Steven Lindell*, Henry Towsner, and Scott Weinstein

Infinitary methods in finite model theory

Abstract: We exhibit the usefulness of infinitary methods in finite model theory.

***Corresponding Author: Steven Lindell:** Haverford College
Henry Towsner, Scott Weinstein: University of Pennsylvania

1 Introduction

This paper explores applications of the infinitary techniques of classical model theory in the context of finite model theory, and thus elaborates on themes explored by [13]. We are grateful for the opportunity to celebrate Jouko's inclusive vision of logic in this way.

2 Preliminaries

We suppose the reader is familiar with the notion of a relational structure of similarity type τ. By a class of structures, we understand a collection of relational structures of common similarity type which is closed under isomorphism. We use FO to denote the collection of sentences of first order logic.

Definition 2.1. Let K be a class of structures of similarity type τ and let Γ be a set of first-order sentences.

1. $\mathsf{Th}(K) = \{\varphi \in \mathsf{FO} \mid (\forall A \in K)A \models \varphi\}$.
2. $\mathsf{Mod}(\Gamma) = \{A \mid A \models \Gamma\}$.
3. Let A be a structure of similarity type τ. A is a *pseudo-K structure* if and only if $A \in (\mathsf{Mod}(\mathsf{Th}(K)) - K)$.

An alternative formulation of Definition 2.13 is: A is a pseudo-K structure if and only if $A \not\in K$ but for every first order sentence φ, if $A \models \varphi$, then for some $B \in K$, $B \models \varphi$. Intuitively, a structure A is pseudo-K just in case it is indistinguishable from structures in K by first order sentences.

Let F_τ be the class of finite structures of similarity type τ. A structure A of similarity type τ is *pseudo-finite* if and only if A is pseudo-F_τ. In general, if there is no finite upper bound on the size of structures in $K \subseteq F_\tau$, then a structure A is pseudo-K just in case A satisfies the *limit theory of* K ($\mathrm{Th}_\infty(K)$) defined as follows.

$$\mathrm{Th}_\infty(K) = \{\varphi \in \mathrm{FO} \mid (\exists n \in \omega)(\forall B \in K)(|B| \geq n \to B \models \varphi)\}.$$

Suppose F is a class of structures, not necessarily finite, and $K \subseteq F$. We say that K is first order definable over F if and only if there is a first order sentence φ such that for all $A \in F$, $A \models \varphi$ if and only if $A \in K$.

3 An easy inexpressibility result – connectivity

In this section, we give some examples which apply the notion of pseudo-K structure to inexpressibility results in finite model theory. The source of these results is the following simple observation.

Proposition 3.1. *Let $K \subseteq F_\tau$. Suppose there is a structure A of similarity type τ which is both pseudo-K and pseudo-($F_\tau - K$). Then K is not first order definable over F_τ.*

The first example is due to [9].

Proposition 3.2. *Connectivity is not definable in first order logic over finite simple graphs.*

Proof: Let K be the class of connected finite simple graphs and let G be the class of all finite simple graphs. By Proposition 3.1, it suffices to exhibit a structure A which is both pseudo-K and pseudo-($G - K$). And for this, it suffices to show that there is a structure A which is both pseudo-C_1 and pseudo-C_2, where C_1 is the class of simple cycles and C_2 is the class of pairs of disjoint simple cycles. We claim that

3.3. *every 2-regular, acyclic simple graph is both pseudo-C_1 and pseudo-C_2.*

Let T be the first order theory which specifies that a graph is 2-regular, acyclic and simple. T can be axiomatized by the following infinite set of first order sentences.

1. $(\forall x)\neg Exx \wedge (\forall x)(\forall y)(Exy \to Eyx)$
2. $(\forall x)(\exists y)(\exists z)(y \neq z \wedge (\forall w)(Exw \leftrightarrow (w = y \vee w = z))$
3. for each $n \geq 3$ the sentence

$$\neg(\exists x_1)\ldots(\exists x_n)(\bigwedge_{1\leq i<j\leq n} x_i \neq x_j \wedge \bigwedge_{1\leq i<n} Ex_ix_{i+1} \wedge Ex_nx_1)$$

It is easy to see that every model of T consists of a disjoint union of bi-infinite simple chains. It follows at once that for every $\kappa > \aleph_0$, T is κ-categorical. Therefore, by the Łos-Vaught Test, T is complete. Note that every finite sub-theory of T has a model C which consists of a single simple cycle and a model D which consists of a disjoint pair of simple cycles. It follows at once, by the Compactness Theorem, that the same holds for any sentence φ which is a consequence of T; this establishes 3.3, thereby concluding the proof. □

In the next section, we give a "classical" proof that connectivity is not monadic-Σ_1^1, as a corollary to the Hanf threshold-locality theorem.

Other investigators have noted the effectiveness of classical model-theoretic techniques in establishing limits on the expressive power of first order logic over finite structures. Antedating the classical proof of [9], [12] showed, using ultraproducts of 2-regular finite graphs, that connectivity is not first-order definable over finite simple graphs; extending this, [13] uses ultraproducts to show that connectivity is not monadic-Σ_1^1 over finite graphs. [13] also uses ultraproducts to show that even cardinality of finite structure is not first order definable; [11] gives a beautiful proof of this result using initial segments of "non-standard models" of the first order theory of $\langle \omega, < \rangle$. Moreover, [6] used classical model theoretic techniques to derive the 0-1 law for first-order logic with respect to finite graphs with constant edge probability, from which a raft inexpressibility results follows. Very recently, Philip Dittmann has applied ultraproducts to obtain a variety of results in finite model theory, as reported in [2].

