3. Connectivity

A graph consists of a set of vertices, some pairs of which are "joined" or "connected" by edges. In this way we see that edges provide units of connectivity. The fundamental question we would like to understand is how the units of connectivity accumulate.

As one example, consider the internet and social media. These can be modelled by graphs. Accumulations of connectivity in this setting corresponds to "communities" that develop. Once these communities are identified, this is powerful information that can be used for good or evil.

But first things first — we had better start from the basics. We already know from Section 1 what it means for a graph to be *connected*, and what the *components* of a graph are.

A connected graph can still be pretty fragile. Consider a country that has earthquakes from time to time. Each earthquake might destroy a road or two. Ideally we would like to keep a connected road or train network, even after an earthquake. If we want this, we had better avoid bridges (in the graph-theoretic sense) in our graph — ask the residents of Kaikoura. The upshot is that we don't just want to know if graphs are connected or not — we want to know *how* well-connected they are. In this section, we will introduce the notion of a k-connected graph, where the larger k is, the more well-connected the graph is.

Given the existence of bridges, we know that it is possible to delete an edge from a connected graph and lose connectivity. Remarkably the same cannot happen for contraction.

Lemma 3.1. *Let G be a connected graph, and let e be an edge of G. Then G/e is connected.*

Proof. [redacted] $\square$

2-connected graphs. A graph is 2-*connected* if it is connected, it has at least three vertices, and has no cut vertices.

→ For a 2-connected graph, not only is it connected, but we have a guarantee that we cannot disconnect the graph by deleting a single vertex.

Why do we require that a 2-connected graph has at least three vertices? The reason is that a graph on at most two vertices is too small to have a cut vertex, so in this case "having no cut vertices" becomes less meaningful.

Although we can lose the property of being 2-connected by deleting a vertex, we always retain the property of being connected after deleting a single vertex.

Exercise 3.2. *Give an example of a 2-connected graph G with at least four vertices, and a vertex v such that $G - v$ is not 2-connected.*

Exercise 3.3. *Prove that for a 2-connected graph G, the graph $G - v$ is connected for any $v \in V(G)$.*

How about when we delete or contract an edge in a 2-connected graph? We saw in Lemma 3.1 that we can never lose the property of being connected by contracting an edge. Unfortunately the same does not hold for the property of being 2-connected. We can lose 2-connectivity by deleting an edge, or by contracting an edge. However, it is not the case that both possibilities occur for the same edge.

Theorem 3.4. *Let G be a 2-connected graph with at least four vertices, and let e be an edge of G. Then either $G\backslash e$ or G/e is 2-connected.*

Proof. Say $e = uv$. Assume that $G\backslash e$ is not 2-connected. We work towards showing that G/e is 2-connected. First, we consider what we can say about the structure of G using the fact that $G\backslash e$ is not 2-connected.

As G has no cut vertices, e is not a bridge by Lemma 2.18, so $G\backslash e$ is connected. Since $G\backslash e$ is connected, but not 2-connected, it has a cut vertex y. By Lemma 2.19, there is a partition $\{A, B\}$ of $V(G\backslash e) \setminus \{y\}$ such that every path in $G\backslash e$ from a vertex in A to a vertex in B contains the vertex y. Since y is not a cut vertex in G, there is a path from a vertex in A to a vertex in B that passes through e and not through y. Therefore, one end of e is in A and the other end is in B. Without loss of generality, we assume that $u \in A$ and $v \in B$. Note also that $y \neq u$ and $y \neq v$, since otherwise $G - y = (G\backslash e) - y$ in which case y would be a cut vertex of G.

Now consider G/e. Here u and v coalesce to form a single vertex, which we shall call w. We proceed by proving the following claim:

3.4.1. *If G/e is not 2-connected, then w is a cut vertex of G/e.*

Subproof. Suppose that G/e is not 2-connected. By Lemma 3.1, G/e is connected. So G/e has a cut vertex z. Since z is a cut vertex of G/e, using Lemma 2.19 we obtain a partition $\{C, D\}$ of $V(G/e) \setminus \{z\}$ such that every path of G/e from a vertex in C to a vertex in D uses the vertex z. We arbitrarily choose a vertex $c \in C$, and a vertex $d \in D$.

