Math 3322: Graph Theory[1] — Mikhail Lavrov

Lecture 5: Proofs by induction

August 31, 2021 — *Kennesaw State University*

1 The logic of induction

In the Towers of Hanoi puzzle (which we discussed in the first lecture) you have three pegs, and some number of disks of different sizes stacked on the pegs. Initially, all the disks are placed on one peg, sorted by size (with the smallest disk on top):

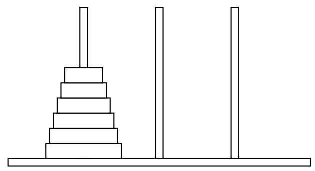

You are allowed to make the following moves in this puzzle: lift the top disk on a peg, and put it down on another peg. However, you cannot place a larger disk on top of a smaller one.

Theorem 1.1. *No matter how many disks there are, it is possible to move all the disks from one peg to another.*

If you try to solve this puzzle, you realize that the hard part is getting to the point where you can move the bottom disk. In order to be able to do this, all the other disks have to be off the first peg. Not only that, but a second peg has to be free, because the largest disk can't go on top of anything else! So in order to move the bottom disk from the first peg to the second, we have to have moved all the other disks onto the third peg.

This means that in order to solve the problem for n disks, we need to know how to solve it for $n-1$ disks. That's the perfect setup for an induction proof!

Proof of Theorem 1.1. We induct on n, the number of disks. When $n = 1$, we can move the disk from one peg to another in a single step. So the base case holds.

Assume that there is a way to move $n-1$ disks from one peg to another. Then there is also a solution to the n-disk puzzle:

1. Using the $(n-1)$-disk solution, move the first $n-1$ disks from the first peg to the second.
2. Move the largest disk from the first peg to the third.
3. Using the $(n-1)$-disk solution again, move the first $n-1$ disks from the second peg to the third.

By induction, there is a solution for all n. □

[1] This document comes from the Math 3322 course webpage: http://facultyweb.kennesaw.edu/mlavrov/courses/3322-fall-2021.php

How does all this work? You can think of proofs by induction as a template for infinitely many proofs. Imagine that insted of Theorem 1.1, we had a whole bunch of lemmas:

Lemma 1.1. *The Towers of Hanoi puzzle with* 1 *disk has a solution.*

Lemma 1.2. *The Towers of Hanoi puzzle with* 2 *disks has a solution.*

Lemma 1.3. *The Towers of Hanoi puzzle with* 3 *disks has a solution.*

Lemma 1.4. *The Towers of Hanoi puzzle with* 4 *disks has a solution.*

Our proof contains a proof of Lemma 1.1: that was the base case. It also contains a proof of Lemma 1.2: take the induction step (replacing n by 2) and use Lemma 1.1 when we need to know that the 1-disk puzzle has a solution.

It also contains a proof of Lemma 1.3: take the induction step (replacing n by 3) and use Lemma 1.2 when we need to know that the 2-disk puzzle has a solution. Similarly, all the other lemmas have proofs.

The reason that we can give these infinitely many proofs all at once is that they all have similar structure, relying on the previous lemma. And that's all that induction is.

2 The hypercube graph

Let Q_n be the hypercube graph. Its vertices are $\{0,1\}^n$: n-tuples $(x_1, x_2, \ldots, x_n)$, where each x_i is either 0 or 1. It has an edge between two vertices that differ in exactly one coordinate, and agree in the $n-1$ others.

The cube graph Q_3 is shown below on the left. It looks like a cube.

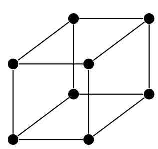

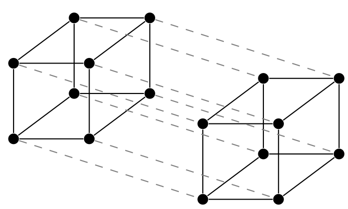

Above on the right, we see Q_4. The dashed edges aren't any different from the others, but I've drawn them this way to highlight some structure of Q_n:

- Take the subgraph induced by the 2^{n-1} vertices with a 0 in the last position. This looks exactly like Q_{n-1}: if all these vertices have the same last coordinate, we can just ignore that coordinate for adjacency purposes.
- Take the subgraph induced by the 2^{n-1} vertices with a 1 in the last position. This, too, looks exactly like Q_{n-1}.
- In Q_n, these two subgraphs are joined by 2^{n-1} edges (which are the dashed edges in the diagram). These are the edges between $(x_1, x_2, \ldots, x_{n-1}, 0)$ and $(x_1, x_2, \ldots, x_{n-1}, 1)$ for all $x_1, \ldots, x_{n-1}$: they join the "corresponding" vertices of the two subgraphs.

