Excerpt from course reader for Stanford CS 103:
Keith Schwarz

Chapter 2 Mathematical Proof

Last chapter we concluded with Cantor's theorem, the fact that the cardinality of the power set of a set S is always greater than the cardinality of the set S itself. Although we worked through a strong argument that this should be true, did we really "prove" it? What does it mean to prove something, at least in a mathematical sense?

Proofs are at the core of the mathematical foundations of computing. Without proofs we couldn't be certain that any of our results were correct, and our definitions would be little better than an intuition to guide us. Accordingly, before we attempt to explore the limits of computation, we first need to build up the machinery necessary to reason about and firmly establish mathematical results.

Proofs are in many ways like programs – they have their own vocabulary, terminology, and structure, and you will need to train yourself to think differently in order to understand and synthesize them. In this chapter and the ones that follow, we will explore proofs and proof techniques, along with several other concepts that will serve as a "proving ground" for testing out these proof ideas.

One quick note before we continue – because this chapter focuses on how to structure mathematical proofs, some of the results that we'll be proving early on will be pretty trivial. I promise you that the material will get a lot more interesting toward the end of the chapter, and once we make it into Chapter Three the results we will be proving will be much more involved and a lot more interesting. Please don't get the impression that math is painfully boring and pedantic! It's a really fascinating subject, but we need to build up a few techniques before we can jump into the real meat of the material.

2.1 What is a Proof?

In order to write a proof, we need to start off by coming up with some sort of definition of the word "proof." Informally, a mathematical proof is a series of logical steps starting with one set of assumptions that ends up concluding that some statement must be true. For example, if we wanted to prove the statement

$$\text{If } x + y = 16, \text{ then either } x \geq 8 \text{ or } y \geq 8$$

then we would begin by assuming that $x + y = 16$, then apply sound logical reasoning until we had arrived at the conclusion that $x \geq 8$ or $y \geq 8$. Similarly, if we wanted to prove that

$$\text{For any set } S, |S| < |\wp(S)|$$

(as we started doing last chapter), we would take as our starting point all of the definitions from set theory – what the power set is, what it means for one set to have smaller cardinality than another, etc. - and would proceed through logical steps to conclude that $|S| < |\wp(S)|$.

Writing a proof is in many ways like writing a computer program. You begin with some base set of things that you know are true (for example, how primitive data types work, how to define classes, etc.), then proceed to use those primitive operations to build up something more complicated. Also like a program, proofs have their own vocabulary, language, structure, and expectations. Unfortunately, unlike programs, there is no "compiler" for proofs that can take in a proof and verify that it's a legal mathematical proof.* Consequently, learning how to write proofs takes time and effort.

---

* Technically speaking such programs exist, but they require the proof to be specified in a very rigid format that is almost never used in formal mathematical proofs.

In this chapter, we will introduce different types of proofs by analyzing real proofs and seeing exactly how they work. We'll also see what *doesn't* work and the sort of logical traps you can easily fall into when writing proofs.

2.1.1 Transitioning to Proof-Based Mathematics

The math that you're probably most familiar with is the style of math that you saw in high school. Typically, high school mathematics focuses on *calculation*. For example, you'll probably get problems like this one:

> Two trains are 50 miles apart from one another and on parallel tracks. One train takes off from the station heading at 24 miles per hour, while the other train takes off heading 12 miles per hour. The two trains are driving toward one another. How long will it take for the two trains to pass one another?

Or perhaps something like this:

> Evaluate the integral $\int_0^5 x\sqrt{1+x^2}\,dx$

Or perhaps something like this:

> A square is circumscribed inside a circle of radius 15. The square is then divided into four right triangles whose right angle is the center of the square. What is the sum of the perimeters of these triangles?

Problems like these have exact numeric values associated with them, and the goal of the problem is to work through a calculation to come up with some number. You can memorize different approaches for solving these sorts of problems and build up a repertoire of techniques for solving them. The focus on these problems is determining which calculations to perform.

The math that we will be doing in this course is of a totally different flavor. We will be interested in proving general properties of different types of objects (numbers, data structures, computer programs, etc.) In doing so, we may perform some small calculations from time to time, but our main objective will be establishing a clear line of reasoning that rigorously demonstrates a result. In fact, you'll find that most proofs that we'll be writing read like essays with some mathematical jargon, since the goal will be to describe a rigorous line of reasoning rather than to compute a value.

2.1.2 What Can We Assume?

One of the most difficult aspects of writing a proof is determining what you can assume going into the proof. In journals, proofs often assume that the reader is familiar with important results, and often cite them without reviewing why they're true. For our purposes, though, we will deliberately play dumb and start with a very weak set of assumptions. We will prove pretty much everything we need, even if it seems completely obvious, in order to see how to formalize intuitive concepts with a level of mathematical rigor.

In this book, we will assume that whoever is reading one of our proofs knows

1. All definitions introduced so far,
2. All theorems introduced so far, and
3. Basic algebra.

We will not assume anything more than this. For example, we're fine assuming that if $x < y$ and $y < z$, then $x < z$, but we will not assume that for any set S, $S \cap \varnothing = \varnothing$ even though this seems "obvious." As we build

up our mathematical repertoire, the set of assumptions we can make will grow, and it will become easier and easier to prove more elaborate results. This is similar to writing libraries in computer programs – although it's difficult and a bit tedious to implement standard data structures like **`ArrayList`** and **`HashMap`**, once you've put in the work to do so it becomes possible to build up off of them and start writing much more intricate and complex programs.

2.2 Direct Proofs

Just as it's often easiest to learn how to program by jumping into the middle of a "Hello, World!" program and seeing how it works, it's useful to jump right into some fully worked-out mathematical proofs to see how to structure general proofs.

To begin our descent into proofs, we'll introduce two simple definitions, then see how to prove results about those definitions.

An integer x is called ***even*** if there is some integer k such that $x = 2k$.

An integer x is called ***odd*** if there is some integer k such that $x = 2k + 1$.

For example, 4 is even since $4 = 2 \times 2$. 8 is even as well, since $8 = 4 \times 2$. 9 is odd, because $9 = 4 \times 2 + 1$. We consider 0 to be an even number, since $0 = 0 \times 2$.

Given this, let's prove the following result:

Theorem: If x is even, then x^2 is even.

Proof: Let x be any even integer. Since x is even, there is some integer k such that $x = 2k$. This means that $x^2 = (2k)^2 = 4k^2 = 2(2k^2)$. Since $2k^2$ is an integer, this means that there is some integer m (namely, $2k^2$) such that $x^2 = 2m$. Thus x^2 is even. ■

Let's look at how this proof works. The proof proceeds in several clean logical steps, each of which we can analyze independently.

First, note how the proof starts: "Let x be any even integer." The goal of this proof is to show that if x is even, then x^2 is even as well. This proof should work no matter what choice of x we make – whether it's 0, 2, 4, 6, 8, etc. In order for our proof to work in this general setting, the proof will proceed by using x as a placeholder for whatever even number we're interested in. If we wanted to convince ourselves that some particular even number has a square that's also even, we could just plug that even number into the proof wherever we were using x. For example, if we want to prove that 12^2 is even, the proof would go like this:

Proof: Since ***12*** is even, there is some integer k such that ***12*** $= 2k$. (This integer k is the integer 6). This means that ***12***$^2 = (2 \times 6)^2 = 4 \times 6^2 = 2(2 \times 6^2) = 2 \times 72$. Since 72 is an integer, this means that there is some integer m (namely, 72) such that ***12***$^2 = 2m$. Thus ***12***^{2} is even. ■

All that we've done here is substitute in the number 12 for our choice of x. We could substitute in any other even number if we'd like and the proof would still hold. In fact, that's why the proof works – we've shown that no matter what choice of an even number you make for x, you can always prove that x^2 is even as well.

Let's continue dissecting this proof. After we've decided to let x be a placeholder for whatever even number we'd like, we then say

> Since x is even, there is some integer k such that $x = 2k$

What does this statement mean? Well, we know that x is an even number, which means that it must be twice some other number. We can't really say what that number is, since we don't know what our choice of x is. However, we can say that there is *some* number such that x is twice that number. In order to manipulate that number in this proof, we'll give this number a name (in this proof, we call it k). Interestingly, note that nowhere in this sentence do we actually say how to figure out what this value of k is; we just say that it has to exist and move forward. From a programming perspective, this may seem strange – it seems like we'd have to show how to find this number k in order to assert that it exists! But it turns out that it's perfectly fine to just say that it exists and leave it at that. Our definition of an even number is an integer that is equal to twice some other number, so we know for a fact that because x is even, this number k must exist.

At this point, we know that $x = 2k$, and our goal is to show that x^2 is even. Let's think about how to do this. To show that x^2 is even, we will need to find some integer m such that $x^2 = 2m$. Right now, all that we know is that x is even and, as a result, that $x = 2k$ for some choice of k. Since we don't have much else to go on right now, let's try seeing if we can describe x^2 in terms of x and k. Perhaps doing this will lead us to finding some choice of m that we can make such that $x^2 = 2m$. This leads to the next part of the proof:

> This means that $x^2 = (2k)^2 = 4k^2 = 2(2k^2)$. Since $2k^2$ is an integer, this means that there is some integer m (namely, $2k^2$) such that $x^2 = 2m$

The first of these two sentences is a simple algebraic manipulation. We know that $x = 2k$, so $x^2 = (2k)^2$. If we simplify this, we get $x^2 = 4k^2$, which is in turn equal to $2(2k^2)$. This last step – factoring out the two from the expression – then makes it clearer that x^2 is twice some other integer (specifically, the integer $2k^2$). We can then conclude that there is some natural number m such that $x^2 = 2m$, since we've found specifically what that value was. Because we've done this, we can conclude the proof by writing:

> Thus x^2 is even. ■

This holds because the definition of an even number is one that can be written as $2m$ for some integer m. Notice that we've marked the end of the proof with the special symbol ■, which serves as a marker that we're done. Sometimes you see proofs ended with other statements like “This completes the proof” or “QED” (from the Latin “quod erat demonstrandum,” which translates roughly to “which is what we wanted to show”). Feel free to end your own proofs with one of these three endings.

Let's take a look at an example of another proof:

Theorem: If m is even and n is odd, then mn is even.

