

Lecture 28: Menger's theorem

November 21, 2024

Kennesaw State University

1 The plan

Our goal for this lecture is to prove Menger's theorem:

Theorem 1.1 (Menger). *Let s, t be two non-adjacent vertices of G . Then G contains a collection of $\kappa(s, t)$ internally disjoint $s - t$ paths, where $\kappa(s, t)$ is the size of a smallest $s - t$ cut in G .*

Many proofs of Menger's theorem are known. None of them are particularly simple, but we will look at a proof where part of the complexity is hidden: it is pushed off to König's theorem, which we proved earlier in the semester. König's theorem will let us prove a special case of the theorem, and then we will reduce other cases to this special case.

The relationship between these two theorems is actually important in practice, where it is used in a different way: it lets us reduce the problem of finding bipartite matchings (which König's theorem deals with) to the problem of finding internally disjoint $s - t$ paths in Menger's theorem. This is useful, because linear programming and network flow methods let us solve the second problem efficiently.

A final thought: how do multigraphs and directed graphs interact with Menger's theorem?

It turns out that going from simple graphs to multigraphs doesn't really change anything about the problem. If we want to find internally disjoint $s - t$ paths, then adding loops or parallel edges will not make our task easier: even if there many edges between vertices v and w , we can only use one of them, because we can only use vertices v and w once. Loops and parallel edges will also not affect $s - t$ cuts: deleting a vertex also destroys all the edges using that vertex, no matter how many parallel edges there are.

There is a directed version of Menger's theorem. For directed graphs, a directed $s - t$ cut is specifically a set of vertices whose removal destroys all directed $s - t$ paths. (This is different from a directed $t - s$ cut, which destroys all directed $t - s$ paths!) If our internally disjoint $s - t$ paths are also directed, then Menger's theorem continues to hold.

We will not think about directed graphs when we prove Menger's theorem, but with a bit of care in the details, our proof will also work in that case.

¹This document comes from an archive of the Math 3322 course webpage: <http://misha.fish/archive/3322-fall-2024>


2 König's theorem and $s - t$ cuts

Recall that:

- $\alpha'(G)$ is the number of edges in a largest matching in G : a largest set of edges that share no endpoints.
- $\beta(G)$ is the number of vertices in a smallest vertex cover in G : a smallest set of vertices that contains at least one endpoint of every edge.

König's theorem says that when G is bipartite, $\alpha'(G) = \beta(G)$.

To see the connection to Menger's theorem, take the following graph as an example:



Every $s - t$ path in this graph must use at least one edge of the form $v_i w_j$. After all, it has to get from the top half of the graph to the bottom half somehow! And once we decide on an edge of this form to use, we might as well simply make the path go (s, v_i, w_j, t) . Even though there are very indirect paths like $(s, v_1, w_1, v_3, w_2, w_3, t)$, every such path contains a short $s - t$ path, plus some extra vertices.

As a result:

- A set of vertices is an $s - t$ cut if and only if it destroys all paths of the form (s, v_i, w_j, t) . We cannot remove either s or t , so an $s - t$ cut must contain either v_i or w_j for every edge $v_i w_j$.

This is exactly the same as a vertex cover of the bipartite graph consisting of all edges between $\{v_1, v_2, v_3\}$ and $\{w_1, w_2, w_3, w_4\}$.

- A collection of internally disjoint $s - t$ paths might as well be reduced to a collection of internally disjoint paths of the form (s, v_i, w_j, t) . These paths are internally disjoint if they do not share any of the v and w vertices.

This happens exactly when the $v_i w_j$ edges used by the path form a matching of the bipartite graph.

So we see that for a graph like this one, Menger's theorem says exactly the same thing as König's theorem!

What is "a graph like this one"? We can apply this argument whenever we can split $V(G) - \{s, t\}$ into a set $\{v_1, v_2, \dots, v_k\}$ (the neighbors of s) and a set $\{w_1, w_2, \dots, w_\ell\}$ (the neighbors of t). In other words, when every vertex of G other than s and t is adjacent to one of s or t , but not both.

3 Dealing with other cases

Now we just have to handle all the other cases, when the graph does not have this structure.

Our reasoning here will be that in all other cases, we can keep reducing the problem to a simpler problem: we can turn G into a smaller graph H such that proving the theorem for H will also prove it for G . If we can always do this, then it means that the case we've already solved is in some sense the only hard case.

You can think of this as a complicated induction on the number of vertices in G , where the base case is the case we proved in the previous section.

3.1 Vertices adjacent to both s and t

One of the ways G can fail to have the “König-structure” is if there is a vertex u adjacent to both s and t .

In this case, u must be part of every $s-t$ cut: if u is not deleted, there is an $s-t$ path (s, u, t) .

Also, the path (s, u, t) is going to be internally disjoint to every $s-t$ path not containing u .

So in this case, let $H = G - u$. If we could prove Menger's theorem for H , we would get $\kappa_H(s, t)$ internally disjoint $s-t$ paths in H . What's more, we know that $\kappa_H(s, t) = \kappa_G(s, t) - 1$, because we've deleted a vertex which is part of every $s-t$ cut, so we now have $\kappa_G(s, t) - 1$ internally disjoint $s-t$ paths.

We can get the last path by taking the path (s, u, t) , which is guaranteed to be internally disjoint from every $s-t$ path in H . Therefore if Menger's theorem is true for H , it is also true for G .

3.2 Vertices not part of a minimum cut

In this section and the next, we'll deal with vertices u that are adjacent to neither s nor t .

These vertices are very easy to deal with if they are not part of *any* minimum $s-t$ cut. In that case, once again, we'll let $H = G - u$, and show that Menger's theorem for H implies Menger's theorem for G .