Towards a contradiction, assume that $z \neq w$, so z is unambiguously a vertex of G. As G is 2-connected, the graph $G - z$ is connected (see Exercise 3.3). Thus there is a path P in G from c to d that does not use the vertex z. If P contains neither u nor v, then, by Lemma 1.9, P is also a path in G/e from c to d, which is a contradiction. We will argue in the same manner when P contains at least one of u or v. There are three possibilities to consider:

(1) e is in the path P. Then the path obtained from P by replacing u, e, v with w is a path in G/e from c to d that does not use z.
(2) P contains at most one of u and v. In this case by replacing either u or v with w in P, we obtain a path in G/e from c to d that does not use z.
(3) P contains both u and v, but not e. View P as a subgraph of G. Adding e to the connected graph P gives a connected graph, which we can think of as the edge-induced subgraph $G[E(P) \cup \{e\}]$. By Lemma 3.1, we can contract e and stay connected; i.e., the graph $G[E(P) \cup \{e\}]/e$ is connected. Thus there is a path from c to d in $G[E(P) \cup \{e\}]/e$, and hence in G/e, and this path does not use the vertex z.

In all cases there is a path in G/e from c to d that does not use the vertex z. From this contradiction we deduce that $z = w$. $\lhd$

Using 3.4.1, we now only need to prove that w cannot be a cut vertex of G/e. We first prove:

3.4.2. *For all $s, t \in V(G) \setminus \{u, v\}$, there is a walk from s to t that avoids both u and v.*

Subproof. Recall that y is a cut vertex of $G \backslash e$ such that $y \notin \{u, v\}$, and there is a partition $\{A, B\}$ of $V(G \backslash e) \setminus \{y\}$ such that every path in $G \backslash e$ from a vertex in A to a vertex in B contains the vertex y. Consider $G - u$. This graph is connected, so there is a path from any vertex $r \in A \setminus \{u\}$ to y that avoids u.

Note that the vertices of this path do not contain any vertex in B as otherwise y would not be a cut vertex of $G\backslash e$. In particular, this path avoids v.

Similarly, there is a path from any vertex $r \in B \setminus \{v\}$ to y that avoids both u and v. Taken together, we see there is a path in G from any vertex $r \in V(G) \setminus \{u, v, y\}$ to y that avoids both u and v.

Consider two vertices $s, t \in V(G) \setminus \{u, v\}$. Clearly 3.4.2 holds if one of these vertices is y, so we may assume $s, t \in V(G) \setminus \{u, v, y\}$. Thus there is a path from s to y and a path from y to t that avoid u and v; combining these, we get a walk from s to t that avoids both u and v. $\triangleleft$

Let s and t be vertices of G that are distinct from u and v. By 3.4.2, there is a walk from s to t that avoids both u and v. Thus, there is a walk in G/e from s to t that avoids w, by Lemma 1.9. We deduce that for all vertices s and t of G/e other than w, there is a walk from s to t that does not contain w. Hence the graph obtained from G/e by deleting the vertex w is connected, so w is not a cut vertex of G/e.

We conclude that G/e does not have any cut vertices and hence G/e is 2-connected. $\square$

→ When G is 2-connected, the previous theorem tells us that it is not possible that neither $G\backslash e$ nor G/e is 2-connected. However, it could be that both $G\backslash e$ and G/e are 2-connected. Note that when we say "either A or B", we allow for the possibility that both A and B occur. (To indicate mutually exclusive possibilities, we could instead say "precisely one of A or B holds".)

Recall that a graph is connected if there is a walk between every pair of vertices. There is a characterisation of 2-connected graphs that is of a somewhat similar style. Before we can prove this (as Theorem 3.7), we require the following lemma.

Lemma 3.5. *Let G be a 2-connected graph. If e and f are parallel edges in G, then $G\backslash e$ is 2-connected.*

Exercise 3.6. *Prove Lemma 3.5*

An *isolated vertex* of a graph is a vertex of degree zero. In other words, an isolated vertex is not incident with any edges.

Theorem 3.7. *Let G be a loopless graph with at least three vertices, and no isolated vertices. Then G is 2-connected if and only if, for every pair $\{e, f\}$ of edges of G, there is a cycle of G that contains both e and f.*

Proof. ($\Rightarrow$) Suppose G is 2-connected. Then G is connected, and has no cut vertices. If G has precisely three vertices, then G can be obtained from K_3 by adding parallel edges. It is easy to check that such a graph has the property that every pair of edges is contained in a cycle. Thus, this direction holds when G has precisely three vertices.

To prove this direction in general, we use induction on the number of edges. For the base case, observe that G has at least three edges, and if G has precisely three edges, then G has three vertices, which we handled in the previous paragraph.

Assume now that G has at least four edges and, for induction, that this direction holds for loopless 2-connected graphs on at least three vertices with fewer edges than G. We may assume that G has at least four vertices (since the case where G has precisely three vertices was handled in the first paragraph). Let e and f be two arbitrarily chosen edges of G. Then it suffices to find a cycle containing e and f. Let h be another edge of G, so $h \notin \{e, f\}$. By Theorem 3.4, either $G \backslash h$ or G/h is 2-connected.

