THE IMMERSION-MINIMAL INFINITELY EDGE-CONNECTED GRAPH

PAUL KNAPPE AND JAN KURKOFKA

ABSTRACT. We show that there is a unique immersion-minimal infinitely edge-connected graph: every such graph contains the halved Farey graph, which is itself infinitely edge-connected, as an immersion minor.

By contrast, any minimal list of infinitely edge-connected graphs represented in all such graphs as topological minors must be uncountable.

1. INTRODUCTION

The Farey graph, shown in Figure 1 and surveyed in [1, 3], 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 in [7] that the Farey graph plays a central role in graph theory too:

Theorem 1.1. *The Farey graph is one of two infinitely edge-connected graphs such that every infinitely edge-connected graph contains at least one of the two as a minor.*
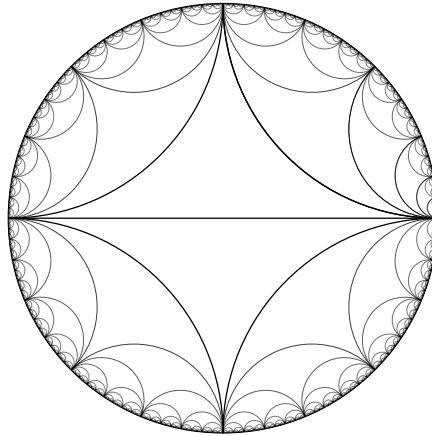

FIGURE 1. The Farey graph

Naturally, this result raises the question of whether there exist similar infinitely edge-connected graphs for graph relations other than the minor relation which determine infinite edge-connectivity by forming minimal lists in this way. In this paper, we address two graph relations that harmonise particularly well with edge-connectivity: the topological minor relation and the immersion relation.

A *weak immersion* of a graph H in a graph G is a map α with domain $V(H) \sqcup E(H)$ that embeds $V(H)$ into $V(G)$ and sends every edge $uv \in H$ to an $\alpha(u)$–$\alpha(v)$ path in G which is edge-disjoint from every other such path. The map α is a *strong immersion* of H in G if additionally all paths $\alpha(e)$ for $e \in E(H)$ have no internal vertices in $\alpha[V(H)]$. The vertices of G that lie in the image $\alpha[V(H)]$ are the *branch vertices* of this immersion.

We say that H is *strongly immersed* in G, or that H is a *strong immersion minor* of G, if there is a strong immersion of H in G. Similarly, we define *weakly immersed* and *weak immersion minor*. Robertson and Seymour showed that the weak immersion relation well-quasi-orders the finite graphs, just like the
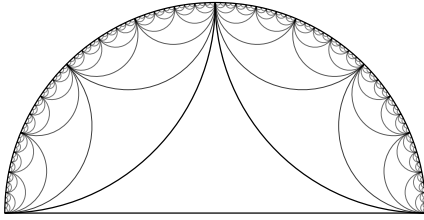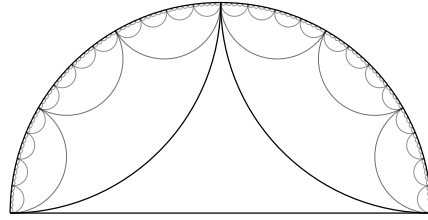
1

FIGURE 2. The halved Farey graph

FIGURE 3. A generalised halved Farey graph

minor relation, and they believe that so does the strong immersion relation [8, §1]. In this paper, we will focus on strong immersions. For brevity, we will often refer to strong immersions simply as *immersions*.

Any infinitely edge-connected graphs that form a minimal list as discussed earlier must be countable, because in every infinitely edge-connected graph we can greedily find a countable infinitely edge-connected subgraph. The countable graphs, however, are not known to be well-quasi-ordered by the minor relation or either of the immersion relations. It is therefore not clear that any immersion-minimal set of infinitely edge-connected graphs must be finite, nor even that such a minimal set exists.

A greedy argument shows that every infinitely edge-connected graph contains the countably infinite complete graph by weak immersion. So for weak immersion we have a minimal list formed by this graph alone. But this is no longer true if we replace 'weak immersion' with 'immersion'. Indeed, we show that the countably infinite complete graph is not immersed in the Farey graph (Theorem 4.1 (i)). Then, does every infinitely edge-connected graph contain the Farey graph by immersion? Perhaps surprisingly, the answer is no: the Farey graph is not an immersion minor of the *halved* Farey graph shown in Figure 2 (Theorem 4.1 (ii)), although the halved Farey graph is infinitely edge-connected. As our main result we show that the halved Farey graph is immersed in every infinitely edge-connected graph, and hence forms the desired list for immersion, again all by itself. Two graphs that are immersed in each other are called *immersion-equivalent*.

Theorem 1. *Up to immersion equivalence, the halved Farey graph is the unique infinitely edge-connected graph that is immersed in every infinitely edge-connected graph.*

Neither the minor relation nor the immersion relation is stronger than the other. Therefore, Theorem 1.1 and Theorem 1 do not obviously imply each other. The topological minor relation, however, is stronger than both the minor relation and the immersion relation. This raises the question of whether the two theorems can be unified by finding a short list of topological-minor-minimal infinitely edge-connected graphs. Our second result, Theorem 2 below, shows that this is impossible. Hence, Theorem 1.1 and Theorem 1 are best possible in this sense.

Let $\mathcal{G}$ be a class of graphs and $\leqslant$ a relation on $\mathcal{G}$. We say that a class $\mathcal{H} \subseteq \mathcal{G}$ is *typical* for $\mathcal{G}$ with regard to $\leqslant$ if for every graph $G \in \mathcal{G}$ there exists a graph $H \in \mathcal{H}$ with $H \leqslant G$. The relation of 'being typical for' with regard to $\leqslant$ is transitive on the subclasses of $\mathcal{G}$. We recall that in every infinitely edge-connected graph we can greedily find a countable infinitely edge-connected subgraph. Combining this with the fact that every countable graph is isomorphic to a subgraph of $K^{\aleph_0}$, and using the transitivity of the 'being typical for'-relation, we find that the class of infinitely edge-connected graphs includes typical sets of graphs with regard to the topological minor relation that are no larger than the continuum. However, we show that none of these typical sets is countable, let alone finite. Recall that a graph is *outerplanar* if it has a drawing in which every vertex lies on the unit circle and every edge is contained in the unit disc. Indeed:

Theorem 2. *Every set of graphs that is typical for the infinitely edge-connected graphs with regard to the topological minor relation, or even just for the outerplanar ones, is uncountable.*

Our proof of Theorem 2 builds on a construction of generalised halved Farey graphs. An instance of a generalised halved Farey graph is shown in Figure 3.

This paper is organised as follows. In Section 2 we introduce the tools and terminology that we need. In Section 3 we define the generalised halved Farey graphs and prove Theorem 2. In Section 4 we prove Theorem 1.

2. Tools and terminology

We use the notation of Diestel's book [2]. Recall that 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. Given a graph H, we call P an *H-path* if P is non-trivial and meets H exactly in its endvertices. An H-path is a $V(H)$-path; the converse is true unless the $V(H)$-path has just one edge and this edge is in H. We write $G[X]$ for the subgraph of G induced by the vertex set X. Given a path P that contains two vertices u and v, we write uPv for the subpath of P from u to v.

Whenever an x–y path P is introduced, we denote by $\leqslant_P$ the linear order on its vertices given by the way that P directed from x to y traverses them. For a partial order $\mathcal{L} = (L, \leqslant_L)$ and a subset M of L, we write $\mathcal{L} \upharpoonright M$ for the partial order obtained from $\mathcal{L}$ by restricting it to M.

2.1. Farey graph.
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 [3].

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 $e \in \breve{F}_n$, joining every v_e precisely to the endvertices of e by two blue edges, and recolouring all the edges of $\breve{F}_{n+1}$ belonging to $\breve{F}_n$ black. The *halved Farey graph* $\breve{F} := \bigcup_{n \in \mathbb{N}} \breve{F}_n$ is the union of all these $\breve{F}_n$ without their edge-colourings (cf. Figure 2), 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 $G_1 \cap G_2 = \breve{F}_0$ (cf. Figure 1).

The (halved) Farey graph and infinite edge-connectivity have been studied and used in [4–7].

2.2. Grain lines.
Suppose that x and y are two vertices in a graph G such that no finite set of edges separates x and y in G. Then we greedily find a sequence $(P_n : n \in \mathbb{N})$ of infinitely many pairwise edge-disjoint x–y paths in G. Since these paths are only edge-disjoint, they can meet in vertices other than x and y. Let us say that two x–y paths are *order-compatible* if they traverse their common vertices in the same order. Is it always possible to choose the paths P_n so that they are pairwise order-compatible? Perhaps surprisingly, the answer is no: in [6], a countable planar graph is constructed that is infinitely edge-connected, but which does not contain infinitely many edge-disjoint pairwise order-compatible paths between any two of its vertices.

Fortunately, not all is lost. While we cannot always choose the paths P_n so that they are pairwise order-compatible, we can always choose them so that they satisfy a slightly weaker form of order-compatibility which is still strong enough for our purpose. Roughly speaking, we will be able to choose the paths P_n so

that they induce a linear order on their limit. This limit will be the set of all vertices that eventually appear on all paths P_n. If two vertices u and v are in the limit, there will be a first path P_n which contains both u and v. The x–y path P_n linearly orders its vertex set from x to y; in particular, it orders u and v. This order on u and v might disagree with a later path that is not order-compatible with P_n. However, we shall achieve that all later paths will order u and v in the same way, and we will use this ordering of u and v in our limit instead of the ordering induced by P_n. This informal idea has been formalised as 'grain lines' in [5], whose definition we recall now.

