THE CHARACTERISTIC SEQUENCE OF A FIRST-ORDER FORMULA

M. E. MALLIARIS

Abstract. For a first-order formula $\varphi(x;y)$ we introduce and study the characteristic sequence $\langle P_n : n < \omega \rangle$ of hypergraphs defined by $P_n(y_1, \dots, y_n) := (\exists x) \bigwedge_{i \leq n} \varphi(x; y_i)$. We show that combinatorial and classification theoretic properties of the characteristic sequence reflect classification theoretic properties of φ and vice versa. The main results are a characterization of NIP and of simplicity in terms of persistence of configurations in the characteristic sequence. Specifically, we show that some tree properties are detected by the presence of certain combinatorial configurations in the characteristic sequence while other properties such as instability and the independence property manifest themselves in the persistence of complicated configurations under localization.

1. Introduction

This article introduces and develops the theory of characteristic sequences. Fix T a background theory and φ a formula. For each $n < \omega$, set $P_n(y_1, \dots y_n) := (\exists x) \bigwedge_{i \leq n} \varphi(x; y_i)$. Then the characteristic sequence of φ is the countable sequence of hypergraphs $\langle P_n : n < \omega \rangle$, each defined on the parameter space of φ. The aim of these investigations is to show that classification-theoretic complexity of φ can naturally be described in terms of graph-theoretic complexity of the hypergraphs P_n. Specifically, we show that the complexity of consistent partial φ-types is reflected in the complexity of graphs which persist around a graph corresponding to the φ-type under analysis (in a sense made precise below). One result of this paper is a characterization: φ has NIP iff, for each n and each complete graph A, the only finite graphs persistent around A are complete, Theorem 5.17.

These simple combinatorial questions relating classification theory and graph theory have deep roots and a surprisingly far reach. At first glance, the characteristic sequence gives a transparent language for arguments which often occur in analyzing the fine structure of dividing or thorn-dividing. However, the power of this basic object lies in the fact that each P_n can be simultaneously seen as:

(1) a graph, admitting graph-theoretic analysis;
(2) a definable set in models of T;
(3) an incidence relation on the parameter space of φ.

So one has a great deal of leverage. As a key example, Keisler asked about the structure of a preorder on countable theories which compares the difficulty of saturating their regular ultrapowers [4]. "Keisler's order" on NIP theories is understood by work of Shelah ([8], Chapter VI), but its structure on theories with the independence property has remained open for forty years. In [6], [7] we showed that the structure of Keisler's order on unstable theories, and thus on theories with the independence property, depends on a classification of φ-types and specifically on an analysis of characteristic sequences.

A second motivation is that these methods establish that certain model-theoretic questions about the complexity of types can be naturally studied in a graph-theoretic framework via the characteristic sequence. Thus deep graph-theoretic structure theorems, such as Szemerédi regularity, can be applied to these graphs to give model-theoretic information [7].

Acknowledgments. This is work from my doctoral thesis at Berkeley under the direction of Thomas Scanlon. I would also like to thank Leo Harrington for many interesting discussions and John Baldwin for helpful remarks on an early version. Part of this work was done as a Mathlogaps fellow in Lyon and I thank the logic group for their warm hospitality. Thanks to the referee for careful comments which improved the presentation of this article.

2. The characteristic sequence

In this section we fix notation, then define the characteristic sequence (Definition 2.2) and establish its basic properties.

Definition 2.1. (Notation and conventions)

1. *Throughout this article, if a variable or a tuple is written x or a rather than $\overline{x}, \overline{a}$, this does not necessarily imply that $\ell(x), \ell(a) = 1$.*
2. *Unless otherwise stated, T is a complete theory in the language $\mathcal{L}$.*
3. *$\mathcal{P}_{\aleph_0}(X)$ is the set of all finite subsets of X.*
4. *A graph in which no two distinct elements are connected is called an* empty graph. *A pair of distinct elements which are not connected is an* empty pair. *When R is an n-ary edge relation, to say that some X is an R-*empty graph *means that R does not hold on any n-tuple of distinct elements of X. X is an R-*complete graph *if R holds on every n-tuple from X.*
5. *Write $\varphi_n(x; y_1, \dots y_n)$ for the formula $\bigwedge_{i \leq n} \varphi(x; y_i)$.*
6. *In discussing graphs we will typically write concatenation for union, i.e. Ac for $A \cup \{c\}$.*
7. *A set is k-*consistent *if every k-element subset is consistent, and it is k-*inconsistent *if every k-element subset is inconsistent.*
8. *A formula $\psi(x; y)$ of $\mathcal{L}$ will be called* dividable *if there exists an infinite set $C \subset P_1$ and $k < \omega$ such that $\{\psi(x; c) : c \in C\}$ is 1-consistent but k-inconsistent. (Thus, by compactness, some instance of ψ divides.) When it is important to specify the arity k, write k-*dividable.

To each formula φ we associate a countable sequence of hypergraphs, the *characteristic sequence*, which describe incidence relations on the parameter space of φ.

Definition 2.2. (Characteristic sequences) *Let T be a first-order theory and φ a formula of the language of T.*

- *For $n < \omega$, $P_n(z_1, \dots z_n) := \exists x \bigwedge_{i \leq n} \varphi(x; z_i)$.*
- *The* characteristic sequence *of φ in T is $\langle P_n : n < \omega \rangle$.*
- *Write $(T, \varphi) \mapsto \langle P_n \rangle$ for this association.*
- *Convention: we assume that $T \vdash \forall y \exists z \forall x (\varphi(x; z) \leftrightarrow \neg\varphi(x; y))$. If this does not already hold for some given φ, replace φ with $\theta(x; y, z) = \varphi(x; y) \wedge \neg\varphi(x; z)$.*

Convention 2.3. We will ask a series of questions about whether certain, possibly infinite, configurations appear as subgraphs of the P_n. Note that for our purposes, the existence of these configurations is a property of T. That is, we may, as a way of speaking, ask if some configuration X appears, or is persistent, inside of some P_n; however, we will always mean *whether or not it is consistent with* T that there are witnesses to X inside of P_n interpreted in some sufficiently saturated model. Thus, the formulas P_n will often be identified with their interpretations in some monster model.

Observation 2.4. (Basic properties) *Let* $\langle P_n : n < \omega \rangle$ *be the characteristic sequence of* (T, φ)*. Then, regardless of the choice of* T *and* φ*, we will have:*

(1) *(Reflexivity)* $\forall x(P_1(x) \rightarrow P_n(x, \dots x))$*. In general, for each* $\ell \leq m < \omega$*,*

$$\forall z_1, \dots z_\ell, y_1, \dots y_m \Big((\{z_1, \dots z_\ell\} = \{y_1, \dots y_m\}) \implies (P_\ell(z_1, \dots z_\ell) \iff P_m(y_1, \dots y_m)) \Big)$$

(2) *(Symmetry) For any* $n < \omega$ *and any bijection* $g : n \rightarrow n$*,*

$$\forall y_1, \dots y_n \Big(P_n(y_1, \dots y_n) \iff P_n(y_{g(1)}, \dots y_{g(n)}) \Big)$$

(3) *(Monotonicity) For each* $\ell \leq m < \omega$*,*

$$\forall z_1, \dots z_\ell, y_1, \dots y_m \Big((\{z_1, \dots z_\ell\} \subseteq \{y_1, \dots y_m\}) \implies (P_m(y_1, \dots y_m) \implies P_\ell(z_1, \dots z_\ell)) \Big)$$

So in particular, if $\models P_m(y_1, \dots y_m)$ *and* $\ell < m$ *then* P_ℓ *holds on all* ℓ*-element subsets of* $\{y_1, \dots y_m\}$*. The converse is usually not true; see Remark 2.7.*

(4) *(Dividing) Suppose that for some* $n < \omega$*, it is consistent with* T *that there exists an infinite subset* $Y \subset P_n$ *such that* Y^k *is a* P_n*-empty graph. Then in any sufficiently saturated model of* T*, some instance of the formula* $\varphi_n(x; y_1, \dots y_n) = \bigwedge_{i<n} \varphi(x; y_i)$ k*-divides.*

(5) *(Consistent types) Let* $A \subset P_1$ *be a set of parameters in some* $M \models T$*. Then* $\{\varphi(x; a) : a \in A\}$ *is a consistent partial* φ*-type iff* $A^n \subset P_n$ *for all* $n < \omega$*.*

Proof. (4) By compactness, there exists an infinite indiscernible sequence of n-tuples $C = \langle c^i_1, \dots c^i_n : i < \omega \rangle$ such that $C^k \cap P_{nk} = \emptyset$. The set $\{\varphi_n(x; c^i_1, \dots c^i_n) : i < \omega\}$ is therefore k-inconsistent. However, it is 1-consistent: for each $c^i_1, \dots c^i_n \in C$, $M \models P_n(c^i_1, \dots c^i_n)$, so $M \models \exists x \varphi_n(x; c^i_1, \dots c^i_n)$. $\square$

A key object of study in this article will be the structure of certain complete graphs.

Definition 2.5. (Complete graphs)

(1) *Following Observation 2.4(5), a complete* P_∞*-graph will be called a* positive base set *to emphasize its representation of a consistent partial* φ*-type under analysis.*
(2) P_∞ *is shorthand for "*P_n*...for all* n*," e.g.* A *is a* P_∞*-complete graph meaning* A *is a* P_n*-complete graph for all* n*.*
(3) *Say that* $a \in P_1$ *is a* one-point extension *of the* P_n*-complete graph* A *just in case* Aa *is also a* P_n*-complete graph. In most cases,* n *will be* ∞*.*

When is an initial segment of the characteristic sequence enough to see the essential complexity? Recall that a formula $\varphi(x;y)$ has the *finite cover property* if for arbitrarily large $n < \omega$ there exist $a_0, \dots a_n$ such that $\{\varphi(x;a_0), \dots \varphi(x;a_n)\}$ is n-consistent but $(n+1)$-inconsistent.

Definition 2.6. *The sequence $\langle P_n \rangle$* has support k *if: $P_n(y_1, \dots y_n)$ iff P_k holds on every k-element subset of $\{y_1, \dots y_n\}$.*

Remark 2.7. *The following are equivalent, for $(T, \varphi) \mapsto \langle P_n \rangle$:*

(1) *There is $k < \omega$ such that the sequence $\langle P_n \rangle$ has support k.*
(2) *φ does not have the finite cover property.*

Note that the finite cover property can sometimes be avoided by a judicious choice of formula. For instance, given any unstable formula φ, some fixed finite conjunction of instances of φ has the finite cover property ([8].II.4). Nonetheless, it may happen in unstable theories that there is a set $\Sigma \subset \mathcal{L}$ of formulas without the fcp such that $M \models T$ is λ^+-saturated iff M realizes all φ_0-types over sets of size λ for all $\varphi_0 \in \Sigma$. This is true, for instance, of $\Sigma = \{\psi(x;y,z) := xRy \land \neg xRz\}$ in the random graph, and of $\Sigma = \{\psi(x;y,z) := y < x < z\}$ in $(\mathbb{Q}, <)$. In these cases, it suffices to analyze a sequence with finite support.

3. Static configurations

We now show how certain configurations in the characteristic sequence reflect classification-theoretic complexity of φ, namely the order property, the independence property, the tree property and SOP_2. "Static" means there is no attempt to change or definably restrict the vertex set P_1.

