A STUDY IN COMBINATORICS AND MODEL THEORY

ESME BAJO

ABSTRACT. We explore graphs with large homogeneous subgraphs via the Erdős-Hajnal conjecture, the proposition that for any graph H, there is a constant $\delta(H) > 0$ such that every H-free graph on n vertices has a homogeneous subgraph of size $n^{\delta(H)}$. We focus on the conjecture for halfgraphs, and then transition to a model-theoretic generalization of the result. We show an analog to Ramsey's theorem and a combinatorial instantiation of the theorem that stable theories have large indiscernible sets, where stable theories extend halfgraph-free graphs and indiscernible sets extend homogeneous subgraphs.

Contents

1. Introduction

This paper examines the interactions between finite combinatorics and model theory. We start with a preliminary discussion of Ramsey's Theorem. Our use of type trees in the proof will motivate an analysis of graphs that admit type trees that imply the existence of especially large homogeneous subgraphs. In particular, we explore graphs satisfying the following conjecture, made by Paul Erdős and András Hajnal in 1989:

Conjecture 1.1 (Erdős-Hajnal Conjecture)**.** *For every graph H, there is a constant $\delta(H) > 0$ such that every graph G on n vertices with no induced subgraph isomorphic to H has a homogeneous set of size at least $n^{\delta(H)}$.*

For reference, Maria Chudnovsky's survey provides an excellent breakdown of the graphs for which the conjecture has been formally proven [1]. We discuss the result for some small graphs, complete and empty graphs, and halfgraphs.

Throughout this paper, we will move freely between two notions of a graph: the combinatorial notion (as used implicitly in the Erdős-Hajnal Conjecture above) and the model-theoretic notion. We reconcile the two now:

Definition 1.2. A (combinatorial) *graph* G is a pair (V, E), where V is a set of vertices and $E \subseteq \binom{V}{2} = \{\{u, v\} : u, v \in V \text{ and } u \neq v\}$ defines an adjacency relation on V. We say $u, v \in V$ are adjacent, denoted $u \sim v$, precisely when $\{u, v\} \in E$.

Definition 1.3. A (model-theoretic) *graph* is a model M equipped with a binary relation R. We say that elements $u, v \in M$ are adjacent precisely when $R(u, v)$.

A combinatorial graph $G = (V, E)$ can be identified with a model $M = (V; R)$, where $R(u, v) \iff \{u, v\} \in E$. This model-theoretic notion is emphasized at the end of the paper, where we generalize the notion of a homogeneous subgraph to the notion of an indiscernible set in a model, generalize the notion of being halfgraph-free to the notion of stability, and work towards the following result:

Theorem 1.4. *If a theory T is stable in λ, $|I| > \lambda \geqslant |A|$, then there is $J \subseteq I$, $|J| > \lambda$, which is an indiscernible sequence over A.*

Section 6 assumes some knowledge of basic model theory, covered in [3]. However, Section 7 attempts to present the main result of this section (Theorem 1.4) in a purely combinatorial manner, assuming no model theory.

2. Ramsey Theory and Type Trees

We first introduce Ramsey theory as a tool for measuring the large homogeneous subgraph we can find in *any* large graph. For both infinite and finite graphs, we want to be explicit about the notion of largeness we wish to maintain. In the infinite case, we turn to ultrafilters for this notion.

Definition 2.1. An *ultrafilter* $\mathcal{D}$ on a set X is a subset of the power set of X such that:
(i) $X \in \mathcal{D}$ and $\varnothing \notin \mathcal{D}$;
(ii) $\mathcal{D}$ is upward closed, i.e. $A \in \mathcal{D}$ and $A \subseteq B \implies B \in \mathcal{D}$;
(iii) $\mathcal{D}$ is closed under finite intersections;
(iv) For any $A \subseteq X$, either $A \in \mathcal{D}$ or $X \backslash A \in \mathcal{D}$.

We say that an ultrafilter $\mathcal{D}$ on X is *principal* if it is generated by a single $x \in X$, i.e. if $\mathcal{D} = \{A \subseteq X : x \in A\}$, and we say that $\mathcal{D}$ is *nonprincipal* otherwise. Note that nonprincipal ultrafilters on X contain all cofinite subsets of X. We use the existence of nonprincipal ultrafilters on $\mathbb{N}$ without proof.

We will also use the fact that given an ultrafilter $\mathcal{D}$ on X and a finite partition $X = X_1 \sqcup X_2 \sqcup ... \sqcup X_k$, at least one of the X_i is in $\mathcal{D}$. We omit a rigorous proof for brevity, but it is true by induction on condition (iv). At this point, we are prepared to prove the infinite version of Ramsey's Theorem.

Theorem 2.2 (Ramsey's Theorem, Infinite Version)**.** *Let V be an infinite set and let $k \geqslant 2$. Every k-coloring of the edges of the complete graph on V has an infinite monochromatic subgraph.*

Proof. Fix $k \geqslant 2$. Reduce V to a countably infinite subset $V' \subseteq V$, and for convenience of notation, let $V = V'$. Denote the edge set of the complete graph on V by E. Let $C : E \to \{c_1, ..., c_k\}$ be a k-coloring of the edges and fix a non-principal ultrafilter $\mathcal{D}$ on V.

We define a function $\varphi : V \to \{c_1, ..., c_k\}$ as follows. For each $v \in V$, partition $V \backslash \{v\} = S_1 \sqcup ... \sqcup S_k$, where $u \in S_i$ if and only if $C(\{u, v\}) = c_i$. Exactly one of $S_1, ..., S_k$, say S_j, is in $\mathcal{D}$. Let $\varphi(v) = c_j$.

Now let $V = T_1 \sqcup ... \sqcup T_k$ be the partition of V such that $v \in T_i$ if and only if $\varphi(v) = c_i$. Exactly one of $T_1, ..., T_k$, say T_j, is in $\mathcal{D}$. We claim that V has an infinite monochromatic subset $W \subseteq V$ such that $\binom{W}{2}$ is monochromatic.

We construct such a W now. Pick any $v_0 \in T_j$ and add v_0 to W. Denote the set of vertices u such that $\{v_0, u\}$ is colored with c_j by A_0. Since $v_0 \in T_j$, we have that $A_0 \in \mathcal{D}$. Ultrafilters are closed under finite intersections, so $T_j \cap A_0 \in \mathcal{D}$. So we can pick $v_1 \in T_j \cap A_0$ and add v_1 to W. Next denote the set of vertices u such that $\{v_1, u\}$ is colored with c_j by A_1. We have $A_1 \in \mathcal{U}$ since $v_1 \in T_j$, and so $T_j \cap A_0 \cap A_1 \in \mathcal{D}$. Proceed inductively. Suppose we have added $v_0, v_1, ..., v_n$ to W and chosen $A_0, A_1, ..., A_n \in \mathcal{D}$ such that $T_j \cap A_0 \cap A_1 \cap ... \cap A_n \in \mathcal{D}$. Pick $v_{n+1} \in T_j \cap A_0 \cap A_1 \cap ... \cap A_n$ and add it to W. Let A_{n+1} be the set of vertices u such that $\{v_{n+1}, u\}$ is colored with c_j. Then $T_j \cap A_0 \cap ... \cap A_n \cap A_{n+1} \in \mathcal{D}$, and the induction can continue at every finite step. $\square$

The finite version follows from the infinite version by Łoś's Theorem, detailed in [3], but we instead suggest a proof analogous to that of Theorem 2.2, as it will provide us with an upper bound on the smallest integer m such that any graph on m vertices necessarily has a homogeneous subgraph of a specified size. We first develop a parallel notion of largeness:

Definition 2.3. For any $k \geqslant 1$, we say that $A \subseteq X$, $|X| < \infty$, is *k-large* if

$$|A| \geqslant \left\lceil \frac{|X|}{k} \right\rceil.$$

Note that for any choice of $k \geqslant 1$ and any partition of a finite set X into k disjoint blocks, at least one block must be k-large.