4 Locality Theorems

Locality theorems characterize various senses in which properties of a relational structure which are definable in a particular way depend on "local information" about the structure, typically, isomorphism types of neighborhoods in the Gaifman-graph of the structure. The next definition formulates some topological notions requisite for describing locality precisely.

Definition 4.1. Let G be a graph, $a, b \in G$ and $r \geq 0$.
1. We write $\delta_G(a, b)$ for the distance between a and b in G.
2. $\mathsf{B}_r(G, a) = \{c \in G \mid \delta_G(a, c) \leq r\}$, for $r \in \omega$. $\mathsf{B}_\infty(G, a) = \bigcup_{r \in \omega} \mathsf{B}_r(G, a)$.
3. Let A be a relational structure of similarity type τ. The *Gaifman-graph* of A (denoted G_A) is the simple graph whose node set is A and whose edge relation holds between a pair of nodes a, b with $a /= b$ if and only if for some $R \in \tau$ and tuple $\overline{c} \in R^A$ there are i, j with $c_i = a$ and $c_j = b$.
4. We say a structure A has *bounded degree* if and only if there is a $d \in \omega$ such that for every $a \in A$ the degree of a in G_A is at most d; we say A has *finite degree* if and only if for every $a \in A$, the degree of a in G_A is finite.

5. Let $r \in \omega$ or $r = \infty$, and let $\overline{a}$ be an n-tuple of elements of A.

$$A^r(\overline{a}) = \langle (A \restriction \bigcup_{1\le i\le n} \mathsf{B}_r(G_A, a_i)), \overline{a}\rangle.$$

The structure $A^r(\overline{a})$ is often called the r-neighborhood of $\overline{a}$ in A.

Note that for any finite relational signature τ, there is a first-order formula $\delta^r_\tau(x, y)$ such that for every τ-structure A and $a, b \in A$, $A \models \delta^r_\tau[a, b]$ if and only if the distance between a and b in G_A is at most r. If $\varphi(x_1, \ldots, x_k)$ is a formula, we write $\varphi^r(x_1, \ldots, x_k)$ for the formula which results from $\varphi(x_1, \ldots, x_k)$ by relativizing all its quantifiers (Qy) to the formula $\delta^r_\tau(x_1, y) \vee \ldots \vee \delta^r_\tau(x_k, y)$. We call such a formula an *r-local formula*. If $\overline{a}$ is a k-tuple of elements of A, the *local type* of $\overline{a}$ in A is the set of formulas $\varphi^r(\overline{x})$, for all $r \in \omega$, such that $A(\overline{a}) \models \varphi^r(\overline{a})$. An n-ary τ-query is an isomorphism invariant map φ from τ-structures A to n-ary relations on A; we write $\varphi[A]$ for the φ image of A.

4.1 Gaifman locality

The next definition formulates a first notion of locality.

Definition 4.2. Suppose φ is an n-ary τ-query, $n > 0$. We say that φ is *Gaifman r-local* if and only if for every τ-structure A and n-tuples $\overline{a}, \overline{b}$ of elements of A, if $A^r(\overline{a}) \cong A^r(\overline{b})$, then $\overline{a} \in \varphi[A] \iff \overline{b} \in \varphi[A]$. A query is *Gaifman local* if and only if it is Gaifman r-local for some $r \in \omega$.

Theorem 4.3. *Every first order definable query is Gaifman local.*

Proof: Suppose that $\varphi(\overline{x})$ is a first order formula defining an n-ary query on τ-structures which is not Gaifman local.

4.4. *Then for every $r \in \omega$ there is a τ-structure A and $\overline{a}, \overline{b} \in A$ such that*
1. $\overline{a} \in \varphi[A]$ *and* $\overline{b} \notin \varphi[A]$, *and*
2. $A^r(\overline{a}) \cong A^r(\overline{b})$.

Expand the signature τ with $2n$ new constants $\overline{a}$ and $\overline{b}$ and a unary function symbol f. Choose, for each $r \in \omega$, a first order formula ψ^r in the expanded signature to uniformly express the condition that f is an isomorphism from $A^r(\overline{a})$ onto $A^r(\overline{b})$. Let $T = \{\varphi(\overline{a}), \neg\varphi(\overline{b})\} \cup \{\psi^r \mid r \in \omega\}$. It follows at once from 4.41 and 2, via the compactness theorem, that T is satisfiable. Let $A' \models T$ and let A be the reduct of A' to τ. It follows at once that there is an isomorphism g mapping $A^\infty(\overline{a})$ onto $A^\infty(\overline{b})$. We extend g to an automorphism h of A thereby contradicting $A \models \varphi[\overline{a}]$ and $A \not\models \varphi[\overline{b}]$.

