CS103 Handout 13
Spring 2015 April 24, 2015

Problem Set 4

This fourth problem set explores functions, the pigeonhole principle, and binary relations. It will be the final problem set in our exploration of discrete mathematics, and by the time you're done with it you'll be ready to start tackling the limits of computability!

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

This problem set has 34 possible points. It is weighted at 5% of your total grade.

Good luck, and have fun!

Checkpoint due Monday, April 27 at the start of lecture.

Remaining problems due Monday, May 4 at the start of lecture.

Write your solutions to the following problems and submit them online by Monday, April 27th at the start of class. These problems will be graded on a 0/1/2 scale based on whether you have attempted to solve all the problem, rather than on correctness. We will try to get these problems returned to you with feedback on your proof style this Wednesday, April 29th.

Checkpoint Problem: Paradoxical Sets (2 Points)

What happens if we take *absolutely everything* and throw it into a set? If we do, we would get a set called the *universal set*, which we denote $\mathcal{U}$:

$$\mathcal{U} = \{ x \mid x \text{ exists} \}$$

Absolutely everything would belong to this set: $1 \in \mathcal{U}$, $\mathbb{N} \in \mathcal{U}$, philosophy $\in \mathcal{U}$, CS103 $\in \mathcal{U}$, etc. In fact, we'd even have $\mathcal{U} \in \mathcal{U}$, which is strange but not immediately a problem.

Unfortunately, the set $\mathcal{U}$ doesn't actually exist, as its existence would break mathematics.

i. Prove that if A and B are sets where $A \subseteq B$, then $|A| \le |B|$. Although this probably makes intuitive sense, to formally prove this result, you need to find an injection $f : A \to B$ and prove that your function is injective.

ii. Using your result from (i), prove that if $\mathcal{U}$ exists at all, then $|\wp(\mathcal{U})| \le |\mathcal{U}|$.

iii. Using your result from (ii) and Cantor's Theorem, prove that $\mathcal{U}$ does not exist.

The result you've proven shows that there is a collection of objects (namely, the collection of everything that exists) that cannot be put into a set. This goes against our intuition of what a set is. When this was discovered at the start of the twentieth century, it caused quite a lot of chaos in the math world and led to a reexamination of logical reasoning itself and a more formal definition of what objects can and cannot be gathered into a set. If you're curious to learn more about what sets can and cannot be created, take Math 161 (Set Theory).

Problem One: Properties of Functions (4 Points)

Consider the following Venn diagram:

