BISIMULATION INVARIANCE AND FINITE MODELS

MARTIN OTTO

Abstract. We study bisimulation invariance over finite structures. This investigation leads to a new, quite elementary proof of the van Benthem-Rosen characterisation of basic modal logic as the bisimulation invariant fragment of first-order logic. The ramification of this characterisation for the finer notion of global two-way bisimulation equivalence is based on bisimulation respecting constructions of models that recover in finite models some of the desirable properties of the usually infinite bisimilar unravellings.

§1. Introduction and preliminaries. Model theory—understood in the broad sense of the study of the expressive power of logical languages and definability of structural properties, or the interplay between syntax (*language*) and semantics (*structural properties*)—has applications in many areas. Many of these fall outside the realm of 'classical' logic and model theory, because any reasonably mature application area comes with its own pre-defined notion of what are its *relevant structures* and which are the *relevant structural properties* for its purposes. Decisions about syntax and semantics should be subordinate to the mathematical modelling that captures the essence of the application by delineating the class of structural properties that matter.

First and foremost these decisions concern the class of structures to be considered (certain classes of algebras or relational structures, possibly just finite structures of a certain kind, etc.); and a notion of equivalence between structures (isomorphism, elementary equivalence, bisimulation equivalence, etc.). While the class of structures depends on *what* is to be modelled, the accompanying notion of equivalence depends on the *coarseness* of this modelling, or its level of abstraction. Classical algebra, for instance, would typically consider the class of all algebras of a certain type, and analyse them up to isomorphism. Classical model theory is concerned with classes

Key words and phrases. Finite model theory, modal logic, guarded fragment, bisimulation, preservation and characterisation theorems.

Research carried out at the Computer Science Department, University of Wales Swansea, United Kingdom; partially supported by EPSRC grant GR/R11896/01.

Meeting
Edited by Unknown

of all structures of a certain type, considered up to elementary equivalence. In fact, classical model theory to a large extent can be understood as an exploration of elementary equivalence (first-order logic) and its relationship to isomorphism (algebra). Finite model theory on the other hand restricts its domain to just the finite structures of certain type, under a variety of equivalences that are coarser than elementary equivalence. Elementary equivalence is of no particular interest since over finite structures it coincides with isomorphism. Apart from equivalences like bounded-variable elementary equivalence, which are partly motivated as appropriate technical substitutes for elementary equivalence in this setting, other notions of equivalence are motivated by specific requirements from application domains in the above sense. One such domain is the theoretical study of the behaviour of computational processes, which has long been a fruitful domain of interaction between logic and computer science.

Transition systems are edge and vertex coloured directed graphs, i.e. relational structures $\mathfrak{A} = (A, E^{\mathfrak{A}}, \dots, P^{\mathfrak{A}}, \dots)$ in a vocabulary consisting of unary predicates P and binary relations E. Intuitively the elements of the universe A of $\mathfrak{A}$ model the states, the unary P describe atomic properties of states, and the binary edge relations E formalise directed transitions from state to state.[1] In as far as possible *behaviours* are concerned, one is interested in transition systems not up to isomorphism but up to bisimulation equivalence. In fact bisimulation equivalence was, apparently independently, isolated as an adequate notion of equivalence both in the context of process analysis [11, 20] (where the term bisimulation was coined), and in the Ehrenfeucht-Fraïssé analysis of Kripke structures for propositional modal logics [23, 24] (under the name of zig-zag equivalence). The common domain of *modal model theory* is that of transition systems (or Kripke structures) up to bisimulation equivalence, or the *model theory of bisimulation invariance.* Of course this can be further ramified, according to whether, for instance, we want to consider the domain of all transition systems, or restrict to the subclass of finite transition systems (process behaviours in finite state systems), or other subclasses (e.g., connected, finite and connected, etc). Generally one can expect any such domain to have its own characteristic flavour of model theory, equally dependent on the class of structures as on the notion of equivalence. The modal domain and its ramifications provide a particularly successful example of specific model theoretic techniques and results, both classically (all structures) and in the sense of finite model theory (just finite structures). An aspect particularly relevant to computer science applications is that within the modal

[1]Formally there is no distinction between this setting and that of relational formalisations of Kripke structures for propositional modal logic with atomic propositions P and accessibility relations E between possible worlds $a \in A$.

domain there are logics whose expressive power goes far beyond that of first-order logic while still being algorithmically very tractable. Tractability here mainly concerns decidability of the satisfiability problem and feasible model checking complexities. Indeed, bisimulation invariance itself and some of its immediate model theoretic consequences, like the tree model property, can account for these good algorithmic features.

A fundamental issue in every ramification of model theory for a particular choice of structures and equivalence is the quest for *expressive completeness*. A logic is expressively complete for a class of structural properties if it precisely expresses just these structural properties. Basic modal logic is expressively complete for the class of bisimulation invariant first-order properties over the class of all transition systems (van Benthem's theorem, Theorem 2.1 below). In other words, basic modal logic is the 'first-order logic for the modal domain'. One of the most important process logics, the modal μ-calculus which extends basic modal logic by means of least and greatest fixed points is similarly known to be expressively complete for precisely the bisimulation invariant monadic second-order properties, by a well known theorem of Janin and Walukiewicz [14].

On the technical side, a model theory for a specific domain will always have its characteristic techniques for constructing and transforming models, in accordance with the underlying equivalence. Bisimulation preserving constructions in the domain of all, finite and infinite, transition systems mostly center on tree models that are naturally obtained as bisimilar unravellings of arbitrary transitions systems. Most of our technical contributions here revolve around bisimulation preserving model constructions in finite models, where unravellings no longer work. It is a characteristic feature of finite model theory in general that model constructions and transformations tend to be combinatorially more involved simply because the result is required to remain finite. This is also true of modal finite model theory.

In the rest of this section, we provide some of the basic definitions and fix some notation. Much of it can safely be skipped by readers familiar with the fundamentals of Ehrenfeucht-Fraïssé games and equivalences, bisimulation games and equivalences, and with basic modal logic.

We typically write $\mathfrak{A} = (A, E^{\mathfrak{A}}, \dots, P^{\mathfrak{A}}, \dots)$ for a structure with universe A, binary predicates $E^{\mathfrak{A}}$, ... , and unary predicates $P^{\mathfrak{A}}$, ... ; superscripts $^{\mathfrak{A}}$ are dropped where inessential. Often a distinguished node or parameter $a \in A$ is indicated as in $\mathfrak{A}, a$.

We want to see bisimulation equivalence both as a notion of behavioural equivalence between transition systems or—more classically—as the equivalence induced by the infinite Ehrenfeucht-Fraïssé game whose moves capture the relativised pattern of modal quantification.

1.1. Ehrenfeucht-Fraïssé games. Recall the classical Ehrenfeucht-Fraïssé characterisation of elementary equivalence (see [7, 6, 21] for text-book treatments). The game is played by two players over the two structures, $\mathfrak{A}$ and $\mathfrak{B}$, that are to be compared. In each round of the game the first player marks an element of one of the structures (his choice) and the second player responds by selecting an element in the opposite structure. After k rounds, k elements will have been individually marked in each structure: we denote the resulting configuration as $\mathfrak{A}; a_1, \dots, a_k$ versus $\mathfrak{B}; b_1, \dots, b_k$, where (a_i, b_i) is the pair of elements marked in round i. It is the second player's task to maintain similarity between the two sides: she loses as soon as the correspondence $\{(a_1, b_1), \dots, (a_k, b_k)\}$ stops being a local isomorphism.[2]

The well known Ehrenfeucht-Fraïssé Theorem says that the second player has a winning strategy for q rounds in this game iff $\mathfrak{A}$ and $\mathfrak{B}$ cannot be distinguished by FO-sentences of quantifier rank q, $\mathfrak{A} \equiv_q \mathfrak{B}$. Algebraically or combinatorially the existence of a winning strategy for the second player in the q round game, $\mathfrak{A} \simeq_q \mathfrak{B}$, is captured in a *back-and-forth system* of depth q of local isomorphisms.

A winning strategy in the unbounded game with an infinite succession of rounds in which the second player needs to maintain a partial isomorphism characterises equivalence in the infinitary logic $\mathrm{L}_{\infty\omega}$, $\mathfrak{A} \equiv_{\infty\omega} \mathfrak{B}$ (Karp's Theorem). The corresponding combinatorial notion of partial isomorphy between structures, $\mathfrak{A} \simeq_{\mathrm{part}} \mathfrak{B}$, precisely captures the existence of such a winning strategy as a single back-and-forth system of local isomorphisms. All these notions naturally extend to structures with parameters, which may be thought of as elements marked before the start of the game.

It is natural to introduce bisimulation equivalence in this game context.