For all $a \in A$, if $a \notin A^\infty(\overline{a}) \cup A^\infty(\overline{b})$, we let $h(a) = a$. For all $a \in (A^\infty(\overline{b}) - A^\infty(\overline{a}))$, there is a unique $b \in A^\infty(\overline{a})$ and $m \leq n$ such that $b \notin \operatorname{ran}(g)$ and $g^m(b) = a$. We let $h(a) = b$. It is easy to verify that h so defined is an automorphism. □

We make use of the following lemma in our proof of Theorem 4.7.

Lemma 4.5. *Let $\overline{a}$ and $\overline{b}$ be non-empty tuples of the same local type in ω-saturated models A and B respectively. Then $A^\infty(\overline{a}) \equiv B^\infty(\overline{b})$.*

Proof: We use the Ehrenfeucht-Fraïssé game to show that $A^\infty(\overline{a}) \equiv B^\infty(\overline{b})$. We proceed by induction on the number of moves. In particular, at the k-th round of the game, we maintain the property that $\langle a_1, \ldots, a_k, \overline{a}\rangle$ and $\langle b_1, \ldots, b_k, \overline{b}\rangle$ have the same local type.

Basis: For $k = 0$, $\langle \overline{a}\rangle$ and $\langle \overline{b}\rangle$ have the same local type in A and B by hypothesis.

Induction Step: By induction hypothesis, $\langle a_1, \ldots, a_{k-1}, \overline{a}\rangle$ and $\langle b_1, \ldots, b_{k-1}, \overline{b}\rangle$ have the same local type $\Lambda(x_1, \ldots, x_{k-1}, \overline{x})$. We will show that the local type $\Lambda'(x_1, \ldots, x_k, \overline{x})$ of every extension in $A^\infty(\overline{a})$ is also realized in $B^\infty(\overline{b})$, and vice versa. Suppose without loss of generality we are given a_k in $A^\infty(\overline{a})$ (resp. b_k in $B^\infty(\overline{b})$) such that $A^\infty(\overline{a}) \models \Lambda'[a_1, \ldots, a_k, \overline{a}]$. We want to find b_k in $B^\infty(\overline{b})$ (resp. a_k in $A^\infty(\overline{a})$) such that $B^\infty(\overline{b}) \models \Lambda'[b_1, \ldots, b_k, \overline{b}]$. Take any finite collection of local formulas $\Phi \subseteq \Lambda'$. Their conjunction is equivalent to an r-local formula $\varphi^r(x_1, \ldots, x_k, \overline{x})$ by taking r to be the maximum of the radii in Φ. Let d be the distance from a_k to $\overline{a}$ and observe that $\psi(x_1, \ldots, x_{k-1}, \overline{x}) = (\exists x_k)[\delta^d(\overline{x}, x_k) \wedge \varphi^r(x_1, \ldots, x_k, \overline{x})]$ is $(d + r)$-local by the triangle inequality. So $A^{d+r}(a_1, \ldots, a_{k-1}, \overline{a}) \supseteq A^r(a_1, \ldots, a_k, \overline{a}) \models \varphi[a_1, \ldots, a_k, \overline{a}]$ implies $A^{d+r}(a_1, \ldots, a_{k-1}, \overline{a}) \models \psi[a_1, \ldots, a_{k-1}, \overline{a}]$, and hence $\psi \in \Lambda$. This means $B^\infty(\overline{b}) \models \psi[b_1, \ldots, b_{k-1}, \overline{b}]$, and hence there is a b_k of distance d from $\overline{b}$ such that $B^r(b_1, \ldots, b_k, \overline{b}) \models \varphi[b_1, \ldots, b_k, \overline{b}]$, and hence $B^\infty(\overline{b}) \models \varphi^r[b_1, \ldots, b_k, \overline{b}]$, which makes $\Lambda'(b_1, \ldots, b_{k-1}, x_k, \overline{b})$ finitely consistent in $B^\infty(\overline{b})$. So by ω-saturation, Λ' is realized by a single b_k such that $B^\infty(\overline{b}) \models \Lambda'[b_1, \ldots, b_{k-1}, b_k, \overline{b}]$. But notice that since the distance from b_k to $\overline{b}$ must be d, it is in fact in $B^\infty(\overline{b})$. This completes the induction.

To finish the proof, notice that if $\langle a_1, \ldots, a_k, \overline{a}\rangle$ and $\langle b_1, \ldots, b_k, \overline{b}\rangle$ have the same local type, then the radius 1 formulas in that type witness a partial isomorphism. Therefore, $A^\infty(\overline{a}) \equiv B^\infty(\overline{b})$, since we have shown we can play the game between them for any number of moves. □

The following definition and theorem are due to [8].

