

LTCC Course: Graph Theory 2022/23

Lecture 2: Graphs on Surfaces; Graph Minors

Jan van den Heuvel and Yani Pehova

14 November 2022

Chapter 4 of Diestel is good for planar graphs, and Section 1.7 covers the notions of minor and topological minor. Section V.3 of Bollobás covers graphs on surfaces, and colourings thereof.

The definitive textbook for graphs on surfaces is *Graphs on Surfaces*, by Bojan Mohar and Carsten Thomassen, Johns Hopkins (2001); www.fmf.uni-lj.si/~mohar/Book.html.

Most of what we will be discussing in the lectures and notes regarding graph minors can also be found in Chapter 12 of Diestel.

We continue to assume that our graphs are finite and simple (no loops or multiple edges). Much of the material could be adapted for non-simple graphs, but many things will go horribly wrong if we allow our graphs to be infinite.

It is also convenient to assume throughout that a graph has at least one vertex. (To paraphrase Frank Harary: “A graph without vertices is a pointless concept.”)

2.1 Surfaces and Embeddings

A (*closed*) *surface* is a compact connected 2-manifold (i.e. every point has a neighbourhood homeomorphic to the open disc in \mathbb{R}^2). Surfaces can be classified as *orientable* and *non-orientable*. Moreover, each orientable surface is homeomorphic to one of the surfaces S_k , $k \geq 0$, where S_k is a “sphere with k handles”. The sphere itself is S_0 ; the torus is S_1 .

The surface S_k , for $k \geq 1$, can be constructed as follows. Take a convex region in the plane whose boundary is a $4k$ -gon. Label the boundary segments consecutively as

$$\overrightarrow{a_1}, \overrightarrow{b_1}, \overleftarrow{a_1}, \overleftarrow{b_1}, \overrightarrow{a_2}, \overrightarrow{b_2}, \overleftarrow{a_2}, \overleftarrow{b_2}, \dots, \overrightarrow{a_k}, \overrightarrow{b_k}, \overleftarrow{a_k}, \overleftarrow{b_k}.$$

Now identify the pairs of segments labelled $\overrightarrow{a_i}$ and $\overleftarrow{a_i}$, for each i , preserving the orientations given by the arrows, and do likewise for the $\overrightarrow{b_i}$ and $\overleftarrow{b_i}$. It's easy to see that this process identifies all the corners of the $4k$ -gon into one point.

The *genus* of S_k is k , and its *Euler characteristic* χ is $2 - 2k$.

There is a similar construction giving all the non-orientable surfaces: N_k is formed from a $2k$ -gon labelled as

$$\overrightarrow{a_1}, \overrightarrow{a_1}, \overrightarrow{a_2}, \overrightarrow{a_2}, \dots, \overrightarrow{a_k}, \overrightarrow{a_k}.$$

The non-orientable surface N_k ($k \geq 1$) has genus k and Euler characteristic $2 - k$. The first two non-orientable surfaces in the list are the *projective plane* N_1 and the *Klein bottle* N_2 .

More details can be found in Bollobás, for instance.

An *embedding* of a graph $G = (V, E)$ on a surface S is a function taking each vertex x of G to a point $\varphi(x)$ of S , and each edge xy of G to a *Jordan curve* in S , with endpoints $\varphi(x)$ and $\varphi(y)$,

in such a way that the only intersections between the points and curves in the surface are those corresponding to incidences between edges and vertices of G . This all means what you think it ought to mean, and this is exactly how we think of graphs being drawn on a surface: vertices are drawn as different points, edges are curves (not necessarily straight lines!) connecting their endvertices, and two edges don't cross.

A graph can be embedded on the sphere S_0 if and only if it can be embedded on the plane, in which case it is called a *planar graph*.

We'll skip the work required to develop enough machinery to prove anything rigorously. (E.g., the "Jordan curve theorem" says that a closed Jordan curve in the plane has an inside and an outside. Although intuitively this is obvious, it is actually easy to prove.) The result we need is that, if we remove the image of G from the surface S , we are left with a number of connected components called the *faces* of the embedding. The embedding is a *2-cell embedding* if each face is homeomorphic to the open unit disc.

In the case of the sphere S_0 , the only way an embedding of a connected graph can fail to be a 2-cell embedding is if the graph has no vertices. For surfaces with more interesting topology, this is a non-trivial condition.

Two central questions of the subject are: (i) Given a surface S , which graphs can be embedded on S ? (ii) Given a graph embedded on S , what can we say about its chromatic number? We discussed (ii) a bit in the previous lecture. Here, we'll concentrate on (i).

2.2 The Euler-Poincaré Formula

The Euler-Poincaré Formula states that, if we have a 2-cell embedding of a graph on a surface S , then

$$v - e + f = \chi,$$

where v is the number of vertices of the graph, e is the number of edges of the graph, f is the number of faces of the embedding, and χ is the Euler characteristic of the surface S .

We'll just prove this in the case where S is the plane, whose Euler characteristic is 2.

Theorem 1 (Euler's Formula).

Let G be a connected graph with at least one vertex, embedded in the plane. Then $v - e + f = 2$, where $v = |V(G)|$, $e = |E(G)|$, and f is the number of faces of the embedding.

