Elementary Proof of the van Benthem-Rosen Characterisation Theorem

Martin Otto *

July 2003; revised May 04

Abstract

This note presents an elementary proof of the well-known characterisation theorem that associates propositional modal logic with the bisimulation invariant fragment of first-order logic. The classical version of this theorem is due to Johann van Benthem [2], its finite model theory analogue to Eric Rosen [8].

1 Introduction

The present proof of the van Benthem/Rosen characterisation theorem is uniformly applicable in both the classical and in the finite model theory scenario. While it is broadly based on Rosen's proof, it reduces the technical input from classical logic and the model theory of modal logics strictly to the use of Ehrenfeucht-Fraïssé games (for first-order, and for the modal variant). Furthermore the proof is constructive and the model constructions and accompanying analysis of games in the expressive completeness argument yield an optimal bound on the modal nesting depth in terms of the first-order quantifier rank. Despite this strengthening, the material becomes presentable in a highly self-contained manner, and can be covered even at the level of an introductory undergraduate course on logic and semantic games that covers the basic Ehrenfeucht-Fraïssé techniques.

Elsewhere this approach has been shown to extend and generalise to characterisations involving stricter forms of bisimulation (global and two-way, and to guarded bisimulation equivalence in transition systems) and corresponding extensions of basic modal logic in [6, 7]. A brief discussion is provided in Section 4.

* Mathematical Logic and Foundations of Computer Science, Department of Mathematics, Darmstadt University of Technology, otto@mathematik.tu-darmstadt.de

2 Notation & Preliminaries

2.1 Kripke structures and basic modal logic

2.1.1 Kripke structures

Consider Kripke structures or transition systems over a finite relational vocabulary consisting (w.l.o.g. for this note) of a single binary relation E and finitely many unary predicates $\mathbf{P} = (P_1, \ldots)$. We write $\mathcal{A} = (A, E^{\mathcal{A}}, \mathbf{P}^{\mathcal{A}})$ for a Kripke structure of this type, and typically indicate a distinguished element as in $\mathcal{A}, a$.

For $a \in A$, let $E^{\mathcal{A}}[a] = \{a' \in A \colon (a, a') \in E_j^{\mathcal{A}}\}$.

2.1.2 Basic modal logic

Denote as ML, or more specifically as ML$[E; \mathbf{P}]$ propositional modal logic over this vocabulary. The formulae of ML$[E; \mathbf{P}]$ are generated from $\top$, $\bot$ and the P in $\mathbf{P}$ allowing Boolean connectives and the modal quantifiers $\Box$ and $\Diamond$. The semantics is the usual one, with

$$\begin{array}{rlcl} & \mathcal{A}, a \models \Box\varphi & \text{iff} & \mathcal{A}, a' \models \varphi \text{ for all } a' \in E^{\mathcal{A}}[a], \\ \text{and (dually)} & \mathcal{A}, a \models \Diamond\varphi & \text{iff} & \mathcal{A}, a' \models \varphi \text{ for some } a' \in E^{\mathcal{A}}[a]. \end{array}$$

We regard ML as a fragment of FO (or indeed FO^2, first-order logic with just two distinct variable symbols, x and y) via the standard translation based on

$$\begin{array}{rcl} [\Box\varphi]^*(x) & = & \forall y\big(Exy \rightarrow [\varphi]^*(y)\big), \\ [\Box\varphi]^*(y) & = & \forall x\big(Eyx \rightarrow [\varphi]^*(x)\big). \end{array}$$

We let ML$_\ell$ stand for the fragment of modal logic consisting of formulae whose nesting depth w.r.t. $\Box/\Diamond$ is at most ℓ. Note that the modal nesting depth coincides with the FO quantifier rank in terms of the standard translation.

2.1.3 Tree structures and their local relatives

A Kripke structure with distinguished element, $\mathcal{A}, a$, is called a *tree structure* if the underlying graph $(A, E^{\mathcal{A}})$ is a directed tree with root a in the graph theoretic sense: E is loop-free and every node is reachable from a on a unique E-path.