Proof: Let m be any even number and n be any odd number. Then $m = 2r$ for some integer r, and $n = 2s + 1$ for some integer s. This means that $mn = (2r)(2s + 1) = 2(r(2s + 1))$. This means that $mn = 2k$ for some integer k (namely, $r(2s + 1)$), so mn is even. ■

The structure of this proof is similar to the previous proof. We want to show that the claim holds for any choice of even m and odd n, so we begin by letting m and n be any even and odd number, respectively. From there, we use the definition of even and odd numbers to write $m = 2r$ and n as $2s + 1$ for some integers r and s. As with the previous proof, we don't actually know what these numbers are, but they're guaranteed to exist. After doing some simple arithmetic, we end up seeing that $mn = 2(r(2s + 1))$, and since $r(2s + 1)$ is an integer, we can conclude that mn is twice some integer, and so it must be even.

The above proofs are both instances of ***direct proofs***, in which the proposition to be proven is directly shown to be true by beginning with the assumptions and ending at the conclusions.

2.2.1 Proof by Cases

Let's introduce one new definition, which you may be familiar with from your programming background:

The ***parity*** of an integer is whether it is odd or even. Two numbers have the same parity if they are both odd or both even.

For example, 1 and 5 have the same parity, because both of the numbers are odd, and 4 and 14 have the same parity because both 4 and 14 are even. However, 1 and 2 have opposite parity, because 1 is odd and 2 is even.

The following result involves the parity of integers:

Theorem: If m and n have the same parity, then $m + n$ is even.

Before we try to prove this, let's check that it's actually correct by testing it on a few simple examples. We can see that $2 + 6 = 8$ is even, and $1 + 5 = 6$ is even as well. But how would we prove that this is true in the general case? In a sense, we need to prove two separate claims, since if m and n have the same parity, then either both m and n are even or both m and n are odd. The definitions of odd numbers and even numbers aren't the same, and so we have to consider the two options separately. We can do this cleanly in a proof as follows:

Proof: Let m and n be any two integers with the same parity. Then there are two cases to consider:

Case 1: m and n are even. Then $m = 2r$ for some integer r and $n = 2s$ for some integer s. Therefore, $m + n = 2r + 2s = 2(r + s)$. Thus $m + n = 2k$ for some integer k (namely, $r + s$), so $m + n$ is even.

Case 2: m and n are odd. Then $m = 2r + 1$ for some integer r and $n = 2s + 1$ for some integer s. Therefore, $m + n = 2r + 1 + 2s + 1 = 2r + 2s + 2 = 2(r + s + 1)$. Thus $m + n = 2k$ for some integer k (namely, $r + s + 1$), so $m + n$ is even. ■

Note how this proof is structured as two cases – first, when m and n are even, and second, when m and n are odd. This style of proof is sometimes called a **proof by cases** or a **proof by exhaustion** (because we've exhausted all possibilities and found that the claim is true, not because we're tired of writing proofs). Each of the branches of the proof reads just like a normal proof, but is individually insufficient to prove the general result. Only when we show that in both of the possible cases the result holds can we conclude that the claim is true in general.

When writing a proof by exhaustion, it's critically important to remember to check that you have covered all possible cases! If you have a proof where four options are possible and you only prove three cases, your proof is likely to be incorrect.

Let's see another example of a proof by cases:

Theorem: If n is even and m is an integer, then $n + m$ has the same parity as m.

Before proving this, it's always good to check that it works for a few test cases. If we let $n = 4$, then we can see that

- $4 + 3 = 7$, and 7 has the same parity as 3.
- $4 + 6 = 10$, and 10 has the same parity as 6.

Let's see a proof of this result:

Proof: Consider any even integer n. Now, consider any integer m and the sum $n + m$. We consider two possibilities for m:

Case 1: m is even. Then m and n have the same parity, so by our previous result (if m and n have the same parity, then $m + n$ is even) we know that $m + n$ is even. Therefore m and $m + n$ have the same parity.

Case 2: m is odd. Since n is even, $n = 2r$ for some integer r, and since m is odd, $m = 2s + 1$ for some integer s. Then $n + m = 2r + 2s + 1 = 2(r + s) + 1$. This means that $n + m = 2k + 1$ for some k (namely, $r + s$), so $n + m$ is odd. Therefore m and $m + n$ have the same parity. ■

This proof is interesting for two reasons. First, notice that in proving that Case 1 is true, we used the result that we have proven previously: if n and m have the same parity, then $n + m$ is even. This means that we didn't have to try writing $n = 2r$ and $m = 2s$, and we ended up saving a lot of space in our proof. Whenever you're writing a proof, feel free to cite any result that you have previously proven. In CS103, it's perfectly

fine to cite proofs from lecture, this book, or the problem sessions, as long as you make it clear what result you're using.

Second, notice that in this proof the cases resulted from the parity of just one of the variables (m). We knew that the parity of n must be even, and the only thing that was unclear was whether m was odd or even.

2.2.1.1 A Quick Aside: Choosing Letters

If you'll notice, the proofs that we've done so far use lots of letters to stand for numbers: m, n, k, r, s, etc. In general, when writing a proof, you should feel free to choose whatever letters you think will make the proof flow most cleanly. However, you should make an effort to pick a consistent naming convention, much in the same way that you should adopt a naming convention when writing computer programs.

In this set of course notes, I will typically use single capital letters (S, T, U) to represent sets. I tend to use the letters m, n, and k to represent natural numbers and x, y, z to represent integers. If I run out of letters, I might borrow others from other parts of the alphabet (for example, r, s, and t for natural numbers) if we exhaust our normal supply.

When working with values in a sequence (more on that later), I'll tend to use subscripted symbols like x_1, x_2, …, x_i. In those cases, the letters i, j, and k will typically refer to variable indices, and n will represent quantities like the total number of elements in the sequence.

2.2.2 Proofs about Sets

In the last chapter, we explored sets and some of the operations on them. You have already seen one theorem about sets (specifically, Cantor's theorem). But what else can we prove about sets? And how do we prove them?

Let's begin with a very simple proof about sets:

> ***Theorem:*** For any sets A and B, $A \cap B \subseteq A$.

This theorem intuitively makes sense. We can think of $A \cap B$ as the set of things that A and B have in common. In other words, we're filtering down the elements of A by just considering those elements that also happen to be in B. As a result, we should end up with a set that's a subset of the set A. So how do we prove this? As you will see, the proof works similarly to our proof about odd and even numbers: it calls back to the definitions of intersection and subset, then proceeds from there.

> ***Proof***: Consider any sets A and B. We want to show that $A \cap B \subseteq A$. By the definition of subset, this means that we need to show that for any $x \in A \cap B$, $x \in A$. So consider any $x \in A \cap B$. By the definition of intersection, $x \in A \cap B$ means that $x \in A$ and $x \in B$. Therefore, if $x \in A \cap B$, $x \in A$. Since our choice of x was arbitrary, $A \cap B \subseteq A$. ■

Let's examine the structure of the proof. We initially wanted to prove that $A \cap B \subseteq A$. To do this, we said something to the effect of "okay, I need to prove that $A \cap B \subseteq A$. What does this mean?" By using the definition of subset, we were able to determine that we needed to prove that for any choice of $x \in A \cap B$, it's true that $x \in A$. Again we ask – so what does it mean for $x \in A \cap B$? Again we call back to the definition: $x \in A \cap B$ means that $x \in A$ and $x \in B$. But at this point we're done – we needed to show that any $x \in A \cap B$ also satisfies $x \in A$, but the very definition of $A \cap B$ guarantees this to us!

This proof illustrates a crucial step in many proofs – if you are unsure about how to proceed, try referring to the definitions of the terms involved. Often this simplification will help you make progress toward the ultimate proof by rewriting complex logic in terms of something similar.

Let's do another simple proof:

> ***Theorem:*** For any sets A and B, $A \subseteq A \cup B$.

This result says that if we take any collection of things (the set A) and combine it together with any other set of things (forming the set $A \cup B$), then the original set is a subset of the resulting set. This seems obvious – after all, if we mix in one set of things with another, that initial set is still present! Of course, it's good to formally establish this, which we do here:

> **Proof:** Consider any sets A and B. We want to show that $A \subseteq A \cup B$. To do this, we show that for any $x \in A$, that $x \in A \cup B$ as well. Note that by definition, $x \in A \cup B$ iff $x \in A$ or $x \in B$.
>
> Consider any $x \in A$. It is therefore true that $x \in A$ or $x \in B$, since we know $x \in A$. Consequently, $x \in A \cup B$. Since our choice of x was arbitrary, this shows that any $x \in A$ also satisfies $x \in A \cup B$. Consequently, $A \subseteq A \cup B$, as required. ■

Again, notice the calling back to the definitions. To prove $A \subseteq A \cup B$, we argue that every $x \in A$ also satisfies $x \in A \cup B$. What does it mean for $x \in A \cup B$? Well, the definition of $A \cup B$ is the set of all x such that either $x \in A$ or $x \in B$. From there we can see that we're done – if $x \in A$, then it's also true that $x \in A$ or $x \in B$, so it's true that $x \in A \cup B$.

Let's do another proof, this time proving a slightly more complex result:

> ***Theorem:*** For any sets A, B, and C, we have $C - (A \cap B) \subseteq (C - A) \cup (C - B)$

As an example, let's take $A = \{1, 2, 3\}$, $B = \{3, 4, 5\}$, and $C = \{1, 2, 3, 4, 5\}$. Then we have that

$$C - (A \cap B) = \{1, 2, 3, 4, 5\} - \{3\} = \{1, 2, 4, 5\}.$$

We also have that

$$(C - A) \cup (C - B) = \{4, 5\} \cup \{1, 2\} = \{1, 2, 4, 5\}.$$

Thus in this single case, $C - (A \cap B) \subseteq (C - A) \cup (C - B)$, since the two sets are equal.

This theorem worked out in the above case, but it's not at all clear exactly why this would be true. How, then, could we prove this?

Whenever you need to write a proof, *always be sure that you understand why what you're proving is true before you try to prove it!* Otherwise, you're bound to get stuck. So before we start to work through the actual proof, let's try to build up an intuition for why this result is true. To do so, let's turn to Venn diagrams, which are surprisingly useful when proving results like these.