An x–y *grain line* between two distinct vertices x and y is an ordered pair $(\mathcal{L}, \mathcal{P})$ where $\mathcal{L} = (L, \leqslant_L)$ is a linear order with $\min_{\leqslant_L} L = x$ and $\max_{\leqslant_L} L = y$, 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 conditions are satisfied:

(GL1) $L = \{v : \{n \in \mathbb{N} : v \in V(P_n)\}$ is a final segment of $\mathbb{N}\}$;

(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;

(GL3) for every $n \geqslant 1$, the linear order $\leqslant_{P_n}$ given by P_n and $\leqslant_L$ induce the same linear order on the vertex set $L_{<n}$, where we set $L_{<n} := L \cap \bigcup_{m<n} V(P_m)$.

Example 2.1. The halved Farey graph defines a grain line, as follows. Let x and y be the two vertices of $\breve{F}_0$. For every $n \in \mathbb{N}$, let P_n be the blue Hamilton path of $\breve{F}_n$, and let us view each P_n as an x–y path in $\breve{F}$. Then letting $L := V(\breve{F})$, $\leqslant_L := \bigcup_{n \in \mathbb{N}} \leqslant_{P_n}$ and $\mathcal{P} := (P_n : n \in \mathbb{N})$ results in an x–y grain line.

Lemma 2.2. *Let x and y be any two distinct vertices of a graph G, and let $\mathcal{Q}$ be any set of infinitely many pairwise edge-disjoint x–y paths in G. Then there exists an x–y grain line $(\mathcal{L}, \mathcal{P})$ in G such that all the paths in $\mathcal{P}$ are in $\mathcal{Q}$.*

Proof. The proof of [5, Theorem 5.4] shows this. $\square$

Whenever a grain line is introduced as $(\mathcal{L}, \mathcal{P})$, we tacitly assume that $\mathcal{L} = (L, \leqslant_L)$ and $\mathcal{P} = (P_n : n \in \mathbb{N})$. We write $\bigcup \mathcal{P} := \bigcup_{n \in \mathbb{N}} P_n$ for the graph *defined* by the grain line $(\mathcal{L}, \mathcal{P})$. A $\mathcal{P}$-*segment* is a subpath uP_dv of some path P_d in $\mathcal{P}$ with $d \geqslant 1$, if u and v are in $L_{<d}$ and v is the successor of u in $\mathcal{L} \restriction L_{<d}$. We follow the convention that P_0 also is a $\mathcal{P}$-segment. We refer to d as the $\mathcal{P}$-*depth* of the $\mathcal{P}$-segment uP_dv, and the $\mathcal{P}$-depth of P_0 is 0.

Example 2.3. Every $\mathcal{P}$-segment of a grain line has at least one edge. It is possible for $\mathcal{P}$-segments to have only one edge: for example, if all paths $P_n \in \mathcal{P}$ are internally disjoint, and all paths have two edges except one path which has exactly one edge.

We introduce the concept of depth in a sequence $\mathcal{P}$ of paths. The path P_d is the path in $\mathcal{P}$-*depth* d. For every $d \in \mathbb{N}$, we abbreviate the sequence $(P_n : n \geqslant d)$ of paths in depth at least d as $\mathcal{P}_{\geqslant d}$. Similarly, we define $\mathcal{P}_{>d}$, $\mathcal{P}_{\leqslant d}$ and $\mathcal{P}_{<d}$. The $\mathcal{P}$-*depth* of a vertex $v \in \bigcup \mathcal{P}$ is defined as $\min \{n : v \in V(P_n)\}$. The $\mathcal{P}$-*depth* of an edge $e \in \bigcup \mathcal{P}$ is $\min \{n : e \in E(P_n)\}$. The *depth* of a vertex or an edge of $\bigcup \mathcal{P}$ in $\mathcal{P}$ is its $\mathcal{P}$-depth.

Let $(\mathcal{L}, \mathcal{P})$ be a grain line. We remark that $L_{<d}$ is the set of vertices in L whose depth is less than d in $\mathcal{P}$, and the depth of an edge in $\mathcal{P}$ is at least the depth of its endvertices in $\mathcal{P}$. For two vertices $u, v \in L$ with $u <_L v$, we write $u\mathcal{P}v$ for the subsequence $(uP_nv : n > d)$ where d is the maximum of the depths of u and v in $\mathcal{P}$.

The following structural properties have been introduced in [5]. A grain line $(\mathcal{L}, \mathcal{P})$ is *wild* if $\mathcal{L}$ is order-isomorphic to $\mathbb{Q} \cap [0, 1]$. It is *wildly presented* if, for every $n \geqslant 1$, whenever $u <_L v$ are elements of $L_{<n}$ then an internal vertex of uP_nv is in the interval $(u, v)_{\mathcal{L}}$.

3. Typical sets with regard to the topological minor relation

In this section, we show the following generalisation of Theorem 2. Missing definitions follow.

Theorem 3.1. *For every countable set of infinitely edge-connected graphs, there exists an outerplanar infinitely edge-connected graph that contains none of them as a topological minor. Moreover, there is such a graph which is a generalised halved Farey graph.*

A *generalised halved Farey graph of order* 0 is a non-trivial path with its edges coloured blue. Inductively, a *generalised halved Farey graph of order* $n+1$ is an edge-coloured graph that is obtained from a generalised Farey graph G_n of order n by adding, for every blue edge $e = uv \in G_n$, a blue-coloured u–v path P_e of length at least two, which is internally disjoint from G_n and every other $P_{e'}$, and recolouring all the edges of G_n black in G_{n+1}. Let $(G_n : n \in \mathbb{N})$ be any sequence obtained by this construction. Then the union of all these G_n without their edge-colourings is a *generalised halved Farey graph*. We remark that every generalised halved Farey graph is outerplanar.

Just like the halved Farey graph in Example 2.1, we may interpret a generalised halved Farey graph as a grain line. This grain line then satisfies the following stronger versions of (GL2) and (GL3):

(GL2') $L = \bigcup_{n \in \mathbb{N}} V(P_n)$;

(GL3') $\leqslant_L = \bigcup_{n \in \mathbb{N}} \leqslant_{P_n}$.

It is immediate to see that the converse is also true:

Lemma 3.2. *A graph G is a generalised halved Farey graph if and only if G is defined by a grain line $(\mathcal{L}, \mathcal{P})$ which satisfies (GL2') and (GL3').* $\qquad\qquad\square$

The generalised halved Farey graphs which we will construct in the proof of Theorem 3.1 are of the following type. Let $\ell \colon \mathbb{N} \to \mathbb{N}$ be a function such that $\ell(0) \geqslant 1$ and, for every $n \geqslant 1$, $\ell(n) \geqslant 2$. Then ℓ *induces* an (up to isomorphism) unique grain line $(\mathcal{L}, \mathcal{P})(\ell)$ which satisfies (GL2'), (GL3') and, for every $n \in \mathbb{N}$, every $\mathcal{P}$-segment in $\mathcal{P}$-depth n is a path of length $\ell(n)$. The *generalised halved Farey graph $\breve{F}(\ell)$ induced by ℓ* is the graph defined by $(\mathcal{L}, \mathcal{P})(\ell)$. Note that by Lemma 3.2, $\breve{F}(\ell)$ is a generalised halved Farey graph.

Example 3.3. The halved Farey graph is the generalised halved Farey graph induced by ℓ with $\ell(0) := 1$ and $\ell(n) := 2$ for all natural numbers $n \geqslant 1$. The generalised halved Farey graph shown in Figure 3 is induced by the function $\ell \colon \mathbb{N} \to \mathbb{N}$ which maps n to $n+1$.

To prove Theorem 3.1, we need two lemmas about grain lines, which are motivated by the proof of Theorem 3.1. So we prove Theorem 3.1 first, giving the statements of the two lemmas where we need them, and then proceed to prove the two lemmas afterwards.

Proof of Theorem 3.1. We have to show that, for every countable set $\mathcal{H}$ of infinitely edge-connected graphs, there exists an outerplanar infinitely edge-connected graph G which contains no subdivision of a graph in $\mathcal{H}$. For this, let $\mathcal{H}$ be any countable set of infinitely edge-connected graphs. By applying Lemma 2.2 in each graph in $\mathcal{H}$, we find a countably infinite set $\{(\mathcal{M}^{(i)}, \mathcal{Q}^{(i)}) : i \in \mathbb{N}\}$ of grain lines such that every graph in $\mathcal{H}$ contains one of these. For every $k \in \mathbb{N}$, we let $\ell(k)$ be one greater than the maximum of the lengths of the paths $Q_j^{(i)}$ with $0 \leqslant i, j \leqslant 2k$. We remark that $\ell(0) \geqslant 1$ and $\ell(1) \geqslant 2$, and that the function ℓ is increasing (though not necessarily strictly so).

Let $(\mathcal{L}, \mathcal{P}) := (\mathcal{L}, \mathcal{P})(\ell)$ be the grain line induced by ℓ, and let $G := \breve{F}(\ell)$ be the generalised halved Farey graph induced by ℓ. In particular, $(\mathcal{L}, \mathcal{P})$ satisfies (GL2') and (GL3'), and G is outerplanar and infinitely edge-connected. We claim that G contains no subdivision of a graph in $\mathcal{H}$.

