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

February 27, 2009

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 frames, finite rooted frames, finite transitive frames, well-founded 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. Over most of the classes considered we obtain finite model theory analogues of the classically expected characterisations, with new proofs also for the classical setting. The class of transitive frames is a notable exception, with a marked difference between the classical and the finite model theory of bisimulation invariant first-order properties. Over the class of all finite transitive frames in particular, we find that monadic second-order logic is no more expressive than first-order as far as bisimulation invariant properties are concerned – though both are more expressive here than basic modal logic. We obtain ramifications of the de Jongh–Sambin theorem and a new and specific analogue of the Janin–Walukiewicz characterisation of bisimulation invariant monadic second-order for finite transitive frames.

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* invariance. Any modal formula is naturally preserved under bisimulations of Kripke structures. Important characterisation theorems have been established in this context, in particular that a formula of first-order logic in a single free variable, interpreted on pointed Kripke structures, is invariant under bisimulations if, and only if, it is logically equivalent to a modal formula (van Benthem [17]) and that a formula 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 [10]).

We are interested in corresponding characterisations of modal fragments of first-order logic over 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 rooted (i.e. where all states are reachable from the root), symmetric or transitive, often combined with a condition of finiteness. While transitivity and symmetry are elementary conditions, the class of all *finite* frames is non-elementary, and so is, for instance, the class of all finite transitive frames. On the other hand, rootedness is non-elementary whether or not infinite structures are permitted. Rosen [14] 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. Rootedness 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 rooted 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 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 modal characterisation theorems 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.

For finite transitive frames, and for what we here call finite transitive tree-like frames (which can represent any finite transitive frame up to bisimulation) a rather interesting phenomenon emerges – one that was overlooked due to an error in the forerunner of this paper [5]. While van Benthem's characterisation of basic modal logic remains valid in restriction to the class of all (finite and infinite) transitive frames, the finite model theory analogue fails. The modal translation of first-order formulae that are bisimulation invariant over all *finite* transitive frames in general requires a new modality, which, furthermore, is only compatible with bisimulation in restriction to transitive frames without strict infinite paths (i.e., infinite irreversible paths, or infinite paths whose edges are all one-way). To our knowledge, this is the first instance of a significant discrepancy of this kind between the classical and the finite model theory of basic modal logic. A major extension of [5] moreover concerns the adaptation of the decomposition techniques employed in connection with transitive frames to cover monadic second-order rather than just first-order logic. We obtain results to the effect that bisimulation invariant monadic second-order logic collapses to bisimulation invariant first-order logic over the class of all finite transitive frames and over related classes of (not necessarily finite) transitive trees with certain well-foundedness conditions. Over these classes, the extension of basic modal logic by the new modality, which is first-order but in modal terms falls within the alternation free fragment of the μ-calculus, captures not only bisimulation invariant first-order but monadic second-order.

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[1] 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 isomorphism and ℓ-partial isomorphism 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.

[1] Such systems are also known as '*zig-zag*' in modal logic literature.

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 $R_1, \dots, R_m, P_1, \dots, P_n$ is a relational structure $\mathfrak{A} = (A, R_1^{\mathfrak{A}}, \dots, R_m^{\mathfrak{A}}, P_1^{\mathfrak{A}}, \dots, P_m^{\mathfrak{A}})$, with $R_i^{\mathfrak{A}} \subseteq A \times A$ and $P_j^{\mathfrak{A}} \subseteq A$. The underlying *frame* is the $R_1, \dots, R_m$ reduct $(A, R_1^{\mathfrak{A}}, \dots, R_m^{\mathfrak{A}})$. By a *pointed Kripke structure*, $\mathfrak{A}, a$, we mean a structure $\mathfrak{A}$ with a distinguished element $a \in A$.

Consider 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. A (modal) *back-and-forth system* of depth $\ell \in \mathbb{N}$ between $\mathfrak{A}$ and $\mathfrak{B}$ is given by a family $(Z_k)_{0 \leqslant k \leqslant \ell}$ 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 interpreted over $\mathfrak{A}$ as $(R_i^{-1})^{\mathfrak{A}} := (R_i^{\mathfrak{A}})^{-1} = \{(a, a') \colon (a, a') \in R_i^{\mathfrak{A}}\}$.

For the corresponding choice of relations R, the *back-and-forth conditions w.r.t. to* R for $(Z_k)_{0 \leqslant k \leqslant \ell}$ require that Z' has back-and-forth-continuations along R for all pairs in Z, where $Z = Z_k$ and $Z' = Z_{k-1}$ for $1 \leqslant k \leqslant \ell$:

Z' has back-and-forth-continuations along R for all pairs in Z:

(*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 $(a', b') \in Z'$.

A single $Z \subseteq A \times B$ is similarly called a back-and-forth system (of unbounded depth) between $\mathfrak{A}$ and $\mathfrak{B}$ if it respects the unary P_j and if $Z' = Z$ satisfies the corresponding back-and-forth for Z. It is this kind of modal back-and-forth systems Z that capture the usual notion of bisimulation, while families $(Z_k)_{0 \leqslant k \leqslant \ell}$ describe finite approximations of depth ℓ, or ℓ-bisimulations, discussed below.

Definition 2.1. Let $\mathfrak{A}$ and $\mathfrak{B}$ be Kripke structures of type $R_1, \dots, R_m, P_1, \dots, P_n$. 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 with respect to each R_i.

A bisimulation is a *two-way bisimulation* if it also satisfies the back-and-forth conditions with respect to the inverses R_i^{-1}.

A (forward or two-way) bisimulation Z between $\mathfrak{A}$ and $\mathfrak{B}$ is *global* if it relates every node of $\mathfrak{A}$ to some node of $\mathfrak{B}$ and vice versa, i.e., if $\pi_1(Z) = A$ and $\pi_2(Z) = B$.

Note in connection with global bisimulations that the condition is equivalent to requiring the back-and-forth conditions also with respect to the universal relations $U^{\mathfrak{A}} = A \times A$ and $U^{\mathfrak{B}} = B \times B$.

We write $Z \colon \mathfrak{A} \sim \mathfrak{B}$ to indicate that Z is a bisimulation between $\mathfrak{A}$ and $\mathfrak{B}$; $Z \colon \mathfrak{A} \sim_{\leftrightarrows} \mathfrak{B}$ 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 $(a, b) \in Z$.

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. 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. Either player loses when stuck for a move; player II also loses in any position in which the pebbled pair does not respect all unary P_j. Player II wins the (infinite) game if she can 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 the 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 corresponding kind (forward, two-way, ...).

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 $\mathfrak{B}$ with respect to the ordinary forward back-and-forth conditions, and such that $(a, b) \in Z_\ell$. The two-way variant $\sim^{\ell}_{\leftrightarrows}$ is similarly defined with reference to backward and forward moves.

For corresponding global variants, $\mathfrak{A}, a \sim^{\ell}_{\forall} \mathfrak{B}, b$ or $\mathfrak{A}, a \approx^{\ell} \mathfrak{B}, b$, one also requires back-and-forth conditions with respect to unrestricted global moves. But this is equivalently covered by the requirement that $\pi_1(Z_{\ell-1}) = A$ and $\pi_2(Z_{\ell-1}) = B$ (note the role of $Z_{\ell-1}$ rather than Z_ℓ).

If we refer to global ℓ-bisimulation equivalence between structures without distinguished elements, $\mathfrak{A} \sim^{\ell}_{\forall} \mathfrak{B}$ or $\mathfrak{A} \approx^{\ell} \mathfrak{B}$, however, we want this to mean that there is a corresponding back-and-forth system $(Z_k)_{0 \leqslant k \leqslant \ell}$ such that $\pi_1(Z_\ell) = A$ and $\pi_2(Z_\ell) = B$.

Definition 2.2. $\mathfrak{A}, a \sim^{\ell} \mathfrak{B}, b$ or $\mathfrak{A}, a \sim^{\ell}_{\leftrightarrows} \mathfrak{B}, b$ if there is a back-and-forth system $(Z_k)_{0 \leqslant k \leqslant \ell}$ of the corresponding kind (i.e., forward or two-way), with $(a, b) \in Z_\ell$.

$\mathfrak{A}, a \sim^{\ell}_{\forall} \mathfrak{B}, b$ or $\mathfrak{A}, a \approx^{\ell} \mathfrak{B}, b$ if there is a (forward or two-way, respectively) back-and-forth system $(Z_k)_{0 \leqslant k \leqslant \ell}$ with $(a, b) \in Z_\ell$ and $\pi_1(Z_{\ell-1}) = A$ and $\pi_2(Z_{\ell-1}) = B$.

Without distinguished parameters, $\mathfrak{A} \sim^{\ell}_{\forall} \mathfrak{B}$ or $\mathfrak{A} \approx^{\ell} \mathfrak{B}$, if there is a (forward or two-way, respectively) back-and-forth system $(Z_k)_{0 \leqslant k \leqslant \ell}$, with $\pi_1(Z_\ell) = A$ and $\pi_2(Z_\ell) = B$.

Other variants of bisimulation equivalence will be introduced and treated in exactly the same spirit, as needed in technical context. In particular we shall – for technical purposes – look at refinements of these notions of bisimulation to variants that count up to some finite threshold value (as associated with graded modalities) in section 3; in parts of section 4 we shall look at a customised refinement of ordinary bisimulation that involves certain distinctions as to reflexive versus irreflexive nodes in transitive frames.

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 $R_1, \dots, R_m, P_1, \dots, P_n$, 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 $R = R_1, \dots, R_m$. 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 with respect to the inverses of the given accessibility relation and with respect to the universal accessibility relation. These accessibility relations give rise to backward and global modalities, respectively.

In connection with finite transitive frames, whose accessibility relation R is neither required to be irreflexive nor to be reflexive, we shall also look at the extension of modal logic with extra modalities $\Diamond^*$ and $\Diamond^*_p$ for sets of propositional types p in the unary predicates P_j. Here $\Diamond^*$ is the modality associated with the derived accessibility relation $\{(a,a')\colon (a,a'),(a',a') \in R^{\mathfrak{A}}\}$; its first-order translation is

$$(\Diamond^*\varphi)(x) \;:=\; \exists y\big(Rxy \wedge Ryy \wedge \varphi(y)\big).$$

A form of this modality, equivalent with the given definition over finite transitive R-frames and bisimulation-invariant over all structures, asserts the existence of an infinite R-path along which φ is true infinitely often. This is no longer first-order definable.

The more complex modalities $\Diamond^*_p$ will only be needed in transitive R-frames, in which over and above reflexive nodes also R-cliques of more than one element are allowed. $\Diamond^*_p\varphi$ then asserts the accessibilty of a node where φ holds and that is part of an R-clique in which all propositional types in p are realised. To be precise, let $S_1, \dots, S_s \subseteq \{P_1, \dots, P_n\}$ be the propositional types in p and for each $1 \leqslant i \leqslant s$, let $\zeta_i(x)$ denote the formula $\bigwedge_{P \in S_i} Px \wedge \bigwedge_{Q \notin S_i} Qx$. Then, in first-order terms:

$$(\Diamond^*_p\varphi)(x) \;:=\; \exists y_0 \exists y_1 \dots \exists y_s\big(Rxy_0 \wedge Ry_0y_0 \wedge \varphi(y_0) \wedge \bigwedge_{0 \leqslant i,j \leqslant s} Ry_iy_j \wedge \bigwedge_{1 \leqslant i \leqslant s} \zeta_i(y_i)\big).$$

In particular, for $p = \emptyset$ this reduces to the above $\Diamond^*$. Again, there is a non-first-order, bisimulation invariant generalisation of this definition, which is equivalent with the given one over finite transitive R-frames.[2]

Definition 2.3. 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.
We write ML^* for the extension of ML by the new modalities $\Diamond^*_p$ (including $\Diamond^* = \Diamond^*_{\emptyset}$).

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.[3] For $\ell \in \mathbb{N}$ and a logic $\mathcal{L}$ like $\mathrm{ML}, \mathrm{ML}^{\forall}, \mathrm{ML}^{-\forall}, \mathrm{ML}^*$ 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

[2] As we shall discuss later, it is not finiteness but the absence of certain strict infinite paths that matter for the equivalence of the given definitions – or the bisimulation invariance of the first-order variants.

[3] Incidentally, this counts the modalities $\Diamond^*_p$ with depth 1, too.

$\mathcal{L}$-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 [7] or [6]. For accounts specifically in connection with the modal case see for instance [3] and also [13] or [9].

Observation 2.4. *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 $\ell \in \mathbb{N}$:*

$\mathfrak{A}, a \sim^\ell \mathfrak{B}, b$ *iff* $\mathfrak{A}, a \equiv_{\mathcal{L}^\ell} \mathfrak{B}, b$ *for* $\mathcal{L} = \mathrm{ML}$,
$\mathfrak{A}, a \sim^\ell_\forall \mathfrak{B}, b$ *iff* $\mathfrak{A}, a \equiv_{\mathcal{L}^\ell} \mathfrak{B}, b$ *for* $\mathcal{L} = \mathrm{ML}^\forall$,
$\mathfrak{A}, a \sim^\ell_{\leftrightarrows} \mathfrak{B}, b$ *iff* $\mathfrak{A}, a \equiv_{\mathcal{L}^\ell} \mathfrak{B}, b$ *for* $\mathcal{L} = \mathrm{ML}^{-}$,
$\mathfrak{A}, a \approx^\ell \mathfrak{B}, b$ *iff* $\mathfrak{A}, a \equiv_{\mathcal{L}^\ell} \mathfrak{B}, b$ *for* $\mathcal{L} = \mathrm{ML}^{-\forall}$,
$\mathfrak{A}, a \sim^\ell_* \mathfrak{B}, b$ *iff* $\mathfrak{A}, a \equiv_{\mathcal{L}^\ell} \mathfrak{B}, b$ *for* $\mathcal{L} = \mathrm{ML}^*$.

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.

Let us use $\leftrightarrows$ as generic notation for the back-and-forth equivalence corresponding to the infinite game (for instance, $\sim$ or $\approx$) and $\leftrightarrows^\ell$ for the finite approximation corresponding to the ℓ-round game (for instance, $\sim^\ell$ or $\approx^\ell$) and $\mathcal{L}$ for the corresponding logic (for instance, ML or $\mathrm{ML}^{-\forall}$). 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 Karp's theorem says that partial isomorphism (the back-and-forth equivalence $\leftrightarrows$ associated with the infinite variant of the classical first-order Ehrenfeucht-Fraïssé game) is captured by equivalence in $\mathrm{L}_{\infty\omega}$ (the infinitary variant of FO). For ω-saturated structures, moreover, elementary equivalence (the corresponding $\leftrightarrows^\omega$) coincides with partial isomorphism.

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 [17] is the paradigmatic result of this kind.

Theorem 2.5 (van Benthem)**.** *The following are equivalent for $\varphi = \varphi(x) \in \mathrm{FO}$:*

(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 $\tilde{\varphi} \in \mathrm{ML}$.*

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 (pointed) 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 classes of (pointed) Kripke structures which we consider here are all defined in terms of the underlying (pointed) frames. We therefore also speak of characterisations over corresponding classes $\mathcal{C}$ of frames when we mean characterisations over the class of Kripke structures over such frames.

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 [14]. For variations on its proof and ramifications for stronger variants of bisimulation equivalence see also [12] and [13].

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, we treat the following classes of frames, which will be formally introduced in the following section:

- (finite) rooted frames;
- (finite) trees;
- (finite) equivalence frames;
- (finite) irreflexive or reflexive transitive trees;
- transitive frames without infinite paths (Löb frames);
- path-finite reflexive transitive frames paths (Grzegorczyk frames);
- path-finite transitive frames;
- (finite) rooted transitive frames.

The corresponding main results are to be found in section 3 for the first three classes, and in section 4 for the latter three.

2.3 Some relevant classes of frames

Obvious constraints on the accessibility relations in Kripke structures concern reflexivity, symmetry and transitivity (combinations of which feature prominently in classical modal logic, especially in *correspondence theory*, see, e.g., [3]). Recall that R° stands for the reflexive part of the relation R; $R^\mathfrak{A}$ is irreflexive if $(R^\circ)^\mathfrak{A} = \emptyset$ and reflexive if $(R^\circ)^\mathfrak{A} = \{(a,a)\colon a \in A\}$. Apart from these elementary constraints we shall be particularly interested in reachability issues, in particular rootedness and tree-likeness.

Definition 2.6. A *directed path of length* m from a to b in $\mathfrak{A}$ is a sequence of nodes $\sigma = a_0, a_1, \ldots, a_m$ starting at $a_0 = a$ and ending in $a_m = b$ in $\mathfrak{A}$ such that for all $k < m$ the pair (a_k, a_{k+1}) is in at least one of the binary relations R_i of $\mathfrak{A}$.

A *directed labelled path of length* m from a to b in $\mathfrak{A}$ is a sequence of nodes and relations $\sigma = a_0, R_{i_0}, a_1, R_{i_1}, \ldots, R_{i_m}, a_m$ such that $(a_k, a_{k+1}) \in R_{i_k}^\mathfrak{A}$. *Undirected labelled paths* are similarly defined, by allowing both the R_i and their inverses R_i^{-1} for links between consecutive nodes.

Where edge labels are not important, we may drop them and speak of just paths rather than labelled paths.

Definition 2.7. A (labelled) *undirected cycle of length* m at a in $\mathfrak{A}$ is an undirected (labelled) path of length m from a to a. A cycle is non-trivial if in its labelled version there are no consecutive traversals of the same edge in opposite directions ($a_{k+2} = a_k$ implies $R_{i_{k+1}} \neq R_{i_k}^{-1}$).
$\mathfrak{A}$ is *acyclic* if it has no non-trivial undirected cycles, and k*-acyclic* for some $k \geqslant 1$, if it has no non-trivial undirected cycles of length up to k.

Undirected paths and cycles may be viewed as paths or cycles in the associated Gaifman graph of $\mathfrak{A}$, which is an undirected graph (cf. Definition 2.14 below).

We note that the strong notion of acyclicity defined here does forbid reflexive edges (loops) as well as symmetric edges or multiple edges (edges in $R_i \cap R_j$ or in $R_i \cap R_j^{-1}$ for $i \neq j$ in the multi-modal case).

We also note that the class of k-acyclic frames is definable in FO by a single sentence $\varphi_k \in$ FO for each k; correspondingly the class of acyclic frames is elementary, defined by the FO-theory $\{\varphi_k\colon k \geqslant 1\}$ though not by any single FO-sentence.

We turn to classes of frames that will be central to our investigations. The key conditions center on notions of reachability and transitivity.

Rootedness refers to the existence of a root node from which all other nodes are reachable on directed paths; with this root as the distinguished node we obtain the class of rooted frames and Kripke structures, as special classes of pointed frames or pointed Kripke structures.

Tree-likeness comes in two rather different flavours, one graph theoretic and one order theoretic. In the usual graph theoretic sense a (directed) tree is a rooted frame in which different paths from the root do not meet (i.e. the induced undirected graph is acyclic); in the order theoretic sense, a tree is a partial order (hence in particular it is transitive) with a unique minimal element (the root) and without meets between incomparable elements (i.e. no forward confluence). One may look both at the irreflexive and at the reflexive variants ($\prec$-trees or $\preccurlyeq$-trees); technically we shall treat the strict or irreflexive variant of a partial order encoding of a tree as the basic notion. The treatment of the more liberal class of transitive tree-like frames, which need neither be reflexive nor irreflexive, turns out to be rather different from that of (ir-)reflexive trees, and will provide the crucial stepping stone towards the important class of all (finite) rooted transitive frames.

As far as infinite tree frames of the order theoretic kind are concerned, it is standard to require the set of predecessors of any node to be well-ordered. In most instances where we shall have occasion to consider infinite order theoretic trees, these will arise through a process of tree unravelling (compare section 2.5) and thus even satisfy the stronger requirement that the predecessor sets of all nodes are finite. Another, converse well-foundedness property, which plays a role from the point of view of modal logic, is the one that rules out infinite R-paths, i.e., well-foundedness w.r.t. R^{-1}. Transitive frames with this property are known as Löb frames; their tree counterparts are irreflexive transitive trees without infinite paths. The reflexive version cannot forbid infinite R-paths, but only infinite paths with respect to the irreflexive part $R \setminus R^{\circ}$ of R (we write R° for the set of reflexive R-arcs or loops). Reflexive transitive frames without infinite paths w.r.t. $R \setminus R^{\circ}$, i.e., the reflexive closures of Löb frames, are known as Grzegorczyk frames [3]; their tree counterparts are the reflexive closures of irreflexive transitive trees without infinite paths. Another, even weaker well-foundedness condition in transitive frames, in which neither reflexivity nor irreflexivity is prescribed, forbids infinite paths w.r.t. $R \setminus R^{-1}$ (infinite *irreversible* or *strict* paths): we shall speak of *path-finite* transitive frames. Note that the class of all path-finite transitive frames is a proper extension of the class of all finite transitive frames, just as the classes of Löb and Grzegorczyk frames are proper extensions of the classes of finite transitive irreflexive frames and of the class of finite transitive reflexive trees, respectively.

Definition 2.8. A pointed Kripke structure $\mathfrak{A}, a$ and its underlying pointed frame is

- (i) *rooted* if every node of A is reachable on a (labelled) directed path from the distinguished element a.
- (ii) a *tree* (a directed tree in the graph theoretic sense) if it is rooted and every node is reachable on a unique labelled directed path from the root.

A pointed Kripke structure $\mathfrak{A}, a$ with a single binary relation R and its underlying frame $(A, R^{\mathfrak{A}}, a)$ is

- (iii) an *irreflexive transitive tree* (a tree in the strict order theoretic sense), if it is rooted, and if R is transitive, irreflexive and such that the R-predecessors of any node are well-ordered by R.[4] In this case we write $x \prec y$ instead of Rxy.
- (iv) a *reflexive transitive tree* (a tree in the reflexive order theoretic sense) if it is rooted, and if R is transitive, reflexive and such that the R-predecessors of any node are well-ordered by the irreflexive part of R, $R \setminus R^{\circ}$. In this case, we write $x \preccurlyeq y$ instead of Rxy.
- (v) *transitive tree-like* if it is rooted, and if R is transitive and if the irreflexive part of R, $R \setminus R^{\circ}$ forms an irreflexive transitive tree.

[4] I.e., linearly ordered and without infinite descending sequences. The linear order condition rules out confluence of paths, while the well-foundedness condition rules out infinite R^{-1}-paths.

(vi) *weak transitive tree-like* if it is rooted, transitive and confluence-free, in the sense that $\forall xyz\big((\neg Rxy \wedge \neg Ryx) \rightarrow \neg(Rxz \wedge Ryz)\big)$.

We shall often just refer to $\prec$-trees and $\preccurlyeq$-trees for irreflexive and reflexive transitive trees, respectively. Clearly $\prec$- and $\preccurlyeq$-trees are transitive tree-like, and transitive tree-like frames are weak transitive tree-like.

Weak transitive tree-like frames are more general than transitive tree-like frames even at the level of finite frames, because R is not required to be acyclic.

Unless explicitly stated otherwise, we reserve the term *tree* or *rooted tree* for trees in the graph theoretic sense; transitive trees or transitive tree-like structures are not trees in this sense.

We note that the class of weak transitive tree-like frames is elementary (FO definable) just like the class of rooted transitive frames, because rootedness is trivially FO definable for transitive relations. The classes of rooted frames or of trees (in the graph theoretic or in the order theoretic sense) or of transitive tree-like frames on the other hand, are not FO definable, since rootedness of graphs or well-ordering conditions are not first-order. In restriction to just finite frames, the order theoretic trees do become FO definable, while the graph theoretic trees remain undefinable (rootedness, or connectivity, not being FO definable even in finite graphs).

Equivalence frames are frames in which the accessibility relations are equivalence relations. This elementary class of frames is prominent in epistemic logic, and is – via modal correspondence theory – linked with the modal logic S5.

Definition 2.9. A Kripke structure $\mathfrak{A}$ and its underlying frame is called an *equivalence frame* or an *S5 frame* if each accessibility relation $R_i^{\mathfrak{A}}$ is an equivalence relation.

All of the frame classes in Definition 2.8 are rooted. Rooted frames are arguably the most commonly intended models in many 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 2.10. *The $\sim$ class of a finite pointed 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 ensure 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 rooted 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.

Observation 2.11. *For all rooted frames $\mathfrak{A}, a$ and $\mathfrak{B}, b$: $\mathfrak{A}, a \sim \mathfrak{B}, b$ iff $\mathfrak{A}, a \sim_{\forall} \mathfrak{B}, b$.*

Indeed, as every node a' of $\mathfrak{A}$ for instance is reachable from the root a in a finite sequence of transitions, corresponding responses by player II to this sequence of moves played by I will lead to a game position (a', b') such that $\mathfrak{A}, a' \sim \mathfrak{B}, b'$.

Hence $\mathrm{ML}^{\forall}$ rather than ML becomes the right candidate for capturing bisimulation invariant FO-properties. That is, in section 3.1 we show that $\mathrm{FO}/{\sim} \equiv \mathrm{ML}^{\forall}$ over the class of arbitrary rooted frames and in section 3.2 we show that the same holds over finite rooted frames (see Theorems 3.4 and 3.5, respectively). We remark that, in restriction to rooted frames, the unsatisfactory situation of Observation 2.10 disappears. The $\sim$ class of any finite rooted structure is definable by a formula of $\mathrm{ML}^{\forall}$ within the class of all rooted structures.

Other characterisation theorems we obtain are that $\mathrm{FO}/{\sim_{\forall}} \equiv \mathrm{ML}^{\forall}$ over (finite) equivalence structures (see Theorem 3.14); $\mathrm{FO}/{\sim} \equiv \mathrm{ML}$ over the classes of finite or of not necessarily finite irreflexive transitive trees (see Theorems 4.11 and 4.12), over the class of transitive structures without infinite paths (so-called Löb-frames, see Theorem 4.13), and over the classes of finite or of not necessarily finite reflexive transitive trees (see Theorems 4.14 and 4.21); and $\mathrm{FO}/{\sim} \equiv \mathrm{ML}^{*}$ over the classes of finite transitive tree-like structures (see Theorem 4.27), of finite rooted transitive structures and of all finite transitive structures (see Theorem 4.40).

2.4 Model theoretic methods

For the following discussion let again $\leftrightarrows$ be a generic 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 $\mathcal{L}^{\ell}$.

(iii) (relevant for the classical argument only:) the common refinement $\leftrightarrows^{\omega} := \bigcap_{\ell} \leftrightarrows^{\ell}$ coincides with $\leftrightarrows$ over ω-saturated structures.

Classical methods. Classical first-order proofs of characterisation theorems of the kind considered here, typically involve a compactness argument to show that any $\leftrightarrows$ invariant first-order formula is equivalently expressible in $\mathcal{L}$. In our framework this is equivalent to showing that $\leftrightarrows$ invariance implies $\leftrightarrows^{\ell}$ invariance at some finite level ℓ (compare [13, 9] for a more comprehensive discussion). We present an outline of an indirect argument, based on compactness and saturation, to this effect.

We reason towards a contradiction. Suppose that $\varphi(x) \in \mathrm{FO}$ is $\leftrightarroweq$ invariant but not expressible in the fragment $\mathcal{L} \subseteq \mathrm{FO}$.

Consider first a fixed level ℓ, starting for instance with $\ell = 0$. By assumptions (i) and (ii) there is a finite family of formulae $\chi_i^\ell \in \mathcal{L}^\ell$ such that each χ_i^ℓ defines an equivalence class of $\leftrightarroweq^\ell$ and $\models \bigvee_i \chi_i^\ell$. As φ is not expressible in $\mathcal{L}$, φ must be inexpressible in restriction to at least one of these equivalence classes. Otherwise, $\varphi \wedge \chi_i^\ell \equiv \psi_i \wedge \chi_i^\ell$ for each i would imply $\varphi \equiv \bigvee_i (\psi_i \wedge \chi_i^\ell)$. It follows that $\varphi \wedge \chi_i^\ell$ is not expressible in $\mathcal{L}$ for at least one i. Replacing φ by $\varphi \wedge \chi_i^\ell$ for one such i and proceeding inductively from ℓ to $\ell + 1$, we obtain a sequence of formulae $(\chi^\ell)_{\ell \in \mathbb{N}}$ such that $\chi^\ell \in \mathcal{L}^\ell$ defines an equivalence class of $\leftrightarroweq^\ell$, $\chi^{\ell+1} \models \chi^\ell$ and φ is not equivalent to a formula of $\mathcal{L}$ in restriction to the equivalence class defined by χ^ℓ. By compactness both $\{\varphi\} \cup \{\chi^\ell \colon \ell \in \mathbb{N}\}$ and $\{\neg\varphi\} \cup \{\chi^\ell \colon \ell \in \mathbb{N}\}$ must be satisfiable. By construction, $\{\chi^\ell \colon \ell \in \mathbb{N}\}$ is a complete theory in the logic $\mathcal{L} = \bigcup_i \mathcal{L}^\ell$. If $\mathfrak{A}, a \models \{\varphi\} \cup \{\chi^\ell \colon \ell \in \mathbb{N}\}$ and $\mathfrak{B}, b \models \{\neg\varphi\} \cup \{\chi^\ell \colon \ell \in \mathbb{N}\}$, then

$$\mathfrak{A}, a \equiv_{\mathcal{L}} \mathfrak{B}, b \quad \text{but} \quad \mathfrak{A}, a \models \varphi \text{ and } \mathfrak{B}, b \models \neg\varphi.$$

Passing to ω-saturated elementary extensions $\mathfrak{A}^*$ and $\mathfrak{B}^*$ of $\mathfrak{A}$ and $\mathfrak{B}$ respectively, which are available by compactness, and using (iii) one finds

$$\mathfrak{A}^*, a \leftrightarroweq \mathfrak{B}^*, b \quad \text{but} \quad \mathfrak{A}^*, a \models \varphi \text{ and } \mathfrak{B}^*, b \models \neg\varphi,$$

which contradicts $\leftrightarroweq$ invariance of φ.

This classical proof naturally relativises to any elementary class of structures. If $\mathcal{C}$ is the class of models of the FO theory T, then all of the above steps can be repeated in restriction to $\mathcal{C}$, if we replace the notions of equivalence, consequence, definability etc., by equivalence, consequence, definability etc. under T. More precisely, if φ were $\leftrightarroweq$ invariant over $\mathcal{C}$, but not logically equivalent to a formula of $\mathcal{L}$ over $\mathcal{C}$, then we would obtain first a sequence of formulae $\chi^\ell \in \mathcal{L}^\ell$ such that χ^ℓ defines an equivalence class of $\leftrightarroweq^\ell$ (within $\mathcal{C}$), $T \models \chi^{\ell+1} \rightarrow \chi^\ell$, and such that $\varphi \wedge \chi^\ell$ is not logically equivalent to a formula of $\mathcal{L}$ over $\mathcal{C}$, i.e., not $T \models \varphi \leftrightarrow \psi$ for any $\psi \in \mathcal{L}$. By compactness, both $T \cup \{\varphi\} \cup \{\chi^\ell \colon \ell \in \mathbb{N}\}$ and $T \cup \{\neg\varphi\} \cup \{\chi^\ell \colon \ell \in \mathbb{N}\}$ must be satisfiable, and we proceed as above to obtain $\mathfrak{A}^*, a \leftrightarroweq \mathfrak{B}^*, b$ with $\mathfrak{A}^*, a \models \varphi$ and $\mathfrak{B}^*, b \models \neg\varphi$, where now both structures are from $\mathcal{C}$, contradicting the assumption that φ is $\leftrightarroweq$ invariant over $\mathcal{C}$.

This immediately gives for instance the following ramifications of Theorem 2.5. Compare section 2.3 for the relevant classes of frames. Recall that we say that $\mathcal{L}$ captures bisimulation invariant first-order logic over $\mathcal{C}$, $\mathrm{FO}/{\sim} \equiv \mathcal{L}$ over $\mathcal{C}$, if the following are equivalent for any $\varphi(x) \in \mathrm{FO}$:

(i) φ is bisimulation invariant over $\mathcal{C}$: $\mathfrak{A}, a \sim \mathfrak{B}, b$ for $(\mathfrak{A}, a), (\mathfrak{B}, b) \in \mathcal{C}$ implies that $\mathfrak{A}, a \models \varphi$ iff $\mathfrak{B}, b \models \varphi$.

(ii) φ is equivalent over $\mathcal{C}$ to a formula $\tilde{\varphi} \in \mathcal{L}$.

Theorem 2.12. ML *captures bisimulation invariant first-order logic over any elementary class of Kripke frames. In particular, for instance, over*

(a) *the classes of acyclic frames;*

(b) *the class of (irreflexive or reflexive) transitive frames;*

(c) *the class of rooted (irreflexive or reflexive) transitive frames; and*

(d) *the class of equivalence frames.*

For (c) note again that rootedness is first-order in conjunction with transitivity, while it is not elementary on its own.

For restrictions to classes of finite structures, however, none of these classical techniques seem to be available. And the same applies to other interesting non-elementary classes of not necessarily finite frames, like trees, transitive trees, rooted frames, or the class of transitive frames without infinite paths (Löb frames).

A models-for-games method. In contrast to the classical methods, and similar to the finite model theory proofs in [14, 13], we proceed along an orthogonal direction to establish a characterisation result $\mathrm{FO}/\simeq \equiv \mathcal{L}$ in restriction to a class of structures $\mathcal{C}$ which in particular need not be first-order definable. Again the crux of the expressive completeness claim is to establish that $\simeq$ invariance implies $\simeq^\ell$ invariance for some finite level ℓ, over the class of structures $\mathcal{C}$.

The alternative approach for establishing this is based on a gradation of elementary equivalence $\equiv$ into successively refined approximations $\overset{n}{\equiv}$ for $n \in \mathbb{N}$.