Let's start with this Venn diagram for three sets:*

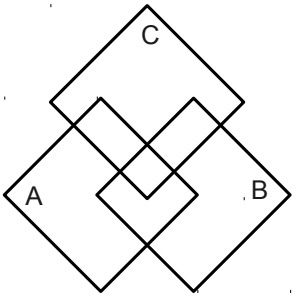

If we highlight the set $A \cap B$, we get this region of the diagram:

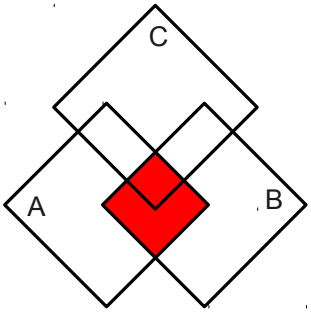

Given this, we can see that $C - (A \cap B)$ corresponds to this region:

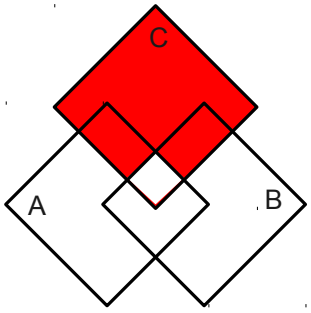

Now, let's take a look at $(C - A) \cup (C - B)$. If we highlight $C - A$ and $C - B$ separately, we get these regions:

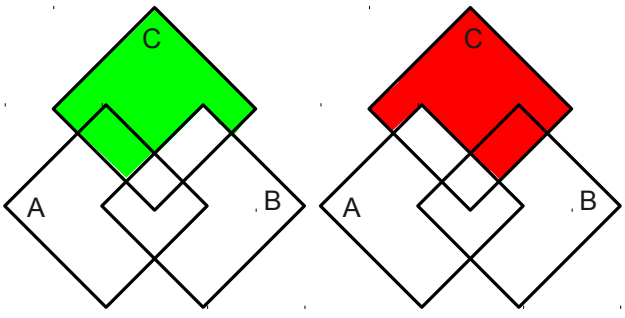

* I'm using diamonds instead of circles here because my drawing program (LibreOffice) makes it tricky to fill in circular regions. If you have any suggestions on how to draw better Venn diagrams, please let me know!

Consequently, their union $(C - A) \cup (C - B)$ is this region here:

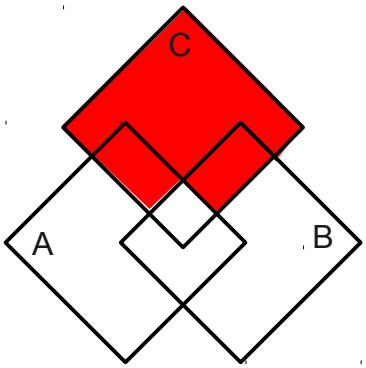

Now it's a bit clearer why this result should be true – the two sets are actually equal to one another! Moreover, it's easier to see why. To build up the sets $C - (A \cap B)$, we can construct the sets $C - A$ and $C - B$, then combine them together.

That said, the above picture isn't really a mathematical proof in the conventional sense. We still need to write out the proof longhand. To do this, we'll try to translate the above pictorial intuition into words. Specifically, we can work as follows. If we take any element of $C - (A \cap B)$, then (as you can see above) it belongs to at least one of $C - A$ or $C - B$. We can therefore write a proof by cases and show that regardless of which of these two sets our element belongs to, we know that the element must belong to $(C - A) \cup (C - B)$.

This is formalized below:

Proof: Consider any sets A, B, and C. We will show $C - (A \cap B) \subseteq (C - A) \cup (C - B)$. By definition, this is true if for any $x \in C - (A \cap B)$, we also have $x \in (C - A) \cup (C - B)$. So consider any $x \in C - (A \cap B)$. By the definition of set difference, this means that $x \in C$ and $x \notin A \cap B$. Since $x \notin A \cap B$, we know that it is not the case that both $x \in A$ and $x \in B$. Consequently, it must be true that either $x \notin A$ or $x \notin B$. We consider these two cases individually:

Case 1: $x \notin A$. Since we know that $x \in C$ and $x \notin A$, we know that $x \in C - A$. By our earlier result, we therefore have that $x \in (C - A) \cup (C - B)$.

Case 2: $x \notin B$. Since we know that $x \in C$ and $x \notin B$, we know that $x \in C - B$. By our earlier result, we therefore have that $x \in (C - A) \cup (C - B)$.

In either case we have that $x \in (C - A) \cup (C - B)$. Since our choice of x was arbitrary, we have that $C - (A \cap B) \subseteq (C - A) \cup (C - B)$ as required. ■

Notice that in the course of this proof, we ended up referring back to the proof we did above in which we claimed that for any sets A and B, $A \subseteq A \cup B$. Using this theorem, we were able to conclude that if $x \in C - A$, then $x \in (C - A) \cup (C - B)$. This is extremely common in mathematics. We begin with a few simple terms and definitions, then build up progressively more elaborate results from simpler ones. Most major results do not work from first principles, but instead build off of earlier work by combining well-known results and clever insights.

2.2.3 Lemmas

Let's think about the simple result that $A \subseteq A \cup B$. In itself, this isn't very surprising. The proof is simple and straightforward, and in the end we don't end up with anything particularly complex. However, as you saw above, this simple result can be used as a building block for proving more elaborate results.

A result that is primarily used as a small piece in a larger proof is sometimes called a **lemma**. Lemmas are distinguished from theorems primarily by how they're used. Some lemmas, such as the pumping lemma (which you'll learn more about later) are actually quite impressive results on their own, but are mostly used as a step in more complex proofs. Other lemmas, like the one you just saw, are simple but necessary as a starting point for future work.

When proving results about sets, lemmas like $A \subseteq A \cup B$ are often useful in simplifying more complex proofs. In fact, many seemingly obvious results about sets are best proven as lemmas so that we can use them later on.

The first lemma that we'll actually treat as such is the following result, which helps us prove that two sets are equal to one another:

> ***Lemma:*** If A and B are sets, then $A = B$ if and only if $A \subseteq B$ and $B \subseteq A$.

Note the use of the phrase ***if and only if*** in this lemma. The phrase "P if and only if Q" means that whenever P is true, Q is true, and whenever Q is true, P is true. In other words, "P if and only if Q" means that P and Q have the same truth value – either both P and Q are true, or both P and Q are false. The statement "if and only if" is a very strong assertion – it says that any time we'd like to speak about whether P is true or false, we can instead speak of whether Q is true or false.

As long as we're on the subject, you sometimes see the word ***iff*** used to mean "if and only if." This is a term that we'll use throughout this text, as it's widely used in the mathematical world. Consequently, we might rewrite the above lemma as

> ***Lemma:*** If A and B are sets, then $A = B$ iff $A \subseteq B$ and $B \subseteq A$.

Note that "iff" is read aloud as "if and only if." That way, we don't need to try to differentiate between "if" and "iff" by listening to how long the speaker draws out the final "f." This means that the above lemma would still be read aloud as "for any sets A and B, A equals B if and only if A is a subset of B and B is a subset of A."

Now, let's get down to business – what does this lemma say? The above lemma tells us that two sets A and B are equal to one another if and only if (in other words, precisely when) A is a subset of B and vice-versa. Recall that two sets are equal when they have exactly the same elements; it doesn't matter how we describe or construct the sets, just that they have the same elements. The above lemma states that if we want to show that two sets are equal, all we need to do is show that all of the elements of one set are contained in the other and vice-versa.

So how exactly do we go about proving this lemma? So far, all of the proofs that we've seen have taken the form "if P, then Q." If we want to prove a statement of the form "P iff Q," then we need to prove two things – first, if P is true, then Q is true; second, if Q is true, then P is true as well. In other words, both P and Q imply one another.

Given this setup, here is one proof of this result:

Proof: We prove both directions of implication.

($\Rightarrow$) First, we show that, for any sets A and B, if $A = B$, then $A \subseteq B$ and $B \subseteq A$. If $A = B$, consider any $x \in A$. Since $A = B$, this means that $x \in B$. Since our choice of x was arbitrary, any $x \in A$ satisfies $x \in B$, so $A \subseteq B$. Similarly, consider any $x \in B$, then since $A = B$, $x \in A$ as well. Since our choice of x was arbitrary, any $x \in B$ satisfies $x \in A$, so $B \subseteq A$.

($\Leftarrow$) Next, we prove that if $A \subseteq B$ and $B \subseteq A$, then $A = B$. Consider any two sets A and B where $A \subseteq B$ and $B \subseteq A$. We need to prove that $A = B$. Since $A \subseteq B$, for any $x \in A$, $x \in B$ as well. Since $B \subseteq A$, for any $x \in B$, $x \in A$ as well. Thus every element of A is in B and vice-versa, so the two sets have the same elements. ■

This proof looks different from the ones we've seen so far both structurally and notationally. This proof is essentially two separate proofs that together prove a larger result; the first half proves that if two sets are equal each is a subset of the other, and the second half proves that if two sets are subsets of one another they are equal. This is because in order to prove the biconditional, we need to prove two independent results, which together combine to prove the biconditional. Within each piece of the proof, notice that the structure is similar to before. We call back to the definitions of subset and set equality in order to reason about how the elements of the sets are related to one another.

You probably noticed the use of ($\Rightarrow$) and ($\Leftarrow$) here. In proofs involving biconditionals, it's common to split the proof up into two logically separate sections, one for the forward direction and one for the reverse direction. To demarcate the sections, we often use the symbols ($\Rightarrow$) and ($\Leftarrow$). The ($\Rightarrow$) part of the proof denotes the forward direction: when proving A iff B, this is the "if A, then B" part. Similarly, the ($\Leftarrow$) part of the proof denotes the "if B, then A" part. As a courtesy to the proof reader, it usually is a good idea to briefly restate what it is that you're going to be proving before you jump into the proof.

Now that we have this lemma, let's go and use it to prove some Fun and Exciting Facts about set equality! Let's begin with a simple result that teaches something about how symmetric difference works:

Theorem: For any sets A and B, $(A \cup B) - (A \cap B) = A \Delta B$.

Intuitively, this says that we can construct the symmetric difference of A and B (that is, the set of elements that are either in A or B, but not both) as follows. First, combine the two sets A and B together into the larger set $A \cup B$. Next, take out from that set all of the elements that are in the intersection of A and B. The remaining elements form the set $A \Delta B$.