Theorem 2.4 (Ramsey's Theorem, Finite Version). *For every $n, k \in \mathbb{N}$, there exists $m \in \mathbb{N}$ such that every k-coloring of the edges of the complete graph K_m on m vertices has a monochromatic subgraph of size n.*

Proof. This proof is nearly identical to the infinite proof via ultrafilters, so we sketch it for brevity. Fix any k and any n, and start with any k-coloring C of a complete graph on at least k^{kn} vertices. Pick any vertex v_1. The k-coloring of the edges incident with v_1 partition the other vertices. Some color must induce a k-large block, from which we can pick another vertex and proceed inductively. By k-largeness, we will be able to proceed for $\geqslant kn$ steps. At this point, we have vertices $v_1, ..., v_{kn}$ and colors $c_1, ..., c_{kn-1}$ such that

$$(\forall i \leqslant kn - 1)(\forall j > i)(C(\{v_i, v_j\}) = c_i).$$

By k-largeness again, there is some color that appears in our list at least n times, and the corresponding vertices form a monochromatic subgraph of that color. $\square$

We now provide an alternate proof of the finite version. It is ultimately a restructuring of the proof of Theorem 2.4, but it provides context for an important concept: a type tree. In our definition of a type tree, we use the following notation. For binary strings p, q, and r, we say that $p = q^\frown r$ if p is the concatenation of q and r (in that order), and we say that $p \succ q$ if $p = q^\frown r$ for some r.

Definition 2.5. A *type tree* of a graph $G = (V, E)$ is a partial binary tree T, the nodes of which are labeled with paths to them, and a bijection $\varphi : V \to N(T)$, where $N(T) \subseteq 2^{<\omega}$ is the set of labels of the nodes of T, such that:

(i) for all $u, v \in V$ such that $\varphi(u) \succ \varphi(v)^\frown 0$, u and v are not adjacent in G, and
(ii) for all $u, v \in V$ such that $\varphi(u) \succ \varphi(v)^\frown 1$, u and v are adjacent in G.

We will usually just refer to a type tree of $G = (V, E)$ as a tree labeled with the elements of V, where the function from the vertices to the paths of V is implied.

Example 2.6. We provide an example of a type tree $T(B)$ of the bull graph B. Note that this tree is the unique type tree created by adding the vertices of B in the order c, b, e, a, d. We place c at the base, add b to its right since $b \sim c$, add e to the right of c since $e \sim c$ and to the left of b since $e \not\sim b$, etc.

$B:$

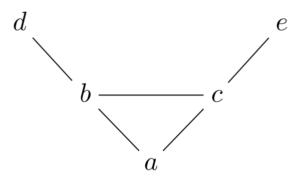

$T(B):$

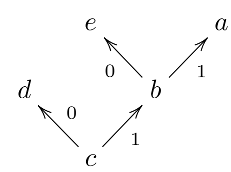

We associate to c to the empty string $\langle\rangle$, to b the string 1, to e the string 10, etc.

We prove only the finite version of Ramsey's Theorem for two colors using type trees, but the proof does extend to any finite number of colors and to hypergraphs.

Theorem 2.7 (Ramsey's Theorem, Finite Version for $k = 2$)**.** *For every $n \in \mathbb{N}$, there exists $m \in \mathbb{N}$ such that every graph on $\geqslant m$ vertices contains a homogeneous subgraph of size n.*

Proof. Let $[N] = \{1, .., N\}$. For any graph $G = ([N], E)$, we construct a type tree $T = T(G)$ with nodes $1, ..., N$ in stages. At the end of stage $s \leqslant N$, we denote the tree constructed so far by $T[s]$ and we let $T = T[N]$. Let $T[0]$ be the tree with the single node labeled 1. Suppose that at the beginning of stage s, $1 < s \leqslant N$, we have constructed the tree $T[s-1]$ with nodes labeled $1, ..., s-1$. At stage s we add a node labeled s to the tree.

Note that every finite path through the partial type tree $T[s-1]$ can be represented by a binary string $\sigma = x_1 ... x_r$, where $x_i = 0$ if σ branches to the left at level i and $x_i = 1$ if σ branches to the right at level i. There is a unique binary string $\sigma = x_1 .... x_{r-1}$, passing through vertices $1 = v_1, v_2, ..., v_r$ (in that order), such that for every i, $1 \leqslant i \leqslant r-1$, $x_i = 0$ if $\{s, v_i\} \notin E$ and $x_i = 1$ if $\{s, v_i\} \in E$ and such that σ is of maximal length.

The existence of such a σ is evident as we can proceed along a branch, going to the left if s is not adjacent to the current node and to the right otherwise, until we reach a point at which no node is present in the direction we wish to proceed. It must be unique since if there are two such strings σ_1, σ_2, they must diverge on some bit x_i but we cannot have both $\{s, v_i\} \in E$ and $\{s, v_i\} \notin E$. Now we can complete the construction of $T[s]$. From the node labeled v_r, place s to the left of v_r if $\{s, v_r\} \notin E$ and place s to the right of v_r if $\{s, v_r\} \in E$.

By induction we can construct $T = T[N]$. We claim that T must have a path of length at least $\log_2 N$. Otherwise, T has at most $\log_2(N) - 1$ levels, and thus the number of nodes in T is at most

$$2^0 + 2^1 + 2^2 + ... + 2^{\log_2(N)-2} = 2^{\log_2(N)-1} - 1 < N,$$

a contradiction. Let this path be given by the binary string $\tau = x_1 .... x_\ell$, $\ell = \log_2 N$. At least $\ell/2$ of the x_i are 0s or at least $\ell/2$ of the x_i are 1s. Suppose $\ell/2$ of the x_i

are 0s. We claim that the corresponding vertices on the path form an anticlique. Denote the $\ell/2$ vertices by $v_1, ..., v_{\ell/2}$. Then $v_2, ..., v_{\ell/2}$ are all to the left of v_1, and there are no edges in the subgraph generated by $v_1, ..., v_{\ell/2}$ that are incident with v_1. Similarly, for every i, $v_i, ..., v_{\ell/2}$ are to the left of v_{i-1} and there are no edges in the subgraph that are incident with v_{i-1}. Hence $v_1, ..., v_{\ell/2}$ form an anticlique. Similarly, if at least $\ell/2$ of the x_i are 1s, we get a clique in the subgraph on the corresponding vertices v_i.