Indeed, suppose for a contradiction that some graph H in $\mathcal{H}$ is a topological minor of G. Then there is a natural number i such that $\bigcup \mathcal{Q}^{(i)}$ is a topological minor of G. In Section 3.1, we will see the crucial fact that if a generalised halved Farey graph contains a grain line as a topological minor, then it actually contains the grain line as a subgraph (up to some finite error):

Key Lemma 3.4. *Let $(\mathcal{M}, \mathcal{Q})$ be a grain line such that a subdivision of $\bigcup \mathcal{Q}$ is contained in a generalised halved Farey graph. Then there exists a number d such that no edge of $\bigcup \mathcal{Q}_{\geqslant d}$ is subdivided.*

Hence, by Key Lemma 3.4, there is a number d such that $(\mathcal{M}, \mathcal{Q}^{(i)}_{\geqslant d})$ is a grain line in G. In Section 3.2, we will see that while we 'dive deeper' into the grain line $(\mathcal{M}, \mathcal{Q}^{(i)}_{\geqslant d})$, we 'dive' with at least the same speed into the grain line $(\mathcal{L}, \mathcal{P})$:

Key Lemma 3.5. *Let $(\mathcal{M}, \mathcal{Q})$ be a grain line in the graph $\bigcup \mathcal{P}'$ defined by a grain line $(\mathcal{L}', \mathcal{P}')$ satisfying (GL2') and (GL3'). Then there are a natural number q, a strictly increasing sequence $(p_k : k \in \mathbb{N})$ of natural numbers and a sequence of nested intervals $([u_k, v_k]_{\mathcal{M}} : k \in \mathbb{N})$ of $\mathcal{M}$ such that, for every $k \in \mathbb{N}$, the path $u_{k+1} P'_{p_k} v_{k+1}$ is a $\mathcal{P}'$-segment and a subpath of $u_k Q_{q+k} v_k$.*

Thus, by Key Lemma 3.5 applied to $(\mathcal{L}', \mathcal{P}') := (\mathcal{L}, \mathcal{P})$, there is a natural number $q \geqslant d$, a strictly increasing sequence $(p_k : k \in \mathbb{N})$ of natural numbers and a sequence of nested intervals $([u_k, v_k]_{\mathcal{M}^{(i)}} : k \in \mathbb{N})$ of $\mathcal{M}^{(i)}$ such that, for every $k \in \mathbb{N}$, the path $u_k Q^{(i)}_{q+k} v_k$ contains a $\mathcal{P}$-segment $u_{k+1} P_{p_k} v_{k+1}$. In particular, by the definition of $(\mathcal{L}, \mathcal{P})$, the path $Q^{(i)}_{q+k}$ has length at least $\ell(p_k) \geqslant \ell(p_0 + k) \geqslant \ell(k)$ for every $k \in \mathbb{N}$. For $k = \max\{q, \lceil \frac{i}{2} \rceil\}$, we have $i \leqslant 2 \cdot \lceil \frac{i}{2} \rceil \leqslant 2k$ and $q + k \leqslant 2k$. Thus, by the definition of $\ell(k)$, the path $Q^{(i)}_{q+k}$ is at least one longer than itself, a contradiction. $\qquad\square$

Proof of Theorem 2. Theorem 3.1 implies Theorem 2. $\qquad\square$

So to complete the proof of Theorem 3.1, it is left to show Key Lemma 3.4 and Key Lemma 3.5. We will prove them in Section 3.1 and in Section 3.2, respectively.

3.1. **Proof of Key Lemma 3.4.** A grain line $(\mathcal{L}, \mathcal{P})$ is *well-structured* if, for every $\mathcal{P}$-segment $u P_d v$, we have the inclusion $V(u P_d v) \cap L \subseteq [u, v]_{\mathcal{L}}$. We say that a grain line $(\mathcal{L}, \mathcal{P})$ is *free* if $\bigcup \mathcal{P}$ is infinitely edge-connected. Obviously, the following assertions are equivalent:

$$(1)\ (\mathcal{L}, \mathcal{P})\ \text{is free};\quad (2)\ L = V(\textstyle\bigcup \mathcal{P});\quad (3)\ \text{no vertex of}\ \textstyle\bigcup \mathcal{P}\ \text{has degree two}.$$

We remark but will not use that free and well-structured grain lines are wild.

The grain line that defines the halved Farey graph in Example 2.1, for instance, is both free and well-structured. More general examples which may be helpful to have in mind while thinking about free and well-structured grain lines are grain lines which satisfy (GL2') and (GL3'). In fact, the graphs which we constructed in the proof of Theorem 3.1 are also defined by such a grain line, i.e., they are generalised halved Farey graphs by Lemma 3.2. We remark that these are not only wild but also wildly presented.

A crucial property of any well-structured grain line is that the deletion of any internal vertex destroys the infinite edge-connectivity between its startvertex and endvertex (see Figure 4):

Lemma 3.6. *Let $(\mathcal{L}, \mathcal{P})$ be a well-structured grain line and let $u <_L v <_L w$ be three vertices in L. Then the vertex v together with the set F of all the finitely many edges in $\mathcal{P}$-depth at most the $\mathcal{P}$-depth of v separate u and w in $\bigcup \mathcal{P}$. That is, v together with F separates all vertices in $[x, v)_{\mathcal{L}}$ from all vertices in $(v, y]_{\mathcal{L}}$, where x and y are the minimum and maximum of $\mathcal{L}$, respectively.*
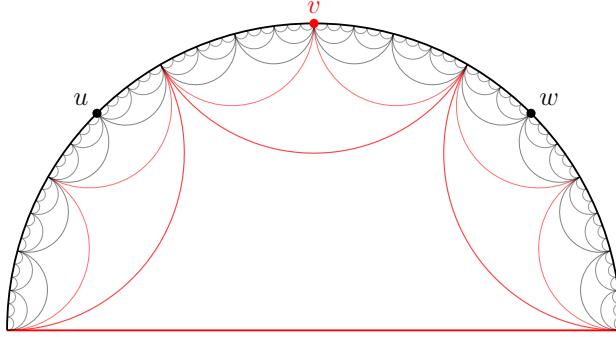
FIGURE 4. A generalised halved Farey graph in which its vertex v together with the edges (red) in depth at most the depth of v separate u and w.

Proof. Let d be the depth of v in $\mathcal{P}$. Since $(\mathcal{L}, \mathcal{P})$ is well-structured, we have for every path P_n with $n > d$ the two inclusions

$$V(xP_nv - v) \cap L \subseteq [x, v)_{\mathcal{L}} \quad \text{and} \quad V(vP_ny - v) \cap L \subseteq (v, y]_{\mathcal{L}}.$$

Hence, every $[x, v)_{\mathcal{L}}$–$(v, y]_{\mathcal{L}}$ path in $\bigcup \mathcal{P}$ which avoids v has to contain at least one of the finitely many edges in depth at most d in $\mathcal{P}$. □

It follows that every well-structured grain line $(\mathcal{L}, \mathcal{P})$ imposes the order of $\mathcal{L}$ on every grain line in $\bigcup \mathcal{P}$:

Lemma 3.7. *Let $(\mathcal{M}, \mathcal{Q})$ be a grain line in the graph $\bigcup \mathcal{P}$ defined by a grain line $(\mathcal{L}, \mathcal{P})$. If $(\mathcal{L}, \mathcal{P})$ is well-structured, then M is an interval of $\mathcal{L}$ and $\leqslant_M$ is either the linear order on M induced by $\leqslant_L$ or the reverse of it.*

Proof. First, we note that M is a subset of L: By (GL2), any vertex of $\bigcup \mathcal{P}$ which is not in L has degree two; but every vertex of M has infinite degree in $\bigcup \mathcal{Q} \subseteq \bigcup \mathcal{P}$ by (GL1). We claim that it suffices to show the following assertion (∗).

> *For every three vertices $u, w \in M$ and $v \in L$ with $u \leqslant_L v \leqslant_L w$, the vertex v is contained in M and the vertex v lies in between u and w with respect to $\leqslant_M$.* (∗)

Let x and y be the endvertices of $(\mathcal{M}, \mathcal{Q})$ named so that $x <_L y$. First, let us deduce $M = [x, y]_{\mathcal{L}}$ from (∗). By possibly reversing $\leqslant_M$, we may assume without loss of generality that also $x <_M y$. For the forward inclusion $M \subseteq [x, y]_{\mathcal{L}}$, let $v \neq x, y$ be any element of M. We show $x <_L v$ and $v <_L y$ separately. To see that $x <_L v$, we suppose for a contradiction that $v <_L x$. Then we apply Lemma 3.6 to the grain line $(\mathcal{L}, \mathcal{P})$ and $v <_L x <_L y$ to find a finite set F_x of edges of $\bigcup \mathcal{P}$ such that x together with F_x separates v from y in the graph $\bigcup \mathcal{P}$. But $v \mathcal{Q} y$ is an infinite system of pairwise edge-disjoint v–y paths that avoid x (since $x <_M v$), contradicting that F_x is finite. Hence we have $x <_L v$. An analogue argument shows $v <_L y$. Hence, $v \in [x, y]_{\mathcal{L}}$. Applying (∗) to x and y and each $v \in [x, y]_{\mathcal{L}}$ yields the backward inclusion $M \supseteq [x, y]_{\mathcal{L}}$. Second, we verify that $\leqslant_M$ and $\leqslant_L$ agree on M. For every two vertices u, v in M with $u \leqslant_L v$, it follows from $M \subseteq [x, y]_{\mathcal{L}}$ that $x \leqslant_L u \leqslant_L v$ and, since $x \leqslant_M v$, we obtain from (∗) that $x \leqslant_M u \leqslant_M v$.

Now suppose for a contradiction that (∗) does not hold, i.e., suppose that there are three vertices $u, w \in M$ and $v \in L$ with $u \leqslant_L v \leqslant_L w$ such that v does not lie between u and w with respect to $\leqslant_M$. Thus, $u <_L v <_L w$. Then the sequence $(u Q_n w : n > d)$ with d larger than the $\mathcal{Q}$-depths of u, v and w is an

infinite system of pairwise edge-disjoint u–w paths in $\bigcup\mathcal{P}$ which avoid v, contradicting Lemma 3.6 applied to the grain line $(\mathcal{L},\mathcal{P})$ and $u <_L v <_L w$. $\square$

Next, we show that if $(\mathcal{L},\mathcal{P})$ is not only well-structured but also free, then $(\mathcal{L},\mathcal{P})$ also imposes its freedom on every grain line in $\bigcup\mathcal{P}$ (up to some finite and hence negligible error).