We assume, as above, that $\simeq$, $\simeq^\ell$, $\mathcal{L}$ and $\mathcal{L}^\ell$ are such that

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

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

In addition, let now $\overset{n}{\equiv}$ be a family of equivalences that approximate elementary equivalence, such that

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

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

$$\mathfrak{A}, a \simeq^{f(n)} \mathfrak{B}, b \;\Rightarrow\; F_n(\mathfrak{A}, a) \overset{n}{\equiv} F_n(\mathfrak{B}, b).$$

We may view point (iv) as a mechanism for *upgrading* suitable finite levels $\simeq^\ell$ of approximate $\simeq$ equivalence to the necessary finite level of approximate first-order equivalence, through model transformations within $\mathcal{C}$ that fully preserve $\simeq$.[5]

A characterisation theorem is immediate under these circumstances.

Proposition 2.13. *Let $\simeq$ and $\simeq^\ell$, $\mathcal{L}$ and $\mathcal{L}^\ell$ be as described in (i) and (ii). If suitable $\overset{n}{\equiv}$ satisfy (iii) and (iv), then* $\mathrm{FO}/\simeq \equiv \mathcal{L}$ *over $\mathcal{C}$.*

Proof. Let $\varphi(x) \in \mathrm{FO}$ be $\simeq$ invariant over $\mathcal{C}$. Choose n such that φ is preserved under $\overset{n}{\equiv}$. For structures in $\mathcal{C}$, $\mathfrak{A}, a \simeq^{f(n)} \mathfrak{B}, b$ implies that

$$\mathfrak{A}, a \simeq F_n(\mathfrak{A}, a) \overset{n}{\equiv} F_n(\mathfrak{B}, b) \simeq \mathfrak{B}, b,$$

and hence $\mathfrak{A}, a \models \varphi$ iff $\mathfrak{B}, b \models \varphi$. So φ is preserved under $\simeq^\ell$ for $\ell = f(n)$ over $\mathcal{C}$, and hence equivalent over $\mathcal{C}$ to a disjunction of characteristic formulae according to (ii). Indeed, if the $\simeq^\ell$ classes are defined by $\mathcal{L}^\ell$-formulae $(\chi_i^\ell)_{i \in I}$, let $I_\varphi := \{i \in I \colon \chi_i \models \varphi\}$. As $\simeq^\ell$ has finite index, I and I_φ are finite and $\varphi \equiv \bigvee_{i \in I_\varphi} \chi_i^\ell \in \mathcal{L}^\ell$ is as desired. $\square$

[5] A similar notion of an *upgrading* of, say, modal equivalence to elementary equivalence in bisimilar companion structures, has previously been discussed in [2] (in the classical context, and without the finitary gradation).

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 as its robustness under relativisation is concerned. As discussed in [12, 9], 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 transformations F_n (certain locally acyclic covers in the case of [13], cf. section 2.5 here).

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.

After outlining some general purpose constructions in section 2.5, the main part of the paper is organised according to the above 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 in epistemic logic.

Section 4 develops the models-for-games technique with an emphasis on *decomposition arguments* with respect to the first-order Ehrenfeucht-Fraïssé games over suitable classes of finite frames, most notably for finite transitive trees and finite transitive tree-like structures. These techniques are applied to cover the cases of all finite transitive frames and of all finite transitive rooted frames as well. Decomposition techniques are of course well suited to the treatment of monadic second-order logic and we correspondingly extend some key results concerning bisimulation invariant first-order logic to bisimulation invariant monadic second-order logic, over several relevant frame classes in section 4.7. These techniques also yield further insights into the relationship between first-order and monadic second-order equivalence and bisimulation equivalence over certain classes of transitive trees, including some classes of not necessarily finite transitive trees.

2.5 Constructing bisimilar companions

Many of the constructions used to prove the results in sections 3 and 4 involve building bisimilar companions of a given structure $\mathfrak{A}$. These are structures that are bisimilar to $\mathfrak{A}$ that also satisfy additional requirements. While many of these constructions are specialised to the requirements of the particular result being proved, we collect here some useful techniques that are used more than once.

Tree unravellings Let $\mathfrak{A}, a$ be a pointed structure. We define the *tree unravelling* of $\mathfrak{A}, a$ from a, denoted $\mathfrak{A}^*_a$, as follows. The universe of $\mathfrak{A}^*_a$ consists of all labelled directed paths $\sigma = a_0, R_{i_0}, a_1, R_{i_1}, \dots, R_{i_n}, a_n$ from $a_0 = a$. For the interpretation of the binary accessibility relations R_i and of the unary predicates P, we put

– an R_i-edge in $\mathfrak{A}^*_a$ from $\sigma = a, R_{i_0}, \dots, a_n$ to precisely all its extensions by one R_i-edge in $\mathfrak{A}$, $\sigma' = a, R_{i_0}, \dots, a_n, R_i, a_{n+1}$;

– $\sigma = a, R_{i_0}, \dots, R_{i_n}, a_n$ in P in $\mathfrak{A}^*_a$ iff $a_n \in P^{\mathfrak{A}}$.

The following observations are straightforward.

(1) $\mathfrak{A}^*_a$ is a tree rooted at a, in general an infinite one.

(2) If $\mathfrak{A}$ is finite and acyclic, then $\mathfrak{A}^*_a$ is a finite rooted tree.

(3) $\mathfrak{A}, a \sim \mathfrak{A}^*_a, a$.

Boosting multiplicities For a structure $\mathfrak{A}$ and a positive integer q, we define the structure $\mathfrak{A} \otimes q$ obtained from $\mathfrak{A}$ by boosting multiplicities to q to be the structure with universe $A \times \{0, \dots, q-1\}$ and

– an R_i-edge from (a, j) to (b, k) in $\mathfrak{A} \otimes q$ iff $(a, b) \in R_i^{\mathfrak{A}}$; and

– $(a, j) \in P^{\mathfrak{A} \otimes q}$ iff $a \in P^{\mathfrak{A}}$.

It is easily seen that $\mathfrak{A}, a \sim \mathfrak{A} \otimes q, (a, i)$, for all $a \in A$, $0 \leqslant i < q$.

In order to simplify notation, we shall often identify a distinguished node $a \in A$ with its copy $(a, 0)$ in $\mathfrak{A} \otimes q$ where this does not cause confusion. In this sense, we may write, for instance, $\mathfrak{A} \otimes q, a \sim \mathfrak{A}, a$.

Note that $\mathfrak{A} \otimes q$ is usually not rooted, even if $\mathfrak{A}, a$ is (in this case, $\mathfrak{A} \otimes q$ is rooted at $(a, 0)$ just in case a is on a directed cycle in $\mathfrak{A}$). In general, however, we may pass to the substructure of $\mathfrak{A} \otimes q$ generated by $(a, 0)$, consisting of just the nodes reachable from $(a, 0)$ in $\mathfrak{A} \otimes q$. This structure, which we denote $(\mathfrak{A} \otimes q)_a$, is trivially rooted at $(a, 0)$, has multiplicities boosted by q, and clearly still $(\mathfrak{A} \otimes q)_a, (a, 0) \sim \mathfrak{A}, a$.

Also, neither $\mathfrak{A} \otimes q$ nor $(\mathfrak{A} \otimes q)_a$ will typically be trees, even if $\mathfrak{A}, a$ is a rooted tree. In order to boost multiplicities in trees, therefore, we combine $\otimes$ with an unravelling. The tree unravelling of $\mathfrak{A} \otimes q$ from $(a, 0)$, $(\mathfrak{A} \otimes q)^*_a$, is a tree rooted at $(a, 0)$, with multiplicities boosted by q, and we still have $(\mathfrak{A} \otimes q)^*_a, (a, 0) \sim \mathfrak{A}, a$.

Importantly, in $\mathfrak{A} \otimes q$ nodes (a, j) and (a, k) for any $0 \leqslant j, k < q$ are related by an automorphism that fixes all other nodes. The same is true in $(\mathfrak{A} \otimes q)_a$ for nodes other than the root. In the rooted tree $(\mathfrak{A} \otimes q)^*_a$ similarly, any node other than the root possesses at least q siblings related by automorphisms that fix their common predecessor; in particular, every subtree rooted at some node other than the root a is one of at least q isomorphic and disjoint sibling subtrees.

The usefulness of boosting multiplicities in this fashion becomes apparent in the connection between ℓ-bisimilarity and (local) FO equivalence in trees established in Lemma 2.15 below. In the sequel, we generally write a instead of $(a, 0)$ for the root of $(\mathfrak{A} \otimes q)_a$, etc., where no confusion would arise.

We introduce Gaifman graphs, distances, and neighbourhoods, which are more generally useful for arbitrary relational structures. These notions will be crucial for our locality based techniques in section 3 below. For Kripke structures, Gaifman distance is just the ordinary graph theoretic distance in the undirected graph induced by the union of the accessibility relations. The lemma below could be formulated in terms of trees of bounded depth, but we phrase it in terms of Gaifman neighbourhoods with a view to further developments in section 3.

Definition 2.14. 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 $\mathfrak{A}$.

Gaifman distance in $\mathfrak{A}$ is the natural graph theoretic distance in the Gaifman graph of $\mathfrak{A}$, which we denote $d(.,.)$.

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: $N^\ell(a) = \{a' \in A : d(a, a') \leqslant \ell\}$.

Recall that $\equiv_q$ stands for FO equivalence up to quantifier rank q.

Lemma 2.15. *Let $q \in \mathbb{N}$, $\mathfrak{A}, a$ and $\mathfrak{B}, b$ be rooted trees.*
If $\mathfrak{A}, a \sim^\ell \mathfrak{B}, b$, then $(\mathfrak{A} \otimes q)^*_a{\restriction}N^\ell(a), a \equiv_q (\mathfrak{B} \otimes q)^*_b{\restriction}N^\ell(b), b$. In fact, player II has a winning strategy in the q-round Ehrenfeucht-Fraïssé game on $(\mathfrak{A} \otimes q)^*_a{\restriction}N^\ell(a), a$ versus $(\mathfrak{B} \otimes q)^*_b{\restriction}N^\ell(b), b$ which preserves distances.

Proof. For the analysis of the q-round Ehrenfeucht-Fraïssé game on $(\mathfrak{A} \otimes q)_a{\restriction}N^\ell(a), a$ versus $(\mathfrak{B} \otimes q)_b{\restriction}N^\ell(b), b$ consider configurations with pebbled tuples $\mathbf{a} = (a, a_1, \ldots, a_s)$ and $\mathbf{b} = (b, b_1, \ldots, b_s)$, where $s \leqslant q$. Let $\pi : (\mathfrak{A} \otimes q)^*_a \to \mathfrak{A}$ and $\pi : (\mathfrak{B} \otimes q)^*_b \to \mathfrak{B}$ be the natural projections. We note that for any node a' of $(\mathfrak{A} \otimes q)^*_a$, $(\mathfrak{A} \otimes q)^*_a, a' \sim \mathfrak{A}, \pi(a')$. We also note that π preserves distances from the root; in particular π restricts to a surjective projection $\pi : (\mathfrak{A} \otimes q)^*_a{\restriction}N^\ell(a) \to \mathfrak{A}{\restriction}N^\ell(a)$. The corresponding projection in $\mathfrak{B}$ is also denoted π in the following.

Let $T(\mathbf{a})$ and $T(\mathbf{b})$ be the tree structures induced on the sets of nodes on the shortest paths from the roots a and b to the a_i and b_i, respectively. We want player II to maintain the following isomorphism condition through all q rounds:

$$\begin{array}{l} T(\mathbf{a}), \mathbf{a} \simeq T(\mathbf{b}), \mathbf{b} \quad \text{via an isomorphism } \rho \text{ such that} \\ \mathfrak{A}, \pi(a') \sim^{\ell - r} \mathfrak{B}, \pi(\rho(a')) \;\; \text{for} \;\; r = d(a, a'). \end{array}$$

This condition is clearly satisfied at the start of the game. In order to see that it can be maintained in a single round from current positions $\mathbf{a} = (a, a_1, \ldots, a_s)$ and $\mathbf{b} = (b, b_1, \ldots, b_s)$, where $s < q$, consider a move played by player I for instance to a new element a' in $(\mathfrak{A} \otimes q)^*_a{\restriction}N^\ell(a)$.

If $a' \in T(\mathbf{a})$, player II can respond according to the existing isomorphism ρ between $T(\mathbf{a})$ and $T(\mathbf{b})$. Otherwise the path from a to a' has a unique last element a'_0 within $T(\mathbf{a})$, and we let $b'_0 := \rho(a'_0)$. By assumption $\mathfrak{A}, \pi(a'_0) \sim^{\ell - r} \mathfrak{B}, \pi(b'_0)$ for $r := d(a, a'_0) = d(b, b'_0)$.

As the branching degree of $T(\mathbf{a}) \simeq T(\mathbf{b})$ is less than q, we find a path from b'_0 in $(\mathfrak{B} \otimes q)^*_b$, disjoint from $T(\mathbf{b})$ beyond b'_0 and matching the path from a'_0 in accordance with the requirement that corresponding nodes at distance d from the roots are equivalent in the sense of $\sim^{\ell - d}$. The correspondence along these paths extends the previous isomorphism ρ to an isomorphism ρ' that meets the requirements.

The isomorphisms ρ constructed for player II's strategy in this fashion are fully distance preserving, since they preserve distances from the roots and are isomorphisms of induced subtrees. ☐

Transitive closure For a structure $\mathfrak{A}$ with one accessibility relation R, we write $\mathrm{TC}(\mathfrak{A})$ to denote the structure with the same elements as A, the same interpretation of the unary predicates, but whose accessibility relation is the transitive closure of R. Forming the transitive closure of the accessibility relations in a structure will not, in general, yield a bisimilar structure. However, we will find the following observation useful.

Lemma 2.16. *If $\mathfrak{A}, a$ is a rooted transitive structure and $\mathfrak{A}, a \sim \mathfrak{B}, b$ then $\mathfrak{A}, a \sim \mathrm{TC}(\mathfrak{B}), b$.*

Proof. Let Z be a bisimulation witnessing $\mathfrak{A}, a \sim \mathfrak{B}, b$. We claim that Z is also a bisimulation witnessing $\mathfrak{A}, a \sim \mathrm{TC}(\mathfrak{B}), b$. Indeed, it is easy to see that Z satisfies the forth condition for the accessibility relation R. For the back condition, suppose $(b, b') \in R^{\mathrm{TC}(\mathfrak{B})}$. Then, there is a directed path $b_0, \dots, b_m$ from $b = b_0$ to $b' = b_m$ in $\mathfrak{B}$. Since Z is a bisimulation, this means we can find a path $a_0, \dots, a_m$ from a to some $a' = a_m$ in $\mathfrak{A}$ such that $(a_i, b_i) \in Z$ for each i. By the transitivity of $R^{\mathfrak{A}}$ we have $(a, a') \in R^{\mathfrak{A}}$ and since $(a', b') \in Z$, we are done. $\square$

The following combination of transitive closure with tree unravelling is useful in the analysis of not necessarily finite (rooted) transitive frames.

Observation 2.17. *If $\mathfrak{A}, a$ is transitive, then the transitive closure of its tree unfolding from a, $\mathrm{TC}(\mathfrak{A}^*_a)$, is a bisimilar companion of $\mathfrak{A}, a$ that is an irreflexive transitive tree with root a in which the predecessor set of any node is finite. The same holds for the variants with boosted multiplicities, $\mathrm{TC}((\mathfrak{A} \otimes q)^*_a)$.*

Bisimilar covers We focus on global two-way bisimulations induced by a homomorphism onto a given structure.

Definition 2.18. 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_i and R_i^{-1} in $\hat{\mathfrak{A}}$ are the same as those at $\pi(\hat{a})$ in $\mathfrak{A}$, for every binary R_i; π is faithful if it is faithful in all $\hat{a} \in \hat{\mathfrak{A}}$.

Faithful covers are meant to provide unique lifts of labelled undirected paths from a in $\mathfrak{A}$ at any $\hat{a} \in \pi^{-1}(a)$. As far as (locally) acyclic structures are concerned, faithful covers provide (locally) isomorphic covers in the sense of the observation below.

Note that in- and out-degrees of nodes are considered with respect to each individual binary relation separately. (Also note that a loop at a node a, $(a, a) \in R_i$, would contribute 1 to both the in- and the out-degree of R_i at a; a symmetric edge, $(a, b), (b, a) \in R_i$, contributes 1 to both the in- and the out-degree of R_i at both ends.)

Observation 2.19. *Let $\pi \colon \hat{\mathfrak{A}} \to \mathfrak{A}$ be a faithful bisimilar cover. If $\mathfrak{A}{\restriction}N^\ell(a)$ is acyclic and $\hat{a} \in \pi^{-1}(a)$, then π restricts to an isomorphism between $\hat{\mathfrak{A}}{\restriction}N^\ell(\hat{a})$ and $\mathfrak{A}{\restriction}N^\ell(a)$. $\hat{\mathfrak{A}}{\restriction}\pi^{-1}(N^\ell(a))$ is a disjoint union of isomorphic copies of $\mathfrak{A}{\restriction}N^\ell(a)$.*

Proof. First observe that $\hat{\mathfrak{A}}{\restriction}N^\ell(\hat{a})$ is acyclic as any undirected cycle in $\hat{\mathfrak{A}}$ projects to an undirected cycle in $\mathfrak{A}$ under π. Since $\mathfrak{A}{\restriction}N^\ell(a)$ is acyclic, the projection of a non-trivial cycle would be non-trivial. Injectivity of $\pi{\restriction}\hat{\mathfrak{A}}{\restriction}N^\ell(\hat{a})$ follows from faithfulness. $\square$

Note that $\mathfrak{A} \otimes q$ is a cover of $\mathfrak{A}$ via the natural projection onto the first component; clearly this cover is not faithful for $q > 1$.

A simple generalisation of tree unravellings to two-way or undirected tree unravellings would easily provide faithful covers by acyclic structures, albeit typically infinite ones unless the given structure was acyclic.

For the following recall Definition 2.7. In connection with our locality based analysis in section 3 we shall want to achieve acyclicity in ℓ-neighbourhoods, and therefore forbid undirected cycles of lengths up to $2\ell + 1$. The following is the key result of [13] based on a product of the given finite $\mathfrak{A}$ with Cayley graphs of large girth (without short cycles).

Proposition 2.20. *For all $\ell \geqslant 3$: 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 ℓ-acyclic $\hat{\mathfrak{A}}$.*

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 [13]; their variations obtained here find their key application in characterisation theorems for classes of (finite) rooted frames. With a further generalisation of these techniques we obtain a characterisation theorem over finite equivalence frames (S5 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 [8, 6]. Recall, in connection with Definition 2.14 above, that the Gaifman graph of a Kripke structure $\mathfrak{A}$ is just the undirected graph generated (through symmetrisation) by the accessibility relations. Gaifman distance and neighbourhoods (Definition 2.14) correspondingly, are definable in terms of minimal lengths of undirected paths between nodes. In addition to ℓ-neighbourhoods we need the following notion of ℓ-scattered sets.

Definition 3.1. 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: $a_1, \ldots, a_m$ such that $d(a_i, a_j) > 2\ell$ for $i \neq j$ and $\mathfrak{A}{\restriction}N^\ell(a_i), a_i \models \psi$ for all i.

The desired local approximations to elementary equivalence are the following equivalences $\equiv^{(\ell)}_{q,n}$:

Definition 3.2. $\mathfrak{A}, a \equiv^{(\ell)}_{q,n} \mathfrak{B}, b$ if

(i) $\mathfrak{A}{\restriction}N^\ell(a), a \equiv_q \mathfrak{B}{\restriction}N^\ell(b), b$, i.e., a and b are indistinguishable in their respective ℓ-neighbourhoods by FO-formulae of quantifier rank q, and

(ii) $\mathfrak{A}$ 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 $\psi(x) \in \mathrm{FO}$ 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 [8, 6] 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.

A first illustration of the idea of locality-based upgrading is provided in the following treatment of the class of all (finite and infinite) rooted trees (trees in the graph theoretic sense, cf. Definition 2.8) and arbitrary rooted frames. Recall the boosting of multiplicities, specifically for trees, as discussed before Lemma 2.15 above.

3.1 Trees and rooted frames

Lemma 3.3. *For rooted trees $\mathfrak{A}, a$ and $\mathfrak{B}, b$ and $\ell, q \geqslant 1$:*

$$\mathfrak{A}, a \sim_{\forall}^{2\ell+2} \mathfrak{B}, b \quad \Rightarrow \quad (\mathfrak{A} \otimes (q+1))^*_a, a \equiv^{(\ell)}_{q,q+1} (\mathfrak{B} \otimes (q+1))^*_b, b.$$

Before proving the lemma, let us see how it applies to characterise FO/$\sim$ over the class of all trees. Let $\varphi(x) \in$ FO be $\sim$ invariant over the class of all rooted trees and let q and ℓ be such that φ is preserved under $\equiv^{(\ell)}_{q,q+1}$. We claim that φ must be $\sim_{\forall}^{2\ell+2}$ invariant over the class of all rooted trees, and hence equivalent to a formula of $\mathrm{ML}^{\forall}$ of nesting depth $2\ell + 2$.

Indeed, if $\mathfrak{A}, a \sim_{\forall}^{2\ell+2} \mathfrak{B}, b$, then $(\mathfrak{A} \otimes (q+1))^*_a, a \equiv^{(\ell)}_{q,q+1} (\mathfrak{B} \otimes (q+1))^*_b, b$ by the lemma. Therefore

$$\begin{array}{rcll} \mathfrak{A}, a \models \varphi & \Leftrightarrow & (\mathfrak{A} \otimes (q+1))^*_a, a \models \varphi & ((\mathfrak{A} \otimes (q+1))^*_a, a \sim \mathfrak{A}, a, \text{ both rooted trees}) \\ & \Leftrightarrow & (\mathfrak{B} \otimes (q+1))^*_b, b \models \varphi & (\varphi \text{ preserved under } \equiv^{(\ell)}_{q,q+1}) \\ & \Leftrightarrow & \mathfrak{B}, b \models \varphi. & ((\mathfrak{B} \otimes (q+1))^*_b, b \sim \mathfrak{B}, b, \text{ both rooted trees}) \end{array}$$

In fact, this argument immediately extends to arbitrary rooted frames, because every rooted frame $\mathfrak{A}, a$ is bisimilar to its tree unravelling $\mathfrak{A}^*, a$, which is a rooted tree. Also, the transformation $\mathfrak{A}, a \mapsto (\mathfrak{A} \otimes (q+1))^*_a, a$ maps finite rooted trees to finite rooted trees. So the lemma proves the following.

Theorem 3.4. $\mathrm{ML}^{\forall}$ *captures bisimulation invariant first-order logic over*

(a) *the class of all rooted trees.*

(b) *the class of all rooted frames.*

(c) *the class of all finite rooted trees.*

Proof of Lemma 3.3. Let $\mathfrak{A}, a \sim_{\forall}^{2\ell+2} \mathfrak{B}, b$, both structures rooted trees. By Lemma 2.15,

$$(\mathfrak{A} \otimes (q+1))^*_a \restriction N^{2\ell+2}(a), a \equiv_{q+1} (\mathfrak{B} \otimes (q+1))^*_b \restriction N^{2\ell+2}(b), b$$

via a strategy that preserves distances. Therefore in particular also

$$(\mathfrak{A} \otimes (q+1))^*_a \restriction N^{\ell}(a), a \equiv_{q+1} (\mathfrak{B} \otimes (q+1))^*_b \restriction N^{\ell}(b), b.$$

In order to show that both structures satisfy the same FO formulae of quantifier rank q in ℓ-neighbourhoods, we first observe that any ℓ-neighbourhood of an element $a' \in (\mathfrak{A} \otimes (q+1))^*_a$ is contained within the 2ℓ-neighbourhood of the root in a subtree rooted at $a'' \in (\mathfrak{A} \otimes (q+1))^*_a$ for a suitable choice of a''. For any such a'', we find – through $\mathfrak{A} \sim^{2\ell}_{\forall} \mathfrak{B}$ – a matching b'' such that $(\mathfrak{A} \otimes (q+1))^*_a, a'' \sim^{2\ell} (\mathfrak{B} \otimes (q+1))^*_b, b''$.

It follows with Lemma 2.15 that the 2ℓ-neighbourhoods of the roots a'' and b'' in their respective subtrees are equivalent in the sense of $\equiv_{q+1}$ in a distance preserving manner. If b' is a suitable response to player I's first move to a' in the q-round game on these subtrees, it follows that $(\mathfrak{A} \otimes (q+1))^*_a{\restriction}N^\ell(a'), a' \equiv_q (\mathfrak{B} \otimes (q+1))^*_b{\restriction}N^\ell(b'), b'$.

Finally, to deal with the condition on scattered sets, we observe that it suffices to show the following, for all $k \leqslant \ell$ and $m \leqslant q$:

(i) for every $\psi = \psi(x) \in \mathrm{FO}$ of quantifier-rank up to q, if $a' \in (\mathfrak{A} \otimes (q+1))^*_a \setminus N^k(a)$ is such that $(\mathfrak{A} \otimes (q+1))^*_a{\restriction}N^k(a'), a' \models \psi$, then there is a corresponding element $b' \in (\mathfrak{B} \otimes (q+1))^*_b \setminus N^k(b)$ such that $(\mathfrak{B} \otimes (q+1))^*_b{\restriction}N^k(b'), b' \models \psi$, and vice versa.

(ii) for every $b' \in (\mathfrak{B} \otimes (q+1))^*_b \setminus N^k(b)$, there are q more elements $b'_1, \ldots, b'_q$ such that $\{b', b'_1, \ldots, b'_q\}$ is a k-scattered set of $q+1$ elements and $(\mathfrak{B} \otimes (q+1))^*_b{\restriction}N^k(b_i), b'_i \simeq (\mathfrak{B} \otimes (q+1))^*_b{\restriction}N^k(b'), b'$ for $i = 1, \ldots, q$. [Similarly in $(\mathfrak{A} \otimes (q+1))^*_a$.]

On the basis of (i) and (ii), it is clear that $(\mathfrak{A} \otimes (q+1))^*_a$ and $(\mathfrak{B} \otimes (q+1))^*_b$ agree on k-scattered subsets as required. For size one subsets ($m = 1$) this was already dealt with above. For $m \geqslant 2$ we note that at least one of two disjoint k-neighbourhoods must be disjoint from the root, whence (i) and (ii) apply to show there are corresponding k-scattered sets of size $m = q+1$ in both structures.

For (i): We see from (the proof of) Lemma 2.15 that the $\sim^{2k+1}$-type of $a'' \in (\mathfrak{A} \otimes (q+1))^*_a$ determines whether there is a directed path of length $k+1$ to some a' whose k-neighbourhood satisfies ψ. In other words, there is a formula $\chi \in \mathrm{ML}_{2k+1}$ such that $(\mathfrak{A} \otimes (q+1))^*_a, a'' \models \chi$ iff there is a directed path of length $k+1$ to some a' whose k-neighbourhood satisfies ψ, and similarly in $(\mathfrak{B} \otimes (q+1))^*_b$. Now let $a' \in (\mathfrak{A} \otimes (q+1))^*_a \setminus N^k(a)$ be such that its k-neighbourhood satisfies ψ. Then there is some a'' such that $(\mathfrak{A} \otimes (q+1))^*_a, a'' \models \chi$. Since $(\mathfrak{A} \otimes (q+1))^*_a \sim^{2\ell+1}_{\forall} (\mathfrak{B} \otimes (q+1))^*_b$, there is some b'' such that and $(\mathfrak{B} \otimes (q+1))^*_b, b'' \models \chi$, whence we find some $b' \in (\mathfrak{B} \otimes (q+1))^*_b \setminus N^k(b)$ whose k-neighbourhood satisfies ψ.

For (ii): Look at some b'' such that the given k-neighbourhood of b' is fully contained in the $2k$-neighbourhood of the subtree rooted at b''. As b'' is distinct from the root b of $(\mathfrak{B} \otimes (q+1))^*_b$, it is the root of one of $q+1$ isomorphic and disjoint sibling subtrees, in which we find the desired q further isomorphic copies of $(\mathfrak{B} \otimes (q+1))^*_b{\restriction}N^k(b')$. $\square$

3.2 Finite rooted frames

In this section we focus on the finite model theory variant of Theorem 3.4, whose proof is similar in spirit but somewhat more involved than that for Theorem 3.4 where we could use unravellings into infinite trees.

Theorem 3.5. *Over the class of all finite rooted frames, bisimulation invariance of a first-order formula is captured by* $\mathrm{ML}^{\forall}$*:*

$$\mathrm{FO}/{\sim} = \mathrm{FO}/{\sim_{\forall}} \equiv \mathrm{ML}^{\forall} \quad \textit{over finite rooted frames.}$$

For this section we restrict ourselves to finite frames, but point out that all arguments equally apply in the case of infinite frames. 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 $\mathfrak{A} = (A, R^{\mathfrak{A}})$ and 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 [13] is necessary in the context of (finite) rooted structures. It shows that within this class, $\sim_{\forall}^{2\ell}$ can be upgraded to $\approx^{\ell}$ in $\sim_{\forall}$ equivalent companion structures.

Lemma 3.6. *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 rooted trees.*

For the proof we need some auxiliary notions and an argument adapted from [13]. We let $\mathrm{tp}^{\ell}_{\mathfrak{A}}(a)$ denote the ℓ-bisimulation type ($\sim^{\ell}$-type) of a in $\mathfrak{A}$. Semantically, $\mathrm{tp}^{\ell}_{\mathfrak{A}}(a)$ precisely determines the $\sim^{\ell}$ equivalence class of $\mathfrak{A}, a$ and is definable by a formula of ML^{ℓ}. For any fixed finite vocabulary, $\sim^{\ell}$ has finite index. We similarly write $\mathrm{tp}_{\mathfrak{A}}(a)$ for the full bisimulation type of a. With a directed path $a_0, \dots, a_k$ in $\mathfrak{A}$ we associate the string consisting of the ℓ-bisimulation types $(\mathrm{tp}^{\ell}_{\mathfrak{A}}(a_i))_{0 \leqslant i \leqslant \ell}$ along this path.

Definition 3.7. A string $\mathrm{tp}^{\ell}_{\mathfrak{A}}(a_0), \dots, \mathrm{tp}^{\ell}_{\mathfrak{A}}(a_k)$ associated with a directed path $a_0, \dots, a_k$ is an *ℓ-history* of $a = a_k$ if either $k = \ell$ (we refer to a *proper ℓ-history*), or $k < \ell$ and the path is not backward extendible, i.e., a_0 has in-degree zero (we refer to a *short ℓ-history*).

Note also that for $\ell' \leqslant \ell$, the ℓ-histories of a node determine its ℓ'-histories.

In general a node may have several different ℓ-histories, depending on the choice of directed paths leading up to it. Nodes in trees, however, always have unique ℓ-histories. Whenever a' has a unique ℓ-history, we denote it as $\mathrm{hist}^{\ell}(a')$.

Proof of Lemma 3.6. Let $\mathfrak{A}, a \sim_{\forall}^{2\ell} \mathfrak{B}, b$ be finite, rooted structures; $\mathfrak{A}^*, a$ and $\mathfrak{B}^*, b$ their tree unravellings from the roots. Note that $\mathfrak{A}^*, a \sim_{\forall} \mathfrak{A}, a$ and $\mathfrak{B}^*, b \sim_{\forall} \mathfrak{B}, b$, whence in particular $\mathfrak{A}^*, a \sim_{\forall}^{2\ell} \mathfrak{B}^*, b$. However, $\mathfrak{A}^*, a$ and $\mathfrak{B}^*, b$ will typically be infinite. We want to replace $\mathfrak{A}^*, a$ and $\mathfrak{B}^*, b$ by globally bisimilar companions $\hat{\mathfrak{A}}, a$ and $\hat{\mathfrak{B}}, b$ that are finite but still guarantee unique ℓ-histories in all their nodes.

We explicitly construct $\hat{\mathfrak{A}}, a$ from $\mathfrak{A}^*, a$; then $\hat{\mathfrak{B}}, b$ is analogously obtained from $\mathfrak{B}^*, b$.

Let $H^{\ell}(\mathfrak{A})$ be the finite set of all proper ℓ-histories realised in $\mathfrak{A}$, which is the same as the set of all ℓ-histories of nodes at depths $\geqslant \ell$ in $\mathfrak{A}^*$. Let $T(\mathfrak{A})$ the set of all $\sim$ bisimulation types realised in $\mathfrak{A}$ (or in $\mathfrak{A}^*$), finite because $\mathfrak{A}$ is finite.

With a node a' of $\mathfrak{A}^*$ whose distance from the root is at least ℓ, associate the pair

$$\eta(a') = (\mathrm{hist}^{\ell}_{\mathfrak{A}^*}(a'), \mathrm{tp}_{\mathfrak{A}^*}(a')) \in H(\mathfrak{A}) \times T(\mathfrak{A}).$$

Note that there are only finitely many distinct η-values. We may therefore choose an $m > \ell$ such that all η-values that are realised in $\mathfrak{A}^*$ at nodes outside $N^{\ell}(a)$ are realised in $N^m(a) \setminus N^{\ell}(a)$. We let $\hat{\mathfrak{A}}$ be an extension of $\mathfrak{A}^*{\restriction}N^m(a)$ with extra edges to replace those that are cut off in this truncation. If a' is a leaf node in $\mathfrak{A}^*{\restriction}N^m(a)$ that has an R-successor a'' in $\mathfrak{A}^*$, let a''' be a node in $N^m(a) \setminus N^{\ell}(a)$ with $\eta(a''') = \eta(a'')$ and put an R-edge from a' to a'''. The operation of identifying a'' with a''' is compatible with ℓ-histories and with bisimulation types in $\mathfrak{A}^*$.