Therefore if we pick $m \geqslant 2^{2n}$, any graph on m vertices yields a homogeneous subset of size at least $\frac{\log_2 m}{2} \geqslant \frac{2n}{2} = n$. ☐

3. The Erdős-Hajnal Conjecture: K_n and $\overline{K_n}$

Ramsey's Theorem yields homogeneous sets larger than $c \cdot \log N$ for some $c > 0$ in *any* graph on N vertices, but Erdős and Hajnal conjectured that there are much larger homogeneous subgraphs in graphs that are H-free for any graph H.

Definition 3.1. Let G be a graph and let $W \subseteq V(G)$. The *induced subgraph* of G on W, denoted $G{\restriction}_W$, is the graph with vertex set W and edge set

$$\{\{u, v\} : \{u, v\} \in E(G) \text{ and } u, v \in W\}.$$

Definition 3.2. Let H be any graph. The graph G is H*-free* if for every $W \subseteq V(G)$, $G{\restriction}_W$ is not isomorphic to H.

Erdős and Hajnal proposed the following conjecture, which remains open:

Conjecture 3.3. *For every graph H, there is a constant $\delta(H) > 0$ such that every graph G with no induced subgraph isomorphic to H has a homogeneous set of size at least $|V(G)|^{\delta(H)}$.*

Definition 3.4. We say that a graph H satisfies the *Erdős-Hajnal property* if there is a constant $\delta(H) > 0$ such that every H-free graph on n vertices has a homogeneous set of size $n^{\delta(H)}$.

For every n, the complete graph K_n on n vertices and the empty graph $\overline{K_n}$ on n vertices satisfy the Erdős-Hajnal property. To see this, we define the *nondiagonal* Ramsey numbers. Let $R(s, t)$ be the least number such that any graph on $R(s, t)$ vertices contains either an induced subgraph isomorphic to K_s or an induced subgraph isomorphic to $\overline{K_t}$.

Proposition 3.5.

$$R(s, t) \leqslant \min\{s^{t-1},\ t^{s-1}\}.$$

Proof. It follows by induction and Pascal's identity that $R(s, t) \leqslant \binom{s+t-2}{s-1}$ [6]. Then note that

$$\binom{s+t-2}{s-1} = \frac{(s+t-2)(s+t-3)...(t)}{(s-1)(s-2)...(1)} = \prod_{i=0}^{s-2} \frac{(s-1-i)+(t-1)}{s-1-i} \leqslant t^{s-1}.$$

Similarly, we have that $R(s, t) \leqslant s^{t-1}$. If a graph G on n vertices were K_s-free and $\overline{K_{n^{1/(s-1)}}}$-free, then $R(s, n^{1/(s-1)}) > n$. This would imply $(n^{1/(s-1)})^{s-1} > n$, a contradiction. Hence, $\delta(K_s) = 1/(s-1)$ suffices, as does $\delta(\overline{K_t}) = 1/(t-1)$. ☐

4. The Erdős-Hajnal Conjecture: Some Small Graphs

In this section, we look at some small graphs with the Erdős-Hajnal property. In fact, we define a substitution operation under which the Erdős-Hajnal property is invariant, which will allow us to build several graphs with the property.

Denote the class of all graphs on n vertices by G_n. It is immediate that all graphs in G_2 satisfy the Erdős-Hajnal property, and we now show that all graphs in G_3 do too. In the previous section, we established that K_3 and $\overline{K_3}$ do. We are left with P_3 (the path of length two) and its complement.

Lemma 4.1. *P_3 satisfies the Erdős-Hajnal property.*

Proof. Let G be a P_3-free graph on n vertices. If no two vertices in G are adjacent, then we have an empty graph of size n and we are done. So assume there exist $u, v \in V(G)$ that are adjacent. For every $w \in V(G)$ such that $w \neq u, v$, we have that $G\restriction_{\{u,v,w\}}$ cannot be isomorphic to P_3. Thus, either w must be connected to both u and v or to neither. Hence, G is the disjoint union of cliques.

Partition its vertex set $V(G) = C_1 \sqcup ... \sqcup C_k$, where $G\restriction_{C_i}$ is a nonempty clique for every i, $1 \leqslant i \leqslant k$. We claim that G has a homogeneous subgraph of size $n^{1/2}$. If some C_i has size $\geqslant n^{1/2}$, we are done. Otherwise,

$$n = |C_1| + ... + |C_k| < kn^{1/2} \implies k > n^{1/2}.$$

Take any sequence of vertices $v_1 \in C_1$, $v_2 \in C_2$,..., $v_k \in C_k$. The graph $G\restriction_{\{v_1,...,v_k\}}$ is an empty graph of size greater than $n^{1/2}$. $\square$

Now suppose G does not contain $\overline{P_3}$. Then its complement G^c does not contain P_3 and has a homogeneous subgraph of size $|V(G)|^{1/2}$ by Lemma 4.1. Then G^c does as well, so we have now shown that all graphs in G_3 satisfy the Erdős-Hajnal property. To build up further, we resort to the following operation from [7]:

Definition 4.2. Let H_1 and H_2 be graphs on disjoint vertex sets V_1 and V_2, respectively. Suppose $|V_1|, |V_2| \geqslant 2$ and let $v_1 \in V_1$. We say that the graph H, with vertex set V, is *obtained from H_1 by substituting H_2 for v_1* if the following conditions hold:
(i) $V = (V_1 \cup V_2)\backslash\{v_1\}$;
(ii) $H\restriction_{V_2} = H_2$;
(iii) $H\restriction_{V_1\backslash\{v_1\}} = H_1\restriction_{V_1\backslash\{v_1\}}$;
(iv) $u \in V_1$ is adjacent in H to $w \in V_2$ if and only if u is adjacent in H_1 to v.

Example 4.3. Let v be the sole vertex of degree 2 in P_3. The graph H shown below is obtained from P_3 by substituting C_2 for v:

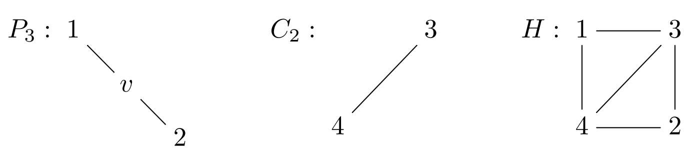

Theorem 4.4. *Suppose that graphs H_1 and H_2 with vertex sets V_1 and V_2, respectively, satisfy the Erdős-Hajnal property. If H is obtained from H_1 by substituting H_2 for $v_1 \in V_1$, then H satisfies the Erdős-Hajnal property.*

Proof. Let G be any H-free graph, and suppose G is a graph on N vertices. Assume for contradiction that G does not contain a large homogeneous subgraph. Informally, we use the Erdős-Hajnal property of H_1 to get many more copies of H_1 in

G than there are embeddings $H_1 \backslash \{v_1\} \hookrightarrow G$. Thus we get a large set of preimages of v_1. This set must be H_2-free since G is H-free, so the Erdős-Hajnal property of H_2 gives us a large homogeneous subset of G.

We formalize this as follows. By assumption, H_1 and H_2 have the Erdős-Hajnal property, so there are constants $\delta(H_1)$ and $\delta(H_2)$ such that any H_1-free graph on n vertices has a homogeneous subgraph of size $n^{\delta(H_1)}$ and any H_2-free graph on n vertices has a homogeneous subgraph of size $n^{\delta(H_2)}$. Suppose that G does not have a homogeneous set of size $N^{c\delta(H_1)}$, for some constant c dependent on H that we define later. We attempt to get a contradiction.