Lemma 3.8. *Let $(\mathcal{M},\mathcal{Q})$ be a grain line in the graph $\bigcup\mathcal{P}$ defined by a grain line $(\mathcal{L},\mathcal{P})$. If $(\mathcal{L},\mathcal{P})$ is free and well-structured, then only finitely many vertices of $\bigcup\mathcal{Q}$ lie outside of M. In particular, there is a number d such that $(\mathcal{M},\mathcal{Q}_{\geqslant d})$ is free.*

Proof. Let U be the set of all vertices of $\bigcup\mathcal{Q}$ that lie outside of M. Since $(\mathcal{L},\mathcal{P})$ is free, we have $L = V(\bigcup\mathcal{P})$, so $U \cup M \subseteq L$. By Lemma 3.7, there exist $x,y \in L$ with $x <_L y$ such that $[x,y]_{\mathcal{L}} = M$, and we may assume without loss of generality that $\leqslant_M$ is induced by $\leqslant_L$ (possibly after reversing $\leqslant_M$). Hence U is included in $[a,x)_{\mathcal{L}} \cup (y,b]_{\mathcal{L}}$, where a and b denote the minimum and maximum of $\mathcal{L}$, respectively.

Let us suppose for a contradiction that U is infinite. Without loss of generality, the intersection of U with $[a,x)_{\mathcal{L}}$ is infinite. Since there are only finitely many vertices of $\bigcup\mathcal{Q}$ at $\mathcal{Q}$-depth d for each $d \in \mathbb{N}$, there exists an infinite subset $U' \subseteq U$ such that every two distinct vertices in U' have distinct $\mathcal{Q}$-depths. Hence the paths $uQ_{d(u)}y$ with $u \in U'$, where $d(u)$ denotes the $\mathcal{Q}$-depth of u, are pairwise edge-disjoint. Since x is the minimum of $\mathcal{M}$ and since all u are distinct from x by definition of U, all these paths avoid x. But by Lemma 3.6 applied to $u <_L x <_L y$, the vertex x together with some finitely many edges separates u and y in $\bigcup\mathcal{P}$, a contradiction. $\square$

Since subdividing vertices have degree two, we conclude from Lemma 3.8 that if a free and well-structured grain line $(\mathcal{L},\mathcal{P})$ contains another grain line as a topological minor, then $(\mathcal{L},\mathcal{P})$ actually contains it as a subgraph (again up to some finite and hence negligible error):

Lemma 3.9. *Let $(\mathcal{L},\mathcal{P})$ and $(\mathcal{M},\mathcal{Q})$ be grain lines such that $\bigcup\mathcal{P}$ contains a subdivision of $\bigcup\mathcal{Q}$. If $(\mathcal{L},\mathcal{P})$ is free and well-structured, then there is a number d such that $(\mathcal{M},\mathcal{Q}_{\geqslant d})$ is a grain line in $\bigcup\mathcal{P}$.* $\square$

Proof of Key Lemma 3.4. Since a generalised halved Farey graph is defined by a free and well-structured grain line, this follows directly from Lemma 3.9. $\square$

3.2. **Proof of Key Lemma 3.5.** We need the following three lemmas.

Lemma 3.10. *Let $(\mathcal{L},\mathcal{P})$ be a grain line and P a path in $\bigcup\mathcal{P}$. If P contains at least one edge in $\mathcal{P}$-depth greater than the $\mathcal{P}$-depth of its endvertices, then it already contains a $\mathcal{P}$-segment in depth equal to the maximum of the $\mathcal{P}$-depths of the edges of P.*

Proof. Let d be the maximum of the depths of the edges of P in $\mathcal{P}$. Let uPv be a subpath of P that is a $(\bigcup\mathcal{P}_{<d})$-path; this subpath exists because an edge of P has $\mathcal{P}$-depth d and the endvertices of P are contained in $\bigcup\mathcal{P}_{<d}$ by assumption. It follows directly from the choice of uPv that all of its edges have $\mathcal{P}$-depth at least d, and hence uPv actually is a subpath of P_d. So both u and v lie in L by (GL1), and thus in $L_{<d}$. Hence, uPv is the desired $\mathcal{P}$-segment. $\square$

Lemma 3.11. *Let $(\mathcal{L},\mathcal{P})$ be a grain line and let P be a u–v path in $\bigcup\mathcal{P}$ containing an edge $e = ww'$ such that $u \leqslant_L w <_L w' \leqslant_L v$ and such that the $\mathcal{P}$-depth d of e is greater than the maximum of the $\mathcal{P}$-depths of u and v. Let Q be the subpath of P which contains e and is a $(\bigcup\mathcal{P}_{<d})$-path. If $(\mathcal{L},\mathcal{P})$ is well-structured, then $V(Q) \cap L \subseteq [u,v]_{\mathcal{L}}$.*

Proof. We denote the minimum and maximum of $\mathcal{L}$ by a and b, respectively. It suffices to show that Q avoids both $[a, u)_\mathcal{L}$ and $(v, b]_\mathcal{L}$. By symmetry, it suffices to show that Q avoids $[a, u)_\mathcal{L}$. Let F denote the set of edges of $\bigcup \mathcal{P}_{<d}$. By Lemma 3.6, the vertex u separates $[a, u)_\mathcal{L}$ from $(u, b]_\mathcal{L}$ in the graph $(\bigcup \mathcal{P}) - F$. The path Q contains the vertex $w' \in (u, b]_\mathcal{L}$; it avoids F because Q is a $(\bigcup \mathcal{P}_{<d})$-path; and it does not contain u as an internal vertex because Q is a subpath of the u–v path P. Therefore, Q avoids $[a, u)_\mathcal{L}$. $\square$

Lemma 3.12. *Let $(\mathcal{M}, \mathcal{Q})$ be a grain line in the graph $\bigcup \mathcal{P}$ defined by a free, well-structured and wildly presented grain line $(\mathcal{L}, \mathcal{P})$. Then there are a natural number q, a strictly increasing sequence $(p_k : k \in \mathbb{N})$ of natural numbers and a sequence of nested intervals $([u_k, v_k]_\mathcal{M} : k \in \mathbb{N})$ of $\mathcal{M}$ such that, for every $k \in \mathbb{N}$, the path $u_{k+1} P_{p_k} v_{k+1}$ is a $\mathcal{P}$-segment and it is a subpath of $u_k Q_{q+k} v_k$.*

Proof. By Lemma 3.8, we may assume by choosing q large enough that, without loss of generality, $(\mathcal{M}, \mathcal{Q})$ is free. Let $u_0, v_0 \in \bigcup \mathcal{P}$ be the two vertices for which $(\mathcal{M}, \mathcal{Q})$ is a u_0–v_0 grain line. By Lemma 3.7 and possibly interchanging u_0 and v_0, we may assume without loss of generality that $\leqslant_L$ and $\leqslant_M$ agree on $M \subseteq L = V(\bigcup \mathcal{P})$. Moreover, the lemma also ensures that M is an interval of $\mathcal{L}$. Let p be the maximum of the $\mathcal{P}$-depths of u_0 and v_0. Since there are only finitely many edges in $\mathcal{P}$-depth at most p, we can pick q large enough so that Q_q contains an edge of $\bigcup \mathcal{P}$ in $\mathcal{P}$-depth greater than p. Then by Lemma 3.10, there is a number $p_0 > p$ such that Q_q contains a $\mathcal{P}$-segment $u_1 P_{p_0} v_1$. Since $(\mathcal{M}, \mathcal{Q})$ is free, u_1 and v_1 are in M. Thus, $[u_1, v_1]_\mathcal{M} \subseteq [u_0, v_0]_\mathcal{M}$ because u_0 and v_0 are the minimum and maximum of $\mathcal{M}$.

Since $(\mathcal{L}, \mathcal{P})$ is wildly presented, there is an internal vertex w_1 of $u_1 P_{p_0} v_1$ which is contained in $(u_1, v_1)_\mathcal{L}$. Since $(\mathcal{M}, \mathcal{Q})$ is free and $\leqslant_L$ and $\leqslant_M$ agree on M, the vertex w_1 is contained in M and $w_1 \in (u_1, v_1)_\mathcal{M}$. Thus by (GL1), Q_{q+1} contains all three vertices u_1, v_1, w_1 and, by (GL3), w_1 is an internal vertex of $u_1 Q_{q+1} v_1$. Since $u_1 P_{p_0} v_1$ is a $\mathcal{P}$-segment, the $\mathcal{P}$-depth of the internal vertex w_1 of $u_1 P_{p_0} v_1$ is p_0. Since Q_q already contains the two edges of P_{p_0} that are incident with w_1, and since w_1 lies on no path P_p with $p < p_0$, the path $u_1 Q_{q+1} v_1$ contains an edge e_{q+1} incident with w_1 in $\mathcal{P}$-depth at least $p_0 + 1$. Let Q'_{q+1} be the subpath of $u_1 Q_{q+1} v_1$ which contains e_{q+1} and is a $(\bigcup \mathcal{P}_{\leqslant p_0})$-path. By Lemma 3.10, $Q'_{q+1} \subseteq u_1 Q_{q+1} v_1$ contains a $\mathcal{P}$-segment $u_2 P_{p_1} v_2$ with $p_1 > p_0$. By Lemma 3.11 and since $(\mathcal{L}, \mathcal{P})$ is free, $V(Q'_{q+1}) \subseteq [u_1, v_1]_\mathcal{M}$. Since $u_2 P_{p_1} v_2$ is a subpath of Q'_{q+1}, we have $[u_2, v_2]_\mathcal{M} \subseteq [u_1, v_1]_\mathcal{M}$.

Next, we repeat the previous step where we replace $u_1 P_{p_0} v_1$, q and p_0 with $u_2 P_{p_1} v_2$, $q + 1$ and p_1, and iterate in this way to conclude the proof. $\square$

Proof of Key Lemma 3.5. Since a grain line which satisfies (GL2') and (GL3') is free, well-structured and wildly presented, this follows directly from Lemma 3.12. $\square$

3.3. **Excluding k-bounded Farey graph minors.** In [5], it was shown that the Farey graph is uniquely determined by its connectivity, as follows. A Π-*graph* is an infinitely edge-connected graph such that no two of its vertices are linked by infinitely many pairwise internally disjoint paths. A κ-*bounded* minor, for a cardinal κ, is a minor with branch sets of size less than κ. A Π-graph is κ-*typical* if it occurs as a κ-bounded minor in every Π-graph. Note that any two κ-typical Π-graphs are κ-bounded minors of each other; we call such graphs κ-*boundedly minor-equivalent*.

