A Lindström characterisation of the guarded fragment and of modal logic with a global modality [1]

Martin Otto and Robert Piro

ABSTRACT. We establish a Lindström type characterisation of the extension of basic modal logic by a global modality (ML[∀]) and of the guarded fragment of first-order logic (GF) as maximal among compact logics with the corresponding bisimulation invariance and the Tarski Union Property.

Keywords: Lindström theorems, modal logic, bisimulation invariance, global modality, guarded fragment, Tarski union property.

1 Introduction

This investigation is motivated by a recent Lindström theorem for basic modal logic (ML) by van Benthem [4] and related investigations in [5]. It is shown in [4] that no logic that is compact, bisimulation invariant and has the relativisation property can properly extend ML. This characterisation itself may be seen as a methodological improvement on an earlier Lindström characterisation of ML by de Rijke [13], which explicitly stipulated a finite depth (or locality) condition as a crucial criterion. [A formula φ (over pointed Kripke structures, say) is called r-local if whether or not a pointed τ-structure $(\mathfrak{M}, w)$ satisfies φ only depends on the substructure induced on the r-neighbourhood of w (the set of elements accessible from w in at most r steps).]

The proof of van Benthem's characterisation in [4] does not carry over to the interesting case of the guarded fragment GF, indeed not even to the extension of basic modal logic by a global (or universal) modality ML[∀]. Crucially, the finite depth criterion is still instrumental in that proof, though instead of being stipulated as a condition it is shown to be a consequence of the combination of compactness and relativisation for any logic invariant under ordinary bisimulation. But locality, or the finite depth criterion, fail for GF and even for ML[∀]. Neither global nor guarded bisimulation invariance implies locality. We therefore switch to an alternative characterisation crucially based on the Tarski Union Property (TUP), which is another natural model theoretic criterion that has been studied in abstract model theory [2]. Just as a variant characterisation of FO can be based on compactness,

[1] This paper summarises results from the second author's diploma thesis [12], which was supervised by the first author.

Advances in Modal Logic, Volume 7.

TUP and invariance under partial isomorphy, we here characterise ML[∀] and GF as maximally expressive among compact logics that are invariant under the appropriate notion of bisimulation (global or guarded bisimulation, respectively) and satisfy TUP. Some discussion of the role of TUP can be found in the concluding section 4. An analogous characterisation of basic modal logic ML itself is of course also available.

Part of the point of such investigations, as expounded in particular in [5], is a new interest in the abstract model theory of logics well below first-order logic. Many of the techniques and constructions that are available in the more classical investigations into abstract model theory, which is aimed at levels above first-order, [2], are no longer available or meaningful for corresponding investigations at levels below FO. Much of the usual coding machinery relies on first-order interpretations of, for instance, embedded substructures or systems of partial isomorphism, etc., which are not generally available at the level of logics of a typically modal character.

We point out that it remains open whether for instance ML[∀] is also maximal in the class of compact logics with the relativisation property that are invariant under global bisimulation.

In the following we presuppose some familiarity with basic model theoretic notions from modal logic (syntax and semantics of basic modal logic, Kripke structures, bisimulation relations and the basic bisimulation game, etc.) as presented in various textbooks and, for instance, in [6]. Corresponding variations for ML[∀] and GF will be reviewed where they arise.

2 Characterisation of ML[∀]

We summarise some standard notions that are important throughout the paper and for which the reader may also want to compare the classical setting for abstract model theory in [2]. A *logic* $\mathcal{L}$ is a pair $(L, \models_{\mathcal{L}})$, where L is a function that maps signatures σ to the sets $L(\sigma)$ of $\mathcal{L}$-formulae over σ. $\models_{\mathcal{L}}$ is a relation between structures and formulae of $\mathcal{L}$. We tacitly assume that structures are of an appropriate type also w.r.t. accommodating any 'free variables' as appropriate for $\mathcal{L}$: speaking of 'structures' we allow structures with parameters, like pointed structures; this latitude is included without explicit mention in (1) below. Any logic $\mathcal{L}$ is assumed to satisfy the following:

1. If $\mathfrak{M} \models_{\mathcal{L}} \varphi$, then $\mathfrak{M}$ is a σ-structure such that $\varphi \in L(\sigma)$.
2. If $\sigma \subseteq \sigma'$, then $L(\sigma) \subseteq L(\sigma')$.
3. If $\sigma \subseteq \sigma'$, $\varphi \in L(\sigma)$ and $\mathfrak{M}$ a σ'-structure, then $\mathfrak{M}{\restriction}\sigma \models \varphi$ iff $\mathfrak{M} \models \varphi$.
4. If $\mathfrak{M}$ is isomorphic to $\mathfrak{N}$, then $\mathfrak{M} \models \varphi$ iff $\mathfrak{N} \models \varphi$.

For the sake of simplicity we denote the set of $\mathcal{L}$-formulae over σ by $\mathcal{L}(\sigma)$, and mostly write just $\models$ for $\models_{\mathcal{L}}$.

A logic $\mathcal{L}'$ is at *least as expressive* as a logic $\mathcal{L}$, if for every signature σ and every formula $\varphi \in \mathcal{L}(\sigma)$ there is a formula $\varphi' \in \mathcal{L}'(\sigma)$ such that $\mathfrak{M} \models \varphi$

iff $\mathfrak{M} \models \varphi'$ for every σ-structure $\mathfrak{M}$. If $\mathcal{L}$ is at least as expressive as $\mathcal{L}'$ and vice versa, then $\mathcal{L}$ and $\mathcal{L}'$ are *equi-expressive* or *equivalent*. Since syntactic variations are immaterial for our purposes, we regard equivalent logics as equal. For the same reason, if $\mathcal{L}'$ is at least expressive as $\mathcal{L}$, we may assume that $\mathcal{L}(\sigma) \subseteq \mathcal{L}'(\sigma)$ for every signature σ. We thus write just $\mathcal{L} \subseteq \mathcal{L}'$ and $\mathcal{L} = \mathcal{L}'$ for the corresponding relations between logics.

We call a logic $\mathcal{L}$ *compact*, if it satisfies the following, for every set $\Psi \subseteq \mathcal{L}(\sigma)$: existence of models for every finite $\Psi_0 \subseteq \Psi$ implies the existence of a model of Ψ (any finitely satisfiable set of $\mathcal{L}$-formulae is satisfiable). Note that we make no restriction on the cardinality of the sets Ψ under consideration (full compactness).

A modal signature (for ML or ML[∀] and similar extensions) consists of a pair of sets (τ, Φ), where the $\alpha \in \tau$ label the binary accessibility relations R_α and the $P \in \Phi$ correspond to unary predicates interpreting basic propositions in Kripke structures of type (τ, Φ). The class of Kripke structures of this type is denoted $\mathrm{Mod}(\tau, \Phi)$. In the sense of the general stipulation above, the relevant signatures σ are thus of the form $\sigma = (\tau, \Phi)$ and we stick to notation like $\mathcal{L}(\tau, \Phi)$. Also, σ-structures are pointed (τ, Φ)-structures, according to the natural semantics of modal logics like ML and ML[∀]. Where appropriate, we shall make this implicit as usual, in notation as in $(\mathfrak{M}, w) \models \varphi$.

ML[∀] is the extension of basic modal logic ML by a global modality (corresponding to the full accessibility relation), which we denote as ∀.

We are interested in logics that extend ML[∀] in the following sense.

DEFINITION 1. A logic $\mathcal{L}$ extends ML[∀] if, for every (τ, Φ), $\mathrm{ML}[\forall](\tau, \Phi) \subseteq \mathcal{L}(\tau, \Phi)$ and $\mathcal{L}(\tau, \Phi)$ is closed under $\wedge$, $\neg$ as well as $\exists$ and $\langle\alpha\rangle$ for every $\alpha \in \tau$.

