Olympiad Combinatorics

Pranav A. Sriram

August 2014

Copyright notices

About the Author

Pranav Sriram graduated from high school at The International School Bangalore, India, and will be a Freshman at Stanford University this Fall.

1. ALGORITHMS

Introduction

Put simply, an **algorithm** is a procedure or set of rules designed to accomplish some task. Mathematical algorithms are indispensable tools, and assist in financial risk minimization, traffic flow optimization, flight scheduling, automatic facial recognition, Google search, and several other services that impact our daily lives.

Often, an algorithm can give us a deeper understanding of mathematics itself. For instance, the famous Euclidean algorithm essentially lays the foundation for the field of number theory. In this chapter, we will focus on using algorithms to prove combinatorial results. We can often prove the existence of an object (say, a graph with certain properties or a family of sets satisfying certain conditions) by giving a procedure to explicitly construct it. These proofs are hence known as constructive proofs. Our main goals in this chapter will be to study techniques for designing algorithms for constructive proofs, and proving that they actually work.

In this chapter, and throughout the book, the emphasis will be on *ideas*. What can we *observe* while solving a given problem? How can disparate ideas and observations be pieced together cohesively to motivate a solution? What can we learn from the solution of one problem, and how may we apply it to others in the future? Each problem in this book is intended to teach some lesson - this may be a combinatorial trick or a new way of looking at problems. We suggest that you keep a log of new ideas and insights into combinatorial structures and problems that you encounter or come up with yourself.

Greedy Algorithms

Be fearful when others are greedy and greedy when others are fearful - Warren Buffet

Greedy algorithms are algorithms that make the best possible short term choices, hence in each step maximizing short term gain. They aren't always the optimal algorithm in the long run, but often are still extremely useful. The idea of looking at extreme elements (that are biggest, smallest, best, or worst in some respect) is central to this approach.

Example 1
In a graph $\boldsymbol{G}$ with n vertices, no vertex has degree greater than Δ. Show that one can color the vertices using at most Δ+1 colors, such that no two neighboring vertices are the same color.

Answer:
We use the following greedy algorithm: arrange the vertices in an arbitrary order. Let the colors be 1, 2, 3… Color the first vertex with color 1. Then in each stage, take the next vertex in the order and color it with the smallest color that has not yet been used on any of its neighbors. Clearly this algorithm ensures that two

adjacent vertices won't be the same color. It also ensures that at most Δ+1 colors are used: each vertex has at most Δ neighbors, so when coloring a particular vertex *v*, at most Δ colors have been used by its neighbors, so at least one color in the set {1, 2, 3, ..., Δ+1} has **not** been used. The minimum such color will be used for the vertex *v*. Hence all vertices are colored using colors in the set {1, 2, 3,..., Δ+1} and the problem is solved. ■

Remark: The "greedy" step here lies in always choosing the color with the smallest number. Intuitively, we're saving larger numbers only for when we really need them.

Example 2 [Russia 2005, Indian TST 2012, France 2006]
In a 2 x *n* array we have positive reals such that the sum of the numbers in each of the *n* columns is 1. Show that we can select one number in each column such that the sum of the selected numbers in each row is at most $(n+1)/4$.

0.4	0.7	0.9	0.2	0.6	0.4	0.3	0.1
0.6	0.3	0.1	0.8	0.4	0.6	0.7	0.9

Figure 1.1: 2xn array of positive reals, n=8

Answer:
A very trivial greedy algorithm would be to select the smaller number in each column. Unfortunately, this won't always work, as can easily be seen from an instance in which all numbers in the top row are 0.4. So we need to be more clever. Let the numbers in the top row in **non-decreasing order** be a_1, a_2, ..., a_n and the corresponding numbers in the bottom row be b_1, b_2, ..., b_n (in non-increasing order, since $b_i = 1 - a_i$). Further suppose that the sum of the numbers in the top row is less than or equal to that of the bottom row. The idea of ordering the variables is frequently used, since it provides some structure for us to work with.

Our algorithm is as follows: Starting from a_1, keep choosing the smallest remaining element in the top row as long as possible. In other words, select $a_1, a_2, \ldots, a_k$ such that $a_1 + a_2 + \ldots + a_k \leq \frac{n+1}{4}$ but $a_1 + a_2 + \ldots + a_k + a_{k+1} > \frac{n+1}{4}$. Now we cannot select any more from the top row (as we would then violate the problem's condition) so in the remaining columns choose elements from the bottom row. We just need to prove that the sum of the chosen elements in the bottom row is at most $\frac{n+1}{4}$. Note that a_{k+1} is at least the average of $a_1, a_2, \ldots, a_k, a_{k+1}$ which is more than $\frac{n+1}{4(k+1)}$.

Hence $b_{k+1} = (1 - a_{k+1}) < 1 - \frac{n+1}{4(k+1)}$. But b_{k+1} is the largest of the chosen elements in the bottom row. So the sum of the chosen elements in the bottom row cannot exceed $(1 - \frac{n+1}{4(k+1)}) \times (n-k)$. We leave it to the reader to check that this quantity cannot exceed $(n+1)/4$. ■

Remark: One of the perks of writing a book is that I can leave boring calculations to my readers.

Example 3

In a graph ***G*** with V vertices and E edges, show that there exists an induced subgraph ***H*** with each vertex having degree at least E/V. (In other words, a graph with average degree d has an induced subgraph with minimum degree at least $d/2$).

Answer:

Note that the average degree of a vertex is $2E/V$. Intuitively, we should get rid of 'bad' vertices: vertices that have degree $< E/V$. Thus a natural algorithm for finding such a subgraph is as follows: start with the graph ***G***, and as long as there exists a vertex with degree $< E/V$, delete it. However, remember that while deleting a vertex we are also deleting the edges incident to it, and in the process vertices that were initially not 'bad' may become bad in the subgraph formed. What if we end up with a graph with all vertices bad? Fortunately, this won't happen: notice that the *ratio* *of edges/vertices is strictly increasing* (it started at E/V and each time we deleted a vertex, less than E/V edges were deleted by the condition of our algorithm). Hence, it is impossible to reach a stage when only 1 vertex is remaining, since in this case the edges/vertices ratio is 0. So at some point, our *algorithm must terminate*, leaving us with a graph with more than one vertex, all of whose vertices have degree at least E/V. ■

Remark: This proof used the idea of monovariants, which we will explore further in the next section.

The next problem initially appears to have nothing to do with algorithms, but visualizing what it actually means allows us to think about it algorithmically. The heuristics we develop lead us to a very simple algorithm, and proving that it works isn't hard either.

Example 4 [IMO shortlist 2001, C4]
A set of three nonnegative integers $\{x, y, z\}$ with $x < y < z$ satisfying $\{z\text{-}y, y\text{-}x\} = \{1776, 2001\}$ is called a *historic* set. Show that the set of all nonnegative integers can be written as a disjoint union of historic sets.

Remark: The problem is still true if we replace {1776, 2001} with an arbitrary pair of distinct positive integers $\{a, b\}$. These numbers were chosen since IMO 2001 took place in USA, which won independence in the year 1776.

Answer:
Let 1776 = a, 2001 =b. A historic set is of the form $\{x, x+a, x+a+b\}$ or $\{x, x+b, x+a+b\}$. Call these *small sets* and *big sets* respectively. Essentially, we want to cover the set of nonnegative integers using historic sets. To construct such a covering, we visualize the problem as follows: let the set of nonnegative integers be written in a line. In each move, we choose a historic set and cover these numbers on the line. Every number must be covered at the end of our infinite process, but no number can be covered twice (the

historic sets must be disjoint). We have the following *heuristics*, or intuitive guidelines our algorithm should follow:

Heuristic 1: At any point, the smallest number not yet covered is the most "unsafe"- it may get trapped if we do not cover it (for example, if x is the smallest number not yet covered but $x+a+b$ has been covered, we can never delete x). Thus in each move we should choose x as the smallest uncovered number.

Heuristic 2: From heuristic 1, it follows that our algorithm should prefer small numbers to big numbers. Thus it should prefer small sets to big sets.

Based on these two simple heuristics, we construct the following greedy algorithm that minimizes short run risk: in any move, choose x to be the **smallest number not yet covered**. Use the small set if possible; only otherwise use the big set. We now show that this simple algorithm indeed works:

Suppose the algorithm fails (that is, we are stuck because using either the small or big set would cover a number that has already been covered) in the (n+1)th step. Let x_i be the value chosen for x in step i. Before the (n+1)th step, x_{n+1} hasn't yet been covered, by the way it is defined. $x_{n+1} + a + b$ hasn't yet been covered since it is larger than all the covered elements ($x_{n+1} > x_i$ by our algorithm). So the problem must arise due to $x_{n+1} + a$ and $x_{n+1} + b$. Both of these numbers must already be covered. Further, $x_{n+1} + b$ must have been the largest number in its set. Thus the smallest number in this set would be $x_{n+1} + b - (a+b) = x_{n+1} - a$. But at this stage, x_{n+1} was not yet covered, so the small set should have been used and x_{n+1} should have been covered in that step. This is a contradiction. Thus our supposition is wrong and the algorithm indeed works. ■

Remark: In an official solution to this problem, the heuristics would be skipped. Reading such a solution would leave you thinking "Well that's nice and everything, but how on earth would anyone come up with that?" One of the purposes of this book is to show that Olympiad solutions don't just "come out of nowhere". By including heuristics and observations in our solutions, we hope that readers will see the motivation and the key ideas behind them.

Invariants and Monovariants

Now we move on to two more extremely important concepts: invariants and monovariants. Recall that a monovariant is a quantity that changes monotonically (either it is non-increasing or non-decreasing), and an invariant is a quantity that doesn't change. These concepts are especially useful when studying combinatorial processes. While constructing algorithms, they help us in several ways. Monovariants often help us answer the question "Well, what do we do now?" In the next few examples, invariants and monovariants play a crucial role in both constructing the algorithm and ensuring that it works.

Example 5 [IMO shortlist 1989]
A natural number is written in each square of an m x n chessboard. The allowed move is to add an integer k to each of two adjacent numbers in such a way that nonnegative numbers are obtained (two squares are adjacent if they share a common side). Find a necessary and sufficient condition for it to be possible for all the numbers to be zero after finitely many operations.

Answer:
Note that in each move, we are adding the same number to 2 squares, one of which is white and one of which is black (if the chessboard is colored alternately black and white). If S_b and S_w denote the sum of numbers on black and white squares respectively, then $S_b - S_w$ is an invariant. Thus if all numbers are 0 at the end, $S_b - S_w = 0$ at the end and hence $S_b - S_w = 0$ in the

beginning as well. This condition is thus necessary; now we prove that it is sufficient.

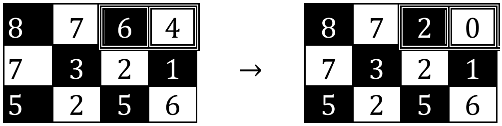

Figure 1.2: A move on the mxn board

Suppose a, b, c are numbers in cells A, B, C respectively, where A, B, C are cells such that A and C are both adjacent to B. If $a \leq b$, we can add ($-a$) to both a and b, making a 0. If $a \geq b$, then add (a-b) to b and c. Then b becomes a, and now we can add ($-a$) to both of them, making them 0. Thus we have an algorithm for reducing a positive integer to 0. Apply this in each row, making all but the last 2 entries 0. Now all columns have only zeroes except the last two. Now apply the algorithm starting from the top of these columns, until only two adjacent nonzero numbers remain. These last two numbers must be equal since $S_b = S_w$. Thus we can reduce them to 0 as well. ■

The solution to the next example looks long and complicated, but it is actually quite intuitive and natural. We have tried to motivate each step, and show that each idea follows quite naturally from the previous ones.

Example 6 [New Zealand IMO Training, 2011]
There are $2n$ people seated around a circular table, and m cookies are distributed among them. The cookies can be passed under the following rules:
(a) Each person can only pass cookies to his or her neighbors
(b) Each time someone passes a cookie, he or she must also eat a cookie
Let A be one of these people. Find the least m such that no matter how m cookies are distributed initially, there is a strategy to pass cookies so that A receives at least one cookie.

Answer:

We begin by labeling the people A_{-n+1}, A_{-n+2}, …., A_0, A_1, A_2, …, A_n, such that $A = A_0$. Also denote $A_{-n} = A_n$. We assign weight $1/2^{|i|}$ to each cookie held by person A_i. Thus for example, if A_3 passes a cookie to A_2, that cookie's weight increases from $1/8$ to $1/4$. Note that A_3 must also eat a cookie (of weight $1/8$) in this step. Thus we see in this case the sum of the weights of all the cookies has remained the same. More precisely, if A_i has a_i cookies for each i, then the total weight of all cookies is

$$W = \sum_{i=-n+1}^{n} \frac{a_i}{2^{|i|}}$$

Whenever a cookie is passed towards A_0 (from $A_{\pm i}$ to $A_{\pm(i-1)}$ for i positive) one cookie is eaten and another cookie doubles its weight, so the total weight remains invariant. If a cookie is passed away from A, then the total weight decreases. Thus the total weight is indeed a monovariant.

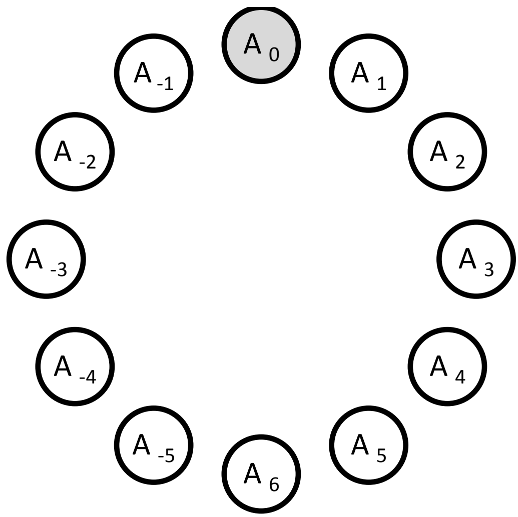

Figure 1.3: Labeling scheme to create a monovariant (n=5)

If $m < 2^n$, then if all the cookies are initially given to A_n, the initial total weight is $m/2^n < 1$. Therefore the total weight is always less than 1 (since it can never increase), so A_0 cannot receive a cookie (if A_0 received a cookie it would have weight 1). Thus we must have $m \geq 2^n$.

We now show that for $m \geq 2^n$, we can always ensure that A_0 gets a cookie. Intuitively, we have the following heuristic:

Our algorithm should never pass away from A_0, otherwise we will decrease our monovariant. Thus in each step we should pass towards A_0.

This heuristic, however, does not tell us which way A_n should pass a cookie, as both directions are towards A_0 (A_n and A_0 are diametrically opposite). This leads us to consider a new quantity in order to distinguish between the two directions that A_n can pass to. Let W_+ be the sum of the weights of cookies held by A_0, A_1, A_2, …., A_n and let W_- be the sum of the weights of cookies held by A_0, A_{-1}, A_{-2}, …., A_{-n}. Assume WLOG $W_+ \geq W_-$. Then this suggests that we should make A_n pass cookies only to A_{n-1} and that we should only work in the semicircle containing nonnegative indices, since this is the semicircle having more weight. Thus our algorithm is to make A_n pass as many cookies as possible to A_{n-1}, then make A_{n-1} pass as many cookies as possible to A_{n-2}, and so on until A_0 gets a cookie. But this works if and only if $W_+ \geq 1$: $W_+ \geq 1$ is certainly necessary since W_+ is a monovariant under our algorithm, and we now show it is sufficient.

Suppose $W_+ \geq 1$. Note that our algorithm leaves W_+ invariant. Suppose our algorithm terminates, that is, we cannot pass anymore cookies from any of the A_i's with i positive, and A_0 doesn't have any cookies. Then A_1, A_2, …., A_n all have at most 1 cookie at the end (if they had more than one, they could eat one and pass one and our algorithm wouldn't have terminated). Then at this point $W_+ \leq \frac{1}{2} + \frac{1}{4} + ... + 1/2^n < 1$, contradicting the fact that W_+ is invariant and ≥ 1. Thus $W_+ \geq 1$ is a sufficient condition for our algorithm to work.

Finally, we prove that we indeed have $W_+ \geq 1$. We assumed $W_+ \geq W_-$. Now simply note that each cookie contributes at least $1/2^{n-1}$ to the sum $(W_+ + W_-)$, because each cookie has weight at least $1/2^{n-1}$ except for cookies at A_n. However, cookies at A_n are counted twice since they contribute to both W_+ and W_-, so they also contribute $1/2^{n-1}$ to the sum. Hence, since we have at least 2^n cookies, $W_+ + W_- \geq 2$, so $W_+ \geq 1$ and we are done. ■

The next example demonstrates three very useful ideas: monovariants, binary representation and the Euclidean algorithm. All of these are very helpful tools.

Example 7 [IMO shortlist 1994, C3]
Peter has 3 accounts in a bank, each with an integral number of dollars. He is only allowed to transfer money from one account to another so that the amount of money in the latter is doubled. Prove that Peter can always transfer all his money into two accounts. Can he always transfer all his money into one account?

Answer:
The second part of the question is trivial - if the total number of dollars is odd, it is clearly not always possible to get all the money into one account. Now we solve the first part. Let A, B, C with $A \leq B \leq C$ be the number of dollars in the account 1, account 2 and account 3 respectively at a particular point of time. If $A = 0$ initially, we are done so assume $A > 0$. As we perform any algorithm, the values of A, B and C keep changing. Our aim is to monotonically strictly decrease the value of *min* (A, B, C). This will ensure that we eventually end up with *min* $(A, B, C) = 0$ and we will be done. Now, we know a very simple and useful algorithm that monotonically reduces a number- the Euclidean algorithm. So let $B = qA + r$ with $0 \leq r < A$. Our aim now is to reduce the number

of dollars in the second account from B to r. Since $r < A$, we would have reduced $min\ (A, B, C)$, which was our aim.

Now, since the question involves doubling certain numbers, it is a good idea to consider binary representations of numbers. Let $q = m_0 + 2m_1 + ... + 2^k m_k$ be the binary representation of q, where $m_i = 0$ or 1. To reduce B to r, in step i of our algorithm, we transfer money to account 1. The transfer is from account 2 if $m_{i-1} = 1$ and from account 3 if $m_{i-1} = 0$. The number of dollars in the first account starts with A and keeps doubling in each step. Thus we end up transferring $A(m_0 + 2m_1 + ... + 2^k m_k) = Aq$ dollars from account 2 to account 1, and we are left with $B - Aq = r$ dollars in account 2. We have thus succeeded in reducing $min\ (A, B, C)$ and so we are done. ■

Now we look at a very challenging problem that can be solved using monovariants.

Example 8 [APMO 1997, Problem 5]

n people are seated in a circle. A total of nk coins have been distributed among them, but not necessarily equally. A *move* is the transfer of a single coin between two adjacent people. Find an algorithm for making the minimum possible number of moves which result in everyone ending up with the same number of coins.

Answer:

We want each person to end up with k coins. Let the people be labeled from 1, 2, …, n in order (note that n is next to 1 since they are sitting in a circle). Suppose person i has c_i coins. We introduce the variable $d_i = c_i - k$, since this indicates how close a person is to having the desired number of coins. Consider the quantity

$$X = |d_1| + |d_1 + d_2| + |d_1 + d_2 + d_3| + \dots + |d_1 + d_2 + \dots + d_{n-1}|$$

Clearly $X = 0$ if and only if everyone has k coins, so our goal is to make $X = 0$. The reason for this choice of X is that moving a coin between person j and person $j + 1$ for $1 \leq j \leq n - 1$ changes X by exactly 1 as only the term $|d_1 + d_2 + \ldots + d_j|$ will be affected. Hence X is a monovariant and is fairly easy to control (except when moving a coin from 1 to n or vice versa). Let $s_j = d_1 + d_2 + \ldots + d_j$.

We claim that as long as $X > 0$ it is always possible to reduce X by 1 by a move between j and $j + 1$ for some $1 \leq j \leq n - 1$. We use the following algorithm. Assume WLOG $d_1 \geq 1$. Take the first j such that $d_{j+1} < 0$. If $s_j > 0$, then simply make a transfer from j to $j + 1$. This reduces X by one since it reduces the term $|s_j|$ by one. The other possibility is $s_j = 0$, which means $d_1 = d_2 = \ldots = d_j = 0$ (recall that d_{j+1} is the first negative term). In this case, take the first $m > i+1$ such that $d_m \geq 0$. Then $d_{m-1} < 0$ by the assumption on m, so we move a coin from m to (m-1). Note that all terms before d_m were either 0 or less than 0 and $d_{m-1} < 0$, so s_{m-1} was less than 0. Our move has increased s_{m-1} by one, and has hence decreased $|s_{m-1}|$ by one, so we have decreased X by one.

Thus at any stage we can always decrease X by at least one by moving between j and $j + 1$ for some $1 \leq j \leq n - 1$. We have not yet considered the effect of a move between 1 and n. Thus our full algorithm is as follows: At any point of time, if we can decrease X by moving a coin from 1 to n or n to 1, do this. Otherwise, decrease X by 1 by the algorithm described in the above paragraph. ■

Sometimes while creating algorithms that monotonically decrease (or increase) a quantity, we run into trouble in particular cases and our algorithm doesn't work. We can often get around these difficulties as follows. Suppose we want to monotonically decrease a particular quantity. Call a position *good* if we can decrease the monovariant with our algorithm. Otherwise, call the position *bad*. Now create an algorithm that converts bad positions into good positions, without increasing our monovariant. We use the first algorithm when possible, and then if we are stuck in a bad position, use the second algorithm to get back to a good position. Then we can again use the first algorithm. The next example (which is quite hard) demonstrates this idea.

Example 9 [USAMO 2003-6]
At the vertices of a regular hexagon are written 6 nonnegative integers whose sum is 2003. Bert is allowed to make moves of the following form: he may pick a vertex and replace the number written there by the absolute value of the difference between the numbers at the neighboring vertices. Prove that Bert can make a sequence of moves, after which the number 0 appears at all 6 vertices.

Remark: We advise the reader to follow this solution with a paper and pen, and fill in the details that have been left for the reader. We first suggest that the reader try some small cases (with 2003 replaced by smaller numbers).

Answer:
Our algorithm uses the fact that 2003 is odd. Let the sum of a position be the sum of the 6 numbers and the maximum denote the value of the maximum of the 6 numbers. Let A, B, C, D, E, F be the numbers at the 6 vertices in that order. Our aim is to monotonically decrease the maximum. Note that the maximum can never increase.

We need two sub-algorithms:

(i) "Good position" creation: from a position with odd sum, go to a position with exactly one odd number
(ii) Monovariant reduction: from a position with exactly one odd number, go to a position with odd sum and strictly smaller maximum, or go to the all 0 position.

For (i), since (A + B + C + D + E + F) is odd, assume WLOG that A + C + E is odd. If exactly one of A, C, E is odd, suppose A is odd. Then make the following sequence of moves: B, F, D A, F (here we denote a move by the vertex at which the move is made). This way, we end up with a situation in which only B is odd and the rest become even (check this), and we are done with step (i). The other possibility is that all of A, C and E are odd. In this case make the sequence of moves (B, D, F, C, E). After this only A is odd (check this).

Now we are ready to apply step (ii), the step that actually decreases our monovariant. At this point, only one vertex contains an odd number; call this vertex A. Again we take two cases. If the maximum is even, then it is one of B, C, D, E or F. Now make moves at B, C, D, E and F in this order. (The reader should check that this works, that is, this sequence of moves decreases the maximum and ensures that the sum is odd). If the maximum is odd, then it is A. If C = E = 0, then the sequence of moves (B, F, D, A, B, F) leaves us with all numbers 0 and we are done. Otherwise, suppose at least one of C and E is nonzero so suppose C > 0 (the case E > 0 is similar). In this case, make the moves (B, F, A, F). The reader can check that this decreases the maximum and leaves us with odd sum.

Thus starting with odd sum, we apply (i) if needed, after which we apply (ii). This decreases the maximum, and also leaves us again with odd sum (or in some cases it leaves us with all 0s and we are done), so we can repeat the entire procedure until the maximum eventually becomes 0. ■

Miscellaneous Examples

Now we look at a few more problems involving moves that don't directly use monovariants or greedy algorithms. These problems can often be solved by algorithms that build up the required configuration in steps. Sometimes, the required algorithm becomes easier to find after making some crucial observations or proving an auxiliary lemma. But in lots of cases, all a combinatorics problem needs is patience and straightforward logic, as the next example shows. Here again the solution looks long but most of what is written is just intended to motivate the solution.

Example 10 [China 2010, Problem 5]
There are some (finite number of) cards placed at the points A_1, A_2, …, A_n and O, where $n \geq 3$. We can perform one of the following operations in each step:

(1) If there are more than 2 cards at some point A_i, we can remove 3 cards from this point and place one each at A_{i-1}, A_{i+1} and O (here $A_0 = A_n$ and $A_{n+1} = A_1$)
(2) If there are at least n cards at O, we can remove n cards from O and place one each at $A_1, A_2, \ldots, A_n$.

Show that if the total number of cards is at least n^2+3n+1, we can make the number of cards at each vertex at least $n + 1$ after finitely many steps.

Answer:
Note that the total number of cards stays the same. We make a few observations:

(a) We should aim to make the number of cards at each A_i equal or close to equal, since if in the end some point has lots of cards, some other point won't have enough.

(b) We can make each of the A_i's have 0, 1 or 2 cards.
Proof: repeatedly apply operation (1) as long as there is a point with at least 3 cards. This process must terminate, since the number of coins in O increases in each step but cannot increase indefinitely. This is a good idea since the A_i's would now have a 'close to equal' number of coins, which is a good thing by observation a).