Proof. We work by induction on the number f of faces. When $f = 1$, the graph has no cycles, so is a tree, and $v = e + 1$, which is consistent with the formula.

For $f \geq 2$, we suppose the result is true for embeddings with at most $f - 1$ faces, and take an embedding of a graph with f faces. Choose an edge separating two different faces, and delete it. The graph remains connected: the number of faces has decreased by one, as has the number of edges, while the number of vertices is unchanged. By the induction hypothesis, Euler's Formula holds for the new embedding. Thus it holds for our embedding. Thus, by induction, the formula is valid for all embeddings. \square

Euler's Formula is often quoted as referring to the number of vertices, edges and faces of a *convex polyhedron* in 3-space. The formula for polyhedra follows from the theorem for graphs, as a convex polyhedron can be "drawn in the plane" so that the notions of vertex, edge and face are preserved.

Euler's Formula is often used in conjunction with a "double-counting" of the edges in an embedding. For a face F which is homeomorphic to an open disc, let the *degree* $d(F)$ be the number of edges we encounter on the boundary walk of F . Note that if an edge appears twice in such a

walk, we count it double. This means, for instance, that an embedding of an n -vertex tree in the plane yields one face with degree $2n - 2$.

If \mathcal{F} is the set of faces of the embedding, then we note that $\sum_{F \in \mathcal{F}} d(F)$ counts the total number of edges on the boundaries of all the faces, and that each edge is counted exactly twice by this sum. So we have

$$\sum_{F \in \mathcal{F}} d(F) = 2e.$$

Since all our graphs are finite, each face has at least 3 edges on its boundary, and hence in particular we see that $2e \geq 3f$. So using Euler's Formula now gives that for a planar graph we have

$$e \leq 3v - 6.$$

Notice that this bound makes no mention of the embedding, so it gives a *necessary* condition for a graph to be planar. (Indeed, the same argument gives an upper bound on the number of edges of a graph that can be embedded on any given surface.)

By the Handshaking Lemma, we know that for any graph G we have $\sum_{v \in V(G)} d(v) = 2e$. Combining that with the formula above means $\sum_{v \in V(G)} d(v) \leq 6v - 12$. This means that the average degree of any planar graph is strictly less than 6, so that any planar graph contains a vertex of degree at most 5. Hence $\deg(G) \leq 5$ for any planar graph, which implies that $\chi(G) \leq \text{ch}(G) \leq 6$. We saw last week that this can be improved!


Today, we head in a different direction. From the above inequality, we see that a planar graph on 5 vertices has at most 9 edges. This means that the complete graph K_5 is not planar. Also the complete bipartite graph $K_{3,3}$ is not planar. To see this, notice that, in any embedding of a bipartite graph on a surface, all faces have an even number of sides, so in particular at have degree at least 4. Thus we have $2e = \sum_{F \in \mathcal{F}} d(F) \geq 4f$. Combining that with Euler's Formula leads to $e \leq 2v - 4$, and this is false for $K_{3,3}$.

2.3 Subgraphs and minors


Now we know that K_5 and $K_{3,3}$ are not planar, we can deduce that any graphs "containing" them are not planar. For sure, this is true if our notion of containment is containment as a subgraph, but in fact we can make stronger statements by introducing more general notions of containment.

Let G be a graph. We define the following operations:

- *Removing a vertex* means removing that vertex from the vertex set of G and also removing all edges that vertices is incident with from the edge set.
- *Removing an edge* means removing that edge from the edge set of G .
- *Suppressing a vertex of degree two* means removing that vertex and adding an edge between its two neighbours, provided that edge is not already present (if the edge is already there, we don't add a new one).



- *Contracting an edge*: If $e = xy$ is an edge of G , then contracting e means removing x and y , adding a new vertex z which is adjacent to all vertices that were adjacent to x or y , after which multiple edges are removed.



Let H and G be two graphs.

- H is an induced subgraph of G , or G has H as an induced subgraph, notation $H \leq_I G$, if H can be obtained by a sequence of vertex removals.
- H is a subgraph of G , or G has H as a subgraph, notation $H \leq_S G$, if H can be obtained from G by a sequence of vertex and edge removals.
- H is a topological minor of G , or G has H as a topological minor, notation $H \leq_T G$, if H can be obtained from G by a sequence of vertex removals, edge removals, and suppression of vertices of degree two.
- H is a minor of G , or G has H as a minor, notation $H \leq_M G$, if H can be obtained from G by a sequence of vertex removals, edge removals, and edge contractions.

Note that in the definitions above we allow the sequences to have length zero, so every graph is a subgraph, etc., of itself.

There is a clear hierarchy of the order relations above:

$$H \leq_I G \implies H \leq_S G \implies H \leq_T G \implies H \leq_M G.$$

The following useful result, whose proof is an exercise, gives an alternative characterisation of the minor relation.

Theorem 2.

The following two statements are equivalent for all graphs H, G :

- H is a minor of G .
- For each $u \in V(H)$, there exists a subset $V_u \subseteq V(G)$ of vertices from G so that
 - the sets $\{V_u \mid u \in V(H)\}$ are disjoint;
 - each set $V_u, u \in V(H)$, induces a connected subgraph of G ; and
 - for all $u, v \in V(H)$ with $uv \in E(H)$, there are vertices $x \in V_u, y \in V_v$ with $xy \in E(G)$.