Convention 3.1. (T_0-configurations) *Throughout this article, let T_0 denote the incomplete theory in the language $\mathcal{L}_0 := \{P_n : n < \omega\} \cup \{=\}$ which describes (1)-(3) of Observation 2.4. Blueprints for hypergraphs in the language $\mathcal{L}_0$ which are consistent with T_0 will be called T_0-configurations. That is:* a finite T_0-configuration *is a pair $X = (V_X, E_X)$ where $V_X = n < \omega$, $E_X \subseteq \mathcal{P}(n)$ and the following is consistent with T_0:*

$$(1) \qquad (\exists x_1, \dots x_n)\ (\forall \sigma \subseteq n, |\sigma| = i, \sigma = \{\ell_1, \dots \ell_i\})\ \Big(P_i(x_{\ell_1}, \dots x_{\ell_i}) \iff \sigma \in E_X\Big)$$

In general, the domain of a T_0-configuration may be infinite; we simply require that its restriction to every finite subdomain satisfy (1). *These are the graphs which can consistenly occur as finite subgraphs of* some *characteristic sequence.*

That every such graph appears in some sequence follows from Example 5.30 below.

Convention 3.2. (T_1-configurations) *Fix T, φ, and the associated sequence $\langle P_n : n < \omega \rangle$. Let $M \models T$; there is a unique expansion of M to $\mathcal{L}_0 = \{P_n : n < \omega\} \cup \{=\}$. Throughout this article, whenever $T, \varphi, \langle P_n \rangle$ are thus fixed, let T_1 denote the complete theory of M in the language $\mathcal{L}_0$. As the characteristic sequence is definable in T, when T is complete this will not depend on the model chosen.*

*Hypergraphs in the language $\mathcal{L}_0$ which are consistent with T_1 will be called T_1-*configurations*.*

We recall several important tree properties from [9], [5].

Definition 3.3. (Tree properties) *Let $\subseteq$ indicate initial segment. To simplify notation, say that the nodes $\rho_1, \rho_2 \in \omega^{<\omega}$ are* tree-incomparable *if*

$$\neg(\rho_1 \subseteq \rho_2) \wedge \neg(\rho_2 \subseteq \rho_1) \wedge \neg(\exists \nu \in \omega^{<\omega}, i, j \in \omega)(\rho_1 = \nu^\frown i, \rho_2 = \nu^\frown j)$$

i.e., if they do not lie along the same branch and are not immediate successors of the same node.

Then the formula φ has:

- *the k-*tree property*, where $k < \omega$, if there is an $\omega^{<\omega}$-tree of instances of φ where paths are consistent and the immediate successors of any given node are k-inconsistent, i.e. $X = \{\varphi(x; a_\eta) : \eta \in \omega^{<\omega}\}$, and:*
  (1) *for all $\nu \in \omega^\omega$, $\{\varphi(x; a_\eta) : \eta \subseteq \nu\}$ is a consistent partial type;*
  (2) *for all $\rho \in \omega^{<\omega}$, $\{\varphi(x; a_{\rho^\frown i}) : i < \omega\}$ is k-inconsistent.*
  *Call any such X a φ-*tree*, or if necessary a φ-k-tree.*
- *the* tree property *if it has the k-tree property for some $2 \leq k < \omega$.*
- *the* strict tree property*, also known as TP_1 or SOP_2, if there exists a φ-tree with $k = 2$ and for which, moreover:*
  $(3)_1$ *if $\rho_1, \rho_2 \in \omega^{<\omega}$ are tree-incomparable, then $\neg\exists x(\varphi(x; a_{\rho_1}) \wedge \varphi(x; a_{\rho_2}))$.*
- *the* non-strict tree property *TP_2 if there exists a φ-tree with $k = 2$ and for which, moreover:*
  $(3)_2$ *if $\rho_1, \rho_2 \in \omega^{<\omega}$ are tree-incomparable, then $\exists x(\varphi(x; a_{\rho_1}) \wedge \varphi(x; a_{\rho_2}))$.*

Theorem A. (Shelah; see [8].III.7*)*

- *T is simple iff no formula φ of T has the tree property, iff no φ has the 2-tree property.*
- *If φ has the 2-tree property then either φ has TP_1 or φ has TP_2.*

We now describe a series of model-theoretically meaningful T_1-configurations. We may fix a monster model M from which the parameters are drawn, following Convention 2.3.

Definition 3.4. (Diagrams, arrays, trees) *Suppose $\aleph_0 \geq \lambda \geq \mu$. Write $\subseteq$ to indicate initial segment. Say that the sequence $\langle P_n \rangle$ has:*

(1) an $(\omega, 2)$-diagram *if there exist elements $\{a_\eta : \eta \in 2^{<\omega}\} \subseteq P_1$ such that*
  - *for all $\eta \in 2^{<\omega}$, $\neg P_2(a_{\eta^\frown 0}, a_{\eta^\frown 1})$, and*
  - *for all $n < \omega$ and $\eta_1, \ldots \eta_n \in 2^{<\omega}$, we have that $\eta_1 \subseteq \cdots \subseteq \eta_n \implies P_n(a_{\eta_1}, \ldots a_{\eta_n})$*

  That is, sets of pairwise comparable elements are P_∞-consistent, while immediate successors of the same node are P_2-inconsistent.

(2) a $(\lambda, \mu, 1)$-array *if there exists $X = \{a^m_l : l < \lambda, m < \mu\} \subset P_1$ such that:*
  - *$P_2(a^{m_1}_{l_1}, a^{m_2}_{l_2}) \iff (l_1 = l_2 \rightarrow m_1 = m_2)$*
  - *For all $i < \omega$,*

$$P_n(a^{m_1}_{l_1}, \ldots a^{m_n}_{l_n}) \iff \bigwedge_{1 \leq i,j \leq n} P_2(a^{m_i}_{l_i}, a^{m_j}_{l_j})$$

  The 1 refers to the fact that any $C \subset X$ is a P_∞-graph iff it contains no more than one *element from each column.*

(3) *Arrays will be revisited in Definitions 5.1 and 5.7. Arrays with only two parameters (λ, μ) refer to those more general objects from §5.*

(4) a (λ, μ)-tree *if there exist elements* $\{a_\eta : \eta \in \mu^{<\lambda}\} \subset P_1$ *such that*
   - *for all* $\eta_2, \eta_2 \in \mu^{<\lambda}$,

$$P_2(a_\eta, a_\nu) \iff (\eta_1 \subseteq \eta_2 \vee \eta_2 \subseteq \eta_1)$$

   i.e. only if *the nodes are comparable; and*
   - *for all* $n < \omega, \eta_1, \dots \eta_n \in \mu^{<\lambda}$,

$$\eta_1 \subseteq \dots \subseteq \eta_n \implies P_n(a_{\eta_1}, \dots a_{\eta_n})$$

Remark 3.5. *Diagrams are prototypes which can give rise to either arrays or trees, in the case where the unstable formula* φ *has the independence property or* SOP_2*, respectively.*

Claim 3.6. *Let* φ *be a formula of* T *and set* $\theta(x; y, z) = \varphi(x; y) \wedge \neg\varphi(x; z)$*. Let* $\langle P_n \rangle$ *be the characteristic sequence of* (T, θ)*. The following are equivalent:*

(1) $\langle P_n \rangle$ *has an* $(\omega, 2)$*-diagram.*
(2) $R(x = x, \varphi(x; y), 2) \geq \omega$*, i.e.* φ *is unstable.*
(3) $R(x = x, \theta(x; yz), 2) \geq \omega$*, i.e.* θ *is unstable.*

Proof. (2) $\to$ (1): We have in hand a tree of partial φ-types $\mathcal{R} = \{p_\nu : \nu \in 2^\omega\}$, partially ordered by inclusion, witnessing that $R(x = x, \varphi, 2) \geq \omega$. Let us show that we can build an $(\omega, 2)$-diagram. That is, we shall choose parameters $\{a_\eta : \eta \in 2^{<\omega}\} \subset P_1$ satisfying Definition 3.4(1).

First, by the definition of the rank R, which requires the partial types to be explicitly contradictory, we can associate to each ν an element $c_\nu \in M$, $\ell(c_\nu) = \ell(y)$ such that:

- $\varphi(x; c_\nu) \in p_{\nu^\frown 1} \setminus p_\nu$, and
- $\neg\varphi(x; c_\nu) \in p_{\nu^\frown 0} \setminus p_\nu$.

i.e., the split after index ν is explained by $\varphi(x; c_\nu)$.

Second, choose a set of indices $\mathcal{S} \subseteq 2^{<\omega}$ such that:

- $(\forall \eta \in 2^{<\omega})\ (\exists s \in \mathcal{S})(\eta \subsetneq s)$
- $(\forall s_1 \subsetneq s_2 \in \mathcal{S})\ (\exists \eta \notin \mathcal{S})\ (s_1 \subsetneq \eta \subsetneq s_2)$

It will suffice to define $a_{s^\frown i}$ for $s \in \mathcal{S}$, $i \in \{0, 1\}$. (The sparseness of S ensures the chosen parameters for φ will not overlap, which will make renumbering straightforward.) Recall that the a_η will be parameters for $\theta(x; y, z) = \varphi(x; y) \wedge \neg(x; z)$. So we define:

- $a_{s^\frown 0} = (c_{s^\frown 0}, c_s)$;
- $a_{s^\frown 1} = (c_s, c_{s^\frown 1})$.

The consistency of the paths through our $(\omega, 2)$-diagram is inherited from the tree $\mathcal{R}$ of consistent partial types. However, $\neg P_2(a_{s^\frown 0}, a_{s^\frown 1})$ because these contain an explicit contradiction:

$$\neg\exists x \left((\varphi(x; c_{s^\frown 0}) \wedge \neg\varphi(x; c_s)) \wedge (\varphi(x; c_s) \wedge \neg(\varphi(x; c_{s^\frown 1}) \right)$$

(1) $\to$ (3): Reading off the parameters from the diagram we obtain a tree of consistent partial θ-types $\{p_\eta : \eta \in 2^{<\omega}\}$, partially ordered by inclusion. For any $\eta \in 2^{<\omega}$, $\neg P_2(a_{\eta^\frown 0}, a_{\eta^\frown 1})$, i.e. $\neg\exists x(\theta(x; a_{\eta^\frown 0}) \wedge \theta(x; a_{\eta^\frown 1}))$. Furthermore, $\theta(x; a_{\eta^\frown 0}) \in p_{\eta^\frown 0} \setminus p_\eta$, while $\theta(x; a_{\eta^\frown 1}) \in p_{\eta^\frown 1} \setminus p_\eta$. So there is no harm in making the types explicitly inconsistent, as the rank R requires, by adding $\neg\, \theta(x; a_{\eta^\frown i})$ to $p_{\eta^\frown j}$ for $i \neq j < 2$.

(2) $\leftrightarrow$ (3): for all A, $|A| \geq 2$, $|S_\varphi(A)| = |S_\theta(A)|$. $\square$

Claim 3.7. *Let φ be a formula of T and set $\theta(x;y,z) = \varphi(x;y) \wedge \neg\varphi(x;z)$. Let $\langle P_n \rangle$ be the characteristic sequence of (T, θ). The following are equivalent:*

1. *$\langle P_n \rangle$ has an $(\omega, 2, 1)$-array.*
2. *φ has the independence property.*
3. *θ has the independence property.*

Proof. (1) $\rightarrow$ (3): Let $A = \{a^t_n : t < 2, n < \omega\}$ be such an array and set $A_0 := \{a^0_n : n < \omega\}$. We verify that φ has the independence property on A_0. Let σ, τ be any two disjoint finite subsets of ω. Let $B \subset A$ be $\{a^0_s : s \in \sigma\} \cup \{a^1_t : t \in \tau\}$. Then B is a positive base set. Any realization c of $\{\varphi(x;b) : b \in B\}$ will satisfy $\varphi(c;a^0_s)$ for $s \in \sigma$ and also $\neg\varphi(c;a^0_t)$ for $t \in \tau$ because $\neg P_2(a^0_t, a^1_t)$.

