

Universal planar graphs for the topological minor relation

Florian Lehner*

June 9, 2023

Abstract

Huynh et al. recently showed that a countable graph G which contains every countable planar graph as a subgraph must contain arbitrarily large finite complete graphs as topological minors, and an infinite complete graph as a minor. We strengthen this result by showing that the same conclusion holds, if G contains every countable planar graph as a topological minor. In particular, there is no countable planar graph containing every countable planar graph as a topological minor, answering a question by Diestel and Kühn.

Moreover, we construct a locally finite planar graph which contains every locally finite planar graph as a topological minor. This shows that in the above result it is not enough to require that G contains every locally finite planar graph as a topological minor.

1 Introduction

Call a graph U *universal* for a graph class \mathcal{G} , if it contains every element of \mathcal{G} , and let us say that \mathcal{G} *contains a universal element* if there is a universal graph U for \mathcal{G} which is contained in \mathcal{G} . Depending on the precise definition of containment, this leads to different notions of universality and different universal graphs.

Clearly, these notions are not independent, for instance, if U is universal for \mathcal{G} with respect to subgraph containment, then it is also universal with respect to minor containment because every subgraph is also a minor. Research has mostly focused on the strongest notions of universality, that is universality with respect to subgraph or induced subgraph containment [1, 3, 7, 9, 13, 14, 15]. However, considering weaker universality notions for graph classes which do not admit universal elements in this strong sense also leads to interesting questions and beautiful results [4, 6, 11, 12].

In the present paper, we are interested in universal elements for classes of planar graphs. A classic result by Pach [14] states that the class of all planar graphs does not

*Much of the research leading to the results presented in this paper was carried out while the author was supported by the Austrian Science Fund (FWF) Grant no. P31889-N35

contain a universal element with respect to the subgraph relation, thereby providing a negative answer to a question which Pach attributes to Ulam.

In contrast to this result, Diestel and Kühn [4] show that there is a countable planar graph containing all countable planar graphs as minors. This immediately leads to the question for which notions of containment in between the subgraph relation and the minor relation the class of planar graphs contains a universal element. In particular, Diestel and Kühn ask [4, Problem 6] whether the class of planar graphs contains a universal element with respect to the topological minor relation. Our main result provides a negative answer to this question.

Theorem 1.1. *The class of countable planar graphs does not contain a universal graph with respect to the topological minor relation.*

We point out that Theorem 1.1 has been proved independently by Krill in his master's thesis [10], see also the preprint [11]. While our proof is longer and more involved than the one presented in [10], it in fact also shows a significantly stronger result.

In a recent preprint, Huynh, Mohar, Šámal, Thomassen, and Wood [8] investigate how sparse a graph which contains all planar graphs as subgraphs can be. In other words, rather than relaxing the notion of containment, they ask how much the requirement of planarity of the universal graph needs to be relaxed in order to get a different answer to Ulam's question. They obtain two complementary results. On the one hand, they show that there are universal graphs for the class of planar graphs which share some key properties with planar graphs such as linear colouring numbers, linear expansion, and balanced separators of size $O(\sqrt{n})$ in every n -vertex subgraph. On the other hand, they prove that a universal graph for the class of countable planar graphs with respect to the subgraph relation is in some sense very far from being planar: it contains arbitrarily large complete graphs as topological minors, and the countably infinite complete graph as a minor. In Section 3, we prove the following strengthening of the latter result which immediately implies Theorem 1.1.

Theorem 1.2. *Let G be a countable graph containing every countable planar graph as a topological minor. Then*

1. *G contains an infinite complete minor, and*
2. *G contains arbitrarily large finite complete topological minors.*

In Section 4, we turn our attention to locally finite graphs, that is, graphs in which every vertex has only finitely many neighbours. The main result of this section shows that the conclusion of Theorem 1.2 no longer holds if we only require G to contain all locally finite planar graphs as topological minors.

Theorem 1.3. *The class of locally finite planar graphs contains a universal element with respect to the topological minor relation.*

To fully appreciate this result, it is worth mentioning that the aforementioned universality results concerning the subgraph relation and the minor relation are unaffected

by restricting to locally finite graphs. Pach's proof from [14] in fact shows that there is no planar graph containing every locally finite planar graph as a subgraph, and the conclusion of the result by Huynh et al. from [8] mentioned above still holds for graphs containing all locally finite planar graphs as subgraphs. Moreover, the universal planar graph for the minor relation constructed in [4] is in fact locally finite. In light of this, it is perhaps surprising that local finiteness makes a big difference when considering the topological minor relation.

2 Preliminaries

The purpose of this section is to recall basic definitions and set up some notation. For graph theoretic notions not explicitly defined, we follow [2].

A graph G consists of a set $V(G)$ of vertices and a set $E(G)$ of edges. Given two graphs G and H , we denote by $G \cup H$ the graph with vertex set $V(G) \cup V(H)$ and edge set $E(G) \cup E(H)$. Similarly define $G \cap H$. We stress that the graphs in a union do not have to be disjoint, in particular we will often consider unions of graphs which have vertices in common.

It will be convenient to consider cycles in graphs as cyclic sequences of vertices. Given a sequence $X = (x_1, \dots, x_n)$, a sequence of the form $(x_k, \dots, x_n, x_1, \dots, x_{k-1})$ is called a *cyclic shift* of X . Call two sequences *cyclically equivalent* if they are cyclic shifts of one another. Call an equivalence class with respect to this relation a *cyclic sequence* and denote the cyclic sequence containing (x_1, \dots, x_n) by $[x_1, \dots, x_n]$. A *directed cycle* in a graph G is a cyclic sequence $C = [v_1, \dots, v_n]$ of vertices such that $v_i v_{i+1} \in E(G)$ for every $i < n$ and $v_1 v_n \in E(G)$. Note that this assigns a direction to the cycle even though the graphs we consider are usually undirected.

Let us call a sequence $Y = (y_1, \dots, y_k)$ *cyclically ordered* with respect to a cyclic sequence $X = [x_1, \dots, x_n]$, if some representative of X contains Y as a subsequence. If X is clear from the context, we simply call Y cyclically ordered. Clearly, if Y is cyclically ordered, then so is any sequence which is cyclically equivalent to Y ; we may thus extend this notion to the case where Y is a cyclic sequence.

All (cyclic) sequences considered from now on will be sequences of vertices of graphs. For such a (cyclic) sequence X , let us denote the set of all vertices appearing in X by $V(X)$. Let X_1 and X_2 be cyclic sequences of vertices, let $Y \subseteq V(X_1)$, and let (y_1, \dots, y_n) be a cyclically ordered enumeration of Y . We say that a function $\phi: Y \rightarrow V(X_2)$ *preserves* the cyclic order if $(\phi(y_1), \dots, \phi(y_k))$ is cyclically ordered with respect to X_2 . We say that such a function *reverses* the cyclic order if $(\phi(y_k), \dots, \phi(y_1))$ is cyclically ordered with respect to X_2 .

A plane embedding of a graph G assigns to each vertex $v \in V(G)$ a point $\iota(v) \in \mathbb{R}^2$ and to each edge $e = uv \in E(G)$ a polygonal arc $\iota(e)$ in \mathbb{R}^2 connecting $\iota(u)$ to $\iota(v)$ such that the map $\iota: V(G) \rightarrow \mathbb{R}^2$ is injective, and for any two distinct edges e and f the arcs $\iota(e)$ and $\iota(f)$ are internally disjoint. By a slight abuse of notation we write $\iota(G)$ for $\bigcup_{v \in V(G)} \{\iota(v)\} \cup \bigcup_{e \in E(G)} \iota(e)$. Let us call a graph *planar* if it has a plane embedding. Note that we do not forbid that $\iota(G)$ has accumulation points.

The following theorem by Dirac and Schuster [5] gives a necessary and sufficient condition for the planarity of countable graphs.

Theorem 2.1. *A countable graph is planar if and only if all of its finite subgraphs are planar.*

A similar condition for arbitrary graphs has been given by Wagner [16].

Theorem 2.2. *An arbitrary graph is planar if and only if it has at most continuum many vertices and at most countably many vertices of degree greater than 2, and all of its finite subgraphs are planar.*

Similarly to embeddings of graphs in the plane, we can also define embeddings of graphs in other graphs. For two graphs G and H , a G -embedding ι of H assigns to every $v \in V(H)$ a vertex $\iota(v) \in V(G)$ and to every edge $e = uv \in E(H)$ a $\iota(u)$ - $\iota(v)$ -path $\iota(e)$ in G such that the map $\iota: V(H) \rightarrow V(G)$ is injective, and for distinct edges e and f the paths $\iota(e)$ and $\iota(f)$ are internally disjoint. We call H a *topological minor* of G if there is a G -embedding of H .

