of these equivalence relations $\equiv_{\mathcal{L}^\ell}$ has finite index over any class of pointed Kripke structures of fixed finite type, whence each equivalence class is definable by a single formula of $\mathcal{L}^\ell$.

At the level of full rather than finite bisimulation, like $\sim$ rather than $\sim^\ell$, the correspondence is with equivalence in infinitary variants $\mathcal{L}_\infty$ of the corresponding logic $\mathcal{L}$ which have conjunctions and disjunctions over arbitrary sets of formulae.

If we write $\leftrightarrows$ for the back-and-forth equivalence corresponding to the infinite game, $\leftrightarrows^\ell$ for the finite approximation corresponding to the ℓ-round game, then $\equiv_{\mathcal{L}^\ell}$ is captured by $\leftrightarrows^\ell$. It follows that $\equiv_\mathcal{L}$ is captured by the common refinement $\leftrightarrows^\omega := \bigcap_\ell \leftrightarrows^\ell$, while $\leftrightarrows$ itself captures $\equiv_{\mathcal{L}_\infty}$. For classical model theoretic methods it is essential that over suitably saturated structures (e.g. ω-saturated structures suffice, but weaker notions of saturation can be used in the modal setting), $\leftrightarrows^\omega$ coincides with $\leftrightarrows$, and $\equiv_\mathcal{L}$ coincides with $\equiv_{\mathcal{L}_\infty}$. For classical logic this is Karp's theorem on partial isomorphy and equivalence in $\mathrm{L}_{\infty\omega}$.

2.2. Modal characterisation theorems

The characterisation theorems we consider here establish a correspondence between bisimulation invariance and expressibility in certain modal logics, over certain classes of Kripke structures. Without any restriction of the class of Kripke structures, and for ordinary bisimulation and basic modal logic, van Benthem's Theorem [10] is the paradigmatic result of this kind.

Theorem 4 (van Benthem). *The following are equivalent for $\varphi = \varphi(x) \in$ FO:*

(a) *φ is bisimulation invariant: $\mathfrak{A}, a \sim \mathfrak{B}, b$ implies that $\mathfrak{A}, a \models \varphi$ iff $\mathfrak{B}, b \models \varphi$.*

(b) *φ is logically equivalent to a formula $\tilde{\varphi} \in$ ML.*

In the sequel we use notation like $\mathrm{FO}/\!\sim\ \equiv \mathrm{ML}$ to indicate a correspondence as expressed in the theorem, and say that ML *captures bisimulation invariant first-order logic*. All our characterisation theorems are of the form

$$\mathrm{FO}/\!\leftrightarrows\ \equiv \mathcal{L} \ \text{ over } \mathcal{C},$$

where $\mathcal{C}$ is a class of Kripke structures. This means that the following are equivalent for $\varphi(x) \in \mathrm{FO}$:

(a) φ is $\leftrightarrows$ invariant in restriction to $\mathcal{C}$.

(b) φ is logically equivalent to a formula $\tilde{\varphi} \in \mathcal{L}$ in restriction to $\mathcal{C}$.

The analogue of van Benthem's theorem in finite model theory, $\mathrm{FO}/\!\sim\ \equiv \mathrm{ML}$ over the class of all finite Kripke structures, is due to Rosen [9]. For variations on its proof and ramifications for stronger variants of bisimulation equivalence see also [7] and [8].

Our main theorems provide analogues of these characterisations, both classical and finite model theory, for important classes of frames – or over classes of (finite) Kripke structures over restricted classes of frames; for instance, the following main results from sections 3 and 4, respectively:

$$\mathrm{FO}/\!\sim\ \equiv \mathrm{ML}^{\forall} \begin{cases} \text{over (finite) rooted frames: Thm 11.} \\ \text{over (finite) equivalence frames: Thm 17.} \end{cases}$$

$$\mathrm{FO}/\!\sim\ \equiv \mathrm{ML} \begin{cases} \text{over finite trees: Thm 28.} \\ \text{over (finite) transitive frames: Thm 29.} \end{cases}$$

2.3. Model theoretic methods

For the following discussion let again $\leftrightarrows$ be a back-and-forth equivalence whose approximations $\leftrightarrows^{\ell}$ have finite index and capture levels of equivalence $\equiv_{\mathcal{L}^{\ell}}$ in a fragment $\mathcal{L} = \bigcup_{\ell} \mathcal{L}^{\ell}$ of first-order logic:

(i) $\mathcal{L}^{\ell}$ is $\leftrightarrows^{\ell}$ invariant: $\mathfrak{A}, a \leftrightarrows^{\ell} \mathfrak{B}, b$ implies that $\mathfrak{A}, a \models \varphi$ iff $\mathfrak{B}, b \models \varphi$, for all $\varphi \in \mathcal{L}^{\ell}$.

(ii) each $\leftrightarrows^{\ell}$ class is definable by a formula of $\mathcal{L}^{\ell}$.

Classical methods. The classical first-order proofs of characterisation theorems essentially involve an indirect argument with application of compactness, along the following lines (compare [8, 5] for a more comprehensive discussion).

Suppose $\varphi(x) \in \mathrm{FO}$ is $\leftrightarrows$ invariant but not expressible in the fragment $\mathcal{L} \subseteq \mathrm{FO}$. By compactness one finds structures $\mathfrak{A}, a \equiv_{\mathcal{L}} \mathfrak{B}, b$ which are distinguished by φ. Passing to ω-saturated elementary extensions of $\mathfrak{A}$ and $\mathfrak{B}$ (by compactness), one finds $\mathfrak{A}^*, a \leftrightarrows \mathfrak{B}^*, b$ but still distinguished by φ; but this is impossible by $\leftrightarrows$ invariance of φ.

This classical proof naturally relativises to elementary classes, for instance to the classes of Kripke structures over symmetric (and/or reflexive, and/or transitive) frames. An alternative approach for some of these can be based on model theoretic reductions via interpretations. This method also applies for instance to some interesting variations of the Janin-Walukiewicz characterisation at the level of monadic second-order. For restrictions to classes of finite structures, however, none of these techniques seems to be available.

A models-for-games method. In contrast to the classical methods, and similar to the finite model theory proofs in [9, 8], we proceed along an orthogonal direction to establish a characterisation result $\mathrm{FO}/\!\leftrightarrows\ \equiv \mathcal{L}$ in restriction to a class of structures $\mathcal{C}$ which in particular need not be first-order definable. This approach is based on a gradation of elementary equivalence $\equiv$ into successively refined approximations $\overset{n}{\equiv}$ for $n \in \mathbb{N}$, such that

(iii) every first-order formula is preserved by $\overset{n}{\equiv}$ for all sufficiently large n.