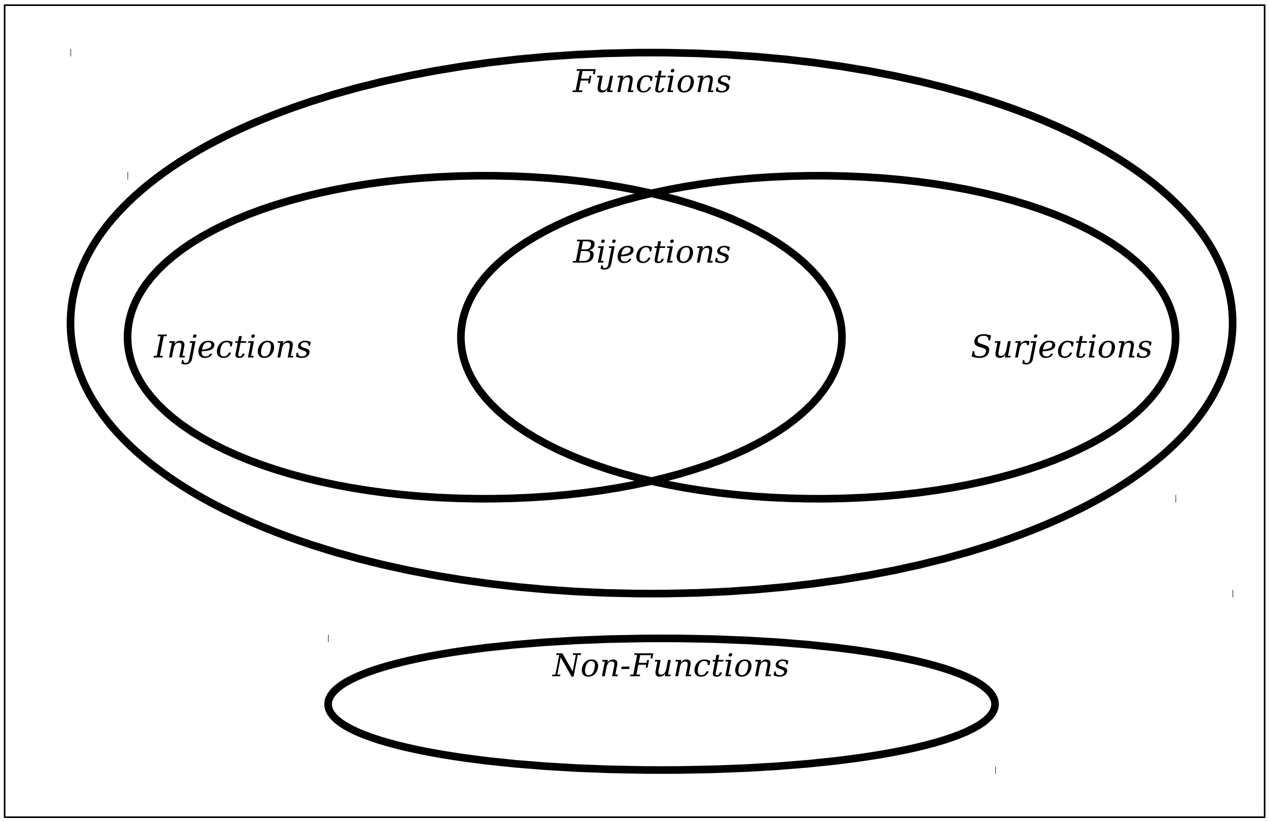

Below is a list of purported functions. For each of those purported functions, determine where in this Venn diagram that object goes. No justification is necessary.

1. $f : \mathbb{N} \to \mathbb{N}$ defined as $f(n) = n^2$
2. $f : \mathbb{Z} \to \mathbb{N}$ defined as $f(n) = n^2$
3. $f : \mathbb{N} \to \mathbb{Z}$ defined as $f(n) = n^2$
4. $f : \mathbb{Z} \to \mathbb{Z}$ defined as $f(n) = n^2$
5. $f : \mathbb{R} \to \mathbb{N}$ defined as $f(n) = n^2$
6. $f : \mathbb{N} \to \mathbb{R}$ defined as $f(n) = n^2$
7. $f : \mathbb{N} \to \mathbb{N}$ defined as $f(n) = \sqrt{n}$.
8. $f : \mathbb{R} \to \mathbb{R}$ defined as $f(n) = \sqrt{n}$.
9. $f : \mathbb{R} \to \{ x \in \mathbb{R} \mid x \geq 0 \}$ defined as $f(n) = \sqrt{n}$.
10. $f : \{ x \in \mathbb{R} \mid x \geq 0 \} \to \{ x \in \mathbb{R} \mid x \geq 0 \}$ defined as $f(n) = \sqrt{n}$.
11. $f : \mathbb{N} \to \wp(\mathbb{N})$, where f is some injective function.
12. $f : \{a, b, c\} \to \{x, y\}$, where f is some surjective function.
13. f : {breakfast, lunch, dinner} $\to$ {shakshuka, soondubu, maafe}, where f is an injection.

Problem Two: Cartesian Products and Cardinalities (5 Points)

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 $|A \times B| = |C \times D|$. If you try this out for a few finite sets, this might seem obvious, though it's less obviously correct for infinite sets. To prove that this result is true for all sets, we'll use the formal definition of equal cardinalities. Since we know that $|A| = |C|$, there has to be some bijection $g : A \to C$. Similarly, since we know that $|B| = |D|$, there has to be some bijection $h : B \to D$. Our goal will be to define a bijection $f : A \times B \to C \times D$.

