Graph Connectivity

1 Introduction

We have seen that trees are minimally connected graphs, i.e., deleting any edge of the tree gives us a disconnected graph. What makes trees so susceptible to edge deletions? One intuitive reason is that trees have a an almost constant average degree (if $T = (V, E)$ then the average degree is $2(|V| - 1)/|V| < 2$). What happens as the average degree increases? Do we get graphs that have better connectivity? In this lecture we study the connectivity of graphs, and cover Mander's theorem (which roughly states the converse, that is high connectivity implies large minimum degree) and Menger's theorem. Before we proceed, we need to introduce some notation.

¶1. Notation: Throughout this lecture we will use $G = (V, E)$ for an undirected graph. A **subgraph** $G' = (V', E')$ of $G = (V, E)$ is a graph for which $V' \subseteq V$ and $E' \subseteq E$. Given $V' \subseteq V$, the subgraph of G that has its vertex set V' and the edges with all their vertices in V' as its edge set is called the **subgraph induced** by V' , and denoted as $G[V']$. We can similarly define a subgraph induced by a subset of edges.

Given a subset $U \subseteq V$, $G - U$ denotes the graph obtained by deleting the vertices U from V and *all the edges incident on* U from E; when U is a singleton $\{v\}$, we simply write $G - v$, and informally say "removing the vertex v from G". Given a subset X of edges, by $G - X$ we mean the graph $(V, E \setminus X)$, i.e., the graph obtained by deleting the edge associations in X, but not deleting the vertices; for a single edge e, we simply write $G - e$, and it informally means "removing an edge e from G".

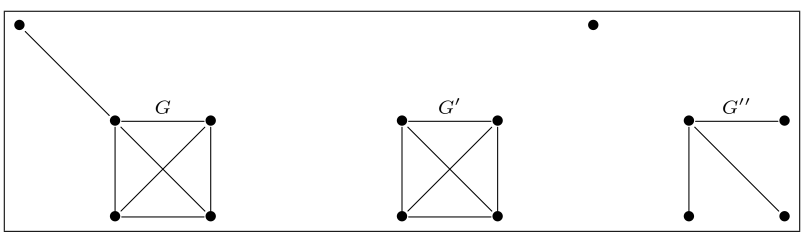

Figure 1: A graph G and two subgraphs G', G''; G' is an induced subgraph but not G''.

¶2. Vertex Cuts: A **vertex** v **separates** G if removing v from G causes the component of G containing v to become disconnected. In other words, there is a pair of vertices in G such that all the paths between these two vertices pass through v. We call v a **cut vertex**. A graph G is said to be k**-vertex-connected** if for all subsets $U \subseteq V$, $|U| < k$, $G - U$ remains connected. The maximum value of k for which G remains k-vertex-connected is called its **vertex connectivity**, $\kappa(G)$ [1]. So, for the complete graph K_n, $\kappa(K_n) = n - 1$. Also, $\kappa(G) = 0$ iff G is already disconnected or it is a single vertex.

¶3. Edge Cuts: Similar to above, we say that an **edge** e **separates** G if $G - e$ causes the component of G containing e to become disconnected. In other words, there is a pair of vertices in G such that all the paths between these two vertices pass through e. We call e a **bridge**. It is not hard to see that bridges

[1] An alternative definition is a minimization one: the minimum number of vertices that can be deleted to make G disconnected; see Bondy-Murty.

cannot be contained in any cycle (in fact, that is one way to characterize them). From this characterization it follows that another way to characterize a tree is a connected graph with all edges as bridges. We can represent an edge-cut for a graph as $[S, \overline{S}]$, where the edges are contained in the set $S \times \overline{S}$. A graph G is said to be **k-edge-connected** if for all subsets $X \subseteq E$, $|X| < k$, $G - X$ remains connected. The maximum value of k for which G remains k-edge-connected is called its **edge connectivity**, $\kappa'(G)$. What is $\kappa'(K_n)$? It is $n-1$ again, because every vertex has degree $n-1$ and to disconnect a vertex we have to remove these edges.

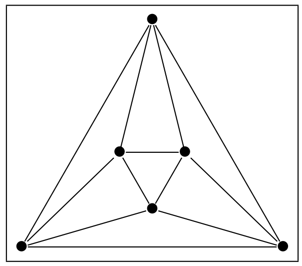

Figure 2: The octahedron graph with $\kappa = \kappa' = 4$.

The following theorem shows the relation between vertex-connectivity, edge-connectivity and minimum degree of the graph.