(iv) there is a transformation $F\colon \mathcal{C} \to \mathcal{C}$ compatible with $\leftrightarrows$, in the sense that $F(\mathfrak{A}, a) \leftrightarrows \mathfrak{A}, a$, and such that for some function $f\colon \mathbb{N} \to \mathbb{N}$:

$$\mathfrak{A}, a \leftrightarrows^{f(n)} \mathfrak{B}, b \ \Rightarrow\ F(\mathfrak{A}, a) \overset{n}{\equiv} F(\mathfrak{B}, b).$$

A characterisation theorem is immediate under these conditions.

Proposition 5. *Let $\leftrightarrows$ and $\leftrightarrows^{\ell}$, as well as $\mathcal{L}$ and $\mathcal{L}^{\ell}$ be as described in (i) and (ii). If suitable $\overset{n}{\equiv}$ satisfy (iii) and (iv), it follows that* $\mathrm{FO}/\!\leftrightarrows\ \equiv \mathcal{L}$ *over* $\mathcal{C}$.

Indeed, let $\varphi(x) \in \mathrm{FO}$ be $\leftrightarrows$-invariant over $\mathcal{C}$. Choose n such that φ is preserved under $\overset{n}{\equiv}$. For structures in $\mathcal{C}$, $\mathfrak{A}, a \leftrightarrows^{f(n)} \mathfrak{B}, b$ implies that $\mathfrak{A}, a \leftrightarrows F(\mathfrak{A}, a) \overset{n}{\equiv} F(\mathfrak{B}, b) \leftrightarrows \mathfrak{B}, b$, and hence $\mathfrak{A}, a \models \varphi$ iff $\mathfrak{B}, b \models \varphi$. So φ is preserved under $\leftrightarrows^{f(n)}$ and equivalent over $\mathcal{C}$ to a disjunction of formulae according to (ii).

It should be stressed that this alternative approach applies equally well in the classical case, where it sheds new light on van Benthem's characterisation, in particular as far its robustness under relativisation is concerned. As discussed in [7, 5], the alternative proofs, for instance of $\mathrm{FO}/\!\sim\ \equiv \mathrm{ML}$, automatically relativise to any bisimulation closed class of frames, or even to any class of (finite) frames that is closed under the particular model construction involved in the transformation F (certain locally acyclic covers in the case of [8], cf. section 3 here). Examples to be treated elsewhere comprise the case of (finite) strongly connected frames ("cyclic return"), or frames of fixed or bounded in- and/or out-degrees.

In some cases, $\overset{n}{\equiv}$ may be taken to be just $\equiv_n$, FO-equivalence up to quantifier rank n, in other cases a finer gradation based on quantifier rank *and* Gaifman locality parameters is shown to be appropriate. We organise the rest of the paper by these methodological distinctions. Section 3 deals with applications of the models-for-games technique that involve an upgrading of levels of bisimulation equivalence to levels of first-order equivalence based on *locality criteria*. These techniques in particular cover connectivity constraints; they do not cover transitivity constraints as transitivity trivialises locality. However, we show that this technique can also be employed for the important class of multi-modal frames of equivalence relations, where it works at the level of the accessibility pattern between classes. Apart from giving a characterisation theorem for multi-modal $S5$, this application yields bisimulation equivalent companion structures for finite multi-$S5$ frames that may be of independent interest e.g. in the analysis of such systems for knowledge representation.

Section 4 develops the models-for-games technique with an emphasis on *decomposition arguments* w.r.t. the first-order Ehrenfeucht-Fraïssé games over suitable classes of frames, most notably for tree-like structures. These techniques are applied to cover the cases of transitive frames and some variations. They also yield further insights into the relationship between first-order equivalence and bisimulation equivalence over the important class of transitive trees.

3. Locality based techniques

In this section we explore model constructions that allow us to upgrade finite bisimulation equivalence to finite local first-order equivalence. These techniques build on methods from [8]; their variations obtained here find their key application in characterisation theorems for classes of (finite) rooted connected frames. With a further generalisation of these techniques we obtain a characterisation theorem over (finite) multi-modal $S5$ frames (equivalence frames).

For all results of this section we upgrade finite approximations of the respective bisimulation equivalence to suitable levels of *local* approximations to elementary equivalence. Locality refers to Gaifman locality and Gaifman distance, compare [4, 2].

The *Gaifman graph* of a relational structure $\mathfrak{A}$ is an undirected graph with vertex set A and an edge between a and a' if $a \neq a'$ are among the components of some tuple $\mathbf{a}$ in one of the relations of $\mathfrak{A}$. For a Kripke structure $\mathfrak{A}$, its Gaifman graph is just the combined symmetrisation of the binary accessibility relations. *Gaifman distance* in $\mathfrak{A}$ is the natural graph theoretic distance in the Gaifman graph of $\mathfrak{A}$.

We write $N^{\ell}(a)$ for the *ℓ-neighbourhood* of a node a in $\mathfrak{A}$ consisting of all those elements at Gaifman distance up to ℓ from a. A subset of $\mathfrak{A}$ is *ℓ-scattered* if the ℓ-neighbourhoods of any two distinct members of this set are disjoint. An ℓ-scattered subset for $\psi(y)$ is an ℓ-scattered subset whose members each satisfy ψ in their ℓ-neighbourhoods. The desired local approximations to elementary equivalence are the following equivalences $\equiv^{(\ell)}_{q,n}$:
$\mathfrak{A}, a \equiv^{(\ell)}_{q,n} \mathfrak{B}, b$ if

- $\mathfrak{A}\restriction N^{\ell}(a), a \equiv_q \mathfrak{B}\restriction N^{\ell}(b), b$, i.e., a and b are indistinguishable in their respective ℓ-neighbourhoods by FO-formulae of quantifier rank q,
- $\mathfrak{A}$ and $\mathfrak{B}$ realise exactly the same quantifier rank q formulae in k-scattered sets of size m for $k \leqslant \ell$ and $m \leqslant n$. I.e., for any $\psi(x) \in$ FO of quantifier rank q and any $m \leqslant n$, $k \leqslant \ell$: $\mathfrak{A}$ has a k-scattered subset of size m for ψ if and only if $\mathfrak{B}$ has.

It is a consequence of Gaifman's theorem [4, 2] that every first-order formula $\varphi(x)$ in a single free variable is invariant under $\equiv^{(\ell)}_{q,n}$ for some ℓ, q, n. Indeed, $\equiv^{(\ell)}_{q,n}$ is precisely designed so as to capture indistinguishability by means of formulae in Gaifman local form that involve constituent k-local formulae for $k \leqslant \ell$ of quantifier rank up to q and quantifications over scattered sets of size up to n.