Theorem 3.13. [5] *Up to $\aleph_0$-bounded minor-equivalence, the Farey graph is the unique $\aleph_0$-typical Π-graph.*

A referee of [5] asked whether $\aleph_0$ is best possible for the above theorem; that is, can $\aleph_0$ be replaced in Theorem 3.13 with some $\kappa < \aleph_0$ so that the theorem's statement remains true? Using generalised halved Farey graphs, we can show that the answer is in the negative:

Lemma 3.14. *For every $k \in \mathbb{N}$ there exists a Π-graph that does not contain the Farey graph as a k-bounded minor and that is a generalised halved Farey graph.*

Proof. Let $k \in \mathbb{N}$ be given. On the one hand, as the Farey graph contains a triangle, every graph that contains the Farey graph as a k-bounded minor must contain a cycle of length at most $3k$. On the other hand, every generalised halved Farey graph clearly is a Π-graph. So it suffices to find a generalised halved Farey graph of girth at least $3k + 1$. For example, we could take $\breve{F}(\ell)$ for $\ell \colon \mathbb{N} \to \mathbb{N}$ given by $\ell(0) := 1$ and $\ell(n) := 3k$ when $n \geqslant 1$. $\qquad\square$

4. Typical graph with regard to the immersion relation

In this section, we prove Theorem 1: we show that the halved Farey graph is immersed in every infinitely edge-connected graph. Before we do this, however, we take a step back to verify that the other two obvious candidates, namely $K^{\aleph_0}$ and the Farey graph, are not immersed in every infinitely edge-connected graph. For readers who are familiar with [7], we remark that the second infinitely edge-connected graph mentioned in Theorem 1.1 is immersion-equivalent to $K^{\aleph_0}$; in particular, it is not immersed in every infinitely edge-connected graph either.

Theorem 4.1.
  (i) *$K^{\aleph_0}$ is not immersed in the Farey graph.*
  (ii) *Neither $K^{\aleph_0}$ nor the Farey graph is immersed in the halved Farey graph.*

Proof. (i): Suppose for a contradiction that there is an immersion of $K^{\aleph_0}$ in the Farey graph F. Let U be the set of branch vertices of this immersion. By the definition of the Farey graph, we may write F as the union of two halved Farey graphs which intersect in the Farey graph of order 0. Hence, the two linear orders on the vertex sets of these two halved Farey graphs introduced in Example 2.1 induce a cyclic order on the vertex set of F. Now, there are two branch vertices $u, v \in U$ such that the open intervals (u, v) and (v, u) with respect to this cyclic order intersect U non-emptily. Since $K^{\aleph_0}$ is infinitely edge-connected even after deleting the two vertices u and v, the immersion yields infinitely many edges between (u, v) and (v, u) in F. This contradicts the fact that, by the construction of F, there are only finitely many of these.

  (ii): Since the halved Farey graph $\breve{F}$ is a subgraph of the Farey graph, $K^{\aleph_0}$ cannot be immersed in $\breve{F}$ by (i). Suppose for a contradiction that there is an immersion of the Farey graph F in the halved Farey graph $\breve{F}$. Let U be the set of branch vertices of this immersion. Let x, y be the two vertices in the halved Farey graph of order 0 in the halved Farey graph $\breve{F}$. There is a branch vertex $w \in U$ such that the two intervals $[x, w)$ and $(w, y]$ with respect to the linear order on $V(\breve{F})$ introduced in Example 2.1 intersect U non-emptily. Since the Farey graph F is infinitely edge-connected even after deleting the single vertex w, the immersion yields infinitely many edges between $[x, w)$ and $(w, y]$ in $\breve{F}$. This contradicts the fact that, by the construction of $\breve{F}$, there are only finitely many of those. $\qquad\square$

  Now we tend to the proof of Theorem 1.

4.1. **Overview of the proof of Theorem 1.** Our aim for the remainder of this paper is to show that every infinitely edge-connected graph contains the halved Farey graph as an immersion minor. For this, it obviously suffices to consider only infinitely edge-connected graphs with no $K^{\aleph_0}$-immersion-minors.

  As our first step, we will transfer the notion of cutvertices from vertex-connectivity to infinite edge-connectivity by introducing 'compound-cutvertices', and we will introduce what could be considered an

analogue of the block-cutvertex theorem for compound-cutvertices and blocks of infinite edge-connectivity (Key Lemma 4.5). Unlike the tree-structure of a graph imposed by its cutvertices, the structure imposed by its compound-cutvertices will be tree-like but not in general a tree-decomposition. Indeed, the tree-like structure that we will obtain from the compound-cutvertices of a graph can exhibit $(\omega + 1)$-chains, and for some graphs like the halved Farey graph, their tree-like structure can even be order-isomorphic to $\mathbb{Q}$.

As our second step, we will show that an infinitely edge-connected graph contains the halved Farey graph as an immersion minor if the tree-like structure given by its compound-cutvertices is 'wild' in that it exhibits an interval which is order-isomorphic to $\mathbb{Q}$ (Key Lemma 4.10).

So as our third and final step, we will deal with the case in which the tree-like structure given by the compound-cutvertices is not wild. Roughly, we will employ the compound-cutvertices to find infinitely many infinitely edge-connected subgraphs which are almost vertex-disjoint and which have no compound-cutvertices themselves. Then we will choose an arbitrary vertex in each subgraph that lies in no other of these subgraphs, and we will link the chosen vertices up with paths to obtain a $K^{\aleph_0}$-immersion-minor. As this contradicts our initial assumption that no $K^{\aleph_0}$-immersion-minor is present, the proof will be concluded.

4.2. **Compound separations.** Recall that a *separation of a set* V is a set $\{A, B\}$ such that $A \cup B = V$. We call A and B the *sides* of this separation. The separation is *proper* if $A \smallsetminus B$ and $B \smallsetminus A$ are non-empty. Now let G be any infinite graph. A *compound-separation* of G is a proper separation $\{A, B\}$ of $V(G)$ such that the *separator* $A \cap B$ is finite and G has only finitely many edges between $A \smallsetminus B$ and $B \smallsetminus A$. Then the cardinality of the separator $A \cap B$ is the *order* of the compound-separation $\{A, B\}$. If the separator of a compound-separation $\{A, B\}$ is a singleton $\{u\}$, then we also refer to the vertex u as the *separator* of $\{A, B\}$, and we say that $\{A, B\}$ is *unitary*. A vertex u is a *compound-cutvertex* of G if there exists a compound-separation of G with separator u. A compound-separation $\{A, B\}$ of G *separates* two vertices u, v of G if u is contained in $A \smallsetminus B$ and v is contained in $B \smallsetminus A$, or vice versa. It separates u and v *minimally* if no compound-separation $\{C, D\}$ with $C \cap D \subsetneq A \cap B$ separates u and v in G. For example, if $\{A, B\}$ is a unitary compound-separation of an infinitely edge-connected graph with $u \in A \smallsetminus B$ and $v \in B \smallsetminus A$, then $\{A, B\}$ minimally separates u and v in that graph. We say that two vertices u and v of G are *k-compound-connected* in G for a natural number k if no compound-separation of G of order less than k separates u and v. If every two vertices of G are k-compound-connected, then G itself is *k-compound-connected*.

If an infinite graph G is k-compound-connected for every $k \in \mathbb{N}$, then G is infinitely vertex-connected; in particular, we greedily find $K^{\aleph_0}$ as a topological minor in G, so the halved Farey graph is immersed in G:

Observation 4.2. *Let G be an infinite graph. If the halved Farey graph is not immersed in G, then there is a pair of vertices of G which is separated by some compound-separation of G.* $\qquad\square$

Every minimally separating compound-separation of an infinitely edge-connected graph G can be used to split G into two infinitely edge-connected subgraphs:

Key Lemma 4.3. *Let G be an infinitely edge-connected graph and $\{A, B\}$ a compound-separation of G that minimally separates two vertices of G. Then $G[A]$ and $G[B]$ are infinitely edge-connected.*

We prepare the proof of Key Lemma 4.3 with the following lemma:

Lemma 4.4. *Let G be an infinitely edge-connected graph and $\{A, B\}$ a compound-separation of G that minimally separates two vertices $u, v \in G$. If $u \in A$, then there is a system of pairwise edge-disjoint u–$(A \cap B)$ paths in $G[A]$ such that, for every vertex $w \in A \cap B$, infinitely many of these paths end in w.*

Proof. It suffices to find for every vertex $w \in A \cap B$ a system $\mathcal{P}$ of infinitely many pairwise edge-disjoint u–w paths in $G[A]$ which are internally disjoint from $A \cap B$. For this, let $w \in A \cap B$ be any vertex. Set $S := (A \cap B) \setminus \{w\}$ and $H := G - S$. Any finite cut $E(C, D)$ of H which separates u and v induces a compound-separation $\{C \cup S, D \cup S\}$ of G with separator $S \subsetneq A \cap B$ which separates u and v. Since the latter does not exist by assumption, u and v cannot be separated by finitely many edges in H. Hence, there are infinitely many pairwise edge-disjoint u–v paths in G that avoid S. Since $\{A, B\}$ is a compound-separation that separates u and v, all but finitely many of these paths have to meet w. Thus, the u–w subpaths of these paths form the desired path system. $\square$

Proof of Key Lemma 4.3. By symmetry, it is enough to prove that $G[A]$ is infinitely edge-connected. Let u and v be two vertices of G which are minimally separated by a compound-separation $\{A, B\}$, named so that $u \in A \setminus B$. To show that $G[A]$ is infinitely edge-connected, it suffices to find infinitely many pairwise edge-disjoint u–a paths in $G[A]$ for every vertex $a \in A$ other than u. For this, let any such vertex a be given. Since G is infinitely edge-connected, there is an infinite system $\mathcal{Q}'$ of pairwise edge-disjoint a–v paths in G. As $\{A, B\}$ is a compound-separation of G, it follows from the pigeonhole principle that there is a vertex $w \in A \cap B$ for which there are infinitely many paths in $\mathcal{Q}'$ whose first vertex in B is w. We denote the system of the a–w subpaths of these paths by $\mathcal{Q}$. Note that all paths in $\mathcal{Q}$ are included in $G[A]$. By Lemma 4.4, there is an infinite system $\mathcal{P}$ of pairwise edge-disjoint u–w paths in $G[A]$. Since $\mathcal{Q}$ and $\mathcal{P}$ both consist of paths in $G[A]$, we can greedily combine these two path systems to obtain infinitely many pairwise edge-disjoint u–a paths in $G[A]$. Thus, $G[A]$ is infinitely edge-connected. $\square$

4.3. **Faithful nested sets of unitary compound-separations.** In this section, we prove the following key lemma which could be viewed as an analogue of the block-cutvertex theorem for infinite edge-connectivity. Missing definitions follow.