For all n, Q_n is bipartite. The idea behind the bipartition is one we've seen before: put vertices with an even number of 1's on one side, and vertices with an odd number of 1's on the other side. If two vertices differ in only one coordinate, the number of 1's in them dffers by 1, so one vertex is on one side, and one is on the other. Here's a drawing of Q_3 in a more conventionally bipartite fashion:

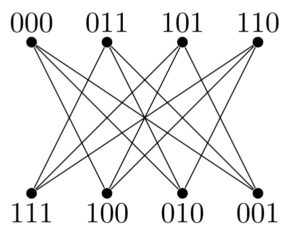

The way that we define Q_n recursively (in terms of Q_{n-1}) makes it easy for us to prove things about it by induction. Here's an example:

Theorem 2.1. *For all $n \geq 1$, the diameter of Q_n is n.*

Proof. We'll actually prove a more specific statement: the diameter of Q_n is n, and any two vertices which don't agree in any coordinate are at distance n from each other.

We induct on n. When $n = 1$, there are only 2 vertices, 0 and 1, and they are in fact at distance 1 from each other (they're not the same vertex, but there is an edge between them). So the statement we want to prove is true here.

Assume that Q_{n-1} has diameter $n-1$. Let's first prove that if u, v are two vertices of Q_n, then $d(u,v) \leq n$. There are two cases:

- If they have the same last coordinate, then they're in a subgraph that looks like Q_{n-1} together. So they're at most $n-1$ steps apart, by the inductive hypothesis.
- Otherwise, let v' be the vertex that differs from v only in the last coordinate; we know $d(u,v') \leq n-1$ by the inductive hypothesis. But we can get from v' to v in 1 more step, so $d(u,v) \leq n$.

Now suppose u and v are opposite vertices, and let's try to prove $d(u,v) = n$. Let v' be the vertex that differs from v only in the last coordinate.

In any $u-v$ path, there must be at least one step where the last coordinate changes. Let's just skip all such steps (and keep the last coordinate the same throughout) getting a $u-v'$ path instead. This $u-v'$ path has length at least $n-1$, because we know $d(u,v') = n-1$ by the inductive hypothesis. The $u-v$ path must have been at least 1 step longer, so $d(u,v) \geq n$. □

Note: this is not the only inductive approach. Another proof strategy you might try is to keep the hypercube Q_n the same, and prove by induction on k that the vertices at distance k from vertex v are exactly the vertices which disagree with v in k coordinates.

3 The induction trap

Let's play the game "what's wrong with this proof?"

Claim 3.1 (False claim)**.** *Suppose every vertex in an n-vertex graph is the endpoint of at least two different edges. Then the graph must have at least $2n - 3$ edges.*

We know this is false, because the cycle graph C_n is an example—and that only has n edges, which is less than $2n - 3$ for large n.

Incorrect proof. For $n = 3$ vertices, the claim holds, because we need all 3 edges in a 3-vertex graph to exist in order for the assumption to hold, and $3 = 2(3) - 3$.

Assume the claim holds for all $(n - 1)$-vertex graphs: if they satisfy the hypothesis, they all have at least $2(n - 1) - 3 = 2n - 5$ edges. To get an n-vertex graph, we add a vertex; to make sure that it's the endpoint of at least two different edges, we need to add at least two new edges. Therefore the n-vertex graph has at least $2n - 5 + 2 = 2n - 3$ edges. □

The problem with this proof is that not all n-vertex graphs where every vertex is the endpoint of at least two edges can be built from $(n - 1)$-vertex graphs with the same property. Consider C_4: a 4-vertex graph where every vertex is the endpoint of exactly two edges. We cannot build C_4 by extending C_3, the only 3-vertex graph with this property! So we've only proven the claim for a subset of all graphs, and that subset does not include the examples with the fewest edges.

To avoid this problem, here is a useful template to use in induction proofs for graphs:

Theorem 3.1 (Template)**.** *If a graph G has property A, it also has property B.*

Proof. We induct on the number of vertices in G. **(Prove a base case here.)**

Assume that all $(n-1)$-vertex graphs with property A also have property B. Let G be an n-vertex graph with property A. Our goal is to show that G also has property B.

Let v be a vertex of G **(usually chosen by some clever rule you'll have to come up with)**. Then $G - v$ (the graph obtained from G by deleting v and all edges out of v) also has property A **(by an argument related to the clever way we chose which vertex to delete)**.

By the inductive hypothesis, $G - v$ also has property B. When we add back the vertex v, G also has property B **(by another argument you'll have to come up with)**.

By induction, all graphs with property A also have property B. □

Why does this work, when our previous argument didn't? The key is the step "Let G be an n-vertex graph with property A." We didn't make any assumptions about G. Rather, we started from an arbitrary graph with property A; to apply the inductive hypothesis, we cooked up a graph $G - v$ which is an $(n - 1)$-vertex graph with property A.