On the modal side, we seek to adapt the construction of (finite) bisimilar coverings from [8] to the restricted classes of frames we are interested in here. We review the basic notions and quote the key result from [8].

Definition 6. A homomorphism $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ is a *bisimilar cover* of $\mathfrak{A}$ by $\hat{\mathfrak{A}}$ if its graph is a global two-way bisimulation. A bisimilar cover $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ is *faithful* in a node $\hat{a} \in \hat{\mathfrak{A}}$ if the incidence degrees of $\hat{a}$ with each R and R^{-1} in $\hat{\mathfrak{A}}$ are the same as those at $\pi(a)$ in $\mathfrak{A}$, for every binary R; π is faithful if it is faithful in all $\hat{a} \in \hat{\mathfrak{A}}$.

Note that faithful bisimilar covers $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ provide unique lifts of $\mathfrak{A}$ to paths in $\hat{\mathfrak{A}}$. If $a_0, \ldots, a_m$ is a path in $\mathfrak{A}$, and if $\hat{a}_0 \in \pi^{-1}(a_0)$, then there is a unique lift $\hat{a}_0, \ldots, \hat{a}_m$ to a path at $\hat{a}_0$ in $\hat{\mathfrak{A}}$. The same applies to *undirected paths* (paths in the Gaifman graph) which may traverse edges backwards as well as forwards.

A cycle in the Gaifman graph of a Kripke structure $\mathfrak{A}$ is an (undirected) cycle in $\mathfrak{A}$; we also regard as (undirected) cycles in $\mathfrak{A}$ any loop $(a, a) \in R_i^{\mathfrak{A}}$ or cycles of length 2 generated by multiple (backward or forward) edges.

For instance two edges like $(a, a') \in R_i^{\mathfrak{A}}$ and $(a, a') \in R_j^{\mathfrak{A}}$ for $j \neq i$, or also $(a, a'), (a', a) \in R_i^{\mathfrak{A}}$ form cycles of length 2. Acyclicity in this sense therefore excludes loops as well as multiple or inverse edges.

Definition 7. A structure $\mathfrak{A}$ is *k-acyclic* if all k-neighbourhoods of $\mathfrak{A}$ are acyclic; equivalently, if $\mathfrak{A}$ has no undirected cycles of length up to $2k$.

The following is the key result of [8].

Theorem 8. *For all k: for every finite $\mathfrak{A}$ there is a faithful bisimilar cover $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ of $\mathfrak{A}$ by a finite k-acyclic simple transition system $\hat{\mathfrak{A}}$.*

3.1. Rooted Frames

Rooted (and hence connected) frames are arguably the most commonly intended models in applications of modal analysis. The possibility that a Kripke structure (transition system) may have unreachable components "in a different universe," would often seem to be a pathology of the model. Interestingly it is also responsible for the following phenomenon.

Observation 9. *The $\sim$-class of a finite Kripke structure is in general not definable by a sentence of* ML *or of* $\mathrm{ML}^{\forall}$.

In fact, because of its lack of global quantification, ML is too weak to make sure that, for the suitable $\ell \in \mathbb{N}$, $\sim$ coincides with $\sim^{\ell}$ over the given structure; but so is $\mathrm{ML}^{\forall}$, because it offers the wrong mode of universal quantification in a potentially disconnected structure. Consider, for instance, the one-element frame with an R-loop in its single node, and an interpretation of all unary predicates P as empty (false in the only node). Then $\sim$ coincides with $\sim^{0}$ over this structure. The bisimulation class of this structure can be characterised by saying that $\neg P$ and $\Diamond\neg P$ hold at its root and at all nodes reachable from this. Bisimilar (finite or infinite) companion structures, however, may have reachable nodes at arbitrarily large distance from the root and also nodes not reachable from the root of arbitrary modal behaviour. Any ML formula of nesting depth q that is true in the given structure fails to enforce the correct behaviour in reachable nodes at distances greater than q; any $\mathrm{ML}^{\forall}$-formula either suffers from that same defect or it wrongly stipulates certain modal behaviours also in nodes not reachable from the root.

Intuitively, one would like to work not in the full bisimulation class, but rather in its restriction to pointed structures in which all nodes are reachable from the distinguished node. In this setting, however, one is forced to look for new characterisation theorems – and non-classical techniques for proving them, as reachability is not first-order. Moreover, in the setting of these intended frames and structures $\sim$ and $\sim_{\forall}$ coincide. Hence $\mathrm{ML}^{\forall}$ rather than ML becomes the right candidate for capturing bisimulation invariant FO-properties.

Definition 10. A pointed Kripke structure $\mathfrak{A}, a$ and its underlying pointed frame are called *rooted* if every node of A is reachable on a directed path from the root a.

Theorem 11. *Over the classes of all rooted frames, and of all finite rooted frames, respectively, bisimulation invariance of a first-order formula is captured by* $\mathrm{ML}^{\forall}$*:* $\mathrm{FO}/{\sim} = \mathrm{FO}/{\sim_{\forall}} \equiv \mathrm{ML}^{\forall}$ *over (finite) rooted frames.*

For the rest of this section we work in restriction to finite frames, but point out that all arguments equally apply in the case of infinite frames although simpler arguments and model constructions (especially unravellings into infinite trees) would usually be available in that case. For simplicity of exposition we also assume just one rather than a finite number of binary relations, even though all results remain valid in the multi-modal case. So all frames are now of the form $\mathfrak{A} = (A, R^{\mathfrak{A}})$, structures are expansions of such frames by (a finite number of) unary predicates.

From forward to two-way bisimilarity. With regard to the relationship between ordinary (forward) and two-way bisimulation, the following adaptation of a lemma from [8] is necessary in the context of (finite) rooted structures. It shows that also within this class, $\sim_{\forall}^{2\ell}$ can be upgraded to $\approx^{\ell}$ in $\sim_{\forall}$ equivalent companion structures. The proof is omitted from the present extended abstract.

Lemma 12. *Let $\mathfrak{A}, a$ and $\mathfrak{B}, b$ be finite and rooted, $\mathfrak{A}, a \sim_{\forall}^{2\ell} \mathfrak{B}, b$. Then there are finite, rooted $\hat{\mathfrak{A}}, \hat{a}$ and $\hat{\mathfrak{B}}, \hat{b}$ with $\hat{\mathfrak{A}}, \hat{a} \sim_{\forall} \mathfrak{A}, a$ and $\hat{\mathfrak{B}}, \hat{b} \sim_{\forall} \mathfrak{B}, b$ such that*

(i) $\hat{\mathfrak{A}}, \hat{a} \approx^{\ell} \hat{\mathfrak{B}}, \hat{b}$.

