8 Infinite Graphs

The study of infinite graphs is an attractive, but often neglected, part of graph theory. This chapter aims to give an introduction that starts gently, but then moves on in several directions to display both the breadth and some of the depth that this field has to offer. Our overall theme will be to highlight the typical kinds of phenomena that will always appear when graphs are infinite, and to show how they can lead to deep and fascinating problems.

Perhaps the most typical such phenomena occur already when the graphs are 'only just' infinite, when they have only countably many vertices and perhaps only finitely many edges at each vertex. This is not surprising: after all, some of the most basic structural features of graphs, such as paths, are intrinsically countable. Problems that become really interesting only for uncountable graphs tend to be interesting for reasons that have more to do with sets than with graphs, and are studied in *combinatorial set theory*. This, too, is a fascinating field, but not our topic in this chapter. The problems we shall consider will all be interesting for countable graphs, and set-theoretic problems will not arise.

The terminology we need is exactly the same as for finite graphs, except when we wish to describe an aspect of infinite graphs that has no finite counterpart. One important such aspect is the eventual behaviour of the infinite paths in a graph, which is captured by the notion of *ends*. The ends of a graph can be thought of as additional limit points at infinity to which its infinite paths converge. This convergence is described formally in terms of a natural topology placed on the graph together with its ends. In Sections 6–8 we shall therefore assume familiarity with the basic concepts of point-set topology; reminders of the relevant definitions will be included as they arise.

8.1 Basic notions, facts and techniques

This section gives a gentle introduction to the aspects of infinity most commonly encountered in graph theory.[1]

After just a couple of definitions, we begin by looking at a few obvious properties of infinite sets, and how they can be employed in the context of graphs. We then illustrate how to use the three most basic common tools in infinite graph theory: Zorn's lemma, transfinite induction, and something called 'compactness'. We complete the section with the combinatorial definition of an end; topological aspects will be treated in Section 8.6.

locally finite

A graph is *locally finite* if all its vertices have finite degrees. An infinite graph (V, E) of the form

$$V = \{x_0, x_1, x_2, \dots\} \qquad E = \{x_0x_1, x_1x_2, x_2x_3, \dots\}$$

rays

is called a *ray*, and a *double ray* is an infinite graph (V, E) of the form

$$V = \{\dots, x_{-1}, x_0, x_1, \dots\} \qquad E = \{\dots, x_{-1}x_0, x_0x_1, x_1x_2, \dots\}\,;$$

in both cases the x_n are assumed to be distinct. Thus, up to isomorphism, there is only one ray and one double ray, the latter being the unique infinite 2-regular connected graph. In the context of infinite graphs, finite paths, rays and double rays are all called *paths*.

path

tail

The subrays of a ray or double ray are its *tails*. Formally, every ray has infinitely many tails, but any two of them differ only by a finite initial segment. The union of a ray R with infinitely many disjoint finite paths having precisely their first vertex on R is a *comb*; the last vertices of those paths are the *teeth* of this comb, and R is its *spine*. (If such a path is trivial, which we allow, then its unique vertex lies on R and also counts as a tooth; see Figure 8.1.1.)

comb

teeth, spine

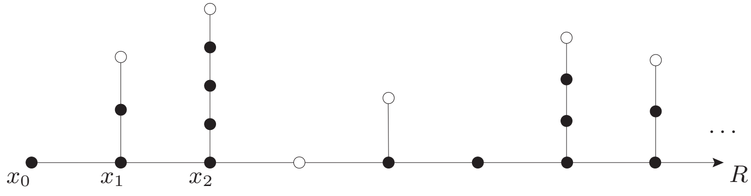

Fig. 8.1.1. A comb with white teeth and spine $R = x_0x_1\dots$

[1] This introductory section is deliberately kept informal, with the emphasis on ideas rather than definitions that do not belong in a graph theory book. A more formal reminder of those basic definitions about infinite sets and numbers that we shall need is given in Appendix A at the end of the book.

Let us now look at a few very basic properties of infinite sets, and see how they appear in some typical arguments about graphs.

An infinite set minus a finite subset is still infinite. (1)

This trivial property is eminently useful when the infinite set in question plays the role of 'supplies' that keep an iterated process going. For example, let us show that if a graph G is infinitely connected (that is, if G is k-connected for every $k \in \mathbb{N}$), then G contains a subdivision of $K^{\aleph_0}$, the complete graph of order $|\mathbb{N}|$. We embed $K^{\aleph_0}$ in G (as a topological minor) in one infinite sequence[2] of steps, as follows. We begin by enumerating its vertices. Then at each step we embed the next vertex in G, connecting it to the images of its earlier neighbours by paths in G that avoid any other vertices used so far. The point here is that each new path has to avoid only finitely many previously used vertices, which is not a problem since deleting any finite set of vertices keeps G infinitely connected. $K^{\aleph_0}$

If G, too, is countable, can we then also find a $TK^{\aleph_0}$ as a spanning subgraph of G? Although embedding $K^{\aleph_0}$ in G topologically as above takes infinitely many steps, it is by no means guaranteed that the $TK^{\aleph_0}$ constructed uses all the vertices of G. However, it is not difficult to ensure this: since we are free to choose the image of each new vertex of $K^{\aleph_0}$, we can choose this as the next unused vertex from some fixed enumeration of $V(G)$. In this way, every vertex of G gets chosen eventually, unless it becomes part of the $TK^{\aleph_0}$ before its time, as a subdividing vertex on one of the paths.

Unions of countably many countable sets are countable. (2)

This fact can be applied in two ways: to show that sets that come to us as countable unions are 'small', but also to rewrite a countable set deliberately as a disjoint union of infinitely many infinite subsets. For an example of the latter type of application, let us show that an infinitely edge-connected countable graph has infinitely many edge-disjoint spanning trees. (Note that the converse implication is trivial.) The trick is to construct the trees simultaneously, in one infinite sequence of steps. We first use (2) to partition $\mathbb{N}$ into infinitely many infinite subsets N_i ($i \in \mathbb{N}$). Then at step n we look which N_i contains n, and add a further vertex v to the ith tree T_i. As before, we choose v minimal in some fixed enumeration of $V(G)$ among the vertices not yet in T_i, and join v to T_i by a path avoiding the finitely many edges used so far.

Clearly, a countable set cannot have uncountably many disjoint subsets. However,

[2] We reserve the term 'infinite sequence' for sequences indexed by the set of natural numbers. (In the language of well-orderings: for sequences of order type ω.)

A countable set can have uncountably many subsets whose pairwise intersections are all finite. (3)

This is a remarkable property of countable sets, and a good source of counterexamples to rash conjectures. Can you prove it without looking at Figure 8.1.4?

Another common pitfall in dealing with infinite sets is to assume that the intersection of an infinite nested sequence $A_0 \supseteq A_1 \supseteq \ldots$ of infinite (or uncountable) sets must still be infinite (or uncountable). It need not be; in fact it may be empty. (Examples?)

Before we move on to our discussion of common infinite proof techniques, let us look at one more type of construction. One often wants to construct a graph G with a property that is in some sense local, a property that has more to do with the finite subgraphs of G than with G itself. Rather than formalize what exactly this should mean, let us consider an example: given two large integers k and g, let us construct a graph G that is k-connected and has girth at least g.[3]

We start with a cycle of length g; call it G_0. This graph has the right girth, but it is not k-connected. To cure this defect for the vertices of G_0, join every pair of them by k new independent paths, keeping all these paths internally disjoint. If we choose the paths long enough, the resulting graph G_1 will again have girth g, and no two vertices of G_0 can be separated in it by fewer than k other vertices. Of course, G_1 is not k-connected either. But we can repeat the construction step for the pairs of vertices of G_1, extending G_1 to G_2, and so on. The limit graph $G = \bigcup_{n \in \mathbb{N}} G_n$ will again have girth g, since any short cycle would have appeared in some G_n on the way. And, unlike all the G_n, it *will* be k-connected: since every two vertices are contained in some common G_n, they cannot be separated by fewer than k other vertices in G_{n+1}, let alone in G.

There are a few basic proof techniques that are found frequently in infinite combinatorics. The two most common of these are the use of Zorn's lemma and transfinite induction. Rather than describing these formally,[4] we illustrate their use by a simple example.

Proposition 8.1.1. *Every connected graph contains a spanning tree.*

First proof (by Zorn's lemma).
Given a connected graph G, consider the set of all trees $T \subseteq G$, ordered by the subgraph relation. Since G is connected, any maximal such tree contains every vertex of G, i.e. is a spanning tree of G.

---

[3] There are finite such graphs, but they are much harder to construct; we shall prove their existence by random methods in Chapter 11.2.

[4] Appendix A offers brief introductions to both, enough to enable the reader to use these tools with confidence in practice.

To prove that a maximal tree exists, we have to show that for any chain $\mathcal{C}$ of such trees there is an upper bound: a tree $T^* \subseteq G$ containing every tree in $\mathcal{C}$ as a subgraph. We claim that $T^* := \bigcup \mathcal{C}$ is such a tree.

To show that T^* is connected, let $u, v \in T^*$ be two vertices. Then in $\mathcal{C}$ there is a tree T_u containing u and a tree T_v containing v. One of these is a subgraph of the other, say $T_u \subseteq T_v$. Then T_v contains a path from u to v, and this path is also contained in T^*.

To show that T^* is acyclic, suppose it contains a cycle C. Each of the edges of C lies in some tree in $\mathcal{C}$. These trees form a finite subchain of $\mathcal{C}$, which has a maximal element T. Then $C \subseteq T$, a contradiction. □

Transfinite induction and recursion are very similar to finite inductive proofs and constructions, respectively. Basically, one proceeds step by step, and may at each step assume as known what was shown or constructed before. The only difference is that one may 'start again' after performing any infinite number of steps. This is formalized by the use of ordinals rather than natural numbers for counting the steps; see Appendix A.

Just as with finite graphs, it is usually more intuitive to construct a desired object (such as a spanning tree) step by step, rather than starting with some unknown 'maximal' object and then proving that it has the desired properties. More importantly, a step-by-step construction is almost always the best way to *find* the desired object: only later, when one understands the construction well, can one devise an inductive ordering (one whose chains have upper bounds) in which the desired objects appear as the maximal elements. Thus, although Zorn's lemma may at times provide an elegant way to wrap up a constructive proof, it cannot in general replace a good understanding of transfinite induction – just as a preference for elegant direct definitions of finite objects cannot, for a thorough understanding, replace the more pedestrian algorithmic approach.

Our second proof of Proposition 8.1.1 illustrates both the constructive and the proof aspect of transfinite induction in the typical intertwined way. We define larger and larger subgraphs $T_\alpha \subseteq G$ inductively. At each step α we prove that T_α is a tree. The definition of T_α will assume and use that subgraphs T_β for all $\beta < \alpha$ have been previously defined, but not only this: it needs to assume that they are nested trees. This fact, therefore, has to be proved along with the recursive definition, always 'just before' it is needed.

Second proof (by transfinite induction).
Let G be a connected graph. We define trees $T_\alpha \subseteq G$ recursively so that

$$T_\beta \subseteq T_\alpha \text{ for all } \beta < \alpha\,. \qquad (*_\alpha)$$

Let T_0 consist of a single vertex. Given a limit ordinal $\alpha > 0$, let $T_\alpha := \bigcup_{\beta<\alpha} T_\beta$. Since the T_β are trees satisfying $(*_\beta)$, our new T_α is also a tree (as in the first proof), and it clearly satisfies $(*_\alpha)$.

Given a successor $\alpha = \beta + 1$, we first check whether $G - T_\beta = \emptyset$. If so, then T_β is a spanning tree and we terminate the recursion. If not, then $G - T_\beta$ has a vertex v_α that sends an edge e_α to a vertex in T_β. Then T_α, obtained from T_β by adding v_α and e_α, is a tree satisfying $(*_\alpha)$.

It remains to check that our recursion does indeed terminate. But if $v_{\beta+1}$ gets defined for all $\beta < \gamma$ then $\beta \mapsto v_{\beta+1}$ is an injective map showing that $|\gamma| \leqslant |G|$. This cannot hold for all ordinals γ; it fails, for example, when γ represents a well-ordering of the power set of $V(G)$. □

Why did these proofs work so smoothly? The reason is that the forbidden substructures, cycles, were finite and therefore could not arise unexpectedly at limit steps. If we wanted to construct a *rayless* spanning tree, on the other hand, one that contains no ray, then the edges of partial finite trees T_β might combine to form a ray in $T_\alpha = \bigcup_{\beta<\alpha} T_\beta$ when α is a limit. And indeed, here lies the challenge in most transfinite constructions: to make the right choices at successor steps to ensure that the structure will also be as desired at limits.

compactness proofs

Our third basic proof technique, somewhat mysteriously referred to as *compactness* (see below for why), offers a formalized way of making the right choices in certain standard cases. These are cases where nothing unexpected happens at limits, but a choice that looks good at the time it is made may lead to a dead end after another *finite* number of steps – unlike the creation of a cycle, which is visible at once.

For example, let G be a graph whose finite subgraphs are all k-colourable. It is natural then to try to construct a k-colouring of G as a limit of k-colourings of its finite subgraphs. Now each finite subgraph will have several k-colourings; will it matter which we choose? Clearly, it will. When $G' \subseteq G''$ are two finite subgraphs and u, v are vertices of G' that receive the same colour in every k-colouring of G'' (and hence also in any k-colouring of G), we must not give them different colours in the colouring we choose for G', even if such a colouring exists. However if we do manage, somehow, to colour the finite subgraphs of G compatibly, we shall automatically have a colouring of all of G.

All compactness proofs deal with situations similar to this. We wish to solve a problem about an infinite structure, and we know how to solve it for all the finite substructures. 'Compactness' enables us to combine these partial solutions to an overall solution if the partial solutions are compatible in the right way.

For countable structures, all this – the choices, the dead ends, the compatibility requirement on the finite solutions, and how they combine to an overall solution – can be made visible in a particularly intuitive way, by a graph:

[8.2.1] [8.2.6] [8.6.1] [8.6.10] [8.7.3] [9.1.3]

Lemma 8.1.2. (König's Infinity Lemma)
Let $V_0, V_1, \dots$ be an infinite sequence of disjoint non-empty finite sets, and let G be a graph on their union. Assume that every vertex v in a set V_n with $n \geqslant 1$ has a neighbour $f(v)$ in V_{n-1}. Then G contains a ray $v_0 v_1 \dots$ with $v_n \in V_n$ for all n.

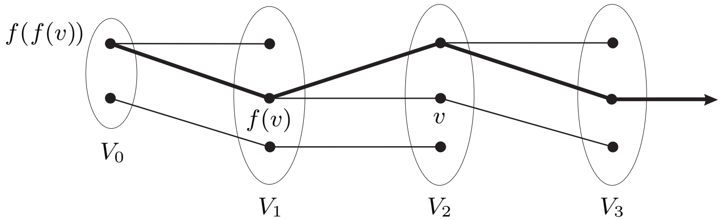

Fig. 8.1.2. König's infinity lemma

Proof. Let $\mathcal{P}$ be the set of all finite paths of the form $v\, f(v)\, f(f(v)) \dots$ ending in V_0. Since V_0 is finite but $\mathcal{P}$ is infinite, infinitely many of the paths in $\mathcal{P}$ end at the same vertex $v_0 \in V_0$. Of these paths, infinitely many also agree on their penultimate vertex $v_1 \in V_1$, because V_1 is finite. Of those paths, infinitely many agree even on their vertex v_2 in V_2 – and so on. Although the set of paths considered decreases from step to step, it is still infinite after any finite number of steps, so v_n gets defined for every $n \in \mathbb{N}$. By definition, each vertex v_n is adjacent to v_{n-1} on one of those paths, so $v_0 v_1 \dots$ is indeed a ray. □

The following 'compactness theorem', the first of its kind in graph theory, answers our question about colourings:

Theorem 8.1.3. (de Bruijn & Erdős, 1951)
Let $G = (V, E)$ be a graph and $k \in \mathbb{N}$. If every finite subgraph of G has chromatic number at most k, then so does G.

First proof (for G countable, by the infinity lemma).
Let $v_0, v_1, \dots$ be an enumeration of V and put $G_n := G[v_0, \dots, v_n]$. Write V_n for the set of all k-colourings of G_n with colours in $\{1, \dots, k\}$. Define a graph on $\bigcup_{n \in \mathbb{N}} V_n$ by inserting all edges cc' such that $c \in V_n$ and $c' \in V_{n-1}$ is the restriction of c to $\{v_0, \dots, v_{n-1}\}$. Let $c_0 c_1 \dots$ be a ray in this graph with $c_n \in V_n$ for all n. Then $c := \bigcup_{n \in \mathbb{N}} c_n$ is a colouring of G with colours in $\{1, \dots, k\}$. □

Applications of the infinity lemma such as this one rely on the fact that a countable graph can be exhausted by a nested sequence of finite subgraphs. Appendix A offers a version of the infinity lemma that works for arbitrary graphs, in which these finite subgraphs need not be sequentially ordered. This general version is still very intuitive and can be used conveniently in many settings, including a proof of Theorem 8.1.3.

The essence of compactness proofs is often encoded more directly, if less graphically, in terms of just sets and functions. The *compactness principle* from Appendix A is such a version that is particularly easy to apply. We illustrate its application by another proof of Theorem 8.1.3:

Second proof (for arbitrary graphs, by the compactness principle). Let $X := V$ and $S := \{1,\dots,k\}$. Let $\mathcal{F}$ be the set of all finite subsets of V. For each $Y \in \mathcal{F}$ let $\mathcal{A}(Y)$ be the set of k-colourings of $G[Y]$. Our proof of Theorem 8.1.3 will be complete once we have found a function $V \to \{1,\dots,k\}$ that induces a k-colouring on every finite subgraph $G[Y]$, as any such function is clearly a k-colouring of G.

By the compactness principle it suffices to show that, given any finite $\mathcal{Y} \subseteq \mathcal{F}$, we can find a function $V \to \{1,\dots,k\}$ that induces a colouring on every $G[Y]$ with $Y \in \mathcal{Y}$. But this is easy: just take a k-colouring of the finite graph $G[\bigcup \mathcal{Y}]$, and extend it arbitrarily to the rest of V. □

Our last proof of Theorem 8.1.3 appeals directly to compactness as defined in topology. Recall that a topological space is *compact* if its closed sets have the 'finite intersection property', which means that the overall intersection $\bigcap \mathcal{A}$ of a set $\mathcal{A}$ of closed sets is non-empty whenever every finite subset of $\mathcal{A}$ has a non-empty intersection. By Tychonoff's theorem of general topology, any product of compact spaces is compact in the usual product topology.

Third proof (for arbitrary graphs, by Tychonoff's theorem). Consider the product space

$$X := \prod_V \{1,\dots,k\} = \{1,\dots,k\}^V$$

of $|V|$ copies of the finite set $\{1,\dots,k\}$ endowed with the discrete topology. By Tychonoff's theorem, this is a compact space. Its basic open sets have the form

$$O_h := \{\, f \in X : f|_U = h \,\},$$

where h is some map from a finite set $U \subseteq V$ to $\{1,\dots,k\}$.

For every finite set $U \subseteq V$, let A_U be the set of all $f \in X$ whose restriction to U is a k-colouring of $G[U]$. These sets A_U are closed (as well as open – why?), and for any finite set $\mathcal{U}$ of finite subsets of V we have $\bigcap_{U \in \mathcal{U}} A_U \neq \emptyset$, because $G[\bigcup \mathcal{U}]$ has a k-colouring. By the finite intersection property of the sets A_U, their overall intersection is non-empty, and every element of this intersection is a k-colouring of G. □

Although our three compactness proofs look formally different, it is instructive to compare them in detail, checking how the requirements in one are reflected in the other (cf. Exercise 18).

As mentioned before, the standard use for compactness proofs is to transfer theorems from finite to infinite graphs, or conversely. This is not always quite as straightforward as above; often, the statement has to be modified a little to make it susceptible to a compactness argument.

As an example – see Exercises 19–30 for more – let us prove the locally finite version of the following famous conjecture. Call a bipartition of the vertex set of a graph *unfriendly* if every vertex has at least as many neighbours in the other class as in its own. Clearly, every finite graph has an unfriendly partition: just take any partition that maximizes the number of edges between the partition classes. At the other extreme, it can be shown by set-theoretic methods that uncountable graphs need not have such partitions. Thus, intriguingly, it is the countable case that has remained unsolved:

Unfriendly Partition Conjecture. *Every countable graph admits an unfriendly partition of its vertex set.*

Proof for locally finite graphs. Let $G = (V, E)$ be an infinite but locally finite graph, and enumerate its vertices as $v_0, v_1, \ldots$. For every $n \in \mathbb{N}$, let $\mathcal{V}_n$ be the set of partitions of $V_n := \{v_0, \ldots, v_n\}$ into two sets U_n and W_n such that every vertex $v \in V_n$ with $N_G(v) \subseteq V_n$ has at least as many neighbours in the other class as in its own. Since the conjecture holds for finite graphs, the sets $\mathcal{V}_n$ are non-empty. For all $n \geqslant 1$, every $(U_n, W_n) \in \mathcal{V}_n$ induces a partition (U_{n-1}, W_{n-1}) of V_{n-1}, which lies in $\mathcal{V}_{n-1}$. By the infinity lemma, there is an infinite sequence of partitions $(U_n, W_n) \in \mathcal{V}_n$, one for every $n \in \mathbb{N}$, such that each is induced by the next. Then $(\bigcup_{n\in\mathbb{N}} U_n, \bigcup_{n\in\mathbb{N}} W_n)$ is an unfriendly partition of G. $\square$

The trick that made this proof possible was to require, for the partitions of V_n, correct positions only of vertices that send no edge out of V_n: this weakening is necessary to ensure that partitions from $\mathcal{V}_n$ induce partitions in $\mathcal{V}_{n-1}$; but since, by local finiteness, every vertex has this property eventually (for large enough n), the weaker assumption suffices to ensure that the limit partition is unfriendly.

Let us complete this section with an introduction to the one important concept of infinite graph theory that has no finite counterpart, the notion of an end. An *end*[5] of a graph G is an equivalence class of rays in G, where two rays are considered equivalent if, for every finite set $S \subseteq V(G)$, both have a tail in the same component of $G - S$. This *end*

5 Not to be confused with the ends, or endvertices, of an edge. In the context of infinite graphs, we use the term 'endvertices' to avoid confusion.

is indeed an equivalence relation: note that, since S is finite, there is exactly one such component for each ray. If two rays are equivalent – and only then – they can be linked by infinitely many disjoint paths: just choose these inductively, taking as S the union of the vertex sets of the first finitely many paths to find the next. The set of ends of G is denoted by $\Omega(G)$, and we write $G = (V, E, \Omega)$ to express that G has vertex, edge and end sets V, E, Ω.

$\Omega(G)$

For example, let us determine the ends of the 2-way infinite ladder shown in Figure 8.1.3. Every ray in this graph contains vertices arbitrarily far to the left or vertices arbitrarily far to the right, but not both. These two types of rays are clearly equivalence classes, so the ladder has exactly two ends. (In Figure 8.1.3 these are shown as two isolated dots – one on the left, the other on the right.)

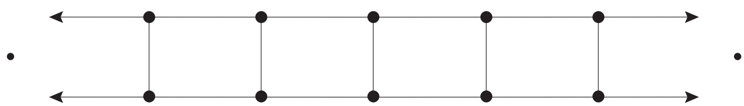

Fig. 8.1.3. The 2-way ladder has two ends

The ends of a tree are particularly simple: two rays in a tree are equivalent if and only if they share a tail, and for every fixed vertex v each end contains exactly one ray starting at v. Even a locally finite tree can have uncountably many ends. The prototype example (see Exercise 39) is the *binary tree* T_2, the rooted tree in which every vertex has exactly two upper neighbours. Often, the vertex set of T_2 is taken to be the set of finite 0–1 sequences (with the empty sequence as the root), as indicated in Figure 8.1.4. The ends of T_2 then correspond bijectively to

binary tree T_2

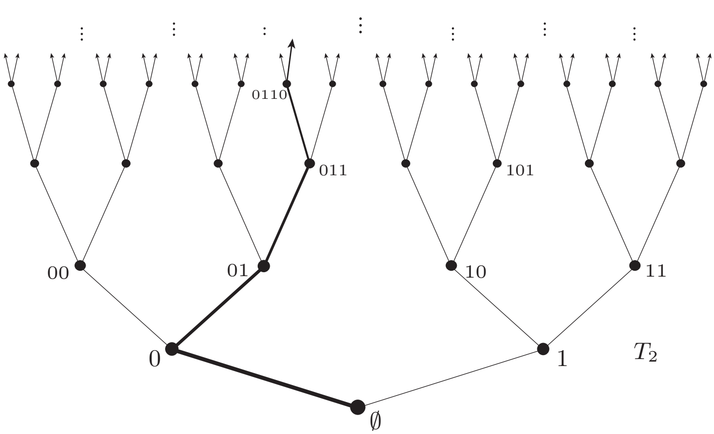

Fig. 8.1.4. The binary tree T_2 has continuum many ends, one for every infinite 0–1 sequence

its rays starting at $\emptyset$, and hence to the infinite 0–1 sequences.

These examples suggest that the ends of a graph can be thought of as 'points at infinity' to which its rays converge. We shall formalize this in Section 8.6, where we define a natural topology on a graph and its ends in which rays will indeed converge to their respective ends.

The maximum number of disjoint rays in an end is the *(combinatorial) vertex-degree* of that end, the maximum number of edge-disjoint rays in it is its *(combinatorial) edge-degree*. These maxima are indeed attained: if an end contains a set of k (edge-) disjoint rays for every integer k, it also contains an infinite set of (edge-) disjoint rays (Exercise 46). Thus, every end has a vertex-degree and an edge-degree in $\mathbb{N} \cup \{\infty\}$. *end degrees*

8.2 Paths, trees, and ends

There are two fundamentally different aspects to the infinity of an infinite connected graph: one of 'length', expressed in the presence of rays, and one of 'width', expressed locally by infinite degrees. The infinity lemma tells us that at least one of these must occur:

Proposition 8.2.1. *Every infinite connected graph has a vertex of infinite degree or contains a ray.*

Proof. Let G be an infinite connected graph with all degrees finite. Let v_0 be a vertex, and for every $n \in \mathbb{N}$ let V_n be the set of vertices at distance n from v_0. Induction on n shows that the sets V_n are finite, and hence that $V_{n+1} \neq \emptyset$ (because G is infinite and connected). Furthermore, the neighbour of a vertex $v \in V_{n+1}$ on any shortest v–v_0 path lies in V_n. By Lemma 8.1.2, G contains a ray. □ (8.1.2)

Often it is useful to have more detailed information on how this ray or vertex of infinite degree lies in G. The following lemma enables us to find it 'close to' any given infinite set of vertices.

Lemma 8.2.2. (Star-Comb Lemma) [8.6.3]
Let U be an infinite set of vertices in a connected graph G. Then G contains either a comb with all teeth in U or a subdivision of an infinite star with all leaves in U.

Proof. As G is connected, it contains a path between two vertices in U. This path is a tree $T \subseteq G$ every edge of which lies on a path in T between two vertices in U. By Zorn's lemma there is a maximal such tree T^*. Since U is infinite and G is connected, T^* is infinite. If T^* has a vertex of infinite degree, it contains the desired subdivided star.

Suppose now that T^* is locally finite. Then T^* contains a ray R (Proposition 8.2.1). Let us construct a sequence $P_1, P_2, \dots$ of disjoint R–U paths in T^*. Having chosen P_i for every $i < n$ for some n, pick $v \in R$ so that vR meets none of those paths P_i. The first edge of vR lies on a path P in T^* between two vertices in U; let us think of P as traversing this edge in the same direction as R, and choose P minimal. Then vP has the form $vRwP$, where $P_n := wP$ is an R–U path. And $P_n \cap P_i = \emptyset$ for all $i < n$, because $P_i \cup Rw \cup P_n$ contains no cycle. □

We shall often apply Lemma 8.2.2 in locally finite graphs, in which case it always yields a comb.

Recall that a rooted tree $T \subseteq G$ is *normal* in G if the endvertices of every T-path in G are comparable in the tree-order of T. If T is a spanning tree, the only T-paths are edges of G that are not edges of T.