Clearly $\hat{\mathfrak{A}}, a \sim_{\forall} \mathfrak{A}^*, a \sim_{\forall} \mathfrak{A}, a$ and $\hat{\mathfrak{A}}{\restriction}N^\ell(a), a$ is a tree with root a. By construction $\hat{\mathfrak{A}}, a$ is rooted and all nodes in $\hat{\mathfrak{A}}$ have unique ℓ-histories.

If $\hat{\mathfrak{B}}$ is analogously obtained from $\mathfrak{B}^*, b$, it follows that $\hat{\mathfrak{A}}, a \sim_{\forall}^{2\ell} \hat{\mathfrak{B}}, b$. Now put, for $k \leqslant \ell$,

$$Z_k := \big\{(a', b') \colon \hat{\mathfrak{A}}, a' \sim^{\ell+k} \hat{\mathfrak{B}}, b' \;\text{ and }\; \mathrm{hist}^k(a') = \mathrm{hist}^k(b')\big\}.$$

We claim that $(Z_k)_{k \leqslant \ell} \colon \hat{\mathfrak{A}}, a \approx^\ell \hat{\mathfrak{B}}, b$.

Clearly $(a, b) \in Z_\ell$. The back-and-forth conditions with respect to forward and backward moves in the bisimulation game are easily verified. For this observe that, under the assumption of unique ℓ-histories, the $(k-1)$-histories of immediate predecessors of nodes that have the same k-history will always have the same $(k-1)$-history, for $k \leqslant \ell$. The same is true for immediate successors that are themselves $(k-1)$-bisimilar.

To argue for the global nature of the two-way ℓ-bisimulation, it suffices to exhibit $(\ell-1)$-bisimilar partner nodes of the same $(\ell-1)$-history in the opposite structure, for every node in $\hat{\mathfrak{A}}$ and $\hat{\mathfrak{B}}$. Consider for instance $a' \in \hat{A}$.

If $\mathrm{hist}^{\ell-1}(a')$ is proper, realised in $\hat{\mathfrak{A}}$ along the path $a_0, \ldots, a_{\ell-1} = a'$, we find b_0 such that $\hat{\mathfrak{A}}, a_0 \sim^{2\ell-1} \hat{\mathfrak{B}}, b_0$ since $\hat{\mathfrak{A}}, a \sim_{\forall}^{2\ell} \hat{\mathfrak{B}}, b$. By repeated application of the forth-property for $\sim^{2\ell-1}$ find b_i such that $\hat{\mathfrak{A}}, a_i \sim^{2\ell-(i+1)} \hat{\mathfrak{B}}, b_i$, for $i < \ell$. Then $a' = a_{\ell-1}$ and $b' := b_{\ell-1}$ have the same $(\ell-1)$-history and are ℓ-bisimilar. So in particular $(a', b') \in Z_{\ell-1}$.

If $\mathrm{hist}^{\ell-1}(a')$ is short, realised by the path $a_0 = a, \ldots, a_k = a'$ where $k < \ell - 1$ we similarly find a matching node in $\hat{\mathfrak{B}}$ following a sequence $b_0 = b, \ldots, b_{\ell-1}$ in $\hat{\mathfrak{B}}$ for which $\hat{\mathfrak{A}}, a_i \sim^{2\ell-i} \hat{\mathfrak{B}}, b_i$ for $i \leqslant k$. For $b' := b_k$ this implies that both $\mathrm{hist}^{\ell-1}(b') = \mathrm{hist}^{\ell-1}(a')$ and $\hat{\mathfrak{A}}, a' \sim^{2\ell-k} \hat{\mathfrak{B}}, b'$, whence in particular $\hat{\mathfrak{A}}, a' \sim^{\ell-1} \hat{\mathfrak{B}}, b'$ as required for $(a', b') \in Z_{\ell-1}$. ☐

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 Proposition 2.20. 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 3-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. This obstacle is due to existing short cycles at the root; the following lemma shows that this is the only obstacle.

Lemma 3.8. *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 $(2k+1)$-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 positive multiple of those in a.*

Proof. Let $\tilde{\pi} \colon \tilde{\mathfrak{A}} \to \mathfrak{A}$ be a faithful $(2k+1)$-acyclic bisimilar cover of $\mathfrak{A}$ as given by Proposition 2.20. Let $\hat{\mathfrak{A}}$ be the result of identifying in $\tilde{\mathfrak{A}}$ all nodes in $\tilde{\pi}^{-1}(a)$ to form a new node $\hat{a}$. We denote the new projection map, which sends $\hat{a}$ to a and otherwise agrees with $\tilde{\pi}$, as $\pi \colon \hat{\mathfrak{A}} \to \mathfrak{A}$.

Clearly $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ is a bisimilar cover.

By Observation 2.19, $\tilde{\mathfrak{A}} \restriction \tilde{\pi}^{-1}(N^k(a))$ is a disjoint union of isomorphic copies of $\mathfrak{A}\restriction N^k(a)$. It follows that $\hat{\mathfrak{A}}$ is still $(2k+1)$-acyclic.

Clearly $\hat{\mathfrak{A}}, \hat{a}$ is also rooted: if $\hat{a}' \neq \hat{a}$ is any other element of $\hat{\mathfrak{A}}$, let $\pi(\hat{a}') = a'$ and consider a path from a to a' in $\mathfrak{A}$. This path has a unique lift in $\tilde{\mathfrak{A}}$ to a path from some element of $\tilde{\pi}^{-1}(a)$ to $\hat{a}'$, which connects $\hat{a}$ to $\hat{a}'$ in $\hat{\mathfrak{A}}$.

The resulting cover is not in general faithful. Incidence degrees at $\hat{a}$ in $\hat{\mathfrak{A}}$ are m times those at a in $\mathfrak{A}$, where m is the cardinality of $\tilde{\pi}^{-1}(a)$ in $\tilde{\mathfrak{A}}$. In all nodes apart from $\hat{a}$, $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ is still faithful. $\square$

Suppose $\mathfrak{A}$ and $\mathfrak{B}$ are k-acyclic for $k = 2\ell + 1$ 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 acyclic structures in such a way that player II in the q-round FO Ehrenfeucht-Fraïssé game can always find matching elements on fresh paths where player I does – similar to the reasoning that supported Lemma 2.15. Because of the additional restrictions imposed on locally acyclic covers by rootedness, the analogue of Lemma 3.3 will have to work with rooted structures in which there are two regions of rather different local behaviour:

- a region close to the root that is a tree, where nodes have in-degree 1 and out-going multiplicities are boosted by q;
- a region beyond this, where in- and out-going multiplicities are boosted by q.

In order to account for the way in which local two-way bisimilarity implies local FO-equivalence up to quantifier rank q in such settings, we use the following notion of two-way ℓ-bisimilarity with *counting up to q*, $\sim_{\leftrightarrows}^{\ell;q}$.

Definition 3.9. $(Z_k)_{0\leqslant k\leqslant \ell}\colon \mathfrak{A} \sim_{\leftrightarrows}^{\ell;q} \mathfrak{B}$ if $(Z_k)_{0\leqslant k\leqslant \ell}\colon \mathfrak{A} \sim_{\leftrightarrows}^{\ell} \mathfrak{B}$ is a back-and-forth system (in the sense of $\sim_{\leftrightarrows}^{\ell}$) satisfying 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 $(a_i, a) \in R^{\mathfrak{A}}$.

The global variant $\approx^{\ell;q}$ is similarly defined.

Note that $\sim_{\leftrightarrows}^{\ell;q}$ and $\approx^{\ell;q}$ are automatically preserved in faithful covers, as these do not affect multiplicities. The proof of the following is then entirely analogous to the proof of Lemma 2.15 (also compare Claim 26 in [13]).

Lemma 3.10. *If $\mathfrak{A}\restriction N^{\ell}(a)$ and $\mathfrak{B}\restriction N^{\ell}(b), b$ are acyclic and $\mathfrak{A}, a \sim_{\leftrightarrows}^{\ell;q} \mathfrak{B}, b$, then $\mathfrak{A}\restriction N^{\ell}(a), a \equiv_q \mathfrak{B}\restriction N^{\ell}(b), b$.*

Corollary 3.11. *If $\mathfrak{A}, a$ and $\mathfrak{B}, b$ are k-acyclic for $k = 2\ell + 1$, then $\mathfrak{A}, a \approx^{\ell+1;q} \mathfrak{B}, b$ implies $\mathfrak{A}, a \equiv_{q,1}^{(\ell)} \mathfrak{B}, b$.*

To upgrade $\approx^\ell$ to $\approx^{\ell;q}$ it essentially suffices to boost all multiplicities that are greater than 1 by q. However, the acyclic neighbourhood of the root requires special treatment.

Lemma 3.12. *Let $\ell > 0$. Let $\mathfrak{A}, a \approx^{2\ell+2} \mathfrak{B}, b$ where $\mathfrak{A}, a$ and $\mathfrak{B}, b$ are both finite and rooted, and $\mathfrak{A}{\restriction}N^\ell(a), a$ and $\mathfrak{B}{\restriction}N^\ell(b), b$ are rooted trees.*

Then there are finite, rooted, $(2\ell+1)$-acyclic $\hat{\mathfrak{A}}, \hat{a}$ and $\hat{\mathfrak{B}}, \hat{b}$ such that

(i) *$\hat{\mathfrak{A}}, \hat{a} \approx \mathfrak{A}, a$ and $\hat{\mathfrak{B}}, \hat{b} \approx \mathfrak{B}, b$.*
(ii) *$\hat{\mathfrak{A}}{\restriction}N^\ell(\hat{a}), \hat{a}$ and $\hat{\mathfrak{B}}{\restriction}N^\ell(\hat{b}), \hat{b}$ are rooted trees.*
(iii) *$\hat{\mathfrak{A}}, N^\ell(\hat{a}), \hat{a} \approx^{\ell+1;q} \hat{\mathfrak{B}}, N^\ell(\hat{b}), \hat{b}$ (we consider the ℓ-neighbourhoods of the roots as marked by new unary predicates interpreted as $N^\ell(\hat{a})$ and $N^\ell(\hat{b})$ as indicated).*

Proof. We concentrate on how to obtain $\hat{\mathfrak{A}}$ from $\mathfrak{A}$.

Let $\mathfrak{A}_0 := \mathfrak{A}{\restriction}N^\ell(a)$ and $\mathfrak{A}_1 := \mathfrak{A}{\restriction}(A \setminus N^\ell(a))$.

Let $\tilde{\mathfrak{A}}_0 := (\mathfrak{A}_0 \otimes q)^*_a$ be the rooted tree obtained from $\mathfrak{A}, a$ by boosting multiplicities as described in section 2.5. Every proper subtree of $\tilde{\mathfrak{A}}_0$ has q isomorphic sibling subtrees. Let $\tilde{\pi} \colon \tilde{\mathfrak{A}}_0 \to \mathfrak{A}_0$ be the natural projection.

Let $\tilde{\mathfrak{A}}_1 := \mathfrak{A}_1 \otimes q$ be the structure obtained by boosting multiplicities in $\mathfrak{A}_1$ by q. We also denote by $\tilde{\pi}$ the natural projection $\tilde{\pi} \colon \tilde{\mathfrak{A}}_1 \to \mathfrak{A}_1$.

$\tilde{A}$ is obtained from the disjoint union of $\tilde{\mathfrak{A}}_0$ and $\tilde{\mathfrak{A}}_1$ by joining $a' \in \tilde{\mathfrak{A}}_0$ and $a'' \in \tilde{\mathfrak{A}}_1$ exactly as ($\tilde{\pi}(a')$ and $\tilde{\pi}(a'')$ are joined in $\mathfrak{A}$. We speak of region 0 and region 1 to distinguish the two parts $\tilde{\mathfrak{A}}_i$ for $i = 0, 1$. Note that only nodes at distance ℓ from the root in region 0 are linked to nodes in region 1. In particular $\tilde{A}{\restriction}N^\ell(a), a \simeq \tilde{A}_0{\restriction}N^\ell(a), a$ is a rooted tree. Clearly $\tilde{A}, a \approx \mathfrak{A}, a$.

If $\tilde{B}$ is obtained from $\mathfrak{B}$ in the same manner, it is also clear that $\tilde{\mathfrak{A}}, a \approx^{2\ell+2} \tilde{\mathfrak{B}}, b$.

That, indeed, we have $\tilde{\mathfrak{A}}, N^\ell(a), a \approx^{\ell+1;q} \tilde{\mathfrak{B}}, N^\ell(b), b$, is shown as follows.

Suppose a', b' are such that $\tilde{\mathfrak{A}}, a' \sim_{\leftrightarrows}^{\ell+1+k} \tilde{\mathfrak{B}}, b'$ for some $0 < k \leqslant \ell$. Then obviously the stronger back-and-forth conditions in the sense of $\sim_{\leftrightarrows}^{\ell+1+k;q}$ are satisfied at (a', b') with respect to forward R-edges (multiplicities of outgoing R-edges have been boosted by q); they are also satisfied with respect to backward R-edges if both a' and b' are from region 0 (which implies that both nodes have in-degree 1, or 0 for the roots); and if both a' and b' are from region 1 (multiplicities of incoming R-edges have been boosted by q). It remains to observe that membership in the respective regions is determined by the $\sim_{\leftrightarrows}^{\ell+1}$-type of nodes: the nodes in region 1 are precisely those satisfying $\langle R^- \rangle^{\ell+1} \top$. It follows that $\tilde{\mathfrak{A}}, N^\ell(a), a \approx^{\ell+1;q} \tilde{\mathfrak{B}}, N^\ell(b), b$.

We now apply Lemma 3.8 to cover both $\tilde{\mathfrak{A}}$ and $\tilde{\mathfrak{B}}$ by rooted, $(2\ell+1)$-acyclic structures $\hat{A}$ and $\hat{\mathfrak{B}}$ such that these covers are faithful outside the roots $\hat{a}$ and $\hat{b}$, and have multiplicities boosted by some factor in the roots. Clearly $\hat{\mathfrak{A}}, \hat{a} \approx \mathfrak{A}, a$ and $\hat{\mathfrak{B}}, \hat{b} \approx \mathfrak{B}, b$. Because multiplicities are preserved (or boosted further at the roots) in these covers, $\tilde{\mathfrak{A}}, a \approx^{\ell+1;q} \tilde{\mathfrak{B}}, b$ implies that $\hat{\mathfrak{A}}, \hat{a} \approx^{\ell+1;q} \hat{\mathfrak{B}}, \hat{b}$. Because of $(2\ell+1)$-acyclicity, the covers $\pi \colon \hat{\mathfrak{A}}, \hat{a} \approx \mathfrak{A}, a$ and $\pi \colon \hat{\mathfrak{B}}, \hat{b} \approx \mathfrak{B}, b$ preserve distances up to $\ell + 1$ from the roots, and therefore $\hat{\mathfrak{A}}, N^\ell(\hat{a}), \hat{a} \approx^{\ell+1;q} \hat{\mathfrak{B}}, N^\ell(\hat{b}), \hat{b}$ follows. ∎

We may now combine the techniques of Lemma 3.6 and Lemma 3.12 with Lemma 3.10 and Corollary 3.11 to upgrade $\sim_{\forall}^{4\ell+4}$ to $\equiv_{q,n}^{(\ell)}$ for any desired q and n. In the light of Proposition 2.13 this proves Theorem 3.5. Figure 1 illustrates the underlying chain of upgradings which is discussed in the proof of the following lemma.

$$
\begin{tikzcd}[column sep=large, row sep=large]
\mathfrak{A}, a \arrow[r, dash, "\sim^{4\ell+4}"] \arrow[d, dash, "\sim_{\forall}"] & \mathfrak{B}, b \arrow[d, dash, "\sim_{\forall}"] & \text{finite, rooted} \arrow[d, "\substack{\text{history-sensitive}\\ \text{partial unfolding, L3.6}}"] \\
\hat{\mathfrak{A}}_1, \hat{a}_1 \arrow[r, dash, "\approx^{2\ell+2}"] \arrow[d, dash, "\approx"] & \hat{\mathfrak{B}}_1, \hat{b}_1 \arrow[d, dash, "\approx"] & \substack{\text{finite, rooted}\\ \text{tree-like to depth } \ell} \arrow[d, "\substack{\text{boosted multiplicities and}\\ \text{acyclic cover, L3.12 \& C3.11}}"] \\
\hat{\mathfrak{A}}_2, N^{\ell}(\hat{a}_2), \hat{a}_2 \arrow[r, dash, "\approx^{\ell+1;q}", "\equiv^{(\ell)}_{q,1}"'] \arrow[d, dash, "\approx"] & \hat{\mathfrak{B}}_2, N^{\ell}(\hat{b}_2), \hat{b}_2 \arrow[d, dash, "\approx"] & \substack{\text{finite, rooted, } (2\ell+1)\text{-acyclic,}\\ N^{\ell}(\cdot) \text{ tree-like and marked}} \arrow[d, "\text{glueing copies at root}"] \\
\hat{\mathfrak{A}}, \hat{a} \arrow[r, dash, "\equiv^{(\ell)}_{q,n}"] & \hat{\mathfrak{B}}, \hat{b} & \text{finite, rooted}
\end{tikzcd}
$$

Figure 1: the upgrading in Lemma 3.13.

Lemma 3.13. *Let $\mathfrak{A}, a \sim_{\forall}^{4\ell+4} \mathfrak{B}, b$ where $\mathfrak{A}, a$ and $\mathfrak{B}, b$ are both finite and rooted. Let $q, n \in \mathbb{N}$.*

Then there are finite, rooted $\hat{\mathfrak{A}}, \hat{a}$ and $\hat{\mathfrak{B}}, \hat{b}$ such that

(i) $\hat{\mathfrak{A}}, \hat{a} \sim_{\forall} \mathfrak{A}, a$ *and* $\hat{\mathfrak{B}}, \hat{b} \sim_{\forall} \mathfrak{B}, b$.

(ii) $\hat{\mathfrak{A}}, \hat{a} \equiv_{q,n}^{(\ell)} \hat{\mathfrak{B}}, \hat{b}$.

Proof. Using Lemma 3.6, we first obtain $\hat{\mathfrak{A}}_1, \hat{a}_1 \sim_{\forall} \mathfrak{A}, a$ and $\hat{\mathfrak{B}}_1, \hat{b}_1 \sim_{\forall} \mathfrak{B}, b$ such that $\hat{\mathfrak{A}}_1, \hat{a}_1 \approx^{2\ell+2} \hat{\mathfrak{B}}_1, \hat{b}_1$ and such that the ℓ-neighbourhoods of the roots are rooted trees. With Lemma 3.12 we then further obtain from these $\hat{\mathfrak{A}}_2, \hat{a}_2 \sim_{\forall} \mathfrak{A}, a$ and $\hat{\mathfrak{B}}_2, \hat{b}_2 \sim_{\forall} \mathfrak{B}, b$ such that $\hat{\mathfrak{A}}_2, N^\ell(\hat{a}_2), \hat{a}_2 \approx^{\ell+1;q} \hat{\mathfrak{B}}_2, N^\ell(\hat{b}_2), \hat{b}_2$.

Let now $\hat{\mathfrak{A}}, \hat{a}$ be the result of gluing n disjoint copies of $\hat{\mathfrak{A}}_2, \hat{a}_2$ in their roots, and similarly for $\hat{\mathfrak{B}}$. Clearly $\hat{\mathfrak{A}}, \hat{a} \sim_{\forall} \mathfrak{A}, a$, $\hat{\mathfrak{B}}, \hat{b} \sim_{\forall} \mathfrak{B}, b$ and $\hat{\mathfrak{A}}, N^\ell(\hat{a}), \hat{a} \approx^{\ell+1;q} \hat{\mathfrak{B}}, N^\ell(\hat{b}), \hat{b}$.

We claim that $\hat{\mathfrak{A}}, \hat{a} \equiv_{q,n}^{(\ell)} \hat{\mathfrak{B}}, \hat{b}$.

It is immediate from Corollary 3.11 that $\hat{\mathfrak{A}} \restriction N^\ell(\hat{a}), \hat{a} \equiv_q \hat{\mathfrak{B}} \restriction N^\ell(\hat{b}), \hat{b}$. For the condition on k-scattered sets of sizes up to n we note that – similar to the reasoning in the proof of Lemma 3.3 – the existence of two disjoint k-neighbourhoods for some formula $\psi(x)$ in either structure implies that there are at least n disjoint pairwise isomorphic k-neighbourhoods for ψ in each structure. Let for instance $a' \in \hat{\mathfrak{A}}$ be such that $\hat{a} \notin N^k(a')$. Picking $b' \in \hat{\mathfrak{B}}$ such that $\hat{\mathfrak{A}}, N^\ell(\hat{a}), a' \approx^{\ell;q} \hat{\mathfrak{B}}, N^\ell(\hat{b}), b'$, it is clear that $d(\hat{b}, b') > k$, i.e., that $\hat{b} \notin N^\ell(b')$. Therefore, $N^k(b')$ is entirely contained within one copy of $\hat{\mathfrak{B}}_2 \setminus \{\hat{b}_2\}$, just like $N^k(a')$ is within one copy of $\hat{\mathfrak{A}}_2 \setminus \{\hat{a}_2\}$. Therefore each of the k-neighbourhoods is one of a family of n many disjoint isomorphic neighbourhoods in its structure. $\square$

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.

3.3 Equivalence frames

We consider Kripke structures over *equivalence frames*, or *equivalence structures*, cf. Definition 2.9. These are structures of the form $\mathfrak{A} = (A, E_1, \ldots, E_m, P_1, \ldots, P_n)$ with binary E_i and unary P_j such that each E_i is an equivalence relation over A. By symmetry of the accessibility relations, $\sim_{\leftrightharpoons}$ 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 rooted (which here is the same as connected) equivalence frames.

Theorem 3.14. $\mathrm{FO}/{\sim_{\forall}} = \mathrm{FO}/{\approx} \equiv \mathrm{ML}^{\forall}$ *over (finite) equivalence structures.*

Compare Theorem 2.12 for the classical version. We now 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 corresponding to different equivalence relations.

Acyclicity criteria can only apply at the level of equivalence classes: any equivalence class being a clique, short cycles within individual classes cannot be ruled out.

Definition 3.15. Let $k \geqslant 2$. An equivalence structure $\mathfrak{A}$ is called *k-acyclic* if every cycle of length up to k lies within a single E_i-class for some i. Equivalently, every labelled cycle $a_0, E_{i_0}, a_1, \dots, E_{i_{n-1}}, a_n$ that is non-trivial with respect to equivalence classes in the sense that $a_{j+1} \neq a_j$ and $E_{i_{j+1}} \neq E_{i_j}$ for all $j \in \mathbb{Z}/n\mathbb{Z}$, must be of a length $n > k$.

Note that 2-acyclicity implies that any two equivalence classes with respect to distinct E_i and E_j intersect in at most one element.

The following serves as an analogue of Proposition 2.20 for finite equivalence frames: finite covers by k-acyclic structures are available for these as well. Recall that faithfulness means that in- and out-degrees are preserved for each individual binary relation separately.

Proposition 3.16. *Let $k \geqslant 2$. Any finite equivalence structure $\mathfrak{A}$ possesses a faithful bisimilar cover $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ by a finite k-acyclic equivalence structure $\hat{\mathfrak{A}}$. Moreover the cover can be chosen such that each equivalence class of $\hat{\mathfrak{A}}$ is an isomorphic copy of an equivalence class of $\mathfrak{A}$.*

For the proof of the proposition and further arguments towards an upgrading of local bisimilarity towards first-order equivalence, we look at a simple interpretation scheme. These interpretations provide translations between equivalence frames and associated non-transitive frames, in which locality arguments can be used more directly.

Let $\mathfrak{A} = (A, E_1, \dots, E_m, P_1, \dots, P_n)$ be an equivalence frame. We denote by $[a]_i$ the $E_i^{\mathfrak{A}}$-equivalence class of $a \in A$, and by E/E_i the quotient of A with respect to E_i. For the following we want to regard the quotients with respect to distinct E_i and E_j as disjoint (of distinct sorts), even if they may not be when regarded as subsets of the power set of A.

With $\mathfrak{A}$ associate the Kripke structure $\mathfrak{A}^+ = (A^+, R, U, Q_1, \dots, Q_m, P_1, \dots, P_n)$ with new binary R, new unary U and one new unary Q_i for each E_i as follows:

$$\begin{array}{rcl}
A^+ &=& A \,\dot\cup\, \dot{\bigcup}_{1 \leqslant i \leqslant m} A/E_i, \\
U^{\mathfrak{A}^+} &=& A, \\
P_j^{\mathfrak{A}^+} &=& P_j^{\mathfrak{A}}, \\
Q_i^{\mathfrak{A}^+} &=& A/E_i, \\
R^{\mathfrak{A}^+} &=& \bigcup_{1 \leqslant i \leqslant m} \big\{ (\alpha, a)\colon a \in \alpha \in A/E_i \big\}.
\end{array}$$

It is clear that $\mathfrak{A}$ is first-order interpretable within $\mathfrak{A}^+$: the nodes of $\mathfrak{A}$ are precisely the elements in U and there is an E_i-edge between two such if there is some element in Q_i to which both are joined by an R-edge. The class of Kripke structures obtained in this manner is easily characterised by the (first-order) conditions that $U, Q_1, \dots, Q_m$ partition the universe, that $P_j \subseteq U$, and that

$$(*) \quad \begin{array}{l}
R \subseteq \bigcup_i Q_i \times U, \text{ and for } 1 \leqslant i \leqslant m\text{:} \\
\text{every } x \in U \text{ has precisely one } R\text{-predecessor in } Q_i, \text{ and} \\
\text{every } x \in Q_i \text{ has at least one } R\text{-successor.}
\end{array}$$

We note that these conditions are preserved in faithful bisimilar covers (in fact they are even preserved under $\approx^{\ell;q}$ for $q, \ell \geqslant 2$).

Observation 3.17. *For equivalence structures $\mathfrak{A}, a$ and $\mathfrak{B}, b$:*

(i) $\mathfrak{A}, a \equiv_q \mathfrak{B}, b \;\Rightarrow\; \mathfrak{A}^+, a \equiv_q \mathfrak{B}^+, b$.

(ii) $\mathfrak{A}^+, a \equiv_{q+1} \mathfrak{B}^+, b \;\Rightarrow\; \mathfrak{A}, a \equiv_q \mathfrak{B}, b$.

(iii) $\mathfrak{A}, a \sim^\ell \mathfrak{B}, b \;\Leftrightarrow\; \mathfrak{A}^+, a \sim_{\leftrightarrows}^{2\ell} \mathfrak{B}^+, b$.

Proof. (i) and (ii) are obvious on the basis of natural FO-translations. For (iii) we observe first that there is an obvious simulation of a move along an E_i-edge from a' to a'' in the bisimulation game on $\mathfrak{A}$ by the succession of the two moves from a' to $[a']_i = [a'']_i$ and from $[a'']_i$ to a'' in $\mathfrak{A}^+$; this immediately implies $\mathfrak{A}^+, a \sim_{\leftrightarrows}^{2\ell} \mathfrak{B}^+, b \;\Rightarrow\; \mathfrak{A}, a \sim^\ell \mathfrak{B}, b$. For the converse, we observe that the bisimulation game on two structures satisfying $(*)$ will necessarily visit the U-parts and Q-parts in strict alternation and that the response by player II in moves from the U-part into the Q-part is always uniquely determined. One infers from this that a winning strategy for II in the 2ℓ round game on $\mathfrak{A}^+$ and $\mathfrak{B}^+, b$ is obtained by a simulation of a winning strategy in the ℓ-round game on $\mathfrak{A}$ and $\mathfrak{B}$ – by looking at every other move and positions in the U-parts. ☐

Regarding acyclicity, we note that non-trivial cycles of length n in $\mathfrak{A}$ correspond to non-trivial cycles of length $2n$ in $\mathfrak{A}^+$. The following is therefore obvious.

Observation 3.18. *For any equivalence structure $\mathfrak{A}$, $k \geqslant 2$: $\mathfrak{A}$ is k-acyclic (in the sense of Definition 3.15) if, and only if, $\mathfrak{A}^+$ is $2k$-acyclic (in the sense of Definition 2.7).*

Proof of Proposition 3.16. Let $\mathfrak{A}$ be a finite equivalence structure, $k \geqslant 2$. We pass to $\mathfrak{A}^+$ and, using Proposition 2.20, find a faithful bisimilar cover $\pi\colon (\mathfrak{A}^+)\hat{} \rightarrow \mathfrak{A}^+$ by some finite $2k$-acyclic $(\mathfrak{A}^+)\hat{}$. By faithfulness and global bisimilarity, $(\mathfrak{A}^+)\hat{}$ also satisfies $(*)$ and therefore is isomorphic to $\hat{\mathfrak{A}}^+$ for an equivalence structure $\hat{\mathfrak{A}}$. The homomorphism π induces a matching homomorphism $\pi\colon \hat{\mathfrak{A}} \rightarrow \mathfrak{A}$, which is easily seen to be a faithful bisimilar cover of $\mathfrak{A}$. For faithfulness observe that there is a strict correspondence between the degree of a with respect to E_i in $\mathfrak{A}$ and the degree of $[a]_i$ with respect to R in $\mathfrak{A}^+$ – both are equal to the size of $[a]_i$.

The isomorphism assertion of the proposition follows from the fact that the isomorphism type of $\mathfrak{A} \restriction [a]_i$ is fully determined by $\mathfrak{A}^+ \restriction N^1([a]_i)$, which is a tree of depth 1 rooted at $[a]_i$. Over such acyclic neighbourhoods of $\mathfrak{A}^+$, the cover $\pi\colon (\mathfrak{A}^+)\hat{} \rightarrow \mathfrak{A}^+$ is a union of isomorphisms and induces corresponding isomorphisms between the E_i classes of $\hat{\mathfrak{A}}$ and of $\mathfrak{A}$. ☐

Towards an upgrading to local FO equivalence recall the operation $\mathfrak{A} \mapsto \mathfrak{A} \otimes q$ for boosting multiplicities as discussed in section 2.5. The class of finite equivalence structures is closed under this operation, which introduces q indistinguishable copies of every node. Similarly, the class of finite equivalence structures 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 $n \geqslant 1$. The following is straightforward.

Observation 3.19. *Let $\mathfrak{A}, a$ and $\mathfrak{B}, b$ be finite equivalence structures, $\ell, q, n \in \mathbb{N}$.*

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

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

In analogy with Lemma 3.10 and Corollary 3.11 the following obtains.

Lemma 3.20. *Let $\ell, q, n \geqslant 1$. Let $\mathfrak{A}, a$ and $\mathfrak{B}, b$ be finite equivalence structures that are k-acyclic for $k = 2\ell + 1$. Then $\mathfrak{A}, a \approx^{\ell+1;q+1} \mathfrak{B}, b$ implies $\mathfrak{A}, a \equiv^{(\ell)}_{q,1} \mathfrak{B}, b$.*

Proof. Passing to $\mathfrak{A}^+$ and $\mathfrak{B}^+$ we find that these are $2k$-acyclic by Observation 3.18, and $\mathfrak{A}^+, a \approx^{2\ell+1;q+1} \mathfrak{B}^+, b$ by Observations 3.17 (iii) (our use of $\approx^{2\ell+1;q+1}$ rather than $\approx^{2\ell+2;q+1}$ reflects the fact that from nodes in the Q-parts we need an extra move back into the U-parts to relate these positions to positions over the original equivalence structures). By Corollary 3.11 therefore, $\mathfrak{A}^+, a \equiv^{(2\ell)}_{q+1,1} \mathfrak{B}^+, b$. By Observations 3.17 (ii), this implies $\mathfrak{A}, a \equiv^{(\ell)}_{q,1} \mathfrak{B}, b$. ☐

We are now ready for the upgrading that proves Theorem 3.14 for finite equivalence structures.

Lemma 3.21. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be finite equivalence structures with $\mathfrak{A}, a \sim^{\ell+1}_{\forall} \mathfrak{B}, b$. 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 $\approx$. By part (i) of Observation 3.19, $\mathfrak{A}_1, a := \mathfrak{A} \otimes (q+1), a$ and $\mathfrak{B}_1, b := \mathfrak{B} \otimes (q+1), b$ satisfy $\mathfrak{A}_1, a \approx^{\ell+1;q+1} \mathfrak{B}_1, b$. Passing to faithful bisimilar covers $\mathfrak{A}_2$ and $\mathfrak{B}_2$ according to Proposition 3.16 leads to bisimilar covers such that, by Lemma 3.20, $\mathfrak{A}_2, a_2 \equiv^{(\ell)}_{q,1} \mathfrak{B}_2, b_2$. Putting $\hat{\mathfrak{A}}, \hat{a} := n \times \mathfrak{A}_2, a_2$ and $\hat{\mathfrak{B}}, \hat{b} := n \times \mathfrak{B}_2, b_2$, we finally have $\hat{\mathfrak{A}}, \hat{a} \equiv^{(\ell)}_{q,n} \hat{\mathfrak{B}}, \hat{b}$ by part (ii) of Observation 3.19. ☐

Together with Proposition 2.13, this lemma proves Theorem 3.14.

4 Decomposition based techniques

As stated in Theorem 2.12, ML characterises bisimulation invariant first-order logic over various classes of rooted transitive frames. However, the methods used to establish those results in the sense of classical model theory are based on compactness and do not yield corresponding characterisations for *finite* rooted transitive frames. Indeed, as we shall see below, bisimulation invariant first-order logic over the latter class is characterised not by ML but by the proper extension ML* of ML, and establishing this requires a careful consideration of the definability of reflexive nodes and R-cliques. On the other hand, natural finite model theory analogues of the expected classical characterisations obtain for

- finite rooted irreflexive transitive trees (finite $\prec$-trees), and for
- finite rooted reflexive transitive trees (finite $\preccurlyeq$-trees).

