

Introduction to Higher Mathematics: Combinatorics and Graph Theory

Melody Chan (modified by Joseph Silverman)

©2017 by Melody Chan
Version Date: June 27, 2018

Contents

1	Combinatorics	1
1.1	The Pigeonhole Principle	1
1.2	Putting Things In Order	3
1.3	Bijections	5
1.4	Subsets	6
1.5	The Principle of Inclusion-Exclusion	11
1.6	The Erdős-Ko-Rado Theorem	15
	Exercises	17
2	Graph Theory	24
2.1	Graphs	25
2.2	Trees	29
2.3	Graph Coloring	32
2.4	Ramsey Theory	38
2.5	The Probabilistic Method	42
	Exercises	45
A	Appendix: Class Exercises	49
A.1	Randomized Playlists	49
A.2	Catalan Bijections	50
A.3	How Many Trees?	51

Chapter 1

Combinatorics

Combinatorics is the study of finite structures in mathematics. Sometimes people refer to it as *the art of counting*, and indeed, counting is at the core of combinatorics, although there's more to it as well.

1.1 The Pigeonhole Principle

Let us start with one of the simplest counting principles. This says that if we put 41 balls into 40 boxes any way we want, then there is some box containing at least two balls. More generally:

Theorem 1.1. (Pigeonhole Principle)¹ *Let b and n be positive integers with $b > n$. If we place b balls into n boxes, then some box must contain at least two balls.*

If the Pigeonhole Principle seems obvious, that's good. But it is worth pausing for a moment and asking ourselves how we would *prove* such a statement, i.e., how could we convince a doubtful person beyond any doubt whatsoever using a reasoned argument?

Proof. We prove the Pigeonhole Principle by contradiction. Suppose that we place b balls into n boxes, but that each box contains at most one ball. Since each box contains at most one ball, and there are at most n boxes, we have most n balls in total. But the number of balls b was assumed strictly greater than n , so we have arrived at a contradiction. Therefore some box contains at least two balls. $\square$

Theorem 1.2. (Pigeonhole Principle, general version) *Suppose b, k, n are positive integers with $b > nk$. If we place b balls into n boxes, then some box must contain at least $k + 1$ balls.*

Proof. We leave the proof as an exercise; see Exercise 1.1. $\square$

¹The Pigeonhole Principle is also known as the *Box Principle*. In German, it's the *Schubfach-Prinzip*, and the French version is the *Principe de Tiroirs*.

This silly-sounding principle is actually quite useful. We just need to free our minds from thinking literally about boxes and balls, or pigeons and pigeonholes, and to think more generally about *placing a finite set of things into a fixed number of categories*. We recall that the set of *integers* is the set of whole numbers $\dots, -2, 1, 0, 1, 2, \dots$.

Example 1.3. No matter how you choose 6 positive integers, two of them will differ by a multiple of 5.

Proof. We observe that two numbers differ by a multiple of 5 precisely when their *remainders* upon division by 5 are the same. There are only 5 possible remainders

$$0, 1, 2, 3, 4,$$

so the Pigeonhole Principle implies that given any six numbers, some pair of them have the same remainder. $\square$

Example 1.4. Given any five points in a square of side length 1, some two of them are at distance < 0.75 of each other.

You might enjoy playing with this example. For example, if you are allowed to choose only *four* points instead of five, you can put them at the four corners of the square. But if you have to choose five points, we're claiming that two of them will be within < 0.75 of each other no matter what. Try it!

How can we prove this using the Pigeonhole Principle? What are the pigeons? What are the pigeonholes? What we would really like is to be able to identify four reasonably small regions of the square that cover the entire square. Then at least two of the five given points lie in one of the regions, so are reasonably close to each other. Here are the details.

Proof. Divide the square into four smaller squares of side length $1/2$, as illustrated in Figure 1.1. Given five points, the Pigeonhole Principle implies that two of the points lie within one of the small squares. Hence the distance between these two points is at most the length of the diagonal of the small squares, which is $\sqrt{2}/2 \approx 0.707$. $\square$




Figure 1.1: A square subdivided into four regions

1.2 Putting Things In Order

Suppose n is a positive integer.

How many different ways can we put n distinct symbols in some order?

Let's try writing down all the ways to order $1, \dots, n$ for some small values of n . Writing down small examples is often a good strategy to get started in solving a math problem.

n	Different ways to order $1, 2, \dots, n$	Number of ways
$n = 1$	1	1
$n = 2$	12, 21	2
$n = 3$	123, 132, 213, 231, 312, 321	6
$n = 4$	1234, 1243, 1324, 1342, 1423, 1432, 2134, 2143, 2314, 2341, 2413, 2431, 3124, 3142, 3214, 3241, 3412, 3421, 4123, 4132, 4213, 4231, 4312, 4321	24

We get 1, 2, 6, 24, $\dots$

Definition. Let n be a positive integer. The *factorial* of n is the product

$$n \cdot (n - 1) \cdot \dots \cdot 3 \cdot 2 \cdot 1.$$

It is denoted by the symbol $n!$.

We observe that $n \cdot (n - 1)! = n!$, which incidentally gives a good reason for why we define $0! = 1$. Our evidence suggests:

Proposition 1.5. *Let n be a positive integer. There are $n!$ ways to put n distinct symbols in order.*

Informally, we might argue as follows: There are n different ways to pick the first symbol. Having done that, there are $n - 1$ unused symbols, so there are $n - 1$ ways to pick the second symbol. Proceeding in this way, we conclude that there are $n(n - 1) \cdots 1 = n!$ ways to put all n symbols in order.

That's pretty good, but it's not quite a proof. What does "proceeding in this way" really mean? How did we know to *multiply* all of the numbers from 1 to n , instead of, say, adding them, or doing something even weirder with them?

Proof. (Proof of Proposition 1.5) By induction on n . When $n = 1$, there is only one ordering possible, so the proposition holds.

Next, assume that $n > 1$ and that we have already shown that there are $(n - 1)!$ ways to put $n - 1$ distinct symbols in order. Now we wish to count the ways to put n distinct symbols in order. Let's group these ways according to what the *first* symbol is. There are n such groups. Within each group, we may chop off the first symbol

and notice that the remaining strings constitute all ways to order the remaining $n - 1$ symbols. Thus each group has size $(n - 1)!$. So, there are

$$\underbrace{(n - 1)! + \cdots + (n - 1)!}_{n \text{ copies of } (n - 1)!} = n \cdot (n - 1)! = n!$$

ways to put n symbols in some order. $\square$

We illustrate the inductive step of the proof. Suppose we already know that there are 6 ways to put 3 symbols in order, and we wish to deduce that there are 24 ways to put that symbols 1, 2, 3, 4 in order. We divide the orderings according to what the first symbol is. There are 4 groups, and each group has size 6.

orderings starting with 1 $\rightarrow$	1 234, 1 243, 1 324, 1 342, 1 423, 1 432
orderings starting with 2 $\rightarrow$	2 134, 2 143, 2 314, 2 341, 2 413, 2 431
orderings starting with 3 $\rightarrow$	3 124, 3 142, 3 214, 3 241, 3 412, 3 421
orderings starting with 4 $\rightarrow$	4 123, 4 132, 4 213, 4 231, 4 312, 4 321

Let's generalize:

Proposition 1.6. (Multiplication rule) *Let $n \geq 1$ and let $m_1, \dots, m_n$ be positive integers. Suppose we have a collection $\mathcal{C}$ of strings of symbols of length n , such that:*

- *there are m_1 symbols that may appear as the first symbol of a string in $\mathcal{C}$;*
- *for each $i = 2, \dots, n$, any initial substring consisting of $i - 1$ symbols appearing in $\mathcal{C}$ may be extended in exactly m_i ways to a substring of i symbols appearing in $\mathcal{C}$.*

Then

$$|\mathcal{C}| = m_1 \cdot m_2 \cdots m_n.$$

Proof. We leave the proof as an exercise; see Exercise 1.2. $\square$

The statement of Proposition 1.6 sounds more complicated than it should. It basically says that if you need to make a sequence of n choices, and the number of choices you have at each juncture doesn't depend on past choices, then all in all, you should multiply the number of choices you have at each juncture. To see this, think of the strings of symbols as all of the possible written records of the choices you made.

Example 1.7. A *binary string* is a sequence of 0s and 1s. How many binary strings of length k are there?

Solution. By the multiplication rule, it's

$$\underbrace{2 \cdot 2 \cdots 2}_{k \text{ factors of } 2} = 2^k.$$

Example 1.8. Suppose that there are 100 students in a class. How many ways are there to choose a president, a vice-president, and a treasurer, if we specify that no student may hold more than one position?

Solution. There are 100 choices of president. Having chosen a president, there are 99 choices of vice-president. Having made those choices, there are 98 choices of treasurer. There are $100 \cdot 99 \cdot 98 = 970200$ choices in all. Notice, by the way, that we can write this answer as

$$100 \cdot 99 \cdot 98 = \frac{100!}{97!}$$

using a whole lot of cancellation.

Generalizing:

Proposition 1.9. Suppose we have n distinct symbols, and let k be a positive integer with $k \leq n$. There are

$$\underbrace{n \cdot (n-1) \cdots (n-k+1)}_{k \text{ factors}} = \frac{n!}{(n-k)!}$$

ways to choose k of these symbols and place them in some order.

Proof. This is a special case of Proposition 1.6. □

Definition. The number

$$P(n, k) = \underbrace{n \cdot (n-1) \cdots (n-k+1)}_{k \text{ factors}} = \frac{n!}{(n-k)!}$$

appearing in Proposition 1.9 is called the *number of permutations of n objects taken k at a time*.

1.3 Bijections

This is as good a time as any to mention that bijections between finite sets are a combinatorialist's favorite tool for counting.

Definition.

1. A function $f: X \rightarrow Y$ is called an *injection* if for all $x, x' \in X$, the equality $f(x) = f(x')$ implies that $x = x'$. In words, “different elements of X are sent to different elements of Y .” Another name for such a function is *one-to-one*.
2. A function $f: X \rightarrow Y$ is called a *surjection* if for every $y \in Y$, there is some $x \in X$ with $f(x) = y$. In words, “every element of Y comes from some element of X .” Another name for such a function is *onto*.

3. A function $f: X \rightarrow Y$ is called a *bijection* if it is both injective and surjective.

Drawing some illustrations of injective/surjective/bijective functions convinces us that *bijections are one-to-one correspondences*. Alternatively, suppose f is both injective and surjective. Surjectivity implies that every element $y \in Y$ is the image under f of *at least one* $x \in X$; injectivity implies that every element $y \in Y$ is the image under f of *at most one* $x \in X$. Hence every element $y \in Y$ is the image under f of *exactly one* $x \in X$, which is the definition of one-to-one correspondence.

In particular, we see that compositions of bijections are bijections. It is a good exercise to convince yourself, and then prove, the following basic fact about bijections:

Proposition 1.10. *A function $f: X \rightarrow Y$ is a bijection if and only if f has a two-sided inverse, i.e., if and only if there is a function $g: Y \rightarrow X$ such that $f \circ g = 1_Y$ and $g \circ f = 1_X$.*

Now here is why bijections are so useful in combinatorics.

Theorem 1.11. (Bijection Correspondence Principle)

If X and Y are finite sets in bijective correspondence then they have the same number of elements.

Bijection correspondence *simply means that there exists a bijection from one set to the other*.

The claim made in Theorem 1.11 is obvious, right? Then we should be able to prove it.

Proof. Let $|X| = n$. What does that actually mean? Presumably, it means that there is a bijection between X and the set $\{1, \dots, n\}$. If so, then by composing that bijection with a bijection between X and Y , we obtain a bijection between Y and $\{1, \dots, n\}$, so $|Y| = n$ also. $\square$

One might view the set $\{1, \dots, n\}$ as the reference n -element set sitting on a shelf at the National Bureau of Standards. The point is that one way to count the elements in a set Y is simply to establish a bijection between Y and some set X whose cardinality you already know. This comes up constantly, and is extremely satisfying! In fact, one could argue that this is the best and most explicit way to count things—the various other methods that we discuss are but tricks that we use when we are unable to produce an explicit bijection.

We will see many examples, starting straightaway in the next section.

1.4 Subsets

Definition. For $n \geq 0$ an integer, we denote the set $\{1, \dots, n\}$ by $[n]$. Somewhat strangely, this convention includes $[0] = \emptyset$.

We ask:

How many subsets of an n -element set are there?

Solution. We claim that the answer is 2^n . By the bijective correspondence principle, it suffices to count subsets of our favorite n -element set, so we'll use $[n] = \{1, \dots, n\}$.

Note that there is a natural bijection between the subsets of $[n]$ and binary strings of length n . Namely, given $A \subseteq [n]$, let e_A be the binary string whose i^{th} digit is 1 if $i \in A$ and 0 if $i \notin A$. Thus one might say that we are viewing e_A as a "membership-recording string." We leave it to you to explain why this is a bijection. Finally, we are done by Example 1.7, because we already counted the binary strings of length n and found that there are 2^n of them.

