

LINDSTRÖM THEOREMS FOR FRAGMENTS OF FIRST-ORDER LOGIC*

JOHAN VAN BENTHEM^a, BALDER TEN CATE^b, AND JOUKO VÄÄNÄNEN^c

^{a,c} ILLC, University of Amsterdam
e-mail address: {johan,vaananen}@science.uva.nl

^b ISLA, University of Amsterdam
e-mail address: balder.tencate@uva.nl

ABSTRACT. Lindström theorems characterize logics in terms of model-theoretic conditions such as Compactness and the Löwenheim-Skolem property. Most existing characterizations of this kind concern extensions of first-order logic. But on the other hand, many logics relevant to computer science are fragments or extensions of fragments of first-order logic, e.g., k -variable logics and various modal logics. Finding Lindström theorems for these languages can be on coding arguments that seem to require the full expressive power of first-order logic.

In this paper, we provide Lindström theorems for several fragments of first-order logic, including the k -variable fragments for $k > 2$, Tarski's relation algebra, graded modal logic, and the binary guarded fragment. We use two different proof techniques. One is a modification of the original Lindström proof. The other involves the modal concepts of bisimulation, tree unraveling, and finite depth. Our results also imply semantic preservation theorems.

1. INTRODUCTION

There are many ways to capture the expressive power of a logical language \mathcal{L} . For instance, one can characterize \mathcal{L} as being a model-theoretically well behaved fragment of a richer language \mathcal{L}' (a *preservation theorem*), or as being maximally expressive while satisfying certain model-theoretic properties (a *Lindström theorem*). The main contribution of this paper is a series of Lindström theorems for fragments of first-order logic. We also show connections between our Lindström theorems and preservation theorems.

The original Lindström theorem for first-order logic, in one of its most widely used formulations, says the following:

1998 ACM Subject Classification: F.4.1, F.4.3.

Key words and phrases: Abstract model theory, Lindström theorems, first-order logic, modal logic, guarded fragment.

* An extended abstract of this paper [20] has appeared in the proceedings of LICS 2007. We would like to thank Martin Otto and the anonymous referees for helpful comments. This work is partially supported by the Netherlands Organization for Scientific Research (NWO) grant 639.021.508.




Figure 1: Fragments of first-order logic

The first-order Lindström Theorem [13] *An extension of first-order logic satisfies Compactness and the Löwenheim-Skolem property if and only if it is no more expressive than first-order logic.*

There are several other versions of the theorem, characterizing first-order logic for instance in terms of Compactness and invariance for potential isomorphisms (the ‘Karp Property’). Analogues of this result have been obtained for various extensions of first-order logic. On the other hand, few Lindström theorems are known for *fragments* of first-order logic. One notable example is Van Benthem’s recent Lindström theorem for modal logic:

The modal Lindström theorem [24] *An extension of basic modal logic satisfies Compactness and Bisimulation Invariance if and only if it is no more expressive than basic modal logic.*¹

We are not aware of a similar characterization for modal logic involving the Löwenheim-Skolem property. Note that first-order logic itself is a compact proper extension of modal logic that has the latter property.

Our primary motivation for considering fragments comes from computer science logic. Many logics relevant to computer science are fragments (or non-FO extensions of fragments) of first-order logic, for example k -variable logics and various modal logics. Finding Lindström theorems for such languages can be a challenging problem, since most techniques used in the past to prove Lindström theorems rely on coding arguments that seem to require the full expressive power of first-order logic. For a recent survey of Lindström theorems in a general setting, see [10]. Another motivation is that, by widening our domain of study to logics not necessarily extending first-order logic, we may come to know more about first-order logic itself.

We follow two global lines of attack. First, we take the original Lindström theorem for first-order logic and generalize the proof as much as possible. In this way, we obtain Lindström theorems for the finite variable fragments FO^k with $k > 2$ and Tarski’s relation algebra. Next, we take the modal Lindström theorem as a starting point, and try to generalize it to richer languages. In this way, we obtain Lindström theorems for graded modal logic (on arbitrary Kripke structures and on trees) and the binary guarded fragment.

Many open questions remain. For example, we have not been able to find Lindström theorems for the two-variable fragment or the full guarded fragment.

¹By *extensions of basic modal logic* we mean language extensions, not axiomatic extensions.

2. FROM FIRST-ORDER LOGIC DOWNWARDS

In this first part, we take the classic Lindström theorem as a starting point, and we show that the proof generalizes to certain fragments of first-order logic.

2.1. A strengthening of the Lindström theorem for first-order logic over binary vocabularies. The first-order Lindström theorem uniquely characterizes first-order logic in terms of Compactness and the Löwenheim-Skolem property *within the class of all its extensions*. As we will show in this section, this result can be improved: first-order logic can be characterized in terms of Compactness and the Löwenheim-Skolem property *within the class of all extensions of the three-variable fragment FO^3* , if we consider vocabularies consisting only of unary and binary relation symbols. The proof is not substantially more difficult than that of the original Lindström theorem, requiring mainly some extra care with coding partial isomorphisms and back-and-forth behaviour. But this simple strengthening will allow us to obtain a number of new results on Tarski's relation algebra, as well as finite variable fragments.

To keep things simple, we will work with a fixed relational signature consisting of a set of unary relation symbols and a set of binary relation symbols, both countably infinite.

By an *abstract logic* we mean a pair $\mathcal{L} = (\text{Fml}_{\mathcal{L}}, \models_{\mathcal{L}})$, where $\text{Fml}_{\mathcal{L}}$ is the set of sentences of \mathcal{L} and $\models_{\mathcal{L}}$ is a binary truth relation between \mathcal{L} -sentences and models. If no confusion arises, we will sometimes write \mathcal{L} for $\text{Fml}_{\mathcal{L}}$ and \models for $\models_{\mathcal{L}}$. We assume that \mathcal{L} -sentences are preserved under isomorphisms, and that \mathcal{L} has the following properties:

- ▶ *closure under Boolean connectives*: for every $\phi \in \mathcal{L}$ there is a sentence $\psi \in \mathcal{L}$ expressing its negation (i.e., for all models M , $M \models \psi$ iff $M \not\models \phi$), and for every $\phi, \psi \in \mathcal{L}$ there is a sentence $\chi \in \mathcal{L}$ expressing the conjunction of ϕ and ψ .
- ▶ *closure under renamings*: for every mapping ρ sending relation symbols to relation symbols of the same arity, and for every sentence $\phi \in \mathcal{L}$, there is a sentence $\psi \in \mathcal{L}$ such that for all models M , $M \models \psi$ iff $\rho(M) \models \phi$.
- ▶ *closure under relativisation by unary predicates*: for every sentence $\phi \in \mathcal{L}$ and unary relation symbol P , there is a sentence $\psi \in \mathcal{L}$ such that for all models M , $M \models \psi$ iff $M^P \models \phi$, with M^P the submodel of M induced by the subset defined by P .

Examples of abstract logics in this sense include first-order logic (FO) and its k -variable fragments (FO^k), with $k \geq 1$.

Given two abstract logics, \mathcal{L} and \mathcal{L}' , we say that \mathcal{L} extends \mathcal{L}' (or, \mathcal{L}' is contained in \mathcal{L} , denoted by $\mathcal{L}' \subseteq \mathcal{L}$), if there is a map $f : \text{Fml}_{\mathcal{L}'} \rightarrow \text{Fml}_{\mathcal{L}}$ preserving truth in the sense that, for all models M and sentences $\phi \in \mathcal{L}'$, $M \models_{\mathcal{L}'} \phi$ iff $M \models_{\mathcal{L}} f(\phi)$.

An abstract logic \mathcal{L} has Compactness if for every set of \mathcal{L} -formulas Σ , if every finite subset of Σ is satisfiable then the entire set Σ is satisfiable. An abstract logic \mathcal{L} has the Löwenheim-Skolem property if every satisfiable set of \mathcal{L} -formulas has a countable model.

First, we show that each compact extension of FO (in fact, already of FO^2) has the “finite occurrence property”.

Lemma 2.1 (Finite occurrence property). *Let \mathcal{L} be any abstract logic extending FO^2 that has Compactness. Then for any $\phi \in \mathcal{L}$ there is a finite set of relation symbols $REL(\phi)$ such that the truth of ϕ in any model is independent of the denotation of relation symbols outside $REL(\phi)$.*

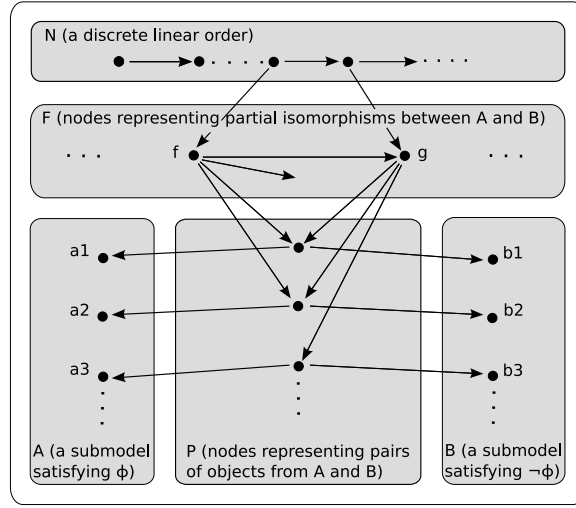


Figure 2: Model from the proof of Lemma 2.2.

Proof. A standard argument, cf. [8]: since our vocabulary contains infinitely many unary and binary relation symbols, there are renamings ρ_1, ρ_2 whose range is disjoint. Take any $\phi \in \mathcal{L}$, and let ϕ_1, ϕ_2 be its renamings according to ρ_1 and ρ_2 . Let $\Sigma = \{\forall x_1 \dots x_k (\rho_1(R)(x_1 \dots x_k) \leftrightarrow \rho_2(R)(x_1 \dots x_k)) \mid R \text{ a } k\text{-ary relation symbol}\}$. Then $\Sigma \models \phi_1 \leftrightarrow \phi_2$, and hence, by Compactness, a finite subset $\Sigma' \subseteq \Sigma$ implies $\phi_1 \leftrightarrow \phi_2$. We can pick for $REL(\phi)$ the relation symbols occurring in Σ' . Note that, in case of binary vocabularies, all formulas in Σ belong to FO^2 . \square

The key to the proof of our improved Lindström theorem is the following observation:

Lemma 2.2. *Let \mathcal{L} be any abstract logic with the Löwenheim-Skolem property and the Finite Occurrence Property, such that \mathcal{L} extends FO^3 and is not contained in FO . Then “ \mathcal{L} can relatively projectively define finiteness”: there is a formula $\psi \in \mathcal{L}$ with a unary predicate N such that, for each $n \in \mathbb{N}$, there is a model of ψ with exactly n elements satisfying N , while no model of ψ has infinitely many elements satisfying N .*

Proof. The basic idea is the same as in traditional proofs of the Lindström theorem (e.g., [8]). Our main contribution is to show that, in the case of binary vocabularies, the coding argument requires only three variables.

Take any $\phi \in \mathcal{L}$ not belonging to FO . Then for each $k \in \mathbb{N}$, there are models $\mathfrak{A}_k \models \phi$ and $\mathfrak{B}_k \not\models \phi$ that are potentially isomorphic up to back-and-forth depth k . At the same time, no potentially isomorphic models can disagree on the sentence ϕ . We can describe this situation *inside* \mathcal{L} . The construction is outlined in Figure 2.

