THE FAREY GRAPH IS UNIQUELY DETERMINED BY ITS CONNECTIVITY

JAN KURKOFKA

ABSTRACT. We show that, up to minor-equivalence, the Farey graph is the unique minor-minimal graph that is infinitely edge-connected but such that every two vertices can be finitely separated.

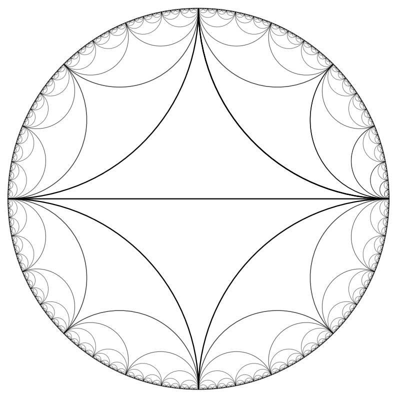

FIGURE 1. The Farey graph

1. INTRODUCTION

The Farey graph, shown in Figure 1 and surveyed in [1,4], plays a role in a number of mathematical fields ranging from group theory and number theory to geometry and dynamics [1]. Curiously, graph theory has not been among these until very recently, when it was shown that the Farey graph plays a central role in graph theory too: it is one of two infinitely edge-connected graphs that must occur as a minor in every infinitely edge-connected graph [5]. Infinite edge-connectivity, however, is only one aspect of the connectivity of the Farey graph, and it contrasts with a second aspect: the Farey graph does not contain infinitely many independent paths between any two of its vertices. In this paper we show that the Farey graph is uniquely determined by these two contrasting aspects of its connectivity: up to minor-equivalence, the Farey graph is the unique minor-minimal graph that is infinitely edge-connected but such that every two vertices can be finitely separated. This is the first graph-theoretic characterisation of the Farey graph.

A Π-*graph* is an infinitely edge-connected graph that does not contain infinitely many independent paths between any two of its vertices. A Π-graph is *typical* if it occurs as a minor in every Π-graph. Note that any two typical Π-graphs are minors of each other; we call such graphs *minor-equivalent*. Our main result reads as follows:

Theorem 1. *Up to minor-equivalence, the Farey graph is the unique typical* Π*-graph.*

We shall see that there exist Π-graphs that contain the Farey graph as a minor but are not minors of the Farey graph (Theorem 3.1).

2020 *Mathematics Subject Classification.* 05C63, 05C40, 05C83, 05C10.

Key words and phrases. infinite graph; Farey graph; characterisation; connectivity; infinite edge-connectivity; infinitely edge-connected graph; minor; typical pi-graph.

Theorem 1 continues to hold if we require all minors to have finite branch sets; see Section 3.2 and Theorem 3.2. This is best possible in the sense that one cannot replace 'minors with finite branch sets' with 'topological minors' (Theorem 3.3).

This paper is organised as follows. Section 2 formally introduces the Farey graph. In Section 3 we prove Theorems 3.1–3.3. We outline the overall strategy of the proof of Theorem 1 in Section 4. We prepare the proof of Theorem 1 in Section 5 and we prove Theorem 1 in Section 6.

2. Preliminaries

We use the notation of Diestel's book [2]. A non-trivial path P is an A-*path* for a set A of vertices if P has its endvertices but no inner vertex in A. Two u–v paths are *order-compatible* if they traverse their common vertices in the same order.

The *Farey graph* F is the graph on $\mathbb{Q} \cup \{\infty\}$ in which two rational numbers a/b and c/d in lowest terms (allowing also $\infty = (\pm 1)/0$) form an edge if and only if $\det\left(\begin{smallmatrix} a & c \\ b & d \end{smallmatrix}\right) = \pm 1$, cf. [1]. In this paper we do not distinguish between the Farey graph and the graphs that are isomorphic to it. For our graph-theoretic proofs it will be more convenient to work with the following purely combinatorial definition of the Farey graph that is indicated in [1] and [4].

The *halved Farey graph* $\breve{F}_0$ of order 0 is a K^2 with its sole edge coloured blue. Inductively, the *halved Farey graph* $\breve{F}_{n+1}$ of order $n+1$ is the edge-coloured graph that is obtained from $\breve{F}_n$ by adding a new vertex v_e for every blue edge e of $\breve{F}_n$, joining each v_e precisely to the endvertices of e by two blue edges, and colouring all the edges of $\breve{F}_n \subseteq \breve{F}_{n+1}$ black. The *halved Farey graph* $\breve{F} := \bigcup_{n \in \mathbb{N}} \breve{F}_n$ is the union of all $\breve{F}_n$ without their edge-colourings, and the *Farey graph* is the union $F = G_1 \cup G_2$ of two copies G_1, G_2 of the halved Farey graph such that $G_1 \cap G_2 = \breve{F}_0$.

