

UBIQUITY IN GRAPHS II: UBIQUITY OF GRAPHS WITH NON-LINEAR END STRUCTURE

NATHAN BOWLER, CHRISTIAN ELBRACHT, JOSHUA ERDE, PASCAL GOLLIN, KARL HEUER,
MAX PITZ, AND MAXIMILIAN TEEGEN

ABSTRACT. A graph G is said to be \preceq -ubiquitous, where \preceq is the minor relation between graphs, if whenever Γ is a graph with $nG \preceq \Gamma$ for all $n \in \mathbb{N}$, then one also has $\aleph_0 G \preceq \Gamma$, where αG is the disjoint union of α many copies of G . A well-known conjecture of Andreae is that every locally finite connected graph is \preceq -ubiquitous.

In this paper we give a sufficient condition on the structure of the ends of a graph G which implies that G is \preceq -ubiquitous. In particular this implies that the full grid is \preceq -ubiquitous.

§1. INTRODUCTION

This paper is the second in a series of papers making progress towards a conjecture of Andreae on the *ubiquity* of graphs. Given a graph G and some relation \triangleleft between graphs we say that G is \triangleleft -ubiquitous if whenever Γ is a graph such that $nG \triangleleft \Gamma$ for all $n \in \mathbb{N}$, then $\aleph_0 G \triangleleft \Gamma$, where αG denotes the disjoint union of α many copies of G . For example, a classic result of Halin [9] says that the ray is \subseteq -ubiquitous, where \subseteq is the subgraph relation.

Examples of graphs which are not ubiquitous with respect to the subgraph or topological minor relation are known (see [2] for some particularly simple examples). In [1] Andreae initiated the study of ubiquity of graphs with respect to the minor relation \preceq . He constructed a graph which is not \preceq -ubiquitous, however the construction relied on the existence of a counterexample to the well-quasi-ordering of infinite graphs under the minor relation, for which only examples of very large cardinality are known [13]. In particular, the question of whether there exists a countable graph which is not \preceq -ubiquitous remains open. Most importantly, however, Andreae [1] conjectured that at least all locally finite graphs, those with all degrees finite, should be \preceq -ubiquitous.

The Ubiquity Conjecture. *Every locally finite connected graph is \preceq -ubiquitous.*

In [2] Andreae proved that his conjecture holds for a large class of locally finite graphs. The exact definition of this class is technical, but in particular his result implies the following.

Theorem 1.1 (Andreae, [2, Corollary 2]). *Let G be a connected, locally finite graph of finite tree-width such that every block of G is finite. Then G is \preceq -ubiquitous.*

Note that every end in such a graph G must have degree¹ one.

Andreae’s proof employs deep results about well-quasi-orderings of labelled (infinite) trees [12]. Interestingly, the way these tools are used does not require the extra condition in Theorem 1.1 that every block of G is finite and so it is natural to ask if his proof can be adapted to remove this condition. And indeed, it is the purpose of the present and subsequent paper in our series, [3], to show that this is possible, i.e. that all connected, locally finite graphs of finite tree-width are \preceq -ubiquitous.

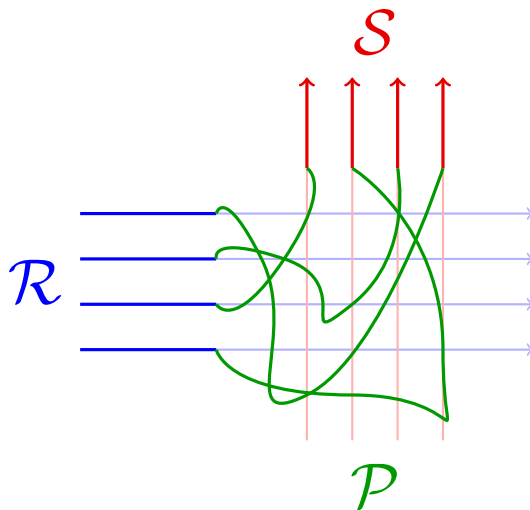


FIGURE 1.1. A linkage between \mathcal{R} and \mathcal{S} .

The present paper lays the groundwork for this extension of Andreae’s result. The fundamental obstacle one encounters when trying to extend Andreae’s methods is the following: Let $[n] = \{1, 2, \dots, n\}$. In the proof we often have two families of disjoint rays $\mathcal{R} = (R_i : i \in [n])$ and $\mathcal{S} = (S_j : j \in [m])$ in Γ , which we may assume all converge¹ to a common end of Γ , and we wish to find a *linkage* between \mathcal{R} and \mathcal{S} , that is, an injective function $\sigma : [n] \rightarrow [m]$ and a set \mathcal{P} of disjoint finite paths P_i from $x_i \in R_i$ to $y_{\sigma(i)} \in S_{\sigma(i)}$ such that the walks

$$\mathcal{T} = (R_i x_i P_i y_{\sigma(i)} S_{\sigma(i)} : i \in [n])$$

formed by following each R_i along to x_i , then following the path P_i to $y_{\sigma(i)}$, then following the tail of $S_{\sigma(i)}$, form a family of disjoint rays (see Figure 1.1). Broadly, we can think of

¹A precise definitions of rays, the ends of a graph, their degree, and what it means for a ray to converge to an end can be found in Section 2.

this as ‘re-routing’ the rays \mathcal{R} to some subset of the rays in \mathcal{S} . Since all the rays in \mathcal{R} and \mathcal{S} converge to the same end of Γ , it is relatively simple to show that, as long as $n \leq m$, there is enough connectivity between the rays in Γ so that such a linkage always exists.

However, in practice it is not enough for us to be guaranteed the existence of some injection σ giving rise to a linkage, but instead we want to choose σ in advance, and be able to find a corresponding linkage afterwards.

In general, however, it is quite possible that for certain choices of σ no suitable linkage exists. Consider for example the case where Γ is the *half grid* (briefly denoted by $\mathbb{Z}\square\mathbb{N}$), which is the graph whose vertex set is $\mathbb{Z} \times \mathbb{N}$ and where two vertices are adjacent if they differ in precisely one co-ordinate and the difference in that co-ordinate is one. If we consider two sufficiently large families of disjoint rays \mathcal{R} and \mathcal{S} in Γ , then it is not hard to see that both \mathcal{R} and \mathcal{S} inherit a linear ordering from the planar structure of Γ , which must be preserved by any linkage between them.

Analysing this situation gives rise to the following definition: We say that an end ϵ of a graph G is *linear* if for every finite set \mathcal{R} of at least three disjoint rays in G which converge to ϵ we can order the elements of \mathcal{R} as $\mathcal{R} = \{R_1, R_2, \dots, R_n\}$ such that for each $1 \leq k < i < \ell \leq n$, the rays R_k and R_ℓ belong to different ends of $G - V(R_i)$.

Thus the half grid has a unique end and it is linear. On the other end of the spectrum, let us say that a graph G has *nowhere-linear end structure* if no end of G is linear. Since ends of degree at most two are automatically linear, every end of a graph with nowhere-linear end structure must have degree at least three.

Our main theorem in this paper is the following.

Theorem 1.2. *Every locally finite connected graph with nowhere-linear end structure is \preceq -ubiquitous.*

Roughly, if we assume that every end of G has nonlinear structure, then the fact that $nG \preceq \Gamma$ for all $n \in \mathbb{N}$ allows us to deduce that Γ must also have some end with a sufficiently complicated structure that we can always find suitable linkages for all σ as above. In fact, this property is so strong that we do not need to follow Andreae’s strategy for such graphs. We can use the linkages to directly build a K_{\aleph_0} -minor of Γ , and it follows that $\aleph_0 G \preceq \Gamma$.

In later papers in the series, we shall need to make more careful use of the ideas developed here. We shall analyse the possible kinds of linkages which can arise between two families of rays converging to a given end. If some end of Γ admits many different kinds of linkages, then we can again find a K_{\aleph_0} -minor. If not, then we can use the results of the present paper to show that certain ends of G are linear. This extra structure allows us to carry out an argument like that of Andreae, but using only the limited collection of these maps

σ which we know to be present. This technique will be key to extending Theorem 1.1 in [3].

Independently of these potential later developments, our methods already allow us to establish new ubiquity results for many natural graphs and graph classes.

As a first concrete example, let G be the full grid, a graph not previously known to be ubiquitous. The *full grid* (briefly denoted by $\mathbb{Z} \square \mathbb{Z}$) is analogously defined as the half grid but with $\mathbb{Z} \times \mathbb{Z}$ as vertex set. The grid G is one-ended, and for any ray R in G , the graph $G - V(R)$ still has at most one end. Hence the unique end of G is non-linear, and so Theorem 1.2 has the following corollary:

Corollary 1.3. *The full grid is \preceq -ubiquitous.*

Using an argument similar in spirit to that of Halin [10], we also establish the following theorem in this paper:

Theorem 1.4. *Any connected minor of the half grid $\mathbb{N} \square \mathbb{Z}$ is \preceq -ubiquitous.*

Since every countable tree is a minor of the half grid, Theorem 1.4 implies that all countable trees are \preceq -ubiquitous, see Corollary 7.4. We remark that while all trees are ubiquitous with respect to the topological minor relation, [5], the problem whether all uncountable trees are \preceq -ubiquitous has remained open, and we hope to resolve this in a paper in preparation [4].

