Session 2

Combinatorics. Part I

Paul Zeitz

Sneak Preview. You may have seen the bumper sticker that says *"Mathematicians Count Too."* That doesn't begin to tell the story. Combining techniques from the "stone age" (when counting was done, well, with stones) with techniques from the "computer age" (when numbers are represented by binary digits), we can boldly say that the subject of *combinatorics* is now in its "golden age."

Through this session you will get a glimpse of the astounding world of counting: *permutations* and *combinations*, *factorials*, "menu-type" encodings, matchmaking and honeymooning, *partitions* and *complements*, *multinomial coefficients* along the Mississippi, dogs and biscuits, balls in urns, sororities of numbers and diagonals refusing to intersect . . . will all swirl in the "magical" pool of combinatorics. Unlike magic, however, revealing the "tricks" will not leave you disappointed. You will want to make these new counting techniques your own and go forth and slay your own dragons.

1. Two Counting Conundrums

"Counting" sounds like a babyish topic: its technical name is *enumerative combinatorics*. This branch of math starts slowly, with no prerequisites at all, but quickly moves into very subtle problems that require good imagination and clarity of thought. Take, for instance, the following classic:

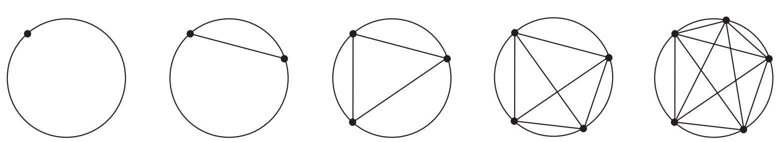

Figure 1. How many regions in the circle?

Problem 1. There are n points in a circle, all joined with line segments. Assume that no three (or more) segments intersect in the same point. How many regions inside the circle are formed this way?

The reader can quickly verify that the answers for 1 through 5 points are correspondingly $1, 2, 4, 8$, and 16 regions, all powers of 2 (cf. Fig. 1). Our intuition is screaming: "the answer for n points is 2^n regions!" Alas, this is false, as one finds out by diligently counting the 31 regions in Figure 2a.[1]

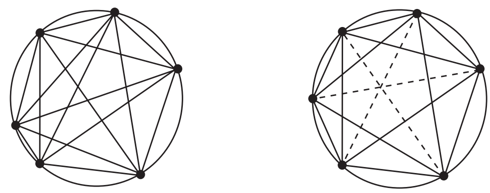

FIGURE 2. Breaking the pattern vs. breaking the rules

As a high school student, I was fascinated by this problem: I couldn't stop thinking about it until I finally solved it! That's how strong a hold a mysterious counting puzzle can have on a curious youth. Let's not spoil the fun for you either: I will leave this problem for you to struggle with and possibly revisit it in a later session.

Meanwhile, our goal for the present session will be the solution of a 1995 problem from the Czech and Slovak National Olympiad.

Problem 2. Do there exist 10,000 10-digit numbers divisible by 7, all of which can be obtained from one another by a reordering of their digits?

Before we can attack this problem (which looks like a number theory question, even though it really isn't), we need to start at the beginning. Counting involves the four basic operations of addition, subtraction, multiplication, and division. In order to count properly, you need to know which operation to do and when to do it. It turns out that multiplication is the easiest of the four, so we will start with it.

2. Multiplication, Menus, and Encoding

2.1. Menus make you multiply. Here's an easy question:

Exercise 1. A taqueria sells burritos with the following fillings: pork, grilled chicken, chicken mole, and beef. Burritos come either small, medium, or large, with or without cheese, and with or without guacamole. How many different burritos can be ordered?

SOLUTION: The answer is, of course, $4 \times 3 \times 2 \times 2 = 48$ because each different burrito is *uniquely* determined by making four decisions: filling, size, cheese, and guacamole, and those decisions involve, respectively, 4, 3, 2, and 2 choices. □

[1] Figure 2b displays 3 diagonals intersecting in a single point, which is not allowed. But even here, you will *not* count the "expected" 32 regions!

The order in which we make the above four decisions does not affect the outcome. We could have chosen the burrito's size first and the filling last, etc. Any time we can use a menu analogy, we multiply. More precisely,

PST 14. If the thing we are counting is the outcome of a *multi-stage process*, then the number of outcomes is *the product of the number of choices for each stage.*

We can solve many problems with this approach, as long as we can carefully formulate the outcomes as a "menu" process. For the following, recall what *sets* and *subsets* are: you can think of them as *collections* and *subcollections* of objects.

Exercise 2. How many subsets does a set with 8 elements have, including the empty set and the whole set itself?

SOLUTION: Can we make choosing a subset act like ordering a burrito? We need to decide which members of the set will be in the subset, and for each element this is a simple yes/no question. If the members of the set are a, b, c, d, e, f, g, and h, we need to merely ask 8 questions:

"Is a in the subset?",

"Is b in the subset?",

"Is c in the subset?",

etc. So the number of ways of performing this 8-stage process is just

$$2 \times 2 \times 2 \times 2 \times 2 \times 2 \times 2 \times 2 = 2^8 = 256.$$ □

Exercise 3. Given a pool of 30 students, how many ways can we choose a 3-person government consisting of a president, vice-president, and treasurer?

SOLUTION: We must make three decisions: who is president, who is vice-president, and who is treasurer. We do not need to decide the three positions in this order; we could instead first choose the treasurer, say. Since the order is not counted, we will fix it using alphabetical order: president, then treasurer, and finally vice-president.

Since there are 30 students, there are 30 possible choices for the first decision (president), but then there are only 29 left for the second decision (treasurer), and finally 28 for vice-president. In all, there will be $30 \times 29 \times 28$ possible governments. □

Exercise 4. In Exercise 3, we tacitly assumed that no one could hold more than one office. Verify that if we allowed this, the answer would be 30^3.

2.2. Permutations and factorials. Exercise 3 used a very common counting entity known as a *permutation*. We denote the *number of permutations of n objects taken k at a time* by the symbol $P(n, k)$. It is the answer to the question:

Problem 3. How many different ways can we choose k different things from a set of n objects, where the order of choice matters?

Solution: Our previous observations generalize to the formula

$$P(n,k) = n(n-1)(n-2)\cdots(n-k+1),$$

because there are n choices for the first thing, $(n-1)$ for the second, etc., for a total of k terms in the product. □

As a reality check, compare with the answer to Exercise 3: the number of governments was simply $P(30,3) = 30(30-1)(30-2)$.

