

Problem Set 5

This problem set – the last one purely on discrete mathematics – is designed as a cumulative review of the topics we’ve covered so far and a proving ground to try out your newfound skills with mathematical induction. The problems here span all sorts of topics – parallel processing, the nature of infinity, tiling problems, and tournament structures – and we hope that it serves as a fitting coda to our whirlwind tour of discrete math!

We recommend that you read the Handout #24, “Guide to Induction,” before starting this problem set. It contains a lot of useful advice about how to approach problems inductively, how to structure inductive proofs, and how to not fall into common inductive traps.

As a note on this problem set – normally, you're welcome to use any proof technique you'd like to prove results in this course. On this problem set, we've specifically requested on some problems that you prove a result inductively. For those problems, you should prove those results using induction or complete induction, even if there is another way to prove the result. (If you'd like to use induction in conjunction with other techniques like proof by contradiction or proof by contrapositive, that's perfectly fine.)

As always, please feel free to drop by office hours, visit Piazza, or send us emails if you have any questions. We'd be happy to help out.

Good luck, and have fun!

**Due Friday, May 12 at the start of class.
There is no checkpoint problem.**

Problem One: Chains and Antichains

Let A be an arbitrary set and $<_A$ be some strict order over A . A **chain in** $<_A$ is a series x_1, \dots, x_k of elements drawn from A such that

$$x_1 <_A x_2 <_A \dots <_A x_k.$$

Intuitively, a chain is a series of values in ascending order according to the strict order $<_A$. The **length** of a chain is the number of elements in that chain.

- i. Consider the \subsetneq relation over the set $\wp(\{a, b, c\})$, where $A \subsetneq B$ means that $A \subseteq B$ but $A \neq B$. What is the length of the longest chain in this strict order? Give an example of a chain with that length. No justification is necessary. (*Hint: Draw the Hasse diagram and see if you can find a visual intuition for the definition of a chain.*)

Now, let's cover a new definition. An **antichain** is a set $X \subseteq A$ such that no two elements in X can be compared by the $<_A$ relation. In other words, a set $X \subseteq A$ is an antichain if

$$\forall a \in X. \forall b \in X. (a \not<_A b \wedge b \not<_A a)$$

The **size** of an antichain X is the number of elements in X .

- ii. Consider the \subsetneq relation over the set $\wp(\{a, b, c\})$. What is the size of the largest antichain in this strict order? Give an example of an antichain with that size. No justification is necessary. (*Hint: Draw the Hasse diagram and see if you can find a visual intuition for the definition of an antichain.*)

Given an arbitrary strictly ordered set, you can't say anything *a priori* about the size of the largest chain or antichain in that strict order. However, you can say that at least one of them must be relatively large relative to the strictly ordered set.

Let $<_A$ be an arbitrary strict order over an arbitrary set A containing exactly n^2+1 elements for some natural number $n \geq 1$. We're going to ask you to prove the following result: either A contains a chain of length $n+1$ or an antichain of size $n+1$ (or both). Following the advice from Handout 14, we'll prove this by instead proving that if A does *not* contain a chain of length $n+1$ or greater, then A must contain an antichain of size $n+1$ or greater.

- iii. For each element $a \in A$, we'll say that the **height** of a is the length of the longest chain whose final element is a . Prove that if A does not contain a chain of length $n+1$ or greater, then there must be at least $n+1$ elements of A at the same height. (*Something to think about: what's the smallest possible height of an element of A ?*)
- iv. Your result from part (iii) establishes that if A does not contain a chain of length $n+1$ or greater, there must be a collection of $n+1$ elements of A at the same height as one another. Prove that if A does not contain a chain of length $n+1$ or greater, then it contains an antichain of size $n+1$ or greater.

Intuitively speaking, if $<_A$ is a strict order over A that represents some prerequisite structure on a group of tasks, a chain represents a series of tasks that have to be performed one after the other, and an antichain represents a group of tasks that can all be performed in parallel (do you see why?) In the context of parallel computing, the result you've proved states that if a group of tasks doesn't contain long dependency chains, that group must have a good degree of parallelism.