Key Lemma 4.5. *Let G be an infinitely edge-connected graph. If $K^{\aleph_0}$ is not immersed in G, then there is a nested set of unitary compound-separations of G which is faithful to G.*

First, we make this statement precise. Let G be an infinitely edge-connected graph. For a separation $\{A, B\}$ of a set V, we recall that the ordered pairs (A, B) and (B, A) are its *orientations*, and (A, B) and (B, A) are *oriented separations* of V. For two oriented separations (A, B) and (C, D), we have that $(A, B) \leqslant (C, D)$ if $A \subseteq C$ and $B \supseteq D$. Two separations are *nested* if they have comparable orientations. Moreover, a set of separations is *nested* if every two separations in it are nested. By choosing precisely one orientation of each unoriented separation in a set S of unoriented separations, we obtain an *orientation* σ of S. The orientation σ is a *star* of separations if for every two distinct $(A, B), (C, D) \in \sigma$ we have $(A, B) \leqslant (D, C)$.

Let w be a compound-cutvertex of G. A nested set N of unitary compound-separations of G is *faithful* to w in G if

- the separators of all separations in N are equal to w,
- there is some orientation of N which is a star of separations,
- and for every pair of vertices which are separated by some compound-separation with separator w there is some compound-separation in N separating them.

A nested set N of unitary compound-separations of G is *faithful* to G if for every compound-cutvertex w of G the subset N_w of N containing all compound-separations of N with separator w is faithful to w in G.

If we can find for each compound-cutvertex w of G a nested set of compound-separations of G which is faithful to w in G, then we can combine these to obtain a nested set which is faithful to G:

Lemma 4.6. *Let G be an infinitely edge-connected graph. If $\{A, B\}$ and $\{C, D\}$ are two unitary compound-separations of G with distinct separators u and v, respectively, such that $u \in C$ and $v \in B$, then $(A, B) \leqslant (C, D)$ and, in particular, $\{A, B\}$ and $\{C, D\}$ are nested.*

Proof. The induced subgraph $G[D]$ shares with B at least the vertex v by assumption, but it does not contain the vertex u because u is contained in $C \smallsetminus \{v\} = V(G) \smallsetminus D$. Since $G[D]$ is infinitely edge-connected by Key Lemma 4.3 and $\{A, B\}$ is a compound-separation with separator u, it follows that $G[D]$ is included in $G[B] - u$. By symmetry, $G[A]$ is included in $G[C] - v$. Hence $(A, B) \leqslant (C, D)$. $\qquad\square$

Hence to find a nested set of compound-separations which is faithful to G, it suffices to find for every compound-cutvertex w of G a nested set of compound-separations which is faithful to w in G.

The *edge-blocks* of a graph G are the classes of the equivalence relation 'not separable by finitely many edges' on the vertex set of G. So a vertex w of G is a compound-cutvertex if and only if $G - w$ is not infinitely edge-connected, if and only if $G - w$ has at least two edge-blocks.

Lemma 4.7. *Let G be an infinitely edge-connected graph and w be a compound-cutvertex of G. If $G - w$ has only finitely many edge-blocks, then there is a nested set N_w which is faithful to w in G.*

Proof. Every edge-block X of $G - w$ induces the separation $\{X \cup \{w\}, V(G) \smallsetminus X\}$ of $V(G)$ with separator w, and the set N_w containing all of these separations is nested. Moreover, if we orient the separation induced by an edge-block X of $G - w$ as $(X \cup \{w\}, V(G) \smallsetminus X)$, then the orientation of N_w containing all of these oriented separations is a star of separations. Since every two vertices of G that are separated by some compound-separation with separator w are contained in distinct edge-blocks of $G - w$, it is enough to show that every separation in N_w is a compound-separation. For this, let X be an edge-block of $G - w$. Then X sends only finitely many edges to each other edge-block of $G - w$. Since $G - w$ has only finitely many edge-blocks, only finitely many edges run between X and $G - w - X$. Hence $\{X \cup \{w\}, V(G) \smallsetminus X\}$ is a compound-separation. $\qquad\square$

For the proof of Key Lemma 4.5, we only need one more ingredient:

Lemma 4.8. *Let G be an infinitely edge-connected graph and X a finite set of vertices of G. Then the following assertions are complementary:*

 (i) *There is an immersion of $K^{\aleph_0}$ in G with at most one branch vertex in every edge-block of $G - X$;*
 (ii) *$G - X$ has only finitely many edge-blocks.*

To prove Lemma 4.8, we need the following tool, and the notion of tree-cut decompositions by Wollan [9]. Recall that a *near-partition* of a set V is a family of pairwise disjoint subsets $X_\xi \subseteq V$, possibly empty, such that $\bigcup_\xi X_\xi = V$. Let G be a graph, T a tree, and let $\mathcal{X} = (X_t)_{t \in T}$ be a family of vertex sets $X_t \subseteq V(G)$ indexed by the nodes t of T. The pair $(T, \mathcal{X})$ is called a *tree-cut decomposition* of G if $\mathcal{X}$ is a near-partition of $V(G)$. The vertex sets X_t are the *parts* of the tree-cut decomposition $(T, \mathcal{X})$. We say that $(T, \mathcal{X})$ is a tree-cut decomposition *into* these parts. Whenever a tree-cut decomposition is introduced as $(T, \mathcal{X})$, we tacitly assume that $\mathcal{X} = (X_t)_{t \in T}$. If $(T, \mathcal{X})$ is a tree-cut decomposition, then we associate with every edge $t_1 t_2 \in T$ its *adhesion set* $E_G(\bigcup_{t \in T_1} X_t, \bigcup_{t \in T_2} X_t)$ where T_1 and T_2 are the two components of $T - t_1 t_2$ with $t_1 \in T_1$ and $t_2 \in T_2$. A tree-cut decomposition has *finite adhesion* if all its adhesion sets are finite.

Theorem 4.9 ([7, Theorem 5.1]). *Every connected graph has a tree-cut decomposition of finite adhesion into its edge-blocks.*

We remark that, since edge-blocks are non-empty, the parts of the tree-cut decomposition in Theorem 4.9 form a partition of the vertex set. Hence, every adhesion set of this tree-cut decomposition is a cut.

Proof of Lemma 4.8. Assertions (i) and (ii) clearly exclude each other, so it suffices to prove ¬(ii)→(i).

We assume that $G - X$ has infinitely many edge-blocks. Then we find an infinite set $\{U_n : n \in \mathbb{N}\}$ of subsets of $V(G) \smallsetminus X$ such that all $\{U_n \cup X, V(G) \smallsetminus U_n\}$ are compound-separations of G (with separator X), as follows. By Theorem 4.9, $G - X$ admits a tree-cut decomposition $(T, \mathcal{X})$ of finite adhesion into its edge-blocks. By König's Infinity Lemma (see e.g. [2, Lemma 8.1.2]), the tree T contains a vertex of infinite degree or a ray. If $t_0, t_1, \ldots$ are infinitely many neighbours of some vertex $t \in T$, then for every $n \in \mathbb{N}$ we let U_n be the side of the finite cut of $G - X$ induced by the edge $t_n t \in T$ that does not include X_t. Otherwise, we find a ray $t_0 t_1 \ldots$ in T. We denote the finite cut of $G - X$ induced by the edge $t_n t_{n+1} \in T$ by $\{A_n, B_n\}$ so that A_n includes X_{t_n}. We set $U_0 := A_0$ and, for every natural number $n \geqslant 1$, we let $U_n := A_n \cap B_{n-1}$. It is straightforward to see that the resulting partition $\{U_n : n \in \mathbb{N}\}$ is as desired in either case.