In a different direction, if G is any locally finite connected graph, then it is possible to show that $G \square \mathbb{Z}$ or $G \square \mathbb{N}$ either have nowhere-linear end structure, or are a subgraph of the half grid respectively. Hence, Theorems 1.2 and 1.4 together have the following corollary.

Theorem 1.5. *For every locally finite connected graph G , both $G \square \mathbb{Z}$ and $G \square \mathbb{N}$ are \preceq -ubiquitous.*

Finally, we will also show the following result about non-locally finite graphs. For $k \in \mathbb{N}$, we let the *k -fold dominated ray* be the graph DR_k formed by taking a ray together with k additional vertices, each of which we make adjacent to every vertex in the ray. For $k \leq 2$, DR_k is a minor of the half grid, and so ubiquitous by Theorem 1.4. In our last theorem, we show that DR_k is ubiquitous for all $k \in \mathbb{N}$.

Theorem 1.6. *The k -fold dominated ray DR_k is \preceq -ubiquitous for every $k \in \mathbb{N}$.*

The paper is structured as follows: In Section 2 we introduce some basic terminology for talking about minors. In Section 3 we introduce the concept of a *ray graph* and *linkages* between families of rays, which will help us to describe the structure of an end.

In Sections 4 and 5 we introduce a pebble-pushing game which encodes possible linkages between families of rays and use this to give a sufficient condition for an end to contain a countable clique minor. In Section 6 we re-introduce some concepts from [5] and show that we may assume that the G -minors in Γ are *concentrated* towards some end ϵ of Γ . In Section 7 we use the results of the previous section to prove Theorem 1.4 and finally in Section 8 we prove Theorem 1.2 and its corollaries.

§2. PRELIMINARIES

In our graph theoretic notation we generally follow the textbook of Diestel [7]. Given two graphs G and H the *cartesian product* $G \square H$ is a graph with vertex set $V(G) \times V(H)$ with an edge between (a, b) and (c, d) if and only if $a = c$ and $(b, d) \in E(H)$ or $(a, c) \in E(G)$ and $b = d$.

Definition 2.1. A one-way infinite path is called a *ray* and a two-way infinite path is called a *double ray*.

For a path or ray P and vertices $v, w \in V(P)$, let vPw denote the subpath of P with endvertices v and w . If P is a ray, let Pv denote the finite subpath of P between the initial vertex of P and v , and let vP denote the subray (or *tail*) of P with initial vertex v .

Given two paths or rays P and Q which are disjoint but for one of their endvertices, we write PQ for the *concatenation of P and Q* , that is the path, ray or double ray $P \cup Q$. Moreover, if we concatenate paths of the form vPw and wQx , then we omit writing w twice and denote the concatenation by $vPwQx$.

Definition 2.2 (Ends of a graph, cf. [7, Chapter 8]). An *end* of an infinite graph Γ is an equivalence class of rays, where two rays R and S are equivalent if and only if there are infinitely many vertex disjoint paths between R and S in Γ . We denote by $\Omega(\Gamma)$ the set of ends of Γ .

We say that a ray $R \subseteq \Gamma$ *converges* (or *tends*) to an end ϵ of Γ if R is contained in ϵ . In this case we call R an ϵ -ray.

Given an end $\epsilon \in \Omega(\Gamma)$ and a finite set $X \subseteq V(\Gamma)$ there is a unique component of $\Gamma - X$ which contains a tail of every ray in ϵ , which we denote by $C(X, \epsilon)$.

For an end $\epsilon \in \Omega(\Gamma)$ we define the *degree* of ϵ in Γ as the supremum of all sizes of sets containing vertex disjoint ϵ -rays. If an end has finite degree, we call it *thin*. Otherwise, we call it *thick*.

A vertex $v \in V(\Gamma)$ *dominates* an end $\epsilon \in \Omega(\Gamma)$ if there is a ray $R \in \epsilon$ such that there are infinitely many $v-R$ -paths in Γ that are vertex disjoint except from v .

We will use the following two basic facts about infinite graphs.

Proposition 2.3. [7, Proposition 8.2.1] *An infinite connected graph contains either a ray or a vertex of infinite degree.*

Proposition 2.4. [7, Exercise 8.19] *A graph G contains a subdivided K_{\aleph_0} as a subgraph if and only if G has an end which is dominated by infinitely many vertices.*

Definition 2.5 (Inflated graph). Given a graph G , we say that a pair (H, φ_H) is an *inflated copy of G* , or an *IG* for short, if H is a graph and $\varphi_H: V(H) \rightarrow V(G)$ is such that:

- For every $v \in V(G)$ the *branch-set* $\varphi_H^{-1}(v)$ induces a non-empty, connected subgraph of H ;
- There is an edge in H between $\varphi_H^{-1}(v)$ and $\varphi_H^{-1}(w)$ if and only if $(v, w) \in E(G)$ and this edge, if it exists, is unique.

When there is no danger of confusion we will simply say that H is an *IG* instead of saying that (H, φ_H) is an *IG*, and denote by $H(v)$ the branch-set of v . By definition, a graph G is a minor of another graph Γ if and only if there is some subgraph $H \subseteq \Gamma$ such that H is an *IG*.

We will say an *IG* (H, φ_H) is *tidy* if H is subgraph-minimal such that the pair (H, φ_H) is an *IG*, that is, $(H', \varphi_H \upharpoonright V(H'))$ is not an *IG* whenever $H' \subsetneq H$. For a given *IG* (H, φ) by Zorn's lemma there is always a subgraph $H' \subseteq H$ such that $(H', \varphi \upharpoonright V(H'))$ is a tidy *IG*, however this choice may not be unique.

We note that in this case the subgraph of H induced on $H(v)$ is a tree for every $v \in V(G)$, and every leaf of this tree is incident with some edge $(u, v) \in E(H)$ between two branch sets. Furthermore, if $d_G(v)$ is finite, then so is $H(v)$. In this paper we will always assume without loss of generality that each *IG* is tidy.

If $G \subseteq G'$ and H is an *IG* we say that an *IG*' $H' \supseteq H$ *extends* H if $H(v) \subseteq H'(v)$ for all $v \in V(G) \cap V(G')$ and we say that H' is an *extension* of H . Note that, since $H \subseteq H'$, for every $(v, w) \in E(G)$ the unique edge between $H'(v)$ and $H'(w)$ is also the unique edge between $H(v)$ and $H(w)$.

Definition 2.6 (Pullback). Let G be a graph, $M \subseteq G$ a subgraph without isolated vertices, and let H be a tidy *IG*. The *pullback of M to H* is the *IM* $(H(M), \varphi_H \upharpoonright_{V(H(M))})$ where $H(M) \subseteq H$ is the unique subgraph such that $(H(M), \varphi_H \upharpoonright_{V(H(M))})$ is a tidy *IM*.

Note that, due to the tidiness requirement, $H(M)$ might be a proper subgraph of $H[\varphi_H^{-1}(V(M))]$. The requirement that M does not contain isolated vertices is necessary to make this subgraph unique, since if M contains an isolated vertex v , then there isn't a

unique choice of a vertex in the branch-set of v to choose. In particular, there two notations $H(M)$ and $H(v)$ for pullback and branch-set respectively are mutually exclusive.

Lemma 2.7. *Let G be a graph and let H be a tidy IG . If $R \subseteq G$ is a ray, then $H(R)$ is also a ray.*

Proof. Let $R = x_1x_2\dots$ be a ray. For each $i \geq 1$ there is a unique edge $(v_i, w_i) \in E(H)$ between $H(x_i)$ and $H(x_{i+1})$. Since $\{w_i, v_{i+1}\} \subseteq H(x_{i+1})$, and H is tidy, there is a unique path P_i between $\{w_i, v_{i+1}\}$ in H .

By minimality of $H(R)$, it follows that $H(R)(x_1) = \{v_1\}$ and $H(R)(x_i) = V(P_i)$ for each $i \geq 2$. Hence $H(R)$ is a ray. \square

§3. THE RAY GRAPH

Definition 3.1 (Ray graph). Given a finite family of disjoint rays $\mathcal{R} = (R_i : i \in I)$ in a graph Γ the *ray graph* $RG_\Gamma(\mathcal{R}) = RG_\Gamma(R_i : i \in I)$ is the graph with vertex set I and with an edge between i and j if there is an infinite collection of vertex disjoint paths from R_i to R_j in Γ which meet no other R_k . When the host graph Γ is clear from the context we will simply write $RG(\mathcal{R})$ for $RG_\Gamma(\mathcal{R})$.