Normal spanning trees are perhaps the single most important structural tool in infinite graph theory. As in finite graphs, they exhibit the separation properties of the graph they span.[6] Moreover, their *normal*
normal ray *rays*, those that start at the root, reflect its end structure:

[8.6.8] **Lemma 8.2.3.** *If T is a normal spanning tree of G, then every end of G contains exactly one normal ray of T.*

(1.5.4) *Proof.* Let $\omega \in \Omega(G)$ be given. Apply the star-comb lemma in T with U the vertex set of any ray $R \in \omega$. If the lemma gives a subdivided star with leaves in U and centre z, say, then the finite down-closure $\lceil z \rceil$ of z in T separates infinitely many vertices $u > z$ of U pairwise in G (Lemma 1.5.4). This contradicts our choice of U.

So T contains a comb with teeth on R. Let $R' \subseteq T$ be its spine. Since every ray in T has an increasing tail (Exercise 4), we may assume that R' is a normal ray. Since R' is equivalent to R, it lies in ω.

Conversely, distinct normal rays of T are separated in G by the (finite) down-closure of their greatest common vertex (Lemma 1.5.4), so they cannot belong to the same end of G. □

Not all connected graphs have a normal spanning tree; complete uncountable graphs, for example, have none. (Why not?) The quest to characterize the graphs that have a normal spanning tree is not entirely over, and it has held some surprises.[7] One of the most useful sufficient conditions is that the graph contains no $TK^{\aleph_0}$; see Theorem 12.6.9. For our purposes, the following result suffices:

[8.7.2] **Theorem 8.2.4.** (Jung 1967)
Every countable connected graph has a normal spanning tree.

[6] Lemma 1.5.4 continues to hold for infinite graphs, with the same proof.

[7] One of these is Theorem 8.6.2; for more see the notes.

Proof. The proof follows that of Proposition 1.5.5; we only sketch the differences. Starting with a single vertex, we construct an infinite sequence $T_0 \subseteq T_1 \subseteq \ldots$ of finite normal trees in G, all with the same root, whose union T will be a normal spanning tree. (1.5.5)

To ensure that T spans G, we fix an enumeration $v_0, v_1, \ldots$ of $V(G)$ and see to it that T_n contains v_n. It is clear that T will be a tree (since any cycle in T would lie in some T_n, and every two vertices of T lie in a common T_n and can be linked there), and clearly the tree order of T induces that of the T_n. Finally, T will be normal, because the endvertices of any edge of G that is not an edge of T lie in some T_n: since that T_n is normal, they must be comparable there, and hence in T.

It remains to specify how to construct T_{n+1} from T_n. If $v_{n+1} \in T_n$, put $T_{n+1} := T_n$. If not, let C be the component of $G - T_n$ containing v_{n+1}. Let x be the greatest element of the chain $N(C)$ in T_n, and let T_{n+1} be the union of T_n and an x–v_{n+1} path P with $\mathring{P} \subseteq C$. Then the neighbourhood in T_{n+1} of any new component $C' \subseteq C$ of $G - T_{n+1}$ is a chain in T_{n+1}, so T_{n+1} is again normal. □

One of the most basic problems in an infinite setting that has no finite equivalent is whether or not 'arbitrarily many', in some context, implies 'infinitely many'. Suppose we can find k disjoint rays in some given graph G, for every $k \in \mathbb{N}$; does G also contain an infinite set of disjoint rays?

The answer to the corresponding question for finite paths (of any fixed length) is clearly 'yes', since a finite path P can never get in the way of more than $|P|$ disjoint other paths. A badly chosen ray, however, can meet infinitely many other rays, preventing them from being selected for the same disjoint set. Rather than collecting our disjoint rays greedily, we therefore have to construct them carefully and all simultaneously.

The proof of the following theorem is a nice example of a construction in an infinite sequence of steps, where the final object emerges only at the limit step. Each of the steps in the sequence will involve a non-trivial application of Menger's theorem (3.3.1).

Theorem 8.2.5. (Halin 1965)

(i) *If an infinite graph G contains k disjoint rays for every $k \in \mathbb{N}$, then G contains infinitely many disjoint rays.*

(ii) *If an infinite graph G contains k edge-disjoint rays for every $k \in \mathbb{N}$, then G contains infinitely many edge-disjoint rays.*

Proof. (i) We construct our infinite system of disjoint rays inductively in ω steps. After step n, we shall have found n disjoint rays $R_1^n, \ldots, R_n^n$ and chosen initial segments $R_i^n x_i^n$ of these rays. In step $n+1$ we choose the rays $R_1^{n+1}, \ldots, R_{n+1}^{n+1}$ so as to extend these initial segments, i.e. so that $R_i^n x_i^n$ is a proper initial segment of $R_i^{n+1} x_i^{n+1}$, for $i = 1, \ldots, n$. (3.3.1)

Then, clearly, the graphs $R_i^* := \bigcup_{n \in \mathbb{N}} R_i^n x_i^n$ will form an infinite family $(R_i^*)_{i \in \mathbb{N}}$ of disjoint rays in G.

For $n = 0$ the empty set of rays is as required. So let us assume
n that $R_1^n, \dots, R_n^n$ have been chosen, and describe step $n+1$. For sim-
R_i, x_i plicity, let us abbreviate $R_i^n =: R_i$ and $x_i^n =: x_i$. Let $\mathcal{R}$ be any set of $|R_1x_1 \cup \ldots \cup R_nx_n| + n^2 + 1$ disjoint rays (which exists by assumption), and immediately delete those rays from $\mathcal{R}$ that meet any of the paths $R_1x_1, \dots, R_nx_n$; then $\mathcal{R}$ still contains at least $n^2 + 1$ rays.

We begin by repeating the following step as often as possible. If there exists an $i \in \{1, \dots, n\}$ such that R_i^{n+1} has not yet been defined and $\mathring{x}_iR_i$ meets at most n of the rays currently in $\mathcal{R}$, we delete those rays from $\mathcal{R}$, put $R_i^{n+1} := R_i$, and choose as x_i^{n+1} the successor of x_i
I on R_i. Having performed this step as often as possible, we let I denote the set of those $i \in \{1, \dots, n\}$ for which R_i^{n+1} is still undefined, and put
m $|I| =: m$. Then $\mathcal{R}$ still contains at least $n^2 + 1 - (n - m)n \geqslant m^2 + 1$ rays. Every R_i with $i \in I$ meets more than $n \geqslant m$ of the rays in $\mathcal{R}$; let
Z z_i be its first vertex on the mth ray it meets. Then $Z := \bigcup_{i \in I} x_iR_iz_i$ meets at most m^2 of the rays in $\mathcal{R}$; we delete all the other rays from $\mathcal{R}$, choosing one of them as R_{n+1}^{n+1} (with x_{n+1}^{n+1} arbitrary).

On each remaining ray $R \in \mathcal{R}$ we now pick a vertex $y = y(R)$ after its last vertex in Z, and put $Y := \{\, y(R) \mid R \in \mathcal{R}\,\}$. Let H be the union of Z and all the paths Ry $(R \in \mathcal{R})$. Then $X := \{\, x_i \mid i \in I\,\}$ cannot be separated from Y in H by fewer than m vertices, because these would miss both one of the m rays R_i with $i \in I$ and one of the m rays in $\mathcal{R}$ that meet $x_iR_iz_i$ for this i. So by Menger's theorem (3.3.1) there
y_i are m disjoint X–Y paths $P_i = x_i \dots y_i$ $(i \in I)$ in H. For each $i \in I$ let R_i' denote the ray from $\mathcal{R}$ that contains y_i, choose as R_i^{n+1} the ray $R_ix_iP_iy_iR_i'$, and put $x_i^{n+1} := y_i$.

(ii) The proof is similar; see Exercise 45 and its hint. □

Does Theorem 8.2.5 generalize to other graphs than rays? Let us call a graph H *ubiquitous* with respect to a relation $\leqslant$ between graphs (such as the subgraph relation $\subseteq$, or the minor relation $\preccurlyeq$) if $nH \leqslant G$
nH for all $n \in \mathbb{N}$ implies $\aleph_0 H \leqslant G$, where nH denotes the disjoint union of n copies of H. Ubiquity appears to be closely related to questions of well-quasi-ordering as discussed in Chapter 12. Non-ubiquitous graphs exist for all the standard graph orderings; see Exercise 49 for an example of a locally finite graph that is not ubiquitous under the subgraph relation.

Ubiquity conjecture. (Andreae 2002)
Every locally finite connected graph is ubiquitous with respect to the minor relation.

A proof of the ubiquity conjecture for trees is indicated in Exercise 50.

Just as in Theorem 8.2.5 one can show that an end contains infinitely many disjoint rays as soon as the number of disjoint rays in it is not finitely bounded, and similarly for edge-disjoint rays (Exercise 46). Hence, the maxima in our earlier definitions of the vertex- and edge-degrees of an end exist as claimed. Ends of infinite vertex-degree are called *thick*; ends of finite vertex-degree are *thin*. *thick/thin*

The $\mathbb{N} \times \mathbb{N}$ *grid*, for example, the graph on $\mathbb{N}^2$ in which two vertices (n, m) and (n', m') are adjacent if and only if $|n - n'| + |m - m'| = 1$, has only one end, which is thick. In fact, the $\mathbb{N} \times \mathbb{N}$ grid is a kind of prototype for thick ends: every graph with a thick end contains it as a minor. This is another classical result of Halin, which we prove in the remainder of this section. *grid*

For technical reasons, we shall prove Halin's theorem for hexagonal rather than square grids. These may seem a little unwieldy at first, but have the advantage that they can be found as topological rather than ordinary minors (Proposition 1.7.3), which makes them much easier to handle. We shall define the hexagonal grid H^∞ so that it is a subgraph of the $\mathbb{N} \times \mathbb{N}$ grid, and it will be easy to see that, conversely, the $\mathbb{N} \times \mathbb{N}$ grid is a minor of H^∞. (See also Exercise 71, Ch. 12.)

To define our standard copy of the *hexagonal quarter grid* H^∞, we delete from the $\mathbb{N} \times \mathbb{N}$ grid H the vertex $(0,0)$, the vertices (n, m) with $n > m$, and all edges $(n, m)(n+1, m)$ such that n and m have equal parity (Fig. 8.2.1). Thus, H^∞ consists of the *vertical rays* H^∞

$$U_0 := H\,[\,\{\,(0, m) \mid 1 \leqslant m\,\}\,]$$
$$U_n := H\,[\,\{\,(n, m) \mid n \leqslant m\,\}\,] \quad (n \geqslant 1)$$

U_n

and between these a set of *horizontal edges*,

$$E := \{\,(n, m)(n+1, m) \mid n \not\equiv m \pmod 2\,\}\,.$$

To enumerate these edges, as $e_1, e_2, \dots$ say, we order them colexicographically: the edge $(n, m)(n+1, m)$ precedes the edge $(n', m')(n'+1, m')$ if $m < m'$, or if $m = m'$ and $n < n'$ (Fig. 8.2.1). $e_1, e_2, \dots$

Theorem 8.2.6. (Halin 1965)
Whenever a graph contains a thick end, it has a TH^∞ subgraph whose rays belong to that end.

Proof. Given two infinite sets $\mathcal{P}, \mathcal{P}'$ of finite or infinite paths, let us write $\mathcal{P} \geqslant \mathcal{P}'$ if $\mathcal{P}'$ consists of final segments of paths in $\mathcal{P}$. (Thus, if $\mathcal{P}$ is a set of rays, then so is $\mathcal{P}'$.) (8.1.2) $\leqslant$

Let G be any graph with a thick end ω. Our task is to find disjoint rays in ω that can serve as 'vertical' (subdivided) rays U_n for our desired grid, and to link these up by suitable disjoint 'horizontal' paths. We ω

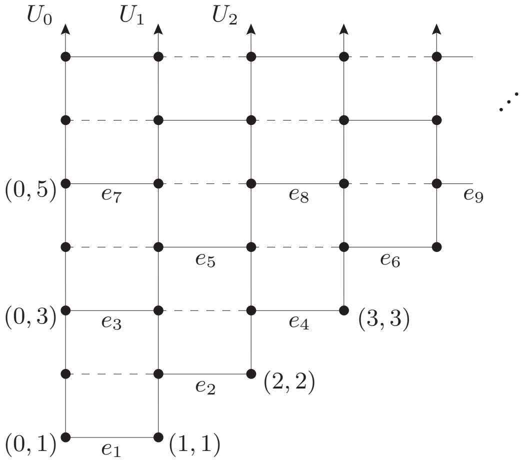

Fig. 8.2.1. The hexagonal quarter grid H^∞.

begin by constructing a sequence $R_0, R_1, \ldots$ of rays (of which we shall later choose some tails R'_n as 'vertical rays'), together with path systems $\mathcal{P}_n$ between the R_n and suitable $R_{p(n)}$ with $p(n) < n$ (from which we shall later choose the 'horizontal paths'). We shall aim to find the R_n in 'supply sets' $\mathcal{R}_0 \geqslant \mathcal{R}_1 \geqslant \ldots$ of unused rays. After the nth construction step we shall have constructed the subgraph $H_n := \bigcup_{i=0}^n \big(R_i \cup \bigcup \mathcal{P}_i\big)$.

H_n

We start with any infinite set $\mathcal{R}$ of disjoint rays in ω; this exists by our assumption that ω is a thick end. Pick $R_0 \in \mathcal{R}$, put $\mathcal{P}_0 := \emptyset$, and let $\mathcal{R}_0 := \mathcal{R} \setminus \{R_0\}$. At step $n \geqslant 1$ of the construction we shall choose the following:

R_0

$\mathcal{R}_0$

(1) a ray $R_n \in \omega$ disjoint from H_{n-1};

(2) an integer $p(n) < n$;

(3) an infinite set $\mathcal{P}_n$ of disjoint R_n–$R_{p(n)}$ paths avoiding every other R_i;

(4) an infinite set $\mathcal{R}_n \leqslant \mathcal{R}_{n-1}$ of disjoint rays in $G - H_n$.

Let $n \geqslant 1$ be given. As a first candidate for R_n consider any ray $R \in \mathcal{R}_{n-1}$. By (4) for smaller values of n, we have both $R \in \omega$ (since $\mathcal{R}_{n-1} \leqslant \ldots \leqslant \mathcal{R}_0 \subseteq \mathcal{R}$) and $R \cap H_{n-1} = \emptyset$, as required for R_n in (1).

R

Next, let us try to find a set $\mathcal{P}_n$ and $p(n)$ to go with R as R_n. Since H_{n-1} contains $R_0 \in \omega$, there exists an infinite set $\mathcal{P}$ of disjoint R–H_{n-1} paths in G. If $\mathcal{P}$ has an infinite subset of paths all ending on the same R_i ($i < n$), we delete all other paths from $\mathcal{P}$. If not, then $\mathcal{P}$ has an infinite subset of paths all ending at an inner vertex of a path in $\mathcal{P}_i$, for the same $i < n$. We extend them back along this path of $\mathcal{P}_i$ until they hit R_i, and delete all other paths from $\mathcal{P}$. In both cases we put $p(n) := i$, satisfying (2), and have found an infinite set $\mathcal{P}$ of disjoint R–$R_{p(n)}$ paths that avoid all the other R_i $(i < n)$. $p(n)$, $\mathcal{P}$

For later use we record:

Every path in $\mathcal{P}$ consists of an R–H_{n-1} path followed by a (possibly trivial) path in H_{n-1}. (∗)

What can still prevent us from choosing R as R_n and $\mathcal{P}$ as $\mathcal{P}_n$ is condition (4): if $\mathcal{P}$ meets all but finitely many rays in $\mathcal{R}_{n-1}$ infinitely, we cannot find an infinite set $\mathcal{R}_n \leqslant \mathcal{R}_{n-1}$ of rays avoiding $\mathcal{P}$.

If this happens, the only option we have is to make a virtue of necessity and use the rich intersection of $\mathcal{P}$ and $\mathcal{R}_{n-1}$ for an altogether different construction of R_n, $\mathcal{P}_n$ and $\mathcal{R}_n$. This will require some work. But we may assume the following, as the result of our effort so far:

Whenever $R' \in \mathcal{R}_{n-1}$ and $\mathcal{P}' \leqslant \mathcal{P}$ is an infinite set of R'–$R_{p(n)}$ paths, there is a ray $R'' \neq R'$ in $\mathcal{R}_{n-1}$ that meets $\mathcal{P}'$ infinitely. (∗∗)

For if (∗∗) failed, we could choose R' as R_n and $\mathcal{P}'$ as $\mathcal{P}_n$, and for $\mathcal{R}_n$ select from every ray $R'' \neq R'$ in $\mathcal{R}_{n-1}$ a tail avoiding $\mathcal{P}'$. This would satisfy conditions (1)–(4) for n.

Consider the paths in $\mathcal{P}$ as linearly ordered by the natural order of their starting vertices on R. This induces an ordering on every $\mathcal{P}' \leqslant \mathcal{P}$. If $\mathcal{P}'$ is a set of R'–$R_{p(n)}$ paths for some ray R', we shall call this ordering of $\mathcal{P}'$ *compatible* with R' if the ordering it induces on the first vertices of its paths coincides with the natural ordering of those vertices on R'.

Starting with $R =: R^0_{n-1}$ and $\mathcal{P} =: \mathcal{P}^0$, let us construct sequences $R^0_{n-1}, R^1_{n-1}, \ldots$ and $\mathcal{P}^0 \geqslant \mathcal{P}^1 \geqslant \ldots$ such that every R^k_{n-1} is a tail of a ray in $\mathcal{R}_{n-1}$ and each $\mathcal{P}^k$ is an infinite set of R^k_{n-1}–$R_{p(n)}$ paths whose ordering is compatible with R^k_{n-1}. The first path of $\mathcal{P}^k$ in this ordering will be denoted by P_k, its starting vertex on R^k_{n-1} by v_k, and the path in $\mathcal{P}^{k-1}$ containing P_k (if $k \geqslant 1$) by P^-_k (Fig. 8.2.2). To define R^k_{n-1} and $\mathcal{P}^k$ for $k \geqslant 1$, we use (∗∗) with $R' \supseteq R^{k-1}_{n-1}$ and $\mathcal{P}' = \mathcal{P}^{k-1}$ to find in $\mathcal{R}_{n-1}$ a ray $R'' \not\supseteq R^{k-1}_{n-1}$ that meets $\mathcal{P}^{k-1}$ infinitely; let R^k_{n-1} be a tail of R'' that avoids the finitely many paths in $\mathcal{P}$ containing $P_0, \ldots, P_{k-1}$. Let P^-_k be a path in $\mathcal{P}^{k-1}$ that meets R^k_{n-1} and let v be its 'highest' vertex on R^k_{n-1}, that is, the last vertex of R^k_{n-1} in $V(P^-_k)$. Replacing R^k_{n-1} with its tail vR^k_{n-1}, we can arrange that P^-_k has only the vertex v on R^k_{n-1}. Then $P_k := vP^-_k$ is an R^k_{n-1}–$R_{p(n)}$ path starting at $v_k = v$. We may now select an infinite set $\mathcal{P}^k \leqslant \mathcal{P}^{k-1}$ of R^k_{n-1}–$R_{p(n)}$ paths compatible with R^k_{n-1} and containing P_k as its first path.

R^0_{n-1}, $\mathcal{P}^0$

P_k, v_k

P^-_k

R^k_{n-1}

$\mathcal{P}^k$

Note that P_k cannot be a subpath of any P_i with $i < k$, since P_k contains $v_k \in R^k_{n-1}$ but $R^k_{n-1} \cap P_i = \emptyset$. As $\mathcal{P}^k \leqslant \mathcal{P}^i$, this means that P_k and P_i are subpaths of disjoint paths in $\mathcal{P}$. Similarly, for $i < k$ the rays R^k_{n-1} and R^i_{n-1} cannot be tails of the same ray in $\mathcal{R}_{n-1}$, and are there-

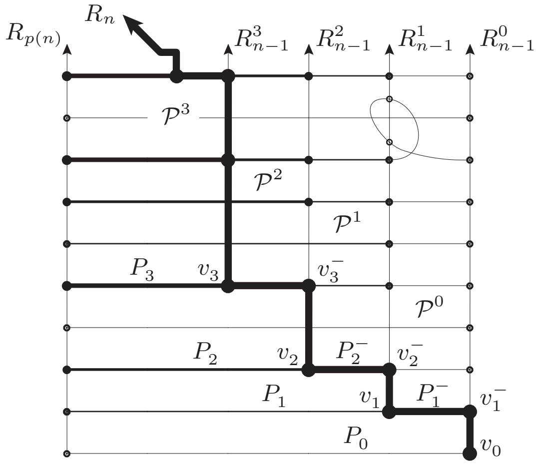

Fig. 8.2.2. Constructing R_n from condition $(**)$

fore tails of disjoint rays in $\mathcal{R}_{n-1}$, because $\bigcup \mathcal{P}^k$ meets R^k_{n-1} infinitely but avoids R^i_{n-1}. (Indeed, as R^k_{n-1} is disjoint from R^{k-1}_{n-1} by definition, the paths in $\mathcal{P}^k$ are proper final segments of paths in $\mathcal{P}^{k-1} \leqslant \mathcal{P}^i$, and the paths in $\mathcal{P}_i$ meet R^i_{n-1} only in their first vertex.)

v_k^- For each k, let v^-_{k+1} denote the starting vertex of P^-_{k+1} on R^k_{n-1}, and put $R^k_n := \mathring{v}^-_{k+1} R^k_{n-1}$. Then let

$$\begin{aligned} R_n &:= v_0 R^0_{n-1} v_1^- P_1^- v_1 R^1_{n-1} v_2^- P_2^- v_2 R^2_{n-1} \dots \\ \mathcal{P}_n &:= \{P_0, P_1, P_2, \dots\} \\ \mathcal{R}_n &:= \{\, R^k_n \mid k \in \mathbb{N}\,\}. \end{aligned}$$

To verify that R_n is indeed a ray, we have to check that the various path segments it is composed of meet only at the vertices at which they are concatenated. We have already noted that for different k the paths P^-_k are final segments of disjoint paths in $\mathcal{P}$ and the rays R^k_{n-1} are tails of disjoint rays in $\mathcal{R}_{n-1}$. Moreover, each R^k_{n-1} by definition avoids the paths in $\mathcal{P}$ containing $P_0, \dots, P_{k-1}$. And each P^-_{k+1} avoids the rays R^i_{n-1} with $i < k$, because P^-_{k+1} lies in $\mathcal{P}^k$, whose paths are proper final segments of R^i_{n-1}–$R_{p(n)}$ paths in $\mathcal{P}^i$. Hence all that remains to check is that, for each k, the segment $v_k R^k_{n-1} v^-_{k+1}$ of R_n meets the previous segment $v^-_k P^-_k v_k$ only in v_k and the next segment $v^-_{k+1} P^-_{k+1} v_{k+1}$ only in v^-_{k+1}. The first of these assertions follows from the definition of R^k_{n-1}, the second by the choice of P^-_{k+1}.

For the same reasons, each P_k meets R_n only in v_k (so $\mathcal{P}_n$ is indeed a set of R_n–$R_{p(n)}$ paths), and the paths in $\mathcal{R}_n$ meet neither R_n nor the paths in $\mathcal{P}_n$. Therefore $\mathcal{R}_n$ satisfies (4), and $\mathcal{P}_n$ satisfies (3); recall that the paths in $\mathcal{P} \geqslant \mathcal{P}_n$ avoid R_i for every $i < n$ other than $p(n)$.

It remains to verify (1). We have $R_n \in \omega$, since $\mathcal{P}_n$ joins R_n disjointly to $R_{p(n)}$, and $R_{p(n)} \in \omega$ by (1). To show that $R_n \cap H_{n-1} = \emptyset$, recall that the 'vertical' segments of R_n lie in rays from $\mathcal{R}_{n-1}$, so by (4) they do not meet H_{n-1}. And a 'horizontal' segment of R_n can meet H_{n-1} only if it does so also in its last vertex v_k, by $(*)$. But v_k also lies on a vertical segment, and hence not in H_{n-1}. Thus, $R_n \cap H_{n-1} = \emptyset$ as desired.

Let us now use our rays R_n and path systems $\mathcal{P}_n$ to construct the desired grid. In the tree on $\mathbb{N}$ defined by joining each n to $p(n)$, apply the infinity lemma (8.1.2) to the distance classes from 0 to find a ray $n_0 n_1 \dots$ with vertices $n_0 < n_1 < \dots$ or, if one of these classes is infinite, a vertex n_0 with infinitely many neighbours $n_1, n_2, \dots$ greater than n_0. We treat these two cases in turn, assuming for notational simplicity that $n_i = i$ for all i. (In other words, we discard any R_n with $n \notin \{n_0, n_1, \dots\}$.)

In the first case, each $\mathcal{P}_n$ is an infinite set of disjoint R_n–R_{n-1} paths. Our aim is to choose tails R'_n of our rays R_n that will correspond to the vertical rays $U_n \subseteq H^\infty$, and paths $S_1, S_2, \dots$ between the R'_n that will correspond to the horizontal edges $e_1, e_2, \dots$ of H^∞. We shall find the paths $S_1, S_2, \dots$ inductively, choosing the R'_n as needed as we go along (but also in the order of increasing n, starting with $R'_0 := R_0$). At every step of the construction, we shall have selected only finitely many S_k and only finitely many R'_n.

Let k and n be minimal such that S_k and R'_n are still undefined. We describe how to choose S_k, and R'_n if the definition of S_k requires it. Let i be such that e_k joins U_{i-1} to U_i in H^∞. If $i = n$, let R'_n be a tail of R_n that avoids the finitely many paths $S_1, \dots, S_{k-1}$; otherwise, R'_i has already been defined, and so has R'_{i-1}. Now choose $S_k \in \mathcal{P}_i$ 'high enough' between R'_{i-1} and R'_i to mirror the position of e_k in H^∞, and to avoid $S_1 \cup \dots \cup S_{k-1}$. Then S_k will also avoid every other R'_j already defined: by (3) for i if $j < i$, and by (1) for j if $j > i$. Since every R'_n is chosen so as to avoid all previously defined S_k, and every S_k avoids all previously defined R'_j (except R'_{i-1} and R'_i), the R'_n and S_k are pairwise disjoint for all $n, k \in \mathbb{N}$, except for the required incidences. Our construction thus yields the desired subdivision of H^∞.

In the second case, every $\mathcal{P}_n$ is a set of disjoint R_n–R_0 paths. We now use only the R_n with $n \geqslant 1$ for vertical rays of H^∞, because R_0 will be needed for the horizontal paths. More precisely, we choose rays $R'_n \subseteq R_n$ for $n \geqslant 1$, and paths S_k between them, inductively as before, except that S_k now consists of three parts: an initial segment from $\mathcal{P}_{i-1}$, followed by a middle segment on R_0, and a final segment from $\mathcal{P}_i$. Such S_k can again be found, since at every stage of the construction only a finite part of R_0 has been used. $\square$

8.3 Homogeneous and universal graphs

Unlike finite graphs, infinite graphs offer the possibility to represent an entire graph property $\mathcal{P}$ by just one specimen, a single graph that contains all the graphs in $\mathcal{P}$ up to some fixed cardinality. Such graphs are called 'universal' for this property.

More precisely, if $\leqslant$ is a graph relation (such as the minor, topological minor, subgraph, or induced subgraph relation up to isomorphism),

universal

we call a countable graph G^* *universal* in $\mathcal{P}$ (for $\leqslant$) if $G^* \in \mathcal{P}$ and $G \leqslant G^*$ for every countable graph $G \in \mathcal{P}$.

Is there a graph that is universal in the class of *all* countable graphs? Suppose a graph R has the following property:

Whenever U and W are disjoint finite sets of vertices in R, there exists a vertex $v \in R - U - W$ that is adjacent in R to all the vertices in U but to none in W. $(*)$

Then R is universal even for the strongest of all graph relations, the induced subgraph relation. Indeed, in order to embed a given countable graph G in R we just map its vertices $v_1, v_2, \ldots$ to R inductively, making sure that v_n gets mapped to a vertex $v \in R$ adjacent to the images of all the neighbours of v_n in $G[v_1, \ldots, v_n]$ but not adjacent to the image of any non-neighbour of v_n in $G[v_1, \ldots, v_n]$. Clearly, this map is an isomorphism between G and the subgraph of R induced by its image.