For every $n \in \mathbb{N}$, we fix a vertex $u_n \in U_n$ and an infinite system $\mathcal{P}_n$ of pairwise edge-disjoint u_n–X paths in G. Since $\{U_n \cup X, V(G) \smallsetminus U_n\}$ is a compound-separation of G, all but finitely many of these paths are contained in $G[U_n \cup X]$; thus, we may assume without loss of generality that all of them are contained in $G[U_n \cup X]$. By the pigeonhole principle, there is a vertex x in the finite vertex set X such that there is an infinite set M of natural numbers n for which infinitely many of the paths in $\mathcal{P}_n$ end in x. Now we find an immersion of $K^{\aleph_0}$ in G with $\{u_n : n \in M\}$ as its set of branch vertices as follows: we greedily connect every two branch vertices u_i and u_j with a path that avoids all other branch vertices and the already chosen finitely many edges using the path systems $\mathcal{P}_i$ and $\mathcal{P}_j$. $\qquad\square$

Proof of Key Lemma 4.5. Let G be an infinitely edge-connected graph and suppose that $K^{\aleph_0}$ is not immersed in G. For every compound-cutvertex w of G, the graph $G - w$ has only finitely many edge-blocks by Lemma 4.8, and so there is a nested set N_w which is faithful to w in G by Lemma 4.7. The union N of all the sets N_w is nested by Lemma 4.6, and N is faithful to G by construction. We remark that, if G has no compound-cutvertex, then $N = \emptyset$ is faithful to G. $\qquad\square$

4.4. **Finding the halved Farey graph in wild structures.** Our next aim is to prove this:

Key Lemma 4.10. *Let G be an infinitely edge-connected graph and S a set of oriented unitary compound-separations with pairwise distinct separators. If $(S, \leqslant)$ is order-isomorphic to $\mathbb{Q}$, then the halved Farey graph is immersed in G.* $\qquad\square$

We have split the proof into two halves which are represented by the following two lemmas:

Lemma 4.11. *The halved Farey graph is immersed in any graph that is defined by a wildly presented grain line.*

Proof. Let $(\mathcal{L}, \mathcal{P})$ be a wildly presented grain line. We define the set U of branch vertices as follows. First, let $U_0 := \{x, y\}$, where x and y are the two endvertices of the grain line. Assume that we have already chosen $U_0, \ldots, U_n$ for some natural number n such that for every $m \in \{1, \ldots, n\}$

  – all vertices in $U_m \smallsetminus U_{m-1}$ are in $\mathcal{P}$-depth m, and

– for every $u \in U_{m-1}$ with successor u' in $\mathcal{L} \restriction U_{m-1}$, there is precisely one $v \in U_m$ such that $u <_L v <_L u'$.

We order U_n according to $\leqslant_L$ as $u_0 <_L u_1 <_L \cdots <_L u_\ell$. Now we obtain U_{n+1} from U_n by adding exactly one vertex $v_k \in L$ for every $k \in \{1, \ldots, \ell\}$ such that v_k is an internal vertex of the subpath $u_{k-1} P_{n+1} u_k$ and $u_{k-1} <_L v_k <_L u_k$: Let u'_k be the $\leqslant_L$-minimal vertex in L in $\mathcal{P}$-depth at most n with $u_{k-1} <_L u'_k \leqslant_L u_k$. Since $(\mathcal{L}, \mathcal{P})$ is wildly presented, there is an internal vertex v_k of $u_{k-1} P_{n+1} u'_k$ with $u_{k-1} <_L v_k <_L u'_k$. In particular, $v_k <_L u_k$ and, by the choice of u'_k, the vertex v_k is in $\mathcal{P}$-depth $n+1$.

Then we set $U := \bigcup_{n \in \mathbb{N}} U_n$. By construction we have that, for every $n \in \mathbb{N}$, the path P_n induces the same linear order on U_n as $\leqslant_L$ and, since $U \smallsetminus U_n$ consists only of vertices in $\mathcal{P}$-depth greater than n, the path P_n is also disjoint from $U \smallsetminus U_n$. Thus, we have found an immersion of the halved Farey graph in $\bigcup \mathcal{P}$ with U as its set of branch vertices. $\qquad \square$

Lemma 4.12. *Let G be an infinitely edge-connected graph. Let S be a set of oriented unitary compound-separations of G with pairwise distinct separators. If $(S, \leqslant)$ is order-isomorphic to $\mathbb{Q} \cap [0, 1]$, then there exists a wildly presented grain line $(\mathcal{L}, \mathcal{P})$ such that $\bigcup \mathcal{P}$ is immersed in G.*

Proof. For every oriented unitary compound-separation s in S, let v_s be the separator of s. Put $L := \{v_s : s \in S\}$. The linear order on S induces the linear order $\leqslant_L$ on L that is defined by letting $v_s \leqslant_L v_t$ whenever $s \leqslant t$. Then $\mathcal{L} := (L, \leqslant_L)$ is order-isomorphic to $(S, \leqslant)$, and in particular to $\mathbb{Q} \cap [0, 1]$.

Let x and y be the minimal and maximal element of $\mathcal{L}$, respectively. Since G is infinitely edge-connected, there is an infinite system $\mathcal{P}''$ of pairwise edge-disjoint x–y paths in G. By Lemma 2.2, we find an x–y grain line $(\mathcal{L}', \mathcal{P}')$ in G such that every path in $\mathcal{P}'$ is also contained in $\mathcal{P}''$. By the definition of x and y, we have that, for every oriented compound-separation $s = (A, B) \in S$, the vertices x and y are contained in A and B, respectively. Thus, only finitely many paths in $\mathcal{P}''$ avoid v_s. Moreover, if s and t are two oriented compound-separations in S with $s \leqslant t$, then, for all but finitely many x–y paths P in $\mathcal{P}''$, the vertex v_s comes before v_t on P. Hence, $L \subseteq L'$ and $\leqslant_L$ agrees with $\leqslant_{L'}$ on L. Since $\mathcal{L}$ is order-isomorphic to $\mathbb{Q} \cap [0, 1]$, we may additionally assume, by passing to a subsequence of $\mathcal{P}'$, that for every $n \geqslant 1$ and every two vertices u and v in $L \cap L'_{<n}$ with $u <_L v$, there is an internal vertex w of $u P'_n v$ which is contained in L and satisfies $u <_L w <_L v$. Let us denote this property by $(*)$.

For every $n \in \mathbb{N}$, let P_n be the path obtained from $P'_n \in \mathcal{P}'$ by replacing each L-path $Q \subseteq P'_n$ that has an internal vertex with an edge $e(Q)$ that joins the ends of Q (this edge need not be an edge in G). By $(*)$, the paths in $\mathcal{P} := (P_n)_{n \in \mathbb{N}}$ are pairwise edge-disjoint, and $(\mathcal{L}, \mathcal{P})$ is a wildly presented grain line. Moreover, $\bigcup \mathcal{P}$ is immersed in $\bigcup \mathcal{P}'$: we take the identity on the vertices and edges, except that we map every edge of the form $e(Q)$ to the path Q. $\qquad \square$

Proof of Key Lemma 4.10. This follows directly from Lemma 4.11 and Lemma 4.12. $\qquad \square$

4.5. **Decomposing along compound-separations.** The following key lemma is the final ingredient for the proof of Theorem 1:

Key Lemma 4.13. *Let G be an infinitely edge-connected graph. Then at least one of the following is true:*

   (i) *The halved Farey graph is immersed in G;*
   (ii) *G is 2-compound-connected;*
  (iii) *there is a unitary compound-separation $\{A, B\}$ of G such that $G[B]$ is 2-compound-connected.*

Proof. Let us assume that (i) does not hold; in particular, $K^{\aleph_0}$ is not immersed in G. Thus, by Key Lemma 4.5, there is a nested set N of unitary compound-separations of G which is faithful to G.

Let $\vec{N}$ be the set consisting of all orientations of the compound-separations in N, equipped with the usual order $\leqslant$ of oriented separations. Note that we may assume that $\vec{N}$ is non-empty: otherwise, G is 2-compound-connected, and thus (ii) holds.

We claim that if $(\vec{N}, \leqslant)$ has no maximal or minimal element, then (iii) holds. So suppose that (A, B) is maximal in $(\vec{N}, \leqslant)$, say (the other case is analogous). Then $G[B]$ is infinitely edge-connected by Key Lemma 4.3. Now if we assume for a contradiction that $G[B]$ is not 2-compound-connected, then there is a unitary compound-separation $\{C, D\}$ of $G[B]$ that separates some two vertices u and v of $G[B]$. After renaming the sides of $\{C, D\}$, we may assume that C contains the separator w of (A, B). Then $\{A \cup C, D\}$ is a unitary compound-separation of G which still separates u and v. However, no compound-separation in N separates u and v, since u and v lie in B and (A, B) is maximal in $(\vec{N}, \leqslant)$. Hence N is not faithful to G, a contradiction.

Therefore, we may assume that $(\vec{N}, \leqslant)$ has no maximal or minimal element. Hence, using Zorn's lemma, we find an infinite subset S' of $\vec{N}$ such that:

- $\leqslant$ orders S' linearly;
- S' is unbounded in $\vec{N}$;
- S' is an interval of $\vec{N}$: for every two $r, t \in S'$ and $s \in \vec{N}$ with $r \leqslant s \leqslant t$, we have $s \in S'$.

Let w be any vertex of G. Since N is faithful to G, it either does not contain a compound-separation with separator w, or, if w is a compound-cutvertex of G, then there is an orientation of the set of all compound-separations in N with separator w which is a star of separations. So it follows from $(S', \leqslant)$ being a linear order that there are at most two oriented compound-separations in S' with separator w. Moreover, if there are two such separations $s_1 < s_2$, then the separator of any oriented compound-separation t with $s_1 < t < s_2$ is a subset of $\{w\}$, thus equal to $\{w\}$, so s_2 is the successor of s_1 in $(S', \leqslant)$. Let S be a subset of S' obtained by deleting exactly one element of every pair of oriented compound-separations in S' with the same separator. Thus, the separators of the oriented compound-separations in S are pairwise distinct.

Let R be the set of all elements of S that have a successor in $(S, \leqslant)$. We claim that R is infinite. To see this, let us assume for a contradiction that R is finite. Then $(S \smallsetminus R, \leqslant)$ is an infinite set equipped with a dense linear order. So we can greedily find a subset of $S \smallsetminus R$ which is ordered by $\leqslant$ like $\mathbb{Q}$, and apply Key Lemma 4.10 to obtain a halved Farey graph immersed in G, contradicting our assumption that (i) fails. Thus, R is infinite.

By passing to a subset of R, we may assume without loss of generality that $(R, \leqslant)$ is order-isomorphic to ω or to the inverse of ω. In the following, we will assume that $(R, \leqslant)$ is order-isomorphic to ω, as the other case is analogous. Let $s_0 < s_2 < s_4 < \ldots$ be an enumeration of R. For every $n \in \mathbb{N}$, let s_{2n+1} be the successor of s_{2n} in S. Moreover, for every $n \in \mathbb{N}$ we let $(A_n, B_n) := s_n$, denote the separator of s_n by v_n, and we let $V_0 := A_0$ as well as $V_{n+1} := B_n \cap A_{n+1}$.

