# Graphs, Groups, Infinity

*Three stories in mathematics*

Travis Dillon

# INTRODUCTION

Good mathematics tells a story.

You wouldn't know it from the classes. Most math classrooms avoid action, plot, and narrative almost as a matter of principle. Nevertheless, this is what I want to convince you of. And I'd like to convince you of more: that mathematics is, at times, closer to art than science, more the poetry of logic than the science of numbers; that a crucial part of mathematics exists beyond its practicality; that it is in turn absurdly delightful and delightfully absurd.

These statements likely stand in stark contrast to your experiences with mathematics. Most students encounter math as a dry, disjointed subject, a sterile collection of truths and procedures passed down from the high priests of mathematics (whoever they might be) to you, to be accepted and memorized. For the most part, these ideas live short and tortured lives the length of the night's homework assignment, being quickly replaced by the next vaguely related lesson and frantically revived only for short periods before the chapter test and the final.

The tragedy is that almost everyone completes more than a decade of mathematical instruction without ever encountering real mathematics. At its core, mathematics is simply the art of explanation. It's what happens when you haven't grown too old to stop asking "Why?" but you're old enough to start investigating those questions yourself. Mathematics is in the business of stimulating and satisfying curiosity; its products are questions and *aha!* moments.

But school mathematics ignores all that. Tossing aside history, context, and content, school mathematics focuses nearly exclusively on the symbols of mathematics. Each lesson is a lecture in the the allowable and forbidden manipulations of increasingly exotic symbols, which students collect like passport stamps as they progress through the curriculum. Milestones are marked and defined by the introduction of new symbols: the four basic operations $+, -, \times, \div$; the square root $\sqrt{\ }$ and the exponent $4^3$; the dreaded variable $x$; functions $f(x)$ and polynomials $x^2 + x + 1$; congruence $\cong$ and similarity $\sim$ and parallel $\|$ and perpendicular $\perp$; logarithms $\log(x)$ and trigonometric functions like $\sin(x)$; and, for a select few initiates, the symbols of calculus, $\frac{d}{dx}$ and $\int$ and $\Sigma$.

This approach mistakes prevalence for importance. Yes, mathematics has evolved a set of symbols for communicating specialized and precise ideas. But they are just that: symbols in service of communication. They're certainly convenient at times, but they're far from the heart of mathematics. Flip through the pages of this book, and you'll likely find many more words and many fewer symbols than you expected. The central focus of mathematics is *ideas*; symbols are just a way to communicate them.

In his essay *A Mathematician's Lament*, Paul Lockhart imagines a parallel universe, where music education takes a similarly symbol-centric approach. Because musicians communicate their ideas with a set of curious symbols that follow certain rules, these symbols and rules are quickly adopted by schools as the "language of music." In classes and in homework, students complete a litany of scale exercises in this language, paying close attention to the spacing of the notes and the shape of their ornaments. Listening to music, let alone playing it, enjoying it, or composing it, has no place in the classroom. Most students, in fact, have no idea that there even *is* any art behind the symbols. They are left completely unaware of the stories that music can share.

This dystopic situation is fortunately a work of fantasy. But it describes mathematics education all too well, and it begs the question: What is the music of mathematics?

Mathematics is the story of ideas, a story told for an exacting and untiringly curious audience. A basic element of storytelling in mathematics is the *proof*, which is at its heart nothing more and nothing less than an explanation of why something is the way it is. Each field of mathematics is a saga, with its own setting, cast of characters, and plot arcs. Each proof is a story, illuminating the relationships between the characters, building or releasing tension, resolving or even thickening the plot. A bad proof is like a bad story: obscure, confusing, pedantic, long-winded. A good proof is like a good story: enjoyable, satisfying, and perhaps even surprising.

Ironically, the proofs you likely remember, two-column proofs from Geometry class, are a quintessential example of bad proofs: tedious, obfuscatory, and no fun whatsoever. Contrary to your experience with these two-column abominations, a proof should be refreshing and enjoyable. The goal of a proof is to convince and enlighten; a good proof is a pleasure to read.

This is best illustrated through an example. Somewhere, in the mythical plains of Doglandia, there is a $5 \times 5$ grid, and in each square sits a dog. At precisely 2:57 PM, each dog will become a bit restless move to an adjacent square (above, below, left, or right; not diagonal). The reasons for this are shrouded in mystery and tradition. The urgent question for these canines is this: Can they coordinate their movements so that, at 2:58 PM, there is still exactly one dog in each square?

The next diagram shows one attempt at coordination. All seems to be going well until we get to the last dog in the lower right corner, who has nowhere to go. But that's just one attempt. Pause for a moment and see if you can find a strategy that works. Seriously: Give it a try.

Once you've tried to construct multiple strategies, as I hope you

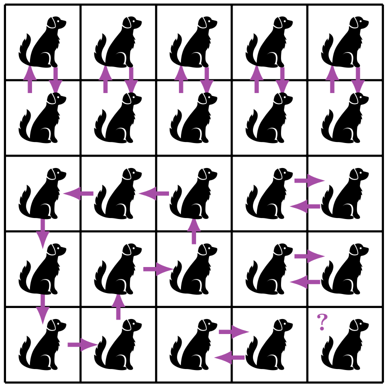


have, and none of them worked, you'll probably begin to wonder: Is the solution somehow extra tricky and difficult to find, or is there actually no solution at all?

What we have here is a genuine mathematical problem: Easy to explain and understand, but without an obvious answer, even after a bit of thought. Now, I could just tell you the answer: There is no solution—it's an impossible problem. But that's not the whole story. In all likelihood, you're not just looking for the answer. You're looking for an *explanation*.

Let's look at the problem again, but this time, let's color the squares like a chessboard:

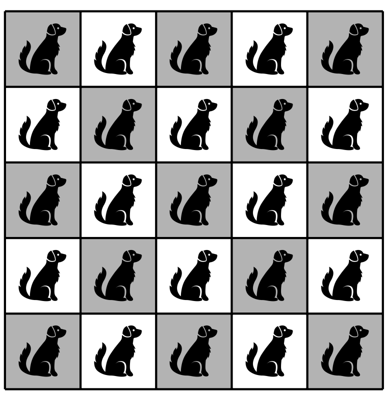

Here's the insight: When a dog on a gray square moves, it always ends up on a white square; and a dog on a white square always moves to a gray square. Since there are 13 gray squares and only 12 white squares, the dogs will always leave one of the gray squares empty when they move at 2:57 PM. No matter how they move, they'll never end up with one dog on each square.

This is mathematical proof at its best: an elegant, deceptively simple idea that casts the problem in a new light, turning what had been mysterious into something nearly obvious. We've all experienced that magical moment when something clicks and it all suddenly makes sense; to do mathematics is to be in the business of creating more of those moments.

And when enlightenment strikes, you can usually see much more than just the problem you first considered. This chessboard-coloring idea, for example, also explains why there's no solution even on a bigger $7\times 7$ grid. When the $7\times 7$ grid is colored like a chessboard, the 49 squares are split as evenly as possible between gray and white—24 white and 25 gray. And once all the dogs move, there are only 24 dogs spread across 25 gray squares. There simply can't be one dog in each square, no matter how the dogs choose to move.

In fact, if you color in a chessboard style *any* square grid whose side length is odd, you end up with one more gray square than white. Which means that not only is the Doglandia problem impossible on $5 \times 5$ and $7 \times 7$ grids, but also on $9 \times 9$, $11 \times 11$, $13 \times 13$, and on and on and on.

At this point, you might be wondering: What if the size of the grid is even? Or what if the sides are different lengths? What if you allow diagonal moves? This is the wonderfully cyclic nature of mathematics: Questions might lead to answers, but at least as often, answers lead to more questions. The story flows onward: Finishing one chapter simply leaves you at the beginning of the next.

The goal of this book is to introduce you to mathematics as a creative, lively, playful endeavor. Its pages are, more than anything else, composed of stories: of people and ideas, puzzles and problems, animals and dinosaurs, and alternate-universe versions of you. A story that starts with bridges in an 18th-century German city and ends with tree houses and salesmen in the 21st. A story of a harried hotel manager coping, just barely, with the absurdities of infinity. And a story about how addition, rotation, and pyramids have a lot more to do with each other than you ever would have thought.

To do mathematics is to be curious, imaginative, creatively frustrated, and inspired. It's one of the most rewarding activities that I know, and I hope that in reading this book, you begin to feel the same.

*A guide to reading this book.* The body of this book is divided into three parts, each a story of one idea. Each comes from a different area of mathematics, so they vary in style and flavor. They're largely independent of one another, so if you read the first chapter of a part and decide it's not your thing, you can skip to the next one. I encourage you to read all three parts, though; each of these stories starts somewhere simple and unfolds into something beautiful and surprising.