(2) $\rightarrow$ (1): Let $\langle i_\ell : \ell < \omega \rangle$ be a sequence over which φ has the independence property. For $t < 2, j < \omega$ set $a^0_j = (i_\ell, i_{\ell+1})$, $a^1_j = (i_{\ell+1}, i_\ell)$. Then $\{a^t_j : t < 2, j < \omega\}$ is an $(\omega, 2, 1)$-array for P_∞.

(3) $\rightarrow$ (2): For any infinite A, $|S_\varphi(A)| = |S_\theta(A)|$, as any type on one side can be presented as a type on the other. The independence property can be characterized in terms of the cardinality of the space of types over finite sets ([8] Theorem II.4.11). $\square$

Claim 3.8. *Let φ be a formula of T and set $\theta(x;y,z) = \varphi(x;y) \wedge \neg\varphi(x;z)$. Let $\langle P_n \rangle$ be the characteristic sequence of (T, θ). Suppose that T does not have SOP_2. Then the following are equivalent:*

1. *$\langle P_n \rangle$ has an $(\omega, \omega, 1)$-array.*
2. *φ has the 2-tree property.*

Proof. (1) $\rightarrow$ (2) Each column (=empty graph) of the array witnesses that φ is 2-dividable, and the condition that any subset of the array containing no more than one element from each column is a P_∞-complete graph ensures that the dividing can happen sequentially.

(2) $\rightarrow$ (1) By Theorem A above, $NSOP_2$ implies φ has TP_2. That is, there is a tree of instances $\{\varphi(x;a_\eta) : \eta \in \omega^{<\omega}\}$ such that first, for any finite n, $\eta_1 \subseteq \cdots \subseteq \eta_n$ implies that the partial type $\{\varphi(x;a_{\eta_1}), \ldots \varphi(x;a_{\eta_n})\}$ is consistent; and second,

$$\neg\exists x \left(\varphi(x;a_\eta) \wedge \varphi(x;a_\nu)\right) \iff (\exists \rho \in \omega^{<\omega})(\exists i \neq j \in \omega)\left(\eta = \rho^\frown i \;\wedge\; \nu = \rho^\frown j\right)$$

Thus the parameters $\{a_\eta : \eta \in \omega^{<\omega}\} \subset P_1$ form an $(\omega, \omega, 1)$-array for P_∞. $\square$

It is straightforward to characterize the analogous k-tree properties in terms of arrays whose columns are k-consistent but $(k+1)$-inconsistent.

Claim 3.9. *The following are equivalent:*

1. *$\langle P_n \rangle$ has an $(\omega, 2)$-tree.*
2. *φ has SOP_2.*

Proof. (2) $\rightarrow$ (1) This is a direct translation of Definition 3.3.

(1) $\rightarrow$ (2) It suffices to show that $\langle P_n \rangle$ has an (ω, ω)-tree. In an $(\omega, 2)$-tree, any two tree-incomparable nodes are inconsistent. Thus we may use the template of an $(\omega, 2)$-tree to produce an (ω, ω)-tree by compactness. $\square$

In the next definition, the power of the classification-theoretic order property on the P_n is magnified when it can be taken to describe the interaction between complete graphs, i.e. base sets for partial φ-types. Compare Remark 5.27.

Definition 3.10. P_∞ *has the* compatible order property *if there exists a sequence* $C = \langle a_i, b_i : i < \omega \rangle \subset P_1$ *such that for any* $n, m < \omega$, $0 < n+m$ *and any* $a_1, \dots a_m, b_1, \dots b_n \subset C$,

$$P_{m+n}(a_1, \dots a_m, b_1, \dots b_n) \iff ((m = 0) \vee (n = 0) \vee (\max\{a_1, \dots a_n\} < \min\{b_1, \dots b_n\}))$$

Observation 3.11. *Suppose* $(T, \varphi) \mapsto \langle P_n \rangle$*, and that* P_∞ *has the compatible order property. Then* φ_2 *has the tree property, and in particular,* SOP_2.

Proof. Let us build an SOP_2-tree $\{\varphi_2(x; a_\eta, b_\eta) : \eta \in \omega^{<\omega}\}$ following Definition 3.3 above by specifying the corresponding tree of parameters $\{c_\eta : \eta \in \omega^{<\omega}\} \subset P_1$, where each c_η is a pair (a_η, b_η). Let $S = \langle a_i b_i : i < \mathbb{Q} \rangle$ be an indiscernible sequence witnessing the compatible order property. We will use several facts in our construction:

(1) Let $\langle a_{i_\ell} b_{j_\ell} : \ell < \omega \rangle$ be any subsequence of S such that $\ell < k \implies a_{i_\ell} < b_{j_\ell} < a_{i_k} < b_{j_k}$. Then $\{\varphi_2(x; a_{i_j}, b_{i_j}) : j < \omega\}$ 2-divides by Observation 4.9.

(2) Let $a_{i_1}, b_{j_1}, \dots a_{i_n}, b_{j_n} \in S$. Then

$$P_n((a_{i_1}, b_{j_1}), \dots (a_{i_n}, b_{j_n})) \iff \max\{i_1, \dots i_n\} < \min\{j_1, \dots j_n\} \tag{3.2}$$

so in particular

$$P_2((a_{i_1}, b_{j_1}), (a_{i_2}, b_{j_2})) \iff \max\{i_1, i_2\} < \min\{j_1, j_2\} \tag{3.3}$$

Let $\eta \in \omega^{<\omega}$ be given and suppose that either c_η has been defined or $\eta = \emptyset$. If c_η has been defined, it will be (a_i, b_j) for some $i < j \in \mathbb{Q}$. Let $\langle k_\ell : \ell < \omega \rangle$ be any ω-indexed subset of $(i, j) \cap \mathbb{Q}$, or of $\mathbb{Q}$ if $\eta = \emptyset$. Define $c_{\eta^\frown \ell} = (a_{k_\ell}, b_{k_{\ell+1}})$. Now suppose we have defined the full tree of parameters c_η in this way. By fact (1) we see that immediate successors of the same node are P_2-inconsistent. By (3.2), paths are consistent, while by (3.3), any two tree-incomparable elements c_ν, c_η are P_2-inconsistent. $\square$

Remark 3.12. *The compatible order property is in fact enough to imply maximality in the Keisler order* [7].

4. Localization and persistence

We now consider "dynamic" arguments. Recall that a goal of these methods is to analyze φ-types by describing the graph-theoretic structure which is, in some sense, definably inseparable from a given positive base set A. This section defines localization, a definable restriction of P_1 in the graph language, and the key notion of persistence, a kind of invariance under localization.

P_n asks about incidence relations on a set of parameters; it will be useful to definably restrict the witness and parameter sets. For instance:

- we may ask that the witnesses lie inside certain instances of φ, e.g. by setting $P_1'(y) = \exists x(\varphi(x; y) \wedge \varphi(x; a))$, i.e. $P_1' = P_2(y, a)$.
- we may ask that the parameters be consistent 1-point extensions (in the sense of some P_n) of certain finite graphs C. For instance, we might define $P_1''(y) = P_1(y) \wedge P_2(y, c_1) \wedge P_3(y, c_2, c_2)$.

The next definition gives the general form.

Definition 4.1. (Localization) *Fix a characteristic sequence* $(T, \varphi) \to \langle P_n \rangle$*, and choose* $B, A \subset M \models T$ *with* A *a positive base set and* $A = \emptyset$ *possible.*

(1) (the localized predicate P^f_n) *A localization P^f_n of the predicate $P_n(y_1, \dots y_n)$ around the positive base set A with parameters from B is given by a finite sequence of triples $f : m \rightarrow \omega \times \mathcal{P}_{\aleph_0}(y_1, \dots y_n) \times \mathcal{P}_{\aleph_0}(B)$ where $m < \omega$ and:*
   - *writing $f(i) = (r_i, \sigma_i, \beta_i)$ and $\check{s}$ for the elements of the set s, we have:*
$$P^f_n(y_1, \dots y_n) := \bigwedge_{i \leq m} P_{r_i}(\check{\sigma}_i, \check{\beta}_i)$$
   - *for each $\ell < \omega$, T_1 implies that there exists a P_ℓ-complete graph C_ℓ such that P^f_n holds on all n-tuples from C_ℓ. If this last condition does not hold, P^f_n is a* trivial localization*. By* localization *we will always mean non-trivial localization.*
   - *In any model of T_1 containing A and B, P^f_n holds on all n-tuples from A.*

   Write $\mathrm{Loc}^B_n(A)$ for the set of localizations of P_n around A with parameters from B (i.e. nontrivial localizations, even when $A = \emptyset$).

(2) (the localized formula φ^f) *For each localization P^f_n of some predicate P_n in the characteristic sequence of φ, define the corresponding formula*
$$\varphi^f_n(x; y_1, \dots y_n) := \varphi_n(x; y_1, \dots y_n) \wedge P^f_n(y_1, \dots y_n)$$
   When $n = 1$, write $\varphi^f = \varphi^f_1$. Let $S^f_\varphi(N)$ denote the set of types $p \in S_\varphi(N)$ such that for all $\{\varphi_{i_1}(x; c_{i_1}), \dots \varphi_{i_n}(x; c_{i_n})\} \subset p$, $P^f_n(c_{i_1}, \dots c_{i_n})$. Then there is a natural correspondence between the sets of types
$$S^f_\varphi(N) \leftrightarrow S_{\varphi^f}(N)$$

(3) (the *localized formula $\varphi^{f+\overline{a}}$) *We have thus far described localizations of the parameters of φ. We will also want to consider restrictions of the possible witnesses to φ by adjoining instances of φ_k. That is, set*
$$\varphi^{f+\overline{a}}(x; y) = \varphi^{f+a_1,\dots a_k}(x; y) := \varphi(x; y) \wedge P^f_1(y) \wedge \varphi_k(x; a_1, \dots a_k)$$
   where, as indicated, $k = \ell(\overline{a})$. The * is to emphasize that this is really the construction from φ of a new, though related, formula, which will have its own characteristic sequence, given by:

(4) (the *localized characteristic sequence $\langle P^{f+\overline{a}}_n : n < \omega \rangle$) *The sequence $\langle P^{f+\overline{a}}_n : n < \omega \rangle$ associated to the formula $\varphi^{f+\overline{a}}$ is given by, for each $n < \omega$,*
$$P^{f+\overline{a}}_n(y_1, \dots y_n) = \bigwedge_{i \leq n} P^f_1(y_i) \wedge P_{n+k}(y_1, \dots y_n, a_1, \dots a_k)$$
   When f or $\overline{a}$ are empty, we will omit them.

Remark 4.2. *Convention 2.3 applies: that is, localization is not essentially dependent on the choice of model M. See Definition 4.14 (Persistence) and the observation following.*

As a first example of the utility of localization, notice that when φ is simple we can localize to avoid infinite empty graphs.

Observation 4.3. *Fix a positive base set A for the formula ψ, possibly empty. When ψ does not have the tree property, then for each $n < \omega$ there is a finite set C over which ψ is not n-dividable. As a consequence, if ψ does not have the tree property, then for each predicate*

P_n there is a localization around A on which there is a uniform finite bound on the size of a P_n-empty graph.

Proof. We proceed by asking: do there exist elements $\langle y_i : i < \omega \rangle$ such that (1) each y_i is a 1-point extension of A and (2) $\langle y_i : i < \omega \rangle$ is an n-empty graph? If not, localize using the finite set of conditions in (1) which prevent (2). Otherwise, let a_1, b_1 be the first pair in any such sequence, set $A_1 := A \cup \{a_1, b_1\}$ and repeat the argument using A_1 in place of A. Simplicity means there is a uniform finite bound on the number of times φ can sequentially n-divide. Condition (1) ensures that the dividing is sequential, corresponding to choosing progressive forking extensions of the partial type corresponding to A. At some finite stage t this will stop, meaning that (1) and (2) fail with A_t in place of A; the finite fragment of (1) which prevents (2) gives the desired localization. □

Remark 4.4. *By applying Observation 4.3 finitely many times, we may choose the localizations so that none of $\psi_1, \ldots \psi_\ell$ are k-dividable for any finite k, ℓ fixed in advance.*

The following important property of formulas was isolated by Buechler [1].

