Problem Set 4

This fourth problem set explores the power of functions, the pigeonhole principle, and structural induction. The problems here include some of my favorite results from discrete mathematics, and I hope that you have fun playing around with the new concepts from this week!

Start this problem set early. It contains five problems (plus one survey question), several of which require a fair amount of thought. I would suggest reading through this problem set at least once as soon as you get it to get a sense of what it covers.

As much as you possibly can, please try to work on this problem set individually. That said, if you do work with others, please be sure to cite who you are working with and on what problems. For more details, see the section on the honor code in the course information handout.

In any question that asks for a proof, you **must** provide a rigorous mathematical proof. You cannot draw a picture or argue by intuition. You should, at the very least, state what type of proof you are using, and (if proceeding by contradiction, contrapositive, or induction) state exactly what it is that you are trying to show. If we specify that a proof must be done a certain way, you must use that particular proof technique; otherwise you may prove the result however you wish.

If you are asked to prove something by induction, you may use weak induction, strong induction, the well-ordering principle, or structural induction. In any case, you should state your base case before you prove it, and should state what the inductive hypothesis is before you prove the inductive step.

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

This problem set has 125 possible points. It is weighted at 7% of your total grade. The earlier questions serve as a warm-up for the later problems, so do be aware that the difficulty of the problems does increase over the course of this problem set.

Good luck, and have fun!

Due <u>Thursday</u> October 27th at 7:00 PM

Problem One: Bijections (10 Points)

In lecture, we claimed that the following function $f : \mathbb{N} \to \mathbb{Z}$ is a bijection from $\mathbb{N}$ to $\mathbb{Z}$, which proves that $|\mathbb{N}| = |\mathbb{Z}| = \aleph_0$:

$$f(n) = \begin{cases} -n/2 & \text{if } n \text{ is even} \\ (n+1)/2 & \text{otherwise} \end{cases}$$

Prove that f is a bijection.

Problem Two: Set Cardinalities (20 Points)

Recall from lecture that two sets A and B have the same cardinality (that is, $|A| = |B|$) if there is a bijection $f : A \to B$. For each of the following, show that the indicated sets have the same cardinality by finding a bijection between them, then proving that your function is a bijection.

i. If $S = \{a\}$ for some a, then $|A \times S| = |A|$

ii. If $|A| = |C|$ and $|B| = |D|$, then $|A \times B| = |C \times D|$. *(Hint: Since you know that* $|A| = |C|$*, there is a bijection* $f : A \to C$*. There is a similar bijection* $g : B \to D$*)*

iii. $|\mathbb{N} \times \{0, 1\}| = |\mathbb{N}|$.

iv. ***(Extra credit)*** $|\mathbb{N}| = |\mathbb{N}^2|$. Note that you have to give an explicit bijection here; drawing a picture is not sufficient.

Problem Three: Pigeonhole Party! (30 Points)

The pigeonhole principle can be used to prove some surprising results about groups of people.

Suppose that you are at a party. Any two people either have met (they are *acquaintances*) or have never met (they are *strangers*). We can therefore think of the party as an undirected graph where each person is a node and there is an edge between two people if they are acquaintances. For example, at this party:

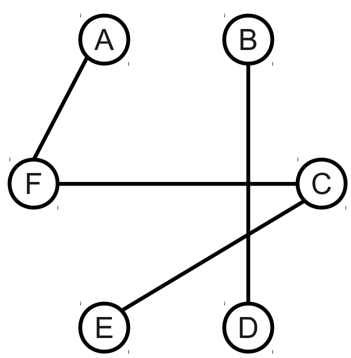

Person A just knows person F, person B knows person D, and person C knows both person E and person F. However, none of A, B, or E know each other.

i. Show that at a party with at least two people present, there are at least two people with the same number of acquaintances at that party. *(Hint: Consider two cases: the case where there are at least two people who know no one else, and the case where at most one person knows no one else.)*

The **generalized pigeonhole principle** says that if there are n objects to be put into k boxes, then there must be some box that contains at least $\lceil n / k \rceil$ objects.*

ii. Prove that this is true. You may want to use the fact that $\lceil n / k \rceil < n / k + 1$.

iii. Show that in any group of six people that there are either three mutual acquaintances or three mutual strangers. *(Hint: Choose any person, then think about how the other people at the party are related to that person)*

Problem Four: Network Fragility (30 Points)

Suppose that you have an undirected graph with n nodes. The **distance** between two nodes is defined as the length of the shortest path between them, where the length of the path is the number of edges it contains. For the purposes of this problem, we'll assume that the graph is connected, so there is at least one path between any two nodes. Denote the distance from x to some node u as d(x, u).

i. Prove that if P is a path from u to v and (x_i, x_j) is an edge in P, then $d(u, x_i) + 1 \geq d(u, x_j)$.

ii. Using (i) as a starting point, prove that if P is a path from u to v, then for any natural number k with $k \leq d(u, v)$, there is some node x in P such that $d(u, x) = k$.