To prove this result, we can use our lemma from above, which says that two sets are equal iff each is a subset of the other. The structure of our proof will thus be as follows – we'll show that each set is a subset of the other, then we'll use the previous lemma to conclude the proof.

Let's begin by showing that $(A \cup B) - (A \cap B) \subseteq A \Delta B$. Since this acts as a stepping stone toward the larger proof, we'll pose it as a lemma.

Lemma 1: $(A \cup B) - (A \cap B) \subseteq A \Delta B$.

How might we prove this lemma? To do so, we'll just call back to the definitions of union, intersection, difference, and symmetric difference:

Proof of Lemma 1: We will show that for any $x \in (A \cup B) - (A \cap B)$, $x \in A \Delta B$. So consider any $x \in (A \cup B) - (A \cap B)$. This means that $x \in A \cup B$, but $x \notin A \cap B$. Since $x \in A \cup B$, we know that $x \in A$ or $x \in B$. Since $x \notin A \cap B$, we know that x is not contained in both A and B. We thus have that x is in at least one of A and B, but not both. Consequently, $x \in A \Delta B$. Since our choice of x was arbitrary, we therefore have that $(A \cup B) - (A \cap B) \subseteq A \Delta B$. ■

The other direction also will be a lemma for the same reasons. Here's the lemma and the proof:

Lemma 2: $A \Delta B \subseteq (A \cup B) - (A \cap B)$

This proof is a little bit more involved because there are two completely separate cases to consider when dealing with elements of $A \Delta B$. The proof is below:

Proof of Lemma 2: We will show that for any $x \in A \Delta B$, $x \in (A \cup B) - (A \cap B)$. Consider any $x \in A \Delta B$. Then either $x \in A$ and $x \notin B$, or $x \in B$ and $x \notin A$. We consider these cases separately:

Case 1: $x \in A$ and $x \notin B$. Since $x \in A$, $x \in A \cup B$. Since $x \notin B$, $x \notin A \cap B$. Consequently, $x \in (A \cup B) - (A \cap B)$.

Case 2: $x \in B$ and $x \notin A$. Since $x \in B$, $x \in A \cup B$. Since $x \notin A$, $x \notin A \cap B$. Consequently, $x \in (A \cup B) - (A \cap B)$.

In either case, $x \in (A \cup B) - (A \cap B)$. Since our choice of x was arbitrary, we have that $A \Delta B \subseteq (A \cup B) - (A \cap B)$. ■

Now that we have these two lemmas, the proof of the general result is surprisingly straightforward:

Proof of Theorem: By Lemma 1, $(A \cup B) - (A \cap B) \subseteq A \Delta B$. By Lemma 2, $A \Delta B \subseteq (A \cup B) - (A \cap B)$. Since each set is a subset of the other, by our earlier lemma we have that $(A \cup B) - (A \cap B) = A \Delta B$. ■

That's all that we have to show!

Before we move on to show more applications of the lemma, let's take a minute to examine the proof of Lemma 2. I've reprinted it below:

Proof of Lemma 2: We will show that for any $x \in A \Delta B$, $x \in (A \cup B) - (A \cap B)$. Consider any $x \in A \Delta B$. Then either $x \in A$ and $x \notin B$, or $x \in B$ and $x \notin A$. We consider these cases separately:

Case 1: $x \in A$ and $x \notin B$. Since $x \in A$, $x \in A \cup B$. Since $x \notin B$, $x \notin A \cap B$. Consequently, $x \in (A \cup B) - (A \cap B)$.

Case 2: $x \in B$ and $x \notin A$. Since $x \in B$, $x \in A \cup B$. Since $x \notin A$, $x \notin A \cap B$. Consequently, $x \in (A \cup B) - (A \cap B)$.

In either case, $x \in (A \cup B) - (A \cap B)$. Since our choice of x was arbitrary, we have that $A \Delta B \subseteq (A \cup B) - (A \cap B)$. ■

Notice the similarity between Case 1 and Case 2. These two cases are virtually identical, except that we've interchanged the role of the sets A and B. If you'll notice, there really isn't anything in the above proof to suggest that set A is somehow "more important" than set B. If we interchange set A and set B, we change the sets $(A \cup B) - (A \cap B)$ and $A \Delta B$ to the sets $(B \cup A) - (B \cap A)$ and $B \Delta A$. But these are exactly the sets we started with! In a sense, because there really isn't an appreciable difference between A and B, it seems silly to have two completely different cases dealing with which sets x is contained in.

This situation – in which multiple parts of a proof end up being surprisingly similar to one another – is fairly common, and mathematicians have invented some shorthand to address it. Mathematicians often write proofs like this one:

Proof of Lemma 2: We will show that for any $x \in A \Delta B$, $x \in (A \cup B) - (A \cap B)$. Consider any $x \in A \Delta B$. Then either $x \in A$ and $x \notin B$, or $x \in B$ and $x \notin A$. Assume without loss of generality that $x \in A$ and $x \notin B$. Since $x \in A$, $x \in A \cup B$. Since $x \notin B$, $x \notin A \cap B$, *so* $x \in (A \cup B) - (A \cap B)$. Since our choice of x was arbitrary, we have that $A \Delta B \subseteq (A \cup B) - (A \cap B)$. ■

Notice the use of the phrase "without loss of generality." This phrase indicates in a proof that there are several different cases that need to be considered, but all of them are identical to one another once we change the names around appropriately. If you are writing a proof where you find multiple cases that seem identical to one another, feel free to use this phrase to write the proof just once. That said, be careful not to claim that you haven't lost generality if the cases are actually different from one another!

As another example of a proof using "without loss of generality," let's consider the following theorem, which has nothing to do with sets:

Theorem: If m and n have opposite parity, $m + n$ is odd.

We can check this pretty easily – 3 + 4 = 7, which is odd, 137 + 42 =179, which is odd, etc. How might we prove this? Well, there are two cases to consider – either m is even and n is odd, or m is odd and n is even. But these two cases are pretty much identical to one another, since $m + n = n + m$ and it doesn't really matter whether it's m or n that's odd. Using this, let's write a quick proof of the above result:

Proof: Without loss of generality, assume that m is odd and n is even. Since m is odd, there exists an integer r such that $m = 2r + 1$. Since n is even, there exists an integer s such that $n = 2s$. Then $m + n = 2r + 1 + 2s = 2(r + s) + 1$. Consequently, $m + n$ is odd. ■

This proof is about half as long as it would be otherwise.

2.2.4 Proofs with Vacuous Truths

To see if we can get some more mileage out of our lemma about set equality, let's try proving some more results about sets. Let's consider the following result:

Theorem: For any sets A and B, if $A \subseteq B$, then $A - B = \varnothing$.

Now, how might we prove this? Right now, the main tool at our disposal for proving two sets are equal is to show that those two sets are subsets of one another. In other words, to prove the above result, we might try proving two lemmas:

Lemma 1: For any sets A and B, if $A \subseteq B$, then $\varnothing \subseteq A - B$.
Lemma 2: For any sets A and B, if $A \subseteq B$, then $A - B \subseteq \varnothing$.

Okay, let's set out to prove them. Let's begin by trying to prove lemma 1. To do this, we need to show that every element of the empty set is also contained in $A - B$. But wait a minute – this doesn't make any sense, since there aren't any $x \in \varnothing$! But not to worry. If you'll recall from Chapter 1, we introduced the idea of a vacuous truth, a statement that is true because it doesn't apply to anything. Fortunately, that's exactly what we have right here – there aren't any elements of the empty set, so it's vacuously true that every element of the empty set is also contained in $A - B$, regardless of what A and B actually are. After all, it's also true that every element of the empty set is made of fire, that every element of the empty set is your best friend,* etc.

How do we formalize this in a proof? Well, we can just say that it's vacuously true! This is shown here:

Proof of Lemma 1: We need to show that every element $x \in \varnothing$ also satisfies $x \in A - B$. But this is vacuously true, as there are no x satisfying $x \in \varnothing$. ■

Well, that was surprisingly straightforward. On to the second lemma!

At first glance, this statement doesn't seem to make any sense. There are no elements of the empty set, so how could something be a subset of the empty set? This would only happen if there are no elements in the first set, since if there were some element $x \in A - B$, then it would have to be true that $x \in \varnothing$, which we know to be impossible. This actually gives us a hint about how to approach the problem. We know that we shouldn't be able to find any $x \in A - B$, so one route for proving that $A - B \subseteq \varnothing$ is to directly show that the statement "for any $x \in A - B$, $x \in \varnothing$" is vacuously true. This is shown below:

* Saying "every element of the empty set is your best friend" is not the same as saying "the set of your best friends is the empty set." The former is a vacuous truth. The latter is a mathematical insult.

Proof of Lemma 2: We need to show that any $x \in A - B$ also satisfies $x \in \varnothing$. Consider any $x \in A - B$. This means that $x \in A$ and $x \notin B$. Since $A \subseteq B$ and since $x \in A$, we know that $x \in B$. But this means simultaneously that $x \in B$ and $x \notin B$. Consequently, there are no $x \in A - B$, so the claim that any $x \in A - B$ also satisfies $x \in \varnothing$ is vacuously true. ■

Notice the structure of the proof. We begin by using definitions to tease apart what it means for an element to be in $A - B$, then show that, in fact, no elements can be in this set. We conclude, therefore, that the entire lemma must be vacuously true.

We can use these two lemmas to complete the proof:

Proof of Theorem: Consider any sets A and B such that $A \subseteq B$. By Lemma 1, we have that $\varnothing \subseteq A - B$. By Lemma 2, we have that $A - B \subseteq \varnothing$. Thus by our earlier lemma, $A - B = \varnothing$ as required. ■

2.3 Indirect Proofs

The proofs that we have done so far have directly shown that a particular statement must be true. We begin with a set of assumptions, then manipulate those assumptions to arrive at a desired conclusion. However, there is an entirely different family of proof techniques called ***indirect proofs*** that indirectly prove that some proposition must be true.

This may seem a bit strange at first, but there are many familiar analogs in real life. For example, suppose that you're biking to class and can't remember whether or not you brought your keys with you. You could directly prove whether you have your keys on you by stopping, getting off your bike, and checking your pockets or purse for your keys. But alternatively, you could use the following line of reasoning. Assuming that you lock your bike (which you should!), you couldn't have unlocked your bike in the first place if you didn't have your keys. Since you definitely unlocked your bike – after all, you're riding it! – you must have your keys with you. You didn't explicitly check to see that you have your keys, but you can be confident that you do indeed have them with you.