Pick $m = \lceil N^c \rceil$. Any m-element subset of $V(G)$ must contain some induced subgraph isomorphic to H_1, or it would necessarily contain a homogeneous subgraph of size $m^{\delta(H_1)} \geqslant N^{c\delta(H_1)}$, a contradiction. But each copy of H_1 can contribute to $\binom{N-|V_1|}{m-|V_1|}$ m-subsets of $V(G)$, so we get a total of $\binom{N}{m}/\binom{N-|V_1|}{m-|V_1|}$ distinct copies of H_1 that are necessarily in G.

However, there are no more than $N(N-1)...(N-|V_1|+2)$ embeddings of $H_1{\restriction}_{V_1\backslash\{v_1\}}$ in G, so by the Pigeonhole Principle, there is some embedding $f : H_1{\restriction}_{V_1\backslash\{v_1\}} \hookrightarrow G$ that can be extended to an embedding of H_1 in M distinct ways, where

$$M = \frac{\binom{N}{m}/\binom{N-|V_1|}{m-|V_1|}}{N(N-1)...(N-|V_1|+2)}.$$

That is, if we let $V_1 = \{v_1, ..., v_{|V_1|}\}$, then there exist $U \subseteq |V(G)|$, $|U| = M$, and $w_2, ..., w_{|V_1|} \in V(G)$ such that for every $u \in U$, the map $f_u : H_1 \hookrightarrow G$ given by

$$f_u(v_1) = u \text{ and } f(v_i) = w_i \text{ for all } i, 2 \leqslant i \leqslant |V_1|$$

is an isomorphic embedding.

Since G is H-free, the induced subgraph $G{\restriction}_U$ must be H_2-free. Thus, $G{\restriction}_U$ contains a homogeneous subgraph of size

$$\begin{aligned} M^{\delta(H_2)} = \left(\frac{\frac{N!}{(N-m)!m!} \cdot \frac{(m-|V_1|)!(N-m)!}{(N-|V_1|)!}}{N(N-1)...(N-|V_1|+2)}\right)^{\delta(H_2)} &= \left(\frac{N-|V_1|+1}{m(m-1)...(m-|V_1|+1)}\right)^{\delta(H_2)} \\ &> (N^{1-c|V_1|})^{\delta(H_2)} \\ &= N^{\delta(H_2)-c|V_1|\delta(H_2)} \end{aligned}$$

By assumption, G does not have a homogeneous subgraph of size $N^{c\delta(H_1)}$, so if we pick c such that $\delta(H_2) - c|V_1|\delta(H_2) > c\delta(H_1)$, we get a contradiction and are done. So just choose c so that

$$c < \frac{\delta(H_2)}{\delta(H_1) + |V_1|\delta(H_2)}.$$

□

Definition 4.5. We say that a graph G is *prime* if it cannot be obtained from two smaller graphs by substitution.

Thus, to show that all graphs in G_4 satisfy the Erdős-Hajnal property, we need only show that all prime graphs on four vertices do. As seen in Example 4.3, K_4 missing one edge is not prime, and similarly, K_4 can be obtained from P_3 and $\overline{K_2}$. Therefore, the only prime graph on four vertices is P_4 (and it is self-complementary so we are not immediately done). We omit a proof that P_4 does indeed satisfy the Erdős-Hajnal property. In order to show that all graphs in G_5 have the EH-property, we again need only check the prime ones. They are $C_5, P_5, \overline{P_5}$, and the bull graph (pictured in Example 2.6). A proof that the bull graph has the Erdős-Hajnal property is in [8], but the conjecture is open even for these three other small graphs.

While the conjecture has not been proved for P_5, in 2014, Bousquet, Lagoutte, and Thomassé proved the weaker statement that for any n, graphs containing neither P_n nor $\overline{P_n}$ have large homogeneous subgraphs [5].

5. The Erdős-Hajnal Conjecture: Halfgraphs

We now turn to type trees to prove the Erdős-Hajnal conjecture for some larger classes of graphs, including the class of all finite halfgraphs.

Definition 5.1. A *k-halfgraph* is a graph on vertices $a_1, ..., a_k$ and $b_1, ..., b_k$ such that

$$a_i \sim b_j \iff i < j$$

Definition 5.2. A graph is *k-stable* if it is k-halfgraph-free.

Note that the definition of a halfgraph imposes no restrictions on the edges $\{a_i, a_j\}$ and the edges $\{b_i, b_j\}$, so a k-stable graph excludes an entire family of induced subgraphs.

In the proof of Theorem 2.7, we saw one method to ensure that a graph G has a large homogeneous subgraph: ensure that a type tree of G has a long branch. We do this for k-stable graphs in two parts. First, we show that k-stable graphs have strict upper bounds on the level to which a type tree of that graph is full. Next, we show that if the fullness is bounded, then some branch in that tree is long. To formalize this, we introduce some notation regarding type trees [2]:

Definition 5.3. Let G be a finite graph. The *tree rank* $t(G)$ is the largest integer t such that there exists some subgraph G' of G and a type tree $2^{<t}$ of G'.

Example 5.4. The graph C_4 (the cycle of length 4, pictured below) has tree rank 2. It does not have enough vertices for each of the 7 nodes of the graph $2^{<3}$, so $t(C_4) \leqslant 2$. And we can arrange the subgraph $C_4{\restriction}_{\{A,C,D\}}$ into a type tree $2^{<2}$ as follows:

$$C_4: \quad \begin{tikzcd} A \arrow[r, dash] \arrow[d, dash] & B \arrow[d, dash] \\ D \arrow[r, dash] & C \end{tikzcd} \qquad\qquad T(C_4): \quad \begin{tikzcd} C & & D \\ & A \arrow[ul, "0"'] \arrow[ur, "1"] & \end{tikzcd}$$

Definition 5.5. Let G be a finite graph. The *tree height* $h(G)$ is the largest integer h such that there exists a branch of length h in some type tree of G.

Example 5.6. For every n, the complete graph K_n and the empty graph $\overline{K_n}$ both have tree height n. Every type tree of K_n has the branch $\underbrace{11\ldots1}_{n \text{ times}}$ and every type tree of $\overline{K_n}$ has the branch $\underbrace{00\ldots0}_{n \text{ times}}$, so $h(K_n) = h(\overline{K_n}) = n$.

We sketch a proof that k-stability forces an upper bound $2^{k+2}-2$ on the tree rank of a graph. If there is some full type tree T of G, we inductively pull out nodes v_i, w_i and decrease the rank of of the tree k times and maintaining $v_i \sim w_j \iff i < j$. In order to do so, we use a lemma: if the nodes of a type tree T of rank $n+m$ are partitioned into two sets A, B, then either A has rank $\geqslant n$ or B has rank $\geqslant m$. This is done by induction on $n+m$. Both proofs are made explicit for the general notion of an order property in Chapter 6 of [11].