2.4 Minors and Embeddings

Suppose that G is a planar graph, and that H is obtained from G by any of the operations of: vertex removal, edge removal, suppression of a vertex of degree 2, and edge contraction. We claim that H is also planar.

The first two of these are obvious. For suppression of vertices of degree 2, we obtain an embedding of H by replacing the two Jordan curves representing the edges removed from G by a single Jordan curve representing the new edge of H . The same also holds if we replace “planar” by “embeddable on the surface S ”, for any S .

For edge contraction, given an embedding of G , and an edge $e = xy$ of G to be contracted, we derive an embedding of H by placing the new vertex z anywhere on the curve representing xy , and extending all the curves incident with x or y inside thin tubes to reach z , following the path of the curve formerly representing xy .

The arguments above have the following consequence.

Theorem 3.

If G can be embedded on a surface S , and G contains H as a minor, then H can be embedded on S .

We say that the family of graphs that can be embedded on a surface S is *minor-closed*: if G is in the family, and H is a minor of G , then H is in the family.

Because of the hierarchy of the operations introduced in the previous subsection, certainly if G can be embedded on a surface, and G contains H in any of the other senses discussed above, then H can be embedded on the surface.

Returning to the planar case, we now have the following results.

Theorem 4.

If G is planar, then:

- (1) G contains neither K_5 nor $K_{3,3}$ as a minor;
- (2) G contains neither K_5 nor $K_{3,3}$ as a topological minor.

Kuratowski's Theorem says that the converses of both (1) and (2) in the previous theorem are true.

Theorem 5 (Kuratowski, 1930).

The following statements are equivalent for all graphs G :

- (a) G is planar.
- (b) G contains neither K_5 nor $K_{3,3}$ as a minor;
- (c) G contains neither K_5 nor $K_{3,3}$ as a topological minor.

We've seen that (a) implies (b), and (b) implies (c), since if G contains one of the graphs as a topological minor, it contains that graph as a minor.

This section contains only a proof that (c) implies (b): if G contains one of K_5 or $K_{3,3}$ as a minor, then G contains one of K_5 or $K_{3,3}$ as a topological minor. It is important to stress that this is *not* a special case of any general result; it's more in the nature of a lucky accident that having one of K_5 and $K_{3,3}$ as a minor implies having one of the two as a topological minor.

Proof of (c) \Rightarrow (b). We use the characterisation of the graph minor relation given in Theorem 2.

Suppose first that G contains $K_{3,3}$ as a minor, and take a collection of six sets $V_u \subseteq V(G)$ as in Theorem 2, one for each $u \in V(K_{3,3})$. For each set V_u , we identify three edges to V_u from the sets V_w , where w is in the opposite class of $K_{3,3}$ from u . These "land" at three, not necessarily distinct, vertices of V_u : call these x, y, z . It is not hard to see that there is some vertex v_u in V_u (possibly equal to one or more of x, y, z) which has disjoint paths to x, y, z (possibly trivial, i.e. of length zero) in $G[V_u]$. (By $G[V_u]$ we denote the subgraph of G induced by the vertex set $V_u \subseteq V(G)$.) The six vertices w_u , $u \in V(K_{3,3})$, together with the various edges and paths we identified above, form a copy of a graph inside G that contains $K_{3,3}$ as a topological minor.

The paragraph above actually shows that if a graph G contains $K_{3,3}$ as a minor, then G contains $K_{3,3}$ as a topological minor. However (and hopefully this gives some insight into how and why the notions of minor and topological minor are different), if G contains K_5 as a minor, then it need not contain K_5 as a topological minor. Indeed, G can have a K_5 minor even if it has maximum degree 3, but a graph with K_5 as a topological minor must have at least five vertices of degree 4.

To complete the proof of (c) \Rightarrow (b), we need to prove that, if G contains K_5 as a minor, then it contains either K_5 or $K_{3,3}$ as a topological minor. Again we use the characterisation from Theorem 2. So suppose that there are five disjoint connected sets V_a, V_b, V_c, V_d, V_e in G , with edges between each pair. The plan is to set off trying to find K_5 as a topological minor. So, for each V_i , $i \in \{a, b, c, d, e\}$, we find the four "landing points" $x_{ij} \in V_i$ of the edges that connect V_i to the other V_j . Either there is a vertex w_i in V_i with four disjoint paths to the x_{ij} , or there are two vertices f and g in V_i , connected by a path, with two of the x_{ij} sending paths to f and the other two to g , all five paths being internally disjoint. If the first case occurs for all V_i , then we found a

topological minor of K_5 in G . If the latter case occurs with any of the V_i , then we change plans: in this case, we divide V_i into two connected parts, one containing f and two of the x_{ij} , and the other containing g and the other two x_{ij} . The six vertex sets now witness that $K_{3,3}$ is also a minor, and therefore a topological minor, of G . \square

Of course, the hard part of Theorem 5 is to prove that (b) and (c) imply (a); given a graph G which doesn't contain K_5 or $K_{3,3}$ as a minor, we want to find a plane embedding of G .