Definition. Let X be any set. The *power set* of X is the set of all subsets of X ; it is denoted $\mathcal{P}(X)$.

For example if $|X| = n$ then we showed above that $|\mathcal{P}(X)| = 2^n$. We next consider the problem of counting subsets of a specified size. Fix $k \geq 0$.

How many subsets of an n -element set have size k ?

As always, we start with an example. How many 3-element subsets of a 4-element set are there? Let's list them:

$$\{1, 2, 3\}, \{1, 2, 4\}, \{1, 3, 4\}, \{2, 3, 4\}.$$

So there are four of them. One way to compute that number, four, is as follows. We already know that there are $4 \cdot 3 \cdot 2 = 24$ ways to choose 3 elements of the 4 *and place them in order*. Now, consider the function

$$\left\{ \begin{array}{l} \text{ordered triples } (a, b, c) \text{ of distinct} \\ \text{elements of the set } \{1, 2, 3, 4\} \end{array} \right\} \longrightarrow \left\{ \begin{array}{l} \text{3-element subsets of} \\ \text{the set } \{1, 2, 3, 4\} \end{array} \right\}$$

that forgets the ordering. (This kind of function really is referred to as a *forgetful function*, since its job is to forget some stuff, in this case, to forget the ordering of the triple.) Here it is, written in full detail:

(1,2,3)	(1,2,4)	(1,3,4)	(2,3,4)
(1,3,2)	(1,4,2)	(1,4,3)	(2,4,3)
(2,1,3)	(2,1,4)	(3,1,4)	(3,2,4)
(2,3,1)	(2,4,1)	(3,4,1)	(3,4,2)
(3,1,2)	(4,1,2)	(4,1,3)	(4,2,3)
(3,2,1)	(4,2,1)	(4,3,1)	(4,3,2)
↓	↓	↓	↓
{1, 2, 3}	{1, 2, 4}	{1, 3, 4}	{2, 3, 4}

In general, each set downstairs has $6 = 3 \cdot 2 \cdot 1$ elements sitting above it, where 6 represents the number of ways to order the three given symbols. So we compute the answer 4 as the solution to the problem $? \cdot 6 = 24$. More generally:

Definition. Let $k, n \geq 0$ be integers with $k \leq n$. We write $\binom{n}{k}$ to denote the number of k -element subsets of an n -element set. Such expressions are called *binomial coefficients*, for reasons that will be clear soon. We note that this quantity is also called the number of combinations of n things taken k at a time and is sometimes denoted by $C(n, k)$, but these days the notation $\binom{n}{k}$ is considered more stylish.

Proposition 1.12. *We have*

$$\binom{n}{k} = \frac{P(n, k)}{k!} = \frac{n!}{k!(n-k)!}, \quad (1.1)$$

where we recall that $P(n, k)$ denotes the number of ways to choose k elements from an n -element set and order them.

Proof. We leave the proof as an exercise; see Exercise 1.15. $\square$

Corollary 1.13. *For each $n \geq 0$, we have the following identity:*

$$2^n = \binom{n}{0} + \binom{n}{1} + \cdots + \binom{n}{n-1} + \binom{n}{n}.$$

Proof. Both sides count the number of subsets of $[n]$. $\square$

There is so much to say about the binomial coefficients! Let's start by arranging them visually on the page in a big triangle, where the n th row goes from $\binom{n}{0}$ to $\binom{n}{n}$.

$$\begin{array}{ccccccc} & & & \binom{0}{0} & & & \\ & & \binom{1}{0} & & \binom{1}{1} & & \\ & \binom{2}{0} & & \binom{2}{1} & & \binom{2}{2} & \\ \binom{3}{0} & & \binom{3}{1} & & \binom{3}{2} & & \binom{3}{3} \end{array}$$

This arrangement is commonly known as *Pascal's triangle*, although it was independently discovered and described by many mathematicians (and poets!) centuries before Pascal's book about it was published in 1665. A portion of Pascal's triangle is illustrated in Figure 1.2.

Looking at Figure 1.2, we make some observations. Note that these are just guesses for the moment:

Conjecture 1.14. *Pascal's triangle has the following properties:*

- (a) *Each number, besides all the 1s, is the sum of the two numbers directly above.*
- (b) *The triangle is left-right symmetric.*
- (c) *Every number down the middle is even.*
- (d) *Every row is unimodal, meaning it first increases and then decreases.*
- (e) *If n is prime, then each $\binom{n}{k}$, other than $\binom{n}{0}$ and $\binom{n}{n}$, is a multiple of n .*

					1					
				1		1				
			1		2		1			
		1		3		3		1		
		1	4		6		4		1	
	1		5		10		10		5	
	1	6		15		20		15		6
	1	7	21		35		35	21	7	1
1		8	28	56		70		56	28	8
1	9	36	84	126	126	84	36	9	1	

Figure 1.2: The first ten rows of Pascal's triangle.

Proof. We prove the first statement of Conjecture 1.14 and leave the other parts as homework; see Exercise 1.16. Conjecture 1.14(a) claim amounts to proving that the identity

$$\binom{n}{k} = \binom{n-1}{k} + \binom{n-1}{k-1}$$

is true for all integers $n \geq 1$ and for all integers k satisfying $1 \leq k \leq n$.

We could hash this out using the expression (1.1), but we'll leave that for you, and we'll take a more combinatorial, less computational, approach. We are going to give a "bijective proof." We consider the set X of k -element subsets of $[n]$. For example, if $n = 4$ and $k = 2$, then

$$X = \{\{1, 2\}, \{1, 3\}, \{1, 4\}, \{2, 3\}, \{2, 4\}, \{3, 4\}\}.$$

We partition the set X into two subsets X_1 and X_2 , where:

- X_1 consists of the sets that do not contain n ;
- X_2 consists of the sets that do contain n .

So in our example with $n = 4$ and $k = 2$, we have

$$X_1 = \{\{1, 2\}, \{1, 3\}, \{2, 3\}\} \quad \text{and} \quad X_2 = \{\{1, 4\}, \{2, 4\}, \{3, 4\}\}.$$

We note that X_1 is in bijection with the set of k -element subsets of $[n-1]$, since, the sets in X_1 are exactly the subsets of $[n-1]$. And X_2 is in bijection with the set of $(k-1)$ -element subsets of $[n-1]$, since each set in X_2 is formed by taking a $k-1$ -element subset of $[n-1]$ and throwing in the element n . Alternatively, we can define a bijection $X_2 \rightarrow [n-1]$ by taking a set in X_2 and *removing the element n* . All in all, we find that


$$\binom{n}{k} = |X| = |X_1| + |X_2| = \binom{n-1}{k} + \binom{n-1}{k-1},$$

which is exactly what we wanted to show. $\square$

Sometimes binomial coefficients are used to count things in unexpected ways:

Example 1.15. You have n numbered boxes. How many ways are there to put k indistinguishable balls into them? Note that “indistinguishable” means that we care only about how many balls go in each box; the balls themselves are identical.

Solution. Convince yourself that you may count instead sequences consisting of k symbols that look like $\bullet$ and $n - 1$ symbols that look like $|$. The circles represent the balls and the vertical bars represent box-separators. Here is an example of a sequence of symbols and how it translates into balls in boxes:



The number of such sequences of circles and bars is $\binom{n-1+k}{k}$, since we can view the problem as that of selecting k of the $n-1+k$ symbols to be circles, and then the other $n-1$ symbols must be bars.

The method that we have used to solve this problem is sometimes referred to as the “stars and bars” method (although we have used circles instead of stars), but at least one mathematician from the Midwest has suggested that it should be called the “cows and fences” method.

Example 1.16. Same question as in Example 1.15, but now we require that every box contains at least one ball.

Solution. If $n > k$, then the answer is clearly 0, so we assume that $n \leq k$. Then we need to start by using k of the balls to put one in each box. This leaves $k-n$ balls to be distributed in any way we please, so we are in the exactly situation of Example 1.15, except we only have $k-n$ balls to put into the n boxes. So the answer is

$$\binom{n-1+k-n}{k-n} = \binom{k-1}{k-n} = \binom{k-1}{n-1},$$

where for the second equality we use the symmetry of Pascal’s triangle.

Finally we come to the binomial theorem, which will also tell us the origin of the name *binomial coefficient*. This theorem explains the pattern of the coefficients when we take powers of $x + y$:

$$\begin{aligned} (x+y)^0 &= 1 \\ (x+y)^1 &= 1x + 1y \\ (x+y)^2 &= 1x^2 + 2xy + 1y^2 \\ (x+y)^3 &= 1x^3 + 3x^2y + 3xy^2 + 1y^3 \\ (x+y)^4 &= 1x^4 + 4x^3y + 6x^2y^2 + 4xy^3 + 1y^4 \\ &\vdots \end{aligned}$$

Theorem 1.17. (Binomial theorem) *Let $n \geq 0$ be an integer. We have the polynomial identity*

$$(x + y)^n = \sum_{k=0}^n \binom{n}{k} x^k y^{n-k}. \quad (1.2)$$

Proof. One way to conceptualize a proof is to completely expand $(x + y)^n$ using the distributive law *but without using commutativity just yet*. For example, we may expand

$$(x + y)(x + y)(x + y) = xxx + xxy + xyx + xyy + yxx + yxy + yyx + yyy.$$

In such an expansion, there are 2^n terms (why?). Now collecting them together using the commutative law, we see that the term $x^k y^{n-k}$ appears a total of $\binom{n}{k}$ times: this is the number of binary sequences of length n made of k symbols x and $n - k$ symbols y . $\square$

Example 1.18. If we evaluate the formula (1.2) at $x = y = 1$, we get a new proof of Corollary 1.13.

Example 1.19. Evaluating the formula (1.2) at $x = 1$ and $y = -1$ gives

$$0 = (1 - 1)^n = \sum_{k=0}^n \binom{n}{k} (-1)^k = \sum_{\substack{k=0 \\ k \text{ even}}}^n \binom{n}{k} - \sum_{\substack{k=0 \\ k \text{ odd}}}^n \binom{n}{k}.$$

This shows that the number of subsets of $[n]$ of even size equals the number of subsets of $[n]$ of odd size. Can you give a bijective proof of this fact?

There is a whole theory of *multinomial coefficients* as well, i.e., the coefficients of

$$(x_1 + \cdots + x_d)^n$$

for some fixed $d > 2$. See if you can compute the multinomial coefficients and generalize the ideas in this section. If you are stuck, look them up!

1.5 The Principle of Inclusion-Exclusion

How many students are there who are on the chess team or the volleyball team? Let X_1 be the set of students on the chess team, and let X_2 be the set of students on the volleyball team. Is the answer to our question simply $|X_1| + |X_2|$?

No, not necessarily. If there are students who are on both teams, then we will have *overcounted*. In order to correct for the overcounting, we need to subtract $|X_1 \cap X_2|$, which is the number of students on both teams. This is called the *Principle of Inclusion/Exclusion*, sometimes abbreviated as PIE.

Theorem 1.20. (Principle of Inclusion/Exclusion for two sets) *If X_1 and X_2 are finite sets, then*

$$|X_1 \cup X_2| = |X_1| + |X_2| - |X_1 \cap X_2|.$$

Proof. Well, you probably already believe it, if you're like 99% of people. The idea is that each element of $X_1 \cup X_2$ is counted the same number of times on both left and right. It is an exercise for you to make that statement entirely precise. $\square$

Similarly, if we have three sets that may have elements in common, then

$$|X_1 \cup X_2 \cup X_3| = |X_1| + |X_2| + |X_3| - |X_1 \cap X_2| - |X_1 \cap X_3| - |X_2 \cap X_3| + |X_1 \cap X_2 \cap X_3|,$$

and in general:

Theorem 1.21. (Principle of Inclusion/Exclusion, general case) *Let $n \geq 1$, and let $X_1, \dots, X_n$ be finite sets. Then*

$$|X_1 \cup \dots \cup X_n| = \sum_{\emptyset \neq I \subseteq [n]} (-1)^{|I|-1} \left| \bigcap_{i \in I} X_i \right|. \quad (1.3)$$

Proof. As in our proof sketch of Theorem 1.20, we verify that every element $x \in X_1 \cup \dots \cup X_n$ ends up contributing exactly one to the expression on the right hand side of (1.3). Let

$$J_x = \{i \in [n] : x \in X_i\}.$$

For example, if x is in X_1 and X_3 and no other sets, then $J_x = \{1, 3\}$. On the right hand side of (1.3), we consider the terms in the sum that have a non-zero from x . These are exactly the terms where I is a subset so that x is in X_i for every $i \in I$, i.e., these are exactly the terms where I is contained in J_x . And these terms are counted with sign $+1$ if $|I|$ is odd and with sign -1 if $|I|$ is even.