Problem Two: Recurrence Relations

A *recurrence relation* is a recursive definition of the terms in a sequence. Typically, a recurrence relation specifies the value of the first few terms in a sequence, then defines the remaining terms from the previous terms. For example, the *Fibonacci sequence* can be defined by the following recurrence relation:

$$\begin{aligned} F_0 &= 0 \\ F_1 &= 1 \\ F_{n+2} &= F_n + F_{n+1} \end{aligned}$$

The first terms of this sequence are $F_0 = 0$, $F_1 = 1$, $F_2 = 1$, $F_3 = 2$, $F_4 = 3$, $F_5 = 5$, $F_6 = 8$, etc.

Some recurrence relations define well-known sequences. For example, consider the following recurrence relation:

$$\begin{aligned} a_0 &= 1 \\ a_{n+1} &= 2a_n \end{aligned}$$

The first few terms of this sequence are 1, 2, 4, 8, 16, 32, ..., which happen to be powers of two. It turns out that this isn't a coincidence – this recurrence relation perfectly describes the powers of two.

- i. Prove by induction that for any $n \in \mathbb{N}$, we have $a_n = 2^n$.

Minor changes to the recursive step in a recurrence relation can lead to enormous changes in what numbers are generated. Consider the following two recurrence relations, which are similar to the a_n sequence defined above but with slight changes to the recursive step:

$$\begin{array}{ll} b_0 = 1 & c_0 = 1 \\ b_{n+1} = 2b_n - 1 & c_{n+1} = 2c_n + 1 \end{array}$$

- ii. Find non-recursive definitions for b_n and c_n , then prove by induction that your definitions are correct.

Finding non-recursive definitions for recurrences (often called “solving” the recurrence) is useful in the design and analysis of algorithms. Commonly, when trying to analyze the runtime of an algorithm, you will arrive at a recurrence relation describing the runtime on an input of size n in terms of the runtime on inputs of smaller sizes. Solving the recurrence then lets you precisely determine the runtime. To learn more, take CS161, Math 108, or consider reading through *Concrete Mathematics* by Graham, Knuth, and Patashnik.

Problem Three: Cartesian Products and Cardinalities

The cardinality of the Cartesian product of two sets depends purely on the cardinalities of those sets, not on what the elements of those sets actually are. This question will ask you to prove this.

Let A , B , C , and D be sets where $|A| = |C|$ and $|B| = |D|$. Our goal is to prove $|A \times B| = |C \times D|$. Since we know $|A| = |C|$, there has to be some bijection $g : A \rightarrow C$. Since we know $|B| = |D|$, there has to be some bijection $h : B \rightarrow D$. Now, consider the function $f : A \times B \rightarrow C \times D$ defined as follows:

$$f(a, b) = (g(a), h(b))$$

That is, the output of f when applied to the pair (a, b) is an ordered pair whose first element is $g(a)$ and whose second element is $h(b)$.

- i. Using the function f defined above, prove that if A , B , C , and D are sets where $|A| = |C|$ and $|B| = |D|$, then we have $|A \times B| = |C \times D|$. Specifically, prove that f is a bijection between $A \times B$ and $C \times D$.

We can define the “Cartesian power” of a set as follows. For any set A and any positive natural number n , we define A^n inductively:

$$A^1 = A$$

$$A^{n+1} = A \times A^n \text{ (for } n \geq 1\text{)}$$

- ii. Using your result from (i) and the above definition, prove by induction that $|\mathbb{N}^k| = |\mathbb{N}|$ for all nonzero $k \in \mathbb{N}$. This result means that for any nonzero finite k , there are the same number of k -tuples of natural numbers as natural numbers. (*Hint: You might want to use a result from lecture.*)

(Here's some justification for the problem you just did. You don't need to read this if you don't want to, but we think you might find it interesting. ☺)

We've defined cardinality on a relative basis, defining $|A| = |B|$, $|A| \neq |B|$, $|A| \leq |B|$, and $|A| < |B|$ in terms of functions between those sets. We specifically didn't talk about cardinalities as actual quantities, since it's tricky to do so. That said, it is possible to treat set cardinalities as actual objects. Intuitively speaking, the cardinality of a set is a measure of how large that set is. For finite sets, the cardinality of that set will be a natural number, and for infinite sets the cardinality of that set will be an infinite cardinality, a generalization of the natural numbers that measure the sizes of infinite sets. For example, \aleph_0 , which we introduced in our first lecture as the cardinality of \mathbb{N} , is an infinite cardinality.

