Preface

Combinatorics is a fascinating branch of mathematics centered around counting various objects and sets. Counting problems make regular appearances on middle and high school mathematics competitions despite the fact that combinatorics is generally covered only very briefly in high school math courses. This is not, however, because combinatorics requires higher level math as a prerequisite; indeed, many counting problems are accessible to anyone with a solid background in arithmetic and some basic algebra.

This book gives students a chance to explore some introductory to intermediate topics in combinatorics. We include chapters featuring tools for solving counting problems, proof techniques, and more to give students a broad foundation to build on. It is worth noting that some sections of this book are significantly more challenging than others. In particular, the chapters on *Invariants*, *Counting in more than one way*, and *Generating functions* cover topics that are considered fairly advanced; readers should not be discouraged if they do not immediately grasp these concepts. Though counting problems in particular are accessible to anyone, that does not mean they are trivial. One of the trickiest aspects of solving a counting problem is determining which tool or trick should be used. To help readers become accustomed to dealing with these subtleties, each section includes several example problems of varying difficulty with solutions to demonstrate how the different techniques may be applied in practice.

Following these topic-based segments we have included several introductory and advanced problems for students to tackle by themselves. These were carefully selected to enable the reader to further hone their problem solving skills based on the material presented in the chapters. Students can check their work in the final part of this book, which includes detailed solutions to these problems.

Several of the problems that appear in this book are pulled from various mathematics competitions worldwide. We would like to express our gratitude to the many writers who have contributed to these contests and provided us with such a rich selection of exercises. We would also like to thank Dr. Titu Andreescu for giving us the opportunity and encouragement to write this book and Dr. Richard Stong, Dr. Branislav Kisacanin, and Dr. Walter Stromquist for their thoughtful feedback, which helped us shape this text to be the absolute best it could be.

We hope you enjoy the problems!

Elizabeth “Lizard” Reiland
Vlad Matei

Contents

Chapter 1

Counting Basics

Before we jump into counting, we will go over some set theory definitions and notation that is important to our study of combinatorics. These are common terms that appear throughout the mathematical literature, so it is good to learn and remember them.

Definition 1. A *set* is a collection of distinct elements whose order is not important. We can specify a set by listing its elements such as $\{1, 2, 4, 8, 16\}$ or $\{3, 5, 7, \ldots, 19\}$. Notice that our definition means that, for example, $\{1, 2, 4\}$, $\{2, 4, 1\}$, and even $\{1, 1, 2, 2, 4\}$ are exactly the same set.

We can also use *set builder* notation where we specify a condition used to determine which elements belong to the set such as $\{x \mid 1 < x < 17,\ x \text{ is an integer}\}$. The bar $\mid$ can be read as "such that," so this set is all values x such that x is an integer and $1 < x < 17$. Thus this set is simply $\{2, 3, \ldots, 16\}$. Another example of set builder notation is $\{(x, y) \mid x \text{ and } y \text{ are real numbers}, y = 3x + 4\}$. Note that this set contains an infinite number of ordered pairs (x, y).

- The *empty set* is the set which contains no elements. We denote it as $\{\ \}$ or $\emptyset$.
- The notation $x \in A$ (read "x is in A" or "x is an element of A") means that the element x is included in the set A. We use the notation $y \notin A$ (read "y is not in A" or "y is not an element of A") to indicate that y is not included in the set A.
- We say that a set A is a *subset* of a set B (denoted $A \subseteq B$) if every element of A is an element of B (i.e., $x \in A$ implies $x \in B$).
- Two sets A and B are *equal* (denoted $A = B$) if they contain exactly the same elements. (One common way to prove $A = B$ is to show that $A \subseteq B$ and $B \subseteq A$. Keep this in mind!)