The bisimulation game of basic modal logic ML may be extended in a natural manner to cover moves that capture the power of the global accessibility relation associated with the extension of ML by ∀. For this, one allows the first player to call 'global rounds' in which both players are allowed to freely relocate pebbles within the respective structure unconstrained by the accessibility relations R_α. For the second player to have a winning strategy in this modified infinite bisimulation game starting from configurations $(\mathfrak{M}, w); (\mathfrak{N}, v)$, which we denote by $(\mathfrak{M}, w) \overset{\forall}{\longleftrightarrow} (\mathfrak{N}, v)$, is the same as to require an ordinary bisimulation relation between $\mathfrak{M}$ and $\mathfrak{N}$ that is global (covers all of $\mathfrak{M}$ and all of $\mathfrak{N}$) and contains the pair (w, v). We speak of global bisimulation equivalence. It is easy to see that global bisimulation equivalence is exactly the right analogue of ordinary bisimulation equivalence that is appropriate for ML[∀]. In particular ML[∀] is invariant under global bisimulation equivalence in the following sense.

DEFINITION 2. A logic $\mathcal{L}$ is invariant under global bisimulation equivalence, or $\overset{\forall}{\longleftrightarrow}$ *invariant*, if for any two globally bisimilar pointed structures $(\mathfrak{M}, w) \overset{\forall}{\longleftrightarrow} (\mathfrak{N}, v)$ of type (τ, Φ) and any formula $\varphi \in \mathcal{L}(\tau, \Phi)$: $(\mathfrak{M}, w) \models \varphi$ iff $(\mathfrak{N}, v) \models \varphi$.

We write $\mathit{Th}(\mathfrak{M}, w)$ for the ML[∀](τ, Φ)-theory and $\mathit{Th}_{\mathcal{L}}(\mathfrak{M}, w)$ for the $\mathcal{L}(\tau, \Phi)$-theory of $(\mathfrak{M}, w)$.

DEFINITION 3. Let $\mathfrak{M}, \mathfrak{N} \in \mathrm{Mod}(\tau, \Phi)$. $\mathfrak{N}$ is an *$\mathcal{L}$-elementary extension* of $\mathfrak{M}$, $\mathfrak{M} \preccurlyeq_{\mathcal{L}} \mathfrak{N}$, if $\mathfrak{M}$ is an induced substructure of $\mathfrak{N}$ and $\mathit{Th}_{\mathcal{L}(\tau,\Phi)}(\mathfrak{M}, w) = \mathit{Th}_{\mathcal{L}(\tau,\Phi)}(\mathfrak{N}, w)$ for all w in $\mathfrak{M}$.

DEFINITION 4. A logic $\mathcal{L}$ is said to have the *Tarski Union Property (TUP)* if for every $\mathcal{L}$-elementary chain, $(\mathfrak{M}_i)_{i \in \mathbb{N}} \colon \mathfrak{M}_0 \preccurlyeq_{\mathcal{L}} \mathfrak{M}_1 \preccurlyeq_{\mathcal{L}} \mathfrak{M}_2 \preccurlyeq_{\mathcal{L}} \ldots$, the union $\bigcup_{i \in \mathbb{N}} \mathfrak{M}_i$ is an $\mathcal{L}$-elementary extension of each $\mathfrak{M}_j$.

OBSERVATION 5. ML[∀] has the Tarski Union Property.

Proof. Let $\mathfrak{M}^* := \bigcup_{i \in \mathbb{N}} \mathfrak{M}_i$. As $\mathfrak{M}_j$ is a substructure of $\mathfrak{M}^*$, atomic formulae are preserved at all w in $\mathfrak{M}_j$. The claim is trivially compatible with boolean operations. It remains to give inductive arguments for the $\langle \alpha \rangle$- and ∃-steps in formula formation.

$\langle \alpha \rangle$. Let $(\mathfrak{M}_j, w) \models \langle \alpha \rangle \varphi$, i.e., $(\mathfrak{M}_j, w') \models \varphi$ for some $(w, w') \in R_\alpha^{\mathfrak{M}_j}$. By the inductive hypothesis for φ, $(\mathfrak{M}^*, w') \models \varphi$, and therefore $(\mathfrak{M}^*, w) \models \langle \alpha \rangle \varphi$, as also $(w, w') \in R_\alpha^{\mathfrak{M}^*}$.

Conversely, if $(\mathfrak{M}^*, w) \models \langle \alpha \rangle \varphi$ through some $(w, w') \in R_\alpha^{\mathfrak{M}^*}$ such that $(\mathfrak{M}^*, w') \models \varphi$, then $w' \in M_k$ and $(w, w') \in R_\alpha^{\mathfrak{M}_k}$ for some $k \in \mathbb{N}$. By the inductive hypothesis for φ, $(\mathfrak{M}_k, w') \models \varphi$, whence $(\mathfrak{M}_k, w) \models \langle \alpha \rangle \varphi$. By the $\mathcal{L}$-elementary nature of the chain, $(\mathfrak{M}_j, w) \models \langle \alpha \rangle \varphi$ for all j such that $w \in M_j$.

∃ is treated analogously. ■

Our goal is the following characterisation of ML[∀] as maximally expressive among a natural class of $\overset{\forall}{\Longleftrightarrow}$ invariant logics.

THEOREM 6. *Any compact $\overset{\forall}{\Longleftrightarrow}$ invariant logic $\mathcal{L}$ with the Tarski Union Property that extends* ML[∀] *is equivalent to* ML[∀] *itself.*

We define several natural notions and provide some lemmas towards the proof.

DEFINITION 7. A set of ML(τ, Φ)-formulae Γ is called an *α-type of* $(\mathfrak{M}, w)$ (or of $\mathit{Th}(\mathfrak{M}, w)$) if $(\mathfrak{M}, w) \models \langle \alpha \rangle \bigwedge \Gamma_0$ for all finite $\Gamma_0 \subseteq \Gamma$.

∃-types of $\mathfrak{M}$ (or of $\mathit{Th}(\mathfrak{M})$) are similarly defined: $\mathfrak{M} \models \exists \bigwedge \Gamma_0$ must apply for all finite subsets Γ_0 of Γ.

DEFINITION 8. An α-type Γ of $(\mathfrak{M}, w)$ is *realised in* $(\mathfrak{M}, w)$ if there is some w' in $\mathfrak{M}$ such that $(w, w') \in R_\alpha^{\mathfrak{M}}$ and $(\mathfrak{M}, w') \models \Gamma$. An ∃-type Γ of $\mathfrak{M}$ is *realised in* $\mathfrak{M}$ if $(\mathfrak{M}, w') \models \Gamma$ for some w' in $\mathfrak{M}$.

A structure is called *saturated* if for all w in $\mathfrak{M}$ and all $\alpha \in \tau$ every α-type of $(\mathfrak{M}, w)$ is realised in $(\mathfrak{M}, w)$ and if every ∃-type of $\mathfrak{M}$ is realised in $\mathfrak{M}$.

The following is the natural variant of the Hennessy–Milner theorem for global bisimulation equivalence and ML[∀] over the class of saturated Kripke structures. In fact it is easily seen via the game that, over saturated Kripke structures, ML[∀]-equivalence induces a global bismulation, see, e.g., [6].

THEOREM 9 (Hennessy–Milner). *If* $\mathfrak{M}, \mathfrak{N} \in \mathrm{Mod}(\tau, \Phi)$ *are saturated, then* $\mathit{Th}(\mathfrak{M}, w) = \mathit{Th}(\mathfrak{N}, v)$ *implies* $(\mathfrak{M}, w) \overset{\forall}{\Longleftrightarrow} (\mathfrak{N}, v)$.

A forest-unfolding $(\mathfrak{M}^F, w)$ of $(\mathfrak{M}, w)$ is the disjoint union of all tree-unravelings in every element of $(\mathfrak{M}, w)$. Since any forest-unfolding $(\mathfrak{M}^F, w)$ is globally bisimilar to its underlying model $(\mathfrak{M}, w)$, they have the same $\mathcal{L}$-theory. Below, we shall use them as a normalised representation of models that allow us to embed one into another.

LEMMA 10. *Let* $(\mathfrak{M}, w)$ *be a forest model with a uniquely assigned propositional letter* $P_{w'}$ *for each element* w' *in* $(\mathfrak{M}, w)$. *Then any* $(\mathfrak{N}, v) \models \mathit{Th}(\mathfrak{M}, w)$ *admits an isomorphic embedding* $\iota : (\mathfrak{M}, w) \hookrightarrow (\mathfrak{N}, v)$.