We give only a brief sketch. The proof goes by induction on $v(G)$. If G is not connected, then we draw its connected components (by induction) disjointly in the plane.¹ If G is connected but not 2-connected, we can do something similar: let v be a vertex such that $G - v$ is disconnected, and let C be a component of $G - v$. We can draw the graphs $G - C$ and $G[C \cup \{v\}]$ in the plane by induction; we put the drawing of the second, without loss of generality with v on the outer face, in a face of the drawing of the first containing v , and identify the vertex v in the two drawings.

To deal with 2-connected but not 3-connected G becomes a little harder. Suppose $\{u, v\}$ is a separating set, that is $G - \{u, v\}$ is disconnected, and let C be a component of $G - \{u, v\}$; let $G_1 = G - C$ and $G_2 = G[C \cup \{u, v\}]$. If uv is an edge of G , then we can repeat the trick of drawing G_1 and G_2 by induction and putting the drawing of G_2 , with uv on an outer face, in a face of G_1 containing uv , then identifying u and v . If uv is not an edge of G , this doesn't work. But the trick is to notice that we can simply add uv to both G_1 and G_2 without creating a K_5 or $K_{3,3}$ topological minor, then use induction. Suppose $G_1 \cup \{uv\}$ has one of these two graphs as a topological minor. The edge uv must be used in this topological minor, since G_1 does not contain $K_{3,3}$ or K_5 as a topological minor. But there is a path in G_2 from u to v : for if not, then $G - u$ would be disconnected. We can replace uv with this path to find $K_{3,3}$ or K_5 as a topological minor in G , a contradiction.

Finally, we can assume G is 3-connected. This is the point where we really need to do some work, and where we will skip most of the details. Note that this sketch is substantially appealing to your intuition on how planar drawings work; if you were forced to actually write down details of how to identify the two copies of uv given planar drawings, it would be rather painful (and long). We can assume that all edges in all our planar drawings are straight lines (this property would easily be preserved by the induction steps above, and the following sketch gives it for the 3-connected case) which makes life rather easier, but still not really easy.

Briefly, to deal with the case G is 3-connected, either we have $G = K_4$ and it is easy to draw a plane embedding (with straight lines) or G has 5 or more vertices. One can fairly easily check that in this case (Lemma 3.2.4 in Diestel) there is some edge xy such that when we contract xy we obtain a graph G' which is still 3-connected. We can draw this in the plane by induction; let z be the vertex obtained by contracting xy . Let C be the face of $G' - z$ containing z . We try to draw x and y inside C , with a straight line joining them, and add more straight lines to form a drawing of C . Now by a case analysis (Lemma 4.4.3 of Diestel) you can check that either this is possible, or you find a K_5 or $K_{3,3}$ minor.

2.5 Orderings and closedness of properties

We will now leave the topic of graphs on surfaces, and examine the notions of graph containment for their own sake. To begin with, we observe that our relations of containment are all *transitive* (if G contains H and H contains J , then G contains J), and so give “orderings” on the set of all graphs. Let us be more precise.

¹Note that this works for the plane; it would not work for the torus or other more complicated surfaces.

If \preceq is a relation on a set X , then (X, \preceq) is called a *quasi-ordering* or *pre-order* if the relation is reflexive ($x \preceq x$ for all $x \in X$) and transitive ($(x \preceq y \wedge y \preceq z) \Rightarrow (x \preceq z)$ for all $x, y, z \in X$). We say that a quasi-ordering (X, \preceq) is *without infinite descent* if there is no infinite strictly decreasing sequence $x_1 \succ x_2 \succ x_3 \succ \dots$ (where $x \succ y$ means $y \preceq x$ and $x \neq y$). It is easy to see that the orderings defined in the previous section on the class \mathcal{G} of all (simple, finite) graphs correspond to quasi-orderings without infinite descent.

A subset $A \subseteq X$ of a quasi-ordering (X, \preceq) is an *antichain* if every two elements from A are incomparable (i.e. if $a, b \in A$ with $a \neq b$, then $a \not\preceq b$ and $b \not\preceq a$).

Proposition 6.

Given a quasi-ordering (X, \preceq) , any infinite sequence a_1, \dots in X either repeats one term infinitely often or contains at least one of the following three things: an infinite strictly decreasing sequence, an infinite strictly increasing sequence, and an infinite antichain.

Proof. This is really a special case of the infinite Ramsey theorem, and we'll prove it in that language. (We shall return to Ramsey theory in the last lecture of this course.) We colour the complete graph on \mathbb{N} with four colours, colouring the edge ij (for $i < j$) with the relation between a_i and a_j (which can be one of $a_i = a_j$, or $a_i \prec a_j$, or $a_j \prec a_i$, or none of the above). We claim that for any 4-colouring of the complete graph on \mathbb{N} there is an infinite set S such that all pairs in S have the same colour; the existence of this set proves the proposition.

To find such a set, we work as follows. Let $s_1 = 1$, and colour s_1 with a colour c_1 appearing on infinitely many edges leaving 1; let T_1 be the set of vertices joined to s_1 with colour c_1 . Now for each $i = 2, \dots$, let s_i be the smallest element of T_{i-1} . Let c_i be a colour appearing on infinitely many edges from s_i to T_{i-1} , and let T_i be the set of vertices in T_{i-1} joined to s_i with colour c_i . We obtain in this way an infinite sequence since by construction each T_i is infinite.