Say $G \backslash h$ is 2-connected. Then, by the induction hypothesis, there is a cycle C of $G \backslash h$ that contains both e and f. But C is also a cycle of G, so we have found the required cycle.

Thus we may assume that $G \backslash h$ is not 2-connected, in which case G/h is 2-connected. If G/h has a loop, h say, then the edges h and h are parallel in G, in which case $G \backslash h$ is 2-connected by Lemma 3.5. So we may also assume that G/h is loopless. Then, by the induction hypothesis, there is a cycle C of G/h containing both e and f. By Lemma 1.12, either C or $C \cup \{h\}$ is a cycle of G. In either case we have a cycle of G containing both e and f.

($\Leftarrow$) [redacted] $\square$

> We had to say that G was "loopless", had "at least three vertices", and had "no isolated vertices" in the above theorem. Why was that?

The previous theorem showed that for any pair of edges in a 2-connected graph, there is a cycle that contains them both. The next corollary shows that we can say the same thing for any pair of *vertices* in a 2-connected graph.

Corollary 3.8. *Let G be a 2-connected graph. Then for all $u, v \in V(G)$, there is a cycle of G containing both u and v.*

Proof. [redacted] □

> Note that Corollary 3.8 applies even when $u = v$. In the proof, a little extra care was required to make sure we also handled this case.

Blocks. When a graph G is not connected, we can decompose G into components, where each component is a maximal connected subgraph.

When a graph G is connected but not 2-connected, it can be useful to decompose G into smaller subgraphs that are "more highly connected". However, things are not quite as simple as for components.

A natural first attempt would be to consider the maximal 2-connected subgraphs of G. However, suppose G is a large tree. Then, although G is connected and not 2-connected, it has no 2-connected subgraphs! The problem here is the presence of bridges. Each bridge corresponds to a K_2 subgraph, which is not contained in any 2-connected subgraph, but is not 2-connected itself because it is too small. We can get around the problem by decomposing G into maximal connected subgraphs that have no cut vertices (with no restriction on their size). These are known as the "blocks" of a graph.

Formally, we define these as follows. Let G be a loopless graph. A graph is *biconnected* if it is connected and has no cut vertices. Note that a 2-connected graph is biconnected, but a biconnected graph can consist only one or two vertices, in which case it is not 2-connected. A *block* of G is a maximal subgraph of G that is biconnected. (In other words, H is a block of G if H is biconnected and there is no biconnected subgraph H' of G that contains H as a subgraph, other than H itself).

For a simple connected graph G with no isolated vertices, each block of G is either 2-connected or isomorphic to K_2.

Blocks have several useful properties.

Lemma 3.9. *Let G be a loopless graph. Then both of the following hold:*

(i) *any two distinct blocks of G have at most one vertex in common, and*
(ii) *every edge of G belongs to a unique block of G.*

In order to prove Lemma 3.9, it is useful to introduce a notion of "graph union". Intuitively, given two graphs G_1 and G_2, we want to construct a new graph by combining G_1 and G_2. More specifically, we want the new graph to have vertex set $V(G_1) \cup V(G_2)$ and edge set $E(G_1) \cup E(G_2)$. For this to be well defined, we require any edge that is in both G_1 and G_2 to be incident to the same set of (one or two) vertices in both G_1 and G_2. We say that graphs G_1 and G_2 are *consistent* if any edge that is in both G_1 and G_2 is incident with the same set of vertices – that is, for every $e \in E(G_1) \cap E(G_2)$, we have $\phi_{G_1}(e) = \phi_{G_2}(e)$. Now, for consistent graphs G_1 and G_2, we define the *union* of G_1 and G_2, denoted $G_1 \cup G_2$, to be the graph with vertex set $V(G_1) \cup V(G_2)$ and edge set $E(G_1) \cup E(G_2)$.

More rigorously, the incidence function of $G_1 \cup G_2$ is the one inherited from G_1 and G_2; that is,

$$\phi_{G_1 \cup G_2}(e) = \begin{cases} \phi_{G_1}(e) & \text{if } e \in E(G_1) \\ \phi_{G_2}(e) & \text{if } e \in E(G_2) \setminus E(G_1). \end{cases}$$

Exercise 3.10. *Prove that if G_1 and G_2 are consistent connected graphs with at least one vertex in common, then $G_1 \cup G_2$ is connected.*

Now, we return to proving Lemma 3.9.

Exercise 3.11. *Prove Lemma 3.9*

Why did we restrict ourselves to loopless graphs when defining blocks? Suppose we use the same definition, but allow graphs with loops. Consider the graph obtained from P_3 by adding a loop to the cut vertex. This graph would have two blocks, both of which contain the loop, so it is not true that the loop edge belongs to a unique block; i.e. Lemma 3.9(ii) fails. (We could get around this by changing the definition of a block to use a slightly different notion of cut vertex, but here we choose just to focus on blocks in loopless graphs.)