Definition 4.6. A *basic local sentence* is one which says there are k disjoint neighborhoods each satisfying the same local formula:

$$(\exists x_1) \ldots (\exists x_k) \bigwedge \{d(x_i, x_j) > 2r \wedge \varphi^r(x_i) \mid 1 \leq i < j \leq k\}.$$

Theorem 4.7. *Every first-order sentence is equivalent to a Boolean combination of basic local sentences.*

Proof: By Corollary 1.0.7 of [4], it suffices to show that any two models A and B which satisfy the same basic local sentences are elementarily equivalent. This is obvious if one of the models is finite, because a basic local sentence can completely describe any finite model up to isomorphism. In case A and B are infinite, we may suppose, without loss of generality, that they are ω-saturated, since every structure has an ω-saturated elementary extension. If A has k distinct isomorphic copies of some component C, choose any element a in C, and let Σ be the infinite set of formulas which says its k variables are infinitely far apart and have the same local type as a:

$$\Sigma(x_1 \dots x_k) = \{d(x_i, x_j) > 2r \wedge \varphi^r(x_i) \mid 1 \leq i < j \leq k; A^r(a) \models \varphi[a]; 1 \leq r < \infty\}.$$

The existential closure of every finite subset of Σ is equivalent to a basic local sentence true in A, hence true in B, by hypothesis. So Σ is finitely consistent over B. But since B is saturated, B satisfies Σ, and therefore has k elements of the same local type that lie in separate components. By Lemma 4.5, these components in B are all elementarily equivalent to each other and to the corresponding components in A. Repeating this argument the other way around, from B to A, we see that each model must have either the same finite number of each component up to elementary equivalence, or an infinite number. Therefore, we can play any finite number of rounds of the Ehrenfeucht-Fraïssé game, and hence $A \equiv B$. □

4.2 Hanf locality

The next definition formulates a second notion of locality.

Definition 4.8. Suppose κ, λ, and μ are cardinals, A and B are τ-structures, and K is a class of τ-structures.

1. $\kappa \sim_\mu \lambda$ if and only if $\kappa = \lambda$ or $\kappa \geq \mu$ and $\lambda \geq \mu$.
2. Let $r \in \omega$ and $t \in \omega \cup \{\omega\}$. A is *Hanf threshold* r, *t-equivalent to* B (written $A \approx_{r,t} B$) if and only if for every pointed τ-structure C,
$$|\{a \in A \mid A^r(a) \cong C\}| \sim_t |\{b \in B \mid B^r(b) \cong C\}|.$$
3. A class of τ-structures J is *Hanf* r, *t-local on* K if and only if for all $A, B \in K$, if $A \approx_{r,t} B$, then $A \in J \iff B \in J$.
4. A class of τ-structures J is *Hanf threshold-local on* K if and only if J is Hanf r, t-local on K for some $r, t \in \omega$.

In the proof of the next theorem, we use the notion of a model pair, which is extensively used by [3] in similar spirit. Suppose A and B are τ-structures and that $A \cap B = \emptyset$. The *model pair* $\langle A, B\rangle$ is the $\tau \cup \{P\}$ structure with universe $A \cup B$, where the extension of the unary predicate P is A and the interpretation of all τ vocabulary is the union of its interpretation of A and B. Observe that for every first order τ sentence φ, there are $\tau \cup \{P\}$ sentences φ_A and φ_B such that

$$\langle A, B\rangle \models \varphi_A \iff A \models \varphi \text{ and } \langle A, B\rangle \models \varphi_B \iff B \models \varphi.$$

The following theorem is due to [7].

Theorem 4.9. *If J is first order definable and K is a class of structures of bounded degree, then J is Hanf threshold-local on K.*

Proof: Let J be class of τ-structures defined by a first order sentence φ, let K be a class of τ-structures of degree bounded by d, and suppose that J is not Hanf threshold-local on K. We formulate a theory T in the language of pairs of τ-structures expanded with a new binary predicate R which describes this situation. Observe that since the structures in K are of uniformly bounded degree, for each $r \in \omega$ there is a first order $\tau \cup \{P\}$ formula $\psi^r(x, y)$ such that for all $a \in A$ and $b \in B$,

$$\langle A, B\rangle \models \psi^r[a, b] \iff A^r(a) \cong B^r(b).$$

Let T be the following set of sentences.

1. $\varphi_A \wedge \neg\varphi_B$
2. the degree of A and of B are bounded by d
3. $(\forall x)(\forall y)(R(x, y) \to \psi^r(x, y))$, for every $r \in \omega$
4. $[(\forall x_1 \dots x_t)(\bigwedge_i P(x_i) \to (\exists y_1 \dots y_t)(\bigwedge_i \neg P(y_i) \wedge \bigwedge_{i<j}(x_i /= x_j \to y_i /= y_j) \wedge \bigwedge_i R(x_i, y_i)))] \wedge$ $[(\forall x_1 \dots x_t)(\bigwedge_i \neg(Px_i) \to (\exists y_1 \dots y_t)(\bigwedge_i P(y_i) \wedge \bigwedge_{i<j}(x_i /= x_j \to y_i /= y_j) \wedge \bigwedge_i R(y_i, x_i)))]$, for every $t \in \omega$

It follows at once from our hypotheses, via the compactness and Löwenheim-Skolem theorems, that T is satisfiable in a countable model pair $\langle A, B\rangle$ expanded by an interpretation F of R which is a bijection of A onto B with the property that