- The *union* of two sets A and B (denoted $A \cup B$) is the set of all elements in either A or B: $\{x \mid x \in A \textit{ or } x \in B\}$. The *intersection* of A and B (denoted $A \cap B$) is the set of all elements belonging to both A and B: $\{x \mid x \in A \textit{ and } x \in B\}$. These definitions can be extended to more than two sets in the intuitive way:

$$S_1 \cup S_2 \cup \cdots \cup S_k = \{x \mid x \in S_i \text{ for some } i, 1 \leq i \leq k\}$$

$$S_1 \cap S_2 \cap \cdots \cap S_k = \{x \mid x \in S_i \text{ for all } i, 1 \leq i \leq k\}$$

- We say two sets A and B are *disjoint* if they have no elements in common (i.e., if $A \cap B = \emptyset$).

- The *set difference* of the set A and the set B (denoted $A \backslash B$) is the set of elements that are in A but not in B. This notation is used even when B is not a subset of A; for example, $\{1, 2, 3\} \setminus \{3, 4\}$ is $\{1, 2\}$.

- If we have a *universal set* U which contains all of the objects we are interested in, we can define the *complement* of a set A (denoted A^c) as the collection of elements not in A (i.e., $A^c = U \backslash A$). For example, if we are working with the set of integers, the complement of the set of even numbers would be the set of odd numbers. (*Note:* we have to have some universal set in order for the idea of a complement to make sense!)

- The *cardinality* or *size* of a set A (denoted $|A|$) is the number of elements in that set.

Though these definitions may seem straightforward, there are some surprisingly subtle issues in set theory. It is possible for elements of a set to be sets in their own right. For example, one could take the set A of all subsets of $\{1, 2, 3\}$ (called the *power set* or A). We have

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

The elements of A are sets. One could iterate this idea to build sets whose elements are sets of sets, and so on. Another interesting example is the set $B = \{\emptyset\}$. Notice that B is not the empty set, but rather the set containing the empty set. The size of the empty set is $|\emptyset| = 0$, but we have $|B| = 1$.

One might then worry about whether a set A could contain itself as an element, $A \in A$. To avoid this one might try to restrict to the set of all sets that do not contain themselves, $B = \{A : A \notin A\}$. Thinking about whether B contains itself will lead you to what is known as Russell's paradox. These issues can be fun but will not be relevant to this text, since our sets will be explicitly defined and usually finite.

As we start thinking about counting, there are two essential rules that will show up in almost every problem you encounter. Once you've done a bit of counting, you'll find yourself using these without even thinking about them. We will state these principles formally in a moment, but first we will examine a simple example.

Example 1. Suppose we are at a clothing store which offers 16 different shirts, 9 different pairs of pants, and 3 different pairs of shoes. How many ways are there to purchase an article of clothing?

Before we discuss the solution to this exercise, note that Example 1 illustrates an important fact about combinatorics problems. It is more fun to phrase combinatorics problems in simple English, and this is the way you will often see them. However, English is not as precise a language as mathematics, and we generally do not want to include long lists of disclaimers and explanations to make the problems technically precise since this would defeat the point of using simple English.

One of the first steps you should take when approaching a combinatorics problem is to decide how you want to interpret the English. For instance, in solving Example 1, we will implicitly assume that the only types of articles of clothing are the three mentioned (shirts, pants, and shoes) and that shoes have to be purchased in a pair. Mathematicians generally agree on how to interpret problems, and this is one of the things you will pick up going through the examples. If you are uncertain how to interpret a problem statement and are unable to ask someone to clarify, make your best attempt at an appropriate interpretation and be sure to note the assumptions you have made in your solution.

Having made these notes, let us now solve Example 1.