A tree structure is of *depth* ℓ if the lengths of paths is bounded by ℓ.

As an intermediary between arbitrary (finite) Kripke structures and tree structures, we consider (finite) structures that look like trees up to a certain depth from the distinguished node.

Generally, in a Kripke structure $\mathcal{A}$ let the *ℓ-neighbourhood* of $a \in A$ be the set $U^\ell(a)$ of all nodes reachable from a on (directed, forward) E-paths of length up to ℓ; $\mathcal{A}{\restriction}U^\ell(a)$ correspondingly denotes the substructure induced in restriction to $U^\ell(a)$.

We say that $\mathcal{A}, a$ is *ℓ-locally a tree structure* iff $\mathcal{A}{\restriction}U^\ell(a), a$ is a tree structure.

2.2 Bisimulation

Bisimulation equivalence between Kripke structures with distinguished nodes is denoted as in $\mathcal{A}, a \sim \mathcal{B}, b$. The corresponding approximations to level ℓ (as induced by the ℓ-round bisimulation game, see below) are denoted as in $\mathcal{A}, a \sim_\ell \mathcal{B}, b$.

The bisimulation game is played by players **I** and **II** over two Kripke structures $\mathcal{A}, a$ and $\mathcal{B}, b$, which each carry a pebble, initially placed on the distinguished elements a and b, respectively. In each round, the challenger, player **I**, moves the pebble in one of the structures forward along an E-edge, and player **II** has to respond by moving the other pebble along an E-edge in the opposite structure. It is player **II**'s task to maintain atomic equivalence throughout: **II** loses as soon as the currently pebbled nodes fail to agree on all monadic predicates (atomic propositions). Apart from that, players lose when they cannot move, for lack of E-edges. We say that **II** has a *winning strategy in the (infinite) bisimulation game* on $\mathcal{A}, a$ and $\mathcal{B}, b$, if she has a strategy to respond to any challenges from **I** without losing, indefinitely; **II** has a *winning strategy in the ℓ-round bisimulation game* on $\mathcal{A}, a$ and $\mathcal{B}, b$, if she has a strategy to respond to any challenges from **I** without losing for ℓ rounds. Then

- $\mathcal{A}, a$ and $\mathcal{B}, b$ are bisimilar, $\mathcal{A}, a \sim \mathcal{B}, b$, iff **II** has a winning strategy in the (infinite) bisimulation game on $\mathcal{A}, a$ and $\mathcal{B}, b$.
- $\mathcal{A}, a$ and $\mathcal{B}, b$ are ℓ-bisimilar, $\mathcal{A}, a \sim_\ell \mathcal{B}, b$, iff **II** has a winning strategy in the ℓ-round bisimulation game on $\mathcal{A}, a$ and $\mathcal{B}, b$.

The standard Ehrenfeucht-Fraïssé analysis of the bisimulation game yields the following.

Lemma 2.1. *Over the class of all Kripke structures of a fixed finite relational type:*

(i) *$\sim_\ell$ has finite index;*

(ii) *$\mathcal{A}, a \sim_\ell \mathcal{B}, b$ iff a in $\mathcal{A}$ and b in $\mathcal{B}$ are indistinguishable in* ML_ℓ*;*

(iii) *each $\sim_\ell$ equivalence class is definable by an* ML_ℓ *formula.*

A further few simple but useful properties of bisimulation equivalence are summarised in the following. In the first lemma we refer to the operation of *disjoint sums* or *disjoint unions* of relational structures: if $\mathcal{A}$ and $\mathcal{C}$ are structures of the same relational type, we denote as $\mathcal{A} + \mathcal{C}$ their disjoint sum (union), whose universe is the disjoint union of the universes and with all relations interpreted as in $\mathcal{A}$ and $\mathcal{C}$, respectively. The second lemma captures the local nature of $\sim_\ell$.