Lemma 2.1. *The halved Farey graph contains the Farey graph as a minor with finite branch sets.*

Proof. If e is a blue edge of $\breve{F}_1$, then the Farey graph is the contraction minor of $\breve{F} - e$ whose sole non-trivial branch set is $V(\breve{F}_0)$, i.e., $(\breve{F} - e)/V(\breve{F}_0) \cong F$. □

3. Atypical Π-graphs and variations of the main result

In this section we provide details on and prove the three Theorems 3.1–3.3 that we briefly mentioned in the introduction.

3.1. Atypical Π-graphs.

Even though every Π-graph contains the Farey graph as a minor by Theorem 1, the converse is generally false:

Theorem 3.1. *There exist* Π*-graphs that contain the Farey graph as a minor but are not minors of the Farey graph.*

Proof. Let the graph G be obtained from some union of uncountably many disjoint copies of the Farey graph by selecting one vertex in every copy and identifying all the selected vertices. Then G is an uncountable Π-graph that contains the Farey graph as a subgraph. However, G is not a minor of the Farey graph, because every minor of the Farey graph must be countable. □

3.2. **Variations of the main result.** To prove Theorem 1 it suffices to show the following theorem. A *tight* minor is a minor with finite branch sets.

Theorem 2. *Every* Π*-graph contains the Farey graph as a tight minor.*

Theorem 2 also implies the following variation of Theorem 1 where all minors are required to have finite branch sets. Two graphs are *tightly* minor-equivalent if they are tight minors of each other. A Π-graph is *tightly* typical if it occurs as a tight minor in every Π-graph.

Theorem 3.2. *Up to tight minor-equivalence, the Farey graph is the unique tightly typical* Π*-graph.* □

This raises the question whether Theorem 1 continues to hold if we require all minors to be topological minors. We answer this question in the negative:

Theorem 3.3. *There is a* Π*-graph that contains the Farey graph as a tight minor but not as a topological minor.*

Proof. By a recent result [7] there exists an infinitely edge-connected graph G that does not contain infinitely many edge-disjoint pairwise order-compatible paths between any two of its vertices; in particular, G is a Π-graph. By Theorem 2, the graph G contains the Farey graph as a tight minor. However, G does not contain a subdivision of the Farey graph because the Farey graph contains infinitely many edge-disjoint pairwise order-compatible paths between any two of its vertices. □

4. Overall proof strategy

Our aim for the remainder of this paper is to prove Theorem 1. As we discussed in the previous section, to prove Theorem 1 it suffices to show that every Π-graph contains the Farey graph as a minor with finite branch sets (Theorem 2). And by Lemma 2.1 in turn it suffices to find a halved Farey graph minor with finite branch sets in any given Π-graph. The key idea of the proof is summarised in Theorem 6.1 which states:

Suppose that G is any subdivided Π-graph and that u, v are two distinct branch vertices of G. Then there exist subgraphs $H_u, H_v \subseteq G$ that satisfy the following conditions:

(i) *$H_u[X] = H_v[X]$ is finite and connected for $X := V(H_u) \cap V(H_v) \neq \emptyset$;*
(ii) *X avoids u and v;*
(iii) *both H_u/X and H_v/X are subdivided Π-graphs in which u, X and v, X are branch vertices, respectively;*
(iv) *uX is an edge of H_u/X and vX is an edge of H_v/X.*

With this theorem at hand, it is straightforward to construct a halved Farey graph minor with finite branch sets in any given Π-graph G: Consider any edge uv of G and apply the theorem in G to u and v to obtain subgraphs H_u, H_v and a non-empty finite connected vertex set $X \subseteq V(G)$. Then the three vertices u, v and X span a triangle $\check{F}_1$ in $(H_u \cup H_v)/X$. And since both H_u/X and H_v/X are subdivided Π-graphs, we can reapply the theorem in H_u/X to u and X, and in H_v/X to v and X. By iterating this process, we obtain a halved Farey graph minor with finite branch sets in the original graph G at the limit, and this will complete the proof. Therefore, it remains to prove Theorem 6.1 on the one hand, and to use it to formally construct a halved Farey graph minor on the other hand. In the next section, we prepare the proof of Theorem 6.1, and in the section after next we prove Theorem 6.1 which we then use to prove Theorems 1 and 2.