Solution. Because an article of clothing is either a shirt, a pair of pants, or a pair of shoes we can simply add up the number of each type of clothes to find $16 + 9 + 3 = 28$ possible ways to buy an article of clothing. □

The counting in this exercise was fairly straightforward, but it illustrates an application of the Sum Rule, a generalized principle which can be used to solve much more complicated problems. The formal statement of the Sum Rule is as follows:

Theorem 1. (Sum Rule) *If $A_1, A_2, \ldots, A_n$ are pairwise disjoint sets (i.e., if no pair of sets have elements in common), then*

$$|A_1 \cup A_2 \cup \cdots \cup A_n| = |A_1| + |A_2| + \cdots + |A_n|.$$

While this may seem like a lot of fancy notation, in practice this rule just tells us that if we are counting the possible ways to pick an object from one of several different sets that do not overlap, we just need to add up the sizes of the individual sets. If the sets do overlap, we will need to be a bit more careful; we discuss how to deal with this possibility in the inclusion exclusion section. Let us give a simple example of applying this principle

Example 2. Let $X = \{1, 2, \ldots, 200\}$. We define

$$S = \{(a, b, c) \mid a, b, c \in X,\ a < b \text{ and } a < c\}.$$

How many elements does S have?

Solution. Note that we can split S up into disjoint set A_k where k is the value of a and $1 \leq k \leq 199$. Note that since $b > k$ and $c > k$ we have $200 - k$ choices for b and $200 - k$ choices for c. Thus $|A_k| = (200 - k)^2$. Using the addition principle we obtain $|S|$. □

Example 3. Suppose we are at a clothing store which offers 16 different shirts and 9 different pairs of pants. How many ways are there to purchase an outfit consisting of one shirt and one pair of pants?

Solution. To help facilitate our counting, let us build a table. Each row of the table will represent a particular shirt, whereas each column will represent a particular pair of pants. A particular cell in the table will correspond to the outfit consisting of the shirt indicated by the row and the pair of pants indicated by the column of that cell. Since each cell will represent one distinct outfit, and every outfit appears in exactly one cell, our number of outfits is simply equal to the number of cells in our table. Since we have 16 shirts and 9 pairs of pants, there are $16 \cdot 9 = 144$ cells in our table, and thus 144 possible outfits we could buy. □

Notice that if we wanted to create an oufit consisting of a shirt *and* a pair of pants *and* a pair of shoes, we could expand on this idea to make a three dimensional table with one coordinate representing shirts, a second representing pants, and the last representing shoes. Similarly, if we had n selections to make, we could imagine counting cells in an n-dimensional table. This brings us to our other basic rule:

Theorem 2. (Product Rule) *If we have a sequence of n choices to make with X_1 possibilities for the first choice, X_2 possibilities for the second choice, and so on up to X_n choices for the nth choice, there are a total of $X_1 \cdot X_2 \cdot \ldots \cdot X_n$ ways to make our choices.*

For most problems we will apply both the Sum Rule and Product Rule to get us to our final solution. By using the Sum Rule we can break problems into a collection of cases where each case is relatively simple to count (generally by employing the Product Rule) as illustrated in the next example.

Example 4. How many three-digit numbers have exactly one even digit?

Solution. We will look at three different cases here: the case where the first digit is even and the other two are odd, the case where the middle digit is even and the other two are odd, and the case where the last digit is even and the other two are odd. Since these cases do not overlap, we can count each individually, then apply the Sum Rule to get our final answer.

We can think of creating a three-digit number as a series of three steps: choosing the first digit, choosing the second digit, and choosing the final digit. In the case where the first digit is even and the other two are odd, there are 4 choices for the first digit $(2, 4, 6, 8)$ since it must be even and cannot be zero (otherwise we would not have a three-digit number). Since the second and third digits are both odd, there are 5 possibilities for each $(1, 3, 5, 7, 9)$. Thus the Product Rule tells us there are $4 \cdot 5 \cdot 5$ three-digit numbers fitting this case. In the case where the middle digit is even and the other two are odd, every digit has 5 possibilities: $1, 3, 5, 7, 9$ for the odd digits and $0, 2, 4, 6, 8$ for the even digit. Overall then, there are $5 \cdot 5 \cdot 5$ three-digit numbers in this case. Similarly, there are $5 \cdot 5 \cdot 5$ numbers satisfying the case where the last digit is even and the other two are odd. Putting these three cases together using the Sum Rule, we have $4 \cdot 5 \cdot 5 + 5 \cdot 5 \cdot 5 + 5 \cdot 5 \cdot 5 = 350$ three-digit numbers with exactly one even digit. □