We get an important special case when $k = n$:

$$P(n,k) = P(n,n) = n(n-1)(n-2)(n-3)\cdots 1.$$

You probably know the more common notation for this: it is $n!$, called n *factorial*. This simply counts the number of ways that n different objects can be arranged in a row.[2] For example, the number of ways we can permute the letters of ERDŐS is 5!, since there are 5 possibilities for the first letter, 4 for the second, etc.

Exercise 5. If you haven't done so already, make a table of $n!$ for $n = 1, 2, \ldots, 10$. You should memorize it at least up to $n = 7$ and passively recognize the rest.

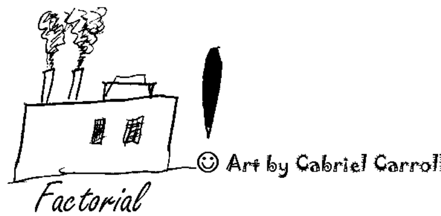

Exercise 6. 10 boys and 9 girls sit in a row of 19 seats. How many ways can this be done if

(a) all boys sit next to each other and all girls sit next to each other?
(b) each child has only neighbors of the opposite sex?

First, we warm up by solving an easier question: how many seating arrangements with no restrictions? This is just 19!. Now for the actual problem.

Solution: We need to make 3 main decisions in part (a). Either the seating will have the boys on the left or on the right; then we need to seat the boys and then the girls. The first decision has 2 possible answers; the second and third have 10! and 9!, respectively. So the answer is $2 \cdot 10! \cdot 9!$. □

Since there are more boys than girls, in part (b) the seating has to be boy-girl-boy-…-girl-boy. But we still need to seat the individual kids. As before, there are 10! possible seating arrangements for the boys and 9! arrangements for the girls. Since we need to seat boys and girls to get a single seating arrangement, this is a 2-stage process with $10! \cdot 9!$ outcomes. □

[2] Such rearrangements will play a prominent role in the Rubik's Cube sessions, where twists of the faces are represented by permutations of the facelets.

2.3. If you encode, are you a spy? I often call the menu/multiplication method *"encoding"*, because it is sometimes very helpful to think like a computer programmer, organizing information compactly using simple coding to represent outcomes.

For example, we could represent each subset in Exercise 2 with an 8-digit *binary*[3] number, with a 1 in position k if and only if the kth element was in the subset. Thus, the empty subset would be encoded by 00000000, while the subset consisting of the last two elements by 00000011.

Exercise 7. How many ways can you choose a team from 11 people where the team must have at least one person and must have a designated captain?

Solution: First, line up the 11 people in an arbitrary way (for example, by height or in alphabetical order). We will encode the team with an 11-digit binary number where digit k is a 1 if and only if person k is on the team. Since the team has at least one person, there has to be at least one digit that is 1. Underline one of these 1's to represent the captain.

How many ways can we perform this encoding? There are just 11 choices for the underlined digit 1. The remaining 10 digits of the 11-digit binary string have no restrictions; each choice of 0's and 1's yields a different subset of non-captains to join the captain. For example, a team where person #4 is the captain and the only person on the team is represented by

$$000\underline{1}0000000,$$

while

$$001\underline{1}0100000$$

encodes a 3-person team with the same captain as before, joined by persons #3 and #6. It should be clear that the number of encodings is 11×2^{10}. □

Exercise 8. In a traditional village, there are 10 young men and 10 young women. The village matchmaker arranges all the marriages. In how many ways can he pair off the 20 young people? Assume (the village is traditional) that all marriages are heterosexual, i.e., a marriage is a union of a male and a female (male-male and female-female unions are not allowed).

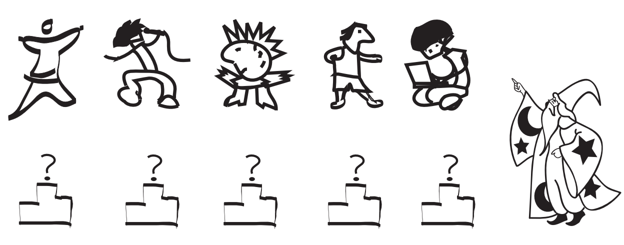

Figure 3. Which bride do I get?

[3] Instead of the 10 decimal digits 0,...,9, the *binary system* has only two digits: 0 and 1.

Solution: The most vivid visualizations are the most helpful with encoding (cf. Fig. 3 for the case of 5 men and 5 women). Imagine lining up all 10 young men, in some arbitrary-but-fixed order, such as height, wealth, looks, math ability, sense of humor, etc. Then the 10 young women line up opposite the men, in any order. Each woman marries the man directly opposite her. Clearly each different ordering of women gives rise to a different matrimonial outcome; the problem is reduced to counting such orderings, but this is just 10!. □

We learn from this and previous solutions that

PST 15. To "keep the chaos down" when counting and to have some control over what is happening, it is helpful to *order all of the involved objects/people* (or some subset of them). This order can be random or according to some specific criterion.

Exercise 9. In a not-so-traditional village, there are 10 young men and 10 young women. The village matchmaker arranges all the marriages. In how many ways can he pair off the 20 men and women, if homosexual marriages (male-male or female-female) as well as heterosexual marriages are allowed?

There are a number of different approaches; we will first give a menu-style solution and return to this problem later with other methods (cf. p. 40).

Solution: Once again, it helps to visualize an effective matching process. In a similar way to Exercise 8, start by lining up all 20 people in some arbitrary order. Everyone has to get married; so the matchmaker points to the first person and chooses for him or her a mate from one of the other people in line. Then these two persons leave (perhaps to go on a honeymoon). This first decision had 19 possible outcomes.

Now there are 18 frightened young people remaining in line. The matchmaker repeats the process, pointing at the first person in line and choosing a mate from the 17 people behind him or her in line. This process goes on until two persons remain: they are automatically mated. There will be

$$19 \cdot 17 \cdot 15 \cdots 3 \cdot 1$$

possible outcomes for marrying off the entire village in this sex-blind way. □

You may wonder – in fact, you *should* wonder – if this method counts all different marriages and counts each outcome once. This is crucial for any counting problem. Here's a way to see why it works. We need some ordering mechanism; let's use alphabetical order. Imagine a particular outcome. This is a collection of 10 pairs of married couples. Within each couple, order the spouses alphabetically. Thus if Pat marries Dana[4], we call this couple "Dana-Pat," *not* "Pat-Dana."

