THE KARP COMPLEXITY OF UNSTABLE CLASSES

M. C. LASKOWSKI AND S. SHELAH

Abstract. A class $\mathbf{K}$ of structures is *controlled* if, for all cardinals λ, the relation of $L_{\infty,\lambda}$-equivalence partitions $\mathbf{K}$ into a set of equivalence classes (as opposed to a proper class). We prove that the class of doubly transitive linear orders is controlled, while any pseudo-elementary class with the ω-independence property is not controlled.

1. Introduction

One of the major accomplishments of model theory has been the discovery of a dividing line between those theories in a countable language whose models can be described up to isomorphism by a reasonable set of invariants and those whose models cannot be so described. Models of *classifiable* theories are described up to isomorphism by an 'independent tree' of countable elementary submodels, while the isomorphism type of any unclassifiable theory cannot be described by any reasonable set of invariants (see [**?**]). Unfortunately, the great majority of classes of structures studied in mathematics are unstable, and thus fall on the 'non-structure' side of this divide. Thus, it is desirable to search for dividing lines between unstable classes of structures. Our thesis is that while an unstable (pseudo-elementary) class necessarily has the maximal number of non-isomorphic models in every uncountable cardinality, it is still possible to assign a *set of invariants* to some unstable classes of structures. In some cases (see e.g., Example 3.6) the large number of non-isomorphic models is due simply to our ability to code arbitrary stationary sets into the skeletons of Ehrenfeucht-Mostowski models. In other words, for some classes of structures the reason for the non-isomorphism of two structures in the class need not be very robust. Indeed, in such cases the structures can be forced to be isomorphic by a forcing that merely adds a new closed, unbounded subset

Date: December 21, 2001.

1991 *Mathematics Subject Classification.* 03C.

Partially supported by NSF Research Grants DMS 9403701 and DMS 9704364.

The authors thank the U.S.-Israel Binational Science Foundation for its support of this project. This is item 560 in Shelah's bibliography.

of some cardinal to the universe. That is, although they are nonisomorphic, the structures are not very different from each other. On the other hand, for other classes of structures (see Theorem 2.5) there are more serious obstructions to a structure theorem.

Our ultimate goal is to determine to which unstable classes of structures one can associate a reasonable set of structural invariants. These invariants need not (and typically will not) determine the structures up to isomorphism. Instead, we ask that any two structures with the same invariants be very much the same. In this paper we focus on $L_{\infty,\lambda}$-equivalence for various cardinals λ and ask which unstable classes are partitioned into only a *set* of equivalence classes (as opposed to a proper class). We call a class $\mathbf{K}$ *controlled* if $\mathbf{K}$ has only a set of $L_{\infty,\lambda}$-equivalence classes for all cardinals λ. Typically, $L_{\infty,\lambda}$-equivalence does not characterize models up to isomorphism even when we fix the cardinalities of the models. (In [**?**] the second author shows that for any unstable pseudo-elementary class and any uncountable regular cardinal λ, there are 2^λ non-isomorphic models of size λ that are $L_{\infty,\lambda}$-equivalent.) However, in some sense two $L_{\infty,\lambda}$-equivalent structures of the same cardinality are very much the same. For instance, if one uses the back-and-forth system witnessing their equivalence as a notion of forcing, then the two structures will become isomorphic in the corresponding forcing extension.

In this paper we obtain two complementary results. On one hand, in Section 3 we analyze the pseudo-elementary class $\mathbf{K}_{2\mathrm{tr}}$ of doubly transitive linear orders. This class is unstable, hence the stigma of non-structure applies. Despite this, we prove that $KC_{\mu^+}(\mathbf{K}_{2\mathrm{tr}}) \leq \omega$ (see Definition 2.3) for all uncountable cardinals μ, hence $\mathbf{K}_{2\mathrm{tr}}$ is controlled. This is one of very few theorems in which an unstable pseudo-elementary class shows any sign of structure. On the other hand, in Section 4 we prove that any pseudo-elementary class with the ω-independence property (see Definition 4.4) is not controlled. In fact, if the language used in describing $\mathbf{K}$ is countable then $KC_\lambda(\mathbf{K}) = \infty$ for all cardinals $\lambda \geq \aleph_3$.

There is still much that we do not know about the notion of control. A fundamental question that remains open is whether there is an unstable *elementary* class that is controlled. We conjecture, and hope to prove, that any pseudo-elementary class with the independence property is not controlled; this would substantially strengthen our second result.

2. Controlled classes

In this section we state a series of definitions that lead to the concept of a class of structures being controlled (see Definition 2.5). We apply these definitions to the theory of dense linear orders to illustrate why it is desirable to consider the λ-Karp complexity of a class for uncountable cardinals λ. We first reintroduce the notion of a partial isomorphism, but with a slight variation. As we are only concerned with the definable subsets of structures (and not their quantifier complexity) we insist that all partial isomorphisms are elementary maps.

Definition 2.1 Given two elementarily equivalent structures M and N in the same language and an infinite cardinal λ, a *λ-partial isomorphism* is a partial elementary map with domain of cardinality less than λ, that is: a function f from a subset D of M into N of size less than λ satisfying

$$M \models \varphi(d_1 \dots d_n) \qquad \text{if and only if} \qquad N \models \varphi(f(d_1), \dots, f(d_n))$$

for all formulas $\varphi(x_1, \dots, x_n)$ of the language and all $d_1, \dots, d_n$ from D. We denote the family of λ-partial isomorphisms by $\mathcal{F}_\lambda(M, N)$. If $M = N$ we simply write $\mathcal{F}_\lambda(M)$.

The complexity of $\mathcal{F}_\lambda(M)$ is a measure of how deeply one needs to look to understand the relationship of a small subset (i.e., of size less than λ) with the rest of the model. In order to measure this depth we endow the family with the following rank.

Definition 2.2 For $f \in \mathcal{F}_\lambda(M, N)$,

(1) $\mathrm{Rank}(f) \geq 0$ always;
(2) For α limit, $\mathrm{Rank}(f) \geq \alpha$ if and only if $\mathrm{Rank}(f) \geq \beta$ for all $\beta < \alpha$;
(3) $\mathrm{Rank}(f) \geq \alpha + 1$ if and only if
  (a) for all $C \subseteq M$ of size less than λ, there is $g \in \mathcal{F}_\lambda(M, N)$ extending f with $C \subseteq \mathrm{dom}(g)$ and $\mathrm{Rank}(g) \geq \alpha$; and
  (b) dually, for all $C \subseteq N$ of size less than λ, there is $g \in \mathcal{F}_\lambda(M, N)$ extending f with $C \subseteq \mathrm{range}(g)$ and $\mathrm{Rank}(g) \geq \alpha$.

The *λ-Karp complexity* $KC_\lambda(M, N)$ of the pair of structures M, N is the least ordinal α such that $\mathrm{Rank}(f) \geq \alpha$ implies $\mathrm{Rank}(f) \geq \alpha + 1$ for all $f \in \mathcal{F}_\lambda(M, N)$. Again, if $M = N$ we simply write $KC_\lambda(M)$.

The λ-Karp complexity of a structure is related to the notions of $L_{\infty,\lambda}$-Scott height and back-and-forth systems. It is a routine diagram-chasing exercise to show that if two structures M and N are $L_{\infty,\lambda}$-equivalent (hence there is a back-and-forth system in $\mathcal{F}_\lambda(M,N)$) then $KC_\lambda(M) = KC_\lambda(N)$.

If one fixes the signature, then for any cardinals κ and λ it is easy to find an ordinal bounding the λ-Karp complexity of any structure of that signature of size at most κ. By contrast, whether or not there is a upper bound on the λ-Karp complexities of all structures in a class $\mathbf{K}$ that does not depend on κ provides a robust dichotomy between classes. This is demonstrated by the following definition and proposition. The reader is referred to [**?**] for the undefined notions.

Definition 2.3 For $\mathbf{K}$ a class of structures, the λ-*Karp complexity* of $\mathbf{K}$, written $KC_\lambda(\mathbf{K})$, is the supremum of the ordinals $KC_\lambda(M)$ among all $M \in \mathbf{K}$ if the supremum exists. Otherwise, we set $KC_\lambda(\mathbf{K}) = \infty$.

Proposition 2.4. *The following conditions are equivalent for a class of structures* $\mathbf{K}$ *and an infinite cardinal* λ*.*

(1) $KC_\lambda(\mathbf{K}) < \infty$;
(2) *The relation of* $L_{\infty,\lambda}$*-equivalence on* $\mathbf{K}$ *has only a set of equivalence classes;*
(3) *There are only a set of* $L_{\infty,\lambda}$*-types of subsets of size less than* λ *realized in elements of* $\mathbf{K}$*;*
(4) *There are only a set of distinct* $L_{\infty,\lambda}$*-Scott sentences among the elements of* $\mathbf{K}$*;*
(5) *There is a cardinal* κ *such that the notions of* $L_{\kappa,\lambda}$*-equivalence and* $L_{\infty,\lambda}$*-equivalence coincide on* $\mathbf{K}$*.*

Proof. The implication (2) $\Rightarrow$ (1) follows from the observation that λ-Karp complexity is preserved under $L_{\infty,\lambda}$-equivalence. The implications (1) $\Rightarrow$ (4) $\Rightarrow$ (5) $\Rightarrow$ (3) $\Rightarrow$ (2) all follow easily. ∎

When $\lambda = \aleph_0$ the λ-Karp complexity often does not yield much information about the inherent complexity of a class $\mathbf{K}$. For example, if $\mathbf{K}$ is the class of models of an $\aleph_0$-categorical theory, then $KC_{\aleph_0}(\mathbf{K}) = 0$ since every model is $\aleph_0$-homogeneous. However, our thesis is that for larger λ, λ-Karp complexity gives a good measure of the complexity of the class. It follows from Proposition 2.4(3) that if $KC_\lambda(\mathbf{K}) = \infty$ for some cardinal λ, then $KC_\kappa(\mathbf{K}) = \infty$ for all larger cardinals κ. This leads us to the crucial definition of the paper.