(ii) $\hat{\mathfrak{A}}{\restriction}N^{\ell}(\hat{a}), \hat{a}$ *and* $\hat{\mathfrak{B}}{\restriction}N^{\ell}(\hat{b}), \hat{b}$ *are directed rooted trees.*

From bisimilarity to FO equivalence. The key to upgrading a finite level of $\approx$-equivalence to local first-order equivalence lies in the construction of locally acyclic bisimilar companion structures as guaranteed by Theorem 8. The case of rooted structures, however, does require some extra care. In fact, a locally acyclic bisimilar cover by a rooted structure is not always available. Consider the example of a structure $(\{0,1,2\}, \{(0,1),(1,2),(0,2)\})$. Any lift of the undirected cycle $0,1,2,0$ from the root $\hat{0}$ of a bisimilar cover that is 2-acyclic would have to end in a node $\tilde{0} \neq \hat{0}$ that is two-way bisimilar to 0. In particular $\tilde{0}$ must have in-degree zero, whence it cannot be reachable from the root $\hat{0}$; so the cover cannot be rooted. However, we are able to show the following, whose proof is omitted for lack of space.

Lemma 13. *Let $k \geqslant 3$, $\mathfrak{A}, a$ be finite, rooted, and such that $\mathfrak{A}{\restriction}N^{k}(a)$ is acyclic. Then there is a bisimilar cover $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ by some k-acyclic finite, rooted $\hat{\mathfrak{A}}, \hat{a}$. π can be chosen such that $\pi^{-1}(a) = \{\hat{a}\}$, and such that π is faithful in all nodes apart from $\hat{a}$, where incidence degrees in $\hat{a}$ are a fixed positive multiple of those in a.*

Suppose $\mathfrak{A}$ and $\mathfrak{B}$ are ℓ-acyclic and $\mathfrak{A} \approx^{\ell} \mathfrak{B}$. Then the acyclic substructures induced on ℓ-neighbourhoods of nodes in $\mathfrak{A}$ and $\mathfrak{B}$ respectively, correspond modulo $\approx^{\ell}$. In order that $\approx^{\ell}$ equivalence guarantees local $\equiv_q$ equivalence of acyclic ℓ-neighbourhoods it suffices to boost or match multiplicities in these tree-like structures in such a way that the second player in the q-round FO Ehrenfeucht-Fraïssé game can always find matching elements on fresh paths where the first player does.

Towards a formal definition consider the following strengthening of $\approx^{\ell}$ to $\approx^{\ell;q}$ involving *counting up to q*. $(Z_k)_{0 \leqslant k \leqslant \ell}\colon \mathfrak{A} \approx^{\ell;q} \mathfrak{B}$ if $(Z_k)_{0 \leqslant k \leqslant \ell}\colon \mathfrak{A} \approx^{\ell} \mathfrak{B}$ is a back-and-forth system for which also satisfies the following stronger back-and-forth conditions, for $(a,b) \in Z_k$, $k > 0$ and any $m \leqslant q$:

(*forth*) for any distinct $a_1, \dots, a_m$ such that $(a, a_i) \in R^{\mathfrak{A}}$ there are distinct $b_1, \dots, b_m$ such that $(b, b_i) \in R^{\mathfrak{B}}$ and $(a_i, b_i) \in Z_{k-1}$; similarly for $(a_i, a) \in R^{\mathfrak{A}}$.

(*back*) for any distinct $b_1, \dots, b_m$ such that $(b, b_i) \in R^{\mathfrak{B}}$ there are distinct $a_1, \dots, a_m$ such that $(a, a_i) \in R^{\mathfrak{A}}$ and $(a_i, b_i) \in Z_{k-1}$; similarly for $(a_i, a) \in R^{\mathfrak{A}}$.

Non-global variants $\sim_{-}^{\ell;q}$ are analogously defined. The proof of the following is an exercise in Ehrenfeucht-Fraïssé games (compare Claim 26 in [8]).

Lemma 14. *If $\mathfrak{A}, a$ and $\mathfrak{B}, b$ are acyclic, then $\mathfrak{A}, a \sim_{-}^{\ell;q} \mathfrak{B}, b$ implies $\mathfrak{A} \restriction N^{\ell}(a), a \equiv_q \mathfrak{B} \restriction N^{\ell}(b), b$. For ℓ-acyclic $\mathfrak{A}, a$ and $\mathfrak{B}, b$ it follows that $\mathfrak{A}, a \approx^{\ell;q} \mathfrak{B}, b$ implies that $\mathfrak{A}, a \equiv_{q,1}^{(\ell)} \mathfrak{B}, b$.*

To achieve $\approx^{\ell;q}$ equivalence rather than just $\approx^{\ell}$ through Lemma 13, it suffices to boost all relevant multiplicities to at least q prior to application of Lemma 13.

Consider finite rooted $\mathfrak{A}, a$ and $\mathfrak{B}, b$, with a and b of in-degree 0. Let $\mathfrak{A} \otimes q$ be the natural two-way globally bisimilar companion with universe $A \times \{0, \dots, q-1\}$ with an R-edge from (a_1, i) to (a_2, j) for all $(a_1, a_2) \in R^{\mathfrak{A}}$. If $\mathfrak{A}, a \approx^{\ell} \mathfrak{B}, b$, then $\mathfrak{A} \otimes q, (a, 0) \approx^{\ell;q} \mathfrak{B} \otimes q, (b, 0)$. However, $\mathfrak{A} \otimes q, (a,0)$ and $\mathfrak{B} \otimes q, (b,0)$ are not rooted, and the ℓ-neighbourhoods of their roots are not acyclic anymore, even if they were in the original structures. Some simple surgery can solve this problem. We present the following lemma without proof.