In this section, we'll build up two indirect proof techniques – *proof by contradiction*, which shows that a proposition has to be true because it can't be false, and *proof by contrapositive*, which proves that P implies Q by proving that an entirely different connection holds between P and Q.

2.3.1 Logical Implication

Before we can move on to talk about proofs by contradiction and contrapositive, we need to discuss logical implication. Many of the proofs that we have done so far are proofs of the form

If P, then Q.

For example, we have proven the following:

If x is even, then x^2 is even.

If m is even and n is odd, then mn is even.

If m and n have the same parity, then $m + n$ is even.

If n is even and m is an integer, then $n + m$ has the same parity as m.

If $A \subseteq B$, then $A - B = \emptyset$.

In structuring each of these proofs, the general format has been as follows: first, we assume that P is true, then we show that given this assumption Q must be true as well. To understand why this style of proof works in the first place, we need to understand what the statement "If P, then Q" means. Specifically, the statement "If P, then Q" means that any time P is true, Q is true as well. For example, consider the statement

If $x \in A$, then $x \in A \cup B$.

This statement says that any time that we find that x is contained in the set A, it will also be contained in the set $A \cup B$. If $x \notin A$, this statement doesn't tell us anything. It's still possible for $x \in A \cup B$ to be true, namely if $x \in B$, but we don't have any guarantees.

Let's try this statement:

If I pet the fuzzy kitty, I will be happy.

This tells us that in the scenario where I pet the fuzzy kitty, it's true that I will be happy. This doesn't say anything at all about what happens if I don't pet the kitty. I still might be happy (perhaps I petted a cute puppy, or perhaps Stanford just won another football game).

The general pattern here is that a statement of the form

If P, then Q.

only provides information if P is true. If P is true, we can immediately conclude that Q must be true. If P is false, Q could be true and could be false. We don't have any extra information.

An important point to note here is that implication deals purely with how the truth or falsity of P and Q are connected, not whether or not there is a causal link between the two. For example, consider this (silly) statement:

If I will it to be true, then $1 + 1 = 2$.

Intuitively, this statement is false: $1 + 1 = 2$ because of the laws of mathematics, not because I consciously wish that it is! But mathematically, the statement is true. If I want $1 + 1 = 2$ to be true, you will indeed find that $1 + 1 = 2$. You'll find $1 + 1 = 2$ regardless of whether or not I want it to be. Consequently, the statement "If I will it to be true, $1 + 1 = 2$" is always true.

Why discuss these (seemingly pedantic) details at all? The reason for this is to make clear what exactly it means for an implication to be true so that we can discuss what it means for an implication to be false. The statement "If P, then Q" is true if whenever we find that P is true, we also find that Q is true. In order for the statement "If P, then Q" to be false, we have to find an example where P is true (meaning that we expect Q to be true as well), but to our surprise found that Q actually is false. For example, if we wanted to disprove the claim

If $x + y$ is even, then x is odd.

we would have to find an example where $x + y$ was even, but x was not odd. For example, we can take $x = 2$ and $y = 2$ as a counterexample, since $x + y = 4$, but x is not odd. However, if we were to take something like $x = 3$ and $y = 2$, it would not be a counterexample: $3 + 2$ is not even, so the above claim says nothing about what's supposed to happen.

It's important to make this distinction, because it's surprisingly easy to think that you have disproven an implication that's perfectly true. For example, consider the statement

If $A \subseteq B$, then $A - B = \emptyset$

What happens if we take the sets A = {1, 2} and B = {3}? Then the statement $A \subseteq B$ is false, as is the statement $A - B = \varnothing$. However, we have not contradicted the above statement! The above statement only tells us something about what happens when $A \subseteq B$, and since A isn't a subset of B here, the fact that $A - B \neq \varnothing$ doesn't matter.

2.3.2 Proof by Contradiction

One of the most powerful tools in any mathematician's toolbox is proof by contradiction. A proof by contradiction is based on the following logical idea: If a statement cannot possibly be false, then it has to be true.

In a proof by contradiction, we prove some proposition P by doing the following:

1. Assume, hypothetically, that P is **not** true. This is the opposite of what we want to prove, and so we want to show that this assumption couldn't possibly have been correct.
2. Using the assumption that P is false, arrive at a *contradiction* – a statement that is logically impossible.
3. Conclude that, since our logic was good, the only possible mistake we could have made would be in assuming that P is not true. Therefore, P absolutely *must* be true.

Let's see an example of this in action. Earlier, we proved the result that if n is even, then n^2 must be even as well. It turns out that the converse of this is true as well:

> ***Theorem:*** If n^2 is even, then n is even.

Empirically, this seems to pan out. 36 is even, and $36 = 6^2$, with 6 even. 0 is even, and $0 = 0^2$, with 0 even as well. But how would we actually prove this? It turns out that this is an excellent use case for a proof by contradiction.

To prove this statement by contradiction, let's assume that it's false, which means that the statement "If n^2 is even, then n is even" is incorrect. As we just saw, this would have to mean that n^2 is even, but n itself is odd. Is this actually possible?

The answer is no – if n were odd, then n^2 would have to be odd as well. However, one of the assumptions we made was that n^2 is even. This contradiction tells us that something is wrong here. The only thing questionable we did was making the assumption that n is odd with n^2 even. Consequently, we know that this combination must be impossible. Therefore, if n^2 is even, we know that n is even as well.

We can formalize this in a proof as follows:

> ***Proof:*** By contradiction; assume that n^2 is even but that n is odd. Since n is odd, $n = 2k + 1$ for some integer k. Therefore $n^2 = (2k + 1)^2 = 4k^2 + 4k + 1 = 2(2k^2 + 2k) + 1$. This means that n^2 is odd, contradicting the fact that we know that n^2 is even. We have reached a contradiction, so our assumption must have been wrong. Therefore, if n^2 is even, n must be even. ■

Let's look at this proof in more depth. First, note how it starts off:

> By contradiction; assume that n^2 is even but that n is odd.

This sets up how we are going to approach the proof. We state explicitly that we are going to attempt a proof by contradiction. We immediately then say what assumption we are going to make. Here, since we want to contradict the statement "If n^2 is even, n is even," we say that the contradiction is that n^2 is even, but n is odd.

Once we have set up the proof by contradiction, the remainder of our proof is a quest to show that this assumption has to have been wrong by deriving a contradiction. The middle section of the proof does just that – it arrives at the conclusion that n^2 has to be both odd and even at the same time.

Now that we have our contradiction, we can finish the proof by stating that this contradiction means that we're done:

> We have reached a contradiction, so our assumption must have been wrong. Therefore, if n^2 is even, n must be even. ■

All proofs by contradiction should end this way. Now that you have the contradiction, explain how it means that the initial assumption was wrong, and from there how this proves the overall result.

Proof by contradiction is a powerful tool. We saw this used in Cantor's theorem in the last chapter (though, admittedly, we haven't seen the formal proof yet), and you will see it used later to prove that several specific important problems cannot be solved by a computer. For now, let's build up some other small examples of how this proof technique can be used.

One interesting application of proofs by contradiction is to show that some particular task cannot be accomplished. Consider the following problem:

> You have 2,718 balls and five bins. Prove that you cannot distribute all of the balls into the bins such that each bin contains an odd number of balls.

This problem seems hard – there are a *lot* of ways to distribute those balls into the bins, though as you'll see there's no way to do it such that every bin has an odd number of balls in it. How might we show that this task is impossible? Using the idea of a proof by contradiction, let's start off by hypothetically assuming that you *can* indeed solve this. Could we then show that this solution leads to some sort of contradiction? Indeed we can. Think of it this way – if we have an odd number of balls in the five bins, then the total number of balls placed into those bins would have to be equal to the sum of five odd numbers. What numbers can you make this way? Well, if we add up two odd numbers, we get an even number (because we know that the sum of two numbers with the same parity is even). If we add up two more of the odd numbers, we get another even number. The sum of those two even numbers is even. If we then add in the last odd number to this even number, we get an odd total number of balls. This is extremely suspicious. We know that the total number of balls has to be odd, because we just proved that it has to. At the same time, we know that there are 2,718 balls distributed total. But this would imply that 2,718 is odd, which it most certainly is not! This is a contradiction, so something we did must have been wrong. Specifically, it has to have been our assumption that we can distribute all of the balls such that each bin has an odd number of balls in it. Therefore, there can't be a solution.

This argument is formalized below as a proof:

Proof: By contradiction; assume that there is a way to distribute all 2,718 balls into five bins such that each bin has an odd number of balls in it. Consider any such way of distributing the balls, and let the number of balls in the five bins be a, b, c, d, and e. Write the sum $a + b + c + d + e$ as $((a + b) + (c + d)) + e$. Since all five numbers have the same parity, both $(a + b)$ and $(c + d)$ are even. Since $(a + b)$ and $(c + d)$ have the same parity, $((a + b) + (c + d))$ must be even. Then, since $((a + b) + (c + d))$ is even, the sum $((a + b) + (c + d)) + e$ must have the same parity as e. Since e is odd, this means that sum of the number of balls in the five bins is odd, contradicting the fact that there are an even number of balls distributed across the bins (2,718). We have reached a contradiction, so our initial assumption must have been wrong and there is no way to distribute 2,718 balls into five bins such that each bin has an odd number of balls. ■

As an aside, I absolutely love this proof. It pulls together our discussion of direct proofs with parities along with proof by contradiction.

Before we move on, though, let's examine the structure of this proof one more time. Note that it has the same shape as the previous proof. We begin by stating that the proof is by contradiction and what that contradiction is. We then derive a contradiction, and conclude by saying that the contradiction proves the original theorem.