Definition 2.5 A class $\mathbf{K}$ of structures is *controlled* if $KC_\lambda(\mathbf{K}) < \infty$ for all infinite cardinals λ.

Note that if a class $\mathbf{K}$ is controlled, then it follows from Proposition 2.4(2) that for every cardinal λ, the relation of $L_{\infty,\lambda}$-equivalence partitions $\mathbf{K}$ into only a set of equivalence classes (as opposed to a proper class). Continuing our example, $KC_{\aleph_0}(DLO) = 0$, as DLO, the theory of dense linear orders with no endpoints is $\aleph_0$-categorical. However, this observation hides the fact that one can code arbitrary ordinals into dense linear orders. This ability to code ordinals implies that the class DLO is not controlled. In fact, $KC_\lambda(DLO) = \infty$ for all uncountable cardinals λ. To see this, fix an uncountable cardinal λ and, for each non-zero ordinal α, let J_α be the linear order with universe $(\eta \cdot \lambda) \cdot \alpha$, where η denotes the order type of the rationals. In light of Proposition 2.4(2) it suffices to show that J_α is not $L_{\infty,\lambda}$-equivalent to J_β whenever $\alpha \neq \beta$. So choose non-zero ordinals α and β such that J_α is $L_{\infty,\lambda}$-equivalent to J_β. Let E be the equivalence relation such that $E(x, y)$ if and only if there are fewer than λ elements between x and y. Since E is expressible in the logic $L_{\infty,\lambda}$, this implies that the condensation J_α/E is $L_{\infty,\lambda}$-equivalent to J_β/E. But $(J_\alpha/E, \leq) \simeq (\alpha, \leq)$, $(J_\beta/E, \leq) \simeq (\beta, \leq)$, and it is readily checked that distinct ordinals are not even $L_{\infty,\omega}$-equivalent. Hence α must equal β.

3. Doubly transitive linear orders

In this section we investigate the class $\mathbf{K}_{2\mathrm{tr}}$ of infinite doubly transitive linear orders. That is, $(I, \leq) \in \mathbf{K}_{2\mathrm{tr}}$ if and only if the linear order I is dense with no endpoints and for all pairs $a < b$, $c < d$ from I, the interval $[a, b]$ is isomorphic to the interval $[c, d]$. Such orders arise naturally: The underlying linear order of any ordered field is necessarily doubly transitive. Clearly, there is only one countable structure in $\mathbf{K}_{2\mathrm{tr}}$ up to isomorphism. The class $\mathbf{K}_{2\mathrm{tr}}$ is a pseudo-elementary (PC) class that is visibly unstable, so by [**?**] there are 2^λ non-isomorphic structures in $\mathbf{K}_{2\mathrm{tr}}$ of size λ for all uncountable cardinals λ. Further, by [**?**], for all uncountable regular cardinals λ there is a family of 2^λ structures in $\mathbf{K}_{2\mathrm{tr}}$ of size λ that are $L_{\infty,\lambda}$-equivalent, yet pairwise non-embeddable.

Nonetheless, the class of doubly transitive linear orders is not entirely without structure. There are natural 'invariants' one can associate with such orders. These invariants will not determine the orders up to isomorphism, but they will be sufficient to demonstrate that the λ-Karp complexity of $\mathbf{K}_{2\mathrm{tr}}$ is bounded for all cardinals λ.

The most natural invariant of a doubly transitive linear order is the isomorphism type of its closed intervals. Accordingly, we call $I_0, I_1 \in$

$\mathbf{K}_{2\mathrm{tr}}$ *locally isomorphic* and write

$$I_0 \sim I_1$$

if $[a,b] \simeq [c,d]$ for $a < b$ from I_0 and $c < d$ from I_1. Evidently local isomorphism is an equivalence relation on $\mathbf{K}_{2\mathrm{tr}}$ and $I \sim J$ for any infinite convex subset $J \subseteq I$, if $I \in \mathbf{K}_{2\mathrm{tr}}$.

The second invariant was developed by Droste and Shelah in [**?**]. The definitions that follow are slight adaptations of similar notions used there. The most notable variation is that in [**?**] there is no bound on the number of levels of the *decomposition tree* and the cardinals λ_η can be any uncountable regular cardinal.

For the whole of this section, fix an uncountable cardinal μ.

Definition 3.1 A *μ-decomposition tree* is a subtree T of $\bigcup\{{}^\alpha\mu : \alpha < \mu^+\}$ satisfying:

(1) T is downward closed, i.e., $\eta \in T$ implies $\eta|\alpha \in T$ for all $\alpha < lg(\eta)$;
(2) If $lg(\eta)$ is a limit ordinal or 0 and $\eta|\alpha \in T$ for all $\alpha < lg(\eta)$ then $\eta \in T$ and η has exactly two immediate successors; more specifically, we require $Succ_T(\eta) = \{\eta\frown\langle 0\rangle, \eta\frown\langle 1\rangle\}$;
(3) If $\eta \in T$ and $lg(\eta)$ is a successor ordinal, then either $Succ_T(\eta) = \emptyset$ or $Succ_T(\eta) = \{\eta\frown\langle\alpha\rangle : \alpha \in C\}$ for some club subset C of a regular cardinal $\lambda_\eta \in [\aleph_1, \mu]$.

Let $T^* = \{\eta \in T : lg(\eta)$ is a successor ordinal$\}$.

We define a linear order on T^* which is a cross between lexicographic and antilexicographic order. To every node η of T^* we first associate a direction $dir(\eta) \in \{\text{LEFT}, \text{RIGHT}\}$. Suppose $lg(\eta) = \delta + n$, where δ is a limit ordinal or 0 and $n \in \omega$. Then

- $dir(\eta) = \text{LEFT}$ if $\eta(\delta) + n$ is even;
- $dir(\eta) = \text{RIGHT}$ if $\eta(\delta) + n$ is odd.

The idea is that if $dir(\eta)$ =LEFT, then the successors of η will all be to the left of η. Each of these successors will have direction RIGHT, so their successors will be to their right and so forth. Formally, the linear order $<^{T^*}$ is defined by the following clauses.

- If $\eta \lhd \nu$ then $\eta <^{T^*} \nu$ if and only if $dir(\eta)$ =RIGHT;
- If η, ν are incomparable, let γ be least such that $\eta(\gamma) \neq \nu(\gamma)$ and let $\rho = \eta|\gamma$.
  - If γ is a limit ordinal or 0 then $\eta <^{T^*} \nu$ if and only if $\eta(\gamma) = 0$ and $\nu(\gamma) = 1$;
  - If γ is a successor ordinal (so $\rho \in T^*$) and $dir(\rho)$ =LEFT then $\eta <^{T^*} \nu$ if and only if $\eta(\gamma) < \nu(\gamma)$;

– If γ is a successor ordinal and $dir(\rho)$ =RIGHT, then $\eta <^{T^*} \nu$ if and only if $\eta(\gamma) > \nu(\gamma)$.

The following definition differs slightly from normal usage as we include the endpoints.

Definition 3.2 For I a dense linear order, the *Dedekind completion* of I is the linear order $(\overline{I}, \leq^{\overline{I}})$ with universe

$$\overline{I} = \{A \subseteq I : A \text{ downward closed with no largest element}\}$$

and $A \leq^{\overline{I}} B$ if and only if $A \subseteq B$. We let $-\infty$ denote the smallest element of $\overline{I}$ and $+\infty$ denote the largest. To simplify notation we identify the element $a \in I$ with $\{x \in I : x < a\} \in \overline{I}$ and write e.g., $I \subseteq \overline{I}$. If J is a convex subset of I, then $\overline{J}$ denotes the smallest closed interval in $\overline{I}$ that contains J and we identify $\overline{J}$ with the Dedekind completion of J.

Definition 3.3 A *μ-representation* of a linear ordering I is a pair (T, g), where T is a μ-decomposition tree and $g : T^* \to \overline{I}$ is an order-preserving function satisfying the following conditions:

1: $g(\langle 0\rangle) = -\infty$, $g(\langle 1\rangle) = +\infty$;
2: If $lg(\eta) = \gamma + 1$, where $\gamma > 0$ is a limit ordinal, let D be the largest interval $[a, b]$ of $\overline{I}$ such that for all successor ordinals $\alpha < \gamma$, D is between $g(\eta|\alpha)$ and $g(\eta|\alpha + 1)$.
  (1) If $\eta(\gamma) = 0$ then $g(\eta) = a$;
  (2) If $\eta(\gamma) = 1$ then $g(\eta) = b$;
  (3) If $a = b$ then we call η *degenerate*.
3: If $dir(\eta)$=LEFT then
  (1) η is maximal in T if and only if one of the three conditions hold:
    (a) η is degenerate;
    (b) $\mathrm{cof}(g(\eta)) = \aleph_0$;
    (c) $\mathrm{cof}(g(\eta)) > \mu$;
  (2) If η is not maximal in T, then $Succ_T(\eta) = \{\eta^\frown\langle\alpha\rangle : \alpha \in C\}$ for some club subset of $\mathrm{cof}(g(\eta))$, and $\{g(\eta^\frown\langle\alpha\rangle) : \alpha \in C\}$ is continuous, strictly increasing, and has supremum $g(\eta)$.
3^***:** If $dir(\eta)$ =RIGHT then
  (1) η is maximal in T if and only if one of the three conditions hold:
    (a) η is degenerate;
    (b) $\mathrm{coi}(g(\eta)) = \aleph_0$;
    (c) $\mathrm{coi}(g(\eta)) > \mu$;

(2) If η is not maximal in T, then $Succ_T(\eta) = \{\eta\frown\langle\alpha\rangle : \alpha \in C\}$ for some club subset of $\mathrm{cof}(g(\eta))$, and $\{g(\eta\frown\langle\alpha\rangle) : \alpha \in C\}$ is continuous, strictly decreasing, and has infimum $g(\eta)$.