Lemma 3.12. *Let G be a loopless connected graph that is not 2-connected, with $|V(G)| \geq 3$. Then every cut vertex of G belongs to at least two distinct blocks.*

Exercise 3.13. *Prove Lemma 3.12*

We can associate with any loopless graph G a bipartite graph $B(G)$ with bipartition $\{\mathcal{B}, S\}$ where $\mathcal{B}$ is the set of blocks of G, and S is the set of cut vertices of G, such that there is an edge in $B(G)$ between $B \in \mathcal{B}$ and $v \in S$ if and only if the block B of G contains the cut vertex v. We call $B(G)$ the *block-cut graph* of G.

Lemma 3.14. *Let G be a loopless graph, and let $B(G)$ be the block-cut graph of G. Then all of the following hold:*

(i) *$B(G)$ is a forest.*
(ii) *If G is connected, then $B(G)$ is connected.*

Exercise 3.15. *Prove Lemma 3.14*

k-connected graphs. We have now seen two notions of connectivity of different strengths: a graph can be connected or 2-connected, where the latter is "more well-connected".

Generalising the definition of a 2-connected graph, we can obtain even stronger notions of connectivity. We first require some definitions.

Let G be a graph, and recall that, for $v \in V(G)$, we let $G - v$ denote the graph obtained from G by removing v and all edges incident with v. Similarly, for a set X of vertices of G, we let $G - X$ denote the graph obtained from G by removing each vertex in X along with all edges incident with at least one vertex in X.

Let G be a graph and let X be a set of vertices of G. We say that X is a *vertex cut* if $G - X$ has more components than G. Usually we are interested in the case that G is connected; in this case, X is a *vertex cut* if $G - X$ is disconnected. We refer to a vertex cut of size k as a *k-vertex cut*.

→ Note the difference between a "vertex cut" and a "cut vertex". A vertex v is a cut vertex if and only if $\{v\}$ is a 1-vertex cut.

→ You may have seen another sort of "cut" in MATH261; unfortunately the term "cut" can get quite overloaded. We will not use "cut" in this sense in MATH361.

Let k be an integer, with $k \geq 2$. A graph is *k-connected* if it is connected, it has at least $k+1$ vertices, and has no vertex cut of size at most $k-1$.

→ Intuitively, a graph is k-connected for "large" k if we cannot easily disconnect the graph by deleting vertices. More specifically, we would need to delete at least k vertices to disconnect a k-connected graph.

We also say that a graph is 1-connected if it is connected. Note that if a graph is k-connected, then it is also j-connected for smaller positive values of j.

Lemma 3.16. *Let G be a graph, and let j and k be positive integers. If G is k-connected, then G is j-connected for any $j \leq k$.*

Exercise 3.17. *Prove Lemma 3.16.*

For a connected graph G, the *connectivity* of G is the smallest k such that G is k-connected. We also say that a disconnected graph has connectivity 0.

→ The notion of connectivity that we have defined is sometimes known as "vertex connectivity", since it concerns the minimum number of vertices that would need to be removed to disconnect it. One could define an analogous notion for edges, in which case a "2-edge-connected" graph would be a connected graph with no bridges.

Although we can lose the property of being k-connected by deleting a vertex, we can't lose the property of being $(k-1)$-connected.

Lemma 3.18. *Let G be a k-connected graph with a vertex v, for some integer $k \geq 2$. Then $G - v$ is $(k-1)$-connected.*

Proof. Since G is k-connected, $|V(G)| \geq k+1$, so $|V(G-v)| \geq k$. Moreover, if $G - v$ is not connected, then v is a cut vertex of G, so G is not 2-connected. But then G is not k-connected (since $k \geq 2$), which is contradictory. So $G - v$ is connected.

It remains to show that $G - v$ has no j-vertex cuts for each positive integer $j < k-1$. Towards a contradiction, suppose that X is a j-vertex cut of $G - v$

for some $j < k-1$. Then $(G-v)-X$ is disconnected. But $(G-v)-X = G-(X \cup \{v\})$, so $X \cup \{v\}$ is a $(j+1)$-vertex cut. As $j+1 < k$, this shows that G is not k-connected, which is a contradiction. So $G-v$ has no j-vertex cuts for $j < k-1$. This shows that $G-v$ is $(k-1)$-connected. □

Similarly, although we can lose the property of being k-connected by deleting an *edge*, we can't lose the property of being $(k-1)$-connected.