[4] Of course, we may assume that no two people have the same name; if they did, we simply change one of them.

Then, we list the couples in alphabetical order of the first name of each couple. For example, if we only had six people named A, B, C, D, E, and F, one possible list might be A-E, B-C, and D-F. Note that the first couple in the list will always be one starting with A, and the next couple will start with the "lowest" unused letter, etc. It should be clear that this is a one-to-one encoding of all possible marriage arrangements. Certainly, each arrangement (for the entire village) will give rise to a different list. And likewise, each list (that obeys the alphabetical order rules) will give rise to a different marriage arrangement of the village. Finally, it should be clear that this encoding method is equivalent to our original argument.

PST 16. When creating and using an encoding method, you must verify that there is a *one-to-one correspondence* between the set of encodings and all possible outcomes in the problem, i.e., that every encoding corresponds to exactly one possible outcome, and every possible outcome corresponds to exactly one encoding.

3. Addition and Partition

3.1. The limitations of multiplication. In Exercise 3, each decision in a multi-stage process influenced the number of decisions for the next stage. This is a pretty common situation, and easy to handle, *as long as the number of decisions required at each stage stays constant.*

Exercise 10. How many even 3-digit numbers have no repeating digits?

Before we solve this problem, let's look at some simpler questions. First of all, how many 3-digit numbers are there? The answer is 900, which we can see in at least two ways:

- Just count the numbers from 100 to 999, inclusive: there is a total of $999 - 100 + 1$ numbers.
- There are 9 choices for the first digit (no zero allowed) and 10 for each of the other two, yielding $9 \times 10 \times 10$.

Next, let's ask the slightly harder question: how many 3-digit numbers have no repeating digits? There are 9 choices for the first digit, as before, but this means that now there will be 9 (instead of 10) choices for the second digit and 8 (instead of 10) choices for the third, yielding $9 \times 9 \times 8 = 648$ numbers.

Finally, let's tackle the original question. How about this argument:

- Start at the *rightmost* digit; to ensure that the number is even, there are 5 choices (0, 2, 4, 6, or 8).
- There will now be 9 (instead of 10) choices for the middle digit.
- There will now be 7 (instead of 9) choices for the first digit.
- So the answer is $5 \times 9 \times 7 = 315$.

This answer is *not* correct, due to a subtle error that stems from a common beginner's mistake: trying to convert every problem into a multiplication. Can you spot the flaw?

It was clever to start with the rightmost digit, but certain decisions will affect the number of choices for future stages in *different* ways. If the rightmost digit is zero, then the first digit will have fewer restrictions, since it is never zero! The correct way to do this is to

PST 17. Break the outcomes into several separate, *mutually exclusive cases*, i.e., every possible outcome must belong to exactly one of these cases.

Solution to Exercise 10: All outcomes boil down to whether 0 is the last digit or not; so only two mutually exclusive cases are needed.

Case 1: The rightmost digit is zero. In this case, we have only two decisions to make. There are 9 numbers left, and we are free to use any of them for the other two digits. So there are 9 choices for (say) the first digit, and then 8 choices for the second digit, a total of $8 \times 9 = 72$ choices.

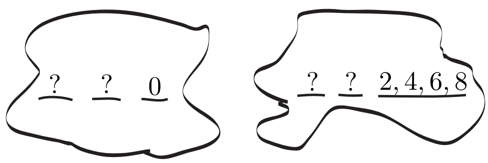

Figure 4. Two mutually exclusive cases

Case 2: The rightmost digit is not zero. Then there are 4 choices for this digit (2, 4, 6, or 8). Again, 9 numbers are left to use (including zero). Since we cannot use zero for the first digit, let's decide the first digit next: there will be 8 choices. Now there are 8 numbers left for the second digit (since zero was left in the pool). So there will be $4 \times 8 \times 8 = 256$ choices in this case.

The answer, then, is the sum $72 + 256 = 328$ numbers. □

3.2. Partition leads to addition. We had to add to get the answer in Exercise 10 above because the outcomes split into two separate cases. A decomposition of the entire set into subsets that are *pairwise disjoint* (as in PST 17) is called a *partition*: imagine partitioning an office room into smaller cubicles.

Partitioned outcomes are counted by adding. More precisely:

PST 18. Whenever we partition the outcomes of something into several cases, each requiring different counting methods, we *add the number of outcomes* in each case to get the total number of outcomes.

Note that this is *not* the same as multiple stages. In the previous exercise each case involved multiple stages. In general, you multiply when each outcome has a menu-inspired encoding. You add if you are forced to partition the outcomes into separate cases; then for each case you may use completely different strategies (for example, some may involve encoding, and others may not).

Exercise 11. An *n-bit string* is an n-digit binary number, i.e., a string of just zeros and ones. How many 10-bit strings contain *exactly* 5 consecutive zeros (no more, no less)? For example, we would not count 0000000111 (too many consecutive zeros), but we would count 1110000011 and 0011000001.

Solution: We will break the problem into 6 cases, depending on where the block of 5 zeros lives inside the larger 10-bit string. The leftmost position of the 5-zero block can be position 1, 2, 3, 4, 5, or 6. In case 1, we start with 5 zeros; then we must have a 1 (to prevent excess zeros), and then we can have anything we want for the remaining 4 digits: $\underline{0}\,\underline{0}\,\underline{0}\,\underline{0}\,\underline{0}\,\underline{1}\,\underline{?}\,\underline{?}\,\underline{?}\,\underline{?}$. So we have just 2^4 "free" choices. Case 6 works exactly the same way, by symmetry.

However, cases 2–5 are different because now the 5-digit zero block "floats" inside the larger string and needs to be surrounded by a 1 on its left and on its right. For instance, case 4 can be represented by $\underline{?}\,\underline{?}\,\underline{1}\,\underline{0}\,\underline{0}\,\underline{0}\,\underline{0}\,\underline{0}\,\underline{1}\,\underline{?}$. The remaining 3 digits are now free, so each of these cases has 2^3 outcomes. Hence the total number of outcomes is the sum

$$2^4 + 2^3 + 2^3 + 2^3 + 2^3 + 2^4 = 64. \qquad \square$$

Note that the cases of your partition cannot overlap at all. If they do, we will need to *employ subtraction to fix the overcounting.* This is the hardest operation of all, which we will discuss in depth in a later session. However, in very simple situations, we can use subtraction as a simple consequence of partitioning.