Proof. $(\mathfrak{M}, w)$ can be embedded in $(\mathfrak{N}, v)$ by an injection $\iota : (\mathfrak{M}, w) \hookrightarrow (\mathfrak{N}, v)$, which is inductively defined (w.r.t. distance from the roots in the component trees of the forest model $\mathfrak{M}$) such that $(\mathfrak{N}, \iota(w')) \models P_{w'}$ for every w' in $\mathfrak{M}$. ■

PROPOSITION 11. *Let* $\mathcal{L}$ *be a compact logic extending* ML[∀]. *Then every forest model* $\mathfrak{M}$ *admits an* $\mathcal{L}$*-elementary extension* $\mathfrak{M}'$ *that realises all* α*-types of* $(\mathfrak{M}, w)$ *(as* α*-types of* $(\mathfrak{M}', w)$*) for all* w *in* $\mathfrak{M}$ *and realises all* $\exists$*-types of* $\mathfrak{M}$ *(as* $\exists$*-types of* $\mathfrak{M}'$*).*

Proof. We introduce new propositional letters to Φ by setting

$$\begin{array}{rclcl} \Psi & := & \Phi & \cup & \{P_w \mid w \in \mathfrak{M}\} \\ & & & \cup & \{P^{\alpha}_{w,\Gamma} \mid w \text{ in } \mathfrak{M},\ \Gamma \text{ an } \alpha\text{-type of } (\mathfrak{M}, w)\} \\ & & & \cup & \{P_\Gamma \mid \Gamma \text{ an } \exists\text{-type of } \mathfrak{M}\} \end{array}$$

for disjoint sets of new unary predicates. Let T' be the following $\mathcal{L}(\tau, \Psi)$-theory (towards an axiomatisation of the (τ, Ψ)-expansion of the desired $\mathfrak{M}'$). T' comprises, for all w in $\mathfrak{M}$, the following $\mathcal{L}(\tau, \Psi)$-formulae:

1. $\exists P_w$.
2. $\forall(P_w \longrightarrow \neg P_{w'})$, for all $w' \neq w$ in $\mathfrak{M}$.
3. $\forall(P_w \longrightarrow \langle\alpha\rangle P_{w'})$, for all $(w, w') \in R^{\mathfrak{M}}_\alpha$, $\alpha \in \tau$.
4. $\forall(P_w \longrightarrow \neg\langle\alpha\rangle P_{w'})$, for all $(w, w') \notin R^{\mathfrak{M}}_\alpha$, $\alpha \in \tau$.
5. $\forall(P_w \longrightarrow \xi)$, for every $\xi \in \mathit{Th}_{\mathcal{L}}(\mathfrak{M}, w)$.
6. $\forall(P_w \longrightarrow \langle\alpha\rangle P^{\alpha}_{w,\Gamma})$, for all $\alpha \in \tau$ and every α-type Γ of $(\mathfrak{M}, w)$.
7. $\exists P_\Gamma$, for every $\exists$-type Γ of $\mathfrak{M}$.
8. $\forall(Q \longrightarrow \xi)$, for all $Q = P^{\alpha}_{w,\Gamma}, P_\Gamma$ in Ψ and every $\xi \in \Gamma$.

T' is finitely satisfiable (in expansions of $\mathfrak{M}$), hence satisfiable by compactness of $\mathcal{L}$. Let $\mathfrak{M}' \models T'$. As $\mathcal{L}$ is invariant under global bisimulation $\overset{\forall}{\longleftrightarrow}$, we may assume w.l.o.g. that $\mathfrak{M}'$ is a forest model. By construction, the (τ, Φ)-reduct of $\mathfrak{M}'$ is isomorphic to an $\mathcal{L}$-elementary extension of the forest model $\mathfrak{M}$: the isomorphism as in the proof of lemma 10 here yields an $\mathcal{L}$-elementary embedding, due to the formulae in (5). W.l.o.g., the forest model $\mathfrak{M}'$ is an $\mathcal{L}$-elementary extension of $\mathfrak{M}$. Moreover, $\mathfrak{M}'$ realises all required types, by (6)–(8). ■

COROLLARY 12. *Let $\mathcal{L}$ be a compact logic extending* ML[∀] *with TUP. Then every forest model* $(\mathfrak{M}, w)$ *possesses a saturated $\mathcal{L}$-elementary extension.*

Proof. Starting with the given model $(\mathfrak{M}, w)$, a repeated application of proposition 11 yields a chain of forest models, in which the (τ, Φ)-reduct of each model is an $\mathcal{L}$-elementary extension of its predecessor (restricted to (τ, Φ)). Since $\mathcal{L}$ has the Tarski Union Property, the limit of this chain is an $\mathcal{L}$-elementary extension for all members of the chain. Every type of $\mathfrak{M}_i$ is realised in $\mathfrak{M}_{i+1}$; so the limit is saturated. ■

The following indicates how to complete an ML[∀](τ, Φ)-theory while maintaining ML[∀]-inexpressibility of a given $\varphi \in \mathcal{L}(\tau, \Phi)$.

LEMMA 13. *Let $\mathcal{L}$ be a logic extending* ML[∀], $\varphi \in \mathcal{L}(\tau, \Phi)$, $T \subseteq$ ML[∀](τ, Φ) *and* $\psi \in$ ML[∀](τ, Φ). *If there is no* $\chi \in$ ML[∀](τ, Φ) *such that* $T \models \varphi \longleftrightarrow \chi$, *then the same is true of at least one of* $T \cup \{\psi\}$ *or* $T \cup \{\neg\psi\}$.

Proof. Assume for both ψ and $\neg\psi$ there were formulae χ and χ' such that $T \models \psi \longrightarrow (\varphi \longleftrightarrow \chi)$ and $T \models \neg\psi \longrightarrow (\varphi \longleftrightarrow \chi')$. Then $T \models \varphi \longleftrightarrow ((\chi \wedge \psi) \vee (\chi' \wedge \neg\psi))$ contradicts our assumptions, since $(\chi \wedge \psi) \vee (\chi' \wedge \neg\psi)$ is in ML[∀](τ, Φ). ■

Compactness of $\mathcal{L}$ guarantees that, in the situation of the lemma, the set of ML[∀](τ, Φ)-theories under which φ is not equivalent to any ML[∀](τ, Φ)-formula is closed under unions of $\subseteq$-chains. By Zorn's lemma, we thus obtain a $\subseteq$-maximal such $T \subseteq$ ML[∀](τ, Φ). By the lemma, such T is a complete ML[∀](τ, Φ)-theory. Moreover, both $T \cup \{\varphi\}$ and $T \cup \{\neg\varphi\}$ are satisfiable, as otherwise φ would be equivalent to $\bot$ or $\top$ under T. We thus get the following.

PROPOSITION 14. *If $\mathcal{L}$ is a compact logic extending* ML[∀] *and for some signature* (τ, Φ) *there is a formula* $\varphi \in \mathcal{L}(\tau, \Phi)$ *not equivalent to any* $\chi \in$ ML[∀](τ, Φ), *then there are two* (τ, Φ)*-models* $\mathfrak{M}$ *and* $\mathfrak{N}$ *such that* $Th(\mathfrak{M}, w) = Th(\mathfrak{N}, v)$ *and* $(\mathfrak{M}, w) \models \varphi$ *while* $(\mathfrak{N}, v) \models \neg\varphi$.