A μ-representation (T, g) partitions $\overline{I}$ into a set of intervals $\{I_\eta : \eta \in T^*\}$ where $I_{\eta\frown\langle 0\rangle} = I_{\eta\frown\langle 1\rangle} = (g(\eta\frown\langle 0\rangle), g(\eta\frown\langle 1\rangle))$ for all $\eta \in T$ of limit length, and if $Succ_T(\eta) = \{\eta\frown\langle\alpha\rangle : \alpha \in C\}$ for a club C then

$$I_{\eta\frown\langle\alpha\rangle} = \begin{cases} (g(\eta\frown\langle\alpha\rangle), g(\eta\frown\langle\alpha^+\rangle)) & \text{if } dir(\eta) = \text{LEFT;} \\ (g(\eta\frown\langle\alpha^+\rangle), g(\eta\frown\langle\alpha\rangle)) & \text{if } dir(\eta) = \text{RIGHT} \end{cases}$$

where α^+ is the least element of C larger than α. It is easily shown by induction that the intervals $\{I_\eta : \eta \in T^* \cap {}^\alpha\mu\}$ are pairwise disjoint for any fixed successor ordinal α.

For any dense linear order I, one can build a μ-representation (T, g) of I level by level by successively choosing a continuous, strictly increasing [or decreasing] sequence $\langle g(\eta\frown\langle\alpha\rangle) : \alpha \in \lambda_\eta\rangle$ from the interval I_η. At first blush, it appears that one has considerable freedom in such a construction. However, our freedom is considerably limited by the following observation.

Observation Let J be any linear order of cofinality $\lambda > \aleph_0$. For any club subsets C_1, C_2 of λ and any two continuous, strictly increasing, cofinal sequences $\langle a_i : i \in C_1\rangle$ and $\langle b_i : i \in C_2\rangle$ in J, the set $D = \{i \in C_1 \cap C_2 : a_i = b_i\}$ is a club subset of λ.

By repeatedly applying this observation to a pair of μ-representations of a linear order, we see that they must 'agree on a club.' More precisely, call a subtree T' of a μ-decomposition tree T a *club subtree* if T' itself is a μ-decomposition tree and, for each $\eta \in T'$ that is not maximal in T', $Succ_T(\eta)$ and $Succ_{T'}(\eta)$ are both indexed by club subsets of the same regular cardinal. If (T_1, g_1) and (T_2, g_2) are two μ-representations of I, then by using the observation above at each node there is a μ-representation (T, g) of I such that T is a club subtree of both T_1 and T_2 with $g(\eta) = g_1(\eta) = g_2(\eta)$ for all $\eta \in T^*$. More generally we have the following definition and lemma.

Definition 3.4 A subset A of a μ-decomposition tree T is *closed* if A is downward closed, (i.e., if $\eta \in A$ then $\eta|\alpha \in A$ for all $\alpha < lg(\eta)$) and A is closed under successor, (i.e., if $\eta \in A$ then $Succ_T(\eta) \subseteq A$).

Note that for any subset $A \subseteq T$ of size at most μ, there is a closed subset $B \supseteq A$ of size at most μ.

Lemma 3.5. *Suppose (T, g) is a μ-representation of I_0, $S \subseteq T$ is closed, and $f_0, f_1 : \overline{I}_0 \to \overline{I}_1$ are order-preserving, continuous partial functions whose domains contain $\{g(\eta) : \eta \in S \cap T^*\}$ that satisfy*

$f_0(-\infty) = f_1(-\infty)$ and $f_0(+\infty) = f_1(+\infty)$. Then there is a club subtree $Y \subseteq T$ such that

$$f_0(g(\eta)) = f_1(g(\eta))$$

for all $\eta \in S \cap Y^*$.

Proof. We construct Y by induction on the levels of T. Assume that we have found Y_γ, a club subtree of $T \cap \bigcup\{{}^\beta\mu : \beta < \gamma\}$ such that $f_0(g(\eta)) = f_1(g(\eta))$ for all $\eta \in S \cap Y^*_\gamma$. If γ is a limit ordinal or 0 then put $Y_{\gamma+1} = Y_\gamma \cup \{\eta \in {}^\gamma\mu\}$ and there is nothing to check. If $\gamma = \delta + 1$ where δ is a limit ordinal or 0, let $Y_{\gamma+1} = Y_\gamma \cup \bigcup\{Succ_T(\eta) : \eta \in Y_\gamma \cap {}^\gamma\mu\}$. Now if $\eta \in S \cap Y_{\gamma+1}$ for some $\eta \in {}^\gamma\mu$, then $\eta|\beta \in S \cap Y_\gamma$ for all $\beta < \delta$, so $f_0(g(\eta|\beta)) = f_1(g(\eta|\beta))$ for all $\beta < \delta$. As both f_0 and f_1 are order-preserving and continuous, it follows that $f_0(g(\eta\frown\langle i\rangle)) = f_1(g(\eta\frown\langle i\rangle))$ for $i = 0, 1$ so our inductive hypothesis is maintained.

Finally, assume $\gamma = \delta + n$, where δ is a limit ordinal or 0 and $n > 1$. Fix $\eta \in Y_\gamma \cap {}^{\delta+n-1}\mu$ and we specify its successors in $Y_{\gamma+1}$:

- If $\eta \notin S$ or if $Succ_T(\eta) = \emptyset$, then let $Succ_{Y_{\gamma+1}}(\eta) = Succ_T(\eta)$ and there is no problem.
- If $\eta \in S$ and $Succ_T(\eta) = \{\eta\frown\langle\alpha\rangle : \alpha \in C\}$ for some club subset of an uncountable regular cardinal λ_η, then our hypotheses imply that $f_0(g(\eta)) = f_1(g(\eta))$ and $\{g(\eta\frown\langle\alpha\rangle) : \alpha \in C\}$ is a continuous, strictly increasing (or decreasing) sequence converging to $g(\eta)$. Thus, as both f_0 and f_1 are order-preserving and continuous, there is a club $C' \subseteq C$ such that $f_0(g(\eta\frown\langle\alpha\rangle)) = f_1(g(\eta\frown\langle\alpha\rangle))$ for all $\alpha \in C'$. So put $Succ_{Y_{\gamma+1}}(\eta) = \{\eta\frown\langle\alpha\rangle : \alpha \in C'\}$. ∎

As noted above, these invariants are not sufficient to determine the isomorphism type of an element of $\mathbf{K}_{2\mathrm{tr}}$. In particular, the second invariant does not specify which elements of the representation are in I (as opposed to $\overline{I}$). This affords considerable freedom in choosing the isomorphism type of the order. The family of structures in the example below was first studied by Conway [?] and was later used as an example by Nadel and Stavi [?].

Example 3.6 There is a family of $2^{\aleph_1}$ locally isomorphic, $L_{\infty,\aleph_1}$-equivalent doubly transitive linear orders of size $\aleph_1$, all of whom have isomorphic $\aleph_1$-representations; yet the orders are pairwise non-embeddable.

Let $\mathcal{S}$ be a collection of $2^{\aleph_1}$ stationary subsets of $\omega_1 \setminus \{0\}$ with $X \setminus Y$ stationary for all distinct $X, Y \in \mathcal{S}$ (see [?] for a construction of such a family). As notation, let $\mathbb{Q}^{\geq 0}$ be the set $\mathbb{Q} \cap [0, \infty)$. For $X \in \mathcal{S}$, let

$$I_X = \sum_{i \in \omega_1} J_i^X \qquad \text{where} \qquad J_i^X = \begin{cases} \mathbb{Q} & \text{if } i \notin X; \\ \mathbb{Q}^{\geq 0} & \text{if } i \in X. \end{cases}$$

Clearly $(a, b) \cong \mathbb{Q}$ for all $a < b$ from I_X, so $I_X \sim I_Y$ for all $X, Y \in \mathcal{S}$. It was first noted by Silver that for any sets $X, Y \in \mathcal{S}$, the set $\mathcal{B}(X, Y)$ of all order-preserving partial functions $f : I_X \to I_Y$, whose domain D is a proper initial segment of I_X such that $I_X \setminus D$ has no least element, and whose range R is a proper initial segment of I_Y such that $I_Y \setminus R$ has no least element, is an $\aleph_1$-back and forth system; hence the orders I_X and I_Y are $L_{\infty,\aleph_1}$-equivalent. As the Dedekind completions of the I_X's are isomorphic we can identify them. After this identification, each of the orders I_X share the same $\aleph_1$-representation, namely (T, g), where $T = \{\langle 0 \rangle, \langle 1 \rangle\} \cup \{\langle 1, \delta \rangle : \delta \in \omega_1\}$ and $g(\langle 1, \delta \rangle)$ is the element of the Dedekind completion realizing the cut preceding J_δ for all $\delta > 0$.

It remains to show that I_X is not embeddable in I_Y whenever $X \neq Y$. (This was proved in [**?**] but is repeated here for convenience.) So fix $X \neq Y$ and assume by way of contradiction that there is an embedding $f : I_X \to I_Y$. It is readily verified that the set

$$C = \{\alpha \in \omega_1 : f(\sum_{i \in \alpha} J_i^X) = \sum_{i \in \alpha} J_i^Y\}$$

is a club subset of ω_1. Thus, since $X \setminus Y$ is stationary, there is an $\alpha \in C \cap X \setminus Y$. But $I_X \setminus \sum_{i \in \alpha} J_i^X$ has a least element, whereas $I_Y \setminus \sum_{i \in \alpha} J_i^Y$ does not, which is a contradiction. ∎

Despite the limitations demonstrated by the example above, the invariants described in this section do allow us to obtain an upper bound on the Karp complexity of $\mathbf{K}_{2\mathrm{tr}}$. The following definitions establish our notation.

Definition 3.7 For $D \subseteq \overline{I}$, a *D-cut* ν is a partition of D into two sets, D_ν^- and D_ν^+ (either may be empty) such that $D_\nu^- \cup D_\nu^+ = D$, $D_\nu^- \cap D_\nu^+ = \emptyset$, and D_ν^- is downward closed. We write $\nu = (D_\nu^-, D_\nu^+)$ and let $I(\nu) = \{x \in I : D_\nu^- < x < D_\nu^+\}$.