3.3. Counting the complement is very different from "counting on a compliment."

Exercise 12. Three different flavors of pie are available, and seven children are each given a slice of pie in such a way that at least two children get different flavors. How many ways can this be done?

Solution: We need to be clear on the meaning of "different" here. In human societies, children tend to be respected as individuals; so let's give each child a number and each flavor a letter. Then we can encode each outcome with a 7-letter string that uses the three letters. For example, $\underline{a}\,\underline{a}\,\underline{a}\,\underline{b}\,\underline{b}\,\underline{c}\,\underline{a}$ represents the outcome where children #1–3 got flavor a, children #4–5 got b, etc. Thus each different 7-letter string uniquely encodes a different outcome.

Let's first not worry about the condition that at least two kids get different flavors. By the simplest menu reasoning, there are 3^7 different strings. This includes outcomes with different flavors, as well as outcomes without different flavors. How many of the latter are there? Well, the only way that no two kids have a different flavor is if all the kids get the same flavor, and there are only 3 such outcomes (namely, $\underline{a}\,\underline{a}\,\underline{a}\,\underline{a}\,\underline{a}\,\underline{a}\,\underline{a}$, $\underline{b}\,\underline{b}\,\underline{b}\,\underline{b}\,\underline{b}\,\underline{b}\,\underline{b}$, $\underline{c}\,\underline{c}\,\underline{c}\,\underline{c}\,\underline{c}\,\underline{c}\,\underline{c}$).

Thus we can break up all outcomes into two cases: all flavors the same, and not all flavors the same. This is a partition of all outcomes, so the total must be 3^7. Consequently, the number of outcomes where at least two kids get different flavors (i.e., not all flavors the same) is $3^7 - 3$. □

PST 19. The method above, called *counting the complement*, is used when we can partition the total set of outcomes into the things we are interested in, and the rest (its "negation" or complement). If the total is easy to count and the negation is easy to count, then we *count the complement* and our answer is just the difference of the two.

Figure 5a depicts the general situation with a subset $\mathcal{A}$ and its complement $\mathcal{A}^c$. When you count the complement $\mathcal{A}^c$, you usually get a number which is easy-to-find or small compared to the total sum. We took advantage of this possibility in Exercise 12, where $|\mathcal{A}^c| = 3$ and $|\mathcal{A}| = 3^7 - 3$.

A different application of partitioning arises when we can divide the set of total outcomes into two sets that have the same number of elements. The following is a nice example, which first appeared as a problem in the Bay Area Math Meet (BAMM).

Exercise 13 (BAMM). How many subsets of the set $\{1, 2, 3, 4, \ldots, 30\}$ have the property that the sum of their elements is greater than 232?

Solution: Notice that in the set $S = \{1, 2, 3, 4, \ldots, 30\}$ the sum of all the elements is $1 + 2 + 3 + \cdots + 30 = 30 \cdot 31/2 = 465$.

Let A be a subset of S, let A^c denote the *complement* of A (the elements of S which are not in A), and let $\Sigma(X)$ denote the sum of the elements of any set X. Then $\Sigma(A)$ plus $\Sigma(A^c)$ must equal 465. Because $465 = 232 + 233$, if $\Sigma(A) > 232$, then $\Sigma(A^c) \leq 232$. For instance, if $A = \{2, 4, 6, \ldots, 30\}$, then $A^c = \{1, 3, 5, \ldots, 29\}$ (cf. Fig. 5b) and $\Sigma(A) = 240 > 232 > 225 = \Sigma(A^c)$.

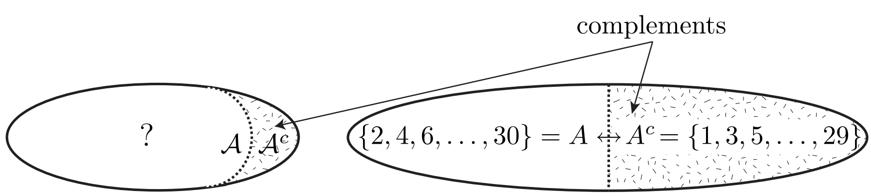

Figure 5. Complements in general and Twin subsets

In other words, there is a one-to-one correspondence between subsets whose element sum is greater than 232 and subsets whose element sum is not (namely, $A \leftrightarrow A^c$). Hence the number of subsets whose element sum is greater than 232 is exactly half of the total number of subsets of S. But the number of subsets of S is 2^{30}; so the answer is 2^{29}. □

4. Division: A Cure for Uniform Overcounting

4.1. The Mississippi formula. The number of different permutations of ERDŐS is 5!, as we saw on page 28. But

Exercise 14. How many different permutations does GAUSS have?

Solution: It is easy to see that 5! is too large by exactly a factor of 2, because GAUSS has two letters that are the same. If we temporarily distinguish the two S's by using subscripts, we see that 5! counts the permutations AS_1GUS_2 and AS_2GUS_1 as two different things when, in fact, they are indistinguishable (since the letters do not actually have subscripts). So each of the $5! = 120$ different (subscripted) permutations can be arranged into two columns where in the first column, S_1 is to the left of S_2, and in the second column it is the other way around. Hence the correct answer is $5!/2 = 60$. □

Suppose we now wished to

Exercise 15. Count the different permutations of RAMANUJAN.

Solution: This 9-letter name has 3 A's and 2 N's, so lots of overcounting occurs if we temporarily distinguish these letters and start with a provisional answer of 9!. Focus on one particular permutation, for example $A_2MA_1RJN_2UN_1A_3$. If we just permute the subscripts for each letter, we get a new rearrangement of the 9 letters, but one that is indistinguishable without subscripts. For example, one such permutation is $A_1MA_3RJN_1UN_2A_2$.

How many such rearrangements are there? All we need to do is permute the 3 subscripts of the A's (in $3! = 6$ ways) and then permute the 2 subscripts of the N's (in $2! = 2$ ways), while leaving all the other letters in place. There are $3! \times 2! = 12$ different subscripted rearrangements that appear the same without subscripts.

Hence the number of *distinguishable* permutations is $\dfrac{9!}{3!2!}$. □

This is called the *Mississippi formula*, because it can be used to count the number of distinguishable permutations of the word MISSISSIPPI. The answer is, of course, $11!/(4!4!2!)$. If we wanted to, we could include the M, which appears only once, and write instead

$$\frac{11!}{4!4!2!1!},$$

which of course leaves the value unchanged. Expressions like the fraction above are called *multinomial coefficients*. We'll see why, shortly.