These form an important stepping stone in the characterisation for finite rooted transitive frames. And, in both cases, entirely different methods from the classical ones are brought to bear. These methods are centred around what we call *decomposition* based techniques. That is, we decompose a transitive tree into an ordered sum of smaller transitive trees and reduce the problem to showing certain properties of finite linearly ordered structures (see [11] for a similar construction in the context of *monadic path logic*). Unlike the locality-based methods of section 3, the methods rely on finiteness properties, in particular finite predecessor sets in transitive trees. Indeed, the crucial Lemma 4.5 below is a property of finite linearly ordered structures that does not hold for infinite linear orders, whether dense or discrete. Finite predecessor sets, however, may, in all cases considered, be achieved by means of an unravelling (cf. Observation 2.17) – and thus we also get in these cases alternative proofs for the results concerning not necessarily finite transitive frames in Theorem 2.12. But for the classes of

- finite transitive tree-like, and, by extension,
- finite rooted transitive trees,

the finite model theory characterisations are not the straightforward analogues of the classical ones.[6] In particular, in order to extend the characterisation from finite (fully irreflexive or fully reflexive) transitive trees to transitive tree-like frames (which need neither be reflexive nor irreflexive), the target logic needs to be ML^*, strictly more expressive than plain ML, as it turns out that an account of bisimulation invariance over finite transitive tree-like frames requires the new modality $\Diamond^*$, which does not materialise in the classical picture. The following example gives an indication of the difference.

Example 4.1. Consider the following FO-property of R-frames:

$$\varphi(x) := \exists y\big(Rxy \wedge Ryy\big).$$

Obviously, φ is not bisimulation invariant over the class of all transitive tree-like frames: if $\mathfrak{A}, a$ is transitive tree-like, then $\mathrm{TC}(\mathfrak{A}^*_a)$ is an irreflexive transitive tree bisimilar to $\mathfrak{A}, a$ – so here φ is false irrespective of its truth value in $\mathfrak{A}, a$. However, if $\mathfrak{A} \models \varphi[a]$, then $\mathrm{TC}(\mathfrak{A}^*_a)$ is necessarily infinite. (Cf. Observation 2.17 for $\mathrm{TC}(\mathfrak{A}^*_a)$.)

Over the class of all *finite* transitive tree-like frames, however, φ is bisimulation invariant: any finite bisimilar companion of a structure with some reflexive R-edge will have to have an R-cycle, and, by transitivity, all nodes along an R-cycle have reflexive R-edges. Yet it is not hard to see that φ is not equivalent to any formula of ML over the class of finite transitive tree-like frames.

From now on we concentrate on transition systems with one accessibility relation R which is transitive. In sections 4.1 and 4.2 we primarily deal with irreflexive transitive R and write $\prec$ for R; in the special case where R is a strict linear ordering (section 4.1) we use $<$. When writing $\prec$ or $<$ for R in these circumstances, we also use $\preccurlyeq$ or $\leqslant$ to denote the reflexive closure of R, e.g., defining $a \preccurlyeq a'$ to mean ($a \prec a'$ or $a = a'$) as usual. Similarly, when dealing with reflexive transitive R in section 4.3, we write $\preccurlyeq$ for R and regard $\prec$ as defined by ($a \preccurlyeq a'$ and $a \neq a'$). In both scenarios we may think of $\prec$ and $\preccurlyeq$ as (irreflexive or reflexive) partial order relations corresponding to R but need to take care which of these two variants governs the semantics of $\Diamond$.

Only in section 4.5 shall we return to transitive R which need neither be reflexive nor irreflexive, and address the interesting difficulty indicated above. As long as we are

[6] It is in this respect that erroneous claims were made in [5], which are corrected here.

still dealing with tree-like frames, which may now have reflexive as well as irreflexive nodes, the simple new modality $\Diamond^*$ with first-order equivalent

$$(\Diamond^*\varphi)(x) \equiv \exists y\big(Rxy \wedge Ryy \wedge \varphi(y)\big)$$

inspired by the above example yields an extension of basic modal logic that is strong enough to capture FO/$\sim$ for instance over the class of all finite transitive tree-like structures.

When passing to finite transitive structures that may also have non-trivial cliques with respect to R, we need to consider the generalisation from $\Diamond^*$ to a family of new modalities $\Diamond^*_p$, one for every set p of propositional types with respect to the unary predicates P_j in the underlying vocabulary σ, as indicated in connection with Definition 2.3. That such further generalisation of $\Diamond^*$ becomes necessary is apparent from the following example.

Example 4.2. The following first-order formula is bisimulation invariant over the class of all finite transitive R-frames with a single unary predicate P:

$$\varphi(x) := \exists y \exists y'\big(Rxy \wedge Ryy' \wedge Ry'y \wedge Py \wedge \neg Py'\big).$$

That this formula is indeed bisimulation invariant over the class of all path-finite transitive R-frames can either be checked directly in an argument similar to the one indicated for Example 4.1 above, or can be inferred from Lemma 4.25 below. We here just note that φ is equivalent to the ML*-formula $\Diamond^*_p\top$ where p consists of the two propositional $\{P\}$-types defined by the formulae P and $\neg P$. We refer to section 4.4 for a thorough discussion of ML*.

4.1 Finite 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 [6]).

We introduce some notation before stating the lemma.

Definition 4.3. *Let σ be a vocabulary consisting of a binary relation symbol $<$ and unary symbols $P_1, \dots, P_k$. A σ-*word *is a finite σ-structure in which $<$ is interpreted as a strict linear order.*

As pointed out above, we write $a \leqslant b$ as an abbreviation for "$a < b$ or $a = b$." We write $\mathfrak{A} \lhd \mathfrak{B}$ for the *ordered sum* of $\mathfrak{A}$ and $\mathfrak{B}$. That is to say that $\mathfrak{A} \lhd \mathfrak{B}$ is the σ-word whose universe consists of the disjoint union of A and B, with the interpretation of the symbols P_i being inherited from $\mathfrak{A}$ and $\mathfrak{B}$ and with $a < b$ if either $a \in A$ and $b \in B$ or $a <^{\mathfrak{A}} b$ or $a <^{\mathfrak{B}} b$. For $a \in A$, we write $\mathfrak{A}_{<a}$ for the substructure of $\mathfrak{A}$ induced by the set of elements less than a; substructures $\mathfrak{A}_{>a}$, $\mathfrak{A}_{\leqslant a}$ and $\mathfrak{A}_{\geqslant a}$ are defined analogously. The following observation is straightforward from the characterisation of $\equiv_r$ in terms of Ehrenfeucht-Fraïssé games.

Observation 4.4. *If $\mathfrak{A}_1$, $\mathfrak{A}_2$, $\mathfrak{B}_1$ and $\mathfrak{B}_2$ are σ-words with $\mathfrak{A}_1 \equiv_q \mathfrak{A}_2$ and $\mathfrak{B}_1 \equiv_q \mathfrak{B}_2$ then $\mathfrak{A}_1 \lhd \mathfrak{B}_1 \equiv_q \mathfrak{A}_2 \lhd \mathfrak{B}_2$.*

We are now ready to state the lemma on ordered structures.

Lemma 4.5. *For a vocabulary σ and $q \geqslant 0$, 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 and containing the last element of $\mathfrak{A}$, such that $\mathfrak{A} \equiv_q \mathfrak{B}$.*

Proof. For fixed σ and q there are, up to logical equivalence, only finitely many first-order formulae $\varphi(x)$ of quantifier rank q with one free variable x. Let f be the number of such formulae and let $N = 2^f$. We will show that if the length of $\mathfrak{A}$ is greater than N then $\mathfrak{A}$ contains a *proper* substructure $\mathfrak{B}$ with $\mathfrak{B}$ containing the last element of $\mathfrak{A}$ and $\mathfrak{A} \equiv_q \mathfrak{B}$, from which the lemma clearly follows by iteration.

Since the length of $\mathfrak{A}$ is greater than 2^f, there must be two distinct elements $a < b$ in $\mathfrak{A}$ satisfying exactly the same formulae $\varphi(x)$ of quantifier rank q. This means that player II wins the q-round Ehrenfeucht-Fraïssé game played on the structures $(\mathfrak{A}, a)$ and $(\mathfrak{A}, b)$. Because of the linear order, II is required to match any move played on an element $\geqslant a$ in the first structure by an element $\geqslant b$ in the second. This implies that $\mathfrak{A}_{\geqslant a} \equiv_q \mathfrak{A}_{\geqslant b}$ and similarly $\mathfrak{A}_{<a} \equiv_q \mathfrak{A}_{<b}$. Because $\mathfrak{A} = \mathfrak{A}_{<a} \lhd \mathfrak{A}_{\geqslant a}$ this gives (using Observation 4.4)

$$\mathfrak{A} \equiv_q \mathfrak{A}_{<a} \lhd \mathfrak{A}_{\geqslant b}.$$

The right hand side is the required substructure $\mathfrak{B}$. ☐

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

4.2 Finite irreflexive transitive trees

We now turn to Kripke structures on frames that are rooted irreflexive transitive trees, or $\prec$-trees, in the sense of Definition 2.8.

We use the symbol $\prec$ for the relation R on irreflexive transitive trees ($\prec$-trees for short), and regard $\preccurlyeq$ as an abbreviation. We aim to use Lemma 4.5 to establish a result about the equivalence of trees. To be precise, we aim to characterise the equivalence relation $\equiv_q$ on $\prec$-trees in terms of the relations $\equiv_{q-1}$ on subtrees along with an equivalence on suitably defined words.

Fix $q > 0$ and a vocabulary σ consisting of a binary symbol $\prec$ 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 $q-1$ or less. Let σ_q be the vocabulary consisting of a binary relation symbol $<$ and unary symbols $P_1, \dots, P_f$. Let $\mathfrak{A}$ be a σ-structure on the frame $(A, \prec, \alpha)$ which is a $\prec$-tree with root α. For any element $a \in A$, we define a σ_q-word $L_q(\mathfrak{A}, a)$ as follows. Let p be the unique maximal path from the root α of $\mathfrak{A}$ to a (which we identify with the set of nodes along the path). Then

- the universe of $L_q(\mathfrak{A}, a)$ consists of the set $p = \{b \in A \colon \alpha \preccurlyeq b \preccurlyeq a\}$.
- $b_1 < b_2$ in $L_q(\mathfrak{A}, a)$ if, and only if, $b_1 \prec b_2$.
- $P_i(b)$ holds if, and only if, the sentence φ_i is true in the induced substructure of $\mathfrak{A}$ on the set

$$S^p(b) := \{b' \in A \colon b \preccurlyeq b' \text{ and } b' \not\succcurlyeq b'' \text{ for any } b \prec b'' \preccurlyeq a\}.$$

For any $b \in A$ we write $\mathfrak{A}_b$ for the subtree of $\mathfrak{A}$ rooted at b. For each b on the path p, we also write $S^p(b)$ for the substructure induced by the set $S^p(b)$. Intuitively, $S^p(b)$ is obtained by removing from $\mathfrak{A}_b$ all elements on the path p after b and all of their descendants.

We can think of $L_q(\mathfrak{A}, a)$ as the linearly ordered set of elements that are on the unique maximal path p 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 $q - 1$ that are true in $S^p(b)$. $L_q(\mathfrak{A}, a)$ 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 $\equiv_{q-1}$. This is illustrated in Figure 2. Here b is a node on the path p from the root α to a and the structure $S^p(b)$ shown hanging below b consists of the subtree rooted at b excluding the part that is hanging from the path p to the right of b. We call this the *slice* of $\mathfrak{A}$ rooted at b. The word $L_q(\mathfrak{A}, a)$ consists of just the path p, with the nodes coloured by the types of the slices hanging below them.

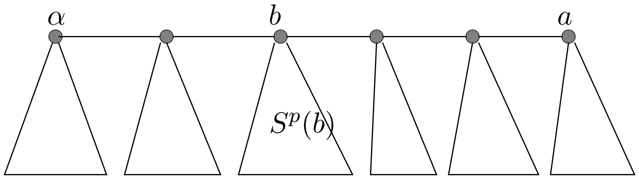

Figure 2: $L_q(\mathfrak{A}, a)$

This representation is not 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 $\equiv_q$ class of $\mathfrak{A}$. To be precise, we aim to prove the following lemma.

Lemma 4.6. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be irreflexive transitive σ-trees, in which every node has finitely many $\prec$-predecessors. If $\mathfrak{A}$ and $\mathfrak{B}$ satisfy*

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

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

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

Proof. A winning strategy for player II on the pair of structures $\mathfrak{A}$ and $\mathfrak{B}$ can be obtained as a composition of the strategy on the structures $L_q(\mathfrak{A}, a)$ and $L_q(\mathfrak{B}, b)$ and winning strategies on pairs of $\equiv_{q-1}$-equivalent subtrees.

To be precise, if player I's first move is in the structure $\mathfrak{A}$ say on a, II responds with the b given by condition (1) and similarly II responds using condition (2) if I's first move is in structure $\mathfrak{B}$. From this point on, the structures $L_q(\mathfrak{A}, a)$ and $L_q(\mathfrak{B}, b)$ are fixed and we describe player II's strategy for the remaining $q - 1$ moves. For each $a' \in \mathfrak{A}$, we write $\pi(a')$ for the unique last element u on the maximal path from the root α of $\mathfrak{A}$ to a such that $u \preccurlyeq a'$. That is, $\pi(a')$ is the element of $L_q(\mathfrak{A}, a)$ such that a' is in $S^p(\pi(a'))$, where p is the path from α to a. Similarly, for $b' \in \mathfrak{B}$, we write $\pi(b')$ to denote the element of $L_q(\mathfrak{B}, b)$ such that b' is in $S^{p'}(\pi(b'))$, where p is the path from the root β of $\mathfrak{B}$ to b. Player II's strategy is to ensure that after r moves have been played, say on $a, a_1, \ldots a_{r-1}$ and $b, b_1, \ldots, b_{r-1}$, the following conditions hold:

1. $L_q(\mathfrak{A}, a), \pi(a_1), \ldots, \pi(a_{r-1}) \equiv_{q-r} L_q(\mathfrak{B}, b), \pi(b_1), \ldots, \pi(b_{r-1})$;

2. for each i, if $a_{j_1}, \ldots, a_{j_k}$ are in $S^p(\pi(a_i))$, then

$$S^p(\pi(a_i)), a_{j_1}, \ldots, a_{j_k} \equiv_{q-r} S^{p'}(\pi(b_i)), b_{j_1}, \ldots, b_{j_k}.$$

These conditions are clearly satisfied when $r = 1$ by the choice of a and b. Suppose now that they hold after r moves, and suppose, without loss of generality, that at move $r+1$ player I plays on an element c of $\mathfrak{A}$. If $S^p(\pi(c))$ already contains an element a_i for some $i < r$, then by condition (2), there is an element $d \in S^{p'}(\pi(b_i))$ which constitutes player II's response in the game played on $S^p(\pi(a_i))$ and $S^{p'}(\pi(b_i))$. It is then easily seen that playing on d preserves both conditions. If, on the other hand, $S^p(\pi(c))$ does not contain any elements previously chosen, let v be player II's response to a move on $\pi(c)$ in the game played on $L_q(\mathfrak{A}, a), \pi(a_1), \ldots, \pi(a_{r-1})$ and $L_q(\mathfrak{B}, b), \pi(b_1), \ldots, \pi(b_{r-1})$. Since $\pi(c)$ and v must have the same atomic types in $L_q(\mathfrak{A}, a)$ and $L_q(\mathfrak{B}, b)$ respectively, we conclude that the respective slices $S^p(\pi(c))$ and $S^{p'}(v)$ satisfy the same sentences of quantifier rank $q-1$. Thus, player II has a winning strategy in the $q-1$ move game played on these subtrees. Let d be the element $S^{p'}(v)$ that gives player II's response to player I's move on c. It is again easily checked that playing on d in the game on $\mathfrak{A}$ and $\mathfrak{B}$ preserves both conditions. $\square$

Saturated trees Our next goal is to define, for any structure built on an irreflexive transitive tree frame, a saturated companion which is bisimilar to the original structure but will enable us to upgrade $\sim^\ell$ to $\equiv_q$ for suitable $\ell = \ell(q)$. For this we again apply the technique of boosting multiplicities.

Let $T = (A, \prec, \alpha)$ be a finite $\prec$-tree with root α and let $q \in \mathbb{N}$. We define $s_q(T)$, the *q-saturated companion* of T to be the $\prec$-tree $(W, \prec, \alpha)$ with

- $W = \{w \in (A \times \{0, \ldots, q-1\})^+ \colon w = (a_1, l_1) \cdots (a_n, l_n)$ with $a_1 = \alpha, l_1 = 0$ and $a_i \prec a_j$ for all $i < j\}$;
- $w_1 \prec w_2$ if, and only if, w_1 is a proper prefix of w_2;
- root $(\alpha, 0)$ identified with α.

If $\mathfrak{A}$ is a structure on the frame T, rooted at α, let $s_q(\mathfrak{A})$ denote the structure on the frame $s_q(T)$ where, for each unary predicate symbol P, $((a_1, l_1) \cdots (a_n, l_n)) \in P^{s_q(\mathfrak{A})}$ if, and only if, $a_n \in P^{\mathfrak{A}}$.

In terms of our considerations in section 2.5

$$s_q(\mathfrak{A}) = \mathrm{TC}((\mathfrak{A} \otimes q)^*_\alpha).$$

According to Observation 2.17 in section 2.5, $s_q(\mathfrak{A}), \alpha$ is a $\prec$-tree such that $s_q(\mathfrak{A}), \alpha \sim \mathfrak{A}, \alpha$. Indeed, the correspondence between $w = \cdots (a, l) \in W$ and a is a natural bisimulation. Moreover, if $\mathfrak{A}$ is finite, then so is $s_q(\mathfrak{A})$. We call $s_q(\mathfrak{A})$ the q-saturated companion of $\mathfrak{A}$.

We recall from Observation 2.17 that, more generally, even for not necessarily finite, rooted transitive structures $\mathfrak{A}, \alpha$, $s_q(\mathfrak{A}, \alpha) = \mathrm{TC}((\mathfrak{A} \otimes q)^*_\alpha) \sim \mathfrak{A}, \alpha$ is still a $\prec$-tree in which the set of $\prec$-predecessors of any node is finite. This will play a role in extensions of the present techniques to not necessarily finite trees and rooted structures.

The key properties of the q-saturated companions are summarised in Lemmas 4.8 and 4.9 below. In order to introduce this, we first establish some notation. Given a structure $s_q(\mathfrak{A})$ on the frame $(W, \prec, \alpha)$ and an element $w \in W$ we write $s_q(\mathfrak{A})_w$ for the subtree rooted at w, i.e., the substructure of $s_q(\mathfrak{A})$ induced by the set of elements $\{v \in W : w \preccurlyeq v\}$. We also write $\mathfrak{A}_a$ for the subtree of the structure $\mathfrak{A}$ rooted at a. Then, if $w = \cdots (a, l)$, by construction we have

$$s_q(\mathfrak{A})_w \simeq \mathrm{TC}((\mathfrak{A} \otimes q)^*_a) \simeq s_q(\mathfrak{A}_a).$$

Given a sequence p of elements $w_1 \prec \cdots \prec w_n$ – not necessarily forming a maximal path in $s_q(\mathfrak{A})$ – we now write $S^p(w_i)$ for the substructure of $s_q(\mathfrak{A})$ induced by the set of elements $\{v \in W : w_i \preccurlyeq v \text{ and } u \not\preccurlyeq v \text{ for any } w_i \prec u \preccurlyeq w_{i+1}\}$. That is, $S^p(w_i)$ is obtained from the subtree rooted at w_i in $s_q(\mathfrak{A})$ by excluding all subtrees rooted on the maximal path containing p. We are now ready to state the key properties of q-saturated companions.

Observation 4.7. *For $w \prec v$ in $s_q(\mathfrak{A})$ such that v is an immediate successor of w, there are q distinct immediate successors of w, $v_0 = v, v_1, \ldots, v_{q-1}$, such that $s_q(\mathfrak{A})_{v_i} \cong s_q(\mathfrak{A})_{v_j}$.*

For this we note that v is of the form $v = w \cdot (a, l)$. The required siblings $v_0, \ldots, v_{q-1}$ are obtained as we let l vary through $\{0, \ldots, q-1\}$ in the last node of v.

Lemma 4.8. *Let $p = w_1 \prec \cdots \prec w_n$ be a path in the q-saturated companion $s_q(\mathfrak{A})$ of $\mathfrak{A}$. Then there is a path $p' = w'_1 \prec \cdots \prec w'_n$ such that $S^p(w_i) \cong S^{p'}(w'_i)$ for each i, $w'_1 = w_1$ and such that w'_i is an immediate $\prec$-successor of w'_{i-1} for $i > 1$ (i.e., the path p' is the maximal path from $w'_1 = w_1$ to w'_n in $s_q(\mathfrak{A})$).*

Proof. Let (a_i, l_i) be the last element of w_i, for $1 \leqslant i \leqslant n$. We define the w'_i inductively as follows: $w'_1 = w_1$ and $w'_{i+1} = w'_i \cdot (a_{i+1}, l_{i+1})$. ☐

Lemma 4.9. *Let $q > 0$, $w \in s_{q+1}(\mathfrak{A})$ be a node on some path p in the $(q+1)$-saturated tree $s_{q+1}(\mathfrak{A})$, $S^p(w)$ the slice of $s_{q+1}(\mathfrak{A})$ rooted at w and $s_{q+1}(\mathfrak{A})_w$ the subtree of $s_{q+1}(\mathfrak{A})$ rooted at w. Then $S^p(w) \equiv_q s_{q+1}(\mathfrak{A})_w$.*

In particular, if w is a node on some path p in $s_{q+1}(\mathfrak{A})$ and w' a node on some path p' in $s_{q+1}(\mathfrak{B})$ such that $s_{q+1}(\mathfrak{A})_w \equiv_q s_{q+1}(\mathfrak{B})_{w'}$, then $S^p(w) \equiv_q S^{p'}(w')$.

Proof. If w is the end point of p, then $S^p(w) = s_{q+1}(\mathfrak{A}_w)$ and the claim is obvious. Otherwise let w' be the immediate successor of w on the maximal path containing p. Thus $w' = w \cdot (a, l)$ for some $a \in \mathfrak{A}$. Then, $S^p(w)$ is obtained from $s_{q+1}(\mathfrak{A})_w$ by removing the subtree rooted at w'. By Observation 4.7 there are at least q other subtrees rooted at immediate successors of w that are isomorphic to $s_{q+1}(\mathfrak{A})_{w'}$. A strategy in the q-round Ehrenfeucht-Fraïssé game on $S^p(w), w$ versus $s_{q+1}(\mathfrak{A})_w, w$ is immediate from this fact. ☐

Our aim is to establish the following claim.

Claim 4.10. *For every vocabulary σ and each q, there is an ℓ such that for all rooted $\prec$-tree structures $\mathfrak{A}, \alpha$ and $\mathfrak{B}, \beta$:*

$$\mathfrak{A}, \alpha \sim^\ell \mathfrak{B}, \beta \quad \Rightarrow \quad s_q(\mathfrak{A}) \equiv_q s_q(\mathfrak{B}).$$

Proof. For a fixed σ, let $f(q)$ be the number of first order formulae of quantifier rank q (up to logical equivalence). We now define ℓ, as a function of q, by induction:

$$\begin{aligned}\ell(0) &= 0\\ \ell(q+1) &= N(q, f(q))(\ell(q)+1)\end{aligned}$$

where N is the function from Lemma 4.5.

The proof proceeds by induction on q. The base case for $q = 0$ is trivial. Suppose now that the claim is true for some value q. This implies that for any tree structure $\mathfrak{A}$, there is a modal formula $\varphi^q_{\mathfrak{A}}$ of modal nesting depth $\ell(q)$ such that for any tree $\mathfrak{B}$ if $\mathfrak{B} \models \varphi^q_{\mathfrak{A}}$ then $s_q(\mathfrak{A}) \equiv_q s_q(\mathfrak{B})$. Moreover, since $\mathfrak{B} \sim s_q(\mathfrak{B})$, this is equivalent to saying that

$$(*) \qquad s_q(\mathfrak{B}) \models \varphi^q_{\mathfrak{A}} \quad \Rightarrow \quad s_q(\mathfrak{A}) \equiv_q s_q(\mathfrak{B}).$$

Now, suppose $\mathfrak{A}$ and $\mathfrak{B}$ are tree structures with $\mathfrak{A}, \alpha \sim^{\ell(q+1)} \mathfrak{B}, \beta$. We wish to show that $s_{q+1}(\mathfrak{A}) \equiv_{q+1} s_{q+1}(\mathfrak{B})$.

To do this, we use Lemma 4.6. By symmetry, it suffices to establish condition (1) of the lemma. Thus, let v be an arbitrary element of $s_{q+1}(\mathfrak{A})$ and consider the structure $L_{q+1}(s_{q+1}(\mathfrak{A}), v)$. We need to find w in $s_{q+1}(\mathfrak{B})$ such that $L_{q+1}(s_{q+1}(\mathfrak{A}), w) \equiv_q L_{q+1}(s_{q+1}(\mathfrak{A}), v)$.

Since $L_{q+1}(s_{q+1}(\mathfrak{A}), v)$ is a linearly ordered structure in a vocabulary with $f(q)$ unary predicates, by Lemma 4.5 there is a subsequence of this linear order of length $N \leqslant N(q, f(q))$ which is $\equiv_q$-equivalent to $L_{q+1}(s_{q+1}(\mathfrak{A}), v)$ and contains v.

Let $p = v_0, v_1, \ldots, v_N$ be the path from the root $\alpha = v_0$ of $s_{q+1}(\mathfrak{A})$ to $v = v_N$ corresponding to this subsequence. By Lemma 4.8 there are v'_i that form a maximal path p' from α to some $v' = v'_N$ such that $S^{p'}(v'_i) \simeq S^p(v_i)$ for each $i \leqslant N$. It follows that $L_{q+1}(s_{q+1}(\mathfrak{A}), v') \equiv_q L_{q+1}(s_{q+1}(\mathfrak{A}), v)$.

For each i, let $\mathfrak{A}_i$ be the subtree $s_{q+1}(\mathfrak{A})_{v'_i}$ of $s_{q+1}(\mathfrak{A})$ rooted at v'_i. Let θ_i be the formula $\varphi^q_{\mathfrak{A}_i}$, i.e. the modal formula of rank $\ell(q)$ that characterises the subtree of $s_{q+1}(\mathfrak{A})$ rooted at v'_i up to $\equiv_q$. Consider now the modal formula

$$\theta_0 \wedge \Diamond(\theta_1 \wedge \Diamond(\theta_2 \cdots \theta_N)),$$

whose nesting depth is bounded by $\ell(q+1)$. This formula is clearly true at the root of $s_{q+1}(\mathfrak{A})$, as witnessed by the path $v'_0, \ldots, v'_N$. Since $s_{q+1}(\mathfrak{A}), \alpha \sim^{\ell(q+1)} s_{q+1}(\mathfrak{B}), \beta$, it is also true at the root of $s_{q+1}(\mathfrak{B})$. Thus, there is a path $w_0, \ldots, w_N$ from the root $\beta = w_0$ to some $w = w_N$ in $s_{q+1}(\mathfrak{B})$ such that $s_{q+1}(\mathfrak{B}), w_i \models \theta_i$. Moreover, by Lemma 4.8, we may assume w.l.o.g. that the path $b_0, \ldots, b_N$ is a maximal path from β to b_N. By the choice of the formulae θ, $s_{q+1}(\mathfrak{B})_{w_i} \equiv_q s_{q+1}(\mathfrak{A})_{v'_i}$. By Lemma 4.9, this implies that the corresponding slices at v'_i in $s_{q+1}(\mathfrak{A})$ and at w_i in $s_{q+1}(\mathfrak{B})$ are $\equiv_q$-equivalent. Therefore, $L_{q+1}(s_{q+1}(\mathfrak{B}), w) \equiv_q L_{q+1}(s_{q+1}(\mathfrak{A}), v')$ and, as $L_{q+1}(s_{q+1}(\mathfrak{A}), v') = L_{q+1}(s_{q+1}(\mathfrak{A}), v)$, we have $L_{q+1}(s_{q+1}(\mathfrak{B}), w) \equiv_q L_{q+1}(s_{q+1}(\mathfrak{A}), v)$ as desired. □

This immediately yields a modal characterisation theorem over the class of finite irreflexive transitive trees.

Theorem 4.11. FO/$\sim$ $\equiv$ ML *over the class of all finite irreflexive transitive trees.*

Proof. Let $\varphi(x) \in$ FO be bisimulation invariant over finite $\prec$-trees, φ of quantifier rank q. Claim 4.10 shows that for some ℓ, φ is $\sim^\ell$ invariant over this class of trees. Therefore φ is equivalent to a formula of ML of nesting depth ℓ over this class. □

Not necessarily finite irreflexive transitive trees. We note that Claim 4.10 works for not necessarily finite irreflexive trees $\mathfrak{A}$ and $\mathfrak{B}$, since the structures $s_q(\mathfrak{A})$ and $s_q(\mathfrak{B})$ are based on unravellings. They therefore satisfy the condition that all nodes have finite predecessor sets (crucial for Lemma 4.6) and in particular the well-ordering condition on trees. For Theorem 4.13 (Löb frames) also note that transitive frames without infinite paths are necessarily irreflexive and that (the transitive closure of) the tree unravelling of such a frame is a $\prec$-tree in which every node has finitely many $\prec$-predecessors. In other words, the upgrading argument of Claim 4.10 is applicable within any class of irreflexive transitive structures that is closed under taking q-saturated companions – and clearly the classes of all irreflexive transitive trees and the class of Löb frames are such. For these frames, the passage to the transitive closure of (irreflexive) unravellings with finitely boosted multiplicities remains within the given class and yields bisimilar companions that are irreflexive transitive trees with finite predecessor sets (cf. Observation 2.17). Thus, one obtains the following characterisations.

Theorem 4.12. FO/$\sim$ $\equiv$ ML *over the class of all (not necessarily finite) irreflexive transitive trees.*

Theorem 4.13. FO/$\sim$ $\equiv$ ML *over the class of all transitive structures without infinite directed paths (Löb frames).*

The same reasoning applies of course to arbitrary irreflexive transitive frames and rooted irreflexive transitive frames, as well as to the class of all transitive frames or all rooted transitive frames. In all these cases, therefore, we obtain alternative proofs for results already obtained in Theorem 2.12 on the basis of a classical argument.

4.3 Finite reflexive transitive trees

We use a simple model theoretic argument based on a natural interpretation of reflexive transitive trees in irreflexive transitive trees to obtain the analogue of Theorem 4.11 also for the class of finite reflexive transitive trees. A simple modification will then also prove the analogous characterisation for not necessarily finite reflexive trees, too.

Theorem 4.14. FO/$\sim$ $\equiv$ ML *over the class of all finite reflexive transitive trees.*

Consider a reflexive transitive tree $\mathfrak{A}$ with accessibility relation $\preccurlyeq$. Its irreflexive variant $\mathfrak{A}_\prec$ is obtained by replacing $\preccurlyeq$ by $\prec$ (that is to say, replacing $R^{\mathfrak{A}}$ by $R^{\mathfrak{A}_\prec} := R^{\mathfrak{A}} \setminus (R^{\mathfrak{A}})^\circ$, where R° stands for the reflexive part of R). As $\mathfrak{A}$ and $\mathfrak{A}_\prec$ are quantifier-free FO-interpretable within one another,

$$\mathfrak{A}, \alpha \equiv_q \mathfrak{B}, \beta \quad \Leftrightarrow \quad \mathfrak{A}_\prec, \alpha \equiv_q \mathfrak{B}_\prec, \beta$$

for any reflexive transitive trees $\mathfrak{A}, \mathfrak{B}$. It is equally obvious that

$$\mathfrak{A}_\prec, \alpha \sim^\ell \mathfrak{B}_\prec, \beta \quad \Rightarrow \quad \mathfrak{A}, \alpha \sim^\ell \mathfrak{B}, \beta.$$

Ideally, if we could show that in the converse direction, for sufficiently large ℓ' also

$$(\dagger) \qquad \mathfrak{A}, \alpha \sim^{\ell'} \mathfrak{B}, \beta \quad \Rightarrow \quad \mathfrak{A}_\prec, \alpha \sim^\ell \mathfrak{B}_\prec, \beta,$$

then a proof of Theorem 4.14 immediately follows from Claim 4.10:

Indeed, let $\varphi = \varphi(x) \in \mathrm{FO}$ be bisimulation invariant over finite $\preccurlyeq$-trees, φ of quantifier rank q and let ℓ be as in Claim 4.10 for q, and ℓ' for ℓ as in (†). To show that φ is $\sim^{\ell'}$ invariant over finite $\preccurlyeq$-trees, let $\mathfrak{A}, \alpha \sim^{\ell'} \mathfrak{B}, \beta$. So $\mathfrak{A}_{\prec}, \alpha \sim^{\ell} \mathfrak{B}_{\prec}, \beta$ and $s_q(\mathfrak{A}_{\prec}), \alpha \equiv_q s_q(\mathfrak{B}_{\prec}), \beta$. Clearly, by construction, $s_q(\mathfrak{A}_{\prec}) \simeq \mathfrak{A}'_{\prec}$ for some $\mathfrak{A}', \alpha \sim \mathfrak{A}, \alpha$, and similarly for $\mathfrak{B}$. So we have $\mathfrak{A}, \alpha \sim \mathfrak{A}', \alpha \equiv_q \mathfrak{B}', \beta \sim \mathfrak{B}, \beta$, whence $\mathfrak{A}, \alpha \models \varphi$ iff $\mathfrak{B}, \beta \models \varphi$.