[11.3.5]

R

Theorem 8.3.1. (Erdős & Rényi 1963)
There exists a unique countable graph R with property $(*)$.

Proof. To prove existence, we construct a graph R with property $(*)$ inductively. Let $R_0 := K^1$. For all $n \in \mathbb{N}$, let R_{n+1} be obtained from R_n by adding for every set $U \subseteq V(R_n)$ a new vertex v joined to all the vertices in U but to none outside U. (In particular, the new vertices form an independent set in R_{n+1}.) Clearly $R := \bigcup_{n \in \mathbb{N}} R_n$ has property $(*)$.

To prove uniqueness, let $R = (V, E)$ and $R' = (V', E')$ be two graphs with property $(*)$, each given with a fixed vertex enumeration. We construct a bijection $\varphi\colon V \to V'$ in an infinite sequence of steps, defining $\varphi(v)$ for one new vertex $v \in V$ at each step.

At every odd step we look at the first vertex v in the enumeration of V for which $\varphi(v)$ has not yet been defined. Let U be the set of those of its neighbours u in R for which $\varphi(u)$ has already been defined. This is a finite set. Using $(*)$ for R', find a vertex $v' \in V'$ outside the image of φ (which is a finite set), so that v' is adjacent in R' to all the vertices in $\varphi(U)$ but to no other vertex in the image of φ. Put $\varphi(v) := v'$.

At even steps in the definition process we do the same thing with the roles of R and R' interchanged: we look at the first vertex v' in the enumeration of V' that does not yet lie in the image of φ, and set

$\varphi(v) = v'$ for a new vertex v that matches the adjacencies and non-adjacencies of v' among the vertices for which φ (resp. φ^{-1}) has already been defined.

By our minimum choices of v and v', the bijection gets defined on all of V and all of V', and it is clearly an isomorphism. □

The graph R in Theorem 8.3.1 is usually called the *Rado graph*, named after Richard Rado who gave one of its earliest explicit definitions. The method of constructing a bijection in alternating steps, as in the uniqueness part of the proof, is known as the *back-and-forth* technique. *Rado graph*

The Rado graph R is unique in another rather fascinating respect. We shall hear more about this in Chapter 11.3, but in a nutshell it is the following. If we generate a countably infinite *random* graph by admitting its pairs of vertices as edges independently with some fixed positive probability $p \in (0,1)$, then with probability 1 the resulting graph has property $(*)$, and is hence isomorphic to R! In the context of infinite graphs, the Rado graph is therefore also called *the* (countably infinite) *random graph*. *'the' random graph*

As one would expect of a random graph, the Rado graph shows a high degree of uniformity. One aspect of this is its resilience against small changes: the deletion of finitely many vertices or edges, and similar local changes, leave it 'unchanged' and result in just another copy of R (Exercise 54).

The following rather extreme aspect of uniformity, however, is still surprising: no matter how we partition the vertex set of R into two parts, at least one of the parts will induce another isomorphic copy of R. Trivial examples aside, the Rado graph is the only countable graph with this property, and hence unique in yet another respect:

Proposition 8.3.2. *The Rado graph is the unique countable graph G other than $K^{\aleph_0}$ and $\overline{K^{\aleph_0}}$ such that, no matter how $V(G)$ is partitioned into two parts, one of the parts induces an isomorphic copy of G.*

Proof. We first show that the Rado graph R has the partition property. Let $\{V_1, V_2\}$ be a partition of $V(R)$. If $(*)$ fails in both $R[V_1]$ and $R[V_2]$, say for sets U_1, W_1 and U_2, W_2, respectively, then $(*)$ fails for $U = U_1 \cup U_2$ and $W = W_1 \cup W_2$ in R, a contradiction.

To show uniqueness, let $G = (V, E)$ be a countable graph with the partition property. Let V_1 be its set of isolated vertices, and V_2 the rest. If $V_1 \neq \emptyset$ then $G \not\cong G[V_2]$, since G has isolated vertices but $G[V_2]$ does not. Hence $G = G[V_1] \cong \overline{K^{\aleph_0}}$. Similarly, if G has a vertex adjacent to all other vertices, then $G = K^{\aleph_0}$.

Assume now that G has no isolated vertex and no vertex joined to all other vertices. If G is not the Rado graph then there are sets U, W for which $(*)$ fails in G; choose these with $|U \cup W|$ minimum.

Assume first that $U \neq \emptyset$, and pick $u \in U$. Let V_1 consist of u and all vertices outside $U \cup W$ that are not adjacent to u, and let V_2 contain the remaining vertices. As u is isolated in $G[V_1]$, we have $G \not\cong G[V_1]$ and hence $G \cong G[V_2]$. By the minimality of $|U \cup W|$, there is a vertex $v \in G[V_2] - U - W$ that is adjacent to every vertex in $U \setminus \{u\}$ and to none in W. But v is also adjacent to u, because it lies in V_2. So U, W and v satisfy $(*)$ for G, contrary to assumption.

Finally, assume that $U = \emptyset$. Then $W \neq \emptyset$. Pick $w \in W$, and consider the partition $\{V_1, V_2\}$ of V where V_1 consists of w and all its neighbours outside W. As before, $G \not\cong G[V_1]$ and hence $G \cong G[V_2]$. Therefore U and $W \setminus \{w\}$ satisfy $(*)$ in $G[V_2]$, with $v \in V_2 \setminus W$ say, and then U, W, v satisfy $(*)$ in G. □

Another indication of the high degree of uniformity in the structure of the Rado graph is its large automorphism group. For example, R is easily seen to be *vertex-transitive*: given any two vertices x and y, there is an automorphism of R mapping x to y.

homogeneous

In fact, much more is true: using the back-and-forth technique, one can easily show that the Rado graph is *homogeneous*: every isomorphism between two finite induced subgraphs can be extended to an automorphism of the entire graph (Exercise 56).

Which other countable graphs are homogeneous? The complete graph $K^{\aleph_0}$ and its complement are again obvious examples. Moreover, for every integer $r \geqslant 3$ there is a homogeneous K^r-free graph R^r, constructed as follows. Let $R^r_0 := K^1$, and let R^r_{n+1} be obtained from R^r_n by joining, for every induced subgraph $H \not\supseteq K^{r-1}$ of R^r_n, a new vertex

R^r

v_H to every vertex in H. Then let $R^r := \bigcup_{n \in \mathbb{N}} R^r_n$. Clearly, as the new vertices v_H of R^r_{n+1} are independent, there is no K^r in R^r_{n+1} if there was none in R^r_n, so $R^r \not\supseteq K^r$ by induction on n. Just like the Rado graph, R^r is clearly universal among the K^r-free countable graphs, and by the back-and-forth argument from the proof of Theorem 8.3.1 it is easily seen to be homogeneous.

By the following deep theorem of Lachlan and Woodrow, the countable homogeneous graphs we have seen so far are essentially all:

Theorem 8.3.3. (Lachlan & Woodrow 1980)
Every countably infinite homogeneous graph is one of the following:

- *a disjoint union of complete graphs of the same order, or the complement of such a graph;*
- *the graph R^r or its complement, for some $r \geqslant 3$;*
- *the Rado graph R.*

To conclude this section, let us return to our original problem: for which graph properties is there a graph that is universal with this property? Most investigations into this problem have addressed it from a more general model-theoretic point of view, and have therefore been based on the strongest of all graph relations, the induced subgraph relation. As a consequence, most of these results are negative; see the notes.

From a graph-theoretic point of view, it seems more promising to look instead for universal graphs for the weaker subgraph relation, or even the topological minor or minor relation. For example, while there is no universal planar graph for subgraphs or induced subgraphs, there is one for minors:

Theorem 8.3.4. (Diestel & Kühn 1999)
There exists a universal planar graph for the minor relation.

This remained the only result about universal graphs for the minor-relation for over 20 years, until Georgakopoulos found $\preccurlyeq$-universal graphs in $\mathrm{Forb}_{\preccurlyeq}(X) = \{\, G \mid X \not\preccurlyeq G \,\}$ when X is K^5, $K_{3,3}$ or $K^{\aleph_0}$.

8.4 Connectivity and matching

In this section we look at infinite versions of Menger's theorem and of the matching theorems from Chapter 2. This area of infinite graph theory is one of its best developed fields, with several deep results. One of these, however, stands out among the rest: a version of Menger's theorem that had been conjectured by Erdős decades ago, and was proved only fairly recently by Aharoni and Berger. The techniques developed for its proof inspired, over the years, much of the theory in this area.

Before we turn to this result, however, let us take a brief look at edge-connectivity. Recall from Section 8.1 that in an infinitely edge-connected countable graph we can easily find infinitely many edge-disjoint spanning trees. Can we still find such trees when the graph is uncountable? We can, but this is not quite as easy to prove (Exercise 62).

The following deep theorem of Laviolette reduces the above problem to its countable case – as it does for many other problems involving edge-connectivity. Let $\mathcal{H}$ be a set of countable graphs forming an edge-decomposition of an arbitrary graph G. Call this decomposition *bond-faithful* if every countable bond of G is contained in some $H \in \mathcal{H}$ and every finite bond of any $H \in \mathcal{H}$ is a bond also of G. Note that the finite bonds of G will be bonds of the $H \in \mathcal{H}$ that contain them. (Why?)

Theorem 8.4.1. (Laviolette 2005)
Every graph has a bond-faithful decomposition into countable graphs.

We shall not be able to prove Laviolette's theorem here. But let us illustrate its power in reducing problems to their countable case by deducing an early classic from the theory of infinite graphs.

even/odd cut

Let us call a cut F in a graph G *even* or *odd* if $|F|$ is finite and even or odd, respectively. Since cycles meet cuts in an even number of edges, graphs with an odd cut cannot have an edge-decomposition into cycles. Suprisingly, all other graphs do:

Theorem 8.4.2. (Nash-Williams 1960)
Every graph with no odd cut has an edge-decomposition into cycles.

Proof. By Theorem 8.4.1 we may assume that our graph G is countable; let us enumerate its edges. We shall find the desired cycles in ω steps. At each step we shall assume inductively that every cut is infinite or even, find a cycle, and delete its edges.

To find that cycle, consider the first remaining edge $e = xy$ in our enumeration. By our inductive assumption e is not a bridge, so there is an x–y path not using e. Together with e this path forms a cycle, whose edges we delete to complete this step. This deletion keeps every finite cut even and every infinite cut infinite.

After at most ω steps no edges remain, so the cycles we found form an edge-decomposition of G. $\square$

Let us now turn to the Aharoni-Berger theorem: the infinite version of Menger's theorem originally conjectured by Erdős. We shall prove this theorem for countable graphs, which will take up most of this section. Although the countable case is much easier, the techniques it requires already give a good impression of the general proof. We then wind up with an overview of infinite matching theorems and a conjecture conceived in the same spirit.

Recall that Menger's theorem, in its simplest form, says that if A and B are sets of vertices in a finite graph G, not necessarily disjoint, and if $k = k(G, A, B)$ is the minimum number of vertices separating A from B in G, then G contains k disjoint A–B paths. (Clearly, it cannot contain more.) The same holds, and is easily deduced from the finite case, when G is infinite but k is still finite:

Proposition 8.4.3. *Let G be any graph, $k \in \mathbb{N}$, and let A, B be two sets of vertices in G that can be separated by k but no fewer than k vertices. Then G contains k disjoint A–B paths.*

(3.3.1) *Proof.* By assumption, every set of disjoint A–B paths has cardinality at most k. Choose one, $\mathcal{P}$ say, of maximum cardinality. Suppose $|\mathcal{P}| < k$. Then no set X consisting of one vertex from each path in $\mathcal{P}$ separates A from B. For each X, let P_X be an A–B path avoiding X. Let H be the union of $\bigcup \mathcal{P}$ with all these paths P_X. This is a finite graph in which no set of $|\mathcal{P}|$ vertices separates A from B. So $H \subseteq G$ contains more than $|\mathcal{P}|$ paths from A to B by Menger's theorem (3.3.1), which contradicts the choice of $\mathcal{P}$. $\square$

When k is infinite, however, the result suddenly becomes trivial. Indeed, let $\mathcal{P}$ be any maximal set of disjoint A–B paths in G. Then the union of all these paths separates A from B, so $\mathcal{P}$ must be infinite. But then the cardinality of this union is no bigger than $|\mathcal{P}|$. Thus, $\mathcal{P}$ contains $|\mathcal{P}| = \big|\bigcup \mathcal{P}\big| \geqslant k$ disjoint A–B paths, as desired.

Of course, this is no more than a trick played on us by infinite cardinal arithmetic: although, numerically, the A–B separator consisting of all the inner vertices of paths in $\mathcal{P}$ is no bigger than $|\mathcal{P}|$, it uses far more vertices to separate A from B than should be necessary. Or put another way: when our path systems and separators are infinite, their cardinalities alone are no longer a sufficiently fine tool to distinguish carefully chosen 'small' separators from unnecessarily large and wasteful ones.

To overcome this problem, Erdős suggested an alternative form of Menger's theorem, which for finite graphs is clearly equivalent to the standard version. Recall that an A–B separator X is said to lie *on* a set $\mathcal{P}$ of disjoint A–B paths if X consists of a choice of exactly one vertex from each path in $\mathcal{P}$. The following so-called *Erdős-Menger conjecture*, now a theorem, influenced much of the development of infinite connectivity and matching theory:

Erdős-Menger conjecture

Theorem 8.4.4. (Aharoni & Berger 2009)
Let G be any graph, and let $A, B \subseteq V(G)$. Then G contains a set $\mathcal{P}$ of disjoint A–B paths and an A–B separator on $\mathcal{P}$.

The next few pages give a proof of Theorem 8.4.4 for countable G.

Of the three proofs we gave for the finite case of Menger's theorem, only the last has any chance of being adaptable to the infinite case: the others were by induction on $|G|$ or $\big|\bigcup \mathcal{P}\big|$, and both these parameters may now be infinite. The third proof, however, looks more promising: recall that, by Lemmas 3.3.2 and 3.3.3, it provided us with a tool to either find a separator on a given system of A–B paths, or to construct another system of A–B paths that covers more vertices in A and in B.

Lemmas 3.3.2 and 3.3.3 (whose proofs work for infinite graphs too) will indeed form a cornerstone of our proof for Theorem 8.4.4. However, it will not do just to apply these lemmas infinitely often. Indeed, although any finite number of applications of Lemma 3.3.2 leaves us with another system of disjoint A–B paths, an infinite number of iterations may leave nothing at all: each edge may be toggled on and off infinitely often by successive alternating paths, so that no 'limit system' of A–B paths will be defined. We shall therefore take another tack: starting at A, we grow simultaneously as many disjoint paths towards B as possible.

To make this precise, we need some terminology. Given a set $X \subseteq V(G)$, let us write $G_{X \to B}$ for the subgraph of G induced by X and all the components of $G - X$ that meet B. $G_{X \to B}$

wave

Let $\mathcal{W} = (\, W_a \mid a \in A\,)$ be a family of disjoint paths such that every W_a starts in a. We call $\mathcal{W}$ an $A \to B$ *wave* in G if the set Z of final vertices of paths in $\mathcal{W}$ separates A from B in G. (Note that $\mathcal{W}$ may contain infinite paths, which have no final vertex.) Sometimes, we shall wish to consider $A \to B$ waves in subgraphs of G that contain A but not all of B. For this reason we do not formally require that $B \subseteq V(G)$.

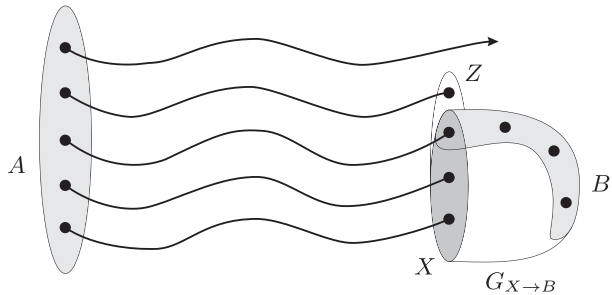

Fig. 8.4.1. A small $A \to B$ wave $\mathcal{W}$ with boundary X

boundary $(\mathcal{W}, X)$

large/small

proper

When $\mathcal{W}$ is a wave, then the set $X \subseteq Z$ of those vertices in Z that either lie in B or have a neighbour in $G_{Z \to B} - Z$ is a minimal A–B separator in G; note that $z \in Z$ lies in X if and only if it can be linked to B by a path that has no vertex other than z on $\mathcal{W}$. We call X the *boundary* of $\mathcal{W}$, and often use $(\mathcal{W}, X)$ as shorthand for the wave $\mathcal{W}$ together with its boundary X. If all the paths in $\mathcal{W}$ are finite and $X = Z$, we call the wave $\mathcal{W}$ *large*; otherwise it is *small*. We shall call $\mathcal{W}$ *proper* if at least one of the paths in $\mathcal{W}$ is non-trivial, or if all its paths are trivial but its boundary is a proper subset of A. Every small wave, for example, is proper. Note that while some $A \to B$ wave always exists, e.g. the family $(\,\{a\} \mid a \in A\,)$ of singleton paths, G need not have a proper $A \to B$ wave. (For example, if A consists of two vertices of $G = K^{10}$ and B of three other vertices, there is no proper $A \to B$ wave.)

$\mathcal{U} + \mathcal{V}$

If $(\mathcal{U}, X)$ is an $A \to B$ wave in G and $(\mathcal{V}, Y)$ is an $X \to B$ wave in $G_{X \to B}$, then the family $\mathcal{W} = \mathcal{U} + \mathcal{V}$ obtained from $\mathcal{U}$ by appending the paths of $\mathcal{V}$ (to those paths of $\mathcal{U}$ that end in X) is clearly an $A \to B$ wave in G, with boundary Y. Note that $\mathcal{W}$ is large if and only if both $\mathcal{V}$ and $\mathcal{U}$ are large. $\mathcal{W}$ is greater than $\mathcal{U}$ in the following sense.

$\leqslant$

limit wave

Given two path systems $\mathcal{U} = (\, U_a \mid a \in A\,)$ and $\mathcal{W} = (\, W_a \mid a \in A\,)$, write $\mathcal{U} \leqslant \mathcal{W}$ if $U_a \subseteq W_a$ for every $a \in A$. Given a chain $(\mathcal{W}^i, X^i)_{i \in I}$ of waves in this ordering, with $\mathcal{W}^i = (\, W^i_a \mid a \in A\,)$ say, let $\mathcal{W}^* = (\, W^*_a \mid a \in A\,)$ be defined by $W^*_a := \bigcup_{i \in I} W^i_a$. Then $\mathcal{W}^*$ is an $A \to B$ wave: any A–B path is finite but meets every X^i, so at least one of its vertices lies in X^i for arbitrarily large $(\mathcal{W}^i, X^i)$ and hence is the final vertex of a path in $\mathcal{W}^*$. Clearly $\mathcal{W}^i \leqslant \mathcal{W}^*$ for all $i \in I$; we call $\mathcal{W}^*$ the *limit* of the waves $\mathcal{W}^i$.

As every chain of $A \to B$ waves is bounded above by its limit wave, Zorn's lemma implies that G has a maximal $A \to B$ wave $\mathcal{W}$; let X be its boundary. This wave $(\mathcal{W}, X)$ forms the first step in our proof for Theorem 8.4.4: if we can now find disjoint paths in $G_{X \to B}$ linking all the vertices of X to B, then X will be an A–B separator on these paths preceded by the paths of $\mathcal{W}$ that end in X. *maximal wave*

By the maximality of $\mathcal{W}$, there is no proper $X \to B$ wave in $G_{X \to B}$. For our proof it will thus suffice to prove the following (renaming X as A):

Lemma 8.4.5. *If G has no proper $A \to B$ wave, then G contains a set of disjoint A–B paths linking all of A to B.*

Our approach to the proof of Lemma 8.4.5 is to enumerate the vertices in $A =: \{a_1, a_2, \dots\}$, and to find the required A–B paths $P_n = a_n \dots b_n$ in turn for $n = 1, 2, \dots$. Since our premise in Lemma 8.4.5 is that G has no proper $A \to B$ wave, we would like to choose P_1 so that $G - P_1$ has no proper $(A \smallsetminus \{a_1\}) \to B$ wave: this would restore the same premise to $G - P_1$, and we could proceed to find P_2 in $G - P_1$ in the same way. $a_1, a_2, \dots$ P_n

We shall not be able to choose P_1 quite like this, but we shall be able to do something almost as good. We shall construct P_1 so that deleting it (as well as a few more vertices outside A) leaves a graph that has a large maximal $(A \smallsetminus \{a_1\}) \to B$ wave $(\mathcal{W}, A')$. We then earmark the paths $W_n = a_n \dots a'_n$ $(n \geqslant 2)$ of this wave as initial segments for the paths P_n. By the maximality of $\mathcal{W}$, there is no proper $A' \to B$ wave in $G_{A' \to B}$. In other words, we have restored our original premise to $G_{A' \to B}$, and can find there an A'–B path $P'_2 = a'_2 \dots b_2$. Then $P_2 := a_2 W_2 a'_2 P'_2$ is our second path for Lemma 8.4.5, and we continue inductively inside $G_{A' \to B}$.

Given a set $\hat{A}$ of vertices in G, let us call a vertex $a \notin \hat{A}$ *linkable* for $(G, \hat{A}, B)$ if $G - \hat{A}$ contains an a–B path P and a set $X \supseteq V(P)$ of vertices such that $G - X$ has a large maximal $\hat{A} \to B$ wave. (The first such a we shall be considering will be a_1, and $\hat{A}$ will be the set $\{a_2, a_3, \dots\}$.) *linkable*

Lemma 8.4.6. *Let $a^* \in A$ and $\hat{A} := A \smallsetminus \{a^*\}$, and assume that G has no proper $A \to B$ wave. Then a^* is linkable for $(G, \hat{A}, B)$.*

Proof of Lemma 8.4.5 (assuming Lemma 8.4.6). Let G be as in Lemma 8.4.5, i.e. assume that G has no proper $A \to B$ wave. We construct subgraphs $G_1, G_2, \dots$ of G satisfying the following statement (Fig. 8.4.2):

G_n contains a set $A^n = \{a^n_n, a^n_{n+1}, a^n_{n+2}, \dots\}$ of distinct vertices such that G_n has no proper $A^n \to B$ wave. In G there are disjoint paths P_i ($i < n$) and W^n_i ($i \geqslant n$) starting at a_i. The P_i are disjoint from G_n and end in B. The W^n_i end in a^n_i and are otherwise disjoint from G_n. (∗)

Clearly, the paths $P_1, P_2, \dots$ will satisfy Lemma 8.4.5.

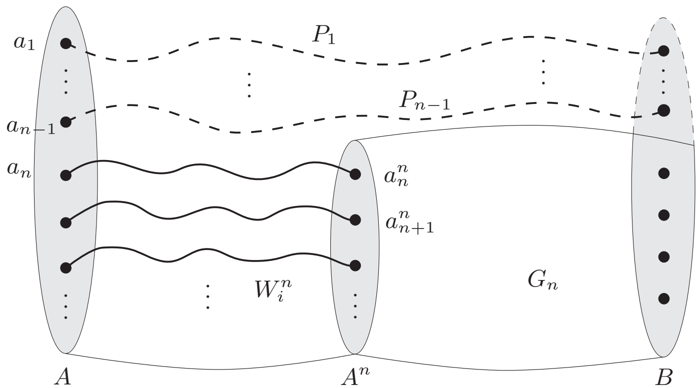

Fig. 8.4.2. G_n has no proper $A^n \to B$ wave

Let $G_1 := G$, and put $a^1_i := a_i$ and $W^1_i := \{a_i\}$ for all $i \geqslant 1$. Since by assumption G has no proper $A \to B$ wave, these definitions satisfy (∗) for $n = 1$. Suppose now that (∗) has been satisfied for n. Put $\hat{A}^n := A^n \setminus \{a^n_n\}$. By Lemma 8.4.6 applied to G_n, we can find in $G_n - \hat{A}^n$ an a^n_n–B path P and a set $X_n \supseteq V(P)$ such that $G_n - X_n$ has a large maximal $\hat{A}^n \to B$ wave $(\mathcal{W}, A^{n+1})$. Let P_n be the path $W^n_n \cup P$. For $i \geqslant n+1$, let W^{n+1}_i be W^n_i followed by the path of $\mathcal{W}$ starting at a^n_i, and call its last vertex a^{n+1}_i. By the maximality of $\mathcal{W}$ there is no proper $A^{n+1} \to B$ wave in $G_{n+1} := (G_n - X_n)_{A^{n+1} \to B}$, so (∗) is satisfied for $n+1$. □

To complete our proof of Theorem 8.4.4, it remains to prove Lemma 8.4.6. For this, we need another lemma:

Lemma 8.4.7. *Let x be a vertex in $G - A$. If G has no proper $A \to B$ wave but $G - x$ does, then every $A \to B$ wave in $G - x$ is large.*

(3.3.2) (3.3.3) *Proof.* Suppose $G - x$ has a small $A \to B$ wave $(\mathcal{W}, X)$. Put $B' := X \cup \{x\}$, and let $\mathcal{P}$ denote the set of A–X paths in $\mathcal{W}$ (Fig. 8.4.3). If G contains an A–B' separator S on $\mathcal{P}$, then replacing in $\mathcal{W}$ every $P \in \mathcal{P}$ with its initial segment ending in S we obtain a small (and hence proper) $A \to B$ wave in G, which by assumption does not exist. By Lemmas 3.3.3

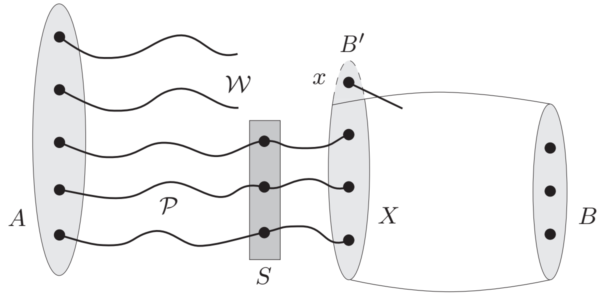

Fig. 8.4.3. A hypothetical small $A \to B$ wave in $G - x$

and 3.3.2, therefore, G contains a set $\mathcal{P}'$ of disjoint A–B' paths exceeding $\mathcal{P}$. The set of last vertices of these paths contains X properly, and hence must be all of $B' = X \cup \{x\}$. But B' separates A from B in G, so we can turn $\mathcal{P}'$ into an $A \to B$ wave in G by adding as singleton paths any vertices of A it does not cover. As x lies on $\mathcal{P}'$ but not in A, this is a proper wave, which by assumption does not exist. □

Proof of Lemma 8.4.6. We inductively construct trees $T_0 \subseteq T_1 \subseteq \dots$ in $G - (\hat{A} \cup B)$ and path systems $\mathcal{W}_0 \leqslant \mathcal{W}_1 \leqslant \dots$ in G so that each $\mathcal{W}_n$ is a large maximal $\hat{A} \to B$ wave in $G - T_n$. $\mathcal{W}_n$

Let $\mathcal{W}_0 := (\{a\} \mid a \in \hat{A})$. Clearly, $\mathcal{W}_0$ is an $\hat{A} \to B$ wave in $G - a^*$, and it is large and maximal: if not, then $G - a^*$ has a proper $\hat{A} \to B$ wave, and adding the trivial path $\{a^*\}$ to this wave turns it into a proper $A \to B$ wave (which by assumption does not exist). If $a^* \in B$, the existence of $\mathcal{W}_0$ makes a^* linkable for $(G, \hat{A}, B)$. So we assume that $a^* \notin B$. Now $T_0 := \{a^*\}$ and $\mathcal{W}_0$ are as desired.