Lemma 15. *Let $\mathfrak{A}, a$ and $\mathfrak{B}, b$ be finite, rooted, and such that $\mathfrak{A} \restriction N^{\ell}(a), a$ and $\mathfrak{B} \restriction N^{\ell}(b), b$ are directed rooted trees. Suppose $\mathfrak{A}, a \approx^{2\ell+1} \mathfrak{B}, b$. Then there are finite, rooted $\hat{\mathfrak{A}}, \hat{a}$ and $\hat{\mathfrak{B}}, \hat{b}$ such that $\hat{\mathfrak{A}}, \hat{a} \approx \mathfrak{A}, a$, $\hat{\mathfrak{B}}, \hat{b} \approx \mathfrak{B}, b$ and*
(i) $\hat{\mathfrak{A}}, \hat{a} \approx^{\ell;q} \hat{\mathfrak{B}}, \hat{b}$.
(ii) *$\hat{\mathfrak{A}} \restriction N^{\ell}(\hat{a}), \hat{a}$ and $\hat{\mathfrak{B}} \restriction N^{\ell}(\hat{b}), \hat{b}$ are directed rooted trees.*

Combining Lemmas 12, 15 and 13, by applying them in this order to finite and rooted $\mathfrak{A}, a$ and $\mathfrak{B}, b$ with $\mathfrak{A}, a \sim_{\forall}^{4\ell+2} \mathfrak{B}, b$, we obtain, for any $q \in \mathbb{N}$, a pair of finite, rooted ℓ-acyclic companion structures for which $\mathfrak{A} \sim_{\forall} \hat{\mathfrak{A}}, \hat{a} \approx^{\ell;q} \hat{\mathfrak{B}}, \hat{b} \sim_{\forall} \mathfrak{B}, b$. By Lemma 14 therefore also $\mathfrak{A}, a \equiv_{q,1}^{(\ell)} \mathfrak{B}, b$. In order to achieve $\equiv_{q,n}^{(\ell)}$ rather than just $\equiv_{q,1}^{(\ell)}$, we apply one further upgrading step that saturates $\hat{\mathfrak{A}}$ and $\hat{\mathfrak{B}}$ w.r.t. to scattered sets.

Lemma 16. *Let $\mathfrak{A}, a$ and $\mathfrak{B}, b$ finite, rooted with $\mathfrak{A}, a \sim_{\forall}^{8\ell+6} \mathfrak{B}, b$. Then for any $q, n \in \mathbb{N}$ there are finite, rooted $\hat{\mathfrak{A}}, \hat{a} \sim_{\forall} \mathfrak{A}, a$ and $\hat{\mathfrak{B}}, \hat{b} \sim_{\forall} \mathfrak{B}, b$ such that $\hat{\mathfrak{A}}, \hat{a} \equiv_{q,n}^{(\ell)} \hat{\mathfrak{B}}, \hat{b}$.*

Proof. Choose $Q \geqslant q$ such that for all (ψ, k, m), where ψ is of quantifier rank q, $k \leqslant \ell$ and $m \leqslant n$, the existence of a k-scattered set consisting of m nodes within distance $\ell + 1$ of x is expressed by an FO formula of quantifier rank Q [this is a $2\ell + 1$-local property of x].

We assume w.l.o.g. that $\mathfrak{A}, a$ and $\mathfrak{B}, b$ themselves are $\approx^{2\ell+1;Q}$ equivalent, and both $(2\ell+1)$-acyclic (by Lemmas 12, 15 and 13). Let $\hat{\mathfrak{A}}$ be the result of gluing n copies of $\mathfrak{A}, a$ in a to form a new root $\hat{a}$; similarly for $\hat{\mathfrak{B}}, \hat{b}$.

It follows from Lemma 14 that also $\hat{\mathfrak{A}}, \hat{a} \equiv_{Q,1}^{(2\ell+1)} \hat{\mathfrak{B}}, \hat{b}$. We claim that $\hat{\mathfrak{A}}, \hat{a} \equiv_{q,n}^{(\ell)} \hat{\mathfrak{B}}, \hat{b}$.

Assume for instance that there is some k-scattered set of size $m \leqslant n$ for a quantifier rank q formula ψ in $\mathfrak{A}$, where $k \leqslant \ell$. Either this set, together with the k-neighbourhoods of all its members lies within $N^{2\ell+1}(\hat{a})$ or there is at least one witness for $\psi(x)$ outside $N^{\ell}(\hat{a})$. In the first case, $\hat{\mathfrak{B}}$ has a corresponding scattered set in $N^{2\ell+1}(\hat{b})$, since $\hat{\mathfrak{A}} \restriction N^{2\ell+1}(\hat{a}), \hat{a} \equiv_Q \hat{\mathfrak{B}} \restriction N^{2\ell+1}(\hat{b}), \hat{b}$. In the second case, let a' be such a witness in $\hat{\mathfrak{A}}$, $a' \notin N^{\ell}(\hat{a})$ and $\hat{\mathfrak{A}} \restriction N^k(a') \models \psi[a']$. If b' in $\hat{\mathfrak{B}}$ is such that $\hat{\mathfrak{A}}, a' \approx^{\ell+1} \hat{\mathfrak{B}}, b'$, then b' will also be outside $N^{\ell}(\hat{b})$, since every element of $N^{\ell}(\hat{b})$ admits a backward path of length up to ℓ into the unique element of in-degree zero. This implies that the isomorphic copies of this witness b' from the n copies of $\mathfrak{B}$ in $\hat{\mathfrak{B}}$ form a k-scattered set of n many points realising the same $\sim_{-}^{\ell;q}$-type as a'. By Lemma 14 these therefore also form a k-scattered witness set for ψ. $\square$

In the light of Proposition 5, this finishes the proof of Theorem 11. One can similarly obtain a characterisation of FO$/\sim_{-}$ over (finite) rooted frames, where FO$/\sim_{-}$ = FO$/\approx$ $\equiv$ $\mathrm{ML}^{-\forall}$. This does not require Lemma 12, but in the remaining upgrading steps one also needs to cover the case of roots of positive in-degree.

All results and proof methods so far apply equally in a multi-modal setting with several binary relations rather than just one. One multi-modal setting of particular interest, for instance in reasoning about knowledge, is considered in the following section.

3.2. Equivalence frames

We consider Kripke structures over *equivalence frames*, i.e., $\mathfrak{A} = (A, E_1, \dots, E_m, P_1, \dots, P_n)$ with binary E_i and unary P_j such that each E_i is an equivalence relation over A. We call such structures or frames *equivalence structures/frames*. By symmetry of the accessibility relations, $\sim_{-}$ and $\approx$ coincide with $\sim$ and $\sim_{\forall}$, respectively. We prove a characterisation theorem for $\sim_{\forall}$. One could similarly obtain corresponding characterisations for $\sim$. Equally, the present characterisation could be combined with the methods from the previous section to give a characterisation in terms of $\sim$ over the subclass of connected equivalence frames.

Theorem 17. $\mathrm{FO}/{\sim_\forall} = \mathrm{FO}/{\approx} \equiv \mathrm{ML}^\forall$ *over (finite) equivalence structures.*

As equivalence frames form an elementary class, the classical version of this characterisation result is obtained by direct relativisation of van Benthem's theorem. We therefore concentrate on the finite case, and only consider finite equivalence structures in the remainder of this section.

