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] □

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. $\lhd$

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? Consider the graph obtained from P_3 by adding a loop to the cut vertex. This graph has 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 using 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 connected 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.*