One more basic but very useful technique to keep in mind is *complementary counting*. Suppose we are interested in determining the size of a set A. If we have a finite universal set U, we know by the Sum Rule that $|A| + |A^c| = |U|$. Rearranging, we find $|A| = |U| - |A^c|$. We can take advantage of this to help us determine the size of A. In particular, we can determine the size of our universal set and the size of the complement of A, then subtract. In some cases this may be significantly easier than trying to directly count A. If you see the words "at least" in a problem, complementary counting will often be a good method to consider.

Example 5. How many four-digit positive integers have at least one digits that is a 2 or a 3?

(2006 AMC 10A)

Solution. Let's first count the total number of four-digit positive integers. The first digit must be from 1 to 9, so we have 9 choices. For each of the three remaining digits, we need a value from 0 to 9 so there are 10 choices each.

Thus in total there are $9 \cdot 10 \cdot 10 \cdot 10 = 9 \cdot 10^3 = 9000$ four-digit positive integers.

Next we count how many four-digit positive integers DO NOT contain a 2 or a 3. Then we have 7 choices for our first digit (1,4,5,6,7,8, or 9) and 8 for the remaining three-digits. This gives a total of $7 \cdot 8^3$ four-digit integers not containing a 2 or a 3. Subtracting this from our total, we conclude that there are $9000 - 7 \cdot 8^3 = 5416$ four-digit positive integers that have at least one digit that is a 2 or a 3. □

Let's look at some examples of problems making use of the techniques we've learned thus far.

Example 6. How many subsets of $\{1, 2, \ldots, n\}$ are there? (Note: This quantity will come up frequently in problems, so it's a useful fact to remember.)

Solution. Consider an element i $(1 \leq i \leq n)$. As we construct a subset S, we have two choices for i: Either it is in S or it is not in S. Since we must make this choice for each of the n elements, by the Product Rule there are 2^n total subsets of $\{1, 2, \ldots, n\}$. □

Example 7. How many subsets S of $\{1, 2, \ldots, n\}$ are there such that $|S|$ is odd?

Solution. For each element i $(1 \leq i \leq n-1)$ we have two choices: Either i is in S or it is not in S. At this point, we consider $|S|$. If $|S|$ is odd, we must not include n in S. On the other hand, if $|S|$ is even (so far), we have to include n in S to satisfy the condition that $|S|$ is odd. In either case, we have only one choice for what to do with n. By the Product Rule, this implies there are $2^{n-1} \cdot 1 = 2^{n-1}$ subsets S of $\{1, 2, \ldots, n\}$ are there such that $|S|$ is odd. □

Notice that this solution does not work when $n = 0$; certainly there are not 2^{-1} subsets of $\{\}$ with an odd number of elements. It is good to get in the habit of watching out for cases like this. If you are writing a solution on an exam, make sure you say that you are assuming $n > 0$.

Example 8. A dessert chef prepares the dessert for every day of a week starting with Sunday. The dessert each day is either cake, pie, ice cream, or pudding. The same dessert may not be served two days in a row. There must be cake on Friday because of a birthday. How many different dessert menus for the week are possible?

(2012 AMC 12B)

Solution. We start with Friday, since we know cake must be served that day. This implies that on Saturday, the dessert served cannot be cake, so we have 3 choices for that day's dessert. Similarly when we work backwards from Friday to Thursday, we see we have 3 choices for the dessert on Thursday (anything but cake). Then Wednesday we may select any of the 3 desserts not served Thursday, and so on back to Sunday. Since we have 3 choices for each day (aside from Friday), by the Product Rule we have $3^6 = 729$ possible menus. $\square$

Note that there is something slightly subtle about our use of the Product Rule in Example 7. The Product Rule only requires that at each step in our chain of choices that we have the same *number* of possible choices at that point in our decision chain. What those specific options are does not matter. In Example 7, though the *set* of desserts allowed might change based on particular choices we make, for each day (besides Friday) the *number* of possible desserts is always exactly 3.

