

Lecture 5: Proofs by induction

August 27, 2024

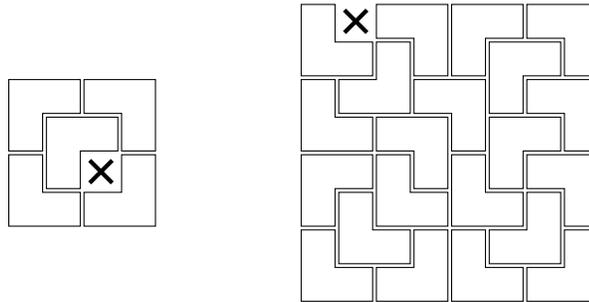
Kennesaw State University

1 The logic of induction

In the first lecture, we discussed a connection between graph theory and tiling puzzles. Here is a simpler tiling puzzle we will not need graph theory to solve.

Suppose that you have an infinite supply of  tiles. We will ask a general problem rather than a concrete one: how many of them can you place on a $2^n \times 2^n$ grid, without overlap?

We will definitely not be able to fill the entire grid. That's because there are 4^n squares in the grid, which is not a multiple of 3. However, in small examples, it looks like we can fill almost the entire grid, with only one space—marked with an X in the diagrams below—not covered by a tile!



Let's prove that this can be done in general. Actually, let's prove something slightly stronger (I'll explain why at the end):

Theorem 1.1. *For all $n \geq 0$, it's possible to tile a $2^n \times 2^n$ grid with  tiles so that only one square of the grid remains uncovered, and the uncovered square is in a corner of the grid.*

Proof. The proof is by induction on n : a technique that I assume you are not seeing for the first time, but also suspect you would not mind reviewing.

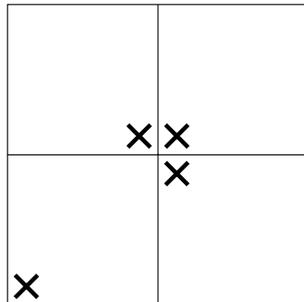
We will begin by proving this statement for $n = 0$. Here, we are tiling a 1×1 grid; to tile it so that only a corner square remains uncovered, we literally have to do nothing.

If this case seems unsatisfying, we can also prove the $n = 1$ case. Here, we are tiling a 2×2 grid; no matter how you place a  in the grid, you will have covered 3 of the 4 squares, with the last being one of the corners.

Now consider $n \geq 2$, and assume that we know a solution for the $2^{n-1} \times 2^{n-1}$ grid. Well, the $2^n \times 2^n$ grid is made up of four quarters, and they each look like the $2^{n-1} \times 2^{n-1}$ grid. Let's apply our

¹This document comes from the Math 3322 course webpage: <http://facultyweb.kennesaw.edu/mlavrov/courses/3322-fall-2024.php>

assumption to each of the quarters, and carefully rotate the $2^{n-1} \times 2^{n-1}$ solutions to make them fit together as follows:



Three of our uncovered squares are right next to each other, so that a single  tile can cover all of them! This leaves one square uncovered, which is in a corner, as promised.

We have shown that if the theorem is true for a $2^{n-1} \times 2^{n-1}$ grid, then it is also true for a $2^n \times 2^n$ grid. By induction, the theorem is true for all $n \geq 0$. \square

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

Lemma 1.2. *We can tile a 2×2 grid with only a corner left uncovered.*

Lemma 1.3. *We can tile a 4×4 grid with only a corner left uncovered.*

Lemma 1.4. *We can tile a 8×8 grid with only a corner left uncovered.*

Lemma 1.5. *We can tile a 16×16 grid with only a corner left uncovered.*

Our proof contains a proof of Lemma 1.2: we proved that as a base case. It can also be used to create a proof of Lemma 1.3: take the induction step (replacing n by 2) and use Lemma 1.2 to tile the four quarters of the grid.

It can also be used to create a proof of Lemma 1.4: take the induction step (replacing n by 3) and use Lemma 1.3 to tile the four quarters. Similarly, every lemma of this type has a proof.

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.

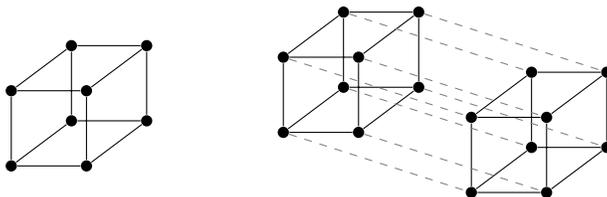
One final question: why did we specify, in our theorem, that it's the corner of the grid that remains uncovered? This is an induction-specific trick called "loading the induction hypothesis".