Next, we show that the tree rank and the tree height cannot both be small. In particular:

Lemma 5.7. *Suppose G is a graph on $n \geqslant 2$ vertices. Let h be the tree height of G and let t be the tree rank of G. Then $h \geqslant (n/t)^{\frac{1}{t+1}}/2$.*

Proof. Fix a type tree $T(G) = \{v_\sigma : \sigma \in 2^{<\omega}\}$ of G (where the v_σ are the vertices $V = V(G)$ of G, labeled to correspond to the nodes of the tree). Define the rank of a vertex v_σ to be the largest r such that there is a full type tree of height r above v_σ, and denote this quantity by $r(v_\sigma)$. By a full type tree of height r above v_σ, we mean a set $W \subseteq \{v_\tau : \tau \geq \sigma\}$, $|W| = 2^r - 1$, such that there is a type tree $2^{<r}$ on $G\restriction_W$. Note that while this could look like a full binary tree originating at that vertex, it does not have to; it may be a tree above that vertex skipping a level or a complete reordering of the vertices succeeding that vertex.

Let L be the length of the longest branch in $T(G)$ and let $R = \max\{r(v_\sigma) : v \in V\}$. Now let $Z_\ell^r = \{v_\sigma \in V : r(v_\sigma) = r \text{ and } |\sigma| = \ell\}$. That is, let Z_ℓ^r be the set of vertices at level ℓ in $T(G)$ that have vertex rank r. Note that

$$n = \sum_{\ell=0}^{L}\sum_{r=0}^{R} |Z_\ell^r|. \tag{0}$$

For every r, ℓ, split Z_ℓ^r into the set X_ℓ^r of vertices in Z_ℓ^r whose immediate predecessors have vertex rank r and the set Y_ℓ^r of vertices in Z_ℓ^r whose immediate predecessors have vertex rank $r+1$. Then

$$|Z_\ell^r| = |X_\ell^r| + |Y_\ell^r|. \tag{1}$$

If some vertex $v_{\sigma^\frown i}$ and its predecessor v_σ both have vertex rank r, then the vertex $v_{\sigma^\frown(1-i)}$ (if it exists) must have rank strictly less than r (or v_σ would then have rank $r+1$). Thus,

$$|X_{\ell+1}^r| \leqslant |Z_\ell^r|. \tag{2}$$

And since each vertex has no more than two immediate successors, we get the bound

$$|Y_{\ell+1}^r| \leqslant 2|Z_\ell^{r+1}|. \tag{3}$$

Combining (1), (2), and (3), we get that:

$$|Z_{\ell+1}^r| = |X_{\ell+1}^r| + |Y_{\ell+1}^r| \leqslant |Z_\ell^r| + 2|Z_\ell^{r+1}|. \tag{4}$$

We use (4) to show by induction that for every r, $0 \leqslant r < R$, and for every ℓ, $0 \leqslant \ell < L$,

$$|Z_{\ell+1}^{R-r}| \leqslant (2(\ell+1))^r. \tag{**}$$

We split this into three cases: $r = 0$, $r = 1$, and $r > 1$. First, consider the case where $r = 0$. If $|Z_{\ell+1}^R| > 1$ for some ℓ, then there is a full type tree above $v_{\langle\rangle}$ with height $R+1$, which is a contradiction. Thus, (**) necessarily holds for $r = 0$.

Next suppose that $r = 1$, and induct on ℓ. The base case $\ell = 0$ is immediate since the total number of vertices at the first level of T is bounded by 2. Suppose that (**) is true for $r = 1$ and for all $k < \ell$ for some $\ell \geqslant 1$. That is, suppose for all $k < \ell$ that $|Z_{k+1}^{R-1}| \leqslant 2(k+1)$. By (4), $|Z_{\ell+1}^{R-1}| \leqslant |Z_{\ell}^{R-1}| + 2|Z_{\ell}^{R}|$. So by our inductive assumption and the $r = 0$ case, $|Z_{\ell+1}^{R-1}| \leqslant 2\ell + 2 = 2(\ell+1)$.

Finally, we show that (**) is true for any $r > 1$. Assume inductively that for some $r > 1$ and for all $q < r$, (**) holds for all ℓ, $0 \leqslant \ell < L$, i.e. that $|Z_{\ell+1}^{R-q}| \leqslant (2(\ell+1))^q$. We show that (**) holds for r and for all ℓ, $0 \leqslant \ell < L$, by another induction, this time on ℓ. The base case $\ell = 0$ is immediate by the same reasoning as in the $r = 1$ case. Assume now that for some $\ell > 0$ and for all $k < \ell$, we have $|Z_{k+1}^{R-r}| \leqslant (2(k+1))^r$. By (4), $|Z_{\ell+1}^{R-r}| \leqslant |Z_{\ell}^{R-r}| + 2|Z_{\ell}^{R-r+1}|$, and so by our inductive assumption,

$$|Z_{\ell+1}^{R-r}| \leqslant (2\ell)^r + 2(2\ell)^{r-1} = (2\ell)^r\left(\frac{\ell+1}{\ell}\right) = 2^r\ell^{r-1}(\ell+1) \leqslant (2(\ell+1))^r.$$

Therefore, (**) is true for every r, $0 \leqslant r < R$, and for every ℓ, $0 \leqslant \ell < L$.

We use can now use this inequality and (0) to solve for the number of vertices:

$$n = \sum_{\ell=0}^{L}\sum_{r=0}^{R}|Z_{\ell}^{r}| \leqslant \sum_{r=0}^{R}|Z_{0}^{r}| + \sum_{\ell=1}^{L}\sum_{r=0}^{R}(2\ell)^r \leqslant 1 + R\sum_{\ell=1}^{L}(2\ell)^R \leqslant R(2L)^{R+1}.$$

Therefore, $L \geqslant (n/R)^{\frac{1}{R+1}}/2$. But $h(G) \geqslant L$ and $t(G) \geqslant R$, so we are done. □

Finally, we formalize the notion that graphs with nice bounds on their tree rank have large homogeneous subgraphs:

Theorem 5.8. *Let G be a graph n vertices such that $t(G) \leqslant t$. Then G contains a homogeneous subgraph of size $(n/t)^{\frac{1}{t+1}}/4$.*

Proof. By Lemma 5.7, $h(G) \geqslant (n/t(G))^{\frac{1}{t(G)+1}}/2$. But by assumption, $t(G) \leqslant t$, so $h(G) \geqslant (n/t)^{\frac{1}{t+1}}/2$. Fix any type tree T of G and a path σ of length $h(G)$ in T. Suppose at least half of the bits in σ are 0s. Suppose the vertices corresponding to these bits are $v_1, ..., v_\ell$. Then by definition of a type tree, for every i, $1 < i \leqslant \ell$, $v_1 \not\sim v_i$, for every j, $2 < j \leqslant \ell$, $v_2 \not\sim v_j$, and so on. Hence, $G\restriction_{\{v_1,\dots,v_\ell\}}$ is empty, and by assumption, $\ell \geqslant (n/t)^{\frac{1}{t+1}}/4$, as desired. Otherwise, at least half of the bits of σ are 1s and we get a large complete subgraph in the same manner. □

Note that this method of proof gives us another proof of the conjecture for complete and empty graphs. K_s-free graphs cannot have the path $\underbrace{11\dots11}_{s\text{ times}}$ and $\overline{K_t}$-free graphs cannot have the path $\underbrace{00\dots00}_{t\text{ times}}$, so they have bounded tree rank, and thus large tree height and large homogeneous subgraphs.