The model depicted in Figure 2 describes two models, connected via a collection of partial isomorphisms, that disagree on the sentence ϕ . The most important feature is that, if N is an infinite set, then the collection of partial isomorphisms constitutes a potential isomorphism, whereas if N is finite (say, of size k), the collection of partial isomorphisms constitutes a potential isomorphism *up to back-and-forth depth k* .

More precisely, A and B are unary predicates defining the domains of two (sub)models, P is a unary predicate whose elements denote pairs from $A \times B$, and the elements of F represent partial isomorphisms (i.e., sets of pairs constituting structure preserving bijections). The arrows represent a binary relation R . For instance, in the given example, f represents the partial isomorphism $\{(a_1, b_1), (a_2, b_2)\}$, and g represents the partial isomorphism that extends f with the pair (a_3, b_3) . The elements of the linearly ordered set N will be used to associate a finite index to each partial isomorphism.

Claim: Each of the following properties of this model can be expressed by a sentence of \mathcal{L} :

- (1) Every $p \in P$ is associated to a pair from $A \times B$.
- (2) Every $f \in F$ is associated to a set of elements of P that form a partial *bijection* between A and B .
- (3) Each such partial bijection preserves structure on the submodels defined by A and B , as far as the finitely many relations occurring in ϕ are concerned (recall that \mathcal{L} has the Finite Occurrence Property).
- (4) Every $f \in F$ has an associated ‘index’ from N .
- (5) N is linearly ordered by R , such that there is a minimal element, and each non-maximal element has an immediate successor (in particular, if N is infinite, then it contains an infinite ascending chain).
- (6) If fRg for $f, g \in F$, this means that g extends f (as a partial bijection), and that the index of g is the successor of the index of f .
- (7) The *back-and-forth* properties hold for partial isomorphisms whose index is not the maximal element of N .
- (8) Some $f \in F$ has as its index the minimal element of N .
- (9) The submodels defined by A and B disagree on ϕ . (Recall that \mathcal{L} is closed under the Boolean connectives and relativisation by unary predicates).

Proof of claim: The first eight properties can already be expressed in FO^3 by a clever re-use of variables, and the ninth property can be expressed in \mathcal{L} by closure under the Boolean connectives and relativisation by unary predicates.

For instance, the third property can be expressed as the conjunction of all FO^3 -formulas of the following forms, where $S \in REL(\phi)$ is a binary relation symbol, and $Q \in REL(\phi)$ is a unary relation symbol:

$$\begin{aligned} \forall xy \left(Px \wedge Py \wedge \exists z (Fz \wedge Rzx \wedge Rzy) \rightarrow \left(\exists z (Rzx \wedge Az \wedge \exists x (Ryx \wedge Ax \wedge Szx)) \right. \right. \\ \left. \left. \leftrightarrow \exists z (Rzx \wedge Bz \wedge \exists x (Ryx \wedge Bx \wedge Szx)) \right) \right) \end{aligned}$$

and

$$\forall x \left(Px \rightarrow \left(\exists z (Rzx \wedge Az \wedge Qz) \leftrightarrow \exists z (Rzx \wedge Bz \wedge Qz) \right) \right)$$

Note that we crucially use the fact that the signature consists of unary and binary relations only. \square

Let χ be the conjunction of all these \mathcal{L} -sentences. By the above assumptions, χ has models in which N has arbitrarily large finite cardinality (this follows from the existence, for each $k \in \mathbb{N}$, of models disagreeing on ϕ that are potentially isomorphic up to back-and-forth depth k). However, there is no model of χ in which N is an infinite set. For, suppose there were. Let θ be a formula expressing that a fresh binary relation symbol S denotes a

linear order without maximal element on the set denoted by N , thus forcing N to denote an infinite set. Then $\chi \wedge \theta$ would also be satisfiable, hence, by the Löwenheim-Skolem property, it would have a countable model. In other words, χ would have a countable model in which N denotes an infinite set. In countable models, being potentially isomorphic means being isomorphic; thus, there would be isomorphic models disagreeing on ϕ , which contradicts \mathcal{L} 's invariance for isomorphisms.

In other words, χ relatively projectively defines finiteness. \square

Theorem 2.3. *An abstract logic extending FO^3 is contained in FO iff it satisfies both Compactness and the Löwenheim-Skolem property.*

Proof. If an abstract logic is contained in FO , then, clearly, it satisfies Compactness and the Löwenheim-Skolem property. If, on the other hand, an abstract logic \mathcal{L} extends FO^3 but is not contained in FO , then it must lack either the Löwenheim-Skolem property or Compactness. For, suppose that \mathcal{L} satisfies the Löwenheim-Skolem property and Compactness. Let $\psi(N)$ be any \mathcal{L} -sentence projectively defining finiteness (cf. Lemma 2.2), and for each $k \in \mathbb{N}$, let χ_k be an FO^3 -formula expressing that there are at least k distinct N -elements (by reusing bound variables as in $\exists x(Nx \wedge \exists y(Ny \wedge Rxy \wedge \exists x(Nx \wedge Ryx \wedge \dots))$). Then every finite subset of $\{\chi_k \mid k \in \mathbb{N}\} \cup \{\psi(N)\}$ has a model while the entire set has no model. This contradicts the Compactness property. \square

The following results can be proved in a similar fashion:

Theorem 2.4. *An abstract logic extending FO^3 is contained in FO iff it satisfies Compactness and invariance for potential isomorphisms.*

Proof sketch. An analogue of Lemma 2.2 can be proved where the Löwenheim-Skolem property is replaced by invariance for potential isomorphisms. Indeed, the proof is even shorter than the proof of Lemma 2.2 itself, as the last step can be dropped. Given this, we can proceed just as in the proof of Theorem 2.3. \square

Theorem 2.5. *A concrete abstract logic extending FO^3 is contained in FO iff it satisfies the Löwenheim-Skolem property and the Finite Occurrence Property, and is recursively enumerable for validity.*

Here, by “concrete” we mean that formulas can be coded as finite strings over some alphabet, in such a way that negation, conjunction, and relativisation are computable operations, and there is a computable translation from FO^3 -formulas to formulas of the logic. Theorem 2.5 follows from Lemma 2.2 by closure under relativisation and the fact that satisfiability of FO^3 formulas on finite models is undecidable [6], hence not recursively enumerable.

Note that these results all rely on our restriction to at most binary relation symbols. In the case with at most k -ary relations ($k \geq 1$), analogous results hold for FO^{k+1} . For the case of $k = 1$, the argument is slightly different. For the sake of completeness, we give an outline of it below (similarly, analogues of Theorems 2.4 and 2.5 can be shown for $k = 1$). See also the Lindström theorem for monadic first-order logic in [21].

Theorem 2.6. *Over unary vocabularies, an abstract logic extending FO^2 is contained in FO iff it satisfies both Compactness and the Löwenheim-Skolem property.*

Proof sketch. If an abstract logic is contained in FO , then, clearly, it satisfies Compactness and the Löwenheim-Skolem property. Conversely, let \mathcal{L} be any abstract logic extending

FO^2 that is not contained in FO . Let $\phi \in \mathcal{L}$ be non-first-order, i.e., for every $k \geq 0$, there are models M, N disagreeing on ϕ such that M and N are potentially isomorphic up to depth k .

By the same argument as in Lemma 2.1, we have that ϕ depends only on finitely many (unary) predicates P_1, \dots, P_n . Moreover, it can be easily seen that two structures are potentially isomorphic up to depth k for P_1, \dots, P_n if and only if for each unary type τ , either τ is realized the same number of times in the two structures, or in both structures τ is realized at least k times. By a *unary type* we mean here a function mapping P_1, \dots, P_n to truth values. Similarly, two structures are potentially isomorphic (for unbounded depth) for P_1, \dots, P_n if and only if for each unary type τ , either τ is realized the same number of times in the two structures, or in both structures τ is realized infinitely many times.

Now, pick fresh unary predicates Q, \bar{Q} and infinitely many fresh unary predicates $P_{\tau i}$, for τ a type and $i \geq 1$. The unary predicates Q, \bar{Q} will be used to define two submodels, which will disagree on ϕ , and the unary predicate $P_{\tau i}$ will be used to code a potential isomorphism. As we pointed out already, in order to test for the existence of a potential isomorphism it is enough to count the number of times each unary type is realized. The predicate $P_{\tau i}$, intuitively, will denote the i -th element of type τ in one of the two substructures, and they are used to compare the cardinalities between the two substructures. More specifically, consider the infinite theory Σ containing the following FO^2 -sentences:

$$\begin{aligned} &\phi^Q \wedge \neg\phi^{\bar{Q}} \wedge \forall x(\bar{Q}x \leftrightarrow \neg Qx) \\ &\forall x(P_{\tau i}(x) \rightarrow \tau(x)) \end{aligned}$$

$$\forall x(P_{\tau i}(x) \rightarrow \neg P_{\tau j}(x)) \quad \text{for } i \neq j$$

$$\begin{aligned} &\forall xy(Qx \wedge P_{\tau i}x \wedge Qy \wedge P_{\tau i}y \rightarrow x = y) \\ &\forall xy(\neg Qx \wedge P_{\tau i}x \wedge \neg Qy \wedge P_{\tau i}y \rightarrow x = y) \end{aligned}$$

$$\begin{aligned} &\exists x.(Qx \wedge P_{\tau(i+1)}x) \rightarrow \exists x.(Qx \wedge x.P_{\tau i}x) \\ &\exists x.(\neg Qx \wedge P_{\tau(i+1)}x) \rightarrow \exists x.(\neg Qx \wedge P_{\tau i}x) \end{aligned}$$

$$\begin{aligned} &\exists x(Qx \wedge Px \wedge \bigwedge_{1 \leq k < i} \neg P_{\tau k}x) \rightarrow \exists x.(Qx \wedge P_{\tau i}x) \\ &\exists x(\neg Qx \wedge Px \wedge \bigwedge_{1 \leq k < i} \neg P_{\tau k}x) \rightarrow \exists x.(\neg Qx \wedge P_{\tau i}x) \end{aligned}$$

$$\exists x(Qx \wedge P_{\tau i}x) \leftrightarrow \exists x(\neg Qx \wedge P_{\tau i}x)$$

for τ any type and $i, j \geq 1$, and where $\tau(x)$ is shorthand for $\bigwedge_{\tau(P)=1} P(x) \wedge \bigwedge_{\tau(P)=0} \neg P(x)$. The first sentence expresses that the two substructures disagree on ϕ . The instances of the 2nd – 9th sentence above force the following: for all types τ and natural numbers $i > 0$, a substructure (induced by Q or by $\neg Q$) contains an element satisfying $P_{\tau i}$ if and only if it contains at least i elements of type τ . The instances of the last sentence above, finally, compare the cardinalities between the two structures. It follows that the theory Σ describes models in which the submodels induced by Q and $\neg Q$ disagree on ϕ and are potentially isomorphic.

Incidentally, note that if there are infinitely many elements of type τ , the theory Σ cannot force each to satisfy some predicate $P_{\tau i}$ (as can be seen by an omitting types argument), but it does force there to be infinitely many elements satisfying some predicate $P_{\tau i}$, which suffices for present purposes.