Lemma 2.2. *Bisimulation equivalence is insensitive to disjoint sums. If $\mathcal{A}, \mathcal{B}, \mathcal{C}$ are of the same relational type, then $\mathcal{A}, a \sim \mathcal{B}, b$ iff $\mathcal{A} + \mathcal{C}, a \sim \mathcal{B}, b$.*

Lemma 2.3. (i) *$\mathcal{A}, a \sim_\ell \mathcal{B}, b$ iff $\mathcal{A} \restriction U^\ell(a), a \sim_\ell \mathcal{B} \restriction U^\ell(b), b$.*

(ii) *if $\mathcal{A}, a$ and $\mathcal{B}, b$ are both tree structures of depth ℓ, then ℓ-bisimulation coincides with bisimulation: $\mathcal{A}, a \sim_\ell \mathcal{B}, b$ iff $\mathcal{A}, a \sim \mathcal{B}, b$.*

The familiar and intuitive process of *unravelling* always guarantees bisimilar companions that are tree structures, albeit typically infinite ones. If only ℓ-local tree likeness is required, finite bisimilar companions are easily constructed for finite structures.

The *tree unravelling* $\mathcal{A}_a^*$ of $\mathcal{A}$ from a, is obtained as follows. The universe of $\mathcal{A}_a^*$ is the set of all (directed, forward) E-paths from a in $\mathcal{A}$. E is interpreted in $\mathcal{A}_a^*$ so that for each $m \in \mathbb{N}$, each path of length $m+1$ is an E-successor of its initial segment of length m. The unary predicates are interpreted in accordance with the projection $\pi \colon \mathcal{A}_a^* \to \mathcal{A}$ that maps each path to its last node.

Lemma 2.4. *Let $\mathcal{A}, a$ be a Kripke structure with distinguished node a.*

(i) *The tree unravelling of $\mathcal{A}$ from a, $\mathcal{A}_a^*$, is a tree structure that is bisimilar to $\mathcal{A}$ via the natural projection $\pi \colon \mathcal{A}_a^*, a \sim \mathcal{A}, a$.*

(ii) *For every $\ell \in \mathbb{N}$, the restriction of the tree unravelling $\mathcal{A}_a^*$ to depth ℓ, is a tree structure of depth ℓ that is ℓ-bisimilar to $\mathcal{A}, a$: $\pi \colon \mathcal{A}_a^* \restriction U^\ell(a), a \sim_\ell \mathcal{A}, a$.*

(iii) *For a finite Kripke structure $\mathcal{A}$ with distinguished node a, and $\ell \in \mathbb{N}$: there is a partial unravelling (to depth ℓ) that yields a finite bisimilar companion that is ℓ-locally a tree structure.*

Proof. (i) is obvious. For (ii) one may appeal to (i) of the previous lemma.

For (iii), take the tree unravelling $\mathcal{A}_a^*$ in restriction to $U^\ell(a)$, and identify each node b^* in $\mathcal{A}_a^* \restriction U^\ell(a)$ at distance ℓ from the root (a leaf in $\mathcal{A}_a^* \restriction U^\ell(a)$) with the node $b = \pi(b^*)$ in a fresh disjoint isomorphic copy of $\mathcal{A}$. $\square$

2.2.1 Bisimulation invariance

Definition 2.5. A formula $\varphi(x) \in \mathrm{FO}[E; \mathbf{P}]$ is *bisimulation invariant* iff, whenever $\mathcal{A}, a \sim \mathcal{B}, b$ then $\mathcal{A}, a \models \varphi$ iff $\mathcal{B}, b \models \varphi$.

All formulae of ML are bisimulation invariant; in fact, $\varphi \in \mathrm{ML}_\ell$ is invariant under $\sim_\ell$. This is an immediate consequence of the modal Ehrenfeucht-Fraïssé analysis, or simply proved directly by syntactic induction on φ.

2.3 Locality

Gaifman's notion of *locality* [5] has been extensively studied in the first-order context, and in particular has proved to be a useful tool in finite model theory [4]. We here only need to make very limited use of the simple concept of ℓ-locality of a first-order formula in one free variable.