THEOREM 1 (Whitney's Theorem (1932)). *For a graph G,*

$$\kappa(G) \leq \kappa'(G) \leq \textit{minimum degree of } G.$$

Proof. Let's prove the easier identity: κ' is at most the minimum degree of the graph. Certainly, if we remove all the edges incident on a vertex with least degree then we get a disconnected graph. Thus we have the upper bound on κ'.

The $\kappa(G) \leq \kappa'(G)$ intuitively makes sense, since "burning the bridges" is the same as "burning the posts on one side of the bridges". Let's be more precise. Suppose $[S, \overline{S}]$ is an edge cut with the smallest size. Our aim is to construct a vertex-cut of a size at most κ'. We consider two cases:

- When all the vertices in S, $\overline{S}$ are neighbours. In this case, the size of the edge-cut is $\kappa' = |S| \cdot |\overline{S}| = k(n-k)$, which is greater than $(n-1)$, for all values of $k = 1, \ldots, n-1$. But, no matter what, $\kappa \leq n-1$. Thus $\kappa \leq \kappa'$.

- Otherwise there is a pair of vertices (x, y) such that $x \in S$, $y \in \overline{S}$ and xy is not an edge. We want to construct a vertex cut that separates these two vertices. The path from x to y has to go through a neighbouring vertex of x. There are two possibilities for the neighbour: it either belongs to $\overline{S}$, or it is in S. Let B be all the neighbours of x of the first type. For x to connect to y via a neighbour in S, the neighbour has to reach one of the other bridge-heads in $S - x$. Let A be all the bridge-heads except for x. Then it is clear that $A \cup B$ forms a vertex-cut for xy. The cardinality of this vertex set is the number of bridges, which is κ'. Thus $\kappa \leq \kappa'$. See Figure 3. Note that we could have also separated x and y by deleting the set $\{x\} \cup A$; however, this is not allowed x has to be present after deletion; also $S \setminus (\{x\} \cup A)$ may be empty.

Q.E.D.

We now further explore what k-connected graphs look like. What is a 1-connected graph? It is simply a connected graph. What is 2-connected, or 3-connected graph?

2 Two-connected Graphs

A connected graph that has no cut vertices is called a **block**. In other words, every block is a 2-connected graph. A **block of a graph** G is a maximal subgraph that is 2-connected, or a bridge, or an isolated vertex.

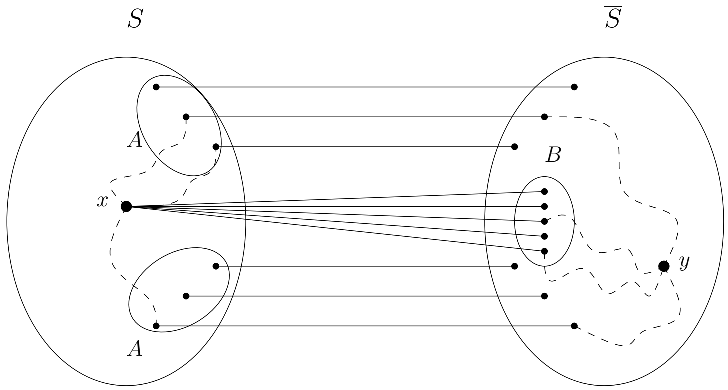

Figure 3: Constructing a vertex cut for the edge xy from the edge-cut $[S, \overline{S}]$. The dashed paths in $\overline{S}$ can crossover to S; to do that, however, they must use an edge ending in a vertex in A.

A graph can thus be represented as a union of its blocks; because of their maximality, two blocks of a graph can only meet at a single vertex, which must necessarily be a cut-vertex for the graph. In this sense, blocks are the 2-connected analogues of components. See Figure 4 for illustration.

