

Counting I

20480135385502964448038	3171004832173501394113017	5763257331083479647409398	8247331000042995311646021
489445991866915676240992	3208234421597368647019265	5800949123548989122628663	8496243997123475922766310
1082662032430379651370981	3437254656355157864869113	6042900801199280218026001	8518399140676002660747477
1178480894769706178994993	3574883393058653923711365	6116171789137737896701405	8543691283470191452333763
1253127351683239693851327	3644909946040480189969149	6144868973001582369723512	8675309258374137092461352
1301505129234077811069011	3790044132737084094417246	6247314593851169234746152	8694321112363996867296665
1311567111143866433882194	3870332127437971355322815	6814428944266874963488274	8772321203608477245851154
1470029452721203587686214	4080505804577801451363100	6870852945543886849147881	8791422161722582546341091
1578271047286257499433886	4167283461025702348124920	6914955508120950093732397	9062628024592126283973285
1638243921852176243192354	423599683112377788211249	6949632451365987152423541	9137845566925526349897794
1763580219131985963102365	4670939445749439042111220	7128211143613619828415650	9153762966803189291934419
1826227795601842231029694	4815379351865384279613427	7173920083651862307925394	9270880194077636406984249
1843971862675102037201420	4837052948212922604442190	7215654874211755676220587	9324301480722103490379204
2396951193722134526177237	5106389423855018550671530	7256932847164391040233050	9436090832146695147140581
2781394568268599801096354	5142368192004769218069910	7332822657075235431620317	9475308159734538249013238
2796605196713610405408019	5181234096130144084041856	7426441829541573444964139	9492376623917486974923202
2931016394761975263190347	5198267398125617994391348	7632198126531809327186321	9511972558779880288252979
2933458058294405155197296	5317592940316231219758372	7712154432211912882310511	9602413424619187112552264
3075514410490975920315348	5384358126771794128356947	7858918664240262356610010	9631217114906129219461111
3111474985252793452860017	5439211712248901995423441	7898156786763212963178679	9908189853102753335981319
3145621587936120118438701	5610379826092838192760458	8147591017037573337848616	9913237476341764299813987
3148901255628881103198549	5632317555465228677676044	8149436716871371161932035	
3157693105325111284321993	5692168374637019617423712	8176063831682536571306791	

Two different subsets of the ninety 25-digit numbers shown above have the same sum. For example, maybe the sum of the numbers in the first column is equal to the sum of the numbers in the second column. Can you find two such subsets? This is a very difficult computational problem. But we'll prove that such subsets must exist! This is the sort of weird conclusion one can reach by tricky use of counting, the topic of this chapter.

Counting seems easy enough: 1, 2, 3, 4, etc. This explicit approach works well for counting simple things, like your toes, and for extremely complicated things for which there's no identifiable structure. However, subtler methods can help you count many things in the vast middle ground, such as:

- The number of different ways to select a dozen doughnuts when there are five varieties available.
- The number of 16-bit numbers with exactly 4 ones.

Counting is useful in computer science for several reasons:

- Determining the time and storage required to solve a computational problem—a central objective in computer science—often comes down to solving a counting problem.

- Counting is the basis of probability theory, which in turn is perhaps the most important topic this term.
- Two remarkable proof techniques, the “pigeonhole principle” and “combinatorial proof”, rely on counting. These lead to a variety of interesting and useful insights.

We’re going to present a lot of rules for counting. These rules are actually theorems, but we’re generally not going to prove them. Our objective is to teach you counting as a practical skill, like integration. And most of the rules seem “obvious” anyway.

1 Counting One Thing by Counting Another


How do you count the number of people in a crowded room? We could count heads, since for each person there is exactly one head. Alternatively, we could count ears and divide by two. Of course, we might have to adjust the calculation if someone lost an ear in a pirate raid or someone was born with three ears. The point here is that we can often *count one thing by counting another*, though some fudge factors may be required. This is the central theme of counting, from the easiest problems to the hardest.