5. Grain lines

It is possible to prove Theorem 6.1 from first principles. In this paper, however, I favour a more methodic proof. The advantage of this proof is that it introduces a new tool, an x–y grain line, that allows one to control infinite systems of edge-disjoint x–y paths even when no two paths in the system are pairwise order-compatible. In this section we introduce the concept of an x–y grain line, we show that these exist whenever it matters (Theorem 5.4) and we show two lemmas that will help us prove Theorem 6.1 using grain lines at the beginning of the next section.

Informally, we may think of an x–y grain line as a pair $(L,\mathcal{P})$ where $\mathcal{P}$ is a sequence of pairwise edge-disjoint x–y paths $P_0,P_1,\ldots$ that need not be pairwise order-compatible but solve all incompatibilities at their linearly ordered 'limit' L. The limit L will not be a graph-theoretic path but will be a linearly ordered set of vertices. We remark, however, that it is possible to use the limit L to define a topological x–y path in a topological extension of any graph containing the grain line, see [6, §6.3].

Here is the formal definition of an x–y grain-line:

Definition 5.1. An x–y *grain line* between two distinct vertices x and y is an ordered pair $(L,\mathcal{P})$ where $L=(L,\le_L)$ is a linearly ordered countable set of vertices with least element x and greatest element y, and $\mathcal{P}=(P_n)_{n\in\mathbb{N}}$ is a sequence of pairwise edge-disjoint x–y paths P_n, such that the following three conditions are satisfied:

(GL1) $L=\left\{ v \;\middle|\; \{\, n\in\mathbb{N} : v\in V(P_n)\,\} \text{ is a final segment of } \mathbb{N} \right\}$;

(GL2) if a vertex of a path P_n is not contained in L, then it is not a vertex of any other path P_m $(m\neq n)$;

(GL3) for all $n\in\mathbb{N}$, the x–y path P_n and the linearly ordered vertex set L induce the same linear ordering on the vertex set $L_{<n}:=L\cap\bigcup_{k<n}V(P_k)$.

We remark that (GL3) allows P_n and L to induce distinct linear orderings on the vertex set $V(P_n)\cap L$ if the inclusion $L_{<n}\subseteq V(P_n)\cap L$ is proper; in particular, P_n and P_{n+1} need not be order-compatible. Allowing this becomes necessary, for example, if an infinitely edge-connected graph does not contain infinitely many edge-disjoint pairwise order compatible paths between x and y, see Example 5.3.

Clearly, $L=\bigcup_n L_{<n}$. Note that if $(L,(P_n)_{n\in\mathbb{N}})$ is a grain line, then a vertex v lies in L if and only if it lies on all paths P_n with $n\ge N$ for N the first number with $v\in P_N$ if and only if it lies on at least two paths P_n,P_m $(n\neq m)$. In particular,

$$V(P_n)\cap\bigcup_{k<n}V(P_k)=L_{<n}\text{ for all }n\in\mathbb{N}.$$

We speak of an x–y grain line $(L,(P_n)_{n\in\mathbb{N}})$ *in* a graph G if $\bigcup_{n\in\mathbb{N}}P_n\subseteq G$ (and hence $L\subseteq V(G)$). Whenever a grain line is introduced as $(L,\mathcal{P})$, we tacitly assume $\mathcal{P}=(P_n)_{n\in\mathbb{N}}$. In general, however, we also allow sequences $\mathcal{P}=(P_n)_{n\ge N}$ whose indexing starts at an arbitrary number $N>0$ in which case the definition of a grain line adapts in the obvious way. We use the interval notation for L as usual, i.e., we write $[\ell_1,\ell_2]_L=\{\,\ell\in L\mid \ell_1\le_L\ell\le_L\ell_2\,\}$ and so on.

Example 5.2. The blue Hamilton paths $P_n \subseteq \breve{F}_n$ are pairwise edge-disjoint and order-compatible, and hence give rise to an x–y grain line in $\breve{F}$ for x and y the two vertices of $\breve{F}_0$. In this case, $L = V(\breve{F})$ is order-isomorphic to $\mathbb{Q} \cap [0,1]$.

Example 5.3. There exists an infinitely edge-connected graph G that does not contain infinitely many edge-disjoint pairwise order-compatible paths between any two of its vertices [7]; in particular, G is a Π-graph. We shall see that the graph G contains a grain line between any two of its vertices because it is infinitely edge-connected; see Theorem 5.4 below. However, since G does not contain infinitely many edge-disjoint pairwise order-compatible paths between any two of its vertices, every grain line $(L, \mathcal{P})$ in G has two paths P_n and P_{n+1} that are not order-compatible; in particular, P_n induces the same linear ordering on $L_{<n} \subsetneq V(P_n) \cap L$ as L does, but disagrees with L on $V(P_n) \cap L = V(P_n) \cap V(P_{n+1})$ because L induces the ordering of P_{n+1} on $V(P_n) \cap V(P_{n+1}) = L_{<n+1}$. This is why we do not strengthen (GL3) to require that P_n and L induce the same linear ordering on $V(P_n) \cap L \supseteq L_{<n}$.