1.2. Bisimulation games. The *bisimulation game* on transition systems $\mathfrak{A}, a$ and $\mathfrak{B}, b$ is played by two players as above, but now only one marked element in each structure is kept. Initially the marked elements are the distinguished nodes a and b. In each round the first player selects one of the structures and one of the transitions available from the marked element in that structure and moves the marker along this transition to its target element; the second player has to respond by moving the marker in the opposite structure along a corresponding transition, and she loses if she cannot move (for lack of an appropriate transition) or if the marked elements do not satisfy exactly the same unary predicates. Clearly the existence of a winning strategy for the second player can be captured by a back-and-forth system of pairs in $A \times B$. More precisely, for the

[2] We follow [21] in using the more natural name of *local isomorphisms* for what the older literature calls *partial isomorphisms*: partial maps that are isomorphisms between the substructures induced on their domain and range.

unbounded game the existence of a strategy corresponds to one system $Z \subseteq A \times B$ with $(a, b) \in Z$ (for initialisation in the distinguished elements), with $a' \in P^{\mathfrak{A}} \Leftrightarrow b' \in P^{\mathfrak{B}}$ for all $(a', b') \in Z$ and all unary predicates P, and satisfying the usual back-and-forth conditions:

forth: for all binary E and for all $(a', b') \in Z$ and all a'' such that $(a', a'') \in E^{\mathfrak{A}}$ there is some b'' such that $(b', b'') \in E^{\mathfrak{B}}$ and $(a'', b'') \in Z$.

back: for all binary E and for all $(a', b') \in Z$ and all b'' such that $(b', b'') \in E^{\mathfrak{B}}$ there is some a'' such that $(a', a'') \in E^{\mathfrak{A}}$ and $(a'', b'') \in Z$.

DEFINITION 1.1. Two structures are bisimilar, $\mathfrak{A}, a \sim \mathfrak{B}, b$, if and only if the second player has a winning strategy in the unbounded game starting from $\mathfrak{A}, a$ and $\mathfrak{B}, b$; equivalently, if there is a back-and-forth system Z with $(a, b) \in Z$.

Similarly we may consider an ℓ-round bisimulation game. The formalisation of a winning strategy takes the form of a depth ℓ stratified back-and-forth system, $(Z_i)_{0 \leqslant i \leqslant \ell}$, with $(a, b) \in Z_\ell$ and back-and-forth conditions that correspondingly guarantee the existence of $(a'', b'') \in Z_{i-1}$ for $(a', b') \in Z_i$. We write $\mathfrak{A}, a \sim^\ell \mathfrak{B}, b$ if there is such, i.e., if the second player has a strategy to survive for ℓ rounds in the game on $\mathfrak{A}, a$ and $\mathfrak{B}, b$.

DEFINITION 1.2. We write $\mathfrak{A}, a \sim^\ell \mathfrak{B}, b$ and speak of *ℓ-bisimulation equivalence* if the second player has a winning strategy in the ℓ-round bisimulation game from $\mathfrak{A}, a$ and $\mathfrak{B}, b$; or, equivalently if there is a stratified back-and-forth systems of depth ℓ, with $(a, b) \in Z_\ell$.

Some variations of basic bisimulation equivalence will be considered in Section 4.

1.3. Basic modal logic. When comparing the bisimulation game with the usual FO Ehrenfeucht-Fraïssé game, we see that the severe restriction of moves in the bisimulation game corresponds to restricted access to elements through quantification. Any move is bound to the current position via a transition along some specific binary edge predicate E, in precisely the way that is captured by the modal operators $\langle E \rangle$ and their duals $[E]$. These may be transcribed into relativised first-order quantifiers according to

$$
\begin{aligned}
\big(\langle E \rangle \varphi\big)(x) &\equiv \exists y \big(Exy \wedge \varphi(y)\big), \\
\big([E] \varphi\big)(x) &\equiv \forall y \big(Exy \rightarrow \varphi(y)\big).
\end{aligned} \tag{1}
$$

For us, basic modal logic ML is exactly the fragment of FO consisting of all FO-formulae in a single free variable in which all quantifications take the form of $\langle E \rangle$ and $[E]$ for binary predicates E.

DEFINITION 1.3. We denote basic modal logic as ML, and regard it as a fragment of first-order logic, syntactically generated from atomic formulae Px, where P is unary, as the closure under Boolean connectives and modal quantification according to (1) for all binary predicates E.

1.4. Unravellings. By the *modal domain* we mean the world of transition systems considered up to bisimulation equivalence. One may pass from any given transition system to one of its bisimilar companions that is most convenient for the task at hand. For a wide variety of such tasks *tree models*—transition systems based on trees rather than arbitrary graphs—are very suitable and it is not hard to obtain bisimilar companion structures that are trees for any given transition system, by a simple process of unravelling.

The unravelling of $\mathfrak{A}, a$ has as its universe all directed finite paths emanating from a in $\mathfrak{A}$; unary predicates are interpreted so as to put a path $a_0, a_1, \dots, a_\ell$, where $a_0 = a$, into P iff $a_\ell \in P^{\mathfrak{A}}$; and there is an E-edge from $a_0, a_1, \dots, a_\ell$ precisely to all $a_0, a_1, \dots, a_\ell, a_{\ell+1}$ for which $(a_\ell, a_{\ell+1}) \in E^{\mathfrak{A}}$. The natural projection that associates $a_0, a_1, \dots, a_\ell \mapsto a_\ell$ induces a bisimulation with $\mathfrak{A}, a$.

Thanks to this universal availability of bisimilar companions that are trees, any bisimulation invariant model theoretic issue about arbitrary transition systems is immediately reduced to the model theory of trees. One can hardly overestimate the scope and success of this approach. To mention but the most obvious application, we may consider satisfiability issues for logics that are preserved under bisimulation (e.g., basic modal logic ML, but also computation tree logic or the modal μ-calculus—the latter two fragments of monadic second-order logic that reach outside FO). By bisimulation invariance a formula is satisfiable if and only if it is satisfiable in a tree model: any bisimulation invariant logic enjoys the *tree model property*. Satisfiability in a tree model can be formalised as a monadic second-order fact about trees, for each of the logics mentioned. Hence, Rabin's Theorem on the decidability of the monadic second-order theory of trees immediately yields decidability results. Quite often the complexity of the satisfiability problem for bisimulation invariant logics can moreover be pinpointed with translations into emptiness problems for suitable tree-automata. As Vardi argues in [25] the tree model property (rather than the finite model property) accounts for the robust decidability of logics for the modal domain (also compare [9]).

Unravellings, however, rarely work for the purpose of finite model theory, simply because any unravelling of a transition system with directed cycles will be infinite; indeed, acyclicity and finiteness are mutually exclusive in bisimilar companions of any system with directed cycles. In section 4 we shall specifically deal with certain qualified analogues of unravellings that work in finite model theory.

1.5. Partial unravellings. One rather simplistic local analogue that can also be useful is provided by partial unravellings that unravel the given structure to a certain depth from its distinguished node, merged into isomorphic copies of the given structure. The *depth ℓ unravelling* of $\mathfrak{A}, a$ provides a structure bisimilar with $\mathfrak{A}, a$ which is tree-like inside a radius ℓ from the distinguished node. We restrict the full unravelling of $\mathfrak{A}, a$ to all nodes whose distance is at most ℓ from the root, and identify any node of the form $a = a_0, a_1, \dots, a_\ell = b$ with the node corresponding to b in a new isomorphic copy of $\mathfrak{A}, b$. Note that the resulting bisimilar companion structure is finite for finite $\mathfrak{A}$, and acyclic at least up to depth ℓ from a.

§2. The van Benthem-Rosen characterisation. The classical version of the characterisation theorem that links ML with bisimulation invariance is the following, due to van Benthem [23, 24]. We choose a formulation that highlights the harder direction of the equivalence, namely the converse of the semantic preservation property.

THEOREM 2.1 (van Benthem). $\mathrm{FO}/{\sim} \;\equiv\; \mathrm{ML}$*: any first-order formula $\varphi(x)$ that is invariant under bisimulation is equivalent to a formula of basic modal logic, and vice versa.*

Clearly this theorem can be read in two different ways. We may take it as a semantic characterisation of the modal fragment inside FO, in the spirit of the classical model theoretic preservation theorems (noting that as with those, the preservation statement is not the interesting direction of the stated equivalence). Alternatively, we may see it as a provision of effective syntax for the bisimulation invariant FO-properties. The full set of FO-formulae that are bisimulation invariant cannot be the answer, since bisimulation invariance of FO formulae is not decidable. The essence of the characterisation theorem, therefore, is that ML is a logic that is *expressively complete* for bisimulation invariant FO.

The finite model theory version of this characterisation result (cf. Theorem 2.4 below) is not an immediate consequence of the classical version, since there are first-order formulae that are bisimulation invariant over finite structures without being bisimulation invariant over all structures. Trivial examples can be generated with the use of some infinity axiom. Let for instance ψ be the first-order sentence that asserts that the binary relation R is a linear ordering without maximal element. Then any formula of the form $\psi \wedge \varphi(x)$ is trivially bisimulation invariant over finite structures, but not bisimulation invariant over infinite structures unless $\varphi(x)$ is unsatisfiable in any model of ψ.

Interestingly, many classical characterisation theorems fail in the sense of finite model theory. For instance, the classical characterisation of FO^2, seemingly a close relative to modal logics [10], as the 2-pebble game invariant fragment of first-order logic, fails in finite model theory. Indeed, the first-order sentence (in three variables) that says of a binary relation R

that it is a linear order of the universe, is invariant under 2-pebble game equivalence in restriction to finite structures (but not in general)—and it is easy to see that no first-order sentence with just two variables is equivalent to it over all finite structures. Compare also [3], and, e.g., Example 1.12 in [16]. Failures like this are only to be expected of course, since the typical tools of classical model theory and most notably the compactness theorem for FO fail in restriction to finite structures [6].