In more formal terms, every counting problem comes down to determining the size of some set. The *size* or *cardinality* of a set S is the number of elements in S and is denoted $|S|$. In these terms, we’re claiming that we can often *find the size of one set S by finding the size of a related set T* . We already have a mathematical tool for relating one set to another: relations. Not surprisingly, a particular kind of relation is at the heart of counting.

1.1 Functions

Functions like $f(x) = x^2 + 1$ and $g(x) = \tan^{-1}(x)$ are surely quite familiar from calculus. We’re going to use functions for counting as well, but in a way that might not be familiar. Instead of using functions that map one real number to another, we’ll use functions that relate elements of finite sets.

In order to count accurately, we need to carefully define the notion of a function. Formally, a *function* $f : X \rightarrow Y$ is a relation between two sets, X and Y , that relates *every* element of X to *exactly one* element of Y . The set X is called the *domain* of the function f , and Y is called the *codomain*. Here is an illustration of a function:



In this example, the domain is the set $X = \{a, b, c, d, e\}$ and the range is the set $Y = \{1, 2, 3, 4, 5\}$. Related elements are joined by an arrow. This relation is a function because *every* element on the left is related to *exactly one* element on the right. In graph-theoretic terms, this is a function because every element on the left has degree *exactly 1*. If x is an element of the domain, then the related element in the codomain is denoted $f(x)$. We often say f *maps* x to $f(x)$. In this case, f maps a to 1, b to 3, c to 4, d to 3 and e to 5. Equivalently, $f(a) = 1$, $f(b) = 3$, $f(c) = 4$, $f(d) = 3$, and $f(e) = 5$.


The relations shown below are *not* functions:




The relation on the left is not a function because a is mapped to *two* elements of the codomain. The relation on the right is not a function because b and d are mapped to *zero* elements of the codomain. Tsk-tsk!

1.2 Bijections

The definition of a function is rather asymmetric; there are restrictions on elements in the domain, but not on elements in the codomain. A *bijection* or *bijjective function* is a function $f : X \rightarrow Y$ that maps exactly one element of the domain to each element of the codomain. In graph terms, *both* domain and codomain elements are required to have degree exactly 1 in a bijective function. Here is an example of a bijection:



In contrast, these relations that are *not* bijections:



The first function is not bijective because 3 is related to *two* elements of the domain. The second function is not bijective because 4 is related to *zero* elements of the domain. The last relation is not even a function, because *b* is associated with zero elements of the domain.

Bijjective functions are also found in the more-familiar world of real-valued functions. For example, $f(x) = 6x + 5$ is a bijective function with domain and codomain \mathbb{R} . For every real number y in the codomain, $f(x) = y$ for exactly one real number x in the domain. On the other hand, $f(x) = x^2$ is not a bijective function. The number 4 in the codomain is related to both 2 and -2 in the domain.

1.3 The Bijection Rule

If we can pair up all the girls at a dance with all the boys, then there must be an equal number of each. This simple observation generalizes to a powerful counting rule:

Rule 1 (Bijection Rule). *If there exists a bijection $f : A \rightarrow B$, then $|A| = |B|$.*

In the example, A is the set of boys, B is the set of girls, and the function f defines how they are paired.


The Bijection Rule acts as a magnifier of counting ability; if you figure out the size of one set, then you can immediately determine the sizes of many other sets via bijections.

For example, let's return to two sets mentioned earlier:


A = all ways to select a dozen doughnuts when five varieties are available

B = all 16-bit sequences with exactly 4 ones

Let's consider a particular element of set A :



We've depicted each doughnut with a 0 and left a gap between the different varieties. Thus, the selection above contains two chocolate doughnuts, no lemon-filled, six sugar, two glazed, and two plain. Now let's put a 1 into each of the four gaps:




We've just formed a 16-bit number with exactly 4 ones—an element of B !