Our first result on grain lines shows that they exist whenever it matters:

Theorem 5.4. *Let x and y be any two distinct vertices of a graph G. Then there exists an x–y grain line in G if and only if G contains infinitely many edge-disjoint x–y paths.*

In the proof we employ inverse systems, and for the sake of convenience we dedicate a paragraph to their definition.

A partially ordered set $(I, \leq)$ is said to be *directed* if for every two $i, j \in I$ there is some $k \in I$ with $k \geq i, j$. Let $(\, X_i \mid i \in I\,)$ be a family of finite sets indexed by some directed poset $(I, \leq)$. Furthermore, suppose that we are given a family $(\, \varphi_{ji} \colon X_j \to X_i\,)_{i \leq j \in I}$ of mappings which are the identity on X_i in case of $i = j$ and which are *compatible* in that $\varphi_{ki} = \varphi_{ji} \circ \varphi_{kj}$ for all $i \leq j \leq k$. Then both families together are said to form an *inverse system* (*of finite sets*), and the maps φ_{ji} are called its *bonding maps*. We denote such an inverse system by $\{X_i, \varphi_{ji}, I\}$ or $\{X_i, \varphi_{ji}\}$ for short if I is clear from context. Its *inverse limit* $\varprojlim X_i = \varprojlim (\, X_i \mid i \in I\,)$ is the set

$$\varprojlim X_i = \{\, (x_i)_{i \in I} \mid \varphi_{ji}(x_j) = x_i \text{ for all } j \geq i \,\} \subseteq \prod_{i \in I} X_i.$$

If every X_i is non-empty, then the inverse limit $\varprojlim X_i$ is non-empty as well. For more details on inverse systems and their more general definition for topological spaces, see [3] or [8].

Proof of Theorem 5.4. Every x–y grain line comes with a system of infinitely many edge-disjoint x–y paths. For the backward implication let x and y be given, and let $\mathcal{Q}$ be any countably infinite collection of edge-disjoint x–y paths in G. Moreover, we let $\mathcal{X}$ be the collection of all finite subsets of the vertex set of the subgraph $\bigcup \mathcal{Q} \subseteq G$, directed by inclusion.

Given $X \in \mathcal{X}$ we write $\mathrm{lin}(X)$ for the finite collection of all linearly ordered subsets of X. Letting, for all $X \subseteq X' \in \mathcal{X}$, the map $\varphi_{X',X} \colon \mathrm{lin}(X') \to \mathrm{lin}(X)$ take every linearly ordered subset of X' to its restriction with respect to X turns the finite sets $\mathrm{lin}(X)$ into an inverse system $\{\, \mathrm{lin}(X),\, \varphi_{X',X},\, \mathcal{X}\,\}$.

Every x–y path $P \in \mathcal{Q}$ naturally induces a linear ordering $\leq_P$ on its vertex set with $x <_P y$, and for every $X \in \mathcal{X}$ we denote by $\leq_P^X$ the linear ordering on

$V(P) \cap X$ induced by $\leq_P$. Then for every $X \in \mathcal{X}$ we define a map $\psi_X \colon \mathcal{Q} \to \operatorname{lin}(X)$ by letting

$$\psi_X(P) := (V(P) \cap X, \leq_P^X)$$

for all $P \in \mathcal{Q}$, and we put

$$\mathcal{L}_X := \{\, \xi \in \operatorname{lin}(X) \mid \psi_X^{-1}(\xi) \subseteq \mathcal{Q} \text{ is infinite} \,\}$$

noting that $\mathcal{L}_X \subseteq \operatorname{lin}(X)$ is non-empty by the pigeonhole principle. Since the maps ψ_X commute with the bonding maps $\varphi_{X',X}$ as pictured in the diagram below,

$$\begin{tikzcd}
 & \mathcal{Q} \arrow[dl, "\psi_X"'] \arrow[dr, "\psi_{X'}"] & \\
\operatorname{lin}(X) & & \operatorname{lin}(X') \arrow[ll, "\varphi_{X',X}"']
\end{tikzcd}$$

the restrictions of these bonding maps to the sets $\mathcal{L}_X$ yield another inverse system, namely $\{\, \mathcal{L}_X,\, \varphi_{X',X} \restriction \mathcal{L}_{X'},\, \mathcal{X} \,\}$. And as the finite sets $\mathcal{L}_X$ are all non-empty, this inverse system has an element $(\, (L_X, \leq_X) \mid X \in \mathcal{X} \,)$ in its limit.