The following lemmas are simple exercises. For a family \mathcal{R} of disjoint rays in G tending to the same end and $H \subseteq \Gamma$ being an IG the aim is to establish the following: if \mathcal{S} is a family of disjoint rays in Γ which contains the pullback $H(R)$ of each $R \in \mathcal{R}$, then the subgraph of the ray graph $RG_\Gamma(\mathcal{S})$ induced on the vertices given by $\{H(R) : R \in \mathcal{R}\}$ is connected.

Lemma 3.2. *Let G be a graph and let $\mathcal{R} = (R_i : i \in I)$ be a finite family of disjoint rays in G . Then $RG_G(\mathcal{R})$ is connected if and only if all rays in \mathcal{R} tend to a common end $\omega \in \Omega(G)$.*

Lemma 3.3. *Let G be a graph, $\mathcal{R} = (R_i : i \in I)$ be a finite family of disjoint rays in G and let H be an IG . If $\mathcal{R}' = (H(R_i) : i \in I)$ are the pullbacks of the rays in \mathcal{R} in H , then $RG_G(\mathcal{R}) = RG_H(\mathcal{R}')$.*

Lemma 3.4. *Let G be a graph, $H \subseteq G$, $\mathcal{R} = (R_i : i \in I)$ be a finite disjoint family of rays in H and let $\mathcal{S} = (S_j : j \in J)$ be a finite disjoint family of rays in $G - V(H)$, where I and J are disjoint. Then $RG_H(\mathcal{R})$ is a subgraph of $RG_G(\mathcal{R} \cup \mathcal{S})[I]$. In particular, if all rays in \mathcal{R} tend to a common end in H , then $RG_G(\mathcal{R} \cup \mathcal{S})[I]$ is connected.*

Recall that an end ω of a graph G is called *linear* if for every finite set \mathcal{R} of at least three disjoint ω -rays in G we can order the elements of \mathcal{R} as $\mathcal{R} = \{R_1, R_2, \dots, R_n\}$ such that for each $1 \leq k < i < \ell \leq n$, the rays R_k and R_ℓ belong to different ends of $G - V(R_i)$.

Lemma 3.5. *An end ω of a graph G is linear if and only if the ray graph of every finite family of disjoint ω -rays is a path.*

Proof. For the forward direction suppose ω is linear and $\{R_1, R_2, \dots, R_n\}$ converge to ω , with the order given by the definition of linear. It follows that there is no $1 \leq k < i < \ell \leq n$ such that (k, ℓ) is an edge in $RG(R_j: j \in [n])$. However, by Lemma 3.2 $RG(R_j: j \in [n])$ is connected, and hence it must be the path $12 \dots n$.

Conversely, suppose that the ray graph of every finite family of ω -rays is a path. Then, every such family \mathcal{R} can be ordered as $\{R_1, R_2, \dots, R_n\}$ such that $RG(\mathcal{R})$ is the path $12 \dots n$. It follows that, for each i , $(k, \ell) \notin E(RG(\mathcal{R}))$ whenever $1 \leq k < i < \ell \leq n - 1$, and so by definition of $RG(\mathcal{R})$ there is no infinite collection of vertex disjoint paths from R_k to R_ℓ in $G - V(R_i)$. Therefore R_k and R_ℓ belong to different ends of $G - V(R_i)$. \square

Definition 3.6 (Tail of a ray after a set). Given a ray R in a graph G and a finite set $X \subseteq V(G)$ the *tail of R after X* , denoted by $T(R, X)$, is the unique infinite component of R in $G - X$.

Definition 3.7 (Linkage of families of rays). Let $\mathcal{R} = (R_i: i \in I)$ and $\mathcal{S} = (S_j: j \in J)$ be families of disjoint rays of Γ , where the initial vertex of each R_i is denoted x_i . A family $\mathcal{P} = (P_i: i \in I)$ of paths in Γ is a *linkage* from \mathcal{R} to \mathcal{S} if there is an injective function $\sigma: I \rightarrow J$ such that

- Each P_i goes from a vertex $x'_i \in R_i$ to a vertex $y_{\sigma(i)} \in S_{\sigma(i)}$;
- The family $\mathcal{T} = (x_i R_i x'_i P_i y_{\sigma(i)} S_{\sigma(i)}: i \in I)$ is a collection of disjoint rays.

We say that \mathcal{T} is obtained by *transitioning* from \mathcal{R} to \mathcal{S} along the linkage. We say the linkage \mathcal{P} *induces* the mapping σ . Given a vertex set $X \subseteq V(G)$ we say that the linkage is *after X* if $X \cap V(R_i) \subseteq V(x_i R_i x'_i)$ for all $i \in I$ and no other vertex in X is used by \mathcal{T} . We say that a function $\sigma: I \rightarrow J$ is a *transition function* from \mathcal{R} to \mathcal{S} if for any finite vertex set $X \subseteq V(G)$ there is a linkage from \mathcal{R} to \mathcal{S} after X that induces σ .

We will need the following lemma from [5], which asserts the existence of linkages.

Lemma 3.8 (Weak linking lemma). *Let Γ be a graph, $\omega \in \Omega(\Gamma)$ and let $n \in \mathbb{N}$. Then for any two families $\mathcal{R} = (R_i: i \in [n])$ and $\mathcal{S} = (S_j: j \in [n])$ of vertex disjoint ω -rays and any finite vertex set $X \subseteq V(G)$, there is a linkage from \mathcal{R} to \mathcal{S} after X .*

§4. A PEBBLE-PUSHING GAME

Suppose we have a family of disjoint rays $\mathcal{R} = (R_i: i \in I)$ in a graph G and a subset $J \subseteq I$. Often we will be interested in which functions we can obtain as transition functions

between $(R_i: i \in J)$ and $(R_i: i \in I)$. We can think of this as trying to ‘re-route’ the rays $(R_i: i \in J)$ to a different set of $|J|$ rays in $(R_i: i \in I)$.

To this end, it will be useful to understand the following pebble-pushing game on a graph.

Definition 4.1 (Pebble-pushing game). Let $G = (V, E)$ be a finite graph. For any fixed positive integer k we call a tuple $(x_1, x_2, \dots, x_k) \in V^k$ a *game state* if $x_i \neq x_j$ for all $i, j \in [k]$ with $i \neq j$.

The *pebble-pushing game (on G)* is a game played by a single player. Given a game state $Y = (y_1, y_2, \dots, y_k)$, we imagine k labelled pebbles placed on the vertices (y_1, y_2, \dots, y_k) . We move between game states by moving a pebble from a vertex to an adjacent vertex which does not contain a pebble, or formally, a Y -*move* is a game state $Z = (z_1, z_2, \dots, z_k)$ such that there is an $\ell \in [k]$ such that $y_\ell z_\ell \in E$ and $y_i = z_i$ for all $i \in [k] \setminus \{\ell\}$.

Let $X = (x_1, x_2, \dots, x_k)$ be a game state. The X -*pebble-pushing game (on G)* is a pebble-pushing game where we start with k labelled pebbles placed on the vertices (x_1, x_2, \dots, x_k) .

We say a game state Y is *achievable* in the X -pebble-pushing game if there is a sequence $(X_i: i \in [n])$ of game states for some $n \in \mathbb{N}$ such that $X_1 = X$, $X_n = Y$ and X_{i+1} is an X_i -move for all $i \in [n-1]$, that is, if it is a sequence of moves that pushes the pebbles from X to Y .

A graph G is k -*pebble-win* if Y is an achievable game state in the X -pebble-pushing game on G for every two game states X and Y .

The following lemma shows that achievable game states on the ray graph $RG(\mathcal{R})$ yield transition functions from a subset of \mathcal{R} to itself. Therefore, it will be useful to understand which game states are achievable, and in particular the structure of graphs on which there are unachievable game states.

Lemma 4.2. *Let Γ be a graph, $\omega \in \Omega(\Gamma)$, $m \geq k$ be positive integers and let $(S_j: j \in [m])$ be a family of disjoint rays in ω . For every achievable game state $Z = (z_1, z_2, \dots, z_k)$ in the $(1, 2, \dots, k)$ -pebble-pushing game on $RG(S_j: j \in [m])$, the map σ defined via $\sigma(i) := z_i$ for every $i \in [k]$ is a transition function from $(S_i: i \in [k])$ to $(S_j: j \in [m])$.*

Proof. We first note that if σ is a transition function from $(S_i: i \in [k])$ to $(S_j: j \in [m])$ and τ is a transition function from $(S_i: i \in \sigma([k]))$ to $(S_j: j \in [m])$, then clearly $\tau \circ \sigma$ is a transition function from $(S_i: i \in [k])$ to $(S_j: j \in [m])$.