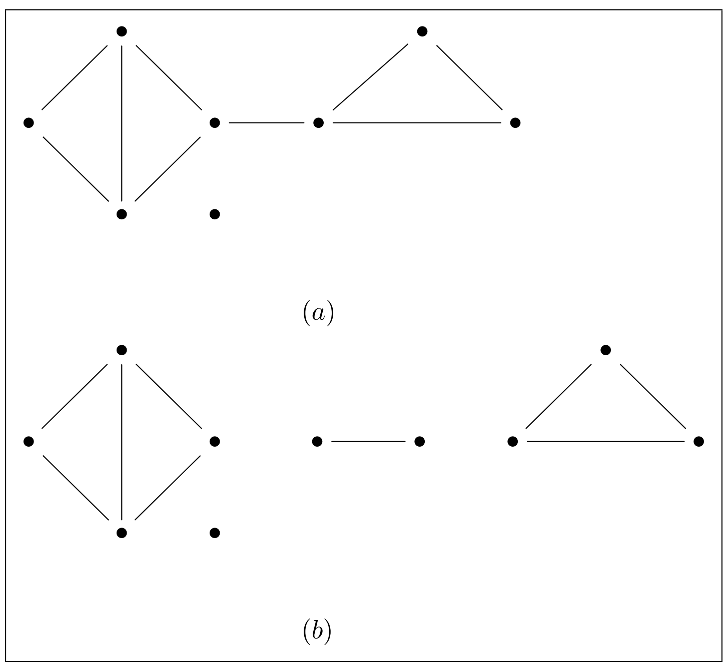

Figure 4: A graph and its block decomposition.

We now present a different characterization of blocks, but to do that we need the following definition: Two paths in a graph connecting the same pair of vertices are said to be **independent** if they do not share any vertices besides the starting and the ending vertex.

THEOREM 2 (Whitney (1932)). *A graph G with three or more vertices is a block iff each pair of vertices in G is connected by at least two independent paths, i.e., every pair of vertices is on a cycle.*

Proof. If any pair of vertices of G is connected by at least two independent paths then clearly there is no cut-vertex separating them. Thus G is 2-connected.

We give two proofs of the converse: one by induction, and other by contradiction.

Let u, v be two vertices in G. Our first proof will do an induction on the distance between u and v. The base case is when they are adjacent. If the only path between u and v is the edge uv, then this edge is a bridge, but this cannot be since G is 2-connected. Now suppose that all pairs of vertices at distance less than $k > 1$ are on a cycle. We then show that any pair of vertices at distance k are also on a cycle. Suppose u, v are two vertices at distance k. Consider a u to v path of distance k, and let w be the vertex that precedes v on this path; let P be that path from u to w. Since u to w are at a distance $k-1$, it follows that there is a u-w path Q independent of P. Since G is two connected, if we consider the graph $G - w$, then there is a path P' that connects u and v. Now there are three cases to consider:

1. P' is independent of P and Q. In that case, the paths Pv and P' are independent.

2. P' is independent of either P or Q, but not both; say it is independent of Q. In that case, the paths Qv and P' are independent.

3. P' is independent of neither P nor Q. Let x be the last common vertex on P' with either P or Q, i.e., the path $xP'v$ does not meet either P or Q. Since P' intersects both P and Q, there must be such a vertex; moreover, $x \neq w$, since P' was a u to v path in $G - w$. Say x is on P. Then the paths $uPxP'v$ and Qv are independent. This is illustrated in Figure 5. Note that if $x = v$, which will happen if Q passes through v, then Pv and uQv are already independent paths.

By way of contradiction, suppose that there are two vertices u, v such that all the paths between them are not independent, or equivalently, they are not on a cycle. Then we want to show that there is a cut vertex separating u and v. To argue about the paths from u to v, we want to first order the vertices. One way to do this is to do a depth first traversal (DFT) of G from u, label all the vertices in a pre-order traversal, and for each vertex w let $a(w)$ be the smallest ancestor that can be reached from w through one of its descendants; note that there can only be back edges in this traversal, since G is undirected. We want to characterize when are two vertices u and v on a cycle in such a DFT. It is clear that if $a(v) = u$ then they are on a cycle, however, this is not necessary. The only other possibility is the following: consider the path between v and $a(v)$; if there is a vertex w on this path such that $a(w) = u$ then u and v are also on a cycle. Thus two vertices are on a cycle iff in the path P from v to $a(v)$ there is a vertex w such that $a(w) = u$. In our setting, however, u and v are not on a cycle. The cut vertex separating u and v is the ancestor $a(w)$, where $w \in P$, closest to u. This is shown in Figure 6.

Q.E.D.