Finally, we define an x–y grain line $(L, \mathcal{P})$, as follows. We let $L := \bigcup_{X \in \mathcal{X}} L_X$ and $\leq_L := \bigcup_{X \in \mathcal{X}} \leq_X$. To obtain $\mathcal{P} = (P_n)_{n \in \mathbb{N}}$ we choose pairwise edge-disjoint x–y paths $P_0, P_1, \ldots$ from $\mathcal{Q}$ inductively, as follows. Choose an enumeration $x_0, x_1, \ldots$ of the countable vertex set $\bigcup \mathcal{X}$ of $\bigcup \mathcal{Q}$. At step 0, we let $X_0 := \{x_0\}$ and choose $P_0 \in \psi_{X_0}^{-1}(L_{X_0})$ arbitrarily (we abbreviate $L_X = (L_X, \leq_X)$). At step $n+1$, we let $X_{n+1} := X_n \cup V(P_n) \cup \{x_{n+1}\}$ and we pick from the infinite preimage $\psi_{X_{n+1}}^{-1}(L_{X_{n+1}})$ a path P_{n+1} other than the previously chosen paths $P_0, \ldots, P_n$. It is straightforward to check that $(L, \mathcal{P})$ is an x–y grain line in G. $\square$

A grain line $(L, \mathcal{P})$ is *wild* if L is order-isomorphic to $\mathbb{Q} \cap [0,1]$. We call a grain line $(L, \mathcal{P})$ *wildly presented* if, for every $n \in \mathbb{N}$, whenever $\ell_1 <_L \ell_2$ are elements of $L_{<n} \subseteq L$ then $\mathring{\ell}_1 P_n \mathring{\ell}_2$ has a vertex in $(\ell_1, \ell_2)_L$. The grain line in Example 5.2 is both wild and wildly presented. Wildly presented grain lines are wild. Conversely, if a grain line $(L, \mathcal{P})$ is wild, then $\mathcal{P} = (P_n)_{n \in \mathbb{N}}$ has a subsequence $(P_{n_k})_{k \in \mathbb{N}}$ such that $(L, (P_{n_k})_{k \in \mathbb{N}})$ is wildly presented.

Lemma 5.5. *Every grain line in a subdivided* Π*-graph is wild; in particular, in a subdivided* Π*-graph every grain line can be chosen to be wildly presented.*

In the proof we use the following properties of grain lines. Given a grain line $(L, \mathcal{P})$ we say that a path P_n does $(L, \mathcal{P})$-*grain* a set U of vertices if, for all $m \geq n$, we have $V(P_m) \cap U = L \cap U$ and the path P_m induces the same linear ordering on this intersection as L does. If $(L, \mathcal{P})$ is clear from context, we also say that P_n *grains* U. Every path P_n grains the union $\bigcup_{k<n} V(P_k)$ by (GL3). And for every finite vertex set X there is a number $n \in \mathbb{N}$ such that P_n grains X. We will use this latter property frequently in the proofs to come.

Proof of Lemma 5.5. Suppose that $(L, \mathcal{P})$ is any grain line in some given subdivided Π-graph G. It suffices to show that $(L, \mathcal{P})$ is wild. For this, consider any two elements $\ell_1, \ell_2 \in L$ with $\ell_1 <_L \ell_2$. Then ℓ_1 and ℓ_2 must have infinite degree in G; in particular, ℓ_1 and ℓ_2 must be branch vertices of G. Since G is a subdivided Π-graph, we find a finite vertex set $S \subseteq V(G) \setminus \{\ell_1, \ell_2\}$ that separates ℓ_1 and ℓ_2 in $G - \ell_1\ell_2$. Then we pick $N \in \mathbb{N}$ such that P_N avoids the edge $\ell_1\ell_2$ and grains the finite vertex set $S \cup \{\ell_1, \ell_2\}$. Now $\ell_1 P_N \ell_2$ must meet S in a vertex s, and then P_N graining $S \cup \{\ell_1, \ell_2\}$ implies $s \in L$ with $\ell_1 <_L s <_L \ell_2$ as desired. □

Grain lines can be restricted such that the restriction is again a grain line, and restricting a grain line preserves wild presentations:

Lemma 5.6. *If $(L, \mathcal{P})$ is a grain line with $\ell_1 <_L \ell_2$ and $N \in \mathbb{N}$ is such that P_N grains $\{\ell_1, \ell_2\}$, then $([\ell_1, \ell_2]_L, (\ell_1 P_n \ell_2)_{n \geq N})$ is an ℓ_1–ℓ_2 grain line that is wildly presented if $(L, \mathcal{P})$ is.*

Proof. First, we show that $([\ell_1, \ell_2]_L, (\ell_1 P_n \ell_2)_{n \geq N})$ is an ℓ_1–ℓ_2 grain line.