Definition 2.6. A property of Kripke structures with distinguished nodes $\mathcal{A}, a$—or a formula $\varphi(x)$ defining such a property—is *ℓ-local* iff whether or not it is satisfied in $\mathcal{A}, a$ only depends on $\mathcal{A} \restriction U^\ell(a), a$:

$$\mathcal{A}, a \models \varphi \quad \Leftrightarrow \quad \mathcal{A} \restriction U^\ell(a), a \models \varphi.$$

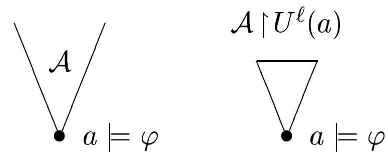

The following is a simple consequence of ℓ-bisimulation invariance, Lemma 2.1, and (i) in Lemma 2.3.

Observation 2.7. *Any $\varphi \in \mathrm{ML}_\ell$ is ℓ-local.*

3 The characterisation theorem

The goal of this note is a simple proof of the following characterisation theorem which goes through uniformly in the sense of finite model theory and classically. As an added benefit we obtain an optimal quantitative bound on quantifier ranks involved.

Theorem 3.1 (van Benthem/Rosen). *The following are equivalent for any $\varphi(x) \in$ FO of quantifier rank q:*

(i) $\varphi(x)$ *is invariant under bisimulation* [in finite Kripke structures].

(ii) $\varphi(x)$ *is logically equivalent* [over finite Kripke structures] *to a formula of* ML_ℓ*, where* $\ell = 2^q - 1$.

Note that the two readings—one classical, one finite model theoretic— really are two distinct theorems, a priori independent of each other. [1] Note that bisimulation invariance in finite structures does not imply bisimulation invariance over all structures: trivial examples are provided by formulae without finite models that happen not to be bisimulation invariant for infinite models.

Our proof proceeds in three stages. Note that even though we do not make this implicit in the statements, each statement is considered in its two readings: classically and in the sense of finite model theory.

Step 1 Any bisimulation invariant $\varphi(x) \in$ FO is ℓ-local for $\ell = 2^q - 1$ where $q = \mathrm{qr}(\varphi)$. This is proved with FO Ehrenfeucht-Fraïssé games, and as far as bisimulation is concerned rests on Lemma 2.2.

Step 2 Any bisimulation invariant $\varphi(x)$ that is ℓ-local, is even invariant under ℓ-bisimulation equivalence $\sim_\ell$. A simple bisimulation argument based on Lemmas 2.3 and 2.4 shows this.

Step 3 Any property invariant under ℓ-bisimulation equivalence is definable in ML_ℓ. This is a direct consequence of the Ehrenfeucht-Fraïssé analysis of bisimulation, based on Lemma 2.1 (iii).

[1] Two-variable first-order logic and two-pebble game equivalence illustrate this point. In that case, the classical characterisation theorem does not hold as a theorem of finite model theory, finite model property of FO^2 notwithstanding.

Proof sketch: step 1 Assume that $\varphi(x) \in \mathrm{FO}$ is bisimulation invariant, let $q = \mathrm{qr}(\varphi)$, and put $\ell := 2^q - 1$. To show that $\varphi(x)$ is ℓ-local, we consider any $\mathcal{A}, a$ and show that $\mathcal{A}, a \models \varphi$ iff $\mathcal{A} \restriction U^\ell(a), a \models \varphi$. As φ is bisimulation invariant, we may w.l.o.g. assume that $\mathcal{A} \restriction U^\ell(a), a \models \varphi$ is a tree of depth ℓ. [We may pass to a (finite partial) unravelling of $\mathcal{A}$, which is a bisimilar companion of $\mathcal{A}, a$ and whose restriction to $U^\ell(a)$ automatically is a bisimilar companion to $\mathcal{A} \restriction U^\ell(a), a \models \varphi$.] Here and in the core argument we use the fact that, due to bisimulation invariance, we may always replace a structure by some bisimilar companion without affecting (the property expressed by) φ.