Suppose now that T_n and $\mathcal{W}_n$ have been defined, and let A_n denote the set of last vertices of the paths in $\mathcal{W}_n$. Since $\mathcal{W}_n$ is large, A_n is its boundary, and since $\mathcal{W}_n$ is maximal, $G_n := (G - T_n)_{A_n \to B}$ has no proper $A_n \to B$ wave (Fig. 8.4.4). A_n G_n

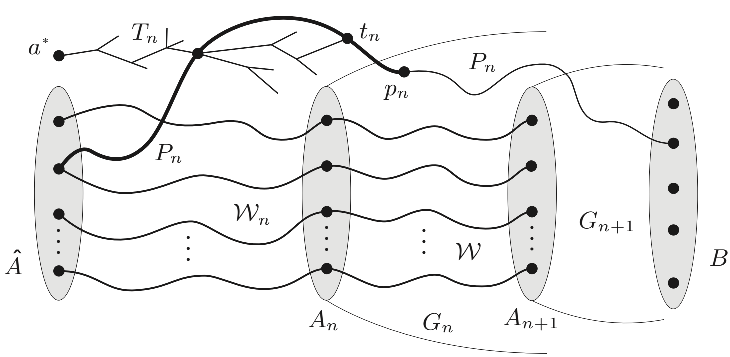

Fig. 8.4.4. As $\mathcal{W}_n$ is maximal, G_n has no proper $A_n \to B$ wave

Note that A_n does not separate A from B in G: if it did, then $\mathcal{W}_n \cup \{a^*\}$ would be a small $A \to B$ wave in G, which does not exist. Hence, $G - A_n$ contains an A–B path P, which meets T_n because $(\mathcal{W}_n, A_n)$ is a wave in $G - T_n$. Let $P =: P_n$ be chosen so that its vertex p_n following its last vertex t_n in T_n is minimal in some fixed enumeration of $V(G)$.

P_n

p_n, t_n

T_{n+1}

Let $T_{n+1} := T_n \cup t_n P_n p_n$ be the tree obtained from T_n by adding the edge $t_n p_n$. If $G_n - p_n$ has no proper $A_n \to B$ wave, then $\mathcal{W}_{n+1} := \mathcal{W}_n$ is large and maximal not only in $G - T_n$ but also in $G - T_{n+1}$. If $G_n - p_n$ has a proper $A_n \to B$ wave, $\mathcal{W}$ say, we choose it maximal. Since G_n has no such wave, $\mathcal{W}$ is large by Lemma 8.4.7. Now $\mathcal{W}_{n+1} := \mathcal{W}_n + \mathcal{W}$ is a large maximal $A_n \to B$ wave in $G - T_{n+1}$. In either case T_{n+1} and $\mathcal{W}_{n+1}$ are as desired, unless $p_n \in B$. But then a^* is linkable for $(G, \hat{A}, B)$ with a^*–B path $a^* T_{n+1} p_n$, wave $\mathcal{W}_{n+1}$ and $X = V(T_{n+1})$.

Put $T^* := \bigcup_{n \in \mathbb{N}} T_n$. Then the $\mathcal{W}_n$ are $\hat{A} \to B$ waves also in $G - T^*$; let $(\mathcal{W}^*, A^*)$ be their limit. Our aim is to show that A^* separates A from B not only in $G - T^*$ but even in G: then $(\mathcal{W}^* \cup \{a^*\}, A^*)$ is a small $A \to B$ wave in G, which by assumption does not exist: a contradiction.

Suppose there exists an A–B path Q in $G - A^*$. Let t be its last vertex in T^*. Since T^* does not meet B, there is a vertex p following t on Q. This p is not among the p_n, since T^* contains those but not p. Let n be large enough that $t \in T_n$ and that p precedes p_n in our fixed enumeration of $V(G)$. (As $p_n \in T_{n+1} - T_n$, the p_n are pairwise distinct.)

The fact that $a^* T_n t Q$ was not chosen as P_n means that its portion pQ outside T_n meets A_n, say in a vertex q. Now $q \notin A^*$ by the choice of Q. Let W be the path in $\mathcal{W}_n$ that joins $\hat{A}$ to q; this path too avoids A^*. But then WqQ contains an $\hat{A}$–B path in $G - T^*$ avoiding A^*, which contradicts the definition of A^*. □

The proof of Theorem 8.4.4 for countable G is now complete.

Turning now to matching, let us begin with a simple problem that is intrinsically infinite. Given two sets A, B and injective functions $A \to B$ and $B \to A$, is there necessarily also a bijection between A and B? Indeed there is – this is the famous Cantor-Bernstein theorem from elementary set theory. Recast in terms of matchings, the proof becomes very simple:

Proposition 8.4.8. *Let G be a bipartite graph, with bipartition $\{A, B\}$ say. If G contains a matching of A and a matching of B, then G has a 1-factor.*

Proof. Let H be the multigraph on $V(G)$ whose edge set is the disjoint union of the two matchings. (Thus, any edge that lies in both matchings becomes a double edge in H.) Every vertex in H has degree 1 or 2. In fact, it is easy to check that every component of H is an even cycle or an infinite path. Picking every other edge from each component, we obtain a 1-factor of G. □

The corresponding path problem in non-bipartite graphs, with sets of disjoint A–B paths instead of matchings, is less trivial. Let us say that a set $\mathcal{P}$ of paths in G *covers* a set U of vertices if every vertex in U is an endvertex of a path in $\mathcal{P}$. *covers*

Theorem 8.4.9. (Pym 1969)
Let G be a graph, and let $A, B \subseteq V(G)$. Suppose that G contains two sets of disjoint A–B paths, one covering A and one covering B. Then G contains a set of disjoint A–B paths covering $A \cup B$.

Some hints for a proof of Theorem 8.4.9 are included with Exercise 70.

Next, let us see how the standard matching theorems for finite graphs – König, Hall, Tutte, Gallai-Edmonds – extend to infinite graphs. For locally finite graphs, they all have straightforward extensions by compactness; see Exercises 26–29. But there are also very satisfactory extensions to graphs of arbitrary cardinality. Their proofs form a coherent body of theory and are much deeper, so we shall only be able to state those results and point out how some of them are related. But, as with Menger's theorem, the statements themselves are interesting too: finding the 'right' restatement of a given finite result to make a substantial infinite theorem is by no means easy, and most of them were found only as the theory itself developed over the years.

Let us start with bipartite graphs. The following Erdős-Menger-type extension of König's theorem (2.1.1) is now a corollary of Theorem 8.4.4:

Theorem 8.4.10. (Aharoni 1984)
Every bipartite graph has a matching, M say, and a vertex cover of its edge set that consists of exactly one vertex from every edge in M.

What about an infinite version of the marriage theorem (2.1.2)? The finite theorem says that a matching exists as soon as every subset S of the first partition class has enough neighbours in the second. But how do we measure 'enough' in an infinite graph? Just as in Menger's theorem, comparing cardinalities is not enough (Exercise 27).

However, there is a neat way of rephrasing the marriage condition for a finite graph without appealing to cardinalities. Call a subset X of one partition class *matchable* to a subset Y of the other if the subgraph spanned by X and Y contains a matching of X. Now if S is *minimal* with $|S| > |N(S)|$ then, by the marriage theorem, S is 'larger' than $N(S)$ also in the sense that S is not matchable to $N(S)$ but $N(S)$ is matchable to S. (Indeed, by the minimality of S and the marriage theorem, any $S' \subsetneq S$ with $|S'| = |S| - 1$ can be matched to $N(S)$. As $|S'| = |S| - 1 \geqslant |N(S)|$, this matching covers $N(S)$.) Thus, if there is any obstruction S of the type $|S| > |N(S)|$ to a perfect matching, there is also one where S is larger than $N(S)$ in this other sense: that S is not matchable to $N(S)$ but $N(S)$ is matchable to S. *matchable*

Rewriting the marriage condition in this way does indeed yield an infinite version of Hall's theorem, which follows from Theorem 8.4.10 just as the marriage theorem follows from König's theorem:

Corollary 8.4.11. *A bipartite graph with bipartition $\{A, B\}$ contains a matching of A unless there is a set $S \subseteq A$ such that S is not matchable to $N(S)$ but $N(S)$ is matchable to S.*

Proof. Consider a matching M and a cover U as in Theorem 8.4.10. Then $U \cap B \supseteq N(A \setminus U)$ is matchable to $A \setminus U$, by the edges of M. And if $A \setminus U$ is matchable to $N(A \setminus U)$, then adding this matching to the edges of M incident with $A \cap U$ yields a matching of A. □

Applied to a finite graph, Corollary 8.4.11 implies the marriage theorem: if $N(S)$ is matchable to S but not conversely, then clearly $|S| > |N(S)|$. Similarly, the finite version of Corollary 8.4.11 implies the finite case of the following sufficient condition for the existence of a matching of A:

Theorem 8.4.12. (Milner & Shelah 1974)
A bipartite graph with bipartition $\{A, B\}$ contains a matching of A if $d(a) \geqslant 1$ for every $a \in A$ and $d(a) \geqslant d(b)$ for every edge ab with $a \in A$.

partial matching

augmenting path

Let us now turn to non-bipartite graphs. If a finite graph has a 1-factor, then the set of vertices covered by any *partial matching* – one that leaves some vertices unmatched – can be increased by an augmenting path, an alternating path whose first and last vertex are unmatched (Ex. 1, Ch. 2). In an infinite graph we no longer insist that augmenting paths be finite, as long as they have a first vertex. Then, starting at any unmatched vertex with an edge of the 1-factor that we are assuming to exist, we can likewise find a unique maximal alternating path that will either be a ray or end at another unmatched vertex. Switching edges along this path we can then improve our current matching to increase the set of matched vertices, just as in a finite graph.

The existence of an inaugmentable partial matching, therefore, is an obvious obstruction to the existence of a 1-factor. The following theorem asserts that this obstruction is the only one:

Theorem 8.4.13. (Steffens 1977)
A countable graph has a 1-factor if and only if for every partial matching there exists an augmenting path.

Unlike its finite counterpart, Theorem 8.4.13 is far from trivial: augmenting a given matching 'blindly' need not lead to a well-defined matching at limit steps, since a given edge may get toggled on and off infinitely often (in which case its status will be undefined at the limit – example?). We therefore cannot simply find the desired 1-factor inductively.

In fact, Theorem 8.4.13 does not extend to uncountable graphs (Exercise 73). However, from the obstruction of inaugmentable partial matchings one can derive a Tutte-type condition that does extend. Given a set S of vertices in a graph G, let us write $\mathcal{C}'_{G-S}$ for the set of factor-critical components of $G-S$, and G'_S for the bipartite graph with vertex set $S \cup \mathcal{C}'_{G-S}$ and edge set $\{ sC \mid \exists c \in C : sc \in E(G) \}$. $\mathcal{C}'_{G-S}$ G'_S

Theorem 8.4.14. (Aharoni 1988)
A graph G has a 1-factor if and only if, for every set $S \subseteq V(G)$, the set $\mathcal{C}'_{G-S}$ is matchable to S in G'_S.

Applied to a finite graph, Theorem 8.4.14 implies Tutte's 1-factor theorem (2.2.1): if $\mathcal{C}'_{G-S}$ is not matchable to S in G'_S, then by the marriage theorem there is a subset S' of S that sends edges to more than $|S'|$ components in $\mathcal{C}'_{G-S}$ that are also components of $G-S'$, and these components are odd because they are factor-critical.

Theorems 8.4.10 and 8.4.14 also imply an infinite version of the Gallai-Edmonds theorem (2.2.3):

Corollary 8.4.15. *Every graph $G = (V, E)$ has a set S of vertices that is matchable to $\mathcal{C}'_{G-S}$ in G'_S and such that every component of $G-S$ not in $\mathcal{C}'_{G-S}$ has a 1-factor. Given any such set S, the graph G has a 1-factor if and only if $\mathcal{C}'_{G-S}$ is matchable to S in G'_S.*

Proof. Given a pair (S, M) where $S \subseteq V$ and M is a matching of S in G'_S, and given another such pair (S', M'), write $(S, M) \leqslant (S', M')$ if

$$S \subseteq S' \subseteq V \setminus \bigcup \{ V(C) \mid C \in \mathcal{C}'_{G-S} \}$$

and $M \subseteq M'$. Since $\mathcal{C}'_{G-S} \subseteq \mathcal{C}'_{G-S'}$ for any such S and S', Zorn's lemma implies that there is a maximal such pair (S, M). S, M

For the first statement, we have to show that every component C of $G-S$ that is not in $\mathcal{C}'_{G-S}$ has a 1-factor. If it does not, then by Theorem 8.4.14 there is a set $T \subseteq V(C)$ such that $\mathcal{C}'_{C-T}$ is not matchable to T in C'_T. By Corollary 8.4.11, this means that $\mathcal{C}'_{C-T}$ has a subset $\mathcal{C}$ that is not matchable in C'_T to the set $T' \subseteq T$ of its neighbours, while T' is matchable to $\mathcal{C}$; let M' be such a matching. Then $(S, M) < (S \cup T', M \cup M')$, contradicting the maximality of (S, M).

Of the second statement, only the backward implication is non-trivial. Our assumptions now are that $\mathcal{C}'_{G-S}$ is matchable to S in G'_S and vice versa (by the choice of S), so Proposition 8.4.8 yields that G'_S has a 1-factor. This defines a matching of S in G that picks one vertex x_C from every component $C \in \mathcal{C}'_{G-S}$ and leaves the other components of $G-S$ untouched. Adding to this matching a 1-factor of $C - x_C$ for every $C \in \mathcal{C}'_{G-S}$ and a 1-factor of every other component of $G-S$, we obtain the desired 1-factor of G. $\square$

Infinite matching theory may seem rather mature and complete as it stands, but there are still fascinating unsolved problems in the Erdős-Menger spirit concerning related discrete structures, such as partially ordered sets or hypergraphs. We conclude with one about graphs.

strongly perfect

weakly perfect

Call an infinite graph G *perfect* if every induced subgraph $H \subseteq G$ has a complete subgraph K of order $\chi(H)$, and *strongly perfect* if K can always be chosen so that it meets every colour class of some $\chi(H)$-colouring of H. (Exercise 75 gives an example of a perfect graph that is not strongly perfect.) Call G *weakly perfect* if the chromatic number of every induced subgraph $H \subseteq G$ is at most the supremum of the orders of its complete subgraphs.

Conjecture. (Aharoni & Korman 1993)
Every weakly perfect graph without infinite independent sets of vertices is strongly perfect.

8.5 Recursive structures

In this section we introduce another tool that is commonly used in infinite graph theory: to define a class of graphs recursively, so as to be able later to prove assertions about these graphs by (transfinite) induction. Rather than attempting a systematic treatment of this technique we give two examples; more can be found in the exercises.

Our first example is very simple: it describes the structure of a tree by recursively pruning away leaves and isolated ends. Let T be any tree, equipped with a root and the corresponding tree-order on its vertices. We recursively label the vertices of T by ordinals, as follows. Given an ordinal α, assume that we have decided for every $\beta < \alpha$ which of the vertices of T to label β, and let T_α be the subgraph of T induced by the vertices that are still unlabelled. Assign label α to every vertex t of T_α whose up-closure $\lfloor t \rfloor_{T_\alpha} = \lfloor t \rfloor_T \cap T_\alpha$ in T_α is a chain. The recursion terminates at the first α not used to label any vertex; for this α we put

T^*

$T_\alpha =: T^*$.

For each α, the vertices labelled α form an up-set in T_α: if $\lfloor t \rfloor_{T_\alpha}$ is a chain, then so is $\lfloor t' \rfloor_{T_\alpha}$ for every $t' \in \lfloor t \rfloor_{T_\alpha}$. Every T_α, therefore, is a down-set in T (induction on α) and hence connected. Thus, T_α is a tree, and the set of vertices labelled α induces in T_α a disjoint union of paths.

recursively prunable

Let us call T *recursively prunable* if every vertex of T gets labelled in this way, i.e., if $T^* = \emptyset$. We may then be able to prove assertions about T, or about graphs containing T as a normal spanning tree, by dealing in turn with those chains as they get deleted. The following proposition shows that the recursively prunable trees form a natural class also in structural terms:

Proposition 8.5.1. *A rooted tree is recursively prunable if and only if it contains no subdivision of the infinite binary tree T_2 as a subgraph.*

Proof. Let T be any rooted tree. Suppose first that T is not recursively prunable, i.e. that $T^* \neq \emptyset$. Since no vertex of T^* gets labelled when the recursion terminates, every $t \in T^*$ has two incomparable vertices of T^* above it. As T^* is connected, it is now easy to find a subdivision of T_2 in T^* inductively, along the levels of T_2.

Conversely, suppose that T contains a subdivision T' of T_2. We shall see in a moment that T' can be chosen 'upwards' in T, that is, in such a way that the tree-order which T induces on its vertices agrees with its own tree-order as induced by T_2. If this is the case, then every vertex of T' has two incomparable vertices of T' above it (in both orders). Hence there can be no minimal ordinal α such that a vertex of T' is labelled α. Thus all of T' remains unlabelled, and $\emptyset \neq T' \subseteq T^*$ as desired.

It remains to show that T' can indeed be chosen in this way. Let T' be any subdivision of T_2 in T, and let u be minimal in the tree-order of T among the vertices of T'. Induction on the levels of the tree $\lfloor u \rfloor_{T'}$ shows that $\leqslant_{T'}$ and $\leqslant_T$ agree on $\lfloor u \rfloor_{T'}$: any upper neighbour in T' of a vertex $t \in \lfloor u \rfloor_{T'}$ must lie above t also in T, since the unique lower neighbour of t in T is either not in T' (if $t = u$), or by induction it is the unique lower neighbour of t also in T'. Pick any branch vertex v of T' in $\lfloor u \rfloor_{T'}$. Then $\lfloor v \rfloor_{T'}$ is the desired subdivision of T_2 in T. □

The charm of the recursive pruning discussed above lies in the fact that it removes the 'messy bits' of a given tree in an automated sort of way: we do not have to know where they are, but if our given tree contains a 'clean' ever-branching subtree, then the recursion will reveal it.

And there is another way of viewing it. We might think of rooted paths (paths with a first vertex, which we take to be the root) as particularly basic objects, and call them *rooted trees of rank* 0. We could then define rooted trees of higher ordinal rank inductively, taking as the rooted trees of rank α those that do not have any rank $\beta < \alpha$ but in which it is possible to delete a path starting at the root so as to leave components that each have some rank $< \alpha$ when taken with the induced tree-order. Then the rooted trees that are assigned a rank in this way are precisely the recursively prunable ones, and those of rank $\leqslant \alpha$ are precisely those whose labels do not exceed α (Exercise 77).

We now apply the same idea to graphs that are not necessarily trees. Let us assign *rank* 0 to all the finite graphs. Given an ordinal $\alpha > 0$, we assign *rank* α to every graph G that does not already have a rank $\beta < \alpha$ and which has a finite set U of vertices such that every component of $G - U$ has some rank $< \alpha$. *rank*

When disjoint graphs G_i have ranks $\alpha_i < \alpha$, their union clearly has a rank of at most α; if the union is finite, it has rank $\max_i \alpha_i$. Induction on α shows that subgraphs of graphs of rank α also have a rank of at most α. Conversely, joining finitely many new vertices to a graph, no matter how, will not change its rank.

Not every graph has a rank. Indeed, the ray cannot have a rank, since deleting finitely many of its vertices always leaves a component that is also a ray. As subgraphs of graphs with a rank also have a rank, this means that only rayless graphs can have a rank. But all these do:

Lemma 8.5.2. *A graph has a rank if and only if it is rayless.*

Proof. Consider a graph G that has no rank. Then one of its components, C_0 say, has no rank; let v_0 be a vertex in C_0. Now $C_0 - v_0$ has a component C_1 that has no rank; let v_1 be a neighbour of v_0 in C_1. Continuing inductively, we find a ray $v_0v_1\ldots$ in G. □

Because of Lemma 8.5.2, we call the ranking defined above the *ranking of rayless graphs*. As an application of this ranking, we now prove the unfriendly partition conjecture from Section 8.1 for rayless graphs.

Theorem 8.5.3. *Every countable rayless graph G has an unfriendly partition.*

Proof. To help with our formal notation, we shall think of a partition of a set V as a map $\pi\colon V \to \{0,1\}$. We apply induction on the rank of G. When this is zero then G is finite, and an unfriendly partition can be obtained by maximizing the number of edges across the partition.
α Suppose now that G has rank $\alpha > 0$, and assume the theorem as true for graphs of smaller rank.

U Let U be a finite set of vertices in G such that each of the components
$C_0, C_1, \ldots$ nents $C_0, C_1, \ldots$ of $G - U$ has rank $< \alpha$. Partition U into the set U_0 of vertices that have finite degree in G, the set U_1 of vertices that have
U_0, U_1, U_2 infinitely many neighbours in some C_n, and the set U_2 of vertices that have infinite degree but only finitely many neighbours in each C_n.

$G_0, G_1, \ldots$ For every $n \in \mathbb{N}$ let $G_n := G[U \cup V(C_0) \cup \ldots \cup V(C_n)]$. This is a graph of some rank $\alpha_n < \alpha$, so by induction it has an unfriendly partition π_n. Each of these π_n induces a partition of U. Let π_U be a partition of U induced by π_n for infinitely many n, say for $n_0 < n_1 < \ldots$. Choose n_0 large enough that G_{n_0} contains all the neighbours of vertices in U_0, and the other n_i large enough that every vertex in U_2 has more
$n_0, n_1, \ldots$ neighbours in $G_{n_i} - G_{n_{i-1}}$ than in $G_{n_{i-1}}$, for all $i > 0$. Let π be the partition of G defined by letting $\pi(v) := \pi_{n_i}(v)$ for all $v \in G_{n_i} - G_{n_{i-1}}$
π and all i, where $G_{n_{-1}} := \emptyset$. Note that $\pi|_U = \pi_{n_0}|_U = \pi_U$.

Let us show that π is unfriendly. We have to check that every vertex is *happy with* π, i.e., that it has at least as many neighbours in the opposite class under π as in its own.[8] To see that a vertex $v \in G - U$ is happy with π, let i be minimal such that $v \in G_{n_i}$ and recall that v was happy with π_{n_i}. As both v and its neighbours in G lie in $U \cup V(G_{n_i} - G_{n_{i-1}})$, and π agrees with π_{n_i} on this set, v is happy also with π. Vertices in U_0 are happy with π, because they were happy with π_{n_0}, and π agrees with π_{n_0} on U_0 and all its neighbours. Vertices in U_1 are also happy. Indeed, every $u \in U_1$ has infinitely many neighbours in some C_n, and hence in some G_{n_i}; let i be minimal with this property. Then u has infinitely many opposite neighbours under π_{n_i} in G_{n_i}, and hence in $G_{n_i} - G_{n_{i-1}}$. Since π_{n_i} agrees with π on both U and $G_{n_i} - G_{n_{i-1}}$, our vertex u has infinitely many opposite neighbours also under π. Vertices in U_2, finally, are happy with every π_{n_i}. By our choice of n_i, at least one of their opposite neighbours under π_{n_i} must lie in $G_{n_i} - G_{n_{i-1}}$. Since π_{n_i} agrees with π on both U_2 and $G_{n_i} - G_{n_{i-1}}$, this gives every $u \in U_2$ at least one opposite neighbour under π in every $G_{n_i} - G_{n_{i-1}}$. Hence u has infinitely many opposite neighbours under π, which clearly makes it happy. □

8.6 Graphs with ends: the complete picture

In this section we shall develop a deeper understanding of the global structure of infinite graphs, especially locally finite ones, that can be attained only by studying their ends. This structure is intrinsically topological, because topology best captures our intuition about convergence.[9]

Our first goal will be to make precise our intuitive idea that the ends of a graph are the 'points at infinity' to which its rays converge. To do so, we shall define a topological space $|G|$ associated with a graph $G = (V, E, \Omega)$ and its ends.[10] By considering topological versions of paths, cycles and spanning trees in this space, we shall then be able to extend to infinite graphs some parts of finite graph theory that would not otherwise have infinite counterparts; see the notes for more examples. Thus, the ends of an infinite graph turn out to be more than a curious phenomenon: they form an integral part of the picture, without which it cannot be properly understood. V, E, Ω

To build the space $|G|$ formally, we start with the set $V \cup \Omega$. For every edge $e = uv$ we add a set $\mathring{e} = (u, v)$ of continuum many points, making these sets $\mathring{e}$ disjoint from each other and from $V \cup \Omega$. We then choose for each e some fixed bijection between $\mathring{e}$ and the real interval $(0, 1)$, and (u, v)

[8] It is only by tradition that such partitions are called 'unfriendly'; our vertices love them.

[9] Only point-set topology is needed for the text. See the exercises for more.

[10] The notation of $|G|$ comes from topology and clashes with our notation for the order of G. But there is little danger of confusion, so we keep both.

$[u,v]$ extend this bijection to one between $[u,v] := \{u\} \cup \mathring{e} \cup \{v\}$ and $[0,1]$. This bijection defines a metric on $[u,v]$; we call $[u,v]$ a *topological edge* with $\mathring{F}$ *inner points* $x \in \mathring{e}$. Given any $F \subseteq E$ we write $\mathring{F} := \bigcup \{\, \mathring{e} \mid e \in F \,\}$. When we speak of a 'graph' $H \subseteq G$, we shall often also mean its corresponding point set $V(H) \cup \mathring{E}(H)$.

Having thus defined the point set of $|G|$, let us choose a basis of open sets to define its topology. For every edge uv, declare as open all subsets of (u,v) that correspond, by our fixed bijection between (u,v) and $(0,1)$, to an open set in $(0,1)$. For every vertex u and $\epsilon > 0$, declare as open the 'open star around u of radius ϵ', that is, the set of all points on edges $[u,v]$ at distance less than ϵ from u, measured individually for each edge in its metric inherited from $[0,1]$. Finally, for every end ω and every finite $C(S,\omega)$ set $S \subseteq V$, there is a unique component $C(S,\omega)$ of $G - S$ that contains rays from ω. Let $\Omega(S,\omega) := \{\, \omega' \in \Omega \mid C(S,\omega') = C(S,\omega) \,\}$. For every $\epsilon > 0$, write $\mathring{E}_\epsilon(S,\omega)$ for the set of all inner points of S–$C(S,\omega)$ edges at distance less than ϵ from their endpoint in $C(S,\omega)$. Then declare as open all sets of the form

$\hat{C}_\epsilon(S,\omega)$
$$\hat{C}_\epsilon(S,\omega) := C(S,\omega) \cup \Omega(S,\omega) \cup \mathring{E}_\epsilon(S,\omega)\,.$$

$|G|$ This completes the definition of $|G|$, whose open sets are the unions of the sets we explicitly chose as open above.

closure $\overline{X}$ The *closure* of a set $X \subseteq |G|$ will be denoted by $\overline{X}$. For example, $\overline{V} = V \cup \Omega$ (because every neighbourhood of an end contains a vertex), and the closure of a ray is obtained by adding its end. More generally, the closure of the set of teeth of a comb contains a unique end, the end of its spine. Conversely, if $U \subseteq V$ and $R \in \omega \in \Omega \cap \overline{U}$, there is a comb with spine R and teeth in U (Exercise 83). In particular, the closure of the subgraph $C(S,\omega)$ considered above is the set $C(S,\omega) \cup \Omega(S,\omega)$.

The subspaces X of $|G|$ we shall be interested in are usually the clo- *standard subspace* sure of a subgraph H of G, i.e., of the form $X = \overline{U} \cup \mathring{D}$ for $H = (U,D)$. We write $V(X)$ for U and $E(X)$ for D, and call such subspaces *standard*. $V(X), E(X)$ We also refer to such X as $\overline{H}$, or even as $\overline{D}$ if H has no isolated vertices, and then say that X is *spanned by* D. Note that the ends in X are always ends of G, not of H; in particular, they need not have a ray in H.