Now pick a colour c which appears infinitely often in the sequence c_1, \dots , and let $S = \{s_i \mid c_i = c\}$. It is easy to check that S is the desired infinite monochromatic clique. Note that this proof doesn't rely on having specifically 4 colours; any finite number would work. □

Let P be a property defined on the elements of X . We say that P is *closed under \preceq* or *\preceq -closed* if for every two elements $x, y \in X$ we have that if x has property P and $y \preceq x$, then y also has property P .

As an example, suppose property P is defined for $G \in \mathcal{G}$ (the set of all graphs) as “ G is bipartite”. This property is closed under both the subgraph and the induced subgraph ordering, but not under the topological minor or the minor ordering. (See Exercises.)

Let (X, \preceq) be a quasi-ordering and suppose P is a \preceq -closed property defined on the elements of X . Then we can talk about the set \mathcal{P} of all elements in X that satisfy property P . And of course we also have the complement $\overline{\mathcal{P}} = X \setminus \mathcal{P}$ of all elements in X that do *not* satisfy property P . Let M be a set of minimal elements of $\overline{\mathcal{P}}$

Since P is assumed to be \preceq -closed, we know that if $x \in \overline{\mathcal{P}}$ and $x \preceq y$, then $y \in \overline{\mathcal{P}}$. This leads to the following crucial observation:

$$x \text{ has property } P \iff \text{there is no } m \in M \text{ with } m \preceq x.$$

In other words: a property that is \preceq -closed is completely determined once we know a set of minimal elements of the set of elements that don't have the property. Such a minimal set is called a *minimal forbidden set* of the property.

The observations above *may* provide a good description of certain properties and *may* provide efficient algorithms to test if a given element satisfies the property. This possible usefulness depends

on the answers to questions like: Can we find a minimal forbidden set? Is this set finite? Is there a good algorithm to test if $x \preceq y$ or not? Etc.


You may wonder why we look at a set of minimal elements of the set $\overline{\mathcal{P}}$ of elements in X that do *not* satisfy property P . Wouldn't it be more natural to look at the set of maximal elements of \mathcal{P} ? Yes, it would be more natural. But for the orderings we are considering such a set of maximal elements usually doesn't exist. The orderings give natural minimal elements, since from every finite graph, we can only have a finite number of descending steps before we have to stop (we've reached a graph with one vertex, say). But in general we won't have maximal elements (except for very special properties).

Here is an example illustrating the concepts above. For a graph $G = (V, E)$, recall that the *line graph* $L(G) = (V_L, E_L)$ is the graph that has the edges of G as vertices: $V_L = E$; and two edges are adjacent in the line graph if they have a common end-vertex in G . A graph H is a line graph if $H \cong L(G)$ for some graph G .

It's easy to see that if H is a line graph, then every induced subgraph of H is also a line graph. Hence the property of "being a line graph", defined on the set \mathcal{G} of graphs, is closed under the induced subgraph ordering \leq_I . For this property, we actually do know the unique set of minimal forbidden elements.

Theorem 7 (Beineke, 1968).

A graph H is a line graph if and only if it does not contain one of the nine graphs below as an induced subgraph.



2.6 Well-quasi-ordering

A quasi-ordering (X, \preceq) is a *well-quasi-ordering* if for every infinite sequence x_1, x_2, \dots of elements from X , there are two indices $i < j$ so that $x_i \preceq x_j$. In particular, Proposition 6 implies the following.

Proposition 8.

The following two properties are equivalent for a quasi-ordering (X, \preceq) :

- (X, \preceq) is a well-quasi-ordering.
- (X, \preceq) is a quasi-ordering without infinite descent and without infinite antichains. □

Requiring a quasi-ordering to be a well-quasi-ordering is a very strong requirement. For instance, for the graph orderings defined in the first section, neither (\mathcal{G}, \leq_I) , nor (\mathcal{G}, \leq_S) , nor (\mathcal{G}, \leq_T) , are well-quasi-orderings. For the induced subgraph ordering and the subgraph ordering, the sequence of cycles C_3, C_4, C_5, \dots forms an infinite sequence that fails the condition in the definition. In one of the exercises you will be asked to find counterexamples yourself for the topological minor ordering.

A fundamental property of well-quasi-orderings is that given a well-quasi-ordering (X, \preceq) , one can build a new well-quasi-ordering as follows. Let $X^{<\infty}$ denote the set of finite subsets of X . Given $A, B \in X^{<\infty}$, we write $A \preceq B$ if there is an injection $f : A \rightarrow B$ such that $a \preceq f(a)$ for each $a \in A$. This turns out to be a well-quasi-order.

Lemma 9 (Higman's Lemma, 1952).

The quasi-ordering $(X^{<\infty}, \preceq)$ is a well-quasi-ordering.

Proof. It's easy to check that since (X, \preceq) is a quasi-ordering, so is $(X^{<\infty}, \preceq)$. So what we need to show is that there is no infinite descending sequence or antichain; in other words, for any sequence A_1, \dots in $X^{<\infty}$, there are $i < j$ such that $A_i \preceq A_j$. Call a sequence with this property *good*, and one which does not have it *bad*. Suppose for contradiction that such bad sequences exist.