Again, we work with locality. As equivalence relations trivialise locality within equivalence classes, locality can only be used to analyse the intersection pattern between classes w.r.t. different equivalence relations. If any two equivalence classes w.r.t. distinct relations can at most intersect in a single node (see definition of *simple* frames below), then a modal quantification in restriction to any one equivalence class is either global quantification within this class or can be eliminated, depending on whether the modality and equivalence class under consideration concern the same E_i or not.

Acyclicity criteria also only apply at the level of equivalence classes, as any equivalence class is a clique, whence short cycles within individual classes cannot be ruled out.

Definition 18. An equivalence structure $\mathfrak{A}$ is called *simple* if any two equivalence classes w.r.t. distinct E_i and E_j intersect in at most one element (in other words: no multiple edges). For $\ell \geqslant 3$, $\mathfrak{A}$ is called *ℓ-acyclic* if it has no non-trivial cycles of length up to ℓ: if $a_0 \neq a_1 \neq \ldots \neq a_{n-1} \neq a_n = a_0$ where $(a_j, a_{j+1}) \in E_{i_j}$ is a cycle such that $E_{i_j} \neq E_{i_{j+1}}$, then $n > \ell$.

Proposition 19. *Simple k-acyclic covers are available for finite equivalence frames. Any finite equivalence structure $\mathfrak{A}$ possesses a bisimilar cover $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ by a finite, simple, k-acyclic equivalence structure $\hat{\mathfrak{A}}$.*

The cover constructed for the proof of this proposition is such that each equivalence class – apart from its intersection pattern with other classes – is an isomorphic copy of an equivalence class of $\mathfrak{A}$.

Towards an upgrading to local FO equivalence consider the following two observations. The operation $\mathfrak{A} \mapsto \mathfrak{A} \otimes q$ is as discussed in connection with Lemma 15 above. The class of finite equivalence structures is closed under this operation, which boosts multiplicities by introducing q indistinguishable copies of every node. Similarly, it is closed under the operation $\mathfrak{A} \mapsto n \times \mathfrak{A}$ of extending $\mathfrak{A}$ by $n-1$ disjoint isomorphic copies of itself, for $n \geqslant 1$.

Observation 20. *Let $\mathfrak{A}, a$ and $\mathfrak{B}, b$ be finite equivalence structures, $\ell, q, n \in \mathbb{N}$.*

(i) $\mathfrak{A}, a \approx^\ell \mathfrak{B}, b \;\Rightarrow\; (\mathfrak{A} \otimes q), (a, 0) \approx^{\ell;q} (\mathfrak{B} \otimes q), (b, 0)$.

(ii) $\mathfrak{A}, a \equiv^{(\ell)}_{q,1} \mathfrak{B}, b \;\Rightarrow\; n \times \mathfrak{A}, a \equiv^{(\ell)}_{q,n} n \times \mathfrak{B}, b$.

Proposition 19, Lemma 14 and these observation yield the following. With Proposition 5, this proves the theorem.

Lemma 21. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be finite equivalence structures with $\mathfrak{A}, a \sim^\ell_\forall \mathfrak{B}, b$. Then for any $q, n \in \mathbb{N}$ there are finite equivalence structures $\hat{\mathfrak{A}}, \hat{a} \approx \mathfrak{A}, a$ and $\hat{\mathfrak{B}}, \hat{b} \approx \mathfrak{B}, b$ such that $\hat{\mathfrak{A}}, \hat{a} \equiv^{(\ell)}_{q,n} \hat{\mathfrak{B}}, \hat{b}$.*

Proof. We reason up to global bisimulation equivalence $\approx$. Passing from $\mathfrak{A}$ and $\mathfrak{B}$ first to $\mathfrak{A} \otimes q$ and $\mathfrak{B} \otimes q$, we may assume that $\mathfrak{A}, a \sim^{\ell;q}_\forall \mathfrak{B}, b$. With Proposition 19 we obtain ℓ-acyclic simple companions for which one directly checks that still $\hat{\mathfrak{A}}, \hat{a} \sim^{\ell;q}_\forall \hat{\mathfrak{B}}, \hat{b}$. So we now assume that $\mathfrak{A}$ and $\mathfrak{B}$ are simple and ℓ-acyclic. It follows from Lemma 14 that $\mathfrak{A}, a \equiv^{(\ell)}_{q,1} \mathfrak{B}, b$. Now we may replace each structure by n disjoint copies to obtain $\mathfrak{A}, a \equiv^{(\ell)}_{q,n} \mathfrak{B}, b$. ☐

4. Decomposition based techniques

We now aim to establish the modal characterisation result for the class of finite transitive frames. That is, we concentrate on transition systems with one accessibility relation R which is transitive. Such a relation (or its reflexive closure) can be thought of as a pre-order and we freely use the notation and terminology of order below. We proceed in a series of steps, beginning by considering finite linear orders and extending to ever wider classes of structures.

Linear Orders. We will need the following useful lemma about structures in which the accessibility relation is a linear order. The statement is a pumping lemma established by standard methods using Ehrenfeucht-Fraïssé games (see [2]).

Lemma 22. *For every vocabulary σ and $r \geqslant 0$, there is an N such that if $\mathfrak{A}$ is a σ-word then it has a substructure $\mathfrak{B} \subseteq \mathfrak{A}$ of length at most N such that $\mathfrak{A} \equiv_r \mathfrak{B}$.*

The number N is a function of r and of k, the number of unary relations in the vocabulary σ. We write $N(r, k)$ when we need to make these parameters explicit.

Note that Lemma 22 remains true if the structures interpreted a strict order $<$ instead of the order $\leqslant$.

Trees A *tree* is a structure in which the unique accessibility relation R is a partial order (i.e. it is transitive, reflexive and anti-symmetric) and, in addition, is tree-like.

Definition 23. *A transitive connected frame (A, R) is* tree-like *if for any distinct $a, b \in A$ which are incomparable (i.e. neither $R(a, b)$ nor $R(b, a)$ holds) there is no c with $R(a, c)$ and $R(b, c)$.*

In the following, we use the symbol $\preccurlyeq$ for the relation R on a tree. We aim to use Lemma 22 to establish a result about the equivalence of trees. To be precise, we aim to characterise the equivalence relation $\equiv_r$ on trees in terms of the relations $\equiv_{r-1}$ on subtrees along with an equivalence on suitably defined words.