Definition 4.5. *The formula φ is* low *if there exists $k < \omega$ such that for every instance $\varphi(x; a)$ of φ, $\varphi(x; a)$ divides iff it r-divides for some $r \leq k$.*

Observation 4.6. *If φ does not have the independence property then φ is low.*

Proof. To show that any non-low formula φ has the independence property, it suffices to establish the consistency of the following schema. For $k < \omega$, Ψ_k says that there there exist $y_1, \ldots y_{2k}$ such that for every $\sigma \subset 2k$, $|\sigma| = k$,

$$\exists x \left(\varphi(x; y_i) \iff i \in \sigma \right)$$

But Ψ_k will be true on any subset of size $2k$ of an indiscernible sequence on which φ is k-consistent but $(k+1)$-inconsistent, and such sequences exist for arbitrarily large k by hypothesis of non-lowness. □

Corollary 4.7. *When the formula ψ of Observation 4.3 is simple* and *low, we can find a localization in which ψ is not k-dividable, for any k.*

In Section 5 we will occasionally assume that both φ and φ_2 are low. The following example justifies the assumption.

Observation 4.8. *$\varphi(x, \overline{y})$ can be low while $\varphi_2(x, \overline{y}, \overline{z}) = \varphi(x, \overline{y}) \wedge \varphi(x, \overline{z})$ is not.*

Proof. The idea is essentially to modify a non-low formula by appending a definable sequence of realizations. Consider the following countable model M. $\mathcal{L}$ contains three unary predicates X, Y, P, three binary predicates E, F, R and a constant c. X, Y are both infinite and partition $M \setminus \{c\}$. E is an equivalence relation on X with infinitely many infinite classes, and F is an equivalence relation on Y with infinitely many infinite classes.

- $\forall ab \, (R(a, b) \implies (a \in X \wedge b \in Y))$
- $\forall ab \, (R(a, b) \implies (\forall x (R(x, b) \implies E(x, a)) \wedge \forall y (R(a, y) \implies F(b, y))))$
- $(\forall x \in X)(\exists y \in Y) R(x, y) \wedge (\forall y \in Y)(\exists x \in X) R(x, y)$

Enumerate the countably many equivalence classes of E and the corresponding classes of F. Choose a_n, b_n in the *nth* classes so that $R(a_n, b_n)$ and set $P(a_n)$ for each n. Now interpret R so that in addition:

- For each n, $(\forall x)((E(x, a_n) \wedge x \neq a_n) \implies \exists^n y(R(x, y)))$
- For each n, $\forall y(F(y, b_n) \implies R(a_n, y))$
- For each n,
  $(\forall y_1 \dots y_n \in Y) \left(\bigwedge_{i \leq n} F(y_i, b_n) \implies (\exists x \in X) \left(x \neq a_n \wedge \bigwedge_{i \leq n} R(x, y_i)\right)\right)$

So $\{R(x, b) : F(b, b_n)\}$ is a consistent partial type but the sole realization is in P, and outside of P this family is strictly n+1-inconsistent. (One could clearly modify this example to add infinitely many realizations in P.) Now, $P(x)$ does not divide. $\{R(x; d_i) : i < \omega\}$ is inconsistent iff $\neg E(d_i, d_j)$ in which case it is 2-inconsistent so R is low.

However, consider the formula $\Theta(x, y, z) := (z = c \wedge R(x, y)) \vee (z \neq c \wedge \neg P(x))$. $\Theta(x, y, z)$ is low because given any instance $\Theta(x, a, b)$, any indiscernible sequence in ab which witnesses dividing will uniformly have $b = c$ or $b \neq c$, so it is not more complex than either R or $\neg P$. But $\Theta(x, y, z) \wedge \Theta(x, y', z')$ can sharply k-divide for any finite k because $\Theta(x, a_n, c) \wedge \Theta(x, b_n, b_n) = R(x, a_n) \wedge \neg P(x)$ which sharply n+1-divides by construction. □

4.1. **Stability in the parameter space.** The classification-theoretic complexity of the formulas P_n is often strictly less than that of the original theory T. Note that the results here refer to the *formulas* $P_n(y_1, \dots y_n)$, not necessarily to their full theory T_1.

Observation 4.9. *Suppose $(T, \varphi) \mapsto \langle P_n \rangle$. If $P_2(x; y)$ has the order property then $\varphi(x; y) \wedge \varphi(x; z)$ is 2-dividable.*

Proof. Let $\langle a_i, b_i : i < \omega \rangle$ be a sequence witnessing the order property for P_2, so $P_2(a_i, b_j)$ iff $i < j$. This means that $\exists x(\varphi(x; a_i) \wedge \varphi(x; b_j))$ iff $i < j$. So $\varphi(x; a_i) \wedge \varphi(x; b_{i+1})$ are consistent for each i, but the set $\{\varphi(x; a_i) \wedge \varphi(x; b_{i+1}) : i < \omega\}$ is 2-inconsistent. □

Remark 4.10. *By compactness, without loss of generality the sequence of Observation 4.9 can be chosen to be indiscernible (in the sense of T, not only T_1), and so to witness the dividing of some instance of φ_2.*

Note that the converse of Observation 4.9 fails: for $\varphi(x; y) \wedge \varphi(x; z)$ to divide it is sufficient to have a disjoint sequence of "matchsticks" in P_2 (i.e. $(a_i, b_i) : i < \omega$ such that $P_2(a_i, b_j)$ iff $i = j$), without the additional consistency which the order property provides.

Nonetheless, work relating the characteristic sequence to Szemerédi regularity illuminates the role of the order [7].

Observation 4.11. *Suppose that $(T, \varphi) \mapsto \langle P_n \rangle$, and for some n, k and some partition of $y_1, \dots y_n$ into k object variables and (n-k) parameter variables, $P_n(y_1, \dots y_k; y_{k+1}, \dots y_n)$ has the order property. Then $\varphi_n(x; y_1, \dots y_n)$ is 2-dividable.*

Proof. The proof is analogous to that of Observation 4.9, replacing the a_i by k-tuples and the b_j by (n-k)-tuples. □

Thus in cases where we can localize to avoid dividing of φ, we can assume any initial segment of the associated predicates P_n are stable:

Conclusion 4.12. *For each formula φ and for all $n < \omega$, if φ_{2n} does not have the tree property, then for each positive base set A there are a finite B and $P_1^f \in \mathrm{Loc}_1^B(A)$ over which $P_2, \dots P_n$ do not have the order property. In particular, this holds if T is simple.*

Proof. Similar to the proof of Observation 4.3 with the following modification: at each stage, ask whether there exist elements $\langle y_i z_i : i < \omega \rangle$ such that (1) each $y_i z_i$ is a 2-point extension of A and (2) $\langle y_i z_i : i < \omega \rangle$ witnesses the order property for P_{2n}. $\square$

By way of motivating the next subsection, let us prove the contrapositive: If the order property in P_2 persists under repeated localization, then φ has the tree property. Compare the proof of Observation 3.11 above. Without the compatible order property, we cannot ensure the tree is strict. While that argument built a tree out of a set of parameters which were given all at once (a so-called "static" argument), the following "dynamic" argument must constantly localize to find subsequent parameters, so cannot ensure that elements in different localizations are inconsistent.

Lemma 4.13. *Suppose that in every localization of P_1 (around $A = \emptyset$), P_2 has the order property. Then φ_2 has the tree property.*

Proof. Let us describe a tree with nodes (c_η, d_η), $(\eta \in \omega^{<\omega})$, such that:

(1) for each $\rho \in \omega^\omega$, $\{c_\eta, d_\eta : \eta \subseteq \rho\}$ is a complete P_∞-graph, where $\subseteq$ means initial segment.
(2) for any $\nu \in \omega^{<\omega}$, $P_2(c_{\eta^\frown i}, d_{\eta^\frown j}) \iff i \leq j$.

For the base case ($\eta \in \omega^1$), let $\langle c_i, d_i : i \in \omega \rangle$ be an indiscernible sequence witnessing the order property (so $P_2(c_i, d_j) \iff i \leq j$) and assign the pair (c_i, d_i) to node i.

For the inductive step, suppose we have defined (c_η, d_η) for $\eta \in \omega^n$. Write $E_\eta = \{(c_\nu, d_\nu) : \nu \leq \eta\}$ for the parameters used along the branch to (c_η, d_η). Using $\check{x}$ to mean the elements of the set x, let $P_1^{f_\eta}$ be given by $P_{n+1}((y,z), \check{E}_\eta)$. Let $\langle a_j, b_j : j \in \omega \rangle$ be an indiscernible sequence witnessing the order property inside this localization, and define $(c_{\eta^\frown i}, d_{\eta^\frown i}) := (a_i, b_i)$.

Finally, let us check that this tree of parameters witnesses the tree property for φ_2. On one hand, the order property in P_2 ensures that for each $n \in \omega^{<\omega}$, the set

$$\{\varphi_2(x; c_{\eta^\frown i}, d_{\eta^\frown i}) : i \in \omega\}$$

is 1-consistent but 2-inconsistent. On the other hand, the way we constructed each localization $P_1^{f_\eta}$ ensured that each path was a complete P_∞-graph, thus naturally a complete P'_∞-graph, where $\langle P'_n \rangle$ is the characteristic sequence of the conjunction φ_2. $\square$

4.2. **Persistence.** Localization, Definition 4.1 above, gives rise to a natural limit question: what happens when certain T_0-configurations persist under all finite localizations?

Definition 4.14. (Persistence) *Fix $(T, \varphi) \mapsto \langle P_n \rangle$, $M \models T$ sufficiently saturated, and a positive base set A, possibly $\emptyset$. Let X be a T_0-configuration, possibly infinite. Then X is* persistent *around the positive base set A if for all finite $B \subset M$ and for all $P_1^f \in Loc_1^B(A)$, P_1^B contains witnesses for X.*

We will write X is A-persistent to indicate that X is persistent around A.

Note 4.15. Persistence asks whether all finite localizations around A contain witnesses for some T_0-configuration X. The predicates P_n mentioned in X are, however, not the localized versions. We have simply restricted the set from which witnesses can be drawn. This is an obvious but important point: for instance, in the proof of Lemma 5.6 below it is important that the sequence of P_2-inconsistent pairs found inside of successive localizations $P_1^{f_n}$ are P_2-inconsistent in the sense of T_1.

Observation 4.16. (Persistence is a property of the theory T) *The following are equivalent, fixing $T, \varphi, \langle P_n \rangle$, $A \subset \mathcal{M}$ a small positive base set in the monster model, and a T_0-configuration X. Write $P_1^f(M)$ for the set which P_1^f defines in the model M.*

(1) *In some sufficiently saturated model $M \models T_1$ which contains A, X is persistent around A in M. That is, for every finite $B \subset M$ and every localization $P_1^f \in \mathrm{Loc}_1^B(A)$, there exist witnesses to X in $P_1^f(M)$.*

(2) *In every model $N \models T_1$, $N \supset A$, for every finite $B \subset N$, every localization $P_1^f \in Loc_1^B(A)$, and every finite fragment X_0 of X, $P_1^f(N)$ contains witnesses for X_0.*

Proof. (2) $\implies$ (1) Compactness.

(1) $\implies$ (2) Suppose not, letting $P_1^f \in Loc_1^B(A)$ and X_0 witness this. To this P_1^f we can associate a T_1-type $p(y_1, \dots y_{|B|}) \in S(A)$ which says that the localization given by f with parameters $y_1, \dots y_{|B|}$ contains A but implies that X_0 is inconsistent. But any sufficiently saturated model containing A will realize this type, and thus contain such a localization. ☐

To reiterate Convention 2.3, then, we may, as a way of speaking, call a configuration "persistent" while working in some fixed sufficiently saturated model, but we always refer to the corresponding property of T.