4.2. Binomial coefficients. If there is such as thing as a *multi*nomial coefficient, it must have a simpler cousin. Let's look at a simple two-letter Mississippi problem:

Exercise 16. How many 10-bit strings have *exactly* 4 zeros?

Solution: This is merely asking us to count the number of ways of arranging 4 zeros and 6 ones in a string. By the Mississippi formula, it is

$$\frac{10!}{4!6!}. \tag{1}$$

Alternatively, we could start with 10 blank places and then choose locations for each zero, with the understanding that the remaining places will be occupied by ones. For the first zero, there are 10 possible places. The next zero has 9 possible locations. The third and fourth zeros have, respectively, 8 and 7 possible slots. By the menu principle, this 4-stage process has $10 \cdot 9 \cdot 8 \cdot 7$, or $P(10, 4)$, possible outcomes. However, the order of choice doesn't matter. For example, we could have picked place numbers $4, 9, 2, 3$, in order, resulting in zeros at those locations. But we could have also chosen $9, 2, 3, 4$, which would give the same result. So the product $10 \cdot 9 \cdot 8 \cdot 7$ has overcounted by a factor of 4!, and the actual answer is

$$\frac{10 \cdot 9 \cdot 8 \cdot 7}{4!}. \tag{2}$$

It is easy to see that this expression is equal to (1) above. □

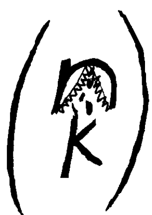

"n chews k"

An unordered pear

☺ Art by Gabriel Carroll

Expressions such as (1) are called *binomial coefficients*, denoted by

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

and pronounced as *"n choose k"*. Exercise 16 demonstrated the following facts about binomial coefficients.

Theorem 1 (Binomial Coefficients). *The quantity* $\binom{n}{k}$

(a) *counts the number of ways to permute an n-bit string consisting of k zeros and $n-k$ ones,*

(b) *also counts the number of different ways to choose k places out of an n-place string where the order of choice doesn't matter, and*

(c) *satisfies* $\binom{n}{k} = \frac{P(n,k)}{k!} = \frac{n(n-1)\cdots(n-k+1)}{k!}$.

Theorem 1(b) introduces a very important concept. The selection of k "places" from a pool of n places where order is not important is also called a *combination*, often to distinguish it from a permutation (where order counts).

This is somewhat old-fashioned terminology, and we have always found it confusing. We prefer the more rigorous interpretation that

Theorem 1′. $\binom{n}{k}$ *counts the number of different k-element subsets chosen from an n-element set.*

Note that for *sets*, order *doesn't* matter. For example, the sets $\{a, b, c\}$ and $\{b, a, c\}$ are the same. This subset interpretation has many interesting consequences. Here are a few for you to verify. If you have trouble, read the solution below. Try to minimize, if you can, use of formulas analogous to (1) on page 36 and algebraic manipulation. Instead, try to rely on Theorem 1′, that is, describe the involved quantities as certain subsets of a bigger set.

Theorem 2 (Properties of Binomial Coefficients).

(a) *For any nonnegative integers n and r with $r \leq n$, $\binom{n}{r} = \binom{n}{n-r}$.*
(b) *Because $\binom{n}{n} = \binom{n}{0} = 1$, the only logical value for $0!$ is 1.*
(c) *For any positive integer n with $r < n$, $\binom{n}{r} + \binom{n}{r+1} = \binom{n+1}{r+1}$.*
(d) $\binom{n}{0} + \binom{n}{1} + \binom{n}{2} + \cdots + \binom{n}{n} = 2^n$ *for any positive integer n.*

CONCRETE "SOLUTION": In part (a), $\binom{10}{3} = \binom{10}{7}$ because picking 3 places in a 10-bit string to be 0's is equivalent to picking 7 places to be 1's. ♢

$\binom{n}{0}$ has to equal 1 in part (b), since it counts the number of 0-element subsets of an n-element set. There is only one, namely the empty set. Plugging $r = 0$ into formula (1) yields $0! = 1$. Note that we are not compelled to define $0!$, but if we do, we must assign it the value of 1 in order to be consistent with formula (1). ♢

Let's keep things concrete with a specific case in part (c). Why is

$$\binom{13}{4} + \binom{13}{5} = \binom{14}{5}?$$

The right-hand side (RHS) counts the number of ways of awarding an ice cream cone to each of 5 lucky children from a class of 14. Suppose one of the children is named Ramanujan. We can partition the outcomes into two cases: either Ramanujan gets ice cream or Ramanujan does not get ice cream.

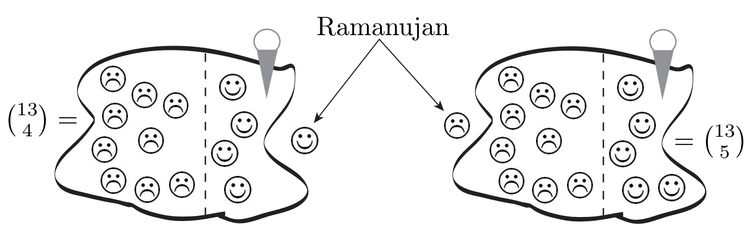

FIGURE 6. Does Ramanujan get ice cream?

For the first case, we need to choose 4 more lucky kids from among the 13 kids remaining. For the second case, we need to eliminate the unlucky

Ramanujan from consideration and choose 5 lucky kids from the remaining 13. These cases correspond to the two terms of the left-hand side (LHS). ♢

We solve a concrete case for part (d) also. If $n = 11$, the RHS represents (by Exercise 2) the number of subsets of an 11-element set, including the empty set and the entire set. The LHS is a sum representing partitioning these subsets into 12 cases: the number of empty subsets, the number of 1-element subsets, the number of 2-element subsets, etc. ♢

4.3. Pascal's Triangle all over. Seeing the properties in Theorem 2, some readers may recognize the connection between the so-called *Pascal's Triangle* and our binomial coefficients. Briefly, Pascal's Triangle is a triangle of numbers: its top row consists of a single 1, and each consecutive row has one more number than the previous; the first and the last number in each row are 1's, while any other number is the sum of the two numbers directly above it.