By definition, $|G|$ is always Hausdorff; indeed one can show that it is normal. When G is connected and locally finite, then $|G|$ is compact:[11]

Proposition 8.6.1. *If G is connected and locally finite, then $|G|$ is a compact Hausdorff space.*

(8.1.2) *Proof.* Let $\mathcal{O}$ be an open cover of $|G|$; we show that $\mathcal{O}$ has a finite subcover. Pick a vertex $v_0 \in G$, write D_n for the (finite) set of vertices at distance n from v_0, and put $S_n := D_0 \cup \ldots \cup D_{n-1}$. For every $v \in D_n$,

[11] Topologists call $|G|$ the *Freudenthal compactification* of G.

let $C(v)$ denote the component of $G - S_n$ containing v, and let $\hat{C}(v)$ be its closure together with all inner points of $C(v)$–S_n edges. Then $G[S_n]$ and these $\hat{C}(v)$ together partition $|G|$. $\hat{C}(v)$

We wish to prove that, for some n, each of the sets $\hat{C}(v)$ with $v \in D_n$ is contained in some $O(v) \in \mathcal{O}$. For then we can take a finite subcover of $\mathcal{O}$ for $G[S_n]$ (which is compact, being a finite union of edges and vertices), and add to it these finitely many sets $O(v)$ to obtain the desired finite subcover for $|G|$.

Suppose there is no such n. Then for each n the set V_n of vertices $v \in D_n$ such that no set from $\mathcal{O}$ contains $\hat{C}(v)$ is non-empty. Moreover, for every neighbour $u \in D_{n-1}$ of $v \in V_n$ we have $C(v) \subseteq C(u)$ because $S_{n-1} \subseteq S_n$, and hence $u \in V_{n-1}$; let $f(v)$ be such a vertex u. By the infinity lemma (8.1.2) there is a ray $R = v_0 v_1 \ldots$ with $v_n \in V_n$ for all n. Let ω be its end, and let $O \in \mathcal{O}$ contain ω. Since O is open, it contains a basic open neighbourhood of ω: there exist a finite set $S \subseteq V$ and $\epsilon > 0$ such that $\hat{C}_\epsilon(S, \omega) \subseteq O$. Now choose n large enough that S_n contains S and all its neighbours. Then $C(v_n)$ lies inside a component of $G - S$. As $C(v_n)$ contains the ray $v_n R \in \omega$, this component must be $C(S, \omega)$. Thus

$$\hat{C}(v_n) \subseteq \hat{C}_\epsilon(S, \omega) \subseteq O \in \mathcal{O},$$

contradicting the fact that $v_n \in V_n$. □

If G has a vertex of infinite degree then $|G|$ cannot be compact. (Why not?) But $\Omega \subseteq |G|$ can be compact; see Exercise 91 for when it is.

What else can we say about the space $|G|$ in general? For example, is it metrizable? Using a normal spanning tree T of G, it is indeed not difficult to define a metric on $|G|$ that induces its topology. But not every connected graph has a normal spanning tree, and it is not easy to determine in graph-theoretical terms which graphs do. Surprisingly, though, it is possible to deduce the existence of a normal spanning tree from that of a defining metric on $|G|$. Thus whenever $|G|$ is metrizable, a metric can be made visible in a natural and structural way.

Theorem 8.6.2. (Diestel 2006; Pitz 2020)
For a connected graph G, the following assertions are equivalent:

(i) *The space $|G|$ is metrizable.*

(ii) *G has a normal spanning tree.*

(iii) *All minors of G have countable colouring number.*

The proof of the equivalence of (i) and (ii) in Theorem 8.6.2 is indicated in Exercises 43 and 92. More on (iii) can be found in the notes.

Our next aim is to review, or newly define, some topological notions of paths and connectedness, of cycles, and of spanning trees. By substituting these topological notions with respect to $|G|$ for the corresponding graph-theoretical notions with respect to G one can extend to locally finite infinite graphs a number of theorems about paths, cycles and spanning trees in finite graphs whose ordinary infinite versions are false. We shall do this, as a case in point, for the tree packing theorem of Nash-Williams and Tutte, Theorem 2.4.1; see the notes for more.

X
connected

Let X be an arbitrary Hausdorff space. (Later, this will be a subspace of $|G|$.) X is *(topologically) connected* if it is not a union of two disjoint non-empty open subsets.[12] Note that continuous images of connected spaces are connected. For example, since the real interval $[0,1]$ is connected,[13] so are its continuous images in X.

arc

A *homeomorphic* image of $[0,1]$ in X is an *arc* in X; it *links* the images of 0 and 1, which are its *endpoints*. Every finite path in G defines an arc in $|G|$ in an obvious way. Similarly, every ray defines an arc linking its starting vertex to its end, and a double ray in G forms an arc with the two ends of its tails if these ends are distinct.

end degrees in subspaces

The *(topological) degree* of an end ω of G *in* a standard subspace X of $|G|$ is the supremum, in fact maximum, of all integers k such that X contains k arcs that end in ω and are otherwise disjoint.

For the remainder of this section let, unless otherwise mentioned,

$G = (V,E,\Omega)$

$G = (V,E,\Omega)$ be a fixed connected locally finite graph.

Unlike ordinary paths, arcs in $|G|$ can jump across a cut without containing an edge from it – but only if the cut is infinite:

[8.7.1]

Lemma 8.6.3. (Jumping Arc Lemma)
Let $F \subseteq E$ be a cut of G with sides V_1, V_2.

(i) *If F is finite, then $\overline{V_1} \cap \overline{V_2} = \emptyset$, and there is no arc in $|G| \setminus \mathring{F}$ with one endpoint in V_1 and the other in V_2.*

(ii) *If F is infinite, then $\overline{V_1} \cap \overline{V_2} \neq \emptyset$, and there will be such an arc if both V_1 and V_2 are connected in G.*

(8.2.2)

Proof. (i) Suppose that F is finite. Let S be the set of vertices incident with edges in F. Then S is finite and separates V_1 from V_2, so for every $\omega \in \Omega$ the connected graph $C(S,\omega)$ misses either V_1 or V_2. But then so does every basic open set of the form $\hat{C}_\epsilon(S,\omega)$. Therefore no end ω lies in the closure of both V_1 and V_2.

As $|G| \setminus \mathring{F} = \overline{G[V_1]} \cup \overline{G[V_2]}$ and this union is disjoint, no connected subset of $|G| \setminus \mathring{F}$ can meet both V_1 and V_2. Since arcs are continuous images of $[0,1]$ and hence connected, there is no V_1–V_2 arc in $|G| \setminus \mathring{F}$.

---

[12] These subsets would be complements of each other, and hence also be closed. Note that 'open' and 'closed' means open and closed *in* X: when X is a subspace of $|G|$ with the subspace topology, the two sets need not be open or closed in $|G|$.

[13] This takes a few lines to prove – can you prove it?

(ii) Suppose now that F is infinite. Since G is locally finite, the set U of endvertices of F in V_1 is also infinite. By the star-comb lemma (8.2.2), there is a comb in G with teeth in U; let ω be the end of its spine. Then every basic open neighbourhood $\hat{C}_\epsilon(S,\omega)$ of ω meets $U \subseteq V_1$ infinitely and hence also meets V_2, giving $\omega \in \overline{V_1} \cap \overline{V_2}$.

To obtain a V_1–V_2 arc in $|G| \setminus \mathring{F}$, all we need now is an arc in $\overline{G[V_1]}$ and another in $\overline{G[V_2]}$, both ending in ω. Such arcs exist if the graphs $G[V_i]$ are connected: we can then pick a sequence of vertices in V_i converging to ω, and apply the star-comb lemma in $G[V_i]$ to obtain a comb whose spine is a ray in $G[V_i]$ converging to ω. Concatenating these two rays yields the desired jumping arc. □

To some extent, arcs in $|G|$ assume the role that paths play in finite graphs. So arcs are important – but how do we find them? It is not always possible to construct arcs as explicitly as in the proof of Lemma 8.6.3 (ii). Figure 8.6.1, for example, shows an arc that goes through continuum many ends; such arcs cannot be constructed greedily by following a ray into its end and emerging from that end on another ray, etc.

There are two basic methods to obtain an arc between two given points, say two vertices x and y. One is to use compactness to obtain, as a limit of finite x–y paths, a *topologial* x–y path, a continuous map $\pi\colon [0,1] \to |G|$ sending 0 to x and 1 to y. A lemma from general topology then tells us that this path can be made injective:

Lemma 8.6.4. *The image of a topological x–y path in a Hausdorff space contains an x–y arc.* [8.7.3]

To illustrate this method, we will use it in the proof of Theorem 8.7.3.

Another method is to prove that the subspace in which we wish to find our x–y arc is topologically connected, and use this to deduce that it contains the desired arc. Our next three lemmas provide the tools needed to implement this approach in practice; we shall then illustrate its use in the proof of Theorem 8.6.9.

Being linked by an arc is an equivalence relation on the points of our Hausdorff space X: every x–y arc A has a first point p on any y–z arc A' (because A' is closed), and the obvious segments Ap and pA' together form an x–z arc in X. The corresponding equivalence classes are the *arc-components* of X. If X has only one arc-component, then X is *arc-connected*. *arc-component* *arc-connected*

Since $[0,1]$ is connected, arc-connectedness implies connectedness. The converse implication is false in general, even for spaces $X \subseteq |G|$ with G locally finite. But it holds in all the cases that matter:

Lemma 8.6.5. *Connected standard subspaces of $|G|$ are arc-connected.* [8.7]

Our proof of Theorem 8.7.3 will show how one can prove Lemma 8.6.5. Two further proofs are indicated in Exercises 94 and 134.

[8.7] **Lemma 8.6.6.** *Arc-components of standard subspaces of $|G|$ are closed.*

Proof. Let A be an arc-component of a standard subspace of $|G|$. Since A is connected, so is its closure $\overline{A}$. If $\overline{A} \setminus A \neq \emptyset$ then its points are limits of vertices in A (why?), so $\overline{A}$ is again standard. Hence $\overline{A}$ is arc-connected, either because $\overline{A} = A$ or by Lemma 8.6.5. But then $\overline{A} = A$, by definition of $\overline{A}$. Hence A is closed, as claimed. □

Connected standard subspaces of $|G|$ containing two given points are much easier to construct than an arc between two points. This has to do with the fact that they can be described in purely graph-theoretical terms, with reference only to finite subgraphs of G rather than to $|G|$. The description can be viewed as a topological analogue of the fact that a subgraph H of G is connected if and only if it contains an edge from every cut of G that separates two of its vertices:

[8.7.1] **Lemma 8.6.7.** *A standard subspace of $|G|$ is connected if and only if it contains an edge from every finite cut of G of which it meets both sides.*

Proof. Let $X \subseteq |G|$ be a standard subspace. For the forward implication, suppose that G has a finite cut $F = E(V_1, V_2)$ such that X meets both V_1 and V_2 but has no edge in F. Then

$$X \subseteq |G| \setminus \mathring{F} = \overline{G[V_1]} \cup \overline{G[V_2]}\,,$$

and this union is disjoint by Lemma 8.6.3 (i). The induced partition of X into non-empty closed subsets of X shows that X is not connected.

The backward implication holds vacuously if X meets more than one component of G; we may therefore assume that G is connected. If X is not connected, we can partition it into disjoint non-empty open subsets O_1 and O_2. As X is standard, $U_i := O_i \cap V(X) \neq \emptyset$ for both i. Let $\mathcal{P}$ be a maximal set of edge-disjoint U_1–U_2 paths in G, and put

$$F := \bigcup \{\, E(P) \mid P \in \mathcal{P} \,\}\,.$$

Then $E(X) \cap F = \emptyset$, and no component of $G - F$ meets both U_1 and U_2. Extending $\{U_1, U_2\}$ to a partition of V in such a way that each component of $G - F$ has all its vertices in one class, we obtain a cut $F' \subseteq F$ of G of which X meets both sides. As $E(X) \cap F = \emptyset$, it thus suffices to show that F is finite.

If F is infinite, then so is $\mathcal{P}$. As G is locally finite, the vertices of each $P \in \mathcal{P}$ are incident with only finitely many edges of G. We can thus inductively find an infinite subset of $\mathcal{P}$ consisting of paths that are not only edge-disjoint but disjoint. As G is connected, the endvertices in U_1 of these paths have a limit point ω in $|G|$ (Proposition 8.6.1), which is also a limit point of their endvertices in U_2. Since both O_1 and O_2 are closed in $|G|$, we thus have $\omega \in O_1 \cap O_2$, contradicting the choice of the O_i. □

A *circle* in a topological space is a homeomorphic image of the unit circle $S^1 \subseteq \mathbb{R}^2$. For example, if G is the 2-way infinite ladder shown in Figure 8.1.3, and we delete all its rungs (the vertical edges), what remains is a disjoint union of two double rays; its closure in $|G|$, obtained by adding the two ends of G, is a circle. Similarly, the double ray 'round the outside' of the 1-way ladder forms a circle together with the unique end of that ladder. *circle*

It is not hard to show that no arc in $|G|$ can consist entirely of ends. This implies that every circle in $|G|$ is a standard subspace; the set of edges spanning it will be called its *circuit*. *circuit*

A more adventurous example of a circle is shown in Figure 8.6.1. Suppose G is the graph obtained from the binary tree T_2 by joining for every finite 0–1 sequence ℓ the vertices $\ell 01$ and $\ell 10$ by a new edge e_ℓ. Together with all the (uncountably many) ends of G, the double rays $D_\ell \ni e_\ell$ shown in the figure form an arc A in $|G|$, whose union with the bottom double ray D is a circle in $|G|$ (Exercise 100). Note that no two of the double rays in A are consecutive: between any two there lies a third (cf. Exercise 101).

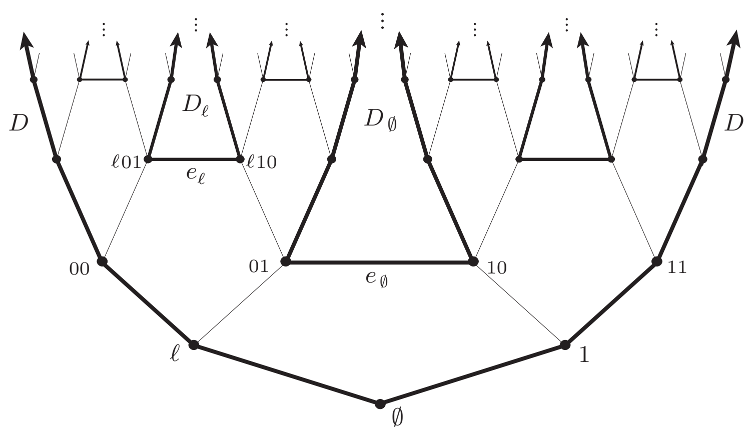

Fig. 8.6.1. The Wild Circle

A *topological spanning tree* of G is a connected standard subspace T of $|G|$ that contains every vertex but contains no circle. Since standard subspaces are closed, T also contains every end, and by Lemma 8.6.5 it *topological spanning tree* is even arc-connected. With respect to the deletion or addition of edges, it is both minimally connected and maximally 'acirclic' (Exercise 105).

One might expect that the closure $\overline{T}$ of an ordinary spanning tree T of G is always a topological spanning tree of $|G|$, but this is not the case: $\overline{T}$ may well contain a circle (Figure 8.6.2). Conversely, a subgraph whose closure is a topological spanning tree may well be disconnected: the 'vertical' rays in the $\mathbb{N} \times \mathbb{N}$ grid, for example, form a topological spanning tree together with the unique end.

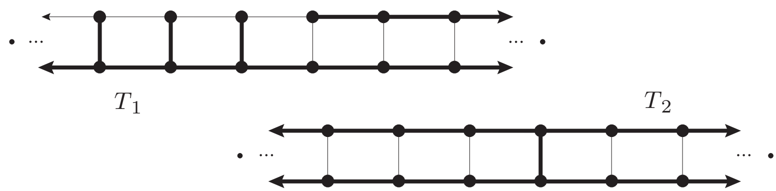

Fig. 8.6.2. T_1 is a topological spanning tree, but T_2 contains three circles.

Topological spanning trees can be constructed much as spanning trees of finite graphs: Lemma 8.6.11 will find one by iteratively deleting edges from $|G|$, but they can also be built up 'from below' (Exercise 108).

(8.2.4) Their mere existence even comes as a corollary of Theorem 8.2.4:

[8.7] **Lemma 8.6.8.** *The closure in $|G|$ of any normal spanning tree of G is a topological spanning tree of G.*

(1.5.4)
(8.2.3) *Proof.* Let T be a normal spanning tree of G. By Lemma 8.2.3, every end ω of G contains a normal ray R of T. Then $R \cup \{\omega\}$ is an arc linking ω to the root of T, so $\overline{T}$ is arc-connected.

It remains to check that $\overline{T}$ contains no circle. Suppose it does, and let A be the u–v arc obtained from that circle by deleting the inner
f points of an edge $f = uv$ it contains. Clearly, $f \in T$. Assume that $u < v$ in the tree-order of T, let T_u and T_v denote the components of $T - f$ containing u and v, and notice that $V(T_v)$ is the up-closure $\lfloor v \rfloor$ of v in T.

Now let $S := \lceil u \rceil$. By Lemma 1.5.4 (ii), $\lfloor v \rfloor$ is the vertex set of a component C of $G - S$. Thus, $V(C) = V(T_v)$ and $V(G - C) = V(T_u)$, so the set $E(C, S)$ of edges between these sets meets $E(T)$ precisely in f. Thus, $\overline{C}$ and $\overline{G - C}$ partition $|G| \setminus \mathring{E}(C, S) \supseteq A$ into two open sets both meeting A. This contradicts the fact that A is topologically connected. □

Note that the proof of Lemma 8.6.8 did not use our assumption that G is locally finite: whenever a graph G has a normal spanning tree T, the closure of T in $|G|$ is an arc-connected subspace that contains no circle.

As a first application of our new concepts, let us now extend the tree packing theorem (2.4.1) of Nash-Williams and Tutte to locally finite graphs. Its naive extension, with ordinary spanning trees, fails. Indeed, for every $k \in \mathbb{N}$ one can construct a $2k$-edge-connected locally finite graph that is left disconnected by the deletion of the edges in any one finite circuit (Exercise 21). Such a graph will have at least $k(\ell - 1)$ edges across any vertex partition into ℓ sets, but it cannot have more than two edge-disjoint spanning trees: adding an edge of one of these to another creates a (finite) fundamental circuit there, whose deletion would disconnect any third spanning tree.

As soon as we replace ordinary spanning trees with topological ones, however, Theorem 2.4.1 does extend:

Theorem 8.6.9. *The following statements are equivalent for all $k \in \mathbb{N}$ and connected locally finite multigraphs $G = (V, E)$:* $G = (V, E)$

(i) *G has k edge-disjoint topological spanning trees.*

(ii) *For every finite partition of V, into ℓ sets say, G has at least $k\,(\ell - 1)$ cross-edges.*

We begin our proof of Theorem 8.6.9 with a compactness extension of the finite theorem. This yields a weaker, 'finitary', statement at the limit (cf. Lemma 8.6.7):

Lemma 8.6.10. *If for every finite partition of V, into ℓ sets say, G has at least $k\,(\ell - 1)$ cross-edges, then G has k edge-disjoint spanning submultigraphs whose closures in $|G|$ are topologically connected.*

Proof. Pick an enumeration $v_0, v_1, \ldots$ of V. For every $n \in \mathbb{N}$ let G_n be the finite multigraph obtained from G by contracting every component of $G - \{v_0, \ldots, v_n\}$ to a vertex, deleting any loops but no parallel edges that arise in the contraction. Then $G[v_0, \ldots, v_n]$ is an induced submultigraph of G_n. Let $\mathcal{V}_n$ denote the set of all k-tuples $(H_n^1, \ldots, H_n^k)$ of edge-disjoint connected spanning submultigraphs of G_n. (2.4.1) (8.1.2)

Since every partition P of $V(G_n)$ induces a partition of V, since G has enough cross-edges for that partition, and since all these cross-edges are also cross-edges of P, Theorem 2.4.1 implies that $\mathcal{V}_n \neq \emptyset$. As every $(H_n^1, \ldots, H_n^k) \in \mathcal{V}_n$ induces an element $(H_{n-1}^1, \ldots, H_{n-1}^k)$ of $\mathcal{V}_{n-1}$, the infinity lemma (8.1.2), yields a sequence $(H_n^1, \ldots, H_n^k)_{n \in \mathbb{N}}$ of k-tuples, one from each $\mathcal{V}_n$, with a limit $(H^1, \ldots, H^k)$ defined by the nested unions

$$H^i := \bigcup_{n \in \mathbb{N}} H_n^i[v_0, \ldots, v_n]\,.$$

These H^i are edge-disjoint for distinct i (because the H_n^i are), but they need not be connected. To show that they have connected closures, it suffices by Lemma 8.6.7 to show that each of them has an edge in every finite cut F of G. Given F, choose n large enough that all the edges of F lie in $G[v_0, \dots, v_n]$. Then F is also a cut of G_n. Now consider the k-tuple $(H_n^1, \dots, H_n^k)$ which the infinity lemma picked from $\mathcal{V}_n$. Each of these H_n^i is a connected spanning submultigraph of G_n, so it contains an edge from F. But H_n^i agrees with H^i on $\{v_0, \dots, v_n\}$, so H^i too contains this edge from F. □

Lemma 8.6.11. *Every connected standard subspace of $|G|$ that contains V also contains a topological spanning tree of G.*

Proof. Let X be a connected standard subspace of $|G|$ containing V. Then G too must be connected, so it is countable. Let $e_0, e_1, \dots$ be an enumeration of $E(X)$, and consider these edges in turn. Starting with $X_0 := X$, define $X_{n+1} := X_n \smallsetminus \mathring{e}_n$ if this keeps X_{n+1} connected; if not, put $X_{n+1} := X_n$. Finally, let $T := \bigcap_{n \in \mathbb{N}} X_n$.

Since T is closed and contains V, it is still a standard subspace. And T has an edge in every finite cut of G, because X does and its last edge in that cut will never be deleted. So T is connected, by Lemma 8.6.7. But T contains no circle: that would contain an edge, which should have got deleted since deleting an edge from a circle cannot destroy connectedness. □

Proof of Theorem 8.6.9. The implication (ii)→(i) follows from our two lemmas. For (i)→(ii), let G have edge-disjoint topological spanning trees $T_1, \dots, T_k$, and consider a partition P of V into ℓ sets. If there are infinitely many cross-edges, there is nothing to show; so we assume there are only finitely many. For each $i \in \{1, \dots, k\}$, let T_i' be the multigraph of order ℓ which the edges of T_i induce on P.

To establish that G has at least $k(\ell - 1)$ cross-edges, we show that the multigraphs T_i' are connected. If not, then some T_i' has a vertex partition crossed by no edge of T_i. This partition induces a cut of G that contains no edge of T_i. By our assumption that G has only finitely many cross-edges, this cut is finite. By Lemma 8.6.7, this contradicts the connectedness of T_i. □

(1.9) 8.7 The topological cycle space

As a more comprehensive application of our new theory, let us now look at how the cycle space theory of finite graphs extends to locally finite graphs $G = (V, E)$ with infinite circuits and topological spanning trees.

$G = (V, E)$

Every two points of a topological spanning tree T are joined by a unique arc in T: existence follows from Lemma 8.6.5, while uniqueness is proved as for finite graphs. Adding a new edge e to T therefore creates a unique circle in $T \cup e$; its edges form the *fundamental circuit* C_e of e with respect to T. Note that C_e can be infinite. *fundamental circuit C_e*

Similarly, for every edge $f \in E(T)$ the space $T \smallsetminus \mathring{f}$ has exactly two arc-components; the set of edges between these is the *fundamental cut* D_f of T. Since the two arc-components of $T \smallsetminus \mathring{f}$ are closed (Lemma 8.6.6) but disjoint, Lemma 8.6.3 (ii) implies that D_f is finite. *fundamental cut D_f*

As in finite graphs, we have $e \in D_f$ if and only if $f \in C_e$, for all $f \in E(T)$ and $e \in E \smallsetminus E(T)$. Topological spanning trees that are the closure of a normal spanning tree, as in Lemma 8.6.8, are particularly useful in this context: their fundamental circuits and cuts are both finite. (8.6.8)

For locally finite graphs there will be two cycle spaces: the usual 'finitary' one from Chapter 1.9, and a new 'topological' one based on topological circuits. The former will be a subspace of the latter, much as the space of all finite cuts is a subspace of the space of all cuts. These four spaces are cross-related by matroid duality in a surprising way; see the notes and Exercise 123.

Call a family $(D_i)_{i \in I}$ of subsets of E *thin* if no edge lies in D_i for infinitely many i. Let the *thin sum* $\sum_{i \in I} D_i$ of this family be the set of all edges that lie in D_i for an odd number of indices i. The *topological cycle space* $\mathcal{C}(G)$ of G is the subspace of its edge space $\mathcal{E}(G)$ consisting of all thin sums of circuits. *thin sum* *topological cycle space $\mathcal{C}(G)$*

We say that a given set $\mathcal{Z}$ of circuits *generates* $\mathcal{C}(G)$ if every element of $\mathcal{C}(G)$ is a thin sum of elements of $\mathcal{Z}$. For example, the topological cycle space of the ladder in Figure 8.1.3 can be generated by all its squares (the 4-element circuits), or by the infinite circuit consisting of all horizontal edges and all squares but one. Similarly, the 'wild circuit' of Figure 8.6.1 is the thin sum of all the finite face boundaries of that graph, which thus generate it. *generates*

Let us use $\mathcal{C}_{\text{fin}}(G)$ to denote the *finitary cycle space* of G as defined in Chapter 1.9: the (finite) sums of its finite circuits. Clearly $\mathcal{C}_{\text{fin}}(G) \subseteq \mathcal{C}(G)$. We shall see later that $\mathcal{C}_{\text{fin}}(G)$ contains all the finite elements of $\mathcal{C}(G)$, but this is not obvious from the definitions; see Exercise 120. When G is finite, however, clearly $\mathcal{C}_{\text{fin}}(G) = \mathcal{C}(G)$. *finitary cycle space $\mathcal{C}_{\text{fin}}(G)$*

As shown in Chapter 1.9, a finite set of edges of G lies in $\mathcal{C}_{\text{fin}}(G)$ if and only if it meets every cut of G evenly, and the fundamental circuits of any ordinary spanning tree generate $\mathcal{C}_{\text{fin}}(G)$ by finite sums: just copy the proofs given there. For $\mathcal{C}(G)$ we have the following topological analogue:

Theorem 8.7.1. *The following statements are equivalent for every set D of edges of a locally finite connected graph G:*

(i) $D \in \mathcal{C}(G)$;

(ii) *D meets every finite cut F of G in an even number of edges;*

(iii) *D is a thin sum of fundamental circuits of any topological spanning tree of G.*

(8.2.4)
(8.6.3)
(8.6.7)

Proof. The implication (iii)→(i) holds by definition of $\mathcal{C}(G)$ and the fact that G has a topological spanning tree (Lemma 8.6.11).

Let us prove (i)→(ii). By assumption, D is a thin sum of circuits. Only finitely many of these can meet F, so it suffices to show that every circuit meets F evenly. This follows from Lemma 8.6.3 (i): given a circle C in $|G|$, the segments of C between its edges in F (if any) are arcs whose vertices all lie on the same side of the cut F. These sides alternate as we follow C round. Therefore, there is an even number of such arcs, and hence also of edges that C has in F.

It remains to prove (ii)→(iii). Write C_e for the fundamental circuit of an edge $e \notin E(T)$, and D_f for the fundamental cut of an edge $f \in E(T)$. Recall that, by Lemma 8.6.3 (ii), these D_f are finite cuts. We show that

$$D = \sum_{e \in D \setminus E(T)} C_e \,. \tag{$*$}$$

This sum is well defined: since $f \in C_e \Leftrightarrow e \in D_f$ and fundamental cuts are finite, the C_e in this sum form a thin family. To prove $(*)$ we show that $D' := D + \sum_{e \in D \setminus E(T)} C_e = \emptyset$.

Note first that $D' \subseteq E(T)$: any chord of T that lies in D also lies in exactly one of the C_e in the sum. Hence any $f \in D'$ is the unique edge of T, and hence of D', in the finite cut D_f, giving $|D' \cap D_f| = 1$. This is a contradiction, since D meets D_f evenly by (ii), and every C_e does by Lemma 8.6.3. □

Corollary 8.7.2. *$\mathcal{C}(G)$ is generated by finite circuits.*

(8.2.4) *Proof.* Apply Theorem 8.7.1 with the closure of a normal spanning tree, which is a topological spanning tree by Lemma 8.6.8. □

Our second aim in this section is to prove the analogue of Proposition 1.9.1 (ii) for the topological cycle space: that its elements D are not only thin sums but even disjoint unions of circuits. For finite graphs, it was easy to find these circuits greedily: we would 'follow the edges of D round' until a circuit was found, delete it, and repeat.