Since, for every $k \geq 0$, there are models disagreeing on ϕ such that the two are potentially isomorphic up to depth k , it can be shown that every finite subset of Σ is satisfiable. It follows that \mathcal{L} cannot have both Compactness and Löwenheim-Skolem, for, the entire theory Σ would have a countable model, and on countable structures, potential isomorphism implies isomorphism. \square

2.2. First application: Tarski's relation algebra. Tarski's relation algebra \mathcal{RA} [17] is an algebraic language in which the terms denote binary relations. It has atomic terms R, S, \dots ranging over binary relations (over some domain), constants δ and \top denoting the identity relation and the total relation, and operators $\cap, -, \bullet, \smile$ for taking the intersection, complement, composition and converse of relations. Thus, the syntax of \mathcal{RA} is:

$$\alpha ::= R \mid \delta \mid \top \mid \alpha \cap \beta \mid \alpha - \beta \mid \alpha \bullet \beta \mid \alpha \smile$$

with R an element from some countably infinite set of variables standing for binary relations. An interpretation for this language is a set X together with an assignment of binary relations over X to the atomic terms. In other words, it is a first-order structure for the vocabulary that contains the atomic terms as binary relation symbols. We write $\alpha \equiv \beta$ if, in each interpretation, α and β denote the same binary relation, and we write $\alpha \subseteq \beta$ if, in each interpretation, α denotes a subrelation of the relation denoted by β .

In this section, we provide a Lindström-theorem for extensions of relation algebra. By an *extended relation algebra* we will mean any language obtained by extending the syntax of relation algebra with zero or more additional logical operations, where a *logical operation* is any operation that takes as input a fixed finite number of binary relations R_1, \dots, R_n (over some set X) and that produces a new binary relation S over the same set. We also require logical operations to respect isomorphisms, and to be *domain independent* in the following sense (familiar from database theory): the output relation S only depends on part of the domain X that participates in at least one pair belonging to an input relation R_i . An example of a domain dependent operation is the (unary) absolute complement operator, while the (binary) relative complement operator is domain independent. Another example of a domain dependent operator is the zero-ary operator that returns the total relation in case the domain X is infinite, and the empty relation otherwise. We disallow domain dependent operations, as they may cause the language to lose closure under relativisation.

Lemma 2.7 (Closure under relativisation). *Let \mathcal{L} be any extended relation algebra, let α be any expression of \mathcal{L} , and let R be any relation variable. Then there is an expression of \mathcal{L} , which we denote by $\alpha^{dom(R)}$ such that in every interpretation M , the denotation of $\alpha^{dom(R)}$ in M is the denotation of α in M' , where M' is the submodel of M induced by the domain of R^M .*

Proof. First, note that the expression $\top_R := R \bullet \top \bullet R \smile$ denote the total relation over the domain of R . Now, $\alpha^{dom(R)}$ is obtained by uniformly replacing every atomic expression α (i.e., the relational variables, as well as \top and δ) by $\alpha \cap \top_R$. A straightforward formula induction shows that $\alpha^{dom(R)}$ satisfies the required conditions. The inductive step for complex expressions uses the fact that all operations of \mathcal{L} are domain independent. \square

One example of an extended relation algebra, \mathcal{RA}_{FO} , is the extension of relation algebra with *all first-order definable logical operations* (see e.g. [25]). Another example is \mathcal{RA}_{TC} , the extension of relation algebra with the *transitive closure* operation [15].

The Compactness and Löwenheim-Skolem properties can be defined for extended relation algebras as usual. For instance, we say that an extended relation algebra \mathcal{L} has the Löwenheim-Skolem property if for every set of \mathcal{L} -expressions Φ , if there is an interpretation under which $\bigcap \Phi$ is a non-empty relation, then there is such an interpretation over a countable domain. As is not hard to see, \mathcal{RA} and \mathcal{RA}_{FO} satisfy both Compactness and the Löwenheim-Skolem property, whereas \mathcal{RA}_{TC} satisfies the Löwenheim-Skolem property but lacks Compactness.

The following Lindström-style theorem shows that all extended relation algebras containing non-first-order definable operations lack either Compactness or the Löwenheim-Skolem property.

Theorem 2.8. *Let \mathcal{L} be any extended relation algebra with the Compactness and Löwenheim-Skolem properties. Then every logical operation of \mathcal{L} is first-order definable.*

Proof. Lemma 2.2 can be adapted to the relation algebra setting, allowing us to show that every extended relation algebra containing a non-elementary logical operation and having the Löwenheim-Skolem property can projectively define finiteness, and hence lacks Compactness. We will not spell out the details, but merely mention the following key points of the proof:

- ▶ For every first-order sentence ϕ containing only three variables, in a signature consisting only of binary relations, there is a relation algebra expression α such that for every model M , $M \models \phi$ iff $\alpha \equiv \top$ holds in M [17].
- ▶ Unary relations can be mimicked by binary relations by systematically intersecting them with the identity relation.
- ▶ In this way, every extended relation algebra gives rise to an abstract logic extending FO^3 . Closure under relativisation of the logic is guaranteed by the domain independence of the logical operations of \mathcal{L} . □

In other words, the extension of relation algebra with *all* elementary operations (a well-known structure in algebraic logic) is the greatest extension that satisfies Löwenheim-Skolem and Compactness. The same holds if we replace the Löwenheim-Skolem property by invariance for potential isomorphisms, or if we replace Compactness by recursive enumerability.

Theorem 2.8 nicely complements a known result: *every extended relation algebra with Craig interpolation can define all first-order definable operations* [18]. Together, these results show that \mathcal{RA}_{FO} is the unique (up to expressive equivalence) extension of \mathcal{RA} satisfying Compactness, Löwenheim-Skolem, and Craig Interpolation.

2.3. Second application: finite variable fragments. In this section, we provide Lindström theorems for the finite variable fragments FO^k with $k > 3$, over vocabularies consisting of unary and binary relation symbols only. It is well known that the finite variable fragments can be characterized as fragments of first-order logic using potential isomorphisms with a restricted number of pebbles:

Definition 2.9. (k -pebble potential isomorphisms) A k -pebble potential isomorphism between M and N is a non-empty family F of finite partial isomorphisms f between M and N with $|\text{dom}(f)| \leq k$ satisfying the following properties:

- **Closure under restrictions:** for all $f \in F$ and $f' \subseteq f$, also $f' \in F$.
- **Forth:** for all $\{(w_1, v_1), \dots, (w_n, v_n)\} \in F$ with $n < k$, and $w \in M$, there is an $v \in N$ such that $\{(w_1, v_1), \dots, (w_n, v_n), (w, v)\} \in F$
- **Back:** for all $\{(w_1, v_1), \dots, (w_n, v_n)\} \in F$ with $n < k$, and $v \in N$, there is an $w \in M$ such that $\{(w_1, v_1), \dots, (w_n, v_n), (w, v)\} \in F$

Fact 2.10. FO^k is (up to logical equivalence) the fragment of FO invariant for k -pebble potential isomorphisms ($k \geq 1$).

For a proof of Fact 2.10 see for instance [1].

Using Lemma 2.2, we can turn this into the following Lindström characterization (remember that we only consider unary and binary relation symbols):

Theorem 2.11 (Lindström theorem for FO^k). *Let $k \geq 3$. An abstract logic extending FO^k satisfies Compactness and invariance for k -pebble potential isomorphisms iff it is no more expressive than FO^k .*

Proof. Consider any abstract logic \mathcal{L} extending FO^3 that has Compactness and is invariant for k -pebble potential isomorphisms. Then in particular it is invariant for potential isomorphisms, and therefore by Theorem 2.4 it must be contained in FO . But then, by Fact 2.10, it must be contained in FO^k . \square

Theorem 2.11 can also be seen as a strengthening of the classical preservation result in Fact 2.10 (for $k \geq 3$, and on binary vocabularies). Indeed, it implies (and in fact is equivalent to, as we will explain) the following “generalized preservation theorem”:

Corollary 2.12. *Let $k \geq 3$. Let \mathcal{L} be any abstract logic extending FO^k with Compactness. Then FO^k is the fragment of \mathcal{L} invariant for k -pebble potential isomorphisms (up to logical equivalence). In particular, FO^k is the fragment of FO invariant for k -pebble potential isomorphisms.*

Proof. Let \mathcal{L} be any compact abstract logic extending FO^k , and let \mathcal{L}' be the fragment of \mathcal{L} invariant for k -pebble potential isomorphisms. Then \mathcal{L}' satisfies all the requirements of abstract logics. For instance, it is closed under relativisation: if $\phi \in \mathcal{L}$ is invariant for k -pebble potential isomorphisms, then so is its relativisation by a unary predicate. Thus, \mathcal{L}' is an abstract logic extending FO^k that is compact and invariant for k -pebble potential isomorphisms. Hence, by Theorem 2.11, it is contained in FO^k . \square

In fact, Corollary 2.12 also implies Theorem 2.11: let \mathcal{L} be any logic extending FO^k satisfying Compactness and invariance for k -pebble potential isomorphisms. By Corollary 2.12, FO^k is the the fragment of \mathcal{L} invariant for k -pebble potential isomorphisms. But since \mathcal{L} itself is invariant for k -pebble potential isomorphisms, this means that \mathcal{L} and FO^k coincide in terms of expressive power. This shows that Corollary 2.12 and Theorem 2.11 are really equivalent to each other: Lindström theorems and preservation theorems sometimes amount to the same thing.

3. FROM MODAL LOGIC UPWARDS

The approach of generalizing the classic Lindström theorem only got us so far. It enabled us to characterize FO^k for $k \geq 3$ but is unlikely to reach much further down, since coding power then falls under the minimum needed to describe partial isomorphisms and their extension properties. Thus, we will now take a different approach, by considering the modal Lindström theorem, and trying to generalize it to richer languages. In particular, we obtain two new Lindström theorems for the graded modal logic.

3.1. The modal Lindström theorem revisited. We recall the proof of the modal Lindström theorem of [24] (which improves on an earlier result in [7]). First, we need to define “abstract modal logics”. As before, we assume a fixed vocabulary, now consisting of a single binary relation symbol R and a countably infinite set of unary relation symbols, also called *proposition letters*. Structures for this vocabulary are usually called *Kripke models* (the restriction to a single binary relation symbol is not essential but is convenient for presentation). We associate to each formula a class of pairs (M, w) , where M is a Kripke model and w is an element of the domain of M . This is because modal formulas are always evaluated *at a point in a model*. We call such pairs (M, w) *pointed Kripke models*. Thus, an *abstract modal logic* is a pair $\mathcal{L} = (Fml_{\mathcal{L}}, \models)$, where $Fml_{\mathcal{L}}$ is the set of formulas of \mathcal{L} and \models is a binary relation between \mathcal{L} -formulas and pointed Kripke models. As before, when no confusion arises, we will write \mathcal{L} for $Fml_{\mathcal{L}}$. Also, as before, we assume that \mathcal{L} -formulas are invariant for isomorphisms, and that \mathcal{L} is closed under the Boolean operations, renaming, and relativisation by unary predicates.²

Examples of abstract modal logics include *basic modal logic*, its extension with counting modalities called *graded modal logic (GML)*, first-order logic (by which we mean the collection of all first-order formulas in one free variable, over the appropriate signature), and the modal μ -calculus. For the syntax and semantics of basic modal logic, see any modal logic textbook (e.g., [5]).

The modal Lindström theorem characterizes basic modal logic in terms of Compactness and invariance for *bisimulations*. A bisimulation between Kripke models M and N is a binary relation Z between the domains of M and N satisfying the following three conditions:

- **Atomic harmony:** if wZv then w and v agree on all proposition letters (unary predicates).
- **Zig:** if wZv and $wR^M w'$, there is a v' such that $vR^N v'$ and $w'Zv'$.
- **Zag:** if wZv and $vR^N v'$, there is a w' such that $wR^M w'$ and $w'Zv'$.