While the pigeonhole principle is often used to guarantee that given n objects and k boxes there is some box with at *least* $\lceil n / k \rceil$ elements in it, the pigeonhole principle can also be used to show that there is some box with at *most* $\lfloor n / k \rfloor$ elements in it. For example, if we try distributing 10 pigeons into 3 boxes, some box must have at most $\lfloor 10 / 3 \rfloor = 3$ pigeons in it.

iii. Prove that this is true. You may want to use the fact that $\lfloor n / k \rfloor > n / k - 1$.

iv. Suppose that there are some nodes u and v in a graph where $d(u, v) > n / 2$. Using (ii) and (iii), prove that there must be some natural number k with $0 < k < d(u, v)$ such that there is only one node x such that $d(u, x) = k$.

v. Using your results from (ii) and (iv), prove that if $d(u, v) > n / 2$, then there is some node x (other than u and v) such that any path from u to v must pass through x. This says that, in a sense, the graph is "fragile" and that if this node were to be removed, u and v would be disconnected.

* I'm not sure why you would be trying to put a whole bunch of pigeons into a small number of pigeonholes, but this result says what would happen if you did. Be nice to animals, folks.

Problem Five: Well-Founded Induction (30 Points)

At this point in the course, you've seen three types of induction:

- Weak induction, which proves properties of the natural numbers,
- Strong induction, which also proves properties of the natural numbers, and
- Structural induction, which proves properties of recursively-defined sets.

All of the inductive techniques we've encountered so far have worked along the following lines. To show that some property P(x) applies to objects in some set S, we first prove that P(x) holds for the smallest elements of that set (for some definition of "smallest.") Next, we assume that for some x, P(x') is true for all x' smaller than x (for some definition of "smaller"), then use this to show that this implies that P(x) is true. If we can do so, we can conclude that P(x) is true for all elements of the set S.

In the case of weak and strong induction, our notion of "smallest" and "smaller" correspond naturally to the ordering of the natural numbers according to the < relation. In the case of structural induction, "smallest" refers to elements that are defined to be in the set. We say that some element of the set x is "smaller than" some element y if x is used to construct y.

Interestingly, all three of these inductive techniques are special cases of a more powerful form of induction called **well-founded induction**. In this problem, you will explore how well-founded induction works, then will use it to prove a result that would otherwise require a much more complex double induction.

Recall from lecture that a *strict order* is a binary relation $<_A$ over a set A such that

- R is **irreflexive**: $\forall a \in A.\ \neg(a <_A a)$
- R is **antisymmetric**: $\forall a \in A.\ \forall b \in A.\ (a <_A b \land b <_A a \rightarrow a = b)$
- R is **transitive**: $\forall a \in A.\ \forall b \in A.\ \forall c \in A.\ (a <_A b \land b <_A c \rightarrow a <_A c)$

Let $<_A$ be a strict order and let $X \subseteq A$ be any arbitrary subset of A. If X is nonempty, we say that an element $x \in X$ is **minimal** if for any $y \in X$, it is not true that $y <_A x$. That is, x is minimal if there is no other element of the set X that is smaller than it. Note that this does *not* mean that x is smaller than every other element in the set X – there could be some other element $y \in X$ such that neither $x <_A y$ nor $y <_A x$ – but rather that nothing else in the set is strictly less than x.

A strict order $<_A$ over a set A is called **well-founded** if every nonempty subset of A has a minimal element. For example, as you proved in the previous problem set, the set $\mathbb{N}$ ordered by < is well-founded because any nonempty subset of $\mathbb{N}$ contains a least element.

i. Prove or disprove: The set $\mathbb{Z}$ ordered by < is well-founded.