Let us say that two G -embeddings ι, ι' of a graph H agree on $S \subseteq V(H) \cup E(H)$, if for any $s \in S$ we have $\iota(s) = \iota'(s)$. By slight abuse of notation we extend this notion to G -embeddings of different graphs as follows. Let $S \subseteq V(H) \cup E(H)$ and let $S' \subseteq V(H') \cup E(H')$. Let $f: S \rightarrow S'$ be a bijection. Let ι be a G -embedding of H , and let ι' be a G -embedding of H' . We say that ι and ι' agree via f on S , if $\iota(s) = \iota'(f(s))$ for all $s \in S$. If H and H' have a subgraph in common and all elements in S are contained in this subgraph, then we omit f and tacitly assume that f is the identity. This in particular includes the case where H and H' are obtained from the same graph by adding some vertices and edges.

We say that a family \mathcal{I} of G -embeddings (possibly of different graphs) agrees on S if any pair ι, ι' agrees on S . We point out that we will only need this notion for sets contained in common subgraphs of all involved graphs. We denote the common image of $s \in S$ under all $\iota \in \mathcal{I}$ by $\mathcal{I}(s)$.

Theorems 2.1 and 2.2 do not extend to this notion of embedding: there are graphs G and H such that all finite subgraphs of H admit a G -embedding, but H does not. For instance, we can let G be the disjoint union of all finite graphs (up to isomorphism), and let H be any graph with at least one infinite component.

Let $(G_n)_{n \in \mathbb{N}}$ be an increasing sequence of graphs, that is, G_n is a subgraph of G_{n+1} for every $n \in \mathbb{N}$. We define $\lim_{n \rightarrow \infty} G_n$ as the graph with vertex set $\bigcup_{n \in \mathbb{N}} V(G_n)$ and edge set $\bigcup_{n \in \mathbb{N}} E(G_n)$.

Lemma 2.3. *Let $(H_n)_{n \in \mathbb{N}}$ be an increasing sequence of graphs and let $H = \lim_{n \rightarrow \infty} H_n$. If there are G -embeddings ι_n of H_n such that ι_{n+1} and ι_n agree on H_n for every n , then there is a G -embedding ι of H which agrees with ι_n on H_n for every n .*

Proof. For $x \in V(H) \cup E(H)$ pick n large enough that $x \in V(H_n) \cup E(H_n)$ and set $\iota(x) = \iota_n(x)$. The conditions of the lemma ensure that this is unambiguous and defines a G -embedding of H . \square



Figure 1: The triangular wedge W and the triple wedge \overline{W} .

Theorems 2.1 and 2.2, and Lemma 2.3 tell us that we can obtain (plane or G -) embeddings of infinite graphs by constructing embeddings of an increasing sequence of finite subgraphs. In the remainder of this section we recall some well known facts and make some easy observations about finite planar graphs.

If ι is a plane embedding of a finite connected graph G , then $\mathbb{R}^2 \setminus \iota(G)$ consists of finitely many open disks and one unbounded region which we call the *faces* of the embedding; the unbounded region is called the *outer face*, all other regions are called *interior faces*. If ι is a plane embedding of an infinite graph G , we still call a connected component of $\mathbb{R}^2 \setminus \iota(G)$ a *face* of the embedding. We note that the complement of an embedding of an infinite planar graph can be much more complicated due to accumulation points of the embedding; in particular, faces of embeddings of infinite graphs are not necessarily homeomorphic to disks.

3 No universal countable planar graph

Before turning to the proof of Theorem 1.2, we briefly give an overview of the key ideas. The proof is similar to the proof of the analogous theorem for subgraph embeddings in [8, Section 3] which can be summarised as follows: take an uncountable family of grid-like planar graphs, prove that any graph containing all of these graphs simultaneously must also contain another grid-like planar graph with certain gadgets attached to it, and then use these gadgets to find a countable complete minor as well as large complete topological minors.

To explain the proof ideas in more detail, we need to describe the grid-like graphs that we are going to use. The triangular wedge W is a ‘triangular quarter grid’, the triple wedge \overline{W} is obtained from the disjoint union of three copies of the triangular wedge by adding edges between these copies, see Figure 1 for a sketch, and Section 3.1 for a precise definition.

We can construct an uncountable family of planar graphs from the triple wedge by adding one of the two diagonals in infinitely many of the faces bounded by cycles of length 4. Denote these graphs by $\overline{W}(\alpha)$ with $\alpha \in \{0, 1\}^{\mathbb{N}}$. For a graph G , let us write G^{\parallel} for the graph obtained from G by replacing every edge by a countably infinite collection of internally disjoint paths with the same endpoints as this edge. We note that each graph $\overline{W}(\alpha)^{\parallel}$ is planar by Theorem 2.1.




Figure 2: A pair of crossing paths attached to the triangular wedge. Note that the paths are disjoint, in particular there is no vertex at the intersection in the drawing

The first and most difficult step in the proof of Theorem 1.2 is to show that a graph which contains every $\overline{W}(\alpha)^{\parallel}$ as a topological minor also contains an embedding some graph W^+ obtained from the triangular wedge W by adding infinitely many pairwise disjoint paths P_i satisfying the following two properties.

- (i) Each P_i starts at a vertex on the leftmost vertical ray of W and ends at a vertex of the bottom horizontal ray of W , and is otherwise disjoint from W .
- (ii) There are infinitely many pairs i, j such that P_i and P_j cross in the sense that the first vertex of P_i lies above the first vertex of P_j and the last vertex of P_i lies left of the last vertex of P_j , see Figure 2.

Once we have established the existence of an embedding of such a graph W^+ , the second step is to show that W^+ contains a set of infinitely many disjoint rays such that each pair of them is connected by infinitely many edges. Clearly, such a set of rays gives a countably infinite complete minor whose branch sets are precisely those rays.

The proof of second part of Theorem 1.2, that is, showing that we can guarantee arbitrarily large finite complete graphs follows the same outline, but requires some small tweaks to both parts of the proof.

As mentioned above, similar ideas are used in [8, Section 3]. We briefly comment on the most important differences between the proof of Theorem 1.2 and the analogous result for subgraph embeddings.

Perhaps the most important difference and the main technical difficulty we need to overcome is that topological minor embeddings are a lot more flexible than subgraph embeddings. Indeed, we note that the graphs $\overline{W}(\alpha)$ are f -isoperimetric plane triangulations in the sense of [8], and thus any graph containing all of them as subgraphs must contain a countably infinite complete minor by [8, Lemma 3.9]. However, each $\overline{W}(\alpha)$ is locally finite, so Theorem 1.3 implies that there is a planar graph containing each $\overline{W}(\alpha)$ as a topological minor. In particular, in the first proof step is not enough to consider the graphs $\overline{W}(\alpha)$ instead of $\overline{W}(\alpha)^{\parallel}$.

The second proof step is similar to the proof of [8, Lemma 3.8]. In fact, if we only want to show the existence of a countably infinite complete minor, then we can directly apply [8, Lemma 3.8] to a suitable minor of W^+ and do not have to explicitly construct any family of rays. However, the modification which yields arbitrarily large finite complete




Figure 3: Triangular wedge. The bold path is the 5-th layer W_5 , the subgraph inside the shaded area is the 2-6-strip $W_{2,6}$.

minors (similarly to [8, Theorem 3.11]) requires a modification of \overline{W} , as well as better control over the family of rays in the graph playing the role of W^+ in the modification.

3.1 Notation and definitions

Before we turn to the proof details, we rigorously define the graphs mentioned in the above proof sketch and set up some notation for them.

The *triangular wedge* W is the graph with vertex set $\mathbb{N}_0 \times \mathbb{N}_0$ and edges between vertices whose coordinates differ by $(1, 0)$, $(0, 1)$ or $(1, -1)$, see Figure 3. The k -th *layer* W_k of W is the subgraph induced by the vertices whose coordinates (i, j) satisfy $i + j = k$. For $0 \leq i \leq k$, let w_k^i be the vertex with coordinates $(i, k - i)$, in other words w_k^i is the i -th vertex of W_k in left-to-right order. For $m < n$, define the m - n -*strip* $W_{m,n}$ as the subgraph of W induced by the vertices in $\bigcup_{m \leq k \leq n} W_k$. We allow $n = \infty$ and define $W_{m,\infty}$ as the subgraph of W induced by the vertices in $\bigcup_{m \leq k} W_k$.

Assume that W is a subgraph of some larger graph G . A *bypass* for $W_{m,n}$ in G is a w_a^0 - w_b^b -path P with $m < a < n, m < b < n$ which meets W only at its endpoints. If we do not care about m and n , we sometimes also simply call P a bypass for W . A bypass P from w_a^0 to w_b^b and bypass P' from $w_{a'}^0$ to $w_{b'}^{b'}$ are said to be *crossing* if either $a < a'$ and $b' < b$, or $a > a'$ and $b' > b$, see Figure 2.