Proof of theorem 6. Assume, for some (τ, Φ) there is a formula $\varphi \in \mathcal{L}(\tau, \Phi)$ which is not equivalent to any formula in ML[∀](τ, Φ). According to proposition 14 there are two (τ, Φ)-models $\mathfrak{M}$ and $\mathfrak{N}$ with $Th(\mathfrak{M}, w) = Th(\mathfrak{N}, v)$ and $(\mathfrak{M}, w) \models \varphi$ and $(\mathfrak{N}, v) \models \neg\varphi$. Since $\mathcal{L}$ is invariant under global bisimulation $\overset{\forall}{\longleftrightarrow}$, we may assume w.l.o.g. that $\mathfrak{M}$ and $\mathfrak{N}$ are forest models. By compactness, corollary 12 yields saturated (τ, Φ)-models $(\mathfrak{M}^*, w)$ and $(\mathfrak{N}^*, v)$ which have the same $\mathcal{L}(\tau, \Phi)$-theories as the originals. Therefore $\mathit{Th}(\mathfrak{M}^*, w) = \mathit{Th}(\mathfrak{N}^*, v)$ and, by theorem 9, $(\mathfrak{M}^*, w) \overset{\forall}{\longleftrightarrow} (\mathfrak{N}^*, v)$. But then $(\mathfrak{M}^*, w) \models \varphi$ and $(\mathfrak{N}^*, v) \models \neg\varphi$ shows that φ is not invariant under global bisimulation, contradicting the assumptions on $\mathcal{L}$. ■

3 Characterisation of GF

First we introduce the relevant basic notions for GF, in analogy with those used for ML[∀] above. For guarded bisimulation invariant candidate logics $\mathcal{L}$, suitable and natural notions of elementary extensions and the Tarski Union Property are presented, before we state the main theorem. For the following we work with arbitrary relational signatures τ.

The guarded fragment $\mathrm{GF}(\tau) \subseteq \mathrm{FO}(\tau)$ is introduced as the restriction of FO that only allows quantification of the following *guarded* format. For a relation symbol $R \in \tau$ (or $=$), we write $R(\bar{x}\bar{y})$ for an R-atom containing all the displayed variables (but not necessarily in this order, and repetitions are also allowed). A quantification $\exists\bar{y}.R(\bar{x}\bar{y}) \wedge \varphi$ is guarded if, and only if, $\mathrm{free}(\varphi) \subseteq \mathrm{free}(R(\bar{x}\bar{y}))$; the R-atom $R(\bar{x}\bar{y})$ is a *guard* in this first-order quantification. Since equality atoms may also serve as guards, $\exists y.y = y \wedge \varphi$ is a formula of $\mathrm{GF}(\tau)$ whenever y is the only free variable in $\varphi \in \mathrm{GF}(\tau)$.

As a fragment of FO, GF is compact.

In a τ-structure $\mathfrak{M}$, a subset $X \subseteq M$ is called *guarded* if it is a singleton set or for some $R \in \tau$ there is a tuple $\bar{a} \in R^{\mathfrak{M}}$ comprising all the elements of X. A tuple $\bar{a}$ is *guarded* in $\mathfrak{M}$ if there is a guarded set in $\mathfrak{M}$ that includes all components of $\bar{a}$. A tuple is called *strictly guarded* if its set of components is precisely the set of components of some $R^{\mathfrak{M}}$-atom, or a singleton set.

We also introduce the following terminology that is suitable for our purposes. A tuple $\bar{a}'$ is an $\exists\bar{y}.R(\bar{x}\bar{y})$-successor of $\bar{a}$ in $\mathfrak{M}$ if the assignment $\bar{x}\bar{y} \mapsto \bar{a}'$ is an extension of the assignment $\bar{x} \mapsto \bar{a}$ such that $\mathfrak{M}, \bar{a}' \models R(\bar{x}\bar{y})$. Clearly guarded quantifications correspond to modal quantifications w.r.t. transitions to $\exists\bar{y}.R(\bar{x}\bar{y})$-successors (which are strictly guarded tuples). This analogy is at the root of the appropriate notion of guarded bisimulation and of guarded tree-unfoldings to be discussed below.

Let $\mathfrak{M}, \mathfrak{N}$ be two τ-structures, possibly with tuples of distinguished parameters $\bar{a}$ and $\bar{b}$ of matching lengths. The guarded bisimulation game on $\mathfrak{M}$ and $\mathfrak{N}$ is played by two players **I** and **II**. A generic configuration of the game consists of designated strictly guarded tuples of the same length, one in each structure, denoted $(\mathfrak{M}, \bar{a}); (\mathfrak{N}, \bar{b})$.[1] **II** will have lost unless the componentwise mapping $\bar{a} \longmapsto \bar{b}$ is a partial isomorphism between $\mathfrak{M}$ and $\mathfrak{N}$.

In each new round player **I** chooses to play in one of the two structures, say $\mathfrak{M}$, and chooses a (possibly empty) subtuple $\bar{a}_0$ of the current tuple $\bar{a}$

[1] In the initial configuration of the game, the given, not necessarily guarded tuples of distinguished parameters are admitted.

that stays fixed, and a completion of $\bar{a}_0$ to some strictly guarded tuple $\bar{a}'$. **II** has to choose a strictly guarded tuple $\bar{b}'$ extending the corresponding $\bar{b}_0$ such that the componentwise mapping $\bar{a}' \longmapsto \bar{b}'$ is again a partial isomorphism between $\mathfrak{M}$ and $\mathfrak{N}$. Note that, if **I** chose an $\exists\bar{y}.R(\bar{x}\bar{y})$-successor of $\bar{a}_0$, then the rules force **II** to do likewise.

II looses the game if she cannot provide an answer that satisfies these constraints. We say that **II** wins the game, if she has a winning strategy which allows her to respond to all challenges of **I** indefinitely. Two τ-structures with designated tuples $\mathfrak{M}, \bar{a}$ and $\mathfrak{N}, \bar{b}$ are *guarded bisimilar*, $\mathfrak{M}, \bar{a} \overset{\mathrm{g}}{\Longleftrightarrow} \mathfrak{N}, \bar{b}$, if **II** wins the game from $(\mathfrak{M}, \bar{a}); (\mathfrak{N}, \bar{b})$.

Guarded bisimulation equivalence is the natural variant of Ehrenfeucht–Fraïssé equivalence associated with GF $\subseteq$ FO. In particular, the semantics of GF is invariant under this equivalence: if $\mathfrak{M}, \bar{a} \overset{\mathrm{g}}{\Longleftrightarrow} \mathfrak{N}, \bar{b}$, then $\mathfrak{M}, \bar{a}$ and $\mathfrak{N}, \bar{b}$ are GF-equivalent, or have the same GF-theories (cf. the following definition).

DEFINITION 15. The GF-*theory* of a tuple $\bar{a}$ in a τ-structure $\mathfrak{M}$ is the set of GF-formulae satisfied by $\bar{a}$ in $\mathfrak{M}$: $\mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{M}, \bar{a}) := \{\varphi \in \mathrm{GF}(\tau) \mid (\mathfrak{M}, \bar{a}) \models \varphi\}$.

DEFINITION 16. Let $\mathfrak{M}$ be a τ-structure. A set $\Gamma \subseteq \mathrm{GF}(\tau)$ is called an $\exists\bar{y}.R(\bar{x}\bar{y})$-*type* of $(\mathfrak{M}, \bar{a})$ if for every finite subset $\Gamma_0 \subseteq \Gamma$ $(\mathfrak{M}, \bar{a}) \models \exists\bar{y}.R(\bar{x}\bar{y}) \wedge \bigwedge \Gamma_0$.

This type is *realised at* $(\mathfrak{M}, \bar{a})$ if there is an $\exists\bar{y}.R(\bar{x}\bar{y})$-successor $\bar{a}'$ of $\bar{a}$ in $\mathfrak{M}$ such that $(\mathfrak{M}, \bar{a}') \models \Gamma$.

A τ-structure $\mathfrak{M}$ is GF-*saturated* if, for every guarded tuple $\bar{a}$ and every $\exists\bar{y}.R(\bar{x}\bar{y})$, all $\exists\bar{y}.R(\bar{x}\bar{y})$-types of $(\mathfrak{M}, \bar{a})$ are realised at $(\mathfrak{M}, \bar{a})$.

The following analogue of the Hennessy–Milner theorem is then immediate.