Lemma 3.19. *Let G be a k-connected graph with an edge e, for some integer $k \geq 2$. Then $G \backslash e$ is $(k-1)$-connected.*

Proof. Let G be a k-connected graph. Then G has no vertex cut of size at most $k-1$. Towards a contradiction, suppose that $G \backslash e$ is not $(k-1)$-connected. Then $G \backslash e$ has a vertex cut S of size at most $k-2$. Say G has c components. Then $G-S$ also has c components, but $(G \backslash e)-S = (G-S) \backslash e$ has more than c components, so e is a bridge in $G-S$. Let $e = uv$. Then (by Lemma 2.18) either $S \cup u$ or $S \cup v$ is a vertex cut in G of size at most $k-1$, a contradiction. We deduce that $G \backslash e$ is $(k-1)$-connected. □

Each time we raise the connectivity of a graph we get stronger properties. For example, in a connected graph there is a path joining every pair of vertices, but in a 2-connected graph we have the stronger property that there is a cycle containing every pair of edges.

A common way to prove theorems in graph theory is to use induction. If we want to use induction to prove something about, say, 2-connected graphs, then we would like to be able to find a slightly smaller minor that is still 2-connected. That's why results like Lemma 3.1 and Theorem 3.4 are so useful.

We now have a trade-off: stronger connectivity gives stronger properties, but stronger connectivity is harder to keep in minors, and we can lose the ability to use induction!

For large k, say $k = 100$, we could certainly prove strong theorems about 100-connected graphs. But such high connectivity is artificial and typical graphs will rarely have that property. For example, consider the road map of a large city. This will be highly connected in some sense, but, for sure there will be some bottlenecks, so it is unlikely to be even 4-connected.

The upshot is that it's not really worth us playing the game of studying graphs of high connectivity. But it does turn out that it really is worthwhile going to 3.

3-connected graphs. Consider what it means to be k-connected when $k = 3$. A graph with at least four vertices is 3-connected if it is connected, it has no cut vertices, and no vertex cuts of size two.

It turns out that being 3-connected is a very useful property to have, but you will have to wait to see why.

Theorem 3.4 tells us that for every edge e of a 2-connected graph, either $G\backslash e$ or G/e is 2-connected.

Exercise 3.20. *Give an example of a 3-connected graph G with an edge e such that neither $G\backslash e$ nor G/e is 3-connected.*

Life in 3-connected graphs is clearly somewhat more difficult; but there is some good news.

Theorem 3.21. *Let G be a 3-connected graph with at least five vertices. Then there exists an edge e such that G/e is 3-connected.*

Such an edge may well be all we need to get induction going if we want to prove things about 3-connected graphs!

Before we prove Theorem 3.21 we need a lemma.

Lemma 3.22. *Let G be a 3-connected graph with at least five vertices and let $e = xy$ be an edge of G such that G/e is not 3-connected. Then $\{x, y, z\}$ is a 3-vertex cut of G.*

Proof. [redacted] □

Proof of Theorem 3.21. Suppose that the theorem fails. Then for any edge e, the graph G/e is not 3-connected. Thus, for any edge $e = xy$, there exists a vertex z such that $G - \{x, y, z\}$ is not connected, by Lemma 3.22.

Key Step: Choose the edge e and the vertex z in such a way that $G - \{x, y, z\}$ has a component F with as many vertices as possible.

Observe that $H = G[V(F) \cup \{x, y\}]$ is 2-connected, as if it has a cut vertex q, then $\{z, q\}$ is a 2-vertex cut of G, which would contradict that G is 3-connected.

Let $f = zu$ be an edge joining z to a vertex in a component of $G - \{x, y, z\}$ different from F. By assumption, G/f is not 3-connected, so by Lemma 3.22 again there is a vertex v such that $\{z, u, v\}$ is a 3-vertex cut of G.

We do not know where v is. But because H is 2-connected, if $v \in V(H)$, then $H-v$ is certainly connected (see Exercise 3.3 or Lemma 3.18). So let $H' = H-v$ if $v \in V(H)$; otherwise let $H' = H$. Then H' is contained in a component of $G-\{z,u,v\}$. But H' contains at least one more vertex than F. This contradicts the choice of e and v in the Key Step.

From this contradiction, we deduce that G has an edge e such that G/e is 3-connected. □

Exercise 3.23. *Give an example of a 3-connected graph that has no edge e such that $G\backslash e$ is 3-connected.*

Exercise 3.24. *Give an example of an infinite family of 3-connected graphs such that, for every edge in each graph, when deleted the resulting graph is not 3-connected.*

Menger's theorem. Here is a very fundamental problem. We have two non-empty sets of vertices of a graph G, say X and Y, such that $|X| = |Y| = k$. When is it possible to find k disjoint paths joining the vertices in X to the vertices in Y?

Note that X and Y need not be disjoint, but if there exists some $v \in X \cap Y$, then v itself is a path from X to Y.

Before we can solve this problem we need to develop some terminology. Two paths in a graph are *vertex disjoint* if they have no vertices in common.