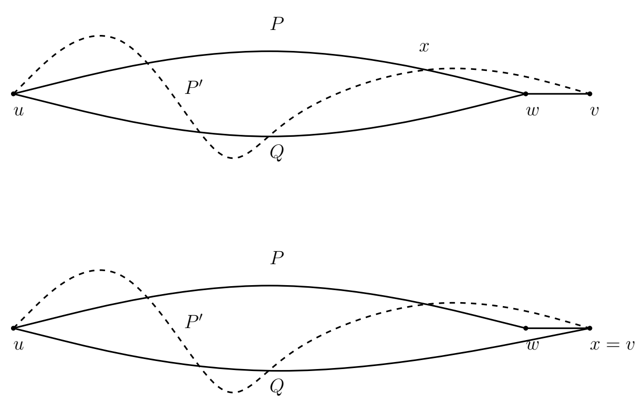

Figure 5: Construction of two independent paths

As a corollary it follows that every pair of vertices must be on a cycle, and hence the degree of every vertex is at least two. But we can say more about the structure of a block, to do that we introduce another definition: for $A, B \subseteq V$, an $A - B$ path is a path that starts from a vertex in A and ends in a vertex in B, and none of whose internal vertices are in $A \cup B$. The sets A, B may not be disjoint; if they are the same and equal to C then we just write C-path ("C" for the intended cycle, though the actual path may not introduce one). We call a graph **two-constructible** if it can be constructed as follows: any cycle is two-constructible;

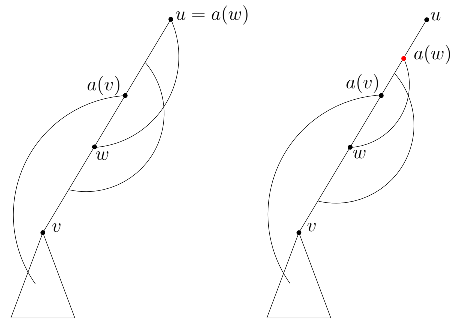

Figure 6: Finding a cut vertex between two vertices that are not on any cycle.

given a two-constructible graph C on, we construct a two-constructible graph from it by adding a C-path between any two vertices of C; such a construction has been illustrated in Figure 7.

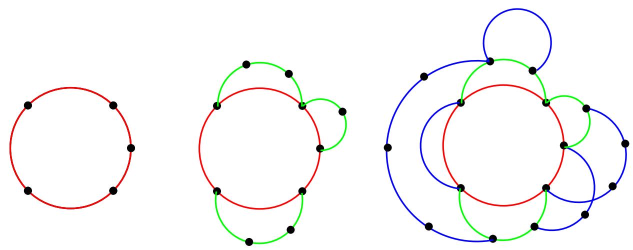

Figure 7: Constructing a two-constructible graph from a cycle

Theorem 3. *A graph is 2-connected iff it is two-constructible.*

Proof. The implication on one side is clear: every two-constructible graph is 2-connected. For the converse, we observe that there is cycle in G and hence it has a maximal subgraph C that is two-constructible. Moreover, C is an induced subgraph since for any pair of vertices $x, y \in C$, the edge xy is in C otherwise it would give us a C-path. We want to show that $C = G$. If not then there is an edge (v, w) such that $v \in G - C$ and $w \in C$. Since G is 2-connected, there must be another path P from v to w going through a different neighbour of v. Clearly, the path $wvPw$ forms a C-path, and is two-constructible. Thus we can extend C to a larger subgraph of G that is two-constructible. But this is a contradiction since C was the maximal subgraph of G with that property. Hence $G = C$. **Q.E.D.**

3 Three Connected Graphs

In Theorem 4 we saw that every 2-connected graph can be constructed in a prescribed manner from any cycle. What is the most basic 3-connected graph? Clearly, K_4 fits the bill. Tutte showed that K_4 plays a fundamental role in the construction of 3-connected graphs. The procedure, however, needs a different device: Given a graph G and an edge $e = xy$ in G, by G/e we mean the graph obtain by **contracting the edge** e, i.e., the graph with a new vertex v adjacent to all the vertices adjacent to x and y.

Lemma 4. *A 3-connected graph G with more than four vertices has an edge e such that G/e is 3-connected.*

Proof. We illustrate the proof in Figure 8.

1. Suppose for all edges $e = xy$, G/e contains a vertex cut of size two. The contracted vertex and another vertex z form a vertex cut in G/e. This implies that the triplet $\{x, y, z\}$ forms a vertex cut in G. Since G is 3-connected every component C of $G - \{x, y, z\}$ must contain a neighbour of x, y, z.

2. Amongst all edges e and vertex z that yield a vertex cut in G, choose an edge xy and a corresponding vertex z such that the component C has the smallest possible number of vertices. We will derive a contradiction to this assumption.

3. Let v be a neighbour of z in C. Note that the neighbourhood of v is strictly contained in C.

4. By the same argument as in (1) we know there is a triplet $\{z, v, w\}$ that forms a vertex cut for G. Moreover, w cannot be either x or y; say it is y then $\{z, v, y\}$ does not form a vertex cut as x is connected to all the components.