There are other ways to solve this problem as well. For example, we could have started with Monday and worked forward. Although this can work, it is much harder and involves some casework. (Try it if you don't believe us.) There are often several correct ways to solve counting problems, and it is always a good idea to consider different possible approaches.

Example 9. A large cube is painted green and then chopped up into 64 smaller congruent cubes. How many of the smaller cubes have at least one face painted green?

(Alabama ARML team selection)

Solution. We use complementary counting and determine how many cubes have no green faces. To have no green faces, a small cube must have been on the interior of the large cube. The large cube is $4 \times 4 \times 4$ with respect to the small cubes, so the interior of this cube is a $2 \times 2 \times 2$ group of small cubes. This is $2^3 = 8$ small cubes with no green faces, so there are $64 - 8 = 56$ small cubes with at least one face painted green. $\square$

Example 10. Suppose $n \geq 2$ is a positive integer with prime factorization $n = p_1^{\alpha_1} p_2^{\alpha_2} \cdots p_k^{\alpha_k}$ where the p_i are prime numbers and the α_i are positive integers. How many factors does n have?

Solution. Recall that a number x is a factor of n if n is divisible by x. In order for this to be the case, the prime factorization of x must be $x = p_1^{\beta_1} p_2^{\beta_2} \cdots p_k^{\beta_k}$ where $0 \leq \beta_i \leq \alpha_i$ for each i. This means we have $\alpha_1 + 1$ choices for the value of β_1, $\alpha_2 + 1$ choices for the value of β_2, and so on. Applying the Product Rule, this tells us that the total number of divisors of n is $(\alpha_1+1)(\alpha_2+1)\cdots(\alpha_k+1)$. As an example, consider $20 = 2^2 \cdot 5^1$.
By our logic, 20 should have $(2+1)(1+1) = 6$ factors. They are $1, 2, 4, 5, 10$, and 20. $\square$

Example 11. Let n and k be positive integers. Count the number of k-tuples $(S_1, S_2, \ldots, S_k)$ of subsets of S_i of $\{1, 2, \ldots, n\}$ subject to each of the following conditions separately (i.e., the three parts are independent problems).

(a) The S_i's are pairwise disjoint.

(b) $S_1 \cap S_2 \cap \cdots \cap S_k = \emptyset$.

(c) $S_1 \cup S_2 \cup \cdots \cup S_k = \{1, \ldots, n\}$.

Solution.

(a) Consider a particular element $j \in \{1, 2, \ldots, n\}$. In order for the S_i's to be pairwise disjoint, j can be in at most one of $S_1, \ldots, S_k$. This is a total of $k+1$ possibilities (one for each subset and one for the possibility of j being in none of the subsets) for each of the n elements, so by the Product Rule there are $(k+1)^n$ k-tuples such that the S_i's are pairwise disjoint.

(b) Again consider a particular element $j \in \{1, 2, \ldots, n\}$. For each of the S_i we have 2 options: either j is in S_i or it is not. Thus there are a total of 2^k possible combinations of the subsets S_i that j could appear in. There is only 1 case that would violate our condition; the case where j is contained in every S_i. Thus there are $2^k - 1$ valid placements for each of the n elements of $\{1, \ldots, n\}$. Thus by the Product Rule there are $(2^k - 1)^n$ k-tuples satisfying $S_1 \cap S_2 \cap \cdots \cap S_k = \emptyset$.

(c) This is actually very similar to the previous part! There is only 1 case that would violate our condition; the case where j is contained in none of the S_i. Thus there are $2^k - 1$ valid placements for each of the n elements of $\{1, \ldots, n\}$. Thus by the Product Rule there are $(2^k - 1)^n$ k-tuples satisfying $S_1 \cup S_2 \cup \cdots \cup S_k = \{1, \ldots, n\}$. □