Here is yet another example of a classic proof by contradiction. Consider a standard 8 × 8 chessboard:

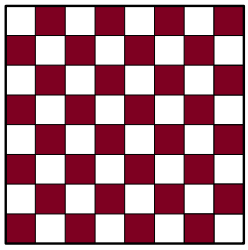

Now, suppose that we cut off two diagonally opposite corners, as shown here:

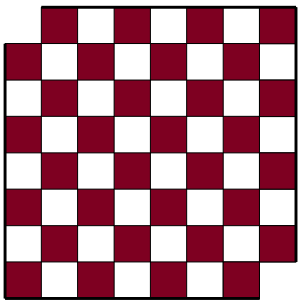

Suppose that we want to cover this chessboard with a set of 2 × 1 dominoes. These dominoes can be positioned horizontally or vertically, but never diagonally. Additionally, we cannot stack the dominoes on top of one another. The question is this – is it possible to cover every square on the modified chessboard with dominoes? Interestingly, the answer is no. It's impossible to do so.

So why is that? Well, let's approach this from the perspective of a proof by contradiction. Suppose, hypothetically, that we can cover the chessboard with dominoes. Since each domino covers two horizontally or vertically adjacent squares, we know for a fact that each domino covers exactly one white square and exactly one black square. Moreover, since no two dominoes can stack atop one another, if we add up the total number of white squares covered by each domino and the total number of black squares covered by each

domino, we should get the total number of white and black squares on the chessboard. But this is where we run into trouble. If each domino covers one white square and one black square, then the total number of white squares and black squares covered should have to be the same. Unfortunately, this isn't true. A standard chessboard has the same number of white and black squares. When we removed two opposite corners, we took away two white squares (check the picture above). This means that there are, in fact, two more black squares than white squares, contradicting the fact that we were supposed to have the same number of white squares and black squares. This means (again!) that our assumption was wrong, and that there must be no solution to this puzzle.

Formalized as a proof, the above argument looks like this:

Theorem: There is no way to tile an 8×8 chessboard missing two opposite corners with dominoes such that each domino is aligned horizontally or vertically and no two dominoes overlap.

Proof: By contradiction; assume that such a tiling exists. Since each domino is aligned horizontally or vertically across two tiles, each domino covers the same number of white and black squares. Since no two dominoes overlap, each square is covered by exactly one domino. Consequently, the number of white squares on the chessboard and the number of black squares on the chessboard should equal the number of dominoes. In turn, this means that the number of white squares and black squares on the chessboard must be equal. But this is impossible – there are 30 white squares and 32 black squares, and $30 \neq 32$. We have reached a contradiction, so our assumption must have been incorrect. Thus there is no solution to the puzzle. ■

2.3.3 Rational and Irrational Numbers

In computer science we commonly work with the natural numbers or integers because our computers are digital. However, the real numbers are quite important in mathematics, and it would be a disservice to them if we didn't spend at least a little time exploring their properties.

To begin with, we should make a distinction between two different types of real numbers – the *rational numbers* and the *irrational numbers*. Intuitively, rational numbers are real numbers that can be expressed as the ratio of two integers. For example, any integer is rational, because the integer x is the ratio $x / 1$. Numbers like $^{7}/_{4}$ and $^{137}/_{42}$ are also rational. Formally, we define the rational numbers as follows:

A real number r is called ***rational*** if there exist integers p and q such that $q \neq 0$ and $p / q = r$.

Let's take a minute to see what this says. We're going to say that a number r is rational if there is some way that we can find two integers p and q where q isn't zero and $p / q = r$. We're going to require that $q \neq 0$ so that we can safely use it as the denominator in the fraction. Note that p can be zero, because we'd like $0 = {}^{0}/_{1}$ to count as a rational number.

An important property of this definition is that we say r is rational if there is some way to write $r = p / q$ for integers p and q. Note that there can actually be *many* different ways to do this. As an example, we can write $2 = {}^{2}/_{1}$, or $2 = {}^{-2}/_{-1}$, or $2 = {}^{6}/_{3}$, etc. For reasons that will become clearer in a moment, we often use the following property of rational numbers:

Any rational number r can be written as $r = p / q$, where p and q are integers, $q \neq 0$, and p and q have no common factors other than ±1.

This statement essentially states that if r is rational, we can write out r in "simplest form" by writing out r as a fraction that cannot be simplified. For example, when we write $1.5 = {}^{6}/_{4}$, we can simplify the fraction down to $1.5 = {}^{3}/_{2}$ because 6 and 4 have 2 as a common factor. However, we can't simplify ${}^{3}/_{2}$, because the only common factors of 3 and 2 are ±1. Note that we could also write $1.5 = {}^{-3}/_{-2}$ and say that it is in "simplest form" because the only common factors of -3 and -2 are ±1. This is certainly a less "pretty" fraction, though according to the above statement we'll consider it to be in simplest form.

One more definition is in order:

The set $\{ r \mid r \in \mathbb{R}$ and r is rational $\}$, the ***set of all rational numbers***, is denoted $\mathbb{Q}$.

From the definition of $\mathbb{Q}$, it's clear that $\mathbb{Q} \subseteq \mathbb{R}$. However, is it true that $\mathbb{Q} = \mathbb{R}$? That is, is every real number rational? It turns out that the answer to this question is "no." There are many ways to show this using advanced mathematics, but one simple solution is to find an explicit example of an irrational number. It's not all that hard to find an example of an irrational number – numbers like e and π are irrational, for example – but to actually prove that these numbers are irrational is surprisingly difficult. Instead, we'll focus on a simple example of a number known to be irrational: $\sqrt{2}$.

Let's go prove the following theorem, which is a beautiful example of a proof by contradiction:

Theorem: $\sqrt{2}$ is irrational.

How exactly can we show this? As you might have guessed from where we are right now, this is a good spot for a proof by contradiction. Let's suppose, for the sake of contradiction, that $\sqrt{2}$ actually is rational. This means that we can find integers p and q such that $q \neq 0$, $p / q = \sqrt{2}$, and p and q have no common factors other than 1 and -1 (that is, they're the "simplest" such p and q that we can use). What to do next? Well, ultimately we're going to try to derive some sort of contradiction. Right now, though, it's not clear what exactly that will be. That's fine, though. Let's just explore around a bit and see if we find anything interesting. If we do, great! We're done. If not, we can back up and try something else.

Let's start off with some simple algebraic manipulations. Since we have that

$$p / q = \sqrt{2}$$

We can square both sides to get

$$p^2 / q^2 = 2$$

If we then multiply both sides by q^2, we get

$$p^2 = 2q^2.$$

What does this tell us? For one thing, we know that p^2 has to be an even number, since q^2 is an integer and p^2 is twice q^2. But if you'll recall, one of the first proofs we did by contradiction was the proof that if n^2 is

even, then n must be even as well. Since p^2 is even, this means that p has to be even as well. This tells us that $p = 2k$ for some integer k.

We've just shown that if $p / q = \sqrt{2}$, then p has to be even. What can we do with this? Looking above, we've shown that $p^2 = 2q^2$. What happens if we plug in $2k$ in place of p? This gives us

$$(2k)^2 = 2q^2$$
$$4k^2 = 2q^2$$
$$2k^2 = q^2$$

This last line tells us that q^2 has to be even as well, since it's twice k^2 and k^2 is an integer. It's at this point that we can see that something unusual is up. Using our previous result, since q^2 is even, q has to be even as well. But then both p and q are even, which means that they have to be divisible by two – contradicting the fact that p and q can't have any divisors other than 1 and -1!

In short, our proof worked as follows. Starting with $p / q = \sqrt{2}$, we showed that p had to be even. Since p was even, q had to be even as well, meaning that p and q weren't simplified as far as possible. In fact, there's no possible way for them to be simplified – we've shown that whatever choice of p and q you make, they can always be simplified further. This contradicts Rule 3 of rational numbers, and so $\sqrt{2}$ has to be irrational. This logic is formalized here in this proof:

> ***Proof:*** By contradiction; assume that $\sqrt{2}$ is rational. Then there exist integers p and q such that $q \neq 0$, $p / q = \sqrt{2}$, and p and q have no common divisors other than 1 and -1.
>
> Since $p / q = \sqrt{2}$, this means that $p^2 / q^2 = 2$, which means that $p^2 = 2q^2$. This means that p^2 is even, so by our earlier result p must be even as well. Consequently, there exists some integer k such that $p = 2k$.
>
> Since $p = 2k$, we have that $2q^2 = p^2 = (2k)^2 = 4k^2$, so $q^2 = 2k^2$. This means that q^2 is even, so by our earlier result q must be even as well. But this is impossible, because it means that p and q have 2 as a common divisor, contradicting the fact that p and q have no common divisors other than 1 and -1.
>
> We have reached a contradiction, so our assumption must have been incorrect. Thus $\sqrt{2}$ is irrational. ■

We now have our first example of a number that we know is not rational. This alone is enough to prove that $\mathbb{Q} \neq \mathbb{R}$. However, is $\sqrt{2}$ the only irrational number? Or are there more irrational numbers like it? It turns out that a great many numbers are irrational; in fact, there are infinitely more irrational numbers than rational numbers! We'll prove this later on in Chapter 6 when we discuss the nature of infinity.

2.3.4 Proof by Contrapositive

There is one final indirect proof technique that we will address right now – proof by contrapositive.

To motivate a proof by contrapositive, let's return to our discussion of mathematical implication. Consider the following statement:

If I close the windows, the velociraptors can't get inside.

This statement says that whenever we know that the windows are closed, we know that the velociraptors won't be able to get inside. Now, let's suppose that we know that, unfortunately, the velociraptors did indeed get inside. What could we conclude from this? We know that I certainly didn't close the windows – if I had closed the window, then the raptors wouldn't be inside in the first place!

Let's try another example. Suppose that we know that

$$\text{If } A \subseteq B, \text{ then } A - B = \varnothing.$$

Suppose we find two sets A and B such that $A - B \neq \varnothing$. What can we conclude? Here, we can say that A is not a subset of B, because if it were, then $A - B$ would have been equal to $\varnothing$.

There seems to be a pattern here. It seems like if we know that the statement "If P, then Q" is true and we know that Q is false, then we know that P must be false as well. In fact, that's exactly correct. Intuitively, the rationale is that if P implies Q and Q is false, P couldn't be true, because otherwise Q would be true. Given any implication "If P, then Q," its **contrapositive** is the statement "If **not** Q, then **not** P." The contrapositive represents the above idea that if Q is false, P has to be false as well.

It's getting a bit tricky to use phrases like "If P, then Q" repeatedly through this text, so let's introduce a bit of notation. We will use the notation $P \to Q$ to mean that P implies Q; that is, if P, then Q. Given an implication $P \to Q$, the contrapositive is **not** $Q \to$ **not** P.