Thus each $x \in X_1 \cup \dots \cup X_n$ contributes exactly

$$|\{I \subseteq J_x : |I| \text{ is odd}\}| - |\{I \subseteq J_x : |I| \text{ is even and } I \neq \emptyset\}| \quad (1.4)$$

to the right hand side of (1.3), and we need to show that (1.4) is equal to 1. But we know from Example 1.19 that

$$|\{I \subseteq J_x : |I| \text{ is odd}\}| = |\{I \subseteq J_x : |I| \text{ is even}\}|,$$

and that gives the desired result, since in 1.4 we haven't included the empty set among the sets with an even number of elements. $\square$

Example 1.22. (The problem of derangements). There are n people sitting on a completely full airplane. The airline decides to wreak maximum havoc by reassigning seats in such a way that no person remains in the same seat. How many ways are there to do this? Such a reassignment is called a *derangement*.

For example, for three people, there are six possible ways to rearrange them, but we find that only two of those six permutations are derangements:

$$\begin{array}{llll} 123 \text{ (all fixed),} & 132 \text{ (1 is fixed),} & 321 \text{ (2 is fixed),} & 213 \text{ (3 is fixed),} \\ & 312 \text{ (derangement),} & 231 \text{ (derangement).} & \end{array}$$

Solution. Let us temporarily assign numbers $1, \dots, n$ to the passengers. With no restrictions, there would be $n!$ ways to reassign their seats. But that's an overcount. For each $i = 1, \dots, n$, the i^{th} passenger remains in their original seat for $(n-1)!$ of the possible $n!$ permutations, so we need to subtract those off. But now we have *undercounted*, because we've subtracted too many times if more than one passenger is in their original seat.

The Principle of Inclusion/Exclusion allows us to organize our successive over-counting and under-counting. For each $i = 1, \dots, n$, let

$$X_i = \{\text{seat reassignments in which the } i^{\text{th}} \text{ person stays in their seat}\}.$$

Note that $\cup X_i$ is the set of all seat reassignments in which *someone* stays in their seat, so the number that we seek is

$$\# \text{ of derangements of } [n] = n! - |X_1 \cup \dots \cup X_n|.$$

We are going to compute $|X_1 \cup \dots \cup X_n|$ using PIE. We note that if $J \subseteq [n]$ is a subset of size j , then

$$\left| \bigcap_{i \in J} X_i \right|$$

is the number of seat reassignments in which the people numbered by J stay in their seat. There are $(n-j)!$ ways for this to happen. Moreover, there are $\binom{n}{j}$ subsets of $[n]$ of size j . Applying PIE while grouping together the subsets of $[n]$ by size, gives that

$$|X_1 \cup \dots \cup X_n| = \sum_{j=1}^n (-1)^{j-1} \binom{n}{j} (n-j)!. \quad (1.5)$$

So the number we want is $n!$ minus the expression in (1.5), which we rewrite as

$$n! - \sum_{j=1}^n (-1)^{j-1} \binom{n}{j} (n-j)! = \sum_{j=0}^n (-1)^j \binom{n}{j} (n-j)! = n! \sum_{j=0}^n (-1)^j \frac{1}{j!}.$$

So we have proved the beautiful formula

$$\# \text{ of derangements of } [n] = n! \sum_{j=0}^n \frac{(-1)^j}{j!}.$$

If you have taken some calculus, you may recognize that the series $\sum_{j=0}^{\infty} (-1)^j / j!$ converges to $e^{-1} \approx 0.368$. So if n is large and you randomly choose a permutation of $[n]$, then there is roughly a 36.8% chance that your permutation will be a derangement.

Example 1.23. (Counting surjective functions). Let n and k be positive integers. How many surjective functions $[n] \rightarrow [k]$ are there?

Solution. This example is similar to Example 1.22. First, the multiplication rule says that there are k^n functions $[n] \rightarrow [k]$. Now for each $i \in \{1, \dots, k\}$, let X_i denote the set of functions $f: [n] \rightarrow [k]$ such that i is not in the *image* of f . In other words, there is no $x \in [n]$ such that $f(x) = i$. The answer we seek, then, is

$$k^n - |X_1 \cup \dots \cup X_k|.$$

Now let's use PIE. We note that if $J \subseteq [k]$ has size j , then

$$\left| \bigcap_{i \in J} X_i \right| = (k - j)^n.$$

So PIE says

$$\begin{aligned} \# \text{ of surjective maps } [n] \rightarrow [k] &= k^n - |X_1 \cup \dots \cup X_k| \\ &= k^n - \sum_{\emptyset \neq J \subseteq [k]} (-1)^{|J|-1} \left| \bigcap_{i \in J} X_i \right| \\ &= k^n - \sum_{j=1}^k \sum_{J \subseteq [k], |J|=j} (-1)^{j-1} (k - j)^n \\ &= k^n - \sum_{j=1}^k (-1)^{j-1} \binom{k}{j} (k - j)^n \\ &= \sum_{j=0}^k (-1)^j \binom{k}{j} (k - j)^n. \end{aligned}$$

That's the answer. An amusing side note is that if $k > n$, then since there are no surjective maps $[n] \rightarrow [k]$, this last sum must equal 0. But this fact is not at all clear if one just looks at the sum.

By the way, there is a closely related quantity that people sometimes refer to as the *Stirling numbers (of the second kind)*. (If you're interested, you can look up the Stirling numbers of the first kind.)

Definition. Let n and k be positive integers. We define the *Stirling numbers of the second kind* to be the quantities

$$S(n, k) = \frac{1}{k!} \sum_{j=0}^k (-1)^j \binom{k}{j} (k - j)^n,$$

So Example 1.23 says that the number of surjective functions $[n] \rightarrow [k]$ is $k! \cdot S(n, k)$. The Stirling number $S(n, k)$ also count something interesting, which in particular, implies that they are integers.

Proposition 1.24. The number of ways to divide a pile of n labeled objects into exactly k piles is the Stirling number $S(n, k)$.

Proof. If the k piles were labeled, this would be exactly the number of surjective functions $[n] \rightarrow [k]$, where the image of each object tells us in which pile it sits, and the surjectivity is required to ensure that we end up with exactly k piles. However, the piles are not supposed to be labeled, so we need to divide by the $k!$ ways that we can rearrange them. Hence the answer is

$$\frac{\# \text{ of surjective maps } [n] \rightarrow [k]}{k!} = S(n, k),$$

where we have used the formula for the number of surjective functions that we derived in Example 1.23. $\square$

1.6 The Erdős-Ko-Rado Theorem

In this section we describe a gem of *extremal combinatorics*. Extremal combinatorics is a very beautiful part of combinatorics. It concerns understanding combinatorial structures that are extremal, i.e., that are maximal or minimal with respect to some given property.

Fix positive integers r and n with $r \leq n$. Your task is to find the largest possible collection of r -element subsets of $\{1, \dots, n\}$ subject to the following condition: every pair of subsets that you choose must have at least one element in common.

Definition. Let $\mathcal{F}$ be a collection of r -element subsets of $[n]$. We say that $\mathcal{F}$ is an *intersecting family* if for every $S, T \in \mathcal{F}$, we have

$$S \cap T \neq \emptyset.$$

We ask:

How large can an intersecting family of r -element subsets of $[n]$ be?

For example, how large can an intersecting family of 2-element subsets of $[4]$ be? Relabeling, we can always start with the subset $\{1, 2\}$. Our next subset must share an element with $\{1, 2\}$, so after relabeling, they share 1 and the new element is 3, i.e., our second subset is $\{1, 3\}$. In order to form a third subset, we could take $\{1, 4\}$, and then every pair of subsets shares 1, or we could take $\{2, 3\}$, in which case we again get an intersecting family. And that's as far as we can go. So the answer to our question when $n = 4$ and $r = 2$ is that an intersecting family of 2-element subsets of $[4]$ contains at most 3 subsets, and that this can be achieved in two distinct ways (up to relabeling the elements),

$$\mathcal{F}_1 = \left\{ \{1, 2\}, \{1, 3\}, \{1, 4\} \right\} \quad \text{and} \quad \mathcal{F}_2 = \left\{ \{1, 2\}, \{1, 3\}, \{2, 3\} \right\} \quad (1.6)$$

We can generalize the construction of $\mathcal{F}_1$ to build an intersecting family,

$$\mathcal{F} = \{S \subset [n] : |S| = r \text{ and } 1 \in S\}. \quad (1.7)$$

It is clear that $\mathcal{F}$ is an intersecting family, since every $S \in \mathcal{F}$ shares the common element 1. Furthermore,

$$|\mathcal{F}| = \binom{n-1}{r-1},$$

since each $S \in \mathcal{F}$ has the form $\{1\} \cup I$ for some $(r-1)$ -element subset of $\{2, \dots, n\}$, and there are $\binom{n-1}{r-1}$ choices for I .

The following theorem says that the family $\mathcal{F}$ in (1.7) is the biggest that an intersecting family can be. This theorem was originally proved by Erdős, Ko, and Rado, but we give here is a later proof due to Katona in 1971.

Theorem 1.25. (Erdős-Ko-Rado theorem) *Let $n \geq 2r$ be positive integers. Any intersecting family of r -elements subsets of $[n]$ has size at most*

$$\binom{n-1}{r-1}.$$

Proof. We define a *cyclic ordering* on $[n]$, loosely, as an ordering of $1, 2, \dots, n$ around a clockwise oriented circle.² For example, there are 6 cyclic orderings of $\{1, 2, 3, 4\}$, as illustrated in Figure 1.3. In general there are $(n-1)!$ cyclic orderings of $[n]$.




Figure 1.3: The 6 cyclic orderings of $\{1, 2, 3, 4\}$

Next we say that a subset of $[n]$ is an *interval* of a cyclic order if its elements appear consecutively in the cyclic order. For example, in Figure 1.3 the set $\{3, 4\}$ is an interval of C_1, C_2, C_4 , and C_6 , while $\{1, 2, 3\}$ is an interval of all six of the C_i . We note that if $n \geq 24$, then a given cycle C has exactly n intervals of length r , which are formed by starting at each element of C and taking the set consisting of r consecutive elements. Thus the cycle C_1 in Figure 1.3 has 4 intervals of length 2, namely $\{1, 2\}$, $\{2, 3\}$, $\{3, 4\}$, and $\{1, 4\}$.

²More precisely, a cyclic ordering is an equivalence class of cyclic permutations of $\{1, \dots, n\}$, under the equivalence relation generated by a cyclic shift.

Let $\mathcal{F}$ an intersecting family of r -element subsets of $[n]$. We are going to count the elements of the following set in two different ways:

$$X = \{(C, S) : C \text{ is a cyclic order and } S \text{ is an interval of } C \text{ with } S \in \mathcal{F}\}.$$

First, we count the elements of X by first choosing a set $S \in \mathcal{F}$, which can be done in $|\mathcal{F}|$ ways, then creating a cycle by putting the elements of S in some order, which can be done in $r!$ ways, and then putting the elements of $[n]$ that are not in S in some order, which can be done in $(n - r)!$ ways. So in total, we find that

$$|X| = |\mathcal{F}| \cdot r! \cdot (n - r)!. \quad (1.8)$$

Now let's count the elements of X by first choosing the cycle C , which as we saw earlier can be done in $(n - 1)!$ ways, and then estimate how many $S \in \mathcal{F}$ can be intervals of the given C . We saw above that C contains n distinct intervals, but since we are restricting to intervals that are in the intersecting family $\mathcal{F}$, we claim that at most r members of $\mathcal{F}$ may appear as intervals of C .

To see why, suppose that I is an interval of C that is in $\mathcal{F}$. We can shift I either clockwise or counterclockwise to get other intervals of C , but if we shift I too far, we'll get an interval that has no elements in common with I , so it can't be in $\mathcal{F}$. More precisely, there are $r - 1$ clockwise shifts of I that have an element in common with I , and similarly there are $r - 1$ counterclockwise shifts of I that have an element in common with I . But a clockwise shift I' and a counterclockwise shift I'' will have no elements in common if the total shift for I' plus I'' is r or greater. Hence the intersecting family $\mathcal{F}$ contains at most r elements that can appear as an interval of C . In symbols, we have shown that

$$|X| \leq (n - 1)! \cdot r. \quad (1.9)$$

Combining the two expressions (1.8) and (1.9) for $|X|$, we get

$$|\mathcal{F}| \cdot r! \cdot (n - r)! \leq (n - 1)! \cdot r,$$

so

$$|\mathcal{F}| \leq \frac{(n - 1)! \cdot r}{r! \cdot (n - r)!} = \frac{(n - 1)!}{(r - 1)!(n - r)!} = \binom{n - 1}{r - 1},$$

which is exactly what we wanted to show. $\square$

Exercises

Section 1.1. The Pigeonhole Principle


1.1. Prove Theorem 1.2, the general version of the Pigeonhole Principle. *Hint.* Try to imitate the proof of Theorem 1.1.