This will still be our overall strategy when G is infinite. But it is no longer straightforward now to isolate a single circuit from D. For example, without using our knowledge that the edge set D of the wild circle in the graph G of Figure 8.6.1 is a circuit, we can see at once that it must lie in $\mathcal{C}(G)$: it is the thin sum of all the finite circuits bounding a face. Our proof must therefore be able to 'decompose' D into disjoint circuits. Since D itself is the only circuit contained in D, the proof thus has to reconstruct the complicated wild circle just from the information that $D \in \mathcal{C}(G)$. And it has to do so generically, without appealing to the special structure of this particular graph.

Theorem 8.7.3. *For every locally finite graph G, every element of $\mathcal{C}(G)$ is a disjoint union of circuits.*

(1.9.1)
(8.1.2)
(8.6.4)

Proof. We may assume that G is connected, and hence countable. Let $D \in \mathcal{C}(G)$ be given, and enumerate its edges. We inductively construct a sequence of disjoint circuits $C \subseteq D$ each containing the smallest edge in our enumeration of D that is not yet contained in the circuits constructed before. Then all these circuits will form the desired partition of D.

Suppose we have already constructed finitely many disjoint circuits all contained in D. Deleting these edges from D leaves a set D' of edges D', e
that is again in $\mathcal{C}(G)$; let e be its smallest edge in our enumeration of D. We shall find a topological path π between the endvertices of e in the standard subspace that $D' \setminus \{e\}$ spans in $|G|$. By Lemma 8.6.4, the image of π will contain an arc A between these vertices, and $A \cup e$ will be the circle defining our next circuit.

Enumerate the vertices of G as $v_0, v_1, \ldots$, with $e = v_0v_1$. Let $e = v_0v_1$
$S_n := \{v_0, \ldots, v_n\}$. For each $n \geqslant 1$, let G_n be the finite multigraph obtained from G by contracting every component of $G - S_n$ to a vertex, deleting any loops but keeping parallel edges that arise in the contraction. Note that both $V(G_n)$ and $E(G_n)$ are finite, and that G_n
$G[S_n] \subseteq G_n$. Let v'_n denote the vertex of $G_{n-1} - S_{n-1}$ whose branch set V_n contains v_n. v'_n, V_n

We may think of $E(G_n)$ as a subset of $E(G)$. Then the cuts of G_n are also cuts of G. By Theorem 8.7.1, D' meets these evenly; in particular, every vertex of G_n is incident with an even number of edges in D'. Hence $D' \cap E(G_n) \in \mathcal{C}(G_n)$, by Proposition 1.9.1, so G_n contains a cycle through e that has all its edges in D'. Let P_n be the unique v_0–v_1 walk in this cycle that does not contain e and does not repeat any vertices.

Let $\mathcal{V}_n$ be the set of all v_0–v_1 walks in $G_n - e$ in which none of the vertices $v_0, \ldots, v_n$, and hence no edge, occurs more than once. Then $P_n \in \mathcal{V}_n \neq \emptyset$, and $\mathcal{V}_n$ is finite. Every walk $W \in \mathcal{V}_n$ with $n \geqslant 2$ induces a walk $W' \in \mathcal{V}_{n-1}$ consisting of the edges that W has in G_{n-1}, traversed in the same order and direction.[14] Thus, W' arises from W by replacing any subwalk of vertices and edges not in G_{n-1} with v'_n. The vertices of any such subwalk of W will be v_n or vertices of $G_n - S_n$ whose branch set is contained in V_n. By the infinity lemma, there exists a choice of walks $W_n \in \mathcal{V}_n$ such that $W'_n = W_{n-1}$ for all $n \geqslant 2$.

Our next aim is to turn these walks W_n into topological paths $\pi_n \colon [0,1] \to |G_n|$ that traverse them from v_0 to v_1 and reflect their compatibility. We shall define these π_n for $n = 1, 2, \ldots$ in turn, as follows.

For $n = 1$, note that W_1 has exactly two edges: at least two, because e has no parallel edge, and at most two, because every edge of G_1

---

[14] These are well defined: every edge $e \in W$ that is an edge of G_{n-1} has at least one endvertex in S_{n-1}, which either precedes it in W or follows it. In W', this vertex will likewise precede or follow e, respectively.

is adjacent to either v_0 or v_1. Let π_1 map $[0, \frac{1}{3}]$ onto the first edge, $[\frac{1}{3}, \frac{2}{3}]$ to the unique inner vertex of W_1, and $[\frac{2}{3}, 1]$ onto the second edge.

For $n \geqslant 2$, assume inductively that π_{n-1} traverses the edges of W_{n-1} in their given order and direction, and that π_{n-1} 'pauses' at each vertex $v \in G_{n-1} - S_{n-1}$ on W_{n-1} for a non-singleton closed interval $I \subseteq [0, 1]$, mapping I constantly to that vertex. (Thus, if W_{n-1} visits v five times, then $\pi_{n-1}^{-1}(v)$ is a disjoint union of five such intervals.) We start our definition of π_n by letting $\pi_n(\lambda) := \pi_{n-1}(\lambda)$ for all λ with $\pi_{n-1}(\lambda) \in |G_n|$.

Every other $\lambda \in [0, 1]$ satisfies $\pi_{n-1}(\lambda) = v_n'$. These λ form a disjoint union of closed intervals, one for every occurrence of v_n' on W_{n-1}. Recall that W_n arises from W_{n-1} by replacing each occurrence of v_n' by a subwalk of W_n whose vertices are either v_n or vertices of $G_n - S_n$ whose branch set is contained in V_n. For every occurrence of v_n' on W_{n-1}, let π_n on the corresponding interval I with $\pi_{n-1}(I) = \{v_n'\}$ traverse this subwalk of W_n, once more pausing for a non-singleton interval at any vertex that this subwalk has in $G_n - S_n$.

These maps π_n tend to a limit $\pi \colon [0, 1] \to |G|$, defined as follows. Let $\lambda \in [0, 1]$ be given. If $\pi_n(\lambda) \in |G|$ for some n, then $\pi_m(\lambda) = \pi_n(\lambda)$ for all $m > n$, and we let $\pi(\lambda) := \pi_n(\lambda)$. Otherwise $\pi_n(\lambda) \in V(G_n) \setminus S_n$ for all n; let U_n be the branch set of this vertex $u_n := \pi_n(\lambda)$ of G_n in G. By our inductive construction of the maps π_n, we have $U_1 \supseteq U_2 \supseteq \ldots$. Since U_n spans a component $C_n = C_n(\lambda)$ of $G - S_n$, we can find a ray in G that has a tail in each C_n; let $\pi(\lambda)$ be the end ω of this ray. Note that ω, and hence $\pi(\lambda)$, is well defined: every end $\omega' \neq \omega$ is separated from ω by some S_n, and then fails to have a ray in C_n.

For a proof that π is our desired topological v_0–v_1 path in $|G|$, we need to check continuity at every λ. If $\pi(\lambda) = \pi_n(\lambda)$ for some n, then π agrees with π_n also in a small neighbourhood of λ, so this follows from the continuity of π_n. Otherwise $\pi(\lambda)$ is an end, ω say. Then ω has a neighbourhood basis in $|G|$ consisting of open sets $\hat{C}_\epsilon(S_n, \omega)$. Here $C(S_n, \omega)$ is the component $C_n(\lambda)$ defined earlier, since ω has a ray in it.

Now λ is an inner point of an interval $I \subseteq [0, 1]$ which π_n maps to the vertex $u_n = \pi_n(\lambda)$. By construction, $\pi(I) \subseteq \overline{C_n(\lambda)} \subseteq \hat{C}_\epsilon(S_n, \omega)$, completing our continuity proof for π. □

Note that in the special case of Theorem 8.7.3 where the cycle space element considered is the entire edge set of G, Theorem 8.4.2 gives the stronger result that $E(G)$ is a disjoint union of finite circuits.

Corollary 8.7.4. *$\mathcal{C}(G)$ is closed under infinite thin sums.*

Proof. Consider a thin sum $\sum_{i \in I} D_i$ of elements of $\mathcal{C}(G)$. By Theorem 8.7.3, each D_i is a disjoint union of circuits. Together, these form a thin family, whose sum lies in $\mathcal{C}(G)$ and equals $\sum_{i \in I} D_i$. □

8.8 Infinite graphs as limits of finite ones

In the last section we saw how the space $|G|$, for a locally finite graph G, seems to appear as a 'limit' of the finite minors G_n of G obtained by contracting the components left on deleting its first n vertices. We now make this relationship between $|G|$ and the G_n more formal. Clarifying this can help a lot with transferring theorems for finite graphs to infinite ones – which, after all, is the idea behind considering $|G|$ in the first place.

Let $(P, \leqslant)$ be a *directed* partially ordered set, one such that for all p, q there exists an r such that $p \leqslant r$ and $q \leqslant r$. A subset $Q \subseteq P$ is *cofinal* in P if for every $p \in P$ there exists some $q \in Q$ with $p \leqslant q$. *directed* *cofinal*

For every $p \in P$ let X_p be a compact Hausdorff topological space; later, these will represent finite graphs. Assume that we have continuous maps $f_{qp}\colon X_q \to X_p$ for all $q > p$, which are *compatible* in that, whenever $r > q > p$, we have $f_{qp} \circ f_{rq} = f_{rp}$. The family $\mathcal{X} = (\, X_p \mid p \in P\,)$, together with these *bonding maps* f_{qp}, is called an *inverse system*. *inverse system* *bonding maps*

The set X of all $x = (\, x_p \mid p \in P\,)$ with $x_p \in X_p$ and $f_{qp}(x_q) = x_p$ for all $p < q$ in P is the *inverse limit* $X = \varprojlim \mathcal{X}$ of $\mathcal{X}$. We give it the subspace topology from the product space $\prod_{p \in P} X_p$ which, like the X_p, is Hausdorff and compact by Tychonoff's theorem. *inverse limit*

The space $X = \varprojlim \mathcal{X}$ is the intersection, over all $q \in P$, of the sets $X_{<q}$ of all $(\, x_p \mid p \in P\,) \in \prod_p X_p$ that satisfy $f_{qp}(x_q) = x_p$ for all $p < q$. Using the fact that the X_p are Hausdorff and the maps f_{qp} are continuous, one can show that these subsets $X_{<q}$ of $\prod_p X_p$ are closed. Thus, $X = \bigcap_{q \in P} X_{<q}$ is closed in the compact space $\prod_p X_p$, and therefore compact.

As P is directed, the sets $X_{<q}$ have the finite intersection property, as long as the X_p are non-empty. Then $X = \bigcap_q X_{<q}$ is also non-empty:

Lemma 8.8.1. *$X = \varprojlim\, (\, X_p \mid p \in P\,)$ is a compact Hausdorff space. It is non-empty if $X_p \neq \emptyset$ for all $p \in P$.* □

Given a graph $G = (V, E, \Omega)$, consider as $P = P(G)$ the set of all finite partitions of V with only finitely many cross-edges. Letting $p \leqslant q$ whenever q refines p makes P into a directed partially ordered set. For each p, let G/p be the finite multigraph on p whose edges are the cross-edges of p.[15] The vertices of G/p that are non-singleton partition classes are its *dummy vertices*. The other vertices of G/p, those of the form $\{v\}$, we consider to be vertices of G and refer to them as v. $G,\ V, E, \Omega$ $P(G),\ P$ G/p *dummy vertices*

On the compact spaces $X_p := |G/p|$ we have compatible quotient maps $f_{qp}\colon X_q \to X_p$ for $q > p$ which send the vertices of G/q to the vertices of G/p that contain them as subsets; which are the identity on the edges $X_p,\ f_{qp}$

[15] If the partition classes $U \in p$ are connected in G, then G/p is the minor of G obtained by contracting them. But we do not require them to be connected.

of G/q that are also edges of G/p; and which send any other edge of G/q to the dummy vertex of G/p that contains both its endvertices in G/q. Let

$\|G\|$

$$\|G\| := \varprojlim (X_p \mid p \in P),$$

with these f_{qp} as bonding maps.

Theorem 8.8.2. *If G is locally finite and connected, then $\|G\|$ is homeomorphic to $|G|$.*

Proof. As $\|G\|$ is compact and $|G|$ is Hausdorff, it suffices to construct a continuous bijection $\sigma \colon \|G\| \to |G|$. Let $x = (x_p \mid p \in P) \in \|G\|$ be given.

If there exists $p \in P$ such that x_p is not a dummy vertex of G/p, then $x_p \in |G| \setminus \Omega$ and we let $\sigma(x) := x_p$. To see that this is well defined, consider two such points x_p and $x_{p'}$ and pick $q > p, p'$. Then x_q is not a dummy vertex either, and $x_p = x_q = x_{p'}$ by the definition of f_{qp} and $f_{qp'}$.

Suppose now that x_p is a dummy vertex for every p. For every $n \in \mathbb{N}$ let S_n be the set of the first n vertices of G in some fixed enumeration, and let $p_n \in P$ consist of the vertices in S_n as singleton partition classes and the vertex sets of the components of $G - S_n$ as the remaining partition classes. This sequence $p_0, p_1, \ldots$ is cofinal in P, since every $p \in P$ is refined by every p_n with n large enough that all the cross-edges of p have their endvertices in S_n.

As $f_{qp}(x_q) = x_p$ whenever $p = p_m < p_n = q$, the connected vertex sets $U_n = x_{p_n}$ form a descending sequence $U_0 \supseteq U_1 \supseteq \ldots$. It is straightforward to construct a ray R in G that has a tail in $G[U_n]$ for every n. Let ω be the end of R.

For every $p \in P$ the set $U = x_p$ contains every U_n with $p < p_n$ as a subset. As the p_n are cofinal in P, every $G[x_p]$ thus contains a tail of R. Conversely, for every end $\omega' \neq \omega$ there is an n such that $G[U_n]$ contains no ray from ω'. Thus, ω is the unique end of G that has a ray in $G[x_p]$ for every $p \in P$. Let $\sigma(x) := \omega$. This completes the definition of σ.

To see that σ is injective, consider distinct points $x, x' \in \|G\|$, differing in their components $x_p \neq x'_p$ say. If p can be chosen so that one of these is not a dummy vertex of G/p, then clearly $\sigma(x) \neq \sigma(x')$. Otherwise $U = x_p$ and $U' = x'_p$ are disjoint vertex sets in G separated by finitely many edges, and $\sigma(x)$ is an end with a ray in $G[U]$ while $\sigma(x')$ is an end with a ray in $G[U']$. Thus again, $\sigma(x) \neq \sigma(x')$.

To see that σ is surjective, let $x \in |G|$ be given. If x is not an end, choose $p(x) \in P$ so as to contain the vertex x, or the endvertices of the edge containing x, as singleton partition classes. For every $q \geqslant p(x)$ in P let $x_q := x$, and for every $p' < q$ for some such q let $x_{p'} := f_{qp'}(x)$. Then $(x_p \mid p \in P)$ is a well-defined point in $\|G\|$ which σ maps to x.

If x is an end, it has a ray in $G[x_p]$ for exactly one dummy vertex x_p of G/p for every $p \in P$. These satisfy $f_{qp}(x_q) = x_p$ whenever $p < q$, so $(x_p \mid p \in P)$ is a point in $\|G\|$ which σ maps to x.

Let us show that σ is continuous at every point $x = (\, x_p \mid p \in P\,)$ of $\|G\|$. If $\sigma(x)$ is not an end, there exists some $p(x) \in P$ such that $\sigma(x) = x_{p(x)}$, which is a point in $X_{p(x)}$ but not a dummy vertex. Then every basic open neighbourhood O of $\sigma(x)$ in $|G|$ is also a basic neighbourhood of this same point $x_{p(x)}$ in $X_{p(x)}$. Then the set $\prod_{p \in P} O_p$ with $O_{p(x)} = O$ and $O_p = X_p$ for all $p \neq p(x)$ is a basic open neighbourhood of x in $\prod_p X_p$. Its intersection with $\|G\|$ is an open neighbourhood of x in $\|G\|$ which σ maps to O.

If $\sigma(x)$ is an end, ω say, consider any basic open neighbourhood $O = \hat{C}_\epsilon(S, \omega)$ of ω in $|G|$. Let $p(\omega) \in P$ be the partition of V into the vertex sets of the components of $G - S$ and the singletons in S. Then $V(C)$ is a dummy vertex of $G/p(\omega)$; call it $x_{p(\omega)}$. Let $O_{p(\omega)} \subseteq X_{p(\omega)}$ consist of $x_{p(\omega)}$ and the inner points in O of any C–S edges; these are also points of $X_{p(\omega)}$. As earlier, x has a basic open neighbourhood $\prod_p O_p$ in $\prod_p X_p$ with $O_p = X_p$ for all $p \neq p(\omega)$, whose intersection with $\|G\|$ maps to O under σ. $\square$

Note that our proof did not use that $|G|$ is compact: we reobtain Proposition 8.6.1 as a corollary.

In the proof of Theorem 8.8.2 we found it convenient to work with a cofinal sequence in P instead of the entire set P. This is justified more generally by the following easy lemma:

Lemma 8.8.3. *Let $(\, X_p \mid p \in P\,)$ be an inverse system of compact spaces, let $Q \subseteq P$ be cofinal in P, and consider $(\, X_p \mid p \in Q\,)$ with the same bonding maps. Mapping every point $(\, x_p \mid p \in P\,)$ to its restriction $(\, x_p \mid p \in Q\,)$ then defines a homeomorphism from $\varprojlim\, (\, X_p \mid p \in P\,)$ to $\varprojlim\, (\, X_p \mid p \in Q\,)$.* $\square$

By Theorem 8.8.2 and this lemma, our familiar $|G|$ for locally finite G is the inverse limit of the finite contraction minors G_n of G defined as in Section 8.6. Indeed, for the cofinal sequence $p_0, p_1, \ldots$ in P defined in the proof of the theorem, we have $G_n = G/p_n$, and by the lemma $|G|$ is the inverse limit of the corresponding compact spaces X_{p_n}.

Just like $|G|$ itself, every standard subspace X' of $X = |G|$ can be obtained as an inverse limit of finite multigraphs. Indeed, the projections $f_p \colon X \to X_p$ defined by $(\, x_p \mid p \in P\,) \mapsto x_p$ are continuous, so their images $X'_p \subseteq X_p$ of X' are compact since X' is, and the f_{qp} send X'_q to X'_p. Thus, $(\, X'_p \mid p \in P\,)$ is an inverse system with bonding maps $f'_{qp} := f_{qp} \restriction X'_q$, and $X' = \varprojlim\, (\, X'_p \mid p \in P\,)$.

More typically, we would like to *find* a standard subspace X' with certain desired properties – for example, a topological spanning tree. We can then try to construct some X'_p whose inverse limit is X'. It may not be straightforward, however, to find such compatible X'_p for all $p \in P$. Here, Lemma 8.8.3 can help: it is only necessary to find them for all p in some cofinal $Q \subseteq P$. For example, we can construct spanning trees inductively in all the G_n by expanding a dummy vertex in the tree $T_n \subseteq G_n$ to a star in $T_{n+1} \subseteq G_{n+1}$. Then our given bonding maps $X_{p_n} \to X_{p_m}$ will map the subspace X'_{p_n} induced by T_n to that induced by T_m, and these X'_{p_n} will have a topological spanning tree in $X = |G|$ as their inverse limit. This construction is possible only because the partition classes of the p_n are connected in G; we could not perform it on all of $P(G)$.

Arcs and circles in $|G|$, or in a standard subspace, can be obtained easily by applying the following *lifting lemma* with $Y = [0,1]$ or $Y = S^1$. Let $(\, X_p \mid p \in P\,)$ be any inverse system of compact spaces, with bonding maps f_{qp} say, and let X be its inverse limit. Let Y be a topological space with continuous *compatible* maps $g_p \colon Y \to X_p$: maps that commute with the f_{qp} in that $g_p = f_{qp} \circ g_q$ whenever $p < q$. Let us call the family $(\, g_p \mid p \in P\,)$ *eventually injective* if for all distinct $y, y' \in Y$ there exists some $p \in P$ with $g_p(y) \neq g_p(y')$.

compatible maps

lifting lemma

Lemma 8.8.4. *There is a unique continuous map $g \colon Y \to X$ that commutes with the projections $f_p \colon X \to X_p$ in that $g_p = f_p \circ g$ for all $p \in P$. If the g_p are eventually injective, then g is injective.* □

For example, suppose we wish to find an arc in X between some points x and y. We can find a topological x–y path $g \colon [0,1] \to X$ by finding topological $f_p(x)$–$f_p(y)$ paths $g_p \colon [0,1] \to X_p$ that commute with the f_{qp}. If we can make these g_p eventually injective, then g will be injective, and its image will be the desired arc.

Similarly, if we can find compatible circles $g_p \colon S^1 \to X_p$ that are eventually injective, whose images contain all the vertices of G/p, and which commute with the f_{qp}, then g will define a *Hamilton circle* of G, a circle in $|G|$ that traverses every vertex.

Hamilton circle

Exercises

1.$^-$ Show that a connected graph is countable if all its vertices have countable degrees.

2.$^-$ Given countably many sequences $\sigma^i = s^i_1, s^i_2, \dots$ $(i \in \mathbb{N})$ of natural numbers, find one sequence $\sigma = s_1, s_2, \dots$ that beats every σ^i eventually, i.e. such that for every i there exists an $n(i)$ such that $s_n > s^i_n$ for all $n \geqslant n(i)$.

3.$^-$ Can a countable set have uncountably many subsets whose intersections have finitely bounded size?

4.$^-$ Let T be an infinite rooted tree. Show that every ray in T has an increasing tail, that is, a tail whose sequence of vertices increases in the tree-order associated with T and its root.

5.$^-$ Let G be an infinite graph and $A, B \subseteq V(G)$. Show that if no finite set of vertices separates A from B in G, then G contains an infinite set of disjoint A–B paths.

6.$^-$ In Proposition 8.1.1, the existence of a spanning tree was proved using Zorn's lemma 'from below', to find a maximal acyclic subgraph. For finite graphs, one can also use induction 'from above', to find a minimal spanning connected subgraph. What happens if we apply Zorn's lemma 'from above' to find such a subgraph?

For the next two exercises it may help to consider the cycle space of the given graph, defined as for finite graphs in Chapter 1.9.

7.$^-$ Show that if a graph has a spanning tree with infinitely many chords then all its spanning trees have infinitely many chords.

8. Show that if a graph contains infinitely many distinct cycles then it contains infinitely many edge-disjoint cycles.

9. Let G be a countable infinitely connected graph. Show that G has, for every $k \in \mathbb{N}$, an infinitely connected spanning subgraph of girth at least k.

10. Construct, for any given $k \in \mathbb{N}$, a planar k-connected graph. Can you construct one whose girth is also at least k? Can you construct an infinitely connected planar graph?

11. Theorem 8.1.3 implies that there exists an $\mathbb{N} \to \mathbb{N}$ function f_χ such that, for every $k \in \mathbb{N}$, every infinite graph of chromatic number at least $f_\chi(k)$ has a finite subgraph of chromatic number at least k. (E.g., let f_χ be the identity on $\mathbb{N}$.) Find similar functions f_δ and f_κ for the minimum degree and connectivity, or show that no such functions exist.

12. Let k be an integer and $\kappa \geqslant \aleph_1$ a regular cardinal. Show that every k-connected graph of order κ contains $TK_{k,\kappa}$ as a topological minor.

13. Let $k \in \mathbb{N}$, and let a, b be two vertices in a graph G.
    (i) Show that there are only finitely many minimal a–b separators of order k, and only finitely many minimal a–b cuts of order k. (A cut is an *a–b cut* if it separates a from b.)
    (ii) Deduce that every edge of G lies in only finitely many bonds of k edges.

14.$^+$ Extend Exercise 42 of Chapter 1 to infinite graphs.

15. Use the infinity lemma to show that a rayless connected graph of minimum degree d has a finite subgraph of minimum degree d.

16. A theorem of Halin says that every graph of chromatic number $\alpha \geqslant \aleph_0$ contains a TK^β for every cardinal $\beta < \alpha$. Prove this for $\alpha = \aleph_0$.

17.$^-$ Prove Theorem 8.1.3 for arbitrary graphs using the generalized infinity lemma from Appendix A.

18. Give a proof of Theorem 8.1.3 for countable graphs that is based on the fact that, in this case, the topological space X defined in the third proof of the theorem is sequentially compact. (Thus, every infinite sequence of points in X has a convergent subsequence: there is an $x \in X$ such that every neighbourhood of x contains a tail of the subsequence.)

19.$^-$ Extend Nash-Williams's tree covering theorem (2.4.3) to infinite graphs.

20.$^+$ Extend the packing-covering theorem (2.4.4) to infinite graphs.

21.$^+$ For every $k \in \mathbb{N}$, construct a k-connected locally finite graph such that the deletion of the edge set of any cycle disconnects that graph. Deduce that the tree packing theorem (2.4.1) of Nash-Williams and Tutte fails for infinite graphs.

(Hint. Start with a k-connected finite graph G_0. If G_0 has a cycle C such that deleting $E(C)$ does not disconnect G_0, graft some more copies of G_0 on to $E(C)$ to give C that property. Continue inductively.)

22. Derive the generalized infinity lemma and the compactness principle in Appendix A from each other.

23. In the text, the unfriendly partition conjecture is proved for locally finite graphs, using the infinity lemma.

    (i) Give an alternative proof using the compactness principle from Appendix A.

    (ii) The proof in the text, by the infinity lemma, required a modification of the statement. Is this still necessary? Which step in the proof using the compactness principle reflects the requirement in the infinity lemma that every admissible partial solution must induce an admissible solution on a smaller substructure? Where is the local finiteness used?

24. (i) Prove the unfriendly partition conjecture for countable graphs with all degrees infinite.

    (ii) Can you adapt the proof to cover also those countable graphs that have finitely many vertices of finite degree?

25. Rephrase Gallai's partition theorem of Exercise 55, Chapter 1, in terms of degrees, and extend the equivalent version to locally finite graphs.

26. Prove Theorem 8.4.10 for locally finite graphs. Does your proof extend to arbitrary countable graphs?

27. Extend the marriage theorem to locally finite graphs, but show that it fails for countable graphs with infinite degrees.

28. Show that every locally finite factor-critical graph is finite.

29.$^+$ Show that a locally finite graph G has a 1-factor if and only if, for every finite set $S \subseteq V(G)$, the graph $G - S$ has at most $|S|$ odd (finite) components. Find a counterexample that is not locally finite.

30.$^+$ Extend Kuratowski's theorem to countable graphs.

31.$^-$ A vertex $v \in G$ is said to *dominate* an end ω of G if any of the following three assertions holds; show that they are equivalent.

(i) For some ray $R \in \omega$ there is an infinite v–$(R-v)$ fan in G.

(ii) For every ray $R \in \omega$ there is an infinite v–$(R-v)$ fan in G.

(iii) No finite subset of $V(G-v)$ separates v from a ray in ω.

32. Show that a graph G contains a $TK^{\aleph_0}$ if and only if some end of G is dominated by infinitely many vertices.

33. Let G be a *finitely separable* graph, one in which any two vertices can be separated by finitely many edges.

(i) Show that any two ends in G that cannot be separated by finitely many edges are dominated by a common vertex.

(ii) Is the assumption of finite separability necessary for (i) to hold?

34. Construct a countable graph with uncountably many thick ends. Can you find a locally finite such graph?

35. Show that a locally finite connected vertex-transitive graph has exactly 0, 1, 2 or infinitely many ends.

36.$^+$ Show that the automorphisms of a graph $G = (V, E)$ act naturally on its ends, i.e., that every automorphism $\sigma: V \to V$ can be extended to a map $\sigma: \Omega(G) \to \Omega(G)$ such that $\sigma(R) \in \sigma(\omega)$ whenever R is a ray in an end ω. Prove that, if G is connected, every automorphism σ of G fixes a finite set of vertices or an end. If σ fixes no finite set of vertices, can it fix more than one end? More than two?