Let X, Y , and Z be isomorphic copies of W . We denote by $X_k, Y_k, Z_k, x_k^i, y_k^i, z_k^i, X_{m,n}, Y_{m,n}$, and $Z_{m,n}$ the respective copies of W_k, w_k^i , and $W_{m,n}$. The *triple wedge* \overline{W} is the graph obtained from the disjoint union of X, Y , and Z by adding edges between x_k^k and y_k^0 , between y_k^k and z_k^0 , and between z_k^k and x_k^0 for all $k \in \mathbb{N}_0$, see Figure 4. The *annulus* $\overline{W}_{m,n}$ is the subgraph of \overline{W} induced by the vertices of $X_{m,n}, Y_{m,n}$, and $Z_{m,n}$. Call a vertex of \overline{W} *enclosed*, if its neighbourhood is entirely contained in one of the three wedges X, Y , or Z .

For $k \in \mathbb{N}$, let C_k be the 4-cycle $x_k^k, y_k^0, y_{k+1}^0, x_{k+1}^{k+1}$ in \overline{W} . For $\alpha \in \{0, 1\}^{\mathbb{N}}$, denote by $\overline{W}(\alpha)$ the graph obtained from \overline{W} by adding the edge from x_k^k to y_{k+1}^0 if $\alpha_k = 0$ and from x_{k+1}^{k+1} to y_k^0 if $\alpha_k = 1$. Note that $\overline{W}(\alpha)$ is planar since each C_k is a facial cycle in the




Figure 4: The triple wedge \overline{W} . The subgraph inside the shaded area is the annulus $\overline{W}_{2,4}$.

embedding of \overline{W} shown in Figure 4. Denote by $\overline{W}_{m,n}(\alpha)$ the subgraph of $\overline{W}(\alpha)$ induced by the vertices of $\overline{W}_{m,n}$.

The *infinitely parallel blow-up* G^{\parallel} of a graph G is the graph obtained from G by replacing every edge e of G by a countably infinite set of paths of length 2 connecting the endpoints of e . Clearly, all vertices of G^{\parallel} have either infinite degree, or degree 2. Note that if G is a countable planar graph, then G^{\parallel} is planar by Theorem 2.1. We call the vertices of infinite degree in (a subdivision of) G^{\parallel} *original vertices*, and the vertices of degree 2 *new vertices*. If no confusion is possible, we identify original vertices with the corresponding vertices in G , in particular we call a pair of original vertices *adjacent* if the corresponding vertices in G are adjacent.

3.2 Finding a triangular wedge with many disjoint bypasses

Our goal in this subsection is to complete the first step in the proof sketch above. In other words, we will show that if all graphs $\overline{W}(\alpha)^{\parallel}$ embed in a graph G , then G also contains an embedding of W with infinitely many pairs of crossing bypasses.

In fact, we prove the following slightly stronger statement. Assume that uncountably many graphs $\overline{W}_{m,\infty}(\alpha)^{\parallel}$ embed in a graph G , and that all of the embeddings coincide on the original vertices in \overline{W}_m , then G contains an embedding of $W_{m,\infty}$ with infinitely many pairs of crossing bypasses, and the embedding of the vertices of W_m coincides with the embedding of the vertices in X_m , Y_m , or Z_m . We point out that this strengthening is necessary in order to be able to use the result in the (slightly more technical) proof of the second part of Theorem 1.2.

To formally state the main result of this subsection, recall that a set \mathcal{I} of embeddings is said to agree on a set S of vertices, if $\iota(s) = \iota'(s)$ for every pair $\iota, \iota' \in \mathcal{I}$ and every $s \in S$. Further recall that in this case we write $\mathcal{I}(s)$ for the image of $s \in S$ under some (and thus every) $\iota \in \mathcal{I}$.

Lemma 3.1. *Let G be a countable graph and let $m > 0$. Let $A \subseteq \{0, 1\}^{\mathbb{N}}$ be an uncountable set, and assume there are G -embeddings ι_α of $\overline{W}_{m,\infty}(\alpha)$ which agree on $V(\overline{W}_m)$. Let $\mathcal{I} = \{\iota_\alpha \mid \alpha \in A\}$.*

There is a graph W^+ obtained from $W_{m,\infty}$ by adding an infinite family of disjoint bypasses containing infinitely many crossing pairs, and a G -embedding ι of W^+ such that there is $\xi \in \{x, y, z\}$ with $\iota(w_m^i) = \mathcal{I}(\xi_m^i)$ for $0 \leq i \leq m$.

We start by showing how adding an edge to \overline{W} creates a pair of crossing bypasses for one of the three wedges X , Y , or Z . By symmetry the following lemma clearly remains true if we swap the roles of X , Y , and Z .

Lemma 3.2. *Let u and v be non-adjacent vertices in $X_{m+1,n-1} \cup Y_{m+1,n-1}$. Assume that u is an enclosed vertex. The graph obtained from $\overline{W}_{m,n}$ by adding the edge uv contains a crossing pair of bypasses for Z .*

Proof. The following notation will be convenient in the proof: If a, b in \overline{W}_k are vertices none of whose neighbourhoods is entirely contained in Z , then there is a unique a - b -path in \overline{W}_k all of whose internal vertices lie in X or Y . Let us denote this path by $P(a, b)$.

First assume that u and v lie in different wedges. By symmetry, we may without loss of generality assume that $u \in X$ and $v \in Y$. There are i, j, k, l such that $u = x_k^i$ and $v = y_l^j$. The concatenation of $P(z_k^k, u)$, the edge uv , and $P(v, z_l^0)$ gives a bypass P for Z . The concatenation of $P(z_m^m, x_m^m)$, the path $x_m^m, x_{m+1}^{m+1}, x_{m+2}^{m+2}, \dots, x_n^n$, and $P(x_n^n, z_n^0)$ gives another bypass P' for Z . Note that since u is enclosed, we know that $0 < i < k$, hence P does not contain x_t^t for any t . Since $m < k, l < n$, we conclude that P and P' are disjoint and crossing.

If u and v lie in the same wedge, then we may without loss of generality assume that they both lie in X . There are i, j, k, l such that $u = x_k^i$ and $v = x_l^j$. Let us assume that $i \leq j$, the case $i \geq j$ is completely analogous. As before, the concatenation of $P(z_k^k, u)$, the edge uv , and $P(v, z_l^0)$ gives a bypass P for Z . Define another bypass P' for Z by concatenation of $P(z_m^m, x_m^{i-1})$, the path

$$x_m^{i-1}, x_{m+1}^{i-1}, x_{m+2}^{i-1}, \dots, x_{k-1}^{i-1}, x_{k-1}^i, x_k^{i+1}, x_{k+1}^{i+1}, x_{k+1}^i, x_{k+1}^{i-1}, x_{k+2}^{i-1}, \dots, x_n^{i-1}$$

and $P(x_n^{i-1}, z_n^0)$; Figure 5 illustrates the restrictions of P and P' to the wedge X . Note that the only vertices x_s^t on P' for which $m < s < n$ and $t > i$ are neighbours of u . If P contained one of these vertices then either v would be a neighbour of u , or $i > j$ both of which contradict our assumptions. $m < k, l < n$, we again conclude that P and P' are disjoint and crossing. \square

Proof of Lemma 3.1. To simplify notation, throughout this proof we let $H = \overline{W}_{m,\infty}$, let $H_n = \overline{W}_{m,n}$, let $H(\alpha) = \overline{W}_{m,\infty}(\alpha)$, and let $H_n(\alpha) = \overline{W}_{m,n}(\alpha)$.

First assume that \mathcal{I} agrees on $V(H)$, that is, each original vertex has the same image under every ι_α . Let $K \subseteq \mathbb{N}$ be the set of all k for which there are $\alpha, \beta \in A$ with $\alpha_k \neq \beta_k$. Since \mathcal{I} is infinite, the set K is infinite as well. Let E^+ be the union of all $E(H(\alpha))$ for $\alpha \in A$. Note that E^+ contains all edges of \overline{W} and both diagonals of the cycle C_k for every $k \in K$. Let H^+ be the graph with vertex set $V(H)$ and edge set E^+ .




Figure 5: Restrictions of the paths P and P' in the last part of the proof of Lemma 3.2 to the wedge X .

We can easily see that there is G -embedding of H^+ as follows. For every edge $uv \in E^+$, the images of the infinitely many paths connecting u to v in $H(\alpha)^\parallel$ under the embedding ι_α are internally disjoint from $\mathcal{I}(V(H))$. Thus we can pick an enumeration of the edges in E^+ and for each edge $uv \in E^+$ pick a $\mathcal{I}(u)$ - $\mathcal{I}(v)$ -path P_{uv} which is disjoint from all paths chosen for the preceding edges.

For every $k \in K$, the graph H^+ contains a pair of crossing bypasses for Z . One is obtained by concatenating the path consisting of z_k^k together with all vertices of X_k , the edge from x_k^k to y_{k+1}^0 , and the path consisting of all vertices of Y_{k+1} and z_{k+1}^0 . The other is obtained by concatenating the path consisting of z_{k+1}^{k+1} and all vertices of X_{k+1} , the edge from x_{k+1}^{k+1} to y_k^0 , and the path consisting of all vertices of Y_k and z_k^0 . Note that among these bypasses there is an infinite disjoint family with infinitely many crossing pairs. Let W^+ be the wedge Z together with such a family of bypasses, and let ι be the restriction of the G -embedding of H^+ to W^+ ; by definition we have $\iota(w_m^i) = \mathcal{I}(z_m^i)$ for $0 \leq i \leq m$. This finishes the proof in case \mathcal{I} agrees on $V(H)$.