This book is an invitation to mathematics through immersion; the guiding philosophy is that there's no better way to know mathematics than to *do* mathematics. This is more challenging than remaining on the sidelines, but it's infinitely more rewarding. Accordingly, various questions and problems are sprinkled throughout the text with the heading PRACTICE. Learning mathematics is hard work, and these problems give you a chance to pause and, as the British say, have a think about things. I highly recommend that when you reach one of these practice problems, you do just that. None of these problems is meant to stymie or confuse, and I think you'll find that working on these problems will help you better internalize the new ideas in each chapter.

At the end of each chapter is a collection of problems. This is a chance to do mathematics yourself, to make your own *aha!* moments. There are few things more satisfying than working at a problem for some time only to have inspiration strike like lightning. I encourage you to work on these problems like mathematicians do: with paper and pencil, over periods of time. Try out some ideas when you first read the problem; if these don't go anywhere, give it some time to marinate. Think about them in idle moments, on walks, or in the shower. Sometimes inspiration strikes only after trying several different approaches.

Other times, though, inspiration seems likely to never come at all. Doing mathematics requires a good amount of patience and acceptance that much of the time, you're not really sure what tactics will pan out. But this shouldn't boil over into frustration, and sometimes, especially when you're first building mathematical intuition, a hint or two is what's called for. You can find a hint at the end of the book for most problems which will hopefully point you in the right direction without taking away the joy of discovery. And remember, this isn't school—none of these problems is mandatory. Find the ones that you enjoy; the ones that, like a catchy tune, your thoughts return to again and again.

At the end of each of the three parts is a chapter titled Wrapping Up. It ties together the material from that portion of the book with a bit of historical context and pointers to places to learn more, if you were particularly taken with the ideas. It also contains an assortment of additional problems to puzzle over that combine ideas from different chapters and explore new directions.

With all that prelude, there's only one more thing to say:

Happy reading!

ACKNOWLEDGEMENTS

I am grateful to the two Sams who read portions of this book and offered many helpful suggestions: Sam "The Elder" Gutekunst, who read Part I and suggested the Doglandia problem in the introduction (which is itself modified from a problem in the textbook *Essential Discrete Mathematics for Computer Science* by Lewis and Zax), and Sam "The Younger" Macdonald, who read Part II. Narmada Varadarajan has been a valuable sounding board at every stage of writing and was always ready to listen to my ideas, be they inchoate or considered, pithy or monologic. Her proposals and counterproposals, if not always serious, were always appreciated.

And finally: Liz Sattler, who has read and improved every part of this book with her unfailingly insightful comments and suggestions. Working with her has been one of the most rewarding experiences I've had in my stint as an undergraduate, and for that, I cannot thank her enough.

# CONTENTS

# PART I

# GRAPHS

# One

## INTRODUCTION TO GRAPHS

### 1. Königsberg, 1735

It's 1735, and you're in love.

Every Sunday afternoon, you walk the streets of the German city Königsberg with your partner, admiring the springtime weather and, naturally, each other. A river flows through the city, carving it in two and leaving a pair of islands in its midst. As fashionable people, this is where you spend your time, wandering the bridges of the Pregel River (Figure 1.1).

A trend has overtaken the city. Those fashionable folk who spend all their time on the bridges of Königsberg have started to notice that they're spending more time on some bridges than others. Like any good people with too much time on their hands, they've turned this into a game. Everyone who's *anyone* is now trying to take a stroll through the city that crosses each bridge exactly once.

This is more trouble than it sounds. The bridges have been walked in all sorts of ways, and no one found a solution yet. You're determined,

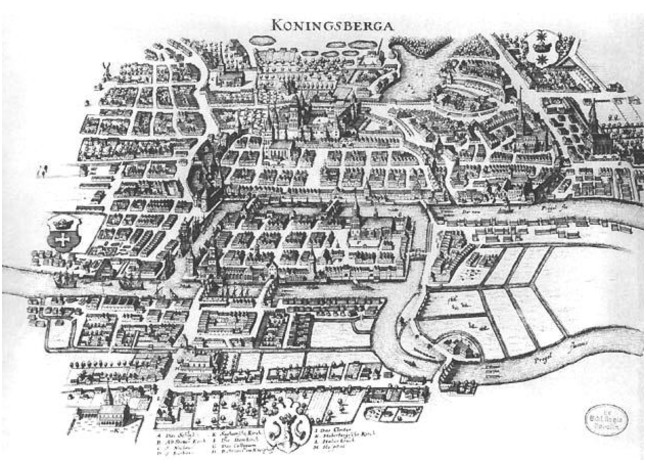


Figure 1.1: A map of Königsberg

though—just because no one has found a solution doesn't mean there isn't one. Just in case, though, you decide to write to your friend, the rockstar mathematician of the $18^{\text{th}}$ century.

Leonhard Euler (pronounced "Oiler") is one of the most prolific mathematicians of all time. His collected mathematical works fill more than 60 volumes, and his written correspondence and notes fill another 15. He's one of the last great mathematical generalists, contributing to nearly every area of mathematics and laying the foundation for several new ones, even after losing vision in both eyes. If anyone could solve this problem, he could.

Before we see what Euler had to say, it's worth trying the puzzle out for yourself. Looking at the map of Königsberg in Figure 1.1, can you find a path through the city that uses each of the seven bridges exactly once? You're allowed to start anywhere, and you don't have to end where you start.

After a few tries, it might begin to seem that there's no way to do it. But is it just that the solution is really hard to find, or is it actually impossible?

Here's one way to tell for sure. Give each bridge a name, and then list all the possible ways to order them. If you give the bridges the creative names A, B, C, D, E, F, and G, then your list might start ABCDEFG, ABCDEGF, ABCDFEG, . . . . You could check each one to see whether it's possible to cross the bridges in that order. If there's one that works, you've solved it! And if there isn't, then the problem is impossible.

In theory, this strategy works, but it would take ages: There are more than 5,000 ways to order the bridges, and you'd have to make sure that you didn't forget any from you list, which is a demanding task in itself.

Instead, we need to analyze the problem a bit more cleverly. Euler's first insight was that the map of Königsberg has a lot more information than we need for the problem. All that really matters is which landmasses each bridge connects. Instead of thinking about the complicated city map, we can focus on a schematic like the one in Figure 1.2.

Each dot represents one of the four landmasses, and each line represents a bridge connecting those two landmasses. Now, instead of thinking about bridges and landmasses, we can instead ask: Is it possible to trace along the lines in Figure 1.2 so that each line is used exactly once? If we can, the corresponding path on the map crosses each bridge exactly once.

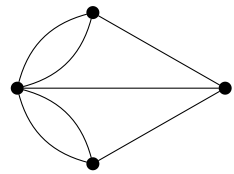

Figure 1.2: A schematic of the Königsberg bridges

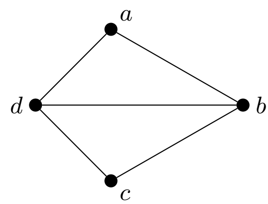


Figure 1.3: A schematic of some fake bridges

The next step Euler took might be called Wishful Thinking. Essentially, he asked: If there *were* a solution to the problem, what would it have to look like? Alternatively, we can think of this as narrowing the field of possibilities. If we found out, for example, that the path had to start at a specific spot, that would make a path significantly easier to find.

It's easier to think about this in an example where we actually know a path exists. So instead of Königsberg, let's consider a different set of bridges shown in Figure 1.3.

Each of the dots has a letter name so that we can represent a path by the sequence of dots that it visits. For example, one path that uses each line exactly once is *dabdcb*, which starts at $d$, moves to $a$, then $b$, then $d$ again, and so on. There are quite a few others; go ahead and find some yourself.

Drawing these paths can give some insight into their properties. Here's a pretty simple observation: With the exception of the very beginning or end of the path, every time we enter a dot, we also leave it. For example, in the path *dabdcb*, we enter $a$ from $d$ and exit toward $b$; after entering $b$, we exit toward $d$; and so on. This is a straightforward observation; it might even feel obvious. But if we're careful, we can use it to get some crucial insight into the problem.

Let's pick a dot that's not where the path begins or ends. The path exits that dot every time it enters it, and the path uses each line exactly once, which means that the dot we chose has to touch an even number of lines—if it were touching an odd number of lines, the path couldn't leave every time it entered. This argument is true for *every* dot that doesn't begin or end the path.

At this point it's good to summarize what we deduced. We started with the assumption that there is a path crosses each line in the drawing exactly once. If that's the case, then every dot, except possibly where the path begins and ends, has to be connected to an even number of lines. We can see that in Figure 1.3, where $b$ and $d$ have an odd number of lines, but $a$ and $c$ have an even number.

In the Königsberg drawing, however, each dot has an odd number of lines. We've determined that if there's a way to trace a path that uses each line exactly once, then at most two of the dots can be connected to an odd number of lines. But that's not the case! Which means that, no matter how you try, you can never walk a path through Königsberg that crosses each bridge exactly once.