Fix $r > 0$ and a vocabulary σ consisting of a binary symbol $\preccurlyeq$ and some collection of unary symbols. Let $\varphi_1, \dots, \varphi_f$ be an enumeration of all first-order σ-sentences (up to equivalence) of quantifier rank $r-1$ or less. Let σ_r be the vocabulary consisting of a binary relation symbol $\leqslant$ and unary symbols $P_1, \dots, P_f$. Let $\mathfrak{A}$ be a σ-structure on the frame $(A, \preccurlyeq, \alpha)$ which is a tree with root α. For any element $a \in A$, we define a σ_r-word $L(\mathfrak{A}, a)$ as follows:

- the universe of $L(\mathfrak{A}, a)$ consists of the set $W = \{b \in A \colon \alpha \preccurlyeq b \preccurlyeq a\}$.
- $b_1 \leqslant b_2$ in $L(\mathfrak{A}, a)$ if, and only if, $b_1 \preccurlyeq b_2$.
- $P_i(b)$ holds if, and only if, the sentence φ_i is true in the substructure of $\mathfrak{A}$ induced by the set $\{b' \in A \colon b \preccurlyeq b'$ and $b' \not\succcurlyeq b''$ for any $b \prec b'' \preccurlyeq a\}$.

We can think of $L_r(\mathfrak{A}, a)$ as the linearly ordered set of elements that are on the unique maximal path from α (the root of $\mathfrak{A}$) to a. Each node b on this path is coloured by the set of sentences of quantifier rank $r-1$ that are true in the subtree rooted at b from which we have excluded those elements on the path to a. $L_r(\mathfrak{A}, a)$ is intended to be a representation of the structure $\mathfrak{A}$ as a linearly ordered sequence of subtrees which are classified only up to equivalence in $\equiv_{r-1}$. This representation is not uniquely determined by $\mathfrak{A}$ but is parametrised by the choice of the element a. The aim of the construction is to establish that the representation captures enough information to determine the $\equiv_r$ class of $\mathfrak{A}$. To be precise, we are able to prove the following lemma.

Lemma 24. *If $\mathfrak{A}$ and $\mathfrak{B}$ are two σ-trees such that:*

(1) *for every $a \in \mathfrak{A}$, there is a $b \in \mathfrak{B}$ such that $L_r(\mathfrak{A}, a) \equiv_{r-1} L_r(\mathfrak{B}, b)$; and*

(2) *for every $b \in \mathfrak{B}$, there is a $a \in \mathfrak{A}$ such that $L_r(\mathfrak{A}, a) \equiv_{r-1} L_r(\mathfrak{B}, b)$,*

then $\mathfrak{A} \equiv_r \mathfrak{B}$.

Once again, we can observe that Lemma 24 could equally well be stated for structures in which the accessibility relation is a *strict* tree-like order.

Saturated Trees Our next goal is to define, for any structure built on a tree frame, a saturated companion which is bisimilar to the original structure but will enable us to upgrade $\sim^l$ to $\equiv_r$ for suitable l and r.

Let $T = (A, R, \alpha)$ be a tree with root α and r be a natural number. We define $s_r(T)$, the *r-saturated companion* of T to be a rooted tree $(W, \preccurlyeq, \lambda)$ defined by

- $W = \{w \in (A \times \{0, \dots, r-1\})^+ \colon w = (a_1, l_1) \cdots (a_n, l_n)$ with $a_i \neq a_j, a_1 = \alpha, l_1 = 0$ and $R(a_i, a_j)$ for all $i < j\}$;
- $w_1 \preccurlyeq w_2$ if, and only if, w_1 is a prefix of w_2;
- the root λ is $(\alpha, 0)$.

$w \prec v$ abbreviates $w \preccurlyeq v \wedge w \neq v$.

If $\mathfrak{A}$ is a structure on the frame T, let $s_r(\mathfrak{A})$ denote the structure on the frame $s_r(T)$ obtained by taking, for each unary predicate symbol P, $P((a_1, l_1) \cdots (a_n, l_n))$ if, and only if, $P^{\mathfrak{A}}(a_n)$. It is easily verified that $\mathfrak{A} \sim s_r(\mathfrak{A})$. Indeed, the relation that relates an element $a \in A$ to every sequence $w \in W$ ending in (a, l) for some l is easily seen to be a bisimulation that, in particular, relates α to λ. By extension, we call $s_r(\mathfrak{A})$ the r-saturated companion of $\mathfrak{A}$.

The key properties of the r-saturated companions are summarised in Lemma 26 below. In order to introduce this, we first establish some notation. Given a structure $s_r(\mathfrak{A})$ on the frame $(W, \preccurlyeq, \lambda)$ and an element $w \in W$ we write $t(w)$ for the subtree rooted at w, i.e. the substructure of $s_r(\mathfrak{A})$ induced by the set of elements $\{v \in W \colon w \preccurlyeq v\}$. We also write $\mathfrak{A}_a$ for the subtree of the structure $\mathfrak{A}$ rooted at a. Given a sequence p of elements $w_1 \prec \cdots \prec w_n$, we write $t_{-p}(w_i)$ for the substructure of $s_r(\mathfrak{A})$ induced by the set of elements $\{v \in W \colon w_i \preccurlyeq v$ and $u \not\preccurlyeq v$ for any $w_i \prec u \preccurlyeq w_{i+1}\}$. That is, $t_{-p}(w_i)$ is the subtree rooted at w_i obtained by excluding all elements that appear on the branch containing the path p. We are now ready to state the key properties of r-saturated companions.

Observation 25. *For $r_1 > r_2$, $s_{r_1}(\mathfrak{A}) \equiv_{r_2} s_{r_2}(\mathfrak{A})$.*

The proof of the following lemma is omitted for reasons of space.

Lemma 26. *For $r > 0$, in $s_r(\mathfrak{A})$ the r-saturated companion of $\mathfrak{A}$:*

(1) *for each v, w, if $w \preccurlyeq v$, then there are r distinct elements $v_1, \dots, v_r$ such that $w \preccurlyeq v_i$ for each i; $v_i \not\preccurlyeq v_j$ for $i \neq j$; and $t(v_i) \cong t(v_j)$.*

(2) *for any path $p = w_1 \prec \cdots \prec w_n$, there is a path $p' = w'_1 \prec \cdots \prec w'_n$ with the property that for each i, $t_{-p}(w_i) \cong t_{-p'}(w'_i)$ and there is no v strictly between w'_i and w'_{i+1}.*

(3) *for any path $p = w_1 \prec \cdots \prec w_n$, if $w_i = (a_1, l_1) \cdots (a_k, l_k)$ then $t_{-p}(w_i) \sim s_{r-1}(\mathfrak{A}_{a_k})$ and $t_{-p}(w_i) \equiv_{r-1} s_{r-1}(\mathfrak{A}_{a_k})$.*

Our aim is to establish the following claim.