If for some $\mathcal{A}', a' \sim \mathcal{A}, a$ and $\mathcal{A}'', a'' \sim \mathcal{A} \restriction U^\ell(a), a$, we can show $\mathcal{A}', a' \equiv_q \mathcal{A}'', a''$, we are done. For then

$$\begin{array}{lll} & \mathcal{A}, a \models \varphi & \\ \text{iff} & \mathcal{A}', a' \models \varphi & \text{(bisimulation invariance)} \\ \text{iff} & \mathcal{A}'', a'' \models \varphi & (\equiv_q \text{ equivalence}) \\ \text{iff} & \mathcal{A} \restriction U^\ell(a), a \models \varphi & \text{(bisimulation invariance)} \end{array}$$

For suitable $\mathcal{A}'$ and $\mathcal{A}''$, the equivalence $\mathcal{A}', a' \equiv_q \mathcal{A}'', a''$ can be established by exhibiting a strategy for player II in the q-round Ehrenfeucht-Fraïssé game.

As companions of $\mathcal{A}, a$ and $\mathcal{A} \restriction U^\ell(a), a$, respectively, we choose structures that are disjoint copies of sufficiently many isomorphic copies of $\mathcal{A}, a$ and $\mathcal{A} \restriction U^\ell(a), a$. Both structures involved will have q isomorphic copies of both $\mathcal{A}, a$ and $\mathcal{A} \restriction U^\ell(a), a$, and only distinguish themselves by the nature of the one extra component, in which live a' and a'', respectively. We indicate the two structures in the diagram below, with distinguished elements a' and a'' marked $\bullet$; the open cones stand for copies of $\mathcal{A}$, the closed cones for copies of $\mathcal{A} \restriction U^\ell(a)$. Clearly the structure on the left is bisimilar to $\mathcal{A}, a$, the one on the right bisimilar to $\mathcal{A} \restriction U^\ell(a), a$, by Lemma 2.2.

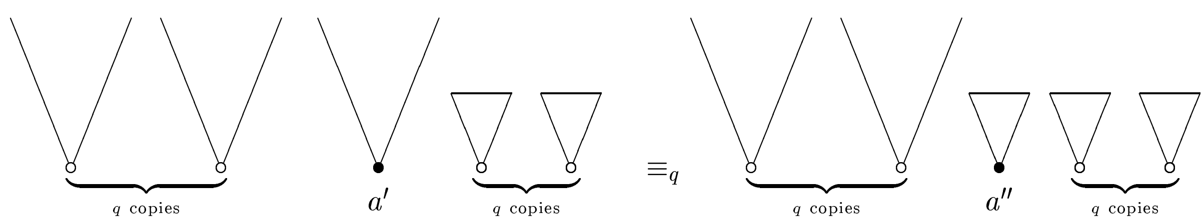

It suffices now to exhibit a strategy for player **II** in q rounds of the game on these structures. The game is started in the configuration with a single pebble in positions marked $\bullet$ in each of the two structures. The description of the strategy makes reference to a *critical distance* d_m, whose value for round m is

$$d_m = 2^{q-m}.$$

Starting with value $d_1 = 2^{q-1} = \lceil \ell/2 \rceil$ in the first round, this critical distance decreases by a factor $1/2$ in each round. Under the proposed strategy, **II** will always play *according* *to local context* if **I**'s move goes to an element within the critical distance from any already marked element; if **I**'s move is further than the critical distance away from all previously marked elements, we let **II** respond by marking *the same* element in one of the isomorphic copies of $\mathcal{A}$ or $\mathcal{A} \restriction \ell$ that has not yet been touched by the game. (There are q many copies of each on each side, hence fresh ones are always available.)

The idea of *local context* works as follows. We think of the pebbles as belonging to disjoint clusters; initially we have just one single cluster consisting of the single elements with pebbles $\bullet$ on each side.

A pebble that is newly placed in round m joins an existing cluster if it is at most distance d_m away from one of the members of that cluster. Note that because of the shrinking d_m, no two clusters can ever be joined. Any elements of different clusters after round m are more than d_m apart.

Our strategy for **II** will have her maintain the condition that after completion of round m

> any two corresponding clusters are linked by an isomorphism that extends to all points within distance d_m of the members of the clusters.