If \mathcal{I} contains an uncountable subfamily which agrees on $V(H)$, then the same argument as above can be applied to this subfamily. Hence from now on assume that there is no such subfamily. In the remainder of the proof, we ignore the additional edges in $H(\alpha)$ and view \mathcal{I} as a family of G -embeddings of H^\parallel .

For each $i \in \mathbb{N}_0$ we recursively define

- an integer n_i ,
- a set M_i of edges, and
- a G -embedding ι_i of the graph H_i^+ obtained from H_{n_i} by adding all edges in M_i

satisfying the following properties:

- (i) $n_0 = m$ and $n_i > n_{i-1}$ for $i > 0$,
- (ii) $M_0 = \emptyset$, and $M_i \setminus M_{i-1} = \{m_i\}$ for $i > 0$; the endpoints of m_i are not adjacent in H , contained in $V(H_{n_i}) \setminus V(H_{n_{i-1}})$, and at least one of them is enclosed,
- (iii) ι_0 agrees with \mathcal{I} on $V(H_m)$, and the restriction of ι_i to H_{i-1}^+ is ι_{i-1} for $i > 0$.

Before carrying out the recursive construction, let us show how the resulting sequences can be used to finish the proof of the lemma. Let $H^+ = \lim_{i \rightarrow \infty} H_i^+$. Since n_i and M_i are strictly increasing, H^+ is the graph obtained from H by adding the (infinite) set

$M = \bigcup_{i \in \mathbb{N}} M_i$. Note that there is an infinite subset $M' \subseteq M$ such that no edge in M' has an endpoint in one of the three wedges X , Y , or Z ; by symmetry we may assume that no edge in M' has an endpoint in Z . By Lemma 3.2 the graph H^+ contains an infinite set of disjoint bypasses for $Z_{m,\infty}$ containing infinitely many crossing pairs. Let W^+ be the union of $Z_{m,\infty}$ and these bypasses. Lemma 2.3 tells us that there is a G -embedding ι of H^+ which agrees with each ι_i on H_i^+ , and thus agrees with \mathcal{I} on $V(H_m)$. The restriction of ι to W^+ is the desired G -embedding of W^+ .

It remains to construct the sequences n_i , M_i , and ι_i . Alongside these sequences, for every i we also construct

- an uncountable family $\mathcal{J}_i \subseteq \mathcal{I}$

such that

- (iv) $\mathcal{J}_i \cup \{\iota_i\}$ agrees on $V(H_{n_i})$, and
- (v) $\iota_i(H_i^+) \cap \kappa(V(H)) \subseteq \kappa(V(H_{n_i}))$ for every $\kappa \in \mathcal{J}_i$.

Note that for any uncountable family $\mathcal{J} \subseteq \mathcal{I}$ and any finite set S of vertices and edges of H^\parallel , there is an uncountable subfamily of \mathcal{J} which agrees on S . This is due to the fact that G is countable and thus there are only countably many possible images of S . This fact will be used at several points in the construction.

Let $n_0 = m$, and let $M_0 = \emptyset$. For any pair a, b of adjacent vertices in H_m pick a path P_{ab} of length 2 connecting them in H_m^\parallel . Let $\mathcal{J}_0 \subseteq \mathcal{I}$ be an uncountable family which agrees on every P_{ab} . Define a G -embedding ι_0 of $H_0^+ = H_m$ by $\iota_0(v) = \mathcal{I}(v)$ for every vertex of H_0^+ , and $\iota_0(ab) = \mathcal{J}_0(P_{ab})$ for every edge of H_0^+ . Properties (i)–(v) are easily seen to hold for n_0 , M_0 , ι_0 , and \mathcal{J}_0 .

The recursive step from i to $i + 1$ rests on the following claim.

Claim. For some $n > n_i$ there are

- (1) an uncountable subfamily $\mathcal{J} \subseteq \mathcal{J}_i$ which agrees on $V(H_n)$,
- (2) $u, v \in V(H_n) \setminus V(H_{n_i})$ which are not adjacent in H such that u is enclosed, and
- (3) a $\mathcal{J}(u)$ – $\mathcal{J}(v)$ -path P such that $P \cap \iota_i(H_i^+)$ is empty, and $P \cap \kappa(V(H)) = \mathcal{J}(\{u, v\})$ for every $\kappa \in \mathcal{J}$.

Let us assume first that the claim is true. We let $n_{i+1} = n$ and $E_{i+1} = E_i \cup \{uv\}$ for the n , u , and v provided by the claim. Let N be the number of vertices contained in $\iota_i(H_i^+) \cup P$. For every pair a, b of adjacent vertices in H_n pick $N + 1$ different a – b -paths of length 2 in H^\parallel , and let $\mathcal{J}_{i+1} \subseteq \mathcal{J}$ be an uncountable set which agrees on all of these paths. Note that these paths are pairwise internally disjoint, and (by the pigeonhole principle) for each pair a, b , the image of at least one such path P_{ab} is internally disjoint from $\iota_i(H_i^+) \cup P$. Let

$$\iota_{i+1}(x) = \begin{cases} \iota_i(x) & \text{if } x \in V(H_i^+) \cup E(H_i^+), \\ \mathcal{J}_{i+1}(x) & \text{if } x \in V(H_{n_{i+1}}) \setminus V(H_{n_i}), \\ \mathcal{J}_{i+1}(P_{ab}) & \text{if } x = ab \in E(H_{n_{i+1}}) \setminus E(H_{n_i}), \\ P & \text{if } x = uv. \end{cases}$$

We briefly check that ι_{i+1} is a G -embedding of H_{i+1}^+ . First note that the image of any edge is a path connecting the images of the respective endpoints. For edges in $E(H_i^+)$ this follows from the induction hypothesis, for edges in $E(H_{n_{i+1}}) \setminus E(H_{n_i})$ it follows from the fact that every element of \mathcal{J}_{i+1} is a G -embedding of H^\parallel which agrees with ι_i on $V(H_i^+)$, and for the edge uv it follows from (3) in the above claim. The path $\iota_{i+1}(uv)$ is internally disjoint from $\iota_i(H_i^+) = \iota_{i+1}(H_i^+)$ by (3). The paths $\mathcal{J}_{i+1}(P_{ab})$ are internally disjoint from P and $\iota_i(H_i^+)$ by definition, and they are internally disjoint from one another since every element of \mathcal{J}_{i+1} is a G -embedding.

Properties (i), (ii), and (iii) follow from the above claim and the resulting definitions of n_{i+1} , M_{i+1} , and ι_{i+1} . For property (iv) note that $\mathcal{J}_{i+1} \cup \{\iota_{i+1}\}$ agrees on $V(H_{n_{i+1}}) \setminus V(H_{n_i})$ by definition of ι_{i+1} and on $V(H_{n_i})$ by the induction hypothesis since $\mathcal{J}_{i+1} \subseteq \mathcal{J}_i$. For (v), note that $\kappa(V(H))$ contains no internal vertex of P by (3) above, and no internal vertex of any image $\mathcal{J}_{i+1}(P_{ab})$ because κ is contained in \mathcal{J}_{i+1} . By the induction hypothesis, $\kappa(V(H))$ also does not contain an internal vertex of $\iota_{i+1}(e) = \iota_i(e)$ for any $e \in E(H_i^+)$. Hence $\iota_{i+1}(H_{n_i}) \cap \kappa(V(H)) \subseteq \iota_{i+1}(V(H_{n_i}))$, and (v) follows from (iv).

It only remains to provide a proof for the above claim. We first show that it suffices to prove that the conclusion of the claim holds when we replace (3) by the weaker condition that there is

(3') a $\mathcal{J}(u)$ - $\mathcal{J}(v)$ -path P such that $P \cap \iota_i(H_i^+)$ is empty, and $P \cap \kappa(V(H_{n+1})) = \mathcal{J}(\{u, v\})$ for every $\kappa \in \mathcal{J}$.

Assume that we have found a family \mathcal{J} , vertices u, v and a $\mathcal{J}(u)$ - $\mathcal{J}(v)$ path P satisfying (1), (2), and (3') for some n . Let P_j be the subpath of P of length j starting at $\mathcal{J}(u)$. Let $\mathcal{K}_j \subseteq \mathcal{J}$ consist of all $\kappa \in \mathcal{J}$ for which $\kappa(V(H))$ does not contain any internal vertex of P_j . Let k be maximal such that \mathcal{K}_k is uncountable.