This is the chain of logic that Euler produced. This type of solution—proving, for now and forever, that something does not and cannot exist—is something that's very peculiar and fascinating about mathematics. Nowhere else can you be sure, absolutely sure, that something does not exist. After all, just because you've never seen Bigfoot doesn't mean he (or she) isn't hiding out there in the trees, biding his (or her) time until a spectacular reveal. This mathematical proof, though, says something much more than observation ever can. It says that no matter how clever someone is, or how much time they spend, or how powerful their computer is, they will *never* find a path—and you can know this with absolute certainty.

These "nonexistence proofs," as they're sometimes called, are found throughout mathematics. In this case, it's perhaps not that impressive. A puzzle doesn't have a solution. If you had a large chunk of free time and the ability to weather astronomical levels of tedium, then you could check every possibility and find out that way that there's no good path. Nonexistence proofs are found throughout mathematics, however, and in many cases it would be impossible to check through every possibility. We'll see situations like this several times in this book, where the unique level of certainty to be found in mathematics truly becomes apparent.

## 2. What is a graph?

A crucial element in solving the Königsberg problem was to determine what information was actually relevant. Once we did that, we were able to work with the simplified model in Figure 1.2. Keying into important information and filtering out the irrelevant parts is a crucial element of any problem-solving method.

Let's give it a go in some other scenarios.

*Question* 1. In a fit of industriousness, you baked four different pies—apple, cherry, peach, and huckleberry—and now you want to distribute them among your four neighbors John, Joan, Jack, and Janice. The catch is that your neighbors only like certain flavors: John and Janice like peach, while Jack likes huckleberry and apple; Janice and Joan have a sweet spot for cherry pie, and everybody likes apple. Is it possible to distribute one pie to each neighbor so that everyone is happy?

This question is obviously contrived, but it's representative of a whole class of problems called *matching* or *assignment* problems. In these sorts of problems, a collection of objects needs to be assigned to some group in a way that satisfies certain conditions. Assigning projects to employees in a way that meets their qualifications is a different example.

For now, let's focus on the simple example in Question 1. In this scenario, the relevant information is clear: which people like which pies. The main issue with the pie assignment problem is that the information is hard to manipulate. One of the best ways of making clunky data easier to work with is to visualize it. So instead of a list of pie preferences,

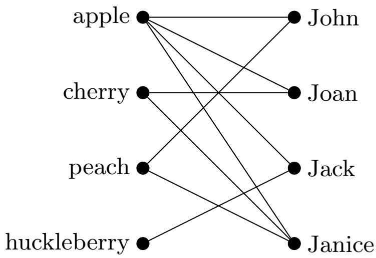


Figure 1.4: Pie preferences amongst your neighbors

we can use the diagram in Figure 1.4, which connects each person with the pie flavors they like.

What we're looking for, then, is a way to match each person with a unique pie using only the lines in the diagram. With this visualization, it's a lot easier to solve the problem.

PRACTICE 1.1. Find a way to distribute the pies in a way that makes your neighbors happy. Then find two more.

Here's a very different question.

*Question* 2. Is it possible to fly from Spokane International Airport to any other commercial airport in the United States with a maximum of two layovers?

This is more difficult to answer, for the simple reason that there are a lot of airports in the United States. But we can still think about it schematically, which, with a computer, would allow us to answer the question. In this problem, there's a lot of potentially irrelevant information—how long the flights are, the size of the planes, and the frequency of the flights, to name a few. In order to answer the problem, all we actually need to know is which pairs of cities are connected by flights. If we had the space, we could draw this as a gigantic dot-and-line diagram as we did for the pie and Königsberg questions. Question 2 then asks whether it's possible to get from the Spokane International Airport dot to any other dot with a path of at most three lines (since each line corresponds to a flight segment).

Although the (hypothetical) dot-and-line drawing of flights in the United States doesn't directly answer the question, it does capture the underlying structure of the problem, just as the dot-and-line drawings did for the pie and Königsberg problems. In the late 19$^{\text{th}}$ and early 20$^{\text{th}}$ centuries, mathematicians began to notice the same thing—that these drawings were popping up in all sorts of seemingly unrelated problems.

When this happens, mathematicians use a technique known as *abstraction*, which is the mathematician's way of getting more for less. Instead of looking at each instance of a dot-and-line drawing in isolation as a particular model, mathematicians throw away the details and study dot-and-line drawings *themselves*, independent of any "real problem." At first glance, this might sound ridiculous—how could severing connections from reality be helpful at all? In fact, it's one of the most powerful techniques in mathematics. By building a body of knowledge about these abstract structures, it's easy to quickly apply this knowledge to each individual instance without having to derive it again each time.

We already have an example of this! When we were analyzing the Königsberg bridge problem, we actually found a condition that *any* dot-and-line drawing has to satisfy in order to have a path that uses each line exactly once: at most two dots can have an odd number of lines. We can use this rule in situations far removed from Königsberg and bridges. Say, for example, a group of flights operates between your four favorite cities, and there's a different specialty snack on each flight, so you'd like to take each flight, but only once—you're not made of money. Is this possible? Well, just make a dot-and-line drawing of the flights between the cities and test it against our condition.

So how do we go about studying abstractly? The first step is to set up the arena: What exactly are we going to study?

Definition 1.1. A *graph* is a set of points, called *vertices*, together with a collection of *edges* that connect pairs of vertices.

In this definition, "vertices" is the formal word for "dots" and "edge" is the formal word for "line." The word "graph" for dot-and-line drawings originates in the late $19^{\text{th}}$ century and has nothing to do with the $xy$-graphs you might remember plotting in high school. We'll actually never need to plot a function in this book, so "graph" will always be used in the sense of Definition 1.1.

A large part of researching something mathematically is figuring out what questions to ask. Once we have an object to study—in this case, abstract graphs—we need to decide what we want to know about it. In this, as with the idea of graphs themselves, inspiration flows from the applications, so it's helpful to have a good number of examples to pull from.

## 3. A world of graphs

A graph is a useful model whenever the important information is the connections or relationships between pairs of things. Because of this generality, graph-like structures can pop up in all sorts of places.

### Maps

A road system in a state might be modelled by a graph, where each town or city is a vertex and edges connect two cities if there is a road that travels directly from one to the other. This is like a terrestrial version of the airport graph discussed in Question 2.

An economic version of this idea might be to make a graph where each country is a vertex and there is an edge from each country to its five largest trading partners. Different cutoffs would give different graphs. An economist might consult a graph like this to determine groups of countries that like to trade among themselves, for example.

Sometimes we need more information than a simple graph can provide. In that case, we can *augment* our model by adding more information to it. For instance, to find the shortest drive between two cities, you need to know not just which cities are connected, but how long the roads are. We can add this information to the graph by labelling each edge with the length of the road it represents. Or we could use a trade graph that connects every country and labels each edge with the net trade between the countries it connects, instead of just including edges that meet a certain threshold. One of the benefits of the graph model is how flexible it is in accommodating adjustments like these. The main question is how much information you want to include. The essential balance is between simplicity and functionality. Ideally, the model is as simple as possible—and no simpler.

## ‖ Networks, the social kind

Graphs can also model more finicky affairs, like human relationships. We might, for example, model the social interactions in a high school by constructing a graph where each student is a vertex and edges connect students who are friends. A graph like this could be used to identify cliques and in-groups or highly popular students. Facebook connections can be modeled in the same way but on a much larger scale.

Both of these graphs change frequently (friendships are ended and formed, and people join or leave the network) and are absolutely massive. For both of these reasons, graphs like these are stored in a computer instead of a sheet of paper. For computers, frequent updates and a massive number of vertices and edges are not a problem.

We can augment these graphs, as well. A more realistic, if perhaps a bit less optimistic, model incorporates the idea that that friendship is not necessarily a two way street. I may think I'm friends with you, even though you don't feel the same. In this case, two things should happen. First, you should become friends with me; I guarantee you'll enjoy it. Second, we can update our graph model by placing a little arrow on each edge to indicate which way the friendship points, like in Figure 1.5. Such a model would also be useful for other social networks, like Twitter, Instagram, and the like, where "following" is not necessarily reciprocal.

## ‖ Networks, the other kind

Back in the days that are olden—but not too olden—before the internet was beamed through the air like a flock of magical invisible messenger owls, computers had to be wired together to communicate. The way the entire network was connected would impact everything from speed of communication to computational ability to vulnerability. Today, the "internet graph" is much larger, but these small networks still exist in

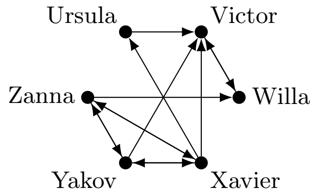


Figure 1.5: Graph of a small social network

various forms, such as when a government sets up servers disconnected from the internet to protect from hacking or establishes a network of secure communication channels among government agents.