(GL1) We have to show the equality

$$[\ell_1, \ell_2]_L = \left\{ v \;\middle|\; \{\, n \in \mathbb{N}_{\geq N} : v \in V(\ell_1 P_n \ell_2) \,\} \text{ is a final segment of } \mathbb{N}_{\geq N} \right\}.$$

We start with the backward inclusion. If a vertex v lies on $\ell_1 P_n \ell_2$ for all n in some final segment of $\mathbb{N}_{\geq N}$ then it lies in L by (GL2) for $(L, \mathcal{P})$, and in particular it also lies on $\ell_1 P_n \ell_2$ when P_n does $(L, \mathcal{P})$-grain $\{\ell_1, v, \ell_2\}$ so $v \in [\ell_1, \ell_2]_L$ follows. Conversely, if v is a vertex in $[\ell_1, \ell_2]_L$ and $k \geq N$ is minimal with $v \in \ell_1 P_k \ell_2$, then P_{k+1} does $(L, \mathcal{P})$-grain $\{\ell_1, v, \ell_2\}$. Therefore, v is contained in $\ell_1 P_n \ell_2$ for all $n \geq k$, and hence $\mathbb{N}_{\geq k}$ witnesses that v is contained in the right hand side of the equation.

(GL2) Consider any vertex $v \in (\bigcup_{n \geq N} \ell_1 P_n \ell_2) - [\ell_1, \ell_2]_L$ and let $k \geq N$ be minimal such that $\ell_1 P_k \ell_2$ contains v. If v is not contained in L, then P_k is the only path from $\mathcal{P}$ containing v, and hence $\ell_1 P_k \ell_2$ is the only path from $(\ell_1 P_n \ell_2)_{n \geq N}$ containing v. Otherwise v is contained in $L \setminus [\ell_1, \ell_2]_L$ so, say, $\ell_2 <_L v$. Then, as P_n with $n > k$ does $(L, \mathcal{P})$-grain $V(P_k)$, the vertex ℓ_2 precedes v on P_n, giving $v \notin \ell_1 P_n \ell_2$ as desired.

(GL3) Consider any $n \geq N$ and write $L'_{<n} := [\ell_1, \ell_2]_L \cap \bigcup_{k=N}^{n-1} V(\ell_1 P_k \ell_2)$. By the already shown (GL1) we have $L'_{<n} \subseteq V(\ell_1 P_n \ell_2)$, so $\ell_1 P_n \ell_2$ does induce a linear ordering on $L'_{<n}$, and it coincides with the linear ordering induced by $[\ell_1, \ell_2]_L$ by (GL3) for $(L, \mathcal{P})$.

Therefore, $([\ell_1, \ell_2]_L, (\ell_1 P_n \ell_2)_{n \geq N})$ is an ℓ_1–ℓ_2 grain line; now we show that it is wildly presented if $(L, \mathcal{P})$ is. For this consider any $n \geq N$ with some two elements $\ell <_L \ell'$ of $L'_{<n}$. Then, as $(L, \mathcal{P})$ is wildly presented and $L'_{<n} \subseteq L_{<n}$, the subpath $\mathring{\ell} P_n \mathring{\ell}'$ of $\ell_1 P_n \ell_2$ has a vertex in $(\ell, \ell')_L$. □

6. Proof of the main result

In this section, we employ our results on grain lines to prove Theorem 6.1, which we then use to prove Theorems 1 and 2.