6. A Model-Theoretic Extension

While Erdős and Hajnal conjectured that the absence of *any* graph imposes enough structure to get a polynomial-sized homogeneous subgraph, there is something noteworthy about the claim for halfgraphs. A halfgraph is an instance of an order property in a model, a stable graph is an instance of a stable model, and these stable models have large indiscernible sets (thereby generalizing the fact that halfgraphs have the Erdős-Hajnal property). We will prove this in the next section, but we first introduce some model-theoretic concepts, namely types, indiscernible sets, and stability. Our definitions will follow those in [4].

The section and the next assume some familiarity with the definition of models, theories of models, elementary equivalences, and the statements of the compactness and completeness theorems, an extensive account of which can be found in [3]. Throughout our work in model theory, we refer to the model of an infinite graph as a motivating example. This is the notion of a model-theoretic graph introduced in Section 1. From now on, we let $\mathscr{L}_G$ be the language consisting of a binary relation R (representing adjacency), let T_G be the theory dictating that R is anti-reflexive and symmetric, and let the phrase "infinite graph" be any infinite model of T_G.

Definition 6.1. Let M be a model in the language $\mathscr{L}$ and let A be a subset of its universe, which we denote by $|M|$. A *type* of M over A is a maximal set of formulas in $\mathscr{L}(A)$. The type of some sequence $\overline{b}$ of elements of $|M|$ over A is

$$\operatorname{tp}(\overline{b}, A, M) = \{\varphi(\overline{x}, \overline{a})^t : \overline{a} \in A, t \in \{0,1\}, \varphi \text{ a formula in } \mathscr{L}, M \models \varphi[\overline{b}, \overline{a}]^t\},$$

where $\varphi(x)^1 = \varphi(x)$ and $\varphi(x)^0 = \neg\varphi(x)$. We often use $p(\overline{x})$ to denote a type in the variables $\overline{x}$.

Definition 6.2. The type $p(\overline{x})$ is *realized* in the model M if there exists a sequence $\overline{b}$ of elements in $|M|$ such that $M \models p(\overline{b})$.

Example 6.3. We turn to the theory of an infinite graph. The language $\mathscr{L}_G$ only has one formula, an adjacency relation, so a type over a set A in an infinite graph is just a partition of A. For instance, the type $\{R(x,a) : a \in A\}$ is realized in a graph G if and only if there is a vertex b in G that is connected to every vertex in A.

Example 6.4. Note that a type need not be realized by some element of the model. Consider the model of $\mathbb{N}$ equipped with the constant 1 and a binary relation $<$, interpreted as the usual linear order. By compactness, the set of sentences $\{\varphi_i(x) : i \in \omega\}$, where $\varphi_i(x) = (\exists y)(y > \underbrace{1 + 1 + ... + 1}_{i \text{ times}})$, is consistent. It can be extended to a maximal consistent set to get a type that is not realized by any element in $\mathbb{N}$.

Definition 6.5. Let M be a model and let $X \subseteq |M|$ be a subset of its universe that is equipped with a relation $<$ that strictly simply orders X. We say that X is an *indiscernible set* (with respect to $<$) over a set A of parameters if for every n and every $x_1 < ... < x_n$ and $y_1 < ... < y_n$ from X,

$$(M, x_1, ..., x_n) \equiv_A (M, y_1, ..., y_n),$$

i.e. the model extended with constants $x_1, ..., x_n$ is elementarily equivalent in $\mathscr{L}(A)$ to the model extended with constants $y_1, ..., y_n$. We simply say that X is an indiscernible set if $A = \varnothing$.

Example 6.6. A homogeneous subset of a graph is an indiscernible sequence. A homogeneous subset of a graph whose vertices all define the same partition of a subset A is an indiscernible set over A.

Example 6.7. Let M be a model of a dense linear order, with the single constant 3. Any increasing sequence of singletons that all satisfy the formula $x > 3$ is an indiscernible set.

We showed in Section 2 that in infinite graphs, we can always find infinite homogeneous subgraphs. Similarly, we show that in theories with infinite models, we can always construct large indiscernible sequences. And later, we show that stability allows us to improve this result.

Proposition 6.8. *Let $\mathscr{L}$ be a language, let T be a theory in $\mathscr{L}$ with infinite models, and let $\langle X, < \rangle$ be any simply ordered set. Then there is a model $M \models T$ with $X \subseteq |M|$ such that X is a set of indiscernibles in M.*

Proof. Extend $\mathscr{L}$ to $\mathscr{L}' = \mathscr{L} \cup \{c_x : x \in X\}$ and extend T to

$$T' = T \cup \{\neg(c_x \equiv c_y) : x, y \in X, x \neq y\} \cup$$

$$\{\varphi(c_{x_1}, ..., c_{x_n}) \leftrightarrow \varphi(c_{y_1}, ..., c_{y_n}) : \varphi \text{ a formula in } \mathscr{L},$$

$$n \in \omega, \text{ and } x_1 < ... < x_n, y_1 < .... < y_n \text{ from } X\}$$

We will show that T' is a consistent set of sentences in $\mathscr{L}'$, so it has some model M'. The reduct $M = M' \restriction_{\mathcal{L}}$ is a model of T. We identify each constant c_x with the corresponding element $x \in X$. Then by construction, for any n-ary formula φ in $\mathscr{L}$ and any sequences $x_1 < ... < x_n$ and $y_1 < ...., y_n$ from X, we have

$$M \models \varphi[x_1, ..., x_n] \iff M \models \varphi[y_1, ..., y_n].$$

Thus, showing that T' is consistent will complete the proof. We do so by compactness. Let N be an infinite model of T and let $I \subseteq |N|$ be countably infinite. Let $<$ be a well-ordering on the elements of I and list $I = \{i_0 < i_1 < ... < i_m < ...\}$. Denote the following by (*):

> For every finite subset Δ of T', there is an infinite subset $J = J(\Delta) \subseteq I$ such that for every infinite subset $j_0 < j_1 < .... < j_m < ...$ of J, the model M extended with constants $\{j_m\}_{m \in \omega}$, denoted by $(M, j_m)_{m \in \omega}$, satisfies Δ.

We show that (*) holds by induction on the number of sentences in Δ. First, note that $T \cup \{\neg(c_x \equiv c_y) : x, y \in X, x \neq y\}$ is consistent already, so we only need to induct on the sentences in T' that are of the form $\varphi(c_{x_1}, ..., c_{x_n}) \leftrightarrow \varphi(c_{y_1}, ..., c_{y_n})$.

Assume that (*) holds for some finite $\Delta \subseteq T'$, i.e. that $J = J(\Delta) \subset I$ is an infinite set such that for every infinite sequence $j_0 < j_1 < ... < j_m < ...$ of J, the model $(M, j_m)_{m \in \omega}$ satisfies Δ. Pick $\varphi(x_1, ..., x_n)$ from $\mathscr{L}$ such that we have not yet added the indiscernibility condition for φ to Δ.