4.10. *for all $a \in A$ and for all $r \in \omega$,*

$$A^r(a) \cong B^r(F(a)).$$

Since A is of bounded degree, for each $r \in \omega$, $A^r(a)$ is finite, hence there are only finitely many isomorphisms from $A^r(a)$ onto $B^r(F(a))$. Observe that the set of all isomorphisms from $A^r(a)$ onto $B^r(F(a))$ for all $r \in \omega$ forms a finitely branching tree under inclusion, hence, by the König Infinity Lemma, this tree has an infinite path which represents an isomorphism from $A^{\infty}(a)$ onto $B^{\infty}(F(a))$. Thus, F is an isomorphism preserving bijection between the *pointed* components of A and B. We say a component C of A has finite character if for some $a \in C$, there are only finitely many $b \in C$ such that $A^{\infty}(a) \cong A^{\infty}(b)$, and similarly for components of B. It is easy to see that F induces a bijection between the components of A and B of finite character. By compactness, we may take countable elementary extensions A' of A and B' of B such that every component of A' which is not of finite character occurs infinitely often in A', up to isomorphism, and similarly for B'. It follows at once that $A' \cong B'$, which contradicts $A' \models \varphi$ and $B' \not\models \varphi$. □

As a corollary to Theorem 4.9, we show that connectivity is not mon-Σ^1_1 definable over the class of finite simple graphs, a result first established by [5]; our proof is an "infinitary" version of an argument presented in [7]. (A formula of second order logic is mon-Σ^1_1 if and only if it is of the form $(\exists\overline{P})\varphi$, where $\overline{P}$ is a sequence of unary predicate variables and φ is a first order formula.)

Theorem 4.11. *Connectivity is not mon-Σ^1_1 definable over the class of finite simple graphs.*

Proof: Suppose that

4.12. *the mon-Σ^1_1 sentence $(\exists\overline{P})\varphi(\overline{P})$ defines connectivity over the class of finite simple graphs.*

Let K_1 be the class of finite $\overline{P}$-colored simple cycles which satisfy $\varphi(\overline{P})$ and let K_2 be the class of pairs of all finite $\overline{P}$-colored simple cycles (a $\overline{P}$-colored simple graph has, in addition to a loop-free undirected edge relation, interpretations of the unary predicates in $\overline{P}$ – we shorten "$\overline{P}$-colored" to "colored" for the remainder). By Theorem 4.9, it follows from 4.12 that for some r_0 and t_0, if $D \approx_{r_0,t_0} C$ for some $C \in K_1$, then $D \in K_1$, for all finite colored simple graphs D. It is easy to see that

4.13. *for every r, t, and d, there is a first-order sentence characterizing graphs of degree at most d up to $\approx_{r,t}$.*

Therefore, in order to derive a contradiction from 4.12, it suffices to show that

4.14. *for every A, if A is pseudo-K_1, then A is pseudo-K_2.*

In order to verify 4.14, suppose that A is pseudo-K_1; by the Löwenheim-Skolem Theorem, we may suppose, without loss of generality, that $|A| > (2^{\aleph_0})^+$. It is easy to see that in this case A consists of a collection of colored bi-infinite simple chains. Moreover, it follows from 4.13 and the fact that A is pseudo-K_1 that for every r, t, there is a colored simple cycle $C_{r,t}$ with $A \approx_{r,t} C_{r,t}$. To complete the verification of 4.14, it therefore suffices to show that there is a colored simple cycle D such that

4.15. $C_{r,t} \oplus D \approx_{r,t} A$.

Since $|A| > (2^{\aleph_0})^+$, it follows at once that there is a colored bi-infinite chain $B \subseteq A$ which occurs infinitely often up to isomorphism in A. It is easy to see, by the pigeon-hole principle, that there is a $2r+1$-neighborhood N in B which occurs infinitely often up to isomorphism. Let N_1 and N_2 be disjoint copies of N occurring in B. It follows at once that every r-neighborhood occurring between them in B occurs infinitely often in A up to isomorphism. Let D be the colored cycle formed by snipping B at the midpoints of N_1 and N_2 and identifying the endpoints of the of the resulting finite simple chain. It is easy to see that this choice of D satisfies 4.15. □

5 Ultraproducts and directed reachability

In this section we present Henry Towsner's proof of the following result which was established by Miklos Ajtai and Ronald Fagin (see [1]).

Theorem 5.1. *Reachability is not mon-Σ^1_1 definable over the class of finite directed graphs.*

Proof: We are considering models of the language $\Rightarrow$ (a binary, asymmetric, irreflexive relation), $=$, and two constant symbols s, t, possibly augmented by finitely many unary predicate symbols. We write $u \Leftrightarrow v$ for $(u \Rightarrow v) \vee (v \Rightarrow u)$. Suppose that

5.2. *the mon-Σ^1_1 sentence $(\exists P_1) \ldots (\exists P_n)\varphi(P_1, \ldots, P_n)$ defines s, t-connectivity over the class of finite directed graphs.*

We derive a contradiction to 5.2 via the following lemmas. The use of the ultraproduct construction obviates the need for any game argument. On a finite structure $(X_n, \Rightarrow_n)$, there is a canonical choice of measure, the normalized counting measure $\mu_n(S) = \frac{|S|}{|X_n|}$. On an ultraproduct of finite structures, as in Lemma 5.3, μ denotes the Loeb-measure, which canonically lifts the normalized counting measure on the finite structures $\mathcal{X}_n$ to a measure on the σ-algebra generated by the internal subsets of their ultraproduct $\mathcal{X}$ (see [10]).