Exercise 3.25. *Let P and Q be vertex-disjoint paths in a graph. Prove that they have no edges in common.*

Our main interest will be in vertex-disjoint paths, however we can also say that two paths are *edge disjoint* if they have no edges in common.

Exercise 3.26. *Give an example of two paths in a simple graph that are edge disjoint but not vertex disjoint.*

Let G be a graph with vertices x and y. An *(x, y)-path* in G is a path whose ends are x and y. More generally, for non-empty sets X and Y of vertices of G, we say that an *(X, Y)-path* in G is a path that begins at a vertex in X and ends at a vertex in Y.

Let G be a graph, let X and Y be non-empty sets of vertices in G, and let S be another (possibly empty) set of vertices in G. We say that S *separates* X *from* Y if every (X, Y)-path contains a vertex in S.

Our goal is to find the maximum number of pairwise vertex-disjoint (X, Y)-paths. We first obtain an upper bound for this number.

Lemma 3.27. *Let G be a graph with non-empty sets $X, Y \subseteq V(G)$, and $S \subseteq V(G)$. Suppose that S separates X from Y, where $|S| = t$. Then there are at most t pairwise vertex-disjoint (X, Y)-paths.*

Proof. [redacted] $\square$

> Earlier we defined when *two paths* are vertex disjoint, but here we are talking about when a *collection* of paths are vertex disjoint. In this case, we can say the paths are *pairwise* vertex disjoint to clarify that any pair of paths in the collection should be vertex disjoint. Henceforth, we will sometimes be a bit lazier: when we refer to a collection of vertex-disjoint paths, we mean that they are pairwise vertex-disjoint paths.

We can finally state and prove Menger's theorem.

Theorem 3.28 (Menger 1927)**.** *Let G be a graph and let $X, Y \subseteq V(G)$ be non-empty. Then the minimum size of a set that separates X from Y is equal to the maximum number of vertex-disjoint (X, Y)-paths.*

Proof. Let S be a set that separates X from Y, where we choose S to have minimum size amongst all such sets. Say $|S| = k$. By Lemma 3.27, the maximum number of vertex-disjoint (X, Y)-paths is at most k. Our task now is to prove that there are in fact k vertex-disjoint (X, Y)-paths.

The proof is by induction on the number of edges. If G has no edges, then each (X, Y)-path is a trivial path of length 0 for some $v \in X \cap Y$. That is, there are $|X \cap Y|$ vertex-disjoint (X, Y)-paths. The set S that separates X from Y must contain each vertex $v \in X \cap Y$, since the trivial path v is an (X, Y)-path. That is $X \cap Y \subseteq S$. Moreover, $X \cap Y$ separates X from Y, so $|X \cap Y| = |S| = k$, as required. This proves the base case, when G has no edges. Thus we may assume that G has at least one edge and, for induction, that the theorem holds for graphs with fewer edges than G. We first handle the case where G has an edge with both ends in X, or both ends in Y.

3.28.1. *If G has an edge $e = uv$ with $\{u, v\} \subseteq X$ or $\{u, v\} \subseteq Y$, then there are k vertex-disjoint (X, Y)-paths in G.*

Subproof. Let $e = uv$ be an edge of G. Without loss of generality, say $\{u, v\} \subseteq X$. Let Z be a set that separates X from Y in $G\backslash e$ where Z is chosen to be of minimum size. Then, by the induction assumption, there are $|Z|$ vertex-disjoint (X, Y)-paths in $G\backslash e$. Hence there are $|Z|$ vertex-disjoint (X, Y)-paths in G. But Z is also a set that separates X from Y in G, so $|Z| \geq k$ (since the minimum size of a set separating X from Y has size k). That is, there are k vertex-disjoint (X, Y)-paths in G, as required. $\triangleleft$

By 3.28.1, we may assume that G has no edges where both ends are in X, or both ends are in Y.

Let $e = uv$ be an edge of G, and consider G/e, where w is the vertex that replaces $\{u, v\}$. If either u or v is in X, then let X' be the set obtained from X by replacing u or v with w; otherwise let $X' = X$. We define Y' in the same way.

We now consider (X', Y')-paths in G/e. Suppose Z' separates X' from Y' in G/e, and Z' is of minimum size. Then, by the induction assumption, there are $|Z'|$ vertex-disjoint (X', Y')-paths in G/e. If $|Z'| = k$, then there are k vertex-disjoint (X', Y')-paths in G/e, and it follows that there are k vertex-disjoint (X, Y)-paths in G, as required. So we may assume that $|Z'| \leq k - 1$.