If in round m, **I** places a new pebble further than d_m away from all previously marked elements, it forms a new cluster on its own, and **II**'s response into a new component makes sure that the same happens on the other side. If **I** places a new pebble to join one of the existing clusters, then **II** uses the isomorphism that comes with that cluster to respond with a matching element to join the corresponding cluster on the other side.

One checks that the above invariant is satisfied initially, and that the prescriptions for **II**'s moves are such that it is maintained in round m through all $m = 1 \ldots, q$.

After q rounds, the local isomorphisms between clusters still guarantee that **II** wins.

Proof sketch: step 2 Let $\varphi(x)$ be ℓ-local and bisimulation invariant. Suppose $\mathcal{A}, a \sim_\ell \mathcal{B}, b$ and $\mathcal{A}, a \models \varphi$. We need to show that then also $\mathcal{B}, b \models \varphi$. Without loss of generality, we may assume that $\mathcal{A}$ and $\mathcal{B}$ are ℓ-locally tree structures. [If they are not, pass to (finite partial) unravellings, Lemma 2.4]

By ℓ-locality, $\mathcal{A}, a \models \varphi$ iff $\mathcal{A} \restriction U^\ell(a), a \models \varphi$. Now $\mathcal{A}, a \sim_\ell \mathcal{B}, b$ iff $\mathcal{A} \restriction U^\ell(a), a \sim_\ell \mathcal{B} \restriction U^\ell(b), b$ iff (as both structures are now trees of depth ℓ) $\mathcal{A} \restriction U^\ell(a), a \sim \mathcal{B} \restriction U^\ell(a)$; Lemma 2.3.

Hence $\mathcal{A} \restriction U^\ell(a), a \models \varphi$ iff $\mathcal{B} \restriction U^\ell(b), b \models \varphi$. Therefore $\mathcal{B}, b \models \varphi$, by ℓ-locality again.

Proof sketch: step 3 If φ is invariant under $\sim_\ell$, we may use the ML_ℓ formulae that define $\sim_\ell$ equivalence classes, according to Lemma 2.1 (iii). Let $\chi_{\ell,\mathcal{A},a} \in \mathrm{ML}_\ell$ be the formula that defines the $\sim_\ell$ equivalence class of $\mathcal{A}, a$. Then φ is equivalent to the disjunction

$$\varphi \equiv \bigvee_{\mathcal{A},a \models \varphi} \chi_{\ell,\mathcal{A},a},$$

which is equivalent to a finite disjunction as $\sim_\ell$ has finite index.

This finishes the proof of the characterisation theorem.

Exercise 3.1. The exponential gap between the first-order quantifier rank q and modal quantifier rank $\ell = 2^q - 1$ cannot be avoided in general, as the example of formulae expressing that "there is an element satisfying p within distance $2^q - 1$" shows. Show that this is expressible in FO_q, but not by any formula in ML_m for $m < 2^q - 1$. [It can be expressed in modal quantifier rank $2^q - 1$.]

Observation 3.2. *There is an exponential succinctness gap between* FO *and* ML, *in expressing bisimulation invariant properties.*

4 Ramifications

In essence the technique outlined for basic modal logic above extends to other settings. We mention three distinct lines of variations and extensions.

Firstly, the classical statement of the theorem relativises to FO-definable classes of structures (as is clear also from the classical proof); but both the classical and the finite model theoretic versions also relativise to arbitrary bisimulation-closed classes.

Secondly, one can treat stronger variants of bisimulation equivalence, and in particular global forms of bisimulation, with a corresponding shift to technically more demanding locality arguments that need to apply uniformly across the entire structure rather than in a neighbourhood of the distinguished node.

Thirdly, one may want to consider other natural, more restricted classes of structures, rather than the class of all (or all finite) Kripke structures. Natural cases of interest include in particular classes of frames defined in terms of connectivity constraints, and in terms of classes of frames corresponding to classical modal theories.