37.$^-$ Show that a locally finite spanning tree of a graph G contains a ray from every end of G.

38. A ray in a graph *follows* another ray if the two have infinitely many vertices in common. Show that if T is a normal spanning tree of G then every ray of G follows a unique normal ray of T.

39. Use normal spanning trees to show that a countable connected graph has either countably many or continuum many ends.

40. Show that the following assertions are equivalent for connected countable graphs G.

(i) G has a locally finite spanning tree.

(ii) For no finite separator $X \subseteq V(G)$ in G does $G - X$ have infinitely many components.

41. Show that every (countable) planar 3-connected graph has a locally finite spanning tree.

42. Prove the following infinite version of the Erdős-Pósa theorem: an infinite graph G either contains infinitely many disjoint cycles or it has a finite set Z of vertices such that $G - Z$ is a forest.

43. Let G be a connected graph. Call a set $U \subseteq V(G)$ *dispersed* if every ray in G can be separated from U by a finite set of vertices. (In the topology of Section 8.6, these are precisely the closed subsets of $V(G)$.)

    (i) Prove that G has a normal spanning tree if and only if $V(G)$ is a countable union of dispersed sets. Show that in this case we can choose any vertex of G as the normal spanning tree's root.

    (ii) Deduce that if G has a normal spanning tree then so does every connected minor of G.

44.$^+$ Show that a connected graph G has a normal spanning tree if it does not contain a subdivision of a *fat* $K^{\aleph_0}$, one in which every edge has been replaced by uncountably many parallel edges.

    (Hint. Apply induction on $|G|$. Given G, construct an increasing sequence $(\, G_\beta \mid \beta < \alpha \,)$ of smaller graphs exhausting G so that the endvertices of any G_β-path in G are joined in G_β by infinitely many independent paths and in G by uncountably many independent paths. Combine suitable normal spanning trees of the components of $G_{\beta+1} - G_\beta$ with the inductively given normal spanning tree of G_β to form a normal spanning tree of $G_{\beta+1}$.)

45.$^+$ Prove Theorem 8.2.5 (ii).

46. (i) Prove that if a given end of a graph contains k disjoint rays for every $k \in \mathbb{N}$ then it contains infinitely many disjoint rays.

    (ii)$^+$ Prove that if a given end of a graph contains k edge-disjoint rays for every $k \in \mathbb{N}$ then it contains infinitely many edge-disjoint rays.

47. Prove that if a graph contains k disjoint double rays for every $k \in \mathbb{N}$ then it contains infinitely many disjoint double rays.

48. Show that, in the ubiquity conjecture, the host graphs G considered can be assumed to be locally finite too.

49. Show that the modified comb below is not ubiquitous with respect to the subgraph relation. Does it become ubiquitous if we delete its 3-star on the left?

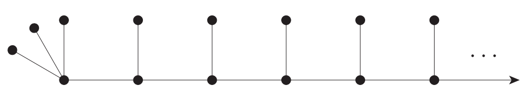

50. Show that every locally finite tree T is minor-ubiquitous in countable graphs, by proving and combining the following statements:

    (i) The $\mathbb{N} \times \mathbb{N}$ grid H satisfies $\aleph_0 H \preccurlyeq H$.

    (ii) $T \preccurlyeq H$.

    (iii)$^+$ Assuming that $nT \preccurlyeq G$ for all $n \in \mathbb{N}$, show that $\aleph_0 T \preccurlyeq G$ unless G has a thick end.

51. Imitate the proof of Theorem 8.2.6 to find a function $f\colon \mathbb{N} \to \mathbb{N}$ such that whenever an end ω of a graph G contains $f(k)$ disjoint rays, G contains a subdivision of the $k \times \mathbb{N}$ hexagonal grid whose rays all belong to ω.

52. Show that there is no universal locally finite connected graph for the subgraph relation.

53. Construct a universal locally finite connected graph for the minor relation. Is there one for the topological minor relation?

54.$^-$ Show that each of the following operations performed on the Rado graph R leaves a graph isomorphic to R:

(i) taking the complement, i.e. changing all edges into non-edges and vice versa;

(ii) deleting finitely many vertices;

(iii) changing finitely many edges into non-edges or vice versa;

(iv) changing all the edges between a finite vertex set $X \subseteq V(R)$ and its complement $V(R) \smallsetminus X$ into non-edges, and vice versa.

55. (i)$^-$ If we delete infinitely many vertices of the Rado graph R leaving an infinite rest, does the rest necessarily induce a copy of R?

(ii) Does R admit a vertex partition into two copies of itself?

(iii) Construct a locally finite tree that has a vertex partition into two copies of itself.

56.$^-$ Prove that the Rado graph is homogeneous.

57. Show that a homogeneous countable graph is determined uniquely, up to isomorphism, by the class of (the isomorphism types of) its finite subgraphs.

58. Recall that subgraphs $H_1, H_2, \ldots$ of a graph G are said to *partition* G if their edge sets form a partition of $E(G)$. Show that the Rado graph can be partitioned into any given countable set of countable locally finite graphs, as long as each of them contains at least one edge.

59. A linear order is called *dense* if between any two elements there lies a third.

(i) Find, or construct, a countable dense linear order that has neither a maximal nor a minimal element.

(ii) Show that this order is unique, i.e. that every two such orders are order-isomorphic. (Definition?)

(iii) Show that this ordering is universal among the countable linear orders. Is it homogeneous? (Supply appropriate definitions.)

60. Given a bijection f between $\mathbb{N}$ and $[\mathbb{N}]^{<\omega}$, let G_f be the graph on $\mathbb{N}$ in which $u, v \in \mathbb{N}$ are adjacent if $u \in f(v)$ or vice versa. Prove that all such graphs G_f are isomorphic.

61. (for set theorists) Show that, given any countable model of set theory, the graph whose vertices are the sets and in which two sets are adjacent if and only if one contains the other as an element, is the Rado graph.

62. Find two proofs that every infinitely edge-connected graph has infinitely many edge-disjoint spanning trees:
   (i)$^-$ Apply Theorem 8.4.1.
   (ii)$^+$ Find a direct argument.

63.$^-$ Given sets A, B of vertices in a graph G, show that either G contains infinitely many edge-disjoint A–B paths or there is a finite set of edges separating A from B in G.

64. Let G be a locally finite graph. Let us say that a finite set S of vertices *separates* two ends ω and ω' if $C(S,\omega) \neq C(S,\omega')$. Use Proposition 8.4.3 to show that if ω can be separated from ω' by $k \in \mathbb{N}$ but no fewer vertices, then G contains k disjoint double rays each with one tail in ω and one in ω'. Is the same true for all graphs that are not locally finite?

65.$^+$ Prove the following more structural version of Exercise 46 (i). Let ω be an end of a graph G. Show that either G contains a $TK^{\aleph_0}$ with all its rays in ω, or there are disjoint finite sets $S_0, S_1, \ldots$ such that, if C_i is the component of $G - (S_0 \cup S_i)$ that contains a tail of every ray in ω, we have for all $i < j$ that $C_i \supsetneq C_j$ and $G[S_i \cup C_i]$ contains $|S_i|$ disjoint S_i–S_{i+1} paths for all $i \geqslant 1$.

66.$^+$ Is there a planar $\aleph_0$-regular graph all whose ends have infinite vertex-degree?

67.$^-$ Let A, B be two vertex sets in a locally finite connected graph G. Can there be an infinite sequence $\mathcal{P}_1, \mathcal{P}_2, \ldots$ of disjoint A–B paths such that each $\mathcal{P}_{n+1}$ arises from $\mathcal{P}_n$ by applying an alternating walk, and such that some edge $e \in G$ lies in $E[\mathcal{P}_n]$ for infinitely many n but not in $E[\mathcal{P}_n]$ for infinitely many other n?

68. Construct an example of a small limit of large waves. Can you find a locally finite one?

69.$^+$ Prove Theorem 8.4.4 for trees.

70.$^+$ Prove Pym's theorem (8.4.9).

71. (i)$^-$ Prove the naive extension of Dilworth's theorem to arbitrary infinite posets P: if P has no antichain of order $k \in \mathbb{N}$, then P can be partitioned into fewer than k chains. (A proof for countable P will do.)

   (ii)$^-$ Find a poset that has no infinite antichain and no partition into finitely many chains.

   (iii) For posets without infinite chains, deduce from Theorem 8.4.10 the following Erdős-Menger-type extension of Dilworth's theorem: every such poset has a partition $\mathcal{C}$ into chains such that some antichain meets all the chains in $\mathcal{C}$.

72. Let G be a countable graph in which for every partial matching there is an augmenting path.
   (i) Find an example of G and a sequence $M_0, M_1, \ldots$ of partial matchings, each obtained from the previous as its symmetric difference with the edge set of an augmenting path, so that for every edge e of G we have $e \in M_{n+1} \setminus M_n$ for infinitely many n.

(ii) Show that for every partial matching M there exists a sequence as in (i) such that $\bigcup_m \bigcap_{n>m} M_n$ is the edge set of a 1-factor.

73. Find an uncountable graph in which every partial matching admits an augmenting path (finite or infinite) but which has no 1-factor.

74.$^-$ Let G be a countable graph whose finite subgraphs are all perfect. Show that G is weakly perfect but not necessarily perfect.

75.$^+$ Let G be the incomparability graph of the binary tree. (Thus, $V(G) = V(T_2)$, and two vertices are adjacent if and only if they are incomparable in the tree-order of T_2.) Show that G is perfect but not strongly perfect.

76.$^+$ (i) Show that the vertices of any infinite connected locally finite graph can be enumerated in such a way that every vertex is adjacent to some later vertex.

(ii) Characterize the class of all these graphs, countable but not necessarily locally finite, by their separation properties.

77. Show that a tree has a rank as defined in the second paragraph after the proof of Proposition 8.5.1 if and only if it is recursively prunable, and that it has rank α if and only if α is the maximum of its pruning labels.

78. Let G be a rayless graph, of rank α say, and let U be a finite set of vertices witnessing this, of minimal order. Show that U is unique.

79. (i) Construct a countable tree that has rank ω in the ranking of rayless graphs. Can you find one such tree that contains all the others?

(ii)$^+$ Is there a tree of rank ω that is a subtree of every such tree?

80. A graph $G = (V, E)$ is called *bounded* if for every vertex labelling $\ell \colon V \to \mathbb{N}$ there exists a function $f \colon \mathbb{N} \to \mathbb{N}$ that exceeds the labelling along any ray in G eventually. (Formally: for every ray $v_1 v_2 \dots$ in G there exists an n_0 such that $f(n) > \ell(v_n)$ for every $n > n_0$.) Prove the following assertions:

(i) The ray is bounded.

(ii) Every locally finite connected graph is bounded.

(iii) A countable tree is bounded if and only if it contains no subdivision of the $\aleph_0$-regular tree $T_{\aleph_0}$.

81.$^+$ Let T be a tree with root r, and let $\leqslant$ denote the tree-order on $V(T)$ associated with T and r. Show that T contains no subdivision of the $\aleph_1$-regular tree $T_{\aleph_1}$ if and only if T has an ordinal labelling $t \mapsto o(t)$ such that $o(t) \geqslant o(t')$ whenever $t < t'$ and no more than countably many vertices of T have the same label.

82. Let G be a countable connected graph with vertices $v_0, v_1, \dots$. For every $n \in \mathbb{N}$ write $S_n := \{v_0, \dots, v_{n-1}\}$. Prove the following statements:

(i) For every end ω of G there is a unique sequence $C_0 \supseteq C_1 \supseteq \dots$ of components C_n of $G - S_n$ such that $C_n = C(S_n, \omega)$ for all n.

(ii) For every infinite sequence $C_0 \supseteq C_1 \supseteq \dots$ of components C_n of $G - S_n$ there exists a unique end ω such that $C_n = C(S_n, \omega)$ for all n.

83. Let G be a graph, $U \subseteq V(G)$, and $R \in \omega \in \Omega(G)$. Show that G contains a comb with spine R and teeth in U if and only if $\omega \in \overline{U}$.

84. Given graphs $H \subseteq G$, let $\eta \colon \Omega(H) \to \Omega(G)$ assign to every end of H the unique end of G containing it as a subset (of rays). For the following questions, assume that H is connected and $V(H) = V(G)$.

(i) Show that η need not be injective. Must it be surjective?

(ii) Investigate how η relates the subspace $\Omega(H)$ of $|H|$ to its image in $|G|$. Is η always continuous? Is it open? Do the answers to these questions change if η is known to be injective?

(iii) A spanning tree is called *end-faithful* if η is bijective, and *topologically end-faithful* if η is a homeomorphism. Show that every connected countable graph has a topologically end-faithful spanning tree.

The *end space* of a graph G is the subspace $\Omega(G)$ of $|G|$.

85. Consider the end space Ω of the binary tree T_2 shown in Figure 8.1.4, in which its vertices are the finite 0–1 sequences.

(i) Show that Ω is homeomorphic to $2^{\mathbb{N}}$, where $2 = \{0, 1\}$ carries the discrete topology and $2^{\mathbb{N}}$ the product topology.

(ii) Identify in Ω every two ends whose infinite binary sequences encode the same rational. Show that the resulting quotient space of Ω is homeomorphic to the real interval $[0, 1]$.

86. Above every horizontal edge of the plane graph shown in Figure 8.6.1 add infinitely many horizontal edges in the plane, so as to turn every pair of rays whose associated 0–1 sequences define the same rational number into a ladder. Prove or disprove that the end space of the resulting graph is homeomorphic to $[0, 1]$.

87. A compact metric space is a *Cantor set* if the singletons are its only connected subsets and every point is an accumulation point.

(i) Characterize the trees whose end space is a Cantor set.

(ii) Show that the end space of a connected locally finite graph is a subset of a Cantor set.

88. (i) Show that if H is a contraction minor of G with finite branch sets, then the end spaces of H and G are homeomorphic.

(ii) Let T_n denote the n-ary tree, the rooted tree in which every vertex has exactly n successors. Show that all these trees have homeomorphic end spaces.

89. Give an independent proof of Proposition 8.6.1 using sequential compactness and the infinity lemma.

90.$^+$ (for topologists) In a locally compact, connected, and locally connected Hausdorff space X, consider sequences $U_1 \supseteq U_2 \supseteq \ldots$ of open, non-empty, connected subsets with compact frontiers such that $\bigcap_{i \in \mathbb{N}} \overline{U_i} = \emptyset$. Call such a sequence *equivalent* to another such sequence if every set of one sequence contains some set of the other sequence and vice versa. Note that this is indeed an equivalence relation, and call its classes the *Freudenthal ends* of X. Now add these to the space X, and define a natural topology on the extended space $\hat{X}$ that makes it homeomorphic to $|X|$ if X is a graph, by a homeomorphism that is the identity on X.

91.$^+$ Let G be a connected countable graph that is not locally finite. Show that $|G|$ is not compact, but that $\Omega(G)$ is compact if and only if for every finite set $S \subseteq V(G)$ only finitely many components of $G - S$ contain a ray.

92.$^+$ Let G be a connected graph. Assuming that G has a normal spanning tree, define a metric on $|G|$ that induces its usual topology. Conversely, use Exercise 43 to show that if $V \cup \Omega \subseteq |G|$ is metrizable then G has a normal spanning tree.

93. Find a graph G for which $|G|$ is not metrizable.

(Hint. Rather than thinking of metrics directly, recall some properties of metric spaces, and construct a graph G without such a property.)

A topological space X is *locally connected* if for every $x \in X$ and every neighbourhood U of x there is an open connected neighbourhood $U' \subseteq U$ of x. A *continuum* is a compact, connected Hausdorff space. By a theorem of general topology, every locally connected metric continuum is arc-connected.

94.$^+$ Show that, for G connected and locally finite, every connected standard subspace of $|G|$ is locally connected. Using the theorem cited above, deduce Lemma 8.6.5.

95.$^+$ Prove Lemma 8.6.6 directly, without relying on Lemma 8.6.5.

96. Let G be a locally finite graph, and X a standard subspace of $|G|$ spanned by a set of at least two edges. Show that X is a circle if and only if, for every two distinct edges $e, e' \in E(X)$, the subspace $X \smallsetminus \mathring{e}$ is connected but $X \smallsetminus (\mathring{e} \cup \mathring{e}')$ is disconnected.

97. Does every infinite locally finite 2-connected graph contain an infinite circuit? Does it contain an infinite bond?

98. Consider a locally finite graph.

(i) Show that every infinite circuit meets some infinite bond in exactly one edge.

(ii) Show that every infinite bond meets some infinite circuit in exactly one edge.

99. Show that the union of all the edges contained in an arc or circle C in $|G|$ is dense in C.

100.$^+$ Prove that the circle shown in Figure 8.6.1 is really a circle, by exhibiting a homeomorphism with S^1.

101. Every arc induces on its points a linear ordering inherited from $[0,1]$. Call an arc in $|G|$ *wild* if it induces on some subset of its vertices the ordering of the rationals: between every two there lies another. Show that every arc containing uncountably many ends is wild.

102. Find a locally finite graph G with a connected standard subspace of $|G|$ that is the closure of a union of disjoint circles.

103. Show that, for G locally finite, a closed standard subspace C of $|G|$ is a circle in $|G|$ if and only if C is connected, every vertex in C is incident with exactly two edges in C, and every end in C has topological degree 2.

104. Let T be a locally finite tree. Construct a continuous map $\sigma\colon [0,1] \to |T|$ that maps 0 and 1 to the root and traverses every edge exactly twice, once in each direction. (Formally: define σ so that every inner point of an edge is the image of exactly two points in $[0,1]$.)

(Hint. Define σ as a limit of similar maps σ_n for finite subtrees T_n.)

105. Let G be a connected locally finite graph. Show that the following assertions are equivalent for a spanning subgraph T of G:

(i) $\overline{T}$ is a topological spanning tree of $|G|$;

(ii) T is edge-maximal with $\overline{T}$ containing no circle;

(iii) T is edge-minimal with $\overline{T}$ a connected subspace of $|G|$.

106.$^-$ (i) Observe that a topological spanning tree need not be homeomorphic to a tree. Is it homeomorphic to the space $|T|$ for a suitable tree T?

(ii) Find a graph G with an ordinary spanning tree T whose closure in $|G|$ is not arc-connected.

107. Let T be an end-faithful spanning tree of a locally finite graph G. (Such trees are defined in Exercise 84.) Is $\overline{T}$ a topological spanning tree of $|G|$?

108. Let G be locally finite and connected, with vertices $v_0, v_1, \ldots$ say. Let G_n be the minor of G obtained by contracting every component of $G - \{v_0, \ldots, v_n\}$ to a vertex. Construct spanning trees T_n of G_n so that $\bigcup_n E(T_n)$ is the edge set of a topological spanning tree of G.

109. Let F be a set of edges in a locally finite connected graph $G = (V, E)$.

(i) Show that F is a circuit if and only if F is not contained in the edge set of any topological spanning tree of G and is minimal with this property.

(ii) Show that F is a finite bond if and only if F meets the edge set of every topological spanning tree of G and is minimal with this property.

110. Extend Exercise 51 of Chapter 1 to characterizations of the bonds, and of the finite bonds, in a locally finite connected graph.

111.$^+$ Prove a topological tree-packing theorem for standard subspaces X of locally finite graphs. Here, a *topological spanning tree of* X is a connected closed subspace of X that contains all its vertices but no circle.

112. To show that Theorem 3.2.6 does not generalize to infinite graphs with its finitary cycle space, construct a 3-connected locally finite planar graph with a separating cycle that is not a finite sum of non-separating induced cycles. Can you find an example where even infinite thin sums of finite non-separating induced cycles do not generate all separating cycles?

113.$^-$ As a converse to Theorem 8.7.1 (iii), show that the fundamental circuits of an ordinary spanning tree T of a locally finite graph G do not generate $\mathcal{C}(G)$ unless $\overline{T}$ is a topological spanning tree.

114.$^-$ Explain why Theorem 8.7.3 is needed in the proof of Corollary 8.7.4: can't we just combine the constituent sums of circuits for the D_i (from our assumption that $D_i \in \mathcal{C}(G)$) into one big family?

115. Deduce Corollary 8.7.4 from Theorem 8.7.1, not using Theorem 8.7.3.

116. If a finite set D of edges meets every finite cut of G evenly, must it also meet every infinite cut evenly?

117. Prove Theorem 8.7.3 by the method used to prove Theorem 8.6.9.

118. Prove Lemma 8.6.5 by the method used to prove Theorem 8.7.3.

For the next ten exercises, let G be a locally finite connected graph. Let $\mathcal{C} = \mathcal{C}(G)$, and define the *cut space* $\mathcal{B} = \mathcal{B}(G)$ of G as in Chapter 1.9. Note that cuts may now be infinite. Define 'generate' for cuts as for circuits, allowing infinite thin sums. Given a set $\mathcal{F} \subseteq \mathcal{E}(G)$, write $\mathcal{F}_{\text{fin}} := \{ F \in \mathcal{F} : |F| < \infty \}$, $\mathcal{F}^\perp := \{ D \in \mathcal{E}(G) : |D \cap F| \in 2\mathbb{N} \ \ \forall F \in \mathcal{F} \}$ and $(\mathcal{F}_{\text{fin}})^\perp =: \mathcal{F}_{\text{fin}}^\perp$.

119. Show that $\mathcal{C}$ and $\mathcal{B}$ are closed in the edge space $\mathcal{E} = \{0,1\}^E$ of G if $\{0,1\}$ carries the discrete topology and $\{0,1\}^E$ the product topology.

120. Show that the definition of $\mathcal{C}_{\text{fin}}$ given above coincides with that given in the text: that the finite elements of $\mathcal{C}$ are finite sums of finite circuits.

121. (i) Show that the fundamental circuits of any ordinary spanning tree of G generate $\mathcal{C}_{\text{fin}}$ by finite sums, but that they need not generate $\mathcal{C}$.

(ii) Show that the fundamental cuts of any topological spanning tree of G generate $\mathcal{B}_{\text{fin}}$ by finite sums, but that they need not generate $\mathcal{B}$.

122. (i)$^-$ Show that $\mathcal{B}$ is a subspace of $\mathcal{E}(G)$ generated by finite cuts.

(ii) Show that every cut is a disjoint union of bonds.

(iii)$^+$ Show that the fundamental cuts of any ordinary spanning tree of G generate $\mathcal{B}$.

(iv)$^+$ Show that $\mathcal{B}$ is closed under infinite thin sums.

123. (i)$^-$ Find in this book a proof, or sketch of a proof, for each of the following two statements: $\mathcal{C} = \mathcal{B}_{\text{fin}}^\perp$ and $\mathcal{B} = \mathcal{C}_{\text{fin}}^\perp$.

(ii)$^+$ Show that $\mathcal{B}^\perp = \mathcal{C}_{\text{fin}}$ and, if G is 2-edge-connected, $\mathcal{C}^\perp = \mathcal{B}_{\text{fin}}$.

124.$^+$ Write $\hat{\mathcal{C}}$ for the set of circuits in G, and $\hat{\mathcal{B}}$ for the set of bonds.

(i) Show that $\hat{\mathcal{C}}^\perp_{\text{fin}} = \mathcal{C}^\perp_{\text{fin}}$ and $\hat{\mathcal{B}}^\perp_{\text{fin}} = \mathcal{B}^\perp_{\text{fin}}$.

(ii) Show that every element of $\hat{\mathcal{C}}^\perp$ is a disjoint union of finite bonds, and that every element of $\hat{\mathcal{B}}^\perp$ is a disjoint union of finite circuits.

(iii) Construct 2-connected graphs with $\mathcal{C}^\perp \subsetneq \hat{\mathcal{C}}^\perp$ or $\mathcal{B}^\perp \subsetneq \hat{\mathcal{B}}^\perp$.

125. Extending Gallai's partition theorem of Exercise 55 (ii), Chapter 1, show that $E(G)$ can be partitioned into a set $C \in \mathcal{C}$ and a set $D \in \mathcal{B}$. (This strengthens Exercise 25.)

126. Let $F \subseteq E(G)$ be a set of edges.

(i)$^-$ Show that F extends to a cut if it contains no odd circuit.

(ii)$^+$ Show that F extends to some $D \in \mathcal{C}$ if it contains no odd bond.

127.$^+$ Let H be an abelian group. The group $\mathcal{C}_H$ of *H-circulations* on $|G|$ consists of the maps $\psi\colon \vec{E} \to H$ that satisfy (F1) and $\psi(X,Y) = 0$ for any finite cut $E(X,Y)$ of G. (See Chapter 6.1 for notation.) Extend Exercise 8 of Chapter 6 to $\mathcal{C}_H$, with $\mathcal{E}_H$ and $\mathcal{D}_H$ as defined there.

128. Let X be a connected standard subspace of $|G|$. Call a continuous map $\sigma\colon S^1 \to X$ a *topological Euler tour* of X if it traverses every edge in $E(X)$ exactly once. (Formally: every inner point of an edge in $E(X)$ must be the image of exactly one point in S^1.) Show that X admits a topological Euler tour if and only if $E(X) \in \mathcal{C}(G)$.

129.$^+$ An *open Euler tour* in an infinite connected graph G is a 2-way infinite walk $\dots e_{-1}v_0e_0\dots$ that contains every edge of G exactly once. Show that G contains an open Euler tour if and only if G is countable, every vertex has even or infinite degree, and any finite cut $F = E(V_1,V_2)$ with both V_1 and V_2 infinite is odd.

130. By Exercise 23 of Chapter 4, every finite 2-connected graph without a K^4 or $K_{2,3}$ minor contains a *Hamilton* cycle, one that contains all its vertices. Show that every locally finite such graph has a *Hamilton circle*, a circle in $|G|$ containing all the vertices (and ends) of G.

131.$^+$ Extend Theorem 2.4.4 to packings and coverings of locally finite graphs with topological spanning trees in appropriate spaces obtained from $|G|$.

132. Where in the proof of Theorem 8.8.2 do we use that G is connected?

133. Use the techniques from Section 8.8 to prove that the wild circle is indeed a circle. Your proof may be informal in its handling of the wild circle graph G – use a picture rather than its formal definition.

134.$^+$ Use the methods from Section 8.8 to prove that connected standard subspaces of $|G|$, for G locally finite, have topological spanning trees: closed connected subspaces containing all their vertices.

135.$^+$ Consider the space $\|G\|$ for the edgeless countably infinite graph G. Have you met this space before in another guise?

Notes

There is no comprehensive monograph on infinite graph theory, but over time several surveys have been published. A relatively wide-ranging collection of survey articles can be found in R. Diestel (ed.), *Directions in Infinite Graph Theory and Combinatorics*, North-Holland 1992. (This has been reprinted as Volume 95 of the journal *Discrete Mathematics*.) Some of the articles there address purely graph-theoretic aspects of infinite graphs, while others point to connections with other fields in mathematics such as differential geometry, topological groups, or logic. A similar collection, edited by R. Diestel, B. Mohar and G. Hahn, appeared in 2011 as a special issue of *Discrete Mathematics (vol. 311)*. It includes a survey of everything to do with $|G|$, the topological space consisting of an infinite graph G and its ends: R. Diestel, Locally finite graphs with ends: a topological approach, arXiv:0912.4213. Another very instructive survey in this volume is by L. Soukup, Elementary submodels in infinite combinatorics, arXiv:1007.4309.

A survey of infinite graph theory as a whole was given by C. Thomassen, Infinite graphs, in (L.W. Beineke & R.J. Wilson, eds.) *Selected Topics in Graph Theory 2*, Academic Press 1983. This also treats a number of aspects of infinite graph theory not considered in our chapter here, including problems of Erdős concerning infinite chromatic number, infinite Ramsey theory (also known as *partition calculus*), and reconstruction. The first two of these topics receive much attention also in A. Hajnal's chapter of the *Handbook of Combinatorics* (R.L. Graham, M. Grötschel & L. Lovász, eds.), North-Holland 1995, which has a strong set-theoretical flavour. Péter Komjáth is currently preparing a monograph about this kind of infinite graph theory. A specific survey on reconstruction by Nash-Williams can be found in the *Directions* volume cited above. R. Halin, Miscellaneous problems on infinite graphs, *J. Graph Theory* **35** (2000), 128–151, contains his legacy of unsolved infinite graph problems.