Suppose $w \notin Z'$. Then Z' is a set that separates X from Y in G, with $|Z'| \leq k - 1$. But then $|Z'| < |S|$, which contradicts that S is a set separating X from Y in G of minimum size. Therefore $w \in Z'$. Let $Z = (Z' \setminus \{w\}) \cup \{u, v\}$. Then Z separates X from Y in G, with $|Z| = k$.

Now consider a set that separates X from Z. Observe that no such set of size at most $k - 1$ exists, as such a set would also separate X from Y, contradicting that S is minimum-sized. But $\{u, v\} \subseteq Z$, so, by 3.28.1, there are k vertex-disjoint (X, Z)-paths in G. Similarly there are k vertex-disjoint (Z, Y) paths in G. We may combine these paths to obtain k vertex-disjoint (X, Y)-paths, as required. $\square$

We can now use Menger's theorem to answer the question from the start of this section.

Corollary 3.29. *Let X and Y be non-empty sets of vertices of a graph G with $|X| = |Y| = k$. Then there are k vertex-disjoint (X, Y)-paths if and only if there is no set of size less than k that separates X from Y.*

→ Menger's theorem is one of the cornerstones of graph theory. In essence, the maximum number of disjoint (X, Y)-paths measures the amount of *communication* that can occur between X and Y and that, one way or another, is one of the most fundamental problems in modern life. Think of the internet, traffic flow, phone networks, etc.

Let G be a graph with distinct vertices x and y. If we are looking for vertex-disjoint $(\{x\}, \{y\})$-paths, the problem is not very interesting: any such path must contain both x and y, so there can be at most one such path (and such a path exists if and only if x and y are in the same component). A more interesting question is to instead ask about the number of so-called internally vertex-disjoint (x, y)-paths. Suppose G has (x, y)-paths P_1 and P_2. The paths P_1 and P_2 are *internally vertex disjoint* if $V(P_1) \cap V(P_2) = \{x, y\}$.

New question: when is it possible to find k internally vertex-disjoint paths joining distinct vertices x and y in a graph?

Let x and y be distinct non-adjacent vertices of G. For a set $S \subseteq V(G) \setminus \{x, y\}$, we say that S *separates* x *from* y if there is no path from x to y in $G - S$.

We have overloaded the term *separates* here. That is, when we say "S separates x from y", we mean something slightly different to "S separates $\{x\}$ from $\{y\}$". The difference is: in the first case, S cannot contain x or y, whereas in the second, $\{x\}$ (or $\{y\}$) is an example of a set that separates $\{x\}$ from $\{y\}$.

→ Although this may seem confusing, it is for good reason: in general, when looking for a set S that separates a set X from a set Y, we are happy for S to contain vertices of X or Y; however, when looking for a set S' that separates a vertex x from a vertex y, this question is only interesting when S' doesn't contain x and y (for the reasons we spoke about just above).

At first glance, it might seem like Theorem 3.28 doesn't help us to answer the above question, but we in fact obtain the answer as a corollary (Corollary 3.30), which we will come to in a moment.

First, a useful definition. Let v be a vertex of a graph G. We define the *neighbourhood* of v in G to be the set of vertices that are adjacent to v, and denote this set as $N_G(v)$, or just $N(v)$ when the graph is clear from context.

Note that $v \in N(v)$ if and only if v is incident to a loop.

Corollary 3.30. *Let G be a graph and let x and y be distinct non-adjacent vertices of G. Then the minimum size of a set that separates x from y is equal to the maximum number of pairwise internally vertex-disjoint (x, y)-paths.*

Proof. Let $X = N(x)$ and $Y = N(y)$. Since x and y are non-adjacent, each (x, y)-path in G contains an (X, Y)-path that avoids $\{x, y\}$, and each (X, Y)-path that avoids $\{x, y\}$ can be extended to an (x, y)-path by appending x at the beginning and y at the end. Moreover, two such (X, Y)-paths are vertex disjoint if and only if the corresponding (x, y)-paths are internally vertex disjoint. It follows that the maximum number of internally vertex-disjoint (x, y)-paths in G equals the maximum number of vertex-disjoint (X, Y)-paths in $G - \{x, y\}$. By Theorem 3.28, this number is equal to the minimum size of a set in $G - \{x, y\}$ that separates X from Y. Let Z be such a set, so $|Z|$ is the minimum size of a set in $G - \{x, y\}$ that separates X and Y.