Definition 3.8 Suppose I and J are two linear orders. If $D \subseteq \overline{I}$ and $f : D \to \overline{J}$ is any order-preserving function then $f(\nu)$ is the $f(D)$-cut $(f(D_\nu^-), f(D_\nu^+))$. A function $f : D \to \overline{J}$ is *proper* if $\{-\infty, +\infty\} \subseteq D$ and f is order-preserving, continuous, $f(-\infty) = -\infty$, $f(+\infty) = +\infty$, and satisfies $d \in I \Leftrightarrow f(d) \in J$ for all $d \in D$.

If $D \subseteq \overline{I}$ and $f : D \to \overline{J}$ is a proper function, then $I \setminus D$ and $J \setminus f(D)$ are partitioned into corresponding families of D-cuts and $f(D)$-cuts. The following definitions measure the similarity of these cuts.

Definition 3.9 Two (possibly empty) linear orders I and J are (μ^+, α)-*equivalent,* written $I \equiv_{\mu^+,\alpha} J$, if I and J are elementarily equivalent and the empty function in $\mathcal{F}_{\mu^+}(I, J)$ has Rank at least α (see Definition 2.2).

By allowing linear orders to be empty and by insisting on elementary equivalence we intend that $I = \emptyset$ if and only if $J = \emptyset$ and $|I| = 1$ if and only if $|J| = 1$ whenever $I \equiv_{\mu+,\alpha} J$ for some ordinal α.

Definition 3.10 If $D \subseteq \overline{I}$ and $f : D \to \overline{J}$ is proper, then f is α-*strong* if $I(\nu) \equiv_{\mu^+,\alpha} J(f(\nu))$ for all D-cuts ν.

If $f \in \mathcal{F}_{\mu^+}(I, J)$ has domain A and has Rank at least 2, then it is easily seen that f is continuous and extends uniquely to a proper function

$$g : A \cup \lim(A) \cup \{-\infty, +\infty\} \to \overline{J},$$

where $\lim(A)$ denotes the set of limit points of A in $\overline{I}$. Also, it is easily established by induction on $\alpha \geq 1$ that if $g : D \to \overline{J}$ is a proper function with domain $D \subseteq \overline{I}$ and the restriction $f = g|(D \cap I)$ is in $\mathcal{F}_{\mu^+}(I, J)$, then g is α-strong if and only if $\mathrm{Rank}(f) \geq \alpha$.

For $\alpha \geq 1$ the class of α-strong proper functions has desirable closure properties. It is routine to show that the restriction of any α-strong proper function to any set that contains $\{-\infty, +\infty\}$ is also proper and α-strong. As well, we have the following lemma, which is proved by a straightforward induction on α.

Lemma 3.11. *Let $\alpha \geq 1$. Suppose that $D \subseteq \overline{I}$, $f : D \to \overline{J}$ is an α-strong proper function, and for each D-cut ν there is a set $E_\nu \subseteq \overline{I(\nu)}$ and an α-strong proper function $g_\nu : E_\nu \to \overline{J(f(\nu))}$. Then $f \cup \bigcup_\nu g_\nu$ is proper and α-strong.*

Lemma 3.12. *Let $I_0, I_1 \in \mathbf{K}_{2tr}$ satisfy $I_0 \sim I_1$ and $I_0 \equiv_{\mu^+,\alpha} I_1$ for some ordinal $\alpha \geq 2$. Assume that $A \subseteq I_0$ is of size at most μ and satisfies*

(1) *A is bounded below or $coi(A) = \aleph_0$; and*
(2) *A is bounded above or $cof(A) = \aleph_0$.*

Then there is an $f \in \mathcal{F}_{\mu^+}(I_0, I_1)$ with domain A of Rank at least α.

Proof. We show that in fact A is contained in an interval of I_0 which is isomorphic to an interval of I_1. This interval will be of the form (a, b), where a is a lower bound for A if one exists, or the symbol $-\infty$, and b is defined similarly. Take as a typical case that in which

$a \in I_0$ and $b = \infty$. Then we claim that the interval (a, ∞) is isomorphic to (a', ∞) for any $a' \in I_1$. The point is that (a, ∞) has cofinality $\aleph_0$, hence (a', ∞) does by (μ^+, α)-equivalence. So we can build the desired isomorphism in a countable sequence of steps, using double transitivity and the local isomorphism of I_0 and I_1.

As well, it follows from the relations $I_0 \sim I_1$ and $I_0 \equiv_{\mu^+,\alpha} I_1$ and another instance of double transitivity that the intervals $(-\infty, a)$ and $(-\infty, a')$ are (μ^+, α)-equivalent. Thus, the the restriction of the isomorphism to A has Rank at least α. ∎

The following Proposition is the key to the proof of Theorem 3.14. Before embarking on it, we introduce some more notation. For $C \subseteq \lambda$, let

$$C_0 = \{\alpha \in C : \alpha \text{ is a limit point of } C \cap \alpha\}$$

and for T a μ-decomposition tree, let T_0 be the club subtree of T satisfying $Succ_{T_0}(\eta) = \{\eta\frown\langle\alpha\rangle : \alpha \in C_0\}$, where $Succ_T(\eta) = \{\eta\frown\langle\alpha\rangle : \alpha \in C\}$ for all non-maximal nodes $\eta \in T_0^*$. Note that if (T, g) is a μ-representation of $\overline{I}$, then $(T_0, g|T_0^*)$ is also a μ-representation of $\overline{I}$ with the additional property that $g(\eta)$ either has cofinality or coinitiality at most μ for all $\eta \in T_0^* \setminus \{\langle 0\rangle, \langle 1\rangle\}$.

Proposition 3.13. *Assume $I_0, I_1 \in \mathbf{K}_{2tr}$, $I_0 \sim I_1$ and $I_0 \equiv_{\mu^+,\omega} I_1$. If $A \subseteq I_0$ and $|A| \le \mu$, then there is a function $f : A \to I_1$ of Rank at least ω.*

Proof. Pick $A \subseteq I_0$ of size at most μ. In order to produce a $h : A \to I_1$ of Rank at least ω, we first construct a desirable proper function $j : D \to \overline{I}_1$. Choose a μ-representation (T, g) of $\overline{I}_0$. By passing to the subtree T_0 in the notation preceding this proposition, we may assume that $g(\eta)$ either has cofinality or coinitiality at most μ for all $\eta \in T^* \setminus \{\langle 0\rangle, \langle 1\rangle\}$. Let $B = B_L \cup B_R$, where

$$\begin{aligned} B_L &= \{\eta \in T^* : dir(\eta) = \text{LEFT and } A \text{ is cofinal in } I_\eta\} \quad \text{and} \\ B_R &= \{\eta \in T^* : dir(\eta) = \text{RIGHT and } A \text{ is coinitial in } I_\eta\}. \end{aligned}$$

We claim that B has size at most μ. To see this, it suffices by symmetry to show that $|B_L| \le \mu$. Recall that for every successor ordinal α, the intervals $\{I_\eta : \eta \in T^* \cap {}^\alpha\mu\}$ are disjoint. Since $\eta \in B_L$ implies $A \cap I_\eta \ne \emptyset$, this implies $|B_L \cap {}^\alpha\mu| \le \mu$ for all successor ordinals α. Further, since $|A| \le \mu$, we can choose a successor ordinal γ so that for every pair $a, a' \in A$, there is $\eta \in T^*$ of length less than γ satisfying $a < g(\eta) < a'$ whenever there is any $\nu \in T^*$ with $a < g(\nu) < a'$. But now, by our choice of γ, if $\nu, \nu' \in B_L$ have length $> \gamma$ and have $\nu|\gamma = \nu'|\gamma$, then

$$g(\nu) = \sup(A \cap I_{\nu|\gamma}) = g(\nu'),$$

so $\nu = \nu'$ and $I_{\nu|\gamma} \cap A \neq \emptyset$. Thus,

$$|\{\nu \in B_L : lg(\nu) > \gamma\}| \leq |\{\eta \in T^* : lg(\eta) = \gamma \text{ and } I_\eta \cap A \neq \emptyset\}| \leq \mu$$

so $|B_L| \leq \mu$.

Let $B' \supseteq B$ be a closed subset of T of size at most μ. As $g(\eta)$ has cofinality or coinitiality at most μ in $\overline{I}_0$ for each $\eta \in B' \setminus \{\langle 0\rangle, \langle 1\rangle\}$, there is a set $X \subseteq I_0$ of size at most μ such that $g(B') \subseteq \lim(X) \cup \{-\infty, +\infty\}$. Since $I_0 \equiv_{\mu^+,\omega} I_1$, for each $n \geq 2$ we can choose an order-preserving $j_n : X \to I_1$ of Rank at least n. As $g(B') \subseteq \lim(X)$, each j_n extends uniquely to a proper function (also called j_n) from $X \cup g(B')$ to $\overline{I}_1$. As $B' \subseteq T$ is closed, by Lemma 3.5 there is a club subtree T'_n for each $n \geq 2$ such that $j_n(g(\eta)) = j_{n+1}(g(\eta))$ for all $\eta \in B' \cap T'_n$. Let $Y = \bigcap_{n \geq 2} T'_n$ and let $D = \{g(\eta) : \eta \in B' \cap Y_0\}$, where Y_0 is the club subtree of Y described in the notation preceding this proposition. As the functions j_n agree on D for all $n \geq 2$, we let $j : D \to \overline{I}_1$ denote this common (proper) function. As each j_n was n-strong, the function j is ω-strong.