Theorem 6.1. *Suppose that G is any subdivided Π-graph and that u, v are two distinct branch vertices of G. Then there exist subgraphs $H_u, H_v \subseteq G$ that satisfy the following conditions:*

(i) *$H_u[X] = H_v[X]$ is finite and connected for $X := V(H_u) \cap V(H_v) \neq \emptyset$;*
(ii) *X avoids u and v;*
(iii) *both H_u/X and H_v/X are subdivided Π-graphs in which u, X and v, X are branch vertices, respectively;*
(iv) *uX is an edge of H_u/X and vX is an edge of H_v/X.*

Proof. Without loss of generality we may assume that uv is not an edge of G. Using that G is a subdivided Π-graph we find a finite vertex set $S \subseteq V(G) \setminus \{u, v\}$ that separates u and v in G. We write C_u and C_v for the two distinct components of $G-S$ that contain u and v respectively. Next, we use Theorem 5.4 and Lemma 5.5 to find a wildly presented u–v grain line $(L,\mathcal{P})$ in G. Without loss of generality we may assume that P_0 grains the finite vertex set S. We let s_u be the first vertex of the u–v path P_0 in S, and we let s_v be the last vertex of P_0 in S. That is to say that s_u and s_v are the least and greatest vertex of L in S. Then, for all $n\in\mathbb{N}$, the paths uP_ns_u and s_vP_nv are contained in $G[C_u+s_u]$ and $G[s_v+C_v]$ respectively.

Next, we let x_u and x_v be the least and greatest vertex of L in $V(\mathring{P}_0)$. Moreover, we let $L_u:=[u,x_u]_L$ and $\mathcal{P}_u:=(uP_nx_u)_{n\geq 1}$, and we let $L_v:=[x_v,v]_L$ and $\mathcal{P}_v:=(x_vP_nv)_{n\geq 1}$. Then $(L_u,\mathcal{P}_u)$ and $(L_v,\mathcal{P}_v)$ are wildly presented u–x_u and x_v–v grain lines in G by Lemma 5.6. We claim that $H_u:=P_0\mathring{v}\cup\bigcup\mathcal{P}_u$ and $H_v:=\mathring{u}P_0\cup\bigcup\mathcal{P}_v$ are the desired subgraphs.

First, we show that $X=V(\mathring{P}_0)$ and that X satisfies (i), (ii) and (iv). For this, it suffices to show that for every $n\geq 1$ the paths uP_nx_u and x_vP_nv are u–$\mathring{P}_0$ and $\mathring{P}_0$–v paths in $G[C_u+s_u]$ and $G[s_v+C_v]$, respectively. The vertex $s_u\in L\cap S\subseteq L\cap V(\mathring{P}_0)$ was a candidate for x_u, implying $x_u\leq_L s_u$, and then for all $n\geq 1$ the path P_n graining $V(P_0)$ gives $uP_nx_u\subseteq uP_ns_u\subseteq G[C_u+s_u]$ on the one hand and that x_u is the first vertex of P_n in $\mathring{P}_0$ on the other hand; for the paths x_vP_nv we employ symmetry.

(iii) follows from the facts that $(L_u,\mathcal{P}_u)$ and $(L_v,\mathcal{P}_v)$ are wildly presented and that all paths uP_nx_u and x_vP_nv $(n\geq 1)$ are u–$\mathring{P}_0$ and $\mathring{P}_0$–v paths respectively. □

Now we have almost all we need to prove Theorems 1 and 2. In the proof of Theorem 2, we will face the construction of a minor with finite branch sets in countably many steps. The following notation and lemma will help us to keep the technical side of this construction to the minimum.

Suppose that G and H are two graphs with H a minor of G. Then there are a vertex set $U\subseteq V(G)$ and a surjection $f\colon U\to V(H)$ such that the preimages $f^{-1}(x)\subseteq U$ form the branch sets of a model of H in G. A *minor-map* $\varphi\colon G\succcurlyeq H$ formally is such a pair (U,f). Given $\varphi=(U,f)$ we address U as $V(\varphi)$ and we write $\varphi=f$ by abuse of notation. Usually, we will abbreviate 'minor-map' as 'map'.

Lemma 6.2. *Let $G_0,G_1,\dots$ and $H_0\subseteq H_1\subseteq\cdots$ be two sequences of graphs $H_n\subseteq G_n$ with maps $\varphi_n\colon G_n\succcurlyeq G_{n+1}$ such that for every vertex $x\in G_{n+1}$ the preimage $\varphi_n^{-1}(x)$ is finite if $x\notin H_n$ and equal to $\{x\}$ if $x\in H_n$. Then G contains $\bigcup_{n\in\mathbb{N}}H_n$ as a minor with finite branch sets.*

Proof. The proof of [5, Lemma 5.12] shows this. □

Proof of Theorem 2. Let G be any Π-graph. We have to find a Farey graph minor in G with finite branch sets. By Lemma 2.1 it suffices to find a halved Farey graph minor with finite branch sets in G.