We now construct a sequence. First, let A_1 be an element of $X^{<\infty}$ with fewest elements such that A_1 is the first term of some bad sequence. Now, for each $i = 2, \dots$, given A_1, \dots, A_{i-1} , let A_i be an element of $X^{<\infty}$ with fewest elements such that A_1, \dots, A_i is an initial segment of some bad sequence. By construction, the result is a bad sequence.

Because $\emptyset \preceq A$ for each $A \in X^{<\infty}$, we have $|A_i| \geq 1$ for each i . Thus for each i we can choose $a_i \in A_i$. Let $B_i = A_i \setminus \{a_i\}$ for each i . Because (X, \preceq) is a well-quasi-ordering, by Proposition 6 there is an infinite (not necessarily strictly) increasing subsequence (a_{i_1}, \dots) of $(a_i)_{i=1}^\infty$. Now consider the sequence $A_1, \dots, A_{i_1-1}, B_{i_1}, B_{i_2}, \dots$. By choice of A_{i_1} this sequence is good. We do not have $A_i \preceq A_j$ for any $i < j$, and since $B_j \preceq A_j$ using the trivial identity injection, we also do not have $A_i \preceq B_j$ for any $i < j$. Thus the reason that $A_1, \dots, A_{i_1-1}, B_{i_1}, B_{i_2}, \dots$ is good is that there exist k and ℓ with $k < \ell$ such that $B_{i_k} \preceq B_{i_\ell}$. Let $\phi : B_{i_k} \rightarrow B_{i_\ell}$ be an injection witnessing $B_{i_k} \preceq B_{i_\ell}$, and let ϕ' be the injection from A_{i_k} to A_{i_ℓ} extending ϕ with $\phi'(a_{i_k}) = a_{i_\ell}$. Now ϕ' witnesses $A_{i_k} \preceq A_{i_\ell}$, a contradiction. \square

Although the whole class of graphs is not well-quasi-ordered under the topological minor ordering, some important subclasses are.

Theorem 10 (Kruskal, 1960).

The class of all trees, with topological minor as the ordering, is well-quasi-ordered.

The proof of this is quite similar to that of Lemma 9, with one extra trick.

Proof. We aim to prove a slightly stronger statement. A *rooted tree* is a tree T with one vertex identified as the *root*. Now a tree T is a topological minor of another tree T' if and only if we can subdivide the edges of T (that is, we are allowed to perform the operation of replacing an edge uv with a new vertex w and the edges uw, vw as often as we like) and find an isomorphism ϕ from the result to a subgraph of T' . If T and T' are rooted trees, we say that in addition ϕ is *order-preserving* if the following condition holds. For any $x, y \in V(T)$ such that the path from the root of T to y goes through x , so the path from the root of T' to $\phi(y)$ goes through $\phi(x)$. We will prove that the class \mathcal{T} of all rooted trees, with order-preserving topological minor as the ordering, is well-quasi-ordered. We denote this ordering by $\preceq_{\mathcal{O}}$.

Let us now show that no bad sequence (as in Lemma 9) of rooted trees exists. Suppose to the contrary that such a bad sequence exists. We construct a sequence T_1, T_2, \dots inductively as follows

(essentially as in the proof of Lemma 9). Let T_1 be a smallest (number of vertices) rooted tree which is in such a bad sequence. Now for each $n \geq 1$ successively, given T_1, \dots, T_n , choose T_{n+1} to be a smallest rooted tree such that T_1, \dots, T_{n+1} is an initial segment of a bad sequence. Let r_1, \dots be the roots of the trees T_1, \dots . Note that trivially each T_i has at least two vertices (actually at least three is also trivial; but we won't need it).

For each i , observe that $T_i - r_i$ forms a finite collection of trees, which we root at the vertices adjacent to r_i in T_i . For each i , let A_i be the set of rooted trees $T_i - r_i$. Now our goal is to show that for some $i < j$ we have $A_i \preceq_O A_j$, where as with Lemma 9 we derive a quasi-ordering on $\mathcal{T}^{<\infty}$ from that on \mathcal{T} . Given such a pair, since $A_i \preceq_O A_j$ we have a collection of maps witnessing this, whose union is a map from $T_i - r_i$ to $T_j - r_j$. Extending this map by adding $r_i \rightarrow r_j$ we have a map from T_i to T_j which witnesses $T_i \preceq_O T_j$; that is the contradiction we want.

The trick now is to show that the subset of \mathcal{T} formed by the rooted trees in $\mathcal{A} := \bigcup_i A_i$ is well-quasi-ordered by \preceq_O . By Lemma 9, this implies that $\mathcal{A}^{<\infty} \subseteq \mathcal{T}^{<\infty}$ is well-quasi-ordered by \preceq_O . Since A_1, \dots is a sequence in $\mathcal{A}^{<\infty}$, we thus get $i < j$ such that $A_i \preceq_O A_j$ as desired.