2.1. Generic classical proof. The classical proof of van Benthem's theorem makes use of compactness and saturation techniques that essentially involve infinite models. Suppose more generally that we are dealing with any equivalence relation $\leftrightharpoons$ between structures (like some unbounded Ehrenfeucht-Fraïssé game equivalence) with finitary approximations $\leftrightharpoons^\ell$ (the corresponding ℓ-round approximants). Let $\mathcal{L} \subseteq \mathrm{FO}$ be a syntactic fragment of first-order logic stratified according to $\mathcal{L} = \bigcup_\ell \mathcal{L}_\ell$, with each $\mathcal{L}_\ell$ closed under disjunction, and related to the $\leftrightharpoons^\ell$ and $\leftrightharpoons$ such that

(i) the $\leftrightharpoons^\ell$ form a chain of successive refinements $\leftrightharpoons^0 \supseteq \leftrightharpoons^1 \supseteq \cdots \supseteq \leftrightharpoons$.
(ii) each $\mathcal{L}_\ell$ is preserved under $\leftrightharpoons^\ell$.
(iii) $\leftrightharpoons^\ell$ has finite index and each $\leftrightharpoons^\ell$ class is definable by a formula of $\mathcal{L}_\ell$.

Typically, the $\mathcal{L}_\ell$ would correspond to levels of $\mathcal{L}$ defined in terms of some suitable notion of quantifier rank, and the formulae for (iii) are the Hintikka formulae associated with the Ehrenfeucht-Fraïssé analysis. In fact (ii) and (iii) imply that any $\mathcal{L}$-fomula is equivalent to a finite disjunction over such Hintikka formulae, thus providing a syntactic normal form.

For $\varphi \in \mathrm{FO}$ consider the following two statements:

$$\varphi \leftrightharpoons \text{ invariant} \iff \varphi \text{ expressible in } \mathcal{L}. \tag{2}$$

$$\varphi \leftrightharpoons \text{ invariant} \implies \varphi \leftrightharpoons^\ell \text{ invariant for some } \ell. \tag{3}$$

$\mathrm{FO}/{\leftrightharpoons} \equiv \mathcal{L}$ precisely means that (2) holds for all φ.

PROPOSITION 2.2. *Given $\leftrightharpoons$, $\leftrightharpoons^\ell$, $\mathcal{L}$ and $\mathcal{L}_\ell$ with (i)–(iii),* (2) *and* (3) *are equivalent for every $\varphi \in \mathrm{FO}$.*

PROOF. (2) $\Rightarrow$ (3) follows from (ii). For the converse implication we only need (3) to establish the crucial direction in (2), that $\leftrightharpoons$ invariance implies $\mathcal{L}$-definability. But if, by (3), the given φ is in fact $\leftrightharpoons^\ell$ invariant for some ℓ, then we may use closure under disjunction and (iii) to equivalently express φ as a disjunction over the defining $\mathcal{L}_\ell$ formulae of all those $\leftrightharpoons^\ell$-classes that satisfy φ. ⊣

Indeed, the characterisation (2) follows outright if $\leftrightharpoons$ and its approximants $\leftrightharpoons^\ell$ further satisfy the following condition (iv), which is best seen as a weak *convergence* of the sequence of the $\leftrightharpoons^\ell$ to $\leftrightharpoons$. Let $\leftrightharpoons^\omega$ stand for the common refinement $\bigcap_\ell \leftrightharpoons^\ell$ of all the finite levels $\leftrightharpoons^\ell$.

(iv) $\leftrightharpoons^\ell$ coincides with $\leftrightharpoons$ in restriction to ω-saturated models.

This condition is naturally satisfied for equivalence relations derived from Ehrenfeucht-Fraïssé games in the indicated manner.

PROPOSITION 2.3. *Let* $\simeq$, $\simeq^\ell$, *and* $\mathcal{L} = \bigcup_\ell \mathcal{L}_\ell$ *be as above with (i)–(iv). Then* $\mathrm{FO}/{\simeq} \equiv \mathcal{L}$.

PROOF. We show (3). Assuming, for the sake of contradiction, that φ were not $\simeq^\ell$ invariant for any ℓ, we obtain a sequence of models $\mathfrak{A}_\ell \models \varphi$ and $\mathfrak{B}_\ell \models \neg\varphi$ with $\mathfrak{A}_\ell \simeq^\ell \mathfrak{B}_\ell$. A simple compactness argument (or a straightforward ultraproduct construction) allows us to construct a pair of models $\mathfrak{A} \models \varphi$ and $\mathfrak{B} \models \neg\varphi$ where $\mathfrak{A} \simeq^\ell \mathfrak{B}$ for all finite ℓ. For instance one can argue for finite satisfiability of the FO-theory $\{\chi^A \leftrightarrow \chi^B : \chi \in \mathcal{L}\} \cup \{\varphi^A, \neg\varphi^B\}$, where superscripts A and B refer to relativisation to new unary predicates A and B that delineate sub-universes for the $\mathfrak{A}$- and $\mathfrak{B}$-parts, respectively. The substructures induced on the A- and B-parts of any model of this theory will provide $\mathfrak{A}$ and $\mathfrak{B}$ as desired.

Passing to ω-saturated elementary extensions $\mathfrak{A}_\infty \succcurlyeq \mathfrak{A}$ and $\mathfrak{B}_\infty \succcurlyeq \mathfrak{B}$, respectively, we have $\mathfrak{A}_\infty \models \varphi$, $\mathfrak{B}_\infty \models \neg\varphi$ even though $\mathfrak{A}_\infty \simeq \mathfrak{A}_\infty$ by (iv), contradicting $\simeq$ invariance of φ. ⊣

The equivalence of (2) and (3) remains valid when interpreted in the sense of finite model theory. Moreover, (iv) holds in particular of finite structures (which all are ω-saturated): $\simeq^\omega$ coincides with $\simeq$ in finite model theory.

Incidentally, this allows us to reformulate the finite model theory version of the characterisation $\mathcal{L} \equiv \mathrm{FO}/{\simeq}$ as an instance of compactness in finite model theory, as follows. In finite model theory, and by the above, $\simeq$ invariance of φ is captured by

$$\{\chi^A \leftrightarrow \chi^B : \chi \in \mathcal{L}\} \models \varphi^A \leftrightarrow \varphi^B.$$

Here we could moreover restrict the formulae χ to the Hintikka formulae characterising the $\simeq^\ell$ classes (see condition (iii) above), for all finite ℓ. Similarly, $\simeq^\ell$ invariance of φ, for fixed level ℓ, is captured (both classically and in finite model theory) by

$$\{\chi^A \leftrightarrow \chi^B : \chi \in \mathcal{L}_\ell\} \models \varphi^A \leftrightarrow \varphi^B.$$

Therefore, the crucial implication (2) ⇒ (3) in finite model theory becomes

$$\begin{aligned} &\{\chi^A \leftrightarrow \chi^B : \chi \in \mathcal{L}\} \models \varphi^A \leftrightarrow \varphi^B \\ \Longrightarrow\quad &\text{for some finite } \ell\text{:}\ \ \{\chi^A \leftrightarrow \chi^B : \chi \in \mathcal{L}_\ell\} \models \varphi^A \leftrightarrow \varphi^B. \end{aligned} \tag{4}$$

Compactness, of course, fails as a general principle in finite model theory. If we look at the (indirect) proof of Proposition 2.3, the structures $\mathfrak{A}$ and $\mathfrak{B}$ obtained there would typically have to be infinite, even if based on sequences of finite $\mathfrak{A}_\ell$ and $\mathfrak{B}_\ell$. And since $\simeq$ invariance can now only be assumed in restriction to finite models, it cannot be brought to bear.

The modal characterisation theorem itself, however, does go through in finite model theory, as shown by Rosen [22]. In other words, with basic modal logic of nesting depth ℓ for $\mathcal{L}_\ell$ and $\mathcal{L} = \mathrm{ML}$, (4) is a valid instance of compactness in finite model theory.

THEOREM 2.4 (Rosen). *Any first-order formula $\varphi(x)$ that is invariant under bisimulation over finite structures is equivalent over finite structures to a formula of basic modal logic, and vice versa.*

While the generic classical proof is rather smooth, it really tells us nothing about the finite model theory of the matter. The rather more constructive argument given by Rosen, however, does equally apply to the classical version, thus providing a new proof there as well. The same is true of our new proof of this theorem and of those ramifications of the van Benthem-Rosen characterisation that will concern us in section 4. Those ramifications in particular will give an exemplary insight into ways in which finite model theory techniques can provide surprising alternatives to classical arguments.

§3. The van Benthem-Rosen theorem reproved. This section provides a self-contained exposition of an elementary proof of the van Benthem-Rosen characterisation, based solely on playing Ehrenfeucht-Fraïssé games. An extended version of the argument can be found in [18].

DEFINITION 3.1. Let $\mathfrak{A}$ be a relational structure. The *Gaifman graph* $G(\mathfrak{A})$ associated with $\mathfrak{A}$ has universe A and edges (a, a') for any pair of distinct elements that coexist in some relational ground atom of $\mathfrak{A}$. *Gaifman distance* d on $\mathfrak{A}$ is the metric induced by ordinary graph distance in $G(\mathfrak{A})$: $d(a, a') = \ell$ if ℓ is the minimal length of a path between a and a' in $G(\mathfrak{A})$.

DEFINITION 3.2. Let $\mathfrak{A}$ be a relational structure. The *ℓ-neighbourhood* of a in $\mathfrak{A}$ is the subset $U^\ell(a) = \{a' \in A \colon d(a, a') < \ell\}$.

DEFINITION 3.3. A formula $\psi(x)$ is *ℓ-local* if it is logically equivalent to its relativisation to $U^\ell(x)$.