Recall from the first problem set the definition of a **lexicographical ordering**. Suppose that $(A, <_A)$ and $(B, <_B)$ are ordered sets such that $<_A$ is a strict order over A and $<_B$ is a strict order over B. Consider the set $A \times B$ and the relationship $<_{lex}$ on $A \times B$ defined as follows: Given pairs (a_1, b_1) and (a_2, b_2):

- If $a_1 <_A a_2$, then $(a_1, b_1) <_{lex} (a_2, b_2)$.
- Otherwise, if $a_1 = a_2$ and $b_1 <_B b_2$, then $(a_1, b_1) <_{lex} (a_2, b_2)$

In the first problem set, you proved that if $<_A$ and $<_B$ are strict orders, then $<_{lex}$ is as well. Interestingly, it's also true that if $<_A$ and $<_B$ are well-founded orders, then so is $<_{lex}$.

ii. Prove that $<_{lex}$ is a well-founded order if both $<_A$ and $<_B$ are.

Now that we've introduced well-founded orders, we are ready to state the **principle of well-founded induction**. Let S be a set with a well-founded order $<_S$. Suppose that we want to show that some property P(s) holds for all s ∈ S. To do so, we can prove it as follows:

- For any minimal element s ∈ S, prove that P(s) is true.
- Assume that for some s ∈ S, P(s') is true for all s' < s. Then prove that P(s) is true.

The first part of the proof asserts the base case, that P(s) is true for all minimal elements of the set S. If we let S be the set of natural numbers, this would just be showing that P(0) is true. In the second step of the proof, we assume that for some s ∈ S, we know that for all smaller s', P(s') is true. We then use this to prove that P(s) is true. This is analogous to how a strong inductive proof assumes that P(n') is true for all 0 ≤ n' ≤ n, then proves that P(n + 1) is true.

Interestingly, it's possible to restate the above proof technique as follows:

Assume that for some s ∈ S, P(s') is true for all s' < s. Then prove that P(s) is true.

This seems odd – what happened to our base case? Don't worry; this is perfectly fine. To see why, suppose that we choose some s ∈ S that is a minimal element. In that case, it's vacuously true that for any s' $<_S$ s, P(s') is true. Consequently, if we're trying to prove that P(s) is true for a minimal element of S, then the assumption that for all s' $<_S$ s, P(s') is true gives us no extra information at all. We therefore need to prove that P(s) holds for this minimal element without any extra information, which amounts to just directly proving that P(s) is true. If you see inductive proofs in later courses, you may see proofs that work this way. The proofs still include base cases, though it may not be explicitly marked as such.

Well-founded induction has many uses cases, one of which is that it lets us simplify many proofs that use double induction. Recall our first example of a double induction, the proof that the Ackermann function is well-defined. The Ackermann function was specified as follows:

$$A(m,n) = \begin{cases} n+1 & \text{if } m=0 \\ A(m-1,1) & \text{if } m>0, n=0 \\ A(m-1, A(m, n-1)) & \text{if } m>0, n>0 \end{cases}$$

The proof in lecture was dense and difficult. However, if we switch to using well-founded induction using the lexicographical ordering of (m, n), the proof can be repeated in about half the space.

Here is how we might prove that A(m, n) is a function (that is, that $A : \mathbb{N}^2 \to \mathbb{N}$):

Theorem: $A : \mathbb{N}^2 \to \mathbb{N}$.

Proof: We need to show that A(m, n) is a natural number for any m, n ∈ ℕ. Let P(m, n) be "A(m, n) is a natural number." We proceed by well-founded induction on $\mathbb{N}^2$ using the lexicographical ordering on pairs of natural numbers to prove that P(m, n) is true for all m, n ∈ ℕ.

Assume that for some (m, n), that for any $(m', n') <_{lex} (m, n)$, P(m', n') is true (A(m', n') is a natural number). We prove that P(m, n) is true (A(m, n) is a natural number). We consider three cases:

Case 1: If m = 0, then by definition A(m, n) = A(0, n) = n + 1, which is a natural number.

Case 2: If m > 0 but n = 0, then by definition A(m, n) = A(m – 1, 1). Note that since m – 1 < m, $(m – 1, 1) <_{lex} (m, n)$. Thus by the inductive hypothesis, A(m – 1, 1) is some natural number, call it k. Thus A(m, n) = A(m – 1, 1) = k, which is a natural number.