Corollary 4.17. *Persistence around the positive base set A remains a property of T in the language with constants for A.*

Finally, let us check the (easy) fact that persistence of some T_0-configuration around $\emptyset$ in some given sequence $\langle P_n \rangle$ implies its persistence around any positive base set A for that sequence. Recall that all localizations are, by definition, non-trivial.

Fact 4.18. *Suppose that X is an $\emptyset$-persistent T_0-configuration in the characteristic sequence $\langle P_n \rangle$ and A is a positive base set for $\langle P_n \rangle$. Then X remains persistent around A.*

Proof. Let $p(x_0, \dots)$ in the language $\mathcal{L}(=, P_1, P_2, \dots)$ describe the type, in V_X-many variables, of the configuration $X = (V_X, E_X)$. Let $q(y) \in S(A)$ be the type of a 1-point P_∞-extension of A in the language $\mathcal{L}_0 = \{P_n : n < \omega\} \cup \{=\}$. We would like to know that $q(x_0), q(x_1), \dots, p(x_0, \dots)$ is consistent, i.e., that we can find, in some given localization, witnesses for X from among the elements which consistently extend A. If not, for some finite subset $A' \subset A$, some $n < \omega$, and some finite fragments q' of $q|_{A'}$ and p' of p,

$$q'(x_0) \cup \dots \cup q'(x_n) \vdash \neg p'(x_0, \dots x_n)$$

But now localizing P_1 according to the conditions on the lefthand side (which are all positive conditions involving the P_n and finitely many parameters A') shows that X is not persistent, contradiction. ☐

Remark 4.19. *Example 5.28 shows that the condition that φ has the tree property is necessary, but not sufficient, for the order property in P_2 to be persistent.*

Question 4.20. *Is SOP_2 sufficient?*

Compare the issue of whether $SOP_2 \implies SOP_3$: see [2], [12].

5. Dividing lines: Stability, Simplicity, NIP

This section gives a new characterization of NIP in terms of persistence in the characteristic sequence, Theorem 5.17. We show that if φ is NIP then, given a positive base set A and an integer n, it is always possible to localize around A so that P_n becomes a complete graph. In other words, even the simplest configuration (a missing P_n-edge) is not persistent. Conversely, if φ has the independence property then there will be many missing P_n-edges in the vicinity of any positive base set. The section concludes with a characterization of simplicity in terms of the persistence of infinite empty graphs, Theorem 5.22.

Recall that a theory T is NIP [10] if no formula of T has the independence property. For more on the interest and importance of this hypothesis, see [13], [3].

5.1. **The case of P_2.** Let us first show that if φ is stable then we can localize around any fixed positive base set so that P_2 is a complete graph. The argument in this technically simpler case will generalize without too much difficulty.

Definition 5.1. ($(\omega, 2)$-arrays revisited)

(1) *The predicate P_n is $(\omega, 2)$ if there is $C := \{a^t_i : t < 2, i < \omega\}$ such that for all $\ell \leq n$, any ℓ-element subset C_0,*

$$P_\ell(C_0) \iff \left(a^t_i, a^s_j \in C_0 \implies (i \neq j) \vee (t = s)\right)$$

(2) *If for all $n < \omega$, P_n is $(\omega, 2)$, we say that P_∞ is $(\omega, 2)$.*

(3) *A* path *through the $(\omega, 2)$-array A is a set $X \subset A$ which contains no more than one element from each column. So paths are positive base sets.*

Remark 5.2. *If P_∞ is $(\omega, 2)$, then φ has the independence property, by the proof of Claim 3.7.*

Lemma 5.3. *(Springboard lemma for 2) If φ is stable then there is a finite localization P^f_1 for which the following are equivalent:*

(1) *There exists $X \subset P^f_1$, X an $(\omega, 2)$-array with respect to P_2*

(2) *There exists $Y \subset P^f_1$, Y an $(\omega, 2)$-array with respect to P_∞*

Proof. Choose the localization P^f_1 according to Observation 4.3 so that neither φ nor φ_2 are dividable using parameters from P^f_1. This is possible because stable formulas are simple and low, and φ stable implies φ_2 stable. Let $Z = \langle c^t_i : t < 2, i < \omega\rangle \subset P^f_1$ be an indiscernible sequence of pairs which is an $(\omega, 2)$-array for P_2. Each of the sub-sequences $\langle c^0_i : i < \omega\rangle$, $\langle c^1_i : i < \omega\rangle$ is indiscernible, so will be either P_2-complete or P_2-empty; by choice of P^f_1, they cannot be empty.

It remains to show that any path through Z is a P_∞-complete graph. Suppose not, and let n be minimal so that the n-type of some increasing sequence of elements $c^{t_1}_1, \ldots c^{t_n}_n$ implies $\neg\exists x(\bigwedge_{i<n} \varphi(x; c^{t_i}_i))$. Choose an infinite indiscernible subsequence of pairs $Z' \subset Z^2$ of the form $\langle c^0_i, c^1_{i+1} : i \in W \subset \omega\rangle$. Then the set $\{\varphi(x; c^0_i) \wedge \varphi(x; c^1_{i+1}) : i \in W\}$ will be 1-consistent by definition but n-inconsistent by assumption (though not necessarily (n-1)-consistent). This contradicts the assumption that φ_2 is not dividable in P^f_1. □

When the formula is low but not necessarily simple, bootstrapping up to P_∞ requires a stronger initial assumption on the array: the finite bound may not be 2.

Corollary 5.4. *Suppose the formulas φ and φ_2 are low. Then there exists $k < \omega$ such that, in any localization P_1^f, the following are equivalent:*

(1) *There exists $X \subset P_1^f$, X an $(\omega, 2)$-array with respect to P_k*
(2) *There exists $Y \subset P_1^f$, Y an $(\omega, 2)$-array with respect to P_∞*

Proof. Let k_0 be a uniform finite bound on the arity of dividing of instances of φ and φ_2, using lowness; by the proof of the previous Lemma, any $k > 2k_0$ will do. $\square$

Recall from Definition 2.1 that an "empty pair" is the T_0-configuration given by $V_x = 2, E_x = \{\{1\}, \{2\}\}$, i.e., a pair y, z such that $P_1(y)$, $P_1(z)$ but $\neg P_2(y, z)$.

Lemma 5.5. *Suppose φ is stable, and that every localization P_1^f around some fixed positive base set A contains an empty pair. Then P_∞ is $(\omega, 2)$.*

Proof. Choose $P_1^{f_0}$ to be a localization given by Lemma 5.3. We construct an $(\omega, 2)$-array as follows.

At stage 0, let c_0^0, c_0^1 be any pair of P_2-incompatible elements each of which is a consistent 1-point extension of A in $P_1^{f_0}$. At stage $n+1$, write C_n for $\{c_i^t : t < 2, i \leq n\}$ and suppose we have defined $P_1^{f_n} \in \mathrm{Loc}_1^{C_n}(A)$. By hypothesis, there are $c_{n+1}^0, c_{n+1}^1 \in P_1^{f_n}$ such that $\neg P_2(c_{n+1}^0, c_{n+1}^1)$ and such that each c_{n+1}^i is a consistent 1-point extension of A (Fact 4.18). Let $C_{n+1} = C_n \cup \{c_{n+1}^0, c_{n+1}^1\}$ and define $P_1^{f_{n+1}} \in \mathrm{Loc}_1^{C_{n+1}}(A)$ by

$$P_1^{f_{n+1}}(y) = P_1^f(y) \wedge P_2(y; c_{n+1}^0) \wedge P_2(y; c_{n+1}^1)$$

Thus we construct an $(\omega, 2)$-array for P_2, as desired. Applying Lemma 5.3 we obtain an $(\omega, 2)$-array for P_∞. $\square$

Conclusion 5.6. *Suppose that φ is stable, $(T, \varphi) \mapsto \langle P_n \rangle$ and A is a positive base set. Then empty pairs are not persistent around A.*

Proof. By stability, we may work inside the localization given by Lemma 5.3. Suppose empty pairs were persistent around A. By Lemma 5.5, P_∞ is $(\omega, 2)$, which by Remark 5.2 implies that φ has the independence property: contradiction. $\square$

This completes the proof for P_2 and φ stable. Corollary 5.4 shows that in order to replace the hypothesis of stable with low, it suffices to replace P_2-consistency in the proof of Lemma 5.5 with P_k-consistency. This will be done in Lemma 5.14.

5.2. **NIP: the case of n.** We now build a more general framework, working towards Theorem 5.15: if T is NIP then no P_n-empty tuple can be persistent. The basic strategy is as follows. Definition 5.7 gives a general definition of (ω, n)-array. Lemma 5.14 shows that if a P_n-empty tuple is persistent, an (ω, n)-array must exist. In this higher-dimensional case, in order to extract the independence property from an (ω, n)-array via Observation 5.9, we need the array to have an additional property called sharpness. The "sharpness lemma," Lemma 5.12, returns an array of the correct form at the cost of possibly adding finitely many parameters. Fact 5.13 then pulls this down to the independence property for φ.

With some care, we are able to get quite strong control on the kind of localization used. When T is stable in addition to NIP, the argument can be done with a uniform finite bound (as a function of n) on the arity of the predicates P_m used in localization.

Definition 5.7. ((ω, n)-arrays revisited) *Assume $n \leq r < \omega$. Compare Definition 3.4; here, the possible ambiguity of the amount of consistency will be important.*

(1) *The predicate P_r is (ω, n) if there is $C = \{c_i^t : t < n, i < \omega\} \subset P_1$ such that, for all $c_{i_1}^{t_1}, \ldots c_{i_r}^{t_r} \in C$,*
   - *r-tuples from r distinct columns are consistent, i.e.*
   $$\bigwedge_{j,k \leq r} i_j \neq i_k \implies P_r(c_{i_1}^{t_1}, \ldots c_{i_r}^{t_r})$$
   - *and no column is entirely consistent, i.e. for all $\sigma \subset r$, $|\sigma| = n$,*
   $$\bigwedge_{j,k \in \sigma} i_j = i_k \implies \neg P_r(c_{i_1}^{t_1}, \ldots c_{i_r}^{t_r})$$

   *Any such C is an (ω, n)-*array*. The precise arity of consistency is not specified, see condition (4).*

(2) *If for all $n \leq r < \omega$, P_r is (ω, n), say that P_∞* is *(ω, n).*

(3) *A* path *through the (ω, n) array C is a set $X \subset C$ which contains no more than n-1 elements from each column.*

(4) *P_r is* sharply *(ω, n) if it contains an (ω, n)-array C on which, moreover, for all $\{c_{i_1}^{t_1}, \ldots c_{i_r}^{t_r}\} \subset C$*
$$P_r(c_{i_1}^{t_1}, \ldots c_{i_r}^{t_r}) \iff \bigwedge_{\sigma \subset r, |\sigma| = n} \left(\bigwedge_{j,k \in \sigma} i_j = i_k \implies \bigvee_{j \neq k \in \sigma} t_j = t_k \right)$$
   i.e., if every path is a P_r-complete graph.