By Lemma 3.11, in order to ascertain the existence of an ω-strong $h : A \to I_1$, it suffices to construct an order-preserving function $f : A \cap I_0(\nu) \to I_1(j(\nu))$ of Rank at least ω for every D-cut ν of I_0. So fix a D-cut $\nu = (D^-_\nu, D^+_\nu)$. We finish the proof by showing that the hypotheses of Lemma 3.12 are satisfied for $I_0(\nu)$ and $I_1(j(\nu))$. As $I_0(\nu)$ and $I_1(j(\nu))$ are convex subsets of I_0 and I_1 respectively, $I_0(\nu) \sim I_1(j(\nu))$. Since j is ω-strong, $I_0(\nu) \equiv_{\mu^+,\omega} I_1(j(\nu))$. Finally, assume by way of contradiction that $A \cap I_0(\nu)$ is unbounded above in $I_0(\nu)$ and has uncountable cofinality. (The case of $A \cap I_0(\nu)$ unbounded below in $I_0(\nu)$ of uncountable cardinality is symmetric.) Let $b = \sup(A \cap I_0(\nu))$ and let $\kappa = \mathrm{cof}(A \cap I_0(\nu))$. We will obtain a contradiction by showing that $b = \sup(D^-_\nu)$, which would make $I_0(\nu)$ empty. First, since Y_0 is a club subtree of T and $b = \inf(D^+_\nu)$, $b = g(\eta)$ for some $\eta \in Y_0$. As we assumed A cofinal below b, $b \in B$ as well. There are now four cases to consider, all of which imply $b = \sup(D^-_\nu)$ or contradict our hypotheses.

Case 1. $dir(\eta)$ =RIGHT and $lg(\eta) = \delta + 1$ where δ is a limit ordinal or 0.

Say $\eta = \rho\frown\langle 0\rangle$. Since $\mathrm{cof}(b) = \kappa > \aleph_0$ there is a strictly increasing sequence of limit ordinals $\langle \gamma_i : i < \kappa\rangle$ such that $b = \sup\{g(\rho|(\gamma_i + 1)) : i < \kappa\}$. Since B' is closed, $\rho|\gamma \in B'$ for all $\gamma < lg(\rho)$, so $g(\rho|\gamma) \in D_\nu$ and $b = \sup(D^-_\nu)$.

Case 2. $dir(\eta)$ =RIGHT and $lg(\eta) = \delta + n$ for some $n > 1$.

Say $\eta = \rho\frown\langle\alpha\rangle$ for some $\alpha \in C_0$, where C is such that $Succ_Y(\rho) = \{\rho\frown\langle\beta\rangle : \beta \in C\}$. As $\mathrm{cof}(b) = \kappa$ there is a continuous, strictly increasing sequence of ordinals $\langle\beta_i : i < \kappa\rangle$ from C with limit α. Again, as B' is closed, $\rho\frown\langle\beta_i\rangle \in B'$ for all $i \in \kappa$. It follows that $\rho\frown\langle\beta_i\rangle \in B' \cap Y_0$ for all limit ordinals $i \in \kappa$, so again $b = \sup(D^-_\nu)$.

Case 3. $dir(\eta)$ =LEFT and η is not maximal in Y_0.

Say $Succ_{Y_0}(\eta) = \{\eta\frown\langle\alpha\rangle : \alpha \in C_0\}$. As A is unbounded below b and $\kappa > \aleph_0$, there is a club $C' \subseteq C_0$ such that A is unbounded below $g(\eta\frown\langle\alpha\rangle)$ for all $\alpha \in C'$. Thus, $\eta\frown\langle\alpha\rangle \in B \cap Y_0^*$ for all $\alpha \in C'$, so again $b = \sup(D^-_\nu)$.

Case 4. $dir(\eta)$ =LEFT and η is maximal in Y_0.

As η maximal in Y_0 implies η maximal in T, it follows from the definition of a μ-representation that $\mathrm{cof}(g(\eta)) = \aleph_0$ or $\mathrm{cof}(g(\eta)) > \mu$. However, we assumed that $\mathrm{cof}(g(\eta)) > \aleph_0$ and A witnesses that $\mathrm{cof}(g(\eta)) \leq \mu$, so both are impossible. ∎

Our theorem now follows easily.

Theorem 3.14. *$KC_{\mu^+}(\mathbf{K}_{2tr}) \leq \omega$ for all uncountable cardinals μ.*

Proof. Fix $I \in \mathbf{K}_{2\mathrm{tr}}$ and an uncountable cardinal μ. Let $f \in \mathcal{F}_{\mu^+}(I)$ have Rank at least ω. We claim that $\mathrm{Rank}(f) \geq \omega + 1$. To see this, it suffices by symmetry to show that if $A \subseteq I$, $|A| \leq \mu$ then there is a function $g \in \mathcal{F}_{\mu^+}(I)$ extending f of Rank at least ω with $A \subseteq \mathrm{dom}(g)$. So fix such a set A and let $\tilde{f}$ denote the proper function extending f with domain $\mathrm{dom}(f) \cup \{-\infty, +\infty\}$. Since $\mathrm{Rank}(f) \geq \omega$, $\tilde{f}$ is ω-strong. Now fix a $\mathrm{dom}(\tilde{f})$-cut ν. Clearly, $I(\nu) \equiv_{\mu^+,\omega} I(f(\nu))$ and $I(\nu) \sim I(f(\nu))$, so it follows from Proposition 3.13 that there is a function $g_\nu : A \cap I(\nu) \to I(\tilde{f}(\nu))$ in $\mathcal{F}_{\mu^+}(I(\nu), I(\tilde{f}(\nu))$ of Rank at least ω. Thus, it follows from Lemma 3.11 that the proper function $g = \tilde{f} \cup \bigcup\{g_\nu : \nu \text{ a } \mathrm{dom}(\tilde{f})\text{-cut}\}$ is ω-strong, hence the restriction of g to $A \cup \mathrm{dom}(f)$ has Rank at least ω. ∎

4. The ω-independence property

This section is devoted to proving that any pseudo-elementary class with the ω-independence property (see Definition 4.4) is not controlled. We begin the section by proving Proposition 4.3, which will provide us with a method for concluding that $KC_\lambda(\mathbf{K}) = \infty$ by looking at the family of λ-partial isomorphisms from one element of $\mathbf{K}$ into another.

Definition 4.1 An *ω-tree* $\mathcal{T}$ is a downward closed subset of ${}^{<\omega}\lambda$ for some ordinal λ. We call $\mathcal{T}$ *well-founded* if it does not have an infinite branch. For a tree $\mathcal{T}$ and $\eta \in \mathcal{T}$, the *depth of $\mathcal{T}$ above η*, $dp_{\mathcal{T}}(\eta)$ is defined inductively by

$$dp_{\mathcal{T}}(\eta) = \begin{cases} \sup\{dp_{\mathcal{T}}(\nu) + 1\} : \eta \lhd \nu\} & \text{if } \eta \text{ has a successor} \\ 0 & \text{otherwise.} \end{cases}$$

and the depth of $\mathcal{T}$, $dp(\mathcal{T}) = dp_{\mathcal{T}}(\langle\rangle)$.

Clearly, $dp(\mathcal{T}) < \infty$ if and only if $\mathcal{T}$ is well-founded. The most insightful example is that for any ordinal α, the tree $des(\alpha)$ consisting of all descending sequences of ordinals $< \alpha$ ordered by initial segment has depth α. The proof of the following lemma is reminiscent of the proof of Morley's Omitting Types Theorem.

Lemma 4.2. *If $\mathcal{T} \subseteq {}^{<\omega}\lambda$ is well-founded and $dp(\mathcal{T}) \geq \kappa^+$, then for any coloring $c : \mathcal{T} \to \kappa$, there is a subtree $\mathcal{S} \subseteq \mathcal{T}$ of depth at least ω such that $c|_{\mathcal{S}\cap {}^n\lambda}$ is constant for each $n \in \omega$.*

Proof. Since $dp_{\mathcal{T}} \geq \kappa^+$, there is an $\eta \in \mathcal{T}$ with $dp_{\mathcal{T}}(\eta) = \kappa^+$. Thus, by concentrating on subtrees extending η, we may assume that $dp(\mathcal{T}) = \kappa^+$.

For each $n \in \omega$ we will name a subset $X_n \subseteq \kappa^+$ of size κ^+ and a function $f_n : X_n \to \mathcal{T} \cap {}^n\lambda$ such that $X_{n+1} \subseteq X_n$, every element of $f_{n+1}(X_{n+1})$ is a successor of an element of $f_n(X_n)$, $dp_{\mathcal{T}}(f_n(\alpha)) \geq \alpha$ and $c|_{f_n(X_n)}$ is constant.

To begin, let $X_0 = \kappa^+$ and let $f_0 : X_0 \to \{\langle\rangle\}$. Given X_n and f_n satisfying our demands, we define X_{n+1} and $f_{n+1} : X_{n+1} \to \mathcal{T} \cap {}^{n+1}\lambda$ as follows. For $\alpha \in X_n$, let β be the least element of X_n greater than α. As $dp_{\mathcal{T}}(f_n(\beta)) \geq \beta$, we can define $f_{n+1}(\alpha)$ to be a successor of $f_n(\beta)$ of depth at least α. Since X_n has size κ^+, let X_{n+1} be a subset of X_n of size κ^+ such that $c|_{f_{n+1}(X_{n+1})}$ is monochromatic.

Now let $R = \{f_n(\beta_n) : n \in \omega\}$, where β_n is the least element of X_n and let $\mathcal{S}$ be the subtree of $\mathcal{T}$ generated by R. ∎

Suppose that $N \equiv M$ and $\emptyset = A_0 \subseteq A_1 \subseteq \ldots \subseteq N$ is an ω-sequence of subsets of N of size less than λ. Let

$$\mathcal{T}_n = \{\text{Range}(f) : f \in \mathcal{F}_\lambda(N, M),\ f \text{ has domain } A_n\}$$

and let $\mathcal{T} = \bigcup\{\mathcal{T}_n : n \in \omega\}$ be a tree under inclusion. Typically $\mathcal{T}$ will be an ω-tree and we can ask whether or not it is well-founded. The relationship between this question and Karp complexity is partially explained by the following proposition.

