

Graph connectivity

Math 314

April 1, 2003

1 Connectivity

Recall that a cut-vertex is a vertex whose removal disconnects a graph. The following definition generalizes this concept.

Definition 1.1 A subset U of vertices of a graph G is called a *cut-set* of G if the graph $G - U$ is disconnected.

Note that we don't require a cut-set to just barely disconnect a graph (in other words, we don't require that any proper subset of a cut-set leave the graph connected). Thus, if U is a cut-set, then any subset of vertices of G containing U is also a cut-set.

Note also that we consider an empty graph K_0 (no vertices, no edges) disconnected since there is no path in this graph. However, a graph K_1 (one vertex, no edges) is considered connected since there is a path P_0 from the graph's only vertex to itself.

It follows that a complete graph has only one cut-set, which consists of all its vertices. Any non-complete graph G has a pair of non-adjacent vertices, so the rest of the vertices of G form a cut-set G (since their removal results in a graph with two isolated vertices). Of course, G may have smaller cut-sets as well.


Definition 1.2 The *connectivity* of a graph G , denoted $\kappa(G)$, is the order of the smallest cut-set of G .

This means that any proof that $\kappa(G) = n$ should have two parts: $\kappa(G) \leq n$ (i.e. there is a set of n vertices whose removal disconnects the graph) and $\kappa(G) \geq n$ (i.e. there is no set of less than n vertices whose removal disconnects the graph).

Example 1.3 A graph G is disconnected if and only if $\kappa(G) = 0$. A nontrivial (not K_0 or K_1) connected graph G has a cut-vertex if and only if $\kappa(G) = 1$. A nontrivial graph G has no cut-vertices if and only if $\kappa(G) \geq 2$. It is also pretty straightforward to determine that

$$\begin{aligned}\kappa(K_n) &= n \text{ for any } n \geq 0, \\ \kappa(P_n) &= 1 \text{ for } n \geq 2 \text{ (recall that } P_0 \cong K_1 \text{ and } P_1 \cong K_2), \\ \kappa(T) &= 1 \text{ for any tree } T \neq K_2, \\ \kappa(C_n) &= 2 \text{ for } n \geq 4 \text{ (recall that } C_3 \cong K_3).\end{aligned}$$

Definition 1.4 A graph G is called n -connected if $\kappa(G) \geq n$, in other words, if removal of any set of less than n vertices leaves the graph connected.

Figure 1: A graph illustrating Menger's theorem: $\kappa(u, v) = 3$, $S = \{w_1, w_2, w_3\}$ separates u and v 

Example 1.5 Any graph is 0-connected. A 1-connected graph is the same as connected graph. A 2-connected graph is any nontrivial graph with no cut-vertices.

There is an obvious upper bound on the connectivity of any non-complete graph G , namely the minimum degree of G . Removing all neighbors of a vertex of degree $\delta(G)$ disconnects it from the rest of the graph.

Fact 1.6 For any non-complete graph G , $\kappa(G) \leq \delta(G)$.

2 Menger's theorem

The notion of a cut-set can be applied not only to a whole graph but also to any two distinct non-adjacent vertices in a graph.

Definition 2.1 Let u and v be two distinct non-adjacent vertices in a graph G . We say that a subset S of vertices of G separates u and v , or is a (u, v) -cut-set, if $G - S$ does not contain any path between u and v . We denote by $\kappa(u, v)$ the order of the smallest (u, v) -cut-set.

Note that $\kappa(u, v) \leq \min(\deg u, \deg v)$ (see Figure 1 for an example where a strict inequality occurs). Note also that $\kappa(u, v) \geq \kappa(G)$ since a set of $\kappa(u, v)$ vertices disconnects G .

We can consider a different measure of how well u and v are connected.

Definition 2.2 A collection of paths is called *internally disjoint* if no interior vertex of any path belongs to another path.

Note that the endpoints of internally disjoint paths are not required to be distinct.

Notation 2.3 For any two u and v are distinct vertices, we denote by $id(u, v)$ the maximum number of internally disjoint u - v paths.

Note that if u and v are adjacent, then the edge uv counts as one of u - v paths.

Fact 2.4 For any distinct non-adjacent vertices u and v , $id(u, v) \leq \kappa(u, v)$.

PROOF. To separate u and v we need to at least disconnect each of the $id(u, v)$ internally disjoint u - v paths. Thus, we need to remove at least 1 vertex on each of those paths, for a total of at least $id(u, v)$ vertices. Therefore, $\kappa(u, v) \geq id(u, v)$. \square

However, a much stronger result, called Menger's theorem, is true.

Theorem 2.5 (Menger) *For any distinct non-adjacent vertices u and v , $id(u, v) = \kappa(u, v)$.*

PROOF. We only need to prove that $id(u, v) \geq \kappa(u, v)$. Our proof will be by induction on the order of the graph and the value of $\kappa(u, v)$.