$$\begin{array}{ccccccccccc} & & & & & 1 & & & & & \\ & & & & 1 & & 1 & & & & \\ & & & 1 & & 2 & & 1 & & & \\ & & 1 & & 3 & & 3 & & 1 & & \\ & 1 & & 4 & & 6 & & 4 & & 1 & \\ 1 & & 5 & & 10 & & 10 & & 5 & & 1 \end{array} \quad \longleftrightarrow \quad \begin{array}{ccccccccccc} & & & & & \binom{0}{0} & & & & & \\ & & & & \binom{1}{0} & & \binom{1}{1} & & & & \\ & & & \binom{2}{0} & & \binom{2}{1} & & \binom{2}{2} & & & \\ & & \binom{3}{0} & & \binom{3}{1} & & \binom{3}{2} & & \binom{3}{3} & & \\ & \binom{4}{0} & & \binom{4}{1} & & \binom{4}{2} & & \binom{4}{3} & & \binom{4}{4} & \\ \binom{5}{0} & & \binom{5}{1} & & \binom{5}{2} & & \binom{5}{3} & & \binom{5}{4} & & \binom{5}{5} \end{array}$$

FIGURE 7. Pascal's Triangle and Binomial coefficients

Figure 7 displays the first 6 rows of Pascal's Triangle and the corresponding binomial coefficients. As one can easily check, the two triangles have identical entries. This correspondence means that whatever we can prove about binomial coefficients can be translated in terms of Pascal's Triangle, and vice versa. Thus, property (a) in Theorem 2 indicates the symmetry of Pascal's Triangle across a vertical line; (b) tells us about the end 1's on each row; (c) is the defining addition property of Pascal's Triangle; and (d) gives the sum of numbers on each row of Pascal's Triangle.

The combinatorics of Pascal's Triangle is very rich, and it deserves a separate session of its own. Without further detours, we shall leave this topic for now and return to our binomial coefficients.

4.4. Some mathematical "etymology." It is time to explain why binomial coefficients have this name. Consider raising the *binomial* $(x + y)$ to the second power. We start with the product

$$(x + y)^2 = (x + y)(x + y).$$

Now let's expand it like a beginning algebra student, writing each "raw" term as it is first spewed out *without combining like terms* or switching the order of multiplications. Our expansion will be $x(x+y)+y(x+y) = xx+xy+yx+yy$. Of course, we then clean this up to get $x^2 + 2xy + y^2$.

Now try it for a higher power. Consider

$$(x+y)^7 = (x+y)(x+y)(x+y)(x+y)(x+y)(x+y)(x+y).$$

If we multiply it all out, without combining like terms, we now get $2^7 = 128$ raw terms: all *monomials* with coefficients of 1, such as $xxxxxxx$ or $xxxyyyx$ or $xxyxyyx$. Notice that the last two monomials are like terms, since both simplify to x^4y^3. Also note that each 7-term string with 4 x's and 3 y's represents a different raw term. For example, the string $xxxxyyy$ arose by multiplying the x's from the first four binomials of the product with the y's from the last three binomials. The string $xyxyxyx$ arose by multiplying the x from the first term with the y from the second term, etc. Thus the number of raw terms is equal to the number of 7-term strings made of x's and y's. For instance, the number of raw terms that simplify to x^4y^3 is, by the Mississippi formula, equal to $\binom{7}{4}$ or, equivalently, $\binom{7}{3}$. Thus the coefficient of x^4y^3 in the expansion will be $\binom{7}{4} = \binom{7}{3}$. Similarly, the coefficients of x^2y^5 and of y^7 will be $\binom{7}{5}$ and $\binom{7}{7}$, respectively. More generally, $x^{7-k}y^k$ will appear with coefficient $\binom{7}{k}$.

We can write out the whole expansion, then, as

$$(x+y)^7 = \binom{7}{0}x^7y^0 + \binom{7}{1}x^6y^1 + \binom{7}{2}x^5y^2 + \cdots + \binom{7}{7}x^0y^7.$$

This is an example, for the exponent 7, of the so-called *Binomial Theorem*[5], which expands the *binomial* expression $(x+y)^n$.

Now you can probably imagine where the name *multinomial* coefficients comes from. For instance, if you expand the 4-variable expression $(m+i+s+p)^{11}$ and combine the like terms, the monomial corresponding to the permutations of *"Mississippi"* will appear as $\frac{11!}{1!4!4!2!}m^1i^4s^4p^2$. With this said, it's time for some popular folklore: a student response to an exam question.

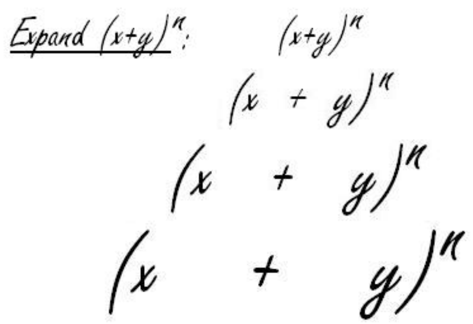

☺Alternate Binomial Theorem

[5]The Binomial Theorem can be succinctly written as $(x+y)^n = \sum_{k=0}^{n}\binom{n}{k}x^{n-k}y^k$.

5. Balls in Urns and Other Applications

Armed with encoding, partitioning, and binomial coefficients, we can now investigate some really interesting problems, including Problem 2. The key is to use creative encoding.

5.1. Choosing a honeymoon location. Let's begin by revisiting Exercise 9, the marriage problem in the non-traditional village. Suppose each couple went on a different honeymoon. We could label each honeymoon location with one of the ten letters $\mathcal{A}$, $\mathcal{B}$, $\mathcal{C}$, ..., $\mathcal{J}$.

Now line up the 20 young people in some arbitrary order. Imagine that the matchmaker has twenty stickers: two labeled "$\mathcal{A}$", two labeled "$\mathcal{B}$", etc. The matchmaker then goes around putting one sticker on each person. That tells that person where he or she will honeymoon. Clearly this is a way to form all possible marriages, and by the Mississippi formula there will be

$$\frac{20!}{\underbrace{2!2!\cdots 2!}_{10}} = \frac{20!}{2^{10}}$$

such marriages. However, we overcounted, since we are distinguishing the honeymoon locations. For example, the outcome where Pat and Dana go to honeymoon $\mathcal{C}$ and the outcome where Pat and Dana go to honeymoon $\mathcal{H}$ are counted as two different outcomes. To cure this overcounting, we need to divide by 10!, since there were 10 different honeymoon locations. So the correct answer is

$$\frac{20!}{10!2^{10}}.$$

Here's another approach. Choose two young people to go on honeymoon $\mathcal{A}$. This can be done in $\binom{20}{2}$ different ways. Now pick two people from the remaining group to go on honeymoon $\mathcal{B}$. This can be done in $\binom{18}{2}$ ways. Continuing this process, the number of ways of arranging marriages where the honeymoons are *distinguishable* will be

$$\binom{20}{2}\binom{18}{2}\binom{16}{2}\cdots\binom{4}{2}\binom{2}{2},$$

and once again we need to divide by 10! to get the answer to the original question. It is a simple and fun exercise to verify that, indeed, the three approaches yield the same number.