4.1 Straightforward relativisations

Let $\mathcal{C}$ be a class of Kripke structures with distinguished elements. The notions of bisimulation invariance and of definability in ML give rise to corresponding notion in restriction to $\mathcal{C}$. For instance, $\varphi(x)$ is bisimulation invariant over $\mathcal{C}$ if for any two structures $\mathcal{A}, a$ and $\mathcal{B}, b$ from $\mathcal{C}$, $\mathcal{A}, a \sim \mathcal{B}, b$ implies that $\mathcal{A}, a \models \varphi$ iff $\mathcal{B}, b \models \varphi$. The classical proof of van Benthem's theorem uses compactness and saturation properties to establish indirectly that bisimulation invariance of $\varphi(x) \in \mathrm{FO}$ implies invariance ℓ-bisimulation for some ℓ. This argument clearly relativises to work within any class $\mathcal{C}$ defined by an FO theory.

The game oriented proof we gave above, on the other hand, is easily seen to work in restriction to (the finite structures within) any class $\mathcal{C}$ that is itself bisimulation closed.

Corollary 4.1. *Let $\mathcal{C}$ be closed under bisimulation, $\mathcal{C}_{fin}$ the class of finite structures within $\mathcal{C}$. Then $\varphi(x)$ is invariant under bisimulation over $\mathcal{C}$* [over $\mathcal{C}_{\mathrm{fin}}$] *iff $\varphi(x)$ is logically equivalent over $\mathcal{C}$* [over $\mathcal{C}_{\mathrm{fin}}$] *to a formula of ML_ℓ, where $\ell = 2^q - 1$.*

4.2 Stronger forms of bisimulation

Variations of this kind have been studied in [6, 7]. The following strengthenings of basic bisimulation equivalence are treated (described here in terms of the modifications in the corresponding bisimulation games):

(i) two-way bisimulation: **I** also has the option to move backward along E-edges, in which case **II** has to respond likewise.

(ii) global bisimulation: **I** can opt to move the pebble to a fresh start node anywhere in the structure, as can **II** in her response to such a move.

(iii) two-way and global bisimulation, $\approx$: both of the above.

It is entirely straightforward to adapt (the classical proof of) the classical characterisation theorem of van Benthem's to cover these variations. One naturally finds that these refined notions of bisimulation characterise within FO the following extensions of basic modal logic:

(i) two-way bisimulation: ML^{-}, ML with backward (past) modalities like $\Diamond^{-}\varphi(x) \equiv \exists y(Eyx \wedge \varphi(y))$.

(ii) global bisimulation: $\mathrm{ML}^{\forall}$, ML with a global modality, corresponding to unrestricted universal/existential quantification as in $\exists x \varphi(x)$ where $\varphi \in \mathrm{ML}$.

(iii) two-way and global bisimulation, $\approx$: $\mathrm{ML}^{-\forall}$, the combined extension by both of the above.

The global variants are technically interesting, because they require non-trivial locality arguments. We discuss the key case of $\approx$ invariance (global two-way bisimulation). On the one hand, the given FO formula $\varphi(x)$ is analysed in terms of Gaifman's locality theorem for FO, [5]. This allows us to determine locality parameters ℓ, m, q from φ such that whenever $\mathcal{A}, a$ and $\mathcal{B}, b$ agree

- on FO properties of quantifier rank q in the ℓ-neighbourhoods of a and b, respectively;
- on the quantifier rank q FO properties of systems of up to m many disjoint ℓ-neighbourhoods anywhere within $\mathcal{A}$ or $\mathcal{B}$, respectively;

then $\mathcal{A}, a \models \varphi$ iff $\mathcal{B}, b \models \varphi$.

In order to show the analogue of the crucial step 2, that then φ actually is invariant under $\approx_\ell$, one can construct, for arbitrary $\mathcal{A}, a \approx_\ell \mathcal{B}, b$, fully $\approx$ equivalent companion structures $\mathcal{A}^*, a \approx \mathcal{A}, a$ and $\mathcal{B}^*, b \approx \mathcal{B}, b$ that agree locally for ℓ, m, q in the above sense.