Let J_n be the set of all n-element sequences $j_1 < ... < j_n$ of J. Since M is maximal consistent, we can write $J_n = A \sqcup B$, where

$$A = \{j_1 < ... < j_n : M \models \varphi[j_1, ..., j_n]\} \text{ and}$$

$$B = \{j_1 < ... < j_n : M \models \neg\varphi[j_1, ..., j_n]\}.$$

By the infinite version of Ramsey's Theorem for hypergraphs, there is an infinite set $K \subseteq J$ such that $[K]^n \subseteq A$ or $[K]^n \subseteq B$. Then for any choice of $i_1 < ... < i_n$ and $j_1 < ... < j_n$ from K, $\varphi(c_{i_1}, ..., c_{i_n}) \leftrightarrow \varphi(c_{j_1}, ..., c_{j_n})$. Thus, (*) holds for $\Delta \cup \varphi$ and by induction, we are done. □

The main result of Section 7 has to do with the notion of model-theoretic stability, which has many characterizations. We provide two equivalent notions. A proof of their equivalence is beyond the scope of this paper, but can be found in [4].

Definition 6.9. Let M be a model in the language $\mathscr{L}$ and let λ be a cardinal. We say that M is λ-*stable* is for all $A \subseteq |M|$, $|A| \leqslant \lambda$, and for all $m < \omega$, the number of types of M in m variables over A is at most λ.

Definition 6.10. A theory T is λ-*stable* if every $M \models T$ is λ-stable.

Our second characterization of stability will use the notion of an order property:

Definition 6.11. Let M be a model in the language $\mathscr{L}$. The formula $\varphi(\overline{x}, \overline{y})$ in $\mathscr{L}$ has the k-*order property* for $k \in \omega$ if there are $\{\overline{a_i} : i < k\}$ and $\{\overline{b_i} : i < k\}$ such that

$$M \models \varphi[\overline{a_i}, \overline{b_j}] \iff i < j.$$

The formula φ has the *order property* if it has the k-order property for every $k \in \omega$.

Definition 6.12. Let T be a theory in the language $\mathscr{L}$. We say that T is *stable* if none of its formulas have the order property.

Example 6.13. The adjacency relation R in $\mathscr{L}_G$ has the k-order property in a graph if and only if that graph contains a k-halfgraph. Hence, if a graph G is halfgraph-free, it is stable.

We can now state a strengthening of Proposition 6.8 for stable theories:

Theorem 6.14. *If the theory T is stable in λ and $|I| > \lambda \geqslant |A|$, then there is $J \subseteq I$, $|J| > \lambda$, which is an indiscernible sequence over A.*

A proof of this is available in Section 1.2 of [4]. We instead prove an entirely combinatorial rendition of the theorem, assuming no knowledge of model theory.

7. A Purely Combinatorial Proof

Essentially, Theorem 6.14 (applied to the theory of an infinite graph) states that if a graph G in which there is some sense of preservation of the local structure has a subgraph W such that few partitions of W are realized in G, then there is a very large homogeneous subset in G whose elements all admit the same partition of W.

First, we formally define a partition and the notion of a graph preserving local structure that we will invoke: the notion of an $\aleph_0$-homogeneous graph (not to be confused with the clique or anticlique we hope to find).

Definition 7.1. A graph G is $\aleph_0$-*homogeneous* if every isomorphism between two finite induced subgraphs of G can be extended to an automorphism of G.

Definition 7.2. Let $G = (V, E)$ be a graph. A *partition* p of $W \subseteq V$ is a subset $W' \subseteq W$ such that there exists a vertex $v \in V \backslash W$ satisfying

$$W' = \{w \in W : v \text{ and } w \text{ are adjacent in } G\}.$$

When we say that a vertex $v \in V$ *admits* or *realizes* a partition p of $W \subseteq V$, we mean that v witnesses the partition p of W. Also, we say that for a subset $U \subseteq V$, a partition *p restricted to U* (denoted by $p\restriction_U$) of W is a partition of W such that $U \backslash W$ contains a vertex that admits the partition.

Now, we can formally state our combinatorial version of Theorem 6.14:

Theorem 7.3. *Let $G = (V, E)$ be an uncountable $\aleph_0$-homogeneous graph. Suppose that given any countable subset $W \subset V$, the vertices $V \backslash W$ admit only countably many distinct partitions of W. Then there is an uncountable homogeneous set $U \subseteq V$ such that every vertex in U admits the same partition of W.*

In our proof, we use the following concepts, instantiated from [4]:

Definition 7.4. Let $G = (V, E)$ be a graph and let p be a partition of $V' \subset V$. We say that p *splits* over V' if there are $v, w \in V$ such that v and w admit the same partition of V' but such that there is some $x \in V$ such that $x \sim v$ and $x \nsim w$.

Definition 7.5. A *p-splitting chain* in an infinite graph $G = (V, E)$ is an increasing sequence $\{W_i\}_{i \leqslant \aleph_0}$ of nested vertex sets of G such that:
(i) p is a partition of $W_{\aleph_0} = \bigcup_{i \leqslant \aleph_0} W_i$, and
(ii) $p{\restriction}_{W_{i+1}}$ splits over W_i for every $i \leqslant \aleph_0$.

Lemma 7.6. *If G satisfies the hypotheses of Theorem 7.3, then G does not contains a splitting chain.*

Proof. Let $G = (V, E)$ be an uncountable $\aleph_0$-homogeneous graph. We show that if a splitting chain were to exist in G, the graph would not satisfy the condition bounding the number of partitions on its countable subgraphs. Let $\{W_i\}_{i \leqslant \aleph_0}$ be a p-splitting chain, where p is a partition of $W_{\aleph_0}$. We build an infinite binary tree T whose nodes are elements of a countable subset of V and whose branches are distinct partitions of that subset. Fix any $v_{\Diamond} \in W_0$ and let this be the first node in T. Since $p{\restriction}_{W_1}$ splits over W_0, there are $v_0, v_1 \in W_1$ that admit the same partition of W_0 (and, in particular, are either both adjacent to or both not adjacent to $v_{\Diamond}$) and there is $w \in V$ such that $w \nsim v_0$ and $w \sim v_1$. Add the vertex v_0 to T to the left of $v_{\Diamond}$ and the vertex v_1 to T to the right of $v_{\Diamond}$.

But now $p{\restriction}_{W_2}$ splits over W_1, so there are $v_{10}, v_{11} \in W_2$ that admit the same partition of W_1 (and, in particular, are either both adjacent to or both not adjacent to v_1) and there is $w \in V$ such that $w \nsim v_{10}$ and $w \sim v_{11}$. Add v_{10} and v_{11} to the nodes corresponding to their indices, i.e. let T be:

$$\begin{tikzcd}
 & v_{10} & & v_{11} \\
v_0 & & v_1 \arrow[ul] \arrow[ur] & \\
 & v_{\Diamond} \arrow[ul] \arrow[ur] & &
\end{tikzcd}$$

By the $\aleph_0$-homogeneity of G, we can also add v_{00} and v_{01} that admit the same partition of W_1 (and, in particular, are either both adjacent or both not adjacent to v_0) but such that there is some $w \in V$ that is not adjacent to v_{00} and that is adjacent to v_{01}. Inductively, we continue in this manner, using the splitting chain $\{W_i\}_{i \leqslant \aleph_0}$ to extend T and the homogeneity of G to fill out T at each level.