1.2. Prove the multiplication rule stated in Proposition 1.6. *Hint.* Give a proof by induction on n , imitating the proof of Proposition 1.5.

1.3. Suppose there are four rows of seats in a classroom, each row containing six seats. How many students are necessary to guarantee that no matter how they seat themselves, some row will be full?

Note that if you think the answer is N , then you need to argue both that N students force a row to be full, but fewer than N students may seat themselves such that no row is full.

1.4. Prove that no matter how you choose five points in an equilateral triangle of side length 1, some pair of them will be at distance at most 0.5 from each other.



1.5. Let us say that two positive integers *have a common factor* if they are a common multiple of some integer > 1 . For example, of the three numbers 6, 10, and 15, each pair has a common factor.


- Can you find 50 numbers between $1, \dots, 100$ such that every pair of them has a common factor?
- Prove that for any 51 numbers between $1, \dots, 100$, some pair will have no common factor.

1.6. Suppose we have 12 dots arranged in a 2×6 square grid, as shown. Prove that no matter how you choose seven of these dots, some three of them are the vertices of an isosceles triangle. (Let us agree that three dots lying on a line do not form a triangle at all.)



1.7. **Explore:** Suppose now that we have a 3×3 square grid of dots. What is the smallest number N such that no matter how you choose N of the dots, some three of them form an isosceles triangle?

What about a 4×4 square grid of dots? (I don't know the answer.)



Section 1.2. Putting Things In Order

1.8. There are 100 students in a class, and we wish to choose a president, vice-president, and treasurer. The only problem is that each student has a nemesis in the class, i.e., the class is comprised of 50 pairs of nemeses, who can't stand each other.

How many ways are there to choose a president, vice-president, and treasurer, so that no two nemeses are chosen?

1.9. This exercise asks you to count numbers having certain properties.

- (a) How many 4-digit numbers are there that are not a multiple of 10?
- (b) How many 4-digit numbers are there whose digits sum to an even number?

Section 1.3. Bijections

1.10. Prove Proposition 1.10, which says that a function $f: X \rightarrow Y$ is a bijection if and only if f has a *two-sided inverse*, i.e., if and only if there is a function $g: Y \rightarrow X$ such that $f \circ g = 1_Y$ and $g \circ f = 1_X$.

1.11. How many bijections $[n] \rightarrow [n]$ are there?

1.12. Let $n > 1$ be an integer.

- (a) How many ways are there to put the numbers $1, \dots, n$ in order such that 2 occurs immediately after 1?
- (b) How many ways are there to put the numbers $1, \dots, n$ in order such that 2 occurs after 1 (but not necessarily immediately after it)?

1.13. Let $a_1, \dots, a_n$ be any positive numbers. Consider all the possible ways of writing a $+$ or $-$ sign before each a_i . Prove that at most 2^{n-1} of these expressions produce a positive sum. For example, if we take 1, 3, and 4, then

$$+1 + 3 + 4 \quad +1 - 3 + 4 \quad -1 + 3 + 4$$

are the three expressions producing a positive sum.

1.14. Let n be a positive integer. Let a_n be the number of ways to write n as a sum of odd positive integers. Let b_n be the number of ways to write n as a sum of distinct positive integers. In each case, order does not matter. For example, if $n = 7$, then we have

$$1 + 1 + 1 + 1 + 1 + 1 + 1, \quad 1 + 1 + 1 + 1 + 3, \quad 1 + 3 + 3, \quad 1 + 1 + 5, \quad 7$$

and

$$1 + 2 + 4, \quad 3 + 4, \quad 2 + 5, \quad 1 + 6, \quad 7$$

so $a_n = b_n = 5$.

- (a) Write out the various decompositions for $n = 5$ and $n = 6$, and verify that $a_5 = b_5$ and $a_6 = b_6$.
- (b) Prove that $a_n = b_n$ for all n .

Section 1.4. Subsets

1.15. Prove Proposition 1.12, which says that the number of k -element subsets of an n element set is given by the formula

$$\binom{n}{k} = \frac{P(n, k)}{k!} = \frac{n!}{k!(n-k)!}.$$

1.16. Prove the remaining parts (b), (c), and (d) of Conjecture 1.14, although you may first need to formulate more rigorous versions.

1.17. In Example 1.19 we used the binomial formula to prove that the number of subsets of $[n]$ of even size equals the number of subsets of $[n]$ of odd size. Give a bijective proof of this fact.

1.18. Let $n \geq 0$ be an integer. Let $A_n \subseteq \mathcal{P}([n])$ be the set of subsets $[n]$ that do not contain any consecutive pairs of numbers. For example,

$$A_3 = \{\emptyset, \{1\}, \{2\}, \{3\}, \{1, 3\}\}.$$

- (a) Compute $A_0, A_1, A_2, A_3, A_4, A_5$.
- (b) Make a conjecture about the sequence of sizes $|A_n|$ for all $n \geq 1$.
- (c) Prove your conjecture.

1.19. Let $n \geq 0$ be an integer. How many *nested pairs* of subsets of $[n]$ are there? In other words, compute the size of the set

$$\{(A, B) \in \mathcal{P}([n]) \times \mathcal{P}([n]) : A \subseteq B\}.$$

Section 1.5. The Principle of Inclusion-Exclusion

1.20. This is a strength-training exercise in counting. Give a brief justification for each answer. Consider functions $f: [6] \rightarrow [3]$.

- (A1) How many functions f are there?
- (A2) How many of them are surjective?
- (A3) How many of them *take the value 1 exactly four times*? By this we mean that there are exactly four numbers $i \in \{1, \dots, 6\}$ such that $f(i) = 1$.
- (A4) How many of them are nondecreasing? By *nondecreasing*, we mean that

$$f(1) \leq f(2) \leq f(3) \leq f(4) \leq f(5) \leq f(6).$$

- (B1) How many of them are surjective and take the value 1 exactly four times?
- (B2) How many of them are surjective and nondecreasing?
- (B3) How many of them take the value 1 exactly four times and are nondecreasing?
- (B4) How many of them are surjective and nondecreasing and take the value 1 exactly four times?
- (C1) How many of them are surjective *or* take the value 1 exactly four times?
- (C2) How many of them are surjective *or* nondecreasing?
- (C3) How many of them take the value 1 exactly four times *or* are nondecreasing?
- (C4) How many of them are surjective *or* nondecreasing *or* take the value 1 exactly four times?

Hint. : You may wish to use the Venn diagram in Figure 1.4 to organize your counting.

1.21. Of the numbers $1, \dots, 225$, how many of them are relatively prime to 225? (Two positive integers are called *relatively prime* if their greatest common divisor is 1.)

1.22. **Explore:** In the setting of Example 1.22, suppose one of the $n!$ possible seat reassignments is chosen uniformly at random. On average, how many people stay in their seat?




Figure 1.4: Venn diagram of functions $[6] \rightarrow [3]$ for Exercise 1.20.

Section 1.6. The Erdős-Ko-Rado theorem

1.23. Fix an integer $n > 0$. Suppose $\mathcal{F} \subseteq \mathcal{P}([n])$ is an *intersecting family* of subsets of $[n]$. In other words, for all $S, T \in \mathcal{F}$ we have $S \cap T \neq \emptyset$.

- (a) Prove that $|\mathcal{F}| \leq 2^{n-1}$.
- (b) Find an example of an intersecting family $\mathcal{F}$ as above that has size exactly 2^{n-1} .

Note: the difference between this problem and the Erdős-Ko-Rado theorem is that we no longer require that all sets in $\mathcal{F}$ have the same fixed size.

1.24. **Explore:** Let $n = 2r$. The Erdős-Ko-Rado theorem says that there are at most $\binom{n-1}{r-1}$ intersecting families of r -elements subsets of $[n]$.

- (a) If $n > 2r$, prove that every intersecting family $\mathcal{F}$ of r -elements subsets of $[n]$ has the property that there is a common element in every $S \in \mathcal{F}$. (*Hint.* Use this stronger assumption $n > 24$ in the proof of the Erdős-Ko-Rado theorem.)
- (b) We saw in (1.6) that (a) is not true for $r = 2$ and $n = 4$, since $\{1, 2\}, \{1, 3\}, \{2, 3\}$ is a family of 2 element subsets of $[4]$ where the subsets do not share a common element. Up to relabeling the elements, find all intersecting families of 3 element subsets of $[6]$. How many are there? Do the same for 4 element subsets of $[8]$.

- (c) Show that every $n = 2r$, there is a family of r -elements subsets of $[n]$ that do not share a common element. Generalize (b) by finding a formula for the number of such families, up to relabeling.
- (d) Same question as (c), but for $n < 2r$.

Chapter 2

Graph Theory

The first order of business is to define our primary object of study.

Definition. A **graph** G is a finite set V , whose elements are the *vertices* of G , together with a set E of unordered pairs of vertices, which are called the *edges* of G .¹

An *unordered pair* of vertices is simply a subset of V of size 2. We sometimes denote the vertex and edge sets of G by $V(G)$ and $E(G)$, respectively, to emphasize that they are associated to the graph G .

To save time and space, we often write an edge $\{i, j\}$ as just ij . Thus, Figure 2.1 shows a drawing of a graph G with vertex set $V = \{1, 2, 3, 4\}$ and edge set

$$E = \{\{1, 2\}, \{2, 3\}, \{2, 4\}, \{3, 4\}\},$$

or $E = \{12, 23, 24, 34\}$ for short.




Figure 2.1: A drawing of a graph G .

You can draw a graph by drawing points to represent the vertices, and line segments connecting the points to represent the edges. (Of course, a given graph may be drawn in many different ways.) Thus in some sense, you have already studied graph theory if you ever did “Connect-the-Dots” drawings as a child. But in saying that, we don’t mean to trivialize the field of graph theory. Although graphs are easy to define,

¹For those readers who are really keeping track of minutiae, we stipulate that $V \cap E = \emptyset$. In practice, it suffices to always take $V = [n]$ for some integer n and then never again have to worry about this issue.

there are plenty of deep theorems about them. We will see a few of these theorems in this chapter.

In addition, graphs are quite useful as structures on which to build models of many real-world phenomena, in particular *phenomena that have to do with pairwise interactions across a collection of agents*. For example, many models of *networks*, such as social networks, computer networks, transportation networks, etc., are built using the data structure of graphs.

2.1 Graphs

In this section we take $G = (V, E)$ to be a graph.

Definition. Let $v, w \in V$ be vertices of G , and let $e \in E$ be an edge of G .

1. The edge e is **incident** to the vertex v if e contains v , in which case we also say that v is an **endpoint** of e .
2. If vw is an edge of G , then v and w are **adjacent** vertices, or **neighbors**, of that edge.
3. The **degree** of the vertex v , denoted $\deg(v)$, is the number of edges incident to v . The degree of v is also sometimes called the **valence** of v .

Example 2.1. Suppose that there are 41 people at a party, and suppose that each pair of people either know each other or don't know each other. Then I claim that *someone at the party knows an even number of people*. How could we possibly deduce such a fact from so little information? The answer is provided by the following little gem from graph theory.

Proposition 2.2. Let $G = (V, E)$ be a graph. If every vertex of a graph has odd degree, then G has an even number of vertices. In other words,

$$\deg(v) \text{ odd for every } v \in V \implies |V| \text{ is even.}$$

Proof. Consider the set

$$H = \{(v, e) \in V \times E : e \text{ is incident to } v\}.$$

A good way to visualize such a pair (v, e) is as a “half-edge” in the graph, i.e., an edge with exactly one of its endpoints, as in the following picture:



We know that each edge has exactly two endpoints, so we can count the number of elements of H by computing

$$|H| = \sum_{e \in E} |\{v \in V : v \text{ is an endpoint of } e\}| = \sum_{e \in E} 2 = 2|E|.$$

In particular, we see that $|H|$ is even.

Alternatively, we can count the elements of H by summing over vertices v and counting how many edges are incident to v . This gives the formula

$$|H| = \sum_{v \in V} |\{e \in E : e \text{ is incident to } v\}| = \sum_{v \in V} \deg(v).$$

We are told that $\deg(v)$ is odd for every vertex v , number, whereas we know that $|H|$ is even. So we have shown that

$$\text{Even Number} = \underbrace{(\text{Odd Number}) + (\text{Odd Number}) + \cdots + (\text{Odd Number})}_{|V| \text{ terms in the sum}}.$$

The only way for this to happen is to have $|V|$ even. $\square$

Actually, we proved something stronger than the statement in Proposition 2.2, namely that in any graph, *the number of odd-degree vertices is even*, which is a consequence of the formula

$$2|E| = \sum_{v \in V} \deg(v) \quad (2.1)$$

that we proved. The proof of (2.1) is a good example of the technique of *double-counting*, that is, counting the same thing two different ways in order to deduce the equality of the two counts.

2.1.1 A Zoo of Graphs

There are many different special kinds of graphs. Here are some nice examples to keep in mind. See Figure 2.2.