If $P_k = P$, then we can simply replace \mathcal{J} by \mathcal{K}_k to satisfy the stronger condition (3). Otherwise, let $t \neq \mathcal{J}(u)$ be the other endpoint of P_k . Since $V(H)$ is countable, there is a vertex $v' \in V(H)$ and an uncountably infinite family $\mathcal{K}' \subseteq \mathcal{K}_k$ such that $\kappa(v') = t$ for every $\kappa \in \mathcal{K}'$. Let n' be such that $v' \in V(H_{n'})$, and let $\mathcal{K}'' \subseteq \mathcal{J}'$ be an uncountably infinite subfamily which agrees on $V(H_{n'})$. The image of $V(H)$ under $\kappa \in \mathcal{K}''$ does not contain any internal vertex of P_k since $\mathcal{K}'' \subseteq \mathcal{K}_k$. Further note that $v' \notin V(H_{n+1})$ by condition (3'), and thus u and v' are not adjacent. Hence \mathcal{K}'' , the pair u, v' , and the path P_k satisfy (1), (2), and (3).

Finally, we need to show how to construct a family \mathcal{J} , vertices u, v and a $\mathcal{J}(u)$ - $\mathcal{J}(v)$ path P satisfying (1), (2), and (3') for some n .

Let $U_1 = V(H_{n_1})$ and recursively define U_k as the union of U_{k-1} and all neighbours (in H) of some enclosed vertex a_k which has at most 2 neighbours outside U_{k-1} . It is not hard to see that we can pick the vertices a_k in this construction such that every vertex of H is contained in some U_k . Recall that we may assume that there is no uncountable subset of \mathcal{I} which agrees on $V(H)$. In particular, since \mathcal{J}_i agrees on $V(H_{n_i})$ but not on $V(H)$ there is some k such that \mathcal{J}_i agrees on U_k but not on U_{k+1} . Let $a = a_k$, let n be large enough that $U_{k+1} \subseteq V(H_n)$, let $\mathcal{K} \subseteq \mathcal{J}_i$ be an uncountable family which agrees on $V(H_n)$, and let $\kappa \in \mathcal{J}_i$ be such that $\mathcal{K} \cup \{\kappa\}$ does not agree on U_{k+1} (and thus does not agree on the neighbours of a).

Suppose first that there are a neighbour x of a and a vertex $y \in V(H_n)$ such that $\kappa(x) = \mathcal{K}(y)$. If y is not adjacent to a , then we can choose an uncountable subfamily $\mathcal{J} \subseteq \mathcal{K}$ which agrees on $V(H_{n+1})$, $u = a$, and $v = y$. Recall that κ is an embedding of H^\parallel , hence G contains infinitely many internally disjoint $\kappa(s)$ – $\kappa(t)$ -paths for any pair of adjacent vertices in H . In particular, there are infinitely many internally disjoint paths connecting $\mathcal{J}(a) = \kappa(a)$ to $\kappa(x) = \mathcal{J}(y)$. Among these paths we find one with the desired properties because $\iota_i(H_i^+)$ and $\mathcal{J}(V(H_{n+1}))$ are finite, and $\iota_i(H_i^+)$ does not contain $\mathcal{J}(a)$ or $\mathcal{J}(y)$ by property (v).

So we may assume that y is adjacent to a ; in this case all neighbours of a except x and y are contained in U_k . Since a is enclosed, there is another common neighbour of a and x besides y , denote this neighbour by z . Note that y and z are not adjacent. At least one of y and z is enclosed because a is enclosed, and the non-enclosed neighbours of any enclosed vertex are adjacent. As above, we can choose an uncountable subfamily $\mathcal{J} \subseteq \mathcal{K}$ which agrees on $V(H_{n+1})$, $u = y$ and $v = z$ (or vice versa), and a $\mathcal{J}(z)$ – $\mathcal{J}(y)$ -path with the desired properties among the infinitely many internally disjoint $\mathcal{J}(z)$ – $\mathcal{J}(y)$ -paths in G .

The above argument only required n to be large enough for $U_{k+1} \subseteq V(H_n)$; in particular, it also works if we replace n by $n + 1$. Hence from now on let us assume that \mathcal{K} agrees on $V(H_{n+1})$ and that no neighbour x of a satisfies $\kappa(x) = \mathcal{K}(y)$ for any $y \in H_{n+1}$. The neighbourhood of a in H either induces a 6-cycle and $\mathcal{K} \cup \{\kappa\}$ agrees on at least 4 vertices of this 6-cycle because they are contained in U_k , or (in case $a \in \overline{W}_m$) it induces a path of length 4 and $\mathcal{K} \cup \{\kappa\}$ agrees on the endpoints of this path because they are contained in U_k . Since $\mathcal{K} \cup \{\kappa\}$ does not agree on all neighbours of a , the neighbourhood of a must contain a path P whose endpoints u and v are not adjacent such that $\mathcal{K} \cup \{\kappa\}$ agrees on u and v but not on the interior points of P . The same argument as above tells us that we can connect the images (under κ) of any two consecutive vertices of P by a path which is disjoint from $\iota_i(H_i^+)$, and does not intersect $\mathcal{K}(V(H_{n+1}))$ except possibly in $\mathcal{K}(u) = \kappa(u)$ or $\mathcal{K}(v) = \kappa(v)$. The union of these paths contains a $\mathcal{K}(u)$ – $\mathcal{K}(v)$ -path Q . The family \mathcal{K} , the pair u, v , and the path Q satisfy (1), (2), and (3'). \square

3.3 Braiding infinite paths

In this section we complete the second step of the above proof sketch. More precisely, we show how to find sets of infinite paths with edges between any two of them in a graph W^+ consisting of a triangular wedge together with a set of infinitely many disjoint bypasses with infinitely many crossing pairs. Similarly to the previous section, we do not merely show that such paths exist, but construct them in a fairly structured way to give us the control we need for the proof of the second statement in Theorem 1.2.

The following observation should be clear, see Figure 6 for a sketch.

Observation 3.3. *Let $m < n$ and let P and P' be disjoint, crossing bypasses for $W_{m,n}$. The following statements hold for any $k \leq m$.*

- (1) *There are disjoint paths in $W_{m,n}$ connecting w_m^i to w_n^i for $0 \leq i \leq k$.*




Figure 6: Routing paths in W with crossing bypasses. Each of the above sketches corresponds to one of the three statements in Observation 3.3 with $m = 4$, $n = 9$ and $k = 4$.

- (2) *There are disjoint paths in $W_{m,n} \cup P$ connecting w_m^i to $w_n^{(i+1 \bmod k)}$ for $0 \leq i \leq k$.*
- (3) *There are disjoint paths in $W_{m,n} \cup P \cup P'$ connecting w_m^0 to w_n^k , w_m^k to w_n^0 , and w_m^i to w_n^i for $0 < i < k$.*

The paths in all three cases can be chosen such that they intersect W_m only in their initial vertices and W_n only in their terminal vertices.

The next lemma is an immediate consequence of the above observations.

Lemma 3.4. *Let $(P_i)_{i \in \mathbb{N}}$ be an infinite family of disjoint bypasses for $W = W_{0,\infty}$ containing infinitely many crossing pairs of bypasses. Denote by $W_{m,n}^+$ the union of $W_{m,n}$ with all P_i that are bypasses for $W_{m,n}$.*

- (1) *For any $k < m$, and any permutation π of $\{0, \dots, k\}$ there are some $n > m$ and a set of disjoint paths in $W_{m,n}^+$ connecting w_m^i to $w_n^{\pi(i)}$ for $0 \leq i \leq k$.*
- (2) *There is an infinite family \mathcal{P} of (infinite) pairwise disjoint paths in $W^+ = W_{0,\infty}^+$ such that every pair of paths in \mathcal{P} is connected by an edge.*

Proof. Note that by reordering the sequence of bypasses we may without loss of generality assume that P_i and P_{i+1} are crossing for infinitely many i .

For the proof of (1), note that the cyclic permutation $x \mapsto x + 1 \bmod k$ and the transposition of 1 and k generate the symmetric group on k elements. Further note that disjointness of the P_i implies that any vertex of W appears as an endpoint of at most one P_i . Hence for any l , all but finitely many P_i lie in $W_{l,\infty}^+$, so there must be some l' such that $W_{l',l'}^+$ contains a pair P_i, P_{i+1} of bypasses satisfying the conditions of Observation 3.3. Thus we can iterate Observation 3.3 and concatenate the corresponding paths to obtain the desired family of paths.

The following recursive construction proves (2). Start with the single path consisting of a vertex v_1^0 . Assume that we have for some m and k constructed disjoint paths in $W_{0,m}^+$ which end at w_m^i for $0 \leq i \leq k$. Applying (1) with different permutations and

concatenating the resulting paths, we obtain $n > m$ and a family of disjoint paths in $W_{0,n}^+$ ending at w_n^i for $0 \leq i \leq k$ such that every pair of them is connected by an edge. Add the path consisting of the single vertex w_n^{k+1} to this family and iterate. In the limit, we get infinitely many disjoint paths each pair of which is connected by an edge. \square

3.4 Finding large complete (topological) minors

In this subsection, we combine the results from the previous two subsections to prove Theorem 1.2. The following two results imply the first and second statement of Theorem 1.2, respectively.