It remains to show that (\mathcal{A}, \preceq_O) is a well-quasi-ordering. To that end, let B_1, \dots be a sequence of trees in \mathcal{A} . For each i , let $n(i)$ be such that $B_i \in A_{n(i)}$, and choose k minimising $n(k)$. Trivially B_k has fewer vertices than $T_{n(k)}$, so by construction of the sequence T_1, \dots , it follows that the sequence $T_1, T_2, \dots, T_{n(k)-1}, B_k, B_{k+1}, \dots$ is good. Let $B \preceq_O B'$ be a pair witnessing that this sequence is good (so B comes before B' in the sequence). Now B cannot be in the initial segment $T_1, \dots, T_{n(k)-1}$, since otherwise either B' is also in this initial segment, contradicting badness of T_1, \dots , or $B' = B_i$ for some i , but then we have $B \preceq_O B' = B_i \preceq_O T_{n(i)}$ and since $n(i) > n(k) - 1$ by choice of k , this too is a contradiction to badness of T_1, \dots . But then we have $B = B_i$ and $B' = B_j$ for some $i < j$; in other words, the sequence B_1, \dots is good, as desired. \square

An interesting point to note here is that this proof (and that of Higman's Lemma 9) really use infinite sets and the Axiom of Choice in a fundamental way; not a very common thing in discrete mathematics. It's easy enough to see that in Theorem 10 we really need an infinite set to guarantee finding two trees where one is a topological minor of another (for any finite n we can just take all non-isomorphic n -vertex trees; the number of such trees tends to infinity with n). But Kruskal's theorem in particular implies the following, observed by Friedman.

Corollary 11.

Let $\text{tree}(n)$ denote the minimum m such that the following statement holds. For every sequence T_1, \dots, T_m of trees, where T_i has $i + n$ vertices, there is some pair with one a topological minor of the other.

Then $\text{tree}(n)$ exists.

Proof. This is a classical *compactness argument*. Suppose $\text{tree}(n)$ does not exist; then for each m there is a counterexample sequence of m trees. Let T_1 be a tree which plays the rôle of T_1 in infinitely many of these sequences, and discard the other sequences. Let T_2 be a tree which plays the rôle of T_2 in infinitely many of these sequences, and discard the other sequences. Repeat forever; the result is an infinite sequence of trees, which by Theorem 10 is good. Fix a pair T_i, T_j with $i < j$ witnessing this, and we see that the supposed counterexample sequence for $m = j$ contains the pair T_i, T_j where T_i is a topological minor of T_j . \square

This proof provides no hint of how fast $\text{tree}(n)$ grows. In fact, it turns out to grow incredibly fast; so fast that it cannot really be described by 'Peano arithmetic' (this is a minimal system of axioms that let you perform arithmetic; normally that's about all we need in discrete mathematics), so that while you can prove any given $\text{tree}(n)$ exists in Peano arithmetic (by brute force enumeration,

say) you cannot prove Corollary 11 in Peano arithmetic (at least, if ZFC is consistent). You really need to assume stronger axioms (like the Axiom of Choice) which normally aren't needed in discrete mathematics. This is also a result of Friedman (indeed, this motivated his observation of Corollary 11); the idea is to show that on the one hand Peano arithmetic cannot prove the existence of very fast-growing functions, and on the other hand that $\text{tree}(n)$ does grow so fast.

We conclude this section with an easy but important proposition.

Proposition 12.

Let (X, \preceq) be a well-quasi-ordering and P a \preceq -closed property on X . Then a minimal forbidden set of P is finite.

Proof. Let \mathcal{F} be a forbidden set of P . If \mathcal{F} is infinite, then by definition there are $F_1 \neq F_2$ such that $F_1 \preceq F_2$. Let $\mathcal{F}' = \mathcal{F} - \{F_2\}$. We claim that \mathcal{F}' is a forbidden set of P . Indeed, if $Q \notin P$, then by definition of \mathcal{F} there exists $F \in \mathcal{F}$ such that $F \preceq Q$. Either $F \neq F_2$, in which case $F \in \mathcal{F}'$, or $F = F_2$, in which case we have $F_1 \preceq F_2 \preceq Q$ and hence $F_1 \preceq Q$. It follows that \mathcal{F} is not a minimal forbidden set of P . □

The importance of this property is that once we know that (X, \preceq) is a well-quasi-ordering, then every property that is \preceq -closed has a finite minimal forbidden set. If we also have, for each $x \in X$, an algorithm which efficiently answers the question ‘Given $y \in X$, is $x \preceq y$?’, then that proves that for any \preceq -closed property P , there exists an efficient algorithm that decides the question ‘Given $y \in X$, is $y \in P$?’.

But note that the proofs we saw so far that some quasi-orderings are well-quasi-orderings are not constructive. In particular, they don't help us find these minimal forbidden sets guaranteed by Proposition 12. So when we can test $x \preceq y$ efficiently, we are in the rather strange position of knowing efficient algorithms exist without being able to find them. We'll see in the next section that this is exactly the case for minor-closed graph properties.

2.7 Minors of graphs

In the early 1980s, Robertson and Seymour announced a proof of a conjecture of Wagner: that the class of finite graphs with the minor ordering is a well-quasi-ordering. The proof spans 20 papers and over 400 pages which appeared over the next 20 years, the ‘Graph Minors’ series (and the series continued for a few more papers to tie up some related problems).

Theorem 13 (Graph Minor Theorem, Robertson & Seymour, 1986–2004).

The class of finite graphs is well-quasi-ordered under the minor ordering.