Hence, it will be sufficient to show the statement holds when σ is obtained from $(1, 2, \dots, k)$ by a single move, that is, there is some $t \in [k]$ and a vertex $\sigma(t) \notin [k]$ such that $\sigma(t)$ is adjacent to t in $RG(S_j: j \in [m])$ and $\sigma(i) = i$ for $i \in [k] \setminus \{t\}$.

So, let $X \subseteq V(G)$ be a finite set. We will show that there is a linkage from $(S_i: i \in [k])$ to $(S_j: j \in [m])$ after X that induces σ . By assumption there is an edge $(t, \sigma(t)) \in E(RG(S_j: j \in [m]))$. Hence, there is a path P between $T(S_t, X)$ and $T(S_{\sigma(t)}, X)$ which avoids X and all other S_j .

Then the family $\mathcal{P} = (P_1, P_2, \dots, P_k)$ where $P_t = P$ and $P_i = \emptyset$ for each $i \neq t$ is a linkage from $(S_i: i \in [k])$ to $(S_j: j \in [m])$ after X that induces σ . \square

We note that this pebble-pushing game is sometimes known in the literature as “permutation pebble motion” [11] or “token reconfiguration” [6]. Previous results have mostly focused on computational questions about the game, rather than the structural questions we are interested in, but we note that in [11] the authors give an algorithm that decides whether or not a graph is k -pebble-win, from which it should be possible to deduce the main result in this section, Lemma 4.9. However, since a direct derivation was shorter and self contained, we will not use their results. We present the following simple lemmas without proof.

Lemma 4.3. *Let G be a finite graph and X a game state.*

- *If Y is an achievable game state in the X -pebble-pushing game on G , then X is an achievable game state in the Y -pebble-pushing game on G .*
- *If Y is an achievable game state in the X -pebble-pushing game on G and Z is an achievable game state in the Y -pebble-pushing game on G , then Z is an achievable game state in the X -pebble-pushing game on G .*

Definition 4.4. Let G be a finite graph and let $X = (x_1, x_2, \dots, x_k)$ be a game state. Given a permutation σ of $[k]$ let us write $X^\sigma = (x_{\sigma(1)}, x_{\sigma(2)}, \dots, x_{\sigma(k)})$. We define the *pebble-permutation group* of (G, X) to be the set of permutations σ of $[k]$ such that X^σ is an achievable game state in the X -pebble-pushing game on G .

Note that by Lemma 4.3, the pebble-permutation group of (G, X) is a subgroup of the symmetric group S_k .

Lemma 4.5. *Let G be a graph and let X be a game state. If Y is an achievable game state in the X -pebble-pushing game and σ is in the pebble-permutation group of Y , then σ is in the pebble-permutation group of X .*

Lemma 4.6. *Let G be a finite connected graph and let X be a game state. Then G is k -pebble-win if and only if the pebble-permutation group of (G, X) is S_k .*

Proof. Clearly, if the pebble-permutation group is not S_k then G is not k -pebble-win. Conversely, since G is connected, for any game states X and Y there is some τ such that Y^τ

is an achievable game state in the X -pebble-pushing game, since we can move the pebbles to any set of k vertices, up to some permutation of the labels. We know by assumption that $X^{\tau^{-1}}$ is an achievable game state in the X -pebble-pushing game. Therefore, by Lemma 4.3 Y is an achievable game state in the X -pebble-pushing game. \square

Lemma 4.7. *Let G be a finite connected graph and let $X = (x_1, x_2, \dots, x_k)$ be a game state. If G is not k -pebble-win, then there is a two colouring $c: X \rightarrow \{r, b\}$ such that both colour classes are non trivial and for all $i, j \in [k]$ with $c(x_i) = r$ and $c(x_j) = b$ the transposition (ij) is not in the pebble-permutation group.*

Proof. Let us draw a graph H on $\{x_1, x_2, \dots, x_k\}$ by letting (x_i, x_j) be an edge if and only if (ij) is in the pebble-permutation group of (G, X) . It is a simple exercise to show that the pebble-permutation group of (G, X) is S_k if and only if H has a single component.

Since G is not k -pebble-win, we therefore know by Lemma 4.6 that there are at least two components in H . Let us pick one component C_1 and set $c(x) = r$ for all $x \in V(C_1)$ and $c(x) = b$ for all $x \in X \setminus V(C_1)$. \square

Definition 4.8. Given a graph G , a path $x_1x_2 \dots x_m$ in G is a *bare path* if $d_G(x_i) = 2$ for all $2 \leq i \leq m - 1$.

Lemma 4.9. *Let G be a finite connected graph with vertex set V which is not k -pebble-win and with $|V| \geq k + 2$. Then there is a bare path $P = p_1p_2 \dots p_n$ in G such that $|V \setminus V(P)| \leq k$. Furthermore, either every edge in P is a bridge in G , or G is a cycle.*

Proof. Let $X = (x_1, x_2, \dots, x_k)$ be a game state. Since G is not k -pebble-win, by Lemma 4.7 there is a two colouring $c: \{x_i: i \in [k]\} \rightarrow \{r, b\}$ such that both colour classes are non trivial and for all $i, j \in [k]$ with $c(x_i) = r$ and $c(x_j) = b$ the transposition (ij) is not in the pebble permutation group. Let us consider this as a three colouring $c: V \rightarrow \{r, b, 0\}$ where $c(v) = 0$ if $v \notin \{x_1, x_2, \dots, x_k\}$.

For every achievable game state $Z = (z_1, z_2, \dots, z_k)$ in the X -pebble-pushing game we define a three colouring c_Z given by $c_Z(z_i) = c(x_i)$ for all $i \in [k]$ and by $c_Z(v) = 0$ for all $v \notin \{z_1, z_2, \dots, z_k\}$. We note that, for any achievable game state Z there is no $z_i \in c_Z^{-1}(r)$ and $z_j \in c_Z^{-1}(b)$ such that (ij) is in the pebble permutation group of (G, Z) . Indeed, if it were, then by Lemma 4.3 $X^{(ij)}$ is an achievable game state in the X -pebble-pushing game, contradicting the fact that $c(x_i) = r$ and $c(x_j) = b$.

Since G is connected, for every achievable game state Z there is a path $P = p_1p_2 \dots p_m$ in G with $c_Z(p_1) = r$, $c_Z(p_m) = b$ and $c_Z(p_i) = 0$ otherwise. Let us consider an achievable game state Z for which G contains such a path P of maximal length.

We first claim that there is no $v \notin P$ with $c_Z(v) = 0$. Indeed, suppose there is such a vertex v . Since G is connected there is some v - P path Q in G and so, by pushing pebbles towards v on Q , we can achieve a game state Z' such that $c_{Z'} = c_Z$ on P and there is a vertex v' adjacent to P such that $c_{Z'}(v') = 0$. Clearly v' cannot be adjacent to p_1 or p_m , since then we can push the pebble on p_1 or p_m onto v' and achieve a game state Z'' for which G contains a longer path than P with the required colouring. However, if v' is adjacent to p_ℓ with $2 \leq \ell \leq m-1$, then we can push the pebble on p_1 onto p_ℓ and then onto v' , then push the pebble from p_m onto p_1 and finally push the pebble on v' onto p_ℓ and then onto p_m .

However, if $Z' = (z'_1, z'_2, \dots, z'_k)$ with $p_1 = z'_i$ and $p_m = z'_j$, then above shows that (ij) is in the pebble-permutation group of (G, Z') . However, $c_{Z'}(z'_i) = c_Z(p_1) = r$ and $c_{Z'}(z'_j) = c_Z(p_m) = b$, contradicting our assumptions on $c_{Z'}$.

Next, we claim that each p_i with $3 \leq i \leq m-2$ has degree 2. Indeed, suppose first that p_i with $3 \leq i \leq m-2$ is adjacent to some other p_j with $1 \leq j \leq m$ such that p_i and p_j are not adjacent in P . Then it is easy to find a sequence of moves which exchanges the pebbles on p_1 and p_m , contradicting our assumptions on c_Z .

Suppose then that p_i is adjacent to a vertex v not in P . Then, $c_Z(v) \neq 0$, say without loss of generality $c_Z(v) = r$. However then, we can push the pebble on p_m onto p_{i-1} , push the pebble on v onto p_i and then onto p_m and finally push the pebble on p_{i-1} onto p_i and then onto v . As before, this contradicts our assumptions on c_Z .

Hence $P' = p_2 p_3 \dots p_{m-1}$ is a bare path in G , and since every vertex in $V - V(P')$ is coloured using r or using b , there are at most k such vertices.