PROPOSITION 17 (Hennessy–Milner). *Let $\mathfrak{M}, \bar{a}$ and $\mathfrak{N}, \bar{b}$ be two* GF*-saturated τ-structures with parameter tuples. If $\mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{M}, \bar{a}) = \mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{N}, \bar{b})$, then $\mathfrak{M}, \bar{a} \overset{g}{\Longleftrightarrow} \mathfrak{N}, \bar{b}$.*

Proof. Indeed, GF-equivalence between (strictly guarded) tuples can be maintained by **II** and thus provides a winning strategy. Assume w.l.o.g. that **I** chooses an $\exists\bar{y}.R(\bar{x}\bar{y})$-successor $\bar{a}'$ of some subtuple $\bar{a}_0$ of $\bar{a}$ in $\mathfrak{M}$. Since $(\mathfrak{M}, \bar{a})$ and $(\mathfrak{N}, \bar{b})$ have the same theory,

$$(\mathfrak{M}, \bar{a}_0) \models \exists\bar{y}.R(\bar{x}\bar{y}) \wedge \bigwedge \Gamma_0 \quad \text{iff} \quad (\mathfrak{N}, \bar{b}_0) \models \exists\bar{y}.R(\bar{x}\bar{y}) \wedge \bigwedge \Gamma_0$$

for all finite $\Gamma_0 \subseteq \mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{M}, \bar{a}')$. So $\mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{M}, \bar{a}')$ is an $\exists\bar{y}.R(\bar{x}\bar{y})$-type of $\mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{N}, \bar{b}_0)$. As $\mathfrak{N}$ is GF-saturated, there is an $\exists\bar{y}.R(\bar{x}\bar{y})$-successor $\bar{b}'$ of $\bar{b}_0$ such that $(\mathfrak{N}, \bar{b}') \models \mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{M}, \bar{a}')$ for **II** to play. ■

DEFINITION 18. A logic $\mathcal{L}$ extends GF if, for every signature τ, $\mathrm{GF}(\tau) \subseteq \mathcal{L}(\tau)$ and $\mathcal{L}$ is closed under boolean operations $\wedge, \neg$ and guarded quantification. Closure under guarded quantification means that for any τ-atom

$R(\bar{x}\bar{y})$ and any $\varphi \in \mathcal{L}(\tau)$ with free variables[2] amongst $\bar{x}\bar{y}$, there is an $\mathcal{L}$-formula φ' with the semantics of $\exists\bar{y}.R(\bar{x}\bar{y}) \wedge \varphi$.

DEFINITION 19. Let $\mathfrak{M}, \mathfrak{N}$ be two τ-structures. $\mathfrak{N}$ is an *$\mathcal{L}$-elementary extension of* $\mathfrak{M}$ (in a guarded sense), abbreviated by $\mathfrak{M} \preccurlyeq_{\mathcal{L}} \mathfrak{N}$, if $\mathfrak{M}$ is an induced substructure of $\mathfrak{N}$ and $\mathit{Th}_{\mathcal{L}(\tau)}(\mathfrak{M}, \bar{a}) = \mathit{Th}_{\mathcal{L}(\tau)}(\mathfrak{N}, \bar{a})$ for all guarded tuples $\bar{a}$ in $\mathfrak{M}$. $\mathfrak{M}, \bar{a} \preccurlyeq_{\mathcal{L}} \mathfrak{N}, \bar{a}$ can be similarly defined for a not necessarily guarded tuple of distinguished parameters.

Recall from definition 4 that $\mathcal{L}$ has the Tarski Union Property (TUP) if the limit of every $\mathcal{L}$-elementary chain $(\mathfrak{M}_i)_{i\in\mathbb{N}}$ is an $\mathcal{L}$-elementary extension of each $\mathfrak{M}_i$.

OBSERVATION 20. GF has TUP.

Proof. Let $\mathfrak{M}_0 \preccurlyeq_{\mathrm{GF}} \mathfrak{M}_1 \preccurlyeq_{\mathrm{GF}} \mathfrak{M}_2 \ldots$ be a chain of GF-elementary extensions and $\mathfrak{M}^* := \bigcup_{i\in\mathbb{N}} \mathfrak{M}_i$. Let furthermore $\mathfrak{M}_j$ be a member of the chain and $\bar{a}$ a guarded tuple in $\mathfrak{M}_j$. It has to be shown that $\mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{M}_j, \bar{a}) = \mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{M}^*, \bar{a})$. The only interesting step in the syntactic induction is that of guarded quantification.

Assume $(\mathfrak{M}_j, \bar{a}) \models \exists\bar{y}.R(\bar{x}\bar{y}) \wedge \varphi$. For some $\exists\bar{y}.R(\bar{x}\bar{y})$-successor $\bar{a}'$ we have $(\mathfrak{M}_j, \bar{a}') \models \varphi$, which entails $(\mathfrak{M}^*, \bar{a}') \models \varphi$ by the induction hypothesis, and hence $(\mathfrak{M}^*, \bar{a}) \models \exists\bar{y}.R(\bar{x}\bar{y}) \wedge \varphi$.

Conversely, if $(\mathfrak{M}^*, \bar{a}) \models \exists\bar{y}.R(\bar{x}\bar{y}) \wedge \varphi$, then there is some $k \in \mathbb{N}$ such that the $\exists\bar{y}.R(\bar{x}\bar{y})$-successor $\bar{a}'$ satisfying φ lives in $\mathfrak{M}_k$. For $m := \max\{j, k\}$ this $\bar{a}'$ is an $\exists\bar{y}.R(\bar{x}\bar{y})$-successor of $\bar{a}$ in $\mathfrak{M}_m$. By the induction hypothesis $(\mathfrak{M}_m, \bar{a}') \models \varphi$ and hence $(\mathfrak{M}_m, \bar{a}) \models \exists\bar{y}.R(\bar{x}\bar{y}) \wedge \varphi$. The latter entails $(\mathfrak{M}_j, \bar{a}) \models \exists\bar{y}.R(\bar{x}\bar{y}) \wedge \varphi$, since $\mathfrak{M}_m$ is an GF-elementary extension of $\mathfrak{M}_j$. ■

In analogy to tree-unravellings of Kripke-structures, every τ-structure is guarded bisimilar to a structure that is tree-like w.r.t. the accessibility relations induced by the $\exists\bar{y}.R(\bar{x}\bar{y})$-successor relations. Technically, these *tree-like structures* are tree-decomposable into substructures consisting of (strictly) guarded tuples, or *guarded tree-decomposable*, cf. the generalised tree model property of [7], and see [8] for details of the analogy with ordinary unravellings. Intuitively, these guarded tree unfoldings are obtained by introducing new disjoint copies of elements along every path of overlapping guarded subsets leading to them. For a logic that is invariant under guarded bisimulation, one may w.l.o.g. restrict attention to the tree-like models thus obtained.

REMARK 21. Let $\mathfrak{M}$ be a tree-like (i.e., guarded tree-decomposable) τ-structure such that for every guarded tuple $\bar{a}$ in $\mathfrak{M}$ there is a predicate $P_{\bar{a}} \in \tau$ with $P_{\bar{a}}^{\mathfrak{M}} = \{\bar{a}\}$. Then $\mathfrak{M}$ can be embedded in any structure $\mathfrak{N}$ that satisfies the following GF-formulae for all $P_{\bar{a}}, P_{\bar{a}'} \in \tau$:

---

[2]Instead of explicitly referring to a notion of free variable for formulae in an abstract logic $\mathcal{L}$, a semantic description can be given: for every $\varphi \in \mathcal{L}(\tau)$ there is $\varphi' \in \mathcal{L}(\tau)$, s.t. $(\mathfrak{M}, \bar{a}) \models \varphi'$ iff there is an $\exists\bar{y}.R(\bar{x}\bar{y})$-successor $\bar{a}'$ of $\bar{a}$ with $(\mathfrak{M}, \bar{a}') \models \varphi$.