Induction base. If $\kappa(u, v) = 0$, then there is no path between u and v , so $id(u, v) = 0 = \kappa(u, v)$. If $\kappa(u, v) = 1$, then at least 1 vertex must be removed to separate u and v , so there is a path between u and v , and hence $id(u, v) \geq 1 = \kappa(u, v)$.

Induction step. Let that $\kappa(u, v) = m \geq 2$. We want to prove that $id(u, v) \geq m$. Assume that the theorem is true for all graphs with less vertices than in our graph (call it G) and any $\kappa(u, v)$, as well as for a graph of any order but lower $\kappa(u, v)$ than in G . We will prove that the theorem is true for G as well.

For our proof we will need the following construction. Let S be any minimal (u, v) -cut-set (so S has $\kappa(u, v) = m$ vertices). Then any u - v path must contain an interior vertex which is in S . Construct two new graphs G_u and G_v as follows. To construct G_u , start by taking the subgraph of G consisting of all the parts of u - v paths from u to some vertex in S . Then adjoin a new vertex v' and edges from v' to each vertex in S . Similarly, to construct G_v , start by taking the subgraph of G consisting of all the parts of u - v paths from v to some vertex in S . Then adjoin a new vertex u' and edges from u' to each vertex in S . Then $\deg_{G_u}(v') = m = \deg_{G_v}(u')$. Clearly, S is a minimal (u, v') -cut-set in G_u as well as a minimal (u', v) -cut-set in G_v , so $\kappa_{G_u}(u, v') = m$ and $\kappa_{G_v}(u', v) = m$. We now consider two cases.

Case 1. Suppose there is a minimal (u, v) -cut-set S such that both G_u and G_v have less vertices than G . Then our induction hypothesis applies, so $id_{G_u}(u, v') = \kappa_{G_u}(u, v') = m$ and $id_{G_v}(u', v) = \kappa_{G_v}(u', v) = m$. Now delete the vertices u' and v' to obtain m internally disjoint paths from u to S and from v to S . Gluing these paths along S , we obtain m internally disjoint paths in G between u and v . Thus, $id(u, v) \geq m$ as desired.

Case 2. Now suppose that for *any* minimal (u, v) -cut-set S , either G_u or G_v (or both) is of the same order as G . Then $G_u \cong G$ and v is adjacent to each vertex in S , or $G_v \cong G$ and u is adjacent to each vertex in S . Remove all edges in G that can be removed without decreasing $\kappa(u, v)$.

Consider the shortest u - v path in this situation. We wish to prove that it has length 2. Suppose this is false and the shortest u - v path has length ≥ 3 , so it is of the form $u-u_1-u_2-\dots-v$. Note that u is not adjacent to u_2 , and v is not adjacent to u_1 , or else this would not be the *shortest* u - v path. We removed all the "slack" edges, so removal of the edge u_1u_2 must decrease $\kappa(u, v)$. Therefore, in $G - u_1u_2$ there is a (u, v) -cut-set T which has $m - 1$ vertices. The vertices u_1 and u_2 cannot be in T , otherwise T will be a (u, v) -cut-set in G with less than $\kappa_G(u, v) = m$ vertices, which is impossible. However, $G - u_1u_2 - T$ is disconnected, so both $G - u_1 - T$ and $G - u_2 - T$ are also disconnected. Therefore, both $T \cup \{u_1\}$ and $T \cup \{u_2\}$ are (minimal) (u, v) -cut-sets. Hence, either u is adjacent to all vertices in $T \cup \{u_1\}$, or v is adjacent to all vertices in $T \cup \{u_1\}$. But v is not adjacent to u_1 , so u is adjacent to all vertices in $T \cup \{u_1\}$, and hence, to all vertices in T . Similarly, either u is adjacent to all vertices in $T \cup \{u_2\}$, or v is adjacent to all vertices in $T \cup \{u_2\}$. But u is not adjacent to u_2 , so v is adjacent to all vertices in $T \cup \{u_2\}$, and hence, to all vertices in T . Thus, both u and v are adjacent to all vertices in T . Now T has $m - 1$ vertices and $m \geq 2$, so $m - 1 \geq 1$. That means u and v have a common neighbor (call it t) and the shortest u - v path, u - t - v , is of length 2.

Consider the graph $G - t$. It has less vertices than G , so $id_{G-t}(u, v) = \kappa_{G-t}(u, v)$ by the induction hypothesis. Now t must be in any (u, v) -cut-set, so $\kappa_{G-t}(u, v) = \kappa_G(u, v) - 1 = m - 1$, so $id_{G-t}(u, v) = m - 1$. In other words, $G - t$ has $m - 1$ internally disjoint u - v paths. Adding t back together with edges ut and vt , we obtain 1 more u - v path internally disjoint from those $m - 1$ paths. Thus, G has at least m internally disjoint u - v paths, so $id_G(u, v) \geq m$ as desired. This ends the proof. \square

Menger's theorem has several useful corollaries.