- The *complete graph on n vertices* is a graph with n vertices such that there is an edge joining every pair of vertices. Typically one uses $[n]$ as the vertex set $[n]$, and then the edge set is²

$$\binom{[n]}{2} := \{S \subseteq [n] : |S| = 2\}.$$

The complete graph on n vertices is denoted K_n .

Question 2.3. How many edges are there in K_n ?

- The *empty graph* on n vertices is a graph with n vertices and no edges.

²Here we have employed a useful expansion of our combinatorial symbol notation. Thus for any set T and integer $k \geq 0$, we write $\binom{T}{k}$ to denote the collection of k -element subsets of T . With this notation, by definition we have $\left| \binom{T}{k} \right| = \binom{|T|}{k}$.

- A *cycle* on n vertices, denoted C_n , is a graph with $n \geq 3$ vertices and with edges that form an n -gon.³ A cycle of length 3 is often called a *triangle*. We note that by definition, a graph cannot contain a cycle of length 1 or 2, since an edge is required to have distinct endpoints, and two vertices can be joined by at most one edge.⁴
- A *path* on n vertices, denoted P_n , is a graph with n vertices that has two endpoints (vertices of valence 1), and every other vertex has valence 2.
- The *complete bipartite* graph $K_{m,n}$ has $m + n$ vertices which are separated into two disjoint subsets, $V = A \cup B$ with $|A| = m$ and $|B| = n$. There is an edge connecting each $a \in A$ with each $b \in B$.
- The *random graph* is denoted $\mathcal{G}(n, p)$. Okay, this isn't really a single graph, but rather it describes a way of randomly picking a graph on n vertices, i.e., it is a *probability distribution* on graphs. We take $[n]$ as our vertex set $[n]$, and we let p be a real number between 0 and 1. For each pair of distinct vertices i and j , we flip a weighted coin, so that heads comes up with probability p and tails with probability $1 - p$. If the flip comes up heads, then we insert an edge from i to j ; if it comes up tails, then we don't. All in all, we flip the coin $\binom{n}{2}$ times.

For example, if $p = 0$, then we always get the empty graph. If $p = 1$, then we always get the complete graph. As we slowly crank p up from 0 to 1, we see graphs that have more and more edges. This turns out to be very interesting. Often, if we are studying a property enjoyed by graphs, such as “connected” or “triangle-free,” there will be some threshold behavior, i.e., as we increase p , the probability that our graph $\mathcal{G}(n, p)$ has the given property suddenly spikes at a particular p depending on n , in a way that can be made precise.




Figure 2.2: Examples of different types of graphs


Just as sets have subsets, graphs have subgraphs, as in the following definition.

Definition. Let $G = (V, E)$ be a graph. A *subgraph* of G is a graph $G' = (V', E')$ with

$$V' \subseteq V \quad \text{and} \quad E' \subseteq E.$$

³A more formal definition is that a cycle is a connected graph in which every vertex has valence 2.

⁴There is a different type of graph that are called *directed graphs*, and directed graph may have 1-cycle and 2-cycles.

Figure 2.3: The graph G from Figure 2.1 and a subgraph G' of G

Note that the endpoints of an edge $e \in E'$ must be vertices in V' , no “edges to nowhere” are allowed. This condition is enforced by the requirement that G' is itself a graph. Therefore, the elements of E' are 2-element subsets of V' . An example is shown in Figure 2.3.

We now introduce a concept that is both simple and deep: **isomorphism**. It is often the case that we are interested in the *structure* of a graph, but we don't care about the exact names of the vertices. For example, at some level we would like to treat the two graphs in Figure 2.4 as the same, even though they are not literally the same, since for instance, their vertex sets are different.




Figure 2.4: Two isomorphic graphs

So we define two graphs to be isomorphic if one is obtained from the other by relabelling the vertices, as is described more precisely in the following definition.

Definition. Let $G = (V, E)$ and $G' = (V', E')$ be graphs. An *isomorphism* $G \rightarrow G'$ is a bijection $f: V \rightarrow V'$ with the property that for $i, j \in V$,

$$ij \in E \quad \text{if and only if} \quad f(i)f(j) \in E'.$$

(We recall that ij is shorthand for $\{i, j\}$, and similarly $f(i)f(j)$ is shorthand for $\{f(i), f(j)\}$.)

If you are given two very small graphs, such as the ones in Figure 2.4, it is easy to tell whether they are isomorphic. But for graphs with many vertices and edges, the general problem of determining whether two graphs are isomorphic is very difficult. This *Graph Isomorphism Problem* is one of the deep problems in computational complexity.

Definition. Let G and G' be graphs. We say that G *contains* G' if G' is isomorphic to a subgraph of G .




Figure 2.5: Examples of Trees

We now have the terminology to define some interesting properties of graphs having to do with whether they contain a fixed subgraph. An important one for our purposes is the property of not containing any cycles.

Definition. A graph G is *acyclic* if it does not contain any cycle C_n for any $n \geq 3$. (We recall that C_1 and C_2 are not actually graphs.)

Now let us make a very simple observation: *If G is acyclic, then it remains acyclic after removing any edges.* Do you agree? This suggests that it could be interesting to study *maximally acyclic* graphs. By this we mean a graph G that is acyclic, but such that if we put in any single additional new edge connecting two of its vertices, then the augmented graph contains a cycle.

A graph that is maximal with respect to the property of being acyclic is called a *tree*. They are of fundamental importance in graph theory. We study trees in the next section, although we define them there somewhat differently.

2.2 Trees

First, we define what it means for a graph to be connected: it means that we can start at any vertex and take a walk in the graph to get to every other vertex.

Definition. A graph G is *connected* if it contains a path between any two vertices.

This is the kind of property that you would want if, say, you ran an airline. You would want your network of cities and direct flights to be connected, so that a customer could travel from any city to any other city on your flights.

A graph T is a *tree* if it is minimally connected.

In other words, a graph T is a *tree* if it is connected, but removing any edge would make it disconnected. Thus it's connected, but just barely, there's no redundancy in the connectivity.⁵ See Figure 2.5 for some examples of trees.

Let's try to understand the structure of trees. First, we observe that *trees are acyclic*. After all, if there were a cycle, then removing an edge e lying in the cycle would not disconnect the graph. (Do you see why?) In fact, we have the following characterization of trees.

⁵And of course, the graph theory name for a graph that is a disjoint union of trees is a *forest*.

Proposition 2.4. A graph G is a tree if and only if G is connected and acyclic.

Proof. If G is a tree, then it is connected by definition, and furthermore we just established that it is acyclic as well.⁶

Suppose G is connected and acyclic. We need to show that G is *minimally* connected, i.e., we need to prove that given any edge ij of G , deleting ij disconnects the graph. We give a proof by contradiction, so we suppose G remains connected if we remove some particular edge ij . This means that G contains a path P from i to j other than the edge ij . But then P together with ij is a cycle, contradicting that G is acyclic. $\square$

There are many equivalent characterizations of trees, some of which you will encounter in the exercises.

Definition. A *leaf* of a tree is a vertex of degree 1.

We now establish that trees have leaves!

Proposition 2.5. Trees have leaves. More precisely, if T is a tree on at least two vertices, then T has at least two leaves.⁷

Proof. First, let's give the idea. Suppose you are sitting around on the middle of an edge ij of T , and you want to find a leaf. What would you do? You would start walking towards i and then keep walking, and walking, until you reach a dead end. That dead end is a leaf. If you go in the other direction toward j instead, you'd get a second leaf. The point is that you cannot loop back to any previously visited vertex, because trees are acyclic.

Now we give the official proof.

The assumption that T has at least 2 vertices and is connected means that T contains at least one edge, so T contains some non-trivial paths. Let P be a maximal path in T , i.e., the path P has the property that it is not contained in a longer path of T .

Let ℓ and m be the endpoints of P . The ℓ is incident to an edge of P , as is m , so in particular $\deg(\ell) \geq 1$ and $\deg(m) \geq 1$.

We claim that the vertices ℓ and m have degree at most 1. To see why, suppose that $\deg(\ell) \geq 2$. We already know that there is a vertex $x \in P$ that is adjacent to ℓ . Suppose that there were some other $y \in T$ that is adjacent to ℓ . If $y \notin P$, then we can extend P to y . This creates a longer path, which contradicts our choice of P . On the other hand, if $y \in P$, then the edge ℓy together with the path in P from y to ℓ is a cycle, contradicting the assumption that T is acyclic. Hence $\deg(\ell) \leq 1$, and the same argument gives $\deg(m) \leq 1$. Therefore ℓ and m are leaves of T . $\square$

Corollary 2.6. Every tree on n vertices has exactly $n - 1$ edges.

⁶Actually, we asked you to establish that it is acyclic; see Exercise 2.3.

⁷Although we would be remiss if we didn't mention that a tree with infinitely many vertices may have 0 or 1 leaf!

Proof. Here's a sketch of a proof: If $n = 1$, there are no edges, so the statement is clear. For larger n , remove a leaf and its incident edge, and use induction. $\square$

How many trees are there on n vertices $\{1, \dots, n\}$?

It's great counting practice to count trees on n vertices for $n = 1, 2, 3, 4, 5, \dots$. What do you get?

If you try to do this, you may start to suspect that in order to count labelled structures, such as trees with labels $1, \dots, n$ assigned to their vertices, it is useful to be able to count unlabelled structures and to understand how many symmetries (self-isomorphisms) they have. This is deep problem. It provides motivation for understanding the idea of *symmetries* and *symmetry groups*, even if all one wants to do is to count stuff. The idea of symmetry features prominently in the Algebra unit of the course.

In any case, if you make a table for the first few values of n , you may be led to guess the following result.

Theorem 2.7. (Cayley's Tree Counting Theorem) *Let $n > 1$ be an integer. There are n^{n-2} trees on the vertex set $\{1, \dots, n\}$.*

There are many wonderful proofs of Cayley's Theorem. The most interesting of them use a little higher algebra, which is not a prerequisite for this chapter. But there is a very nice, self-contained proof that uses *Prüfer codes*.

First, we define a map

$$P_n: \{\text{trees on vertex set } [n]\} \longrightarrow \{\text{length } n-2 \text{ sequences with symbols in } [n]\}$$

according to the following simple procedure. Given a tree T , erase the leaf of T with largest label, and write down *the label of the vertex to which it was adjacent*. Keep doing that until you are left with only two vertices. Call the sequence that you wrote down the *Prüfer code* of T , denoted $P_n(T)$. See Figure 2.6 for an example of a tree T and its Prüfer code.




Figure 2.6: A tree T with Prüfer code $P_7(T) = 37117$

Our goal is to show that P_n provides a bijection between the trees on vertex set $[n]$ and the n^{n-2} sequences of length $n - 2$ with alphabet $[n]$. The key is the following innocuous looking lemma.

Lemma 2.8. Let T be a tree on vertex set $[n]$. Then the leaves of T are precisely the vertices not appearing in the Prüfer code of T .

Proof. If a number i appears in $P_n(T)$, that means that at some point there was an edge incident to i whose other endpoint was a leaf, and the leaf and edge were ripped off. So i is not a leaf.

If i is not a leaf, then it has at least two incident edges. At least one of these edges must have been ripped off at some point. Consider the exact moment when the first of these edges was ripped off; it couldn't have been that i was the leaf, since i had degree at least 2. Therefore a leaf was ripped off of the other end of an edge containing i , so i appears in $P_n(T)$. $\square$

Incidentally, Lemma 2.8 shows that if you are given a Prüfer code $P_n(T)$, but you don't know T , you can at least determine one small thing that's true about T . Namely, if v is the largest vertex not in $P_n(T)$, then the first thing that was done was to rip the leaf v off of the vertex whose label is recorded *first* in $P_n(T)$. Amazingly, this fact is all that we need to give an inductive proof of Cayley's Theorem.

Proof of Cayley's Theorem (Theorem 2.7). Given a sequence $S = s_1 \cdots s_{n-2}$ where each $s_i \in \{1, \dots, n\}$, we want to show that S is the Prüfer code of exactly one tree T . In other words, we need to show that P_n is a bijection. That claim is definitely true for $n = 2$, which furnishes the base case of our induction.

Now suppose that $n > 2$. We will assume that P_{n-1} is a bijection, and we'll deduce that P_n is also a bijection. (This is how induction works.) So, suppose that T is a tree with $P_n(T) = S$. We'd like to show that exactly one such T exists.

Let ℓ be the largest element in $[n]$ not appearing in the Prüfer code $S = s_1 s_2 \cdots s_{n-2}$. We know the following two things:

- Lemma 2.8 says that T has a leaf ℓ attached to the vertex s_1 .
- After deleting the leaf ℓ and its edge, the remaining tree T' on vertices $[n] \setminus \{\ell\}$ has Prüfer code $s_2 \cdots s_{n-2}$.