Here's one possibility. Consider the function $f : A \times B \to C \times D$ 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 $C \times D$. *(Hint: Use the formal definitions of injectivity, surjectivity, and bijectivity.)*

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:

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

ii. Using your result from (i) and the above definition, prove 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. ☺)

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 $\kappa_1 \cdot \kappa_2$ 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 of $|A \times B|$ 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 of $|A \times B|$ 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 (ii) shows that $\aleph_0^n = \aleph_0$ for any positive natural number n. Isn't infinity weird?

Problem Three: Understanding Diagonalization (2 Points)

Proofs by diagonalization are quite tricky and rely on nuanced arguments. In this problem, we'll ask you to review the diagonalization proof we covered in lecture to help you better understand how it works.

i. Consider the function $f : \mathbb{N} \to \wp(\mathbb{N})$ defined as $f(n) = \emptyset$. Trace through our proof of Cantor's theorem with this choice of f in mind. In the middle of the argument, the proof defines some set D in terms of f. Given that $f(n) = \emptyset$, what is that set D? Is it clear why $f(n) \neq D$ for any $n \in \mathbb{N}$?

ii. Repeat part (i) of this problem using the function $f : \mathbb{N} \to \wp(\mathbb{N})$ defined as

$$f(n) = \{ \ m \in \mathbb{N} \mid m > n \ \}$$

Now what do you get for the set D? Is it clear why $f(n) \neq D$ for any $n \in \mathbb{N}$?

Problem Four: Simplifying Cantor's Theorem? (2 Points)

In lecture, we proved Cantor's theorem, that if S is a set, then $|S| < |\wp(S)|$. Our proof used a diagonal argument, which is clever but tricky. Below is a purported proof that $|S| \neq |\wp(S)|$ that doesn't use a diagonal argument:

Theorem: If S is a set, then $|S| \neq |\wp(S)|$.

Proof: Let S be any set and consider the function $f : S \to \wp(S)$ defined as $f(x) = \{x\}$. To see that this is a valid function from S to $\wp(S)$, note that for any $x \in S$, we have $\{x\} \subseteq S$. Therefore, $\{x\} \in \wp(S)$ for any $x \in S$, so f is a legal function from S to $\wp(S)$.

Let's now prove that f is injective. Consider any $x_1, x_2 \in S$ where $f(x_1) = f(x_2)$. We'll prove that $x_1 = x_2$. Because $f(x_1) = f(x_2)$, we have $\{x_1\} = \{x_2\}$. Since two sets are equal if and only if their elements are the same, this means that $x_1 = x_2$, as required.

However, f is not surjective. Notice that $\emptyset \in \wp(S)$, since $\emptyset \subseteq S$ for any set S, but that there is no x such that $f(x) = \emptyset$; this is because $\emptyset$ contains no elements and $f(x)$ always contains one element. Since f is not surjective, it is not a bijection. Thus $|S| \neq |\wp(S)|$ ■

Unfortunately, this proof is incorrect. What's wrong with this proof? Justify your answer.

Problem Five: Coloring a Grid (4 Points)

You are given a 3×9 grid of points, like the one shown below:

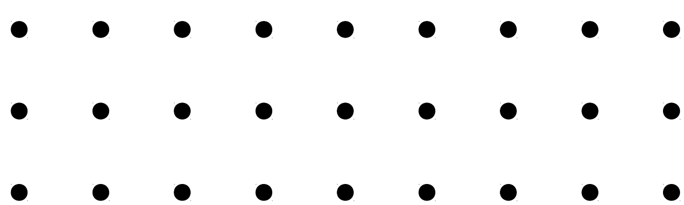

Suppose that you color each point in the grid either red or blue. Prove that no matter how you color those points, you can always find four points of the same color that form the corners of a rectangle.