However, simple examples show that (†) does not hold in general. Consider, for instance, a $\preccurlyeq$-tree consisting of a single reflexive node and its bisimilar companion consisting of two copies of that reflexive node linked by an edge. But it turns out that (†) is true for suitable bisimilar companions $\hat{\mathfrak{A}}$ and $\hat{\mathfrak{B}}$, so that a corresponding upgrading fills the gap.

One obstruction in connection with (†), in terms of the bisimulation game on the $\prec$-trees, is the following. In a position $\mathfrak{A}, a \sim^m \mathfrak{B}, b$ it may be that the only available response b' for a move along $a \prec a'$ in $\mathfrak{A}$ say, such that $\mathfrak{B}, b' \sim^{m-1} \mathfrak{A}, a'$ is $b' = b$. While such a response is good in the game on the $\preccurlyeq$-trees, the game on the $\prec$-trees requires such b' with $b \prec b'$. However, this can only happen if $\mathfrak{A}, a' \sim^{m-1} \mathfrak{A}, a$ and if $\mathfrak{B}, b$ satisfies

$$\xi^b_m := \Box\big(\chi^b_{m-1} \rightarrow \chi^b_m\big).$$

Here we use formulae $\chi^b_n \in \mathrm{ML}_n$ that characterise the $\sim^n$-type of b in $\mathfrak{B}$. Call elements b that satisfy this condition *m-critical.* We need the following facts about critical elements.

Observation 4.15. (i) *If $\mathfrak{B}, b \sim^{m+1} \mathfrak{B}, b'$, then b is m-critical if, and only if, b' is m-critical.*

(ii) *For $b \preccurlyeq b'$ with $\mathfrak{B}, b \sim^m \mathfrak{B}, b'$, if b is m-critical then b' is m-critical.*

In particular, as bisimulations preserve m-critical elements, these can not be eliminated through passage to bisimilar companions. In order to deal with the problem of critical elements, we need to make sure that in a position $\mathfrak{A}, a \sim^m \mathfrak{B}, b$ with m-critical a or b, there are a' with $a \prec a'$ and $a \sim^m a'$ if, and only if, there are b' with $b \prec b'$ and $b \sim^m b'$. To control this requirement we look at the *$\sim^m$-depth* of elements.

In a finite $\preccurlyeq$-tree $\mathfrak{A}$, the $\sim^m$-depth $d^m(a)$ of an element a is defined as the maximal length n of a path $p = a_0 \prec a_1 \prec \cdots \prec a_n$ from $a_0 = a$ for which $\mathfrak{A}, a \sim^m \mathfrak{A}, a_i$ for all i. Note that $d^{m-1}(a) \geqslant d^m(a)$, for all m. For an m-critical element a, moreover, $d^{m-1}(a) = d^m(a)$.

Observation 4.16. *For every finite rooted $\preccurlyeq$-tree structure $\mathfrak{A}, \alpha$, for all $d, m \in \mathbb{N}$: if a in $\mathfrak{A}$ is such that $d^m(a) > d$ there is some a' with $a \prec a'$, then $\mathfrak{A}, a \sim^m \mathfrak{A}, a'$ and $d^m(a') = d$.*

For the desired upgrading to support (†) we want to use structures $\mathfrak{A}$ and $\mathfrak{B}$ that always provide nodes of sufficiently high $\sim^m$-depth in the ℓ-round bisimulation game on $\mathfrak{A}_{\prec}$ versus $\mathfrak{B}_{\prec}$. To this end, we pass from the given $\mathfrak{A}$ to its bisimilar companion $\hat{\mathfrak{A}}$ obtained by replacing each node of $\mathfrak{A}$ by a chain of $\ell + 1$ copies of that node, and taking the reflexive transitive closure of the result. We thus obtain the following as a basis for the desired upgrading.

Observation 4.17. *For every finite rooted $\preccurlyeq$-tree structure $\mathfrak{A}, \alpha$ and $\ell \in \mathbb{N}$ there is a finite rooted $\preccurlyeq$-tree structure $\hat{\mathfrak{A}}, \hat{\alpha}$ such that $\hat{\mathfrak{A}}, \hat{\alpha} \sim \mathfrak{A}, \alpha$ and with the following additional properties:*

(i) $d^{2\ell}(\hat{\alpha}) \geqslant \ell$.

(ii) *for all $d, m \leqslant \ell$ and any pair a, a' in $\hat{\mathfrak{A}}$ such that $a \prec a'$ and $\hat{\mathfrak{A}}, a' \not\sim \hat{\mathfrak{A}}, a$, there is some a'' with $a \prec a''$, $\hat{\mathfrak{A}}, a'' \sim \hat{\mathfrak{A}}, a'$ and $d^m(a'') = d$.*

Let us say that a finite rooted $\preccurlyeq$-tree $\hat{\mathfrak{A}}, \hat{\alpha}$ is *ℓ-good* if it satisfies conditions (i)–(ii) of Observation 4.17.

Claim 4.18. *If $\mathfrak{A}, \alpha$ and $\mathfrak{B}, \beta$ are ℓ-good finite rooted $\preccurlyeq$-tree structures such that $\mathfrak{A}, \alpha \sim^{2\ell+1} \mathfrak{B}, \beta$, then $\mathfrak{A}_\prec, \alpha \sim^\ell \mathfrak{B}_\prec, \beta$.*

Proof. We want to show that the following system $(Z_k)_{k \leqslant \ell}$ satisfies the back-and-forth conditions for $\mathfrak{A}_\prec, a \sim^\ell \mathfrak{B}_\prec, b$. We use the notation $n \stackrel{k}{=} n'$ for ($n = n'$ or $n, n' \geqslant k$).

$$Z_k := \bigcup_{m \geqslant 2k} \left\{ (a,b) \colon \begin{array}{r} \mathfrak{A}, a \sim^m \mathfrak{B}, b \text{ and either both } a, b \text{ not } m\text{-critical,} \\ \text{or both } m\text{-critical and } d^m(a) \stackrel{k}{=} d^m(b) \end{array} \right\}.$$

Clearly $(\alpha, \beta) \in Z_\ell$: note that whether or not the roots are 2ℓ-critical is determined by their $\sim^{2\ell+1}$ type. For the forth property, consider $(a,b) \in Z_k$ for some $k \geqslant 1$ and a move along $a \prec a'$ to a' in $\mathfrak{A}$. Assume that $m \geqslant 2k$ is such that $\mathfrak{A}, a \sim^m \mathfrak{B}, b$ and conditions on m-criticalness and $\sim^m$-depth as in the definition of $(a,b) \in Z_k$ are satisfied.

We distinguish two cases:

(1) $\mathfrak{A}, a \not\sim^{m-1} \mathfrak{A}, a'$, or $\mathfrak{A}, a \sim^{m-1} \mathfrak{A}, a'$ but a and b are not m-critical.

Using either the strategy for $\mathfrak{A}, a \sim^m \mathfrak{B}, b$ or the condition that b is not m-critical, we find b' such that $b \prec b'$, $b' \sim^{m-1} a'$ and $b' \not\sim b$. As $b' \sim^{m-1} a'$, b' is $(m-2)$-critical iff a' is.

If a' and b' are both not $(m-2)$-critical, then $(a', b') \in Z_{k-1}$. If a' and b' are $(m-2)$-critical, then we further need to consider d^{m-2}. If $d^{m-2}(b') > d^{m-2}(a')$ we pass to some b'', $b' \prec b''$ within the same $\sim^{m-2}$ class (hence b'' still $(m-2)$-critical by Observation 4.15 (ii)) such that $d^{m-2}(b'') = d^{m-2}(a')$ (this uses Observation 4.16).

If $d^{m-2}(b') < d^m(a')$ we pass to some b'', $b \prec b'' \prec b'$ within the same $\sim$ class (hence b'' is still $(m-2)$-critical by Observation 4.15 (i)) such that $d^{m-2}(b'') \stackrel{k-1}{=} d^{m-2}(a')$ (this uses Observation 4.17 (ii)).

(2) $\mathfrak{A}, a \sim^{m-1} \mathfrak{A}, a'$, a and b both m-critical. It follows that $\mathfrak{A}, a \sim^m \mathfrak{A}, a'$.

In this case we want to match a' with some b' such that $b' \sim^m b$, $b \prec b'$ of the right $\sim^{m-1}$-depth.

As a is m-critical, $a \sim^m a'$, and a' is $(m-1)$-critical iff a is $(m-1)$-critical iff b is $(m-1)$-critical. So any $b' \sim^m b$ will automatically be $(m-1)$-critical if a' is.

As a and b are m-critical, we have $d^{m-1}(b) = d^m(b) \stackrel{k}{=} d^m(a) = d^{m-1}(a) > d^{m-1}(a')$. So there is some b', $b \prec b'$, with $b' \sim^m b$ and $d^{m-1}(b') \stackrel{k}{=} d^{m-1}(a')$. Hence $(a', b') \in Z_{k-1}$. $\square$

Since every finite $\prec$-tree structure is bisimilar to an ℓ-good one by Observation 4.17, we get the following.

$$
\begin{tikzcd}
\mathfrak{A}, \alpha \arrow[r, dash] \arrow[d, dash, "\sim"] & \sim^{f_0(\ell)} \arrow[r, dash] & \mathfrak{B}, \beta \arrow[d, dash, "\sim"] & \text{finite } \preccurlyeq\text{-trees} \\
\tilde{\mathfrak{A}}, \alpha \arrow[r, dash] \arrow[d, dotted, leftrightarrow] & \sim^{f_0(\ell)} \; \equiv_q & \tilde{\mathfrak{B}}, \beta \arrow[l, dash] \arrow[d, dotted, leftrightarrow] & \text{finite } \ell\text{-good } \preccurlyeq\text{-trees} \\
& \equiv_q \arrow[u, dotted, Rightarrow] & & \\
\tilde{\mathfrak{A}}_{\prec}, \alpha \arrow[r, dash] & \sim^{\ell} \arrow[r, dash] & \tilde{\mathfrak{B}}_{\prec}, \beta & \text{finite } \prec\text{-trees}
\end{tikzcd}
$$

Figure 3: the upgrading in Corollaries 4.19 and 4.20.

Corollary 4.19. *For every fixed finite vocabulary σ there is a function f_0 such that for all ℓ and any finite rooted $\preccurlyeq$-tree structures $\mathfrak{A}, \alpha \sim^{f_0(\ell)} \mathfrak{B}, \beta$, there are bisimilar companions $\tilde{\mathfrak{A}}, \tilde{\alpha} \sim \mathfrak{A}, \alpha$ and $\tilde{\mathfrak{B}}, \tilde{\beta} \sim \mathfrak{B}, \beta$ within the class of finite rooted $\preccurlyeq$-tree structures such that $\tilde{\mathfrak{A}}_{\prec}, \tilde{\alpha} \sim^{\ell} \tilde{\mathfrak{B}} \prec, \tilde{\beta}$.*

Combining this with Claim 4.10, we obtain the following, which then immediately proves Theorem 4.14.

Corollary 4.20. *For $q \in \mathbb{N}$ and fixed finite vocabulary σ there, there is an $\ell \in \mathbb{N}$ such that for any two finite reflexive transitive tree structures $\mathfrak{A}, \alpha \sim^{\ell} \mathfrak{B}, \beta$, there are finite transitive tree structures $\hat{\mathfrak{A}}, \hat{\alpha} \sim \mathfrak{A}, \alpha$ and $\hat{\mathfrak{B}}, \hat{\beta} \sim \mathfrak{B}, \beta$ such that $\hat{\mathfrak{A}}, \hat{\alpha} \equiv_q \hat{\mathfrak{B}}, \hat{\beta}$.*

Not necessarily finite reflexive transitive trees. If we admit infinite transitive trees, a small modification simplifies the argument and yields the corresponding characterisation over the class of all (finite and infinite) reflexive transitive trees. Over these, $\sim^m$-depth can of course be infinite, viz., if the $\sim^m$-class of a node contains arbitrarily long paths. But in fact, any reflexive transitive tree is bisimilar to one in which all nodes have infinite depth with respect to $\sim^m$ for all m and even with respect to $\sim$. For this we replace every (reflexive) node of $\mathfrak{A}$ by an ω-chain of isomorphic copies, and obtain $\hat{\mathfrak{A}}$ as the transitive closure on the resulting frame, giving an infinite reflexive transitive tree. It is clear that, for these infinite bisimilar companions $\hat{\mathfrak{A}}, \hat{\alpha}$ and $\hat{\mathfrak{B}}, \hat{\beta}$ of the given reflexive transitive trees $\mathfrak{A}, \alpha$ and $\mathfrak{B}, \beta$,

$$\mathfrak{A}, \alpha \sim^{\ell} \mathfrak{B}, \beta \quad \Rightarrow \quad \hat{\mathfrak{A}}, \hat{\alpha} \sim^{\ell} \hat{\mathfrak{B}}, \hat{\beta} \quad \Rightarrow \quad \hat{\mathfrak{A}}_{\prec}, \hat{\alpha} \sim^{\ell} \hat{\mathfrak{B}}_{\prec}, \hat{\beta}.$$

We note that the infinite $\prec$-trees $\hat{\mathfrak{A}}_{\prec}$ and $\hat{\mathfrak{B}}_{\prec}$ do not in general have finite predecessor sets as required for Lemma 4.6, but Claim 4.10 applies to arbitrary $\prec$-trees, since finite predecessor sets are guaranteed in the q-saturated companions $s_q(\hat{\mathfrak{A}}_\prec)$ and $s_q(\hat{\mathfrak{B}}_\prec)$, due to unravelling. We thus obtain the following.

Theorem 4.21. FO/$\sim$ $\equiv$ ML *over the class of all (not necessarily finite) reflexive transitive trees.*

We also obtain again the characterisation theorems over the classes of all (not necessarily finite) reflexive transitive frames and rooted reflexive transitive frames. For this one may use the above reasoning in combination with a passage to the reflexive transitive closures of unravellings to get an alternative proof for corresponding results already obtained by classical methods in Theorem 2.12.

Interestingly, the above argument goes through, with little modifications, also for the class of those not necessarily finite $\preccurlyeq$-trees which admit no infinite $\prec$-paths. Up to unravelling, this class of trees represents the class of Grzegorczyk frames. Over these trees $\sim^m$-depth $d^m(a)$ may be defined, in analogy with the finite case, as the ordinal rank of the well-founded relation $\succ$ in restriction to the $\sim^m$-class of a.

Analogues of Observations 4.16 and 4.17 (for finite d) are straightforward. The analogue of Claim 4.18 then provides an upgrading which reduces the following characterisation issue to (the proof) of Theorem 4.13.

Corollary 4.22. *For every fixed finite vocabulary σ there is a function f_0 such that for all ℓ and any rooted $\preccurlyeq$-tree structures $\mathfrak{A}, \alpha \sim^{f_0(\ell)} \mathfrak{B}, \beta$ without infinite $\prec$-paths, there are bisimilar companions $\tilde{\mathfrak{A}}, \tilde{\alpha} \sim \mathfrak{A}, \alpha$ and $\tilde{\mathfrak{B}}, \tilde{\beta} \sim \mathfrak{B}, \beta$ within the class of $\preccurlyeq$-tree structures without infinite $\prec$-paths such that $\tilde{\mathfrak{A}}_\prec, \tilde{\alpha} \sim^\ell \tilde{\mathfrak{B}} \prec, \tilde{\beta}$.*

Theorem 4.23. FO/$\sim$ $\equiv$ ML *over the class of all reflexive transitive structures without infinite paths with respect to $R \setminus R^\circ$ (Grzegorczyk frames).*

4.4 The extension of basic modal logic to ML*

Let us fix some finite vocabulary $\sigma = \{R\} \cup \{P_1, \dots, P_k\}$ with unary predicates $P_1, \dots, P_k$ besides R. We shall now be interested in σ-structures with transitive R only.

Recall from Definition 2.3 the modal logic ML* as well as the discussion in connection with Examples 4.1 and 4.2 above. This extension of basic modal logic has modalities $\Diamond_p^*$ for every set p of propositional types. If p is the set of types defined by propositional formulae $\zeta_1, \dots, \zeta_s$, the first-order transaltion of $\Diamond_p^* \varphi$ at x is

$$\exists y_0 \exists y_1 \dots \exists y_s \big(Rxy_0 \wedge Ry_0y_0 \wedge \varphi(y_0) \wedge \bigwedge_{0 \leqslant i,j \leqslant s} Ry_iy_j \wedge \bigwedge_{1 \leqslant i \leqslant s} \zeta_i(y_i) \big).$$

Having fixed a finite set of unary predicates in σ, we may restrict attention to sets p of *complete* propositional types. We may think of an individual complete propositional type as a subset $S \subseteq \{P_1, \dots, P_k\}$, or as its defining quantifier-free formula $\zeta_S(x) = \bigwedge_{P_i \in S} P_i(x) \wedge \bigwedge_{P_i \notin S} \neg P_i(x)$. A not necessarily complete propositional type would be represented by a disjunction over complete types. The use of sets p of not necessarily complete types in $\Diamond_p^*$-assertions can then be eliminated through case distinctions. For our puposes we shall now assume that the sets p in question are sets of complete types, and thus formalised as subsets $p \subseteq \mathcal{P}(\{P_1, \dots, P_n\})$.

Remark 4.24. *An alternative definition of (a logic equivalent to)* ML^* *can be based on a natural polyadic variant of the plain modality* $\Diamond^* = \Diamond^*_\emptyset$*, along the following lines. Define* $\Diamond^*(\varphi_1, \ldots, \varphi_m)$ *to be true at* a *in* $\mathfrak{A}$ *if an* R*-clique of nodes* $a_1, \ldots, a_m$ *is accessible from* a *such that* $\mathfrak{A}, a_i \models \varphi_i$*.*

We shall only consider ML^* over path-finite transitive R-frames, where ML^* is guaranteed to be bisimulation invariant, see Lemma 4.25 below. Recall that path-finiteness forbids infinite paths with respect to $R \setminus R^{-1}$.

If we restrict attention to transitive tree-like structures, then the modalities $\Diamond^*_p$ are no more powerful than the plain $\Diamond^* = \Diamond^*_\emptyset$. This is because, due to the absence of R-cliques of more than one element,

$$\begin{array}{rrcll} & \Diamond^*_p\varphi & \equiv & \bot & \text{for } |p| > 1, \\ \text{and} & \Diamond^*_p\varphi & \equiv & \Diamond^*(\zeta \wedge \varphi) & \text{for } |p| = 1, \end{array}$$

if ζ defines the single propositional type in p. Over transitive tree-like frames therefore, in which there cannot be any non-trivial R-cliques, the more complex modalities $\Diamond^*_p$ do in fact reduce to the simple version $\Diamond^*_\emptyset = \Diamond^*$, and ML^* becomes a fragment of the extension of ML by a modality $\Diamond^\circ$ for the reflexive part R° of R via $\Diamond^*\varphi \equiv \Diamond\Diamond^\circ\varphi$.

We write $\sim_*$ for the corresponding notion of bisimulation equivalence. Over transitive tree-like structures, we just need to consider back-and-forth conditions with respect to $\Diamond$- and $\Diamond^*$-moves along R-edges; in the more general setting of transitive structures, which may have non-trivial R-cliques, $\Diamond^*_p$-moves for every set of complete propositional types have to considered. Here a $\Diamond^*_p$-move from a in some transitive σ-structure $\mathfrak{A}$ is a move along an $R^{\mathfrak{A}}$-edge to some node a' that is in the same $R^{\mathfrak{A}}$-clique with nodes realising all the propositional types listed in p. A matching response move in some $\mathfrak{B}$ must similarly be to a node b' in whose $R^{\mathfrak{B}}$-clique all the types in p are realised.

The finite approximations to $\sim_*$ are denoted $\sim^\ell_*$ and defined in terms of back-and-forth systems $(Z_k)_{k \leqslant \ell}$ as usual. Characteristic formulae $\chi^a_M \in \mathrm{ML}^*_m$ that characterise the $\sim^m_*$-type are also obtained as usual.

The following lemma implies in particular that ML^* is bisimulation invariant over the class of finite transitive frames. The discussion of Example 4.1 indicated that the same is not true for all transitive frames. The crucial point for the argument, however, is path-finiteness rather than finiteness: the absence of infinite paths with respect to $R \setminus R^{-1}$.

Lemma 4.25. *For path-finite transitive frames* $\mathfrak{A}, a$ *and* $\mathfrak{B}, b$ *(and hence in particular for finite transitive frames):* $\mathfrak{A}, a \sim \mathfrak{B}, b \quad \Leftrightarrow \quad \mathfrak{A}, a \sim_* \mathfrak{B}, b$*.*

This assertion is in fact a direct consequence of the following observation.

Observation 4.26. *Over any path-finite transitive* σ*-structure* $\mathfrak{A}$*,*

$$\mathfrak{A}, a \models \Diamond^*_p\varphi \quad \text{iff} \quad \left\{ \begin{array}{l} \text{there is an infinite } R\text{-path from } a \\ \text{along which } \varphi \text{ as well as all the types in } p \\ \text{are true infinitely often.} \end{array} \right.$$

We note that the path condition on the right is expressible in the alternation-free μ*-calculus (without any restriction on the class of frames).*

Proof. That $\Diamond_p^*\varphi$ implies the existence of a path as indicated is obvious: choose any path that cyclically visits the relevant nodes in an R-clique witnessing the $\Diamond_p^*$ assertion.

The converse follows from the fact that any infinite R-path in $\mathfrak{A}$ must eventually stay within one and the same R-clique of $\mathfrak{A}$, if $\mathfrak{A}$ is path-finite. $\square$

4.5 Finite transitive tree-like structures

Transitive tree-like structures, as defined in Definition 2.8 generalise both $\prec$- and $\preccurlyeq$-trees in allowing reflexive as well as irreflexive nodes. For transitive tree-like $\mathfrak{A}$ with accessibility relation $R^{\mathfrak{A}}$ we also write $\prec$ for the irreflexive part of $R^{\mathfrak{A}}$, $\prec^{\mathfrak{A}} = R^{\mathfrak{A}} \setminus (R^{\mathfrak{A}})^\circ$, and $\preccurlyeq$ for its reflexive closure. We shall also look at the natural irreflexive encoding of the frame $\mathfrak{A} = (A, R^{\mathfrak{A}})$ by the Kripke structure $\mathfrak{A}^\circ_\prec := (A, \prec^{\mathfrak{A}}, (P^\circ)^{\mathfrak{A}})$ with the extra unary predicate $(P^\circ)^{\mathfrak{A}} := \{a \in A \colon (a,a) \in R^{\mathfrak{A}}\}$ marking the reflexive elements.

We know from section 4.4 that $\mathrm{ML}^* = \mathrm{ML}[\Diamond^*]$ over transitive tree-like structures, and that ML^* is bisimulation invariant over the class of all path-finite transitive tree-like structures. We next show the following.

Theorem 4.27. $\mathrm{FO}/{\sim} \equiv \mathrm{ML}^*$ *over the class of all finite transitive tree-like structures.*

We reduce the theorem to Theorem 4.11 along the following lines. We wish to show that any $\varphi = \varphi(x)$ in FO that is bisimulation invariant over the class of finite transitive tree-like structures is $\sim_*^\ell$ invariant over this class, for some ℓ depending on the quantifier rank q of φ (and on its vocabulary).

We use an upgrading from suitable levels of $\sim_*^{\ell'}$ equivalence between finite transitive tree-like $\mathfrak{A}$ and $\mathfrak{B}$ – in bisimilar companions $\hat{\mathfrak{A}}$ and $\hat{\mathfrak{B}}$ – to $\sim^\ell$ equivalence between the irreflexive encodings $\hat{\mathfrak{A}}^\circ_\prec$ and $\hat{\mathfrak{B}}^\circ_\prec$ of these structures. In combination with passing to saturated companions, we thus achieve $\equiv_q$ equivalence between the irreflexive encodings. As $\hat{\mathfrak{A}}^\circ_\prec$ and $\hat{\mathfrak{A}}$ are related by a quantifier free interpretation, we have that $\hat{\mathfrak{A}}^\circ_\prec \equiv_q \hat{\mathfrak{B}}^\circ_\prec$ if, and only if, $\hat{\mathfrak{A}} \equiv_q \hat{\mathfrak{B}}$.

We define m-critical nodes as before, but now with respect to $\sim_*^m$ and $\sim_*^{m-1}$: a node b is m-critical if it satisfies the formula

$$\Box^*\big(\chi^b_{m-1} \rightarrow \chi^b_m\big),$$

where $\chi^b_n \in \mathrm{ML}^*_n$ characterises the $\sim_*^n$-type of b. In other words, b is m-critical if all reflexive successors within its $\sim_*^{m-1}$ class share its $\sim_*^m$-type. Clearly the m-critical nature of a node is determined by its $\sim_*^{m+1}$-type. The analogue of Observation 4.15 is then immediate.

Observation 4.28. (i) *If* $\mathfrak{B}, b \sim_*^{m+1} \mathfrak{B}, b'$, *then b is m-critical if, and only if, b' is.*
(ii) *For* $b \preccurlyeq b'$ *with* $\mathfrak{B}, b \sim_*^m \mathfrak{B}, b'$, *if b is m-critical then b' is m-critical.*

The notion of $\sim_*^m$-depth now needs to take into account reflexive and irreflexive nodes on paths within the $\sim_*^m$ class of an m-critical node. We therefore define two depths, $d_1^m(a)$ and $d_2^m(a)$, as follows.

For a node a let $d_1^m(a)$ be the maximal length n of a path $a \prec a_1 \prec \cdots \prec a_n$ consisting of reflexive nodes $a_i \sim_*^m a$; $d_2^m(a)$ the maximal length n of a path $a \prec a_1 \prec \cdots \prec a_n$ consisting, for $i \geqslant 1$, of irreflexive nodes $a_i \sim_*^m a$ with $d_1^m(a_i) = d_1^m(a)$.

The analogue of Observation 4.16 for d_1^m with respect to $\sim_*^m$ is then immediate.

Observation 4.29. *For every finite transitive tree-like $\mathfrak{A}, \alpha$, for all $d, m \in \mathbb{N}$: if a in $\mathfrak{A}$ is such that $d_1^m(a) > d$, then there is some reflexive $a' \succ a$ such that $\mathfrak{A}, a' \sim_*^m \mathfrak{A}, a'$ and $d_1^m(a') = d$.*

Note that $d_1^m(a) = 0$ if, and only if, a has no reflexive $\prec$-successors of the $\sim_*^m$-type of a.

For an irreflexive node a this is the case if, and only if, $\mathfrak{A}, a \models \neg\Diamond^*\chi_m^a$, where $\chi_m^a \in \mathrm{ML}_m^*$ characterises the $\sim_*^m$-type of a. So the $\sim_*^{m+1}$-type of an irreflexive node a determines whether or not $d_1^m(a) = 0$. Any reflexive node a, on the other hand, satisfies $\mathfrak{A}, a \models \Diamond^*\chi_m^a$ irrespective of $d_1^m(a)$. Indeed, reflexive nodes of arbitrary d_1^m-values may be bisimilar.

Note that, if a is m-critical, then for all $a' \sim_*^{m-1} a$ with $a \preccurlyeq a'$ we have $d_1^{m-1}(a') = d_1^m(a')$: generally $d_1^{m-1}(a') \geqslant d_1^m(a')$, but for m-critical nodes also $d_1^m(a') \geqslant d_1^{m-1}(a')$, because the reflexive successors of a within its $\sim_*^{m-1}$-class are members of its $\sim_*^m$-class. No such correspondence obtains in general between $d_2^m(a)$ and $d_2^{m-1}(a)$ for m-critical a.

We also note that, if $d_1^m(a) > 0$, d_2^m only accounts for the maximal number of steps across irreflexive nodes up to some next reflexive node of the same d_1^m value and within the same $\sim_*^m$ class. If $d_1^m(a) = 0$, then $d_2^m(a)$ accounts for the number of steps across irreflexive nodes within the same $\sim_*^m$ class, and only in this case is d_2^m determined by the $\sim_*$-type as follows.

$$\begin{array}{llcl} d_1^m(a) = 0,\ a \text{ irreflexive:} & d_2^m(a) \geqslant d & \Leftrightarrow & a \models (\Diamond\chi_m^a)^d\ ; \\ d_1^m(a) = 0,\ a \text{ reflexive:} & d_2^m(a) > d & \Leftrightarrow & a \models \Diamond\big(\neg\Diamond^*\chi_m^a\ \wedge\ (\Diamond\chi_m^a)^d\big). \end{array}$$

Here and in the following we use $(\Diamond\chi)^d$ as shorthand for the formula inductively obtained from $(\Diamond\chi)^1 := \Diamond\chi$ through $(\Diamond\chi)^{d+1} := \Diamond(\chi \wedge (\Diamond\chi)^d)$.

To illustrate the need for controlling d_2^m-values, consider the situation indicated in Figure 4. Here $\mathfrak{A} \simeq \mathfrak{B}$ (the actual R is the transitive closure of the indicated edges and there are no propositional symbols) but we look at distinguished nodes a and b such that $d_1^m(a) = d_1^m(b) = 2$, but $d_2^m(a) = 1$ while $d_2^m(b) = 0$ (for any m). Note that in fact all nodes with the exception of the last, which is a dead end, are bisimulation ($\sim$ and $\sim_*$) equivalent; in particular $\mathfrak{A}, a \sim \mathfrak{B}, b$. The indicated sequence of moves (a_i, b_i) in the game on $\mathfrak{A}_\prec^\circ, a$ versus $\mathfrak{B}_\prec^\circ, b$, with player I playing the a_i in $\mathfrak{A}$, shows how II is forced to violate depths in such a way that she loses within 4 rounds. We indicate P° (the reflexive nodes) through fat dots; as I avoids them, II has to avoid them, too.

Suitable bisimilar companions for the upgrading from $\sim_*^{\ell'}$ between $\mathfrak{A}$ and $\mathfrak{B}$ to $\sim^\ell$ between $\hat{\mathfrak{A}}_\prec^\circ$ and $\hat{\mathfrak{B}}_\prec^\circ$ are obtained as follows. We replace each reflexive node of $\mathfrak{A}$ by a chain of $(\ell+1)(\ell+2)$ many isomorphic copies of the node consisting of $\ell+1$ blocks, each of which has $\ell+1$ irreflexive copies followed by one reflexive copy of the node, and then pass to the transitive closure for the resulting R to obtain $\hat{\mathfrak{A}}$. Importantly, irreflexive nodes cannot be duplicated, let alone be given reflexive partner nodes, in

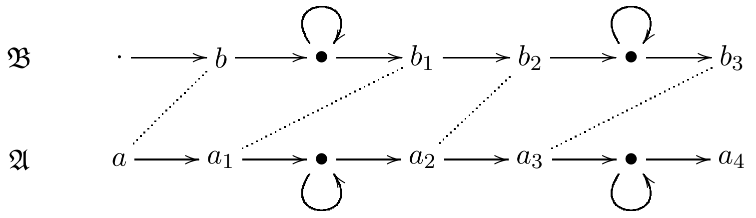

Figure 4: $\sim^\ell$ vs. $\sim^\ell_*$.

bisimilar companions. The chosen duplication of reflexive nodes suffices to give the following, though.

Observation 4.30. *For every finite rooted transitive tree-like structure $\mathfrak{A}, \alpha$ and $\ell \in \mathbb{N}$ there is a finite rooted transitive tree-like structure $\hat{\mathfrak{A}}, \hat{\alpha}$ such that $\hat{\mathfrak{A}}, \hat{\alpha} \sim \mathfrak{A}, \alpha$ and:*

(i) *every reflexive node a has a chain of $\ell - 1$ irreflexive direct predecessors a' that are bisimilar to a; in particular, the root $\hat{\alpha}$ is irreflexive.*

(ii) *for $m = \ell(\ell+1)$, either $d_1^m(\hat{\alpha}) = 0$, or $d_1^m(\hat{\alpha}), d_2^m(\hat{\alpha}) \geqslant \ell$.*

(iii) *for all a in $\hat{\mathfrak{A}}$ with $d_1^m(a) > d$ there is a reflexive $a' \succ a$ such that $a' \sim_*^m a$ and $d_1^m(a') = d$; if $d > 0$, then $d_2^m(a') \geqslant \ell$ can simultaneously be had.*

(iv) *for all a in $\hat{\mathfrak{A}}$ such that $d_1^m(a) > 0$ and $d_2^m(a) > d$ there is an irreflexive $a' \succ a$ such that $a' \sim_*^m a$, $d_1^m(a') = d_1^m(a)$ and $d_2^m(a') = d$.*

(v) *for all $a \prec a'$ such that $a' \not\sim a$, a' reflexive:*
there are reflexive, as well as irreflexive, $a'' \succ a$ with $a'' \sim a'$ and $d_1^m(a''), d_2^m(a'') \geqslant \ell$ for all m.

Again, we call such finite rooted transitive tree-like structures *ℓ-good*.

Regarding the badly controlled behaviour of irreflexive successors of m-critical nodes, we have the following in ℓ-good structures.