## ‖ Other examples

Any type of assignment problem, like Question 1 for example, can be represented as a graph. And the red string on a stereotypical conspiracy theorist's board of truth form the edges of a graph representing hidden connections between events and people (which together form the vertices of the graph).

## ‖ Some concrete families of graphs

Just as various applications are helpful to keep in mind for inspiration, it's useful to have some explicit examples of graphs to build and test conjectures. Here are a few that we'll find useful throughout the book. Each graph has a symbol that will make it easier to refer to.

A *complete graph* has one edge between every pair of vertices. The complete graph with $n$ vertices is denoted $K_n$.

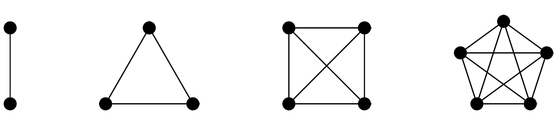

Figure 1.6: Complete graphs on 2, 3, 4, and 5 vertices

A *path graph* is made by connecting some number of vertices in a row, as in Figure 1.7. We write $P_n$ as shorthand for the path graph with $n$ vertices.

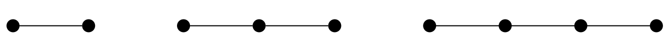

Figure 1.7: Path Graphs on 2, 3, and 4 vertices

A *cycle graph* is a path graph with an additional edge that joins the two endpoints; this graph is denoted $C_n$.

A graph is *bipartite* if the vertices can be divided into two groups so that each edge of the graph connects a vertex in one group to a vertex in the other. For example, the graph in Figure 1.4 is bipartite. The

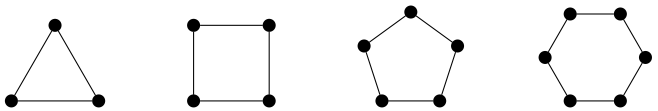

Figure 1.8: Cycle graphs on 3, 4, 5, and 6 vertices

*complete bipartite graph* on $m$ and $n$ vertices, denoted $K_{m,n}$, consists of two groups of vertices with sizes $m$ and $n$ and every edge between the two groups.

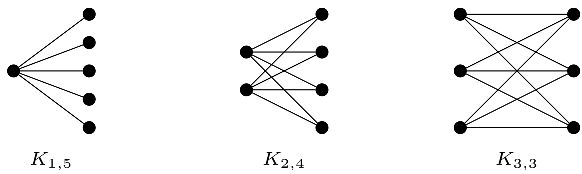


Figure 1.9: Three complete bipartite graphs on 6 vertices

PRACTICE 1.2. Show that $P_n$ is bipartite for every $n$. What values of $n$ make $C_n$ bipartite?

We call a graph *simple* if no two vertices are connected by more than one edge. Another way to make graph examples is to just list all the simple graphs with a given number of vertices. There are 34 simple graphs with 5 vertices and 156 simple graphs with 6 vertices, so this rapidly becomes quite difficult. But it's not too hard to list all graphs with four vertices—they appear in Figure 1.10.

It's important to realize that the position of the vertices in a graph isn't important; only the relationships are. This is what is meant when we say that Figure 1.10 contains "all" graphs with four vertices: By moving its vertices around, any graph with four vertices can be made to look like one of the 11 graphs in Figure 1.10. For example, the graph 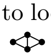 is the same as the second-to-last graph on the bottom row.

## 4. KÖNIGSBERG REVISITED

Now that we've set up a more general theory of graphs, let's take another look at the Königsberg problem. Instead of tackling one bridge diagram at a time, we can search for conditions that guarantee or forbid the existence of a path crossing every edge exactly once in general. To talk about this precisely, it helps to establish some terminology.

DEFINITION 1.2. A *walk* on a graph is made by tracing along consecutive edges of a graph. If a walk uses each edge exactly once, it is called *Eulerian*.

I like to imagine a walk as a minuscule person literally walking along the edges of a graph. For example, one walk in Figure 1.3 consists of starting at vertex $b$, then travelling along the edge to $d$, from there to

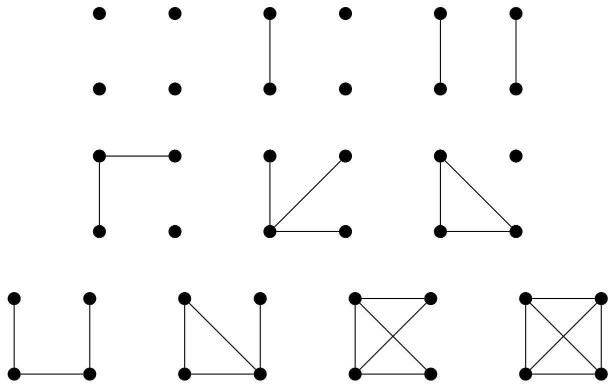

Figure 1.10: All simple graphs with four vertices

$c$, then to $b$, $a$, $d$, and $b$ in turn, and ending at $a$. Usually, walks are written down by the vertices they visit, so this one would be recorded as *bdcbadba*. As you can see, it's perfectly acceptable to cross the same edge twice in a walk. If a walk doesn't repeat any edges, it's called a *trail*. (So, for example, every Eulerian walk is actually a trail.)

DEFINITION 1.3. The *degree* of a vertex is the number of edges connected to it.

In Figure 1.3, for instance, the degree of vertices $b$ and $d$ is 3, and the degree of vertices $a$ and $b$ is 2. For a different example, the degree of every vertex in a cycle graph is 2.

The reasoning discussed in pages 4–5 to solve the Königsberg bridge problem works equally well for any graph. Using the terminology we just introduced, we can state this as a general result. This is our first *theorem*: a proved mathematical result.*

THEOREM 1.4. *Any graph with an Eulerian walk has at most two vertices whose degree is odd.*

Contrary to what it may look like, this is a fairly powerful result. To understand why, take a glance at Figure 1.11. Which graph has an Eulerian walk?

The answer, of course, is that it's a trick question—neither of them does! Likely you solved this by noticing that there are four vertices with odd degree in both graphs; so Theorem 1.4 shows that neither has an Eulerian walk. But now try to imagine how you might solve the problem without Theorem 1.4. Trial and error seems to be the only way forward, which makes it very hard to definitively say that no Eulerian walk exists.

* You'll notice that definitions, theorems, and examples are numbered collectively. This is so that Theorem 3.4 isn't followed by Definition 3.1 and Example 3.8; it keeps everything a bit more ordered. Only the practice problems are numbered separately.

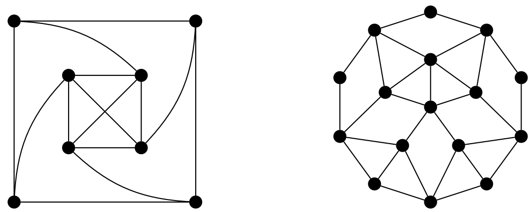

Figure 1.11: Two graphs

And in fact, this theorem can be interpreted as a first powerful nonexistence result: *There is no graph that has both an Eulerian walk and at least three vertices with odd degree.* This is a more impressive result than the simple fact that the Königsberg bridge graph has no Eulerian walk, since there are infinitely many graphs that have more than two vertices with odd degree.

PRACTICE 1.3. Back up this claim by finding infinitely many graphs that have at least three vertices with odd degree. (Think about the graph families introduced in Section 1.3.)

So this is really cool! But in a certain sense, we have only solved half the problem. We know that graphs which don't satisfy the condition in Theorem 1.4 don't have an Eulerian walk, but that doesn't mean that graphs which satisfy the condition *do* have one. Consider the graph in Figure 1.12. (It is in fact one graph, even though it looks like two, since it fits our definition of a graph: vertices, with some pairs connected by edges.) Even though each vertex has even degree, it has no Eulerian cycle. The issue, of course, is that there is no way to get from the square on the left to the square on the right (or vice versa).

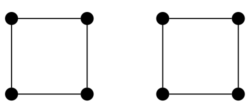

Figure 1.12: A disconnected graph

What this indicates is that something beyond the degree condition is required for a graph to have an Eulerian walk. At the very least, the graph can't be split into different pieces.

DEFINITION 1.5. A graph is *connected* if there's a walk between any pair of its vertices.

This is just a way of writing down formally what you already know "connected" should mean: You can to get from one point to any other. The graph in Figure 1.12 is not connected, because there's no walk from a vertex on the left to a vertex on the right. The graph in Figure 1.3, on the other hand, is connected, since you can get from any vertex to any other.

Does a graph have to be connected to have an Eulerian walk?

> ⚠ A vertex whose degree is 0 is called *isolated*. It may seem a bit silly to have a vertex without any edges connected to it, but sometimes they're useful—a landmass without any bridge connected to it, for example. In any case, if we add isolated vertices to a connected graph, it doesn't change whether that graph has an Eulerian walk—since it doesn't change the edges—but it *does* change whether the graph is connected.