1. $\exists\bar{x}.P_{\bar{a}}\bar{x}$.
2. $\forall\bar{x}.P_{\bar{a}}\bar{x} \longrightarrow \bigwedge_{1\leq i\leq k} P_{a_i}x_i$ where $\bar{a}=(a_1,\dots,a_k)$.
3. $\forall x.P_a x \longrightarrow \neg P_{a'}x$ for all $a \neq a'$ in $\mathfrak{M}$.
4. $\forall\bar{x}.P_{\bar{a}}\bar{x} \longrightarrow \exists\bar{y}.P_{\bar{a}'}(\bar{x}\bar{y})$, where $\bar{y}$ represents the components of $\bar{a}'\setminus\bar{a}$.
5. $\forall\bar{x}.P_{\bar{a}}\bar{x} \longrightarrow R(\bar{x})$ for all atomic formulae $R(\bar{x}) \in \mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{M},\bar{a})$.
6. $\forall\bar{x}.R\bar{x} \longrightarrow \neg(\bigwedge_{i\leq k} P_{a_i}x_i)$
   for all $(a_1,\dots,a_k)\notin R^{\mathfrak{M}}$, $R\in\tau$ of arity k and $(a_1,\dots,a_k)\in M^k$.

Since $\mathfrak{M}$ is tree-like, the embedding can be defined inductively w.r.t. distance from the root in a guarded tree-decomposition (i.e., successively proceeding to $\exists\bar{y}.R(\bar{x}\bar{y})$-successors).

PROPOSITION 22. *Let $\mathcal{L}$ be a compact, guarded bisimulation invariant logic that has TUP. Then every tree-like τ-structure $\mathfrak{M}$ has an $\mathcal{L}$-elementary extension that is* GF*-saturated.*

The proof is based on a chain limit (TUP) of a chain obtained through the following process.

PROPOSITION 23. *Let $\mathfrak{M}$ be a tree-like τ-structure. Then there is a tree-like $\mathcal{L}$-elementary extension $\mathfrak{N}$ of $\mathfrak{M}$ such that, for every guarded tuple $\bar{a}$ in $\mathfrak{M}$ and for every $\exists\bar{y}.R(\bar{x}\bar{y})\in\mathrm{GF}(\tau)$, every $\exists\bar{y}.R(\bar{x}\bar{y})$-type of $\mathfrak{M},\bar{a}$ is realised at $(\mathfrak{N},\bar{a})$.*

Proof. Let $G(\mathfrak{M})$ be the set of all guarded tuples in $\mathfrak{M}$. We extend τ to $\sigma := \tau \,\dot\cup\, \{P_{\bar{a}} \mid \bar{a}\in G(\mathfrak{M})\}$ and σ further to

$$\rho := \sigma \,\dot\cup \bigcup_{\bar{a}\in G(\mathfrak{M})} \{P^{\alpha}_{\bar{a},\Gamma} \mid \alpha = \exists\bar{y}.R(\bar{x}\bar{y}), \Gamma\subseteq\mathrm{GF}(\tau) \text{ an } \alpha\text{-type of } (\mathfrak{M},\bar{a})\}.$$

The set of formulae defined in remark 21 can now be formulated in $\mathrm{GF}(\sigma)$. We extend this set to $T\subseteq\mathcal{L}(\rho)$ by adding the following formulae for every $\bar{a}\in G(\mathfrak{M})$.

7. $\forall\bar{x}.P_{\bar{a}}\bar{x} \longrightarrow \varphi$ for all $\varphi\in\mathit{Th}_{\mathcal{L}(\tau)}(\mathfrak{M},\bar{a})$.
8. $\forall\bar{x}.P_{\bar{a}}\bar{x} \longrightarrow \exists\bar{y}.R(\bar{x}\bar{y}) \wedge P^{\alpha}_{\bar{a},\Gamma}(\bar{x}\bar{y})$
   for every $\alpha = \exists\bar{y}.R(\bar{x}\bar{y})$ and every $\exists\bar{y}.R(\bar{x}\bar{y})$-type Γ of $(\mathfrak{M},\bar{a})$.
9. $\forall\bar{x}.P^{\alpha}_{\bar{a},\Gamma}\bar{x} \longrightarrow \varphi$ for every $\varphi\in\Gamma$ and $P^{\alpha}_{\bar{a},\Gamma}\in\rho$.

A simple compactness argument shows that T is satisfiable; indeed any finite subset of T is satisfiable in an expansion of $\mathfrak{M}$.

Let $\mathfrak{N}$ be a ρ-structure satisfying T. Since $\mathcal{L}$ is bisimulation invariant we may assume that $\mathfrak{N}$ is tree-like. By setting $P^{\mathfrak{M}}_{\bar{a}} := \{\bar{a}\}$ the τ-structure $\mathfrak{M}$ can be extended to a σ-structure, and is embeddable into the σ-reduct $\mathfrak{N}\restriction\sigma$

according to remark 21. Hence we may assume, that $\mathfrak{M}$ is a substructure of $\mathfrak{N}\restriction\tau$.

For every $\bar{a} \in G(\mathfrak{M})$ the formulae in (7) guarantee $(\mathfrak{N}, \bar{a}) \models \mathit{Th}_{\mathcal{L}(\tau)}(\mathfrak{M}, \bar{a})$ and therefore $\mathfrak{N}\restriction\tau$ is an $\mathcal{L}$-elementary extension of $\mathfrak{M}$.

It remains to show, that every $\exists\bar{y}.R(\bar{x}\bar{y})$-Type of $(\mathfrak{M}, \bar{a})$ is realised at $(\mathfrak{N}, \bar{a})$. This is clear from (8) and (9). ■

Proof of proposition 22. Let $\mathfrak{M}_0 := \mathfrak{M}$ and inductively let $\mathfrak{M}_{i+1}$ be a tree-like $\mathcal{L}$-elementary extension that realises every GF-type of $\mathfrak{M}_i$ as obtained through proposition 23. We thus get an $\mathcal{L}$-elementary chain $\mathfrak{M}_0 \preccurlyeq_{\mathcal{L}} \mathfrak{M}_1 \preccurlyeq_{\mathcal{L}} \mathfrak{M}_2 \preccurlyeq_{\mathcal{L}} \ldots$, and since $\mathcal{L}$ has TUP, the limit $\mathfrak{M}^*$ is an $\mathcal{L}$-elementary extension of every member $\mathfrak{M}_i$ and in particular of $\mathfrak{M}$.

$\mathfrak{M}^*$ is GF-saturated: let Γ be an $\exists\bar{y}.R(\bar{x}\bar{y})$-type for some guarded tuple $\bar{a}$ in $\mathfrak{M}^*$. There is some $i \in \mathbb{N}$ such that $\bar{a}$ is a guarded tuple of $\mathfrak{M}_i$. Since $\mathfrak{M}^*$ is an $\mathcal{L}$-elementary extension for $\mathfrak{M}_i$, in particular $\mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{M}_i, \bar{a}) = \mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{M}^*, \bar{a})$. Therefore $(\mathfrak{M}_i, \bar{a})$ has exactly the same $\exists\bar{y}.R(\bar{x}\bar{y})$-types as $(\mathfrak{M}^*, \bar{a})$. All those types are realised in $\mathfrak{M}_{i+1}$ and hence in $\mathfrak{M}^*$ by the $\mathcal{L}$-elementary nature of the extension.

Hence $\mathfrak{M}^*$ is a GF-saturated $\mathcal{L}$-elementary extension of $\mathfrak{M}$. ■

COROLLARY 24. *For every τ-structure $\mathfrak{M}$ with distinguished parameters $\tilde{a}$ there is a* GF*-saturated $\mathcal{L}$-elementary extension $\mathfrak{M}^*$ for $\mathfrak{M}$ such that also $\mathit{Th}_{\mathcal{L}(\tau)}(\mathfrak{M}, \tilde{a}) = \mathit{Th}_{\mathcal{L}(\tau)}(\mathfrak{M}^*, \tilde{a})$.*