They also proved, in this series, that indeed, for any fixed G , one can check if G is a minor of an input graph H in time cubic in $|E(H)|$ (and the algorithm for this is completely explicit), so that for any given minor-closed class of graphs \mathcal{G} , one can decide in cubic time if $H \in \mathcal{G}$. But, unless you are given \mathcal{G} in a way which makes it obvious what the list of forbidden minors is, then you probably cannot find this list in any useful way. Indeed, if \mathcal{G} is given via Turing machines (which is a reasonable enough way), then the problem is known to be undecidable (Fellows and Langston, 1989). To take a more practical view, we know there are 2 forbidden minors for the class of planar graphs (Theorem 5). For the next orientable surface, the torus, there are over 239,000 known forbidden minors (found by computer search); and there is no reason to believe that the total number is close to this.

We can try to sketch, very briefly(!) how the Graph Minor Theorem is proved. For now, consider only planar graphs. Let G_1, \dots be a sequence of planar graphs.

If this sequence contains as minors arbitrarily large ‘grids’, then it is a good sequence. A grid of size k is simply a graph drawn on the integer points $[k] \times [k]$ in the plane by joining each point to the at most four points at unit distance from it. A large grid ‘obviously’ contains any small planar graph (say G_1) as a minor: to see this, enlarge your (straight line) planar drawing of G_1 by a large factor, replace the edges and vertices with small rectangles and circles respectively, and superimpose it on the grid; this gives an ‘obvious’ minor.

On the other hand, the case that all the planar graphs ‘look like a tree’ is easy (at least by comparison to the rest of this monumental proof) to handle; one can start with Theorem 10 and try to generalise it.

To be more precise about what ‘looks like a tree’ means, we need to say what a *tree decomposition* of a graph G is. This is a tree T , whose nodes are sets of vertices of G , with the following two properties. First, every edge of G goes between two vertices which are in some node set of T . Second, every vertex of G is in a collection of node sets of T which form a subtree of T . The *width* of such a tree decomposition is defined to be the size of a largest node set minus one (the reason for subtracting 1 is just to look pretty: so that trees have tree decompositions of width 1). Finally, the *treewidth* of a graph G is the minimum width of any tree decomposition of G . This parameter turns out to be important in a lot of graph theory, especially when one starts doing algorithms. Here, the relevance is that if there is some r such that all graphs G_1, \dots have treewidth at most r , then one can push through the generalisation of Theorem 10 to show that G_1, \dots is a good sequence.

Finally, we are left to deal with sequences G_1, \dots which don’t contain arbitrarily large grid minors, but where there is no bound on the treewidth. Somewhat surprisingly (and the proof is not easy), Robertson and Seymour proved that there is no such sequence: for any k , there is r such that a graph whose treewidth exceeds r necessarily contains a grid of size k as a minor.

Moving to the general case, very vaguely, the idea is as follows. First, we can assume no graph in our sequence contains a K_t minor for $t = v(G_1)$. Then, we prove that any graph which doesn’t contain a K_t minor has a tree-decomposition in which all the node-sets induce ‘nice’ graphs, and we prove that ‘nice’ graphs are well-quasi-ordered, and then we can use something like Kruskal’s theorem to put the pieces together. Unfortunately the definition of ‘nice’ gets rather complicated for the general case; a short version is ‘almost embeds in some surface that K_t does not embed into’, but this short version omits so many technicalities it’s almost a lie. In any case, the point is that there are only finitely many of these surfaces, and a sufficiently complicated graph on a surface S contains all small graphs on S as minors (by the same sort of argument as why a large grid contains all small planar graphs). Of course, this sketch too leaves out so many details it is almost closer to false than true.

Exercises

Exercise 1.

How does Euler's Formula for graphs embedded in the plane need to be modified to handle graphs with c components, where c is not necessarily equal to 1?

Exercise 2.

Consider the oriented surface S_k , $k \geq 1$. Draw a graph in S_k by putting a vertex at each corner of the boundary $4k$ -gon, and an edge along each segment of the boundary, identifying any of these edges and vertices as necessary. Count the number of vertices, edges, and faces in this embedding, and verify the Euler-Poincaré Formula in this case.

Exercise 3.

- (a) Describe the graphs not containing K_3 as a minor.
- (b) Describe the graphs not containing the 4-cycle C_4 as a minor.

Exercise 4.

Let P be the Petersen graph. (If you don't know what this is, find out!)

- (a) Show that P is non-planar in the following three ways:
 - using Euler's Formula;
 - by showing that P contains $K_{3,3}$ as a topological minor;
 - by showing that P contains K_5 as a minor.
- (b) What is the minimum size of a set F of edges of P whose deletion leaves a planar graph?

Exercise 5.

Show that the property of being bipartite is not closed for the topological minor ordering on graphs. (Very easy)

Exercise 6.

Prove that the property of having no connected component with more edges than vertices is minor-closed.

Find as many minimal forbidden minors as you can for this property.

Exercise 7.

Prove that the set of finite simple graphs \mathcal{G} with the topological minor ordering \leq_T is not a well-quasi-ordering. In other words, give an infinite sequence of graphs G_1, G_2, \dots , for which there are no two indices i, j with $i < j$ and $G_i \leq_T G_j$.

(This is a hard question. Feel free to do an Internet search, but you must show that the sequence you give has the desired property.)