Theorem 3.5. *Let G be a countable graph, and assume that there are G -embeddings of $\overline{W}(\alpha)^\parallel$ for every $\alpha \in \{0, 1\}^\mathbb{N}$. Then G contains a countably infinite complete graph as a minor.*

Proof of Theorem 1.2. For each $\alpha \in \{0, 1\}^\mathbb{N}$ let ι_α be a G -embedding of $\overline{W}(\alpha)^\parallel$. Uncountably many of these embeddings agree on $V(\overline{W}_m)$ for any m , so by Lemma 3.1 there is a G -embedding of a graph W^+ consisting of $W_{m,\infty}$ and an infinite set of disjoint bypasses containing infinitely many crossing pairs. The graph W^+ contains an infinite complete minor whose branch sets are the infinite paths obtained by Lemma 3.4 (2), and thus so does G . \square

For each $k \in \mathbb{N}$ and $\alpha \in \{0, 1\}^\mathbb{N}$, define a graph $H_k(\alpha)$ as follows. Take $\overline{W}_{k^2-1,\infty}(\alpha)$, and add $3k$ vertices $x_*^0, \dots, x_*^{k-1}, y_*^0, \dots, y_*^{k-1}, z_*^0, \dots, z_*^{k-1}$. Add edges between x_*^i and x_m^{ik+j} , between y_*^i and y_m^{ik+j} , and between z_*^i and z_m^{ik+j} for $0 \leq j < k$. It is not hard to see that $H_k(\alpha)$ is planar, and thus so is $H_k(\alpha)^\parallel$.

Theorem 3.6. *Let G be a countable graph, and assume that there are G -embeddings of $H_k(\alpha)^\parallel$ for every $\alpha \in \{0, 1\}^\mathbb{N}$. Then G contains K_k as a topological minor.*

Proof. Let $m = k^2 - 1$. As before, let ι_α be a G -embedding of $H_k(\alpha)^\parallel$ for every α , and note that an uncountable family \mathcal{I} of these embeddings agrees on $V(\overline{W}_m)$ and every x_*^i , y_*^i , and z_*^i .

Let G' be the graph obtained from G by removing all images $\mathcal{I}(x_*^i)$, $\mathcal{I}(y_*^i)$, and $\mathcal{I}(z_*^i)$ for $0 \leq i < k$ and note that each embedding $\iota_\alpha \in \mathcal{I}$ gives rise to a G' -embedding of $\overline{W}_{m,\infty}(\alpha)$.

By Lemma 3.1 there is a G' -embedding ι of a graph W^+ consisting of $W_{m,\infty}$ and an infinite set of disjoint bypasses containing infinitely many crossing pairs such that (without loss of generality) $\iota(w_m^i) = \mathcal{I}(x_m^i)$ for $0 \leq i \leq m$.

Let π be a permutation of $\{0, \dots, k^2\}$ such that $\pi(ik + j) = \pi(jk + i) + 1$ holds for every pair i, j satisfying $0 \leq i < j < k$. Lemma 3.4 (1) implies that there is some n such that W^+ contains disjoint paths connecting w_m^ℓ to $w_n^{\pi(\ell)}$ for $0 \leq \ell \leq k^2$. Combining these paths with the edges from $w_n^{\pi(ik+j)}$ to $w_n^{\pi(jk+i)}$ in W gives disjoint paths in W^+ connecting w_m^{ik+j} to w_m^{jk+i} for every pair i, j satisfying $0 \leq i < j < k$. The images of these paths under ι are disjoint $\mathcal{I}(x_m^{ki+j})$ - $\mathcal{I}(x_m^{kj+i})$ -paths P_{ij} in G' .

There are infinitely many internally disjoint paths in G connecting $\mathcal{I}(x_*^i)$ to $\mathcal{I}(x_m^{ki+j})$ since any element of \mathcal{I} is a G -embedding of $H_k(\alpha)^\parallel$. Among these paths we can inductively find internally disjoint $\mathcal{I}(x_*^i)$ - $\mathcal{I}(x_m^{ki+j})$ -paths Q_{ij} for $0 \leq i, j < k$ which are also internally disjoint from all paths P_{ij} .

We can define a G -embedding κ of the complete graph on k vertices v_0, \dots, v_{k-1} by letting $\kappa(v_i) = \mathcal{I}(x_*^i)$, and $\kappa(v_i v_j)$ the union of the paths Q_{ij} , P_{ij} , and Q_{ji} . \square

4 A universal, locally finite, planar graph

In this section, we construct a locally finite graph which is universal for the topological minor relation, thus proving Theorem 1.3. Before doing so, we introduce some notation.

Assume that F is a face of a finite, connected graph. Call a vertex or edge x *incident* to F if $\iota(x)$ lies in the closure of F . By tracing the boundary of F in clockwise direction if F is an interior face, or in anti-clockwise direction if F is the outer face, we obtain a cyclic sequence of vertices incident to this face which we call a *facial sequence*. The reason we treat the outer face differently is that we want to make sure that facial sequences are invariant under making a different face the outer face by applying an appropriate inversion. A facial sequence may contain the same vertex more than once (this happens only for cut vertices). If a facial sequence contains each vertex at most once, then this sequence defines a cycle in the graph which we call a *facial cycle*. Note that since we prescribed a direction on the boundary of F , all facial cycles are in fact directed cycles.

Remark 4.1. We can combine two connected planar graphs into a larger one by identifying facial cycles. More precisely, let G and H be two graphs, and let C and C' be facial cycles of the same length in G and H , and let $\phi: V(C) \rightarrow V(C')$ be an order reversing bijection. It is not hard to see that the graph obtained by identifying each vertex v of C with $\phi(v)$ is again planar, and that this graph has a plane embedding in which all facial sequences of G and H except C and C' are again facial sequences. Note that if we choose ϕ to be order preserving rather than order reversing, then we also obtain a planar graph, but the facial sequences of one of the two graphs are reversed in the combined graph.

Let G be a graph, let $W \subseteq V(G)$ and let $\phi: W \rightarrow V(G)$ be such that $\phi(W) \cap W = \emptyset$. A ϕ -*linkage* in G is a set of disjoint paths in G containing a w - $\phi(w)$ -path P_w for every $w \in W$. Recall that we consider cycles as cyclically ordered sequences of vertices, which in particular implies that every cycle has a direction. Let us call two cycles C_1 and C_2 *well-linked* if for any order reversing injection from $W \subseteq V(C_1)$ to $V(C_2)$ there is a ϕ -linkage whose paths meet C_1 and C_2 only in their respective endpoints. Note that we take order reversing functions ϕ so that two facial cycles in a planar graph can be well linked. An m - n -*mesh* is a triple (G, C, C') where G is a planar graph, C and C' are cycles of lengths m and n which are well linked in G and facial with respect to some plane embedding of G . By a slight abuse of notation, we also call the graph G an m - n -mesh if there exist two such facial cycles with respect to some embedding. It is easy to see that n - m -meshes exist for all m and n , for instance we may start with a




Figure 7: Construction of the graph $\mathbf{M}(5)$. The bold cycles are (from the centre to the outer cycle) C_1 , C_1' , C_2' , and C_2 .

The cycle C_2 bounding the outer face in this drawing is the boundary, the cycles bounding the faces between M_1 and M_2 are the attachment cycles.

Cartesian product $C_N \square P_N$ where N is much larger than n and m and connect cycles of length n and m to the two facial cycles of length N in an appropriate way.

For every $n \in \mathbb{N}$, let (M_1, C_1, C_1') and (M_2, C_2, C_2') be two $(2n)$ - (n^2) -meshes. Denote the facial sequences of the well linked cycles by $C_i = (v_{1,i}, \dots, v_{2n,i})$ and $C_i' = (v'_{1,i}, \dots, v'_{n^2,i})$. Let $\mathbf{M}(n)$ be the graph obtained by adding an additional vertex z , connecting z to every vertex of C_1 , and adding edges between $v'_{kn,1}$ and $v'_{kn,2}$ for $1 \leq k \leq n$, see Figure 7. We call M_1 the *inner mesh*, M_2 the *outer mesh*, and z the *centre* of $\mathbf{M}(n)$. The cycles C_1' and C_2' are called the *inner perimeter* and *outer perimeter*, respectively, and C_2 is called the *boundary* of $\mathbf{M}(n)$. The edges connecting the outer perimeter to the inner perimeter are called the *spokes* of $\mathbf{M}(n)$. Moreover, the n cycles of length $2n + 2$ consisting of a path of length n on C_1' , a path of length n on C_2' , and the two spokes connecting the endpoints of these paths are called the *attachment cycles* of $\mathbf{M}(n)$. Note that $\mathbf{M}(n)$ has a plane embedding such that the boundary of $\mathbf{M}(n)$ and all attachment cycles are facial cycles.