Finally, suppose that there is some edge in P' which is not a bridge of G , and so no edge of P' is a bridge of G . We wish to show that G is a cycle. We first make the following claim:

Claim 4.10. *There is no achievable game state $W = (w_1, w_2, \dots, w_k)$ such that there is a cycle $C = c_1 c_2 \dots c_r c_1$ and a vertex $v \notin C$ such that:*

- *There exist distinct positive integers i, j, s and t such that $c_W(c_i) = r$, $c_W(c_j) = b$ and $c_W(c_s) = c_W(c_t) = 0$;*
- *v adjacent to some $c_v \in C$.*

Proof of Claim 4.10. Suppose for a contradiction there exists such an achievable game state W . Since C is a cycle we may assume without loss of generality that $c_i = c_1, c_s = c_2 = c_v, c_t = c_3$ and $c_j = c_4$. If $c_W(v) = b$, then we can push the pebble at v to c_2 and then to c_3 , push the pebble at c_1 to c_2 and then to v , and then push the pebble at c_3 to c_1 . This contradicts our assumptions on c_W . The case where $c_W(v) = r$ is similar. Finally

if $c_W(v) = 0$, then we can push the pebble at c_1 to c_2 and then to v , then push the pebble at c_4 to c_1 , then push the pebble at v to c_2 and then to c_4 . Again this contradicts our assumptions on c_W . \square

Since no edge of P' is a bridge, it follows that G contains a cycle C containing P' . If G is not a cycle, then there is a vertex $v \in V \setminus C$ which is adjacent to C . However by pushing the pebble on p_1 onto p_2 and the pebble on p_m onto p_{m-1} , which is possible since $|V| \geq k + 2$, we achieve a game state Z' such that C and v satisfy the assumptions of the above claim, a contradiction. \square

§5. PEBBLY ENDS

Definition 5.1 (Pebbly). Let Γ be a graph and ω an end of Γ . We say ω is *pebbly* if for every $k \in \mathbb{N}$ there is an $n \geq k$ and a family $\mathcal{R} = (R_i : i \in [n])$ of disjoint rays in ω such that $RG(\mathcal{R})$ is k -pebble-win. If for some k there is no such family \mathcal{R} , we say ω is *not k -pebble-win*.

The following is an immediate corollary of Lemma 4.9.

Corollary 5.2. *Let ω be an end of a graph Γ which is not k -pebble-win and let $\mathcal{R} = (R_i : i \in [m])$ be a family of $m \geq k + 2$ disjoint rays in ω . Then there is a bare path $P = p_1 p_2 \dots p_n$ in $RG(R_i : i \in [m])$ such that $|[m] \setminus V(P)| \leq k$. Furthermore, either each edge in P is a bridge in $RG(R_i : i \in [m])$, or $RG(R_i : i \in [m])$ is a cycle.*

Hence, if an end in Γ is not pebbly, then we have some constraint on the behaviour of rays towards this ends. In a later paper [3] we will investigate more precisely what can be said about the structure of the graph towards this end. For now, the following lemma allows us to easily find any countable graph as a minor of a graph with a pebbly end.

Lemma 5.3. *Let Γ be a graph and let $\omega \in \Omega(\Gamma)$ be a pebbly end. Then $K_{\aleph_0} \preceq \Gamma$.*

Proof. By assumption, there exists a sequence $\mathcal{R}_1, \mathcal{R}_2, \dots$ of families of disjoint ω -rays such that, for each $k \in \mathbb{N}$, $RG(\mathcal{R}_k)$ is k -pebble-win. Let us suppose that

$$\mathcal{R}_i = (R_1^i, R_2^i, \dots, R_{m_i}^i) \text{ for each } i \in \mathbb{N}.$$

Let us enumerate the vertices and edges of K_{\aleph_0} by choosing some bijection $\sigma : \mathbb{N} \cup \mathbb{N}^{(2)} \rightarrow \mathbb{N}$ such that $\sigma(i, j) > \sigma(i), \sigma(j)$ for every $\{i, j\} \in \mathbb{N}^{(2)}$ and also $\sigma(1) < \sigma(2) < \dots$. For each $k \in \mathbb{N}$ let G_k be the graph on vertex set $V_k = \{i \in \mathbb{N} : \sigma(i) \leq k\}$ and edge set $E_k = \{\{i, j\} \in \mathbb{N}^{(2)} : \sigma(i, j) \leq k\}$.

We will inductively construct subgraphs H_k of Γ such that H_k is an IG_k extending H_{k-1} . Furthermore for each $k \in \mathbb{N}$ if $V(G_k) = [n]$ then there will be tails T_1, T_2, \dots, T_n

of n distinct rays in \mathcal{R}_n such that for every $i \in [n]$ the tail T_i meets H_k in a vertex of the branch set of i , and is otherwise disjoint from H_k . We will assume without loss of generality that T_i is a tail of R_i^n .

Since $\sigma(1) = 1$ we can take H_1 to be the initial vertex of R_1^1 . Suppose then that $V(G_{n-1}) = [r]$ and we have already constructed H_{n-1} together with appropriate tails T_i of R_i^r for each $i \in [r]$. Suppose firstly that $\sigma^{-1}(n) = r + 1 \in \mathbb{N}$.

Let $X = V(H_{n-1})$. There is a linkage from $(T_i: i \in [r])$ to $(R_1^{r+1}, R_2^{r+1}, \dots, R_r^{r+1})$ after X by Lemma 3.8, and, after relabelling, we may assume this linkage induces the identity on $[r]$. Let us suppose the linkage consists of paths P_i from $x_i \in T_i$ to $y_i \in R_i^{r+1}$.

Since $X \cup \bigcup_i P_i \cup \bigcup_i T_i x_i$ is a finite set, there is some vertex y_{r+1} on R_{r+1}^{r+1} such that the tail $y_{r+1} R_{r+1}^{r+1}$ is disjoint from $X \cup \bigcup_i P_i \cup \bigcup_i T_i x_i$.

To form H_n we add the paths $T_i x_i \cup P_i$ to the branch set of each $i \leq r$ and set y_{r+1} as the branch set for $r + 1$. Then H_n is an IG_n extending H_{n-1} and the tails $y_j R_j^{r+1}$ are as claimed.

Suppose then that $\sigma^{-1}(n) = \{u, v\} \in \mathbb{N}^{(2)}$ with $u, v \leq r$. We have tails T_i of R_i^r for each $i \in [r]$ which are disjoint from H_{n-1} apart from their initial vertices. Let us take tails T_j of R_j^r for each $j > r$ which are also disjoint from H_{n-1} . Since $RG(\mathcal{R}_r)$ is r -pebble-win, it follows that $RG(T_i: i \in [m_r])$ is also r -pebble-win. Furthermore, since by Lemma 3.2 $RG(T_i: i \in [m_r])$ is connected, there is some neighbour $w \in [m_r]$ of u in $RG(T_i: i \in [m_r])$.

Let us first assume that $w \notin [r]$. Since $RG(T_i: i \in [m_r])$ is r -pebble-win, the game state $(1, 2, \dots, v-1, w, v+1, \dots, r)$ is an achievable game state in the $(1, 2, \dots, r)$ -pebble-pushing game and hence by Lemma 4.2 the function φ_1 given by $\varphi_1(i) = i$ for all $i \in [r] \setminus \{v\}$ and $\varphi_1(v) = w$ is a transition function from $(T_i: i \in [r])$ to $(T_i: i \in [m_r])$.

Let us take a linkage from $(T_i: i \in [r])$ to $(T_i: i \in [m_r])$ inducing φ_1 which is after $V(H_{n-1})$. Let us suppose the linkage consists of paths P_i from $x_i \in T_i$ to $y_i \in T_i$ for $i \neq v$ and P_v from $x_v \in T_v$ to $y_v \in T_w$. Let

$$X = V(H_{n-1}) \cup \bigcup_{i \in [r]} P_i \cup \bigcup_{i \in [r]} T_i x_i$$

Since u is adjacent to w in $RG(T_i: i \in [m_r])$ there is a path \hat{P} between $T(T_u, X)$ and $T(T_w, X)$ which is disjoint from X and from all other T_i , say \hat{P} is from $\hat{x} \in T_u$ to $\hat{y} \in T_w$.

Finally, since $RG(T_i: i \in [m_r])$ is r -pebble-win, the game state $(1, 2, \dots, r)$ is an achievable game state in the $(1, 2, \dots, v-1, w, v+1, \dots, r)$ -pebble-pushing game and hence by Lemma 4.2 the function φ_2 given by $\varphi_2(i) = i$ for all $i \in [r] \setminus \{v\}$ and $\varphi_2(w) = v$ is a transition function from $(T_i: i \in [r] \setminus \{v\} \cup \{w\})$ to $(T_i: i \in [m_r])$.