Ultimately, T has countably many nodes and uncountably many branches. By construction, for each branch of T, there is a vertex $w \in V$ realizing a distinct partition of the nodes (two partitions differ on the vertex at which their corresponding paths diverge). But this contradicts our assumption about G, so G cannot have such a splitting chain. $\square$

Our next step is to show that given a graph G satisfying the assumptions of Theorem 7.3, we can construct a *non-splitting closure* of each countable subgraph:

Definition 7.7. Let $G = (V, E)$ be an uncountable graph and let W be a countable subset of V. We say that X, Y, and p form a *non-splitting closure* of W if $W \subseteq X \subseteq Y \subset V$, where $|Y| \leqslant \aleph_0$, and p partitions Y such that:

(i) for every Z with $Y \subseteq Z \subset V$ and $|Z| \leqslant \aleph_0$, p extends to a partition q of Z that does not split over X, and
(ii) for every partition of X in G, there is a vertex $y \in Y$ realizing the partition.

Lemma 7.8. *If $G = (V, E)$ satisfies the hypotheses of Theorem 7.3, then any countable $W \subset V$ has a non-splitting closure.*

Proof. Let $W \subset V$ be countable and let p be any partition of W. We define a nested chain of vertex sets $X_0 \subseteq Y_0 \subseteq X_1 \subseteq Y_1 \subseteq ...$ and partitions $p_0 \subseteq p_1 \subseteq ...$ (of $Y_0, Y_1, ...$) such that if no X_i, Y_i, and p_i form a non-splitting closure of W, then $\{X_i\}_{i \leqslant \aleph_0}$ will form a splitting chain in G, contradicting Lemma 7.6. Let $X_0 = W$. Suppose that X_i has been defined for some $i \leqslant \aleph_0$ and suppose that $|X_i| \leqslant \aleph_0$. By the hypotheses of Theorem 7.3, the number of partitions of X_i in G is countable. Therefore, there is a countable set P_i that contains an element realizing every partition of X_i. Let $Y_i = X_i \cup P_i$ and let p_i be any partition of Y_i extending the previous partition. Note that $|Y_i| \leqslant \aleph_0 + \aleph_0 = \aleph_0$ is maintained.

If X_i, Y_i, and p_i form a non-splitting closure of W, we are done. Otherwise, condition (i) in Definition 7.7 must fail. Thus, for every partition p of Y_i that does not split over X_i, we can add vertices to get X_{i+1} such that every partition $q \supseteq p$ of X_{i+1} splits over X_i. We continue this process to get a sequence of length $\aleph_0$, and we claim that $\{X_i\}_{i \leqslant \aleph_0}$ is a splitting chain.

Choose any $v \in V \backslash X_{\aleph_0}$ and let p be the partition of $X_{\aleph_0}$ realized by v. Suppose that $p \restriction_{X_{i+1}}$ does not split over X_i for some i. Consequently, $p \restriction_{Y_i}$ also does not split over X_i. Then there is X_i', which we added to get X_{i+1} in our construction, that contradicts condition (i) of Definition 7.7. By assumption, $p \restriction_{X_{i+1}}$ does not split over X_i, so $p \restriction_{X_i'}$ does not either. However, $p \restriction_{X_i'}$ is realized by $v \in V \backslash X_i'$, so we get a contradiction. Thus, $\{X_i\}_{i \leqslant \aleph_0}$ is a splitting chain and Lemma 7.6 then gives us a contradiction, implying that W has a non-splitting closure. $\square$

Finally, we are ready to use the non-splitting closure of a countable W to create an uncountable homogeneous set, whose every element admits the same partition of W. This will complete the proof of Theorem 7.3:

Proof of Theorem 7.3. Suppose $G = (V, E)$ meets the assumptions of Theorem 7.3. Let W be any countable subset of V. By Lemma 7.8, W has a non-splitting closure, say from X, Y, and p. We inductively define an uncountable $Z \subseteq V$ that satisfies the desired conditions. We choose vertices $\{z_\alpha : \alpha < \aleph_1\}$ in stages, and let $Z_\alpha = Y \cup \{z_\beta : \beta < \alpha\}$. We then let $Z = \bigcup_{\alpha < \aleph_1} Z_\alpha$. By condition (i) of Definition 7.7, we can always find a partition $p_\alpha \supseteq p$ of Z_α that does not split over X.

We claim that for any $\beta < \alpha$, p_α extends p_β. We use part (ii) of Definition 7.7, the condition that Y already realizes all partitions of X in G. Suppose $z \in p_\beta$. Then by (ii), there is some $z' \in Y$ realizing the same partition of X as z. Since p_β does not split over X, $z \in p_\beta$ implies that $z' \in p_\beta$. Since p_β extends p, $z' \in p$ as well. But p_α also extends p, so $z' \in p_\alpha$. However, p_α also does not split over X, so $z \in p_\alpha$. The proof that $z \notin p_\alpha \implies z \notin p_\beta$ is identical.

Then since for every α, p_α does not split over X, any two vertices of Z have the same relation to all others, and Z is consequently homogeneous. $\square$

8. Conclusion

There are many extensions of the results detailed in this paper that we have not yet mentioned. Maria Chudnovsky's survey discusses the Erdős-Hajnal Conjecture for *tournaments*, directed graphs where for any two vertices u, v, exactly one of (u, v) and (v, u) is an edge. She defines the substitution operation for tournaments, and proves the conjecture for several specific tournaments.

Other sources consider the importance of the characterization of a *perfect* graph: a graph G such that for every induced subgraph H of G, the maximum size of a clique in H is equal to the chromatic number of H. An equivalent characterization of a perfect graph is a graph G such that no induced subgraph of G or its complement is an odd cycle of length $\geqslant 5$ [9]. All of the graphs which have been shown to satisfy the Erdős-Hajnal property are perfect. Furthermore, to prove the conjecture for some graph H, we need only prove it for H-free graphs that are not perfect, as all perfect graphs already have especially large homogeneous sets.

In addition, there are other graphs that have been shown to have the Erdős-Hajnal property, for example caterpillars and their complements (shown in [10]). An analysis of these was forfeited in exchange for a foray into model theory. I hope the analogy between halfgraphs and order properties and between homogeneous and indiscernible sets was appreciated, that the combinatorial rendition of Theorem 7.1 was convincing, and that the transition between infinite model theory and finite combinatorics seemed as natural as it felt to me.

Acknowledgments. It is a pleasure to thank my mentor Sarah Reitzes for her input and advice and Professor Peter May for organizing another incredible REU. I would also like to thank Professor Maryanthe Malliaris for directing me to this topic and for her invaluable mentorship throughout the summer.

References

[1] M. Chudnovsky. The Erdős-Hajnal Conjecture–A Survey.
[2] M. Malliaris and C. Terry. On Unavoidable Induced Subgraphs in Large Prime Graphs.
[3] C. C. Chang and H. J. Keisler. Model Theory
[4] S. Shelah. Classification Theory.
[5] N. Bousquet, A. Lagoutte, and S. Thomassé. The Erdős-Hajnal Conjecture for Paths and Antipaths.
[6] Ramsey's Theorem. `https://en.wikipedia.org/wiki/Ramsey%27s_theorem`.
[7] N. Alon, J. Pach, and J. Solymosi. Ramsey-type theorems with forbidden subgraphs.
[8] M. Chudnovsky and S. Safra. The Erdős-Hajnal Conjecture for Bull-free Graphs.
[9] Strong perfect graph theorem. `https://en.wikipedia.org/wiki/Strong_perfect_graph_theorem`.
[10] A. Liebenau and M. Pilipczuk. The Erdős-Hajnal conjecture for caterpillars and their complements.
[11] W. Hodges. Model Theory.