Gaifman distance d is first-order definable, in the sense that for every ℓ there is a first-order formula saying that $d(x, y) \leqslant \ell$. For a vocabulary consisting of just unary and binary predicates in particular, as is the case for transition systems, $d(x, y) \leqslant 1$ is just a disjunction over $x = y$ and all $Exy \vee Eyx$ for binary E; inductively, $d(x, y) \leqslant \ell$ is expressible as $\exists z \big(d(x, z) \leqslant \ell_1 \wedge d(z, y) \leqslant \ell_2\big)$ for any $\ell = \ell_1 + \ell_2$, $\ell_i \geqslant 1$. Generally, Gaifman neighbourhoods are FO-definable, and ℓ-locality has an obvious syntactic counterpart in FO. For a fuller account of decompositions of FO formulae into local formulae see section 5.1, where Gaifman's Theorem (Theorem 5.3) will feature prominently.

DEFINITION 3.4. For two structures $\mathfrak{A}$ and $\mathfrak{B}$ we say that a in $\mathfrak{A}$ and b in $\mathfrak{B}$ are *ℓ-locally q-equivalent* if $\mathfrak{A}{\restriction}U^\ell(a), a$ and $\mathfrak{B}{\restriction}U^\ell(b), b$ are q-elementarily equivalent, i.e., cannot be distinguished by first-order formulae of quantifier rank q. Symbolically we write

$$\mathfrak{A}, a \equiv_q^{(\ell)} \mathfrak{B}, b.$$

Of course $\mathfrak{A}, a \equiv_q^{(\ell)} \mathfrak{B}, b$ may also be characterised by saying that $\mathfrak{A}, a$ and $\mathfrak{B}, b$ agree on all ℓ-local FO-formulae of quantifier rank up to q.

LEMMA 3.5. *Let $\varphi(x) \in$ FO be invariant under disjoint sums. Then φ is ℓ-local for $\ell = 2^{\mathrm{qr}(\varphi)}$.*

PROOF. Let $q := \mathrm{qr}(\varphi)$. It suffices to show that for $\ell := 2^q$

$$\mathfrak{A}, a \equiv_q^{(\ell)} \mathfrak{B}, b \quad \Longrightarrow \quad \Big(\mathfrak{A} \models \varphi[a] \iff \mathfrak{B} \models \varphi[b] \Big).$$

Let $\mathfrak{A}, a \equiv_q^{(\ell)} \mathfrak{B}, b$. As φ is invariant under disjoint sums the desired preservation of φ follows from the following:

$$q \cdot \mathfrak{A} + q \cdot \mathfrak{B} + \mathfrak{A} \models \varphi[a] \quad \iff \quad q \cdot \mathfrak{A} + q \cdot \mathfrak{B} + \mathfrak{B} \models \varphi[b] \tag{5}$$

Here $q \cdot \mathfrak{A} + q \cdot \mathfrak{B}$ stands for the disjoint sum of q copies of $\mathfrak{A}$ and $\mathfrak{B}$ each, also disjoint form the original copies of $\mathfrak{A}$ and $\mathfrak{B}$ in which we take a, respectively b, to live. Let $\mathfrak{C} := q \cdot \mathfrak{A} + q \cdot \mathfrak{B}$ denote the common part of the two structures.

We establish (5) by exhibiting a strategy for the second player in the q-round Ehrenfeucht-Fraïssé game on $\mathfrak{C} + \mathfrak{A}, a$ versus $\mathfrak{C} + \mathfrak{B}, b$.

After k moves in the game, positions are of the form $\mathfrak{C} + \mathfrak{A}, a, a_1, \dots, a_k$ versus $\mathfrak{C} + \mathfrak{B}, b, b_1, \dots, b_k$. Some of the moves will have been played according to a strategy in the game on $\mathfrak{A}{\restriction}U^\ell(a), a$ versus $\mathfrak{B}{\restriction}U^\ell(b), b$, which is available as $\mathfrak{A}, a \equiv_q^\ell \mathfrak{B}, b$. Some other moves will have been made according to an isomorphic copying strategy into appropriate isomorphic partner structures in $\mathfrak{C}$. Let us call moves of the first kind (played according to $\mathfrak{A}, a \restriction U^\ell(a)$ versus $\mathfrak{B}, b \restriction U^\ell(b)$ in the original copies of $\mathfrak{A}$ and $\mathfrak{B}$) *close*, and also refer to the elements a_i and b_i marked in those rounds as *close*. Note that in this terminology a_i is close iff b_i is. We let $(a_i)_{a_i\,\mathrm{close}}$ and $(b_i)_{b_i\,\mathrm{close}}$, respectively, stand for the subtuples consisting of just those elements that were played 'close' in this sense, in either structure. A critical distance related to 'closeness' in round k of the game will be $\ell_k := 2^{q-k}$; note that $\ell_0 = \ell$. Our strategy will be such that the following conditions are maintained. After the k-th round:

(i) all close elements have been played within distance $\ell - \ell_k$ of a or b, respectively,

(ii) $\mathfrak{A}{\restriction}U^\ell(a), a, (a_i)_{a_i\,\mathrm{close}} \equiv_{q-k}^{(\ell)} \mathfrak{B}{\restriction}U^\ell(b), b, (b_i)_{b_i\,\mathrm{close}}$,

(iii) if a_i and b_i are not close, then their distance from a (respectively b) or any other close element is greater than ℓ_k,
(iv) if a_i and b_i are not close then there is an isomorphism ρ_i between the copy of $\mathfrak{A}$ or $\mathfrak{B}$ in which a_i lives and the copy in which b_i lives, such that $\rho_i(a_i) = b_i$,
(v) if a_i and a_j (b_i and b_j) are both not close, and $d(a_i, a_j) \leqslant \ell_k$ (or $d(b_i, b_j) \leqslant \ell_k$) then $\rho_i = \rho_j$.

It is clear that (i)–(v) are satisfied at the start of the game, for $k = 0$. We show how the second player can maintain these conditions in response to any next move by the first player. If these conditions are maintained through all q rounds, then the second player wins: note in particular that the non-close elements are still at a distance greater than $\ell_q = 1$ from the close ones, and hence have no relational edges in common with those, while amongst the close or non-close elements the validity of the strategy inside $\mathfrak{A}, a{\restriction}U^\ell(a)$ and $\mathfrak{B}, b{\restriction}U^\ell(b)$, or the existence of isomorphisms according to (iv) and (v) guarantee that $\{(a, b), (a_1, b_1), \dots (a_q, b_q)\}$ is a local isomorphism.

So, how can (i)–(v) be maintained? W.l.o.g. assume that the first player makes his $(k + 1)$st move into the a-structure, choosing a_{k+1}. In order to advise the second player on her response we distinguish several cases.[3]

Case (A): a_{k+1} is at distance greater than ℓ_{k+1} from a and all previously played elements; a_{k+1} and b_{k+1} are not going to be close. Consider the copy of $\mathfrak{A}$ or $\mathfrak{B}$ in which a_{k+1} lives in the a-structure. As there are q copies of $\mathfrak{A}$ and $\mathfrak{B}$ in $\mathfrak{C}$, at least one hitherto unused copy of the respective kind is still available in the b-structure. Pick an isomorphism onto such a fresh copy and pick b_{k+1} to be the image of a_{k+1} under this isomorphism. (i)–(v) are clearly maintained.

Case (B): a_{k+1} is within distance ℓ_{k+1} of some previously played element that is not close. Again, a_{k+1} and b_{k+1} are not going to be close, but b_{k+1} needs to be chosen according to requirement (v). Let b_{k+1} be the image of a_{k+1} under the isomorphism that works for all the (non-close) previously played elements within distance ℓ_{k+1} of a_{k+1}. (i)–(v) are clearly maintained.

Case (C): a_{k+1} is within distance ℓ_{k+1} of some previously played close element; a_{k+1} and b_{k+1} are going to be close. Note that condition (i) will be satisfied for a_{k+1}: $d(a, a_{k+1}) \leqslant \ell - \ell_k + \ell_{k+1} = \ell - \ell_{k+1}$. b_{k+1} can be selected according to the strategy in $\mathfrak{A}{\restriction}U^\ell(a), a; (a_i)_{a_i \text{ close}} \equiv^{(\ell)}_{q-k} \mathfrak{B}{\restriction} U^\ell(b), b; (b_i)_{b_i \text{ close}}$. It remains to argue that condition (i) also applies to this b_{k+1}. But the game for $\mathfrak{A}{\restriction}U^\ell(a), a; (a_i)_{a_i \text{ close}} \equiv^{(\ell)}_{q-k} \mathfrak{B}{\restriction}U^\ell(b), b; (b_i)_{b_i \text{ close}}$ automatically preserves distances up to ℓ_{k+1} in its first round, as $d(x, y) \leqslant$

[3]The three cases are mutually exclusive; for (B) and (C) this comes from (iii).