Let us take a further linkage from $(T_i: i \in [r] \setminus \{v\} \cup \{w\})$ to $(T_i: i \in [m_r])$ inducing φ_2 which is after $X \cup \hat{P} \cup T_u \hat{x} \cup y_v T_w \hat{y}$. Let us suppose the linkage consists of paths P'_i from $x'_i \in T_i$ to $y'_i \in T_i$ for $i \in [r] \setminus \{v\}$ and P'_v from $x'_v \in T_w$ to $y'_v \in T_v$.

In the case that $w \in [r]$, $w < v$, say, the game state

$$(1, 2, \dots, w - 1, v, w + 1, \dots, v - 1, w, v + 1, \dots, r)$$

is an achievable game state in the $(1, 2, \dots, r)$ -pebble pushing-game and we get, by a similar argument, all $P_i, x_i, y_i, P'_i, x'_i, y'_i$ and \hat{P} .

We build H_n from H_{n-1} by adjoining the following paths:

- for each $i \neq v$ we add the path $T_i x_i P_i y_i T_i x'_i P'_i y'_i$ to H_{n-1} , adding the vertices to the branch set of i ;
- we add \hat{P} to H_{n-1} , adding the vertices of $V(\hat{P}) \setminus \{\hat{y}\}$ to the branch set of u ;
- we add the path $T_v x_v P_v y_v T_w x'_v P'_v y'_v$ to H_{n-1} , adding the vertices to the branch set of v .

We note that, since $\hat{y} \in y_v T_w x'_v$ the branch sets for u and v are now adjacent. Hence H_n is an IG_n extending H_{n-1} . Finally the rays $y'_i T_i$ for $i \in [r]$ are appropriate tails of the used rays of \mathcal{R}_r . \square

As every countable graph is a subgraph of K_{\aleph_0} , a graph with a pebbly end contains every countable graph as a minor. Thus, as $\aleph_0 G$ is countable, if G is countable, we obtain the following corollary:

Corollary 5.4. *Let Γ be a graph with a pebbly end ω and let G be a countable graph. Then $\aleph_0 G \preceq \Gamma$.*

§6. G -TRIBES AND CONCENTRATION OF G -TRIBES TOWARDS AN END

To show that a given graph G is \preceq -ubiquitous, we shall assume that $nG \preceq \Gamma$ holds for every $n \in \mathbb{N}$ and show that this implies $\aleph_0 G \preceq \Gamma$. To this end we use the following notation for such collections of nG in Γ , most of which we established in [5].

Definition 6.1 (G -tribes). Let G and Γ be graphs.

- A G -tribe in Γ (with respect to the minor relation) is a family \mathcal{F} of finite collections F of disjoint subgraphs H of Γ such that each member H of \mathcal{F} is an IG .
- A G -tribe \mathcal{F} in Γ is called *thick*, if for each $n \in \mathbb{N}$ there is a *layer* $F \in \mathcal{F}$ with $|F| \geq n$; otherwise, it is called *thin*.
- A G -tribe \mathcal{F}' in Γ is a G -subtribe¹ of a G -tribe \mathcal{F} in Γ , denoted by $\mathcal{F}' \preceq \mathcal{F}$, if there is an injection $\Psi: \mathcal{F}' \rightarrow \mathcal{F}$ such that for each $F' \in \mathcal{F}'$ there is an injection

¹When G is clear from the context we will often refer to a G -subtribe as simply a subtribe.

$\varphi_{F'}: F' \rightarrow \Psi(F')$ such that $V(H') \subseteq V(\varphi_{F'}(H'))$ for each $H' \in F'$. The G -subtribe \mathcal{F}' is called *flat*, denoted by $\mathcal{F}' \subseteq \mathcal{F}$, if there is such an injection Ψ satisfying $F' \subseteq \Psi(F')$.

- A thick G -tribe \mathcal{F} in Γ is *concentrated at an end* ϵ of Γ , if for every finite vertex set X of Γ , the G -tribe $\mathcal{F}_X = \{F_X: F \in \mathcal{F}\}$ consisting of the layers $F_X = \{H \in F: H \not\subseteq C(X, \epsilon)\} \subseteq F$ is a thin subtribe of \mathcal{F} . It is *strongly concentrated at* ϵ if additionally, for every finite vertex set X of Γ , every member H of \mathcal{F} intersects $C(X, \epsilon)$.

We note that, every thick G -tribe \mathcal{F} contains a thick subtribe \mathcal{F}' such that every $H \in \bigcup \mathcal{F}$ is a tidy IG . We will use the following lemmas from [5].

Lemma 6.2 (Removing a thin subtribe, [5, Lemma 5.2]). *Let \mathcal{F} be a thick G -tribe in Γ and let \mathcal{F}' be a thin subtribe of \mathcal{F} , witnessed by $\Psi: \mathcal{F}' \rightarrow \mathcal{F}$ and $(\varphi_{F'}: F' \in \mathcal{F}')$. For $F \in \mathcal{F}$, if $F \in \Psi(\mathcal{F}')$, let $\Psi^{-1}(F) = \{F'_F\}$ and set $\hat{F} = \varphi_{F'_F}(F'_F)$. If $F \notin \Psi(\mathcal{F}')$, set $\hat{F} = \emptyset$. Then*

$$\mathcal{F}'' := \{F \setminus \hat{F}: F \in \mathcal{F}\}$$

is a thick flat G -subtribe of \mathcal{F} .

Lemma 6.3 (Pigeon hole principle for thick G -tribes, [5, Lemma 5.3]). *Suppose for some $k \in \mathbb{N}$, we have a k -colouring $c: \bigcup \mathcal{F} \rightarrow [k]$ of the members of some thick G -tribe \mathcal{F} in Γ . Then there is a monochromatic, thick, flat G -subtribe \mathcal{F}' of \mathcal{F} .*

Note that, in the following lemma, it is necessary that G is connected, so that every member of the G -tribe is a connected graph.

Lemma 6.4 ([5, Lemma 5.4]). *Let G be a connected graph and Γ a graph containing a thick G -tribe \mathcal{F} . Then either $\aleph_0 G \preceq \Gamma$, or there is a thick flat subtribe \mathcal{F}' of \mathcal{F} and an end ϵ of Γ such that \mathcal{F}' is concentrated at ϵ .*

Lemma 6.5 ([5, Lemma 5.5]). *Let G be a connected graph and Γ a graph containing a thick G -tribe \mathcal{F} concentrated at an end ϵ of Γ . Then the following assertions hold:*

- (1) *For every finite set X , the component $C(X, \epsilon)$ contains a thick flat G -subtribe of \mathcal{F} .*
- (2) *Every thick subtribe \mathcal{F}' of \mathcal{F} is concentrated at ϵ , too.*

Lemma 6.6. *Let G be a connected graph and Γ a graph containing a thick G -tribe \mathcal{F} concentrated at an end $\epsilon \in \Omega(\Gamma)$. Then either $\aleph_0 G \preceq \Gamma$, or there is a thick flat subtribe of \mathcal{F} which is strongly concentrated at ϵ .*

Proof. Suppose that no thick flat subtribe of \mathcal{F} is strongly concentrated at ϵ . We construct an $\aleph_0 G \preccurlyeq \Gamma$ by recursively choosing disjoint IGs H_1, H_2, \dots in Γ as follows: Having chosen H_1, H_2, \dots, H_n such that for some finite set X_n we have

$$H_i \cap C(X_n, \epsilon) = \emptyset$$

for all $i \in [n]$, then by Lemma 6.5(1), there is still a thick flat subtribe \mathcal{F}'_n of \mathcal{F} contained in $C(X_n, \epsilon)$. Since by assumption, \mathcal{F}'_n is not strongly concentrated at ϵ , we may pick $H_{n+1} \in \mathcal{F}'_n$ and a finite set $X_{n+1} \supseteq X_n$ with $H_{n+1} \cap C(X_{n+1}, \epsilon) = \emptyset$. Then the union of all the H_i is an $\aleph_0 G \preccurlyeq \Gamma$. \square

The following lemma will show that we can restrict ourself to thick G -tribes which are concentrated at thick ends.

Lemma 6.7. *Let G be a connected graph and Γ a graph containing a thick G -tribe \mathcal{F} concentrated at an end $\epsilon \in \Omega(\Gamma)$ which is thin. Then $\aleph_0 G \preccurlyeq \Gamma$.*

Proof. Since ϵ is thin, we know by Proposition 2.4 that only finitely many vertices dominate ϵ . Deleting these yields a subgraph of Γ in which there is still a thick G -tribe concentrated at ϵ . Hence we may assume without loss of generality that ϵ is not dominated by any vertex in Γ .

Let $k \in \mathbb{N}$ be the degree of ϵ . By [8, Corollary 5.5] there is a sequence of vertex sets $(S_n : n \in \mathbb{N})$ such that:

- $|S_n| = k$,
- $C(S_{n+1}, \epsilon) \subseteq C(S_n, \epsilon)$, and
- $\bigcap_{n \in \mathbb{N}} C(S_n, \epsilon) = \emptyset$.