But now we are done! By induction, the tree T' with Prüfer code $s_2 \cdots s_{n-2}$ exists and is unique, and the two conditions determine a unique T from the given T' and ℓ , i.e., we must take T' and attach the leaf ℓ to the vertex s_1 . So T also exists and is unique. $\square$

This proof may annoy you a little, because we have not directly described a procedure for taking a Prüfer code and recovering a tree. This is a nice problem for you to try.

2.3 Graph Coloring

Definition. Let G be a graph and $k \geq 1$ an integer. A *proper k -coloring* of G is a function $c: V(G) \rightarrow \{1, \dots, k\}$ such that $ij \in E(G)$ implies $c(i) \neq c(j)$.

We view $\{1, \dots, k\}$ as being a set of “colors,” and c describes a way to assign a color to each vertex. With that in mind, the map $c: V(G) \rightarrow S$ is a proper coloring if adjacent vertices get different colors.

One of the most well-studied numbers that associated to a graph G is its chromatic number.

Definition. The *chromatic number* of G , denoted $\chi(G)$,⁸ is the smallest positive integer k such that G has a proper k -coloring.

Example 2.9. Let $n \geq 3$ be an integer. What is the chromatic number of the cycle C_n ? We claim that the answer is as follows:

$$\chi(C_n) = \begin{cases} 2 & \text{if } n \text{ is even,} \\ 3 & \text{if } n \text{ is odd.} \end{cases}$$

We give an informal argument, just to get a feel for chromatic numbers. Figure 2.7 shows how to 3-color C_5 and how to 2-coloring of C_4 . It is clear how to generalize these examples to 2-color C_{even} and to 3-color C_{odd} . It remains to show that we can’t do better.

First, since $n \geq 3$, we clearly need at least 2 colors, so $\chi(C_n) \geq 2$. This handles the case that n is even. Suppose that n is odd and that we have a 2-coloring of C_n . Start at a vertex that has color 1 and trace the vertices clockwise around the cycle. They must be colored alternately 12121 $\dots$. But since n is odd, when we get back to the original vertex, we’ll be stuck with adjacent vertices that are both colored 1. Hence we need a third color.




Figure 2.7: A 3-coloring of C_5 and a 2-coloring of C_4

Example 2.10. At the other extreme, the complete graph K_n on n vertices has chromatic number $\chi(K_n) = n$. This is clear, since every pair of vertices are connected by an edge, so every vertex needs its own color. On the other hand, the complete bipartite graph $K_{m,n}$ is 2-colorable, that is, $\chi(K_{m,n}) = 2$. This, too, is clear, since taking the vertices to be $V = A \cup B$ as in the definition of $K_{m,n}$, we use one color for the vertices in A and a second color for the vertices in B . See Exercise 2.11 for a converse statement.

⁸This is the Greek letter χ , pronounced “ki”. It is the first letter of the Greek word $\chi\rho\omega\mu\alpha\tau\iota\kappa\omicron\varsigma$ for chromatic.

It is useful to unpack the definition of chromatic number slowly. We observe that for any graph G :

- (a) $\chi(G) \leq k$ holds precisely when G has a proper k -coloring.
- (b) $\chi(G) > \ell$ holds precisely when G does not have any proper ℓ -coloring.

Do you agree? These will be useful statements when proving things about chromatic numbers. To illustrate, we prove the following simple statement.

Proposition 2.11. *If H is a subgraph of G then $\chi(H) \leq \chi(G)$.*

Proof. Write $k = \chi(G)$ for short, and let $c : V(G) \rightarrow [k]$ be a proper k -coloring of G . But then $c : V(H) \rightarrow [k]$ is a proper k -coloring of H , since if two vertices are adjacent in H , then they are also adjacent in G , and hence are assigned different colors by c . Therefore $\chi(H) \leq k = \chi(G)$, which is what we wanted to show. $\square$

A Personal Anecdote. In 2005, Professor Chan was a student at the Kneisel Hall Chamber Music Festival in Blue Hill, Maine. Each of the 55 or so students in the program is assigned to two *chamber groups*, which are small ensembles composed of 3-6 students. Every summer the director has to slot the 27 or so chamber groups into a daily rehearsal schedule. Of course, she can't have the same student assigned to two groups rehearsing at the same time.

Question. Do you see what this scheduling problem have to do with graph coloring?

When Professor Chan arrived that summer, a friend “volunteered” her to write a computer program to solve the Kneisel Hall scheduling problem, so that the director would no longer have to spend hours and hours overnight each summer doing the scheduling by hand. Almost every summer since then, she still uses graph-coloring algorithms to do the scheduling of all the rehearsals and coachings at Kneisel Hall—as a sort of mathematician’s “in-kind donation.”

2.3.1 Planar Graphs and the Four and Five Color Theorems

Let G be a graph.

Definition. We say that G is *planar* if it can be drawn in the plane, i.e., on a sheet of paper, without edge crossings.

We will be somewhat informal about what *exactly* a drawing is, but what you picture in your mind is likely to be correct. Namely, a drawing consists of a placement of the vertices of G at distinct points of your sheet of paper, and edges drawn in a reasonable non-intersecting way, e.g. using a finite sequence of line segments or smooth curves.⁹

It is important to note that planarity is a property of an abstract graph: it asserts the *existence* of a drawing in the plane without edge-crossings. It does not assert that

⁹For those of you who think we are being too uptight about this, we want to inform, or remind, you that *space-filling curves exist*. You can look up space-filling curves.

any particular drawing has no edge-crossings. For example, the graph K_4 is a planar graph, as illustrated in right-hand picture in Figure 2.8, even though if you were to draw K_4 the first way that pops into your head, it might well be as the left-hand picture in Figure 2.8 that has an edge-crossing.




Figure 2.8: Drawing K_4 as a planar graph

Suppose that you want to draw a map of the world, with all the countries colored in various shades. As is usual with maps, you impose the condition that *two bordering countries receive different colors*.

What does this have to do with graphs? Draw a vertex inside each country. Then, if two countries share a border, draw an edge between their vertices crossing the border. The resulting graph is called the *dual graph*, and the map-coloring question becomes a question of properly coloring the dual graph.

Mapmakers over the ages have long known, heuristically at least, that four colors are always enough.

Theorem 2.12. (The Four Color Theorem) *Every planar graph is 4-colorable.*

This theorem is a cornerstone of combinatorial mathematics. It eluded proof for several centuries. The first proof was given by Appel and Haken in 1976, but the proof relied on computers so heavily that it was regarded as not being human-checkable, which upset some people. A simpler proof, but still relying on computers, was obtained decades later by Robertson, Sanders, Seymour, and Thomas, and in 2005, Werner and Gonthier gave a formalized computer-readable proof.

The four color theorem sits at an interesting crossroads in mathematical culture, exposing differences in our perspectives on *what proof really is*. Does a proof need to be entirely human-understandable to really be a proof? After all, humans make mistakes all the time in verifying proofs. But what is a proof other than an argument used by one human being to convince another of a mathematical truth? What do you think?

Unsurprisingly, we will not prove the four color theorem in this text. But we will prove the five color theorem.

Theorem 2.13. (The Five Color Theorem) *Every planar graph is 5-colorable.*

This result was proven by Heawood in 1890 while exposing an error of a purported proof by Kempe of the four color theorem published in 1879. It is remarkable

that it then took more than a century to squeeze that fifth color out of the theorem statement!

Before proving the five color theorem, we first state a result that we will use without proof. First of all, if G is a planar graph, equipped with a fixed drawing on your sheet of paper, we define a *face* to be a region of the complement of G . This includes the infinite “outer” face that surround the entire graph.

Intuitively, if you take scissors and cut your paper along all of the edges of G , then the faces are the pieces of paper that are left over. The formula that we need gives a beautiful relationship satisfies by the vertices, edges, and faces.

Proposition 2.14. (Euler’s Polyhedron Formula) *Let G be a connected graph drawn in the plane, with v vertices, e edges, and f faces. Then*

$$v - e + f = 2.$$

Euler’s Formula is a deep theorem that you will likely discuss in the Geometry unit. By the way, a special case of this formula states that any polyhedron in $\mathbb{R}^3$ at all, e.g., a cube, a tetrahedron, a dodecahedron, etc., satisfies $v - e + f = 2$ for its vertices, edges, and faces. Do you see why this statement about polyhedra follows from Proposition 2.14?

As we now prove, a consequence of Euler’s formula is that every planar graph has a vertex of degree at most 5.

Lemma 2.15. *Every planar graph has a vertex of degree at most 5.*

Proof. We may as well prove that every *connected* planar graph G has a vertex of degree at most 5. Suppose to the contrary that G is a connected planar graph and that every vertex of G has degree 6 or more. Consider a plane drawing of G with v vertices, e edges, and f faces.

We start by double-counting the set of *edge-face incidences*, i.e. the set of ordered pairs

$$EF := \{(b, c) : b \text{ is an edge bounding the face } c\}.$$

We note that each edge bounds exactly 2 faces, and every face is bounded by at least 3 edges (since the boundary of a face is a cycle in G). This allows us to double-count the set EF :

$$\begin{aligned} |EF| &= \sum_{b \in E(G)} |\{c \in F(G) : b \text{ is an edge of } c\}| = 2|E(G)| = 2e. \\ |EF| &= \sum_{c \in F(G)} |\{b \in E(G) : b \text{ is an edge of } c\}| \geq 3|F(G)| = 3f. \end{aligned}$$

Hence

$$3f \leq |EF| = 2e, \quad \text{so} \quad f \leq \frac{2}{3}e.$$

Next we double-count the set VE of *vertex-edge incidences*, i.e. the set of ordered pairs

$$VE := \{(a, b) : a \text{ is a vertex incident to the edge } b\}.$$

Using the fact that every edge has exactly 2 endpoints and the assumption that every vertex is incident to at least 6 edges, we double-count/estimate:

$$\begin{aligned} |VE| &= \sum_{a \in V(G)} |\{b \in E(G) : a \text{ is incident to } b\}| \geq 6|V(G)| = 6v. \\ |VE| &= \sum_{b \in E(G)} |\{a \in V(G) : a \text{ is incident to } b\}| = 2|E(G)| = 2e. \end{aligned}$$

Thus

$$6v \leq |VE| = 2e, \quad \text{so} \quad v \leq \frac{1}{3}e.$$

Using the estimates $f \leq \frac{2}{3}e$ and $v \leq \frac{1}{3}e$, we find that

$$v - e + f \leq \frac{2}{3}e - e + \frac{1}{3}e = 0,$$

which contradicts Euler's Formula (Lemma 2.14). $\square$

We are ready to prove the five color theorem. Our strategy will be to delete a strategically chosen vertex v of G , inductively color the rest of G , and then show that things may be arranged so that there is still a color available for v .

Proof. (Proof of Theorem 2.13, the Five Color Theorem) We wish to show that every planar graph is 5-colorable. We induct on the number of vertices. It is obvious that a 1-vertex graph is 5-colorable; indeed, it is 1-colorable! Now suppose that $n > 1$ and that we have already proved that every graph with $n - 1$ vertices is 5-colorable. Let G be a planar graph on n vertices. We need to show that G is 5-colorable.


It is enough to prove the statement for *connected* planar graphs, since once we do that, then we can 5-color any planar graph by 5-coloring its connected components one at a time. So we may assume that G is connected.

Let $v \in V(G)$ be a vertex of degree at most 5. The existence of such a vertex is asserted in Lemma 2.15. Let G' denote the graph obtained from G by deleting v and the 5 edges incident to v . Then our inductive hypothesis says that we can find a 5-coloring $c : G' \rightarrow [5]$ of G' .

If v has degree less than 5, then we win by coloring v using a color not yet used on any of its neighbors. Note that such a color exists because there are 5 colors available, but v has at most 4 neighbors. So we may assume that $\deg(v) = 5$.

Let $v_1, \dots, v_5$ be the neighbors of v , ordered clockwise around v , as illustrated in Figure 2.9. If 2 or more of those 5 neighbors are the same color, then there is a color left over that we can use for v . So we may assume that $v_1, \dots, v_5$ use up all 5 colors, i.e., $c(v_1), \dots, c(v_5)$ are distinct.

We make the following simple observation: if the proper coloring of G' may be tweaked in some way so that at most 4 colors are used on $v_1, \dots, v_5$, then as above, we win.

Figure 2.9: A degree 5 vertex v with its 5 neighbors

We consider some subgraphs of G' . The first subgraph, which we call G'_{13} , is the largest connected subgraph of G' that contains v_1 and whose vertices are colored using only the two colors $c(v_1)$ and $c(v_3)$. Another way to view G'_{13} is to take the union of all paths that start at v_1 and whose edges have endpoints that are colored $c(v_1)$ and $c(v_3)$. We call such a path an *alternating 1-3-path* for short.