Proposition 4.3. *If $KC_\lambda(\mathbf{K}) < \infty$ then there is an ordinal α^* such that whenever $N \equiv M \in \mathbf{K}$ and $\emptyset = A_0 \subseteq A_1 \subseteq \ldots \subseteq N$ are chosen with $|A_i| < \lambda$, then the induced tree $\mathcal{T}$ either has depth at most α^* or has an infinite branch.*

Proof. If $KC_\lambda(\mathbf{K}) < \infty$ then by Proposition 2.4, there is a cardinal κ bounding the number of $L_{\infty,\lambda}$-types realized in elements of $\mathbf{K}$. We claim that $\alpha^* = \kappa^+$ has the desired property. To see this, choose $N \equiv M$ from $\mathbf{K}$ and $\emptyset = A_0 \subseteq A_1 \subseteq \ldots \subseteq N$ and assume that $dp(\mathcal{T}) \geq \kappa^+$. By Lemma 4.2, there is a subtree $\mathcal{S}$ of $\mathcal{T}$ of depth ω such that the $L_{\infty,\lambda}$-types of the elements of $\mathcal{S}$ depend only on their level in $\mathcal{S}$. In particular, for each n there is an element $B_n \in \mathcal{S}$ at level n that has a successor in $\mathcal{S}$. Consequently, for each $n \in \omega$ the $L_{\infty,\lambda}$-formula

$$\Theta(X_n) = \exists Y_n tp_{\infty,\lambda}(X_n, Y_n) = tp_{\infty,\lambda}(B_{n+1})$$

is implied by $tp_{\infty,\lambda}(B_n)$. Applying this iteratively produces an elementary partial function $f : N \to M$ with domain $\bigcup\{A_n : n \in \omega\}$, so $\mathcal{T}$ has an infinite branch. ∎

Definition 4.4 A class $\mathbf{K}$ of L-structures has the *ω-independence property* if there is a set $\{\varphi_n(\overline{x}_0, \ldots, \overline{x}_{n-1}, \overline{y}_n) : n \in \omega\}$ of L-formulas such that for all $M \in \mathbf{K}$ there is a sequence $\langle \overline{a}_i : i < \omega \rangle$ from M such that for all $n \in \omega$ and all functions $f : n \to \{0, 1\}$ there is a sequence $\langle \overline{b}_i : i < n \rangle$ from M such that for all $i < n$,

$$M \models \varphi_i(\overline{b}_0, \ldots, \overline{b}_{i-1}, \overline{a}_i) \qquad \text{if and only if} \qquad f(i) = 1.$$

As an example, the model completion of the empty theory in the language $L = \{R_n : n \in \omega\}$ consisting of one n-ary relation for every n is a complete, simple theory with the ω-independence property. (In this example, the $\overline{y}_n$'s do not appear.) Clearly, if $\mathbf{K}$ has the ω-independence property, then $\mathbf{K}$ has the independence property. However, the theory of the random graph has the independence property, but fails to have the ω-independence property. We remark that despite this failure, the theory of the random graph is not controlled. We do not attempt to prove this assertion here.

Our interest in the notion of ω-independence is largely captured by the proposition given below.

Definition 4.5 An *ordered multigraph* is a structure $(G, <, R_n)_{n\in\omega}$ where $<$ is interpreted as a linear order and each R_n is a symmetric n-ary relation on G.

Proposition 4.6. *If $L_1 \supseteq L_0$, T_1 is an L_1-theory with Skolem functions and $\mathbf{K}$, the class of reducts of models of T_1 to L_0 has the ω-independence* *property witnessed by* $\{\varphi_n : n \in \omega\}$ *then for every ordered multigraph* $(G, <, R_n)_{n\in\omega}$ *there is a structure* $M_G \in \mathbf{K}$ *and sequences* $\langle \overline{a}_n : n \in \omega\rangle$ *and* $\langle \overline{b}_g : g \in G\rangle$ *from* M_G *such that*

(1) M_G *is the* L_1*-Skolem hull of* $\{\overline{a}_n : n \in \omega\} \cup \{\overline{b}_g : g \in G\}$*;*
(2) *If* $g_1, \ldots, g_n$ *and* $h_1, \ldots, h_n$ *have the same quantifier-free type in* $(G, <, R_n)_{n\in\omega}$ *then the sequences* $\overline{b}_{g_1}, \ldots, \overline{b}_{g_n}$ *and* $\overline{b}_{h_1}, \ldots, \overline{b}_{h_n}$ *have the same type over* $\{\overline{a}_n : n \in \omega\}$ *in* M_G*;*
(3) $M_G \models \varphi_n(\overline{b}_{g_1}, \ldots, \overline{b}_{g_n}, \overline{a}_n)$ *if and only if* $G \models R_n(g_1, \ldots, g_n)$ *for all* n *and all* $g_1, \ldots, g_n$ *from* G*.*

The proof of Proposition 4.6 is word for word like the proof of the existence of Ehrenfeucht-Mostowski models for unstable pseudo-elementary classes (see e.g., Section 11.3 of [**?**]) but with the Nešetřil-Rödl theorem (see [**?**] or [**?**]) in place of Ramsey's theorem.

The following lemma tells us that we need not explicitly consider the constants $\{\overline{a}_n : n \in \omega\}$ in the proof of Theorem 4.9.

Lemma 4.7. *Let* $\mathbf{K}$ *be a class of* L*-structures and let* C *be a set of fewer than* λ *new constant symbols. Let* $\mathbf{K}^*$ *be the class of all expansions of elements of* $\mathbf{K}$ *to* $L \cup C$*-structures. Then* $KC_\lambda(\mathbf{K}^*) \leq KC_\lambda(\mathbf{K})$*.*

Proof. For any $M^* \in \mathbf{K}^*$, let M be its reduct to the language of L. For every partial function $f \in \mathcal{F}_\lambda(M^*)$, let $\tilde{f} \in \mathcal{F}_\lambda(M)$ be the extension of f that is the identity on every element of C^M. It is easy to show by induction that $\mathrm{Rank}_{\mathcal{F}^*}(f) = \mathrm{Rank}_{\mathcal{F}}(\tilde{f})$. Hence, $KC_\lambda(M^*) \leq KC_\lambda(M)$, so $KC_\lambda(\mathbf{K}^*) \leq KC_\lambda(\mathbf{K})$. ∎

The other theorem we will need is that there exist very complicated colorings of a number of cardinals. As notation, for x a finite subset of μ, let x^m denote the m^{th} element of x in increasing order. Following the notation in [**?**], let $Pr_0(\mu, \mu, \aleph_0, \aleph_0)$ denote the following statement:

- There is a symmetric two-place function $c : \mu \times \mu \to \omega$ such that for every $n \in \omega$, every collection of μ disjoint, n-element subsets $\{x_\alpha : \alpha \in \mu\}$ of μ, and every function $f : n \times n \to \omega$, there are $\alpha < \beta < \mu$ such that
$$c(x_\alpha^m, x_\beta^{m'}) = f(m, m')$$
for all $m, m' < n$.

It is shown in [**?**] that $Pr_0(\lambda, \lambda, \aleph_0, \aleph_0)$ holds for an uncountable cardinal λ whenever there exists a nonreflecting stationary subset of λ. (A stationary subset $S \subseteq \lambda$ is nonreflecting if $S \cap \alpha$ is not stationary in α for all limit ordinals $\alpha < \lambda$.) In particular, $Pr_0(\aleph_3, \aleph_3, \aleph_0, \aleph_0)$ holds. More recently, in [**?**] the second author has shown that $Pr_0(\aleph_2, \aleph_2, \aleph_0, \aleph_0)$

holds as well. This suffices for our purpose. See [?] for more of the history of Pr_0 and its cousins.

The following Lemma recasts $Pr_0(\mu,\mu,\aleph_0,\aleph_0)$ into the form we will use in the proof of Theorem 4.9.

Lemma 4.8. *Let $c:[\mu]^2\to\omega$ witness $Pr_0(\mu,\mu,\aleph_0,\aleph_0)$. For every $k,n\in\omega$, every collection $\{x_\alpha:\alpha\in\mu\}$ of μ disjoint, n-element subsets of μ, and every family of colorings $\{f_{i,j}:n^2\to\omega:i<j<k\}$, there are $\beta_0<\beta_1<\cdots<\beta_{k-1}$ such that*

$$c(x^m_{\beta_i},x^{m'}_{\beta_j})=f_{i,j}(m,m')$$

for all $i<j<k$ and all $m,m'<n$.

Proof. Fix k,n, $\{x_\alpha:\alpha\in\mu\}$, and $\{f_{i,j}:i<j<k\}$ satisfying the hypotheses. Without loss, we may assume that $x^{n-1}_\alpha<x^0_{\alpha+1}$ for all α. For α limit, let $y_\alpha=\bigcup\{x_{\alpha+i}:i<k\}$ and let $W_0=\{\alpha\in\mu:\alpha \text{ limit}\}$. By induction on $k'\le k$ we will build a sequence $\beta_0<\beta_1<\cdots<\beta_{k'-1}$ and a subset $W_{k'}$ of size μ such that $W_{k'+1}\subseteq W_{k'}$ and $c(y^{ni+m}_{\beta_i},y^{nj+m'}_\gamma)=f_{i,j}(m,m')$ for all $m,m'<n$, all $i<j<k$ with $i<k'$, and all $\gamma\in W_{k'}$. For $k'=0$ there is nothing to do. Assuming $\beta_0<\cdots<\beta_{k'-1}$ and $W_{k'}$ have been chosen, it follows from $Pr_0(\mu,\mu,\aleph_0,\aleph_0)$ that there is $\beta_{k'}$ such that the set

$$\{\gamma\in W_{k'}:\gamma>\beta_{k'}\text{ and }c(y^{nk'+m}_{\beta_{k'}},y^{jn+m'}_\gamma)=f_{k',j}(m,m')\text{ for }j>k'\}$$

has size μ, hence is a suitable choice for $W_{k'+1}$. (If there were no such $\beta_{k'}$ then one could successively build a subset Z of $W_{k'}$ of size μ on which there would be no $\alpha<\beta$ from Z satisfying the coloring.) ∎