Lemma 5.3. *There exists a sequence of models $\mathcal{X}_i = (X_i, \Rightarrow_i)$ with $|X_i| \to \infty$ such that s, t are connected in each $\mathcal{X}_i$, and in any ultraproduct (by a nontrivial ultrafilter) $\mathcal{X}$,*

- *for μ-almost every u, the $\Leftrightarrow$-connected component of u is a tree (i.e. has no loops),*
- *if $A \subseteq X$ is internal with $\mu(A) > 0$ then for almost every $u \in X$, there are infinitely many $b \in A$ with $u \Leftrightarrow b$.*

Proof: The construction of the models $(\mathcal{X}_i, \Rightarrow_i)$ is essentially the one in Proposition 8.3 of [1]. Suppose that, in $\mathcal{X}$, the set L of u such that the $\Leftrightarrow$-connected component of u contains a loop has measure $\epsilon > 0$. Observe that $L = \bigcup_d L_d$ where L_d is the set of u such that there is a loop in the ball of radius d around u (where distance is $\Leftrightarrow$-distance). So if $\mu(L) > 0$ then there is some d with $\mu(L_d) > 0$. The set L_d is definable, so $L_d = \langle L_{d,i}\rangle$ where $L_{d,i}$ is the set of $u \in X_i$ such that there is a loop in the ball of radius d around u. Therefore there are infinitely many i such that $\mu_i(L_{d,i}) > 0$.

Write $\#(A, u) = |\{b \in A \mid u \Leftrightarrow b\}|$ if this is finite and $\#(A, u) = \infty$ otherwise. Let $A = \langle A_i\rangle$ be internal. Suppose that $\mu(A) > 0$ and the set of u with $\#(A, u) < \infty$ has positive measure. Then for some d, the set of u with $\#(A, u) < d$ has positive measure. Then there must be infinitely many i where the set of u with $\#(A_i, u) < \infty$ has positive measure.

So it suffices to arrange for the sequence of models $\mathcal{X}_i$ to have the following properties:
- for every d and every $\epsilon > 0$ there is an i such that for all $j \geq i$, $\mu_j(L_{d,j}) < \epsilon$, and
- for every d and every $\epsilon > 0$ there is an n such that for all $j \geq i$ and every set $A \subseteq X_j$ with $\mu_j(A) \geq \epsilon$, the set of u with $\#(A, u) < d$ has μ_j-measure $< \epsilon$.

The models $(X_i, \Rightarrow_i)$ will have the following form: $X_i = [1, i]$, $s^{X_i} = 1$, $t^{X_i} = i$, for all $a < i$, $a \Rightarrow_i a + 1$, and in addition we select a small set of pairs $a < b$ so that $b \Rightarrow_i a$. The only discretion we have is in choosing the set of "backwards" arrows.

Consider the random process where for each pair $a < b$, we decide whether $b \Rightarrow_i a$ holds randomly with probability $p(i) = \ln i/i$, with each pair determined independently. (As in [1], what really matters is that $p(i)$ grows faster than k/i for all k and slower than i^s/i for all $s > 0$.) We show that the properties we want hold with high probability, and so certainly at least one model with the desired properties exists.

We recall, for repeated use, the Chernoff bound. Suppose Y is a Bernoulli random variable—that is, $Y = \sum_{i \leq m} Y_i$ where each $Y_i \in \{0, 1\}$ is chosen randomly and independently with the probability $\mathbb{P}[Y_i = 1] = p$. Then the expected value of Y is $\mu = mp$ and for any $y \geq 0$,

$$\mathbb{P}[Y > (1 + y)\mu] \leq e^{-y^2\mu/(2+y)}, \ \mathbb{P}[Y < (1 - y)\mu] \leq e^{-y^2\mu/(2+y)}.$$

Let us write $B_r(u)$ for the ball of radius r around u. Observe that $|B_1(u)| = Y + 3$ (or $Y + 2$ if u is an endpoint) where Y is a Bernoulli random variable with mean $\ln n$, and so for a given u, the probability that $|B_1(u)| \geq 4\ln i = (1+3)\ln i$ is bounded by $e^{-9\ln i/5} = i^{-9/5}$. Therefore, the probability that *some* u has $|B_1(u)| \geq 4\ln i$ is bounded by $i^{-4/5}$, and in particular goes to 0 as $i \to \infty$. It follows that whenever $r \ll i$, with high probability for every u, $|B_r(u)| \leq \ln^k i$ for some constant k depending only on r.