Two pointed Kripke models (M, w) and (N, v) are *bisimilar* if there is a bisimulation Z between M and N with wZv . A formula is *bisimulation invariant* if it does not distinguish bisimilar pointed Kripke models, and an abstract modal logic is bisimulation invariant if all its formulas are.

Given a pointed Kripke model (M, w) , we denote by M_w the submodel of M containing all points that are reachable in finitely many steps from w along the binary relation. Likewise, for $k \in \mathbb{N}$, M_w^k is the submodel of M containing all points reachable from w in at

²So far, there is nothing particularly ‘modal’ about these systems, and we might also speak of “abstract local logics”, or some other appealing terminology. In particular, we do *not* replace isomorphism invariance by ‘modal’ invariance for bisimulations in the definition of these logics. The power of the latter condition is precisely one of the things we want to scrutinize in what follows.

most k steps along the binary relation. We say that a formula ϕ is *invariant for generated submodels* if, for all models M with nodes w , $(M, w) \models \phi$ iff $(M_w, w) \models \phi$. We say that ϕ *has the finite depth property* if there is a $k \in \mathbb{N}$ such that $(M, w) \models \phi$ iff $(M_w^k, w) \models \phi$, for all models M with nodes w . Clearly, the latter implies the former. Also, bisimulation invariance implies invariance for generated submodels, because the natural inclusion map is a bisimulation. An abstract modal logic \mathcal{L} is invariant for generated submodels (or, has the finite depth property), if every $\phi \in \mathcal{L}$ is invariant for generated submodels (respectively, has the finite depth property).

We are now ready to proceed with the proof of the modal Lindström theorem, Theorem 3.3 below. We first prove a finite occurrence property for compact extensions of basic modal logic that satisfy invariance for bisimulation (in fact, invariance for generated submodels will do).

Lemma 3.1 (Finite occurrence property). *Let \mathcal{L} be any abstract modal logic extending basic modal logic that is compact and invariant for generated submodels. Then for each $\phi \in \mathcal{L}$, there is a finite set of proposition letters $PROP(\phi)$ such that the truth of ϕ in any pointed Kripke model is independent of the denotation of proposition letters outside $PROP(\phi)$.*

Proof. Like that of Lemma 2.1. Since the set of proposition letters (unary predicates) is infinite, we can find renamings ρ_1, ρ_2 for them whose range is disjoint. Now, take any $\phi \in \mathcal{L}$, and let ϕ_1, ϕ_2 be its renamings according to ρ_1 and ρ_2 . Let $\Sigma = \{\Box^n(\rho_1(p) \leftrightarrow \rho_2(p)) \mid n \in \mathbb{N} \text{ and } p \text{ a proposition letter}\}$, where \Box^n stands for a sequence of n boxes. It follows from generated submodel invariance that $\Sigma \models \phi_1 \leftrightarrow \phi_2$, hence, by Compactness, a finite subset $\Sigma' \subseteq \Sigma$ implies $\phi_1 \leftrightarrow \phi_2$. We can pick for $PROP(\phi)$ the set of proposition letters occurring in Σ' . \square

Our crucial observation in this Lindström argument is the following:

Lemma 3.2 ([24]). *Let \mathcal{L} be any abstract modal logic extending basic modal logic that is compact and invariant for generated submodels. Then \mathcal{L} has the finite depth property.*

Proof. Let $\phi \in \mathcal{L}$, with p a proposition letter not occurring in ϕ , and let ϕ^p be the relativisation of ϕ by p . By the generated submodel-invariance of \mathcal{L} , $\{p, \Box p, \Box \Box p, \dots\} \models \phi \leftrightarrow \phi^p$. By the compactness of \mathcal{L} , there is an $n \in \mathbb{N}$ such that $\{p, \Box p, \dots, \Box^n p\} \models \phi \leftrightarrow \phi^p$. But this expresses exactly that ϕ has the finite depth property, for depth n . \square

Theorem 3.3 ([24]). *An abstract modal logic extending basic modal logic satisfies Compactness and bisimulation invariance iff it is no more expressive than basic modal logic.*

Proof. Let \mathcal{L} be any abstract modal logic extending basic modal logic and satisfying Compactness and bisimulation invariance. Since bisimulation invariance implies invariance for generated submodels, \mathcal{L} is also invariant for generated submodels. Then, by Lemma 3.1 and Lemma 3.2, it follows that \mathcal{L} has the finite occurrence property and the finite depth property. Next, we use the following well known fact [5]:

Assuming a finite vocabulary, every bisimulation-invariant class of pointed models with the finite depth property is definable by a formula of basic modal logic.

We conclude that \mathcal{L} is contained in basic modal logic. \square

Theorem 3.3 can be seen as a strengthening of van Benthem’s more familiar characterization of basic modal logic as the bisimulation invariant fragment of first-order logic. Indeed, it implies the latter theorem (as observed in [24]), but something even stronger holds, viz. the following “generalized preservation theorem”:

Corollary 3.4. *Let \mathcal{L} be any abstract modal logic extending basic modal logic that has Compactness. Then basic modal logic is the bisimulation invariant fragment of \mathcal{L} (up to logical equivalence). In particular, basic modal logic is the bisimulation invariant fragment of first-order logic.*

Proof. Let \mathcal{L}' be the bisimulation invariant fragment of \mathcal{L} . Then \mathcal{L}' satisfies all criteria for being an abstract modal logic. For instance it is closed under relativisation: whenever $\phi \in \mathcal{L}$ is invariant for bisimulations then the relativisation of ϕ by a unary predicate is also invariant for bisimulations. Likewise for the Boolean connectives.

Hence, \mathcal{L}' is an abstract modal logic extending basic modal logic and it is bisimulation invariant and Compact. Hence, it is no more expressive than basic modal logic. \square

Corollary 3.4 strengthens the traditional bisimulation preservation theorem, as there are compact extensions of basic modal logic not contained in first-order logic. Indeed:

Theorem 3.5. *There is an abstract modal logic extending basic modal logic that is not contained in first-order logic, but still satisfies Compactness, the Löwenheim-Skolem property, invariance for potential isomorphisms, and Craig interpolation. Moreover, it has a finite axiomatization and a PSPACE-complete satisfiability problem.*

The proof is given in the next section. Note that Theorem 3.5 is also interesting for another reason: it shows that Theorem 2.3 no longer holds when FO^3 is replaced by basic modal logic.

By the same argument as before in the case of finite variable logics, Corollary 3.4 not only follows from Theorem 3.3 but also implies it, and hence the two are equivalent. Our results resolve a question posed in [24] concerning the relationship between the modal Lindström theorem and the ‘classic’ bisimulation preservation theorem.

3.2. A well-behaved non-elementary extension of modal logic. In this section, we prove Theorem 3.5: we introduce an extension of modal logic, ML^\bullet , that is non-elementary but at the same time is very well behaved (e.g., it has the Compactness and Löwenheim-Skolem properties). ML^\bullet extends basic modal logic with an extra operator \bullet . Thus, the formulas of ML^\bullet are given by:

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \psi \mid \diamond\phi \mid \bullet\phi$$

The semantics of the newly added operator is as follows: $M, w \models \bullet\phi$ iff w has infinitely many reflexive successors satisfying ϕ (a reflexive node is one related to itself).

The reader may verify that \bullet behaves as any normal modal operator: $\bullet(\phi \vee \psi)$ is equivalent to $\bullet\phi \vee \bullet\psi$, and $\bullet\perp$ is equivalent to \perp . The dual of \bullet is denoted by $\bar{\bullet}$.

Proposition 3.6. *ML^\bullet is a non-elementary abstract modal logic extending basic modal logic.*

Proof. This is an easy verification: e.g., the language is closed under relativisation. \square

Proposition 3.7. *ML^\bullet has the Löwenheim-Skolem property and is invariant for potential isomorphisms.*

Proof. By the containment of ML^\bullet in $L_{\omega_1\omega}$. \square

Before we proceed to prove the remaining properties of ML^\bullet , we will first provide an alternative semantics for the logic. Let K be the class of all bi-modal Kripke frames $\mathfrak{F} = (W, R_\diamond, R_\bullet)$ such that $R_\bullet \subseteq R_\diamond$ and for all $(w, v) \in R_\bullet$, $(v, v) \in R_\diamond$. Kripke models based on such frames can be thought of as quasi-models for ML^\bullet . The next lemma shows that they are adequate as such.

Lemma 3.8. *For all $\phi \in ML^\bullet$, ϕ is satisfiable according to the intended semantics iff ϕ is satisfiable on the class K .*

Proof. In one direction, suppose that $M, w \models \phi$ according to the intended semantics, where $M = (W, R, V)$ is a uni-modal Kripke model. Define an equivalence relation \sim on W by letting $v \sim v'$ iff v and v' assign the same truth value to all subformulas of ϕ . Let $R_\diamond := R$ and $R_\bullet := \{(v, v') \in R \mid (v', v') \in R \text{ and there are infinitely many } v'' \sim v' \text{ such that } (v, v'') \in R \text{ and } (v'', v'') \in R\}$. Then the underlying Kripke frame of M' belongs to K , and a straightforward inductive argument shows that, for all subformulas ψ of ϕ , and for all $v \in W$, $M, v \models \psi$ according to the intended semantics iff $M', v \models \psi$. In particular, $M', w \models \phi$. Incidentally, it is important for the above argument that \sim be an equivalence relation of finite index, so that the pigeon hole principle can be applied in the inductive clause for the new modality.

Conversely, let $M, w \models \phi$, where $M = (W, R_\diamond, R_\bullet, V)$ is a bi-modal Kripke model based on a frame in K . We will construct a uni-modal model M' with a node w' such that $M', w' \models \phi$ according to the intended semantics. Roughly, the construction of M' is based on the following ideas: (i) make sure that all $R_\diamond \setminus R_\bullet$ -successors of a node are irreflexive, by *unraveling*, and (ii) taking infinitely many copies of each R_\bullet successor of each node.

More precisely, as the domain of M' , we choose $W \times \mathbb{N}$. The accessibility relation R is the set of all pairs $((v, n), (v', m))$ satisfying one of the following conditions:

- $m \leq 1$, $m \neq n$, and $(v, v') \in R_\diamond$; or
- $m \geq 2$, and $(v, v') \in R_\bullet$.

Observe that, for all $v \in W$, $(v, 0)$ and $(v, 1)$ are R -irreflexive by construction, whereas (v, n) for $n \geq 2$ is R -reflexive or R -irreflexive, depending on v .

Finally, the valuation function V' for the atomic propositions is defined as usual, by letting $V'(p) = \{(v, n) \mid v \in V(p)\}$. Let $M' = (W \times \mathbb{N}, R, V')$.

A straightforward inductive argument now shows that, for all formulas ψ , nodes $v \in W$ and natural numbers n , $M, v \models \psi$ iff $M', (v, n) \models \psi$ according to the intended semantics. In particular, we see that $M', (w, 0) \models \phi$. \square

In the remainder of this subsection, we use this alternative semantics.

Proposition 3.9. *A complete axiomatization for ML^\bullet can be obtained by extending any standard axiomatization for basic (bi-)modal logic with the following axioms:*

$$\bullet\phi \rightarrow \diamond\phi \quad \text{and} \quad \bar{\bullet}(\phi \rightarrow \diamond\phi)$$

This axiomatization is even strongly complete: each consistent set of formulas is jointly satisfiable.