Suppose there is a thick subtribe \mathcal{F}' of \mathcal{F} which is strongly concentrated at ϵ . For any $F \in \mathcal{F}'$ there is an $N_F \in \mathbb{N}$ such that $H \setminus C(S_{N_F}, \epsilon) \neq \emptyset$ for all $H \in F$ by the properties of the sequence. Furthermore, since \mathcal{F}' is strongly concentrated, $H \cap C(S_{N_F}, \epsilon) \neq \emptyset$ as well for each $H \in F$.

Let $F \in \mathcal{F}'$ be such that $|F| > k$. Since G is connected, so is H , and so from the above it follows that $H \cap S_{N_F} \neq \emptyset$ for each $H \in F$, contradicting the fact that $|S_{N_F}| = k < |F|$. Thus $\aleph_0 G \preccurlyeq \Gamma$ by Lemma 6.6. \square

Note that, whilst concentration is hereditary for subtribes, strong concentration is not. However if we restrict to *flat* subtribes, then strong concentration is a hereditary property.

Let us show see how ends of the members of a strongly concentrated tribe relate to ends of the host graph Γ . Let G be a connected graph and $H \subseteq \Gamma$ an IG. By Lemmas 3.2 and 3.4, if $\omega \in \Omega(G)$ and R_1 and $R_2 \in \omega$ then the pullbacks $H(R_1)$ and $H(R_2)$ belong

to the same end $\omega' \in \Omega(\Gamma)$. Hence, H determines for every end $\omega \in G$ a *pullback end* $H(\omega) \in \Omega(\Gamma)$. The next lemma is where we need to use the assumption that G is locally finite.

Lemma 6.8. *Let G be a locally finite connected graph and Γ a graph containing a thick G -tribe \mathcal{F} strongly concentrated at an end $\epsilon \in \Omega(\Gamma)$ where every member is a tidy IG. Then either $\aleph_0 G \preceq \Gamma$, or there is a flat subtribe \mathcal{F}' of \mathcal{F} such that for every $H \in \bigcup \mathcal{F}'$ there is an end $\omega_H \in \Omega(G)$ such that $H(\omega_H) = \epsilon$.*

Proof. Since G is locally finite and every $H \in \bigcup \mathcal{F}$ is tidy, the branch sets $H(v)$ are finite for each $v \in V(G)$. If ϵ is dominated by infinitely many vertices, then we know by Proposition 2.4 that Γ contains a topological K_{\aleph_0} minor, in which case $\aleph_0 G \preceq \Gamma$, since every locally finite connected graph is countable. If this is not the case, then there is some $k \in \mathbb{N}$ such that ϵ is dominated by k vertices and so for every $F \in \mathcal{F}$ at most k of the $H \in F$ contain vertices which dominate ϵ in Γ . Therefore, there is a thick flat subtribe \mathcal{F}' of \mathcal{F} such that no $H \in \bigcup \mathcal{F}'$ contains a vertex dominating ϵ in Γ . Note that \mathcal{F}' is still strongly concentrated at ϵ , and every branch set of every $H \in \bigcup \mathcal{F}'$ is finite.

Since \mathcal{F}' is strongly concentrated at ϵ , for every finite vertex set X of Γ every $H \in \bigcup \mathcal{F}'$ intersects $C(X, \epsilon)$. By a standard argument, since H as a connected infinite graph does not contain a vertex dominating ϵ in Γ , instead H contains a ray $R_H \in \epsilon$.

For a subgraph $K \subseteq H$ let us define K^\downarrow to be the subgraph of G where $V(K^\downarrow) = \varphi_H(K)$ and $(v, w) \in E(K^\downarrow)$ if and only if K contains the edge in H between $H(v)$ and $H(w)$. Then, since each $H(v)$ is finite, $R_H^\downarrow \subseteq G$ is an infinite subgraph of a locally finite connected graph, and hence includes a ray S_H in G . Furthermore, by construction, $V(H(S_H)) \subseteq V(R_H)$ and so the pullback of S_H tends to ϵ in Γ . Hence, if S_H tends to $\omega_H \in \Omega(G)$ then $H(\omega_H) = \epsilon$. \square

§7. UBIQUITY OF MINORS OF THE HALF GRID

Here, and in the following, we denote by \mathbb{H} the infinite, one-ended, cubic hexagonal half grid (see Figure 7.1). The following theorem of Halin is one of the cornerstones of infinite graph theory.

Theorem 7.1 (Halin, see [7, Theorem 8.2.6]). *Whenever a graph Γ contains a thick end, then $\mathbb{H} \preceq \Gamma$.* \square

In [10], Halin used this result to show that every topological minor of \mathbb{H} is ubiquitous with respect to the topological minor relation \preceq . In particular, trees of maximum degree 3 are ubiquitous with respect to \preceq .




FIGURE 7.1. The hexagonal half grid \mathbb{H} .

However, the following argument, which is a slight adaptation of Halin’s, shows that every connected minor of \mathbb{H} is ubiquitous with respect to the minor relation. In particular, the dominated ray, the dominated double ray, and all countable trees are ubiquitous with respect to the minor relation.

The main difference to Halin’s original proof is that, since he was only considering locally finite graphs, he was able to assume that the host graph Γ was also locally finite.

Lemma 7.2 ([10, (4) in Section 3]). $\aleph_0\mathbb{H}$ is a topological minor of \mathbb{H} .

Theorem 1.4. Any connected minor of the half grid $\mathbb{N}\square\mathbb{Z}$ is \preceq -ubiquitous.

Proof. Suppose $G \preceq \mathbb{N}\square\mathbb{Z}$ is a minor of the half grid, and Γ is a graph such that $nG \preceq \Gamma$ for each $n \in \mathbb{N}$. By Lemma 6.4 we may assume there is an end ϵ of Γ and a thick G -tribe \mathcal{F} which is concentrated at ϵ . By Lemma 6.7 we may assume that ϵ is thick. Hence $\mathbb{H} \leq \Gamma$ by Theorem 7.1, and with Lemma 7.2 we obtain

$$\aleph_0 G \preceq \aleph_0(\mathbb{N}\square\mathbb{Z}) \preceq \aleph_0\mathbb{H} \leq \mathbb{H} \leq \Gamma. \quad \square$$

Lemma 7.3. \mathbb{H} contains every countable tree as a minor.

Proof. It is easy to see that the infinite binary tree T_2 embeds into \mathbb{H} as a topological minor. It is also easy to see that countably regular tree T_∞ where every vertex has infinite degree embeds into T_2 as a minor. And obviously, every countable tree T is a subgraph of T_∞ . Hence we have

$$T \subseteq T_\infty \preceq T_2 \leq \mathbb{H}$$

from which the result follows. □

Corollary 7.4. All countable trees are ubiquitous with respect to the minor relation.

Proof. This is an immediate consequence of Lemma 7.3 and Theorem 1.4. □

§8. PROOF OF MAIN RESULTS

Theorem 1.2. *Every locally finite connected graph with nowhere-linear end structure is \preceq -ubiquitous.*

Proof. Let Γ be a graph such that $nG \preceq \Gamma$ holds for every $n \in \mathbb{N}$. Thus Γ contains a thick G -tribe \mathcal{F} . By Lemmas 6.4 and 6.6 we may assume that \mathcal{F} is strongly concentrated at an end ϵ of Γ and so by Lemma 6.8 we may assume that for every $H \in \bigcup \mathcal{F}$ there is an end $\omega_H \in \Omega(G)$ such that $H(\omega_H) = \epsilon$.