(c) From observation b), we see that it is also possible to make each of the A_i's have 1, 2, or 3 cards (from the stage where each vertex has 0, 1 or 2 cards, just apply operation (2) once). This still preserves the 'close to equal' property, but gives us some more flexibility since we are now able to apply operation 1.

(d) Based on observation c), we make each of the A_i's have 1, 2 or 3 cards. Suppose x of the A_i's have 1 card, y of the A_i's have 2 cards and z of the A_i's have 3 cards. The number of cards at O is then at least $(n^2+3n+1) - (x+2y+3z)$. Since $x + y + z = n$, $(x + 2y + 3z) = (2x + 2y + 2z) + z - x = 2n + z - x \leq 2n$ if $x \geq z$. Thus if $x \geq z$, O will have at least $(n^2+3n+1) - 2n = n^2+n+1$ cards. Now we can apply operation (2) n times. Then all the A_i's will now have at least $n + 1$ cards (they already each had at least 1 card), and O will have at least $n^2 + n + 1 - n^2 = n + 1$ cards and we will be done.

Thus, based on observation d), it suffices to find an algorithm that starts with a position in which each of the A_i's have 1, 2, or 3 cards and ends in a position in which each of the A_i's have 1, 2, or 3 cards but the number of points having 3 cards is not more than the number of points having 1 card. This is not very difficult- the basic idea is to ensure that between any two points having 3 cards, there is a point containing 1 card. We can do this as follows:

If there are consecutive 3's in a chain, like $(x, 3, 3, \ldots, 3, y)$ with $(x, y \neq 3)$, apply operation (1) on all the points with 3 cards to get $(x + 1, 1, 2, 2, \ldots, 2, 1, y+1)$. Thus we can ensure that there are no adjacent 3's. Now suppose there are two 3's with only 2's between them, like $(x, 3, 2, 2, 2, \ldots, 2, 3, y)$ with $x, y \neq 3$. After doing operation (1) first on the first 3, then on the point adjacent to it that has become a 3 and so on until the point before y, we get the sequence $(x+1, 1, 1, \ldots, 1, y+1)$.

Thus we can repeat this procedure as long as there exist two 3's that do not have a 1 between them. Note that the procedure preserves the property that all A_i's have 1, 2 or 3 cards. But this cannot go on indefinitely since the number of coins at 0 is increasing. So eventually we end up with a situation where there is at least one 1 between any two 3's, and we are done. ■

Example 11 [IMO 2010, Problem 5]
Six boxes B_1, B_2, B_3, B_4, B_5, B_6 of coins are placed in a row. Each box initially contains exactly one coin. There are two types of allowed moves:

Move 1: If B_k with $1 \le k \le 5$ contains at least one coin, you may remove one coin from B_k and add two coins to B_{k+1}.
Move 2: If B_k with $1 \le k \le 4$ contains at least one coin, you may remove one coin from B_k and exchange the contents (possibly empty) of boxes B_{k+1} and B_{k+2}.

Determine if there exists a finite sequence of moves of the allowed types, such that the five boxes B_1, B_2, B_3, B_4, B_5 become empty, while box B_6 contains exactly $2010^{2010^{2010}}$ coins.
Note: $a^{b^c} = a^{(b^c)}$

Answer:

Surprisingly, the answer is yes. Let $A = 2010^{2010^{2010}}$. We denote by $(a_1, a_2, \ldots, a_n) \to (a_1', a_2', \ldots, a_n')$ the following: if some consecutive boxes have $a_1, a_2, \ldots, a_n$ coins respectively, we can make them have $a_1', a_2', \ldots, a_n'$ coins by a legal sequence of moves, with all other boxes unchanged.

Observations:

a) Suppose we reach a stage where all boxes are empty, except for B_4, which contains at least $A/4$ coins. Then we can apply move 2 if necessary until B_4 contains exactly $A/4$ coins, and then apply move 1 twice and we will be done. Thus reaching this stage will be our key goal.

b) Move 1 is our only way of increasing the number of coins. Since it involves doubling, we should look for ways of generating powers of 2. In fact, since A is so large, we should try to generate *towers* of 2's (numbers of the form 2^{2^2}).

Based on this, we construct two sub algorithms.

Algorithm 1: $(a, 0, 0) \to (0, 2^a, 0)$ for any positive integer a.
Proof: First use move 1: $(a, 0, 0) \to (a-1, 2, 0)$.
Now use move 1 on the middle box till it is empty: $(a-1, 2, 0) \to (a-1, 0, 4)$
Use move 2 on the first box to get $(a-2, 4, 0)$.
Repeating this procedure (that is, alternately use move one on the second box till it is empty, followed by move one on the first box and so on), we eventually get $(0, 2^a, 0)$.
Now, using this algorithm, we can construct an even more powerful algorithm that generates a large number of coins.

Algorithm 2: Let P_n be a tower of n 2's for each positive integer n (eg. $P_3 = 2^{2^2} = 16$). Then
$(a, 0, 0, 0) \to (0, P_a, 0, 0)$.
Proof: We use algorithm 1. As in algorithm 1, the construction is stepwise. It is convenient to explain it using induction.

We prove that $(a, 0, 0, 0) \to (a-k, P_k, 0, 0)$ for each $1 \le k \le a$. For $k = 1$, simply apply move 1 to the first box. Suppose we have reached the stage $(a-k, P_k, 0, 0)$. We want to reach $(a-(k+1), P_{k+1}, 0, 0)$. To do this, **apply algorithm 1** to get $(a-k, 0, 2^{P_k}, 0)$. Note that $2^{P_k} = P_{k+1}$. So now just apply move 2 to the first box and we get $(a-k-1, P_{k+1}, 0, 0)$. Thus by induction, we finally reach (for $k = a$) $(0, P_a, 0, 0)$.

With algorithm 2 and observation a), we are ready to solve the problem.

First apply move 1 to box 5, then move 2 to box 4, 3, 2 and 1 in this order:
(1, 1, 1, 1, 1, 1) → (1, 1, 1, 1, 0, 3) → (1, 1, 1, 0, 3, 0) → (1, 1, 0, 3, 0, 0) →(1, 0, 3, 0, 0, 0) → (0, 3, 0, 0, 0, 0).

Now we use algorithm 2 twice:
(0, 3, 0, 0, 0, 0) → (0, 0, P_3, 0, 0, 0) → (0, 0, 0, P_{16}, 0, 0).

Now we leave it to the reader to check that $P_{16} > A/4$ (in fact P_{16} is much larger than A). By observation a), we are done.

Remark: In the contest, several contestants thought the answer was no, and spent most of their time trying to prove that no such sequence exists. Make sure that you don't ever jump to conclusions like that too quickly. On a lighter note, in a conference of the team leaders and deputy leaders after the contest, one deputy leader remarked "Even most of us thought that no such sequence existed". To this, one leader replied, "That's why you are deputy leaders and not team leaders!"

We close this chapter with one of the hardest questions ever asked at the IMO. Only 2 out of over 500 contestants completely solved problem 3 in IMO 2007. Yup, that's right- 2 high school students in the entire world.

Example 12 [IMO 2007, Problem 3]
In a mathematical competition some competitors are friends; friendship is always mutual. Call a group of competitors a clique if each two of them are friends. The number of members in a clique is called its size. It is known that the size of the largest clique(s) is even. Prove that the competitors can be arranged in two rooms such that the size of the largest cliques in one room is the same as the size of the largest cliques in the other room.

Answer:
Let $\boldsymbol{M}$ be one of the cliques of largest size, $|\boldsymbol{M}| = 2m$. First send all members of $\boldsymbol{M}$ to Room A and all other people to Room B. Let c(A)

and c(B) denote the sizes of the largest cliques in rooms A and B at a given point in time. Since $\boldsymbol{M}$ is a clique of the largest size, we initially have $c(A) = |\boldsymbol{M}| \geq c(B)$. Now we want to “balance things out”. As long as $c(A) > c(B)$, send one person from Room A to Room B. In each step, c(A) decreases by one and c(B) increases by at most one. So at the end we have $c(A) \leq c(B) \leq c(A) + 1$. We also have $c(A) = |A| \geq m$ at the end. Otherwise we would have at least m+1 members of $\boldsymbol{M}$ in Room B and at most $m-1$ in Room A, implying $c(B) - c(A) \geq (m+1) - (m-1) = 2$.

Clearly if $c(A) = c(B)$ we are done so at this stage the only case we need to consider is $c(B) - c(A) = 1$. Let $c(A) = k$, $c(B) = k+1$. Now if there is a competitor in B, who is also in $\boldsymbol{M}$ but is not in the biggest clique in B, then by sending her to A, c(B) doesn't change but c(A) increases by 1 and we are done. Now suppose there is no such competitor. We do the following: take each clique of size k+1 in B and send one competitor to A. At the end of this process, $c(B) = k$. Now we leave it to the reader to finish the proof by showing that c(A) is still k. (You will need to use the supposition that there is no competitor in B who is also in $\boldsymbol{M}$ but not in the biggest clique of B. This means that every clique in B of size (k+1) contains B∩M). ■

Exercises

1. **[Activity Selection Problem]**

   On a particular day, there are n events (say, movies, classes, parties, etc.) you want to attend. Call the events $E_1, E_2, \ldots, E_n$ and let E_i start at time s_i and finish at time f_i. You are only allowed to attend events that do not overlap (that is, one should finish before the other starts). Provide an efficient algorithm that selects as many events as possible while

satisfying this condition.
(Note: We have not defined what "efficient" here means. Note that this problem can be solved by simply testing all $2n$ possible combinations of events, and taking the best combination that works. However, this uses a number of steps that is exponential in n. By efficient, we mean a procedure that is guaranteed to require at most a number of steps that is polynomial in n).

2. **[Weighted Activity Selection]**
Solve the following generalization of the previous problem: event E_i has now weight w_i and the objective is not to maximize the number of activities attended, but the sum of the weights of all activities attended.

3. **[Russia 1961]**
Real numbers are written in an $m \times n$ table. It is permissible to reverse the signs of all the numbers in any row or column. Prove that after a number of these operations, we can make the sum of the numbers along each line (row or column) nonnegative.

4. Given $2n$ points in the plane with no three collinear, show that it is possible to pair them up in such a way that the n line segments joining paired points do not intersect.

5. **[Czech and Slovak Republics 1997]**
Each side and diagonal of a regular n-gon ($n \geq 3$) is colored blue or green. A move consists of choosing a vertex and switching the color of each segment incident to that vertex (from blue to green or vice versa). Prove that regardless of the initial coloring, it is possible to make the number of blue segments incident to each vertex even by following a sequence of moves. Also show that the final configuration obtained is uniquely determined by the initial coloring.

6. **[Bulgaria 2001]**

Given a permutation of the numbers 1, 2, …, n, one may interchange two consecutive blocks to obtain a new permutation. For instance, 3 5 4 8 9 7 2 1 6 can be transformed to 3 9 7 2 5 4 8 1 6 by swapping the consecutive blocks 5 4 8 and 9 7 2. Find the least number of changes required to change n, n-1, n-2, …, 1 to 1, 2, …, n.

7. **[Minimum makespan scheduling]**

Given the times taken to complete n jobs, t_1, t_2, …, t_n, and m identical machines, the task is to assign each job to a machine so that the total time taken to finish all jobs is minimized. For example, if $n = 5$, $m = 3$ and the times are 5, 4, 4, 6 and 7 hours, the best we can do is make machine 1 do jobs taking 4 and 5 hours, machine 2 do jobs taking 4 and 6 hours, and machine 3 do the job taking 7 hours. The total time will then be 10 hours since machine 2 takes (4 + 6) hours.

Consider the following greedy algorithm: Order the jobs arbitrarily, and in this order assign to each job the machine that has been given the least work so far. Let T_{OPT} be the total time taken by the best possible schedule, and T_A the time taken by our algorithm. Show that $T_A/T_{OPT} \leq 2$; in other words, our algorithm always finds a schedule that takes at most twice the time taken by an optimal schedule. (This is known as a *2-factor approximation algorithm*.)

8. **[USAMO 2011-2]**

An integer is written at each vertex of a regular pentagon. A solitaire game is played as follows: a turn consists of choosing an integer m and two adjacent vertices of the pentagon, and subtracting m from the numbers at these vertices and adding 2m to the vertex opposite them. (Note that m and the vertices chosen can change from turn to turn). The game is said to be won at a vertex when the number 2011 is written at it and the other four vertices have the number 0 written at them. Show that there is exactly one vertex at which the game can be won.

9. [Chvatal's set covering algorithm]

Let S_1, S_2, …, S_k be subsets of {1, 2, …, n}. With each set S_i is an associated cost c_i. Given this information, the minimum set cover problem asks us to select certain sets among S_1, …, S_k such that the union of the selected sets is {1, 2, …, n} (that is, each element is covered by some chosen set) and the total cost of the selected sets is minimized. For example, if $n = 4$, $k = 3$, $S_1 = \{1, 2\}$; $S_2 = \{2, 3, 4\}$ and $S_3 = \{1, 3, 4\}$ and the costs of S_1, S_2 and S_3 are 5, 6 and 4 respectively, the best solution would be to select S_1 and S_3.

Consider the following greedy algorithm for set cover: In each stage of the algorithm, we select the subset S_i which maximizes the value of $\frac{|S_i \cap C'|}{c_i}$, where C' denotes the set of elements not yet covered at that point. Intuitively, this algorithm maximizes (additional benefit)/cost in each step. This algorithm does not produce an optimal result, but it gets fairly close: let C_A be the cost of the selected sets produced by the algorithm, and let C_{OPT} be the cost of the best possible selection of sets (the lowest cost). Prove that $C_A/C_{OPT} \le H_n$, where $H_n = 1 + ½ + … + 1/n$. (In other words, this is an H_n-factor approximation algorithm.)

10. A ***matroid*** is an ordered pair (S, F) satisfying the following conditions:

(i) S is a finite set

(ii) F is a nonempty family of subsets of S, such that if A is a set in F, all subsets of A are also in F. The members of F are called ***independent*** sets

(iii) If A and B belong to F but $|A| > |B|$, then there exists an element $x \in B \setminus A$ such that $A \cup \{x\} \in F$.

For example, if S = {1, 2, 3, 4} and F = {∅, {1}, {2}, {3}, {4}, {1, 2}, {1, 3}, {1, 4}, {2, 3}, {2, 4}, {3,4}}, then you can easily verify that the above properties are satisfied. In general, note that if F contains all subsets of S with k or fewer elements for some $k \leq |S|$, {S, F} will be a matroid.

An independent set A is said to be ***maximal*** if there does not exist any element x in S such that $A \cup \{x\} \in F$. (In other words, adding any element to A destroys its independence.) Prove that ***all maximal independent sets have the same cardinality***.

11. Consider a matroid {S, F} where $S = \{a_1, \ldots, a_n\}$. Let element a_i have weight w_i, and define the weight of a set A to be the sum of the weights of its elements. A problem central to the theory of greedy algorithms is to find an independent set in this matroid of ***maximum weight***. Consider the following greedy approach: starting from the null set, in each stage of the algorithm add an element (that has not been selected so far) with the highest weight possible while preserving the independence of the set of selected elements. When no more elements can be added, stop.
Show that this greedy algorithm indeed produces a ***maximum weight independent set***.

12. [IMO Shortlist 2013, C3]
A crazy physicist discovered a new kind of particle which he called an imon. Some pairs of imons in the lab can be entangled, and each imon can participate in many entanglement relations. The physicist has found a way to perform the following two kinds of operations with these particles, one operation at a time.
(i) If some imon is entangled with an odd number of other imons in the lab, then the physicist can destroy it.
(ii) At any moment, he may double the whole family of imons in the lab by creating a copy I’ of each imon I. During this procedure, the two copies I’ and J’ become entangled if and

only if the original imons I and J are entangled, and each copy I' becomes entangled with its original imon I; no other entanglements occur or disappear at this moment.

Show that after a finite number of operations, he can ensure that no pair of particles is entangled.

13. **[Japan 1998]**

Let n be a positive integer. At each of $2n$ points around a circle we place a disk with one white side and one black side. We may perform the following move: select a black disk, and flip over its two neighbors. Find all initial configurations from which some sequence of such moves leads to a position where all disks but one are white.

14. **[Based on IOI 2007]**

You are given n integers $a_1, a_2, \ldots, a_n$ and another set of n integers $b_1, b_2, \ldots, b_n$ such that for each i, $b_i \leq a_i$. For each $i = 1, 2, \ldots, n$, you must choose a set of b_i distinct integers from the set $\{1, 2, \ldots, a_i\}$. In total, $(b_1 + b_2 + \ldots + b_n)$ integers are selected, but not all of these are distinct. Suppose k distinct integers have been selected, with multiplicities $c_1, c_2, c_3, \ldots, c_k$. Your ***score*** is defined as $\sum_{i=1}^{k} c_k(c_k - 1)$. Give an efficient algorithm to select numbers in order to ***minimize*** your score.

15. **[Based on Asia Pacific Informatics Olympiad 2007]**

Given a set of n distinct positive real numbers $S = \{a_1, a_2, \ldots, a_n\}$ and an integer $k < n/2$, provide an efficient algorithm to form k pairs of numbers $(b_1, c_1), (b_2, c_2), \ldots, (b_k, c_k)$ such that these $2k$ numbers are all distinct and from S, and such that the sum $\sum_{i=1}^{k} |b_i - c_i|$ is minimized.

Hint: A natural greedy algorithm is to form pairs sequentially by choosing the closest possible pair in each step. However, this doesn't always work. Analyze where precisely the problem in this approach lies, and then accordingly adapt this algorithm so that it works.

16. **[ELMO Shortlist 2010]**

You are given a deck of kn cards each with a number in {1, 2, …, n} such that there are k cards with each number. First, n piles numbered {1, 2, …, n} of k cards each are dealt out face down. You are allowed to look at the piles and rearrange the k cards in each pile. You now flip over a card from pile 1, place that card face up at the bottom of the pile, then next flip over a card from the pile whose number matches the number on the card just flipped. You repeat this until you reach a pile in which every card has already been flipped and wins if at that point every card has been flipped. Under what initial conditions (distributions of cards into piles) can you guarantee winning this game?

17. **[Russia 2005]**

100 people from 25 countries, four from each country, sit in a circle. Prove that one may partition them onto 4 groups in such way that no two countrymen, nor two neighboring people in the circle, are in the same group.

18. **[Saint Petersburg 1997]**

An *Aztec diamond of rank n* is a figure consisting of those squares of a gridded coordinate plane lying inside the square $|x| + |y| \leq n+1$. For any covering of an Aztec diamond by dominoes, a move consists of selecting a 2x2 square covered by two dominoes and rotating it by 90 degrees. The aim is to convert the initial covering into the covering consisting of only horizontal dominoes. Show that this can be done using at most $n(n+1)(2n+1)/6$ moves.

Olympiad Combinatorics

Pranav A. Sriram

August 2014

Copyright notices

About the Author

Pranav Sriram graduated from high school at The International School Bangalore, India, and will be a Freshman at Stanford University this Fall.

2. ALGORITHMS – PART II

In this chapter we focus on some very important themes in the study of algorithms: *recursive algorithms, efficiency* and *information*. A recursive algorithm is one which performs a task involving n objects by breaking it into smaller parts. This is known as a “divide and conquer” strategy. Typically, we either do this by splitting the task with n objects into two tasks with $n/2$ objects or by first reducing the task to a task with (n-1) objects. The latter approach, which is essentially induction, is very often used to solve Olympiad problems.

Induction

We first look at two problems which use induction. In the first one, we use the technique of ignoring one object and applying the induction hypothesis on the remaining (n-1) objects. This obviously needs some care: we cannot completely ignore the nth object if it has some effect on the other objects!

Example 1 [China Girls Math Olympiad 2011-7]
There are n boxes B_1, B_2, …, B_n in a row. N balls are distributed amongst them (not necessarily equally). If there is at least one ball

in B_1, we can move one ball from B_1 to B_2. If there is at least 1 ball in B_n, we can move one ball from B_n to B_{n-1}. For $2 \leq k \leq (n-1)$, if there are at least two balls in B_k, we can remove two balls from B_k and place one in B_{k+1} and one in B_{k-1}. Show that whatever the initial distribution of balls, we can make each box have exactly one ball.

Answer:
We use induction and monovariants. The base cases $n = 1$ and 2 are trivial. Suppose we have an algorithm A_{n-1} for n-1 boxes; we construct an algorithm A_n for n boxes. We use two steps. The first step aims to get a ball into B_n and the second uses the induction hypothesis.

Step 1: If B_n contains at least one ball, move to step two. Otherwise, all n balls lie in the first (n-1) boxes. Assign a weight 2^k to box B_k. Now keep moving balls from the boxes B_1, B_2, ..., B_{n-1} as long as possible. This cannot go on indefinitely as the total weight of the balls is a positive integer and strictly increases in each move but is bounded above by $n2^n$. Thus at some point this operation terminates. This can only happen if B_1 has 0 balls and B_2, B_3, ..., B_{n-1} each has at most 1 ball. But then B_n will have at least 2 balls. Now go to step 2.

Step 2: If B_n has $k > 1$ balls, move (k-1) balls from B_n to B_{n-1}. Now B_n has exactly one ball and the remaining (n-1) boxes have (n-1) balls. Color these (n-1) balls red and color the ball in B_n blue. Now we apply the induction hypothesis. Use algorithm A_{n-1} to make each of the first (n-1) boxes have one ball each. The only time we run into trouble is when a move needs to be made from B_{n-1}, because in A_{n-1}, B_{n-1} only needed 1 ball to make a move, but now it needs 2. We can easily fix this. Whenever A_{n-1} says we need to move a ball from B_{n-1} to B_{n-2}, we first move the blue ball to B_{n-1}. Then we move a ball from B_{n-1} to B_{n-2} and pass the blue ball back to B_n. This completes the proof. ■

Example 2 [IMO Shortlist 2005, C1]
A house has an even number of lamps distributed among its rooms in such a way that there are at least three lamps in every room. Each lamp shares a switch with exactly one other lamp, not necessarily from the same room. Each change in the switch shared by two lamps changes their states simultaneously. Prove that for every initial state of the lamps there exists a sequence of changes in some of the switches at the end of which each room contains lamps which are on as well as lamps which are off.

Answer:
Call a room *bad* if all its lamps are in the same state and *good* otherwise. We want to make all rooms good. We show that if $k \geq 1$ rooms are bad, then we can make a finite sequence of switches so that (k-1) rooms are bad. This will prove our result.

Call two lamps connected if they share a switch. Take a bad room R_1 and switch a lamp there. If this lamp is connected to a lamp in R_1, we are done since each room has at least 3 lamps. If this lamp is connected to a lamp in another room R_2, then R_1 becomes good but R_2 might become bad. If R_2 doesn't become bad, we are done. If R_2 does become bad, then repeat the procedure so that R_2 becomes good but some other room R_3 becomes bad. Continue in this manner. If we ever succeed in making a room good without making any other room bad we are done, so assume this is not the case. Then eventually we will reach a room we have already visited before. We prove that at this stage, the final switch we made would not have made any room bad.

Consider the first time this happens and let $R_m = R_n$ for some $m > n$. We claim that R_m is good at this stage. The first time we switched a lamp in R_n, we converted it from bad to good by switching one lamp. Now when we go to R_m (= R_n), we cannot switch the same lamp, since this lamp was connected to a lamp in room R_{n-1}, whereas the lamp we are about to switch is connected to a lamp in R_{m-1}. So two distinct lamps have been switched in R_m and hence R_m is good (since there are at least three lamps, at least one lamp hasn't been switched, and initially all lamps were in the same state since the room was bad before). Thus our final switch has made R_{m-1} good without making R_m bad. Hence we have reduced the number of bad rooms by one, and repeating this we eventually make all rooms good. ■

The next two examples demonstrate how to construct objects inductively.

Example 3:
Given a graph ***G*** in which each vertex has degree at least (n-1), and a tree ***T*** with n vertices, show that there is a subgraph of ***G*** isomorphic to ***T***.

Answer:
We find such a subgraph inductively. Assume the result holds for (n-1); we prove it holds for n. Delete a terminal vertex v from ***T***. By induction we can find a tree ***H*** isomorphic to ***T*** \ {v} as a subgraph of ***G***. This is because ***T*** \ {v} has (n-1) vertices and each vertex in ***G*** has degree at least (n-1) > (n-1) - 1, so we can apply the induction hypothesis. Now suppose v was adjacent to vertex u in ***T*** (remember that v is adjacent to only one vertex). Let w be the vertex in ***G*** corresponding to u. w has at least (n-1) neighbors in ***G***, and at most (n-2) of them are in ***H*** since ***H*** has (n-1) vertices and w is one of them. Thus w has at least 1 neighbor in ***G*** that is not in ***H***, and we take this vertex as the vertex corresponding to v. ■

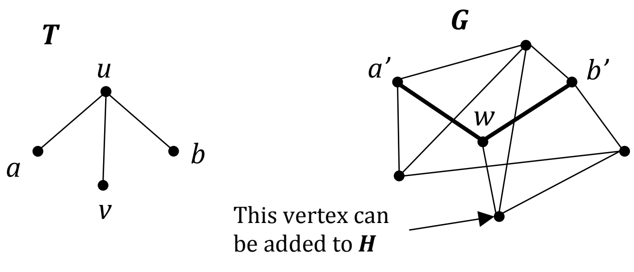

*Figure 2.1: Finding **H** inductively*

Example 4 [USAMO 2002]
Let ***S*** be a set with 2002 elements, and let N be an integer with $0 \leq N \leq 2^{2002}$. Prove that it is possible to color every subset of ***S*** black or white, such that:

a) The union of two white subsets is white
b) The union of two black subsets is black
c) There are exactly N white subsets.