The contrapositive is immensely useful because of the following result:

> ***Theorem:*** If **not** $Q \to$ **not** P, then $P \to Q$.

This theorem is very different from the sorts of proofs that we've done before in that we are proving a result about logic itself! That is, we're proving that if one implication holds, some other implication must hold as well! How might we go about proving this? Right now, we have two techniques at our disposal – we can proceed by a direct proof, or by contradiction. The logic we used above to justify the contrapositive in the first place was reminiscent of a proof by contradiction ("well, if Q is false, then P couldn't be true, since otherwise Q would have been true."). Accordingly, let's try to prove this theorem about the contrapositive by contradiction.

How might we do this? First, let's think about the contradiction of the above statement. Since we are contradicting an implication, we would assume that **not** $Q \to$ **not** P, but that $P \to Q$ is false. In turn we would ask: what does it mean for $P \to Q$ to be false? This would only be possible if P was true but Q was not. So at this point, we know the following:

1. **not** $Q \to$ **not** P.
2. P is true.
3. Q is false.

And now all of the pieces fall into place. Since Q is false, we know that **not** Q is true. Since **not** Q implies **not** P, this means that **not** P is true, which in turn tells us that P should be false. But this contradicts the fact that P is true. We've hit our contradiction, and can conclude, therefore, that if **not** $Q \to$ **not** P, then $P \to Q$.

Here is a formal proof of the above:

Proof: By contradiction; assume that **not** $Q \to$ **not** P, but that $P \to Q$ is false. Since $P \to Q$ is false, we know that P is true but Q is false. Since Q is false and **not** $Q \to$ **not** P, we have that P must be false. But this contradicts the fact that we know that P is true. We have reached a contradiction, so our initial assumption must have been false. Thus if **not** $Q \to$ **not** P, then $P \to Q$. ■

This proof has enormous importance for how we can prove implications. If we want to prove that $P \to Q$, we can always instead prove that **not** $Q \to$ **not** P. This then implies $P \to Q$ is true.

Let's work through an example of this. Earlier we proved the following result:

Theorem: If n^2 is even, then n is even.

Our proof proceeded by contradiction. What if we wanted to prove this result by contrapositive? Well, we want to show that if n^2 is even, then n is even. The contrapositive of this statement is that if n is not even, then n^2 is not even. More clearly, if n is odd, then n^2 is odd. If we can prove that this statement is true, then we will have successfully proven that if n^2 is even, then n is even. Such a proof is shown here:

Proof: By contrapositive; we prove that if n is odd, then n^2 is odd. Let n be any odd integer. Since n is odd, $n = 2k + 1$ for some integer k. Therefore, $n^2 = (2k + 1)^2 = 4k^2 + 4k + 1 = 2(2k^2 + 2k) + 1$. Thus n^2 is odd. ■

Notice the structure of the proof. As with a proof by contradiction, we begin by announcing that we're going to use a proof by contrapositive. We then state the contrapositive of the statement that we want to prove, both so that readers know what to expect and so that we're clear on what we want to show. From there, we proceed just as we would in a normal proof – we need to show that if n is odd, n^2 is odd, and so we assume that n is odd and proceed from there. The result is a remarkably clean and elegant proof.

Here's another example of a proof by contrapositive: suppose that we have 16 objects that we want to distribute into two bins. There are many ways that we might do this – we might split them evenly as an 8/8 split, or might put all of them into one bin to give a 16/0 split, or might have something only a bit lopsided, like a 10/6 split. Interestingly, though, notice that in each case we have at least one bin with at least 8 objects in it. Is this guaranteed to happen? Or is it just a coincidence?

It turns out that this isn't a coincidence, and in fact we can prove the following:

Theorem: If $m + n = 16$, then $m \geq 8$ or $n \geq 8$.

To prove this by contrapositive, we first need to figure out what the contrapositive of the above statement is. Right now, we have the following:

$$m + n = 16 \to m \geq 8 \text{ or } n \geq 8$$

The contrapositive of this statement is

$$\textbf{not } (m \geq 8 \text{ or } n \geq 8) \to \textbf{not } (m + n = 16)$$

Hmmm... that's not very easy to read. Perhaps we can simplify it. Let's start with the right-hand side. We can simplify **not** ($m + n = 16$) to the easier $m + n \neq 16$. This gives

$$\textbf{not } (m \geq 8 \text{ or } n \geq 8) \rightarrow m + n \neq 16$$

But what about the first part? This is a bit more subtle. What is the opposite of $m \geq 8$ or $n \geq 8$? Well, this statement is true if either $m \geq 8$ or $n \geq 8$, so for it to be false we need to ensure that both $m \geq 8$ and $n \geq 8$ are false. This would be true if $m < 8$ and $n < 8$. This gives us the final contrapositive of

$$m < 8 \text{ and } n < 8 \rightarrow m + n \neq 16$$

The important takeaway point from this process is as follows – when determining the contrapositive of a statement, be very careful to make sure that you understand how to negate things properly!

From here, the reason why the initial statement is true should be a bit clearer. Essentially, if both m and n are too small, then their sum can't be 16. This is formalized below:

Proof: By contrapositive; we show that if $m < 8$ and $n < 8$, then $m + n \neq 16$. To see this, note that

$$\begin{aligned} m + n &< 8 + n \\ &< 8 + 8 \\ &= 16 \end{aligned}$$

So $m + n < 16$. Consequently, $m + n \neq 16$. ■

2.4 Writing Elegant Proofs

We've now seen three approaches to writing proofs: direct proofs, which follow a logical train of thought to show a conclusion; proof by contradiction, which assumes the opposite of what it wants to prove and shows that this leads to impossible conclusions; and proof by contrapositive, which proves that P implies Q by proving that not Q implies not P. We will use these proof techniques throughout the rest of our exploration of the mathematical foundations of computing.

Although we've described the technical details of how to write proofs of this sort, we have not talked about how to write *elegant* proofs in these styles. Proofs are like essays – you can express the same ideas in many different ways, and depending on how you do it you can make it easier or harder for your reader to understand what you're doing. This final part of the chapter explores techniques for writing good proofs and how to determine that your proofs are correct.

2.4.1 Treat Proofs as Essays

Proofs are rigorous arguments intended to prove some point. The point of writing a proof is to convey a mathematically rigorous argument. Your goal when writing a proof is not just to write down the argument, but to help someone unfamiliar with the result see why it must be true. Well-written proofs make it easy for others to understand your argument, while poorly-written proofs – even ones that have sound reasoning – can actually *prevent* others from following your reasoning.

Approach writing proofs as you would writing essays. Describe what you're going to talk about before you jump right into it. Give guideposts about where the proof is going so that a reader can tell what you are planning on doing. If you have a long, complicated thought, break it down into smaller pieces (i.e. lemmas) and explain each one individually.

As an example, consider the following proof:

> ***Theorem:*** If n is an even integer, then n^2 is an even integer.

> ***Proof:*** $(2k)^2 = 2(2k^2)$ for any integer k. If we have a particular $2k^2$, we can write $2k^2 = r$ for some integer r. If n is even, then $n = 2k$ for some integer k. Therefore, $n^2 = 2r$ for some integer r. ■

Make sense? Didn't think so. All of the right pieces of the proof are here, but they're so jumbled and disorganized that it's almost impossible to determine what's being said here.

There are lots of things that are confusing about this proof: statements are introduced with no context or justification, the flow isn't clear at all, and the ultimate conclusion only makes sense if you skip around randomly. Does that mean the logic isn't a valid? In this case, no; the logic is perfectly fine. However, the *argument* is extremely weak because it requires an intense effort on the part of the reader to see why it works at all.

The standard proof of this result, which we saw earlier, more clearly lays out the steps in sequence. First, we start by writing $n = 2k$ for some integer k, because that's one of the few things we're assuming. Since we want to talk about n^2, it's reasonable to square both sides to get $n^2 = (2k)^2$. We want to write n^2 as $2r$ for some choice of r, so we can simplify the math by noting $n^2 = (2k)^2 = 4k^2 = 2(2k^2)$. And then we're done, because $2k^2$ is an integer and n^2 is twice that integer. If we were to turn this line of reasoning into a proof, it would be significantly easier to follow because it follows a reasonable narrative – at each point, it makes sense to try the step we're about to try.

Logical flow is the main reason we start all our proofs by contradiction or contrapositive by saying something to the effect of "we proceed by contradiction" or "by contrapositive;" by doing so, we provide context for what will follow so that the reader (often, you!) can tell where we're going.

2.4.2 Avoid Shorthand or Unnecessary Symbols

A lot of math involves working with symbols; at this point, we've seen all of the following symbols (and a few more) used to denote specific concepts or quantities:

$$\cup \; - \; \cap \; \Delta \; \wp \; \varnothing \; \mathbb{N} \; \mathbb{R} \; \mathbb{Z} \; \mathbb{Q} \; \blacksquare \; \rightarrow \; \subseteq$$

Although mathematics relies on using precise symbols to denote precise concepts, that doesn't mean that a proof should consist purely of symbols. Rather, a proof should be a piece of writing that conveys how different concepts relate to one another and how the truth of the overall theorem follows from those relations. When you need to use symbols to convey an object or concept, such as the fact that A is a subset of B, it's perfectly appropriate to write out $A \subseteq B$ in symbolic notation, since the notation is specifically designed to convey this. Similarly, if you want to talk about the set of all elements in at least one of A and B, it's best to just write $A \cup B$. This is appropriate because ultimately you are trying to prove a mathematical result that itself is represented in mathematical notation.

On the other hand, you should try to avoid using mathematical notation to describe the *structure* of the argument. For example, you may have seen the symbols $\therefore$ and $\because$ to mean "therefore" and "because." These symbols are commonly-used shorthands that are great for taking notes or writing something up on a whiteboard. However, they can make proofs extremely dense. For example, consider the following proof, which is technically correct but almost impossible to read.

> ***Theorem:*** If n is an even integer, then n^2 is even.