Exercise 17. Check that $\dfrac{20!}{10!2^{10}} = \dfrac{1}{10!}\binom{20}{2}\binom{18}{2}\cdots\binom{2}{2} = 19\cdot 17\cdots 3\cdot 1.$

5.2. Geometry falls under the spell of combinatorics. Recall the enticing Problem 1, which asked for the number of regions in a circle. No, we shall not break our promise here – this problem is still yours to hack. But there are so many geometry problems of a similar spirit that they all form a separate branch of mathematics named *combinatorial geometry*. This session would be amiss without at least one such example.

Exercise 18. The n vertices of a polygon are arranged on the circumference of a circle so that no three diagonals intersect in the same point.

✎ (a) How many diagonals does the polygon have?
(b) How many intersection points do these diagonals have?

SOLUTION: An easy application of the definition of binomial coefficients tells us that there are $\binom{n}{2}$ pairs of vertices and hence there are that many segments joining them. Thus, there are $\binom{n}{2} - n$ diagonals (we had to subtract the n sides of the polygon). Let's call this number d.

At first you may think that the answer in part (b) is just $\binom{d}{2}$, but not all diagonals intersect. However, each intersection point in the interior is the intersection of two diagonals. The endpoints of these two diagonals are the vertices of a quadrilateral. Each choice of 4 vertices of the polygon gives rise to a different quadrilateral (cf. Fig. 8a). Once you focus on these quadrilaterals, it should be clear that there is a one-to-one correspondence between quadrilaterals formed from vertices of the polygon and interior diagonal intersection points. For instance, look at the pentagon in Figure 1: you can easily count 5 such quadrilaterals, each accounting for one of the 5 interior intersection points. So the answer is $\binom{n}{4}$. □

As a further example, check out the case of a hexagon in Figure 2a: we have $9 = \binom{6}{2} - 6$ diagonals and $15 = \binom{6}{4}$ intersection points of these diagonals.

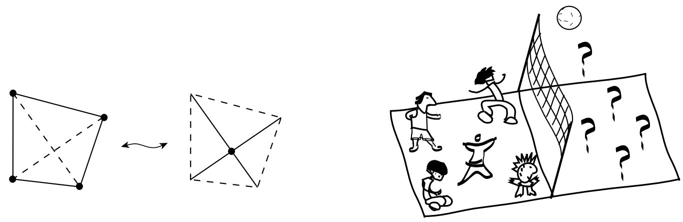

FIGURE 8. Quadrilaterals and Volleyball teams

5.3. Teaming up for volleyball.

✎ **Exercise 19.** How many ways can 10 people form two teams of 5?

SOLUTION: You may think the answer is $\binom{10}{5}$, but this exactly double-counts! To see why, note that $\binom{10}{5}$ counts each 5-person choice as different from the complement, which is incorrect: once you choose a team, its complement is already automatically selected as the other team. Alternatively, one may realize that $\binom{10}{5}$ counts the number of ways of choosing 5 people to play in the *left* half of the court (cf. Fig. 8b). The correct answer is $\binom{10}{5}/2$. □

This is really an easier version of the non-traditional marriage problem (a "marriage" is interpreted as a team of 5 people), but it may seem harder.

5.4. Dogs and biscuits are transformed into "urns" and "balls." The next problem, on the other hand, seems quite innocent, but it is actually rather tricky. Think about it carefully before reading the solution.

Problem 4. How many ways can 7 dogs consume 10 dog biscuits? The dogs are distinguishable; the biscuits are indistinguishable. Dogs do not share.

Solution: Let's arrange the dogs is some order (e.g., alphabetical). The difficulty is that we do not know how many biscuits an individual dog eats. Dog #1 could eat between 0 and 10 biscuits, and each of these 11 outcomes will affect the possibilities for the remaining dogs.

Imagine a sample outcome and try to think of how we can encode generic outcomes. For example,

dog	1	2	3	4	5	6	7
biscuits	0	0	3	1	0	2	4

means that dog #1 ate no biscuits, while dog #7 ate 4, etc. What we want to count is the number of different tables of this sort. Instead of using numerals to indicate the number of biscuits, we can achieve some uniformity by replacing numerals with an equivalent number of symbols. For example, the following string means the same thing as the table above:

`|||bbb|b||bb|bbbbb|`

Instead of the numeral "3", we write 3 `b`'s. We retain the `|` symbol to indicate the boundary of a table cell. This way, if two of these symbols are next to each other with nothing in between, it means that the table cell has a zero in it. For example, the first three symbols `|||` translate into "zeros in the first two cells of the table."

Notice that our encoding method has more symbols than needed. The first and last symbols are always `|`'s. Thus we can eliminate them – they provide no information. So, for example, the string `|b|b|bbb|||bbbbb` is equivalent to the table

dog	1	2	3	4	5	6	7
biscuits	0	1	1	3	0	0	5

In this case, the first `|` is the *rightmost* boundary of the first cell, indicating a zero in the first cell, i.e., no biscuits for the first dog. If you have trouble understanding, here's another way to think about it. Start with the following:

`| | | | | |`

These six vertical lines are the boundary lines of *seven* cells. Then we drop 10 `b`'s into the seven cells (the horizontal line is the "floor" which prevents the `b`'s from falling through). If n `b`'s lie to the left of the first vertical line, it means that the first cell has value n. If m `b`'s lie between the first and second lines, it means that the second cell has the value m, etc. If j `b`'s lie to the right of the rightmost vertical line, then the seventh cell has value j.

We don't need the horizontal line, of course; it is just to make us think of physical cells for the b's to drop down into. All we really are doing is creating a 16-symbol word consisting of 10 b's and 6 |'s. For instance, the last table can be encoded on 16 slots as follows:

```
| b | b | b b b | | | b b b b b
```

By the Mississippi formula, this can be done in $\binom{16}{10} = \binom{16}{6}$ ways. □

This problem of course generalizes and gives rise to a useful formula that is called, among other things, the *Balls in Urns* or *Stars-and-Bars* formula. We prefer the former name because it is easier to remember:

Theorem 3 (Balls in Urns). *The number of ways to distribute b indistinguishable balls among u distinguishable urns is*

$$\binom{b+u-1}{b} = \binom{b+u-1}{u-1}.$$

I remember it by thinking, "It's a hard formula; so the top is not something obvious like $b+u$. For the bottom, just think about *dropping balls.*" You'd never put u there by mistake!