This example suggests a bijection from set A to set B : map a dozen doughnuts consisting of:

c chocolate, l lemon-filled, s sugar, g glazed, and p plain

to the sequence:



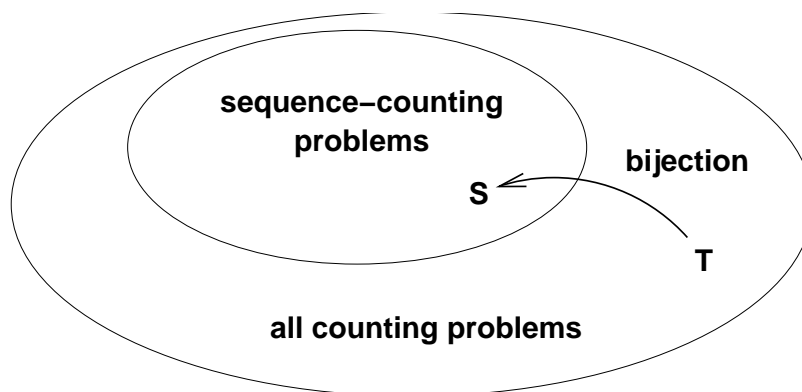
The resulting sequence always has 16 bits and exactly 4 ones, and thus is an element of B . Moreover, the mapping is a bijection; every such bit sequence is mapped to by exactly one order of a dozen doughnuts. Therefore, $|A| = |B|$ by the Bijection Rule!

This demonstrates the magnifying power of the bijection rule. We managed to prove that two very different sets are actually the same size—even though we don't know exactly how big either one is. But as soon as we figure out the size of one set, we'll immediately know the size of the other.

This particular bijection might seem frighteningly ingenious if you've not seen it before. But you'll use essentially this same argument over and over and over, and soon you'll consider it boringly routine.

1.4 Sequences

The Bijection Rule lets us count one thing by counting another. This suggests a general strategy: get really good at counting just a *few* things and then use bijections to count *everything else*:



This is precisely the strategy we'll follow. In particular, we'll get really good at counting *sequences*. When we want to determine the size of some other set T , we'll find a bijection from T to a set of sequences S . Then we'll use our super-ninja sequence-counting skills to determine $|S|$, which immediately gives us $|T|$. We'll need to hone this idea somewhat as we go along, but that's pretty much the plan!

In order to pull this off, we need to clarify some issues concerning sequences and sets. Recall that a *set* is an unordered collection of distinct elements. A set is often represented by listing its elements inside curly-braces. For example, $\{a, b, c\}$ is a set, and $\{c, b, a\}$ is another way of writing the same set. On the other hand, $\{a, b, a\}$ is not a set, because element a appears twice.

On the other hand, a *sequence* is an ordered collection of elements (called *components* or *terms*) that are not necessarily distinct. A sequence is often written by listing the terms inside parentheses. For example, (a, b, c) is a sequence, and (c, b, a) is a different sequence. Furthermore (a, b, a) is a perfectly valid three-term sequence.

The distinction between sets and sequences is crucial for everything that follows. If you don't keep the distinction clear in your mind, you're doomed!

2 Two Basic Counting Rules

We'll harvest our first crop of counting problems with two basic rules.

2.1 The Sum Rule

Linus allocates his big sister Lucy a quota of 20 crabby days, 40 irritable days, and 60 generally surly days. On how many days can Lucy be out-of-sorts one way or another? Let set C be her crabby days, I be her irritable days, and S be the generally surly. In these terms, the answer to the question is $|C \cup I \cup S|$. Now assuming that she is permitted at most one bad quality each day, the size of this union of sets is given by the Sum Rule:

Rule 2 (Sum Rule). If A_1, A_2, \dots, A_n are disjoint sets, then:

$$|A_1 \cup A_2 \cup \dots \cup A_n| = |A_1| + |A_2| + \dots + |A_n|$$

Thus, according to Linus' budget, Lucy can be out-of-sorts for:

$$\begin{aligned} |C \cup I \cup S| &= |C| + |I| + |S| \\ &= 20 + 40 + 60 \\ &= 120 \text{ days} \end{aligned}$$

Notice that the Sum Rule holds only for a union of *disjoint* sets. Finding the size of a union of intersecting sets is a more complicated problem that we'll take up later.