Theorem 2.6 *A nontrivial graph G is n -connected if and only if for any $n + 1$ distinct vertices of G , say u, u_1, \dots, u_n , there is a collection of n internally disjoint paths u - u_1, u - u_2, \dots, u - u_n .*

PROOF. Adjoin a new vertex v to G together with edges vu_i , $1 \leq i \leq n$. Then the new graph G' is n -connected if and only if G is n -connected. Also v is not adjacent to u , therefore $\kappa(u, v) = n$, i.e. $id(u, v) = n$. Since v has n neighbors, u_1, u_2, \dots, u_n , it follows that n internally disjoint u - v paths must contain all neighbors of v , one neighbor on every path. Therefore, there are n internally disjoint u - v paths in G' if and only if there are n internally disjoint paths u - u_i ($1 \leq i \leq n$) in G . \square

The following theorem, due to Whitney, gives a necessary and sufficient condition for a graph to be n -connected.

Theorem 2.7 (Whitney) *A nontrivial graph G is n -connected if and only if for any two distinct vertices u and v in G , there are (at least) n internally disjoint u - v paths.*

Note that Whitney's theorem does not require that the vertices u and v be nonadjacent, only distinct.

PROOF. (\Leftarrow) If u and v are nonadjacent and $id(u, v) \geq n$, then by Menger's theorem, $\kappa(u, v) = id(u, v) \geq n$. Thus, $\kappa(u, v) \geq n$ for any two nonadjacent vertices u and v in G , so $\kappa(G) \geq n$, so G is n -connected.

(\Rightarrow) Let G be n -connected. If u and v are not adjacent and $\kappa(u, v) \geq n$, then by Menger's theorem $id(u, v) = \kappa(u, v) \geq n$. If u and v are adjacent, then they are not adjacent in $G - uv$. Since we deleted only 1 edge, uv , it follows that $\kappa(G - uv) \geq \kappa(G) - 1 \geq n - 1$, so $id_{G-uv}(u, v) = \kappa_{G-uv}(u, v) \geq n - 1$. Thus, there are at least $n - 1$ internally disjoint u - v paths in $G - uv$. Adding back the edge uv , we get 1 more u - v path for a total of at least n internally disjoint u - v paths in G . Thus, $id_G(u, v) \geq n$ as desired. \square

There are many more properties that n -connected graphs have. Here is one of the more surprising ones that ties together seemingly unrelated things.

Theorem 2.8 *If G is n -connected ($n \geq 2$), then any n distinct vertices of G lie in a common cycle, i.e. G has a cycle which passes through all these n vertices.*

Of course, this common cycle may contain other vertices as well.

PROOF. Let U be the set of our n vertices, and let C be a cycle of G that contains the maximum number of vertices in U , say k . We will prove that C contains all of U , i.e. $k = n$.

We will start by proving that C must have at least n vertices. If not, then C has $\leq n - 1$ vertices, and U has a vertex v not in C . Since G is n -connected, Theorem 2.6 implies that there is a collection of internally disjoint vertices from v to every vertex in C . Then we can substitute any edge of C by the two internally disjoint paths from v to its endpoints to make a larger cycle, which contradicts the maximality of C . Thus, C has $\geq n$ vertices.

Suppose now that C contains $k < n$ vertices in U . Then C has a vertex not in U . Say C contains vertices u_1, \dots, u_k, u_{k+1} (among others), where $u_i \in U$ ($1 \leq i \leq k$) and $u_{k+1} \notin U$. Now $k < n$ means that $k + 1 \leq n$, so G contains $k + 1$ internally disjoint paths from v - u_1, v - u_2, \dots, v - u_k, v - u_{k+1} . Let v_i ($1 \leq i \leq k + 1$) be the

first vertex on the $v-u_i$ path from v to u_i that belongs to C (possibly $v_i = u_i$). Let P_i denote the subpaths $v-v_i$ of our paths $v-u_i$. Since C contains exactly k vertices in U , and there are $k+1$ vertices v_i , there are distinct integers $1 \leq s, t \leq k+1$ such that one of the two v_s-v_t paths along C contains no interior vertex belonging to U . Replace that v_s-v_t path by the two paths $v-v_s$ and $v-v_t$ to obtain a cycle that contains at least 1 more vertex in U (namely, v) than C . But this contradicts our definition of C . Therefore, we cannot have $k < n$, so $k = n$ as desired. \square

Finally, we remark that *edge-connectivity*, *edge cut-set* and *internally edge-disjoint* paths may be defined similarly to their vertex analogs discussed above. As might be expected, the edge-analogs of Menger's and Whitney's theorems are also true.