2^m is expressible in quantifier rank m. Again, (i)–(v) are seen to be maintained. ⊣

Corollary 3.6. *Let $\varphi(x) \in$ FO be invariant under bisimulation (over finite structures). Then φ is equivalent (over finite structures) to an* ML*-formula of nesting depth less than $2^{\mathrm{qr}(\varphi)}$. This bound is tight:* FO *can be exponentially more succinct than* ML *in expressing bisimulation invariant properties.*

Proof. Note that bisimulation invariance in particular implies invariance under disjoint sums. So $\varphi(x)$ is ℓ-local for $\ell = 2^{\mathrm{qr}(\varphi)}$, by the previous lemma. We now want to show that

$$\mathfrak{A}, a \sim^{\ell-1} \mathfrak{B}, b \quad \Longrightarrow \quad \Big(\mathfrak{A} \models \varphi[a] \iff \mathfrak{B} \models \varphi[b] \Big).$$

By bisimulation invariance (in finite structures) we may w.l.o.g. assume that $\mathfrak{A}{\restriction}U^\ell(a)$ and $\mathfrak{B}{\restriction}U^\ell(b)$ are trees: if they are not we pass to depth ℓ unravellings [in the classical case we may of course fully unravel even if the resulting structures are infinite]. With this proviso observe that then

$$\begin{array}{rrcl} & \mathfrak{A}, a & \sim^{\ell-1} & \mathfrak{B}, b \\ \Leftrightarrow & \mathfrak{A}{\restriction}U^\ell(a), a & \sim^{\ell-1} & \mathfrak{B}{\restriction}U^\ell(b), b \\ \Rightarrow & \mathfrak{A}{\restriction}U^\ell(a), a & \sim & \mathfrak{B}{\restriction}U^\ell(b), b \\ \Rightarrow & \mathfrak{A}{\restriction}U^\ell(a) \models \varphi[a] & \Leftrightarrow & \mathfrak{B}{\restriction}U^\ell(b) \models \varphi[b] \\ \Rightarrow & \mathfrak{A} \models \varphi[a] & \Leftrightarrow & \mathfrak{B} \models \varphi[b] \end{array}$$

The equivalence between the first two lines, as well the implication to the third are due to the fact that the ℓ-neighbourhoods encompass all elements that are reachable from the source node in $\ell - 1$ rounds of the bisimulation game. The next two implications use first bisimulation invariance, then ℓ-locality of φ.

For the tightness claim it suffices to consider the property

$$\mathrm{Red}_q(x) : \begin{cases} \text{a red element is reachable from } x \\ \text{on an } E\text{-path of length less than } 2^q \end{cases}$$

in structures with binary E and unary R for *red*. Observe that $\mathrm{Red}_q(x)$ is expressible by an FO-formula $\varphi_q(x)$ of quantifier rank q. As auxiliary formulae we use variants of the distance formulae considered in connection with Definition 3.1 above, but now for distances along (directed) E-paths: let $\delta(x,y) \leqslant 2^q$ be shorthand for the quantifier rank q formulae generated from $\delta(x,y) \leqslant 2^0 := x = y \vee Exy$ according to $\delta(x,y) \leqslant 2^{q+1} := \exists z \big(\delta(x,z) \leqslant 2^q \wedge \delta(z,y) \leqslant 2^q\big)$. Now let $\varphi_0(x) = \mathrm{Red}(x)$ and inductively $\varphi_{q+1}(x) := \exists y \big(d(x,y) \leqslant 2^q \wedge \varphi_q(y)\big)$. Clearly $\varphi_q(x)$ is invariant under bisimulation (in fact under $\sim^\ell$ for $\ell = 2^q - 1$), but not invariant under $\sim^\ell$ for any $\ell < 2^q - 1$, hence not expressible in ML at nesting depth less than $2^q - 1$. ⊣

§4. Locally acyclic covers. We look at stronger variants of bisimulation equivalence and in particular at global two-way bisimulations. A bisimulation between $\mathfrak{A}, a$ and $\mathfrak{B}, b$ is called a *two-way bisimulation* if the back-and-forth conditions are also satisfied w.r.t. to the inverse edge relations (backward transitions) $E^{-1} = \{(y,x) \colon (x,y) \in E\}$. A bisimulation is *global* if it pairs every node in $\mathfrak{A}$ with some node in $\mathfrak{B}$, and vice versa.

DEFINITION 4.1. Two-way global bisimulation equivalence, $\mathfrak{A}, a \approx \mathfrak{B}, b$, is defined like bisimulation equivalence, but on the basis of back-and-forth conditions w.r.t. both forward and backward moves along E-edges, and with the additional requirement that for every a' in $\mathfrak{A}$ there is some b' from $\mathfrak{B}$ such that the second player has a winning strategy in the game from $\mathfrak{A}, a'$ and $\mathfrak{B}, b'$, and vice versa.

Two-way global ℓ-bisimulation, $\approx^\ell$, is the corresponding approximation, defined in terms of a winning strategy in the ℓ-round game, or in terms of a stratified back-and-forth system of depth ℓ.

Note that $\approx$ and $\approx^\ell$ are meaningful also without distinguished nodes, as in $\mathfrak{A} \approx \mathfrak{B}$, due to the global nature of these bisimulations.

The passage from ordinary bisimulation to two-way and/or global bisimulation corresponds to the introduction of moves in which markers are moved backwards along edges, and moves in which markers can be freely relocated to any element of their structure. In Ehrenfeucht-Fraïssé terms these extra moves capture the power of inverse and universal modalities, respectively. In first-order terms, unconstrained moves correspond to universal and existential quantification, $\forall x \varphi(x)$ and $\exists x \varphi(x)$, which in the classical modal context may be associated with modal operators $[U]$ and $\langle U \rangle$ where $U = A \times A$ is the full binary relation on $\mathfrak{A}$. Inverse (backward) modalities take the form of

$$\begin{aligned}\big(\langle E^{-1}\rangle\varphi\big)(x) &\equiv \exists y\big(Eyx \wedge \varphi(y)\big),\\ \big([E^{-1}]\varphi\big)(x) &\equiv \forall y\big(Eyx \rightarrow \varphi(y)\big).\end{aligned}$$

DEFINITION 4.2. We denote as $\mathrm{ML}^{-\forall}$ the extension of basic modal logic by global and inverse modalities.

A characterisation theorem analogous to van Benthem's obtains. Indeed, we shall see that this also can be proved in ways that show the same characterisation to be valid in finite model theory: both classically and in finite model theory, the $\approx$ invariant fragment of first-order logic is precisely captured by $\mathrm{ML}^{-\forall}$.

As global quantification clearly allows us to jump anywhere inside the given structure, the tight localisation to a neighbourhood of the distinguished parameter (or free variable) is lost. For model constructions this means that a more global approach, which treats all nodes and their local neighbourhoods on the same footing, becomes necessary. This is where approximate, at least locally good, finite approximations to unravellings come into play. As these considerations are of some technical interest in their own right, we firstly isolate these model construction issues and then return to the proof of the above-mentioned characterisation result as an application.

DEFINITION 4.3. For an onto mapping $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ between transition systems of the same type:

π is a *cover* of $\mathfrak{A}$ by $\hat{\mathfrak{A}}$, if

—for every unary predicate P: $\hat{a} \in P^{\hat{\mathfrak{A}}} \Leftrightarrow \pi(\hat{a}) \in P^{\mathfrak{A}}$,

—for every binary relation E, π provides *lifts*: if $(a, a') \in E^{\mathfrak{A}}$ then there are E-edges in $\hat{\mathfrak{A}}$ out of every node $\hat{a} \in \pi^{-1}(a)$ and into every node $\hat{a}' \in \pi^{-1}(a')$.

π is a *faithful cover* if all these lifts are unique.

π is a *bisimilar cover* if its graph $\{(\hat{a}, a)\colon \pi(\hat{a}) = a\}$ is a global two-way bisimulation between $\hat{\mathfrak{A}}$ and $\mathfrak{A}$.

Note that covers need not be homomorphisms. Unlike bisimilar companions, covers merely provide a kind of one-sided simulation, satisfying the back-requirements, but not necessarily the forth-requirements of bisimulations. Bisimilar covers, however, are homomorphisms by virtue of the forth-requirements. Even bisimilar covers need not be faithful. Faithful covers always allow unique lifts of (undirected) paths to any node in the fibre above any given node in the path.

Recall the bisimilar unravellings discussed in section 1.4. Consider the unravelling of $\mathfrak{A}, a$, based on the set of all paths from a in $\mathfrak{A}$; let π be the natural projection from this unravelling to $\mathfrak{A}$, which maps a path to the last vertex on that path. This projection induces a bisimulation, though typically not a two-way bisimulation, nor a global one unless every node of $\mathfrak{A}$ is reachable from a. Instead of the ordinary (directed) unravelling, however, we may consider a two-way unravelling based on all *undirected* paths, in which edges may be traversed backwards as well as forwards. The corresponding projection π induces a two-way bisimulation. Taking the disjoint union of such two-way unravellings from at least one element from each connected component of $\mathfrak{A}$, π becomes a bisimilar cover; and even a faithful bisimilar cover if we only admit paths that do not traverse the same edge in opposite directions in consecutive steps. In this sense, two-way unravellings generally can provide faithful bisimilar covers, albeit typically infinite ones.

A fuller justification for the following proviso is given, with explicit encoding prescriptions, in [17, 19]. The essential steps of the constructions discussed below remain the same but become somewhat more transparent in the restricted setting. The assumption laid out in the proviso can essentially be made without loss of generality, as the construction of bisimilar covers is compatible with suitable encodings of arbitray edge relations in structures that do conform to the proviso.