Answer
You may have thought of inducting on N, but instead we induct on the number of elements of ***S***. In this problem $|\boldsymbol{S}| = 2002$, but we prove the more general result with $|\boldsymbol{S}| = n$ and $0 \leq N \leq 2^n$. The result trivially holds for $n = 1$, so suppose the result holds for $n = k$. Now we prove the result for $n = k+1$. If $N \leq 2^{n-1}$, note that by induction there is a coloring for the same value of N and $n = k$. We use this coloring for all sets that do not contain the $(k+1)$th element of ***S***, and all subsets containing the $(k+1)$th element of ***S*** (which were not there in the case $|\boldsymbol{S}| = k$) are now colored black. (Essentially, all "new" subsets are colored black while the old ones maintain their original color). Clearly, this coloring works.

If $N \geq 2^{n-1}$, simply interchange the roles of white and black, and then use the same argument as in the previous case. ■

Information, Efficiency and Recursions

The next few problems primarily deal with collecting information and performing tasks efficiently, that is, with the minimum possible number of moves. Determining certain information with the least number of moves or questions is extremely important in computer science.

The next example is a simple and well-known problem in computer science.

Example 5 [Merge Sort Algorithm]
Given n real numbers, we want to sort them (arrange them in non-decreasing order) using as few comparisons as possible (in one comparison we can take two numbers a and b and check whether $a < b$, $b < a$ or $a = b$). Clearly, we can sort them if we make all possible $n(n-1)/2$ comparisons. Can we do better?

Answer:
Yes. We use a recursive algorithm. Let $f(n)$ be the number of comparisons needed for a set of n numbers. Split the set of n numbers into 2 sets of size $n/2$ (or if n is odd, sizes $(n-1)/2$ and $(n+1)/2$. For the rest of this problem, suppose n is even for simplicity). Now sort these two sets of numbers individually. This requires $2f(n/2)$ comparisons. Suppose the resulting sorted lists are $a_1 \leq a_2 \leq \ldots \leq a_{n/2}$ and $b_1 \leq b_2 \leq \ldots \leq b_{n/2}$. Now we want to combine or 'merge' these two lists. First compare a_1 and b_1. Thus after a comparison between a_i and b_j, if $a_i \leq b_j$, compare a_{i+1} and b_j and if $b_j < a_i$, compare b_{j+1} and a_i in the next round. This process terminates after at most n comparisons, after which we would have completely sorted the list. We used a total of at most $2f(n/2) + n$ comparisons, so $f(n) \leq 2f(n/2) + n$.

From this recursion, we can show by induction that $f(2^k) \leq k \times 2^k$ and in general, for n numbers the required number of comparisons is of the order $n\log_2(n)$, which is much more efficient than the trivial bound $n(n-1)/2$ which is of order n^2. ■

Example 6
Suppose we are given n lamps and n switches, but we don't know which lamp corresponds to which switch. In one operation, we can specify an arbitrary set of switches, and all of them will be switched from off to on simultaneously. We will then see which lamps come on (initially they are all off). For example, if $n = 10$ and we specify the set of switches {1, 2, 3} and lamps L_6, L_4 and L_9

come on, we know that switches {1, 2, 3} correspond to lamps L_6, L_4 and L_9 in *some order*. We want to determine which switch corresponds to which lamp. Obviously by switching only one switch per operation, we can achieve this in n operations. Can we do better?

Answer:
Yes. We actually need only $\lceil \log_2(n) \rceil$ operations, where $\lceil\ \rceil$ is the *ceiling* function. This is much better than *n* operations. For example, if *n* is one million, individually testing switches requires 999,999 operations, whereas our solution only requires 20. We give two solutions. For convenience assume *n* is even.

Solution 1:
In the first operation specify a set of $n/2$ switches. Now we have two sets of $n/2$ switches, and we know which $n/2$ lamps they both correspond to. Now we want to apply the algorithm for $n/2$ lamps and switches to the two sets. Hence it initially appears that we have the recursion $f(n) = 2f(n/2)+1$, where $f(n)$ is the number of steps taken by our algorithm for *n* lamps. However, note that we can actually apply the algorithms for both sets simultaneously, since we know which set of switches corresponds to which set of lamps. Thus the actual recursion is $f(n) = f(n/2)+1$. Since $f(1) = 0$, we inductively get $f(n) = \lceil \log_2(n) \rceil$.

Solution 2:
The algorithm in this solution is essentially equivalent to that in solution 1, but the thought process behind it is different. Label the switches 1, 2, …, *n*. Now read their labels in binary. Each label has at most $\lceil \log_2(n) \rceil$ digits. Now in operation 1, flip all switches that have a 1 in the units place of the binary representation of their labels. In general, in operation *k* we flip all switches that have a 1 in the *k*th position of their binary representation. At the end of $\lceil \log_2(n) \rceil$ operations, consider any lamp. Look at all the operations in which it came on. For example, if a lamp comes on in the second, third and fifth operations, but not in the first, fourth and

6^{th} operations, then it must correspond to the switch with binary representation 010110 (1s in the 2^{nd}, 3^{rd} and 5^{th} positions from the right). Thus each lamp can be uniquely matched to a switch and we are done. ■

Example 7 [Generalization of IMO shortlist 1998, C3]
Cards numbered 1 to n are arranged at random in a row with $n \geq 5$. In a move, one may choose any block of consecutive cards whose numbers are in ascending or descending order, and switch the block around. For example, if n=9, then 91 6 5 3 2 7 4 8 may be changed to 91 3 5 6 2 7 4 8. Prove that in at most $2n - 6$ moves, one can arrange the n cards so that their numbers are in ascending or descending order.

Answer:
We use a recursive algorithm relating the situation with n cards to the situation with n-1 cards. Let $f(n)$ be the minimum number of moves required to 'monotonize' any permutation of the n cards. Suppose we have a permutation with starting card k. In $f(n-1)$ moves, we can monotonize the remaining $(n-1)$ cards to get either the sequence $(k, 1, 2, \ldots, k-1, k+1, \ldots, n)$ or $(k, n, n-1, \ldots, k+1, k-1, \ldots, 2, 1)$. In one move, we can make the former sequence $(k, k-1, k-2, \ldots, 1, k+1, k+2, \ldots, n)$ and with one more move we get the sequence $(1, 2, 3, \ldots, n)$ and we are done. Similarly in the latter case we need only two additional moves to get $(n, n-1, \ldots, 1)$. Thus in either case, we can complete the task using $f(n-1) + 2$ moves, so $f(n) \leq f(n-1) + 2$.

Now to prove the bound for general $n \geq 5$, it suffices to prove it for $n = 5$ and then induct using $f(n) \leq f(n-1) + 2$. To prove that $f(5) \leq 4$, first note that $f(3) = 1$ and $f(4) = 3$. With a little case work (do this), we can show that any permutation of 4 cards can be monotonized *either way* in at most 3 moves (thus both {1, 2, 3, 4} and {4, 3, 2, 1} can be reached after at most 3 moves, regardless of the initial permutation). Now given a permutation of {1, 2, 3, 4, 5}, use one move if necessary to ensure that either 1 or 5 is at an extreme position. Now monotonize the remaining 4 numbers in 3 moves, in such a way that the whole sequence is monotonized (we can do this by the previous statement). Hence at most 4 moves are required for 5 cards, and we are done. ■

Remark: Since we wanted a linear bound in this problem, we tried to relate $f(n)$ to $f(n-1)$. However, when we want a logarithmic bound, we generally relate $f(n)$ to $f(n/2)$, or use binary representations. Thus the question itself often gives us a hint as to what strategy we should use.

Example 8 [Russia 2000]
Tanya choses a natural number $X \leq 100$, and Sasha is trying to guess this number. She can select two natural numbers M and N less than 100 and ask for the value of gcd $(X+M, N)$. Show that Sasha can determine Tanya's number with at most seven questions (the numbers M and N can change each question).

Answer:
Since $2^6 < 100 < 2^7$ we guess that more generally $\lceil \log_2(n) \rceil$ guesses are needed, where n is the maximum possible value of X and $\lceil \ \rceil$ is the ceiling function.

Our strategy is to determine the digits of X in binary notation; that is, the bits of X. First ask for gcd $(X+2, 2)$. This will tell us whether X is even or odd, so we will know the units bit of X. If X is even, ask for gcd $(X+4, 4)$. This tells us whether or not X is divisible by 4. Otherwise ask for gcd $(X+1, 4)$. This tells us if X is 1 or 3 mod 4 (if the gcd is 4, then $X+1$ is divisible by 4 and so $X \equiv 3$ mod 4). With this information we can determine the next bit of X. For example, if X is odd and is 3 mod 4, its last two bits will be 11. Now suppose $X = i$ mod 4. To determine the next digit, ask for gcd $(X + (4-i), 8)$. This gcd is either 4 or 8, according as $X = i$ or $4+i$ mod 8. This gives us the next bit. For example, if $X = 3$ mod 4 but $X = 7$ mod 8, then the last 3 bits of X will be 111, but if $X = 3$ mod 8, then the last 3 bits would be 011. Now the pattern is clear. We continue in this manner until we obtain all the bits of X. This takes k

questions, where k is the number of bits of n (since $X \leq n$, we don't have to ask for further bits), which is at most equal to $\lceil \log_2(n) \rceil$. ■

Example 9 [Generalization of Russia 2004, grade 9 problem 3]
On a table there are n boxes, where n is even and positive, and in each box there is one ball. Some of the balls are white and the number of white balls is even and greater than 0. In each turn we are allowed to point to two arbitrary boxes and ask whether there is at least one white ball in the two boxes (the answer is yes or no). Show that after $(2n - 3)$ questions we can indicate two boxes which definitely contain white balls.

Answer:
Label the boxes from 1 to n. Ask for the pairs of boxes $(1, j)$, where $j= 2, 3, \ldots, n$. If at some stage we get a no, this means box 1 contains a black ball. Then for all j such that we got a 'yes' for $(1, j)$, box j contains a white ball and we are done. The only other possibility is if we got a yes for all boxes $(1, j)$, in which case there are 2 possibilities: either box 1 has a white ball or box 1 has a black ball and all the other $(n$-1) boxes have white balls. The latter case is ruled out since we are given that an even number or boxes have black balls, and $(n$-1) is odd. Hence box 1 has a white ball. Now ask for the pairs $(2, j)$ where j =3, 4, …, n. Note that now we have asked a total of $(n-1) + (n-2) = (2n-3)$ questions. Again if we get a 'no' somewhere, then box 2 has a black ball and all yeses tell us which boxes have white balls. In this case we are done. The other case is if all the answers are yes. The same argument we used earlier shows that box 2 has a white ball and we are done. ■

Now we look at a simple problem from computer science (part *a* of the next problem), which also happens to be a game I played when I was a little kid. Part *b* is a little trick question I just came up with.

Example 10 a) [Binary Search Algorithm]

My friend thinks of a natural number X between 1 and 2^k inclusive. In one move I can ask the following question: I specify a number n and he says bigger, smaller or correct according as $n < X$, $n > X$, or $n = X$. Show that I can determine the number X using at most k moves.

Answer:
In my first move I say 2^{k-1}. Either I win, or I have reduced number of possibilities to 2^{k-1}. Repeat this process- in each stage reduce the number of possibilities by a factor of 2. Then in k moves there is only one possibility left.

Example: If $k = 6$, first guess 32. If the answer is "smaller", guess 16. If the answer is now "bigger", guess 24 (the average of 16 and 32). If the answer is now "smaller", guess 20. If the answer is again smaller, guess 18. If the answer is now "bigger", the number is 19.

In general if we replace 2^k with n, we need $\lceil \log_2(n) \rceil$ questions. ■

Example 10 b) [Binary Search with noise – a little trick]
We play the same game, but with a couple changes. First, I can now ask **any** yes or no question. Second, now my friend is allowed to lie - at most once in the whole game though. Now, from part *a*) I can win with $2k + 1$ moves: I simply ask each question twice, and if the answer changes, that means my friend has lied. Then I ask the question again, and this time I get the true answer (he can only lie once). Thus I ask each question twice except for possibly one question which I ask thrice, for a total of $2k + 1$ questions. Can I do better?

Answer:
First I ask about each of the k digits in the binary representation of X. If the game didn't involve lying, I would be done. Now I need to account for the possibility that one answer was a lie. I ask the question, "did you ever lie this game?" If the answer is no, we are done (if he had lied, he would have to say yes now as he can't lie

twice). If the answer is yes, I ask the question, "was the previous answer a lie?" If the answer to this is yes, then that means he never lied in the first k questions and again we are done. If the answer is no, then we can be sure that one of the first k answers we received (about the binary digits) was a lie. Note that he cannot lie anymore. We want to determine which answer was a lie. But using part a), we can do this in at most $\lceil \log_2(k) \rceil$ moves! This is because determining which of k moves was a lie is equivalent to guessing a number X' with $X' \le k$, and for this I use the algorithm in part a). After this, I know which digit in my original binary representation of X is wrong and I change it, and now I am done. I have used $k + 2 + \lceil \log_2(k) \rceil$ questions, which is much less than $2k + 1$ questions for large k.

In general, if 2^k is replaced by n, this algorithm takes $\lceil \log_2(n) \rceil +$ $\lceil \log_2 \lceil \log_2(n) \rceil \rceil + 2$ moves. ■

As in the previous chapter, we end this section with one of the hardest questions ever asked at the IMO. Only 3 out of over 550 contestants managed to completely solve it. However, the difficulty of this now famous problem has been hotly debated on AOPS, with many people arguing that it is a lot easier that the statistics indicate. We'll let the reader be the judge of that. The following solution is based on one found during the contest. The official solution is much more complicated.

Example 11 [IMO 2009, Problem 6]
Let n be a nonnegative integer. A grasshopper jumps along the real axis. He starts at point 0 and makes $n + 1$ jumps to the right with pairwise different positive integral lengths $a_1, a_2, \ldots a_{n+1}$ in an arbitrary order. Let $\boldsymbol{M}$ be a set of n positive integers in the interval $(0, s)$, where $s = a_1 + a_2 + \cdots + a_{n+1}$. Prove that the grasshopper can arrange his jumps in such a way that he never lands on a point from $\boldsymbol{M}$.

Answer:

We construct an algorithm using induction and the extremal principle. The case n = 1 is trivial, so let us assume that n > 1 and that the statement holds for 1, 2, ..., $n-1$. Assume that $a_1 < \cdots < a_n$. Let m be the smallest element of $\boldsymbol{M}$. Consider the following cases:

Case 1: $m < a_{n+1}$: If a_{n+1} does not belong to $\boldsymbol{M}$ then make the first jump of size a_{n+1}. The problem gets reduced to the sequence a_1, ..., a_n and the set $\boldsymbol{M}\backslash\{m\}$, which immediately follows by induction. So now suppose that $a_{n+1} \in \boldsymbol{M}$. Consider the following n pairs: (a_1, a_1+a_{n+1}), ..., (a_n, a_n+a_{n+1}). All numbers from these pairs that are in $\boldsymbol{M}$ belong to the ($n-1$)-element set $\boldsymbol{M}\backslash\{a_n\}$, hence at least one of these pairs, say (a_k, $a_k + a_n$), has both of its members outside of $\boldsymbol{M}$. If the first two jumps of the grasshopper are a_k and $a_k + a_{n+1}$, it has jumped over at least two members of $\boldsymbol{M}$: m and a_{n+1}. There are at most $n-2$ more elements of $\boldsymbol{M}$ to jump over, and $n-1$ more jumps, so we are done by induction.

Case 2: $m \geq a_{n+1}$: Note that it is equivalent to solve the problem in reverse: start from $s = a_1 + a_2 + \cdots + a_{n+1}$ and try to reach 0 without landing on any point in $\boldsymbol{M}$. By the induction hypothesis, the grasshopper can start from s make n jumps of sizes a_1, ..., a_n to the left, and avoid all the points of $\boldsymbol{M}\backslash\{m\}$. If it misses the point m as well, then we are done, since we can now make a jump of size a_{n+1} and reach 0. So suppose that after making the jump a_k the grasshopper landed at m. If it changes the jump a_k to the jump a_n, it will jump past m and all subsequent jumps will land outside of $\boldsymbol{M}$ because m is the left-most point. ■

Exercises

1. **[Spain 1997]**
The exact quantity of gas needed for a car to complete a single loop around a track is distributed among n containers placed along the track. Show that there exists a point from which the car can start with an empty tank and complete the loop (by collecting gas from tanks it encounters along the way). [Note: assume that there is no limit to the amount of gas the car can carry].

2. **[Russia]**
Arutyun and Amayak perform a magic trick as follows. A spectator writes down on a board a sequence of N (decimal) digits. Amayak covers two adjacent digits by a black disc. Then Arutyun comes and says both closed digits (and their order). For which minimal N can this trick always work?

3. **[Generalization of Russia 2005]**
Consider a game in which one person thinks of a permutation of $\{1, 2, \ldots, n\}$ and the other's task is to deduce this permutation (n is known to the guesser). In a turn, he is allowed to select three positions of the permutation and is told the relative order of the three numbers in those positions. For example, if the permutation is 2, 4, 3, 5, 1 and the guesser selects positions 1, 4 and 5, the other player will reveal that 5^{th} number < 1^{st} number < 4^{th} number. Determine the minimum number of moves for the guesser to always be able to figure out the permutation.

4. **[IMO Shortlist 1990]**
Given n countries with three representatives each, m committees $A_1, A_2, \ldots, A_m$ are called a *cycle* if

i. each committee has n members, one from each country;
ii. no two committees have the same membership;
iii. for $1 \leq i \leq m$, committee A_i and committee A_{i+1} have no member in common, where A_{m+1} denotes A_1
iv. if $1 < |i-j| < m-1$, then committees A_i and A_j have at least one member in common.

Is it possible to have a cycle of 1990 committees with 11 countries?

5. **[Canada 2012 – 4]**

A number of robots are placed on the squares of a finite, rectangular grid of squares. A square can hold any number of robots. Every edge of each square of the grid is classified as either passable or impassable. All edges on the boundary of the grid are impassable. A move consists of giving one of the commands up, down, left or right. All of the robots then simultaneously try to move in the specified direction. If the edge adjacent to a robot in that direction is passable, the robot moves across the edge and into the next square. Otherwise, the robot remains on its current square.

Suppose that for any individual robot, and any square on the grid, there is a finite sequence of commands that will move that robot to that square. Prove that you can also give a finite sequence of commands such that all of the robots end up on the same square at the same time.

6. **[IMO Shortlist 2002, C4]**

Let $\boldsymbol{T}$ be the set of ordered triples (x, y, z), where x, y, z are integers with $0 \leq x, y, z \leq 9$. Players A and B play the following guessing game. Player A chooses a triple in $\boldsymbol{T}$ (x, y, z), and Player B has to discover A's triple in as few moves as possible. A *move* consists of the following: B gives A a triple (a, b, c) in $\boldsymbol{T}$, and A replies by giving B the number $|x + y - a - b| + |y + z - b$

$- c| + |z + x - c - a|$. Find the minimum number of moves that B needs in order to determine A's triple.

7. **Given a finite set of points in the plane, each with integer** coordinates, is it always possible to color the points red or white so that for any straight line L parallel to one of the coordinate axes the difference (in absolute value) between the numbers of white and red points on L is not greater than 1?

8. **[Generalization of Russia 1993]**
   There are n people sitting in a circle, of which some are truthful and others are liars (we don't know who is a liar and who isn't though). Each person states whether the person to in front of him is a liar or not. The truthful people always tell the truth, whereas the liars may either lie or tell the truth. The aim is for us to use the information provided to find one person who is definitely truthful. Show that if the number of liars is at most $2\sqrt{n} - 3$, we can always do this.

9. On each square of a chessboard is a light which has two states- on or off. A move consists of choosing a square and changing the state of the bulbs in that square and in its neighboring squares (squares that share a side with it). Show that starting from any configuration we can make finitely many moves to reach a point where all the bulbs are switched off

10. **[Indian Postal Coaching 2011]**
    Let C be a circle, $A_1, A_2, \ldots, A_n$ be distinct points inside C and $B_1, B_2, \ldots, B_n$ be distinct points on C such that no two of the segments $A_1B_1, \ldots, A_nB_n$ intersect. A grasshopper can jump from A_r to A_s if the line segment A_rA_s does not intersect any line segment A_tB_t ($t \neq r, s$). Prove that after a certain number of jumps, the grasshopper can jump from any A_u to any A_v.

11. [USAMO 2000, Problem 3]

You are given R red cards, B blue cards and W white cards and asked to arrange them in a row from left to right. Once arranged, each card receives a score as follows. Each blue card receives a score equal to the number of white cards to its right. Each white card receives a score equal to twice the number of red cards to its right. Each red card receives a score equal to three times the number of blue cards to its right. For example, if the arrangement is Red Blue White Blue Red Blue, the total score will be 9 + 1 + 2 + 0 + 3 + 0 = 15. Determine, as a function of R, B and W, the minimum possible score that can obtained, and find all configurations that achieve this minimum.

12. [IMO Shortlist 2005, C7]

Suppose that $a_1, a_2, \ldots, a_n$ are integers such that $n \mid (a_1 + a_2 + \ldots + a_n)$. Prove that there exist two permutations $b_1, b_2, \ldots, b_n$ and $c_1, c_2, \ldots, c_n$ of $(1, 2, \ldots, n)$ such that for each integer i with $1 \leq i \leq n$, we have $n \mid a_i - b_i - c_i$.

13. [St. Petersburg 2001]

In the parliament of the country of Alternativia, for any two deputies there exists a third who is acquainted with exactly one of the two. There are two parties, and each deputy belongs to exactly one of them. Each day the President (not a member of the parliament) selects a group of deputies and orders them to switch parties, at which time each deputy acquainted with at least one member of the group also switches parties. Show that the President can eventually ensure that all deputies belong to the same party.

Olympiad Combinatorics

Pranav A. Sriram

August 2014

Chapter 3: Processes

Copyright notices

About the Author

Pranav Sriram graduated from high school at The International School Bangalore, India, and will be a Freshman at Stanford University this Fall.

3. PROCESSES

Introduction

In this chapter we analyze combinatorial processes. In Chapter 1 on algorithms, we often encountered combinatorial processes from a different viewpoint. Problems from both chapters are similar in that they typically specify an initial configuration and allowed set of moves. In the chapter on algorithms, we were asked to prove that a certain final configuration could be reached using these moves, and we solved these problems by constructing procedures to reach the desired final configuration. In this chapter, our job is not to construct our own procedures, but rather to analyze given ones. Some questions ask us to determine whether a process terminates, and if it does, what the final configuration looks like. Others may ask us to bound the number of steps it takes for a certain configuration to arise.

Our main tools for this chapter are invariants, the extremal principle, induction and other clever ideas that we will develop as we go further, such as making transformations to a problem that simplify the problem but leave the end result invariant. It must also be stressed that **experimentation, trial and error, observation, intuition and conjectures** play a big role in solving problems related to processes (and combinatorics in general). We remind the reader that the ideas to solve combinatorial problems often arise from experimenting with small values.

Invariants

Our first few examples use *invariants*, a technique we have already used in earlier chapters. The usefulness of invariants while analyzing combinatorial processes can hardly be overstated.

Example 1 [Indian TST 2004]
The game of pebbles is played as follows. Initially there is a pebble at (0, 0). In a move one can remove a pebble from (i, j) and place one pebble each on $(i+1, j)$ and $(i, j+1)$, provided (i, j) had a pebble to begin with and $(i+1, j)$ and $(i, j+1)$ did not have pebbles. Prove that at any point in the game there will be a pebble at some lattice point (a, b) with $a+b \leq 3$.

Answer:
Clearly the pebbles will always be on lattice points in the first quadrant. How can we find an invariant? Just assign a weight of $2^{-(i+j)}$ to a pebble at (i, j). Then in each move one pebble is replaced by two pebbles, each having half its weight. So the total weight of pebbles is invariant. Initially the weight is $2^0 = 1$. Suppose at some stage no pebble is on a point (a, b) with $a+b \leq 3$. Then the maximum possible total weight of all pebbles is the weight of the whole first quadrant minus that of the squares (a, b) with $a+b \leq 3$, which is

$$\left(\sum_{i=0}^{\infty}\sum_{j=0}^{\infty} 2^{-(i+j)}\right) - (1 + 2 \times \tfrac{1}{2} + 3 \times \tfrac{1}{4} + 4 \times \tfrac{1}{8})$$
$$= 4 - (1 + 1 + \tfrac{3}{4} + \tfrac{1}{2}) = \tfrac{3}{4} < 1.$$

This is a contradiction as the weight should always be 1. ■

Remark: The double summation was computed by noticing $\sum_{i=0}^{\infty}\sum_{j=0}^{\infty} 2^{-(i+j)} = \sum_{i=0}^{\infty} 2^{-i} \times \sum_{i=0}^{\infty} 2^{-j} = 2 \times 2 = 4$. The second parenthesis was the weight of all squares a, b with $a+b \leq 3$.

Example 2 [IMO shortlist 2005, C5]
There are n markers, each with one side white and the other side black, aligned in a row so that their white sides are up. In each step, if possible, we choose a marker with the white side up (but not one of the outermost markers), remove it and reverse the closest marker to the left and the closest marker to the right of it. Prove that one can achieve the state with only two markers remaining if and only if $n-1$ is not divisible by 3.