Halin's book, *Graphentheorie* (2nd ed.), Wissenschaftliche Buchgesellschaft 1989, is also good general reference for infinite graphs. A more specific monograph about *simplicial decompositions* of infinite graphs (see Chapter 12) is R. Diestel, *Graph Decompositions*, Oxford University Press 1990. Our Chapter 12.6 closes with two theorems about forbidden minors in infinite graphs.

When sets get bigger than countable, combinatorial set theory offers some interesting ways other than cardinality to distinguish 'small' from 'large' sets. Among these are the use of *clubs* and *stationary sets*, of *ultrafilters*, and of *measure and category*. See P. Erdős, A. Hajnal, A. Máté & R. Rado, *Combinatorial Set Theory: partition relations for cardinals*, North-Holland 1984; W.W. Comfort & S. Negropontis, *The Theory of Ultrafilters*, Springer 1974; J.C. Oxtoby, *Measure and Category: a survey of the analogies between topological and measure spaces* (2nd ed.), Springer 1980.

Infinite matroids, whose study long lay dormant for want of a clear idea of what exactly they should be, were finally axiomatized in H. Bruhn, R. Diestel, M. Kriesell & P. Wollan, Axioms for infinite matroids, *Adv. Math.* **239** (2013), 18–46, arXiv:1003.3919. Since that paper, the theory of infinite matroids has flourished. It draws much inspiration from topological infinite graph theory as presented in Sections 8.6–8.7, but also has its own problems and agenda. Regular updates can be found in Nathan Bowler's blog at http://matroidunion.org/?author=11.

In addition to its main guiding themes, as followed in this chapter and in the sources just mentioned, infinite graph theory has a number of interesting individual results which, as yet, stand essentially by themselves. One such is a theorem of A. Huck, F. Niedermeyer and S. Shelah, Large κ-preserving sets in infinite graphs, *J. Graph Theory* **18** (1994), 413–426, which says that every infinitely connected graph G has a set S of $|G|$ vertices such that $\kappa(G - S') = \kappa(G)$ for every $S' \subseteq S$. Another is Halin's *bounded graph conjecture* and related problems. (See Exercise 80 for the definition of 'bounded' and the tree case of the conjecture.) A proof can be found in R. Diestel & I.B. Leader, A proof of the bounded graph conjecture, *Inv. math.* **108** (1992), 131–162.

König's infinity lemma, sometimes referred to as *König's lemma*, is as old as the first-ever book on graph theory, which includes it: D. König, *Theorie der endlichen und unendlichen Graphen*, Akademische Verlagsgesellschaft, Leipzig 1936. Appendix A gives a generalization of the infinity lemma to structures of any cardinality, which is still very intuitive and graph-like: the compactness theorem for inverse limits of finite sets. Compactness proofs can also come in the guise of Rado's *selection lemma*, or of Gödel's *compactness theorem* from first-order logic. These two, as well as the generalized infinity lemma, are equivalent to the compactness principle as stated in Appendix A, but stronger than the ordinary infinity lemma. They follow from Tychonoff's theorem (which is one of the many statments equivalent to the axiom of choice) but do not imply it. They do however imply the weakening of Tychonoff's theorem that we typically use in compactness proofs, namely, that spaces of the form S^X with S finite are compact in the product topology.

Theorem 8.1.3 is due to N. G. de Bruijn and P. Erdős, A colour problem for infinite graphs and a problem in the theory of relations, *Indag. Math.* **13** (1951), 371–373. The infinite analogue of the weakening of Hadwiger's conjecture that every graph of chromatic number $\alpha \geqslant \aleph_0$ contains a TK_β for every $\beta < \alpha$, whose proof for $\alpha = \aleph_0$ is asked in Exercise 16, is due to R. Halin, Unterteilungen vollständiger Graphen in Graphen mit unendlicher chromatischer Zahl, *Abh. Math. Sem. Univ. Hamburg* **31** (1967), 156–165.

Unlike for the chromatic number, a bound on the colouring number of all finite subgraphs does not extend to the whole graph by compactness. P. Erdős & A. Hajnal, On the chromatic number of graphs and set systems, *Acta Math. Acad. Sci. Hung.* **17** (1966), 61–99, proved that if every finite subgraph of G has colouring number at most k then G has colouring number at most $2k - 2$, and showed that this is best possible. However, as for finite graphs, the colouring number appears to interact better with other graph invariants than the chromatic number does; compare Theorem 8.6.2. For any cardinal κ, the graphs with colouring number at most κ were characterized by forbidden subgraphs by N. Bowler, J. Carmesin and C. Reiher, The colouring number of infinite graphs, *Combinatorica* **39** (2019), 1225–1235, arXiv:1512.02911.

The unfriendly partition conjecture is one of the best-known open problems in infinite graph theory, but there are few results. E.C. Milner and S. Shelah, Graphs with no unfriendly partitions, in (A. Baker, B. Bollobás & A. Hajnal, eds.), *A tribute to Paul Erdős*, Cambridge University Press 1990, construct an uncountable counterexample, but show that every graph has an unfriendly partition into three classes. (The original conjecture, which they attribute to R. Cowan and W. Emerson (unpublished), appears to have asserted for every graph the existence of a vertex partition into any given finite number of classes such that every vertex has at least as many neighbours in other classes as in its own.) Some positive results for bipartitions were obtained by R. Aharoni, E.C. Milner and K. Prikry, Unfriendly partitions of graphs, *J. Comb. Theory, Ser. B* **50** (1990), 1–10. H. Bruhn, R. Diestel, A. Georgakopoulos and Ph. Sprüssel, Every rayless graph has an unfriendly partition, *Combinatorica* **30** (2010), 521–532, arXiv:0901.4858, used rankings such as defined in Section 8.5 to prove that all rayless graphs have unfriendly partitions. E. Berger, Unfriendly partitions for graphs not containing an infinite clique, *Combinatorica* **37** (2017), 157–166, strengthened this to prove that every graph not containing a subdivision of an infinite complete graph has an unfriendly partition.

Theorem 8.2.4 is a special case of the result stated in Exercise 43 (i), which is due to H.A. Jung, Wurzelbäume und unendliche Wege in Graphen, *Math. Nachr.* **41** (1969), 1–22. The graphs that admit a normal spanning tree have been characterized by forbidden minors by M. Pitz, Proof of Halin's normal spanning tree conjecture, *Isr. J. Math.* **246** (2021), 353–370, arXiv:2005.02833. Note that such a characterization is possible only because the class of graphs admitting a normal spanning tree is closed under taking connected minors – a consequence of Jung's theorem (see Exercise 43 (ii)) that is not so easy to prove directly. Pitz's main result is a proof of Halin's conjecture that a connected graph has a normal spanning tree if and only if all its minors have countable colouring number; see Theorem 8.6.2. Exercise 44 says that connected graphs not containing a subdivided fat $K^{\aleph_0}$ have normal spanning trees. The proof from the hint is also due to M. Pitz, Quickly proving Diestel's normal spanning tree criterion, *Electronic. J. Comb.* **28** (2021), P3.59, arXiv:2006.02994.

Theorems 8.2.5 and 8.2.6 are from R. Halin, Über die Maximalzahl fremder unendlicher Wege, *Math. Nachr.* **30** (1965), 63–85. Our proof of Theorem 8.2.5 (i) is due to Andreae (unpublished); the proof of Theorem 8.2.5 (ii) given in the hint for Exercise 45 is new. The analogue of Theorem 8.2.5 (ii) for double rays was proved only recently by N. Bowler, J. Carmesin and J. Pott, Edge-disjoint double rays in infinite graphs: a Halin type result, *J. Comb. Theory, Ser. B* **111** (2015), 1–16.

Our proof of Theorem 8.2.6 is new. Halin's paper also includes a structure theorem for graphs that do not contain infinitely many disjoint rays. Except for a finite set of vertices, such a graph can be written as an infinite chain of rayless subgraphs each overlapping the previous in exactly m vertices, where m is the maximum number of disjoint rays (which exists by Theorem 8.2.5). These overlap sets are disjoint, and there are m disjoint rays containing exactly one vertex from each of them.

The ubiquity conjecture is from Th. Andreae, On disjoint configurations in infinite graphs, *J. Graph Theory* **39** (2002), 222–229. After lying dormant for nearly twenty years, it has seen considerable recent progress; see N. Bowler, C. Elbracht, J. Erde, P. Gollin, K. Heuer, M. Pitz & M. Teegen, Ubiquity in graphs I: Topological ubiquity of trees, *J. Comb. Theory, Ser. B* **157** (2022), 70–95, arXiv:1806.04008, and the references included there.

Universal graphs have been studied mostly with respect to the induced subgraph relation, with numerous but mostly negative results. See G. Cherlin, S. Shelah & N. Shi, Universal graphs with forbidden subgraphs and algebraic closure, *Adv. Appl. Math.* **22** (1999), 454–491, for an overview and a model-theoretic framework for the proof techniques typically applied.

The Rado graph is probably the best-studied single graph in the literature (with the Petersen graph a close runner-up). The most comprehensive source for anything related to it (and far beyond) is R. Fraïssé, *Theory of Relations* (2nd edn.), Elsevier 2000. More accessible introductions are given by N. Sauer in his appendix to Fraïssé's book, and by P.J. Cameron, The random graph, in (R.L. Graham & J. Nešetřil, eds.): *The Mathematics of Paul Erdős*, Springer 1997, and its references.

Theorem 8.3.1 is due to P. Erdős and A. Rényi, Asymmetric graphs, *Acta Math. Acad. Sci. Hung.* **14** (1963), 295–315. The existence part of their proof is probabilistic and will be given in Theorem 11.3.5. Rado's explicit definition of the graph R was given in R. Rado, Universal graphs and universal functions, *Acta Arithm.* **9** (1964), 393–407. However, its universality and that of R^r are already included in more general results of B. Jónsson, Universal relational systems, *Math. Scand.* **4** (1956), 193–208.

Theorem 8.3.3 is due to A.H. Lachlan and R.E. Woodrow, Countable ultrahomogeneous undirected graphs, *Trans. Amer. Math. Soc.* **262** (1980), 51–94. The classification of the countable homogeneous *directed* graphs is much more difficult still. It was achieved by G. Cherlin, The classification of countable homogeneous directed graphs and countable homogeneous n-tournaments, *Mem. Am. Math. Soc.* **621** (1998), which also includes a shorter proof of Theorem 8.3.3. M. Hamann obtained, in his 2014 Habilitation at Hamburg University, the even more difficult classification of the *Connected-homogeneous digraphs*. The thesis is available on the internet and still appearing, spread over a number of papers.

Proposition 8.3.2, too, has a less trivial directed analogue: the countable directed graphs that are isomorphic to at least one of the two sides induced by any bipartition of their vertex set are precisely the edgeless graph, the random tournament, the transitive tournaments of order type ω^α, and two specific orientations of the Rado graph (R. Diestel, I. Leader, A. Scott & S. Thomassé, Partitions and orientations of the Rado graph, *Trans. Amer. Math. Soc.* **359** (2007), 2395–2405.

Theorem 8.3.4 is proved in R. Diestel & D. Kühn, A universal planar graph under the minor relation, *J. Graph Theory* **32** (1999), 191–206. There is no minor-universal graph for embeddability in any closed surface other than the sphere; see the above paper. The results of A. Georgakopoulos mentioned after Theorem 8.3.4 are in arXiv:2212.05498. There is no universal planar graph for the topological minor relation, see T. Krill, Universal graphs for the topological minor relation, *J. Graph Theory* **104** (2023), 683–696, arXiv:2203.12643.

Theorem 8.4.1 is from F. Laviolette, Decompositions of infinite graphs, *J. Comb. Theory, Ser. B* **94** (2005), 259–333. The paper comes in two parts, the second of which builds on the first and characterizes the graphs admitting edge-decompositions into cycles and double rays. Theorem 8.4.1 has been re-proved independently by C. Tomassen, Nash-Williams' cycle-decomposition theorem, *Combinatorica* **37** (2017), 1027–1037, and by L. Soukup in his survey on the use of elementary submodels cited at the start of these notes. The latter proof yields a more general version, in which the $\aleph_0$ implicit in the statement of Theorem 8.4.1 is replaced everywhere by an arbitrary fixed cardinal. While Laviolette's original proof uses Nash-Williams's Theorem 8.4.2, which we derived from it, the two later proofs do not.

Theorem 8.4.2 is from C.St.J.A. Nash-Williams, Decomposition of graphs into closed and endless chains, *Proc. Lond. Math. Soc.* **10** (1960), 221–238. An analogue for directed graphs was conjectured by Thomassen in his above-cited paper, and proved by A. Joó, On partitioning the edges of an infinite digraph into directed cycles, *Adv. Comb.* **18702** (2021), doi.org/10.19086/aic.18702 arXiv:1704.08830.

Lacking the concept of an infinite circuit as we defined it here, Nash-Williams also sought to generalize Theorem 8.4.2 and other theorems about finite cycles by replacing 'cycle' with '2-regular connected graph' (which may be finite or infinite). The resulting statements are not always as smooth as the finite theorems they generalize, but some substantial work has been done in this direction. C.St.J.A. Nash-Williams, Decompositions of graphs into two-way infinite paths, *Can. J. Math.* **15** (1963), 479–485, characterizes the graphs admitting edge-decompositions into double rays.

When Erdős conjectured his extension of Menger's theorem is not known; C.St.J.A. Nash-Williams, Infinite graphs – a survey, *J. Comb. Theory, Ser. B* **3** (1967), 286–301, cites the proceedings of a 1963 conference as its source. Its proof as Theorem 8.4.4 by R. Aharoni and E. Berger, Menger's theorem for infinite graphs, *Inv. math.* **176** (2009), 1–62, arXiv:math/0509397, came as the culmination of a long effort over many years, for the most part also due to Aharoni. Our proof of its countable case is adapted from R. Aharoni, Menger's theorem for countable graphs, *J. Comb. Theory, Ser. B* **43** (1987), 303–313. The theorem has been extended to graphs with ends by H. Bruhn, R. Diestel & M. Stein, Menger's theorem for infinite graphs with ends, *J. Graph Theory* **50** (2005), 199–211.

Theorem 8.4.9 is due to J.S. Pym, A proof of the linkage theorem, *J. Math. Anal. Appl.* **27** (1969), 636–638. The short proof outlined in Exercise 70 can be found in R. Diestel & C. Thomassen, A Cantor-Bernstein theorem for paths in graphs, *Amer. Math. Monthly* **113** (2006), 161–166.

The matching theorems of Chapter 2 – König's duality theorem, Hall's marriage theorem, Tutte's 1-factor theorem, and the Gallai-Edmonds matching theorem – extend essentially unchanged to locally finite graphs by compactness; see e.g. Exercises 26–29. For non-locally-finite graphs, matching theory is considerably deeper. A good survey and open problems can be found in R. Aharoni, Infinite matching theory, in the *Directions* volume cited earlier.

Most of the results and techniques for infinite matching were developed first for countable graphs, by Podewski and Steffens in the 1970s. In the 1980s, Aharoni extended them to arbitrary graphs, where things are more difficult still and additional methods are required. Theorem 8.4.10 is by R. Aharoni, König's duality theorem for infinite bipartite graphs, *J. Lond. Math. Soc.* **29** (1984), 1–12. The proof builds on R. Aharoni, C.St.J.A. Nash-Willaims & S. Shelah, A general criterion for the existence of transversals, *Proc. Lond. Math. Soc.* **47** (1983), 43–68. It is described in detail also in M. Holz, K.P. Podewski & K. Steffens, *Injective choice functions*, *Lecture Notes in Mathematics* **1238**

Springer-Verlag 1987. Theorem 8.4.12 is due to E.C. Milner & S. Shelah, Sufficiency conditions for the existence of transversals, *Can. J. Math.* **26** (1974), 948–961; a short proof was given by H. Tverberg, On the Milner-Shelah condition for transversals, *J. Lond. Math. Soc.* **13** (1976), 520–524. Theorem 8.4.13 can be derived from the material in K. Steffens, Matchings in countable graphs, *Can. J. Math.* **29** (1977), 165–168. Theorem 8.4.14 is due to R. Aharoni, Matchings in infinite graphs, *J. Comb. Theory, Ser. B* **44** (1988), 87–125; a shorter proof was given by F. Niedermeyer and K.P. Podewski, Matchable infinite graphs, *J. Comb. Theory, Ser. B* **62** (1994), 213–227.

The recursive ranking of rayless graphs defined in Section 8.5 was introduced by R. Schmidt, Ein Ordnungsbegriff für Graphen ohne unendliche Wege mit einer Anwendung auf n-fach zusammenhängende Graphen, *Arch. Math.* **40** (1983), 283–288. The paper offers an interesting structure theory for rayless graphs including applications, such as to reconstruction.

The topology on G introduced in Section 8.6 coincides, when G is locally finite, with the usual topology of a 1-dimensional CW-complex. The space $|G|$ is its *Freudenthal compactification*, as suggested by H. Freudenthal, Über die Enden topologischer Räume und Gruppen, *Math. Zeit.* **33** (1931), 692–713; see Exercise 90. Although $|G|$ has been known for so long and is a familiar object in geometric group theory, its fundamental group was characterized only recently by R. Diestel & Ph. Sprüssel, The fundamental group of a locally finite graph with ends, *Adv. Math.* **226** (2011), 2643–2675, arXiv:0910.5647.

For graphs that are not locally finite, the graph-theoretical notion of an end is more general than the topological one; see R. Diestel & D. Kühn, Graph-theoretical versus topological ends of graphs, *J. Comb. Theory, Ser. B* **87** (2003), 197–206. For such graphs it can be natural to consider a coarser topology on $|G|$, obtained by taking as basic open sets $\hat{C}_\epsilon(S,\omega)$ only those with $\epsilon = 1$. Under this topology, $|G|$ is no longer Hausdorff, because every vertex dominating an end ω will lie in the closure of every $\hat{C}(S,\omega)$. But $|G|$ can now be compact, and it can have a natural quotient space – in which ends are identified with vertices dominating them and rays converge to vertices – that is both Hausdorff and compact. For details see R. Diestel, On end spaces and spanning trees, *J. Comb. Theory, Ser. B* **96** (2006), 846–854, where also the equivalence between (i) and (ii) in Theorem 8.6.2 is proved. The fact that $|G|$, and therefore also its closed subspace Ω, is normal also for non-locally-finite G was proved by Ph. Sprüssel, End spaces are normal, *J. Comb. Theory, Ser. B* **98** (2008), 798–804. Further topological aspects of the subspaces Ω and $V \cup \Omega$ were studied by Polat; see e.g. N. Polat, Ends and multi-endings I & II, *J. Comb. Theory, Ser. B* **67** (1996), 56–110.

Arbitrary graphs can be compactified by adding their $\aleph_0$-tangles. Tangles are defined in Chapter 12; the $\aleph_0$-tangles of an infinite graph include its ends. For G locally finite, its tangle-compactification is precisely $|G|$. See R. Diestel, Ends and tangles, Special volume in memory of Rudolf Halin, *Abh. Math. Sem. Univ. Hamburg* **87** (2017), 223–244, arXiv:1510.04050.

Lemma 8.6.4, that the image of any x–y path in a Hausdorff space contains an x–y arc, is from D.W. Hall and G.L. Spencer, Elementary Topology. A locally finite graph G with a connected subset of $|G|$ that is not arc-connected has been constructed by A. Georgakopoulos, Connected but not path-connected subspaces of infinite graphs, *Combinatorica* **27** (2007), 683–698.

The construction of highly connected graphs that are left disconnected by the deletion of any finite circuit (Exercise 21) is due to R. Aharoni and C. Thomassen, Infinite highly connected digraphs with no two arc-disjoint spanning trees, *J. Graph Theory* **13** (1989), 71–74. Its consequence that the infinite analogue of the tree packing theorem fails with ordinary spanning trees had already been established by J.G. Oxley, On a packing problem for infinite graphs and independence spaces, *J. Comb. Theory, Ser. B* **26** (1979), 123–130.

Nash-Williams had conjectured in his original paper that the finite tree packing theorem should extend to infinite graphs verbatim. What Tutte thought about an infinite version is not recorded. In his original paper he does consider the infinite case, but 'backwards': rather than speculating on which infinite graphs might admit edge-decompositions into k spanning trees, he proves that the locally finite graphs satisfying the cross-edges condition decompose into 'semiconnected subgraphs', defined – just for this purpose – as those subgraphs that contain an edge from every finite cut. These subgraphs, of course, are precisely the spanning subgraphs whose closures are connected (Lemma 8.6.7), so Tutte in fact proved Lemma 8.6.10 without having the topological language to express it.

The companion to the finite tree packing theorem, the tree-covering theorem, extends to locally finite graphs verbatim (Exercise 19). The packing-covering theorem, Theorem 2.4.4, extends to infinite graphs separately in two ways: with ordinary spanning trees, and with topological ones. See Exercises 20 and 131, and their hints.

The example of the wild circle C illustrates why end-degrees in subspaces are defined in terms of arcs rather than rays. Every end on C contains only one ray that also lies in C, but it is the endpoint of two otherwise disjoint arcs in C. If we wish to be able to characterize the circles as those subspaces in which every vertex and every end has degree 2 (Exercise 103), we thus need the topological definition of end-degrees in subspaces.

The supremum in the definition of topological end degrees is in fact a maximum. This is a special case of Menger's ω-*Beinsatz* (1927), which he proved in the same paper as his now famous Theorem 3.3.1.

The (combinatorial) vertex-degree of an end used to be called its *multiplicity*. The term 'degree' was introduced by H. Bruhn and M. Stein, On end degrees and infinite circuits in locally finite graphs, *Combinatorica* **27** (2007), 269–291. Their main result was that the (entire) edge set of a locally finite graph lies in its topological cycle space if and only if every vertex has even degree and every end has even edge-degree – with a newly found division of the ends of infinite degree into 'even' and 'odd'. This was later generalized to arbitrary edge sets by E. Berger & H. Bruhn, Eulerian edge sets in locally finite graphs, *Combinatorica* **31** (2011), 21–38.

An interesting new aspect of end degrees is that they can make it possible to study extremal-type problems for infinite graphs that would otherwise make sense only for finite graphs. For example, while finite graphs of large enough minimum degree contain any desired topological minor or minor (see Chapter 7), an infinite graph of large minimum degree can be a tree. The ends of a tree, however, have degree 1. An assumption that the degrees of both vertices and ends of an infinite graph are large can still not force a non-planar minor (because such graphs can be planar), but it does force arbitrarily highly connected subgraphs. See R. Diestel, Forcing finite minors in sparse infinite graphs by large-degree assumptions, *Electronic. J. Comb.* **22** (2015), #P1.43, arXiv:1209.5318, as well as M. Stein, Extremal infinite graph theory, *Discrete Math.* **311** (2011), 1472–1496, arXiv:1102.0697, for this and other results in this vein. Another approach to 'extremal' infinite graph theory, which seeks to force infinite substructures by assuming a lower bound for $\|G[v_1, \dots v_n]\|$ when $V(G) = \{v_1, v_2, \dots\}$, is taken by J. Czipszer, P. Erdős and A. Hajnal, Some extremal problems on infinite graphs, *Publ. Math. Inst. Hung. Acad. Sci., Ser. A* **7** (1962), 441–457.

Our topological notion of the cycle space $\mathcal{C}(G)$ may appear natural in an infinite setting, but historically it is very young. It was developed in order to extend the classical applications of the cycle space of finite graphs, such as in planarity and duality, to locally finite graphs. As in the case of the tree packing theorem, those extensions fail when only finite circuits and sums are permitted, but they do hold for the topological cycle space. Examples include Tutte's theorem (3.2.6) that the non-separating induced cycles generate the whole cycle space; MacLane's (4.5.1), Kelmans's (4.5.2) and Whitney's (4.6.3) characterizations of planarity; and Gallai's partition theorem of Exercise 55, Chapter 1. There are a couple of papers by Diestel and Sprüssel that extend the notion of the topological cycle space to a general homology theory for locally compact spaces: The homology of locally finite graphs with ends, *Combinatorica* **30** (2010), 681–714, and On the homology of locally compact spaces with ends, *Topology and its Applications* **158** (2011), 1626–1639, arXiv:0910.5650.

For graphs that are not locally finite, there seems to be no one notion of topological cycle space that caters for all needs. A promising approach that takes account of this diversity was suggested by A. Georgakopoulos, Graph topologies induced by edge lengths, *Discrete Math.* **311** (2011), 1523–42, arXiv:0903.1744. For locally finite G this approach builds our familiar $|G|$ as the completion, rather than compactification, of G with a suitable metric – a fact to which the title of Section 8.6 alludes. By varying the metric, however, this approach can also yield other kinds of boundaries, such as the hyperbolic boundary of a hyperbolic graph.

Theorems 8.7.1 and 8.7.3 are from R. Diestel & D. Kühn, On infinite cycles I–II, *Combinatorica* **24** (2004), 69–116. While these theorems extend the familiar properties of the cycle space of a finite graph to locally finite infinite graphs, the same can be done for the cut space (Exercise 122). The finitary cycle space $\mathcal{C}_{\text{fin}}$, which is clearly a subspace of the topological cycle space $\mathcal{C}$, is in fact the set of *all* its finite elements (Exercise 120), just as the finitary cut space is the set $\mathcal{B}_{\text{fin}}$ of all the finite elements of the entire cut space $\mathcal{B}$. For 3-connected graphs (but not otherwise, see Exercise 124 (iii)), edge sets orthogonal to all the circuits are in fact in $\mathcal{C}^\perp$, and sets orthogonal to all the bonds are in fact in $\mathcal{B}^\perp$; see R. Diestel & J. Pott, Orthogonality and minimality in the homology of locally finite graphs, *Electronic. J. Comb.* **21** (2014), #P3.36, arXiv:1307.0728.

The orthogonalities between $\mathcal{C}$, $\mathcal{B}$, $\mathcal{C}_{\text{fin}}$ and $\mathcal{B}_{\text{fin}}$ described in Exercise 123 is best captured in terms of matroids; see H. Bruhn & R. Diestel, Infinite matroids in graphs, *Discrete Math.* **311** (2011), 1461–1471, arXiv:1011.4749. Duality for planar infinite graphs is treated in H. Bruhn & R. Diestel, Duality in infinite graphs, *Comb. Probab. Comput.* **15** (2006), 75–90.

Theorem 8.8.2 is folklore – it was probably already known to Freudenthal. As mentioned in the text, one can obtain any standard subspace X' of $X = |G|$ as an inverse limit of the finite subgraphs H_p of G/p induced by the edges in X'. In order to find a topological spanning tree in a standard subspace $X \subsetneq |G|$, however, or even just an arc between two given points (so as to prove Lemma 8.6.5), we will need to construct these H_p by hand: in order to be able to expand a spanning tree or path in H_p to one in H_q for $p < q$, we need to ensure that every dummy vertex of H_p induces a connected subspace in X.

One way to do this is to mimic the G_n inside X: to enumerate $V(X)$, and take as p_n the partition of $V(X)$ consisting of the singleton classes of its first n vertices of X and, as further classes, the vertex sets of the arc-components obtained from X by deleting these n vertices and their incident edges. If G is locally finite and X is connected, then these p_n will be finite partitions, and by the jumping arc lemma they will have only finitely many cross-edges even in G. (One needs to show here that arc-components of standard subspaces are closed; but this is easy, even without Lemma 8.6.5).

One can then show directly that X is the inverse limit of the compact spaces X_n obtained from X by collapsing each of the partition classes in p_n. If one wishes to apply, rather than re-prove, Theorem 8.8.2 to obtain X as such an inverse limit, one has to expand the p_n to partitions $\overline{p}_n$ of $V(G)$, making sure that these $\overline{p}_n$ are cofinal in $P(G)$, and note that the desired X_n arise as the images of X under the given projections f_p. This can be done too, but it needs some care; see the hint for Exercise 134.

For connected graphs G that are not locally finite one can still define $\|G\|$ as in the text. The compact space one obtains can be described as a Hausdorff quotient of the space obtained from G by adding its 'edge ends', the equivalence classes of rays with respect to finite edge separators, minus loops.