Our aim now is to show that ϵ is pebbly, which will complete the proof by Corollary 5.4. So, let us assume for contradiction that there is some $k \in \mathbb{N}$ such that ϵ is not k -pebble-win. Since \mathcal{F} is thick there is some $F \in \mathcal{F}$ which contains $k+2$ disjoint IG s, H_1, H_2, \dots, H_{k+2} . Since by assumption G has nowhere-linear end structure, by Lemma 3.5 for each $i \in [k+2]$ there is a family of disjoint rays $\{R_1^i, R_2^i, \dots, R_{m_i}^i\}$ in G tending to ω_{H_i} whose ray graph in G is not a path. Let

$$\mathcal{S} = (H_i(R_j^i) : i \in [k+2], j \in [m_i]).$$

By construction \mathcal{S} is a disjoint family of rays which tend to ϵ in Γ and by Lemma 3.3 and Lemma 3.4 $RG_\Gamma(\mathcal{S})$ contains disjoint subgraphs K_1, K_2, \dots, K_{k+2} such that $K_i \cong RG_G(R_j^i : j \in [m_i])$. However, by Corollary 5.2, there is a set X of vertices of size at most k such that $RG_\Gamma(\mathcal{S}) - X$ is a bare path P . However, then some $K_i \subseteq P$ is a path, a contradiction. \square

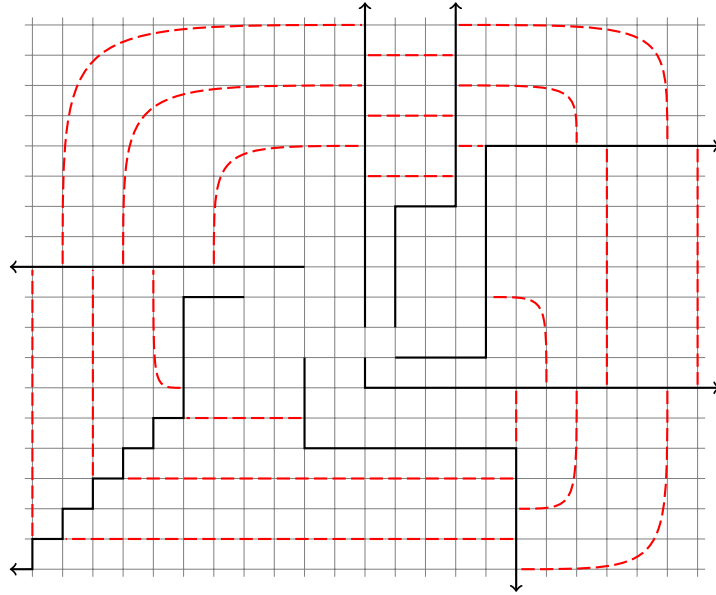


FIGURE 8.1. The ray graphs in the full grid are cycles.

Corollary 1.3. *The full grid is \preceq -ubiquitous.*

Proof. Let G be the full grid. Since $G - R$ has at most one end for any ray $R \in G$, by Lemma 3.2 the ray graph $RG(\mathcal{R})$ is 2-connected for any finite family of three or more rays. Hence, by Theorem 1.2 G is \preceq -ubiquitous \square

Remark 8.1. In fact, every ray graph in the full grid is a cycle (see Figure 8.1).

Theorem 1.5. *For every locally finite connected graph G , both $G \square \mathbb{Z}$ and $G \square \mathbb{N}$ are \preceq -ubiquitous.*

Proof. If G is a path or a ray, then $G \square \mathbb{Z}$ is a subgraph of the half grid $\mathbb{N} \square \mathbb{Z}$ and thus \preceq -ubiquitous by Theorem 1.4. If G is a double ray then $G \square \mathbb{Z}$ is the full grid and thus \preceq -ubiquitous by Corollary 1.3. Otherwise let G' be a finite connected subgraph of G which is not a path. For any end ω of $G \square \mathbb{Z}$ there is a ray R of \mathbb{Z} such that all rays of the form $\{v\} \square R$ for $v \in V(G)$ go to ω . But then G' is a subgraph of $RG_{G \square \mathbb{Z}}(\{(\{v\} \square R)_{v \in V(G')}\})$, so this ray-graph is not a path, hence by Lemma 3.5 $G \square \mathbb{Z}$ has nowhere-linear end structure and is therefore \preceq -ubiquitous by Theorem 1.2. \square

Finally let us prove Theorem 1.6. Recall that for $k \in \mathbb{N}$ let DR_k denote the graph formed by taking a ray R together with k vertices v_1, v_2, \dots, v_k adjacent to every vertex in R .

Theorem 1.6. *The k -fold dominated ray DR_k is \preceq -ubiquitous for every $k \in \mathbb{N}$.*

Proof. Note that if $k \leq 2$ then DR_k is a minor of the half grid, and hence ubiquity follows from Theorem 1.4.

Suppose then that $k \geq 3$ and Γ is a graph which contains a thick DR_k -tribe \mathcal{F} all of whose members are tidy. By Lemma 6.6 we may assume that there is an end ϵ of Γ such that \mathcal{F} is concentrated at ϵ . If there are infinitely many vertices dominating ϵ , then $\aleph_0 DR_k \preceq K_{\aleph_0} \leq \Gamma$ holds by Proposition 2.4. So we may assume that only finitely many vertices dominate ϵ . By taking a thick subtribe if necessary, we may assume that no member of \mathcal{F} contains such a vertex.

As before, if we can show that ϵ is pebbly, then we will be done by Corollary 5.4. So suppose for a contradiction that ϵ is not r -pebble-win for some $r \in \mathbb{N}$.

Let R be the ray as stated in the definition of DR_k and let $v_1, v_2, \dots, v_k \in V(DR_k)$ be the vertices adjacent to each vertex of R . For each $i \in [k]$ let S_i denote the star in $V(DR_k)$ consisting of v_i and all incident edges. For each $H \in \bigcup \mathcal{F}$ and each $i \in [k]$ we have that $H(S_i)$ is a locally finite infinite tree since H is tidy and any vertex of $H(S_i)$ whose degree is infinite would dominate ϵ . So $H(S_i)$ includes a ray, call it $R_{H,i}$,

by Proposition 2.3. Let $R_H = H(R)$ be the pullback of the ray R in H . Now we set $\mathcal{R}_H = (R_{H,1}, R_{H,2}, \dots, R_{H,k}, R_H)$.

Since all leaves of $H(S_i)$ are in branch sets of vertices of R , it follows that in the graph $RG_H(\mathcal{R}_H)$ each $R_{H,i}$ is adjacent to R_H . Hence $RG_H(\mathcal{R}_H)$ contains a vertex of degree $k \geq 3$.

There is some layer $F \in \mathcal{F}$ of size $\ell \geq r + 1$, say $F = (H_i : i \in [\ell])$. For every $i \in [r + 1]$ we set $\mathcal{R}_{H_i} = (R_{H_i,1}, R_{H_i,2}, \dots, R_{H_i,k}, R_{H_i})$. Let us now consider the family of disjoint rays

$$\mathcal{R} = \bigcup_{i=1}^{r+1} \mathcal{R}_{H_i}.$$

By construction \mathcal{R} is a family of disjoint rays which tend to ϵ in Γ and by Lemma 3.3 and Lemma 3.4 $RG_\Gamma(\mathcal{R})$ contains $r + 1$ vertices whose degree is at least $k \geq 3$. However, by Corollary 5.2, there is a vertex set X of size at most r such that $RG_\Gamma(\mathcal{R}) - X$ is a bare path P . But then some vertex whose degree is at least 3 is contained in the bare path, a contradiction. \square

REFERENCES

- [1] T. Andreae. On disjoint configurations in infinite graphs. *Journal of Graph Theory*, 39(4):222–229, 2002.
- [2] T. Andreae. Classes of locally finite ubiquitous graphs. *Journal of Combinatorial Theory, Series B*, 103(2):274–290, 2013.
- [3] N. Bowler, C. Elbracht, J. Erde, P. Gollin, K. Heuer, M. Pitz, and M. Teegen. Ubiquity in graphs III: Ubiquity of locally finite graphs with extensive tree decompositions. in preparation.
- [4] N. Bowler, C. Elbracht, J. Erde, P. Gollin, K. Heuer, M. Pitz, and M. Teegen. Ubiquity in graphs IV: Ubiquity of graphs of bounded tree-width. in preparation.
- [5] N. Bowler, C. Elbracht, J. Erde, P. Gollin, K. Heuer, M. Pitz, and M. Teegen. Ubiquity in graphs I: Topological ubiquity of trees. *arXiv preprint arXiv:1806.04008*, 2018.
- [6] G. Călinescu, A. Dumitrescu, and J. Pach. Reconfigurations in graphs and grids. *SIAM Journal on Discrete Mathematics*, 22(1):124–138, 2008.
- [7] R. Diestel. *Graph Theory*. Springer, 5th edition, 2016.
- [8] P. Gollin and K. Heuer. k -connected sets in infinite graphs. in preparation.
- [9] R. Halin. Über die Maximalzahl fremder unendlicher Wege in Graphen. *Mathematische Nachrichten*, 30(1-2):63–85, 1965.
- [10] R. Halin. A problem in infinite graph-theory. In *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, volume 43, pages 79–84. Springer, 1975.
- [11] D. Kornhauser, G. Miller, and P. Spirakis. Coordinating pebble motion on graphs, the diameter of permutation groups, and applications. In *Proceedings of the 25th Annual Symposium on Foundations of Computer Science, 1984*, SFCS '84, pages 241–250. IEEE Computer Society, 1984.
- [12] R. Laver. On Fraïssé's order type conjecture. *Annals of Mathematics*, pages 89–111, 1971.

- [13] R. Thomas. A counter-example to ‘Wagner’s conjecture’ for infinite graphs. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 103, pages 55–57. Cambridge University Press, 1988.