Answer:
If $n-1$ is not divisible by 3, it is easy to construct an inductive algorithm to make only two markers remain (chapter 2 FTW!). We leave this to the reader (just do it for $n = 5$, $n = 6$ and induct with step 3). Now we do the harder part: if $(n-1)$ is divisible by 3, we need to show that this cannot be done.

Call a marker *black* or B if its black side is up and *white* or W if its white side is up. One invariant we immediately find is that the number of black markers is always even since each move changes the number of black markers by 0, 2 or -2. Now we look for another invariant.

We assign numbers to each white marker. If a white marker has t black markers to its left, we assign the number $(-1)^t$ to it. Let S be the sum of all the labels. Initially all labels are $(-1)^0 = 1$, so $S = n$ initially. The labels may keep changing, but we claim that S stays invariant mod 3. For example, suppose we have the substring …*WWB*… and remove the middle white marker. Then it becomes …*BW*…. If the 2 white markers had t black markers to their left initially, then the white marker now has $(t+1)$ black markers to its left. Thus the two white markers both had labels $(-1)^t$ initially, but now the white marker has label $(-1)^{t+1}$. The sum of the labels has changed by $(-1)^{t+1} - 2(-1)^t = 3(-1)^{t+1} \equiv 0 \mod 3$. The reader can verify that in the other cases (*WWW*, *BWB*, *BWW*) as well the sum of labels S doesn't change mod 3.

Now the rest is easy. If two markers remain, they are either both white or both black (number of black markers must be even). In the first case, both labels are 1 and $S = 2$. In the second case, $S = 0$ as no markers are labeled. So $S = 0$ or 2 at the end and $S = n$ in the beginning. Since S stays invariant mod 3, $n \equiv 0$ or 2 mod 3 and we are done. ■

Example 3 [IMO Shortlist 1998 C7]
A solitaire game is played on an $m \times$ n board with markers having one white side and one black side. Each of the mn cells contains a marker with its white side up, except for one corner square which has a marker with its black side up. The allowed move is to select a marker with black side up, remove it, and turn over all markers in squares sharing a side with the square of the chosen marker. Determine all pairs (m, n) for which it is possible to remove all markers from the board.

Answer:
It is natural (but not essential) to rephrase the problem using graph theory. We take the markers as vertices. Each vertex is black or white. 2 vertices are connected by an edge if and only if the markers lie on adjacent squares. In each move, we are deleting one black vertex and all its incident edges, but all its white neighbors become black and all its white neighbors become black. Suppose in a move we delete a black vertex v and s edges, where s is the degree of v. Suppose w of v's neighbors were white vertices, and $(s–w)$ were black vertices. Then these w vertices become black and $(s–w)$ become white, so the number of white vertices increases by $s – 2w$.

This information alone does not immediately give us an invariant, since the quantity $s – 2w$ is quite random. However, suppose we consider the quantity $W + E$, where W is the total number of remaining white vertices (don't confuse this with w) and E is the number of edges. Then when E reduces by s and W

changes by $s - 2w$, $(W+E)$ decreases by $2w$, which is always even. Hence **the parity of ($W+E$) remains the same**. So if $W+E$ is 0 at the end (when all markers are gone), we need $W+E$ to be even in the beginning. But initially

$$W = mn\text{-}1,\ E = m(n\text{-}1) + n(m\text{-}1), \text{ and}$$

$$W+E = 3mn - m - n - 1 \equiv mn - m - n + 1 \pmod 2 = (m\text{-}1)(n\text{-}1),$$

so at least one of m and n must be odd. In this case the task is indeed possible and we leave it to the reader to find an algorithm. (Assume m is odd and use an inductive procedure that makes each column empty one by one). ■

Good and Bad Objects

Another useful idea while analyzing processes is to distinguish between "good" and "bad" objects. For example, if at the end of a process we want to show that all objects satisfy a certain property, call objects with that property *good* and the other objects *bad*. We will use this idea in different forms several times throughout this chapter. The next example combines this idea with monovariants by showing that the number of "good" objects monotonically increases.

Example 4 [Based on Canada 1994]
There are $2n+1$ lamps placed in a circle. Each day, some of the lamps change state (from on to off or off to on), according to the following rules. On the k^{th} day, if a lamp is in the same state as at least one of its neighbors, then it will not change state the next day. If a lamp is in a different state from both of its neighbors on the k^{th} day, then it will change its state the next day. Show that regardless of the initial states of each lamp, after some point none of the lamps will change state.

Answer:
Call a lamp "good" if it is in the same state as at least one of its neighbors. Once a lamp is good, it will remain good forever (if two adjacent lamps are in the same state on the k^{th} day, they will not change state the next day, and hence both remain good). Hence the number of good lamps never decreases and is a monovariant.

We show that in fact the number of good lamps strictly increases until it reaches $2n+1$. Initially there must be 2 adjacent lamps with the same state since the number of lamps is odd. Suppose at some point there are j good lamps and $2 \le j < 2n+1$. Then there must exist 2 adjacent lamps such that one is bad and one is good. Then the bad lamp will switch states the next day and the good lamp will remain in the same state. Then the bad lamp will now be good, so the number of good lamps has increased (remember all good lamps remain good). So the number of good lamps increases until all lamps are good, and at this point there will be no more changes of state. ■

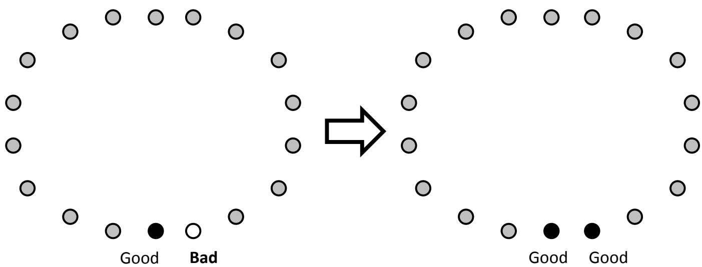

Figure 3.1: Bad lamps next to good lamps become good

Bounds on the number of steps

Now we look at another class of problems, which ask us to bound the number of steps or moves it takes for a process to terminate. To bound the total number of moves, it is often useful to bound the number of times a particular object is involved in a move.

Example 5 [USAMO 2010-2]
There are n students standing in a circle, one behind the other, all facing clockwise. The students have heights $h_1 < h_2 < \ldots < h_n$. If a student with height h_k is standing directly behind a student with height h_{k-2} or less, the two students are permitted to switch places. Prove that it is not possible to make more than ${}^nC_3$ such switches before reaching a position in which no further switches are possible.

Answer:
We bound the number of times an individual student can switch places with another student. Let s_k denote the number of times the student with height k switches with someone shorter. Obviously $s_1 = s_2 = 0$. Now consider the number of people between student k and student $(k-1)$ (along the clockwise direction) of height less than h_{k-1}. This number is at most $(k-2)$. This quantity decreases by 1 each time student k switches with someone shorter and increases by one each time student $k-1$ switches with someone shorter. This quantity doesn't change when students taller than student k switch with either student k or student $(k-1)$. Hence $s_k - s_{k-1}$ denotes the decrease from beginning to end in the number of students between student k and student $(k-1)$, which cannot exceed $(k-2)$. Thus we get the bound $s_k - s_{k-1} \le (k-2)$. Using this recursive bound and the initial values $s_1 = s_2 = 0$, we get $s_3 \le 1$, $s_4 \le 3$, etc. In general it is easy to show by induction $s_k \le \binom{k-1}{2}$. Hence the total number of moves cannot exceed $\sum_{k=1}^{n}\binom{k-1}{2} = {}^nC_3$ and we are done. ■

Note: The last step uses a well-known binomial identity. More generally
$\sum_{k=1}^{n}\binom{k-1}{r} = \binom{n}{r+1}$, where by convention $\binom{j}{r} = 0$ for $j < r$.

Example 6 [Based on Canada 2012, IMO Shortlist 1994 C4]
A bookshelf contains n volumes, labeled 1 to n, in some order. The

librarian wishes to put them in the correct order as follows. The librarian selects a volume that is too far to the right, say the volume with label k, takes it out, and inserts it in the k^{th} position. For example, if the bookshelf contains the volumes 3, 1, 4, 2 in that order, the librarian could take out volume 2 and place it in the second position. The books will then be in the order 3, 2, 1, 4. Show that the sequence $(1, 2, \dots, n)$ is reached in fewer than 2^n moves.

Answer:

We bound the number of times book k can be selected by the librarian. Clearly, book n can never be selected since it will never be too far right. Book 1 can only be selected once, because once selected, it will move to the first position and never move again. Book 2 can be selected twice: it may be selected once and put in the correct position, but then it may move because of book 1.

More generally, let $f(k)$ denote the number of times book k is selected for $1 \le k \le (n-1)$. We have

$$f(k) \le 1 + f(k-1) + f(k-2) + \dots + f(1).$$

This is because once k is in the correct position, it can only be displaced $f(k-1) + f(k-2) + \dots + f(1)$ times, because the only way in which book k can be displaced is if one of the books with number less than k "pushes" k.

For example: If we start from (4, 1, 3, 2, 5) and we choose book 2, it becomes (4, 2, 1, 3, 5). Book 3 was in the correct position, but has been "pushed out" because of book 2 being chosen.

Thus using this recursive bound on $f(k)$ and the fact that $f(1) = 1$, we obtain by a simple induction $f(k) \le 2^{k-1}$. Hence the total number of moves required is at most

$f(1) + f(2) + \dots + f(n-1) \le 1 + 2 + 4 + \dots + 2^{n-2} = 2^{n-1} - 1$. ■

Remark: A solution with monovariants is also possible.

Induction

In the previous section, we essentially broke down the analysis of a process into the analysis of the individual entities involved. To find the total time for the process to terminate, we used recursive bounds to estimate the time a particular object could contribute to the total time. These essential elements of somehow "breaking down" a process and using induction and/or recursion will be central to this section as well. However, rather than the object-centric approach of the previous section, a structure-centric approach will be taken here: the inductive proofs will rely on exploiting the nice combinatorial structure of n x n boards.

Example 7 [Belarus 2001]
Let n be a positive integer. Each unit square of a $(2n\text{-}1) \times (2n\text{-}1)$ square board contains an arrow, either pointing up, down left or right. A beetle sits in one of the squares. In one move, the beetle moves one unit in the direction of the arrow in the square it is sitting on, and either reaches an adjacent square or leaves the board. Then the arrow of the square the beetle left turns 90^{o} clockwise. Prove that the beetle leaves the board in at most $2^{3n\text{-}1}(n\text{-}1)! - 3$ moves.

Answer:
The base case $n = 1$ is trivial since the beetle leaves in the first move. Now suppose the result is true for $n = k$; we prove it for $n = k+1$. It is natural to distinguish between boundary squares (squares on the edge of the board) and interior squares, since the interior squares form a $(2k\text{-}1) \times (2k\text{-}1)$ board and we can use the induction hypothesis on this board. We further distinguish between corner squares and non-corner boundary squares.

Suppose the beetle is still on the board after T moves. We want to show that $T < 2^{3(k+1)\text{-}1}k! - 3$. At this stage, if any non-corner boundary square has been visited 4 times, then one of the four times the arrow would have been pointing out of the board (since its direction changes each time). Similarly if a corner square has been visited 3 times, then at least once it would have pointed out of the board. Hence in each of the cases, the beetle would have left the board, contradiction. Hence the beetle has visited each corner square at most twice and each non-corner boundary square at most thrice. Moreover, the beetle can move at most once from a non-corner boundary square to an interior square. Thus:

i. The beetle has made at most $2 \times 4 + 3(8k - 4) = 24k - 4$ moves from boundary squares to other squares of the board (since there are 4 corner squares and $8k - 4$ non-corner boundary squares).

ii. The beetle has made at most $4(2k - 1) = 8k - 4$ moves from a boundary square to an interior square, since there are $8k - 4$ non-corner boundary squares.

iii. If a beetle is in the interior $(2k - 1) \times (2k - 1)$ square, it can make at most $M = 2^{3k-1}(k - 1)! - 3$ moves before reaching a boundary square, by the induction hypothesis.

iv. From (ii), the beetle can stay in the interior square for at most $8k - 3$ periods of time (once in the beginning, then once for each time it moves from a boundary square back to the interior). Each period lasts at most $2^{3k-1}(k - 1)! - 3$ moves by (iii). Hence the number of moves made from interior squares is at most

$$(2^{3k-1}(k-1)! - 3) \times (8k - 3)$$
$$< 8k \times (2^{3k-1}(k-1)! - 3)$$
$$= 2^{3(k+1)-1}k! - 24k.$$

From (i) and (iv), we see that
$T \leq (24k - 4) + (2^{3(k+1)-1}k! - 24k) = 2^{3(k+1)-1}k! - 4$, as desired. ■

Now we look at another example using induction. This problem is different from the previous one in that we are not asked to bound the number of moves for a process to terminate. However, the idea of inducting by dividing an $n \times n$ board into an $(n-1) \times (n-1)$ sub board and an additional row and column is very similar to the idea in the previous example. This inductive technique is just one of many ways in which the structure of boards can be exploited.

Example 8 [Russia 2010]
On an $n \times n$ chart where $n \geq 4$, stand n '+' signs in cells of one diagonal and a '–' sign in all the other cells. In a move, one can change all the signs in one row or in one column, (– changes to + and + changes to –). Prove that it is impossible to reach a stage where there are fewer than n pluses on the board.

Answer:
Note that operating twice on a row is equivalent to not operating on it at all. So we can assume that each row and column has been operated upon 0 or 1 times. Now we use induction on n. The base case $n = 4$ is not entirely trivial, but is left to the reader in keeping with my general habit of dismissing base cases.

Now passing to the induction step, given an $n \times n$ board there are at least $(n - 1)$ pluses in the bottom right $(n - 1) \times (n - 1)$ square by the induction hypothesis. If we have a plus in the first row or column we are done. Suppose there is no plus in the first column or row. Then either the first row or the first column (but not both) has been operated upon (otherwise the top left square would have a plus). Suppose WLOG the first row has been operated upon. Then columns 2, 3, ..., n have all been operated upon (otherwise row 1 would have a plus). Also no other row has been operated upon (otherwise the first column would have a plus). But in this case, the lower right $(n - 1) \times (n - 1)$ square has had all its columns and none of its rows operated upon, and hence each column has $(n - 2)$ pluses. In total it has $(n - 2)(n - 1) > n$ pluses, so in this case as well we are done. ■

Problem Alteration: Don't play by the rules

Next we look at a very powerful technique of solving problems related to processes. In the next three examples, we will alter the problem statement slightly in such a way that the result we need to show doesn't change, but the process becomes much easier to analyze. In other words, we simplify the process to be analyzed while leaving the aspect of the process that we want to prove something about invariant. This may take some time to understand, so read through the next few examples slowly and carefully, and multiple times if necessary.

Example 9 [warm up for example 11]
There are n ants on a stick of length one unit, each facing left or right. At time $t = 0$, each ant starts moving with a speed of 1 unit per second in the direction it is facing. If an ant reaches the end of the stick, it falls off and doesn't reappear. When two ants moving in opposite directions collide, they both turn around and continue moving with the same speed (but in the opposite direction). Show that all ants will fall off the stick in at most 1 second. (We will use a very similar idea in example 11, so make sure you understand this trick.)

Answer:
The key observation is that the problem doesn't change if we alter it as follows: when two ants moving in opposite directions meet, they simply pass through each other and continue moving at the same speed. Thus instead of rebounding, if the ants pass through each other, the only difference from the original problem is that the identities of the ants get exchanged, which is inconsequential. Now the statement is obvious – each ant is unaffected by the others, and so each ant will fall of the stick of length one unit in at most 1 second. ■

Example 10 [Russia 1993 generalized]
The integers from 1 to n are written in a line in some order. The following operation is performed with this line: if the first number is k then the first k numbers are rewritten in reverse order. Prove that after some finite number of these operations, the first number in the line of numbers will be 1.

Answer:
The base case $n = 1$ is trivial. Suppose the result is true for $(n - 1)$. First observe that if n appears in the first position at some point, then in the next step n will be in the last position and will remain there permanently. Then we can effectively ignore n and we are done by induction. So suppose n never appears in the first position. Let j be the number in the last position. If we switch n and j, it has absolutely no effect on the problem, as j will never appear in the first position (since we assumed n will never appear in the first position). Now n is in the last position and as in the first case, we are done by induction. ■

Remark: Based on the above proof, it is not difficult to show that for $n > 1$ if the first number becomes 1 after at most $f(n)$ operations, we have the recursive bound $f(n+1) \leq 2f(n)+1$. I believe this bound can be further improved for most values of n.

As if "cheating" once isn't bad enough, we'll cheat twice in the next problem. Combining the insights obtained from these two instances of "cheating" will greatly restrict the possible positions of otherwise very chaotic ants.

Example 11 [IMO Shortlist 2011, C5]
Let m be a positive integer and consider a checkerboard consisting of $m \times m$ unit squares. At the midpoints of some of these unit squares there is an ant. At time 0, each ant starts moving with speed 1 parallel to some edge of the checkerboard. When two ants moving in opposite directions meet, they both turn 90° clockwise and continue moving with speed 1. When more than two ants meet, or when two ants moving in perpendicular

directions meet, the ants continue moving in the same direction as before they met. When an ant reaches one of the edges of the checkerboard, it falls off and will not reappear. Considering all possible starting positions, determine the latest possible moment at which the last ant falls off the checkerboard or prove that such a moment does not necessarily exist.

Answer:
After experimenting with small values of m, we conjecture that the answer is $\frac{3m}{2}$ -1. Clearly this is attainable if initially there are only 2 ants, one in the bottom left square facing upwards and one in the top left square facing downward. Now we prove that it is the maximum possible. Let U, D, L, R represent the directions up, down, left and right respectively.

Step 1: We use a modified version of the trick in example 7. Using the same reasoning, we can change the rules so that each ant travels in only two directions- either U and R or D and L. So if an ant travelling R meets an ant travelling L, they now move U and D respectively (even though in the original problem they should now move D and U respectively). This doesn't affect the problem. Now based on their initial direction, each ant can be classified into two types: UR or DL. UR ants can only move up and right the whole time and DL ants only move down and left the whole time. Note that we can ignore collisions between two ants of the same type. From now on, "collision" only refers to collisions between two ants of opposite types.

Step 2: Choose a coordinate system such that the corners of the checkerboard are $(0, 0)$, $(m, 0)$, (m, m) and $(0, m)$. At time t, there will be no UR ants in the region $\{(x, y): x + y < t + 1\}$ and no DL ants in the region $\{(x, y): x + y > 2m-t-1\}$. So if a collision occurs at (x, y) at time t, we have $t + 1 \leq x + y \leq 2m-t-1$.

Step 3: In a similar manner, we can change the rules of the original problem (without affecting it) by assuming that each ant can only move U and L or D and R, so each ant is UL or DR. Using the same logic as in step 2, we get a bound $|x-y| \leq m-t-1$ for each collision at point (x, y) and time t. Thus we have shown that all collisions at time t are within the region bounded by the 4 lines represented by the equations $t + 1 \leq x + y \leq 2m-t-1$ and $|x-y| \leq m-t-1$.

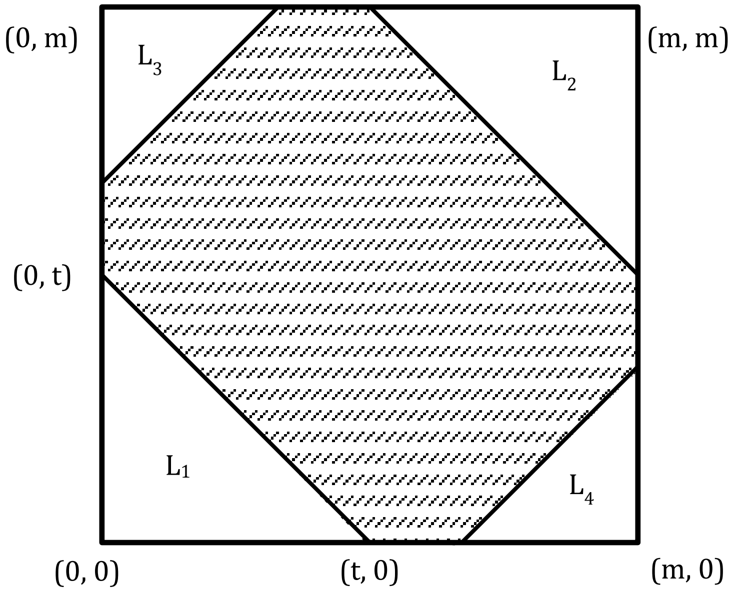

Figure 3.1: All collisions at time t must lie within the shaded region

Step 4: We finish the proof for a UR ant; by symmetry the same final bound will hold for DL ants. Take a UR ant and suppose its last collision is at (x, y) at time t. Adding the bounds $x+y \geq t+1$ and $x - y \geq -(m - t - 1)$, we get $x \geq t + 1 - \frac{m}{2}$. Similarly, $y \geq t + 1 - \frac{m}{2}$. Since this is the last collision, the ant will now move straight to an edge and fall off. This takes at most $m-\min\{x, y\}$ units of time. The total amount of time the ant stays on the board is hence at most

$$t + m - \min\{x, y\} \leq t + m - \{t + 1 - m/2\} = \frac{3m}{2} - 1$$

units of time. ■

Remark: Let's reverse engineer the solution a little bit, to see how the main ideas fit together so nicely – did you notice how the parameter t disappeared so conveniently in the last step? The basic goal in the above solution was to obtain tight bounds on the location of an ant after its last collision because after this the ant travels straight off the board. The intuition behind getting rid of t was that the longer an ant has been wandering around till its last collision, the closer it must be to an edge, and so the less time it will take to fall off now. But for this to work we need the ants to be "well behaved" - and hence the cheating!

Concluding Examples

Our final two examples lie at the heart of this chapter. Example 12 is a particular case of a more general and extensively studied process known as a "chip firing game", and Example 13 is a distant cousin of the chip firing game. Through these problems we introduce some important ideas such as using the extremal principle in different ways and obtaining contradictions, and combine these with ideas we have already seen like invariants and making assumptions that don't affect the problem. In example 12, we use the following idea: if a process never terminates, there must be some object that is moved or operated upon infinitely times. If we can find an object that is only operated upon finitely many times, we may be able to get a contradiction.

Example 12 [IMO shortlist 1994, C5]
1994 girls are seated in a circle. Initially one girl is given n coins. In one move, each girl with at least 2 coins passes one coin to each of her two neighbors.

(a) Show that if $n < 1994$, the game must terminate.
(b) Show that if $n = 1994$, the game cannot terminate.

Answer:

(a) Label the girls G_1, G_2, ..., G_{1994} and let G_{1995} = G_1, G_0 = G_{1994}. Suppose the game doesn't terminate. Then some girl must pass coins infinitely times. If some girl passes only finitely many times, there exist two adjacent girls, one of whom has passed finitely many times and one of whom has passed infinitely many times. The girl who has passed finitely many times will then indefinitely accumulate coins after her final pass, which is impossible. Hence every girl must pass coins infinitely many times.

Now the key idea is the following: For any two neighboring girls G_i and G_{i+1}, let c_i be the first coin ever passed between them. After this, we may assume that c_i always stays stuck between G_i and G_{i+1}, because whenever one of them has c_i and makes a move, we can assume the coin passed to the other girl was c_i. Therefore, each coin is eventually stuck between two girls. Since there are fewer than 1994 coins, this means there exist two adjacent girls who have never passed a coin to each other. This contradicts the result of the first paragraph.

(b) This is simple using invariants. Let a coin with girl i have weight i, and let G_1 have all coins initially. In each pass from G_i to her neighbors, the total weight either doesn't change or changes by ±1994 (if G_1 passes to G_{1994} or vice versa). So the total weight is invariant mod 1994. The initial weight is 1994, so the weight will always be divisible by 1994. If the game terminates, then each girl has one coin, so the final weight is 1+2+3+...+ 1994 = (1994 x 1995)/2 which is not divisible by 1994. Contradiction. ■

Before reading the solution to the next problem, we recommend that the reader experiment with small values of n and try to guess what the final configuration looks like. Several combinatorics problems require experimentation, observation and conjecturing before actually proving anything.

Example 13 [IMO shortlist 2001, C7]
A pile of n pebbles is placed in a vertical column. This configuration is modified according to the following rules. A pebble can be moved if it is at the top of a column which contains at least two more pebbles than the column immediately to its right. (If there are no pebbles to the right, think of this as a column with 0 pebbles.) At each stage, choose a pebble from among those that can be moved (if there are any) and place it at the top of the column to its right. If no pebbles can be moved, the configuration is called a *final configuration*. For each n, show that, no matter what choices are made at each stage, the final configuration obtained is unique. Describe that configuration in terms of n.

Answer:
It is clear that if s_i denotes the number of stones in column i, then in the final configuration $s_{i+1} = s_i$ or s_i -1. After experimenting with small values of n, we are led to the following claim:

Claim: In the final configuration, there is at most one index i such that $s_{i+1} = s_i$ (hence the remaining columns satisfy $s_{j+1} = s_j - 1$).

Proof: Call an index j bad if $s_{j+1} = s_j$. Assume to the contrary that there exist (at least) 2 bad indices in the final configuration. Take k and m ($k > m$) to be consecutive bad indices. Then $s_{k+1} = s_k$, $s_{m+1} = s_m$ and $s_{i+1} = s_i - 1$ for $m < i < k$. Consider the *earliest* configuration, say C, with the 2 bad indices. Now look at the last move before C. Since C is the earliest such configuration, the last move was either from the k^{th} or m^{th} column. But then in either case the configuration before C also had 2 bad indices, contradicting our assumption. This proves the claim.