Proof. Both axioms are Sahlqvist formulas, and together, they define the frame class K . It follows by the Sahlqvist completeness theorem that the axiomatization is strongly complete for K : every consistent set of formulas is jointly satisfiable in a Kripke model based on a frame in K .

It follows by Lemma 3.8 that the axiomatization is also sound and strongly complete with respect to the intended semantics of ML^\bullet (the proof of Lemma 3.8 actually shows that every set of formulas jointly satisfiable on K is jointly satisfiable according to the intended semantics). \square

It follows that ML^\bullet also satisfies Compactness. We add a few further interesting properties.

Lemma 3.10. *The satisfiability problem for ML^\bullet is PSPACE-complete.*

Proof. The lower bound follows from the PSpace-hardness of basic modal logic. As for the upper bound, the standard PSpace-algorithm for basic modal logic due to Ladner can be extended in a straightforward manner to ML^\bullet (using the alternative semantics provided by Lemma 3.8).

Alternatively, one can also perform a direct reduction: for any ML^\bullet formula ϕ , let ϕ^* be the uni-modal formula obtained by replacing each subformula of the form $\bullet\psi$ by $\diamond(r \wedge \psi)$, for r a (fixed) fresh proposition letter. Let $Sub(\phi)$ be the set of subformulas of ϕ , and let $depth(\phi)$ be the modal nesting depth of ϕ .

It can be shown that, for any $\phi \in ML^\bullet$, ϕ is satisfiable on K iff

$$\phi^* \wedge \bigwedge_{\substack{k \leq depth(\phi) \\ \psi \in Sub(\phi)}} \Box^k (r \wedge \psi \rightarrow \diamond\psi)$$

is a satisfiable basic modal formula. Roughly, the idea is that the additional conjuncts in the above formula guarantee that we can make every r -node reflexive in the model without affecting the truth of ϕ . Together with Lemma 3.8 this gives us a polynomial time reduction from ML^\bullet -satisfiability to satisfiability of basic modal formulas. \square

Proposition 3.11. *ML^\bullet has the Craig Interpolation property: for all ML^\bullet -formulas ϕ, ψ , if $\models \phi \rightarrow \psi$ then there is an ML^\bullet -formula χ such that $\models \phi \rightarrow \chi$, $\models \chi \rightarrow \psi$, and all proposition letters occurring in χ occur both in ϕ and in ψ .*

Proof. This follows from a result in [14], according to which every multi-modal logic axiomatized by Sahlqvist formulas with first-order universal Horn correspondents has Craig interpolation in the sense described above (note that we do not require the modal operators occurring in χ to occur both in ϕ and in ψ). The two axioms from Proposition 3.9 are precisely of this form. \square

Note that ML^\bullet is *not* invariant for bisimulations, and it does *not* have the finite model property. It would be of interest to characterize the extensions of basic modal logic satisfying Compactness and the Finite Model Property. Perhaps this could still lead to some sort of modal counterpart to the Compactness-Löwenheim-Skolem version of the Lindström theorem for first-order logic.

3.3. Graded modal logic. Graded modal logic (*GML*) extends basic modal logic with counting modalities: for each formula ϕ and natural number k , $\Diamond_k\phi$ is admitted as a formula, and it says that at least k successors of the current node satisfy ϕ .

GML-formulas are in general not invariant for bisimulations. Still, an important weaker invariance property does hold: *GML* formulas are invariant for *tree unraveling*. A *tree model* is a Kripke model whose underlying frame is a tree (in the graph theoretic sense, possibly infinite, and with a unique root). We will denote tree models by T, T', \dots , and we will use $\text{root}(T)$ to denote the root of the tree model T . Every pointed Kripke model (M, w) can be unraveled into a tree model, by the following standard construction:

Definition 3.12 (Tree unraveling). Given a Kripke model $M = (W, R, V)$ and $w \in W$, the tree unraveling $\text{unr}(M, w)$ is defined as (W', R', V') , where

- ▶ W' consists of all finite paths $\langle w_1, w_2, \dots, w_n \rangle$ satisfying $w_1 = w$ and $w_i R w_{i+1}$.
- ▶ R' contains all pairs of sequences of the form $(\langle w_1, \dots, w_n \rangle, \langle w_1, \dots, w_n, w_{n+1} \rangle) \in W' \times W'$
- ▶ $\langle w_1, \dots, w_n \rangle \in V'(p)$ iff $w_n \in V(p)$.

It is easily seen that, for any pointed Kripke model (M, w) , $\text{unr}(M, w)$ is indeed a tree model, and that $\langle w \rangle$ is its root. *GML*-formulas are invariant for this operation:

Fact 3.13 (*GML* is invariant for tree unravelings). For all pointed Kripke models (M, w) and *GML*-formulas ϕ , $M, w \models \phi$ iff $\text{unr}(M, w), \langle w \rangle \models \phi$.

The proof is a straightforward formula induction.

We will prove two Lindström theorem for *GML*. The first characterizes *GML* on arbitrary structures in terms of Compactness, the Löwenheim-Skolem property and invariance for tree unravelings. It can be seen as a natural generalization of Theorem 3.3. The second theorem, which will be proved in the following section, considers *GML* as a language for describing nodes in *tree models*, and it characterizes *GML* as being maximal with respect to Compactness and the Löwenheim-Skolem property on such structures.

Theorem 3.3 characterized modal logic in terms of Compactness and bisimulation invariance. One might wonder if, likewise, *GML* can be characterized by Compactness and invariance for tree unraveling. The answer is negative: the extension of *GML* with the modal operator \Diamond_{\aleph_1} (“uncountably many successors ...”) still has these properties (even first-order logic extended with the “*uncountably many*” quantifier has (countable) Compactness [9]). Instead, we prove the following:

Theorem 3.14. *An abstract modal logic extending GML satisfies invariance for tree unravelings, Compactness, and the Löwenheim-Skolem property iff it is no more expressive than GML.*

As in the case of modal logic, we obtain the following “generalized preservation theorem” as a corollary (the proof is analogous to that of Corollary 3.4):

Corollary 3.15. *Let \mathcal{L} be any abstract modal logic extending GML and satisfying Compactness and the Löwenheim-Skolem property. Then GML is the tree unraveling invariant fragment of \mathcal{L} (up to logical equivalence). In particular, GML is the tree unraveling invariant fragment of FO.*

The rest of this section is dedicated to the proof of Theorem 3.14. Two easily established facts about *GML* will be used in the proof:




Figure 3: Model from the proof of Lemma 3.18

Fact 3.16 (*GML* has the finite tree model property). Every satisfiable *GML* formula is satisfied at the root of some finite tree model.

Fact 3.17. (*GML* can describe finite tree models up to isomorphism) Assuming a finite vocabulary, for every finite tree model T , there is a *GML*-formula ψ_T such that for every tree model T' , $(T', \text{root}(T')) \models \psi_T$ iff $T' \cong T$.

Fix an abstract modal logic \mathcal{L} extending *GML* and satisfying Compactness and Löwenheim-Skolem, as well as invariance for tree unravelings. By Lemmas 3.1 and 3.2, \mathcal{L} has the finite occurrence property and the finite depth property (note that invariance for tree unraveling implies invariance for generated submodels). The following Lemma shows a kind of “finite width property”:

Lemma 3.18. (Can only count successors up to a finite number) *Let \mathcal{L} be any abstract modal logic extending *GML* and satisfying Compactness, Löwenheim-Skolem, and invariance for tree unravelings. Then for each formula $\phi \in \mathcal{L}$ and finite tree model T , there is a natural number k such that “ ϕ can only count T -successors up to k ”: whenever a tree model contains a node v that has k successor subtrees isomorphic to T , then adding more copies of T will not affect the truth value of ϕ at the root.*

Proof. Since T is a finite tree model, it can be described up to isomorphism by a single *GML*-formula ψ_T . Let Σ be the following set of formulas, where p is a proposition letter not occurring in ϕ , and \Box^n stands for a sequence of n boxes:

$$\{p, \Box^n(\neg p \rightarrow \Box \neg p) \mid n \in \mathbb{N}\}$$

“ p defines an initial subtree”

$$\cup \{\Box^n(p \rightarrow \Box(\neg p \rightarrow \psi_T)) \mid n \in \mathbb{N}\}$$

“the root of every $\neg p$ -subtree satisfies ψ_T ”

$$\cup \{\Box^n(p \wedge \Diamond \neg p \rightarrow \Diamond_m(p \wedge \psi_T)) \mid n, m \in \mathbb{N}\}$$

“every p -node with a $\neg p$ -successor has infinitely many p -successors satisfying ψ_T ”

Whenever a *countable tree model* satisfies Σ at the root, the submodel defined by p is isomorphic to the whole model — isomorphic in the language without p , to be precise (see Figure 3). Since \mathcal{L} satisfies the Löwenheim-Skolem property, invariance for tree unraveling and invariance for isomorphisms, we can conclude that $\Sigma \models \phi \leftrightarrow \phi^p$. But then, by compactness, there is a $k \in \mathbb{N}$ such that $\Sigma_k \models \phi \leftrightarrow \phi^p$, where Σ_k is the following subset of

Σ :

$$\begin{aligned} & \{ p, \Box^n(\neg p \rightarrow \Box\neg p) \mid n \in \mathbb{N} \} \\ \cup & \{ \Box^n(p \rightarrow \Box(\neg p \rightarrow \psi_T)) \mid n \in \mathbb{N} \} \\ \cup & \{ \Box^n(p \wedge \Diamond\neg p \rightarrow \Diamond_m(p \wedge \psi_T)) \mid n, m \in \mathbb{N} \text{ with } m \leq k \} \end{aligned}$$

This shows that the Lemma holds. \square

Proof of Theorem 3.14. Let \mathcal{L} be any abstract modal logic extending *GML*, satisfying Compactness, the Löwenheim-Skolem property, and tree unraveling invariance. Observe that \mathcal{L} still satisfies Compactness and the Löwenheim-Skolem property if we restrict attention to trees (the tree unraveling of a countable model is countable). Also, by Lemmas 3.1 and 3.2, \mathcal{L} has the finite occurrence property and the finite depth property (note that invariance for tree unraveling implies invariance for generated submodels).

Consider any formula $\phi \in \mathcal{L}$. We will construct a set of equivalence relations \sim_i^ϕ for tree models (with $i \geq 0$), satisfying the following two properties:

- (1) $T \sim_i^\phi T'$ implies that the truth value of ϕ at the root of a tree model is not affected if subtrees isomorphic to T at depth i are replaced by copies of T' (or vice versa).
- (2) Each \sim_i^ϕ has only finitely many equivalence classes, and each is definable by a *GML*-formula.

This then implies that ϕ is equivalent to a *GML* formula (take the disjunction of the *GML*-formulas defining the \sim_0^ϕ -equivalence classes that satisfy ϕ).

The claim holds trivially for $i > \text{depth}(\phi)$. Next, assume that the claim holds for $i + 1$. We will show that it also holds for i . Let K_1, \dots, K_n be the (finitely many) \sim_{i+1}^ϕ -equivalence classes, and for each $\ell \leq n$, pick a finite representative $T_\ell \in K_\ell$ (using Proposition 3.16). It follows from Lemma 3.18 that there is a $k \in \mathbb{N}$ such that, for all $\ell \leq n$: “ ϕ can only count T_ℓ -successors up to k ”, and hence, by \sim_{i+1}^ϕ -equivalence, “ ϕ can only count K_ℓ -successors up to k , at depth i ”. But then, it follows that there are at most $k^n \cdot 2^{\text{PROP}(\phi)}$ many \sim_i^ϕ -equivalence classes. Moreover, all these classes are explicitly definable by *GML*-formulas (in fact, by Boolean combinations of proposition letters and formulas of the form $\Diamond_m\psi$ with $m \leq k$ and ψ a *GML*-formula defining some \sim_{i+1}^ϕ -equivalence class).