We create a new coloring for $c' : G' \rightarrow [5]$ by swapping the colors of the vertices in G_{13} and leaving all of the other vertices unchanged. We can describe c' explicitly by the rule

$$c'(v) = \begin{cases} c(v_1) & \text{if } v \in G'_{13} \text{ and } c(v) = c(v_3), \\ c(v_3) & \text{if } v \in G'_{13} \text{ and } c(v) = c(v_1), \\ c(v) & \text{if } v \notin G'_{13}. \end{cases}$$

Note that c' is also a proper 5-coloring of G' , since the vertices that are adjacent to vertices of G'_{13} have colors different from $c(v_1)$ and $c(v_3)$.

If the vertex v_3 is not in the subgraph G_{13} , then the color of c_3 did not get swapped. Thus using our new coloring c' , the vertices v_1 and v_3 have the same color, and we're done.

But what do we do if the vertex v_3 is in the subgraph G_{13} ? Note that this happens precisely when there is an alternating 1-3 path that starts at v_1 and ends at v_3 . In that case, we repeat the above construction using the subgraph G'_{24} , which is the largest connected subgraph of G' that contains v_2 and whose vertices are colored using only the two colors $c(v_2)$ and $c(v_4)$. As above, if the vertex v_4 is not in the subgraph G'_{24} , then we can swap the colors of the vertices in G'_{24} and win.

So the proof is done unless the graph G' contains both an alternating 1-3 path that starts at v_1 and ends at v_3 and an alternating 2-4 path that starts at v_2 and ends at v_4 . But it is not possible for both such paths to exist. Indeed, they would have to cross at some vertex! We have completed the proof of the five color theorem. $\square$

2.4 Ramsey Theory

There are 6 people at a party. To simplify matters, we assume that any pair of people either know each other or don't. We claim that either there are three people who all

know each other, or there are three people who all don't know each other. This is an example of *Ramsey theory*, which is a beautiful topic that we explore in this section.

We can formulate this statement in terms of coloring the edges of K_6 , the complete graph on 6 vertices.

Definition. Let G be a graph G , and let $k \geq 1$ be an. A k -edge-coloring of G is a function

$$E(G) \rightarrow [k]$$

that assigns a color to each edge of G .

As usual, you may want to visualize an edge-coloring as an assignment of k actual colors to the edges of the graph. We can now formulate the previous claim as follows:

Proposition 2.16. Every 2-edge-coloring of K_6 contains a monochromatic triangle.

Remember that a *triangle* just means a C_3 subgraph, which is the same as a K_3 subgraph. Figure 2.10 shows some examples of 2-edge colorings of K_6 with two or three monochromatic triangles as indicated. Can you find a coloring of K_6 that has only one monochromatic triangle?

Do you see why Proposition 2.16 is a rephrasing of the original claim? Let the 6 people be the vertices of K_6 , and color an edge blue if its endpoint people know each other and red if they do not. Then a blue triangle gives three people who know each other, and a red triangle give three people who don't know each other.




Figure 2.10: Some 2-edge colored K_6 graphs with indicated monochromatic triangles

Proof. For psychological ease, we label the vertices $1, \dots, 6$, and call the colors red and blue. There are five edges incident to vertex 1, and only two colors, so the Pigeonhole Principle tells us that three of these edges are the same color, which we may as well assume is red. So there are edges $1i$, $1j$, and $1k$ that are red.

If any of the three edges ij , ik , or jk are red, then we have found an edge that completes a monochromatic red triangle. On the other hand, if ij , ik , or jk are all blue, then they form a monochromatic blue triangle. So in either case we get a monochromatic triangle, and we are done. $\square$

We mention that Proposition 2.16 cannot be improved for monochromatic triangles in 2-edge-coloring, since it is possible to 2-edge-color the complete graph K_5 such that there are no monochromatic triangles. See Exercise 2.15. Thus the number 6 is some kind of threshold, since it is the smallest number of vertices which *forces* the existence of a red K_3 or a blue K_3 .

Let's think bigger. Is there any party that we could throw that's big enough to guarantee that either there are 100 people who all know each other or there are 1000 people who are all strangers to one another? Note that we definitely need to invite at least 1000 people, since otherwise we could invite 999 pairwise-strangers. But it's also clear that 1000 isn't going to be enough. What do you think?

The question is whether there is always a threshold, like the number 6 that we found earlier, or whether instead it is possible to have bigger and bigger parties such that the people who know each other, and conversely the people who don't, do not come in large clumps.

Definition. Let $k, \ell \geq 1$ be integers. We define the (k, ℓ) -*Ramsey number*, denoted $R(k, \ell)$, to be the smallest number n , if such a number exists, such that every 2-edge-coloring of K_n using colors red and blue contains either a red K_k , or a blue K_ℓ , or both.

So we have rephrased the question as: Do Ramsey numbers exist? For example, Proposition 2.16 says that $R(3, 3)$ exists and satisfies $R(3, 3) \leq 6$, and combined with Exercise 2.15 we see that $R(3, 3) = 6$.

Proposition 2.17. For all integers $k, \ell \geq 1$, the Ramsey number $R(k, \ell)$ exists.

This is rather amazing. We model the proof on our proof that $R(3, 3) \leq 6$.

Proof. We first note that for every $k \geq 1$ and $\ell \geq 1$, we have

$$R(1, \ell) = R(k, 1) = 1.$$

This is clear, since a monochromatically edge-colored K_1 is a vacuous condition, amounting to asking for the existence of at least one vertex.

Now, suppose we have established that $R(k', \ell')$ exists whenever $k' \leq k$ and $\ell' \leq \ell$, and suppose further that at least one of these inequalities is strict. We wish to show that $R(k, \ell)$ exists. (Thus, we are running an induction argument on $\mathbb{N}_{\geq 0}^2$. Do you agree with the validity of this induction?) In fact, we shall show that $R(k, \ell)$ exists and that

$$R(k, \ell) \leq R(k-1, \ell) + R(k, \ell-1). \quad (2.2)$$

In other words, we shall show that if $n = R(k-1, \ell) + R(k, \ell-1)$, then any 2-edge-coloring of K_n has a red K_k or a blue K_ℓ .

Indeed, given a 2-edge-coloring of K_n , consider the

$$n-1 = R(k-1, \ell) + R(k, \ell-1) - 1$$

edges incident to vertex 1. Among these, the Pigeonhole Principle implies that either there are $R(k-1, \ell)$ red edges or there are $R(k, \ell-1)$ blue edges. Let us suppose

that there are $R(k-1, \ell)$ red edges incident to vertex 1; the other case can be argued analogously.

We consider the $R(k-1, \ell)$ vertices at the other ends of these red edges. By the inductive hypothesis, there is either a red K_{k-1} or a blue K_ℓ among them. In the second case, we win. In the first case, we also win, since the red K_{k-1} together with the vertex 1 forms a red K_k . $\square$

Three colors. Why stop at two colors? Suppose that you are given three integers $k, \ell, m \geq 1$. Does there exist an n such that every 3-edge-coloring of K_n contains a red K_k , a blue K_ℓ , or a green K_m ?

Definition. Let $k, \ell, m \geq 1$ be integers. We define the (k, ℓ, m) -Ramsey number, denoted $R(k, \ell, m)$, to be the smallest number n , if such a number exists, such that every 3-edge-coloring of K_n using colors red, blue, and green contains a red K_k , a blue K_ℓ , or a green K_m .

Proposition 2.18. For all integers $k, \ell, m \geq 1$, the Ramsey number $R(k, \ell, m)$ exists.

Proof. We claim that $R(k, \ell, m)$ exists and satisfies

$$R(k, \ell, m) \leq R(k, R(\ell, m)).$$

The idea is to suppose that we are temporarily unable to distinguish between blue and green.

Let $n = R(k, R(\ell, m))$ and suppose we have a 3-edge-coloring of K_n . Consider it as 2-edge-coloring of K_n using colors “red” and “blue-green.” Then Proposition 2.17 says that either there is a red K_k or a blue-green $K_{R(\ell, m)}$. In the first case we win. In the second case, remembering again the difference between blue and green, we see that the original coloring has a $K_{R(\ell, m)}$ -subgraph which is 2-edge-colored blue and green. By definition of $R(\ell, m)$, this subgraph contains either a blue K_ℓ or a green K_m , and again we are done. $\square$

This same idea may be extended to arbitrarily many colors.

Proposition 2.19. Let $n \geq 1$ be an integer. For integers $m_1, \dots, m_n \geq 1$, there exists a number N with the following property: for every n -edge-coloring of the complete graph K_N , there is some $i \in \{1, \dots, n\}$ so that K_N contains a complete subgraph K_{m_i} whose edges are all colored i .

Proof. Exercise, imitating the proof of Proposition 2.18. $\square$

Definition. We write $R(m_1, \dots, m_n)$ for the smallest such N such that Proposition 2.19 is true. It is called the $(m_1, \dots, m_n)$ -Ramsey number.

What about lower bounds on Ramsey numbers? Let’s unpack this. The statement

$$R(k, \ell) > n$$

means that *there exists a 2-edge-coloring of K_n with no red K_k and no blue K_ℓ .* For example, in Exercise 2.15 you will show that $R(3, 3) > 5$ by producing an appropriate 2-edge-coloring of K_5 .

Question 2.20. What is the best lower bound for $R(3, 4)$ that you can find?

Imagine trying to find a lower bound for $R(100, 100)$. It seems that you would have to painstakingly construct a 2-edge-coloring of a truly enormous graph and then demonstrate in some way that the graph contains no monochromatic K_{100} subgraphs. That sounds hard!

It turns out that there is an amazing way to get around this difficulty. It hinges on a crucial idea that has been a theme of this text: it is possible to prove that something exists *without actually constructing it*. It's better, of course, to give a constructive proof, but sometimes that can be very hard. For example, we have seen several arguments by contradiction that show that something exists without necessarily giving a way to find it. Instead, we simply assume that it does not exist, and then derive an absurdity.

In the next section we discuss a remarkable idea, called the probabilistic method, that can sometimes be used to prove that something exists without actually constructing it. We will apply this method to prove some lower bounds on Ramsey numbers $R(k, \ell)$.

2.5 The Probabilistic Method

Suppose that we want to demonstrate that there is at least one dog living in Boston. The best thing to do would be to provide a constructive proof, i.e., actually produce a specific dog who lives in Boston. But we can't always achieve a constructive proof, even though we'd like to.

Suppose instead that we have a mechanism, such as an alien spaceship, for scooping up a random creature who resides in Boston. And suppose further—and this is the mysterious part—that we can somehow separately demonstrate that *there is a positive probability that the Boston creature who we scoop up is a dog*. Then we have proved that there exists a dog living in Boston! Why? Well, if there were no dogs living in Boston, then the aforementioned probability would be zero.

The above example was chosen to sound maximally ludicrous, even though the method itself, when applied correctly, is not. We are simply hoping to impress upon you the remarkable fact that the probabilistic method works, and indeed is an important method to working combinatorialists today.

In this section we follow the discussion on page 1 of Alon and Spencer's book *The Probabilistic Method* [1]. It leads to a lower bound for Ramsey numbers $R(k, k)$, discovered by Erdős in 1947.

Proposition 2.21. Let $k \geq 3$ be an integer. If n is a positive integer satisfying

$$\binom{n}{k} \cdot 2^{1-\binom{k}{2}} < 1, \quad (2.3)$$

then $R(k, k) > n$.

Before proving the proposition, let's assume that it's true and get a sense for how good a bound it gives for Ramsey numbers.

Corollary 2.22. We have $R(k, k) > 2^{k/2}$.

Proof of Corollary 2.22, assuming Proposition 2.21. We want to show that $n = \lfloor 2^{k/2} \rfloor$ satisfies inequality (2.3). We shall use the following useful bound on sizes of binomial coefficients:

$$\binom{n}{k} = \frac{n \cdot (n-1) \cdot (n-k+1)}{k!} \leq \frac{n^k}{k!}.$$

So, if $n = \lfloor 2^{k/2} \rfloor$, then

$$\binom{n}{k} \cdot 2^{1-\binom{k}{2}} \leq \frac{(2^{k/2})^k}{k!} \cdot 2^{1-\binom{k}{2}} = \frac{1}{k!} \cdot 2^{\frac{k^2}{2}+1-\binom{k}{2}} = \frac{1}{k!} \cdot 2^{1+\frac{k}{2}} < 1.$$

(We leave it for you to check that the final inequality is true, e.g., by induction.) This is the desired result, and we're done. $\square$

Now onward to the proof of Proposition 2.21.

Proof of Proposition 2.21. Let n be a positive integer satisfying the inequality (2.3). We want to show that there is a 2-edge-coloring of K_n with no monochromatic K_k . There are finitely many 2-edge-colorings of K_n . We pick one at random and shall demonstrate that the probability that our randomly chosen coloring contains no monochromatic K_k is positive. That's the probabilistic method.