Now it is easy to see that the claim uniquely determines the final configuration. For example, for n = 17 the final heights would be (5, 4, 3, 2, 2, 1). ■

Exercises

1. **[Austrian-Polish Mathematical Competition 1997]**
The numbers $49/k$ for k = 1, 2, …, 97 are written on a blackboard. A move consists of choosing two number a and b, erasing them and writing the number $2ab - a - b + 1$ in their place. After 96 moves, only one number remains. Find all possible values of this number.

2. We have $n(n+1)/2$ stones in k piles. In each move we take one stone from each pile and form a new pile with these stones (if a pile has only one stone, after that stone is removed the pile vanishes). Show that regardless of the initial configuration, we always end up with n piles, having 1, 2, …, n stones respectively.

3. **[ELMO Shortlist 2013, C9, generalized]**
There are n people at a party. Each person holds some number of coins. Every minute, each person who has at least $(n - 1)$ coins simultaneously gives one coin to every other person at the party. (So, it is possible that A gives B a coin and B gives A a coin at the same time.) Suppose that this process continues indefinitely. That is, for any positive integer m, there exists a person who will give away coins during the m^{th} minute. What is the smallest number of coins that could be at the party?

4. **[China TST 2003]**
There is a frog in every vertex of a regular $2n$-gon ($n \geq 2$). At a certain time, all frogs jump simultaneously jump to one of their neighboring vertices. (There can be more than one frog in one vertex). Suppose after this jump, no line connecting any two distinct vertices having frogs on it after the jump passes through the circumcentre of the $2n$-gon. Find all possible values of n for which this can occur.

5. [Chip firing lemma]

Let $\boldsymbol{G}$ be a connected graph with m edges. Consider $2m+1$ frogs, each placed on some vertex of $\boldsymbol{G}$. At each second, if a vertex v contains at least d_v frogs, then d_v of the frogs on v jump, one on each of the d_v adjacent vertices. Show that every vertex will be visited by a frog at some point.

6. [IMO 1986, Problem 3]

An integer is assigned to each vertex of a regular pentagon, and the sum of all five integers is positive. If three consecutive vertices are assigned the numbers x, y, z respectively, and $y < 0$, then the following operation is allowed: x, y, z are replaced by $x+y$, $-y$, $z+y$ respectively. Such an operation is performed repeatedly as long as at least one of the five numbers is negative. Determine whether this procedure necessarily comes to an end after a finite number of steps.

7. [Russia 1997]

There are some stones placed on an infinite (in both directions) row of squares labeled by integers. (There may be more than one stone on a given square). There are two types of moves:

(i) Remove one stone from each of the squares n and $n - 1$ and place one stone on $n + 1$

(ii) Remove two stones from square n and place one stone on each of the squares $n + 1$ and $n - 2$.

Show that at some point no more moves can be made, and this final configuration is independent of the choice of moves.

8. [APMO 2007, Problem 5]

A regular 5×5 array of lights is defective, so that toggling the switch for one light causes each adjacent light in the same row and in the same column as well as the light itself to change state, from on to off, or from off to on. Initially all the lights are switched off. After a certain number of toggles, exactly one light is switched on. Find all the possible positions of this light.

9. [IMO Shortlist 2007, C4]

Let $A_0 = \{a_1, a_2, ..., a_n\}$ be a finite sequence of real numbers. For each $k \geq 0$, from the sequence $A_k = \{x_1, x_2, ..., x_n\}$ we construct a new sequence A_{k+1} in the following way:

(i) We choose a partition $\{1, 2, ..., n\} = I \cup J$, where I and J are two disjoint sets, such that the expression $|\sum_{i \in I} x_i - \sum_{j \in J} x_j|$ is minimized. (We allow I or J to be empty; in this case the corresponding sum is 0.) If there are several such partitions, one is chosen arbitrarily.

(ii) We set $A_{k+1} = \{y_1, y_2, ..., y_n\}$, where $y_i = x_i + 1$ if $i \in I$, and $y_i = x_i - 1$ if $i \in J$.

Prove that for some k, the sequence A_k contains an element x such that $|x| \geq n/2$.

10. [Romanian TST 2002]

After elections, every Member of Parliament (MP) has his own absolute rating. When the parliament is set up, he enters a group and gets a relative rating. The relative rating is the ratio of its own absolute rating to the sum of all absolute ratings of the MPs in the group. An MP can move from one group to another only if in his new group his relative rating is greater. In a given day, only one MP can change the group. Show that only a finite number of group moves is possible (that is, the process eventually terminates).

11. [ELMO Shortlist 2013, C10]

Let $N > 1$ be a fixed positive integer. There are $2N$ people, numbered 1, 2, ..., $2N$, participating in a tennis tournament. For any two positive integers i, j with $1 \leq i < j \leq 2N$, player i has a higher skill level than player j. Prior to the first round, the players are paired arbitrarily and each pair is assigned a unique court among N courts, numbered 1, 2, ..., N.

During a round, each player plays against the other person assigned to his court (so that exactly one match takes place per court), and the player with higher skill wins the match (in other words, there are no upsets). Afterwards, for $i = 2, 3, \ldots, N$, the winner of court i moves to court $(i - 1)$ and the loser of court i stays on court i; however, the winner of court 1 stays on court 1 and the loser of court 1 moves to court N.

Find all positive integers M such that, regardless of the initial pairing, the players 2, 3, …, N+1 all change courts immediately after the M^{th} round.

12. [IMO 1993, Problem 3]
On an infinite chessboard, a solitaire game is played as follows: at the start, we have n^2 pieces occupying a square of side n. The only allowed move is to jump over an occupied square to an unoccupied one, and the piece which has been jumped over is removed. For which n can the game end with only one piece remaining on the board?

13. [South Korea TST 2009]
2008 white stones and 1 black stone are in a row. A move consists of selecting one black stone and change the color of its neighboring stone(s). The goal is to make all stones black after a finite number of moves. Find all possible initial positions of the black stone for which this is possible.

14. [IMO Shortlist 1996, C7]
A finite number of coins are placed on an infinite (in both directions) row of squares. A sequence of moves is performed as follows: at each stage a square containing more than one coin is chosen. Two coins are taken from this square; one of them is placed on the square immediately to the left while the other is placed on the square immediately to the right of the chosen square. The sequence terminates if at some point there is at most one coin on each square. Given some initial

configuration, show that any legal sequence of moves will terminate after the same number of steps and with the same final configuration.

15. [IMO Shortlist 2010, C6]

Given a positive integer k and other two integers $b > w > 1$. There are two strings of pearls, a string of b black pearls and a string of w white pearls. The length of a string is the number of pearls on it. One cuts these strings in some steps by the following rules. In each step:

i. The strings are ordered by their lengths in a non-increasing order. If there are some strings of equal lengths, then the white ones precede the black ones. Then k first ones (if they consist of more than one pearl) are chosen; if there are less than k strings longer than 1, then one chooses all of them.

ii. Next, one cuts each chosen string into two parts differing in length by at most one. The process stops immediately after the step when a first isolated white pearl appears.

Prove that at this stage, there will still exist a string of at least two black pearls.

Olympiad Combinatorics

Pranav A. Sriram

August 2014

Copyright notices

About the Author

Pranav Sriram graduated from high school at The International School Bangalore, India, and will be a Freshman at Stanford University this Fall.

4. EXISTENCE

The devil's finest trick is to persuade you that he does not exist.
-Charles Baudelaire

Introduction

In this chapter, we focus on problems asking us to determine whether objects satisfying certain conditions exist. We encountered these types of problems in Chapter One, and solved them by creating algorithms that explicitly constructed the required objects. First of all, note that this approach does not give us any way to solve problems that ask us to prove that something does **not** exist. In addition, even when we want to prove existence, it may not always be possible to explicitly construct the required object. In these situations, we turn to less direct proof techniques, which are *existential* rather than *constructive*.

Some of the ideas in the first two chapters, such as induction, invariants and the extremal principle, can be adapted to provide non-constructive proofs. We will also introduce several new techniques in this chapter, including discrete continuity, divide and conquer strategies, the "hostile neighbors" trick, injective mappings and two very powerful variants of the extremal principle. A key theme that will pervade the examples in this chapter is the notion of proofs by contradiction, which the mathematician G. H. Hardy described as "one of a mathematician's finest weapons".

Induction

Our first example lies somewhere in between the inductive constructions of Chapter Two and the purely existential arguments of the rest of this chapter.

Example 1 [IMO Shortlist 1985]
A set of 1985 points is distributed around the circumference of a circle and each of the points is marked with 1 or -1. A point is called "good" if the partial sums that can be formed by starting at that point and proceeding around the circle for any distance in either direction are all strictly positive. Show that if the number of points marked with -1 is less than 662, there must be at least one good point.

Answer:
Note that 1985 = 3 x 661 + 2. This suggests that we try to show that for any n, if we have $3n$+2 points and at most n (-1)s, then there will be a good point. The result is true for n = 1. Assume it is true for k. Now we are given 3(k+1) + 2 points, of which at most (k+1) are (-1)s. Take a chain of consecutive (-1)s, having at least one (-1) and surrounded by two 1s. For example, (1, -1, -1, -1, -1, 1) or 1, -1, 1. Such a chain exists unless there are no (-1)s at any point, in which case we are trivially done. Now delete one (-1) from the chain as well as the bordering 1s. For example, 1, -1, -1, -1, 1 becomes -1, -1. Now we have $3k$+2 points and at most k (-1)s, so by induction there is a good point P. Note that P is obviously not part of the chain of (-1)s. Hence P is good in our original configuration as well, since after we add back the deleted points, each partial sum starting from P either doesn't change or increases by 1. ■

This was an example of "top-down" induction: we started from a configuration of 3(k + 1) + 2 points, then reduced it to a configuration of $3k$+2 points by *deleting* 3 points. We saw top down induction in the chapter on processes as well, where we broke down $n \times n$ boards into $(n-1) \times (n-1)$ boards with an extra row and column in order to use induction. On the other hand, if in the above example we had started with $3k+2$ points and then *added* 3 points, it would be "bottom up" induction. When applying induction, often one of the two approaches will work much better than the other. In the above example a top down approach works better since we can choose which points to delete, whereas in a bottom up approach we wouldn't be able to choose where to add the points (as this would lead to a loss of generality).

The next example uses a potent combination of induction, the pigeonhole principle and contradiction: we will essentially use the pigeonhole principle to inductively construct a contradiction.

Example 2 [IMO shortlist 1990]
Assume that the set of all positive integers is decomposed into r (disjoint) subsets $A_1 \cup A_2 \cup \ldots \cup A_r = \boldsymbol{N}$. Prove that one of them, say A_i has the following property: There exists a positive m such that for any k one can find numbers $a_1, a_2, \ldots, a_k$ in A_i with $0 < a_{j+1} - a_j \le m$ $(1 \le j \le k-1)$.

Answer:
Call a set with the given property *good*. Assume to the contrary that none of the sets is good. We will use this assumption to prove by induction that for each $s \le r$, $A_s \cup A_{s+1} \cup \ldots A_r$ contains arbitrarily long sequences of consecutive integers. For $s = r$ this will imply that A_r is good, contradicting our assumption.

A_1 is not good, so for every k there exist k consecutive numbers not in A_1. This means that $A_2 \cup A_3 \cup \ldots \cup A_r$ contains arbitrarily long sequences of consecutive integers. Now suppose we have shown that $A_s \cup A_{s+1} \cup \ldots A_r$ contains arbitrarily long sequences of consecutive integers. Since A_s is not good, for each m there exists a number k_m such that A_s doesn't contain a sequence of k_m integers with consecutive terms differing by at most m. Now take mk_m consecutive integers in $A_s \cup A_{s+1} \cup \ldots A_n$. If A_s contains fewer than

k_m of these numbers, then by the pigeonhole principle there are m consecutive numbers in $A_{s+1} \cup A_{s+2} \cup \dots A_n$, proving the inductive step. Otherwise, if A_s contains at least k_m of the numbers, by the definition of k_m some two of them differ by at least m. The m numbers in between then belong to $A_{s+1} \cup A_{s+2} \cup \dots A_n$. Since m is arbitrary, this proves the inductive step. By the first paragraph, the proof is complete. ■

The Extremal Principle and Infinite descent

The extremal principle basically says that any finite set of real numbers has a smallest and a largest element. "Extremal arguments" in combinatorics come in various forms, but the general idea is to look at objects that are extreme in some sense: smallest or largest numbers in a finite set, leftmost or rightmost points in combinatorial geometry problems, objects that are best or worst in a sense, etc. This provides a good starting point to solving complicated problems, since extremal objects are likely to satisfy certain restrictions and conditions that make them easy to analyze.

Example 3 [France 1997]
Each vertex of a regular 1997-gon is labeled with an integer such that the sum of the integers is 1. Starting at some vertex, we write down the labels of the vertices reading counterclockwise around the polygon. Is it always possible to choose the starting vertex so that the sum of the first k integers written down is positive for k = 1, 2, 3,…, 1997?

Some Intuition: Let the vertices be V_1, V_2, …, V_{1997} in anticlockwise order. Suppose we place V_1 at sea level, and for each V_i, define the altitude of V_i be equal to the altitude of V_{i-1} plus the number at V_i. Then, if we start at V_j and walk around the polygon, the sum of all the integers we have encountered is the net gain or loss in our altitude. Obviously, if we start at the **lowest** point, we can never have a net loss in altitude! In other words, the sum of numbers encountered can never be negative. (Note that this argument does not break down even when we cross from V_{1997} to V_1, because since the sum of all numbers is 1, the sum of encountered numbers is actually even more than the net altitude gain.) Below we convert this intuitive proof to a more formal proof.

Answer:
Yes. Starting from V_1, let the sum of the labels at the first k vertices in anticlockwise order be b_k. Let m be the minimum of all the b_k. Then take k such that $b_k = m$ (if there are many such k, take the largest such k). We claim that the vertex V_{k+1} satisfies the required conditions. Indeed, if the sum of the labels from V_{k+1} to V_j for some $j > k+1$ is negative, then the sum of the labels from V_1 to V_j is strictly less than m, since the sum from V_1 to V_j = sum from V_1 to V_k + (sum from V_k to V_j) = m + (a negative number). This contradicts the minimality of m. The only other case is if $j < k+1$, in which case we get a similar contradiction after using the fact that the sum of the 1997 labels is positive (since it is given to be 1). ■

Remark: Another intuitive interpretation of this solution is as follows: If we have had extremely bad luck until V_k, then by the "law of averages", we must have pretty good luck from there onwards.

The existence of extremal objects enables us to reach contradictions using a technique known as **infinite descent**, which you may have seen in the famous proof of the irrationality of $\sqrt{2}$. This technique works as follows: suppose we want to show that no number in a finite set $\boldsymbol{S}$ satisfies a certain property P. We assume to the contrary that some number in $\boldsymbol{S}$ does satisfy P, and use this assumption to show that there exists an even smaller number in $\boldsymbol{S}$ satisfying P.

This immediately yields a contradiction as follows: the argument shows that for any x_1 in $\boldsymbol{S}$ satisfying P, we can find x_2 in $\boldsymbol{S}$

satisfying P with $x_1 > x_2$. But repeating the same argument, we get a number x_3 in $\boldsymbol{S}$ satisfying P with $x_2 > x_3$. We can continue this indefinitely, to get numbers $x_1 > x_2 > x_3 > x_4 > \ldots$ all in $\boldsymbol{S}$ and satisfying P. But $\boldsymbol{S}$ is finite, so we cannot have infinitely many numbers in $\boldsymbol{S}$; contradiction.

In the next example, infinite descent will provide a key lemma needed to solve the problem.

Example 4 [Indian TST 2003]
Let n be a positive integer and {A, B, C} a partition of (1, 2, 3, …, 3n) such that |A| = |B| = |C| = n. Prove that there exist $x \in$ A, $y \in$ B, $z \in$ C such that one of x, y, z is the sum of the other two.

Answer:
Assume to the contrary that there exists a partition that does not have this property. WLOG suppose that 1 ϵ A. Let k be the smallest number not in A, and suppose WLOG that $k \in$ B. Hence 1, …, k – 1 are all in A and k is in B. Hence:
(i) No elements from C and A can differ by k
(ii) No elements from B and C can differ by less than k, since 1, 2, …, k-1 are in A. In particular no elements from B and C can differ by 1.

Let m be any element in C. By (ii), m-1 is not in B. What happens if m-1 is in C? First, m-k is not in A by (i). Further, m-k is not in B, since (m-1) – (m-k) = k-1, which is in A. So m-k must be in C. Also, m-k-1 is not in A, since (m-1) – (m-k-1) = k. By (ii), m-k-1 is not in B since m-k is in C. Hence m-k-1 is also in C.

Thus starting with any pair of consecutive numbers in C, namely (m, m-1) we get a smaller pair, namely (m-k, m-k-1). This leads to an infinite descent, which is a contradiction. Hence if m is in C, m-1 has to be in A. Hence we have an *injective* correspondence between elements in C and elements in A. This correspondence must also be *bijective* (otherwise |A| > |C|, but we are given that |A| = |C| = n). Thus if $t \in$ A, t+1 ϵ C. So 1 ∈ A implies 2

∈ C. This is a contradiction since we assumed that the smallest number not in A belongs to B. ■

Let us analyze the above proof. Essentially, we were given an abstract problem about sets and we simplified things by making a few assumptions. A natural starting point was to place 1 in one of the sets. Then using the extremal principle by assuming k was the least element not in A gave us some more structure. The infinite descent we used was almost accidental – even if we were not deliberately looking for an extremal argument, we were fortunate to find that given a pair of consecutive numbers in C, we could find a smaller such pair. These “accidental” infinite descents pop up very frequently in combinatorics, graph theory, number theory and algebra problems. So keep your eyes open while solving Olympiad problems – keep making observations, and you might just walk right into the solution!

Example 5 [ELMO Shortlist 2012]
Find all ordered pairs of positive integers (m, n) for which there exists a set $\boldsymbol{C} = \{c_1, c_2, \dots c_k\}$ $(k \geq 1)$ of colors and an assignment of colors to each of the mn unit squares of an $m \times n$ grid such that for every color c_i and unit square S of color c_i, exactly two direct (non-diagonal) neighbors of S have color c_i.

Answer:
If m and n are both even, then we can partition the board into 2×2 squares. Then color each 2×2 square with a different color. This clearly satisfies the problem’s conditions.

Now suppose at least one of m and n is odd. WLOG suppose the width of the board is odd. Consider a horizontal chain of squares of the same color in the top row of the board of **odd length**. Define a *good chain* as a chain of squares of the same color C of odd length. For example, if the width is 7 and the top row consists of colors c_1, c_1, $\boldsymbol{c_4}$, $\boldsymbol{c_4}$, $\boldsymbol{c_4}$, c_5, c_5 then the 3 c_4’s form a good chain. A good chain in the top row must exist since the width of the board is odd. Note that there can be no good chain of length 1 in the top row, since then that square will have at most one neighbor with the same color (the square below it).

Now look at the set of squares below the good chain in the top row. Let the good chain in the top row be x, c_i, c_i, … c_i, y where there are an odd number of c_i's flanked by two colors that are different from c_i (or by edges of the board). There are no squares above this chain. Thus there are **no squares of color c_i directly above, left or right of the chain**. The leftmost and rightmost c_i have only one neighboring square that is of color c_i; hence the squares below these two squares must have color c_i. The squares below the other squares in our chain of c_i's cannot have color c_i (since these squares of our chain already have exactly 2 neighbors with color c_i). Thus the set of squares below our row of c_i's must be of the form c_i, X, Y, Z, …, W, c_i where W, X, Y, Z stand for any colors other than c_i. An example is shown below.

$\boldsymbol{X}$	$\boldsymbol{c_i}$	$\boldsymbol{c_i}$	$\boldsymbol{c_i}$	$\boldsymbol{c_i}$	$\boldsymbol{c_i}$	$\boldsymbol{c_i}$	$\boldsymbol{c_i}$	$\boldsymbol{Y}$
	$\boldsymbol{c_i}$	c_5	$\boldsymbol{c_4}$	$\boldsymbol{c_4}$	$\boldsymbol{c_4}$	c_3	$\boldsymbol{c_i}$	

There are an odd number of squares between the two c_i's in the second row. Hence among these squares we can find a chain of odd length of squares having the same color c_k (different from c_i). Furthermore this chain is of **smaller length than our original chain**. Since all the squares above this chain are of color c_i, which is different from c_k, the new chain is bordered on 3 sides by squares not having color c_k, which is just like the first good chain. Hence we can repeat the above argument to obtain an even smaller good chain, obtaining a descent. We obtain smaller and smaller good chains until finally we get a good chain of length 1. This is a contradiction, because we would then have a single square bordered on three sides by squares of other colors, and it would hence have at most one neighbor of the same color. ■

Optimal Assumption

Assume first, ask questions later

Now we turn to an idea related to the extremal principle that I call the "optimal assumption" method. Suppose we want to find a set of size *at least* X satisfying some condition. Instead of constructing such a set using an algorithm, we merely prove its existence. We take the **largest** set $\boldsymbol{S}$ satisfying the required condition, and then **use the assumption** that $\boldsymbol{S}$ is as large as possible to prove that $|\boldsymbol{S}|$ must be at least X.

Here is a simple example to demonstrate this idea.

Example 6

In a graph G, suppose all vertices have degree at least δ. Show that there exists a path of length at least δ + 1.

Answer:

Take the longest possible path (optimal assumption) and let v be its last vertex. By the assumption that this is the longest possible path, we cannot extend the path any further. This means that all of v's neighbors must already lie in the path. But v has at least δ neighbors. Thus the path must contain at least δ + 1 vertices (v and all of its neighbors). ■

The next example shows the true power of this approach.

Example 7 [Italy TST 1999]

Let $\boldsymbol{X}$ be an n-element set and let $A_1, A_2, \ldots, A_m$ be subsets of $\boldsymbol{X}$ such that:

(i) $|A_i| = 3$ for $i = 1, 2, \ldots, m$

(ii) $|A_i \cap A_j| \leq 1$ for any two distinct indices i, j.

Show that there exists a subset of $\boldsymbol{X}$ with at least $\lfloor\sqrt{2n}\rfloor$ elements which does not contain any of the A_i's. (Note: Here $\lfloor\,\rfloor$ denotes the floor function).

Answer:
Call the elements of $\boldsymbol{X}$ $b_1, b_2, \ldots, b_n$. Let $\boldsymbol{S}$ be the largest subset of $\boldsymbol{X}$ not containing any of the A_i's. Let $|\boldsymbol{S}| = k$. We want to show that $k \geq \lfloor\sqrt{2n}\rfloor$. Now comes the crucial observation. For any element x in $\boldsymbol{X}$ but not in $\boldsymbol{S}$, there exists a pair of elements $\{y, z\}$ in $\boldsymbol{S}$ such that $\{x, y, z\} = A_i$ for some i. Otherwise we could add x to $\boldsymbol{S}$, and the new set would still not contain any set A_i, contradicting our assumption that $\boldsymbol{S}$ is the largest set satisfying this property.

Thus we can construct a mapping from elements in $\boldsymbol{X} \setminus \boldsymbol{S}$ to pairs of elements in $\boldsymbol{S}$ such that the element in $\boldsymbol{X} \setminus \boldsymbol{S}$ together with the pair of elements it is mapped to forms one of the sets A_i. Moreover, it cannot happen that two distinct elements in $\boldsymbol{X} \setminus \boldsymbol{S}$ are mapped to the same pair of elements. If this happened, say x_1 and x_2 were both mapped to $\{y, z\}$, then $\{x_1, y, z\} = A_i$ and $\{x_2, y, z\} = A_j$ for some i and j, and then $|A_i \cap A_j| = 2$. This violates condition 2 of the problem. Thus the mapping we have constructed is **injective**. This implies that the number of elements in $\boldsymbol{X} \setminus \boldsymbol{S}$ is cannot exceed the number of pairs of elements in $\boldsymbol{S}$. Hence we get $(n - k) \leq \binom{k}{2}$. This simplifies to $k^2 + k \geq 2n$, and from this the result easily follows (remember that k is an integer). ■

Example 8
Show that it is possible to partition the vertex set $\boldsymbol{V}$ of a graph $\boldsymbol{G}$ on n vertices into two sets $\boldsymbol{V}_1$ and $\boldsymbol{V}_2$ such that any vertex in $\boldsymbol{V}_1$ has at least as many neighbors in $\boldsymbol{V}_2$ as in $\boldsymbol{V}_1$, and any vertex in $\boldsymbol{V}_2$ has at least as many neighbors in $\boldsymbol{V}_1$ as in $\boldsymbol{V}_2$.

Answer:
What properties would such a partition have? Intuitively, such a partition would have lots of 'crossing edges', that is, edges joining a vertex in $\boldsymbol{V}_1$ to a vertex in $\boldsymbol{V}_2$. This suggests the following idea:

Take the partition *maximizing the number of crossing edges*. We claim that such a partition satisfies the problem conditions.

Suppose it doesn't. Suppose there is a vertex v in $\boldsymbol{V}_1$ that has more neighbors in $\boldsymbol{V}_1$ than in $\boldsymbol{V}_2$. Consider a new partition $\boldsymbol{V}_1' = \boldsymbol{V}_1 \setminus \{v\}$, $\boldsymbol{V}_2' = \boldsymbol{V}_2 \cup \{v\}$ (in other words, we have just moved v from $\boldsymbol{V}_1$ to $\boldsymbol{V}_2$). This has more crossing edges than the original partition by the assumption on v. This contradicts our earlier assumption that we took the partition maximizing the number of crossing edges. Hence the initial partition indeed works. ■

Remark 1: A partition of the vertices into two sets is known as a *cut.* The partition maximizing the number of crossing edges is known as a *max cut*.