5. Since xy is an edge and $\{z, v, w\}$ forms a vertex cut there must be a component D in $G - \{z, v, w\}$ that does not contain x, y (otherwise $\{z, v, w\}$ do not form a vertex cut, since there is just one component containing x and y).

6. Now D contains a neighbourhood of v, and hence must be contained strictly inside C. That is $|D| < |C|$. But this gives us an edge zv, a vertex w and a component D, that is smaller than C, which is a contradiction to how we had chosen C.

Q.E.D.

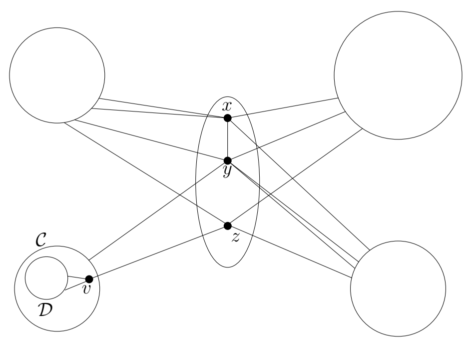

Figure 8: Every 3-connected graph has a contractible edge yielding a 3-connected graph.

We now give a characterization of 3-connected graphs by Tutte. To do that we need the following definition: A graph is **3-constructible** if there is a sequence of graphs $G_0, \ldots, G_n$ such that $G_0 = K_4$, and G_{i+1} is obtained by picking a vertex v of G_i, splitting it into two vertices x, y connected with an edge such that the both the vertices are connected to a subset of the neigbours of v, their degree is at least three, and each neighbour of v is a neighbor of at least one of x or y. Note: the resulting may be 4-connected, but it is still 3-connected.

Theorem 5 (Tutte 1961). *A graph is 3-connected iff it is 3-constructible*

Proof. If G is 3-connected then by the lemma above there always is a contractable edge. So there exists a sequence of graphs starting from G.

Suppose G_i is 3-connected but G_{i+1} is not then it has a vertex cut S of size two. Since xy is an edge in G_{i+1} there are two cases to consider:

1. Both x, y belong to the same connected component of $G_{i+1} - S$. But then $G_i = G_{i+1}/xy$ is not 3-connected, a contradiction.

2. Either x or y belongs to S. Then contracting xy still gives us a graph G_i that has a vertex cut of size two, which is again a contradiction.

So G_{i+1} is 3-connected. **Q.E.D.**

4 The Fundamental Theorem of Graph Connectivity

Menger's theorem was proposed by him in the context of topological studies in defining a curve (see Kurventheorie, and his description of the discovery of the theorem). However, the original proof had a gap, which was later corrected by Nöbling and independently by Menger. It was D. König who introduced the result in his book. Harary later called this result as the fundamental theorem of graph connectivity. The theorem gives us an alternate definition of k-connectedness. It describes the connections between the vertices in a k-connected graph.

THEOREM 6 (Menger's Theorem 1927). *Let G be a graph and $A, B \subseteq V$. Then the minimum number of vertices separating A from B in G is equal to the maximum number of vertex disjoint (or independent) $A-B$ paths in G.*

Let $k(A, B)$ be the minimum number of vertices needed to separate A from B in G. We prove the following stronger claim:

THEOREM 7. *Given $k = k(A, B)$, and $n < k$ disjoint $A - B$ paths $P_1, \ldots, P_n$, there exists $n + 1$ vertex disjoint $A - B$ paths $Q_1, \ldots, Q_{n+1}$ whose starting and end vertices contain the starting and end vertices of $P_1, \ldots, P_n$.*

Proof. The proof is by McCuaig and is a clever induction on $|G| - |B|$. Base case: $B = G$. Then $A \subseteq B$. and $k(A, B) = |A|$. The n independent paths $P_1, \ldots, P_n$ are simply n vertices of A. Since $n < |A|$, we know there must be at least one more vertex v separate from $P_1, \ldots, P_n$. These $n + 1$ vertices then are the $n + 1$ independent $A - B$ paths.