This is really a technicality, though: If isolated vertices are disallowed, then any graph with an Eulerian walk is connected (try to prove this!). It turns out that adding this condition is enough to guarantee an Eulerian walk: Any connected graph with no isolated vertices and at most two vertices with odd degree has an Eulerian walk. We won't prove it here, but you can poke around the web if you're interested.

Now we know statements that point in both directions:

1. If a graph with no isolated vertices has an Eulerian walk, then it is connected and has at most two vertices with odd degree.
2. If a graph with no isolated vertices is connected and has at most two vertices with odd degree, then it has an Eulerian walk.

Together, these two statements say that "having an Eulerian walk" is the same as "being connected and having at most two vertices with odd degree." If either one of these is true, then the other one must be, as well; and if one is false, so is the other. Mathematicians call conditions like these *equivalent*, and they're one of the more satisfying flavors of theorem, since they show that two seemingly different conditions are actually the very same thing. We'll see this again when we look at trees in the problem section at the end of this chapter.

Great! We've completely solved the problem. This is the point where research mathematics is the most philosophically different from school math. In school, this is the end of the process. You solve the problem, turn it in, and collect a gold star or whatever it is they're handing out these days—and that's that. A mathematician views the solution of this problem as a jumping off point, as inspiration for further investigation.

This requires a shift from observing to participating in mathematics, but the gain is huge. Instead of answering a host of boring questions just 'cause, well, you have to, you choose questions that are interesting to you and figure out how to answer them. So get curious!

Here's one general question to ask: What if something changed? Here's an example: In the Königsberg problem, it would be nice to have a path that not only crossed every bridge exactly once but also ended on the same landmass where it began. When is it possible to do this?

Let's call an Eulerian walk that begins and ends at the same vertex an *Eulerian cycle*. Since an Eulerian cycle is a special type of Eulerian walk, any graph with an Eulerian cycle (and no isolated vertices) is connected and has at most two vertices with odd degree. But is this enough? The next exercise shows that it's not.

PRACTICE 1.4. Find a graph that has an Eulerian walk but not an Eulerian cycle.

PRACTICE 1.5. Modify our proof of Theorem 1.4 to show that the degree of *every* vertex in a graph with an Eulerian cycle is even.

It turns out that the condition in Practice 1.5, along with connectedness, is equivalent to having an Eulerian cycle.

Sometimes the inspiration for new problems is not a specific question but a call to investigate. To wit: What's up with odd-degree vertices? They play a special role in determining whether a graph has an Eulerian walk or cycle, and it's not entirely clear why. Do they have any other significance for graphs? What sort of patterns do they follow?

With this sort of question, there's no clear way forward. It's like deciding to embark on a quest of your own devising. You have to just venture forth and see what happens.

A good way to start is to collect data. Let's get our hands dirty.

PRACTICE 1.6. Examine the graphs in Figures 1.2–1.12 and gather some data about degrees. For each graph: How many vertices of each degree are there? How many vertices with odd degree are there? How many edges does the graph have? How many vertices? Once you have some data, see if you can find any patterns. Make a guess about something that's true for every graph!

There are quite a few things you might have noticed. Here's one: Each of these graphs has an even number of vertices whose degree is odd. Some graphs have zero vertices with odd degree, some have two, others have four. But none have one or three. So from the data we have conjured a concrete question: Does every graph have an even number of vertices with odd degree?

Play around with this a bit! See if you can draw a graph with one, or three, or five vertices with odd degree. If you can do it, then the answer to the question is "No." Otherwise, try to give a reason why the answer should be "Yes." When they're working, mathematicians don't know the answers to the questions in advance. They have to work both angles at once to solve the problem. Once you've given it a go (and do give it a try), look at the next paragraph for the answer.

I'm going to guess that you didn't draw a graph with an odd number of vertices with odd degree. If you did, (1) check your graph again, and (2) check your graph again—you definitely didn't draw one. How can I be so confident? I have a proof on my side.

The idea for our proof comes from the process of constructing a graph. Imagine that you're given a graph, and you want to build it up, one edge at a time. While you're doing this, keep track of the number of vertices whose degree is odd at each step; let's call this number $d$ for convenience. Before we add any edges, the degree of each vertex is 0, so $d = 0$. Now we add one edge. This makes two vertices with degree

1, and the rest still have degree 0; so $d = 2$. Now let's consider what happens when you add an edge at the next step:

- If the edge connects two vertices that had even degrees, then both vertices now have odd degrees, so $d$ increases by 2.
- If the edge connects two vertices that had odd degrees, then they now both have even degrees, so $d$ decreases by 2.
- If the edge connects a vertex with odd degree to a vertex with even degree, then the odd degree becomes even and the even degree becomes odd, so $d$ doesn't change.

In each case, $d$ is even after you add the edge. As you keep adding edges, $d$ will change, but it stays even at each step. In particular, after the final edge of the graph is added, $d$ is even; so in the graph you built, there are an even number of vertices with odd degree. Since this argument doesn't place any restrictions on the graph we start with, this is true for every graph.

This gives us another interesting nonexistence result: There is no graph with exactly three vertices whose degrees are odd. (Of course, "three" can be replaced by any odd number.) This isn't a statement that you can check by hand, since there are too many—*infinitely* many—graphs.

PRACTICE 1.7. Just as we did with Eulerian walks, we can ask whether the graphs we haven't ruled out actually exist. Devise a method to draw, for each even number $k$, a graph with exactly $k$ vertices whose degrees are odd.

Before we close out this section, let's look at one last result from Euler. Euler was famous for keeping long tables with tons of numbers and being able to sort through them to find patterns. One time, he guessed that a sum he was calculating was $\pi^2/6$ just because the decimal approximation "looked like it"—and he was right. This is a way of saying that we're about to look at something that might seem a little strange.

Go back to the data you collected in Practice 1.6, and add up the degree values for each graph. Notice any pattern?*

THEOREM 1.6. *The sum of the degrees in a graph is equal to twice the number of edges.*

*Proof.* An *incidence* is a place where a vertex connects to an edge. The degree of a vertex is the number of incidences at that vertex, so the sum of all the degrees is a graph is the total number of incidences. On the other hand, each edge has two incidences, one at either end, so there are twice as many incidences as there are edges. Since these two numbers—sum of degrees and twice the number of edges—both count the number of incidences, they must be equal. □

The technique used in this proof is called *double counting*, and it's used in a variety of situations to show that two expressions are actually

* This theorem illustrates the typical typesetting of proofs in this book: closely following the theorem, with a box □ to indicate where the proof ends.

equal. In each case, the basic idea is the same: You show that both expressions count the same thing in two different ways and therefore must be equal. Our proof of Theorem 1.6, for example, counts the number of incidences in two different ways.

Mathematicians have a sense of aesthetic about proofs. A good proof doesn't just prove that a statement *is* true; it makes you understand *why* it's true. Aesthetic is subjective, so not all mathematicians agree on the "artistic value" of a proof. Almost universally, though, mathematicians agree that this proof of Theorem 1.6 is "slick." It's a way of saying that the proof is clean, concise, and enlightening. We all aspire to write slick proofs.

PRACTICE 1.8. Use Theorem 1.6 to find a different proof that every graph has an even number of vertices with odd degree.

## 5. Induction

This section is a diversion that seems to have nothing to do with graphs—at first. It introduces a new way to prove things that's actually very useful in graph theory. (You can see this in some of the problems at the end of this chapter.) For now, let's jump into some puzzles.

### Tile patterns

*Question* 3. How many ways are there to fill an $8 \times 8$ grid with black and white tiles, like in Figure 1.13?

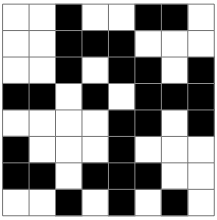

Figure 1.13: One potential tiling choice

Perhaps you're doing a home renovation project and you want to know how many options are available to you. In any case, you want to know all the possibilities.

This is a hard problem simply because there are *so many* possibilities. In situations like these, mathematicians often turn to a straightforward strategy: *Solve a simpler problem.* The hard part about this tiling problem is that it's too big to get an intuitive feeling about what's going on. To get that feeling, we can make up a similar problem that's similar but hopefully easier to solve.

In this case, we can try to solve the same problem but with a smaller grid. You're likely thinking, "Okay, I can probably manage a $3 \times 3$ or

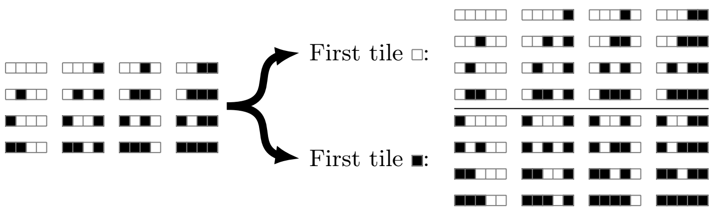


Figure 1.14: Constructing the list of all five-tile patterns using the list of all four-tile patterns

$2 \times 2$ grid," but I encourage you to go all the way: What's the answer for a $1 \times 1$ grid?