Remark 2: The algorithmic problem of efficiently finding maximum or minimum cuts in general graphs is very difficult. Algorithms for finding approximate solutions to these and related problems have been extensively studied, and a rich combinatorial theory surrounding cuts, flows (the "dual" of a cut) and multicuts and multiway cuts (generalizations of cuts) has been developed. Several problems in this field remain open.

Invariants

(Again. Some things just don't change.)

Example 9 [Italy TST 1995]
An 8 × 8 board is tiled with 21 trominoes (3 × 1 tiles), so that exactly one square is not covered by a tromino. No two trominoes can overlap and no tromino can stick out of the board. Determine all possible positions of the square not covered by a tromino.

Answer:
The idea is to color the board in 3 colors, such that each tromino covers one square of each color. The figure shown below demonstrates such a coloring, where 1, 2, 3 denote 3 colors.

1	2	3	1	2	3	1	2
2	3	1	2	3	1	2	3
3	1	**<u>2</u>**	3	1	**<u>2</u>**	3	1
1	2	3	1	2	3	1	2
2	3	1	2	3	1	2	3
3	1	**<u>2</u>**	3	1	**<u>2</u>**	3	1
1	2	3	1	2	3	1	2
2	3	1	2	3	1	2	3

Figure 4.1: Coloring of the board

Since any tromino covers one square of each color, in total exactly 21 squares of each color will be covered. However, in the figure there are 22 2s, 21 1s and 21 3s. So the uncovered square would contain a 2. Now for my favorite part: symmetry. Suppose we take the initial coloring and create a new coloring by reflecting the board across its vertical axis of symmetry. For example, the top row of the board would now be colored 2, 1, 3, 2, 1, 3, 2, 1- the same coloring "backwards".

In the new coloring also, the uncovered square should be colored with the number 2. So the uncovered square should be colored by a 2 in both colorings. The only such squares are the ones underlined in the figure, since when one of these 2s is reflected in the vertical axis the image is on another 2.

Thus we have 4 possible positions of the uncovered square. To construct a tiling that works for these positions, first tile only the inner 4 × 4 square with one corner missing, and then tile the outer border. ■

The Hostile Neighbors Principle

(Yes, I made that name up)

Suppose we have n objects, A_1, A_2, …, A_n. Suppose some of these objects are of type one, and the rest are of type two. Further suppose that there is at least one object of each type. Then there exists an index i such that A_i and A_{i+1} are of opposite type. This statement is obvious, but as the next two examples demonstrate, it is surprisingly powerful.

Example 10 [Redei's theorem]

A *tournament* on n vertices is a directed graph such that for any two vertices u and v, there is either a directed edge from u to v or from v to u. Show that in any tournament on n vertices, there exists a (directed) Hamiltonian path.

(Note: a *Hamiltonian path* is a path passing through all the vertices. In other words we need to show that we can label the vertices v_1, v_2, …, v_n such that for each i, $1 \le i \le n$-1, there is a directed edge from v_i to v_{i+1}.)

Answer:

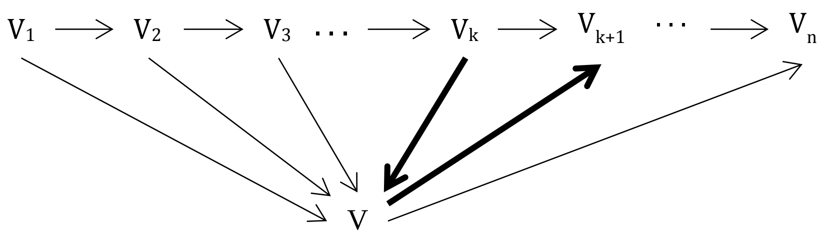

Figure 4.2: Illustration of how to extend the path to include V

We use induction on n, with the base cases n = 1, 2 and 3 being trivial. Suppose the result is true for n-1 vertices. Delete a vertex and form a Hamiltonian path with the remaining n-1 vertices. Let the path be $v_1 \rightarrow v_2 \rightarrow v_3 \rightarrow \ldots \rightarrow v_{n-1}$. Let the remaining vertex be v.

If $v_{n-1} \rightarrow v$, we are done, since we get the path $v_1 \rightarrow v_2 \rightarrow v_3 \rightarrow \ldots \rightarrow v_{n-1} \rightarrow v$. Similarly if $v \rightarrow v_1$ we are done. So suppose $v \rightarrow v_{n-1}$ and $v_1 \rightarrow v$. Hence there must be an index k such that $v_k \rightarrow v$ and $v \rightarrow v_{k+1}$. Then the path $v_1 \rightarrow v_2 \rightarrow \ldots v_k \rightarrow v \rightarrow v_{k+1} \ldots \rightarrow v_{n-1}$ is a Hamiltonian path and we are done. ■

The next example demonstrates the true power of this idea.

Example 11 [IMO shortlist 1988]

The numbers 1, 2, … , n^2 are written in the squares of an $n \times n$ board, with each number appearing exactly once. Prove that there exist two adjacent squares whose numbers differ by at least n.

Answer:

Assume to the contrary that there exists a labeling such that the numbers in any pair of adjacent squares differ by at most n-1.

Let $S_k = \{1, \ldots, k\}$ for each $k \geq 1$. Let $N_k = \{k+1, k+2, \ldots, k+n-1\}$. These are the numbers that can possibly neighbor a number in S_k. Let $T_k = \{k+n, k+n+1, \ldots, n^2\}$. No number from S_k can be next to a number from T_k.

For each k, since $|N_k| = n-1$, there exists a row that contains no element of N_k. Similarly there exists a column containing no element of N_k. The union of this row and this column must contain either **only elements from S_k or only elements from T_k,** otherwise some element of S_k would be next to an element of T_k. Call the union of this row and column a *cross.*

For $k = 1$, the cross cannot contain only elements from S_k (since there are $2n$-1 squares in the cross and only one element in S_1). Thus this cross contains only elements from T_k. But for $k = n^2-n$, the cross will contain only elements from S_k, as T_{n^2-n} has only one element. Hence from some j with $1 \leq j < n^2 - n$, the cross formed due to N_j will have elements only from T_k but the cross formed due to N_{j+1} will have elements only from S_{j+1}. But these crosses intersect at two squares. The numbers in these two squares belong to both S_{j+1} and T_j. This is a contradiction since $S_{j+1} \cap T_j = \emptyset$. ■

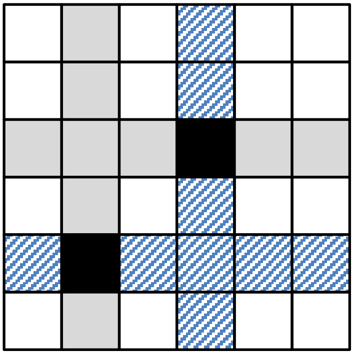

Figure 4.3: The black squares illustrate a contradiction as they cannot simultaneously belong to T_j and S_{j+1}.

Divide and Conquer

In the next example, we use the following idea: we are asked the minimum number of tiles needed to cover some set of squares. What we do is that we take a certain subset of these squares, such that no tile can cover more than one of these squares. Then clearly we need at least as many tiles as the number of squares in our subset, which gives us a good bound. This type of idea is frequently used in tiling problems as well as in other combinatorics problems asking for bounds of some sort.

Example 12 [IMO shortlist 2002, C2]
For n an odd positive integer, the unit squares of an $n \times n$ chessboard are colored alternately black and white, with the four corners colored black. An *L-tromino* is an L-shape formed by three connected unit squares. For which values of n is it possible to

cover all the black squares with non-overlapping L-trominoes? When it is possible, what is the minimum number of L-trominoes needed?

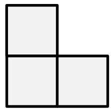

Figure 4.4: An L-tromino

Answer:

Let $n = 2k+1$. Consider the black squares at an odd height (that is, in rows 1, 3, 5, …, n). The key observation is that each L-tromino can cover at most one of these squares. There are $(k+1)^2$ such squares, so at least $(k+1)^2$ L-trominoes are needed. These L-trominoes cover a total of $3(k+1)^2$ squares. For n = 1, 3 or 5 this exceeds n^2 so we require $n \geq 7$ and at least $(k+1)^2 = \frac{(n+1)^2}{4}$ L-trominoes.

To construct tilings, induct with step 2. The base case 7 is left to the reader (do it systematically: start with the corners and then keep covering black squares of odd height). Given a $2k+1 \times 2k+1$ board, divide it into the top left $(2k-1) \times (2k-1)$ board along with a border of thickness 2. The $(2k-1) \times (2k-1)$ board can be tiled with k^2 tiles by induction. Now tile the border with $2k+1$ squares (this is left to the reader again). This shows that $(k+1)^2$ L-trominoes are sufficient, so the answer is $n \geq 7$ and $(k+1)^2$ L-trominoes are necessary and sufficient.

Discrete Continuity

The following example uses an idea known as discrete continuity that is very similar to the hostile neighbors principle. Discrete continuity is a very intuitive concept: basically, suppose in a sequence of integers, each pair of consecutive terms differ by at most 1. Then if a and b are members of the sequence, all integers between a and b will necessarily be members of the sequence. For instance, if 2 and 5 are in the sequence, then 3 and 4 must be as well. In particular, if we have such a sequence containing both a positive and a negative term, the sequence must contain 0 at some point. Such sequences where consecutive terms differ by at most one arise very often in combinatorics, and several problems can be solved by exploiting this "discrete continuity".

Example 13 [USAMO 2005, Problem 5]
Let n be an integer greater than 1. Suppose $2n$ points are given in the plane, no three of which are collinear. Suppose n of the given $2n$ points are colored blue and the other n colored red. A line in the plane is called a *balancing line* if it passes through one blue and one red point and, for each side of the line, the number of blue points on that side is equal to the number of red points on the same side.
Prove that there exist at least two balancing lines.

Answer:
Take the convex hull of the $2n$ points. If it contains points of both colors, then there will be two pairs of adjacent points in the hull that are of different colors. Take the two lines through these two pairs of points. There will be 0 points on one side and n-1 points of each color on the other side, so we are done. From now suppose the convex hull contains points of only 1 color, WLOG blue.

Take a point P that is part of this convex hull. Take a line L through P, such that all other points lie on the same side of L (this is possible since P lies on the convex hull). Now rotate L clockwise and let R_1, R_2, … R_n be the red points in the order in which they are encountered. Let b_i be the number of blue points encountered before R_i (excluding P) and r_i be the number of red points encountered before R_i (hence $r_i = i-1$). Let $f(i) = b_i - r_i$ and note that $f(i) = 0$ if and only if PR_i is a balancing line. Also $f(1) = b_1 - 0 \geq 0$ and $f(n) = b_n - (n-1) \leq 0$, since b_n is at most n-1. Thus $f(i)$ goes from nonnegative to nonpositive as i goes from 1 to n.

Furthermore, f can decrease by at most 1 when going from i to i+1, since r_i increases by only 1. Hence at some point f becomes 0, and we get a balancing line through P.

Repeating this argument for each point on the convex hull, we get balancing lines for each point on the convex hull, so we get at least 3 balancing lines in this case (there are at least 3 points on the convex hull), so we are done. ■

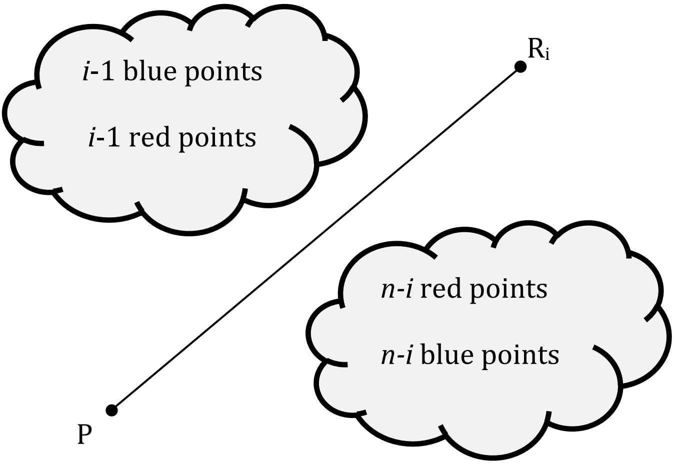

Figure 4.5: A balancing line for f(i) = 0

Miscellaneous Examples

(Because I ran out of imaginative names)

Example 14 [Romania 2001]

Three schools each have 200 students. Every student has at least one friend in each school (friendship is assumed to be mutual and no one is his own friend). Suppose there exists a set $\boldsymbol{E}$ of 300 students with the following property: for any school $\boldsymbol{S}$ and two students $x, y \in \boldsymbol{E}$ who are not in school $\boldsymbol{S}$, x and y do not have the same number of friends in $\boldsymbol{S}$. Prove that there exist 3 students, one in each school, such that any two are friends with each other.

Answer:
Let S_1, S_2, S_3 be the sets of students in the three schools. Since there are 300 students in $\boldsymbol{E}$, one of the schools must have at most 300/3 = 100 students in $\boldsymbol{E}$. WLOG let $|S_1 \cap \boldsymbol{E}| \leq 100$. Then consider the 200 or more students in $\boldsymbol{E} \setminus S_1$. Each of these students has at least one at and most 200 friends in S_1, and moreover no two of them have the same number of friends in S_1 (by the conditions of the problem and the condition on $\boldsymbol{E}$). This implies that exactly one of them has 200 friends in S_1. Let this student by X, and assume WLOG that X is in S_2. Then X has a friend Y in S_3 and Y has a friend Z in S_1 (everyone has at least one friend in each school). But Z and X are friends since Z is friends with everyone in S_1. So (X, Y, Z) is our required triple and we are done. ■

Example 15 [IMO shortlist 1988]
Let n be an even positive integer. Let $A_1, A_2, \ldots, A_{n+1}$ be sets having n elements each such that any two of them have exactly one element in common while every element of their union belongs to at least two of the given sets. For which n can one assign to every element of the union one of the numbers 0 and 1 in such a manner that each of the sets has exactly $n/2$ zeros?

Answer:
Let $n = 2k$. Observe that any set A_j has $2k$ elements and intersects each of the other $2k$ sets in exactly one element. Hence each of the $2k$ elements in A_j belongs to at least one of the other $2k$ sets but each of the $2k$ sets contains at most one element from A_j. This implies that each of the $2k$ elements of A_j belongs to *exactly one* other set. This holds for each j, so every element in the union of the sets belongs to exactly two sets and any two sets intersect in exactly one element.

Now suppose we count the number of elements labeled 0. Each set contains k zeros and there are $2k$+1 sets. But each element labeled 0 is in two sets, and if we simply multiplied k and $2k$+1 we would be counting each element twice. So the total number of elements labeled 0 will be $k(2k+1)/2$. This quantity must be an integer, so k must be divisible by 2. Hence n must be divisible by 4.

To show that such a coloring indeed exists when n is divisible by 4, incidence matrices provide an elegant construction. Incidence matrices will be introduced in the chapter on counting in two ways, and the rest of the proof of this example is left as an exercise in that chapter. ■

Example 16 [IMO shortlist 2001 C5]
Find all finite sequences $x_0, x_1, x_2, \dots, x_n$ such that for every j, $0 \le j \le n$, x_j equals the number of times j appears in the sequence.

Answer:
The terms of such a sequence are obviously nonnegative integers. Clearly $x_0 > 0$, otherwise we get a contradiction. Suppose there are m nonzero terms in the sequence. Observe that the sum $x_1 + x_2 \dots + x_n$ counts the total number of nonzero terms in the sequence; hence $x_1 + \dots + x_n = m$. One of the nonzero terms is x_0, so there are exactly m-1 nonzero terms among $x_1, x_2, \dots, x_n$. These m-1 nonzero terms add up to m, so m-2 of these terms are equal to 1 and one term is equal to 2. This means that no term of the sequence is greater than two, except possibly x_0. Hence at most one of $x_3, x_4, \dots$ can be positive (For example, if $x_0 = 4$, then x_4 will be positive since 4 appears in the sequence). Thus the only terms that can be positive are x_0, x_1, x_2 and at most one x_k with $k > 2$. It follows that $m \le 4$. Also $m = 1$ is impossible. So we have 3 cases:

(i) $m = 2$. Then there are m-2 = 0 1s and one 2 among the terms $x_1, x_2, \dots x_n$. Hence $x_2 = 2$ (as $x_1 = 2$ is impossible) and the sequence is (2, 0, 2, 0).

(ii) $m = 3$. Either $x_1 = 2$ or $x_2 = 2$. These cases give the sequences (1, 2, 1, 0) and (2, 1, 2, 0, 0) respectively.

(iii) $m = 4$. Then the positive terms are x_0, x_1, x_2 and x_k for some $k > 2$. Then $x_0 = k$ and $x_k = 1$. There are m-2 = 2 1s so $x_1 = 2$, and hence $x_2 = 1$. The final sequence is (k, 2, 1, 0, …., 0, 1, 0, 0, 0), where there are k 0s between the two 1s.

Hence the sequences listed in (i), (ii) and (iii) are the only possible sequences and we're done. ■

Exercises

1. **[Russia 2001]**
Yura put 2001 coins of 1, 2 or 3 kopeykas in a row. It turned out that between any two 1-kopeyka coins there is at least one coin; between any two 2-kopeykas coins there are at least two coins; and between any two 3-kopeykas coins there are at least 3 coins. Let n be the number of 3-kopeyka coins in this row. Determine all possible values of n.

2. **[Indian TST 2001]**
Given that there are 168 primes between 1 and 1000, show that there exist 1000 consecutive numbers containing exactly 100 primes.

3. **[Canada 1992]**
$2n+1$ cards consists of a joker and, for each number between 1 and n inclusive, two cards marked with that number. The $2n+1$ cards are placed in a row, with the joker in the middle. For each k with $1 \le k \le n$, the two cards numbered k have exactly $(k-1)$ cards between them. Determine all the values of n not exceeding 10 for which this arrangement is possible. For which values of n is it impossible?

4. **[IMO 1997-4]**
An $n \times n$ matrix whose entries come from the set $\boldsymbol{S} = \{1, 2, \ldots, n\}$ is called a *silver matrix* if, for each $i = 1, 2, \ldots, n$, the i-th row and the i-th column together contain all elements of $\boldsymbol{S}$. Show that:

a) there is no silver matrix for $n = 1997$;
b) silver matrices exist for infinitely many values of n.

5. [Russia 1996]
Can a 5 × 7 board be tiled by *L-trominoes* (shown in the figure below) with overlaps such that no L-tromino sticks out of the board, and each square of the board is covered by the same number of L-trominoes?

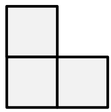

An L-tromino

6. IMO Shortlist 2011, C2]
Suppose that 1000 students are standing in a circle. Prove that there exists an integer k with $100 \le k \le 300$ such that in this circle there exists a contiguous group of $2k$ students, for which the first half contains the same number of girls as the second half.

7. [Bulgaria 2001]
Let n be a given integer greater than 1. At each lattice point (i, j) we write the number k in $\{0, 1, \ldots, n\text{-}1\}$ such that $k \equiv (i+j) \bmod n$. Find all pairs of positive integers (a, b) such that the rectangle with vertices $(0,0)$, $(a, 0)$, (a, b) and $(0, b)$ has the following properties:
(i) Each number 0, 1, …, n-1 appears in its interior an equal number of times
(ii) Each of these numbers appear on the boundary an equal number of times

8. [Russia 1998]
Each square of a board contains either 1 or -1. Such an arrangement is called *successful* if each number is the product of its neighbors. Find the number of successful arrangements.

9. [IMO Shortlist 2010, C3]
2500 chess kings have to be placed on a 100×100 chessboard so that

i. no king can capture any other one (i.e. no two kings are placed in two squares sharing a common vertex);
ii. each row and each column contains exactly 25 kings.

Find the number of such arrangements. (Two arrangements differing by rotation or symmetry are supposed to be different.)

10. [Russia 2011]
There are some counters in some cells of 100×100 board. Call a cell *nice* if there are an even number of counters in adjacent cells. Is it possible for there to exist exactly one nice cell?

11. [Bulgaria 1997]
A triangulation of a convex n-gon is a division of the n-gon into triangles by diagonals with disjoint interiors. Call a triangulation even if each vertex of the n-gon is the endpoint of an even number of diagonals. Determine all natural numbers n for which an even triangulation of an n-gon exists.

12. [India Postal Coaching 2011]
On a circle there are n red and n blue arcs given in such a way that each red arc intersects each blue one. Prove that there exists a point contained by at least n of the given colored arcs.

13. Call a rectangle *integral* if at least one of its dimensions is an integer. Let R be a rectangle such that there exists a tiling of R with smaller integral rectangles with sides parallel to the sides of R. Show that R is also integral.

14. [IMO Shortlist 1999, C6]
Suppose that every integer has been given one of the colors red, blue, green or yellow. Let x and y be odd integers so that $|x| \neq |y|$. Show that there are two integers of the same color

whose difference has one of the following values: x, y, $(x+y)$ or $(x-y)$.

15. [China TST 2011]
Let l be a positive integer, and let m, n be positive integers with $m \geq n$, such that $A_1, A_2, \ldots, A_m, B_1, B_2, \ldots, B_n$ are $(m+n)$ pairwise distinct subsets of the set $\{1, 2, \ldots, l\}$. It is known that $A_i \Delta B_j$ are pairwise distinct, for each $1 \leq i \leq m$, $1 \leq j \leq n$, and run over all nonempty subsets of $\{1, 2, \ldots, l\}$. Find all possible values of (m, n).

16. [IMO 1996, Problem 6]
Let p, q, n be three positive integers with $p+q < n$. Let $(x_0, x_1, \ldots, x_n)$ be an $(n+1)$-tuple of integers satisfying the following conditions: :
(a) $x_0 = x_n = 0$, and
(b) For each i with $1 \leq i \leq n$, either $x_i - x_{i-1} = p$ or $x_i - x_{i-1} = -q$.

Show that there exist indices $i < j$ with $(i, j) \neq (0, n)$ such that $x_i = x_j$.

17. [IMO 2004, Problem 3]
Define a *hook* to be a figure made up of six unit squares as shown below in the picture, or any of the figures obtained by applying rotations and reflections to this figure. Determine all $m \times n$ rectangles that can be tiled with hooks without gaps, overlaps, and parts of any hook lying outside the rectangle.

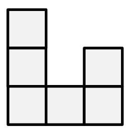

A hook

Olympiad Combinatorics

Pranav A. Sriram

August 2014

Copyright notices

About the Author

Pranav Sriram graduated from high school at The International School Bangalore, India, and will be a Freshman at Stanford University this Fall.

5. COMBINATORIAL GAMES

I'm afraid that sometimes,
you'll play lonely games too.
Games you can't win
'cause you'll play against you.

-Doctor Seuss

Introduction

In this chapter, we study combinatorial games involving two players. Typical problems describe a game, and ask us to find a winning strategy for one of the players or determine whether one exists.

As in problems in the chapters on algorithms and processes, games are specified by a starting position, allowed move(s) and the final or winning position. These and other attributes make problems in this chapter seem superficially similar to those in the first three chapters, but a fundamental difference is that our algorithms or strategies now need to compete against those of an opponent. Some of the techniques we will develop in this chapter are hence significantly different from earlier ones. One simple yet surprisingly powerful technique we use is exploiting *symmetry*.

Others include coloring and invariants, recursion, induction, parity and a very important technique we will introduce known as positional analysis.

The term "combinatorial games" is generally used to describe games with the following characteristics:

(i) There are no elements of luck or chance (so no dice rolling or coin flipping involved).
(ii) There is usually perfect information, unlike in card games, where I cannot see your hand (unless I peak).
(iii) Players typically move alternately. In this chapter, all our games are played by Alice and Bob, whom we refer to as A and B respectively. In general A starts.
(iv) There is no cheating (unfortunately).

One more thing we will use frequently is that if a game is finite and does not have ties, then someone has to eventually lose. Because of the perfect information property and the absence of probabilistic elements, one player will hence have a winning strategy. In particular, one concept we will use frequently to develop strategies is that if one player can ensure that he always "stays alive", the other player will eventually lose.

Symmetry, Pairing and Copying

They said no cheating, but we can still copy!

This section is devoted to the very important technique of "copycat strategies" and pairing techniques. The next few examples should illustrate what these techniques are about.

Example 1
A and **B** each get an unlimited supply of identical circular coins. **A**

and **B** take turns placing the coins on a finite square table, in such a way that no two coins overlap and each coin is completely on the table (that is, it doesn't stick out). The person who cannot legally place a coin loses. Assuming at least one coin can fit on the table, prove that **A** has a winning strategy.

Answer:
A first places a coin such that its center is at the center of the table. Then whenever **B** places a coin with center at a point X, **A** places a coin with center at the point X', where X' is the reflection of X in the center of the table. This ensures that after each of **A**'s moves, the board is completely symmetrical. Thus if **B** can make a legal move, then by symmetry, **A**'s next move is also legal. Since the area of the table is finite, eventually the game must terminate and someone must lose. Since **A** can always "stay alive", **B** loses. ■

Example 2 [Saint Petersburg 1997]
The number N is the product of k different primes ($k \geq 3$). **A** and **B** take turns writing composite divisors of N on a board, according to the following rules. One may not write N. Also, there may never appear two coprime numbers or two numbers, one of which divides the other. The first player unable to move loses. If **A** starts, who has the winning strategy?