The Balls in Urns formula incidentally solves another equivalent, very frequently occurring problem.

Exercise 20. How many ways can we choose 7 people out of 26 people, allowing for repetitions?

Solution: Our intuition may suggest that the 26 people are the "balls," being "dropped" into 7 winning slots. However, our intuition is wrong. Imagine that we are giving 7 identical awards to the 26 people, and the number of times a person is chosen corresponds to the number of awards he/she receives. Thus, the 7 awards are the "balls," and the 26 people are the "urns." For instance, if the people have names starting with the different 26 Latin alphabet letters, then in Figure 9, $\mathcal{A}my$ received three awards, $\mathcal{B}ob$ – two awards, while $\mathcal{I}van$ and $\mathcal{L}ily$ – one award each. We can string together a name from these letters, namely, $\mathcal{ALI}$ $\mathcal{BABA}$, but of course, the order of the letters does not matter.

So, the answer is $\binom{7+26-1}{7} = \binom{32}{7} = 3{,}365{,}856.$ □

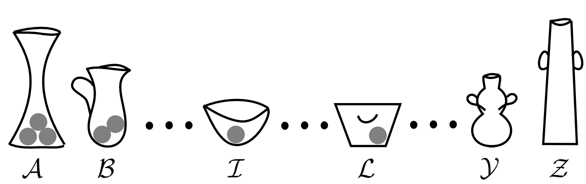

Figure 9. $\mathcal{ALI}$ $\mathcal{BABA}$

6. Sororities of Numbers: A Promise Fulfilled

Now we can finally tackle Problem 2. Recall that this problem asked whether there exist 10,000 10-digit numbers divisible by 7, all of which can be obtained from one another by a reordering of their digits.

6.1. Can "0" be a leader? Let us first solve a slightly simpler modified version of the problem: we shall allow "10-digit" numbers to start with the digit 0, e.g., we shall consider the number 1102 as a 10-digit number and write it as the string 0000001102. Thus, we have a total of 10^{10} numbers to work with.

Ignore the issue of divisibility by 7 for a moment, and concentrate on understanding sets of numbers that "can be obtained from one another by a reordering of their digits." Let us call two 10-digit numbers that can be obtained from one another in this way "sisters," and let us call a set that is as large as possible whose elements are all sisters a "sorority." For example, 1,111,233,999 and 9,929,313,111 are sisters who belong to a sorority with $\frac{10!}{4!3!2!}$ members, since the membership of the sorority is just the number of ways of permuting the digits. The sororities have vastly different sizes. The most "exclusive" sororities have only one member – for example, the sorority consisting entirely of 6,666,666,666; yet one sorority has 10! members – the one containing 1,234,567,890 (cf. Fig. 10).

In order to solve this problem, we need to show that there is a sorority with at least 10,000 members that are divisible by 7. One approach is to look for big sororities (like the one with 10! members), but it is possible (even though it seems unlikely) that somehow most of the members will not be multiples of 7. Instead,

PST 20. The crux idea is that *it is not the size of the sororities that really matters, but how many sororities there are.* If the number of sororities is fairly small, then even if the multiples of 7 are dispersed very evenly, "enough" of them will land in some sorority.

Let's make this more precise. Suppose it turned out that there were only 100 sororities (of course there are more than that). The number of multiples of 7 is $\lfloor (10^{10} - 1)/7 \rfloor + 1 = 1{,}428{,}571{,}429$. By the Pigeonhole Principle,[6] at least one sorority will contain $\lceil 1{,}428{,}571{,}429/100 \rceil$ multiples of 7, which is way more than we need.[7] In any event, we have the penultimate step to work toward: compute (or at least estimate) the number of sororities.

We can compute the exact number. Each sorority is uniquely determined by its collection of 10 digits *where repetition is allowed*. For example, one (highly exclusive) sorority can be named "ten 6's," while another is called

[6] Part II of the Proofs session discusses the Pigeonhole Principle in detail.

[7] $\lfloor a \rfloor$ rounds a down to the closest integer $\leq a$, while $\lceil a \rceil$ rounds a up to the closest integer $\geq a$, e.g., $\lfloor 5.7 \rfloor = 5$ and $\lceil 3.2 \rceil = 4$.

"three 4's, a 7, two 8's, and four 9's." So now the question becomes: In how many different ways can we choose 10 digits, with repetition allowed? Just as in Exercise 20, this is equivalent to putting 10 "balls" (one for each digit *place*) into 10 "urns" (one for each *decimal* digit $0, 1, 2, \ldots, 9$). By the Balls in Urns formula, this is $\binom{19}{10} = 92{,}378$.

Finally, we conclude that there will be a sorority with at least

$$\lceil 1{,}428{,}571{,}429/92{,}378 \rceil = 15{,}465$$

members divisible by 7. This is greater than 10,000. □

FIGURE 10. Sororities of various sizes: from smallest to largest

6.2. "0" is not born to lead. Now back to our original problem. A couple of the calculations above must be adjusted, but the main idea of the proof remains the same. Since we do not allow a 10-digit number to begin with 0, we are dealing only with the numbers from 1,000,000,000 to 9,999,999,999. So there will be $\lfloor (10^{10} - 1)/7) \rfloor - \lfloor (10^9 - 1)/7 \rfloor = 1{,}285{,}714{,}286$ of them divisible by 7.

Even though it has no bearing on our calculations, it is worth observing that the big sorority containing 1,234,567,890 will now have "only" $10! - 9! = 9(9!)$ members, as we must exclude all 9! strings starting with 0.

One would also expect big changes in the number of sororities, but . . . we "lose" only one sorority, namely, the sorority corresponding to the 0-string 0000000000 (cf. Fig. 10). Thus, we have $\binom{19}{10} - 1 = 92{,}377$ sororities total.

To complete the solution, we only have to divide:

$$\lceil 1{,}285{,}714{,}286/92{,}377 \rceil = 13{,}919.$$

Thus, some sorority must contain at least 13,919 members divisible by 7, so the answer is Yes. □