There are a total of $2^{\binom{n}{2}}$ distinct 2-edge-colorings of K_n , since K_n has $\binom{n}{2}$ edges. Let's start by focusing our attention on one particular K_k subgraph—for instance, the one on vertices $\{1, \dots, k\}$ —and asking in how many of the 2-edge colorings of K_n is this particular K_k monochromatic. The answer is

$$2 \cdot 2^{\binom{n}{2} - \binom{k}{2}},$$

since there are two ways to monochromatically 2-color K_k , multiplied by all the different ways to 2-color the other $\binom{n}{2} - \binom{k}{2}$ edges. And of course, this number does not depend on the particular K_k subgraph of K_n that we chose.

In order to figure out exactly how many of the 2-edge-colorings of K_n have *at least one* monochromatic K_k subgraph, we could do some counting via inclusion/exclusion. But let's not do that. Instead, we make the simple observation that since the number of K_k subgraphs in K_n is $\binom{n}{k}$, the number of 2-edge colorings in which some K_k is monochromatic is *at most*

$$\left(\begin{array}{c} \text{number of } K_k \\ \text{subgraphs of } K_n \end{array} \right) \cdot \left(\begin{array}{c} \text{number of ways to 2-edge color} \\ K_n \text{ so that a specified } K_k \\ \text{subgraph is monochromatic} \end{array} \right) = \binom{n}{k} \cdot 2^{\binom{n}{2} - \binom{k}{2} + 1}.$$

We conclude that *the probability that a randomly chosen 2-edge-coloring of K_n has a monochromatically colored K_k is at most*

$$\frac{\binom{n}{k} \cdot 2^{\binom{n}{2} - \binom{k}{2} + 1}}{2^{\binom{n}{2}}},$$

which is equal to

$$\binom{n}{k} \cdot 2^{1 - \binom{k}{2}}.$$

This last quantity is smaller than 1 by hypothesis. In other words, if n satisfies the stated inequality (2.3), then *some* 2-edge-coloring of K_n has no monochromatic K_k . So $R(k, k) > n$. $\square$

To this day, the lower bounds on Ramsey numbers proven by the probabilistic method are better than those that have been given in a constructive manner, but this is a topic of current research. It is also interesting that the only values of $R(k, k)$ that we know exactly are $R(1, 1) = 1$, $R(2, 2) = 2$, $R(3, 3) = 6$, and $R(4, 4) = 18$. For higher values of k there are upper and lower bounds, for example $43 \leq R(5, 5) \leq 48$ and $102 \leq R(6, 6) \leq 165$.

In fact, something tantalizing can be gleaned from the bound that we derived for $R(k, k)$. We showed that if $n = \lfloor 2^{k/2} \rfloor$, then the probability that a random 2-edge-coloring of K_n has no monochromatic K_k is $\frac{1}{k!} \cdot 2^{1+k/2}$. Not only is that number smaller than 1 when $k \geq 3$, it is *much, much* smaller than 1 when k is large. For example, when $k = 10$ it is roughly 0.0000176, and when $k = 100$, it is smaller than 10^{-142} .

This says that a great way to construct a 2-edge coloring of K_n with no monochromatic K_k is to pick a coloring at random. You should then check that the one you picked is a good one, but it's very likely to be good. (The reason that this isn't the same as a constructive lower bound on Ramsey numbers is that it only deals with one k at a time, and not all k at once.)

With hindsight, it's reasonable that we appealed to a probabilistic method to obtain a lower bound on Ramsey numbers, since a lower bound asserts the existence of a structure that is rather *disordered*, i.e., there are no large clumps of red and no large clumps of blue. Roughly speaking, if you try to write down 2-edge-colorings in some systematic way, you will tend to find these large monochromatic clumps. In some vague philosophical sense, that's *due to the very fact that you are being systematic*. On the other hand, a random 2-edge-coloring of $K_{\lfloor k/2 \rfloor}$ is likely to be not very clumpy and thus do the trick.

This is a demonstration, at least morally, of an aptly named principle that is found all over mathematics and computer science, the so-called difficulty of **finding hay in a haystack**.¹⁰ In practice, it can indeed be hard to find hay in a haystack, despite the fact that 99.9% (or whatever) of the haystack is made of hay, since every time you systematically reach into the haystack and pull something out, you may find that it is a needle!


¹⁰This is from a 2017 talk "Finding Hay in Haystacks: The Power and Limits of Randomness" by mathematician/computer scientist Avi Wigderson that may be viewed at <https://simons.berkeley.edu/events/theoretically-speaking-avi-wigderson>.

One final remark: you might complain that we haven't really used *probability*, as opposed to just *counting*, in a meaningful way. Indeed, in our proof of Proposition 2.21, we counted the number of 2-edge-colorings of K_n not containing a monochromatic K_k and showed that this number is smaller than the total number of 2-edge-colorings of K_n . Your complaint is a good one. In this particular instance, everything we did can be phrased simply in terms of counting and without mention of probabilities. However, in many situations the perspective of probability and the probabilistic method, and the techniques they bring to bear, are very useful, even when applied in purely combinatorial situations. That is a story for another time, which you can read about in [1]. For further reading, see Bóna's excellent textbook [2] on combinatorics and Diestel's graduate level textbook [3] on graph theory.

Exercises

Section 2.1. Graphs

- 2.1. Prove that in any graph, there are two vertices having equal degree. (*Hint.* Assume not, and think about the set of vertex degrees.)
- 2.2. Let G and G' be the two 6-vertex graphs shown below. How many *isomorphisms* (see page 28) from G to G' are there? Explain your answer briefly.



Section 2.2. Trees

- 2.3. Let G be a graph that contains a cycle. Prove that G is not minimally connected, i.e., prove that it is possible to remove an edge from G and still be left with a connected graph.
- 2.4. Prove that a graph G is a tree if and only if it is maximally acyclic. In other words, G is a tree T if and only the following two facts are true:
- T is acyclic.
 - Adding a new edge connecting two vertices of T creates a cycle.
- 2.5. Draw six nonisomorphic trees on vertex set $\{1, 2, 3, 4, 5, 6\}$, and compute each of their Prüfer codes.
- 2.6. Let $n \geq 2$ be an integer. Of the n^{n-2} trees on vertex set $[n]$, prove that exactly $2n^{n-3}$ of them contain the edge $\{1, 2\}$. Here are two hints suggesting possible ways to solve this problem.
- (i) Double-count the elements of the set

$$\{(T, \{i, j\}) : T \text{ is a tree with vertices } [n] \text{ and } \{i, j\} \text{ is an edge of } T\}.$$

- (ii) Alternatively, prove carefully that a tree contains the edge $\{1, 2\}$ if and only if its Prüfer code ends in 1 or 2.
- 2.7. Let $n \geq 2$ be an integer. Of the n^{n-2} trees on vertex set $[n]$, in how many of them is vertex n a leaf?
- 2.8. Let $m > n \geq 2$ be integers. Of the n^{n-2} trees on vertex set $[n]$, how many have exactly m leaves?
- 2.9. **Explore:** Let $n \geq 3$ be an integer. Among the n^{n-2} trees on vertex set $[n]$, what is the average number of leaves?

Section 2.3. Graph coloring

- 2.10. Let G be a graph, and let D be the maximum degree of its vertices. Prove that

$$\chi(G) \leq D + 1.$$

- 2.11. Let G be a graph that is 2-colorable, i.e., $\chi(G) = 2$. Prove that G is a subgraph of a complete bipartite graph $K_{m,n}$ with $m + n = |V|$.
- 2.12. Let G be a graph on n vertices. The *complement* of G , denoted $\overline{G}$, is the graph obtained from G by changing edges to non-edges and non-edges to edges. In other words, $ij \in E(G)$ if and only if $ij \notin E(\overline{G})$. Prove that

$$\chi(G) \cdot \chi(\overline{G}) \geq n.$$

- 2.13. The four-color theorem asserts that *every planar graph is 4-colorable*. Is the converse true? That is, is every 4-colorable graph planar?
- 2.14. **Explore:** Lemma 2.15 says that every planar graph has a vertex of degree at most 5. Try to find a planar graph such that every vertex has degree exactly 5. *Hint.* Use Euler's formula with a double count of edge-face incidences and vertex-edge incidences, as in the proof of Lemma 2.15, to find the minimum number of vertices, edges, and faces that might work.

Section 2.4. Ramsey theory

- 2.15. Find a 2-edge-coloring of the complete graph K_5 such that there are no monochromatic triangles.
- 2.16. Let $k, \ell \geq 1$ be positive integers. Prove that

$$R(k, \ell) \leq \binom{k + \ell - 2}{k - 1}.$$

- 2.17. Consider an $N \times N$ square grid of dots, with rows and columns labeled $1, \dots, N$. For $n \leq N$, let $S, T \subseteq [N]$ be subsets of size n . An $n \times n$ *square subgrid* is the set of the n^2 dots in rows S and columns T . In particular, square subgrids don't have to be contiguous.
- (a) Given $n \geq 1$, is it true that there exists a number N such that every 2-coloring of the dots in an $N \times N$ square grid contains an $n \times n$ monochromatic square subgrid?
- (b) If $n = 2$, can you calculate the smallest N with the property of part (a)?

Appendix A

Appendix: Class Exercises

These are the problems that we may work on during class.

A.1 Randomized Playlists

Suppose you have 15 songs stored on a computer. A music service creates a 15-song playlist for you by choosing one song at random from your collection for each of the 15 slots. (So, for example, it is possible that it plays the same song 15 times, although that is exceedingly unlikely.)


- (a) How many possible playlists are there?
- (b) Of the playlists in part (a), what percentage of them repeat at least one song? First take a wild guess; then figure it out.
- (c) Of the playlists in part (a), what percentage of them repeat at least one song twice in a row? First take a wild guess; then figure it out.
- (d) **Explore:** Of the playlists in part (a), what percentage of them repeat at least one song k times in a row, when $k = 3$? What about $k = 4, \dots, 15$?

Note: On the whole, people tend to significantly underestimate the likelihood that a song repeats, or even repeats twice in a row. For this reason, music services like Spotify or iTunes initially fielded many complaints from users that their playlist algorithms were defective. In some cases, these companies decided to rewrite their algorithms to make their playlists feel more random, even though the new methods were actually less random!

A.2 Catalan Bijections

(a) Let n be a positive integer. Construct bijections between the following four sets.¹

- (1) The set X_n of walks from $(0, 0)$ to $(2n, 0)$ and consisting of $2n$ steps, each step in direction $(1, 1)$ or $(1, -1)$, with the property that the walk never crosses below the x -axis. For example:



- (2) The set Y_n of sequences of n open parenthesis symbols $($ and n close parenthesis symbols $)$ that are *correctly parenthesized*, meaning that in any initial subsequence, one never encounters more close parentheses $)$ than open parentheses $($. For example, here are two valid sequences and one that is not valid:

$((()()))$
a valid sequence


$((()())())$
a valid sequence

$((()))(())$
an invalid sequence

- (3) The set Z_n of bijective assignments of the numbers $1, \dots, 2n$ to a $2 \times n$ grid of boxes such that the rows increase from left to right and the columns increase from top to bottom. For example:

1	3	4
2	5	6

- (4) The set W_n of *triangulations of a convex polygon on $n + 2$ vertices*. A *triangulation* of a convex polygon P is a collection of noncrossing diagonals in P that divide it into triangles. For example:



- (b) Compute the size of X_4 and X_5 .
(c) Make a conjecture about the sequence of sizes of the sets $X_1, X_2, X_3, \dots$ and prove your conjecture.

¹This exercise is inspired by Exercise 6.19 in Richard Stanley's classic text *Enumerative Combinatorics, Volume II*. [4] However, instead of having only 4 sets, Stanley lists 66 sets asks the reader to come up with $\binom{66}{2}$ bijections!

A.3 How Many Trees?

1. How many trees are there on vertex set $\{1\}$?
2. How many trees are there on vertex set $\{1, 2\}$?
3. How many trees are there on vertex set $\{1, 2, 3\}$?
4. How many trees are there on vertex set $\{1, 2, 3, 4\}$?
5. How many trees are there on vertex set $\{1, 2, 3, 4, 5\}$?
6. How many trees are there on vertex set $\{1, 2, 3, 4, 5, 6\}$?

Hint. You probably won't want to actually try to draw all of them when there are 5 or 6 vertices!

References

- [1] N. Alon and J. H. Spencer. *The probabilistic method*. Wiley Series in Discrete Mathematics and Optimization. John Wiley & Sons, Inc., Hoboken, NJ, fourth edition, 2016.
- [2] M. Bóna. *A walk through combinatorics*. World Scientific Publishing Co., Inc., River Edge, NJ, 2002. An introduction to enumeration and graph theory, With a foreword by Richard Stanley.
- [3] R. Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997. Translated from the 1996 German original.
- [4] R. P. Stanley. *Enumerative combinatorics. Vol. 2*, volume 62 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1999. With a foreword by Gian-Carlo Rota and appendix 1 by Sergey Fomin.
- [5] R. P. Stanley. *Enumerative combinatorics. Volume 1*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition, 2012.