Answer:
A has a winning strategy. **A** first writes pq for some primes p and q dividing N. Then all the subsequent numbers written must be of the form pm or qm for some m dividing N, by the conditions of the problem. Whenever **B** writes qm, **A** writes qn. This "copying strategy" ensures that **A** always has a move. Since the game is finite (N has a finite number of divisors), **A** will eventually win. ■

Example 3 [USAMO 2004 – 4]
Alice and Bob play a game on a 6 × 6 grid. On his or her turn, a player chooses a rational number not yet appearing in the grid and writes it in an empty square of the grid. Alice goes first and then the players alternate. When all of the squares have numbers

written in them, in each row, the square with the greatest number in that row is colored black. Alice wins if she can then draw a path from the top of the grid to the bottom of the grid that stays in black squares, and Bob wins if she can't. (A path is a sequence of squares such that any two consecutive squares in the path share a vertex). Find, with proof, a winning strategy for one of the players.

Answer:

X	X	X			
X	X	X			
X	X			X	
			X	X	X
			X	X	X
			X	X	X

Figure 5.1

B has a winning strategy. The idea is to ensure that the largest number in each row is in one of the squares marked X in Figure 5.1. Then clearly there will be no path and **B** will win.

To do this, **B** pairs each square marked with an X with a square not marked with an X in the same row. Whenever **A** plays in a square marked with an X, **B** writes a smaller number in the paired square. Whenever **A** writes a number in a square not having an X, **B** writes a larger number in the paired square. This ensures that after each of **B**'s moves, the largest number in each row is in one of the squares marked with an X. ■

Example 4 [Tic Tac Toe for mathematicians]
On a 5 × 5 board, **A** and **B** take turns marking squares. **A** always writes an X in a square and **B** always writes O. No square can be marked twice. **A** wins if she can make one full row, column or diagonal contain only Xs. Can **B** prevent **A** from winning?

Answer:
Yes. Mark the board as shown.

5	9	11	9	6
8	1	1	2	7
12	4		2	12
8	4	3	3	7
6	10	11	10	5

Figure 5.2: Squares with equal numbers are paired

Since each number from 1 to 12 appears in 2 squares, **B** can ensure that he always marks at least one square of each number (if **A** marks an X in a square with number i, **B** puts an O in the other square marked i). Observe that each row, column and diagonal contains a pair of squares having equal numbers. Since **B** ensures that at least one number in each pair is marked, each row, column and diagonal would have been marked by **B** at least once. ■

Remark 1: This is one of those solutions that initially appears to come out of nowhere, but actually has a very intuitive explanation. Here's one way of thinking about it. There are 5 columns, 5 rows and 2 diagonals, for a total of 12 "winning lines". We can only make 12 pairs of numbers (since there are 25 squares and 12 x 2 = 24). Thus our idea is to construct a pairing strategy such that each row, column and diagonal is covered by exactly one pair. In this example, a natural way to construct the pairing is to start with the inner 3 x 3 square, covering its border, and then fill the outer layer, ensuring that the remaining winning lines are taken care of.

Remark 2: The game of tic-tac-toe can be generalized to getting k squares in a row on an $m \times n$ board (regular tic-tac-toe has $k = m = n = 3$ and this problem has $k = m = n = 5$). For most values of k, m and n, the question of whether there exists a winning strategy for the starting player remains an open problem. In addition, the game of tic-tac-toe in higher dimensions (such as on 3-D cubes or in general, n-dimensional hyperspaces) has strong connections to a branch of extremal combinatorics known as Ramsey theory.

Example 5 [IMO Shortlist 1994, C1]
A and **B** play alternately on a 5 × 5 board. **A** always enters a 1 into an empty square, and **B** always enters a 0 into an empty square. When the board is full, the sum of the numbers in each of the nine 3 × 3 squares is calculated and **A**'s score S is the largest such sum. What is the largest score **A** can make, regardless of the responses of **B**?

Answer:

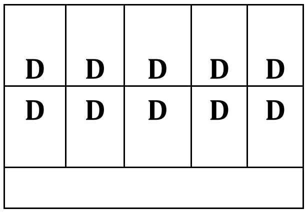

Figure 5.3: A pairing strategy for B

Tile the board with 10 dominoes (marked **D** in the figure), leaving the bottom row untiled. **B** can ensure that each domino contains at least one 0, since whenever **A** plays in a domino **B** writes a 0 in the other square. Each 3 × 3 square has at least 3 full dominoes, and hence will have at least three 0s. Thus **B** can ensure that S is at most 6. Now we leave it to the reader to show that **A** can ensure that S is at least 6. ■

Parity Based Pairing

We solved the last few examples by explicitly constructing pairing strategies. But all we actually need to do sometimes is prove the *existence* of a pairing strategy. This is often remarkably easy- observe that if there are an even number of objects, there exists a way to pair them up.

Example 6 [Based on Italy TST 2009, problem 6]
A and **B** play the following game. First **A** writes a permutation of the numbers from 1 to *n*, where *n* is a fixed positive integer greater than 1. In each player's turn, he or she must write a sequence of numbers that has not been written yet such that either:

a) The sequence is a permutation of the sequence the previous player wrote, OR
b) The sequence is obtained by deleting one number from the previous player's sequence

For example, if **A** first writes 4123, **B** could write 3124 or 413. The player who cannot write down a sequence loses. Determine who has the winning strategy.

Answer:
If $n = 2$, **B** wins: after **A**'s first move, **B** deletes the number 2 and is left with the sequence {1}. Then **A** has no move.

The idea now is to construct an inductive strategy for **B**. Suppose **B** wins for $n = k$; we now want a strategy for $n = k+1$. **B**'s aim is to make **A** be the first player to delete a number from the sequence. Then from this point the game is reduced to a game with k numbers, and **B** will win this by induction. But this is very easy to do – whenever **A** writes a sequence of k+1 numbers, there will always exist at least one permutation of the k+1 numbers that has not been written yet, simply because the total number of permutations is even ((k+1)! is even). ■

Remark: You can also use an explicit pairing strategy. One example is that whenever **A** writes a sequence, **B** writes the same sequence backwards.

Example 7 [USAMO 1999-5]
The *Y2K Game* is played on a 1 × 2000 grid as follows. Two players in turn write either an S or an O in an empty square. The first player who produces three consecutive boxes that spell SOS wins. If all boxes are filled without producing SOS then the game is a draw. Prove that the second player has a winning strategy.

Answer:
Call an empty square *bad* if playing in that square will let the other player form SOS in the next turn. We claim that an empty square is bad **if and only if** it is in a block of 4 squares of the form S_ _ S.

It is easy to see that both these empty squares are bad, as playing either S or O will allow the other player to form SOS. Conversely, if a square is bad, then playing an O in it will allow the other player to win, so it must have an S next to it and an empty square on the other side. Also, playing an S in a bad square allows the other player to win, so there must be another S beyond the empty square. This forces the configuration to be S_ _S, proving our claim.

Now after **A**'s first move, **B** writes an S at least 4 squares away from either end of the grid and **A**'s first move. On **B**'s second move, he writes S three squares away from his first S so that the two squares in between are empty. These two squares are bad. Note that at any point in the game there will always be an even number of bad squares (since they come in pairs, by our claim above). So whenever it is **B**'s turn, an odd number of moves would have been made, so an odd number of squares would be empty, of which an even number would be bad. Hence there will always be at least one square that is not bad on **B**'s turn, so **B** won't lose. Eventually the game will end since there are at least 2 bad squares, so **B** must win. ■

Positional Analysis

Heads I win, tails you lose.

For the rest chapter, we will use the following convention: a position in a game is a ***P-position*** if the player who has just played can force a win (that is, if he has a winning strategy). A position is called an ***N-position*** if the player whose turn it is can force a win. P and N respectively stand for Player and Next player.

This means that the starting player has a winning strategy if and only if the initial position of a game is an N-position (he is the "next player" as it is his turn at the start of the game). The second player has a winning strategy if and only if the initial position is a P-position (even though the game has not yet begun, by convention, he is the player who has "just played", since it is not his turn).

Our definition of P- and N-positions also implies the following: From an N-position, the player whose turn it is can move into a P-position. In other words, the player who has a winning strategy can move to a position in which he still has a winning strategy. From a P-position, the player whose turn it is must move into an N-position. **The winning strategy for a player is to always move to P-positions.**

Example 8

A and **B** play a game as follows. First **A** says 1, 2 or 3. Then **B** can add 1, 2 or 3 to the number the first player said. The game continues with the players playing alternately, in each turn adding 1, 2 or 3 to the previous number. For example, **A** can say 2, then **B** can say 5, then **A** could say 6, and so on. The player who says 100 wins. Who has the winning strategy?

Answer:
Observe that **B** can always say a multiple of 4 in his turn. For example, consider the following sequence of moves: **A**-1; **B**-4; **A**-6; **B**-8 and so on. Regardless of what **A** says, **B can always say a multiple of 4**. This is fairly obvious, but if you want to be more rigorous you can prove it by induction – if **B** says $4k$ then **A** says $4k+1$, $4k+2$ or $4k+3$, and in all cases **B** can say $4k+4$. Hence **B** will say 100 and will win. ■

Remark 1: Let's analyze this proof. First of all, how would one come up with it? One idea is to work backwards. Clearly the player who says 100 wins. Hence 100 is a P-position – the person who has just played wins. Then 99, 98 and 97 are N-positions as the next player can reach 100 from these positions. But 96 is a P-position, since from 96 only N-positions can be reached (97, 98 and 99). Continuing in this manner, we see that every multiple of 4 is a P-position, so the winning strategy is to always play multiples of 4. This type of analysis of P- and N-positions will be the central idea in the rest of this chapter.

Remark 2: This game can be generalized- instead of 1, 2, or 3 we can allow a player to increase the number by 1, 2, …, or k for any positive integer k and we can replace 100 by n. If $k+1$ divides n, then all multiples of $k+1$ are P-positions, so 0 is a P-position and **B** has a winning strategy. Otherwise, suppose $n \equiv r \bmod (k+1)$. Then all numbers that are congruent to r mod $k+1$ are P-positions, and **A** has a winning strategy by saying r in her first move.

Note: In several problems, we will use arguments of the form "If one player does this, then the other player does that, then the first player does this…". In order to avoid referring to the players as "one" or "the other" or the particularly ambiguous "the first player", we will use X and Y in these situations. X could refer to either of the players A and B, and Y refers to the other one. (On the other hand this precludes the usage of personal pronouns… writing a book is just frustrating sometimes.)

Example 9 [Lithuania 2010]

In an $m \times n$ rectangular chessboard there is a stone in the lower leftmost square. **A** and **B** move the stone alternately, starting with **A**. In each step one can move the stone upward or rightward any number of squares. The player who moves it into the upper rightmost square wins. Find all (m, n) such that **A** has a winning strategy.

Answer:

For convenience, lets flip the board upside down so that the stone starts in the upper rightmost square and the player who moves it to the lower leftmost square wins. All moves are now down or left. Now we label the squares by their coordinates, so that the lower leftmost square is $(1, 1)$ and the upper rightmost square is (m, n). Note that the game must terminate – this is a key ingredient for positional analysis to work.

						P		Start
					P			
				P				
			P					
		P						
	P							
P								

Figure 5.4: The P-positions

The crucial observation is that the P-positions are squares with **equal coordinates**, that is, squares with coordinates (k, k) for some k. These are marked **P** in the figure. This is because if X moves to a square marked P, Y cannot reach any other P square in Y's turn. Then X can move back to a P-position. This continues – X keeps moving to P-positions, and eventually reaches $(1, 1)$ and wins.

Thus if $m \neq n$, the initial position is N, so **A** wins by always moving to a square marked P. If $m = n$, the initial position is P, so **B** wins. Hence **A** wins if and only if $m \neq n$. ■

The previous two examples illustrate the general approach to problems in this section: characterize the N- and P-positions. If the starting position is an N-position, **A** wins; otherwise **B** wins. The next example is similar to example 8, but the inductive proof for characterizing N-positions is slightly trickier. There is also an important difference between this example and the previous two, discussed in the remark at the end.

Example 10 [Saint Petersburg 2001]
The number 1,000,000 is written on a board. **A** and **B** take turns, each turn consisting of replacing the number n on the board with $n - 1$ or $\lfloor (n+1)/2 \rfloor$. The player who writes the number 1 wins. Who has the winning strategy?

Answer:
Note that the game eventually ends, so one of the players must have a winning strategy. After experimenting with 1,000,000 replaced with smaller values, we see that **A** wins when the starting number is 2, 4, 6 and 8 and conjecture that **A** wins for any even starting number. In other words, we claim that all even numbers are N-positions. We prove this by induction.

2 is an N-position, so now suppose that all even numbers less than $2k$ are N-positions. We show that $2k$ is also an N-position. If k is a P-position, **A** can write k in her first move and win. Otherwise **A** first writes $2k - 1$. Then **B** must write $2k - 2$ or k, both of which are N-positions by assumption, so **A** wins in this case as well. Thus all even numbers are N-positions, so 1,000,000 is an N-position and **A** has a winning strategy. ■

Remark: Note that unlike the solutions in examples 8 and 9, this solution does not characterize *all* positions as P or N. We proved all even numbers are N-positions, but didn't prove anything about odd numbers since we didn't need to. All we really cared about was the number 1,000,000, so characterizing the positions of even numbers was sufficient. Another interesting thing to observe is the following: suppose N = 50. **A**'s strategy says to write 25 **if 25 is a P-position**, and otherwise write 49. But **we don't know whether 25 is a P-position or an N-position!** Does that mean our solution is incomplete or wrong? No. The problem only asked *who* has the winning strategy, not *what* the winning strategy is. We have guaranteed *the existence* of a winning strategy for **A**, without explicitly finding all its details – that's Alice's problem! (Well, actually, it's kind of your problem too- exercise 6 asks you to characterize the odd positions of this problem.)

An idea that has started to recur in this book is to take the techniques used in constructive proofs and use them for existential proofs. The next proof is a beautiful example of this idea, in which we use positional analysis to prove the existence of a winning strategy for one player by contradiction.

Example 11 [Russia 2011, adapted]

There are $N > n^2$ stones on a table. **A** and **B** play a game. **A** begins, and then they alternate. In each turn a player can remove k stones, where k is a positive integer that is either less than n or a multiple of n. The player who takes the last stone wins. Prove that **A** has a winning strategy.

Answer:

The game is finite and deterministic so some player must have a winning strategy. Suppose to the contrary that **B** can win. **A** first removes kn stones. If **B** removes jn stones for some j, then that means $N - (k + j)\,n$ is a P-position (since we are assuming **B** is playing his winning strategy). In this case **A** could have removed $(k + j)\,n$ stones in her first move and won, meaning that **A** has a winning strategy, which is a contradiction.

Thus if **A** removes kn stones in the first move, where $1 \le k \le n$, let $f(k)$ denote the number of stones **B** takes in response according to his winning strategy. By the first paragraph, $1 \le f(k) \le n - 1$. Hence by the pigeonhole principle, for some distinct k and j, $f(k) = f(j)$. This means that both $N - kn - f(k)$ and $N - jn - f(k)$ are P-positions, since these are the positions that arise after **B**'s move.

Now WLOG let $k < j$. **A** first removes kn stones, then **B** removes $f(k)$ stones. Now **A** removes $(k - j)n$ stones. There are now $N - kn - f(k)$ stones remaining, which is a P-position by the second paragraph. Hence **A** will win, contradicting our assumption that **B** has a winning strategy. ■

In slightly harder problems recursion often proves to be a useful technique. This means that we relate the outcome of a game of size N to games of smaller size.

Example 12 [IMO Shortlist 2004, C5]
A and **B** take turns writing a number as follows. Let N be a fixed positive integer. First **A** writes the number 1, and then **B** writes 2. Hereafter, in each move, if the current number is k, then the player whose turn it is can either write $k + 1$ or $2k$, but no player can write a number larger than N. The player who writes N wins. For each N, determine who has a winning strategy.

Answer:
Step 1: We quickly observe that if N is odd, **A** can win. **A** can always ensure that she writes an odd number, after which **B** would have to write an even number. Hence **B** cannot say N so **A** wins. Now suppose N is even.

Step 2: The key observation is that all **even numbers greater than $N/2$ are P-positions**. This is because after this point **neither player can double the number** (otherwise it will exceed N). Hence they both must keep adding 1 in their turns, and one player will keep writing even numbers and the other will keep writing odd numbers. The player who wrote the even number greater than $N/2$ will hence write N since N is even.

Step 3: **If $N = 4k$ or $N = 4k+2$, then k is a P-position**. This is because if X writes k, Y must write $k+1$ or $2k$. Then X writes $2k+2$ if Y writes k and X writes $4k$ if Y writes $2k$. X has thus written an even number greater than $N/2$ and by step 2, X wins.

Steps 2 and 3 now give us the final critical lemma.

Step 4: **If X has a winning strategy for $N = k$, then X has a winning strategy for $N = 4k$ and $N = 4k+2$**.
Proof: Consider a game where $N = 4k$ or $4k + 2$. If X writes k at some point during the game, we are done by step 3. So X's aim is to write k, so X starts implementing the winning strategy for $N = k$. How can Y prevent X from writing k? By "jumping over" k at some point: after X says some number j with $k/2 < j < k$, Y doubles it, resulting in a number 2j with $(k + 1) \leq 2j \leq N/2$. But then X simply doubles this number, resulting in an even number at least equal to $2(k+1) > N/2$. So X wins by step 2.

Finally, we have a recursive method of determining the answer for even N. The answer for N is the same as that for $\lfloor N/4 \rfloor$. To convert this recursion into an explicit answer, write N in base 4. The function $\lfloor N/4 \rfloor$ is equivalent to removing the last digit of N in base 4. Starting from the base 4 representation of N, keep removing the rightmost digit. The resulting numbers will all be winning for the same player by our recursion. If at some point we obtain an odd number, then **A** wins for this number and hence **A** wins for N. Hence if N has an odd digit in base 4, then **A** wins. Otherwise, suppose **N** has only 0s and 2s in its base 4 representation. Then applying our procedure we eventually end up with the number 2, and since **B** wins for 2, **B** wins for N in this case. ■

The final example of this chapter is different in a few ways. First, it is asymmetrical, in the sense that the two players have different (in fact, opposite) objectives. Such games are typically called "maker-breaker" games. We cannot define P- and N-positions the way we did before. However, similar ideas of analyzing positions based on outcome still apply, and winning strategies are still based on always sticking to some particular type of position. These positions are typically characterized by some specific property or invariant, as the next example shows.

Example 13 [IMO Shortlist 2009, C5]
Five identical empty buckets of 2-liter capacity stand at the vertices of a regular pentagon. Cinderella and her wicked Stepmother go through a sequence of rounds: At the beginning of every round the Stepmother takes one liter of water from the nearby river and distributes it arbitrarily over the five buckets. Then Cinderella chooses a pair of neighboring buckets, empties them into the river, and puts them back. Then the next round begins. The Stepmother's goal is to make one of these buckets overflow. Cinderella's goal is to prevent this. Can the wicked Stepmother enforce a bucket overflow?

Answer:
Let the volume of water in the buckets be B_1, B_2, B_3, B_4 and B_5. Indices will be taken mod 5. Clearly if both B_i and B_{i+2} are greater than one before one Cinderella's moves, she cannot empty both of them and hence the stepmother will win enforce an overflow in her turn. Thus Cinderella's aim is to prevent this from happening. To do this, clearly it is sufficient to ensure that $B_i + B_{i+2}$ is at most one for each i after each of Cinderella's turns. Call such a situation *good*.

Suppose after some round we have a good situation: two buckets are empty, say $B_4 = B_5 = 0$. Then $B_1 + B_3 \leq 1$ and $B_2 \leq 1$ (since $B_2 + B_4 \leq 1$). After the stepmother's turn, we will have $B_1 + B_3 + B_4 + B_5 \leq 2$. Hence either $B_5 + B_3 \leq 1$ or $B_4 + B_1 \leq 1$. WLOG $B_5 + B_3 \leq 1$. Then Cinderella empties B_1 and B_2. Now observe that the new configuration is still good, since $B_4 \leq 1$ and $B_5 + B_3 \leq 1$.

Hence starting from a good configuration, Cinderella can ensure that at the end of the round the new configuration is still good. Initially the buckets are all empty, so this configuration is good. Hence Cinderella can prevent an overflow by always staying in a good position. ■

Exercises

1. [Cram]
A and **B** take turns placing dominoes on an $m \times n$ rectangular grid, where mn is even. **A** must place dominoes vertically and **B** must place dominoes horizontally, and dominoes cannot overlap with each other or stick out of the board. The player who cannot make any legal move loses. Given m and n, determine who has a winning strategy, and find this strategy.

2. [Double Chess]
The game of double chess is played like regular chess, except each player makes two moves in their turn (white plays twice, then black plays twice, and so on). Show that white can always win or draw.

3. [Russia 1999]
There are 2000 devices in a circuit, every two of which were initially joined by a wire. The hooligans Vasya and Petya cut the wires one after another. Vasya, who starts, cuts one wire on his turn, while Petya cuts two or three. A device is said to be disconnected if all wires incident to it have been cut. The player who makes some device disconnected loses. Who has a winning strategy?

4. [IMO Shortlist 2009, C1]
Consider 2009 cards, each having one gold side and one black side, lying on parallel on a long table. Initially all cards show their gold sides. Two players, standing by the same long side of the table, play a game with alternating moves. Each move consists of choosing a block of 50 consecutive cards, the leftmost of which is showing gold, and turning them all over,

so those which showed gold now show black and vice versa. The last player who can make a legal move wins.

a) Does the game necessarily end?

b) Does there exist a winning strategy for the starting player?

5. [Russia 1999]
There are three empty jugs on a table. Winnie the Pooh, Rabbit, and Piglet put walnuts in the jugs one by one. They play successively, with the order chosen by Rabbit in the beginning. Thereby Winnie the Pooh plays either in the first or second jug, Rabbit in the second or third, and Piglet in the first or third. The player after whose move there are exactly 1999 walnuts in some jug loses. Show that Winnie the Pooh and Piglet can cooperate so as to make Rabbit lose.

6. Solve the problem in example 10 with 1,000,000 replaced by n, an arbitrary odd number. Use this complete characterization of positions to provide a complete description of the winning strategy. If you have some programming experience, you could also write a program to play this game against you.
Remark (for programmers): You could also write a program to solve this probem, that is, to determine for each n who has a winning strategy. A simple dynamic programming approach would run in $O(n)$ time. Using this as a subroutine, the program to play the game against you would take $O(n)$ time for each move. However, if you found the characterization of positions on your own first, the program to play the game would only take $O(\log n)$ time for each move.

7. [Bulgaria 2005]
For positive integers t, a, b, a *(t, a, b)-game* is a two player game defined by the following rules. Initially, the number t is written on a blackboard. In his first move, the first player replaces t with either $t - a$ or $t - b$. Then, the second player subtracts either a or b from this number, and writes the result on the blackboard, erasing the old number. After this, the first player once again erases either a or b from the number

written on the blackboard, and so on. The player who first reaches a negative number loses the game. Prove that there exist infinitely many values of t for which the first player has a winning strategy for all pairs (a, b) with $(a + b) = 2005$.

8. [Rioplatense Math Olympiad 2010]
Alice and Bob play the following game. To start, Alice arranges the numbers 1, 2, …, n in some order in a row and then Bob chooses one of the numbers and places a pebble on it. A player's *turn* consists of picking up and placing the pebble on an adjacent number under the restriction that the pebble can be placed on the number k at most k times. The two players alternate taking turns beginning with Alice. The first player who cannot make a move loses. For each positive integer n, determine who has a winning strategy.

9. [Russia 2007]
Two players take turns drawing diagonals in a regular $(2n+1)$-gon $(n > 1)$. It is forbidden to draw a diagonal that has already been drawn or intersects an odd number of already drawn diagonals. The player who has no legal move loses. Who has a winning strategy?

10. [Indian Practice TST 2013]
A marker is placed at the origin of an integer lattice. Calvin and Hobbes play the following game. Calvin starts the game and each of them takes turns alternatively. At each turn, one can choose two (not necessarily distinct) integers a and b, neither of which was chosen earlier by any player and move the marker by a units in the horizontal direction and b units in the vertical direction. Hobbes wins if the marker is back at the origin any time after the first turn. Determine whether Calvin can prevent Hobbes from winning.
Note: A move in the horizontal direction by a positive quantity will be towards the right, and by a negative quantity will be towards the left (and similarly in the vertical case as well).

11. [Based on South Korea 2009, Problem 5]

Consider an $m \times (m+1)$ grid of points, where each point is joined by a line segment to is immediate neighbors (points immediately to the left, right, above or below). A stone is initially placed on one of the points in the bottom row. **A** and **B** alternately move the stone along line segments, according to the rule that no line segment may be used more than once. The player unable to make a legal move loses. Determine which player has a winning strategy.

12. [IMO Shortlist 1994, C6]

Two players play alternatively on an infinite square grid. The first player puts an X in an empty cell and the second player puts an O in an empty cell. The first player wins if he gets 11 adjacent X's in a line - horizontally, vertically or diagonally. Show that the second player can always prevent the first player from winning.

13. [Nim]

There are k heaps of stones, containing a_1, a_2, …, a_k stones respectively, where the a_i's are positive integers. Players **A** and **B** play alternately as follows: in each turn, a player chooses one non-empty heap and removes as many stones as he or she wants. The person who takes the last stone wins. Determine when each player has a winning strategy, and find this winning strategy.