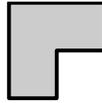
Given cardinalities κ_1 and κ_2 , we define the product of those two cardinalities, $\kappa_1 \cdot \kappa_2$, to be $|A \times B|$, where A and B are any sets where $|A| = \kappa_1$ and $|B| = \kappa_2$. For example, $4 \cdot 3$ is *by definition* the cardinality $|A \times B|$ for any set A of cardinality 4 and any set B of cardinality 3. Similarly, *by definition* the value of $\aleph_0 \cdot \aleph_0$ is the cardinality $|A \times B|$ for any sets A and B of cardinality \aleph_0 .

To make sure that this definition is legal, we have to make sure that the cardinality of the Cartesian product depends purely on the cardinalities of the two sets, not their contents. For example, this definition wouldn't give us a way to compute $4 \cdot 3$ if the cardinality of the Cartesian product of a set of four apples and three oranges was different than the cardinality of the Cartesian product of a set of four unicorns and three ponies. We need to show that for any sets A , B , C , and D , that if $|A| = |C|$ and $|B| = |D|$, then $|A \times B| = |C \times D|$. That way, when determining the value of $\kappa_1 \cdot \kappa_2$, it doesn't matter which sets of cardinality κ_1 and κ_2 we pick; any choice works. Your proof from (i) filled in this step.

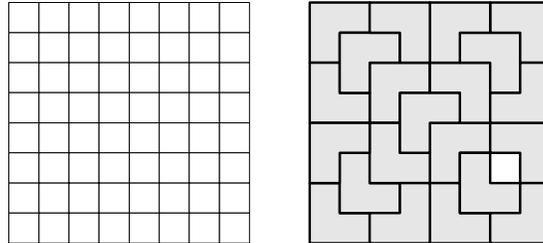
Your result from part (ii) shows that $\aleph_0^n = \aleph_0$ for any positive natural number n . Isn't infinity weird?

Problem Four: Tiling with Triominoes

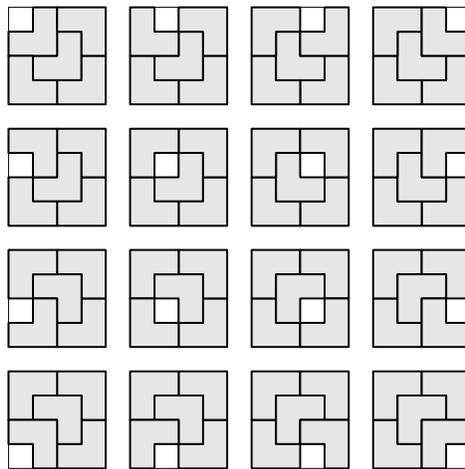
Recall from Problem Set One that a *right triomino* is an L-shaped tile that looks like this:



Suppose that you are also given a square grid of size $2^n \times 2^n$ and want to *tile* it with right triominoes by covering the grid with triominoes such that all triominoes are completely on the grid and no triominoes overlap. Here's an attempt to cover an 8×8 grid with triominoes, which fails because not all squares in the grid are covered:



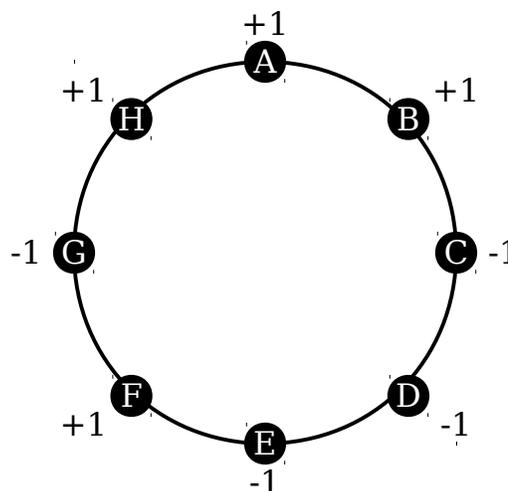
Amazingly, it turns out that it is always possible to tile any $2^n \times 2^n$ grid that's missing exactly one square with right triominoes. It doesn't matter what n is or which square is removed; there is always a solution to the problem. For example, here are all the ways to tile a 4×4 grid that has a square missing:



Prove by induction that any $2^n \times 2^n$ grid with any one square removed can be tiled by right triominoes.