If u is not part of any minimum $s-t$ cut, that exactly means $\kappa_H(s, t) = \kappa_G(s, t)$. If we deleted a vertex that is not part of some efficient way to disconnect s from t , we have not made any progress.

Therefore if Menger's theorem is true for H , it gives us $\kappa_G(s, t)$ internally disjoint $s-t$ paths in H , and that proves the theorem for G even without having to use u .

3.3 Nontrivial minimum cuts

The last case of the theorem is the trickiest. Here, we have to deal with vertices u not adjacent to s or t , but that are still important—they're part of some minimum $s-t$ cut.

In such a case, we'll take a step back and pick some $s-t$ cut U , with $|U| = \kappa_G(s, t)$, that does not just consist of s 's neighbors or t 's neighbors.

Because s is not connected to t in $G - U$, we can split up G into two pieces:


- G_s is the induced subgraph “between s and U ”. It includes all vertices that can be reached by a path from s without going through a vertex in U (but including the vertices in U themselves).

Because U is an $s - t$ cut, G_s does not contain the vertex t ; it may be missing some other vertices.

- G_t is the induced subgraph “between U and t ”. It includes U , and all the vertices that *cannot* be reached by a path from s without going through a vertex in U .

The only vertices G_s and G_t have in common are the vertices of U .


Here’s an illustration in the case $|U| = 2$ (where $U = \{u_1, u_2\}$):



Now we reduce the problem not to one smaller problem but two.

- In the first graph, H_1 , leave G_s alone and replace G_t by only the single vertex t adjacent to u_1, \dots, u_k .
- In the second graph, H_2 , leave G_t alone and replace G_s by only the single vertex s adjacent to u_1, \dots, u_k .

Here is an illustration of what happens to our graph in this case:



It’s important to know that the set U does not just consist of neighbors of s , or not just of neighbors of t . This means that both H_1 and H_2 actually have fewer vertices than G : we really are reducing to a smaller problem.

What do $s - t$ cuts in H_1 and H_2 look like? We can think of an $s - t$ cut in H_1 as being a set of vertices that *prevents us from leaving G_s* (because if we leave G_s , we end up in t). Deleting those same vertices in G will also prevent us from leaving G_s (because up until we leave G_s , the two graphs are the same). So an $s - t$ cut in H_1 is also an $s - t$ cut in G . We conclude that $\kappa_{H_1}(s, t) \geq \kappa_G(s, t)$. In the same way, we prove $\kappa_{H_2}(s, t) \geq \kappa_G(s, t)$.

Actually, we must have $\kappa_{H_1}(s, t) = \kappa_G(s, t) = \kappa_{H_2}(s, t)$ for the following reason: U continues to be an $s - t$ cut in all three graphs!

So if we prove Menger’s theorem for both H_1 and H_2 , we can get $\kappa_G(s, t)$ $s - t$ paths in H_1 , and also $\kappa_G(s, t)$ $s - t$ paths in H_2 . These can be glued together to get k internally disjoint $s - t$ paths


in G ; here's how.

For each i , take the $s - t$ path we found in H_1 that goes through u_i : $(s, (\text{some stuff}), u_i, t)$. Also, take the $s - t$ path we found in H_2 that goes through u_i : $(s, u_i, (\text{some other stuff}), t)$. Join them together at u_i :


$$(s, (\text{some stuff}), u_i, (\text{some other stuff}), t).$$

These are internally disjoint: they can't intersect in G_s (because otherwise the corresponding paths in H_1 would also intersect) and they can't intersect in G_t (because otherwise the corresponding paths in H_2 would also intersect).

Here's an illustration of how to do this in our example. Here are two $s - t$ paths in our H_1 and H_2 :



When we combine them, we get two $s - t$ paths in G :



So once again, proving Menger's theorem for the two smaller cases H_1, H_2 proves Menger's theorem for G .

We've successfully done this reduction in all cases that we can't prove by using König's theorem, so we have a complete proof of Menger's theorem!

4 Practice problems

There are many extensions of k -connectivity that we don't have time to cover. Here are a few to think about.

1. Prove the following generalization of Dirac's fan lemma (see the practice problems from the previous lecture):

Let s be a vertex in G and let U be a subset of $V(G) - s$ with $|U| \geq \kappa(G)$. Then we can find a subset $T \subseteq U$ with $|T| = \kappa(G)$ such that G contains an $s - T$ fan which does not share any other vertices with U .

(Hint: start by forgetting the clause "which does not share any other vertices with U ".)

2. Prove that if G is 3-connected, then for any three vertices u_1, u_2, u_3 , G contains a cycle that passes through all three of them.

(Hint: the previous problem can help.)

More generally, it is true that if G is k -connected, then for any k vertices, G contains a cycle through all k of them (in some order). If you're brave, try proving this by induction on k .

3. In a graph G , an $s - t$ edge cut is a set of edges whose removal disconnects s from t . We write $\kappa'(s, t)$ for the least number of edges in an $s - t$ edge cut. (This is now always defined, whether or not s and t are adjacent.)

Here, too, we can prove a lower bound on $\kappa'(s, t)$ by finding a collection of $s - t$ paths. What condition should replace "internally disjoint" here, and why?

(By the way, Menger's theorem also holds for edge cuts.)

4. There is a surprising connection between connectivity and planar graphs.

In general, we can define the dual graph of a plane embedding, but the dual graph of a planar graph is not well-defined: different plane embeddings can give different, non-isomorphic dual graphs.

However, if G is a planar graph and $\kappa(G) \geq 3$, then this does not happen: two dual graphs that come from different plane embeddings of G are necessarily isomorphic.

Proving this is hard. For this exercise, just give an example of a planar graph G with $\kappa(G) = 2$ in which two different embeddings give non-isomorphic dual graphs, and think about how the 2-vertex cut is "responsible" for this problem.