It remains to show that $|Z|$ is also the minimum size of a set in G that separates x from y. Suppose there were a smaller set S that separates x from y. Then every (x, y)-path contains a vertex in S. As every (X, Y)-path in $G - \{x, y\}$ extends to an (x, y)-path (by appending x at the beginning and y at the end, neither of which is in S), we have that every (X, Y)-path in $G - \{x, y\}$ also contains a vertex in S. That is, S separates X from Y in $G - \{x, y\}$. But $|S| < |Z|$, contradicting that the minimum size of a set is $|Z|$. We deduce that $|Z|$ is the minimum size of a set that separates x from y in G, as required. $\square$

Note that Corollary 3.30, stated above, is sometimes referred to as Menger's theorem, rather than Theorem 3.28. It gives us another way of characterising k-connected graphs, as shown by the next corollary:

Corollary 3.31. *Let G be a graph, let k be a positive integer, and suppose G has at least $k + 1$ vertices. The graph G is k-connected if and only if, for every pair of distinct vertices u and v in G, there are at least k internally vertex-disjoint paths between u and v.*

Proof. [redacted] □

Menger's theorem is also a powerful tool for proving other theorems in graph theory.

Exercise 3.32. *Prove the "only if" direction of Theorem 3.7 using Menger's theorem. That is, prove that if G is a 2-connected graph with no loops, then, for all $e, f \in E(G)$, there is a cycle that contains both e and f.*

Separations. We have seen that vertex cuts are a useful notion: we used these to define the connectivity of a graph, thereby obtaining a measure of "how well connected" a graph is. We have also seen the notion of when a set S separates two sets of vertices X and Y, which we used to describe Menger's Theorem. These two notions are related, as shown by the next lemma.

Lemma 3.33. *Let G be a connected graph, and let S be a subset of $V(G)$. Then S is a vertex cut of G if and only if there exist non-empty sets $X, Y \subseteq V(G) \setminus S$ such that S separates X from Y.*

Proof. [redacted] □

Note that it is important that G is connected in the statement of Lemma 3.33, for the "if" direction to hold. To see this, suppose that G has two components on vertex sets A and B. Then $\emptyset$ separates A from B, but $\emptyset$ is not a vertex cut. More generally, after choosing any $A' \subseteq A$ and $B' \subseteq B$ and $S = V(G) \setminus (A' \cup B')$, we have that S separates X' from Y', but S is not a vertex cut.

Another way to describe not only a vertex cut, but also the vertex sets that are separated by the cut, is using the notion of a "separation".

Let $G = (V, E)$ be a graph, and let A and B be subsets of V. Then the pair $\{A, B\}$ is a *separation* of G if

(i) $A \cup B = V$, and
(ii) there is no edge joining a vertex in $A \setminus B$ to $B \setminus A$.

- The *boundary* of the separation $\{A, B\}$ is the set $A \cap B$.
- The *order* of the separation $\{A, B\}$ is $|A \cap B|$.
- A separation is *proper* if both $A \setminus B$ and $B \setminus A$ are non-empty.

Lemma 3.34. *Let G be a connected graph and let X be a non-empty set of vertices of G, with $|X| = k$. The graph G has a proper separation of order k with boundary X if and only if G has a k-vertex cut X.*

Proof. [redacted] □

The previous lemma shows that in a connected graph, k-vertex cuts and proper separations of order k are really two different ways of thinking about the same thing. However, a separation $\{A, B\}$ describes not only a k-vertex cut (i.e. $A \cap B$), but also which vertices get separated by that cut (i.e. $A \cap B$ separates $A \setminus B$ and $B \setminus A$).

→ Remember Lemma 2.19, from when we introduced cut vertices? It is just a special case of Lemma 3.34: i.e. the ($\Leftarrow$) direction in the case that $k = 1$.

The next two characterisations show that we can think of connectivity purely in terms of separations.

Lemma 3.35. *Let G be a non-empty graph. The graph G is connected if and only if G has no proper separations of order 0.*

Proof. Suppose G has a proper separation $\{A, B\}$ of order 0. Then $\{A, B\}$ is a partition of $V(G)$ and there are no edges between A and B. Pick $a \in A$ and $b \in B$. Then there is no path from a to b in G, so G is disconnected.

Now suppose G is disconnected. Let C be the vertex set of one component of G. Then $\{C, V(G) \setminus C\}$ is a proper separation of order 0. □

Next we prove a corollary of Lemmas 3.34 and 3.35.

Corollary 3.36. *Let k be an integer with $k \geq 2$, and let G be a graph with at least $k+1$ vertices. The graph G is k-connected if and only if G has no proper separation of order less than k.*

Proof. Suppose G is not k-connected. Then either G is not connected, or it has a vertex cut of size at most $k-1$. In the former case, G has a proper separation of order 0, by Lemma 3.35; while in the latter case, G has a proper separation of order at most $k-1$ by Lemma 3.34. Either way, G has a proper separation of order less than k.

Conversely, suppose G has a proper separation of order j, where $j < k$. If $j = 0$, then G is not connected, by Lemma 3.35. If $j \geq 1$, then G has a j-vertex cut, by Lemma 3.34. In either case, G is not k-connected. $\square$