Problem Five: The Circle Game

Here's a game you can play. Suppose that you have a circle with $2n$ arbitrarily-chosen points on its circumference. n of these points are labeled $+1$, and the remaining n are labeled -1 . One sample circle with eight points, of which four are labeled $+1$ and four are labeled -1 , is shown to the right.



Here's the rules of the game. First, choose one of the $2n$ points as your starting point. Then, start moving clockwise around the circle. As you go, you'll pass through some number of $+1$ points and some number of -1 points. You lose the game if at any point on your journey you pass through more -1 points than $+1$ points. You win the game if you get all the way back around to your starting point without losing.

For example, if you started at point A, the game would go like this:

Start at A: $+1$.
 Pass through B: $+2$.
 Pass through C: $+1$.
 Pass through D: 0 .
 Pass through E: -1 . (*You lose.*)

If you started at point G, the game would go like this:

Start at G: -1 (*You lose.*)

However, if you started at point F, the game would go like this:

Start at F: $+1$.
 Pass through G: 0 .
 Pass through H: $+1$.
 Pass through A: $+2$.
 Pass through B: $+3$.
 Pass through C: $+2$.
 Pass through D: $+1$.
 Pass through E: $+0$.
 Return to F. (*You win!*)

Interestingly, it turns out that no matter which n points are labeled $+1$ and which n points are labeled -1 , there is always at least one point you can start at to win the game.

Prove, by induction, that the above fact is true for any $n \geq 1$.

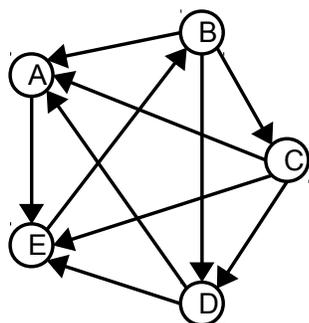
Problem Six: Nim

Nim is a family of games played by two players. The game begins with several piles of stones, each of which has zero or more stones in it, that are shared between the two players. Players alternate taking turns removing any nonzero number of stones from any single pile of their choice. If at the start of a player's turn all the piles are empty, then that player loses the game.

Prove, by induction, that if the game is played with just two piles of stones, each of which begins with exactly the same number of stones, then the second player can always win the game if she plays correctly.

Before trying to write up your answer to this question, we recommend playing this game with a partner until you can find a winning strategy.

Problem Seven: Tournament Graphs



In Problem Set Two, you explored tournaments, contests among $n \geq 0$ players where each player plays exactly one game against each other player and no game ends in a tie.

From a graph-theoretic perspective, we can treat tournaments as graphs. Each player is a node, and there's exactly one edge between each pair of nodes. The edge is directed such that the edge goes from the winner of the game to the loser of the game. For example, in the graph to the left, player *D* beat players *A* and *E*, but lost to players *B* and *C*. These sorts of graphs are called *tournament graphs*.

Let's introduce a new definition. If T is a tournament graph, a *subtournament* of T is a tournament graph T' formed by picking a subset of the players in T and considering just the games those players played against each other. (Note that, just like any set is a subset of itself, any tournament is a subtournament of itself).

Now, let's introduce one final definition. A tournament graph T is called *transitive* if it contains no cycles. Prove by induction that for any natural number n , if T is a tournament graph with exactly 2^n nodes, then T contains a transitive subtournament T^* with at least $n + 1$ nodes.

As a hint, you might want to use the *bowtie decomposition*. Given any nonempty tournament T , if you choose a player p from T , you can split the remaining players in T into two subtournaments T_0 and T_1 consisting of just the players who beat p and just the players who lost to p , respectively.

Extra Credit Problem: My Old CS103 Midterm (1 Point Extra Credit)

I learned proofwriting and discrete math when I took CS103 here way back in Winter 2007. I still remember the midterm exam for that class – it probably the single hardest exam I've ever taken. Here's a question that was actually asked on the midterm that quarter. And don't worry... I would never ask something like this on a CS103 midterm!

Imagine you have a collection of n planes in three-dimensional space. The planes are chosen so that any two planes will intersect each other at a line, any three planes will intersect at a single point, but no combination of four planes will intersect at the same point. These planes will split three-dimensional space into a collection of different regions.

How many of those regions will there be, as a function of n ? Prove it by induction.