

PROFESSOR: So let's get set to state Hall's theorem in a way that doesn't mention boys and girls. But let's remember the girl-boy setup to start. So the general setup is a bipartite graph H . And a bipartite graph has two sets of vertices-- the girl vertices and the boy vertices.

Formally, there's a set L of H , called the left vertices of H , and a set R of H , called the right vertices of H . The vertices of H altogether are L of H union R of H . Both of these are non-empty. And they don't overlap.

Then the edges of H have the property that they only go between L of H and R of H . That is the definition of a bipartite graph. Now, we're interested in a matching in a bipartite graph. So let's be precise again of what's a matching without having to mention boys and girls and like.

A match is a total injective function from the left vertices to the right vertices. So that means every L vertex, or girl, has a match m of L that is on the other side. And we need this total injective function, match function m , to follow the edges.

What does follow the edges mean precisely? It simply means that the edge $l m$ of l is a legitimate edge of H . Another way to say that is that the graph of this total injective function is a subset of the edges of H .

So we can state the no-bottleneck condition, Hall's condition, as follows. If the size of S is less than or equal to its image under the edges, for every set of left vertices H , that's called Hall's condition, that the size of S is less than or equal to the image of S for every S , then there is a match. So this is a way of precisely stating Hall's theorem.

No boys and girls mentioned. We'll be comfortable going back to the boy-girl language, because it's more memorable and easier to talk about. But just for the record, we have now formally stated, defined bipartite graphs, matches in bipartite graphs, and Hall's theorem really that there is a match when Hall's condition is satisfied without mentioning boys and girls.

The puzzle is, how do you verify that there are no bottlenecks? It's pretty hard. The bottleneck condition involves checking that every subset S of L of H satisfies this cardinality condition. And there are a lot of subsets of L of H relative to the size of the graph, exponentially many.

So we can ask, how do you verify that there are no bottlenecks? Well, it turns out in algorithms classes, you will learn a fairly efficient matching procedure, runs about quadratically, for finding

perfect matches when they exist. But there is one particular special case that ensures no bottlenecks that we'll be making use of several times in the term.

And that special situation is when it turns out that every girl likes at least d boys and every boy likes at most d girls. This is called a degree-constrained graph. If a graph is degree constrained, then there are no bottlenecks. So this condition that each girl likes at least d boys and every boy likes at most d girls could have been rephrased as saying that the minimum degree of the girls is greater than or equal to the maximum degree of the boys. But I think expressing it in terms of d maybe makes it a little clearer.

Let's prove this. I'm going to prove that, if you have a degree-constrained bipartite graph, then it, in fact, satisfies Hall's condition. There are no bottlenecks. So we have that, if every girl likes at least d boys and every boy likes to most d girls, there are no bottlenecks. The proof goes as follows.

Let's look at some arbitrary set S of girls. And suppose that the total number of edges that are incident to S has cardinality e . There are e edges that touch vertices in S . Well, that tells us that since every vertex in S has at least d edges out-- maybe more-- d times the size of S has to be less than or equal to the total number of edges coming out of S .

Likewise, all of the edges that come out of S go, by definition, into E of S . And each vertex in E of S can only absorb d edges, d girls. So that means that the total number of edges absorbed by E of S -- that is, at most d times E of S -- had better be bigger than the number of edges that have to be absorbed e . So we get that d times the size of S is less than or equal to the total number of edges incident to S , which is less than or equal to d times the total number vertices in E of S .

Well, I'll cancel the d 's. And we have that the size of S is less than or equal to the size of its image. That's the definition of a bottleneck-- The violation of that would be a bottleneck. This says, there's no bottlenecks. And we have proved the degree-constrained condition is sufficient to verify Hall's condition. And by Hall's theorem, it's sufficient to guarantee that there is a match.

Now, there's a lot of graphs with matches that are not degree-constrained. This is not an if and only if theorem. It's just a sufficient condition that comes up often enough that it's worth mentioning. Degree-constrained implies Hall's condition is satisfied, which implies that there's

a perfect match.

So let's turn now to the general case of Hall's theorem. And let's talk about proving Hall's theorem. Hall's theorem says that if there's no bottleneck, then there is a match.

So let's begin by setting up a lemma that will play a crucial role. The strategy for proving it is basically going to be to try to break the problem of finding a match for a large set of vertices into sub problems of finding matches for smaller sets of vertices. And by strong induction, Hall's condition will apply to the subsets and will win. Let's look at that more carefully.

So let's begin by supposing that there are no bottlenecks in some bipartite graph H . Well, in particular, if there's no bottlenecks anywhere, then if you restrict yourself to some set S of girls, no subset of that set S is going to have a bottleneck, obviously, because a bottleneck within S would be a bottleneck within the whole graph. So that part's trivial.

What's not trivial and takes a little bit of arguing is that suppose you have a set of girls with the property that the number of boys that are compatible with that set of girls is exactly the same as the number of girls. So technically, the size of S is equal to the size of E of S . In that case, we can argue that there aren't any bottlenecks within the complement of S and the complement of E of S either.

Well, let's look at a picture to illustrate what's going on. So here we have a bipartite graph. And there's S and its image on the right E of S . And we know that a bottleneck here would imply a bottleneck in the whole graph.