Call a graph a *foresighted* halved Farey graph of order $n\in\mathbb{N}$ if it is the edge-disjoint union of $\breve{F}_n$ with subdivided Π-graphs A_{uv}, one for every blue edge $uv\in\breve{F}_n$, such that:

- each A_{uv} meets $\breve{F}_n$ precisely in u and v but $uv\notin A_{uv}$;
- u and v are branch vertices of A_{uv};
- every two distinct A_e and $A_{e'}$ meet precisely in the intersection $e\cap e'$ of their corresponding edges (viewed as vertex sets).

To find a halved Farey graph minor with finite branch sets in G, it suffices by Lemma 6.2 to find a sequence $G =: H_0, H_1, \ldots$ of foresighted halved Farey graphs of orders $0, 1, \ldots$ with maps $\varphi_n \colon H_n \succcurlyeq H_{n+1}$ such that $\varphi_n^{-1}(x)$ is finite for all $x \in H_{n+1} - \breve{F}_n$ and $\varphi_n^{-1}(x) = \{x\}$ for all $x \in \breve{F}_n$.

To get started, pick any edge e of G, and note that $G = H_0$ is a foresighted halved Farey graph of order 0 with $A_e = G - e$ when we rename e to the edge of which $\breve{F}_0 = K^2$ consists.

At step $n+1$ suppose that we have already constructed $H_n \supseteq \breve{F}_n$ and consider the subdivided Π-graphs A_e that were added to $\breve{F}_n$ to form H_n. Theorem 6.1 yields in each A_e two subgraphs H_u^e, H_v^e for $e = uv$ that satisfy the following conditions:

(i) $H_u^e[X^e] = H_v^e[X^e]$ is finite and connected for $X^e := V(H_u^e) \cap V(H_v^e) \neq \emptyset$;
(ii) X^e avoids u and v;
(iii) both H_u^e/X^e and H_v^e/X^e are subdivided Π-graphs in which u, X^e and v, X^e are branch vertices, respectively;
(iv) uX^e is an edge of H_u^e/X^e and vX^e is an edge of H_v^e/X^e.

Then we let $A_{uv_e} := H_u^e/X^e$ and $A_{v_e v} := H_v^e/X^e$ for every blue edge $uv \in \breve{F}_n$, where we recall that v_e is the vertex $v_e \in \breve{F}_{n+1} - \breve{F}_n$ that arises from $uv \in \breve{F}_n$ in the recursive definition of $\breve{F}_{n+1}$. After renaming the vertex X^e to v_e in both A_{uv_e} and $A_{v_e v}$, we let

$$H_{n+1} := \breve{F}_{n+1} \cup \bigcup \{\, A_f \mid f \in \breve{F}_{n+1} \text{ is a blue edge} \,\}$$
$$V(\varphi_n) := V(\breve{F}_n) \cup \bigcup \{\, V(H_u^e) \cup V(H_v^e) \mid e = uv \in \breve{F}_n \text{ is a blue edge} \,\}$$

and we let $\varphi_n \colon V(\varphi_n) \to V(H_{n+1})$ send w to v_e if $w \in X^e$ for some blue edge $e \in \breve{F}_n$ and $\varphi_n(w) := w$ otherwise. This completes the proof. $\square$

Proof of Theorem 1. Theorem 2 implies Theorem 1. $\square$

References

[1] M. Clay and D. Margalit, *Office Hours with a Geometric Group Theorist*, Princeton University Press, 2017. MR3645425 ↑1, 2
[2] R. Diestel, *Graph Theory*, 5th, Springer, 2016. ↑2
[3] R. Engelking, *General Topology*, 2nd ed., Sigma Series in Pure Mathematics, vol. 6, Heldermann Verlag, Berlin, 1989. MR1039321 ↑5
[4] A. Hatcher, *Topology of numbers*, Book in preparation (2017). Available online. ↑1, 2
[5] J. Kurkofka, *Every infinitely edge-connected graph contains the Farey graph or $T_{\aleph_0} * t$ as a minor* (2020), available at arXiv:2004.06710. Submitted. ↑1, 6
[6] ———, *On the tangle compactification of infinite graphs* (2017), available at arXiv:1908.10212. ↑5
[7] ———, *Ubiquity and the Farey graph* (2019), available at arXiv:1912.02147. Submitted. ↑3.2, 5.3
[8] L. Ribes and P. Zalesskii, *Profinite Groups*, Springer, 2010. MR2599132 ↑5

University of Hamburg, Department of Mathematics, Bundesstrasse 55 (Geomatikum), 20146 Hamburg, Germany

E-mail address: jan.kurkofka@uni-hamburg.de