Claim 27. *For every vocabulary σ and each r, there is an ℓ such that if $\mathfrak{A}$ and $\mathfrak{B}$ are two structures with tree frames such that $\mathfrak{A} \sim^{\ell} \mathfrak{B}$ then $s_r(\mathfrak{A}) \equiv_r s_r(\mathfrak{B})$.*

This immediately yields the modal characterisation theorem for the class of finite trees.

Theorem 28. $\mathrm{FO}/{\sim} \equiv \mathrm{ML}$ *over the class of all finite trees.*

We note once again that nothing in the proof depends on the assumption that the accessibility relation in trees is reflexive. The argument goes through unchanged if the partial order is strict.

Transitive Structures If a structure $\mathfrak{A}$ is on a frame that is irreflexive and transitive then it is necessarily bisimilar to a tree (with a strict order) by a simple unravelling. Thus, Theorem 28 extends naturally to well-founded, transitive frames (so-called Löb frames).

We wish now briefly to consider transitive frames which are not necessarily acyclic. It is straightforward to see that any transitive structure is bisimilar to one on a tree-like frame, i.e. one which satisfies Definition 23.

Define the equivalence relation $\simeq$ on a frame (A, R) by $a \simeq b$ if $R(a, b)$ and $R(b, a)$. Any tree-like frame can be seen as a tree of $\simeq$ equivalence classes. We use this observation to interpret any tree-like structure in a tree.

Fix a vocabulary σ consisting of the binary relation symbol R and a set $P_1, \ldots, P_k$ of unary relation symbols. Consider a transition system $\mathfrak{A}$ over this vocabulary with a tree-like frame. We associate with $\mathfrak{A}$ a tree $T(\mathfrak{A})$ over a vocabulary $T(\sigma)$ consisting of R and a unary relation symbol S for each set $S \subseteq \{P_1, \ldots, P_k\}$. The nodes of $T(\mathfrak{A})$ are the equivalence classes induced by $\simeq$ on $\mathfrak{A}$. Writing $[a]$ for the equivalence class of a, we have $R^{T(\mathfrak{A})}([a], [b])$ if, and only if, $R^{\mathfrak{A}}(a, b)$. This is easily seen to be well defined due to the transitivity of $R^{\mathfrak{A}}$. We also let $S([a]) \Leftrightarrow \exists\, b(b \simeq a \wedge S = \{P_i \colon P_i(b)\})$. That is, S holds at $[a]$ if there is some element b in the same equivalence class as a for which S is exactly the set of unary predicates that hold.

There is one point which needs care. If $[a]$ is an equivalence class that contains more than one element, or if $R(a, a)$ in $\mathfrak{A}$, then we will have $R([a], [a])$ in $T(\mathfrak{A})$, otherwise $R([a], [a])$ will not hold. This means that the tree $T(\mathfrak{A})$ is neither reflexive nor irreflexive in the sense we have used above. This does not pose a serious problem to the application of the above construction. There is a standard translation that interprets loops $R([a], [a])$ by means of an additional unary relation. This translation preserves first-order equivalence but it does not preserve bisimulation. However, if we define r-saturated structures to also have loops only on those points w which consist of paths ending at (a, l) where $R(a, a)$, all stages in the construction still work.

We define a reverse translation from trees to tree-like transitive relations. Take any structure $\mathfrak{A}$ on the vocabulary $T(\sigma)$ such that (A, R) is a tree and define $\tau(\mathfrak{A})$ to be the tree-like structure in the vocabulary σ obtained from $\mathfrak{A}$ by expanding each element a to an equivalence class (under $\simeq$) with one element a_S for each unary S such that $S(a)$ holds and with $P(a_S) \Leftrightarrow P \in S$.

The following observations are straightforward.

(1) $\mathfrak{A} \sim \mathfrak{B}$ if, and only if, $T(\mathfrak{A}) \sim T(\mathfrak{B})$.

(2) $\tau(T(\mathfrak{A})) \sim \mathfrak{A}$.

(3) If $\mathfrak{A} \equiv_{r+1} \mathfrak{B}$, then $T(\mathfrak{A}) \equiv_r T(\mathfrak{B})$.

(4) If $\mathfrak{A} \equiv_{r+1} \mathfrak{B}$, then $\tau(\mathfrak{A}) \equiv_r \tau(\mathfrak{B})$.

These enable us to transfer the characterisation theorem from trees to arbitrary transitive frames.

Theorem 29. $\mathrm{FO}/{\sim} \;\equiv\; \mathrm{ML}$ *over the class of all (finite) transitive frames.*

Since all constructions in above proof preserve reachability, the following is also easily established.

Theorem 30. $\mathrm{FO}/{\sim} \;\equiv\; \mathrm{ML}$ *over the class of all (finite) rooted transitive frames.*

5. Conclusion

For classes $\mathcal{C} \subseteq \mathcal{C}'$ of structures, a characterisation theorem on $\mathcal{C}'$ does not necessarily yield a similar theorem on $\mathcal{C}$ as this involves a weakening of both the hypothesis and conclusion of the expressive completeness (the difficult direction of the characterisation). Whether or not the characterisation transfers from $\mathcal{C}'$ to $\mathcal{C}$ depends on other factors such as whether $\mathcal{C}$ is closed under the model constructions employed in the proof as well as under the equivalences one is attempting to characterise. For this reason we have focussed on the novelty of the methodologies and constructions we use in establishing characterisations of modal logic on various classes of (especially finite) structures of interest.

References

[1] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.

[2] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 2nd edition, 1999.

[3] H.-D. Ebbinghaus, J. Flum, and W. Thomas. *Mathematical Logic*. Springer, 1994.

[4] H. Gaifman. On local and nonlocal properties. In J. Stern, editor, *Logic Colloquium '81*, pages 105–135. North Holland, 1982.

[5] V. Goranko and M. Otto. Model Theory of Modal Logic. In preparation for: Handbook of Modal Logic, 2005.

[6] D. Janin and I. Walukiewicz. On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic. In *Proceedings of 7th International Conference on Concurrency Theory CONCUR '96*, number 1119 in Lecture Notes in Computer Science, pages 263–277. Springer-Verlag, 1996.

[7] M. Otto. Elementary proof of the van Benthem-Rosen characterisation theorem. Technical Report 2342, Department of Mathematics, Darmstadt University of Technology, 2004.

[8] M. Otto. Modal and guarded characterisation theorems over finite transition systems. *Annals of Pure and Applied Logic*, 130:173–205, 2004.

[9] E. Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6:427–439, 1997.

[10] J. van Benthem. *Modal Logic and Classical Logic*. Bibliopolis, Napoli, 1983.