(5) *P_∞ is sharply (ω, n) if P_r is sharply (ω, n) for all $n \leq r < \omega$. By compactness and Convention 2.3, we may assume that there exists a single array C witnessing this simultaneously for all P_r.*

(6) *An (ω, n)-array is* sharp for P_k *if it satisfies the conditions of (4) for all $r \leq k$. An array is simply* sharp *if it is sharp for all P_r, $r < \omega$.*

Remark 5.8. (1) *For any $n > 1$, if P_n is $(\omega, 2)$ then it is by definition sharply $(\omega, 2)$.*

(2) *Suppose P_∞ has an (ω, n)-array; this does not necessarily imply that P_∞ has an (ω, m)-array for $m < n$, because m elements from a single column need not be inconsistent, e.g. if the (ω, n)-array is sharp.*

Observation 5.9. *If P_∞ is sharply (ω, k) then φ_{k-1} has the independence property.*

Proof. Let $X = \langle a_i^1, \ldots a_i^k : i < \omega \rangle$ be the array in question; then φ_{k-1} has the independence property on any maximal path, e.g. $B := \langle a_i^1, \ldots a_i^{k-1} : i < \omega \rangle$. To see this, fix any $\sigma, \tau \subset \omega$ finite disjoint; then by the sharpness hypothesis $\{a_i^1, \ldots a_i^{k-1} : i \in \sigma\} \cup \{a_j^2, \ldots a_j^k : j \in \tau\}$ is a P_∞-complete graph and thus corresponds to a consistent partial φ-type q. But any realization α of q cannot satisfy $\varphi(x; a_j^1)$ for any $j \in \tau$, because P_k does not hold on the columns. A fortiori $\neg\varphi_k(\alpha; a_j^1, \ldots a_j^k)$. □

Let us write down some conventions for describing types in an array.

Definition 5.10. *Let x_i^t, x_j^s be elements of some (ω, n)-array X.*

(1) *Let $[x_i^t] = \{x_j^s \in X : j = i\}$, i.e. the elements in the same column as x_i^t.*

(2) *Let* $X_0 = \{x^{t_1}_{i_1}, \dots x^{t_\ell}_{i_\ell}\} \subset X$ *be a finite subset. The* column count *of* $\{x^{t_1}_{i_1}, \dots x^{t_\ell}_{i_\ell}\}$ *is the unique tuple* $(m_1, \dots m_\ell) \in \omega^\ell$ *such that:*
- $m_i \geq m_{i+1}$ *for each* $i \leq \ell$
- $\Sigma_i \, m_i = \ell$
- *if* $Y_0 = \{y_1, \dots y_r\}$ *is a maximal subset of* X_0 *such that* $y, z \in Y_0, y \neq z \;\rightarrow\; y \notin [z]$*, then some permutation of*

$$\left(\Big|[y_1] \cap X_0\Big|, \dots, \Big|[y_r] \cap X_0\Big|\right)$$

is equal to $(m_1, \dots m_\ell)$.

In other words, we count how many elements have been assigned to each column, and put these counts in descending order of size. Write col-ct$(\overline{x})$ *for this tuple.*

(3) *Let* $\leq$ *be the lexicographic order on column counts, i.e.* $(1, 1, \dots) < (2, 1, \dots)$. *This is a discrete linear order, so we can define* $(m_1, \dots m_\ell)^+$ *to be the immediate successor of* $(m_1, \dots m_\ell)$ *in this order. Define* gap$((m_1, \dots m_\ell)) = m_i$ *where* $((n_1, \dots n_\ell)^+ = (m_1, \dots m_\ell)$ *and* $\forall j \neq i \; m_j = n_j$, *i.e. the value which has just incremented.*

Lemma 5.11. (Springboard lemma) *Fix* $2 \leq n < \omega$, *and let* $\langle P_r \rangle$ *be the characteristic sequence of* (T, φ). *Suppose that the formulas* $\varphi, \varphi_2, \dots \varphi_{2n-2}$ *are low. Then there exist* k_0, $1 \leq k_0 < \omega$ *and a localization* P^f_1 *of* P_1 *in which the following are equivalent:*

(1) P^f_1 *contains a sharp* (ω, n)*-array for* P_μ, *where* $\mu = (2n-2)k_0$.
(2) P^f_1 *contains a sharp* (ω, n)*-array for* P_∞.

Proof. Assume (1), so let $C = \{c^t_i : t < n, i < \omega\} \subset P^f_1$ be sharply (ω, n) for P_μ, chosen without loss of generality to be an indiscernible sequence of n-tuples. The hypothesis is that any subset of this array C, of any size, which does not include all the elements of any column (in other words: any path) is a complete P_μ-graph. We would like to show that any path is in fact a complete P_∞-graph.

Suppose not, so fix a path $Y = y_1, \dots y_m$ of minimal size m such that $\neg P_m(y_1, \dots y_m)$. Notice that $m \geq \mu \geq n$ by definition. Let $S := \{c^0_i, \dots c^{n-1}_i, c^1_{i+1}, \dots c^n_{i+1} : i < \omega\} \subset C^{2n-2}$ be a sequence of pairs of offset $(n-1)$-tuples.

Note that S is 1-consistent as we assumed (1).

On the other hand, C is indiscernible, so any increasing sequence of m elements from S will cover all the possible m-types from C. Since Y is inconsistent, this implies that S is m-inconsistent. These m elements will be distributed over at least $\frac{m}{2n-2}$ instances of φ_{2n-2}; by inductive hypothesis, one fewer element, thus one fewer instance, would be consistent. Thus φ_{2n-2} is sharply m'-dividable for some $m' \geq k_0$.

The appropriate k_0 is thus a strict upper bound on the possible arity of dividing of each of the formulas $\{\varphi_{2\ell-2} : 1 \leq \ell \leq k_0\}$, which exists by lowness. When T is low but possibly unstable, determining k_0 is the important step; no localization is then necessary. When T is stable, however, w.l.o.g. $k_0 = 2n-2$ as by Corollary 4.7 we can simply choose a localization in which the $2n-2$ formulas are not k-dividable for any k. ☐

We next give a lemma which will extract a sharp array from an array. Recall that $P^{\overline{a}}_\infty$ is the *localized *sequence* from Definition 4.1, i.e. the characteristic sequence of the formula $\varphi(x;y) \wedge \bigwedge_{a \in \overline{a}} \varphi(x;a)$.

Lemma 5.12. (Sharpness lemma) *Let $\overline{a} \subset P_1$ be finite, $n < \omega$. Suppose that $P_\infty^{\overline{a}}$ contains an (ω, n)-array. Then there exist $\overline{a}'$, ℓ with $\overline{a} \subseteq \overline{a}' \subset P_1$ and $2 \leq \ell \leq n$ such that $P_\infty^{\overline{a}'}$ contains a sharp (ω, ℓ)-array.*

Proof. Let us show that, given an (ω, n)-array for $P_\infty^{\overline{a}}$, either

- there is a *sharp* (ω, n) array for $P_\infty^{\overline{a}}$, or else
- by adding no more than finitely many parameters we can construct an (ω, ℓ)-array for $P_\infty^{\overline{a}'}$ and some $\ell < n$.

Note that the second is nontrivial by Remark 5.8. As an $(\omega, 2)$-array is automatically sharp, we can then iterate the argument to obtain the lemma.

We have, then, some (ω, n)-array C in hand. Without loss of generality C is an indiscernible sequence of n-tuples. Now, by definition, every subset Y of C which contains no more than one element from each column is a $P_\infty^{\overline{a}}$-complete graph. Consider what happens when we add elements to Y while still not choosing an entire column, i.e. ensuring the larger set remains a path. If every path through C is a $P_\infty^{\overline{a}}$-complete graph, then by definition C is a sharp (ω, n)-array and we are done. Otherwise, some path will not be a $P_\infty^{\overline{a}}$-complete graph. More precisely, by compactness, there will be some finite $X \subset C$, which we can write more precisely as $X := Z \cup Y$ where:

(1) Z is a path
(2) $Y \subset C$, possibly empty, is such that
  (i) $Y \cap [Z] = \emptyset$
  (ii) if Y is nonempty, then $y_1, y_2 \in Y \implies [y_1] \cap [y_2] = \emptyset$
  but $Z \cup Y$ is not a $P_\infty^{\overline{a}}$-complete graph.

There may be many such Z; if so, choose any one with minimal column count (possible as this is well ordered).

The assumption that C is not sharp gives an unspecified finite bound on $|Z|$; in fact the springboard lemma gives a more informative bound $k \geq 2n - 2$. On the other hand, by definition of (ω, n)-array, any such Z must contain at least two elements from the same column, so $|Z| > 1$ and we can isolate our witness by systematically checking for inconsistency as column count grows (to ensure that behavior below the witness is as uniform as possible). Because C is an indiscernible sequence of n-tuples, we may assume that the elements of Z are in columns which are arbitrarily far apart (or indeed, by compactness, infinitely far apart). Finally, if $Z_0 \subsetneq Z$, then $\text{col-ct}\,(Z_0) < \text{col-ct}\,(Z)$. So for any $W \subset C$ satisfying conditions (2)(i)-(ii) just given, $Z_0 \cup W$ is a $P_\infty^{\overline{a}}$-complete graph.

In particular, we can choose a partition $X = X_0 \cup X_1$ where

(I) $X_0 \cap X_1 = \emptyset$ and $\emptyset \subsetneq X_1 \subset X$
(II) $x, x' \in X_1 \implies [x] = [x']$
(III) $n > \ell := |X_1| = gap(\text{col-ct}(Z)) > 1$
(IV) For any $W \subset C$ satisfying (2)(i)-(ii), we have that $X_0 \cup W$ is a $P_\infty^{\overline{a}}$-complete graph.

To finish, let $a' = a \cup X_0$ and let $C' \subset C$ be an infinite sequence of ℓ-tuples which realize the same type as X_1 over $a \cup X_0$. (For instance, restrict C' to the rows containing elements of X_1 and to infinitely many columns which do not contain elements of X_0.) Since Z was chosen to be a path, $\ell < n$ (condition (III)) and $|a'| < |X| < \omega$. By condition (2), $\neg P_\ell^{\overline{a}'}(\overline{c})$ for any column $\overline{c}$ of C'. On the other hand, by condition (IV) any subset of C' containing no more than one element from each column is a $P^{\overline{a}'}_\infty$-complete graph. Thus C' is an (ω, ℓ)-array for $P^{\overline{a}'}_\infty$, as desired. If it is not sharp, repeat the argument. □

Fact 5.13. *The following are equivalent for a formula $\varphi(x; y)$.*

(1) *φ has the independence property.*
(2) *For some $n < \omega$, φ_n has the independence property.*
(3) *For every $n < \omega$, φ_n has the independence property.*
(4) *Some *localization $\varphi^{\overline{a}}$ has the independence property.*

Proof. (1) → (3) → (2) → (1) → (4) are straightforward: use the facts that the formulas φ_i, φ_j generate the same space of types, and that the independence property can be characterized in terms of counting types over finite sets ([8]:II.4). Finally, (4) → (2) as we have simply specified some of the parameters. □

Lemma 5.14. *Suppose that for some $n < \omega$, every localization of P_1 around some fixed positive base set A contains an n-tuple on which P_n does not hold. Then P_∞ is (ω, n), though not necessarily sharply (ω, n).*

Proof. Let us show that P_k is (ω, n) for any $k \geq n$. This suffices as, by Convention 2.3, we may apply compactness.

Fix $k \geq n$ and let P^f_1 be any localization, for instance that of Lemma 5.11.

At stage 0, let $c^0_0, c^1_0, \dots c^{n-1}_0 \subset P^{f_0}_1 := P^f_1$ be an n-tuple of elements on which P_n does not hold, chosen by Fact 4.18 so that each c^i_0 is a consistent 1-point extension [in the sense of P_n] of A. Let $X_0 = \{\{c^i_0\} : i \leq n\}$ be the set of these singletons. Write $\check{x}$ to denote the elements of x. Define

$$P^{f_1}_1(y) = P^{f_0}_1(y) \wedge \bigwedge_{x \in X_0} P_2(y; \check{x})$$

which includes A by construction.

