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 Section 5 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 8.5.

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

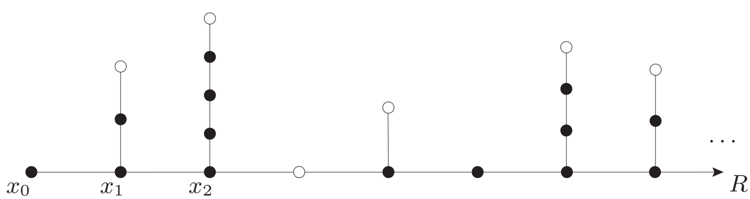

Fig. 8.1.1. A comb with white teeth and spine $R = x_0x_1 \dots$

[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 an appendix 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 uncountable sets must still be uncountable. It need not be; in fact it may be empty. (Example?)

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 frequently found 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] The appendix offers brief introductions to both, enough to enable the reader to use these tools with confidence in practice.

[8.2.1] [8.2.6] [8.5.1] [8.5.8] [9.1.3]

Lemma 8.1.2. (König's Infinity Lemma)
Let $V_0, V_1, \ldots$ 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 $v_0 v_1 \ldots$ with $v_n \in V_n$ for all n.

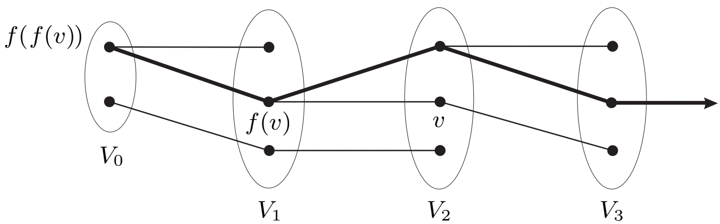

Fig. 8.1.2. König's infinity lemma

Proof. Let $\mathcal{P}$ be the set of all finite paths of the form $v\, f(v)\, f(f(v)) \ldots$ 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 $v_0 v_1 \ldots$ is indeed a ray. $\square$

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 $v_0, v_1, \ldots$ be an enumeration of V and put $G_n := G[v_0, \ldots, v_n]$. Write V_n for the set of all k-colourings of G_n with colours in $\{1, \ldots, k\}$. 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 $\{v_0, \ldots, v_{n-1}\}$. Let $c_0 c_1 \ldots$ 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 $\{1, \ldots, k\}$. $\square$

The particular intuitive appeal of the infinity lemma is made possible by the fact that a countable graph can be exhausted by a single nested sequence of finite subgraphs. For graphs of arbitrary cardinality this is not possible, but there are other standard ways in which compactness

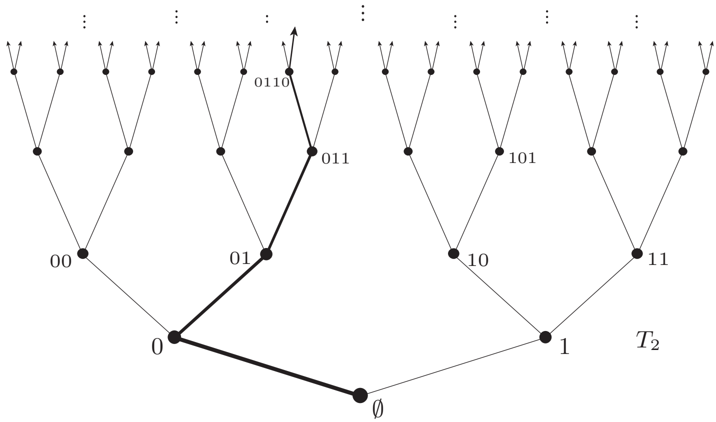

Fig. 8.1.4. The binary tree T_2 has continuum many ends, one for every infinite 0–1 sequence

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 (8.1.2)
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. □

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) [8.5.3]
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^*.

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 $v_1, v_2, \dots$ 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 $G[v_1, \dots, v_n]$ but not adjacent to the image of any non-neighbour of v_n in $G[v_1, \dots, v_n]$. Clearly, this map is an isomorphism between G and the subgraph of R induced by its image.

[11.3.5] **Theorem 8.3.1.** (Erdős and Rényi 1963)

R *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 41).

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 only 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 $G \not\simeq G[V_2]$, since G has isolated vertices but $G[V_2]$ does not. Hence $G = G[V_1] \simeq \overline{K^{\aleph_0}}$. 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 $G \not\simeq G[V_1]$ and hence $G \simeq G[V_2]$. 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, $G \not\simeq G[V_1]$ and hence $G \simeq G[V_2]$. 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 42).

R^r

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 $R_0^r := K^1$, and let R_{n+1}^r be obtained from R_n^r by joining, for every subgraph $H \not\simeq K^{r-1}$ of R_n^r, a new vertex v_H to every vertex in H. Then let $R^r := \bigcup_{n \in \mathbb{N}} R_n^r$. Clearly, as the new vertices v_H of R_{n+1}^r are independent, there is no K^r in R_{n+1}^r if there was none in R_n^r, 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 it is clearly 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. Unfortunately, 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.

So far, this theorem is the only one of its kind. But it should be possible to find more. For instance: for which graphs X is there a minor-universal graph in the class $\mathrm{Forb}_{\preccurlyeq}(X) = \{\, G \mid X \not\preccurlyeq G \,\}$?

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 and was proved only recently by Aharoni and Berger. The techniques developed for its proof inspired, over the years, much of the theory in this area.

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 8.4.1. *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.*

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 (3.3.1) 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 $\mathcal{P}$. □

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 $|\mathcal{P}| = \left|\bigcup \mathcal{P}\right| \geqslant k$ 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 8.4.2. (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 8.4.2 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 $|\mathcal{P}|$ or on $|G| + \|G\|$, 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 8.4.2. 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