In the classical case, infinite companions $\mathcal{A}^*$ and $\mathcal{B}^*$ are admissible, and one can resort to bisimilar tree models obtained as suitable two-way unravellings, over which ℓ-two-way-bisimilarity then enforces local first-order equivalence.

In the finite model theory instance of the argument, one similarly seeks companion structures that are, at least ℓ-locally tree-like (acyclic), but at the same time need to be kept finite. This can be achieved with a construction of locally acyclic bisimilar coverings as developed in [6, 7]. These techniques have also been shown to extend to guarded bisimulation equivalence and to a characterisation of the guarded fragment GF of first-order logic, [1], over relational structures of width 2.

The general case for guarded bisimulation and the guarded fragment GF in relational structures with predicates of higher arities is the theme of ongoing investigations. In fact, it is currently open, whether the characterisation of GF as the guarded bisimulation invariant fragment of FO, due to Andréka, van Benthem and Németi [1], also obtains in the context of finite model theory. For this it would seem to be necessary to lift essential features of the construction of finite locally acyclic covers from the graph theoretic setting of bisimulations to the hypergraph theoretic setting of guarded bisimulations.

Further ramifications currently under investigation concern counting bisimulations (where the number of available successors of a certain kind matters) and modal logics with graded modalities.

4.3 Other natural classes of frames

These variations look at characterisation theorems for modal logics, of the above kind, over still more restricted classes of (finite) Kripke structures. Of particular interest from a transition system point of view are connected systems, as the existence of disconnected components (unreachable states) is often counterintuitive. Consider, for instance, the class of Kripke structures $\mathcal{A}, a$ that are *connected* in the sense that each node is reachable on some (forward, directed) E-path from a. In restriction to such connected frames, bisimulation equivalence coincides with global bisimulation equivalence, and correspondingly one expects a characterisation of the following kind.

Proposition 4.2. *The following are equivalent for every* $\varphi(x) \in \mathrm{FO}$*, both classically and in the sense of finite model theory:*

(i) *φ is invariant under bisimulation over the class of all* [finite] *connected Kripke structures.*

(ii) *φ is invariant under global bisimulation over the class of all* [finite] *connected Kripke structures.*

(iii) *φ is equivalent to a formula of* $\mathrm{ML}^{\forall}$ *over the class of all* [finite] *connected Kripke structures.*

Indeed, the constructive approach outlined above adapts to these settings to prove this proposition, as well as several other natural characterisation theorem of this kind. Interestingly, there seems to be no straightforward classical proof along the standard lines, as the underlying class of connected Kripke structures is not an elementary class, and compactness arguments are not directly available.

A number of other ramifications related to, for instance, frame conditions dealing with symmetry or transitivity requirements are being considered in ongoing joint work with A. Dawar [3]. Transitivity in particular is interesting from a technical point of view, as locality cannot be used in a straightforward manner.

References

1. H. Andréka, J. van Benthem, and I. Németi, *Modal languages and bounded fragments of predicate logic*, Journal of Philosophical Logic, 27 (1998), pp. 217–274.

2. J. van Benthem, *Modal Logic and Classical Logic*, Bibliopolis, Napoli, 1983.

3. A. Dawar and M. Otto, *Modal Characterisation Theorems over Connected Frames*, in preparation/unpublished note, 2003.

4. H.-D. Ebbinghaus and J. Flum, *Finite Model Theory*, Springer, 2nd ed., 1999.

5. H. Gaifman, *On local and nonlocal properties*, in Logic Colloquium '81, J. Stern, ed., North Holland, 1982, pp. 105–135.

6. M. Otto, *Bisimulation-invariant Ptime and Higher-Dimensional μ-calculus*, *Proc. 4th Symp. on Logical Foundations of Computer Science*, Lecture Notes in Computer Science, vol. 1234, Springer, 1997, pp. 294–305.

7. M. Otto, *Bisimulation-invariant Ptime and Higher-Dimensional μ-calculus*, to appear in TCS special issue, 2004.

8. E. Rosen, *Modal logic over finite structures*, Journal of Logic, Language and Information, 6 (1997), pp. 427–439.