This seems a little silly—of course there are only 2 ways to fill a $1 \times 1$ grid, with a black tile or with a white tile. The benefit of doing absurdly easy examples, even though it feels a bit ridiculous, is that we can use them to gather data pretty quickly.

With that in mind, instead of jumping to a $2 \times 2$ grid with four spots to fill, let's see what happens with only two spots. It's still pretty easy: There are 4 options: □□, ■□, □■, and ■■. And what about three spots? You can list them out; there are 8 options. If you're a bit more persistent, you can list all 16 options for the $2 \times 2$ grid with four spots to fill. We can combine these numbers in a table like this:

| number of tiles | 1 | 2 | 3 | 4 |
|---|---|---|---|---|
| number of options | 2 | 4 | 8 | 16 |

The benefit of a table is that we can look for patterns. As the number of tiles increases, how does the number of options change? If we can find an answer to this, then we can use the pattern we identify to find an answer for the $8 \times 8$ grid.

Looking back at the numbers in the second row, here's one pattern: Each numbers is twice the size of the one before it. That is, 4 is $2+2$, and 8 is $4+4$, and 16 is $8+8$. So now we have a *conjecture*, an educated guess about the behavior of these numbers: Adding one tile doubles the number of options.

A conjecture is not good enough, though. How do we know the pattern doesn't just stop somewhere on down the line? This, of course, is why we look for a proof.

Let's see how it plays out with five tiles. After we place the first tile, there are four left; but we've already solved that problem—there are 16 ways to choose four tiles. So if the first tile is white, there are 16 possibilities for the remaining tiles; and there are another 16 possibilities when the first tile is black (see Figure 1.14). Altogether, then, there are $16 + 16$, or 32, options for five tiles.

And the same idea works for six tiles: There are 32 options if the first tile is white and another 32 if the first tile is black; so there are $32+32 = 64$ options for six tiles. For the same reason, there are $64 + 64 = 128$ options for seven tiles, and so on.

This is great, since we know how to calculate these numbers, but the problem is that it will take a while to get all the way up to the 64th number, which is what we need (an $8 \times 8$ grid has 64 tiles).

What we need is a short, punchy notation, and exponents have our back. In short, an exponent means "multiply this number by itself some number of times," with the number of times written to the upper right of the main number. For example,

$$2^n = \underbrace{2 \cdot 2 \cdots 2}_{n \text{ times}},$$

which is a convenient way to express repeated doubling.* So we can write that there are $2^1$ options with one tile, $2^2$ options with two tiles, $2^3$ options with three tiles, and so on, leading all the way up to 64 tiles, where there are $2^{64}$ options. This is easy to plug into a calculator: $2^{64} = 18{,}446{,}744{,}073{,}709{,}551{,}616$. To put this in perspective, this is more than the number of grains of rice that the world produces in an entire year. In fact, it would take more than 500 years to produce this much rice. So you have a lot of options for your tiling project.

The type of reasoning we used in this problem, moving sequentially from one number to the next, is called *induction*. The idea is that instead of showing something all at once, we prove it for one number, then use that to prove it for the next, then the next, and so on, building up to whatever number you want. It's a powerful idea, and it's used throughout mathematics. As we'll see, it's especially useful for proving things about graphs.

## ‖ The case of the missing tile

Let's look at a different problem. You're looking at an $8 \times 8$ square grid again, this time with a bag of L-shaped pieces in your hand; they look like . What you'd like to do is to cover the entire grid with these tiles (some possibly rotated) so that none of the tiles overlap. Unfortunately, that's impossible: Since each L-shaped tile covers three spots, you can only cover a number of spots that's a multiple of 3, and 64 isn't. But 63 is, so you wonder: Is it possible to cover every spot but one? Yes; Figure 1.15 shows one way.

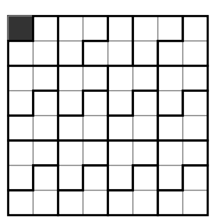

Figure 1.15: An arrangement that covers every tile but one

* This book uses $\cdot$ instead of $\times$ to express multiplication. This is common in mathematics, since $\times$ is used for other things, like the cross product in Parts II and III.

But what if one of the inside spots were missing instead?

*Question* 4. If one square is removed from an $8 \times 8$ grid, is it always possible to cover the remaining tiles with L-shaped pieces?

Like before, let's go to an easier problem. Since the shape of the grid is important in this problem, let's keep it a grid and just reduce the size to $2 \times 2$. In this case, no matter which spot is missing, we can fill in the remaining tiles.

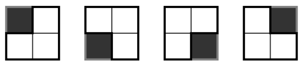

Figure 1.16: Covering all but one tile in a $2 \times 2$ grid

It turns out that $3 \times 3$ is a no-go, but maybe we can do something with $4 \times 4$.

Practice 1.9. Show that you can't cover the remaining spots in a $3 \times 3$ grid after removing one.

With a $4 \times 4$ grid, though, we can do something clever. First, we divide it into four smaller $2 \times 2$ grids, like this:

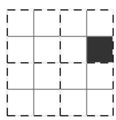

One of these four $2 \times 2$ grids has a missing tile (in this case, the upper right one). We can place an L-shaped tile to cover one tile in each of the remaining sections, like this:

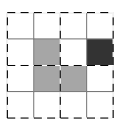

And now we have four different $2 \times 2$ grids which we know how to fill. Filling these in, we get

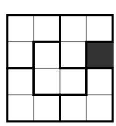

The same idea works for any missing tile in a $4 \times 4$ grid.

Practice 1.10. Use the same method to find a tiling for this grid:

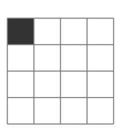

At this point, we might have an inkling that we could extend this to an $8 \times 8$ grid: divide it into four $4 \times 4$ grids, place a single L-shaped tile, and fill in each of the four sections (since we know that we can fill any $4 \times 4$ grid with a missing space). For example:

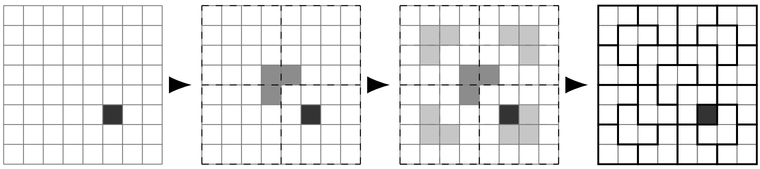

The same process works for any $8 \times 8$ grid, so the answer to our original question is: Yes, no matter which spot you choose to remain uncovered, you can cover the remaining spots with L-shaped tiles.

In fact, this same method works whenever we double the size of the grid. So we can apply the same divide-and-conquer idea to a $16 \times 16$ grid, $32 \times 32$ grid, and so on. In general, if you remove one spot from a $2^n \times 2^n$ grid, then the remaining spots can be covered by L-shaped tiles. This is another example of induction: We use the $2 \times 2$ solution to solve the $4 \times 4$ one, the $4 \times 4$ one to solve the $8 \times 8$ one, and so on.

Of course, there's a question of which tile results in the most aesthetically pleasing configuration. This is, of course, personal taste. The patterns do quickly get quite elaborate, though. Here's one tiling of a $32 \times 32$ grid:

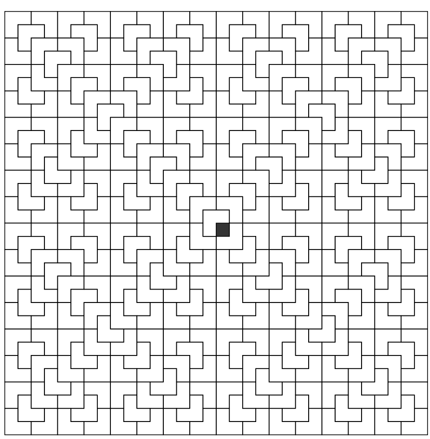

As a side note, these L-shaped pieces are pretty interesting themselves. For example, you can decompose one of these tiles into four smaller copies:

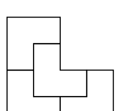

Another way to think of this is that if we scale up the L-shaped tile by a factor of 2, then we can cover it with four L-tiles of the original size.

PRACTICE 1.11. Show that an L-tile scaled up by a factor of 3 can be covered with nine L-tiles of the original size. What if the scale is 4? (The picture below might help with scratch work.)

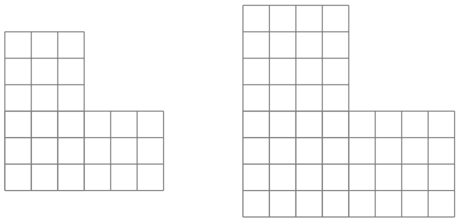

In fact, L-shaped tiles are very special: This works for any value of $n$.

THEOREM 1.7. *If an L-shaped tile is scaled up by a factor of $n$, then it can be covered with exactly $n^2$ L-shaped tiles of the original size.*

A proof of this is more challenging than the ones we've covered, but this statement, too, can be proved with the same inductive principles we've been discussing.