PROVISO 4.4. *For the remainder of this section we assume that there is just one single, strictly asymmetric edge relation.*

We are interested in bisimilar covers of finite transition systems by finite transition systems that are at least locally acyclic, i.e., do not have short (undirected) cycles. An *undirected cycle* of length $k \geqslant 3$ in $\mathfrak{A}$ is a sequence of nodes $a_0, a_1, \dots, a_{k-1}$ (cyclically indexed by $\mathbb{Z}_k$) such that for every i either (a_i, a_{i+1}) or (a_{i+1}, a_i) is an edge and $a_i \neq a_{i+2}$.

DEFINITION 4.5. (i) The *girth* of a vertex a in a transition system $\mathfrak{A}$, $\mathrm{girth}(a, \mathfrak{A})$, is the minimal length of an undirected cycle through a.
(ii) The *girth* of a transition system $\mathfrak{A}$, $\mathrm{girth}(\mathfrak{A})$, is the minimal length of any cycle in $\mathfrak{A}$.
(iii) $\mathfrak{A}$ is *k-acyclic* if its girth is at least k.

Note that k-acyclicity is related to *local acyclicity* in the sense that the restrictions of the Gaifman graph of $\mathfrak{A}$ to ℓ-neighbourhoods of any $a \in \mathfrak{A}$ will be acyclic if $\mathfrak{A}$ is k-acyclic for some $k \geqslant 2\ell$: ℓ-locally one will not see any (undirected) cycles. It is necessary for our purposes to consider undirected cycles, precisely because we shall need local acyclicity in the Gaifman graph, which corresponds to the symmetrised version of the transition relations. It may be worth pointing out that short directed cycles are much more easily avoided. In fact one would merely have to use a truncated partial unravelling to sufficient depth and introduce new edges from the frontier of the unravelled part to suitable nodes closer to the root. The point of the rather more elaborate proofs required for the following theorem is that we want to avoid even short undirected cycles, or short cycles in the Gaifman graph of the desired companion structure. Both of the constructions given below supersede the construction from [17], which is hyper-exponential in nature.

THEOREM 4.6. *For any fixed $k \geqslant 3$: Any finite transition system $\mathfrak{A}$ possesses a faithful bisimilar cover $\pi \colon \hat{\mathfrak{A}} \to \mathfrak{A}$ by a finite k-acyclic transition system $\hat{\mathfrak{A}}$. For fixed k, the size of $\hat{\mathfrak{A}}$ can be polynomially bounded in terms of the size of $\mathfrak{A}$.*

It is clear that the size of any cover as stated in the theorem has to be exponential in k: if $\ell \leqslant k/2$ then the ℓ-neighbourhood of any $\hat{a}$ in $\hat{\mathfrak{A}}$ is isomorphic with the ℓ-neighbourhood of $\pi(a)$ in a two-way unravelling of $\mathfrak{A}$ from $\pi(a)$ (since both neighbourhoods are acyclic and faithful bisimilar covers of each other); the size of the latter is in general exponential in ℓ.

4.1. A game oriented construction. We reduce the task of finding a bisimilar cover $\hat{\mathfrak{A}}$ to that of finding faithful covers $\mathfrak{A}^*$ that are suitable locally in the vicinity of individual nodes. This construction is inspired by the powerset construction in the determinisation of non-deterministic automata. In the associated bisimilar cover $\hat{\mathfrak{A}}$ we keep track of all possible lifts of paths to the faithful cover $\mathfrak{A}^*$ to make sure that a path in $\hat{\mathfrak{A}}$ cannot cycle back as long as at least one of the lifts into $\mathfrak{A}^*$ does not. In other words we need only make sure that at least one 'leaf' of the faithful cover $\mathfrak{A}^*$ above every node is locally acyclic. This construction is exponential and does not yield a polynomial size cover; the alternative group theoretic construction indicated below, however, will yield a polynomial size cover.

Let $\rho\colon \mathfrak{A}^* \to \mathfrak{A}$ be a faithful cover. It is easy to see that the cardinality of the fibres $\rho^{-1}(a)$ must be constant within each connected component of $\mathfrak{A}$. Indeed, the set of all lifts to $\mathfrak{A}^*$ of a single edge (a, a') of $\mathfrak{A}$ induces a bijection between $\rho^{-1}(a)$ and $\rho^{-1}(a')$. For the rest of the construction we assume w.l.o.g. that $\mathfrak{A}$ is weakly connected. Let $\rho\colon \mathfrak{A}^* \to \mathfrak{A}$ be a faithful cover of finite multiplicity N. We obtain a bisimilar cover for $\mathfrak{A}$ as follows.

$$\hat{A} := \big\{(s\colon N \to A^*)\colon s \text{ one-to-one}, \rho \circ s \text{ constant}\big\}.$$

The conditions on s say that it is a bijection between N and some fibre $\rho^{-1}(a)$; intuitively each s can be used as an instantaneous description in tracking all *possible* positions in a one-sided simulation of $\mathfrak{A}$-behaviours in $\mathfrak{A}^*$. We let $\pi\colon \hat{A} \to A$ be the natural projection induced by ρ.

In $\hat{\mathfrak{A}}$ we now put an edge from s to s' iff there is an edge from $a := \pi(s)$ to $a' := \pi(s')$ and if

$$s' \circ s^{-1}\colon \rho^{-1}(a) \to \rho^{-1}(a')$$

is the bijection induced by the lifts of (a, a') to $\mathfrak{A}^*$. It is now easy to check that $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ is a faithful bisimilar cover.

Moreover, the girth of an element s of $\hat{\mathfrak{A}}$ above $a = \pi(s)$ can be bounded as follows. If $s_0, s_1, \dots, s_{k-1}$ forms an undirected cycle of length k at $s_0 = s$ in $\hat{\mathfrak{A}}$, then $s_0(m), s_1(m), \dots, s_{k-1}(m)$ is an undirected cycle in $\mathfrak{A}^*$, for every $m \in N$. This is because each edge $(s_i(m), s_{i+1}(m))$ (or its inverse) is a lift of the edge $(\pi(s_i), \pi(s_{i+1}))$ (or its inverse) to $\mathfrak{A}^*$. As the $s_0(m)$ list all the elements in $\pi^{-1}(a)$, the girth of s in $\hat{\mathfrak{A}}$ is at lesat as large as the girth of any of the elements of $\pi^{-1}(s)$ in $\mathfrak{A}^*$:

$$\text{girth}(s, \hat{\mathfrak{A}}) \geqslant \max\big\{\text{girth}(a^*, \mathfrak{A}^*)\colon \rho(a^*) = \pi(s)\big\}.$$

It remains to obtain faithful covers $\mathfrak{A}^*$ that possess for each $a \in \mathfrak{A}$ at least one covering element a^* above a whose ℓ-neighbourhood is acyclic (where $\ell = \lceil k/2 \rceil$ to make $\hat{\mathfrak{A}}$ k-acyclic). Since we are only after covers, not bisimilar covers, it suffices to cover any complete tournament of the same size as $\mathfrak{A}$ in the required manner—this cover will automatically also be good for $\mathfrak{A}$, up to re-direction of edges where required. We therefore fix $\mathfrak{A}_n$ to be a complete tournament on n vertices, $\mathfrak{A}_n = \big(n = \{0, \dots, n-1\}, E = <\big)$.

LEMMA 4.7. *For every ℓ, $\mathfrak{A}_n$ admits a finite faithful cover in which every vertex is covered by at least one vertex whose ℓ-neighbourhood is acyclic.*

PROOF. Let $\mathfrak{A} := \mathfrak{A}_n$ and fix ℓ. For every $i \in n$ let $\mathfrak{A}^{(i)}$ be the restriction of the two-way unravelling of $\mathfrak{A}$ from i to the $(\ell+1)$-neighbourhood of the root i. We let $\rho\colon \mathfrak{A}^{(i)} \to \mathfrak{A}$ be the natural projection. By the *rim* of $\mathfrak{A}^{(i)}$ we mean the set of vertices at distance ℓ from the root vertex i. The desired $\mathfrak{A}^*$ will simply be the disjoint union of the $\mathfrak{A}^{(i)}$ for $i \in n$ with additional edges joining rim vertices. This immediately guarantees that the ℓ-neighbourhood of the root i in $\mathfrak{A}^{(i)}$ is isomorphic to the ℓ-neighbourhood of i in $\mathfrak{A}^*$, and therefore acyclic. We also denote by ρ the union of the individual ρ, defined on the union of the $\mathfrak{A}^{(i)}$.

The missing edges on the rim can be filled in arbitrarily in such a way that for every $i < j$: every rim element in $\rho^{-1}(i)$ which is not already linked (inside its own $\mathfrak{A}^{(i)}$ that is) to an element in $\rho^{-1}(j)$, is joined to precisely one other rim element in $\rho^{-1}(j)$ which is not already linked to an element in $\rho^{-1}(i)$. This is possible as there are exactly the same number of rim elements of each kind across all the $\mathfrak{A}^{(i)}$, by symmetry. This turns $\rho\colon \mathfrak{A}^* \to \mathfrak{A}$ into a faithful cover. $\dashv$

4.2. A simple group theoretic construction. Instead of adherence to a (bi-)simulation game intuition we may directly construct a bisimilar cover as a bundle over $\mathfrak{A}$ for whose fibres we use an arbitrary abstract group. Using Cayley groups of large girth will produce covers of large girth. This construction has the benefit of simplicity and, using groups constructed by Margulis and Imrich, can also be kept polynomial for each fixed value of k.