Induction hypothesis: True for all $\beta < |G| - |B|$. Let $a_j \in A$ be the starting vertex of P_j and $b_j \in B$ the ending vertex of P_j. Since $n < k$ the vertices $b_1, \ldots, b_j$ do not separate A from B. Thus there must be a $A - B$ path R that must avoid all the b_j's. If R is vertex disjoint from $P_1, \ldots, P_n$ then our desired set of $n + 1$ paths is $P_1, \ldots, P_n, R$. Otherwise, let x be the last vertex that is on R and another path P_j; by reordering the paths we can assume that this path is P_n (therefore, $x \neq b_n$, but may be $x = a_n$, i.e., R is the edge $a_n b_n$); see Figure 9. Let C be vertices on the path $x b_n$ and D the vertices on the path xR. Define $B' := B \cup C \cup D \cup \{x\}$; thus $k(A, B) \leq k(A, B')$. Then $P_1, \ldots, P_{n-1}, a_n P_n x$ are vertex disjoint $A - B'$ paths and since $n < k(A, B')$ by our inductive hypothesis we know that there are $n + 1$ vertex disjoint $A - B'$ paths $Q_1, \ldots, Q_{n+1}$ whose starting vertices contain $a_1, \ldots, a_n$ and ending vertices contain $b_1, \ldots, b_{n-1}, x$ and a new vertex $y \neq \{b_1, \ldots, b_{n-1}, x\}$. By reordering we can assume that Q_n is the path that ends in x and Q_{n+1} the path that ends in y. Since $y \neq x$, where can it come from. There are three cases to consider, and in each case we create two new $A - B$ paths that are vertex disjoint from $Q_1, \ldots, Q_{n-1}$:

1. $y \in C = xP_n$; note $b_n \in C$. In this case, we extend Q_n, which ends in x, to take xR, whereas Q_{n+1}, which ends in y, is extended to continue on the path yP_n and end in b_n.

2. $y \in D$. In this case, we extend Q_n, which ends in x, to take xP_n, whereas Q_{n+1}, which ends in y, is extended to continue on the path yR.

3. $y \in B \setminus \{C \cup D\}$. In this case y is a vertex in $B \setminus \{b_1, \ldots, b_n\}$. Thus we keep the path Q_{n+1} as such, but extend Q_n by xP_n to an $A - B$ path.

In all the three cases we have $n + 1$ disjoint $A - B$ paths as desired. See Figure 10 for an illustration.

Q.E.D.

As a corollary, we have the following local phenomenon:

Corollary 8. *If ab is not an edge, then the minimum number of vertices (distinct from a, b) separating a from b in G is equal to the maximum number of independent a-b paths.*

The result follows by applying the theorem to the neighborhoods of a and b. The following theorem captures the overall global phenomenon, and gives us an equivalent way of defining k-vertex-connectedness:

THEOREM 9. *A graph is k-vertex-connected iff it contains k independent paths between any two vertices.*

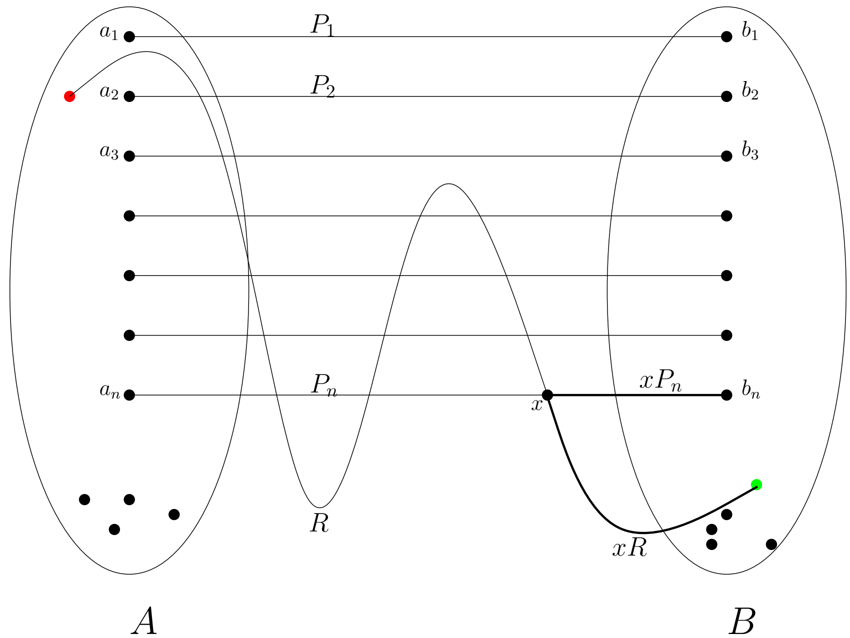

Figure 9: The path R and its intersection with P_n at the vertex x.

Proof. If every pair of vertices has k independent paths between them then we cannot separate them with fewer than k vertices, so k-connectedness is clear.