At stage $m+1$, write C_m for $\{c^t_i : t < n, i \leq m\}$ and consider the localized set of elements $P^{f_m}_1 \in \mathrm{Loc}^{C_m}_1(A)$. Let

$$X_m := \{x \subset C_m : \ |x| = k-1 \text{ and for all i < m, } |x \cap (C_{i+1} \setminus C_i)| \leq 1\}$$

i.e. sets which choose no more than one element from each stage in the construction.

By hypothesis, there are $c^0_{m+1}, \dots c^{n-1}_{m+1} \in P^{f_m}_1$ such that $\neg P_n(c^0_{m+1}, \dots c^{n-1}_{m+1})$ and such that for all $x \in X_m$, each c^i_{m+1} is a consistent 1-point extension of $A \cup x$, in the sense of P_n. Let $C_{m+1} = C_m \cup \{c^0_{m+1}, \dots c^{n-1}_{m+1}\}$, and let X_{m+1} be the sets from C_{m+1} which choose no more than one element from each stage in the construction. We now define $P^{f_{m+1}}_1 \in \mathrm{Loc}^{C_{m+1}}_1(A)$ by

$$P^{f_{m+1}}_1(y) = P^{f_m}_1(y) \wedge \bigwedge_{x \in X_{m+1}} P_k(y; \check{x})$$

(If $m < k$, the parameters from $\check{x}$ need not necessarily be distinct.) Again, this localization contains A by construction. Thus we construct an (ω, n)-array for P_k, as desired. As k was arbitrary, we finish. □

Recall that a P_n-empty tuple is any T_0-configuration for which $X = n$ and $\{1, \dots n\} \notin E_x$, that is, $y_1, \dots y_n \in P_1$ such that $\neg P_n(y_1, \dots y_n)$. We are now in a position to prove:

Theorem 5.15. *Suppose that φ is NIP and A is a positive base set for φ. Then for each $n < \omega$ we have that P_n-empty tuples are not persistent around A.*

Proof. By lowness, we work inside the localization $P_1^f \supset A$ given by the Springboard Lemma 5.11. Suppose that P_n-empty tuples are persistent for some $n < \omega$. Apply Lemma 5.14 to obtain an (ω, n)-array for P_∞, which is not necessarily sharp. The Sharpness Lemma 5.12 then gives a sharp (ω, ℓ)-array for $P_\infty^{\overline{a}}$, where $\overline{a} \subset P_1$ is a finite set of parameters. The sequence $\langle P_n^{\overline{a}} \rangle$ is just the characteristic sequence of the *localized formula $\varphi^{\overline{a}}$, that is, $\varphi(x;y) \wedge \varphi_m(x;\overline{a})$, where $m = |\overline{a}|$. By Observation 5.9 this means $\varphi_{\ell-1}^{\overline{a}}$ has the independence property. Now by (2) $\rightarrow$ (1) of Fact 5.13 applied to $\varphi^{\overline{a}}$, we see that $\varphi^{\overline{a}}$ has the independence property. By (4) $\rightarrow$ (1) of the same fact, φ must also have the independence property, contradiction. □

Corollary 5.16. *Suppose that φ is NIP, $(T, \varphi) \mapsto \langle P_n \rangle$ and A is a positive base set for φ. Then for each $n < \omega$, there is a localization P_1^f such that:*

(1) $A \subset P_1^f$
(2) $\{y_1, \ldots y_n\} \subset P_1^f \implies P_n(y_1, \ldots y_n)$

In other words, when φ is NIP, given a positive base set A there is, for each n, a definable restriction of P_1 containing A on which P_n is a complete graph. In the other direction, if φ has the independence property then P_∞ is $(\omega, 2)$ by Claim 3.7, so in particular P_1 is not a complete graph. In fact:

Theorem 5.17. *Let φ be a formula of T and $\langle P_n : n < \omega \rangle$ its characteristic sequence. Then the following are equivalent for any positive base set A:*

(1) *There exists a localization φ^f of φ such that φ^f is NIP and $P_1^f \supset A$.*
(2) *There exists a localization φ^g of φ such that φ^g is stable and $P_1^g \supset A$.*
(3) *For every $n < \omega$, there exists a localization $P_1^{f_n} \supset A$ which is a P_n-complete graph.*

Proof. Note that $\varphi^f = \varphi$ and $\varphi^g = \varphi$ are possible.

(2)$\rightarrow$(1) because stable implies NIP.

(1)$\rightarrow$(3) is just Theorem 5.15: no P_n-empty tuple is persistent, so eventually one obtains a localization which is a complete graph.

(3)$\rightarrow$(2) By Claim 3.6, if φ^g has the order property its associated P_1^g contains a diagram in the sense of Definition 3.4. Thus it contains an empty pair, and so a fortiori a P_n-empty tuple, for each n. □

Example 5.18. *Consider $(\mathbb{Q}, <)$, let $\varphi(x; y, z) = y > x > z$ and let the positive base set A be given by concentric intervals $\{(a_i, b_i) : i < \kappa\} \subset P_1$. Then there is indeed a P_2-empty pair $(c_1, c_2), (d_1, d_2)$ which are each consistent 1-point extensions of A – namely, any pair of disjoint intervals lying in the cut described by the type corresponding to A. Localizing to require consistency with any such pair amounts to giving a definable complete graph containing A, i.e. realizing the type.*

5.3. **Simplicity.** We have seen that the natural first question for persistence, whether there exist persistent empty tuples, characterizes stability: Theorem 5.17. Here we will show that a natural next question, whether there exist persistent infinite empty graphs, characterizes simplicity. Recall that a formula is simple if it does not have the tree property; see [5], [14].

Notice that we have an immediate proof of this fact by Observation 4.3, which appealed to finite $D(\varphi, k)$-rank for simple formulas to conclude that infinite empty graphs are not persistent. Let us sketch the framework for a different proof by analogy with the previous section. This amounts to deriving Observation 4.3 directly in the characteristic sequence.

Remark 5.19. *In the case of stability, much of the work came in establishing sharpness of the (ω, ℓ)-array. Here, since the persistent configuration is infinite, we have compactness on our side; we may in fact always choose the persistent empty graphs to be indiscernible and uniformly k-consistent but $(k+1)$-inconsistent, for some given $k < \omega$.*

Observation 5.20. *Suppose that $(T, \varphi) \mapsto \langle P_n \rangle$. Then the following are equivalent:*

(1) *there is a set $T = \{a_\eta : \eta \in 2^{<\omega}\} \subset P_1$ such that, writing $\subseteq$ for initial segment:*
   (a) *For each $\nu \in 2^\omega$, $\{a_\eta : \eta \subset \nu\}$ is a complete P_∞-graph.*
   (b) *For some $k < \omega$, and for all $\rho \in \omega^{<\omega}$, the set $\{a_{\rho^\frown i} : i < \omega\} \subset P_1$ is a P_k-empty graph.*

(2) *φ has the k-tree property.*

Proof. This is a direct translation of Definition 3.3. □

Lemma 5.21. *Let X_k be the T_0-configuration describing a strict $(k+1)$-inconsistent sequence, i.e. $V_{X_k} = \omega$ and $E_{X_k} = \{\sigma : \sigma \subset \omega, |\sigma| \leq k\}$. Suppose that for some fixed $k < \omega$ and some formula φ, X_k is persistent in the characteristic sequence $\langle P_n \rangle$ of φ. Then φ is not simple.*

Proof. Let us show that φ has the tree property, around some positive base set A if one is specified. At stage 0, by hypothesis there exists an infinite indiscernible sharply $(k+1)$-inconsistent sequence $Y_0 \subset P_1$, each of whose elements can be chosen to be a consistent 1-point extension of A in the sense of P_∞ by Fact 4.18. Set a_i to be the ith element of this sequence, for $i < \omega$.

At stage $t+1$, suppose we have constructed a tree of height n, $T_n = \{a_\eta : \eta \in \omega^{\leq n}\}$ such that, writing $\subseteq$ for initial segment:

- every path is a consistent n-point extension of A, that is, for each $\nu \in \omega^n$ $A \cup \{a_\eta : \eta \subseteq \nu\}$ is a complete P_∞-graph, and
- for all $0 \leq k < n$ and all $\eta \in \omega^k$, we have that $\{a_{\eta^\frown i} : i < \omega\}$ is P_k-complete but P_{k+1}-empty.

We would like to extend the tree to level $n+1$, and it suffices to show that the extension of any given node a_ν (for $\nu \in \omega^n$) can be accomplished. But this amounts to repeating the argument for stage 0 in the case where $A = A \cup \{a_\eta : \eta \subseteq \nu\}$. By assumption and Fact 4.18, this remains possible, so we continue.

Notice that the threat of all possible localizations is what makes continuation possible. That is, the schema which says that "x is a 1-point extension of A" simply says that x remains (along with witnesses for X_k) in each of an infinite set of localizations of P_1 with parameters from A. If this schema is inconsistent, there will be a localization contradicting the hypothesis. □

We can now characterize simplicity in terms of persistence:

Theorem 5.22. *Let φ be a formula of T and $\langle P_n \rangle$ its characteristic sequence.*

(1) *If the localization φ^f of φ is simple, then for each P_∞-graph $A \subset P_1^f$ and for each $n < \omega$, there exists a localization $P_1^{f_n} \supset A$ of P_1^f in which there is a uniform finite bound on the size of a P_n-empty graph, i.e. there exists m_n such that $X \subset P_1^f$ and $X^n \cap P_n = \emptyset$ implies $|X| \leq m_n$.*
(2) *If localization φ^g of φ is not simple, then for all but finitely many $r < \omega$, P_1^g contains an infinite $(r+1)$-empty graph.*

In other words, the following are equivalent for any positive base set A:

(i) *There exists a localization φ^f of φ (with $\varphi^f = \varphi$ possible) such that φ^f is simple and $P_1^f \supset A$.*
(ii) *For each $n < \omega$, there exists a localization $P_1^{f_n} \supset A$ in which there is a uniform finite bound on the size of a P_n-empty graph.*

Proof. It suffices to show the first two statements. (1) is Lemma 5.21 applied to the formula φ^f. (2) is the second clause of Observation 5.20, where "almost all" means for r above k, the arity of dividing. □

Appendix: Some examples

This section works out several motivating examples. Recall that localization and persistence are described in Definitions 8 and 12 respectively, and that (η, ν)-arrays and trees appear in Definition 3.4.

Example 5.23. *(The random graph)*

T is the theory of the random graph, and R its binary edge relation. Let $\varphi(x; y, z) = xRy \wedge \neg xRz$, with $(T, \varphi) \mapsto \langle P_n \rangle$. Then:

- $P_1((y,z)) \iff y \neq z$.
- $P_n((y_1, z_1), \dots (y_n, z_n)) \iff \{y_1, \dots y_n\} \cap \{z_1, \dots z_n\} = \emptyset$.

Notice:

(1) The sequence has support 2.
(2) There is a uniform finite bound on the size of an empty graph $C \subset P_1, C^2 \cap P_2 = \emptyset$: an analysis of the theory shows that φ is not dividable, and inspection reveals this bound to be 3.
(3) P_n does not have the order property for any n and any partition of the $y_1, \dots y_n$ into object and parameter variables. (Proof: The order property in P_n implies dividability of φ_{2n} by Observation 4.11. But none of the φ_ℓ are dividable, as inconsistency only comes from equality.)
(4) Of course, the formula φ has the independence property in T. We can indeed find a configuration in P_2 which witnesses this: any C which models the T_0-configuration having $V_X = \omega$ and $\{i, j\} \notin E_X \iff \exists n(i = 2n \wedge j = 2n+1)$. Note that φ will have the independence property on any infinite P_2-complete subgraph of the so-called $(\omega, 2)$-array C (see Observation 5.9 below).
(5) As φ is unstable, φ-types are not necessarily definable in the sense of stability theory. However, we can obtain a kind of definability "modulo" the independence property, or more precisely, definability over the name for a maximal consistent subset of an $(\omega, 2)$-array as follows:

Definable types modulo independence. Let $p \in S(M)$ be a consistent partial φ-type presented as a positive base set $A \subset P_1$. Let us suppose $p \vdash \{xRc : c \in C\} \cup \{\neg xRd : d \in D\} \vdash p$, so that $A \subset M^2$ is a collection of pairs of the form (c, d) which generate the type.