Using these graphs, we construct a graph \mathbf{G} recursively as follows. In each iteration, we have a graph $\mathbf{G}(n)$ and a set \mathcal{C}_n of pairwise disjoint facial cycles with respect to some embedding of $\mathbf{G}(n)$ such that each $C \in \mathcal{C}_n$ has length $2(n + 1)$.

We start the inductive construction by letting G_1 be a cycle of length 4, and choosing \mathcal{C}_1 as the set consisting of this cycle. In each subsequent step, for each cycle $C \in \mathcal{C}_n$ we take a copy of $\mathbf{M}(n + 1)$ and identify its boundary with C as indicated in Remark 4.1. Note that apart from the boundaries, all facial cycles in all copies of $\mathbf{M}(n + 1)$ are facial cycles of $\mathbf{G}(n + 1)$; in particular, all attachment cycles are facial cycles of length $2(n + 2)$. Let the set \mathcal{C}_{n+1} consist of all attachment cycles of all copies of $\mathbf{M}(n + 1)$. Define $\mathbf{G} = \lim_{n \rightarrow \infty} \mathbf{G}(n)$. This graph is planar by Theorem 2.1 since every finite subgraph of \mathbf{G} is a subgraph of some $\mathbf{G}(n)$ and planarity is preserved under taking subgraphs.

Note that the above construction is not unique (even when $\mathbf{M}(n)$ is fixed for every n) since there are different, non-isomorphic ways of identifying the facial cycles. However, this non-uniqueness will not be an issue; we will show that any choice leads to a universal planar graph with respect to the topological minor relation.

Theorem 4.2. *Any connected, locally finite, planar graph has a \mathbf{G} -embedding.*

Proof. Let G be a connected, locally finite, planar graph, and let H be the graph obtained from G by subdividing every edge. Clearly, any \mathbf{G} -embedding of H gives rise to a \mathbf{G} -embedding of G ; the path corresponding to an edge $e \in E(G)$ is simply the union of the two paths corresponding to the edges obtained by subdividing e . In particular, it suffices to show that H has a \mathbf{G} -embedding.

We partition the vertices of H into *original vertices*, that is, vertices corresponding to vertices of G , and *subdivision vertices*, that is, vertices added to subdivide an edge. Any subdivision vertex has precisely two neighbours both of which are original, and any original vertex only has subdivision neighbours.

Let $(v_n)_{n \in \mathbb{N}}$ be an enumeration of the original vertices such that the subgraph of G induced by the vertices corresponding to v_1, \dots, v_n is connected for every $n \in \mathbb{N}$. Let $V_n = \{v_k \mid k \leq n\}$. Let H_n be the subgraph of H induced by V_n and all neighbours of V_n . Let H'_n be the subgraph of H induced by V_n and all subdivision vertices both of whose neighbours are in V_n . Subdivision vertices with exactly one neighbour in V_n are called *loose ends* of H_n , they have degree 1 in H_n and are the only vertices of H_n which are not contained in H'_n . Note that all graphs H_n and H'_n are connected by our choice of the enumeration (v_n) .

Since G is planar, so is H ; for the remainder of the proof we fix an arbitrary plane embedding ι of H . Restricting this embedding to H_n or H'_n clearly gives plane embeddings for all $n \in \mathbb{N}$, by a slight abuse of notation we denote these embeddings by ι as well. When referring to faces of H_n or faces of H'_n we tacitly assume that these are faces with respect to the embedding ι . We say that a loose end v of H_n belongs to a face F of H_n if $\iota(v)$ is contained in the closure of F . Note that any loose end belongs to exactly one face since it has degree 1 and thus it has a neighbourhood whose intersection with $\mathbb{R}^2 \setminus \iota(H_n)$ is connected. Further note that any loose end appears precisely once in the boundary sequence of the face it belongs to. Denote by $L(F)$ the set of loose ends belonging to the face F .

We now inductively construct $\mathbf{G}(m)$ -embeddings of H_n for appropriate choices of m and n . Call a $\mathbf{G}(m)$ -embedding ϕ of H_n *good* if there is an injective map assigning to each face F of H_n a cycle $C_F \in \mathcal{C}_m$ such that the restriction of ϕ to $L(F)$ is an order preserving injection from $L(F)$ to C_F (with respect to the cyclic orders given by the boundary sequence of F and C_F , respectively).

Our inductive construction rests on the following two claims whose proofs are fairly straightforward. Let ϕ be a good $\mathbf{G}(m)$ -embedding of H_n and assume that H_n has at most m loose ends.

Claim 1. There is a good $\mathbf{G}(m+1)$ -embedding ψ of H_n such that ϕ and ψ agree on H'_n




Figure 8: Sketch of the constructions used to prove Claims 1 and 2.

Claim 2. If H_{n+1} has at most m loose ends, then there is a good $\mathbf{G}(m+1)$ -embedding ψ of H_{n+1} such that ϕ and ψ agree on H'_n .

Before proving the two claims, we show that they can be applied inductively to obtain a sequence of embeddings of H'_n into \mathbf{G} satisfying the conditions of Lemma 2.3, thereby finishing the proof of Theorem 4.2.

For the base case, note that H_1 is a star with centre v_1 where all leaves are loose ends. For any $m \geq \deg(v_1)$, by choosing an appropriate linkage in the inner mesh we can construct a $\mathbf{M}(m)$ -embedding of H_1 mapping v_1 to the centre of $\mathbf{M}(m)$ and all loose ends to vertices on the same attachment cycle of $\mathbf{M}(m)$. Viewing $\mathbf{M}(m)$ as a subgraph of $\mathbf{G}(m)$ this defines a good $\mathbf{G}(m)$ -embedding of H_1 .

For the inductive step assume that we have a good $\mathbf{M}(m)$ -embedding of H_n , and that the number of loose ends of H_n is at most m . We can recursively apply Claim 1 to get a good $\mathbf{M}(m')$ -embedding of H_n where $m' \geq m$ is at least as big as the number of loose ends of H_{n+1} , and then apply Claim 2 to obtain a good $\mathbf{M}(m')$ -embedding of H_{n+1} .

It remains to prove the two claims. In both claims, for any vertex or edge x of H'_n , we must have $\psi(x) = \phi(x)$. Thus the maps ϕ and ψ in the above claims only differ in the embeddings of the loose ends, their incident edges, and potentially the additional vertex v_{n+1} and its incident edges. We refer to Figure 8 for a sketch of how the embeddings of loose ends and their incident edges are extended into the copies of $\mathbf{M}(m+1)$ that were added in the construction of $\mathbf{G}(m+1)$ from $\mathbf{G}(m)$.

For a formal proof, consider the following setup. Let F be an arbitrary face of H_n . As before, denote by $L(F)$ the loose ends belonging to F , and for $l \in L(F)$ let e_l be the unique edge incident to l . Let M_F be the copy of $\mathbf{M}(m+1)$ whose boundary was identified with C_F in the construction of $\mathbf{G}(m+1)$. Consider $\phi(L(F))$ as vertices on the boundary of M_F , ignoring the embedding of the rest of H_n . Recall that the boundary was identified with C_F using an order reversing bijection, so the restriction of ϕ to $L(F)$ is an order reversing injection from $L(F)$ to the boundary of M_F .

For the proof of Claim 1 we apply the following construction for each face F of H_n , see the left half of Figure 8. Pick an arbitrary attachment cycle C'_F of M_F . Let Y be the set of vertices on C'_F which are contained in the outer mesh of M_F ; note that $|Y| = m + 2 \geq |L(F)|$. Let $\xi: \phi(L(F)) \rightarrow Y$ be an order reversing injection. Since M_F is a mesh, we can find a ξ -linkage in the outer mesh of M_F whose paths intersect the boundary of M_F only in their endpoints. Note that since the image of ξ is contained in $Y \subseteq C'_F$, the paths in this linkage connect $\phi(L(F))$ to C'_F .

Let $\psi(e_l)$ be the concatenation of $\phi(e_l)$ and the $\phi(l)$ - $\xi(\phi(l))$ -path in this linkage, let $\psi(l) = \xi(\phi(l))$, and let $\psi(x) = \phi(x)$ for every vertex or edge x of H'_n . It is easy to check that ψ is a $\mathbf{G}(m + 1)$ -embedding of H_n . By construction, the images of all loose ends belonging to F lie on C'_F . Moreover, the cyclic orders of the $\phi(v_i)$ on C_F and the $\psi(x_i)$ on C'_F coincide since the composition of two order reversing maps is an order preserving map. For any two faces F_1 and F_2 we have $C_{F_1} \neq C_{F_2}$ and thus C'_{F_1} and C'_{F_2} are attachment cycles of different copies of $\mathbf{M}(m + 1)$, so the function mapping F to C'_F is an injection.

Thus ψ is a good $\mathbf{G}(m + 1)$ -embedding of H_n which coincides with ϕ on H'_n . This proves Claim 1.