Conversely, suppose G is k-vertex-connected and there exists a pair of vertices v, w with fewer than k independent paths. By the corollary above, we know that vw is an edge in G. Let G' be the graph obtained by deleting the edge vw from G. Then there are at most $k-2$ independent paths connecting v, w, and hence a vertex cut X of size at most $k-2$ separating v from w in G'. Let a be any other vertex in G'; there is one because $|G| > k$. Then a cannot be connected to both v and w in $G' - X$, but it must be connected to one of them, say w, otherwise, X separates a from v and w by a vertex cut of size $\leq k-2$. If a is connected to w and not to v, then $X \cup \{w\}$ separates a from v, but this is a contradiction since we have a vertex cut of size $\leq k-1$ in a k-vertex-connected graph. **Q.E.D.**

Menger's theorem applies to edge-connectivity as well: A graph is k-edge-connected iff there are k edge disjoint paths between any two vertices.

The algorithmic aspect of finding a vertex cut, or independent paths reduces to solving a linear program.

¶4. Mixed Connectivity: A set $S \subseteq V$ and $T \subseteq E$ forms a **disconnecting pair** for vertices v, w if the two vertices are disconnected in $G - \{S \cup T\}$. The vertices v, w are (k, ℓ) connected if there is disconnecting pair (S, T) with $|S| < k$ and $|T| \leq \ell$ or $|S| \leq k$ and $|T| < \ell$. The Beineke-Harary conjecture from 1967 states the following:

Conjecture 1. *If vertices v, w are (k, ℓ)-connected in G then G has a system of $k + \ell$ edge-disjoint v-w paths of which k are vertex disjoint.*

Figure 11 shows an example of a $(2, 2)$ connected pair.

5 Applications of Menger's Theorem

In this section we use Menger's theorem to derive other fundamental results.

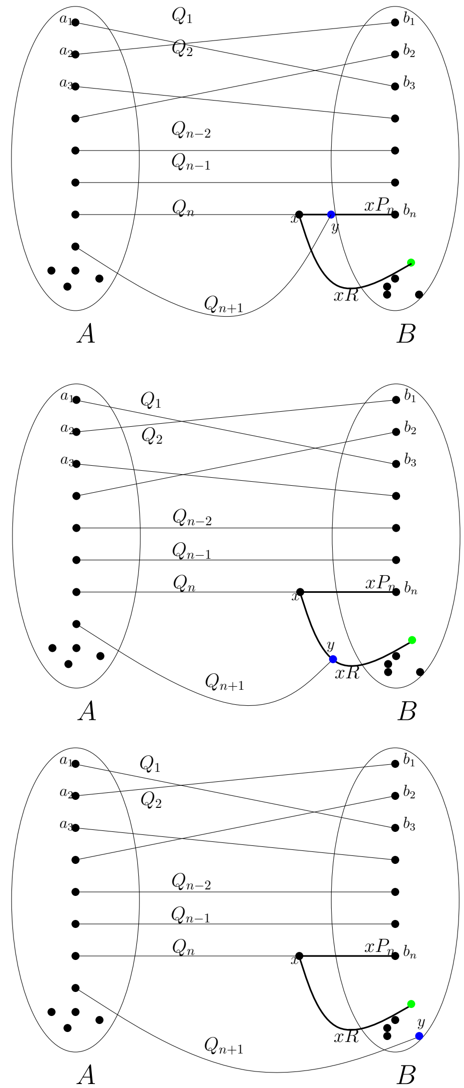

Figure 10: Three cases to consider when constructing $n+1$ vertex-disjoint $A-B$ paths from $Q_1, \ldots, Q_{n+1}$.

5.1 Hall's Theorem

We show the the sufficiency, that is, in a bipartite graph (A, B) if for all $S \subseteq A$, $|N(S)| \geq |S|$ then A can be matched in B. What we want to show is that there are $|A|$ independent (vertex disjoint) edges between A and B. Given any $S \subseteq A$, let (S, T) be a minimum vertex cut for (A, B), where $T \subseteq B$. Then we know that $N(\overline{S}) \subseteq T$, which implies

$$|T| \geq |N(\overline{S})| \geq |\overline{S}|,$$

where the last inequality follows from the assumption that $|N(\overline{S})| \geq |\overline{S}|$. But, $|\overline{S}| = |A| - |S|$. Therefore, we obtain

$$|T| + |S| \geq |A|.$$

However, the LHS above is the size of the vertex cut, which by Menger's theorem is equal to the maximum number of independent $A-B$ paths (in this case edges). Thus, to paraphrase, there are at least $|A|$ number of independent $A-B$ paths, and hence a matching for A in B.