Observation 4.31. *In ℓ-good structures, if $a \models \Diamond(\chi^a_{m-1} \wedge \neg\chi^a_m \wedge \Diamond^*\chi^a_{m-1})$ for some m-critical node a, then $d_1^m(a) \geqslant \ell$.*

Proof. Note first that any $a' \succcurlyeq a$ such that $a' \models \chi^a_{m-1} \wedge \neg\chi^a_m \wedge \Diamond^*\chi^a_{m-1}$ can only be irreflexive, since $a \models \Box^*(\chi^a_{m-1} \to \chi^a_m)$. As $a' \models \Diamond^*\chi^a_{m-1}$, there is some reflexive $a'' \succ a'$ such that $a'' \sim_*^{m-1} a$, whence $a'' \sim_*^m a$ and $a' \not\sim_*^m a$ imply that $a' \not\sim a''$. By Observation 4.30 (v), there is some $a''' \succ a'$ such that $a''' \sim a''$ and $d_1^m(a''') \geqslant \ell$. As $a \sim_*^m a'''$, we have $d_1^m(a) \geqslant d_1^m(a''')$ and the claim follows. ☐

Claim 4.32. *Let $\ell > 2$. For finite transitive tree-like structures $\mathfrak{A}, \alpha$ and $\mathfrak{B}, \beta$ that are ℓ-good:*

$$\mathfrak{A}, \alpha \sim_*^{\ell^2+\ell+1} \mathfrak{B}, \beta \quad \Rightarrow \quad \mathfrak{A}^\circ_\prec, \alpha \sim^\ell \mathfrak{B}^\circ_\prec, \beta.$$

Proof. We show that the following system $(Z_k)_{k \leqslant \ell}$ satisfies the back-and-forth conditions for $\mathfrak{A}^\circ_\prec, a \sim^\ell \mathfrak{B}^\circ_\prec, b$. Again, $n \stackrel{k}{=} n'$ is shorthand for ($n = n'$ or $n, n' \geqslant k$). The following notion of m-safe pairs $(a, b) \in A \times B$ serves as a convenient abbreviation in the definition of the Z_k:

$$(a,b)\ m\text{-safe} \quad :\Leftrightarrow \quad \begin{cases} \mathfrak{A}, a \sim_*^m \mathfrak{B}, b, \\ a \text{ reflexive iff } b \text{ reflexive}, \\ a\ m\text{-critical iff } b\ m\text{-critical, and} \\ \text{if } a, b \text{ are } m\text{-critical and reflexive: } d_1^m(a) \stackrel{k}{=} d_1^m(b) \\ \text{if } a, b \text{ are } m\text{-critical and irreflexive: } d_i^m(a) \stackrel{k}{=} d_i^m(b) \text{ for } i = 1, 2. \end{cases}$$

We then put

$$Z_k := \bigl\{(a,b) \in A \times B \colon (a,b)\ m\text{-safe for some } m \geqslant k(\ell+1)\bigr\}.$$

Note that the roots of ℓ-good structures are irreflexive, and that the $\sim_*^n$-type of the root for $n = \ell^2 + \ell + 1$ determines whether $d_1^m = 0$ or $d_{1/2}^m \geqslant \ell$ for $m = \ell(\ell+1)$. It follows that $(\alpha, \beta) \in Z_\ell$, as (α, β) is m-safe for $m = \ell(\ell+1)$.

The verification of the back-and-forth property is analogous to, but slightly more involved than in the case of Claim 4.18. We consider the forth condition corresponding to a move along

$$a \prec a' \qquad \text{in } \mathfrak{A}^\circ_\prec,$$

from a position $(a,b) \in Z_k$, which is m-safe for some $m \geqslant k(\ell+1)$. We consider two main cases:

Case (1) $\quad \mathfrak{A}, a \not\sim_*^{m-1} \mathfrak{A}, a'$, or $\mathfrak{A}, a \sim_*^{m-1} \mathfrak{A}, a'$ but a and b are not m-critical.

Case (2) $\quad \mathfrak{A}, a \sim_*^{m-1} \mathfrak{A}, a'$, a and b are m-critical.

Case (1) We want to find b' such that $b \prec b'$ and such that (a', b') is m'-safe for $m' = m - \ell$, so that $(a', b') \in Z_{k-1}$ follows.

The strategy for $\mathfrak{A}, a \sim_*^m \mathfrak{B}, b$ will immediately give us some $b' \succcurlyeq b$ such that $b' \sim_*^{m-1} a'$. But we additionally need to make sure that $b \prec b'$, that b' is reflexive iff a' is, and that b' is $(m-\ell)$-critical iff a' is $(m-\ell)$-critical, and of matching depths in that case. This obtains for different reasons in the following sub-cases (1a)–(1c):

(1a) $\quad a \sim_*^{m-1} a'$, but a and b not m-critical.
(1b) $\quad a \not\sim_*^{m-1} a'$ and $a \models \Diamond^* \chi^{a'}_{m-1}$.
(1c) $\quad a \not\sim_*^{m-1} a'$ and $a \models \neg\Diamond^* \chi^{a'}_{m-1}$.

Note that the conditions in cases (1b) and (1c) automatically transfer to b, b' for any choice of $b' \succcurlyeq b$ with $b' \sim_*^{m-1} a'$; in these two cases, moreover, $b' \succ b$ is automatically guaranteed, as $b' \sim_*^{m-1} a' \not\sim_*^{m-1} a \sim_*^m b$ implies $b' \neq b$. In each one of the three cases, we first find some $b' \sim_*^{m-1} a'$, $b' \succ b$, b' reflexive iff a' reflexive. In particular, $b' \sim_*^{m-1} a'$ also implies both that $b' \sim_*^{m-\ell} a'$ and that b' is $(m-\ell)$-critical iff a' is, since this is determined by their $\sim_*^{m-\ell+1}$-type. The adjustment of the depths of a' and b', where necessary, will be addressed in a second step. The depths of a and b are irrelevant for case (1).

(1a) $\quad a \sim_*^{m-1} a'$; a, b not m-critical.

As b is not m-critical, we find some reflexive $b' \succcurlyeq b$ such that $b' \sim_*^{m-1} b$ but $b' \not\sim_*^m b$. In particular, $b' \not\sim b$ and $b \prec b'$. This implies, by Observation 4.30 (v), that we may replace b' by some $b'' \sim b'$ which is reflexive or irreflexive as desired in relation to a', and of depths $\geqslant \ell$ (to be adjusted later if necessary).

(1b) $a' \not\sim_*^{m-1} a$ and there is some a'' such that $a \prec a''$, a'' reflexive and $a'' \sim_*^{m-1} a'$.
In this case we find some b'' as a response to the $\Diamond^*$-move from a to a''. This guarantees that b'' is reflexive and $b'' \sim_*^{m-1} a'$, whence in particular also $b'' \succ b$ (since $b'' \sim_*^{m-1} a'$ implies $b'' \not\sim_*^{m-1} b$). By Observation 4.30 (v), we may replace such b'' by some $b' \sim b''$ which is reflexive or irreflexive as desired and of depths $\geqslant \ell$ (to be adjusted later if necessary).

(1c) $a' \not\sim_*^{m-1} a$ and there is no a'' such that $a \prec a''$, a'' reflexive and $a'' \sim_*^{m-1} a'$.
As $a \models \neg\Diamond^*\chi_{m-1}^{a'}$ and, as $a \sim_*^m b$, also $b \models \neg\Diamond^*\chi_{m-1}^{a'}$. In particular b' obtained as a response to the $\Diamond$-move from a to a' will automatically be irreflexive, too. As noted before, $b \succ b'$ is guaranteed as $b \not\sim_*^{m-1} b'$.

In each of the cases (1a)–(1c), since $b' \sim_*^{m-1} a'$, we know that b' is $(m-\ell)$-critical if, and only if, a' is. If both are not $(m-\ell)$-critical, then (a', b') is m'-safe for $m' := m - \ell$ already, and hence $(a', b') \in Z_{k-1}$. It remains to deal with the case that a', and therefore also b', are $(m-\ell)$-critical. For the necessary adjustment of depths, we deal with cases (1a) and (1b) together, and then with case (1c).

Cases (1a) and (1b). We have $d_1^{m-\ell}(b'), d_2^{m-\ell}(b') \geqslant \ell$ by construction so far.
We may use Observation 4.30 (iii) or (iv) to reduce depths to match those of a' in the sense of $\overset{k-1}{=}$ as follows.
For reflexive a' and b', reduce $d_1^{m-\ell}(b')$ if necessary by moving b' forward in its $\sim_*^{m-\ell}$ class to match $d_1^{m-\ell}(a')$.
For irreflexive a' and b' such that $d_1^{m-\ell}(a') > 0$ first adjust $d_1^{m-\ell}(b')$ according to Observation 4.30 (iii), then $d_2^{m-\ell}(b')$ according to Observation 4.30 (iv). For irreflexive a' and b' such that $d_1^{m-\ell}(a') = 0$ we claim that $a' \sim_*^{m-1} b'$ implies matching depths as required. Indeed,

$$\begin{aligned} &d_1^{m-\ell}(a') = 0 \;\Leftrightarrow\; a' \models \neg\Diamond^*\chi_{m-\ell}^{a'} \;\Leftrightarrow\; b' \models \neg\Diamond^*\chi_{m-\ell}^{a'} \;\Leftrightarrow\; d_1^{m-\ell}(b') = 0, \quad \text{and} \\ &d_2^{m-\ell}(a') \geqslant d \;\Leftrightarrow\; a' \models (\Diamond\chi_{m-\ell}^{a'})^d \;\Leftrightarrow\; b' \models (\Diamond\chi_{m-\ell}^{a'})^d \;\Leftrightarrow\; d_2^{m-\ell}(a') \geqslant d \quad \text{for } d \leqslant \ell - 1 \end{aligned}$$

imply the desired matches between $\sim_*^{m-\ell}$-depths.[7]
So (a', b') is m'-safe for $m' = m - \ell$, and hence in Z_{k-1}.

Case (1c). We know that a' and b' are irreflexive. As $a' \sim_*^{m-1} b'$, we know that $d_1^{m-\ell}(a') > 0$ iff $d_1^{m-\ell}(b') > 0$ as above. Moreover, in case $d_1^{m-\ell}(a') = d_1^{m-\ell}(b') = 0$, the above reasoning establishes matching $d_2^{m-\ell}$ as well. If $d_1^{m-\ell}(a'), d_1^{m-\ell}(b') > 0$, then we may first replace b' by some b'' such that $b \prec b''$, b'' reflexive, $b'' \sim_*^{m-\ell} b'$ but $b'' \not\sim b$. By Observation 4.30 (v) we may further replace b'' by some $b''' \sim b''$ which is irreflexive and of depths $\geqslant \ell$, which can then be adjusted to match those of a' by Observation 4.30 (iii) and (iv). Again (a', b') is m'-safe for $m' = m - \ell$, hence in Z_{k-1}.

We note that in case (1a), the assumption that $a' \sim_*^{m-1} a$ and that a and b are not m-critical, allowed us to find b' such that (a', b') is m'-safe for $m' = m - \ell$. It

[7] See page 45 for the shorthand $(\Diamond\chi)^d$.

is obvious from the construction that $a' \sim_*^{m-2} a$ for a and b that are not $(m-1)$-critical, would still yield b' such that $(a', b') \in Z_{k-1}$ on the basis of being M'-safe for $m' = m - \ell - 1 \geqslant (k-1)(\ell+1)$. In this sense, we could now narrow case (2) to the situation where a and b are m-critical *and* $(m-1)$-critical.

Case (2) $\mathfrak{A}, a \sim_*^{m-1} \mathfrak{A}, a'$, a and b both m-critical. Where this is useful we may, by the above remark, assume that a and b are also $(m-1)$-critical.

We distinguish sub-cases, depending on reflexivity and depths of the target node a':

(2a) a' reflexive.
(2b) a' irreflexive and $d_1^{m-1}(a') > 0$.
(2c) a' irreflexive and $d_1^{m-1}(a') = 0$.

In case (2a), where $a' \sim_*^m a$ is an immediate consequence of the m-critical nature of a, we match a' with some reflexive $b' \succ b$ such that $b' \sim_*^m b$ and such that $(a', b') \in Z_k$ on the basis of being m'-safe for $m' = m - 1$. Note that, since $b' \sim_*^m a'$, b' is guaranteed to be $(m-1)$-critical iff a' is. Similar reasoning is available in case (2b), only that $b' \sim_*^{m-1} a'$ may be too weak to deal with the condition on $(m-1)$-critical nature; here we can invoke the extra assumption that a and b are also $(m-1)$-critical, which implies that a' and b' are both $(m-1)$-critical. In case (2c) we may need to descend to m'-safe matches for $m' = m - \ell$ or $m' = m - s - 1$ for some $1 \leqslant s < \ell$, because the $\sim_*^m$-type of a and b does not give any guarantees even as to $d_2^{m-1}(a)$ or $d_2^{m-1}(b)$.

(2a) a' reflexive.

This is the most straightforward case: since $a' \succ a$ and $a' \sim_*^m a$, we know that $d_1^m(a) > 0$, whence $d_1^m(b) \overset{k}{=} d_1^m(a) > 0$ too. So there is a matching reflexive $b' \succ b$ for which $b' \sim_*^m a'$ and $d_1^m(b') \overset{k}{=} d_1^m(b)$. As a and b are m-critical, so are a' and b' and $d_1^{m-1}(a') = d_1^m(a') \overset{k}{=} d_1^m(b') = d^{m-1}(b')$ imply that $\sim_*^{m-1}$-depths match. Also, as not only $a' \sim_*^{m-1} b'$ but in fact $a' \sim_*^m b'$, b' is $(m-1)$-critical iff a' is. Therefore, (a', b') is m'-safe for $m' = m - 1$, hence in Z_{k-1}.

(2b) a' irreflexive and $d_1^{m-1}(a') > 0$.

For parts of the argument we use (w.l.o.g.) the additional assumption that a and b are also $(m-1)$-critical.

We first observe that $a' \sim_*^m a$. As $d_1^{m-1}(a') > 0$, there is some reflexive $a'' \succ a'$ for which $a'' \sim_*^{m-1} a'$ and hence $a'' \sim_*^{m-1} a$. As a is m-critical and a'' is reflexive, $a'' \sim_*^m a$. But $a'' \sim_*^m a$ where $a \prec a' \prec a''$ implies that any $\sim_*^{m-1}$-type realised by (reflexive) successors of a is also realised by (reflexive) successors of a'' and hence of a'. So $a' \sim_*^m a$ follows. This further implies that $d_1^{m-1}(a') = d_1^m(a')$.

How to find a matching b', depends on whether a and b are reflexive or irreflexive.

(i) Let a and b be reflexive. As $d_1^m(a) \overset{k}{=} d_1^m(b)$, we find in the $\sim_*^m$-class of b some irreflexive b' such that $d_1^m(b') \overset{k}{=} d_1^m(a')$ and $d_2^m(b') \geqslant \ell$. Now $d_1^{m-1}(b') = d_1^m(b') \overset{k}{=} d_1^m(a') = d_1^{m-1}(a')$, by the m-critical nature of a and b. As these depths are non-zero, $d_2^{m-1}(b') \geqslant d_2^m(b') \geqslant \ell$ may be adjusted downward by Observation 4.30 (iv) if necessary.

(ii) Let a and b be irreflexive. In this case $d_1^m(a) \overset{k}{=} d_1^m(b)$ and $d_2^m(a) \overset{k}{=} d_2^m(b)$ guarantee the existence of some irreflexive $b' \succ b$ such that $b' \sim_*^m b$ and $d_1^m(b') \overset{k}{=} d_1^m(a')$ and $d_2^m(b') \overset{k}{=} d_2^m(a')$. Now $d_1^{m-1}(b') = d_1^m(b')$ as $b' \sim_*^m b$ and as b is m-critical. So

$d_1^{m-1}(b') \stackrel{k}{=} d_1^{m-1}(a')$ is automatic. The adjustment of $d_2^{m-1}(b')$ to $d_2^{m-1}(a')$, however, is not immediate. Clearly $d_2^{m-1}(b') \geqslant d_2^m(b')$, but it could be that $d_2^{m-1}(a) > d_2^m(a)$, even though a is m-critical. But this can only be the case if the $\sim_*^{m-1}$ class of a' contains irreflexive members between a' and some next reflexive member that are not $\sim_*^m$ equivalent. In this case, Observation 4.31 tells us that $d_1^m(a') \geqslant \ell$ and therefore also $d_1^m(a) \geqslant \ell$ and $d_1^m(b) \geqslant k$. In this case, therefore, we find a new $b' \succ b$ such that $b' \sim_*^m b$ and $d_1^{m-1}(b') \geqslant k-1$ and $d_2^{m-1}(b') \geqslant \ell$, which can then be adjusted downward to match the $\sim_*^{m-1}$-depths of a' as above.

It remains to argue that b' is $(m-1)$-critical iff a' is. For this we use the additional assumption that a and b are also $(m-1)$-critical. This, together with $a' \sim_*^{m-1} a$, implies that a' and b' are both $(m-1)$-critical too. Consider a', for instance: $a' \models \diamond^*(\chi_{m-2}^{a'} \wedge \neg\chi_{m-1}^{a'})$ clearly implies the same for a. Similar reasoning shows that b' is $(m-1)$-critical. Therefore (a', b') is M'-safe for $m' = m-1$, so $(a', b') \in Z_{k-1}$.

(2c) a' irreflexive and $d_1^{m-1}(a') = 0$.

We again want to match a' with some irreflexive $b' \sim^{m'} a'$ such that (a', b') is m'-safe for some $m' \geqslant m-\ell-1$.

We treat the situation differently, depending on whether or not $d_1^{m-\ell}(a') = 0$.

(i) Assume first that $d_1^{m-\ell}(a') = 0$. We use $m' = m-\ell$. If $d_2^{m-\ell}(a') = d < k-1$, then

$$a \models \diamond\big(\chi_{m-\ell}^{a'} \wedge \neg\diamond^*\chi_{m-\ell}^{a'} \wedge (\diamond\chi_{m-\ell}^{a'})^d \wedge \neg(\diamond\chi_{m-\ell}^{a'})^{d+1}\big).$$

As $b \sim_*^m a$ and as the nesting depth of this formula is at most m, the same applies to b and we find some $b' \succcurlyeq b$ with

$$b' \models \chi_{m-\ell}^{a'} \wedge \neg\diamond^*\chi_{m-\ell}^{a'} \wedge (\diamond\chi_{m-\ell}^{a'})^d \wedge \neg(\diamond\chi_{m-\ell}^{a'})^{d+1}.$$

It follows that $b' \succ b$, as either b is irreflexive or b is reflexive and hence $b \models \diamond^*\chi_{m-\ell}^{a'}$. Clearly b' is $(m-\ell)$-critical iff a' is: this is determined by whether or not $b' \models \diamond^*\chi_{m-\ell-1}^{a'}$. It follows that (a', b') is m'-safe for $m' = m-\ell$.

If $d_2^{m-\ell}(a') = d \geqslant k-1$, the formula $\diamond\big(\chi_{m-\ell}^{a'}\neg\diamond^*\chi_{m-\ell}^{a'} \wedge \diamond^{k-1}\chi_{m-\ell}^{a'}\big)$ can be used instead.

(ii) Assume now that $d_1^{m-\ell}(a') > 0$. We find the desired level m' as $m' = m-s-2$ where s is determined as the least $1 \leqslant s < \ell$ such that $d_1^{m-s-1}(a') > 0$. As $0 = d_1^{m-s}(a') < d_1^{m-s-1}(a')$ it is clear that a' is not $(m-s)$-critical. As $a' \sim_*^{m-1} a \sim_*^m b$, both a and b are not $(m-s)$-critical. It follows that there is some reflexive $b'' \succ b$ such that $b'' \sim_*^{m-s-1} b$ but $b'' \not\sim_*^{m-s} b$. Just as in case (1a) we find an irreflexive $b' \succ b$ such that $b' \sim_*^{m-s-1} b$ and all depths $\geqslant \ell$. As $b' \sim_*^{m-s-1} b \sim_*^m a \sim_*^{m-s-1} a'$, b' is $(m-s-2)$-critical iff a' is. If a' and b' are $(m-s-2)$-critical, we may finally adjust depths as follows: first adjust $d_1^{m-s-2}(b')$ to match $d_1^{m-s-2}(a') \geqslant d_1^{m-s-1}(a') > 0$ according to to Observation 4.30 (iii), then adjust $d_2^{m-s-2}(b')$ to match $d_2^{m-s-2}(a')$ according to Observation 4.30 (iv).

So we get $(a', b') \in Z_{k-1}$, as (a', b') is m'-safe for $m' = m-s-2 \geqslant m-\ell-1 \geqslant (k-1)(\ell+1)$. □

As every finite transitive tree-like structure is bisimilar to an ℓ-good one by Observation 4.30, we get the following.

Corollary 4.33. *For every fixed finite vocabulary σ there is a function f_1 such that for all ℓ and any finite rooted transitive tree-like structures $\mathfrak{A}, \alpha \sim_*^{f_1(\ell)} \mathfrak{B}, \beta$, there are bisimilar companions $\tilde{\mathfrak{A}}, \tilde{\alpha} \sim \mathfrak{A}, \alpha$ and $\tilde{\mathfrak{B}}, \tilde{\beta} \sim \mathfrak{B}, \beta$ within the class of finite rooted transitive tree-like structures such that $\tilde{\mathfrak{A}}^\circ_\prec, \tilde{\alpha} \sim^\ell \tilde{\mathfrak{B}}^\circ_\prec, \tilde{\beta}$.*

Combining Claims 4.10 and Corollary 4.33, we obtain the following upgrading from $\sim_*^\ell$ to $\equiv_q$, which in particular proves Theorem 4.27.

Corollary 4.34. *For $q \in \mathbb{N}$ and fixed finite vocabulary σ there, there is an $\ell \in \mathbb{N}$ such that for any two finite transitive tree-like structures $\mathfrak{A}, \alpha \sim_*^\ell \mathfrak{B}, \beta$, there are finite transitive tree-like structures $\hat{\mathfrak{A}}, \hat{\alpha} \sim \mathfrak{A}, \alpha$ and $\hat{\mathfrak{B}}, \hat{\beta} \sim \mathfrak{B}, \beta$ such that $\hat{\mathfrak{A}}, \hat{\alpha} \equiv_q \hat{\mathfrak{B}}, \hat{\beta}$.*

Path-finite transitive tree-like frames. Note that path-finite transitive tree-like frames are those transitive tree-like frames whose irreflexive part is a Löb frame. In transitive tree-like frames, $R \cap R^{-1} = R^\circ$ and $R \setminus R^{-1}$ is the irreflexive part of R, i.e., we forbid infinite $\prec$-paths.

For not necessarily finite but path-finite transitive tree-like frames, $\sim_*^m$-depths d_1^m and d_2^m are definable as ordinal-valued ranks of $\succ$ (in restriction to the reflexive members of the $\sim_*^m$ class of a for $d_1^m(a)$, and in restriction to those irreflexive members that have the same d_1^m-value as a, for $d_2^m(a)$). The transformation to a bisimilar ℓ-good structure described in connection with Observation 4.30 can similarly be applied to not necessarily finite transitive tree-like structures without infinite $\prec$-paths and produces a bisimilar companion within this class. The analogue of Claim 4.32 then yields analogues of Corollaries 4.34 and 4.33 over the class of not necessarily finite transitive tree-like structures without infinite $\prec$-paths. In particular we find the following.

Corollary 4.35. *For every fixed finite vocabulary σ there is a function f_1 such that for all ℓ and any path-finite transitive tree-like structures $\mathfrak{A}, \alpha \sim_*^{f_1(\ell)} \mathfrak{B}, \beta$, there are bisimilar companions $\tilde{\mathfrak{A}}, \tilde{\alpha} \sim \mathfrak{A}, \alpha$ and $\tilde{\mathfrak{B}}, \tilde{\beta} \sim \mathfrak{B}, \beta$ within the class of path-finite transitive tree-like structures such that $\tilde{\mathfrak{A}}^\circ_\prec, \tilde{\alpha} \sim^\ell \tilde{\mathfrak{B}}^\circ_\prec, \tilde{\beta}$. It follows that for every q there is an ℓ such that $\sim_*^\ell$ can be upgraded to $\equiv_q$ also within the class of path-finite transitive tree-like structures.*

Corollary 4.36. $\mathrm{FO}/{\sim} \equiv \mathrm{ML}^*$ *over the class of path-finite transitive tree-like structures.*

4.6 Finite transitive structures

If a (finite) structure $\mathfrak{A}$ is on a frame that is irreflexive and transitive, in other words it is a strict partial order, then it is necessarily bisimilar to a (finite) $\prec$-tree by a simple unravelling.

We wish now to consider finite transitive frames which are not necessarily acyclic. Recall from Definition 2.8 (vi) that a rooted transitive frame with accessibility relation R is *weak transitive tree-like* if it satisfies the condition $(\neg Rxy \wedge \neg Ryx) \rightarrow \neg(Rxz \wedge Ryz)$. While a finite rooted transitive structure cannot be bisimilar to any *finite* transitive tree-like structure if it contains non-trivial R-cycles (or non-trivial R-cliques), any finite

$$
\begin{tikzcd}
\mathfrak{A}, \alpha \arrow[r, dash] & \sim_*^{f_1(\ell)} \arrow[r, dash] & \mathfrak{B}, \beta & \text{finite transitive tree-like} \\
\sim \arrow[u, dash] \arrow[d, dash] & & \sim \arrow[u, dash] \arrow[d, dash] & \\
\tilde{\mathfrak{A}}, \alpha \arrow[r, dash] & \sim_*^{f_1(\ell)} \arrow[r, dash] & \tilde{\mathfrak{B}}, \beta & \text{finite } \ell\text{-good transitive tree-like} \\
& \equiv_q & & \\
& \equiv_q \arrow[u, dotted, Rightarrow] & & \\
\tilde{\mathfrak{A}}^{\circ}_{\prec}, \alpha \arrow[r, dash] \arrow[uuu, dotted, leftrightarrow] & \sim^{\ell} \arrow[r, dash] & \tilde{\mathfrak{B}}^{\circ}_{\prec}, \beta \arrow[uuu, dotted, leftrightarrow] & \text{finite } \prec\text{-trees}
\end{tikzcd}
$$

Figure 5: the upgrading in Corollaries 4.33 and 4.34.

rooted transitive structure is bisimilar to a finite *weak* transitive tree-like structure. Hence, every finite pointed transitive structure is bisimilar to a finite weak transitive tree-like structure.

We now want to associate to any weak transitive tree-like structure a transitive tree-like structure in a manner that is compatible with finiteness, bisimulation equivalence, and with FO definability.

Define $E^{\mathfrak{A}} := \big\{(a,b) \colon a = b \text{ or } (a,b),(b,a) \in R^{\mathfrak{A}}\big\}$. If $R^{\mathfrak{A}}$ is transitive, then $E^{\mathfrak{A}}$ is an equivalence relation. Equivalence classes of nodes $a \in \mathfrak{A}$ w.r.t. E will be denoted as

$$[a] := \{a\} \cup \{a' \in A \colon (a,a'),(a',a) \in R^{\mathfrak{A}}\}.$$

Any weak transitive tree-like frame can be seen as a transitive tree-like frame of E equivalence classes – and this observation is used to interpret any weak transitive tree-like structure in a tree-like structure as follows. In order to have a unique identification of the root in terms of the accessibility relation R, let us assume that the root is irreflexive (which it is, up to bisimulation equivalence, without loss of generality).

Fix a vocabulary σ consisting of the binary relation symbol R and a set $P_1, \ldots, P_k$ of unary predicates. Consider a rooted transitive structure $\mathfrak{A}, \alpha$ over this vocabulary with irreflexive root α.

We associate with $\mathfrak{A}$ a structure $t(\mathfrak{A})$ over a vocabulary $t(\sigma)$ consisting of R and a unary predicate S for each set $S \subseteq \{P_1, \ldots, P_k\}$. Each such S encodes a complete propositional type with respect to $P_1, \ldots, P_k$, which is also represented by a modal formula ζ_S of nesting depth 0. For an element of the given structure $\mathfrak{A}$ we also let $S(a)$ denote the complete propositional type of a, $S(a) = \{P_i \colon a \in P_i^{\mathfrak{A}}\}$.

For a rooted transitive σ-structure $\mathfrak{A}, \alpha$ with irreflexive root α define the $t(\sigma)$-structure

$$\mathfrak{A}/E := \big(A/E^{\mathfrak{A}}, R^{\mathfrak{A}/E}, (S^{\mathfrak{A}/E})\big),$$

on the set of equivalence classes of A with respect to $E^{\mathfrak{A}}$ as follows.

We put $([a],[b]) \in R^{\mathfrak{A}/E}$ if, and only if, $(a,b) \in R^{\mathfrak{A}}$. This is easily seen to be well defined due to the transitivity of $R^{\mathfrak{A}}$. Note that $[a]$ is a reflexive node of $\mathfrak{A}/E$ if, and only if, a is a reflexive node of $R^{\mathfrak{A}}$ (which in particular is true of any non-singleton $[a]$, i.e., of non-trivial cliques with respect to $R^{\mathfrak{A}}$). Note that the root $[\alpha]$ is irreflexive, since α was; we identify $[\alpha] = \{\alpha\}$ with α itself.

For the unary predicates, we put $[a] \in S^{\mathfrak{A}/E}$ precisely for those S for which $S = S(a')$ for some $a' \in [a]$. That is, S holds at $[a]$ if some element in the same equivalence class as a has propositional type S. For later use we also define

$$p(a) = p([a]) := \{S(a') \colon a' \in [a]\} \subseteq \mathcal{P}(\{P_1, \dots, P_k\}),$$

so that $[a] \in S^{\mathfrak{A}/E}$ iff $S \in p(a)$.

Note that precisely one $S = S(\alpha)$ will hold at the irreflexive root α.

It is clear that $\mathfrak{A}/E$ is transitive tree-like for any finite weak transitive tree-like $\mathfrak{A}$ with irreflexive root α. For an arbitrary pointed transitive structure $\mathfrak{A}, \alpha$ (with an irreflexive distinguished node α), we may similarly obtain a transitive tree-like $t(\mathfrak{A}, \alpha)$ as follows.

Definition 4.37. For a pointed transitive σ-structure $\mathfrak{A}, \alpha$ with irreflexive α let $t(\mathfrak{A})$ be the following transitive tree-like $t(\sigma)$-structure:

$$t(\mathfrak{A}, \alpha) := \begin{cases} \mathfrak{A}/E & \text{if } \mathfrak{A}, \alpha \text{ is weak transitive tree-like,} \\ \mathrm{TC}\big((\mathfrak{A}/E)^*_\alpha\big) & \text{otherwise.} \end{cases}$$

The structure $\mathrm{TC}\big((\mathfrak{A}/E)^*_\alpha\big)$ stipulated in the second case above is obtained in the following steps:

$$\mathfrak{A}, \alpha \longmapsto \mathfrak{A}/E \longmapsto \big((\mathfrak{A}/E)^\circ\big)^*_\alpha \longmapsto \mathrm{TC}\big((\mathfrak{A}/E)^*_\alpha\big).$$

The first transformation describes passage to the quotient frame with respect to E with the natural interpretation of new unary $S \in t(\sigma)$ on equivalence classes as explained above (resulting in a transitive $t(\sigma)$-structure). The second step is an unravelling with respect to the irreflexive part of $R^{\mathfrak{A}/E}$ in the presence of a marker predicate P° for the reflexive nodes with respect to $R^{\mathfrak{A}/E}$ (resulting in a $t(\sigma) \cup \{P^\circ\}$-tree structure). In the third step we pass to the transitive closure with respect to the edge relation (this would give a $\prec$-tree of type $t(\sigma) \cup \{P^\circ\}$) and re-institute reflexive R-edges instead of the marker predicate P° (resulting in a transitive tree-like $t(\sigma)$-structure).

It is obvious that $t(\mathfrak{A})$ is finite for finite $\mathfrak{A}$, and is path-finite if $\mathfrak{A}$ is.

Since α was irreflexive in $\mathfrak{A}$, its equivalence class is a singleton, which remains irreflexive in $t(\mathfrak{A}, \alpha)$ and is contained in precisely one of the new unary predicates S of $t(\sigma)$, which we denote by $S(\alpha)$. Any node a of $t(\mathfrak{A}, \alpha)$ satisfies at least one of the predicates S, and any node satisfying more than one S is reflexive. Call a transitive tree-like $t(\sigma)$-structure *consistent* if it satisfies these conditions.

In the converse direction, we may now associate a weak transitive tree-like structure $\tau(\mathfrak{A}, \alpha)$ to any consistent transitive tree-like $t(\sigma)$-structure as follows.

Definition 4.38. Let $\mathfrak{A}, \alpha$ be a consistent transitive tree-like $t(\sigma)$-structure. Let the universe of $\tau(\mathfrak{A})$ be the set $A^{\tau(\mathfrak{A})} := \{(a, S) \colon a \in S^{\mathfrak{A}}, \} \subseteq A \times \mathcal{P}(\{P_1, \dots, P_k\})$. Put

$$\begin{aligned} \tau(\mathfrak{A}, \alpha) := \ & \big(A^{\tau(\mathfrak{A})}, R^{\tau(\mathfrak{A})}, P_1^{\tau(\mathfrak{A})}, \dots, P_k^{\tau(\mathfrak{A})}\big) \\ \text{where} \ & R^{\tau(\mathfrak{A})} = \{((a, S), (a', S')) \colon (a, a') \in R^{\mathfrak{A}}\}, \\ & P_i^{\tau(\mathfrak{A})} = \{(a, S) \colon P_i \in S\}. \end{aligned}$$

Note that the interpretation of the $P_i^{\tau(\mathfrak{A})}$ is precisely such that the propositional type of (a, s) in $\tau(\mathfrak{A})$ is S.

One checks that $\tau(\mathfrak{A}), \alpha$ (where we identify α with the pair $(\alpha, S(\alpha))$) is a weak transitive tree-like σ-structure. Clearly $\tau(\mathfrak{A})$ is finite for finite $\mathfrak{A}$, and the passage from $\mathfrak{A}$ to $\tau(\mathfrak{A})$ cannot introduce infinite strict paths.