Problem Six: Properties of Relations (5 Points)

In this problem, we'll ask you to give examples of relations with various properties. You can define a relation either by giving an existing example of a relation (for example, $\leq$, $\subseteq$, etc.) or by drawing a graph that defines the relation (where each node is some element and an edge from u to v indicates that u is related to v). You might find the online relation editor useful for finding relations with each property, but you'll still need to formally prove why your relations have the desired properties.

This problem is designed to make sure you understand the definitions of symmetry, reflexivity, transitivity, antisymmetry, partial orders, and equivalence relations. We strongly recommend reviewing those definitions before working through these problems.

i. Give an example of a relation that is neither an equivalence relation nor a partial order. Prove that your relation meets these criteria.

ii. Give an example of a relation that is both a partial order and an equivalence relation (yes, this is possible!) Prove that your relation meets these criteria. *(Hint: Think about the graph of a relation. If that relation is an equivalence relation, how many edges can there be between any pair of distinct nodes? If that relation is a partial order, how many edges can there be between any pair of distinct nodes?)*

In what follows, if p is a person, let $H(p)$ denote that person's height, rounded to the nearest centimeter, and let P be the set of all people.

iii. Define the relation $=_H$ over P as follows: if x and y are people, then $x =_H y$ if $H(x) = H(y)$. Is $=_H$ an equivalence relation? If so, prove it. If not, prove why not.

iv. Define the relation $\leq_H$ over P as follows: if x and y are people, then $x \leq_H y$ if $H(x) \leq H(y)$. Is $\leq_H$ a partial order? If so, prove it. If not, prove why not. *(Hint: Read the definition of a partial order very, very carefully!)*

Problem Seven: Meet Semilattices (5 Points)

Most compilers these days don't just compile code; they improve it as well. Compilers that improve the programs they generate are called *optimizing compilers*.

Many compiler optimizations are based on a mathematical structure called a ***meet semilattice***. A meet semilattice is an ordered pair $(D, \wedge)$, where D is a set of values and $\wedge$ is a binary operator called a ***meet operator*** that can be applied to pairs of those values. Although the symbol $\wedge$ is the same one that we will use for logical AND, in the context of meet semilattices we use $\wedge$ to represent the meet operator. For example, we would read $x \wedge y$ as "x meet y" rather than "x and y."

In order for $(D, \wedge)$ to be a meet semilattice, the following properties must hold of D and $\wedge$:

- D must be ***closed under*** $\wedge$: If $x \in D$ and $y \in D$, then $x \wedge y \in D$.
- $\wedge$ must be ***idempotent***: If $x \in D$, then $x \wedge x = x$.
- $\wedge$ must be ***commutative***: If $x \in D$ and $y \in D$, then $x \wedge y = y \wedge x$.
- $\wedge$ must be ***associative***: If $x \in D$, $y \in D$, and $z \in D$, then $x \wedge (y \wedge z) = (x \wedge y) \wedge z$.

As an example, the function **min** over $\mathbb{R}$, which takes in two real numbers and returns the smaller one, is a meet semilattice. The set intersection operator $\cap$ is also meet semilattice over the set $\wp(\mathbb{N})$.

Amazingly, these four rules about the behavior of $\wedge$, which say nothing about how elements of D rank against one another, allow us to define a partial order over the elements of D. Given a meet semilattice $S = (D, \wedge)$, define the relation $\leq_S$ over D as follows:

$$x \leq_S y \quad \text{if} \quad x \wedge y = x$$

This might all seem really abstract at this point, so let's make things more concrete.

i. Suppose that D is the set of all real numbers and $x \wedge y = \mathbf{min}\{x, y\}$. What is the relation $\leq_S$ in this case? Briefly justify your answer; no proof is necessary.

ii. Suppose that D is the set $\wp(\mathbb{N})$ and that $x \wedge y = x \cap y$. What is the relation $\leq_S$ in this case? Briefly justify your answer; no proof is necessary.