Let G be any group and suppose the edge set E of $\mathfrak{A}$ is injectively embedded into a subset of G not containing both g and g^{-1} for any $g \in G$. For simplicity we identify E with a corresponding subset $E \subseteq G$. Let

$$\hat{A} := A \times G, \quad \pi\colon \hat{A} \to A \text{ the natural projection.}$$

In $\hat{\mathfrak{A}}$ we now put an edge from (a, g) to (a', g') iff $(a, a') = e \in E$ and $g' = g \cdot e$. It is obvious that $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ is a faithful bisimilar cover. The girth of $\hat{\mathfrak{A}}$ cannot be less than the girth of the Cayley graph associated with $E \subseteq G$. This is the graph whose vertices are elements of the subgroup $H \subseteq G$ generated by E in G, with an edge between g and g' if $g' = g \cdot h$ for some $h \in E \cup E^{-1}$. Again, this relationship between the girths stems from the fact that any (undirected) cycle in $\hat{\mathfrak{A}}$ induces a cycle in the Cayley graph.

So in this case it remains to provide Cayley graphs of large girth, an issue which has independently been studied in the context of algebraic methods in combinatorial graph theory [1]. Alon [1] attributes to Biggs [4] the following simple construction of n-regular Cayley graphs of exponential size for fixed girth. For our purposes, n is the number of generators and corresponds to $|E|$. Let T_n^ℓ be the complete symmetric tree graph of degree n and depth ℓ, with a node 0 at the root. Let the edges of T_n^ℓ be coloured $1, \dots, n$ such that every internal node is incident with exactly one edge of each colour. With colour i associate a permutation $i \colon T_n^\ell \to T_n^\ell$, which swaps all vertex pairs (t, t') that form an edge of colour i. Let G be the permutation group generated by these n generators. Any non degenerate composition of ℓ generators takes the root to a leaf, whence any non-degenerate generator sequence that fixes the root must have length greater than 2ℓ. In other words, the girth of the Cayley graph of G with respect to generators $1, \dots, n$ is greater than 2ℓ.

Slightly more involved explicit group theoretic constructions, due to Imrich [13] and Margulis [15], yield smaller Cayley graphs of given degree and girth. In fact, Imrich's construction realises an asymptotically near optimal dependence of the girth on degree and size. In our context, these Cayley graphs give rise to bisimilar covers whose size is polynomial in the size of the given structure, for any fixed acyclicity k.

THEOREM 4.8 (Margulis, Imrich). *For every n and ℓ there are n-regular Cayley graphs of girth greater than 2ℓ in size $\mathcal{O}(n^{c\ell})$, c a fixed constant.*

§5. Characterising $\approx$ invariant first-order logic.

THEOREM 5.1. *Both in the sense of classical and of finite model theory:* $\mathrm{FO}/{\approx} \equiv \mathrm{ML}^{-\forall}$. *I.e., for any $\varphi(x) \in \mathrm{FO}$: $\varphi(x)$ is invariant under global two-way bisimulation (over finite transition systems) if and only it is equivalently expressible in* $\mathrm{ML}^{-\forall}$.

While this characterisation result is exactly as (classically) expected, its proof illustrates several interesting features of the approach inspired by finite model theory: the more combinatorial and also rather more constructive nature of the argument; and a strategy that is exactly orthogonal to the one in the generic classical proof (as outlined in Section 2.1 above; that proof applies verbatim to the classical case of this theorem).

5.1. Gaifman locality. Recall the definition of Gaifman distance, Gaifman neighbourhoods and ℓ-local formulae from Definitions 3.1, 3.2, and 3.3.

DEFINITION 5.2. A subset of $\mathfrak{A}$ is *ℓ-scattered* if the mutual distance between any two distinct members of the set is at least 2ℓ.

A *basic local sentence*—for some locality radius ℓ, an ℓ-local formula $\psi(x)$, and some $n \geqslant 1$—is a sentence that asserts the existence of an ℓ-scattered set of size n whose elements all satisfy ψ.

The following classical result of Gaifman [8] says that every first-order formula is essentially local. See [6] for a textbook treatment.

THEOREM 5.3 (Gaifman). *Any first-order formula $\varphi(x)$ is equivalent to one that is a Boolean combination of local formulae and basic local sentences.*

Let us say that a first-order formula $\varphi(x)$ has *locality rank ℓ*, *local quantifier rank q*, and *scattering rank n*, if it has a representation according to Gaifman's theorem in which all constituent local formulae are ℓ-local, of quantifier rank q, and all its basic local sentences speak about k-scattered sets of sizes m, where $m \leqslant n$ and $k \leqslant \ell$.

Note that Lemma 3.5, in this terminology, says that any $\varphi(x)$ that is invariant under disjoint sums has locality rank 2^q, where $q = \mathrm{qr}(\varphi)$ (= local quantifier rank), and scattering rank 0 (i.e., no basic local sentences are required).

Global (two-way) bisimulation, however, does not imply invariance under disjoint sums. But it still implies invariance under disjoint copies, i.e., $q \cdot \mathfrak{A}, a \approx \mathfrak{A}, a$, where as above $q \cdot \mathfrak{A}$ stands for the q-fold sum of (isomorphic copies of) $\mathfrak{A}$. The following is the natural ramification of Gaifman's theorem for this setting. A detailed proof can be found in [17, 19]; it applies Gaifman's theorem and shows how to eliminate constituents that make assertions about non-trivial scattered sets.

PROPOSITION 5.4. *Both in the sense of classical and finite model theory: if $\varphi(x) \in \mathrm{FO}$ is invariant under disjoint copies, then φ has scattering rank 1, i.e., the only basic local sentences needed are plain existentially quantified local formulae.*

In the following paragraph we use Gaifman locality of a given $\approx$ invariant formula $\varphi(x)$ to capture an approximate, local level of elementary equivalence strong enough to preserve φ.

5.2. Upgrading ℓ-bisimulation towards elementary equivalence. Let $\dot{\equiv}$ be some equivalence relation between structures, coarser than elementary equivalence $\equiv$ and serving as an approximation to $\equiv$. We say that $\mathrel{\underline{\leftrightarrow}}^\ell$ can be *upgraded* to $\dot{\equiv}$ modulo $\mathrel{\underline{\leftrightarrow}}$, if for any (finite) $\mathfrak{A}, a \mathrel{\underline{\leftrightarrow}}^\ell \mathfrak{B}, b$ there are (finite) $\tilde{\mathfrak{A}}, a \mathrel{\underline{\leftrightarrow}} \mathfrak{A}, a$ and $\tilde{\mathfrak{B}}, b \mathrel{\underline{\leftrightarrow}} \mathfrak{B}, b$ such that $\tilde{\mathfrak{A}}, a \mathrel{\dot{\equiv}} \tilde{\mathfrak{B}}, b$:

$$
\begin{array}{ccc}
\mathfrak{A}, a & \overset{\underline{\leftrightarrow}^\ell}{\rule{3cm}{0.4pt}} & \mathfrak{B}, b \\
\Big| \underline{\leftrightarrow} & & \Big| \underline{\leftrightarrow} \\
\tilde{\mathfrak{A}}, a & \overset{\dot{\equiv}}{\rule{3cm}{0.4pt}} & \tilde{\mathfrak{B}}, b
\end{array}
$$

If $\dot{\equiv}$ is sufficiently strong to preserve $\varphi(x)$, then the diagram shows that $\mathrel{\underline{\leftrightarrow}}$ invariance of φ implies $\mathrel{\underline{\leftrightarrow}}^\ell$ invariance of φ. This is exactly as required for Theorem 5.1 according to the discussion in section 1; see in particular the crucial 'compactness' properties (3) and (4) there.

The approximation levels of $\equiv$ needed in our applications are the following. The parameters ℓ, q, and n in $\equiv^{(\ell)}_{q,n}$ are modelled after the levels of locality rank, local quantifier rank, and scattering rank of formulae in Gaifman form.

Definition 5.5. $\mathfrak{A}, a \equiv^{(\ell)}_{q,n} \mathfrak{B}, b$ if for every $k \leqslant \ell$, for every k-local formula $\psi(x)$ of quantifier rank q, and for every $m \leqslant n$:
(i) $\mathfrak{A} \models \psi[a] \Leftrightarrow \mathfrak{B} \models \psi[b]$.
(ii) $\mathfrak{A}$ has a k-scattered subset of size m for ψ iff $\mathfrak{B}$ has.

Note that $\equiv^{(\ell)}_{q}$ of Definition 3.4 is recovered in the special case of $\equiv^{(\ell)}_{q,0}$. Since our $\varphi(x)$ is $\approx$ invariant, and hence invariant under disjoint copies, we may use Proposition 5.4 to work with scattering rank 1, and in fact with $\equiv^{(\ell)}_{q,1}$ where ℓ is the locality rank, q the local quantifier rank of φ.

Lemma 5.6. *Both classically and in finite models, $\approx^{\ell}$ can be upgraded modulo $\approx$ to $\equiv^{(\ell)}_{q,1}$, for any q.*

Proof. The main step in upgrading is achieved by using ℓ-locally acyclic covers, which guarantee acyclicity of the relevant ℓ-neighbourhoods in both covers. There remains an issue of multiplicities, however: a first-order formula of quantifier rank q can distinguish multiplicities less than q. We therefore precede the acyclic cover construction by a process that boosts all multiplicities to at least q.