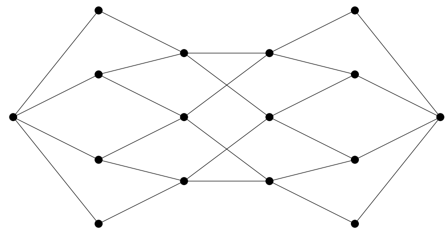

Figure 11: A graph that has mixed connectivity $(2, 2)$

6 Edge-disjoint Spanning Trees and Arboricity of Graphs

The edge-connectedness of a graph G ensures that between any pair of vertices in G there are k edge disjoint paths. What if we simultaneously want k edge-disjoint paths between all pairs of vertices? Clearly, we will have such paths if we can find k edge-disjoint spanning trees for G. This is a stronger condition that just edge-connectedness, and there are graphs for which high edge-connectivity does not imply existence of such spanning trees. Here we give a necessary and sufficient condition by Nash-Williams and Tutte (both gave it independently in 1961) for a graph G to have k edge-disjoint spanning trees.

The necessary condition is an obvious one. Consider any partition P of V into r sets. Then all the k spanning trees must connect these partitions; each tree must use at least $r-1$ edges to connect them, let's call these edges **cross-edges**; moreover, since the spanning trees are edge-disjoint no two share the same cross-edge; thus the number of cross-edges in the graph are at least $k(r-1)$. As is the case, this obvious condition is also sufficient. The theorem is derived for multigraphs

THEOREM 10 (Nash-Williams, Tutte 1961). *A multigraph contains k edge-disjoint spanning tress iff for every partition P of its vertices it has at least $k(|P|-1)$ cross-edges.*

As a corollary we have:

Corollary 11. *Every $2k$-edge-connected graph has k edge-disjoint spanning trees.*

Proof. Let $P_1, \ldots, P_r$ be a partition of the vertices. Then there are at least $2k$ cross-edges connecting P_i to the remaining partitions. Thus the number of cross-edges are at least $\sum_{i=1}^{r} 2k/2 = kr > k(r-1)$; note the two in the denominator is because each partition counts a cross-edge twice. **Q.E.D.**

Let $F := (F_1, \ldots, F_k)$ be a k-tuple of edge-disjoint forests, such that each F_i spans the vertices of the graph. Furthermore, let F be such that is has the maximal number of edges, i.e., $||F|| := \sum_i |E(F_i)|$ is as large as possible. Let $\mathcal{F}$ be the set of all such tuples. For our proof, we want to show the existence of such a k-tuple where each F_i is a tree.

Given an $F = (F_1, \ldots, F_k)$, let $e \in E \setminus E[F]$. Then for all i, $F_i + e$ must contain a cycle; otherwise F is not maximal. Moreover, the endpoints u, v of e are connected in every F_i. The following lemma states that there is a set U, containing u, v, of vertices that are connected in all F_i's.

LEMMA 12. *For every edge $e \in E \setminus E[F]$ there is a subset $U \subseteq V$ that is connected in every F_i and U contains the endpoints of e.*

Given the lemma above, we now prove the theorem by induction on $|G|$. By assumption, we know that for every partition P of V there are at least $k(|P|-1)$ edges. In particular, if we take P to be the individual vertices then there are at least $k(|V|-1)$ edges in G. Let $F = (F_1, \ldots, F_k) \in \mathcal{F}$. If all the F_i's are trees then we are done. Otherwise,

$$\|F\| = \sum_i \|F_i\| < k(|V|-1).$$

Thus the number of edges in $E[F]$ are strictly smaller than the number of edges in G. So there is an edge $e \in E \setminus E[F]$. From the lemma above we know that there is a corresponding subset U of vertices common to all F_i's.

Consider the contracted graph G/U. Any partition P' of G/U induces a partition P of G where the vertices U always belong to the same partition. Moreover, the number of cross-edges in P' and P are the same. Thus by our assumption, every partition of P' has $k(|P|-1) \geq k(|P'|-1)$ cross-edges and hence by the induction hypothesis G/U has k edge-disjoint spanning trees $T'_1, \ldots, T'_k$. Let v_U be the vertex corresponding to the contracted U in G/U. To construct edge disjoint spanning trees T_i's from T'_i's we do the following: in T'_i replace v_U by the sub-tree $F_i[U]$, which exists by the lemma above. Why are the resulting trees are edge-disjoint? Because, F_i's are edge-disjoint; the remaining edges of T'_i are not perturbed and hence they remain edge-disjoint by construction.