The transformation τ is compatible with first-order equivalence, since it is given by a quantifier free first-order interpretation. Further, τ is compatible with $\sim$. And, up to bisimulation, τ acts as an inverse to t. We sum up these crucial properties in the following observation.

Observation 4.39.
For any consistent transitive tree-like $t(\sigma)$-structures $\mathfrak{A}, \alpha$ and $\mathfrak{B}, \beta$:
(i) $\mathfrak{A}, \alpha \equiv_q \mathfrak{B}, \beta \ \Rightarrow \ \tau(\mathfrak{A}), \alpha \equiv_q \tau(\mathfrak{B}), \beta$.
(ii) $\mathfrak{A}, \alpha \sim \mathfrak{B}, \beta \ \Rightarrow \ \tau(\mathfrak{A}), \alpha \sim \tau(\mathfrak{B}), \beta$.
For any transitive rooted σ-structure $\mathfrak{A}, \alpha$:
(iii) $\mathfrak{A}, \alpha \sim \tau(t(\mathfrak{A})), \alpha$.

Note for the last point, that a bisimulation is induced by the natural map that sends $a \in A$ to the pair $([a], S(a))$ in $\tau(t(\mathfrak{A}))$.

Moreover, the composition $\tau \circ t$ is a projection from the class of all rooted transitive σ-structures onto some subclass of weak transitive tree-like σ-structures.

Up to bisimulation equivalence, we may thus replace any rooted transitive structure $\mathfrak{A}, \alpha$ by its image $\tau(t(\mathfrak{A}))$ under this projection and in this situation then assume that $\mathfrak{A}, \alpha \simeq \tau(t(\mathfrak{A}))$. In fact we shall in the following simplify matters notationally by assuming actual equality:

$$\begin{aligned} \mathfrak{A} &= \tau(t(\mathfrak{A})), \\ a &= ([a], S(a)). \end{aligned} \qquad (\dagger)$$

We can then liberally pass between an element $a \in \mathfrak{A}$ and its E-class $[a] \in t(\mathfrak{A})$, which together with the propositional σ-type $S(a)$ of a identifies a.

In this view, or for $\mathfrak{A}$ pre-processed in this manner, t just acts like a quotient operation with respect to E, the reflexive closure of the symmetric part of R.

For the desired upgrading, we now seek a translation within the class of such weak transitive tree-like σ-structures, $\mathfrak{A}, \alpha \longmapsto \hat{\mathfrak{A}}, \alpha$ such that:

$$\begin{aligned} \mathfrak{A}, \alpha \ &\sim \ \hat{\mathfrak{A}}, \alpha, \\ \mathfrak{A} \text{ finite/ path-finite} \ &\Rightarrow \ \hat{\mathfrak{A}} \text{ finite/ path-finite}, \\ \mathfrak{A}, \alpha \sim_*^{f_2(\ell)} \mathfrak{B}, \beta \ &\Rightarrow \ t(\hat{\mathfrak{A}}), \alpha \sim_*^{\ell} t(\hat{\mathfrak{B}}), \beta \quad \text{for suitable } f_2(\ell). \end{aligned}$$

$$
\begin{tikzcd}
t(\hat{\mathfrak{A}}), \alpha \arrow[rr, dash, "\sim^{\ell}_{*}"] \arrow[d, dash, "\sim"] & & t(\hat{\mathfrak{B}}), \beta \arrow[d, dash, "\sim"] \\
\big(t(\hat{\mathfrak{A}})\big)^{\sim}, \alpha \arrow[rr, dash, "\equiv_q"] & & \big(t(\hat{\mathfrak{B}})\big)^{\sim}, \beta
\end{tikzcd}
\qquad
\begin{tikzcd}
\hat{\mathfrak{A}}, \alpha \arrow[rr, dash, "\sim^{f_2(\ell)}_{*}"] \arrow[d, dash, "\sim"] & & \hat{\mathfrak{B}}, \beta \arrow[d, dash, "\sim"] \\
\tilde{\mathfrak{A}}, \alpha \arrow[rr, dash, "\equiv_q"] & & \tilde{\mathfrak{B}}, \beta
\end{tikzcd}
$$

Figure 6: the upgrading for Theorems 4.40 and 4.41.

Corollary 4.34 (for finite $\mathfrak{A}$) and Corollary 4.35 (for path-finite $\mathfrak{A}$) imply that, for sufficiently large $\ell = \ell(q)$ there are bisimilar companions $(t(\hat{\mathfrak{A}}))^{\sim}$ and $(t(\hat{\mathfrak{B}}))^{\sim}$ of $t(\hat{\mathfrak{A}})$ and $t(\hat{\mathfrak{B}})$ as in Figure 6 on the left.

We now pass to the τ-images of these $t(\sigma)$-structures, in a situation where $\mathfrak{A} \sim^{f_2(\ell)}_{*} \mathfrak{B}$, and hence $\hat{\mathfrak{A}} \sim^{f_2(\ell)}_{*} \hat{\mathfrak{B}}$, and both structures have been pre-processed such that $\tau(t(\mathfrak{A})) = \mathfrak{A}$ and $\tau(t(\mathfrak{B})) = \mathfrak{B}$. The resulting situation is depicted on the right in Figure 6. Here $\tilde{\mathfrak{A}}$ and $\tilde{\mathfrak{B}}$ denote the τ-images of the bisimilar companions $\big(t(\hat{\mathfrak{A}})\big)^{\sim}$ and $\big(t(\hat{\mathfrak{B}})\big)^{\sim}$. Observation 4.39 (ii) guarantees that $\tilde{\mathfrak{A}}, \alpha \sim \tau(t(\hat{\mathfrak{A}})), \alpha = \hat{\mathfrak{A}}, \alpha \sim \mathfrak{A}, \alpha$ and similarly for $\hat{\mathfrak{B}}$. Part (i) guarantees that $\tilde{\mathfrak{A}}, \alpha \equiv_q \tilde{\mathfrak{B}}, \beta$.

The resulting upgrading of some level $\sim^{f_2(\ell)}_{*}$ to $\equiv_q$ in bisimilar companions within the classes of finite/path-finite rooted transitive σ-structures then yields the following.

Theorem 4.40. $\mathrm{FO}/{\sim} \equiv \mathrm{ML}^*$ *over the class of all finite rooted transitive frames and the class of all finite transitive frames.*

Note that the extension from rooted to arbitrary finite transitive frames is trivial, by passage to the substructure of reachable nodes (generated subframe).

Contrast this theorem with Theorem 2.12, which says that $\mathrm{FO}/{\sim} \equiv \mathrm{ML}$ over the classes of all (rooted) transitive frames, admitting infinite frames.

The analogue of Theorem 4.40 for path-finite transitive frames is similarly obtained.

Theorem 4.41. $\mathrm{FO}/{\sim} \equiv \mathrm{ML}^*$ *over the class of path-finite transitive structures.*

It remains to exhibit the translation $\mathfrak{A} \longmapsto \hat{\mathfrak{A}}$. This requires some preparation, for which we introduce the following terminolgy.

Reduced bisimulation types. In the passage from a to its E-class $[a]$ we shall want to strip this type of its direct propositional information and thus retain just the information concerning elements reachable by one or more steps from a. To this end we introduce *reduced bisimulation types* and their characteristic formulae, comprising all but the local propositional information about the bisimulation type. Reduced $\sim^{\ell}_{*}$-types, for $\ell \geqslant 1$, are defined by nesting depth ℓ formulae $\chi^a_{0;\ell} \in \mathrm{ML}^*_\ell$ in such a manner that $\chi^a_{0;\ell}$ consists of just the $\Box$-ed, $\Box^*_q$-ed, $\Diamond$-ed and $\Diamond^*_p$-ed conjuncts of χ^a_ℓ and such that $\chi^a_\ell = S(a) \wedge \chi^a_{0;\ell}$. We write

$$a \sim^{0;\ell}_{*} a' \quad \text{and} \quad a \sim^{0;\infty} a'$$

to indicate that a and a' have the same reduced $\sim_*^\ell$-type or reduced $\sim$-type. Only if additionally $S(a) = S(a')$ do these imply $a \sim_*^\ell a'$ and $a \sim a'$ as well. All elements of $[a]$ share the same reduced $\sim$-type and reduced $\sim_*^\ell$-type for every ℓ:

$$[a] = [a'] \quad \Rightarrow \begin{cases} a \sim^{0;\infty} a' & \text{and} \\ a \sim_*^{0;\ell} a' & \text{for all } \ell. \end{cases}$$

Saturation with respect to p-values. Recall that we need to deal just with transitive rooted σ-structures $\mathfrak{A}, \alpha = \tau(t(\mathfrak{A})), \alpha$ via the identification of $a \in A$ with the pair $([a], S(a))$ recording the E equivalence class of a in $\mathfrak{A}$ and its propositional type $S(a)$. We also assume irreflexive roots α.

The set of atomic propositions $S \in t(\sigma)$ that are true at a node $[a]$ of $t(\mathfrak{A})$ precisely records $p(a) = p([a]) = \{S(a') \colon a' \in [a]\}$. But $p(a)$ is not determined even by the full bisimulation type of a. This is what makes the translation non-trivial and requires an extra step of finitary saturation.

Example 4.42. Consider the following modification of a given structure $\mathfrak{A}$. Let $[a] \in t(\mathfrak{A})$ be such that $|p([a])| > 1$, so that there is some non-trivial subset $p' \subsetneq p([a])$. (Since $|p([a])| > 1$, $[a]$ and a must be refelxive nodes in $t(\mathfrak{A})$ and $\mathfrak{A}$.) Replace the substructure induced by the E-class $[a]$ in $\mathfrak{A}$, which is single maximal R-clique, by the succession of two cliques: one realising just the propositional types in p' followed by an isomorphic copy of the old clique $[a]$, which realises all types in $p([a])$. The resulting structure is bisimilar to the original $\mathfrak{A}$, while a new element of different propositional type p' has been inserted in its t-image before the element $[a]$.

While $p([a])$ is not determined by the bisimulation type of a, maximal p-values for reachable realisations of the bisimulation type of a are. Moreover, if $\mathfrak{A}$ is path-finite, then there is a unique such maximal p.

Observation 4.43. *Consider $a \in \mathfrak{A}$ for some transitive path-finite structure $\mathfrak{A}$. If a is reflexive, then there is a unique maximal set p of propositional σ-types for which there is an infinite path from a along which all $S \in p$ as well as the reduced $\sim^{0;\infty}$-type of a occur infinitely often. We let $p^\infty(a)$ stand for this p.*[8]

We show that there cannot be two distinct maximal such p. Suppose $p_1 \neq p_2$ can both be realised infinitely often along paths of the required kind, but, since both are maximal, $p_1 \cup p_2$ cannot. By path-finiteness, we can reach from a an R-clique in which all the propositional types from p_i and the reduced type of a are realised, for $i = 1$ and another one for $i = 2$. These R-cliques must be disjoint from each other and at least one of them must therefore be disjoint from $[a]$. So we find a realisation a' of the reduced type of a for which $(a, a') \in R \setminus R^{-1}$. Iteration of this step would yield an infinite path with respect to $R \setminus R^{-1}$, contradicting path-finiteness of $\mathfrak{A}$.

We now fix a parameter $n \in \mathbb{N}$ and construct $\hat{\mathfrak{A}}$ from $\mathfrak{A}$ and $t(\mathfrak{A})$ as follows: we first create a new $t(\sigma)$-structure $t_n(\mathfrak{A})$ and then put $\hat{\mathfrak{A}} := \tau(t_n(\mathfrak{A}))$.

[8] For irreflexive a it may be that there is no such p, if there is no path visiting the $\sim^{0;\infty}$-type of a infinitely often, but uniqueness still obtains.

The idea of the construction of $t_n(\mathfrak{A})$ is to insert, in front of any reflexive node $[a]$ with non-trivial $p([a])$ all possible successions of piecewise monotone chains of cliques $[a']$ with $p([a']) \subseteq p([a])$ that reach up to the maximal p-value $p([a])$ up to n times.

We give an explicit descripton of the structure $t_n(\mathfrak{A})$.

Construction of $t_n(\mathfrak{A})$ and $\hat{\mathfrak{A}}$. We assume that $\mathfrak{A}, \alpha$ has an irreflexive root α and that $\mathfrak{A} = \tau(t(\mathfrak{A}))$. The construction of the auxiliary $t_n(\mathfrak{A})$ is based on the transitive tree-like $t(\sigma)$-structure $t(\mathfrak{A})$. In accordance with the discussion in (†) above, we use the assumption that $\mathfrak{A} = \tau(t(\mathfrak{A}))$ for an explicit represention of the node set A of $\mathfrak{A} = \tau(t(\mathfrak{A})$ by the set of pairs

$$\big\{([a], S)\colon [a] \in S^{t(\mathfrak{A})}\big\} \subseteq t(\mathfrak{A}) \times \mathcal{P}(\{P_1, \dots, P_k\})$$

with the natural identifications

$$a \longmapsto ([a], S(a)) \quad \text{and} \quad ([a], S) \longmapsto \textit{the } a \in [a] \text{ with } \mathrm{atp}(a) = S,$$

and irreflexive root $\alpha = ([\alpha], S(\alpha))$.

We let $\mathcal{P} := \mathcal{P}\big(\mathcal{P}(\{P_1, \dots, P_k\})\big) \setminus \{\emptyset\}$ stand for the set of possible p-values.

From $t(\mathfrak{A})$ to $t_n(\mathfrak{A})$. Then $t_n(\mathfrak{A})$ is obtained as the transitive unravelling of the following irreflexive transitive structure

$$\big(\Gamma, \prec, P^\circ, (S)_{S \subseteq \{P_1,\dots,P_k\}}\big)$$

in vocabulary $t(\sigma) \cup \{P^\circ\}$. We write $\prec$ for the irreflexive, transitive edge relation, and use P° as a marker predicate for reflexive nodes. The universe of this structure is the following set $\Gamma \subseteq t(\mathfrak{A}) \times \{0, \dots, n\} \times \mathcal{P}$,

$$\Gamma := \big\{([a], m, p)\colon [a] \in t(\mathfrak{A}), \emptyset \neq p \subseteq p([a]), 0 \leqslant m \leqslant n \text{ and } m \neq 0 \text{ iff } [a] \text{ reflexive } \big\}$$

$$\begin{array}{lrcl} \text{with associated projection} & \pi\colon \quad \Gamma & \longrightarrow & t(\mathfrak{A}) \\ & ([a], m, p) & \longmapsto & [a] \\ \text{and distinguished root} & \multicolumn{3}{l}{\alpha := ([\alpha], 0, S(\alpha)) \in \Gamma.} \end{array}$$

Recall that the root α of $\mathfrak{A}$ was irreflexive, whence $[\alpha]$ is a singleton with $p(\alpha) = \{S(\alpha)\}$, which induces this unique new α. For the irreflexive transitive edge relation we put

$$\begin{array}{llll} ([a], m, p) \prec ([a'], m', p') & \text{if} & [a] \prec [a'] & \\ & \text{or} & [a] = [a'] \text{ and } & m' < m \\ & & & \text{or } m = m' \text{ and } p' \supsetneq p. \end{array}$$

The marker predicate for reflexive nodes is interpreted as

$$P^\circ := \{([a], m, p) \in \Gamma\colon [a] \text{ reflexive } \}.$$

Finally put

$$([a], m, p) \in S \text{ if } S \in p,$$

so that the propositional $t(\sigma)$-type of $([a], m, p)$ is p.

Now $t_n(\mathfrak{A})$ is the transitive tree-like $t(\sigma)$-structure obtained as the transitive unravelling of this structure from the irreflexive new root $\alpha = ([\alpha], 0, S(\alpha))$, with reflexive nodes according to P°.

Explicitly, the node set of $t_n(\mathfrak{A})$ is the set Γ^*_α of all finite $\prec$-paths from this new root α in $(\Gamma, \prec, P^\circ, (S))$. We again denote as π the natural projection

$$\pi \colon \Gamma^*_\alpha \longrightarrow t(\mathfrak{A}),$$

which now maps a path $\gamma = ([\alpha], 0, \{S(\alpha)\}) \cdots ([a], m, p) \in \Gamma^*_\alpha$ to $[a]$.

The transitive accessibility relation $\preccurlyeq$ of $t_n(\mathfrak{A})$ is the irreflexive prefix relation among these paths with additional reflexive edges at all nodes in $\pi^{-1}(P^\circ)$.

Predicates S are interpreted in $t_n(\mathfrak{A})$ in the natural fashion so that $\gamma \in S$ if $\pi(\gamma) \in S^{t(\mathfrak{A})}$.

From $t_n(\mathfrak{A})$ to $\hat{\mathfrak{A}}$. We now put $\hat{\mathfrak{A}} := \tau(t_n(\mathfrak{A}))$. Similar to the above representation of $\mathfrak{A} = \tau(t(\mathfrak{A}))$, we explicitly represent the node set of $\hat{\mathfrak{A}}$ by

$$\big\{(\gamma, S) \colon \gamma \in S^{t_n(\mathfrak{A})}\big\} \subseteq \Gamma^*_\alpha \times \mathcal{P}(\{P_1, \ldots, P_k\}).$$

In these terms, the accessibility relation of $\hat{\mathfrak{A}}$ is defined as $\big((\gamma, S), (\gamma', S')\big) \in R$ iff $\gamma \preccurlyeq \gamma'$ in $t_n(\mathfrak{A})$, and the atomic propositions are such that the propositional σ-type of (γ, S) is S.

One checks that the map

$$\begin{array}{rcl} \hat{\pi} \colon \hat{\mathfrak{A}} & \longrightarrow & \mathfrak{A} \\ (\gamma, S) & \longmapsto & (\pi(\gamma), S) \end{array}$$

induces a bisimulation between $\hat{\mathfrak{A}} = \tau(t_n(\mathfrak{A}))$ and $\mathfrak{A} = \tau(t(\mathfrak{A}))$. Note also that $t(\hat{\mathfrak{A}}) = t_n(\mathfrak{A})$. Clearly $t_n(\mathfrak{A})$ and $\hat{\mathfrak{A}}$ are finite for finite $\mathfrak{A}$, and path-finite for path-finite $\mathfrak{A}$. We next discuss the special properties of this bisimilar companion $\hat{\mathfrak{A}}$ that will be crucial for the game argument in the desired upgrading of some $\sim^\ell_*$ to $\equiv_q$. These properties play a role analogous to that encountered as ℓ-goodness in connection with Observation 4.17.

Special properties of $\hat{\mathfrak{A}}$. Let $\mathfrak{A}$ and $\hat{\mathfrak{A}}$ as above, both path-finite. Let $\hat{a} \in \hat{\mathfrak{A}}$ and consider, for some $m \geqslant 1$, the characteristic formula of its reduced $\sim^m_*$-type,

$$\chi := \chi^{\hat{a}}_{0;m}.$$

Let $p \in \mathcal{P}$ be some $\subseteq$-maximal element for which $\hat{a} \models \Diamond^*_p \chi$. Note that, since $\hat{a} \models \chi$, we have $\hat{a} \models \Diamond^*_{p(\hat{a})} \chi$ for reflexive $\hat{a}$.

Let $a := \hat{\pi}(\hat{a})$, and $[a] \in t(\mathfrak{A})$ the corresponding node of $t(\mathfrak{A})$.

Let $[a^+] \in t(\mathfrak{A})$ be any $\prec$-maximal reflexive node reachable from $[a]$ such that $p([a^+]) = p$ and $a^+ \models \chi$ (i.e., a^+ maximal in $\mathfrak{A}$ with respect to $R \setminus R^{-1}$ with $a^+ \models \Diamond^*_p \chi$).

Then $p^\infty(a^+) = p(a^+) = p$. The reduced type of a^+ implies χ and p was $\subseteq$-maximal for realisations of χ; therefore $p^\infty(a^+) \subseteq p$. On the other hand, $p^\infty(a^+) \supseteq p(a^+) = p$ since a^+ is reflexive.

In the passage from $t(\mathfrak{A})$ to $t_n(\mathfrak{A})$, the reflexive node $[a^+]$ gives rise to strict chains $\gamma_1 \prec \cdots \prec \gamma_n$ in $t_n(\mathfrak{A})$ of the form

$$\begin{array}{lcll} \gamma_1 & := & \gamma_0 \cdot ([a^+], n, p) \in \Gamma^*_\alpha, & \\ \gamma_{i+1} & := & \gamma_i \cdots ([a^+], n-i, p) \in \Gamma^*_\alpha, & \text{for } i < n. \end{array}$$

This chain consists of the E-classes of nodes that all realise the same reduced type as a^+; and γ_n is a $\prec$-maximal E-class realising $\Diamond^*_p\chi$.

Here $\gamma_0 \cdot ([a^+], n, p) \in \Gamma^*_\alpha$ can be any path from the root α in Γ to a node of the indicated form; similarly $\gamma_i \cdots ([a^+], n-i, p) \in \Gamma^*_\alpha$ is any continuation of the path γ_i to a node of the indicated form. There may be intermediate nodes γ' along this path between γ_i and γ_{i+1} with $p' \subsetneq p$ if $|p| > 1$, but none with $p' = p$. All of the nodes of $\hat{\mathfrak{A}}$ represented by any of the nodes within this chain share the same reduced type with a^+.

Definition 4.44. For $\hat{c} \in \hat{\mathfrak{A}}$ such that $\hat{c} \models \Diamond^*_p\chi$, define $d^\chi_p(\hat{c}) \in \{0, \ldots, n\}$ to be the maximal length of any strict R-path $\hat{c}, \hat{a}_1, \ldots, \hat{a}_d$ in $\hat{\mathfrak{A}}$ where $\hat{a}_i \models \Diamond^*_p\chi$ and $p(\hat{a}_i) = p$; and $d^\chi_p(\hat{c}) = n$ if there is such a path of length n.

Since $d^\chi_p(\hat{c})$ only depends on the reduced type of $\hat{c}$, the value is the same within $[\hat{c}]$, and we may regard d^χ_p as a function on $t_n(\mathfrak{A}) = t(\hat{\mathfrak{A}})$. The strict R-paths in the definition of d^χ_p translate into (or stem from) $\prec$-paths in $t_n(\mathfrak{A})$.

For the above $\prec$-path $\gamma_1 \prec \cdots \prec \gamma_n$ in $t_n(\mathfrak{A})$, we have $d^\chi_p(\gamma_i) = n - i$:

$d^\chi_p(\gamma_i) \geqslant n - i$ is obvious from the R-path $\gamma_{i+1}, \ldots, \gamma_n$;

$d^\chi_p(\gamma_i) \leqslant n - i$ follows from the construction of $\hat{\mathfrak{A}}$, since there are no realisations of $\Diamond^*_p\chi$ beyond $[a^+]$ or γ_n.

By varying the path beyond γ_i for some fixed $i < n$ we also find $\gamma' \succ \gamma_i$, representing elements $\hat{a}' \models \Diamond^*_p\chi$, with $d^\chi_p(\gamma') = n - i$ and $p(\gamma') = p'$ for any given $\emptyset \neq p' \subsetneq p$. Note that, since $t_n(\mathfrak{A})$ is an unfolding of Γ, the subtrees rooted at $\gamma \cdot ([a^+], n, p)$ and $\gamma' \cdot ([a^+], n, p)$ (paths ending in the same node $([a^+], n, p)$ of Γ) are isomorphic.

The key issue for our considerations will be which nodes of the form $\gamma \cdot ([a^+], k, p') \in t_n(\mathfrak{A})$ are reachable from a given node $[\hat{c}]$, from which the original $[\hat{a}]$ is reachable.[9] The answer is determined by $d^\chi_p(\hat{c})$ as follows.

Observation 4.45. *Let $\hat{a}$ be a reflexive node in $\hat{\mathfrak{A}} := \tau(t_n(\mathfrak{A}))$, $[\hat{a}]$ the corresponding clique in $t(\hat{\mathfrak{A}}) = t_n(\mathfrak{A})$. Let $m \geqslant 1$, $\chi := \chi^{\hat{a}}_{0;m}$, and p maximal such that $\hat{a} \models \Diamond^*_p\chi$. Let $\gamma_1 \prec \cdots \prec \gamma_n$ be chosen as described above, stemming from some maximal $[a^+] \in t(\mathfrak{A})$ with $a^+ \models \Diamond^*_p\chi$ that is reachable from $[\hat{\pi}(\hat{a})]$. Let $\hat{c} \in \hat{\mathfrak{A}}$ with corresponding $[\hat{c}] \in t_n(\mathfrak{A})$ such that $[\hat{a}]$, and thus also $\gamma_n = \gamma \cdot ([a^+], 1, p)$, are reachable from $[\hat{c}]$.*
In this situation:

(a) $[\hat{c}] \prec \gamma_1$ *iff* $d^\chi_p(\hat{c}) = n$.
If $d^\chi_p(\hat{c}) = n$, then $\hat{c}$ has reflexive R-successors $\hat{a}' \models \Diamond^*_p\chi$ for every combination – $p(\hat{a}') = p'$ and $d^\chi_p(\hat{a}') = d'$, for $\emptyset \neq p' \subseteq p$, $d' < n$.

(b) *Otherwise there is a unique maximal $i \geqslant 1$ such that $\gamma_i \preccurlyeq [\hat{c}]$. This value of i is characterised by $d^\chi_p(\hat{c}) = n - i$.*

[9] It will be the second player's task to match moves $\hat{a} \to \hat{c}$ over $\hat{\mathfrak{A}}$ with moves in $t_n(\mathfrak{A})$, which requires compatibility with p-values and, indirectly, with d^χ_p-values.

In this case $\hat{c}$ has reflexive R-successors $\hat{a}' \models \Diamond_p^* \chi$ such that
– $p(\hat{a}') := p'$ *and* $d_p^\chi(\hat{a}') = d_p^\chi(\hat{c})$, *for* $p(\hat{c}) \subseteq p' \subsetneq p$ *or* $p' = p$ *if* $p(\hat{c}) = p$,
– $p(\hat{a}') := p'$ *and* $d_p^\chi(\hat{a}') := d'$, *for* $\emptyset \neq p' \subseteq p$ *and* $d' < d_p^\chi(\hat{c})$.

(c) *The first case,* $d_p^\chi(\hat{c}) = n$, *must in particular apply*
  (i) *if $\hat{c}$ is irreflexive or if there is some irreflexive $\hat{a}' \models \Diamond_p^* \chi$ reachable from $\hat{c}$;*
  (ii) *if $\hat{c}$ has some successor $\hat{a}'$ with $\hat{a}' \models \Diamond_p^* \chi$, $d_p^\chi(\hat{a}') = d_p^\chi(\hat{c})$ and $p(\hat{a}') \subsetneq p(\hat{c})$;*
  (iii) *if $\hat{c}$ and $\hat{a}$ do not share the same reduced bisimulation type.*

For (a)–(c) note that (a) all nodes inserted in $t_n(\mathfrak{A})$ are reflexive and produce reflexive nodes in $\hat{\mathfrak{A}}$, (b) nodes inserted in $t_n(\mathfrak{A})$ are monotone increasing with respect to p-value within each one of the n blocks, and (c) nodes inserted in $t_n(\mathfrak{A})$ always reproduce the reduced type of the node they replace.

Towards the upgrading, the following lemma is the main goal.

Lemma 4.46. *Let $n := 2\ell$. For any path-finite (finite) rooted transitive structures*

$$\mathfrak{A}, \alpha \sim_*^{n+1} \mathfrak{B}, \beta$$

with irreflexive roots α and β, for which w.l.o.g. $\mathfrak{A} = \tau(t(\mathfrak{A}))$ and $\mathfrak{B} = \tau(t(\mathfrak{B}))$, let $\hat{\mathfrak{A}} := \tau(t_n(\mathfrak{A}))$ and $\hat{\mathfrak{B}} := \tau(t_n(\mathfrak{B}))$. Then $\hat{\mathfrak{A}}, \alpha \sim \mathfrak{A}, \alpha$ and $\hat{\mathfrak{B}}, \beta \sim \mathfrak{B}, \beta$ are path-finite (finite) rooted transitive structures such that

$$t_n(\mathfrak{A}), \alpha = t(\hat{\mathfrak{A}}), \alpha \;\sim_*^\ell\; t(\hat{\mathfrak{B}}), \beta = t_n(\mathfrak{B}).$$

The game argument. We extract a strategy for $t(\hat{\mathfrak{A}}), \alpha \sim_*^\ell t(\hat{\mathfrak{B}}), \beta$ from the strategy for $\hat{\mathfrak{A}}, \alpha \sim_*^{2\ell+1} \hat{\mathfrak{B}}, \beta$, essentially by lifting the game on the equivalence classes $[\hat{a}] \in t(\hat{\mathfrak{A}}) = t_{2\ell}(\mathfrak{A})$ and $[\hat{b}] \in t(\hat{\mathfrak{B}}) = t_{2\ell}(\mathfrak{B})$ to suitable representatives $\hat{a} \in \hat{\mathfrak{A}}$ and $\hat{b} \in \hat{\mathfrak{B}}$.

As in other cases before, certain critical nodes and associated depths will be crucial. There is considerable similarity with the game argument in the upgrading from irreflexive transitive trees to transitive tree-like structures in section 4.5. Some things are simpler here because we need not worry about reflexive versus irreflexive nodes; instead, the complication comes from the necessity to find responses of matching p-values, since these values determine the propositional type of $[\hat{a}]$ in $t(\hat{\mathfrak{A}})$.

Essentially we want to simulate a $\Diamond$- or $\Diamond^*$-move in the game $(t(\hat{\mathfrak{A}}); t(\hat{\mathfrak{B}}))$ from $[\hat{a}]$ to $[\hat{a}']$ with the help of a $\Diamond_p^*$-move in the game on $(\hat{\mathfrak{A}}; \hat{\mathfrak{B}})$ from $\hat{a}$ to $\hat{a}'$, for $p = p(\hat{a}')$; the matching response in $\hat{\mathfrak{B}}$, however, may be any $\hat{b}'$ of the right $\sim_*^k$-type, but only $p(\hat{b}') \supseteq p(\hat{a}')$ rather than equality is guaranteed.

In other words, $\Diamond_p^*$ assertions in ML^* only put p as a lower bound on the actual p-value we find. It is precisely for this purpose that the $n = 2\ell$-fold downward saturation with respect to all possible under-approximations was built into $t_n(\mathfrak{A})$ and $t_n(\mathfrak{B})$.

Critical nodes and their depths. Call $\hat{a} \in \hat{\mathfrak{A}}$ and the corresponding $[\hat{a}] \in t(\hat{\mathfrak{A}})$ *m-critical* if

$$\hat{\mathfrak{A}}, \hat{a} \models \Box^* \big(\chi_{0;m-1}^{\hat{a}} \rightarrow \chi_{0;m}^{\hat{a}}\big).$$

It is clear from the definition that the status of $\hat{a}$ with respect to this notion depends only on its reduced $\sim_*^{m+1}$-type, hence is the same throughout $[\hat{a}]$ and is inherited from the corresponding elements in the original $\mathfrak{A}$.

As a finite analogue, at level m, of the maximal p-value $p^\infty(a)$ among all reachable cliques containing a realisation of the reduced type of a, we now define, for reduced $\sim^m_*$-types, the following:

$$P^m(a) := \{p \colon p \subseteq\text{-maximal with } \mathfrak{A}, a \models \Diamond^*_p \chi^a_{0;m}\}.$$

$P^m(a)$ may be empty for irreflexive a. Note that $P^m(a)$ is determined by the reduced $\sim^{m+1}_*$-type of a and only depends on $[a]$. Unlike the situation for p^∞, $P^m(a)$ may in general have more than one element; but not so for m-critical a.

Observation 4.47. *If a is m-critical in a path-finite transitive structure $\mathfrak{A}$, then $P^m(a)$ has at most one element.*

If $P^m(a) \neq \emptyset$, we denote the unique element $p^m(a)$.

Proof. The reason for uniqueness is similar to the one given for p^∞ in Observation 4.43. Assume that $p_1 \neq p_2$ were both maximal clique-types for $\chi^a_{0;m}$. As $p_i \in P^m(a)$, $\mathfrak{A}, a \models \Diamond^*_{p_i} \chi^a_{0;m}$ for $i = 1, 2$, and $p := p_1 \cup p_2 \notin P^m(a)$ implies that $\mathfrak{A}, a \models \neg\Diamond^*_p \chi^a_{0;m}$. Since a is m-critical, these assertions are equivalent to $\mathfrak{A}, a \models \Diamond^*_{p_i} \chi^a_{0;m-1}$ for $i = 1, 2$, and $\mathfrak{A}, a \models \neg\Diamond^*_p \chi^a_{0;m-1}$, and as such they are themselves implied by $\chi^a_{0;m}$. Now any two R-cliques witnessing the positive assertions $\Diamond^*_{p_i} \chi^a_{0;m-1}$ at a must be disjoint from each other, whence at least one is disjoint from a. Iteration of this argument shows that $\mathfrak{A}$ has an infinite path of disjoint R-cliques, contradicting path-finiteness. ◻

For m-critical $\hat{a}$ in $\hat{\mathfrak{A}}$ and the associated $[\hat{a}] \in t(\hat{\mathfrak{A}})$ we now define m-depth as follows.

If $P^m(\hat{a}) = \emptyset$ (which implies that $\hat{a}$ and $[\hat{a}]$ are irreflexive), put $d^m(\hat{a}) := d^m([\hat{a}]) := 0$.

In all other cases, $p^m(\hat{a})$ is defined and we let $d^m(\hat{a}) = d^m([\hat{a}])$ be the maximal length of a strict R-path from $\hat{a}$ consisting of R-cliques of type $p^m(\hat{a})$ in which $\chi^a_{0;m}$ (or, equivalently, $\chi^{\hat{a}}_{0;m-1}$) is realised; and $d^m(\hat{a}) := n$ if at least this length can be had. This is just d^χ_p as defined in Definition 4.44 for $\chi = \chi^{\hat{a}}_{0;m}$ and $p = p^m(\hat{a})$.