Case 3: If m > 0 and n > 0, then by definition A(m, n) = A(m – 1, A(m, n – 1)). Because m = m and n – 1 < n, $(m, n – 1) <_{lex} (m, n)$. Thus by the inductive hypothesis, A(m, n – 1) is some natural number; call it k. This means that A(m, n) = A(m – 1, A(m, n – 1)) = A(m, n) = A(m – 1, k). Since m – 1 < m, $(m – 1, k) <_{lex} (m, n)$. Therefore, by the inductive hypothesis, A(m – 1, k) is a natural number, call it r. So A(m, n) = A(m – 1, A(m, n – 1)) = A(m – 1, k) = r, which is a natural number.

In each case, A(m, n) is a natural number, so P(m, n) is true, completing the induction. ■

Compare this proof to the original double-induction proof that we used in lecture. This version of the proof is much shorter, much cleaner, and (assuming that you're familiar with well-founded induction) much easier to read. It also gives a better intuition for why the proof works – each of the recursive calls made is to arguments that are, in a sense, "smaller," and they can't keep getting smaller forever.

iii. This proof does not have an explicit base case. Where does it handle the case where (m, n) is a minimal element of $\mathbb{N}^2$?

To give you a chance to play around with well-founded induction, let's consider the following game for two players, which is played on an m-by-n chessboard, where m > 0 and n > 0. Play begins with a stone in the upper-left corner of the board. Each turn, the current player either moves the stone an odd number of steps to the right or an odd number of steps down. If at the start of a player's turn the stone is in the bottom-right corner, that player wins. For example, here's a game played on a 5x5 board:

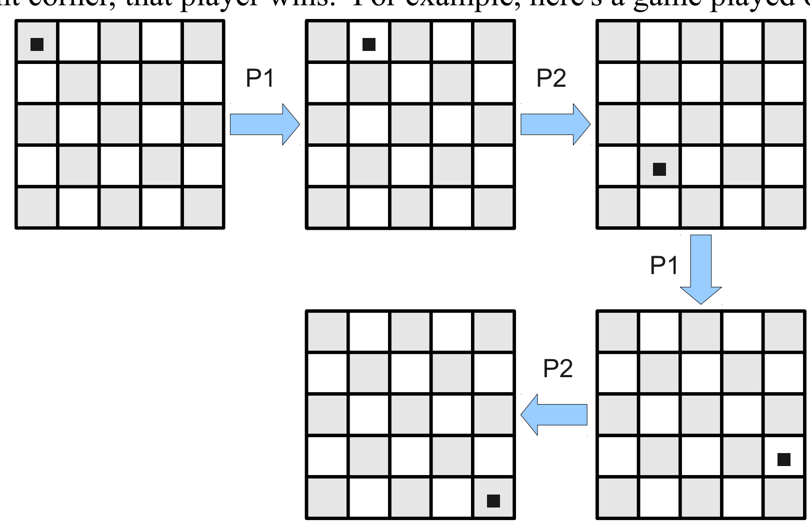

At this point, since the stone is in the bottom-right corner and it's the first player's turn, the first player wins the game.

iv. Suppose that we number the squares so that (0, 0) is the bottom-right corner and (m – 1, n – 1) is the top-left corner. Using well-founded induction and the $<_{lex}$ relation, prove that if at the start of some player's turn the stone is at position (*row*, *col*) where *row* and *col* are either both odd or both even, that player can always win the game. There are many other ways that you could prove this result, but to gain a familiarity with well-founded induction you must use that proof technique here.

Problem Six: Course Feedback (5 Points)

We want this course to be as good as it can be, and we'd really appreciate your feedback on how we're doing. For a free five points, please answer the following questions. We'll give you full credit no matter what you write (as long as you write something!), but we'd appreciate it if you're honest about how we're doing.

i. How hard did you find this problem set? How long did it take you to finish?

ii. Does that seem unreasonably difficult or time-consuming for a five-unit class?

iii. How is the pace of this course so far? Too slow? Too fast? Just right?

iv. Is there anything in particular we could do better? Is there anything in particular that you think we're doing well?

Submission Instructions

There are three ways to submit this assignment:

1. Hand in a physical copy of your answers at the midterm.

2. Hand in a physical copy of your answers at the start of lecture (if you do this on Friday, you'll be using a late day.)

3. Submit a physical copy of your answers in the filing cabinet in the open space near the handout hangout in the Gates building. If you haven't been there before, it's right inside the entrance labeled "Stanford Venture Fund Laboratories." There will be a clearly-labeled filing cabinet into which you can submit your homework.

4. Send an email with an electronic copy of your answers to **cs103@cs.stanford.edu**