Proof. We expand the signature τ to $\dot{\tau} := \tau \cup \{P_{\tilde{a}}\}$ and set $P_{\tilde{a}}^{\mathfrak{M}} := \{\tilde{a}\}$. Using proposition 22 for the $\dot{\tau}$-structure $\mathfrak{M}$, we obtain a GF-saturated $\dot{\tau}$-structure $\mathfrak{M}^*$ such that $\mathfrak{M} \preccurlyeq_{\mathcal{L}} \mathfrak{M}^*$ w.r.t. $\dot{\tau}$. So $\mathit{Th}_{\mathcal{L}(\dot{\tau})}(\mathfrak{M}, \bar{a}) = \mathit{Th}_{\mathcal{L}(\dot{\tau})}(\mathfrak{M}^*, \bar{a})$ for every guarded tuple $\bar{a}$ in $\mathfrak{M}$ and especially for $\tilde{a}$, which was guarded by $P_{\tilde{a}} \in \dot{\tau}$. Since $\mathcal{L}$ is compatible with reducts (cf. condition (3) on abstract logics in section 2), we get $\mathit{Th}_{\mathcal{L}(\tau)}(\mathfrak{M}, \bar{a}) = \mathit{Th}_{\mathcal{L}(\tau)}(\mathfrak{M}^*, \bar{a})$ for every guarded tuple, including $\tilde{a}$. Therefore $\mathfrak{M}\restriction\tau \preccurlyeq_{\mathcal{L}} \mathfrak{M}^*\restriction\tau$ and $\mathit{Th}_{\mathcal{L}(\tau)}(\mathfrak{M}, \tilde{a}) = \mathit{Th}_{\mathcal{L}(\tau)}(\mathfrak{M}^*, \tilde{a})$. Since the τ-reduct of $\mathfrak{M}^*$ remains GF-saturated, $\mathfrak{M}^*\restriction\tau$ is the model we are looking for. ■

THEOREM 25. *Any compact $\overset{g}{\Longleftrightarrow}$ invariant logic $\mathcal{L}$ with the Tarski Union Property that extends GF is equivalent to GF itself.*

Proof. Assume $\mathcal{L}$ were more expressive than GF. Then there would be a signature τ and a formula $\varphi \in \mathcal{L}(\tau)$ that is not equivalent to any formula in GF(τ). Since $\mathcal{L}$ is compact, there are two τ-structures $\mathfrak{M}, \bar{a}$ and $\mathfrak{N}, \bar{b}$ with parameters such that $\mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{M}, \bar{a}) = \mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{N}, \bar{b})$ yet $\mathfrak{M}, \bar{a} \models \varphi$ and $\mathfrak{N}, \bar{b} \models \neg\varphi$ (cf. proposition 14).

According to corollary 24 there are two GF-saturated structures $\mathfrak{M}^*$ and $\mathfrak{N}^*$ such that $\mathit{Th}_{\mathcal{L}(\tau)}(\mathfrak{M}, \bar{a}) = \mathit{Th}_{\mathcal{L}(\tau)}(\mathfrak{M}^*, \bar{a})$ and $\mathit{Th}_{\mathcal{L}(\tau)}(\mathfrak{N}, \bar{b}) = \mathit{Th}_{\mathcal{L}(\tau)}(\mathfrak{N}^*, \bar{b})$. It follows that $\mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{M}^*, \bar{a}) = \mathit{Th}_{\mathrm{GF}(\tau)}(\mathfrak{N}^*, \bar{b})$. As both structures are GF-saturated, $\mathfrak{M}^*, \bar{a}$ and $\mathfrak{N}^*, \bar{b}$ are guarded bisimilar, by the Hennessy–Milner theorem. But by the $\mathcal{L}$-elementary nature of the extensions, still $\mathfrak{M}^*, \bar{a} \models \varphi$ while $\mathfrak{N}^*, \bar{b} \models \neg\varphi$, contradicting invariance under guarded bisimulation. ■

4 Concluding remarks

We have shown that ML[∀] and GF are maximally expressive logics whose semantics is invariant under the corresponding notion of bisimulation among compact logics with the Tarski Union Property (TUP).

The choice of TUP as a leading criterion may deserve some comment. Any choice of model theoretic criteria in a characterisation of the proposed kind has to be argued in the light of the question "*What constitutes a good Lindström characterisation?*". Clearly Lindström characterisations are very sensitive to the particular conditions imposed; the setting of the stage involves a critical choice as to *which competing logics* are admitted. While this may partly be a matter of taste or of tradition, it is also clear that a proposal is the more creditable, the wider the class of competitors is a priori, and the more fundamental the individual constraints are in the broader context of abstract model theory. There is no claim that the choices we made here are optimal in any sense. The following discussion is merely meant to indicate the setting in which this choice is being made, and thus points to some considerations that led us to favour the Tarski Union Property as a reasonably natural choice in a situation where compactness and basic semantic invariance conditions alone are at least not known to suffice to pin down the logics in question.

As pointed out in [4, 5], Lindström characterisations are closely related to semantic characterisation theorems in the tradition of classical preservation results. For basic modal logic, this companion/precursor is van Benthem's classical characterisation of ML as the bisimulation invariant fragment of first-order logic [3]; see [10, 9, 6] for a locality based account of, for instance, ML and ML[∀] as bisimulation invariant fragments of FO. The key difference is, of course, that in a Lindström characterisation we usually do not want to assume any a priori inclusion in some background logic, certainly not inclusion in FO. Modal logics like basic modal logic ML itself, or its extension with a global modality ML[∀], or the guarded fragment GF, are to be characterised not just *as fragments* of FO, but rather within the family of all logics that respect the same fundamental semantic invariance condition (bisimulation, global bisimulation, or guarded bisimulation invariance), including in particular candidate logics that are incomparable with FO. In this context, it is useful to recall Karp's theorem. All the semantic invariances considered concern equivalences that are

(a) bounded by partial isomorphism $\simeq_{\text{part}}$.[3]

(b) game based in the sense that equivalence corresponds to the existence of a winning strategy in infinite plays of some Ehrenfeucht–Fraïssé type model theoretic game.[4]

By Karp's theorem, $\simeq_{\text{part}}$ coincides with $\equiv_{\infty\omega}$, i.e., with equivalence in the

[3]We say that one equivalence relation is bounded by another, if the classes of the former are unions of classes of the latter, i.e., if the former can only be coarser than the latter.

[4]In the cases at hand: the ordinary bisimulation game, its variant with global moves, or the guarded bisimulation game.

infinitary logic $L_{\infty\omega}$, the extension of FO that allows infinite disjunctions and conjunctions rather than just finite ones. Therefore, (a) implies that equivalence $\equiv_{\mathcal{L}}$ for the candidate logics $\mathcal{L}$ is bounded by $\equiv_{\infty\omega}$.

Any infinitary game based equivalence as in (b), on the other hand, has natural finite approximations induced by existence of strategies in truncated games with a fixed finite number of rounds. By the corresponding variant of the classical Ehrenfeucht–Fraïssé theorem, these finitary approximate levels correspond – for finite vocabularies at least – to equivalences in fragments of the target logic that are finite up to logical equivalence and definable in FO. In the examples mentioned, these are the fragments of fixed finite nesting depths of the logics under consideration. In such a situation, a Hennessy–Milner–Karp connection between equivalence w.r.t. the common refinement of the finite approximation levels of the infinitary equivalence and full infinitary equivalence comes into focus. For basic modal logic, for instance, the classical Hennessy–Milner theorem tells us that equivalence in basic modal logic, i.e., equivalence w.r.t. all finite levels of n-bisimulation, guarantees full bisimulation equivalence for instance over the class of all finitely branching Kripke structures, and more generally over modally saturated Kripke structures.

Following an approach outlined in [11], one could sum up the key parameters in this setting as follows:

We consider a target logic $L = \bigcup_{\ell\in\omega} L_\ell \subseteq \mathrm{FO}$, stratified into syntactic levels L_ℓ such that equivalence w.r.t. L_ℓ has finite index and is captured by the ℓ-round game equivalence $\leftrightharpoons^\ell$, which is in FO for every fixed finite vocabulary. It follows that equivalence w.r.t. L, $\equiv_L$, is captured by the common refinement of all its finite approximations, $\leftrightharpoons^\omega := \bigcap_\ell \leftrightharpoons^\ell$. Moreover, $\leftrightharpoons^\omega$ coincides with $\leftrightharpoons$ at least over ω-saturated structures, since here 'good responses in every finite game' constitute (partial) types, whose realisations yield 'good responses for the infinite game'. (This reasoning is pursued in [11] in terms of first-order interpretations translating between the games for $\leftrightharpoons$ and the ordinary bisimulation game, also at the level of their finite approximations.)