The condition to be maintained through the game is expressed in terms of constraints on positions $(\hat{a}, \hat{b}) \in \hat{\mathfrak{A}} \times \hat{\mathfrak{B}}$. With k rounds still to be played, we require the pair $(\hat{a}, \hat{b})$ be *m-safe* for some $m \geqslant 2k$ in the following sense:

$$m\text{-safe:} \quad \begin{cases} \hat{a} \sim^m_* \hat{b} \\ p(\hat{a}) = p(\hat{b}) \\ \hat{a}\ m\text{-critical} \;\Leftrightarrow\; \hat{b}\ m\text{-critical} \\ \text{and for } m\text{-critical } \hat{a}, \hat{b}\colon\ d^m(\hat{a}) = d^m(\hat{b}) \text{ or } d^m(\hat{a}), d^m(\hat{b}) \geqslant m. \end{cases}$$

Note that, although the conditions are expressed in terms of $\hat{a}, \hat{b}$, all but the first clause are manifestly dependent only on $[\hat{a}]$ and $[\hat{b}]$. This first clause, $\hat{a} \sim^m_* \hat{b}$, should be seen as a combination of the requirements $\hat{a} \sim^{0;m}_* \hat{b}$ and $S(\hat{a}) = S(\hat{b})$. The reduced type match again only depends on $[\hat{a}]$ and $[\hat{b}]$, and the additional equality of propositional σ-types can be achieved by a matching choice of representatives from $[\hat{a}]$ or $[\hat{b}]$, since $p(\hat{a}) = p(\hat{b})$ is also required.

Base case. At the start of the game, with ℓ rounds to be played from the initial configuration (α, β): $\mathfrak{A}, \alpha \sim_*^{2\ell+1} \mathfrak{B}, \beta$ implies $\hat{\mathfrak{A}}, \alpha \sim_*^{2\ell+1} \hat{\mathfrak{B}}, \beta$; put $m := 2\ell$ to have $\alpha \sim_*^m \beta$ and α m-critical iff β is m-critical. Since α and β are irreflexive, $p(\alpha) = \{S(\alpha)\} = \{S(\beta)\} = p(\beta)$ follows from just $\sim_*^0$ equivalence. By $\sim_*^{m+1}$ equivalence moreover $P^m(\alpha) = P^m(\beta) =: P^m$. In case α and β are m-critical, we distinguish two cases.

(1) $P^m = \emptyset$ and hence $d^m(\alpha) = d^m(\beta) = 0$.

(2) $P^m = \{p^m\}$ and we claim that then $d^m(\alpha), d^m(\alpha) \geqslant 2\ell$ by construction. This follows from Observation 4.45 (c)(i), since α and β are irreflexive.

One round in the game. Consider a move played on the A-side in the game on $(t(\hat{\mathfrak{A}}); t(\hat{\mathfrak{B}}))$ from position $([\hat{a}], [\hat{b}])$. Suppose there are k rounds still to be played and that the pair $(\hat{a}, \hat{b})$ is m-safe for some $m \geqslant 2k$. Let player I make a move from $[\hat{a}]$ to $[\hat{a}']$. We need to find b' such that $(\hat{a}', \hat{b}')$ is m'-safe for some $m' \geqslant m - 2$.

Whenever $[\hat{a}']$ is reflexive, we may assume w.l.o.g. that the move is played by I as a $\diamond^*$-move; if on the other hand $[\hat{a}']$ is irreflexive, only a plain $\diamond$-move can be responsible.

Case 1: a $\diamond^*$-move to $[\hat{a}']$, $\hat{a}' \not\sim_*^{0;m-1} \hat{a}$. By a matching $\diamond_p^*$-move where $p := p(\hat{a}')$, we find a reflexive $\hat{b}_0' \sim_*^{m-1} \hat{a}'$ with $p(\hat{b}_0') \supseteq p(\hat{a}')$. For any such $\hat{b}_0' \sim_*^{m-1} \hat{a}'$:

– $\hat{b}_0' \sim_*^{m-2} \hat{a}'$, trivially;
– $\hat{b}_0'$ is $(m-2)$-critical iff $\hat{a}'$ is;
– $\hat{b}_0'$ and $\hat{b}$ do not share the same reduced type, since even $\hat{b}_0' \not\sim_*^{0;m-1} b$.

Let $\hat{p} \supseteq p(\hat{b}_0')$ be maximal such that $\hat{b}_0' \models \diamond_{\hat{p}}^* \chi_{0;m-1}^{\hat{a}'}$. Since $\hat{b}_0' \not\sim^{0;\infty} \hat{b}$, Observation 4.45 (c)(iii) implies that $d_{\hat{p}}^{\chi}(\hat{b}) = n$ for $\chi := \chi_{0;m-1}^{\hat{a}'}$. We thus find $\hat{b}'$ reachable from $\hat{b}$ and realising $p \subseteq \hat{p}$ and $\chi_{0;m-1}^{\hat{a}'}$, as desired.

If $\hat{a}'$ is $(m-2)$-critical, we work with $\chi_{0;m-2}^{\hat{a}'}$ and $p^{m-2}(\hat{a}')$ instead of $\hat{p}$, and find, by Observation 4.45, a matching $\hat{b}'$ of the right value for both $p(\hat{b}') = p(\hat{a}')$ and $d^{m-2}(\hat{b}') = d^{m-2}(\hat{a}')$ if $d^{m-2}(\hat{a}') < n$; or $d^{m-2}(\hat{b}') = n - 1$ if $d^{m-2}(\hat{a}') = n$.

By the match in p-values, we find a representative $\hat{b}'' \in [\hat{b}']$ such that $\hat{b}'' \sim_*^{m-2} \hat{a}'$.

Case 2: a $\diamond^*$-move to $[\hat{a}']$, $\hat{a}' \sim_*^{0;m-1} \hat{a}$, but $\hat{a}$ and $\hat{b}$ not m-critical. Similar to the above, consider a matching $\diamond_p^*$-move from $\hat{a}$ to $\hat{a}'$ with response $\hat{b}_0'$. As $\hat{b}$ is not m-critical, there is $\hat{b}_0'' \sim_*^{0;m-1} \hat{b}_0' \sim_*^{0;m-1} \hat{a}'$ such that $\hat{b}_0'' \not\sim_*^{0;m} \hat{b}$ and therefore $\hat{b}_0'' \not\sim^{0;\infty} \hat{b}$. From this $\hat{b}_0''$ we may obtain a suitable $\hat{b}'$ as in case 1.

Case 3: a $\diamond^*$-move to $[\hat{a}']$, $\hat{a}' \sim_*^{0;m-1} \hat{a}$, with $\hat{a}$ and $\hat{b}$ both m-critical. By m-safety, $d^m(\hat{a}) = d^m(\hat{b})$. Because $\hat{a}$ and $\hat{b}$ are m-critical, we have that $\hat{a}' \sim_*^{0;m} \hat{a}$; $p^m(\hat{a}) = p^{m-1}(\hat{a}) = p^{m-1}(\hat{b}) = p^m(\hat{b})$ well defined; since $\hat{a}' \sim_*^{0;m} a \sim_*^{0;m} \hat{b}$, $\hat{a}'$ is $(m-1)$-critical iff $\hat{a}$ and $\hat{b}$ are, and in this case $d^{m-1}(\hat{a}) = d^m(\hat{a}) = d^m(\hat{b}) = d^{m-1}(\hat{b})$.

Case 3.1: $d^m(\hat{a}) = d^m(\hat{b}) = 0$. Since $p^m(\hat{a}) = p(\hat{a})$ and $[\hat{a}]$ admits a $\diamond^*$-move to $[\hat{a}']$ where $\hat{a}' \sim_*^{0;m} \hat{a}$, it follows that $[\hat{a}'] = [\hat{a}]$. This can be matched by $[\hat{b}'] = [\hat{b}]$, and we also find a representative $\hat{b}' \in [\hat{b}]$ for which $S(\hat{b}') = S(\hat{a}')$. If $\hat{a}'$ is $(m-1)$-critical, so are $\hat{a}, \hat{b}$ and $\hat{b}'$, and in this case $d^{m-1} = d^m$ guarantees the required match for these.

Case 3.2: $d^m(\hat{a}) = d^m(\hat{b}) > 0$ and $d^m(\hat{a}') = d^m(\hat{a}) < m$. Since $d^m(\hat{a}) < n$, Observation 4.45 (c)(ii) implies that $p(\hat{a}) \subseteq p(\hat{a}') \subseteq p^m(\hat{a})$, and we find a matching $\hat{b}' \sim_*^{0;m-1} \hat{a}'$

with $p(\hat{b}') = p(\hat{a}')$ and $d^m(\hat{b}') = d^m(\hat{b}) = d^m(\hat{a}) = d^m(\hat{a}')$. Again, $\hat{b}'$ is $(m-1)$-critical iff $\hat{a}'$ is; and as in this case $d^{m-1} = d^m$ these depths match too.

Case 3.3: $d^m(\hat{a}) = d^m(\hat{b}) > 0$ and $d^m(\hat{a}') = d^m(\hat{a}) \geqslant m$. Similar to the last case, only that $p(\hat{a}') \subsetneq p(\hat{a})$ may now occur. In this case one finds a suitable $\hat{b}'$ not necessarily of the same d^m, but with $d^m(\hat{b}') = d^m(\hat{a}') - 1$, which is still good enough to ensure $d^{m-1}(\hat{b}'), d^{m-1}(\hat{a}') \geqslant m-1$ in case $\hat{a}'$ and $\hat{b}'$ are $(m-1)$-critical.

Case 3.4: $d^m(\hat{a}) = d^m(\hat{b}) > 0$ and $d^m(\hat{a}') < d^m(\hat{a})$. Similar to the above, find $\hat{b}'$ in the appropriate interval along a path of length $d^m(\hat{b}) = d^m(\hat{a})$ of R-cliques as in Observation 4.45.

Case 4: a $\Diamond$-move to irreflexive $[\hat{a}']$. Consider the corresponding $\Diamond$-move from $\hat{a}$ to $\hat{a}'$ in $\hat{\mathfrak{A}}$. This will directly produce some response $\hat{b}'_0$ such that $\hat{b}'_0 \sim_*^{m-1} \hat{a}'$, whence $\hat{b}'_0$ is $(m-2)$-critical iff $\hat{a}'$ is, and such that $S(\hat{b}'_0) = S(\hat{a}')$. So only $p(\hat{a}') = \{S(\hat{a}')\} \subseteq p(\hat{b}'_0)$ is immediate. Moreover, in the case of $(m-2)$-critical $\hat{a}'$, $d^{m-2}(\hat{b}')$ needs to be adjusted.

Case 4.1: $\hat{a} \models \neg\Diamond^*\chi^{\hat{a}'}_{m-1}$. In this case $p(\hat{a}') = \{S(\hat{a})\} = p(\hat{b}'_0)$, since any $p(\hat{b}'_0)$ must also be a singleton set.

If $\hat{a}'$ is $(m-2)$-critical, then so is $\hat{b}' := \hat{b}'_0$. In this case, we let $\chi := \chi^{\hat{a}'}_{0;m-2}$ and distinguish cases as to whether $\hat{a}' \models \Diamond^*\chi$ or not; by $\sim_*^{m-1}$ equivalence, $\hat{b}' \models \Diamond^*\chi$ $\hat{a}' \models \Diamond^*\chi$:

– if $\hat{a}' \models \neg\Diamond^*\chi$, then $d^{m-2}(\hat{a}') = 0$ and the same is true of $\hat{b}'$.

– if $\hat{a}' \models \Diamond^*\chi$, then by Observation 4.45 (c)(i) $d^{m-2}(\hat{a}') = n$ as $\hat{a}'$ is irreflexive; again, the same is true of $\hat{b}'$, since $\hat{b}' \sim_*^{0;m-1} \hat{a}'$ and as $\hat{b}'$ is $(m-2)$-critical, too.

Case 4.2: $\hat{a} \models \Diamond^*\chi^{\hat{a}'}_{m-1}$. If $\hat{a}' \not\sim_*^{0;m-1} \hat{a}$ or if $\hat{a}' \sim_*^{0;m-1} \hat{a}$ but $\hat{a}$ and $\hat{b}$ not m-critical, then we find a reflexive response (as if to a $\Diamond^*$-move to some reflexive variant of $\hat{a}'$; see cases 1 and 2).

If $\hat{a}$ and $\hat{b}$ are m-critical and $\hat{a}' \sim_*^{0;m} \hat{a}$, then Observation 4.45 (c)(i) implies that $d^{m-1}(\hat{a}') = d^m(\hat{a}) = n$, because $\hat{a}'$ is irreflexive. We then proceed in analogy with case 3.3. above.

This completes the proof of Lemma 4.46.

For future reference we state the upgrading assertion obtained with the construction in Figure 6 on the basis of Lemma 4.46 as a corollary. This in turn proves Theorems 4.40 and 4.41.

Corollary 4.48. *For every finite vocabulary σ and all q there is an ℓ such that any two finite/path-finite pointed transitive σ-structures $\mathfrak{A}, \alpha \sim_*^{\ell} \mathfrak{B}, \beta$ possess bisimilar companions $\hat{\mathfrak{A}}, \hat{\alpha} \sim \mathfrak{A}, \alpha$ and $\hat{\mathfrak{B}}, \hat{\beta} \sim \mathfrak{B}, \beta$ within the class of finite/path-finite rooted transitive structures such that $\hat{\mathfrak{A}}, \hat{\alpha} \equiv_q \hat{\mathfrak{B}}, \hat{\beta}$.*

4.7 Beyond FO in transitive frames

We address the extension of several of the previous results of this section to MSO in place of FO. Indeed, MSO rather than FO is the natural focus of decomposition based methods. While some of the characterisations obtained for not necessarily finite frames do not generalise to corresponding characterisations of the bisimulation invariant fragment of MSO, surprisingly all those that we formulated for various classes of *finite transitive frames* or *path-finite transitive frames* turn out to be characterisations of MSO/$\sim$. This also implies the collapse of MSO to FO for bisimulation invariant properties over various classes of transitive frames. For instance, MSO/$\sim$ $\equiv$ FO/$\sim$ $\equiv$ ML over the class of all finite irreflexive transitive trees ($\prec$-trees) – a result, which has also been obtained by ten Cate, Fontaine and Litak [16], via a clever combination of the Janin–Walukiewicz characterisation of bisimulation invariant MSO, [10] (also valid in restriction to the class all finite trees in the graph theoretic sense), and the de Jongh–Sambin theorem (see e.g. [15] for a model theoretic discussion and proof), which implies that the μ-calculus L_μ collapses to basic modal logic over $\prec$-trees without infinite paths.

Theorem 4.49. MSO/$\sim$ $\equiv$ FO/$\sim$ $\equiv$ ML *over the class of all finite irreflexive transitive frames; and similarly over the class of all transitive frames without infinite paths (Löb frames).*

We here find an alternative proof of this theorem based on a natural generalisation of the composition and upgrading arguments that support Claim 4.10. We shall then see how the further upgrading arguments developed in the previous sections allow us to lift this characterisation also to characterisations of MSO/$\sim$ over finite reflexive transitive trees, finite transitive tree-like frames and finite transitive trees.

The starting point is the following analogue of Claim 4.10 for finite $\prec$-trees and, more generally, for $\prec$-trees without infinite paths (cf. Löb frames). We fix a finite vocabulary σ, consisting of finitely many unary predicates beside the accessibility relation constituting the frame. As a default, $\prec$-tree structures, $\preccurlyeq$-tree structures, etc. now always refer to expansions of corresponding frames to σ-structures for this fixed finite σ. The bounds on nesting depths expressed in the following all depend on σ, but we do not make this dependency explicit.

Lemma 4.50. *There is a function L such that for all rooted $\prec$-tree structures $\mathfrak{A},\alpha$ and $\mathfrak{B},\beta$ without infinite paths and for all q the following holds for all sufficiently large Q:*

$$\mathfrak{A},\alpha \sim^{L(q)} \mathfrak{B},\beta \quad\Rightarrow\quad s_Q(\mathfrak{A}) \equiv_q^{\mathrm{MSO}} s_Q(\mathfrak{B}).$$

Note that passing from $\prec$-trees to their Q-saturated companions is admissible both within the class of finite $\prec$-trees and within the class of $\prec$-trees without infinite paths.

Proof. Let $\Theta(q)$ be the finite set of all complete MSO_q-types of rooted $\prec$-tree structures of vocabulary σ, $N = N(q) := |\Theta(q)|$ the number of such rooted types. For $a \in \mathfrak{A}$, $\mathfrak{A}$ a $\prec$-tree structure, let $\mathfrak{A}_a$ stand for the $\prec$-tree structure induced on the subtree rooted at a in $\mathfrak{A}$ and let $\angle(a) := \{\mathrm{tp}_q^{\mathrm{MSO}}(\mathfrak{A}_b, b)\colon a \prec b\} \subseteq \Theta(q)$ be the set of those rooted MSO_q-types realised in $\mathfrak{A}_a \setminus \{a\}$.

It follows from the usual MSO composition arguments, which are easily proved on the basis of the Ehrenfeucht–Fraïssé game for MSO, that $\mathrm{tp}_q^{\mathrm{MSO}}(\mathfrak{A}_a, a)$ is fully determined by the atomic type $\mathrm{atp}(a) := \{P \in \sigma\colon a \in P^{\mathfrak{A}}\}$ of a in $\mathfrak{A}$ and the tuple of multiplicities of rooted MSO_q-types of subtrees rooted in the immediate successors of a,

$$\Big(\big|\{b\colon b \text{ an immediate successor of } a, \mathrm{tp}_q^{\mathrm{MSO}}(\mathfrak{A}_b, b) = \theta\}\big|\Big)_{\theta\in\Theta(q)}.$$

Let Q be sufficiently large (in relation to σ and q) such that exact multiplicities of rooted MSO_q-types do not matter beyond Q.

For the Q-saturated companion $s_Q(\mathfrak{A})$ in place of $\mathfrak{A}$ itself, we then find that types $\mathrm{tp}_q^{\mathrm{MSO}}(s_Q(\mathfrak{A})_a, a)$ are determined by the atomic type $\mathrm{atp}(a)$ and the *set*

$$\begin{array}{rl} & \{\mathrm{tp}_q^{\mathrm{MSO}}(s_Q(\mathfrak{A})_b, b)\colon b \text{ an immediate successor of } a\} \\ = & \{\mathrm{tp}_q^{\mathrm{MSO}}(s_Q(\mathfrak{A})_b, b)\colon a \prec b\} = \angle(a), \end{array}$$

where $\angle(a)$ now refers to the set of rooted types in $s_Q(\mathfrak{A})$.

This is because multiplicities do not matter beyond threshold Q, and because, due to unravelling in the passage from $\mathfrak{A}$ to $s_Q(\mathfrak{A})$, every rooted MSO_q-type in $\angle(a)$ is represented at immediate successors of a.

Claim 4.51. *For any $s \subseteq \Theta(q)$ and $\theta \in \Theta(q)$ there are* ML*-formulae ξ_s and $\xi_{\theta/s}$ of nesting depth $|s|+1$ such that*

$$\begin{array}{rrcl} & s_Q(\mathfrak{A}), a \models \xi_s & \Leftrightarrow & \angle(a) = s \\ \textit{and} & s_Q(\mathfrak{A}), a \models \xi_{\theta/s} & \Leftrightarrow & \angle(a) = s \ \textit{ and } \ \mathrm{tp}_q^{\mathrm{MSO}}(a) = \theta. \end{array}$$

Formulae ξ_s and $\xi_{\theta/s}$ are obtained by induction with respect to $|s|$.

For $s = \emptyset$, $\xi_\emptyset := \Box\bot$ is as required, and $\xi_{\theta/s}$ can be chosen as $\xi_{\theta/s} = \xi_s \wedge \bigvee_i \chi_i$ where the χ_i are a collection of formulae defining those atomic types which, taken as the type of the single node in a trivial tree frame, satisfy θ.

Generally, we obtain $\xi_{\theta/s}$ from ξ_s as $\xi_{\theta/s} = \xi_s \wedge \bigvee_i \chi_i$ where the χ_i describe those atomic types $\mathrm{atp}(a)$ which, for a with $\angle(a) = s$, determine $\mathrm{tp}_q^{\mathrm{MSO}}(s_Q(\mathfrak{A})_a, a)$ to be θ.

For $|s| > 0$ inductively put

$$\begin{array}{rcl} \xi_s & := & \Box\Big(\bigwedge_{s' \subsetneq s}(\xi_{s'} \rightarrow \bigvee_{\theta \in s} \xi_{\theta/s'}) \ \wedge \ \big((\bigwedge_{s' \subsetneq s} \neg\xi_{s'}) \rightarrow \bigvee \chi_i\big)\Big) \\ & & \wedge \bigwedge_{\theta \in s} \bigvee_{s' \subsetneq s} \Diamond \xi_{\theta/s'}, \end{array}$$

where $\bigvee_i \chi_i$ collects those atomic types which, at nodes b with $\angle(b) = s$ determine some type in s.

Clearly any a with $\angle(a) = s$ satisfies ξ_s. This is obvious for the $\Box$-part. For the other conjunct look, for a given $\theta \in s$, at some $\prec$-maximal realisation b of θ in $s_Q(\mathfrak{A})_a \setminus \{a\}$. By $\prec$-maximality, $\angle(b) \subseteq s \setminus \{\theta\} \subsetneq s$, and b satisfies a corresponding $\xi_{\theta/s'}$.

Conversely, let a satisfy ξ_s. The second part of ξ_s enforces $s \subseteq \angle(a)$. Supposing that $\angle(a) \not\subseteq s$, consider a $\prec$-maximal node b in $s_Q(\mathfrak{A})_a \setminus \{a\}$ whose type is not a member of s. Then $\angle(b) \subseteq s$ by $\prec$-maximality of b. If $\angle(b) = s$, then the last conjunct inside the $\Box$-part forces $\mathrm{tp}_q^{\mathrm{MSO}}(s_Q(\mathfrak{A})_b, b) \in s$ by the choice of $\bigvee_i \chi_i$; if $\angle(b) = s' \subsetneq s$, then the corresponding conjunct in the $\Box$-part guarantees $\mathrm{tp}_q^{\mathrm{MSO}}(s_Q(\mathfrak{A})_b, b) \in s$.

From the formulae $\xi_{\theta/s}$ we obtain ML-formulae ξ_θ of nesting depth $L(q) := N(q)+1$ defining each rooted MSO_q-type θ. For all $\prec$-trees $\mathfrak{A}, \alpha$ without infinite paths, and for Q as above:

$$s_Q(\mathfrak{A}), \alpha \models \bigvee_s \xi_{\theta/s} \quad \Leftrightarrow \quad \mathrm{tp}_q^{\mathrm{MSO}}(s_Q(\mathfrak{A}), \alpha) = \theta.$$

This proves the lemma for $L(q) = N(q)+1$. $\Box$

Lemma 4.50 immediately proves Theorem 4.49, along exactly the same lines as before. Using the lemma, we may upgrade $\sim^{L(q)}$ to $\equiv_q^{\mathrm{MSO}}$ within the class of all finite rooted $\prec$-tree structures as well as within the class of not necessarily finite rooted $\prec$-tree structures without infinite paths. Therefore, any MSO formula φ of quantifier rank q that is invariant under bisimulation over one of these classes is in fact invariant under $\sim^{L(q)}$ over that class and hence equivalent to an ML formula of nesting depth $L(q)$.

We next translate this characterisation from irreflexive to reflexive transitive trees and Grzegorczyk frames, by application of the same translation through interpretation as in section 4.3; then to transitive tree-like and finally to transitive frames, just as we did for FO/$\sim$ in sections 4.5 and 4.6. Indeed the upgrading arguments remain the same, only that the core upgrading from $\sim^\ell$ to $\equiv_q$ (Claim 4.10) gets replaced by an upgrading to $\equiv_q^{\mathrm{MSO}}$ as provided by Lemma 4.50.

Reflexive transitive trees. Considering the natural translation from $\preccurlyeq$-trees to $\prec$-trees that maps $\mathfrak{A}$ to $\mathfrak{A}_\prec$ as discussed in section 4.3, the combination of Corollary 4.19 or 4.22 with Lemma 4.50 yields the following (cf. Figure 3). Due to the quantifier-free first-order nature of the translation (and its inverse), we clearly have

$$\mathfrak{A}, \alpha \equiv_q^{\mathrm{MSO}} \mathfrak{B}, \beta \quad \Leftrightarrow \quad \mathfrak{A}_\prec, \alpha \equiv_q^{\mathrm{MSO}} \mathfrak{B}_\prec, \beta$$

for any $\preccurlyeq$-tree structures $\mathfrak{A}, \alpha$ and $\mathfrak{B}, \beta$.

Corollary 4.52. *Any finite rooted $\preccurlyeq$-tree structures $\mathfrak{A}, \alpha \sim^{f_0(L(q))} \mathfrak{B}, \beta$, have bisimilar companions $\tilde{\mathfrak{A}}, \tilde{\alpha} \sim \mathfrak{A}, \alpha$ and $\tilde{\mathfrak{B}}, \tilde{\beta} \sim \mathfrak{B}, \beta$, such that $\tilde{\mathfrak{A}}, \tilde{\alpha} \equiv_q^{\mathrm{MSO}} \tilde{\mathfrak{B}}, \tilde{\beta}$. Similarly for not necessarily finite $\preccurlyeq$-trees without infinite $\prec$-paths.*

Theorem 4.53. MSO/$\sim$ $\equiv$ FO/$\sim$ $\equiv$ ML *over the class of all finite reflexive transitive frames as well as over the class of all not necessarily finite reflexive transitive frames without infinite paths with respect to the irreflexive part of R (Grzegorczyk frames) and the corresponding class of $\preccurlyeq$-trees.*

Transitive tree-like structures. The variation for transitive tree-like frames is similarly obtained, but based on the translation that associates to a (finite) transitive tree-like structure $\mathfrak{A}, \alpha$ the (finite) irreflexive transitive tree structure $\mathfrak{A}^\circ_\prec, \alpha$, in which $R^{\mathfrak{A}}$ is replaced by its irreflexive part and a new unary predicate P° marks the reflexive nodes.

Again, the translation and its inverse are governed by quantifier free FO interpretations, whence it is compatible with levels of MSO equivalence:

$$\mathfrak{A}, \alpha \equiv_q^{\mathrm{MSO}} \mathfrak{B}, \beta \quad \Leftrightarrow \quad \mathfrak{A}^\circ_\prec, \alpha \equiv_q^{\mathrm{MSO}} \mathfrak{B}^\circ_\prec, \beta.$$

Combining Corollaries 4.33 and 4.35 with Lemma 4.50, we thus get the following (cf. Figure 5).

Corollary 4.54. *Any finite rooted transitive tree-like structures $\mathfrak{A}, \alpha \sim_*^{f_1(L(q))} \mathfrak{B}, \beta$, have bisimilar companions $\tilde{\mathfrak{A}}, \tilde{\alpha} \sim \mathfrak{A}, \alpha$ and $\tilde{\mathfrak{B}}, \tilde{\beta} \sim \mathfrak{B}, \beta$, such that $\tilde{\mathfrak{A}}, \tilde{\alpha} \equiv_q^{\mathrm{MSO}} \tilde{\mathfrak{B}}, \tilde{\beta}$. Similarly for not necessarily finite path-finite transitive tree-like frames.*

Theorem 4.55. MSO/$\sim$ $\equiv$ FO/$\sim$ $\equiv$ ML* *over the class of all finite transitive tree-like structures.*

Weak transitive tree-like and transitive structures. Similarly Corollary 4.48 can be boosted by the use of Lemma 4.50 instead of Claim 4.10 to yield the following. It is important here that the transformation τ preserves $\equiv_q^{\mathrm{MSO}}$ as well as $\equiv_q$ (cf. Observation 4.39 (i)).

Corollary 4.56. *For any finite rooted weak transitive tree-like structures* $\mathfrak{A}, \alpha \sim_*^{f_2(L(q))} \mathfrak{B}, \beta$ *have bisimilar companions* $\hat{\mathfrak{A}}, \hat{\alpha} \sim \mathfrak{A}, \alpha$ *and* $\hat{\mathfrak{B}}, \hat{\beta} \sim \mathfrak{B}, \beta$ *within the class of finite rooted weak transitive tree-like structures such that* $\hat{\mathfrak{A}}, \hat{\alpha} \equiv_q^{\mathrm{MSO}} \hat{\mathfrak{B}}, \hat{\beta}$. *Similarly for path-finite weak transitive tree-like structures.*

We summarise the characterisation and collapse results of this section.

Theorem 4.57. $\mathrm{MSO}/{\sim} \equiv \mathrm{FO}/{\sim} \equiv \mathrm{ML}^*$ *over*

(1) *the class of all finite rooted weak transitive tree-like structures,*
(2) *the class of all path-finite weak transitive tree-like structures,*
(3) *the class of all finite rooted transitive structures,*
(4) *the class of all rooted path-finite transitive structures,*
(5) *the class of all finite transitive structures,*
(6) *the class of all path-finite transitive structures.*

We therefore also have the following ramifications and extensions of the de Jongh–Sambin theorem over several wider classes of structures than Löb frames, and with ML^* definability of fixed points instead of ML-definability in the key cases. Note that ML^* uniformly translates in fact into the alternation free fragment of L_μ over path-finite transitive frames, where $\Diamond_p^*\varphi \equiv \nu_X\big(\Diamond(\varphi \wedge X) \wedge \bigwedge_{\zeta \in p} \Diamond(\zeta \wedge X)\big)$.

Corollary 4.58. $\mathrm{L}_\mu \equiv \mathrm{ML}^*$ *(uniformly) over all path-finite transitive frames.* $\mathrm{L}_\mu \equiv \mathrm{ML}$ *(uniformly) over all reflexive transitive frames without infinite paths with respect to the irreflexive part of* R *(Grzegorczyk frames), just as over Löb frames and over all transitive frames without infinite paths.*

We remark that results by Alberucci and Facchini [1] give a corresponding collapse to the alternation free fragment even over the class of *all* transitive frames. The collapse of L_μ not just to the alternation free fragment but to FO over *finite* transitive frames is also covered by recent work of D'Agostino and Lenzi [4].

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 focused 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. The last section in particular shows a rather interesting new phenomenon regarding the behaviour of basic modal logic and bisimulation invariance over the natural class of transitive frames. Comparing Theorem 2.12 with Theorem 4.40, we see that the restriction to finite structures necessitates a new modality. Without the new modality $\Diamond^*$, ML is expressively complete for bisimulation invariant FO properties over transitive frames, but not for bisimulation invariant FO properties over finite transitive frames. Adding the modality $\Diamond^*$ to ML, we obtain a logic which is expressively complete for bisimulation invariant FO properties over *finite* transitive frames, but not bisimulation invariant over the class of all transitive frames.

Simultaneous characterisations of FO/$\sim$ and MSO/$\sim$ over various classes of path-finite transitive frames, and in particular over the class of all finite transitive frames provide analogues of the Janin–Walukiewicz result over these classes – in particular a finite model theory analogue over the class of all finite transitive frames. Since MSO/$\sim$ and hence the μ-calculus collapse to FO/$\sim$ over the classes of transitive frames considered in this context, we also obtained ramifications of the de Jongh–Sambin collapse results for the μ-calculus over these classes of path-finite transitive frames.

Acknowledgement We would like to thank Rebecca Lukas for a careful reading of an earlier version which led to a number of corrections and improvements; we are also grateful to Johan van Benthem and Balder ten Cate as well as the anonymous referees for valuable comments and suggestions including the pointer to [16], which led to the ramifications presented in section 4.7. An error in an earlier version of section 4.6 was pointed out through an instructive counterexample (Example 4.2) by Giovanna D'Agostino and Giacomo Lenzi – thanks for their watchfulness. We would also like to thank the Isaac Newton Institute in Cambridge where some of this work was carried out, during the programme on Logic and Algorithms.

References

[1] L. Alberucci and A. Facchini. *The modal μ-calculus hierarchy over restricted classes of transition systems.* To appear in The Journal of Symbolic Logic.

[2] H. Andréka, J. van Benthem, and I. Németi. *Modal Languages and Bounded Fragments of Predicate Logic.* Journal of Philosophical Logic, 27:217–274, 1998.

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

[4] G. D'Agostino and G. Lenzi. *On the μ-calculus over transitive and finite transitive frames.* Preprint 2008.

[5] A. Dawar and M. Otto. *Modal characterisation theorems over special classes of frames.* In *Proceedings of 20th IEEE Symposium on Logic in Computer Science, LICS'05*, pages 21–30, 2005.

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

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

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

[9] V. Goranko and M. Otto. *Model Theory of Modal Logic.* In Blackburn, van Benthem, Wolter, editors, *Handbook of Modal Logic*, pages 249–330. Elsevier 2006.

[10] 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.

[11] F. Moller and A. Rabinovich. *On the expressive power of CTL**. In *Proceedings of 14th IEEE Symposium on Logic in Computer Science, LICS'99*, pages 360–369,1999.

[12] M. Otto. *Elementary proof of the van Benthem–Rosen characterisation theorem.* Technical Report 2342, Department of Mathematics, Technische Universität Darmstadt, 2004.

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

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

[15] C. Smorynski. *Beth's theorem and self-referential sentences.* In A. Macintyre, L. Pacholski, and J. Paris (editors), *Logic Colloquium '77*, pages 253–260, North Holland, 1978.

[16] B. ten Cate, G. Fontaine, and T. Litak. *Some modal aspects of XPath*, preprint, in preproceedings of Methods for Modalities, 2007.

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