We show that, with high probability, there are few cycles. Specifically, for fixed $r, d, \epsilon > 0$, we show that as i goes to ∞, with high probability the set of v such that $B_r(v)$ contains a loop of length $\leq d$ has density $< \epsilon$. Consider the following method for generating $\Rightarrow_i$. We first define a relation $\Rightarrow_i'$. We have $a \Rightarrow_i' a+1$ for all a, and we place the pairs $a < b$ in an arbitrary order, and successively decide in this order whether $b \Rightarrow_n' a$ will hold with probability $p(i)$, except that if $b \Rightarrow_i' a$ would create a loop of length $\leq d$, we do not include an edge. We then extend $\Rightarrow_i'$ to $\Rightarrow_i$ by deciding, for each pair we skipped in the creation of $\Rightarrow_i'$, whether $\Rightarrow_i$ should hold between this pair. This is equivalent to the original method for generating $\Rightarrow_i$ (we are making the same choices with the same probability, and the choices are independent; the only difference is the order in which we make the choices). Further, $\Rightarrow_i$ has a loop of length $\leq d$ exactly if we add an edge in the second stage. By taking i big enough, with high probability we have $|B_d(u)| \leq \ln^k i$ for all u. Therefore, for each u, we consider at most $\ln^k i$ possible edges, so there are $i\ln^k i$ possible edges added in the second stage, each with probability $\ln i/i$. The number of edges added is a Bernoulli random variable, and by the Chernoff bound again, when i is large, the probability is close to 1 that we add fewer than $4\ln^{k+1} i$ edges. Every loop in the graph contains one of these edges; for each of these edges, taking u to be an arbitrary endpoint, $|B_r(u)| \leq \ln^{k'} i$. Since $B_r(v)$ contains a loop of length $\leq d$ iff $v \in B_r(u)$ where u is one of the chosen endpoints, there are at most $4\ln^{k'(k+1)} i$ vertices v such that $B_r(v)$ contains a loop of length $\leq d$.

Now suppose we fix a set $A \subseteq X_i$ with $|A| \geq \epsilon i$ and fix some u and then generate $\Rightarrow_i$. Then $\#(A,u)$ is a binomial random variable with expected value $\epsilon \ln i$, so letting i be large enough so that $\epsilon\ln i \geq 2d$,

$$\mathbb{P}(\#(A,u) \leq d) \leq e^{-3\epsilon\ln i/8}.$$

Let $\delta = e^{-3\epsilon\ln i/8}$; observe that $\delta \to 0$ as $i \to \infty$. For any fixed set B with $|B| \geq \epsilon i$, the probability that for every $u \in B$, $\#(A,u) \leq d$ is bounded by $\delta^{|B|} \leq \delta^{\epsilon i}$. There are 2^i choices for A and 2^i possible choices for a set B, so the probability that there is any A and B with both $|A| \geq \epsilon i$, $|B| \geq \epsilon i$, and for every $u \in B$, $\#(A,u) \leq d$, is bounded by $2^i 2^i \delta^{\epsilon i} = e^{n(\ln 4 + \epsilon\ln\delta)}$. By taking i sufficiently large, and so δ sufficiently small, we may ensure that this probability approaches 0. Therefore, with high probability, for every set A with $|A| \geq \epsilon i$, $|\{u \mid \#(A,u)\}| < \epsilon i$. □

We may make the following simple observation:

Lemma 5.4. *Every $\Leftrightarrow$-connected component has measure* 0.

Proof: Suppose some connected component A has positive measure. Since almost every point belongs to a component with no loops, A must contain no loops. A itself is not internal, but it is a union of internal sets: pick any $u \in A$; then $A = \bigcup_r B_r(u)$. In particular, some $B_r(u)$ must have positive measure; let r be least such that $B_r(u)$ has positive measure. Since $B_r(u) \setminus B_{r-1}(u)$ is an internal set of positive measure, almost every v satisfies $\#(B_r(u) \setminus B_{r-1}(u), v) = \infty$. In particular, we may pick two points $a_0, a_1 \in B_r(u)$ such that $a_i \Leftrightarrow v$. Further, we may assume $v \notin B_{r-1}(u)$. This means a_0, a_1 each have paths to u not including v or a_{1-i}, and therefore combining these paths with v gives a loop in A, which is the desired contradiction. □

It follows from Lemma 5.2, 5.3, and Łoś's Theorem that there are internal sets $A_1, \dots, A_n$ such that the expansion $\mathcal{X} = (X, \Rightarrow, A_1, \dots, A_n)$ satisfies $\varphi(P_1, \dots, P_n)$.

Definition 5.5. If ϕ is a formula with a single free variable, we write $\mu(\phi)$ for $\mu(\{a \in X \mid \mathcal{X} \vDash \phi(a)\})$.

We say $tp(x)$ is *wide* if whenever $\phi \in tp(x)$, $\mu(\phi) > 0$. We say x is *central* if for every $\phi \in tp(x)$, $\mu(\{a \mid a < x \wedge \mathcal{X} \vDash \phi(a)\}) > 0$ and $\mu(\{a \mid x < a \wedge \mathcal{X} \vDash \phi(a)\}) > 0$.

Clearly central elements have wide type.

Lemma 5.6. *Almost every x is central.*

Proof: We show that for every ϵ, the non-central elements are contained in a set of measure $\leq \epsilon$. Let $\phi_1, \dots, \phi_n, \dots$ be an enumeration of the formulas with exactly one free variable (and, ambiguously, the subsets of $\mathcal{X}$ they define). We will define $U = \bigcup U_i$ where $\mu(U_i) \leq \epsilon 2^{-i}$. If $\mu(\phi_i) = 0$, set $U_i = \phi_i$. If $\mu(\phi_i) > 0$, define

$$U_i = \{x \mid \mu(\{a \mid a < x \wedge \mathcal{X} \vDash \phi_i(a)\}) < \epsilon 2^{-(i+1)}\}$$
$$\cup \ \{x \mid \mu(\{a \mid x < a \wedge \mathcal{X} \vDash \phi_i(a)\}) < \epsilon 2^{-(i+1)}\}.$$

That is, U_i is the "bottom" and "top" $\epsilon 2^{-(i+1)}$ of the set defined by ϕ_i.