2.2 The Product Rule

The product rule gives the size of a product of sets. Recall that if P_1, P_2, \dots, P_n are sets, then

$$P_1 \times P_2 \times \dots \times P_n$$

is the set of all sequences whose first term is drawn from P_1 , second term is drawn from P_2 and so forth.

Rule 3 (Product Rule). If P_1, P_2, \dots, P_n are sets, then:

$$|P_1 \times P_2 \times \dots \times P_n| = |P_1| \cdot |P_2| \cdot \dots \cdot |P_n|$$

Unlike the sum rule, the product rule does not require the sets P_1, \dots, P_n to be disjoint. For example, suppose a *daily diet* consists of a breakfast selected from set B , a lunch from set L , and a dinner from set D :

$$\begin{aligned} B &= \{\text{pancakes, bacon and eggs, bagel, Doritos}\} \\ L &= \{\text{burger and fries, garden salad, Doritos}\} \\ D &= \{\text{macaroni, pizza, frozen burrito, pasta, Doritos}\} \end{aligned}$$

Then $B \times L \times D$ is the set of all possible daily diets. Here are some sample elements:

$$\begin{aligned} &(\text{pancakes, burger and fries, pizza}) \\ &(\text{bacon and eggs, garden salad, pasta}) \\ &(\text{Doritos, Doritos, frozen burrito}) \end{aligned}$$

The Product Rule tells us how many different daily diets are possible:

$$\begin{aligned} |B \times L \times D| &= |B| \cdot |L| \cdot |D| \\ &= 4 \cdot 3 \cdot 5 \\ &= 60 \end{aligned}$$

2.3 Putting Rules Together

Few counting problems can be solved with a single rule. More often, a solution is a flurry of sums, products, bijections, and other methods. Let's look at some examples that bring more than one rule into play.

Passwords

The sum and product rules together are useful for solving problems involving passwords, telephone numbers, and license plates. For example, on a certain computer system, a valid password is a sequence of between six and eight symbols. The first symbol must be a letter (which can be lowercase or uppercase), and the remaining symbols must be either letters or digits. How many different passwords are possible?

Let's define two sets, corresponding to valid symbols in the first and subsequent positions in the password.

$$F = \{a, b, \dots, z, A, B, \dots, Z\}$$

$$S = \{a, b, \dots, z, A, B, \dots, Z, 0, 1, \dots, 9\}$$

In these terms, the set of all possible passwords is:

$$(F \times S^5) \cup (F \times S^6) \cup (F \times S^7)$$

Thus, the length-six passwords are in set $F \times S^5$, the length-seven passwords are in $F \times S^6$, and the length-eight passwords are in $F \times S^7$. Since these sets are disjoint, we can apply the Sum Rule and count the total number of possible passwords as follows:

$$\begin{aligned} |(F \times S^5) \cup (F \times S^6) \cup (F \times S^7)| &= |F \times S^5| + |F \times S^6| + |F \times S^7| && \text{Sum Rule} \\ &= |F| \cdot |S|^5 + |F| \cdot |S|^6 + |F| \cdot |S|^7 && \text{Product Rule} \\ &= 52 \cdot 62^5 + 52 \cdot 62^6 + 52 \cdot 62^7 \\ &\approx 1.8 \cdot 10^{14} \text{ different passwords} \end{aligned}$$

Subsets of an n -element Set

How many different subsets of an n element set X are there? For example, the set $X = \{x_1, x_2, x_3\}$ has eight different subsets:

$$\begin{array}{cccc} \{\} & \{x_1\} & \{x_2\} & \{x_1, x_2\} \\ \{x_3\} & \{x_1, x_3\} & \{x_2, x_3\} & \{x_1, x_2, x_3\} \end{array}$$