> ***Proof:*** $\because n$ even, $\exists k \in \mathbb{Z}$ s.t. $n = 2k$. $\therefore n^2 = (2k)^2 = 2(2k^2)$. $n^2 = 2(2k^2) \Rightarrow \exists r \in \mathbb{Z}$ s.t. $n^2 = 2r$ ($\because r = 2k^2$). $\therefore n^2$ even. ■

If you haven't seen some of the symbols or shorthands in this proof before, you can appreciate just how inscrutable this proof is. (In case it helps: $\exists$ means "there exists," the $\Rightarrow$ symbol means "implies," and s.t. is shorthand for "such that.") If you have seen these symbols, hopefully your initial response is "wow, that's really, really dense."

Perhaps the best way to summarize the problem with this proof is that it's not *inviting*. This proof isn't trying to help you understand what's going on; it's trying to save on space. You can think of it as "dehydrated proof:" all the essential parts are there, but you're going to have to add water and stir for a long time before you can rehydrate it back into an argument like this one:

> ***Theorem:*** If n is an even integer, then n^2 is even.

> *Proof:* Because n is even, there is some integer k such that $n = 2k$. Therefore, $n^2 = (2k)^2 = 2(2k^2)$. Consequently, there is an integer r (namely, $2k^2$) such that $n^2 = 2r$. Therefore, n^2 is even. ■

The argument has the same structure as the one given above, but it's much, much easier to read.

2.4.3 Write Multiple Drafts

Building off our proofs-are-essays metaphor, it's extremely valuable to sketch out a draft of a proof before writing up your final version. If you start a proof without a clear sense of where you're going with it, you will probably end up with a very convoluted line of reasoning that goes down a lot of unnecessary paths before arriving at the result. Consequently, we recommend writing out a draft of the proof, looking over it critically, then rewriting it to be as clean as possible (and, possibly, repeating this process.)

When you have a first draft of a proof, we suggest tracing through the line of reasoning and checking it for correctness (we'll discuss this in more detail a bit later on). If you're sure the proof is correct, we then recommend taking a look at what you've written and seeing how much of it is actually necessary to show the overall result. Did you prove something you thought would be useful but didn't actually need? Leave it out of the next draft! You can often shrink proofs significantly by using this approach, and the resulting proof will be a lot more focused, easier to read, and less likely to contain extraneous (and potentially incorrect!) details.

2.4.4 Avoid "Clearly" and "Obviously"

When you're writing a proof, you are trying to show a logical train of thought that establishes a result. Key to doing so is the fact that you can back up all of the assertions you're making along the way. Usually, this means manipulating definitions or combining together steps you took earlier in the proof.

When writing proofs, it is extremely common to arrive at a point where you need to show a result that just feels *obvious*, only to find that it's surprisingly hard to show. When this happens, it can be tempting to justify the result by writing "clearly, X is true" or "X is obviously true." *Resist the temptation to do this!* When someone is reading over your proof, words like "clearly" and "obviously" come across as belittling. The reader might look at it and say "wow, apparently it's *clearly* true, but I guess I'm not smart enough to see why..." or "something must be wrong with me, because *I* don't see this as obvious..." Given that the main reason why you might want to mark something as "clearly" or "obviously" true is that it just *seems* true even though you can't prove it, this sort of writing can often be perceived as insulting or condescending.

A good test for whether something is "clearly" or "obviously" true is the following: if you want to write something to the effect of "clearly, X is true," grab a pen or pencil and try to write out a proof of X on a sheet of paper. If you can immediately sketch out a proof, then *write that proof* instead of claiming that X is obvious so that others can follow how you know X is true. If you can't, then probably X isn't as obvious as you might think it is, and you should take the time to work through the proof of X!

2.5 Chapter Summary

- A mathematical proof is a series of logical steps starting from a basic set of assumptions and arriving at a conclusion. Assuming the assumptions are valid and the logic is sound, the result is incontrovertibly true.
- Proofs often involve *lemmas*, smaller proofs of intermediate results which then build into the overall proof.
- Proofs often involve *cases*, branches in the proof that cover different possibilities.
- The *parity* of an integer is whether it is even or odd. Parity interacts in interesting ways with addition and multiplication.
- Two sets are equal if and only if each set is a subset of the other.
- Logical implications are statements of the form "If P, then Q." We denote this $P \to Q$. Such a statement means that whenever P is true, Q must be true as well, but say nothing about causality or correlation.
- To disprove an implication, one finds a way for P to be true and Q to be false.
- A *proof by contradiction* works by assuming the opposite of what is to be shown, then deriving a *contradiction*, a logically impossible statement.
- A number is called *rational* if it is the ratio of two integers, the second of which is not zero, which share no common factors other than ± 1.
- The *contrapositive* of the implication "If P, then Q" is the statement "If not Q, then not P." A statement is logically equivalent to its contrapositive.
- A *proof by contrapositive* proves an implication by proving its contrapositive instead.

2.6 Chapter Exercises

1. Let's define the function $max(x, y)$ as follows: if $x < y$, then $max(x, y) = y$; else, $max(x, y) = x$. For example, $max(1, 3) = 3$, $max(2, 2) = 2$, and $max(-\pi, 137) = 137$. Prove that the following holds for any x, y, and z: $max(x, max(y, z)) = max(max(x, y), z)$.
2. Let's define the absolute value function $|x|$ as follows: if $x < 0$, then $|x| = -x$; otherwise, $|x| = x$. Prove that $|xy| = |x||y|$.
3. Prove that mn is odd iff m is odd and n is odd.
4. Prove that if n is an integer and n is a multiple of three (i.e. $n = 3k$ for some integer k), then n^2 is a multiple of three.
5. A number n called *congruent to one modulo three* iff $n = 3k + 1$ for some integer k and is called *congruent to two modulo three* iff $n = 3k + 2$ for some integer k. Every integer is either a multiple of three, congruent to one modulo three, or congruent to two modulo three. Prove that if n is an integer and n^2 is a multiple of three, then n is a multiple of three.
6. Prove that $\sqrt{3}$ is irrational.
7. A triple of positive natural numbers (a, b, c) is called a ***Pythagorean triple*** if there is a right triangle whose sides have length a, b, and c. Formally, $a^2 + b^2 = c^2$. Some examples of Pythagorean triples include (3, 4, 5), (5, 12, 13), and (7, 24, 25). Prove that if (a, b, c) is a Pythagorean triple, then at least one of a, b, or c must be even.
8. Prove that if (a, b, c) is a Pythagorean triple, then $(a + 1, b + 1, c + 1)$ is *not* a Pythagorean triple.
9. Prove that (a, a, b) is *never* a Pythagorean triple.
10. A natural number n is called a multiple of four iff there is some $k \in \mathbb{N}$ such that $n = 4k$. For every natural number n, exactly one of n, $n + 1$, $n + 2$, or $n + 3$ is a multiple of four. Prove that for any natural number n, that either n^2 or $n^2 + 3$ is a multiple of four.
11. According to the World Bank, the population of Canada in 2011 was 34,482,779.* Prove that there are no natural numbers m and n such that $m^2 + n^2 = 34{,}482{,}779$.
12. Prove or disprove: if r is rational and s is irrational, then $r + s$ is irrational.
13. Prove or disprove: r is rational iff $-r$ is rational.
14. Prove or disprove: if r is irrational and s is irrational, then $r + s$ is irrational. ★
15. Prove or disprove: if r is irrational and s is irrational, then r^s is irrational. ★
16. Suppose you are having dinner with nine friends and want to split the bill, which is $44. Everyone pays in dollar bills. Prove that at least two people in your group paid the same amount of money.
17. Suppose that A, B, and C are sets. Prove that $(C - B) - A = C - (B \cup A)$.
18. Prove that for any sets A, B, and C, that $(A \Delta B) \Delta C = A \Delta (B \Delta C)$. This shows that symmetric difference is ***associative***.
19. Prove that $A \cap (B \cup C) = (A \cap B) \cup (A \cap C)$.
20. Prove that $A \cup (B \cap C) = (A \cup B) \cap (A \cup C)$.
21. Prove or disprove: If $A = B$, then $\wp(A) = \wp(B)$.

* Source: http://www.google.com/publicdata/explore?ds=d5bncppjof8f9_&met_y=sp_pop_totl&idim=country:CAN&dl=en&hl=en&q=population+of+canada, as of September 30, 2012.

22. Prove or disprove: If $\wp(A) = \wp(B)$, then $A = B$.

23. Prove that if $A \subseteq B$ and $B \subseteq C$, then $A \subseteq C$. This shows that $\subseteq$ is *transitive*.

24. A *right triomino* is an L-shaped tile made of three squares. Prove that it is impossible to tile an 8×8 chessboard missing two opposite corners using right triominoes.

25. Is it *ever* possible to tile an $n \times n$ chessboard missing two opposite corners with right triominoes, assuming $n \geq 3$? If so, find an n for which it's possible. If not, prove that it's always impossible for any $n \geq 3$.

26. Suppose that you have twenty-five balls to place into five different bins. Eleven of the balls are red, while the other fourteen are blue. Prove that no matter how the balls are placed into the bins, there must be at least one bin containing at least three red balls.

27. Suppose that you have a 3" × 3" × 3" cube of cheese (or 3cm × 3cm × 3cm if you live in a place that uses a sane, rational system of measurements) that's subdivided into 27 1" × 1" × 1" smaller cubes of cheese. A mouse wants to eat the cube of cheese and does so as follows: she first picks any cube to eat first, then moves to an adjacent cube of cheese (i.e. a cube that shared a face with the cube that was just eaten) to eat next. Is it possible for the mouse to eat the center cube of cheese last? If so, show how. If not, prove it's impossible.[*] ★

28. Consider the quadratic question $ax^2 + bx + c = 0$, where a, b, and c are integers. Prove that if a, b, and c are odd, then $ax^2 + bx + c = 0$ has no rational roots (that is, there are no rational values of x for which $ax^2 + bx + c = 0$). As a hint, proceed by contradiction; assume that $x = p / q$ for some p and q, then think about the parities of p and q. ★

29. A *Latin square* is an $n \times n$ grid filled with the natural numbers 1 through n such that every row and every column contains each number exactly once. A *symmetric Latin square* is one where the square is symmetric across the main diagonal; that is, the number at position (i, j) is equal to the number at position (j, i).

    Prove that in every symmetric Latin square, every natural number between 1 and n appears exactly once on the main diagonal.

* Adapted from Problem 4E of *A Course in Combinatorics*, *Second Edition* by Lint and Wilson.