Now, we'd like you to prove some properties about the $\leq_S$ relation.

iii. Let $S = (D, \wedge)$ be an arbitrary meet semilattice. Prove that the relation $\leq_S$ is a partial order over D.

The $\leq_S$ relation defined above interacts with $\wedge$ in interesting ways.

iv. Prove for all $x, y \in D$ that $x \wedge y \leq_S x$ and $x \wedge y \leq_S y$. This proves the value $x \wedge y$ is a *lower bound* of x and y.

v. Prove for all $x, y, z \in D$ that if $z \leq_S x$ and $z \leq_S y$, then $z \leq_S x \wedge y$. This proves the value $x \wedge y$ is the *greatest lower bound* of x and y.

In the context of program analysis, semilattices give a mathematical model of what information is known about a program at a particular point. Larger values according to $\leq_S$ indicate more precise information about the program, while smaller values indicate less precise information. The meet operator then gives a mechanism for combining pieces of information together in a way that preserves as much information as possible. If you're curious how semilattices are used this way, consider taking CS243.

Problem Eight: Chains and Antichains (5 Points)

Let A be an arbitrary set and $\leq_A$ be an arbitrary partial order over A. We'll say that a ***chain*** in A is a series of distinct elements $x_1, x_2, \ldots, x_k$ such that

$$x_1 \leq_A x_2 \leq_A \ldots \leq_A x_k.$$

Intuitively, a chain is a series of values in ascending order according to the partial order $\leq_A$. The ***length*** of an chain is the number of elements in that chain.

i. Consider the $\subseteq$ relation over the set $\wp(\{a, b, c\})$. What is the length of the longest chain in this partial order? Give an example of a chain with that length. No justification is necessary.

An ***antichain*** is a set $X \subseteq A$ such that no two elements in X can be compared by the $\leq_A$ relation. In other words, a set $X \subseteq A$ is an antichain if for any $a, b \in X$, both $a \leq_A b$ and $b \leq_A a$ are false. The ***size*** of an antichain X is the number of elements in X.

ii. Consider the $\subseteq$ relation over the set $\wp(\{a, b, c\})$. What is the size of the largest antichain in this partial order? Give an example of an antichain with that size. No justification is necessary.

Let $\lceil x \rceil$ denote the smallest integer greater than or equal to x, so $\lceil 1 \rceil = 1$, $\lceil 1.37 \rceil = 2$, and $\lceil \pi \rceil = 4$. 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.

Given an arbitrary partially ordered set, you can't say anything a priori about the size of the largest chain or antichain in that partial order. However, you can say that at least one of them must be relatively large relative to the partially ordered set. Let r and s be natural numbers. We're going to ask you to prove the following result: if $|A| = rs+1$, then either A contains a chain of length $r+1$ or an antichain of size $s+1$. The propositional equivalence $P \lor Q \equiv \neg P \to Q$ will be useful here. To prove at least one of P or Q is true, you can instead prove that if P is false, then Q is true. For the purposes of this problem, we're going to prove that if A does ***not*** contain an chain of length at least $r+1$, it does contain an antichain of size at least $s+1$.

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 $r+1$ or greater, then there must be at least $s+1$ elements of A at the same height. *(Hint: Use the generalized pigeonhole principle.)*

iv. Your result from part (iii) establishes that there must be a collection of $s+1$ elements of A at the same height as one another. Let X be any set of $s+1$ such elements. Prove that X must be an antichain. *(Hint: The height of an element is the length of the longest possible chain ending at that element, so you know that if element x has height h, there is some chain of length h ending in x and there are no chains of length greater than h ending in x.)*

Intuitively speaking, if $\leq_A$ is a partial 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.

Extra Credit Problem: *k*-Regular Graphs (1 Point Extra Credit)

An undirected graph G is called ***k-regular*** if every node in G has degree exactly k. The ***girth*** of a graph is the length of the shortest simple cycle in G. If G has no cycles, its girth is infinite.

Prove that any k-regular graph with girth five has at least $k^2 + 1$ nodes.