Let $\mathfrak{A}, a \approx^{\ell} \mathfrak{B}, b$ be given. Let $q \otimes \mathfrak{A}$ be the transition system over universe $q \times A$ with $((i, a'), (j, a'')) \in E^{q \otimes \mathfrak{A}}$ iff $(a', a'') \in E^{\mathfrak{A}}$. Clearly $q \otimes \mathfrak{A} \approx \mathfrak{A}$. Moreover, all in-degrees and out-degrees in $q \otimes \mathfrak{A}$ are multiples of q, and this carries over to any faithful ℓ-acyclic bisimilar cover $\tilde{\mathfrak{A}}$ of $q \otimes \mathfrak{A}$, $\tilde{\pi} \colon \tilde{\mathfrak{A}} \to q \otimes \mathfrak{A}$. Let $\pi \colon \tilde{\mathfrak{A}} \to \mathfrak{A}$ be the composition of $\tilde{\pi}$ with the natural projection from $q \otimes \mathfrak{A}$ to $\mathfrak{A}$; we identify a with some element from $\pi^{-1}(a)$ and clearly have $\mathfrak{A}, a \approx \tilde{\mathfrak{A}}, a$. If $\tilde{\mathfrak{B}}$ is similarly obtained from $\mathfrak{B}$, we claim that $\tilde{\mathfrak{A}}, a \equiv^{(\ell)}_{q,1} \tilde{\mathfrak{B}}, b$. Indeed, it suffices to show that $\tilde{\mathfrak{A}}, a' \approx^{\ell} \tilde{\mathfrak{B}}, b'$ implies that $\tilde{\mathfrak{A}} \restriction U^{\ell}(a'), a' \equiv_q \tilde{\mathfrak{B}} \restriction U^{\ell}(b'), b'$. But this follows with a simple Ehrenfeucht-Fraïssé argument. $\tilde{\mathfrak{A}} \restriction U^{\ell}(a')$ and $\tilde{\mathfrak{B}} \restriction U^{\ell}(b')$ are acyclic with multiplicities greater than q, and the second player will always find suitable branches (new where needed) to respond for q rounds. $\dashv$

We remark that explicit appeal to Proposition 5.4 can be avoided in the proof of Theorem 5.1, in favour of a stronger upgrading as follows.

Lemma 5.7. *For any n, $\approx^{\ell}$ can be upgraded to $\equiv^{(\ell)}_{q,n}$ modulo $\approx$, classically and in finite models.*

Proof. Let $\mathfrak{A}, a \approx^{\ell} \mathfrak{B}, b$. With Lemma 5.6 we find (finite) $\hat{\mathfrak{A}} \approx \mathfrak{A}$ and $\hat{\mathfrak{B}} \approx \mathfrak{B}$ such that $\hat{\mathfrak{A}}, a \equiv_{q,1}^{(\ell)} \hat{\mathfrak{B}}, b$. It follows that $\tilde{\mathfrak{A}}, a \equiv_{q,n}^{(\ell)} \tilde{\mathfrak{B}}, b$, if we let $\tilde{\mathfrak{A}} := n \cdot \hat{\mathfrak{A}}$ and $\tilde{\mathfrak{B}} := n \cdot \hat{\mathfrak{B}}$. Clearly still $\tilde{\mathfrak{A}}, a \approx \mathfrak{A}, a$ and $\tilde{\mathfrak{B}}, b \approx \mathfrak{B}, b$. ⊣

It is interesting to observe how the constructive proof of our characterisation theorem lined out above is truly orthogonal to the generic classical proof as sketched in Section 2.1. Both proofs establish that for any first-order $\varphi(x)$:

$$\varphi \approx \text{invariant} \implies \varphi \approx^{\ell} \text{invariant for some } \ell,$$

cf. (3) in Section 2. The classical proof would achieve this (indirectly) through a compactness argument that upgrades the sequence of the $\approx^{\ell}$, first to its limit $\approx^{\omega}$ and then to $\approx$, and preserves φ in the process on the basis of its first-order nature. The alternative proof which also works in finite model theory, on the other hand, upgrades $\approx^{\ell}$ (where ℓ is the locality rank of φ) to an approximation of elementary equivalence that is sufficient to preserve φ, and preserves φ in the process because the entire construction fully respects $\approx$.

§6. Outlook and further ramifications. The proof technique illustrated above for the characterisation of the first-order logic of properties invariant under two-way global bisimulation invariance, can be adapted to work for other related characterisation results.

Almost as it stands, it can be used to infer a similar characterisation of properties of (finite) transition systems invariant under *guarded bisimulation*. As shown in [17], the classical characterisation theorem for the *guarded fragment*, due to Andréka, van Benthem and Németi [2], is also valid over finite transition systems. It remains open, however, whether a similar finite model theory characterisation obtains also in the setting of general relational structures with predicates of arities greater than 2.

In the modal setting, some modifications are necessary in order to obtain a similar characterisation for the interesting class of properties that are invariant under global (but forward) bisimulation. As one might expect, these are captured precisely by the extension of basic modal logic by a universal modality. The proof however has to take into account the apparent mismatch between the inherently symmetric notion of locality in Gaifman's Theorem and the forward direction of the notion of bisimulation. This extension is presented in [19].

Another natural class of transition systems is that of (possibly just finite) connected structures $\mathfrak{A}, a$, i.e., those in which every node is reachable from the source node a. In restriction to such structures, the distinction between global and ordinary bisimulation equivalence is lost. Correspondingly, even the extension of ML by global quantification is preserved under ordinary bisimulation over connected structures. A number of related characterisation theorems—also w.r.t. other natural frame conditions, and in particular transitivity—are explored in [5].

The construction of certain kinds of finite bisimilar covers has also been successfully carried out in the general guarded scenario, where at least one aspect of the usually infinite guarded unravellings of relational structured is recovered in a finite model construction. It remains open how much acyclicity (in the hypergraph theoretical rather than the graph theoretical sense) can be achieved in finite guarded covers. The construction given in [12] breaks up cliques in the Gaifman graph exactly as an unravelling would. See [12] also for interesting applications of that construction not just to the finite model theory of guarded logics, but also to issues of classical model theory to do with extension properties for partial isomorphisms in finite structures.

REFERENCES

[1] N. Alon, *Tools from higher algebra*, ***Handbook of combinatorics*** (R. Graham et al., editor), vol. II, North-Holland, 1995, pp. 1749–1783.

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

[3] J. Barwise and J. van Benthem, *Interpolation, preservation, and pebble games*, ***The Journal of Symbolic Logic***, vol. 64 (1999), pp. 881–903.

[4] N. Biggs, *Cubic graphs with large girth*, ***Annals of the New York Academy of Sciences*** (G. Blum et al., editor), vol. 555, 1989, pp. 56–62.

[5] A. Dawar and M. Otto, *Modal characterisation theorems over special classes of frames*, preprint, 2005.

[6] H.-D. Ebbinghaus and J. Flum, ***Finite model theory***, 2nd ed., Springer, 1999.

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

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

[9] E. Grädel, *Why are modal logics so robustly decidable?*, ***Bulletin of the European Association for Theoretical Computer Science***, vol. 68 (1999), pp. 90–103.

[10] E. Grädel and M. Otto, *On Logics with Two Variables*, ***Theoretical Computer Science***, vol. 224 (1999), pp. 73–113.

[11] M. Hennessy and R. Milner, *Algebraic laws of indeterminism and concurrency*, ***Journal of the ACM***, vol. 32 (1985), pp. 137–162.

[12] I. Hodkinson and M. Otto, *Finite conformal hypergraph covers and Gaifman cliques in finite structures*, submitted, 2002.

[13] W. Imrich, *Explicit construction of regular graphs without small cycles*, ***Combinatorica***, vol. 4 (1984), pp. 53–59.

[14] D. Janin and I. Walukiewicz, *On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic*, ***Proceedings of 7th International Conference on Concurrency Theory CONCUR '96***, Lecture Notes in Computer Science, no. 1119, Springer-Verlag, 1996, pp. 263–277.

[15] G. Margulis, *Graphs without short cycles*, ***Combinatorica***, vol. 2 (1982), pp. 71–78.

[16] M. Otto, ***Bounded variable logics and counting***, Lecture Notes in Logic, vol. 9, Springer, 1997.

[17] ———, *Modal and guarded characterisation theorems over finite transition systems*, ***Proceedings of 17th Annual IEEE Symposium on Logic in Computer Science LICS '02***, 2002, pp. 371–380.

[18] ———, *An elementary proof of the van Benthem–Rosen characterisation theorem*, TUD online preprint no. 2342, 2004.

[19] ———, *Modal and guarded characterisation theorems over finite transition systems*, ***Annals of Pure and Applied Logic***, vol. 130 (2004), pp. 173–205, extended journal version of [17].

[20] D. Park, *Concurrency and automata on infinte sequences*, ***Proceedings of 5th GI Conference***, Springer-Verlag, 1981, pp. 176–183.

[21] B. Poizat, ***A course in model theory***, Springer-Verlag, 2000.

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

[23] J. van Benthem, *Modal correspondence theory*, ***Ph.D. thesis***, University of Amsterdam, 1976.

[24] ———, ***Modal logic and classical logic***, Bibliopolis, Napoli, 1983.

[25] M. Vardi, *Why is modal logic so robustly decidable?*, ***Descriptive complexity and finite models*** (N. Immerman and P. Kolaitis, editors), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 31, AMS, 1997, pp. 149–184.

FACHBEREICH MATHEMATIK
TECHNISCHE UNIVERSITÄT DARMSTADT
64289 DARMSTADT, GERMANY
E-mail: otto@mathematik.tu-darmstadt.de
URL: www.mathematik.tu-darmstadt.de/~otto