There is a natural bijection from subsets of X to n -bit sequences. Let x_1, x_2, \dots, x_n be the elements of X . Then a particular subset of X maps to the sequence (b_1, \dots, b_n)

where $b_i = 1$ if and only if x_i is in that subset. For example, if $n = 10$, then the subset $\{x_2, x_3, x_5, x_7, x_{10}\}$ maps to a 10-bit sequence as follows:

$$\begin{array}{l} \text{subset: } \{ \quad x_2, \quad x_3, \quad x_5, \quad x_7, \quad x_{10} \} \\ \text{sequence: } (0, \quad 1, \quad 1, \quad 0, \quad 1, \quad 0, \quad 1, \quad 0, \quad 0, \quad 1) \end{array}$$

We just used a bijection to transform the original problem into a question about sequences—*exactly according to plan!* Now if we answer the sequence question, then we've solved our original problem as well.

But how many different n -bit sequences are there? For example, there are 8 different 3-bit sequences:

$$\begin{array}{cccc} (0, 0, 0) & (0, 0, 1) & (0, 1, 0) & (0, 1, 1) \\ (1, 0, 0) & (1, 0, 1) & (1, 1, 0) & (1, 1, 1) \end{array}$$

Well, we can write the set of all n -bit sequences as a product of sets:

$$\underbrace{\{0, 1\} \times \{0, 1\} \times \dots \times \{0, 1\}}_{n \text{ terms}} = \{0, 1\}^n$$

Then Product Rule gives the answer:

$$\begin{aligned} |\{0, 1\}^n| &= |\{0, 1\}|^n \\ &= 2^n \end{aligned}$$


This means that the number of subsets of an n -element set X is also 2^n . We'll put this answer to use shortly.

3 More Functions: Injections and Surjections

Bijjective functions are incredibly powerful counting tools. A few other kinds of functions are useful as well; we'll look at two now and one more next time. A function $f : X \rightarrow Y$ is:

- *surjective* if every element of Y is mapped to *at least once*
- *injective* if every element of Y is mapped to *at most once*
- *bijective* if every element of Y is mapped to *exactly once*

We've repeated the definition of a bijective function for comparison. Notice that these definitions immediately imply that a function is bijective if and only if it is both injective and surjective. Now the names "surjective" and "injective" are hopelessly unmemorable and nondescriptive. Some people prefer the terms *onto* and *into*, respectively, perhaps on the grounds that these are hopelessly unmemorable and nondescriptive—but shorter. Anyway, here are a couple examples:



The function on the left is surjective (every element on the right is mapped to at least once), but not injective (element 3 is mapped to twice). The function on the right is injective (every element is mapped to at most once), but not surjective (element 4 is mapped to zero times).

Earlier, we observed that two sets are the same size if there is a bijection between them. Similarly, surjections and injections imply certain size relationships between sets.

Rule 4 (Mapping Rule).

1. If $f : X \rightarrow Y$ is surjective, then $|X| \geq |Y|$.
2. If $f : X \rightarrow Y$ is injective, then $|X| \leq |Y|$.
3. If $f : X \rightarrow Y$ is bijective, then $|X| = |Y|$.

3.1 The Pigeonhole Principle

Here is an old puzzle:

A drawer in a dark room contains red socks, green socks, and blue socks. How many socks must you withdraw to be sure that you have a matching pair?


For example, picking out three socks is not enough; you might end up with one red, one green, and one blue. The solution relies on the Pigeonhole Principle, which is a friendly name for the contrapositive of part (2) of the Mapping Rule. Let's write it down:

If $|X| > |Y|$, then no function $f : X \rightarrow Y$ is injective.

Now let's rewrite this a second time to eliminate the word "injective" since, by now, there's not a ghost of a chance that you remember what that means:

Rule 5 (Pigeonhole Principle). *If $|X| > |Y|$, then for every function $f : X \rightarrow Y$ there exist two different elements of X that are mapped to the same element of Y .*

Perhaps the relevance of this abstract mathematical statement to selecting footwear under poor lighting conditions is not obvious. However, let A be the set of socks you pick out, let B be the set of colors available, and let f map each sock to its color. The Pigeonhole Principle says that if $|A| > |B| = 3$, then at least two elements of A (that is, at least two socks) must be mapped to the same element of B (that is, the same color). For example, one possible mapping of four socks to three colors is shown below.