Theorem 4.9. *Let $L_1\supseteq L_0$ be first order languages, let T_1 be an L_1-theory and let $\mathbf{K}$ denote the class of reducts of models of T_1 to L_0. If $\mathbf{K}$ has the ω-independence property then $\mathbf{K}$ is not controlled. More precisely, if a cardinal $\mu>|T_1|$ is regular and there is a coloring of $[\mu]^2$ satisfying $Pr_0(\mu,\mu,\aleph_0,\aleph_0)$, then $KC_\lambda(\mathbf{K})=\infty$ for all cardinals $\lambda>\mu$.*

Proof. First, by adding countably many constants to the language L_0 and invoking Lemma 4.7, we may assume that the ω-independence of $\mathbf{K}$ is witnessed by formulas $\varphi_n(\overline{x}_0,\ldots,\overline{x}_{n-1})$ with no additional constants. Second, by considering M^{eq} in place of M for each $M\in\mathbf{K}$, we may assume that each $\overline{x}$ is a singleton. Third, by expanding T_1 if necessary, we may assume that it has built-in Skolem functions. Fix a coloring $c:[\mu]^2\to\omega$ that witnesses $Pr_0(\mu,\mu,\aleph_0,\aleph_0)$ and fix an ordinal α^*. We will use the coloring to define two rather complicated ordered multigraphs I and J and then use Proposition 4.6 to get Ehrenfeucht-Mostowski models $M, N \in \mathbf{K}$ that are built from I and J respectively. We will find a tree of λ-partial isomorphisms from N into M that is well-founded, yet has depth at least α^*. Since α^* was arbitrary, it follows immediately from Proposition 4.3 that $KC_\lambda(\mathbf{K}) = \infty$. So, let

$$des(\alpha^*) = \{\text{strictly decreasing sequences of ordinals } < \alpha^*\}$$

and let $(I, <)$ be the linear order with universe $\mu \times des(\alpha^*)$, ordered lexicographically. Let $(J, <)$ be the linear order with universe $\mu \times \{\rho_n : n \in \omega\}$, where $\rho_n = \langle 0, -1, -2, \ldots, -n+1\rangle$, also ordered lexicographically.

As notation, for finite sequences η, ν we write $\eta \lhd \nu$ when η is a proper initial segment of ν. For $t \in I \cup J$, let $t = (\alpha^t, \eta^t)$, where $\alpha^t \in \mu$ and η^t is a finite, decreasing sequence. For $s, t \in I \cup J$, we write $s \lhd^* t$ when $\eta^s \lhd \eta^t$. Fix, for the whole of this section, a partition of $\omega \setminus \{0\}$ into disjoint, infinite sets $\{Z_n : n \in \omega\}$.

We expand $(I, <)$ into an ordered multigraph $(I, <, R_n)_{n\in\omega}$ as follows: We posit that R_0 holds, $R_1(t)$ holds for all $t \in I$, and for $n > 1$, $R_n(t_0, \ldots, t_{n-1})$ holds if and only if for some permutation $\sigma \in Sym(n)$,

- $\eta^{t_{\sigma(0)}} \lhd \cdots \lhd \eta^{t_{\sigma(n-1)}}$;
- $lg(\eta^{t_{\sigma(i)}}) = i$ for all i;
- $\alpha^{t_i} \neq \alpha^{t_j}$ and $c(\alpha^{t_i}, \alpha^{t_j}) \in Z_n$ for all $i < j < n$; and
- $c(\alpha^{t_i}, \alpha^{t_j}) = c(\alpha^{t_k}, \alpha^{t_l})$ for all $i, j, k, l < n$ with $i \neq j$ and $k \neq l$.

Similarly, expand $(J, <)$ to an ordered multigraph $(J, <, R_n)_{n\in\omega}$ by positing that R_0 holds, $R_1(t)$ holds for all $t \in J$, and for all $n > 1$ $R_n(t_0, \ldots, t_{n-1})$ holds if and only if for some $\sigma \in Sym(n)$,

- $\eta^{t_{\sigma(i)}} = \rho_i$ for all $i < n$;
- $\alpha^{t_i} \neq \alpha^{t_j}$ and $c(\alpha^{t_i}, \alpha^{t_j}) \in Z_n$ for all $i < j < n$; and
- $c(\alpha^{t_i}, \alpha^{t_j}) = c(\alpha^{t_k}, \alpha^{t_l})$ for all $i, j, k, l < n$ with $i \neq j$ and $k \neq l$.

Now build Ehrenfeucht-Mostowski models $M, N \in \mathbf{K}$ from I and J respectively that satisfy Conditions 1-3 of Proposition 4.6. To avoid wanton use of nested subscripts, we identify the elements $b_g \in M$ and $g \in I$ (and similarly for N and J).

For each $n \in \omega$ let $A_n = \{t \in N : lg(\eta^t) < n\}$ and let $\mathcal{T}_n = \{\text{Range}(f) : f \in \mathcal{F}_\lambda(N, M) \text{ has domain } A_n\}$. We will show that $\mathcal{T} = \bigcup\{\mathcal{T}_n : n \in \omega\}$ is both well-founded and has depth α^*. As noted above, this is sufficient to conclude that $KC_\lambda(\mathbf{K}) = \infty$. If we assume that $\mathcal{T}$ is well-founded then the family of maps

$$f_\eta : A_{lg(\eta)} \to M$$

for $\eta \in des(\alpha^*)$ defined by $f_\eta((\alpha, \eta_i)) = (\alpha, \eta|i)$ witness that the depth of $\mathcal{T}$ is at least α^*.

So it remains to show that $\mathcal{T}$ is well-founded. The obvious distinction between the ordered multigraphs I and J is that J has an infinite, strictly increasing sequence $\langle \eta_n : n \in \omega \rangle$, whereas I does not. Suppose that an elementary map $g : \bigcup\{A_n : n \in \omega\} \to M$ is given. We will obtain a contradiction by constructing an infinite strictly increasing sequence in $des(\alpha^*)$. The construction of this sequence proceeds in three stages. First, since $\mu > |T_1|$ is regular, for every $l \in \omega$ there is an integer $n(l)$, an L_1-term $\tau_l(x_1, \dots, x_{n(l)})$, a subset X_l of μ of size μ, and functions $t_{l,m} : X_l \to I$ such that for each $\beta \in X_l$

$$g((\beta, \eta_l)) = \tau_l(d_l(\beta)),$$

where $d_l(\beta) = \langle t_{l,1}(\beta), \dots, t_{l,n(l)}(\beta) \rangle$. As notation, let $W = \{(l, m) : l \in \omega \text{ and } m \in [1, \dots, n(l)]\}$ and for each $(l, m) \in W$, let $\alpha_{l,m}$ and $\eta_{l,m}$ be the functions with domain X_l satisfying

$$t_{l,m}(\beta) = (\alpha_{l,m}(\beta), \eta_{l,m}(\beta)).$$

Next, we state two claims, whose proofs we defer until the end of the argument.

Claim 1. There is a sequence $\langle Y_l : l \in \omega \rangle$ such that each $Y_l \subseteq X_l$ has size μ and for each $k \in \omega$

$$\text{qftp}(d_0(\beta_0), \dots, d_{k-1}(\beta_{k-1})) = \text{qftp}(d_0(\beta'_0), \dots, d_{k-1}(\beta'_{k-1}))$$

in the structure $(I, <, \lhd^*)$ for all sequences $\beta_0 < \dots < b_{k-1}$, $\beta'_0 < \dots < b'_{k-1}$ with $\beta_l, \beta'_l \in Y_l$ for each $l < k$.

Claim 2. For every $k > 1$ there is a sequence $\langle m_l : l < k \rangle$ and a permutation σ of k such that

$$t_{\sigma(0), m_{\sigma(0)}}(\beta_{\sigma(0)}) \lhd^* t_{\sigma(1), m_{\sigma(1)}}(\beta_{\sigma(1)}) \lhd^* \dots \lhd^* t_{\sigma(k-1), m_{\sigma(k-1)}}(\beta_{\sigma(k-1)})$$

for every sequence $\beta_0 < \dots < \beta_{k-1}$ with $\beta_l \in Y_l$ for each $l < k$.

Given these two claims, it follows from König's Lemma (and the fact that the permutation σ is uniquely determined by the lengths of the η^t's) that there is an infinite sequence $\langle m_l : l \in \omega \rangle$ and a permutation $\sigma \in Sym(\omega)$ such that, letting $\eta_l = \eta^{t_{\sigma(l), m_{\sigma(l)}}}$ for each $l \in \omega$,

$$\eta_0(\beta_{\sigma(0)}) \lhd \eta_1(\beta_{\sigma(1)}) \lhd \dots$$

for all sequences $\beta_0 < \beta_1 < \dots$ satisfying $\beta_l \in Y_l$ for each $l \in \omega$. But the existence of such a sequence is clearly impossible as each $\eta_l(\beta) \in des(\alpha^*)$. ∎

Thus, to complete the proof of the theorem it suffices to prove the claims. The proof of Claim 1 is tedious, but straightforward. First, by trimming each of the sets X_l we may assume that for each $(l, m) \in W$,

(1) $\alpha_{l,m}$ is constant on X_l;
(2) $\alpha_{l,m}(\beta) = \beta$ for all $\beta \in X_l$; or
(3) $\{\alpha_{l,m}(\beta) : \beta \in X_l\}$ is strictly increasing and disjoint from X_l.

We call (l, m) *α-constant* if (1) holds and call (l, m) *α-trivial* if (2) holds. Similarly, we may assume that for each $(l, m) \in W$,

- $lg(\eta_{l,m}(\beta))$ is constant for all $\beta \in X_l$ and
- $\eta_{l,m}$ is constant on X_l or else $\{\eta_{l,m}(\beta) : \beta \in X_l\}$ is strictly increasing (in lexicographic order).