## ‖ Writing induction proofs

Mathematical induction, this principle of building up, is an excellent way to prove things. Just like any other proof method or technique, it doesn't work everywhere, but induction is a quite common one and definitely a useful one to have in your toolbox.

The issue is that the justifications for the previous two problems don't really rise to the standard of proofs. Essentially, what we've done is solve it for a few small cases and say "And the rest are all like these." For these problems, it's pretty believable, and yet—how do you know that all the rest actually *are* the same? Maybe somewhere along the line, something changes and things break down. We haven't actually checked them all, so how can you be sure?

Historically, mathematicians didn't care. Induction, for them, actually was checking a few small cases and saying the rest were all the same. Soon enough, though, people realized that such a carefree attitude was an easy way to run into problems exactly like the ones outlined above.

So is there nothing to be salvaged in induction, destined for the garbage heap of failed proof methods in mathematics? Of course not—it says no more than three paragraphs ago that induction is a useful proof method. The question is how to make rigorous.

Here's the idea. We have some general statement that we want to prove, like "There are $2^n$ ways to place black and white tiles into $n$ spots." This is really infinitely many statements bundled into one:

- There are 2 ways to place black and white tiles into 1 spot.
- There are 4 ways to place black and white tiles into 2 spots.
- There are 8 ways to place black and white tiles into 3 spots.
- There are 16 ways to place black and white tiles into 4 spots.

$\vdots \qquad \vdots$

A formal proof by induction can be broken down into two steps:

1. Prove that the first statement in this list is true. (The **base case**)
2. Prove that, if you know that the first $k$ statements in the list are true, then the $(k+1)$th statement is also true. (The **inductive step**)

If you can prove these two things, then you've proved every statement in the list.

Why is that? Well, you know the first statement is true because you proved it in the base case. And the inductive step says that if the first statement is true, so is the second. It also says that if the first two statements are true (which we now know they are), then the third is, too. And if that's the case, then so is the fourth, and the fifth, and the sixth, and . . .

Here's a more visual way to think about it. Imagine you have a ladder that stretches up to infinity, and at each rung is an item from the list of statements we want to prove. Step 1 of the inductive method is like saying "I can climb onto the first rung of the ladder." Step 2 is like saying "I know that once I've climbed the first $k$ rungs, I can always climb the next one." Take these together, and you can climb to any point on the ladder: Every statement in the list is true.

Let's do some test runs with this and give true inductive proofs of the two puzzles we looked at before. Here's the first.*

PROPOSITION 1.8. *There are exactly $2^n$ ways to place black and white tiles into $n$ empty spots.*

*Proof.* To prove this by induction, we start with the base case, $n = 1$. If there is only one spot to fill, then there are only two possibilities: ■ and □. ✓

Now we prove the inductive step. What we're really doing here is constructing a template for a whole set of arguments. Instead of writing each particular argument individually, since they all follow the same pattern, we use the variable $k$ to bundle them all together. Then, if you wanted to extract one specific argument, you can just plug in a particular value of $k$. This is how it pans out.

Suppose we know that Proposition 1.8 is true with 1 empty spot, 2 empty spots, and so on, all the way up to $k$ empty spots. What happens if we have $k + 1$ empty spots? We can split all of the possibilities into two camps: One where the first tile is black, and one where the first tile

* A *proposition* is fundamentally the same thing as a theorem: A proven mathematical result. The difference is a matter of emphasis. Big, important, or otherwise notable results are called theorems; other results are called propositions. The cut-off between these levels varies among mathematicians; in this book, results that play a relatively minor role in the overall story are labelled propositions, while the results that function as more major plot points are labelled theorems.

is white. If we choose a color for the first tile, then there are $k$ empty spots yet to fill—and since we know that the proposition is true for $k$ empty spots, this means there are $2^k$ possible ways to fill the remaining spots. So there are $2^k$ possibilities if the first tile is black and $2^k$ more possibilities if the first tile is white. Altogether, that gives

$$2^k + 2^k = 2 \cdot 2^k = 2 \cdot \underbrace{2 \cdot 2 \cdots 2}_{k \text{ times}} = 2^{k+1}$$

possible ways to fill $k + 1$ spots. In other words, if Proposition 1.8 is true for the first $k$ values of $n$, then it's also true for the $(k + 1)$th. ✓

Since we've proven both the base case and the inductive step, we're done! □

You can see that this proof isn't actually so different from the informal one we discussed before. This time, though, the argument focuses on the key inductive step instead of a handful of cases.

Here's the second proof.

Proposition 1.9. *If one square of a $2^n \times 2^n$ grid is removed, the remaining squares can be tiled by copies of the L-shaped tile* .

*Proof.* Figure 1.16 shows that the base case, $n = 1$, is true. ✓

For the inductive step, suppose that Proposition 1.9 is true when $n$ is 1, or 2, or anything up to $k$, and we have a $2^{k+1} \times 2^{k+1}$ grid with one square removed. By drawing the two lines that split the grid in half vertically and horizontally, we can divide the grid into four square $2^k \times 2^k$ regions. The missing tile lies in one of these regions, and we can place one L-shaped tile so that it covers one square in each of the three other regions. Each of these regions is now a $2^k \times 2^k$ grid with one square missing (either because it was originally missing or because it's covered by the one tile we've placed). We already know that any grid of this size can be tiled by copies of —that's how we started this paragraph. And once we've covered each of the four regions, we've covered the entire $2^{k+1} \times 2^{k+1}$ grid. So Proposition 1.9 is true when $n$ is $k + 1$, as well, which finishes the inductive step. ✓ □

If you look back at page 15, where we proved that there's always an even number of vertices whose degree is odd, you'll notice that this proof, too, uses the inductive method. We just didn't have terminology like "base case" and "inductive step" at that point.

That's basically it—you're licensed to induct! It's nice to take a step back and realize that the induction technique solved three fairly different problems: One of them (Proposition 1.8) asked "How many?", another (Proposition 1.9) asked "Is it possible?", and the last (on odd degrees) was a mix of both. You'll see, especially in the exercises at the end of this chapter, that induction can be used in many different settings.

One more note: Sometimes you need to prove the first handful of statements in the list before you can prove the inductive step. In other words, you need to prove multiple base cases, sort of like giving the ladder-climber a boost. To see something like this in action, flip to problem 1.11.

## ‖ Why formal proofs?

It may seem that writing out a formal induction proof is more trouble than it's worth. After all, we already knew that Propositions 1.8 and 1.9 were true, so why go to the trouble? Let's go through a few more patterns.

Example 1.10. If you take a circle and mark a few points evenly around its boundary and then connect all these points, how many regions are there? (See Figure 1.17.) The first few answers are $1, 2, 4, 8, 16$. Notice a pattern? ♢

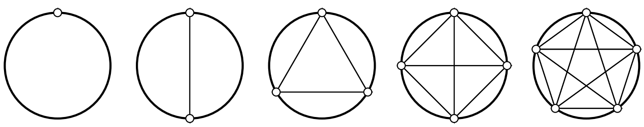

1 2 4 8 16

Figure 1.17: Circle divisions

Example 1.11. Add the first few odd numbers together:

$$1 = 1 \qquad 1+3+5 = 9$$
$$1+3 = 4 \qquad 1+3+5+7 = 16$$

You get the square numbers $1^2, 2^2, 3^2, 4^2, \ldots$ ♢

Example 1.12. A *prime number* is a counting number greater than 1 that is only a multiple of 1 and itself. Since every number is at least a multiple of 1 and itself, prime numbers can be thought of as the "indivisible" numbers. The first few prime numbers are 2, 3, 5, 7, 11, 13, 17, 19, and 23. Also, each of the numbers 31, 331, 3331, 33331, 333331, and 3333331 is prime. ♢

Example 1.13. Substitute numbers for $n$ in the expression $n^2 - n + 41$. The first few values are $1^2 - 1 + 41 = 41$, $2^2 - 2 + 41 = 43$, and $3^3 - 3 + 41 = 47$; it then continues $53, 61, 71, 83, 97, 113, \ldots$, all of which are prime. ♢

These examples lead us to guess the following results:

1. If $n$ points are evenly placed on a circle and all lines between these points are drawn, the circle is divided into $2^{n-1}$ regions.
2. The sum of the first $n$ odd numbers is $n^2$.
3. The numbers $3 \cdots 31$ are all prime.
4. Every value of $n^2 - n + 41$ is prime.

Unfortunately, they're all false. Well, all but one, just to keep you on your toes. This is the issue we run into when induction is thought of as trying a few small cases and extrapolating from them—the extrapolation might be wrong. For example, 33333331 is a multiple of 17, and $41^2 - 41 + 41 = 41^2$ isn't prime.