Thus, \mathcal{L} is not more expressive than *GML* on tree models. It follows that \mathcal{L} is not more expressive than *GML* on arbitrary Kripke models: consider any \mathcal{L} -formula ϕ , and let ψ be any *GML*-formula equivalent to ϕ on tree models. If $\phi \leftrightarrow \psi$ were falsifiable on an arbitrary model, then, by unraveling, it could be falsified on a tree, which, by assumption, is not the case. Thus, ϕ and ψ are equivalent on all Kripke models. \square

3.4. Graded modal logic on trees. In this section, rather than assuming invariance for tree unraveling, we consider only tree models from the start. That is, we view *GML* as a language for describing nodes of tree models. From this perspective, *GML* has three distinctive limiting features: (i) when evaluated in a node, formulas can only see the subtree starting from that node; (ii) when evaluated at a node, each formula can only look finitely deeply into the subtree starting at that node; (iii) each formula can only count successors up to a finite natural number (depending on the largest index occurring in the formula).

Fact 3.19. On trees, *GML* has the Compactness and Löwenheim-Skolem properties.

Proof. This follows from the fact that *GML* has these properties on arbitrary structures (recall that the tree unraveling of a countable model is still countable). \square

We will turn this into a Lindström characterization for *GML* on trees. However, before we proceed, two technical issues need to be discussed.

The first concerns *closure under relativisation*. Recall from Sect. 3.1 that all abstract modal logics are assumed to satisfy this property. But what does it mean for a logic to be closed under relativisation if we consider only trees? Note that a submodel of a tree is not necessarily a tree. We solve this problem as follows. Given a tree model T containing a node n , and a unary predicate p true at n , we define $Subtree(T, n, p)$ to be the largest subtree of T that contains n and contains only nodes satisfying p . Note that n is not necessarily the root of $Subtree(T, n, p)$. We say that an abstract modal logic \mathcal{L} is closed under relativisation on trees, if for every formula $\phi \in \mathcal{L}$ and proposition letter p , there is a formula $\psi \in \mathcal{L}$ such that for all pointed tree models (T, n) , $(T, n) \models \psi$ iff $(T, n) \models p$ and $(Subtree(T, n, p), n) \models \phi$. In the case of *GML*, we can simply pick ψ to be the syntactic relativisation of ϕ by p , i.e., the formula obtained from $\phi \wedge p$ by replacing all subformulas of the form $\diamond_n \psi$ by $\diamond_n(p \wedge \psi)$.

Secondly, we need to make an extra assumption, namely that the extensions \mathcal{L} we consider are *closed under substitution*. Intuitively, this means that \mathcal{L} allows us to uniformly substitute formulas for proposition letters. More precisely, \mathcal{L} is closed under substitution if for all formulas $\phi, \psi \in \mathcal{L}$ and proposition letters p , there is a formula χ such that for all pointed (tree) models (M, w) , $(M, w) \models \chi$ iff $(M^{[p \mapsto \{v \mid (M, v) \models \psi\}]}, w) \models \phi$, where $M^{[p \mapsto \{v \mid (M, v) \models \psi\}]}$ is obtained from M by changing the valuation of p to $\{v \mid (M, v) \models \psi\}$. If we would not assume closure under substitution, there would be proper extensions of *GML* with Compactness and Löwenheim-Skolem. Indeed, the extension of *GML* with formulas of the form $\diamond^- p$ (“the current node has a parent that satisfies p ”) for p an atomic proposition, and closed under the Boolean connectives, is an example. This logic is not closed under substitution, as p may not be replaced by a complex formula in $\diamond^- p$.

Theorem 3.20. *Let \mathcal{L} be an abstract modal logic closed under substitution and extending *GML* on trees. \mathcal{L} satisfies Compactness and the Löwenheim-Skolem property on trees iff it is no more expressive than *GML* on trees.*

This is remarkable, since Compactness and Löwenheim-Skolem are also the characterizing features of first-order logic in the classic Lindström theorem. Note that first-order logic lacks Compactness on trees (due to the connectedness of trees).

Now for the proof. Let \mathcal{L} be an abstract modal logic satisfying the conditions of Theorem 3.20. Lemma 3.22 below shows that \mathcal{L} -formulas can only look downwards in the tree.

Lemma 3.21. *If \mathcal{L} is not invariant for generated submodels, then there is a formula $\chi \in \mathcal{L}$ containing a unary predicate p such that the following two conditions hold:*

- (1) *for all pointed trees (T, n) , $(T, n) \models \chi$ implies that n has a parent satisfying p*
- (2) *there is a pointed tree (T, n) satisfying χ in which p is true only at the parent of n .*

Proof. Since \mathcal{L} is not invariant for generated submodels, there is a formula $\phi \in \mathcal{L}$ and a pointed tree (T, n) such that $(T, n) \models \phi$ and $(T_n, n) \not\models \phi$, or vice versa. Since \mathcal{L} is closed under negation, we may assume w.l.o.g. that the former applies. Moreover, since \mathcal{L} is closed under renamings and the set of proposition letters is infinite, we may assume that there are

infinitely many proposition letters not occurring in ϕ , in the sense that their interpretation does not influence truth of ϕ at any state. We will refer to these proposition letters as being “fresh”.

Pick two distinct fresh proposition letters p, q , let ϕ^p and ϕ^q be the relativisations of ϕ to p and to q , and let Σ be the following set of \mathcal{L} -formulas:

$$\Sigma = \{\phi^p, \neg\phi^q, \Box^k(p \wedge q) \mid k \geq 0\}$$

Observe that Σ is satisfiable: it is true at (T, n) when we make p true at all nodes, and q only at n and its descendants. We claim that truth of Σ at a node n in a tree implies that n has a parent and it satisfies either p or q . For, if not then the submodels $\text{Subtree}(T, n, p)$ and $\text{Subtree}(T, n, q)$ would coincide, and hence ϕ^p and ϕ^q would have to have the same truth value at n .

Next, we will use Compactness to obtain a finite subset of Σ that implies that the current node has a parent satisfying $p \vee q$. First, we ‘redescribe’ the situation encoded by Σ from the perspective of the parent node. Let Σ' be the following set of \mathcal{L} -formulas, where r is another fresh proposition letter:

$$\Sigma' = \{\Diamond(\phi^p \wedge \neg\phi^q \wedge r), \Box(r \rightarrow \Box^k(p \wedge q)) \mid k \geq 0\}$$

By the previous observations, $\Sigma' \models p \vee q$. Hence, by Compactness, there is a $\ell \in \mathbb{N}$ such that $\Diamond(\phi^p \wedge \neg\phi^q \wedge r) \wedge \bigwedge_{k \leq \ell} \Box(r \rightarrow \Box^k(p \wedge q)) \models p \vee q$. Going back to the perspective of the node n , if we define ψ to be the formula $\phi^p \wedge \neg\phi^q \wedge \bigwedge_{k \leq \ell} \Box^k(p \wedge q)$, then ψ is satisfiable, and it implies the existence of a parent node satisfying $p \vee q$. To see this, suppose for the sake of contradiction that ψ is satisfiable in a pointed tree (T', n') and n' does *not* have a parent satisfying $p \vee q$. There are two cases. If n' has a parent (satisfying $\neg p \wedge \neg q$), we immediately derive a contradiction, by our earlier arguments. If n' does not have a parent, i.e., is the root of T' , we may change the tree by adding a parent above n' satisfying $\neg p \wedge \neg q$, without affecting the truth of ψ at n' , as follows from the definition of ψ and the way we defined closure under relativisation. Therefore, we obtain again a contradiction.

Finally, we take two more fresh proposition letters, s and t , and we use the fact that \mathcal{L} is closed under substitution: we define χ to be $s \wedge \psi[p/(p \wedge (\Diamond s \rightarrow t)), q/(q \wedge (\Diamond s \rightarrow t))]$.

On the one hand, truth of χ at a node implies it has a parent satisfying either $(p \wedge (\Diamond s \rightarrow t))$ or $(q \wedge (\Diamond s \rightarrow t))$, and hence t . On the other hand, there is a pointed tree satisfying χ in which t is only true at the parent node: take (T, n) and extend the valuation by making s true only at n , and t at its parent. \square

Lemma 3.22. *\mathcal{L} is invariant for generated submodels.*

Proof. Suppose not. Let $\chi(p) \in \mathcal{L}$ be as described by Lemma 3.21. By a “fresh renaming” of χ we will mean a copy in which all proposition letters have been renamed to fresh ones, and which has been relativised by an additional fresh proposition letter. For the reasons explained in the proof of Lemma 3.21, we may assume that $\chi(p)$ has infinitely many fresh renamings $(\chi_i(p_i))_{i \in \mathbb{N}}$.

Finally, we define Σ to be the set of \mathcal{L} -formulas $\{\chi_1(p_1), \chi_1(\chi_2(p_2)), \chi_1(\chi_2(\chi_3(p_3))), \dots\}$. Every finite subset of Σ is satisfiable. Indeed, a satisfying model may be constructed by “overlying” different copies of the model (T, n) from Lemma 3.21, clause 2. On the other hand, if a node would satisfy all formulas in Σ at once, its ancestors would form an infinite ascending chain, which contradicts the well-foundedness property of trees. \square

The remainder of the proof of Theorem 3.20 is along the same lines as in Sect. 3.3: first we prove that \mathcal{L} has the finite occurrence, finite depth and finite width properties on trees (using the fact that it is invariant for generated submodels), and then we derive the Lindström theorem by the same argument used in the proof of Theorem 3.14.

3.5. The guarded fragment. The *guarded fragment* GF forms a second extension of modal logic, incomparable to graded modal logic. It allows for arbitrary quantifications of the form $\exists \vec{y}(G(\vec{x}, \vec{y}) \wedge \phi(\vec{x}, \vec{y}))$, where \vec{x} and \vec{y} are tuples of variables, and the *guard* G is an atomic formula containing the variables in \vec{x} and \vec{y} . The guarded fragment is decidable and has many ‘modal’ meta-properties, due to its invariance for *guarded bisimulations* [1, 23], see below for the definition.

Because of its modal character, GF seems an obvious case for a Lindström-style analysis like the one we have given for modal logic and graded modal logic. However, there are some technical difficulties, and we have not been able to prove an analogue of Theorem 3.3 for the guarded fragment yet. In this section, we focus on a special case, in fact the same special case as in the first half of the paper, namely for vocabularies with only unary and binary predicates. The *Binary Guarded Fragment* (GF_{bin}) has the following syntax:

$$\phi ::= R\vec{x} \mid x = y \mid \neg\phi \mid \phi \vee \psi \mid \exists y(G(x, y) \wedge \psi(x, y))$$

where R is a unary or binary atomic predicate, x and y are distinct variables, the *guard* G is an atomic formula containing both x and y (in any order) and ψ contains no free variables besides (possibly) x and y . Note that, by this definition, unary guards such as Py , Ryy and $y = y$ are not allowed, and also unguarded quantification over a single remaining free variable, as in $\exists y.\phi(y)$, can no longer be expressed by $\exists y.(y = y \wedge \phi(y))$. Thus, every GF_{bin} -formula contains at least one free variable. This prohibition on unary guards will be important, as it implies that the truth of a formula only depends on a local neighborhood of the elements assigned to the free variables.