Therefore, four socks are enough to ensure a matched pair.

Not surprisingly, the pigeonhole principle is often described in terms of pigeons: if more than n pigeons fly into n pigeonholes, then at least two pigeons must fly into some hole. In this case, the pigeons form set A , the pigeonholes are set B , and f describes the assignment of pigeons to pigeonholes.

Mathematicians have come up with many ingenious applications for the pigeonhole principle. If there were a cookbook procedure for generating such arguments, we'd give it to you. Unfortunately, there isn't one. One helpful tip, though: when you try to solve a problem with the pigeonhole principle, the key is to clearly identify three things:

1. The set A (the pigeons).
2. The set B (the pigeonholes).
3. The function f (the rule for assigning pigeons to pigeonholes).

Hairs on Heads

There are a number of generalizations of the pigeonhole principle. For example:

Rule 6 (Generalized Pigeonhole Principle). *If $|X| > k \cdot |Y|$, then every function $f : X \rightarrow Y$ maps at least $k + 1$ different elements of X to the same element of Y .*

For example, if you pick two people at random, surely they are extremely unlikely to have *exactly* the same number of hairs on their heads. However, in the remarkable city of Boston, Massachusetts there are actually *three* people who have exactly the same number of hairs! Of course, there are many bald people in Boston, and they all have zero hairs. But I'm talking about non-bald people.

Boston has about 500,000 non-bald people, and the number of hairs on a person's head is at most 200,000. Let A be the set of non-bald people in Boston, let $B = \{1, \dots, 200,000\}$, and let f map a person to the number of hairs on his or her head. Since $|A| > 2|B|$, the Generalized Pigeonhole Principle implies that at least three people have exactly the same number of hairs. I don't know who they are, but I know they exist!

Subsets with the Same Sum

We asserted that two different subsets of the ninety 25-digit numbers listed on the first page have the same sum. This actually follows from the Pigeonhole Principle. Let A be the collection of all subsets of the 90 numbers in the list. Now the sum of any subset of numbers is at most $90 \cdot 10^{25}$, since there are only 90 numbers and every 25-digit number is less than 10^{25} . So let B be the set of integers $\{0, 1, \dots, 90 \cdot 10^{25}\}$, and let f map each subset of numbers (in A) to its sum (in B).

We proved that an n -element set has 2^n different subsets. Therefore:

$$\begin{aligned} |A| &= 2^{90} \\ &\geq 1.237 \times 10^{27} \end{aligned}$$

On the other hand:

$$\begin{aligned} |B| &= 90 \cdot 10^{25} + 1 \\ &\leq 0.901 \times 10^{27} \end{aligned}$$

Both quantities are enormous, but $|A|$ is a bit greater than $|B|$. This means that f maps at least two elements of A to the same element of B . In other words, by the Pigeonhole Principle, two different subsets must have the same sum!

Notice that this proof gives no indication *which* two sets of numbers have the same sum. This frustrating variety of argument is called a *nonconstructive proof*.

Sets with Distinct Subset Sums

How can we construct a set of n positive integers such that all its subsets have *distinct* sums? One way is to use powers of two:

$$\{1, 2, 4, 8, 16\}$$

This approach is so natural that one suspects all other such sets must involve larger numbers. (For example, we could safely replace 16 by 17, but not by 15.) Remarkably, there are examples involving *smaller* numbers. Here is one:

$$\{6, 9, 11, 12, 13\}$$

One of the top mathematicians of the century, Paul Erdős, conjectured in 1931 that there are no such sets involving *significantly* smaller numbers. More precisely, he conjectured that the largest number must be $\Omega(2^n)$. He offered \$500 to anyone who could prove or disprove his conjecture, but the problem remains unsolved.