For the proof of Claim 2, we first note that any face F of H_{n+1} which is not incident to v_{n+1} is also a face of H_n . We can thus apply the same construction as above to F to obtain an attachment cycle of M_F in $\mathbf{G}(m + 1)$ and an appropriate embedding of the loose ends belonging to F and their incident edges.

It remains to provide a construction for faces F incident to v_{n+1} . A sketch of this construction is shown in the right half of Figure 8.

Formally, let F_0 be the face of H_n which contains the embedding of v_{n+1} , and let l_1, \dots, l_k be an enumeration of the loose ends $L(F_0)$, cyclically ordered according to the boundary sequence of F_0 . Let i_1, \dots, i_r be the indices such that l_{i_j} is incident to v_{n+1} , for convenience set $i_{r+1} = i_1$. Let $B(F_0)$ be the boundary sequence of F_0 and let $B_j(F_0)$ be the part of $B(F_0)$ strictly between l_{i_j} and $l_{i_{j+1}}$; if v_{n+1} is incident to a unique $l \in L(F_0)$, then $B_1(F_0)$ is the whole boundary sequence without l , cyclically permuted so it starts with the successor of l . Clearly, $B_j(F_0)$ contains at most $|L(F_0)| - 1$ vertices, all of which are loose ends belonging to the same face of H_{n+1} . For $j \neq j'$, the loose ends in $B_j(F_0)$ and $B_{j'}(F_0)$ belong to different faces of H_{n+1} .

Let Y be the set of vertices in the outer perimeter of M_{F_0} in clockwise cyclic order (that is, we consider the outer perimeter as a face of the outer mesh). Pick an order reversing injection $\xi: \phi(L(F_0)) \rightarrow Y$ such that every $\phi(l_{i_j})$ is incident to a spoke, and the image of each $B_j(F_0)$ is completely contained in an attachment cycle C_j . This is possible, because M_{F_0} has at least $m \geq |L(F_0)|$ spokes, and between any two spokes we can find $m - 1 \geq |L(F_0)| - 1$ vertices belonging to the same attachment cycle.

Next, let N be the set of neighbours of v_{n+1} in H_{n+1} . Going around $\iota(v_{n+1})$ in clockwise direction defines a cyclic order on N . The restriction of this cyclic order to the vertices l_{i_j} agrees with the restriction of the boundary sequence of F_0 to these vertices, otherwise the embedding would not be planar. Let Z be the set of vertices on the inner perimeter of M_{F_0} in clockwise cyclic order. Note that every neighbour of v_0 in H_{n+1} is either a loose end of H_n , or a loose end of H_{n+1} . Thus N consists of at most $2m$ vertices,

at most m of which are loose ends of H_{n+1} . Hence we can find an order preserving map $\eta: N \rightarrow Z$ such that $\eta(l_{i_j})$ is joined to $\xi(\phi(l_{i_j}))$ by a spoke, and the vertices between l_{i_j} and $l_{i_{j+1}}$ are all mapped to the attachment cycle C_j defined above.

Now define the embedding ψ as follows. Let $\psi(v_{n+1})$ be the centre z of M_{F_0} . For $x \in N$ let $\psi(x) = \eta(x)$. Disjoint z - $\eta(x)$ -paths for the images $\psi(v_{n+1}x)$ can be constructed from a linkage between $\eta(N)$ and the neighbours of z . Such a linkage exists in the inner mesh of M_{F_0} because v_{n+1} has $N < |L(F_0)| + |\{\text{loose ends of } H_{n+1}\}| \leq 2m$ neighbours in H_{n+1} , so there is an injection from $\eta(N)$ to the cycle consisting of the neighbours of z . Next fix a ξ -linkage in the outer mesh of M_{F_0} whose paths intersect the boundary of M_{F_0} only in their endpoints. For $l \in L(F_0) \setminus N$ we set $\psi(l) = \xi(\phi(l))$ and let $\psi(e_l)$ be the concatenation of $\phi(e_l)$ with the $\phi(l)$ - $\xi(\phi(l))$ -path in this linkage. For $l \in N$, let $\psi(e_l)$ be the concatenation of $\phi(e_l)$ with the $\phi(l)$ - $\xi(\phi(l))$ -path in this linkage and the incident spoke of M_{F_0} ; note that $\psi(l) = \eta(l)$ is the other endpoint of this spoke. Finally, let $\psi(x) = \phi(x)$ for every vertex or edge x of H'_n .

This clearly gives a $\mathbf{G}(m+1)$ -embedding of H_{n+1} . By definition, ϕ and ψ coincide on H'_n . To see that ψ is a good embedding, note that the boundary sequence of each face F of H_{n+1} incident to v_{n+1} has the form

$$l_{i_j}, B_j(F_0), l_{i_{j+1}}, v_{n+1}, l'_1, v_{n+1}, l'_2, v_{n+1} \dots v_{n+1} l'_s$$

for some $s \geq 0$, where l'_1, \dots, l'_s is the reversal of (possibly empty) sequence of neighbours of v_{n+1} appearing between l_{i_j} and $l_{i_{j+1}}$ in the cyclic order. The order of the loose ends in this sequence coincides with the cyclic order of their embeddings on the cycle C_j . \square

Theorem 1.3 is now an easy consequence of the above result and Theorem 2.2.

Proof of Theorem 1.3. Let G be the disjoint union of the following graphs:

1. countably many copies of \mathbf{G} ,
2. continuum many copies of the cycle C_k for every $k \in \mathbb{N}$,
3. continuum many double rays.

This graph is planar by Theorem 2.2, and it is locally finite since all of the constituent graphs are locally finite.

If H is a locally finite planar graph, then by Theorem 2.2 there are at most countably many components of H containing a vertex of degree 3 or more. Each such component can be embedded into a different copy of \mathbf{G} in G . There are at most continuum many other components all of which are either cycles or (possibly infinite) paths. Each of these components can be embedded into a different copy of some cycle C_k or double ray. \square

Remark 4.3. Theorem 4.2 shows that the class of connected locally finite planar graphs also contains a universal element with respect to the topological minor relation. The above proof of Theorem 1.3 can be easily adapted to yield the same conclusion for the class of countable, locally finite planar graphs.

Remark 4.4. The construction of \mathbf{G} can be modified to give a graph with maximum degree d for any $d \geq 3$. The graphs $\mathbf{M}(n)$ can be built in a way that every vertex except the centre has degree at most 3 (replace vertices of higher degree by appropriate cycles) and the centre has degree d (do not connect it to all vertices on the cycle of length $2n$). Using this modified construction, it is straightforward to check that the above proof shows that for every $d \in \mathbb{N}$, the class of planar graphs with maximum degree at most d contains a universal graph with respect to the topological minor relation (the cases $d \leq 2$ are trivial), and the same is true for the class of such graphs which are connected.

References

- [1] R. Diestel. On universal graphs with forbidden topological subgraphs. *Eur. J. Comb.*, 6:175–182, 1985.
- [2] R. Diestel. *Graph theory*. Graduate Texts in Mathematics. Springer, Berlin, fifth edition, 2017.
- [3] R. Diestel, R. Halin, and W. Vogler. Some remarks on universal graphs. *Combinatorica*, 5:283–293, 1985.
- [4] R. Diestel and D. Kühn. A universal planar graph under the minor relation. *J. Graph Theory*, 32(2):191–206, 1999.
- [5] G. A. Dirac and S. Schuster. A theorem of Kuratowski. *Nederl. Akad. Wet., Proc., Ser. A*, 57:343–348, 1954.
- [6] A. Georgakopoulos. On graph classes with minor-universal elements. 2022. Preprint, arXiv:2212.05498.
- [7] C. W. Henson. A family of countable homogeneous graphs. *Pac. J. Math.*, 38:69–83, 1971.
- [8] T. Huynh, B. Mohar, R. Šámal, C. Thomassen, and D. R. Wood. Universality in minor-closed graph classes. 2021. Preprint, arXiv:2109.00327.
- [9] P. Komjáth and J. Pach. Universal elements and the complexity of certain classes of infinite graphs. *Discrete Math.*, 95(1-3):255–270, 1991.
- [10] T. Krill. On universal graphs with forbidden substructures. Master’s thesis, Universität Hamburg, 2021.
- [11] T. Krill. Universal graphs for the topological minor relation. 2022. Preprint, arXiv:2203.12643.
- [12] D. Kühn. Minor-universal planar graphs without accumulation points. *J. Graph Theory*, 36(1):1–7, 2001.

- [13] F. Lehner. A note on classes of subgraphs of locally finite graphs. *J. Comb. Theory, Ser. B*, 161:52–62, 2023.
- [14] J. Pach. A problem of Ulam on planar graphs. *Eur. J. Comb.*, 2:357–361, 1981.
- [15] R. Rado. Universal graphs and universal functions. *Acta Arith.*, 9:331–340, 1964.
- [16] K. Wagner. Fastplättbare Graphen. *J. Comb. Theory*, 3:326–365, 1967.