GF_{bin} can be seen as an abstract logic contained in FO^2 , provided that the definition of abstract logics is adapted in order to take into account formulas with free variables, in the natural way. Below, we will assume such a generalized notion of abstract logic. When interpreted over Kripke models (and considering only formulas with one free variable), the language GF_{bin} also constitutes an abstract modal logic extending basic modal logic.

Guarded bisimulations admit a natural adaptation to this restricted version of the guarded fragment, which we call GF_{bin} -bisimulations, defined below.

Theorem 3.23. *An abstract logic extending GF_{bin} satisfies Compactness and Invariance for GF_{bin} -bisimulation iff it is no more expressive than GF_{bin} .*

The proof is along the same lines as for modal logic and graded modal logic: using Compactness, we prove a finite occurrence property and a finite depth property (where depth is now measured as distance in the ‘Gaifman graph’). We then use a tree-unraveling argument to show that GF_{bin} can express all properties invariant for guarded bisimulations and having the finite depth property. It is exactly in this last step that the restriction to unary and binary relation symbols, as well as the restrictions on the allowed guards, turn out to be crucial. Roughly, these restrictions allows us to relate distance in the unraveled tree to Gaifman distance in the original structure.

In the remainder of this section, we present the proof in more detail. We start with some modal features that hold for GF in general. First, there is a natural syntactic notion

of formula depth, whose inductive definition counts the above polyadic quantifiers as single steps:

$$\begin{aligned}
\text{depth}(Px) &= 0 \\
\text{depth}(\neg\phi) &= \text{depth}(\phi) \\
\text{depth}(\phi \vee \psi) &= \max(\text{depth}(\phi), \text{depth}(\psi)) \\
\text{depth}(\exists \vec{y}(G(\vec{x}, \vec{y}) \wedge \phi(\vec{x}, \vec{y}))) &= \text{depth}(\phi) + 1
\end{aligned}$$

Call a set of elements of M “*guarded*” if some tuple belonging to some atomic relation contains all these elements. We define distance for points in models (M, \vec{s}) , where \vec{s} is a tuple of nodes, as follows: $\text{dist}(\vec{s}, s_i, 0)$ holds for $s_i \in \vec{s}$, and $\text{dist}(\vec{s}, t, n+1)$ holds if there is a u with $\text{dist}(\vec{s}, u, n)$ and $\{t, u\}$ is guarded. Note that, by this definition, dist describes “distance at most” rather than “exact distance”.

We write $\text{Cut}((M, \vec{s}), n)$ for the submodel $\{t \in (M, \vec{s}) \mid \text{dist}(\vec{s}, t, n)\}$ consisting of all points t in M lying at distance at most n from s . The following result shows that GF , like basic modal logic, satisfies a *finite depth property*, suitably defined:

Lemma 3.24 (Distance-Depth Lemma). *Let ϕ be any guarded formula of depth n , and let (N, \vec{s}) be any submodel of (M, \vec{s}) containing all of $\text{Cut}((M, \vec{s}), n)$. Then $(M, \vec{s}) \models \phi$ iff $(N, \vec{s}) \models \phi$.*

Next comes a generalization of modal bisimulation. A guarded bisimulation is a non-empty set F of finite partial isomorphisms between two models M and N which has the following back-and-forth conditions: given any function $f : X \rightarrow Y$ in F , (i) for all guarded $Z \subseteq M$, there is a $g \in F$ with domain Z such that g and f agree on the intersection $X \cap Z$, (ii) for all guarded $W \subseteq N$, there is a $g \in F$ with range W such that the inverses g^{-1} and f^{-1} agree on $Y \cap W$.

Also, we have “*rooted*” guarded bisimulations F running between models (M, s_1, \dots, s_n) and (N, t_1, \dots, t_n) with given initial objects, where one requires that $\{(s_i, t_i) \mid 1 \leq i \leq n\}$ is a partial isomorphism in F . By a simple inductive argument, one then shows:

Fact 3.25. *GF -formulas are invariant for rooted guarded bisimulations: whenever there is a rooted guarded bisimulation between models (M, s_1, \dots, s_n) and (N, t_1, \dots, t_n) , then for all GF -formulas $\phi(x_1, \dots, x_n)$, $M \models \phi[s_1, \dots, s_n]$ iff $N \models \phi[t_1, \dots, t_n]$.*

Andréka, van Benthem and Németi [1] show that, conversely, GF consists, up to logical equivalence, of just those first-order formulas which are invariant for guarded bisimulations.

Another ‘modal’ use of guarded bisimulation in the same paper is *model unraveling*. This is like standard modal unraveling, but the construction is a bit more delicate:

Definition 3.26 (Tree unravelings for GF). Let M be any model. A *guarded path* will be a non-empty finite sequence of guarded sets in M . We say that an element of M is *active* in a guarded path $\pi = \langle S_1, \dots, S_n \rangle$ if it belongs to S_n and not to S_{n-1} (or $n = 1$), and *passive* if it belongs to both S_n and S_{n-1} . The intuition is when an element occurs actively, a new copy is created in the unraveling, whereas when it occurs passively, the previous copy is used. Correspondingly, let \equiv be the equivalence relation generated by $(\pi_1, d) \sim (\pi_2, d)$ for π_2 a guarded path extending π_1 with a single guarded set, and d passive in π_2 . Now, the domain of the tree unraveling $\text{unr}_{GF}(M)$ consists of all pairs (π, d) where π is a guarded path and d is active in π , and the interpretation of predicate symbols Q is as follows. $I(Q)$ contains the tuple $\langle (\pi_1, d_1), \dots, (\pi_k, d_k) \rangle$ if and only if (d_1, \dots, d_k) belongs to the interpretation of Q in M , and there is a guarded path π (which will in fact be a

common extension of π_1, \dots, π_n , such that d_1, \dots, d_k all belong to the final set of π , and $(\pi, d_i) \sim (\pi_i, d_i)$ for all $1 \leq i \leq n$.

For a model with parameters (M, \vec{s}) this generalizes as follows. $\text{unr}_{GF}(M, \vec{s})$ has paths π all starting from the initial set \vec{s} , but then continuing with guarded sets only. The objects (π, d) are defined as before.

The point here is that the set F of all restrictions of the finite maps sending (π_i, d_i) to d_i for all guarded finite domains in $\text{unr}_{GF}(M, \vec{s})$ is a rooted guarded bisimulation between (M, \vec{s}) and $\text{unr}_{GF}(M, \vec{s})$. Checking the zigzag conditions for the bisimulation will reveal the reason for the above definition of the predicate interpretations $I(Q)$. For other, essentially equivalent formulations of this notion, cf. [1, 11].

Now we have the generalities in place for our Lindström Theorem, but it remains to make some adjustments. Firstly, the notion of guarded bisimulation needs to be slightly adjusted. GF_{bin} -bisimulations are defined like guarded bisimulations, except that the back-and-forth conditions are only required to hold when $X \cap Z \neq \emptyset$ and $Y \cap W \neq \emptyset$, respectively. It can be shown that GF_{bin} is precisely the GF_{bin} -bisimulation invariant fragment of FO (over vocabularies consisting of unary and binary relation symbols only).

Secondly, in the definition of tree unravelings, we make one simple change for the binary case:

The finite paths of guarded sets always introduce one new object at each stage. At each step, one takes a new object related to that new object.

This allows paths starting with object a and then continuing with Rab, Qcb, \dots , while ruling out paths like Rab, Qac . But the final atom is not omitted from the unraveled model, since one can have paths starting with a and then placing Qac immediately. Thus, even with these restricted paths, we still have a GF_{bin} -bisimulation between tree unravelings and their original models. The real point of this adjustment is the following observation:

The definition of predicates for path objects makes binary relations hold only between objects $(\pi_1, d_1), (\pi_2, d_2)$ where π_2 is a one-step continuation of the path π_1 , or vice versa. But then, counting distance as before,

The new object at the end of a path of length k lies at distance k from the initial object of the path.

Put in more vivid terms, ‘*tree distance is true distance*’ in the original model. This is a non-obvious fact. E.g., with ternary guards $Rayz$, objects at the end of a path may keep links to the initial object a which might recur at any finite depth in the guarded sets along the path.³ Having a direct correspondence between the length of the path and the distance in the tree unraveling is essential, as it will allow us to establish a finite depth property analogous to Lemma 3.2 for the basic modal language.

We will write $\text{unr}_{GF_{\text{bin}}}(M)$ and $\text{unr}_{GF_{\text{bin}}}(M, \vec{s})$ for this new type of unraveling. Now everything is in place for our final argument:

Proof of Theorem 3.23. Let \mathcal{L} be any abstract logic extending GF_{bin} and satisfying Compactness and invariance for GF_{bin} -bisimulations. As before, it suffices to show that every formula $\phi \in \mathcal{L}$ is invariant for models that are equivalent for all GF_{bin} -formulas up to some finite depth n .

First, largely as in the earlier modal proof of Sect. 3.1, we use the Compactness of \mathcal{L} , together with its *Relativization Closure*, to show that ϕ must have the *Finite Occurrence*

³This observation is due to Martin Otto, p.c. to the authors.

Property and a *Finite Distance Property* for some level n . Before, universal prefix formulas $\Box^k p$ (for all finite k) made sure that p holds in the generated submodel at the current node. This time, one uses all nested sequences of universal guarded quantifiers up to depth k , requiring that some new predicate P holds for all objects reached at the end. The n thus found for the local depth of the formula ϕ is the same n as needed for the following semantic invariance:

Given the above unraveling construction and invariance for GF_{bin} -bisimulation for \mathcal{L} , we may then assume, without loss of generality, that we have the following situation:

(a) $\text{unr}_{GF_{\text{bin}}}(M, \vec{s}) \models \phi$

(b) $\text{unr}_{GF_{\text{bin}}}(M, \vec{s})$ and $\text{unr}_{GF_{\text{bin}}}(N, \vec{t})$ satisfy the same GF_{bin} -formulas up to depth n

Our aim is to show that $(\text{unr}_{GF_{\text{bin}}}(N, \vec{t})) \models \phi$.

We start by cutting the tree models to depth n , as before in our modal argument, obtaining $\text{Cut}(\text{unr}_{GF_{\text{bin}}}(M, \vec{s}), n)$ and $\text{Cut}(\text{unr}_{GF_{\text{bin}}}(N, \vec{t}), n)$. Since tree depth is true depth, this does not change truth values of ϕ in either model.

Then we define a GF_{bin} -bisimulation between these models. First of all, note that all elements in the first one are of the form $(\langle \{\vec{s}\}, \{d_1, d_2\}, \dots, \{d_{k-1}, d_k\}\rangle, d_k)$ with $d \in \vec{s}$ and $0 \leq k \leq n$ where each set in the sequence is guarded. Likewise for the second model. Secondly, note that the only guarded subsets in these models are singleton sets and sets of the form $\{(\pi, d), (\langle \pi, \{d, e\}\rangle, e)\}$. Now, let the relation Z be the set consisting of $(\{\vec{s}\}, \{\vec{t}\})$, as well as all sets of pairs

$$\{(\langle \pi, d \rangle, \langle \pi', d' \rangle), (\langle \langle \pi, \{d, e\} \rangle, e \rangle, \langle \langle \pi', \{d', e'\} \rangle, e' \rangle)\}$$

with π, π' sequences of some length $k \leq n$ for which it holds that, for all $i \leq k$, the guarded sets π_i and π'_i satisfy the same GF_{bin} -formulas of quantifier depth $n - i$.