If x is not central then there is some ϕ_i witnessing this, and by construction, we have included x in $U_i \subseteq U$. □

Now we argue as follows. Consider those x such that x and $x + 1$ are both central, and such that the $\Leftrightarrow$-connected component of x contains no loops; this is a set of measure 1, so we may take some c with these properties.

Lemma 5.7. *There are infinitely many a with $tp(a) = tp(c + 1)$ and $c \Rightarrow a$.*

Proof: For any $\phi \in tp(c+1)$, since $c+1$ is central, $\mu(\{a \mid a < c+1 \wedge \mathcal{X} \vDash \phi(a)\}) > 0$. Therefore, by choice of $\mathcal{X}$, there are infinitely many a such that $\mathcal{X} \vDash \phi(a)$, $a < c$, and $c \Leftrightarrow a$; but this implies $c \Rightarrow a$. Since this holds for each ϕ, saturation ensures that there are actually infinitely many a with $tp(a) = tp(c+1)$ and $c \Rightarrow a$. □

Similarly,

Lemma 5.8. *There are infinitely many b with $tp(b) = tp(c)$ and $c+1 \Rightarrow b$.*

Lemma 5.9. *There is a d with $tp(d) = tp(c)$ not in the same $\Leftrightarrow$-connected component as c.*

Proof: For every r $\mu(B_r(c)) = 0$, while for each $\phi \in tp(c)$, $\mu(\phi) > 0$, so $\mu(\phi \setminus B_r(c)) > 0$ as well. Since this holds for each r and each $\phi \in tp(c)$, saturation implies that there is a d with $tp(d) = tp(c)$ and $d \notin \bigcup_r B_r(c)$. □

Let $\mathcal{X}'$ be the model identical to $\mathcal{X}$ except that $c \not\Rightarrow c+1$. It is immediate that $\mathcal{X}'$ is an ultraproduct of structures $\mathcal{X}'_i$ almost all of which are not s,t-connected; hence, $\mathcal{X}' \not\vDash \varphi(P_1, \ldots, P_n)$. On the other hand, we claim $\mathcal{X}$ is isomorphic to $\mathcal{X}'$, which completes the *reductio*.

We define $\pi : \mathcal{X} \to \mathcal{X}'$ to be the identity outside the connected components of c and d.

We define π on the connected component of c so $\pi(c) = c$ but the image does not include $c+1$: let $A = \{a \mid tp(a) = tp(c+1) \wedge c \Rightarrow a\}$, and we define π"$A = A \setminus \{c+1\}$, which we can do since A is infinite. Since the connected component of c is a tree and the type of a point determines the type of its neighbors, we can define π on the rest of the $\Leftrightarrow$-connected component of c.

On the connected component of d, we do the reverse: let $B = \{b \mid tp(b) = tp(c+1) \wedge d \Rightarrow b\}$, and define π"$B = B \cup \{c+1\}$. □

Acknowledgement: Henry Towsner is partially supported by NSF grant DMS-1340666. We are grateful to Roman Rabinovich and Marlene Gregorek for valuable comments on an earlier draft of this paper.

Bibliography

[1] Miklos Ajtai and Ronald Fagin. Reachability is harder for directed than for undirected finite graphs. *Journal of Symbolic Logic*, 55:113–150, 1990.

[2] P. Dittmann. Ultraproducts as a tool for first-order inexpressibility in the finite and infinite. arXiv:1310.3137, 2013.

[3] K. Doets. Uniform short proofs for classical theorems. *Notre Dame Journal for Formal Logic*, 42:121–127, 2001.

[4] Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite model theory*. Springer, 1995.

[5] R. Fagin. Monadic generalized spectra. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 21:89–96, 1975.
[6] R. Fagin. Probabilities on finite models. *Journal of Symbolic Logic*, 41:50–58, 1976.
[7] R. Fagin, L. Stockmeyer, and M. Y. Vardi. On monadic NP vs. monadic co-NP. *Information and Computation*, 120(1):78–92, 1995.
[8] H. Gaifman. On local and nonlocal properties. In J. Stern, editor, *Logic Colloquium '81*, pages 105–135. North Holland, 1982.
[9] H. Gaifman and M. Y. Vardi. A simple proof that connectivity is not first-order definable. *Bulletin of the European Association for Theoretical Computer Science*, 26:43–45, 1985.
[10] Peter A. Loeb. Conversion from nonstandard to standard measure spaces and applications in probability theory. *Transactions of the American Mathematical Society*, 211:113–122, 1975.
[11] Martin Otto. Model theoretic methods for fragments of FO and special classes of (finite) structures. In Javier Esparza, Christian Michaux, and Charles Steinhorn, editors, *Finite and Algorithmic Model Theory*, pages 271–341. Cambridge University Press, 2011.
[12] G. Turan. On the definability properties of finite graphs. *Discrete Mathematics*, 49:291–302, 1984.
[13] J. Väänänen. Pseudo-finite model theory. *Matematica Contemporanea*, 24:169–183, 2003.