We want to characterise L as maximal among certain well-behaved logics $\mathcal{L}$ whose semantics is invariant under $\leftrightharpoons$, the full infinitary game equivalence. Setting aside the issue of finite vocabularies (a finite occurrence property may be stipulated explicitly, or may be derivable from compactness assumptions), there are the following two essential hurdles in showing that L is maximally expressive among all $\leftrightharpoons$ invariant logics $\mathcal{L} \supseteq L$ satisfying some additional model theoretic criteria:

(i) the gap between $\leftrightharpoons^\omega$ (which is the same as $\equiv_L$) and $\leftrightharpoons$ (which we know to be a refinement of $\equiv_{\mathcal{L}}$, by the fundamental assumption of $\leftrightharpoons$ invariance of $\mathcal{L}$). Here Hennessy–Milner–Karp is useful: $\equiv_L$ ($\leftrightharpoons^\omega$) coincides with $\leftrightharpoons$ and hence with $\equiv_{\mathcal{L}}$ at least for ω-saturated structures.

(ii) the gap between showing that $\equiv_{\mathcal{L}}$ is bounded by $\equiv_L$ and showing that $\mathcal{L} \subseteq L$: this gap can typically be bridged by a compactness argument. Clearly compactness is a most natural criterion in the context of a

Lindström characterisation of a compact logic.

It follows that, in the light of (i), any additional assumption on $\mathcal{L}$ that guarantees the existence of $\mathcal{L}$-equivalent companions that are ω-saturated allows us to upgrade $\equiv_L$ to $\equiv_{\mathcal{L}}$. If F is some model transformation that preserves $\mathcal{L}$ and hence in particular $\equiv_L$ and produces ω-saturated companions, then $\mathfrak{M} \equiv_L \mathfrak{N}$ implies $F(\mathfrak{M}) \equiv_L F(\mathfrak{N})$, and hence $F(\mathfrak{M}) \equiv_{\mathcal{L}} F(\mathfrak{N})$, since $\leftrightarrows^{\omega}$ coincides with $\leftrightarrows$ in ω-saturated structures. So $\mathfrak{M} \equiv_{\mathcal{L}} F(\mathfrak{M}) \equiv_{\mathcal{L}} F(\mathfrak{N}) \equiv_{\mathcal{L}} \mathfrak{N}$ shows that $\mathfrak{M} \equiv_{\mathcal{L}} \mathfrak{N}$ for any $\mathfrak{M} \equiv_L \mathfrak{N}$.

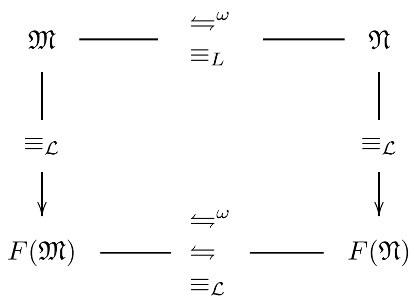

Clearly, ω-saturation can, in typical concrete instances, be replaced by weaker, specifically adapted notions of saturation (for instance, modal saturation suffices in the case of basic modal logic and bisimulation). Indeed, we here used the Tarski Union Property and compactness to establish the availability of not necessarily ω-saturated companions, but companions that are 'sufficiently saturated in the sense of $\leftrightarrows$', as limits of suitable elementary chains. Here 'sufficiently saturated in the sense of $\leftrightarrows$' really was sufficiency for the purpose of a Hennessy–Milner argument for the passage from $\leftrightarrows^{\omega}$ to $\leftrightarrows$.

Of course, other model theoretic criteria can serve the same purpose. In particular, in the spirit of de Rijke [14], preservation of $\mathcal{L}$ under (countable) ultrapowers produces ω-saturated (even ω_1-saturated) companions and hence immediately clinches the argument. If moreover, preservation under ultraproducts is assumed, then even compactness follows and needs not be stipulated separately. But even just preservation under countable ultrapowers immediately shows that $\equiv_{\mathcal{L}}$ is bounded by elementary equivalence. Since passage to ultrapowers in this case preserves both $\mathcal{L}$ and FO, it also gives us an upgrading of $\equiv$ to $\simeq_{\text{part}}$ and hence to $\equiv_{\mathcal{L}}$. In a sense, therefore, preservation assumptions of this calibre go some way in reducing Lindström characterisations to the semantic characterisations of fragments of FO in the style of a classical preservation theorem.

Our Lindström characterisations are meant to be different in this respect: we did not want to assume any a priori guarantees for the logics $\mathcal{L}$ under consideration to be fragments of FO or even for $\equiv_{\mathcal{L}}$ to be bounded by elementary equivalence. (As pointed out above, though, all equivalences considered are bounded by $\simeq_{\rm part}$ and hence by $\equiv_{\infty\omega}$.) Nevertheless, our arguments are also based on an upgrading of $\equiv_L$ ($\rightleftharpoons^\omega$) to $\equiv_{\mathcal{L}}$ ($\rightleftharpoons$) in suitably saturated companion structures $F(\mathfrak{M}) \equiv_{\mathcal{L}} \mathfrak{M}$, but since our model transformations, unlike ultrapower constructions, do not guarantee $F(\mathfrak{M}) \equiv \mathfrak{M}$, the same argument does not upgrade elementary equivalence $\equiv$ to $\equiv_{\mathcal{L}}$. In other words, $\mathcal{L}$-equivalence is only seen to be coarser than elementary equivalence a posteriori, because the target logic L happens to be a fragment of FO.

As already mentioned, it remains open whether a characterisation in the style of van Benthem's [4] is also available for ML[∀] or GF. We do not even know whether compactness, the appropriate notion of bisimulation invariance and possibly the relativisation property (and/or some other innocuous condition) might imply TUP.

BIBLIOGRAPHY

[1] H. Andréka, J. van Benthem, and I. Németi: Modal languages and bounded fragments of predicate logic. Journal of Philosophical Logic, 27(3), pp. 217-274, 1998.

[2] J. Barwise and S. Feferman (eds.): Model-Theoretic Logics. Springer, 1985.

[3] J. van Benthem: Modal Logic and Classical Logic. Bibliopolis, Napoli, 1983.

[4] J. van Benthem: A new modal Lindström theorem. Logica Universalis, 1(1), pp. 125-138, 2007.

[5] J. van Benthem, B. ten Cate, and J. Väänänen: Lindström theorems for fragments of first-order logic. Proceedings of the 22nd IEEE Symp. on Logic in Computer Science (LICS 07), pp. 280-292, 2007.

[6] V. Goranko and M. Otto: Model theory of modal logic, in Handbook of Modal Logic, P. Blackburn et al. (eds.), Elsevier, 2006.

[7] E. Grädel: On the restraining power of guards. Journal of Symbolic Logic, 64, pp. 1719-1742, 1999.

[8] E. Grädel, C. Hirsch, and M. Otto: Back and forth between guarded and modal logics. ACM Transactions on Computational Logic, 3, pp. 418-463, 2002.

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

[10] M. Otto: Elementary proof of the van Benthem–Rosen characterisation theorem. Technical report 2342, Fachbereich Mathematik, Technische Universität Darmstadt, 2004.

[11] M. Otto: Model theoretic methods for fragments of FO and special classes of (finite) structures. Draft of survey on the basis of a tutorial given at the 2006 Durham workshop on Finite and Algorithmic Model Theory (Isaac Newton Institute, Logic and Algorithms, 2006), 2007.

[12] R. Piro: Lindströmsche Sätze für modale Logiken. Diploma thesis, Technische Universität Darmstadt, 2008.

[13] M. de Rijke: Extending Modal Logic. Dissertation, Institute for Logic, Language and Computation, University of Amsterdam, 1993.

[14] M. de Rijke: A Lindström theorem for modal logic. In: Modal Logic and Process Algebra – A Bisimulation Perspective, A. Ponse, M. de Rijke and Y. Venema (eds), CSLI Publications, 1995, pp. 217–230.

Martin Otto and Robert Piro
Department of Mathematics
Technische Universität Darmstadt
64289 Damstadt, Germany
otto@mathematik.tu-darmstadt.de and robert.piro@gmx.de