It can be shown (the argument involves a case distinction) that Z is a GF_{bin} -bisimulation between $\text{Cut}(\text{unr}_{GF_{\text{bin}}}(M, \vec{s}), n)$ and $\text{Cut}(\text{unr}_{GF_{\text{bin}}}(N, \vec{t}), n)$. Hence, the truth of ϕ transfers from $\text{Cut}(\text{unr}_{GF_{\text{bin}}}(M, \vec{s}), n)$ to $\text{Cut}(\text{unr}_{GF_{\text{bin}}}(N, \vec{t}), n)$, and hence to (N, \vec{t}) . \square

As indicated, this argument hits a barrier when ternary predicate are allowed, and hence we leave a similar Lindström characterization of the full guarded fragment GF as an open problem. Recently [16], a Lindström-style characterization was obtained for the full GF along slightly different lines, by means of adding a further model-theoretic property, cf. the discussion below.

4. DISCUSSION

To conclude, we identify a few lines of research suggested by our results.

4.1. Charting the landscape of first-order fragments. We have seen how many fragments of first-order logic have Lindström characterizations, but with very different proof techniques. ‘Top-down’, we showed how careful coding can get the original proof down to cover the 3-variable fragment, while ‘bottom up’, we gave a modal technique which lifts to various richer modal-like extensions.

There are more positive results than we have presented here: e.g., our coding technique can easily be adapted to the *bounded fragment* of first-order logic (cf. [19]), showing that *every abstract logic extending the bounded fragment and having Compactness and the Löwenheim-Skolem property is contained in FO*, and hence also *the bounded fragment is*

maximal with respect to Compactness, the Löwenheim-Skolem property, and invariance for generated submodels (a restriction to binary vocabularies is not necessary here).

However, there is also a ‘gap’ between the two proof techniques, and logics like the 2-variable fragment, or modal logic with an added ‘universal modality’ (“true in all nodes, accessible or not”) pose a challenge. We may need new proof techniques here, or indeed, a rethinking of what a Lindström theorem should be in such cases. Recently, Otto and Piro [16] have characterized modal logic extended with the universal modality as being a maximal system \mathcal{L} with respect to Compactness, *invariance for global bisimulations* and also, the *Tarski union property* (i.e., preservation of \mathcal{L} -sentences under limits of chains of \mathcal{L} -elementary embeddings), and similarly for the guarded fragment. Thus, additional preservation properties may be needed.

4.2. Re-positioning first-order logic in this broader setting. Lindström’s theorem in its classical formulation characterizes first-order logic only within the class of its extensions. A natural question to ask in our wider setting is within which broader classes of languages we can characterize first-order logic. Our Theorem 2.3 gave a partial answer: it characterized first-order logic within the class of extensions of FO^3 . This at once suggests many new questions. For instance, are there extensions of FO^2 not contained in FO that satisfy Compactness and the Löwenheim-Skolem property? ⁴

4.3. Down and then up again: adding fixed-point operators. It is well-known that fragments of first-order logic can behave very differently from the full system when it comes to adding transitive closure, or even complete fixed-point extensions. In particular, the *modal μ -calculus* remains decidable, just like the basic modal logic, whereas the logic $LFP(FO)$ extending first-order logic with arbitrary fixed-point operators becomes not recursively enumerable, and indeed Π_1^1 complete. While fixed-point logics of both sorts are natural from a computational point of view, they have resisted Lindström-style analysis so far. Here is one question, out of many which suggest themselves following our earlier results:

Question 4.1. Can the modal μ -calculus be characterized in terms of bisimulation invariance and the finite model property? ⁵

4.4. Characterizing logics on specific classes of structures. No Lindström characterizations are known for first-order logic on finite structures, or on trees. Compactness fails for first-order logic on such structures, and, on finite structures, the Löwenheim-Skolem property becomes meaningless. In this paper, we proved one positive result: we showed that *GML* behaves on trees as first-order logic does on arbitrary structures: it is maximal with respect to Compactness and the Löwenheim-Skolem property. In general, however, this area remains underexplored: [12] raise the issue of the ‘missing Lindström theorem’ for finite model theory. Here we state just one related technical question. Given a logic \mathcal{L} interpreted on finite structures and a logic \mathcal{L}' interpreted on arbitrary structures (both

⁴Theorem 3.5 gave an extension of *modal logic* not contained in first-order logic which has these properties.

⁵This formulation naturally arises in our setting since so many of our modal arguments involved finite reachability of nodes from the origin, a typical non-first-order fixed-point notion.

satisfying the usual conditions such as closure under the Boolean operations), let us say that \mathcal{L}' is a conservative extension of \mathcal{L} if \mathcal{L} is simply the restriction of \mathcal{L}' to finite structures.

Question 4.2. Are there extensions of first-order logic on finite structures that have a conservative extension to infinite structures satisfying Compactness and Löwenheim-Skolem?

4.5. Relating different characterizations: the case of interpolation. Finally, taking stock of our results, there is also the more elusive, but still interesting question of “what is a Lindström theorem?”. As we already indicated in the introduction, there are many natural ways to characterize the expressive power of a logic, including Lindström theorems, preservation theorems, and characterizations involving Craig interpolation.

And even the Lindström theorems in this paper fall into different classes. Some use Compactness and Löwenheim-Skolem, or some other size-restricting principle of model existence, while others use Compactness and Invariance for Potential Isomorphisms, or some other semantic transfer property of the language. While these formulations both capture first-order logic, we have seen that they diverge for basic modal logic, where we have a characterization of the latter type but not of the former. The precise extent of this phenomenon remains to be understood.

While the picture of ‘natural properties’ is not quite clear yet, it is noticeable that Lindström theorems seem to go hand in hand with *interpolation results* in abstract model theory and related areas. Among the proper extensions of FO the main examples of interpolation are $L_{\omega_1\omega}$, its countable admissible fragments L_A , and second order logic SO. Barwise [3] gives a maximality characterization of $L_{\omega_1\omega}$ and its countable admissible fragments L_A in terms of a (generalized) recursion theoretic criterion called *strict absoluteness*. The proof uses interpolation.⁶ Also second-order logic has a maximality characterization: SO is the maximal extension of FO in which every definable model class has a “flat” definition in set theory [22]. Again this result is intimately connected with the (trivial) way in which SO satisfies (single-sorted) interpolation. Finally, the interpolation theorem of $L_{\infty\omega}$ for entailment along potential isomorphism in [2] is related to a maximality characterization of this logic: $L_{\infty\omega}$ is the maximal extension of FO which is bounded and has the Karp Property [4]. We feel that this link between Lindström theorems and interpolation theorems requires further analysis.

Moreover, this connection makes sense for our study of fragments as well. Some of our results used interpolation properties already, and in particular, our paradigmatic modal logics have them. Indeed, among fragments, even further types of interpolation property emerge, which fail for first-order logic. Thus, the two major modal logics for which we managed to obtain Lindström theorems, namely basic modal logic and graded modal logic, both possess *uniform interpolation*, a strong form of Craig interpolation where the interpolant can be chosen uniformly for all consequents sharing some specified vocabulary with the antecedent. We would like to understand the impact of this condition on abstract logics in conjunction with our Lindström-style analysis in this paper.

⁶It is arguable whether this characterization of $L_{\omega_1\omega}$ (or L_A) should be considered a Lindström type results as strict absoluteness is not a model theoretic condition. However the boundary between generalized recursion theory and infinitary model theory is not very strict.

REFERENCES

- [1] Hajnal Andréka, Johan van Benthem, and István Németi. Modal logics and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27(3):217–274, 1998.
- [2] Jon Barwise and Johan van Benthem. Interpolation, preservation, and pebble games. *J. Symbolic Logic*, 64(2):881–903, 1999.
- [3] K. Jon Barwise. Absolute logics and $L_{\infty\omega}$. *Ann. Math. Logic*, 4:309–340, 1972.
- [4] K. Jon Barwise. Axioms for abstract model theory. *Ann. Math. Logic*, 7:221–265, 1974.
- [5] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal logic*. Cambridge University Press, Cambridge, UK, 2001.
- [6] Egon Börger, Erich Grädel, and Yuri Gurevich. *The Classical Decision Problem*. Springer-Verlag, 1997.
- [7] Maarten de Rijke. A Lindström theorem for modal logic. In A. Ponse, M. de Rijke, and Y. Venema, editors, *Modal Logic and Process Algebra: A Bisimulation Perspective*, pages 217–230. CSLI Publications, 1995.
- [8] Jörg Flum. Characterizing logics. In Jon Barwise and Solomon Feferman, editors, *Model-Theoretic Logics*. Springer, 1985.
- [9] Gebhard Fuhrken. Skolem-type normal forms for first-order languages with a generalized quantifier. *Fundamenta mathematicae*, 54:291–302, 1964.
- [10] Marta García-Matos and Jouko Väänänen. Abstract model theory as a framework for universal logic. In J.-Y. Beziau, editor, *Logica Universalis*. Birkhäuser, 2005.
- [11] Erich Grädel. On the restraining power of guards. *Journal of Symbolic Logic*, 64(4):1719–1742, 1999.
- [12] Phokion G. Kolaitis and Jouko A. Väänänen. Generalized quantifiers and pebble games on finite structures. *Annals of Pure and Applied Logic*, 74(1):23–75, 1995.
- [13] Per Lindström. On extensions of elementary logic. *Theoria*, 35:1–11, 1969.
- [14] Maarten Marx and Yde Venema. *Multi-dimensional Modal Logic*. Applied Logic Series. Kluwer Academic Publishers, 1997.
- [15] Kan Ching Ng and Alfred Tarski. Relation algebras with transitive closure. *Notices of the AMS*, 24:A29–A30, 1977.
- [16] Martin Otto and Robert Piro. A Lindström characterisation of the guarded fragment and of modal logic with a global modality. In Carlos Areces and Robert Goldblatt, editors, *Advances in Modal Logic* 7, pages 273–187. College Publications, 2008.
- [17] Alfred Tarski and Steven Givant. *A Formalization of Set Theory without Variables*, volume 41 of *Colloquium Publications*. American Mathematical Society, 1987.
- [18] Balder ten Cate. Interpolation for extended modal languages. *Journal of Symbolic Logic*, 70(1):223–234, 2005.
- [19] Balder ten Cate. *Model theory for extended modal languages*. PhD thesis, University of Amsterdam, 2005. ILLC Dissertation Series DS-2005-01.
- [20] Balder ten Cate, Johan van Benthem, and Jouko Väänänen. Lindström theorems for fragments of first-order logic. In *Proceedings of LICS 2007*, pages 280–292. IEEE Computer Society, 2007.
- [21] Leslie Tharp. Which logic is the right logic. *Synthese*, 31:1–21, 1975.
- [22] Jouko Väänänen. Abstract logic and set theory. I. Definability. In *Logic Colloquium '78 (Mons, 1978)*, volume 97 of *Stud. Logic Foundations Math.*, pages 391–421. North-Holland, Amsterdam, 1979.
- [23] Johan van Benthem. Guards, bounds, and generalized semantics. *Journal of Logic, Language, and Information*, 14:263–279, 2005.
- [24] Johan van Benthem. A new modal Lindström theorem. *Logica Universalis*, 1:125–148, 2007.
- [25] Yde Venema. Relational games. In Hajnal Andréka, Donald Monk, and István Németi, editors, *Algebraic Logic*, 1991.