Additionally, we may assume that for each pair $(l, m), (l, m') \in W$ with the same l, the truth values of

- "$\alpha_{l,m}(\beta) < \alpha_{l,m'}(\beta)$";
- "$\eta_{l,m}(\beta) \lhd \eta_{l,m'}(\beta)$";
- "$\eta_{l,m}(\beta) <_{\text{lex}} \eta_{l,m'}(\beta)$"; and hence of
- "$t_{l,m}(\beta) < t_{l,m'}(\beta)$"

are constant for all $\beta \in X_l$. By trimming each X_l further, we may additionally assume that for all pairs $m, m' \in [1, \ldots n(l)]$, the truth values of

- "$\alpha_{l,m}(\beta_1) < \alpha_{l,m'}(\beta_2)$";
- "$\eta_{l,m}(\beta_1) \lhd \eta_{l,m'}(\beta_2)$";
- "$\eta_{l,m}(\beta_1) <_{\text{lex}} \eta_{l,m'}(\beta_2)$"; and hence of
- "$t_{l,m}(\beta_1) < t_{l,m'}(\beta_2)$"

are constant for all pairs $\beta_1 < \beta_2$ from X_l.

So far, each of our trimmings has concentrated on a single set X_l. However, to complete the proof of the claim, we must consider pairs of sets as well. Fortunately, this presents no problem. We illustrate one such reduction and leave the other (virtually identical) reductions to the reader. We claim that there are subsets $Y_l \subseteq X_l$, each of size μ, such that for all $(l_1, m_1), (l_2, m_2) \in W$ the truth value of

$$\text{“}\alpha_{l_1,m_1}(\beta_1) < \alpha_{l_2,m_2}(\beta_2)\text{ ”} \qquad (*)$$

is constant for all pairs (β_1, β_2) satisfying $\beta_1 \in Y_{l_1}$, $\beta_2 \in Y_{l_2}$, and $\beta_1 < \beta_2$. To see this, let C be the α-constant pairs $(l, m) \in W$ and let $\delta < \mu$ be the supremum of all $\alpha_{l,m}(\beta)$ for $(l, m) \in C$, $\beta \in X_l$. By removing fewer than μ elements from each X_l, we may assume that $\alpha_{l,m}(\beta) > \delta$ for all non-α-constant $(l, m) \in W$ and all $\beta \in X_l$. It is now routine to inductively construct the sets $\{Y_l : l \in \omega\}$ in μ steps so as to ensure

$$\alpha_{l_1,m_1}(\beta_1) < \alpha_{l_2,m_2}(\beta_2)$$

whenever $(l_1, m_1), (l_2, m_2)$ are not α-constant, $l_1 < l_2$, $\beta_1 \in Y_1$, $\beta_2, \in Y_2$, and $\beta_1 < \beta_2$. Combining this with the earlier trimmings of the X_l's establish $(*)$.

Finally, we prove Claim 2. This is the heart of the argument and is where properties of the coloring c are used. Fix an integer $k > 1$. In light of Claim 1, it suffices to find a sequence $\langle m_l : l < k\rangle$ and a permutation σ of k such that $t_{\sigma(0),m_{\sigma(0)}}(\beta_{\sigma(0)}) \lhd^* \cdots \lhd^* t_{\sigma(k-1),m_{\sigma(k-1)}}(\beta_{\sigma(k-1)})$ for *some* sequence $\beta_0 < \cdots < \beta_{k-1}$ with $\beta_l \in Y_l$ for each $l < k$. Consequently, we can trim the sets Y_l still further. As notation, let W_k denote the *finite* set of all pairs $(l, m) \in W$ with $l < k$. For each $l < k$, let h_l enumerate Y_l, i.e., $h_l(\delta) =$ the δ^{th} element of Y_l.

By trimming each Y_l for $l < k$, we may additionally assume that:

- The sets Y_l are disjoint and $\alpha_{l,m}(\beta) \notin \bigcup_{l<k} Y_l$ unless (l, m) is α-trivial;
- $\delta_1 < \delta_2$ implies $h_l(\delta_1) < h_{l'}(\delta_2)$ for all $l, l' < k$;
- For all pairs $(l, m), (l', m') \in W_k$ with (l, m) α-constant, there is an integer $c^*(l, m, l', m') \in \omega$ such that
$$c(\alpha_{l,m}(\beta), \alpha_{l',m'}(\beta')) = c^*(l, m, l', m')$$
for all distinct β, β' from Y_l, Y_l' respectively.

Let C^* denote the (finite) set of all integers $c^*(l, m, l', m')$, where the pairs $(l, m), (l', m')$ are from W_k and (l, m) is α-constant. Choose integers $p \in Z_k \setminus C^*$ and $q \in Z_r$ for some $r > |W_k|$. As notation, for each ordinal $\delta \in \mu$, let

$$B_l(\delta) = \{\alpha_{l,m}(h_l(\delta)) : (l, m) \in W,\ (l, m) \text{ not } \alpha\text{-constant}\} \cup \{h_l(\delta)\}$$

For $\overline{\delta} = \delta_0 < \delta_1 < \cdots < \delta_{k-1}$, let $B(\overline{\delta}) = \bigcup_{l<k} B_l(\delta_l)$. By trimming the sets Y_l, $l < k$ still further, we may assume that the order type of $B(\overline{\delta})$ is constant among all increasing k-tuples $\overline{\delta}$. Thus, by employing Lemma 4.8, we can choose two increasing k-tuples $\overline{\delta}^0$ and $\overline{\delta}^1$ satisfying:

- $c(\alpha, \beta) = q$ for all $\alpha, \beta \in B(\overline{\delta}^0)$; and
- $c(\alpha, \beta) = q$ for all $\alpha, \beta \in B(\overline{\delta}^1)$ EXCEPT that $c(h_i(\delta^1_i), h_j(\delta^1_j)) = p$ for all $i \neq j$.

As notation, let $\nu_l = h_l(\delta^0_l)$, $\bar{\nu} = \nu_0 < \cdots < \nu_{k-1}$, and $D(\bar{\nu}) = \{t_{l,m}(\nu_l) : (l, m) \in W_k\}$. Dually, let $\beta_l = h_l(\delta^1_l)$, $\bar{\beta} = \beta_0 < \cdots < \beta_{k-1}$, and $D(\bar{\beta}) = \{t_{l,m}(\beta_l) : (l, m) \in W_k\}$.

Now, working in the multigraph J,

$$J \models \neg R_k((\nu_0, \rho_0), \ldots, (\nu_{k-1}, \rho_{k-1})) \wedge R_k((\beta_0, \rho_0), \ldots, (\beta_{k-1}, \rho_{k-1})),$$

so

$$N \models \neg\varphi_k((\nu_0, \rho_0), \dots, (\nu_{k-1}, \rho_{k-1})) \wedge \varphi_k((\beta_0, \rho_0), \dots, (\beta_{k-1}, \rho_{k-1})).$$

Hence, by the elementarity of the map g,

$$M \models \neg\varphi_k(\tau_0(d_0(\nu_0)), \dots, \tau_{k-1}(d_{k-1}(\nu_{k-1}))) \text{ and}$$
$$M \models \varphi_k(\tau_0(d_0(\beta_0)), \dots, \tau_{k-1}(d_{k-1}(\beta_{k-1}))).$$

It follows from Proposition 4.6 that the discrepancy in φ_k implies that $\mathrm{qftp}(D(\bar{\nu})) \neq \mathrm{qftp}(D(\bar{\beta}))$ in the ordered multigraph $(I, <, R_n)_{n\in\omega}$. However, since $\mathrm{qftp}(D(\bar{\nu})) = \mathrm{qftp}(D(\bar{\beta}))$ in the structure $(I, <, \lhd^*)$, the sets $D(\bar{\nu})$ and $D(\bar{\beta})$ must differ on some R_n. This difference can only be explained by a discrepancy of the function c on some pairs of elements from the sets $B(\bar{\nu})$ and $B(\bar{\beta})$. Since c can only attain the values of p and q on pairs from $B(\bar{\nu})$ and $B(\bar{\beta})$, our choice of p and q implies that R_k is the only relation that can differ between $D(\bar{\nu})$ and $D(\bar{\beta})$. Thus, there are sequences $s_0, \dots, s_{k-1} \in D(\bar{\beta})$ and $s'_0, \dots, s'_{k-1} \in D(\bar{\nu})$ of corresponding elements such that

$$I \models R_k(s_0, \dots, s_{k-1}) \wedge \neg R_k(s'_0, \dots, s'_{k-1}).$$

In particular, $s_0 \lhd^* \cdots \lhd^* s_{k-1}$ and for all $i < j < k$ we have $\alpha^{s_i} \neq \alpha^{s_j}$ and $c(\alpha^{s_i}, \alpha^{s_j}) = p$. As each $s_i \in D(\bar{\beta})$, there are functions l, m with domain k such that

$$s_i = t_{l(i),m(i)}(\beta_{l(i)}).$$

Now fix $i < k$. Since $k > 1$ and $c(\alpha^{s_i}, \alpha^{s_j}) = p \notin C^*$ for any $j \neq i$, the pair $(l(i), m(i))$ is not α-constant. As well, the choice of the coloring of $B(\bar{\beta})$ ensures that $\alpha^{s_i} = h_r(\delta^1_r) \in Y_r$ for some $r < k$. Thus, the disjointness of the Y_l's imply that $r = l(i)$ and the pair $(l(i), m(i))$ is α-trivial. That is, $\alpha^{s_i} = \beta_{l(i)}$. Further, since $\alpha^{s_i} \neq \alpha^{s_j}$ whenever $i \neq j$, the function l must be a permutation of the set k. So, letting $m_i = m(l^{-1}(i))$ and $\sigma = l$, the sequence $\langle m_i : i < k\rangle$ and permutation σ are as desired. ∎

Department of Mathematics, University of Maryland, College Park, MD 20742

E-mail address: `mcl@math.umd.edu`

Department of Mathematics, Hebrew University of Jerusalem and Department of Mathematics, Rutgers University