It seems like we're only making our lives harder by giving ourselves a stronger claim to prove: after all, it's easier to tile the grid if you don't care *which* square gets left uncovered. However, in induction, that's not always so clear-cut. Making our claim stronger also meant that when we assumed the $n - 1$ case of the theorem to prove the n case, our assumption was stronger—and having a stronger assumption makes our lives easier!

2 The hypercube graph

Let Q_n be the hypercube graph. Its vertices are $\{0, 1\}^n$: n -tuples (x_1, x_2, \dots, 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.

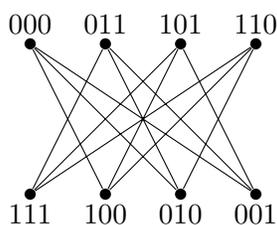
The cube graph Q_3 is shown below on the left. It looks like a cube.



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 $(x_1, x_2, \dots, x_{n-1}, 0)$ and $(x_1, x_2, \dots, x_{n-1}, 1)$ for all x_1, \dots, 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 differs 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—another instance of “loading the induction hypothesis”! We'll prove that the diameter of Q_n is n , and also that 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 $d(u, v) \geq n$. \square

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

Now 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. \square

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 3.2 (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**).

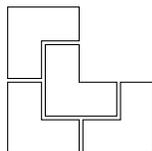
By induction, all graphs with property A also have property B . □

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 .

4 Practice problems

1. Let's take a closer look at Theorem 1.1.

- (a) Draw a diagram of the tiling that the proof gives us for an 8×8 grid.
- (b)  is what is called a **rep-tile**²: we can put four tiles together to make a scaled copy of the tile that's twice as large. Here's how:



Use this fact to give a completely different inductive proof of Theorem 1.1.

- (c) It is possible to make our claim even stronger: if an X is placed in an arbitrary square of a $2^n \times 2^n$ grid, then we can cover all the squares except that one with  tiles.

Modify the proof (either version) to prove the stronger claim.

2. Prove by induction on n that for all n , the graph Q_n has $n \cdot 2^{n-1}$ edges.

(In the next lecture, we will see a way to prove this directly.)

3. To prove lower bounds on distance and diameter, induction is often useful. As an example, consider the graph P_n :



Prove by induction on k that $d(v_1, v_k) \geq k - 1$, and conclude that $\text{diam}(P_n) \geq n - 1$.

(Throughout your proof, n should be fixed; only k is changing. The base case is $k = 1$.)

4. In addition to the graph operations we learned last time, there are several notions of products of two graphs. One of them is the **box product** $G \square H$. It is defined as follows:

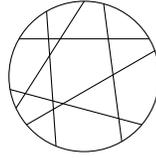
- The vertices of $G \square H$ are pairs (v, w) , where v is a vertex of G and w is a vertex of H .
- Two vertices (v, w) and (v', w') are adjacent in two cases: if $v = v'$ and w is adjacent to w' in H , or if $w = w'$ and v is adjacent to v' in G .

Let's look at some consequences of this definition.

- (a) Draw the graph $P_6 \square P_6$.
- (b) Let G be a graph with n vertices and m edges. How many vertices and edges does $G \square G$ have?
- (c) Let G and H be two connected graphs. Prove that $G \square H$ is connected.
- (d) The graph $K_2 \square K_2 \square \dots \square K_2$, with n "factors" of K_2 , also has a different name. What is it, and why?

²Mathematical humor at its finest.

5. Suppose that we take a circular disk and cut it along straight lines (as though we were cutting cake, but very badly). Here is an example:



- (a) This divides the circle into several regions. Prove that no matter how you cut, those regions can always be colored black and white so that whenever two regions share a side, they are different colors.
- (b) The claim in part (a) is equivalent to saying that a certain graph is bipartite. How is that, and what is the definition of the graph?
6. A **threshold graph** is a graph you can build from K_1 (a graph with only one vertex) by repeatedly doing one of two operations: (1) add a vertex adjacent to all other vertices and (2) add a vertex not adjacent to any vertices. (Each time you add a vertex, you make a separate decision about which operation to do.)

- (a) As a warm-up, find all 8 different threshold graphs on 4 vertices.
- (b) Prove the following by induction: if G is a threshold graph, you can give every vertex v a number $0 < x_v < 1$ such that two vertices v, w are adjacent if and only if $x_v + x_w \geq 1$.

(Hint: use the induction template. If G is a threshold graph, the natural choice of v to delete so that $G - v$ is also a threshold graph is the vertex we last added in the process of building G .)