What we want to argue is that a bottleneck in the complement of S would imply a bottleneck in the complement of E of S . So there's the complement of E of S . And there's the complement of S .

Now, notice that some edges that come out of the complement of S may very well go into the E of S . That is, this would be a point that's both in E of S and in E of S complement. But we're not allowed to use that point, because we're trying to find a match only within S bar and E of S bar. And what we're really trying to argue is that a bottleneck here would imply a bottleneck in the whole graph.

So let's see if we can argue this. We want to claim that there's no bottleneck between S bar any E of S bar given that there's no bottleneck anywhere and when S and E of S are the same size. So let's suppose that we had a set T over here that was a subset of S bar. And let's look

at its image over here in orange.

So we've got that the image of this set T , when we restrict it just to the part that's in the complement of S -- notice I'm leaving out this point here. I'm not taking the image of T . I'm taking the image of T restricted to the points that are not in the image of S . And what I want to know is, could that cause a bottleneck?

After all, I've left out some points that used to be included in the image of T . And so the image of T might have been bigger than T because of those extra points that I'm leaving out, I have to be sure that that doesn't happen. Could there be a bottleneck that has been created by these points that I've left out? Could there be a bottleneck in E of T intersection, a complement to E of S . Could that orange guy be a bottleneck?

Well, if it was a bottleneck, then I'd have a bottleneck in the whole graph. And that's because S and E of S are the same size, because all I do is, if I had a bottleneck there, I would take this set T along with S and this set orange along with E of S , and that would be a bottleneck. Because now, if you look at the size of T , I've added the same amount to the size of T as I've added to that orange set, because E of S and S are the same size.

So if this orange part was smaller than T , then orange part along with E of S is smaller than T union S . And that defines a bottleneck in the whole graph. So again, if there was a bottleneck in here caused by restricting images to the complement of E of S , it means there was a bottleneck in the whole graph. If there's no bottleneck at all, then indeed, there's no bottleneck in this other part of the complement of S and the complement of E of S .

And this gives me a hook on proving Hall's theorem, because that's basically the way I'm going to split the problem into two separate matching parts. So let's now proceed to prove that if there are no bottlenecks in a graph, then there's a perfect match. And it's going to be by strong induction on the number of girls.

So case one is the case that we just examined, that there's a non-empty subset of girls-- a proper subset, not all the girls-- some subset of girls who-- another way to say it is there's a subset of girls that's not empty and its complement is not empty either. And the size of S is the same as the size of E of S . That is, the number of boys that are compatible with this set of girls is exactly the same size as this set of girls.

Well, if there's going to be a match, there has to be one in that subset. And by the lemmas, if

there's no bottlenecks anywhere, there's no bottlenecks when I restrict myself to just S and E of S . And also by the previous Lemma, there's no bottleneck in the complement of S restricted to the vertices in E of S .

So this is a new bipartite graph, where I'm using S for the left vertices and E of S for the right vertices. And here, I'm using the complement of S for the left vertices and the complement of E of S for the right vertices. And the previous lemma says that if there's no bottlenecks overall, there's no bottlenecks in either of these two restricted graphs.

So what I can do now is, by induction, since there's no bottlenecks in S , E of S , I can find a match there. So I can match up all of the girls in S with compatible boys in E of S . And I can match up all of the girls that are not in S , in the complement of S , with boys that I haven't used already. And there will be a perfect match there. And then these two separately will provide a match for the entire set of left vertices-- S union the complement of S will be all of the vertices.

And these two matchings don't overlap. So combining them gives a complete matching. That knocks off this case. That's the nice case, where I can find some subset with the number of boys is exactly equal to the number of girls, which means that that has to work by itself. And I worked that one by itself and what's left over by itself. And that's how I can find an overall match.

The second case is that there isn't any set of girls, which are compatible with exactly the same number of boys. Well, since Hall's condition holds, that means that every set of girls actually has to be compatible with a larger set of boys than the number of girls. That is, for every set S , the size of S is strictly less than the size of E of S for every non-empty and proper subset S .

Well in that case, actually finding a match is easy, because now I've got slack. What I can do is just pick an arbitrary girl-- call her g -- and she's got to be compatible with at least one boy, because there's no bottlenecks. In fact, in this case, she'll actually be compatible with these two boys. But all we need is one boy.

Since there's no bottlenecks, if I pick a girl, she's got to be compatible with some boy. And I'm just going to arbitrarily match her with that boy-- match g with b . Now, the consequence of that is that, if I now remove g and b from the graph, I am left with a graph in which at worst I've shrunk the set E of S by that girl that I've matched with B , which means that E of S , which used to be strictly greater than S has gone down by at most one. It's still greater than or equal to E of S .

So I'm left that in the graph after I've removed this g paired with an arbitrary b that she's compatible with, I'm left with a graph that satisfies Hall's condition that has one fewer girl vertex. And so it has no bottlenecks. And therefore, it has a match.

And that match in the sub graph, where I've taken out g and b , when I put back the match, when I add the part of g matching with b , it's a match for the whole graph. And that's how I get it. That completes the proof of Hall's theorem.