We claim that, for every natural number $n \geqslant 1$,

(1) either v_{n-1} and v_n are equal or $G[V_n]$ is infinitely edge-connected, and
(2) if n is odd then v_{n-1} and v_n are distinct and 2-compound-connected in $G[V_n]$.

Let $n \geqslant 1$ be any natural number. If v_{n-1} and v_n are distinct, then applying Key Lemma 4.3 twice yields that $G[V_n]$ is infinitely edge-connected. Thus, (1) holds. Now we additionally assume that n is

odd. Then s_n is the successor of s_{n-1} in S; in particular, s_{n-1} and s_n are distinct. Thus, since the separators of the oriented compound-separations in S are pairwise distinct, v_{n-1} and v_n are distinct. We claim that there is no unitary compound-separation of $G[V_n]$ that separates v_{n-1} and v_n. Indeed, suppose for a contradiction that there is a unitary compound-separation $\{A, B\}$ of $G[V_n]$ with separator w such that $v_{n-1} \in A \smallsetminus B$ and $v_n \in B \smallsetminus A$. Then $\{A_{n-1} \cup A, B \cup B_n\}$ is a unitary compound-separation of G which separates v_{n-1} and v_n. Since N is faithful to G, there is a compound-separation s with the same separator w in N which separates v_{n-1} from v_n, and by Lemma 4.6 we have that $s_{n-1} < s < s_n$ after suitably orienting s. As S' is an interval of $\vec{N}$, we have $s \in S'$. By the construction of S, there is a unique oriented compound-separation $t \in S$ with separator w. Since we have $s_{n-1} < s < s_n$ and since w is distinct from the separators v_{n-1} and v_n of s_{n-1} and s_n, Lemma 4.6 implies $s_{n-1} < t < s_n$. This contradicts the fact that s_n is the successor of s_{n-1} in S. So (2) holds.

To conclude the proof, we construct an immersion of $K^{\aleph_0}$ in G, as follows. For every even natural number $n \geqslant 1$, if v_{n-1} is distinct from v_n then we pick an infinite system $\mathcal{P}_n$ of infinitely many pairwise edge-disjoint v_{n-1}–v_n paths in $G[V_n]$, which exists by (1). For every odd natural number $n \geqslant 1$, we pick an arbitrary vertex u_n in $V_n \smallsetminus \{v_{n-1}, v_n\}$ and—using that $G[V_n]$ is 2-compound-connected by (2)—we greedily pick an infinite system $\mathcal{P}_n$ of pairwise edge-disjoint v_{n-1}–v_n paths in $G[V_n]$ such that infinitely many of these contain u_n and infinitely many do not. Note that a vertex u_n can only lie on paths in $\mathcal{P}_n$ and not on paths in any other $\mathcal{P}_m$. Let $U := \{u_n : n \in 2\mathbb{N} + 1\}$. Finally, we greedily connect every two distinct vertices $u_i, u_j \in U$ with a path avoiding all other u_n and the already used finitely many edges by suitably concatenating paths from the sets $\mathcal{P}_m$ with $m \geqslant 1$. This yields an immersion of $K^{\aleph_0}$ in G with U as its set of branch vertices, and thus (i) holds. $\qquad \square$

4.6. The halved Farey graph is everywhere.
Finally, we are ready to prove Theorem 1:

Proof of Theorem 1. Let G be any infinitely edge-connected graph, and let us assume for a contradiction that the halved Farey graph is not immersed in G. We will recursively construct infinitely edge-connected subgraphs $G_0, G_1, \ldots$ and $H_0, H_1, \ldots$ of G such that, for every $n \in \mathbb{N}$,

(i) $\{G_0, \ldots, G_n\} \cup \{H_n\}$ is a set of pairwise edge-disjoint subgraphs of G,
(ii) if $n \geqslant 1$ then $\{V(G_n), V(H_n)\}$ is a compound-separation of H_{n-1} with $V(G_n) \cap V(H_n) \neq \emptyset$, and
(iii) the intersection graph of $\{V(G_m) : m \leqslant n\}$ is connected.

Later, we will continue working with the graphs G_n, while the graphs H_n only serve as 'reservoirs' during the construction of these two sequences, which we explain next.

To get started, suppose that $n = 0$, so we have to define G_0 and H_0. By Observation 4.2, there is a compound-separation $\{A_0, B_0\}$ of G that separates two vertices of G minimally. By Key Lemma 4.3, $G[A_0]$ and $G[B_0]$ are infinitely edge-connected. Hence, $G_0 := G[A_0]$ and $H_0 := G[B_0] - E_G(A_0 \cap B_0)$ are the desired subgraphs, where $E_G(A_0 \cap B_0)$ denotes the set of all the finitely many edges of G that run between the vertices in $A_0 \cap B_0$.

Suppose next that $n \geqslant 1$ and that we have already chosen $G_0, \ldots, G_{n-1}$ and $H_0, \ldots, H_{n-1}$. By Observation 4.2, there is a compound-separation $\{A_n, B_n\}$ of H_{n-1} that separates two vertices of H_{n-1} minimally. Since the intersection $V(G_{n-1}) \cap V(H_{n-1})$ is nonempty by (ii) for $n - 1$, we can rename the sides of $\{A_n, B_n\}$ so that A_n meets $V(G_{n-1})$. By Key Lemma 4.3, both $H_{n-1}[A_n]$ and $H_{n-1}[B_n]$ are infinitely edge-connected. Hence $G_n := H_{n-1}[A_n]$ and $H_n := H_{n-1}[B_n] - E_{H_{n-1}}(A_n \cap B_n)$ also are infinitely edge-connected. (i) holds by construction. Since H_{n-1} is infinitely edge-connected, the

intersection $A_n \cap B_n$ is nonempty, so (ii) holds. Finally, (iii) holds because we named the sides of $\{A_n, B_n\}$ so that A_n meets $V(G_{n-1})$.

In the next step, we will obtain from the sequence $(G_n)_{n \in \mathbb{N}}$ of infinitely edge-connected graphs a countably infinite collection $\mathcal{U}$ of pairwise edge-disjoint infinitely edge-connected subgraphs of G such that

    (1) infinitely many graphs in $\mathcal{U}$ are 2-compound-connected,

    (2) each graph in $\mathcal{U}$ has a vertex that is contained in no other member of $\mathcal{U}$, and

    (3) the intersection graph of $\mathcal{U}$ is connected.

We construct $\mathcal{U}$ from the empty set by adding one or two graphs for each $n \in \mathbb{N}$, as follows. Given n, we apply Key Lemma 4.13 to G_n, which tells us that G_n is 2-compound-connected or that there is a unitary compound-separation $\{X, Y\}$ of G_n such that $G_n[Y]$ is 2-compound-connected (as the halved Farey graph is not immersed in $G_n \subseteq G$ by assumption). In the former case, we add G_n to $\mathcal{U}$. In the latter case, we add $G_n[X] - E_{G_n}(X \cap Y)$ and $G_n[Y]$ to $\mathcal{U}$.

The 2-compound-connected graphs in $\mathcal{U}$ are infinitely edge-connected by definition, and the graphs of the form $G_n[X] - E_{G_n}(X \cap Y)$ are infinitely edge-connected by Key Lemma 4.3. While properties (1) and (3) follow immediately from the above construction and (iii), property (2) follows from (ii): For every $n \in \mathbb{N}$, (ii) implies that only finitely many vertices of G_n occur in any later G_m with $m > n$. Hence every G_n has only finitely many vertices that occur in any G_m with $m \neq n$. Thus, every graph in $\mathcal{U}$ shares only finitely many vertices with other members of $\mathcal{U}$. As the graphs in $\mathcal{U}$ are infinite, (2) follows.

To obtain a contradiction, we construct an immersion of $K^{\aleph_0}$ in G, as follows. In each 2-compound-connected graph in $\mathcal{U}$ we select one vertex that lies in no other member of $\mathcal{U}$, which is possible by (2). In total, $\aleph_0$ vertices are select in this way, by (1). The selected vertices will be the branch vertices of the immersion. Now (1) and (3) combined with the infinite edge-connectivity of the members of $\mathcal{U}$ allow us to greedily connect every two branch vertices with a path that avoids all other branch vertices and the finitely many edges used by the previously chosen paths. $\square$

References

[1] M. Clay and D. Margalit, *Office Hours with a Geometric Group Theorist*, Princeton University Press, 2017.

[2] R. Diestel, *Graph Theory*, 5th ed., Springer, 2016.

[3] A. Hatcher, *Topology of numbers*, Book in preparation (2017). Available online.

[4] J. Kurkofka, R. Melcher, and M. Pitz, *A strengthening of Halin's grid theorem*, Mathematika **68** (2022), 1009–1013. Available on arXiv:2104.10672.

[5] J. Kurkofka, *The Farey graph is uniquely determined by its connectivity*, Journal of Combinatorial Theory, Series B **151** (2021), 223–234. Available on arXiv:2006.12472.

[6] _____, *Ubiquity and the Farey graph*, European Journal of Combinatorics **95** (2021), no. 103326. Available on arXiv:1912.02147.

[7] _____, *Every infinitely edge-connected graph contains the Farey graph or $T_{\aleph_0} * t$ as a minor*, Mathematische Annalen **382** (2022), 1881–1900. Available on arXiv:2004.06710.

[8] N. Robertson and P. D. Seymour, *Graph minors XXIII. Nash-Williams' immersion conjecture*, Journal of Combinatorial Theory, Series B **100** (2010), no. 2, 181–205.

[9] P. Wollan, *The structure of graphs not admitting a fixed immersion*, Journal of Combinatorial Theory, Series B **110** (2015), 47–66. Available on arXiv:1302.3867.

Universität Hamburg, Department of Mathematics, Bundesstrasse 55 (Geomatikum), 20146 Hamburg, Germany
Email address: paul.knappe@uni-hamburg.de

University of Birmingham, Birmingham, UK
Email address: j.kurkofka@bham.ac.uk