There is no definable (in T with or without parameters, so in particular not from P_2) extension of the type A, so we cannot expect to find a localization of P_1 around A which is a P_2-complete graph. However:

Claim 5.24. *In the theory of the random graph, with $\varphi(x; y, z) = xRy \wedge \neg xRz$ as above, for any positive base set $A \subset P_1$ there exist a definable $(\omega, 2)$-array $W \subset P_1$, a solution S of W and an S-definable P_∞-graph containing A.*

Proof. Work in P_1. Fix any element (a, b) with $a, b \notin C, D$ and set $W_0 := \{(y, z) \in P_1 : \neg P_2((y, z), (a, b))\}$. Thus $W_0 = \{(b, z) : z \neq b\} \cup \{(y, a) : y \neq a\}$. So the only P_2-inconsistency among elements of W_0 comes from pairs of the form $(b, c), (c, a)$; thus, writing Greek letters for the elements of P_1,

$$(\forall \eta \in W_0)(\exists \nu \in W_0)(\forall \zeta \in W_0)\,(\neg P_2(\eta, \zeta) \rightarrow \zeta = \nu)$$

In other words, $W := W_0 \setminus \{(b, a)\}$ is an $(\omega, 2)$-array (Definition 3.4). Moreover:

(1) $(y, z), (w, v) \in W$ and $\neg P_2((y, z), (w, v))$ implies $y = v$ or $z = w$, and
(2) for any $c \neq a, b$, there are $d, e \in M$ such that $(d, c), (c, e) \in W$. Thus:
(3) we may choose a maximal complete P_2-subgraph C of W such that CA is a complete P_∞-graph. For instance, let C be any maximal complete extension of $\{(b, d) : d \in D\} \cup \{(c, a) : c \in C\}$. Call any such C a *solution* of the array W.

Let S be a new predicate which names this solution C of W. Then $\{y \in P_1 : z \in S \rightarrow P_2(y, z)\} \supset A$ is a P_2-complete graph, definable in $\mathcal{L} \cup \{S\}$. Support 2 implies that it is a P_∞-graph. Notice that by (2), we have in fact chosen a maximal consistent extension of A (i.e. a complete global type). ☐

Remark 5.25. *The idiosyncracies of this proof, e.g. the choice of a* definable *$(\omega, 2)$-array, reflect an interest in structure which will be preserved in ultrapowers.*

Example 5.26. *(Coding complexity into the sequence)*

It is often possible to choose a formula φ so that some particular configuration appears in its characteristic sequence. For instance, by applying the template below when φ has the independence property, we may choose a simple unstable θ whose P_2 is universal for finite bipartite graphs (X, Y), provided we do not specify whether or not edges hold between $x, x' \in X$ or between $y, y' \in Y$. Nonetheless, Conclusion 4.12 below will show this is "inessential" structure in the case of simple theories: whatever complexity was added through coding can be removed through localization.

The construction. Fix a formula φ of T. Let $\theta(x; y, z, w) := (z = w \wedge x = y) \vee (z \neq w \wedge \varphi(x; y))$. Write $(y, *)$ for (y, z, w) when $z = w$, and $(y, -)$ for (y, z, w) when $z \neq w$. Let $\langle P_n \rangle$ be the characteristic sequence of θ, $\langle P_n^\varphi \rangle$ be the characteristic sequence of φ, and $\langle P_n^= \rangle$ be the characteristic sequence of $x = y$. Then P_n can be described as follows:

- $P_n((y_1, -), \ldots (y_n, -)) \leftrightarrow P_n^\varphi(y_1, \ldots y_n)$.
- $P_n((y_1, *), \ldots (y_n, *)) \leftrightarrow P_n^=(y_1, \ldots y_n)$.

- Otherwise, the n-tuple $y := ((y_1, z_1, w_1), \dots (y_n, z_n, w_n))$ can contain (up to repetition) at most one $*$-pair, so $z_i = z_j = * \to y_i = y_j$. In this case the unique y_* in the $*$-pair is the realization of some φ-type in the original model M of T, and $P_{n+1}((y_*, *), (y_1, -), \dots (y_n, -))$ holds iff $M \models \bigwedge_{j \leq n} \varphi(y^*; y_j)$.

Remark 5.27. *This highlights an important distinction: the fact that a characteristic sequence may contain a bipartite graph is not as powerful as the fact of containing a random graph. See Example 5.30 below. In the coding just given we could not choose how elements within each side of the graph interrelated. This is quite restrictive, and eludes our coding for deep reasons. For instance, applying a consistency result of Shelah on the Keisler order one can show that the order property in the characteristic sequence cannot imply the* compatible *order property in the characteristic sequence, Definition 3.10. See also* [7].

Example 5.28. *(A theory with TP_2)*

TP_2 is Shelah's tree property of the second kind, to be defined and discussed in detail in Definition 3.3. Let T be the model completion of the following theory [12]. There are two infinite sorts X, Y and a single parametrized equivalence relation $E_x(y, z)$, where $x \in X$, and $y, z \in Y$. Let $\varphi_{eq} := \varphi(y; xzw) = E_x(y, z) \wedge \neg E_x(z, w)$. Then:

- $P_1((xzw)) \iff negE_x(z, w)$.
- $P_2((x_1 z_1 w_1), (x_2 z_2 w_2)) \iff$ each triple is in P_1 and furthermore:

$$(x_1 = x_2) \to (\ E_x(z_1, z_2) \wedge \bigwedge_{i \neq j \leq 2} \neg E_x(w_i, z_j))$$

The sequence has support 2. There are many empty graphs; these persist under localization (Theorem 5.22). One way to see the trace of TP_2 is as follows. Fixing α, choose a_i $(i < \omega)$ to be a set of representatives of equivalence classes in E_α, and choose b such that $\neg E_\alpha(a_i, b)$ $(i < \omega)$. Then $\{(\alpha, a_i, b) : i < \omega\} \subset P_1$ is a P_2-empty graph. We in fact have arrays $\{(\alpha^t, a_i^t, b^t) : i < \omega, t < \omega\}$ whose "columns" (fixing t) are P_2-empty graphs and where every path which chooses exactly one element from each column is a P_2-complete graph, thus a P_∞-complete graph. The parameters in this so-called (ω, ω)-array describe TP_2 for φ_{eq} (Claim 3.8).

Note that a gap has appeared between the classification-theoretic complexity of φ, which is not simple, and that of the formula P_2:

Claim 5.29. *P_2 does not have the order property.*

Proof. This is essentially because inconsistency requires the parameters x to coincide. Suppose that $\langle \gamma_s, \delta_s : s < \omega \rangle$ were a witness to the order property for P_2. Fix any $a_s = (\alpha_s, a_s, d_s)$. Now $\neg P_2(b_t, a_s)$ for $t < s$, where $b_t = (\beta_t, b_t, c_t)$. P_2-inconsistency requires $\alpha_s = \beta_t$. As this is uniformly true, $\alpha_s = \alpha_t = \beta_s = \beta_t$ for all $s, t < \omega$ in the sequence. But now that we are in a single equivalence relation E_α, transitivity effectively blocks order: $\neg P_2(\gamma_s, \delta_t) \leftrightarrow \neg E_\alpha(a_s, b_t)$. Depending on whether at least one of the a- or b-sequences is an empty graph, we can find a contradiction to the order property with either three or four elements. ☐

Example 5.30. *(A maximally complicated theory)*

In this example the sequence is universal for finite T_0-configurations (Convention 3.1), a natural sufficient condition for "maximal complexity."

Let the elements of M be all finite subsets of ω; the language has two binary relations, $\subseteq$ and $=$, with the natural interpretation. Set $T = Th(M)$.

Choose $\varphi_\subseteq := \varphi(x; y, z) = x \subseteq y \wedge x \not\subseteq z$. Then:

- $P_1((y, z)) \iff \emptyset \subsetneq y \not\subseteq z$.
- $P_n((y_1, z_1), \ldots (y_n, z_n)) \iff \emptyset \subsetneq \bigcap_{i \leq n} y_i \not\subseteq \bigcup_{i \leq n} z_i$.

The sequence does not have finite support. Moreover:

Claim 5.31. *Let $\langle P_n \rangle$ be the characteristic sequence of $\varphi_\subseteq$, $k < \omega$, and let X be a finite T_0-configuration. Then there exists a finite $A \subseteq P_1$ witnessing X.*

Proof. Write the elements of P_1 as $w_i = (y_i, z_i)$; it suffices to choose the positive pieces y_i first, and afterwards take the z_i to be completely disjoint. More precisely, suppose X is given by $V_X = m$ and $E_X \subset \mathcal{P}(m)$. We need simply to choose $y_1, \ldots y_m$ such that for all $\sigma \subseteq m$,

$$\left(\bigcap_{j \in \sigma} y_j \neq \emptyset \right) \iff \sigma \in E_X$$

which again, is possible by the downward closure of E_X. □

Corollary 5.32. *This characteristic sequence is universal for finite T_0-configurations.*

Remark 5.33. *That the sequence is universal for finite T_0-configurations is sufficient, though not necessary, for maximal complexity in the Keisler order. By [8]. VI.3, $\varphi(x; y, z) = y < x < z$ in $Th(\mathbb{Q}, <)$ is maximal. Its characteristic sequence has support 2, but its P_2 is clearly not universal.*

References

[1] Buechler, "Lascar Strong Types in Some Simple Theories." Journal of Symbolic Logic 64(2) (1999), 817–824.

[2] Džamonja and Shelah, "On $\triangleleft^*$-maximality." Annals of Pure and Applied Logic 125 (2004) 119–158.

[3] Hrushovski, Peterzil, and Pillay, "Groups, measures, and the NIP." J. Amer. Math. Soc. 21 (2008), no. 2, 563–596.

[4] Keisler, "Ultraproducts which are not saturated." Journal of Symbolic Logic, 32 (1967) 23–46.

[5] Kim and Pillay, "From stability to simplicity." Bull. Symbolic Logic 4 (1998), no. 1, 17–36.

[6] Malliaris, "Realization of φ-types and Keisler's order." Annals of Pure and Applied Logic 157 (2009) 220–224.

[7] Malliaris, "Persistence and Regularity in Unstable Model Theory." Ph.D. thesis, University of California, Berkeley, 2009.

[8] Shelah, *Classification Theory and the number of non-isomorphic models*, rev. ed. North-Holland, 1990.

[9] Shelah, "Toward classifying unstable theories." Annals of Pure and Applied Logic 80 (1996) 229–255.

[10] Shelah, "Classification theory for elementary classes with the dependence property – a modest beginning." Scientiae Mathematicae Japonicae, 59 (2004) 265–316, Special issue on set theory and algebraic model theory.

[11] Shelah, "Dependent theories and the generic pair conjecture." (2008) arXiv:math.LO/0702292.

[12] Shelah and Usvyatsov, "More on SOP_1 and SOP_2." Annals of Pure and Applied Logic 155 (2008), no. 1, 16–31.

[13] Usvyatsov, "On generically stable types in dependent theories." Journal of Symbolic Logic 74 (2009), 216-250.

[14] Wagner, *Simple theories*. Mathematics and Its Applications 503. Kluwer Academic Publishers, 2000.

Group in Logic, University of California at Berkeley, 910 Evans Hall, Berkeley, CA 94720
E-mail address: mem@math.berkeley.edu