Even if we can come up with a plausible explanation for the behavior we observe, we need to check that the explanation matches reality. It sounds reasonable to say that each region is divided into two when you add a new point to the circle's boundary, which would explain the doubling phenomenon, but a quick check of Figure 1.17 shows that's not the case. In fact, if you count the number of regions in the circle when six points are placed on the boundary, you'll notice that there are 30 regions, not 32 as the statement suggests.

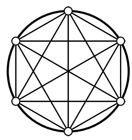

30

A formal proof helps ensure that mistakes like these aren't made, by shifting the focus from a handful of specific cases to verifiable statements and the relationship *between* the cases.

The sum of the first $n$ odd numbers, by the way, *is* always $n^2$; problem 1.10 suggests one way to prove this with induction. Speaking of which, the best way to get a handle on how induction works is to do some problems. There are a handful in the next section—go try a few!

## 6. Problems

The problems in this section and throughout the book will—by design—vary in difficulty. A handful are about the same difficulty as a practice problem; most will require more thought, some quite a significant amount over a period of time. The more difficult problems are marked with an asterisk as a heads-up. This isn't meant to scare you off but to alert you to the possibility that these problems will require more time.

Of course, any system like this has some standard disclaimers: Not all starred problems are equally difficult, and not all unstarred problems are equally difficult. You may even find that some starred problems are easier for you than some unstarred ones; in the wise words of Kurt Vonnegut, so it goes. The goal of the system is more a rough-and-ready guide than anything else.

That's enough chitchat. Let's get down to mathematics!

**1.1.** Draw all four simple graphs that have 3 vertices.

∗ **1.2.** Draw all 34 simple graphs that have 5 vertices.

**1.3.** Draw a simple graph with six vertices, three that have degree 4 and three that have degree 2.

**1.4.** How many edges does the graph $C_{72}$ have? (See Figure 1.8 for a reminder on cycle graphs.)

## ‖ Degrees

**1.5.** A simple graph is called *$k$-regular* if every vertex has degree $k$.
- (a) Is there a 5-regular graph with 7 vertices?
- (b) How many graphs with 4 vertices are 2-regular? How many with 5? How many with 51 vertices?
- (c) Find a 4-regular graph with 5 vertices and a 4-regular graph with 6 vertices.
- * (d) Describe a procedure to draw a 4-regular graph with any number of vertices.
- (e) Find a 3-regular graph with 4 vertices and a 3-regular graph with 6 vertices.
- (f) Is there a 3-regular graph with an odd number of vertices?
- * (g) Describe a procedure to draw a 3-regular graph with any even number of vertices.

**1.6.** The *degree sequence* of a graph is a list of the degrees of its vertices arranged in decreasing order. For example, the degree sequence of the graph in Figure 1.3 is $(3, 3, 2, 2)$.
- (a) Draw a simple graph with degree sequence $(3, 2, 2, 1)$.
- (b) Draw a simple graph with degree sequence $(4, 3, 3, 2, 2)$.
- (c) Is there a simple graph with the degree sequence $(4, 3, 2, 1)$?
- (d) Is there a simple graph with the degree sequence $(4, 3, 2, 1, 1)$?
- (e) Is there a simple graph with degree sequence $(3, 3, 2, 2, 2, 0)$?

**1.7.** In a group of seven people, is it possible for every person to shake hands with exactly five other people?

**1.8.** Is there a graph with 9 vertices and 9 edges so that exactly seven vertices have degree 1?

## ‖ Induction

Induction can be difficult to get the hang of when you first encounter it. Give these problems an honest try, and if you're stuck, take a gander at the hints. Only turn to them after you've worked at a problem for a while—at least 20 minutes, or maybe more—without solving it. You'll remember it better that way; plus, any hard problem requires some time to wrestle with it. An important part of mathematical thinking is making peace with uncertainty. On the other hand, induction is tricky, and it can help to read a few more example proofs. So if you're truly stuck, even after reading the hint, you can check the end of the Chapter 1 hints for full solutions to these problems.

Enough philosophy. Here are some problems:

**1.9.** Use induction to prove Theorem 1.6 in a new way. Which proof do you prefer?

**1.10.** Consider the statement: $1 + 3 + \cdots + (2n - 1) = n^2$. (In other words, the sum of the first $n$ odd numbers is $n^2$.) What does Figure 1.18 have to do with it? Provide an inductive proof of the statement.

**1.11.** Suppose you have an infinite supply of 3-cent and 5-cent coins. You can *make change* for a certain amount of money if it's pos-

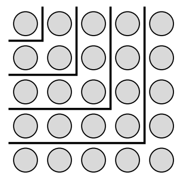

Figure 1.18: A square of circles

sible to make that exact amount of money using the coins. For example, you can make change for 18 cents (as three 5-cent coins and one 3-cent coin), but not for 4 cents. Prove that it's possible to make change for any amount of money that's at least 8 cents.

**1.12.** Remember that $K_n$ is the complete graph with $n$ vertices.
(a) If you delete one vertex along with all the edges that touch it from $K_n$, what graph is left?
(b) Use induction to prove that $K_n$ has $1+2+\cdots+(n-1)$ edges for any number $n$ that's at least 2.

**1.13.** Use induction to prove that the sum of the first $n$ numbers is $n(n+1)/2$. (In other words, prove that $1+2+\cdots+n = n(n+1)/2$.)

We can combine these two problems: The number of edges in the complete graph $K_n$ is the sum of the first $n-1$ numbers, which the next problem says is $(n-1)n/2$ edges.

**1.14.** The *Fibonacci numbers*, named after the 12th-century mathematician with the same name, start $1, 1, 2, 3, 5, 8, 13, \ldots$: Each number is the previous two added together. We write the $n$th Fibonacci number as $F_n$, so $F_1 = 1$, $F_3 = 2$, and $F_6 = 8$, for example. Unrelatedly, let's use $D(n)$ to mean the number of ways to arrange $2\times 1$ dominoes in a $2\times n$ grid. For example, $D(2) = 2$: There are two ways to put dominoes in a $2\times 2$ grid, 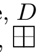 and . And $D(3) = 3$; the three ways are , , and 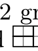.
(a) Find $D(4)$ and $D(5)$.
(b) Guess the value of $D(6)$ and see if you're right. Guess a pattern for the numbers $D(n)$.
(c) Show that there are exactly $D(n-1)$ different arrangements of dominoes in a $2\times n$ grid where the leftmost domino is vertical.
(d) Show that there are exactly $D(n-2)$ different arrangements of dominoes in a $2\times n$ grid that begin with .
(e) Prove that $D(n) = D(n-1) + D(n-2)$.
(f) What's the relationship between the domino numbers and the Fibonacci numbers?

**1.15.** We let $s(n)$ denote the number of sequences of $n$ letters, either $a$ or $b$, which do not have consecutive $a$'s. So $s(3) = 5$, since *aba*, *abb*, *bab*, *bba*, and *bbb* are all the 3-letter sequences of $a$'s and $b$'s that don't have two $a$'s next to each other.
(a) Find $s(2)$ and $s(4)$.
(b) Find a general pattern for the value of $s(n)$ and prove it.

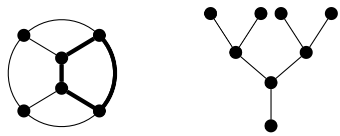

Figure 1.19: Two graphs. The left contains several cycles, one of which is highlighted; the right is a tree

## ‖ Trees

Trees are a special type of graph that are often useful in constructing algorithms for graphs. They form a kind of backbone for any connected graph, and they have many interesting properties that are far from obvious just by looking at the definition.

Definition 1.14. A *cycle* in a graph is a walk with no repeating edges whose starting and ending vertices are the same. A graph is called a *tree* if it is both connected and does not contain a cycle. (See Figure 1.19.)

These graphs will play a very important role in Section 3.1.

**1.16.** How many trees with 4 vertices are there? (Consult Figure 1.10.) How many trees with 5 vertices?

**1.17.** Draw all six trees with 6 vertices.

∗ **1.18.** Show that any graph where the degree of every vertex is at least 2 contains a cycle. Why does this mean that any tree has at least one vertex with degree 1? (A degree-1 vertex is descriptively called a *leaf*.)

**1.19.** Use induction to prove that any tree with $n$ vertices has exactly $n-1$ edges.

∗ **1.20.** Prove that any connected graph (1) is a tree or (2) you can delete some edges from the graph to make a tree that uses every vertex. (This is called a *spanning tree* of the graph.)

**1.21.** Prove that any connected graph with $n$ vertices and $n-1$ edges is a tree.

**1.22.** Prove that any graph with $n$ vertices and $n-2$ or fewer edges is not connected. (In other words, any connected graph has at least $n-1$ edges.)

This is some weird stuff: You know whether a connected graph has a cycle just by counting the number of edges it has. It seems like the arrangement of the edges should play a larger role, but, surprisingly, it doesn't.
For a challenge, prove that:

∗ **1.23.** Any connected graph with $n$ vertices and $n$ edges has exactly one cycle.