14. [The name of this problem would give the answer away]

There is one pile of N counters. **A** and **B** play alternately as follows. In the first turn of the game, **A** may remove any positive number of counters, but not the whole pile. Thereafter, each player may remove at most twice the number of counters his opponent took on the previous move. The player who removes the last counter wins. Who has the winning strategy?

Olympiad Combinatorics

Pranav A. Sriram

August 2014

Copyright notices

About the Author

Pranav Sriram graduated from high school at The International School Bangalore, India, and will be a Freshman at Stanford University this Fall.

6. COUNTING IN TWO WAYS

Introduction

Several combinatorics problems ask us to count something – for example, the number of permutations of the numbers from 1 to n without fixed points, or the number of binary strings of length n with more 1s than 0s. What's interesting is that the techniques used to solve counting or enumeration problems can be applied to problems that *don't* ask us to count anything. Problems in fields such as combinatorial geometry, graph theory, extremal set theory and even number theory can be solved by clever applications of counting – *twice*.

The basic idea underlying this chapter is to compute or estimate some quantity Q (which will depend on the problem and information given to us) by counting in two different ways. We hence obtain two different expressions or bounds for Q. For instance, we may obtain $E_1 \geq Q$ and $E_2 = Q$. This allows us to conclude that $E_1 \geq E_2$, which may have been very difficult to prove directly. The role of counting in this approach is thus to allow us to convert complicated combinatorial information into convenient algebraic statements. The main challenge lies in choosing Q appropriately, so that we use all the information given to us and derive an algebraic conclusion relevant to what we are trying to prove.

Incidence Matrices

Let $A_1, A_2, \ldots, A_n$ be subsets of $\boldsymbol{S} = \{1, 2, \ldots, m\}$. A convenient way to express this information is by drawing an $n \times m$ matrix, with the n rows representing $A_1, A_2, \ldots, A_n$ and the m columns representing the elements of $\boldsymbol{S}$. Entry $a_{ij} = 1$ if and only if element j belongs to A_i. Otherwise, $a_{ij} = 0$. The idea of counting the total number of 1s in an incidence matrix is very useful.

Example 1
Let $A_1, A_2, \ldots, A_6$ be subsets of $\boldsymbol{S} = \{1, 2, \ldots, 8\}$. Suppose each set A_i has 4 elements and each element in $\boldsymbol{S}$ is in m of the A_i's. Find m.

Answer:
We draw an incidence matrix with six rows, representing the subsets $A_1, A_2, \ldots, A_6$ and eight columns representing the elements of $\boldsymbol{S}$. The entry in the ith row and jth column is 1 if and only if the element j belongs to A_i. Otherwise the entry is 0. Since $|A_i| = 4$, each row contains four 1s. There are 6 rows, so the total number of 1s in our matrix is $6 \times 4 = 24$.

Now each element of $\boldsymbol{S}$ is in m of the A_i's. Thus each column of our matrix contains m 1s. So the total number of 1s in the matrix is $8m$, since there are 8 columns. Thus $24 = 8m$, so $m = 3$. ■

1	1	0	1	1	0	0	0

Figure 6.1: A sample 6 x 8 incidence matrix with one row filled in, illustrating set A_3 containing elements 1, 2, 4 and 5.

Counting Pairs and Triples

What we are actually doing in the above proof is counting pairs of the form (*element*, *set*) where the set contains the element. Each 1 in the matrix corresponds to such a pair. If we choose the set first and then the element, there are 6 choices for the set and then 4 for the element, for a total of 24 pairs. We can also choose the element first (8 choices), and then choose the set (m choices, since each element belongs to m sets) for a total of $8m$ pairs. Equating the two answers, $8m = 24$, so $m = 3$.

More generally, we have the following result: If $A_1, A_2, \ldots, A_m$ are subsets of $\{1, 2, \ldots, n\}$ and each element j belongs to d_j of the subsets, then

$$\sum_{i=1}^{m} |A_i| = \sum_{j=1}^{n} d_j.$$

Both sides count the total number of 1s in the matrix, which is the number of pairs (*set*, *element*). The left side counts this quantity by picking the set first and the right side counts it by picking the element first. Note that both sides are also the sum of 1s in the incidence matrix.

In the first example, we counted pairs of the form (*set*, *element*) where the element belongs to the set. There are a few important variations of this technique:

(i) Count *triples* of the form (*set*, *set*, *element*) where the two sets are distinct both contain the element. This is especially useful if we are given information about the **intersection size** of any two sets. These triples can be counted either by first fixing the two sets and then picking the element from their intersection, or by fixing the element and then picking two sets to which it belongs.

Note that counting triples of the form (*set*, *set*, *element*) is equivalent to counting the number of **pairs of 1s that are in the same column** in the incidence matrix representation.

(ii) Count triples of the form (*element*, *element*, *set*) where the two elements both belong to the set. This is useful if we are given information about how many sets two elements appear together in. These triples can be counted in two ways: you can either fix the two elements first or you can fix the set first.

Note that counting triples of the form (*element*, *element*, *set*) is equivalent to counting the number of **pairs of 1s in the same row** in the incidence matrix representation.

The next example demonstrates (ii) in part (a), and the original idea of counting pairs (*set, element*) in part (b).

Example 2 [Balanced block designs]
Let $\boldsymbol{X} = \{1, 2, \dots, v\}$ be a set of elements. A (v, k, λ) *block design* over $\boldsymbol{X}$ is a collection of distinct subsets of $\boldsymbol{X}$ (called blocks) such that:
(i) Each block contains exactly k elements of $\boldsymbol{X}$
(ii) Every pair of distinct elements of $\boldsymbol{X}$ is contained in exactly λ blocks

Let b be the number of blocks. Prove that:
(a) Each element of $\boldsymbol{X}$ is contained in exactly $r = \lambda\,(v - 1)/(k - 1)$ blocks. (In particular, this means that each element is in the same number of blocks, which is initially not obvious)
(b) $r = b\,k\,/\,v$

Answer:
(a) Consider an element s in $\boldsymbol{X}$. We count in two ways the number of triples $(s, u, \boldsymbol{B})$ where u is an element (different from s) and $\boldsymbol{B}$ is a block containing s and u. The first way we count will be to fix $\boldsymbol{B}$ and then u, and the second way will do the reverse. If s is in r blocks, then there are r ways to choose $\boldsymbol{B}$, and then $(k - 1)$ ways to choose u from B. This gives a total of $r\,(k - 1)$. On

the other hand, there are $(v - 1)$ ways to choose u first, and then λ ways to choose $\boldsymbol{B}$ such than $\boldsymbol{B}$ contains both u and s (by condition (ii)). This gives a total of $\lambda (v - 1)$. Hence $r (k - 1) = \lambda (v - 1)$, which is what we wanted.

(b) We count in two ways the number of pairs $(x, \boldsymbol{B})$ where x is an element in a block $\boldsymbol{B}$. There are v ways to choose x, and then r ways to choose $\boldsymbol{B}$. This gives $v\, r$ pairs. On the other hand, there are b ways to choose $\boldsymbol{B}$ first, and then k ways to choose x since $|\boldsymbol{B}| = k$. This gives $b\, k$ pairs. Hence $b\, k = v\, r$. ■

The real power of counting in two ways lies in proving *inequalities*. Typically, we count the number of pairs P (or triples T) of some objects in two ways. At least one of the two counting procedures should give us a bound on P (or T, as may be the case). To do this, we need to cleverly exploit information given to us in the problem statement. The next example is fairly simple, as we use ideas we have already seen in preceding examples.

Example 3 [USA TST 2005]
Let n be an integer greater than 1. For a positive integer m, let $\boldsymbol{S}_m = \{1, 2, \ldots, mn\}$. Suppose that there exists a $2n$-element set $\boldsymbol{T}$ such that
(a) each element of $\boldsymbol{T}$ is an m-element subset of $\boldsymbol{S}_m$
(b) each pair of elements of $\boldsymbol{T}$ shares at most one common element; and
(c) each element of $\boldsymbol{S}_m$ is contained in exactly two elements of $\boldsymbol{T}$.

Determine the maximum possible value of m in terms of n.

Remark: Make sure you understand the problem – the "elements" of $\boldsymbol{T}$ are actually sets, that is, $\boldsymbol{T}$ is actually a family of subsets of $\boldsymbol{S}_m$.

Answer:
Let $A_1, A_2, \ldots, A_{2n}$ be the elements of $\boldsymbol{T}$. Let S be the number of triples (x, A_i, A_j) where x is an element of S_m belonging to both sets A_i and A_j. If we choose x first, there is only one choice for the pair

(A_i, A_j) since x belongs in exactly two elements of $\boldsymbol{T}$ by (c). This gives $S = mn$ (the number of choices for x). If we select A_i and A_j first, there is at most one choice for x by (b). Thus $S \leq \binom{2n}{2}$, the number of ways of choosing the pair (A_i, A_j). Hence

$$mn = S \leq \binom{2n}{2} \;\Rightarrow\; m \leq 2n - 1.$$

To give a construction when $m = 2n$-1, simply take $2n$ lines in the plane, no three of which concur and no two of which are parallel. There will be $\binom{2n}{2} = mn$ intersection points formed. The $2n$ lines are the $2n$ elements of $\boldsymbol{T}$, and the mn points are the elements of $\boldsymbol{S}_n$. The conditions of the problem are satisfied, since each point lies on exactly two lines, each two lines meet at exactly one point and each line contains $m = 2n - 1$ points since it meets the other $2n$-1 lines once each. ■

Slightly harder problems require a clever choice of *what pairs or triples to count,* and how to use the information in the problem to get the bounds we want. This comes with practice. One general principle to note is to pay attention to key phrases in the problem like "at most" and "at least". These pieces of information often give a good idea of what we should count.

Example 4 [IMO 1998, Problem 2]
In a competition, there are a contestants and b judges, where $b \geq 3$ is an odd integer. Each judge rates each contestant as either *pass* or *fail*. Suppose k is a number such that for any two judges, their ratings coincide for at most k contestants. Prove that $\frac{k}{a} \geq \frac{b-1}{2b}$.

Answer:
Let T be the number of triples (judge, judge, contestant) where the two judges both gave the same rating to the contestant. We can select the two judges in $\binom{b}{2}$ ways, and then select the contestant in **at most k** ways by the condition of the problem. Hence $T \leq k\binom{b}{2}$.

Now take any individual contestant, and suppose the number of judges who rated her "pass" is p and the number who rated her "fail" is $b - p$. The number of triples containing this candidate is $\binom{p}{2} + \binom{b-p}{2} \geq \binom{(b+1)/2}{2} + \binom{(b-1)/2}{2} = (b - 1)^2/4$. Here we used convexity and the fact that b is odd.

Thus each candidate is in at least $(b - 1)^2/4$ triples, so $T \geq a(b - 1)^2/4$. Combining this with our earlier estimate,

$$a(b - 1)^2/4 \leq kb(b - 1)/2 \Rightarrow \frac{k}{a} \geq \frac{b-1}{2b} \blacksquare$$

Unlike the previous example, the next problem offers us no clues that lead us to guess what we should count. However, we can exploit the geometry of the problem to our advantage.

Example 5 [Iran 2010]
There are n points in the plane such that no three of them are collinear. Prove that the number of triangles whose vertices are chosen from these n points and whose area is 1 is not greater than $\frac{2}{3}(n^2 - n)$.

Answer:
Let the number of such triangles be k. We count pairs (*edge*, *triangle*) such that the triangle contains the edge. If the number of such pairs is P, then clearly $P = 3k$, since each triangle has 3 edges.

On the other hand, for any edge AB, there are at most four points such that the triangles they form with A and B have the same area. This is because those points have to be the same distance from line AB, and no three of them are collinear. Hence P is at most 4 times the number of edges, which is at most $\binom{n}{2}$. Thus $P \leq 4\binom{n}{2}$. This gives

$$3k \leq 4\binom{n}{2} \Rightarrow k \leq \frac{2}{3}(n^2 - n) \blacksquare$$

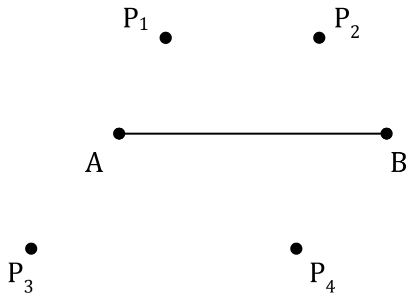

Figure 6.2: At most four points P_1, P_2, P_3, P_4 can form a triangle of unit area with segment AB

Remark: Whenever I'm faced with a combinatorial geometry problem that involves proving an inequality, like the above problem, I use the following principle: use the geometry of the situation to extract some combinatorial information. After that, ignore the geometry completely and use the combinatorial information to prove the inequality. We use this principle in the next example as well.

Example 6 [IMO 1987]

Let n and k be positive integers and let $\boldsymbol{S}$ be a set of n points in the plane such that:

(i) No three points of $\boldsymbol{S}$ are collinear

(ii) For every point P in $\boldsymbol{S}$, there are at least k points in $\boldsymbol{S}$ equidistant from P.

Prove that $k < \frac{1}{2} + \sqrt{2n}$.

Answer:

Condition (ii) implies that for each point P_i in $\boldsymbol{S}$, there exists a circle C_i with center P_i and passing through at least k points of $\boldsymbol{S}$.

Now we count pairs (P_i, P_j) such that P_i and P_j are points in $\boldsymbol{S}$.

Obviously the number of such pairs is $\binom{n}{2}$. On the other hand, each circle C_i has k points on its circumference, which give rise to $\binom{k}{2}$ pairs of points. Thus the n circles in total give us $n\binom{k}{2}$ points. However, there is over counting, since some pairs of points may belong to two circles. Since any two circles meet in at most 2 points, the number of pairs of points that we have counted twice is at most equal to the number of pairs of circles, which is $\binom{n}{2}$. Hence the total number of pairs of points in $\boldsymbol{S}$ is at least $n\binom{k}{2}$ - $\binom{n}{2}$. This implies

$$n\binom{k}{2} - \binom{n}{2} \le \binom{n}{2} \Rightarrow n\binom{k}{2} \le 2\binom{n}{2} \Rightarrow k^2 + k - (n-1) \le 0.$$

Solving this quadratic inequality, noting that k and n are integers, gives us the desired result. ■

The next example again requires a good choice of what to count, in order to capture all the given information.

Example 7 [IMO Shortlist 2004, C1]

There are 10001 students at a university. Some students join together to form several *clubs* (a student may belong to different clubs). Some clubs join together to form several *societies* (a club may belong to different societies). There are a total of k societies. Suppose that the following conditions hold:

(i) Each pair of students is in exactly one club.

(ii) For each student and each society, the student is in exactly one club of the society.

(iii) Each club has an odd number of students. In addition, a club with $2m + 1$ students (m is a positive integer) is in exactly m societies.

Find all possible values of k.

Answer:

In order to use all the information in the question, we count triples (a, C, S), where a is a student, C is a club and S is a society, where $a \in C$ and $C \in S$. Let the number of such triples be T.

Suppose we first fix a, then S, then C. We can choose a in 10001 ways, S in k ways and then finally C in only one way (by condition (ii)). Hence $T = 10001\, k$.

Now suppose we fix C first. There are $|C|$ ways of doing this. Then by condition (iii), there are $(|C| -1)/2$ ways to choose S. Finally there is only one way to choose a, by (ii). This gives

$$T = \sum_{C \text{ is a club}} |C|(|C| - 1)/2 = \sum_{C \text{ is a club}} \binom{|C|}{2}$$

On the other hand, the sum $\sum_{C \text{ is a club}} \binom{|C|}{2}$ is actually equal to the number of pairs of students. This is because each pair of students is in exactly one club by (i), so each pair of students is counted exactly once. Hence this sum is equal to $\binom{10001}{2}$, so putting everything together

$$\binom{10001}{2} = T = 10001\, k \Rightarrow \boldsymbol{k = 5000}.$$

Finally, to construct a configuration for $k = 5000$, let there be only one club C containing all students and 5000 societies all containing only one club (C). It's easy to see that this works. ■

Counting with Graphs

In the next few examples, we show how to use counting in two ways to solve some problems on graphs. Modeling situations using graphs is very useful, since graphs are very convenient to work with while counting in two ways. For example, suppose we want to count pairs of people such that the two people are friends. If we draw a graph with vertices representing people and an edge between two people if and only if they are friends, then the problem is equivalent to counting the number of edges in the graph.

Some useful properties of graphs

Let ***G*** be a graph with n vertices $v_1, v_2, \ldots, v_n$. Let d_i be the degree of v_i, ***E*** be the set of edges and $|E| = k$. All summations without indices are assumed to be from 1 to *n*. We have the following useful properties:

Lemma 1: $\sum d_i = 2k$ (this is because the LHS counts each edge of the graph twice)

Lemma 2: $\sum d_i^2 \geq \frac{(\sum d_i)^2}{n}$ (By Cauchy-Schwarz)

$$\Rightarrow \sum d_i^2 \geq \frac{4k^2}{n}$$

Lemma 3: $\sum \binom{d_i}{2} \geq \frac{2k^2}{n} - k$

Proof: $\binom{d_i}{2} = \frac{d_i^2 - d_i}{2}$. Using lemma 1 and lemma 2 produces the result.

Lemma 4: $\sum_{v_i v_j \in E}(d_i + d_j) = \sum_{i=1}^{n} d_i^2$

Proof: Each term d_i appears in the sum on the LHS d_i times (once for each of the neighbors of v_i). Thus the total sum will be the sum of $d_i \times d_i = d_i^2$ for each *i*, which is the RHS.

There are also some important results on directed graphs, especially *tournaments.* A tournament on *n* vertices is a directed graph such that between any two vertices *u* and *v*, there is either a directed edge from *u* to *v* or a directed edge from *v* to *u*. One can interpret these graphs as follows: the *n* vertices stand for participants in a tournament, and each two players play a match.

There are no ties. If v beats u, then there is a directed edge from v to u.

Let P_1, P_2, …, P_n be the n participants. Let w_i and l_i denote the number of wins and losses of P_i. Clearly $w_i + l_i = (n - 1)$ for each i, because each person plays against $(n - 1)$ others. Also, $\sum w_i = \sum l_i$ since each match has a winner and a loser, and so contributes 1 to both sides. Hence in fact both sides are equal to $\binom{n}{2}$, the total number of matches. We have another interesting but less obvious result:

Lemma 5: $\sum w_i^2 = \sum l_i^2$.

Proof: Define a *noncyclic triple* to be a set of 3 players A, B and C such that A beat both B and C and B beat C. Call A the *winner* of the triplet and C the *loser* of the triplet. If we count noncyclic triplets by winners, the sum would be $\sum\binom{w_i}{2}$, since after choosing the winner there are $\binom{w_i}{2}$ ways to choose the other two players who he beat. If we count by losers, the sum is $\sum\binom{l_i}{2}$, since after choosing the loser there are $\binom{l_i}{2}$ ways to choose the other two players. Hence $\sum\binom{w_i}{2} = \sum\binom{l_i}{2}$. Combining this with $\sum w_i = \sum l_i$ we get the result.

Remark: Whenever you see expressions of the type in this lemma, like a sum of squares, try to interpret them combinatorially. For instance, it is often useful to convert x^2 to $2\binom{x}{2} + 2x$. Allow these sums to give you hints as to what to count. In the proof of lemma 5, the term $\binom{w_i}{2}$ gives us a hint to count triples of the form (X, Y, Z) such that X beat both Y and Z. This almost automatically leads us to the solution.

Example 8 [APMO 1989] (U*)

Show that a graph with n vertices and k edges has at least $k(4k - n^2)/3n$ triangles.

Note: The symbol **U*** in brackets next to a problem indicates that it is a useful result and should be remembered.

Answer:
We count pairs (*edge*, *triangle*) where the triangle contains the edge. Consider an edge v_iv_j. How many triangles have v_iv_j as an edge? v_i is joined to $(d_i - 1)$ vertices other than v_j, and v_j is joined to d_j vertices other than v_i. There are only $n - 2$ vertices other than v_i and v_j. Hence at least $(d_i - 1) + (d_j - 1) - (n - 2) = (d_i + d_j - n)$ vertices are joined to *both* v_i and v_j. Each of these gives one triangle. Hence each edge v_iv_j is in at least max $\{0, (d_i + d_j - n)\}$ triangles.

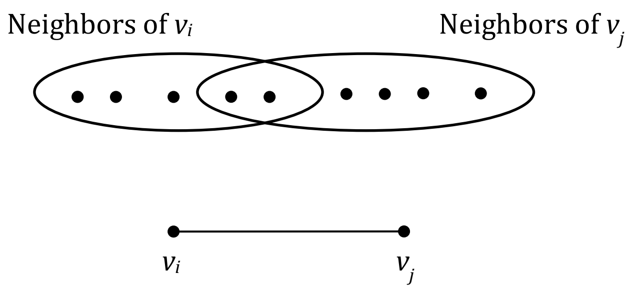

Figure 6.3: The set of vertices neighboring both v_i and v_j must contain at least $(d_i + d_j - n)$ vertices

Thus the total number of triangles is at least

$$\frac{1}{3}\sum_{v_iv_j \in E}(d_i + d_j - n) = \frac{1}{3}\sum_{i=1}^{n} d_i^2 - \frac{nk}{3} \quad \text{(Using lemma 4)}$$

$$\geq \frac{1}{3} \times \frac{4k^2}{n} - \frac{nk}{3} = \frac{k(4k-n^2)}{3n}. \quad \text{(Using lemma 3)}$$

Note that we divided by 3 because otherwise each triangle would be counted thrice (once for each edge). ■

Corollary 1 (U*)
A graph with no triangles has at most $\lfloor n^2/4 \rfloor$ edges. Equality is achieved only by bipartite graphs with an equal or almost equal number of vertices in each part. This is an extremely useful result, and is a special case of **Turan's theorem**, which will be discussed in the exercises of chapter 8.

Example 9 [Indian TST 2001] (U*)
Let $\boldsymbol{G}$ be a graph with E edges, n vertices and no 4-cycles. Show that $E \leq \frac{n}{4}(1 + \sqrt{4n-3})$.

Answer:
Let the vertices be $\{v_1, \ldots, v_n\}$ and let the degree of v_i be d_i. Let T be the number of "*V-shapes*": that is, triples of vertices (u, v, w) such that v and w are both adjacent to u. The vertices v and w may or may not be adjacent and triples $\{u, v, w\}$ and $\{u, w, v\}$ are considered the same.

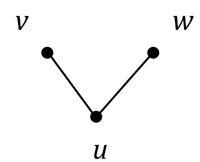

Figure 6.4: A "V shape"

The reason for this choice of T is that if we first select v and w, then there is at most one u such that $\{u, v, w\}$ is a triple in T. Otherwise there would be a 4 cycle. Hence we get $\boldsymbol{T} \leq \binom{\boldsymbol{n}}{\boldsymbol{2}}$, since for each of the $\binom{n}{2}$ ways of choosing v and w, there is at most one way to choose u.

If we choose u first, there are $\binom{d_u}{2}$ ways of choosing v and w, where d_u is the degree of u. Summing over all choices for u, and then using lemma 3, we get

$$T = \sum_{i=1}^{n} \binom{d_i}{2} \geq \frac{2E^2}{n} - E$$

Combining this with $T \leq \binom{n}{2}$,

$$\frac{2E^2}{n} - E \leq \frac{n(n-1)}{2}.$$

This reduces to a quadratic inequality in E, which yields the desired bound. ■

Sometimes, when we need to bound or count the number of objects satisfying some property, it is easier or more convenient to count the number of objects *not* satisfying the property. Then we can subtract this from the total number of objects to get the result.

Example 10 [USAMO 1995]
Suppose that in a certain society, each pair of persons can be classified as either *amicable* or *hostile*. We shall say that each member of an amicable pair is a *friend* of the other, and each member of a hostile pair is a *foe* of the other. Suppose that the society has n people and q amicable pairs, and that for every set of three persons, at least one pair is hostile. Prove that there is at least one member of the society whose foes include $q\,(1 - 4q/n^2)$ or fewer amicable pairs.

Answer:
We naturally rephrase the problem in graph theoretic terms, with vertices representing people and an edge joining two vertices if and only if they form an amicable pair. The graph has no triangles by assumption, n vertices and q edges. We wish to estimate the number of edges containing 2 foes of X, where X is a vertex. To do this, we first count P, the number of pairs (E, X), where E is an edge containing X or a friend of X.

First we count the number of pairs (X, E) where E is an edge containing a neighbor of X but not containing X. This quantity will be equal to T, the number of triples (X, Y, Z) such that XY and YZ are edges. (X, Y, Z) is considered different from (Z, Y, X). Note that XZ cannot be an edge by the condition that there are no triangles. To compute T, we count by Y. The number of triples containing Y is $d_Y(d_Y - 1)$, so the total number of triples is $\sum_{i=1}^{n} d_i(d_i - 1)$.

Now clearly the number of pairs (X, E) where E is an edge containing X is given by the sum $\sum_{i=1}^{n} d_i$. If we add this summation to the previous summation, we would have counted the number of pairs (X, E) where E is an edge containing X OR a friend of X but not X. Thus the total number of such pairs is

$$P = \sum_{i=1}^{n} d_i(d_i - 1) + \sum_{i=1}^{n} d_i = \sum_{i=1}^{n} d_i^2 \geq 4q^2/n,$$

by lemma 2.

Hence by averaging, there is some X such that there at least $4q^2/n^2$ pairs (X, E), where E is an edge X or at least one neighbor of X. Thus the number of edges joining two foes of X is at most $q - 4q^2/n^2 = q(1 - 4q/n^2)$.

Example 11 [Generalization of Iran TST 2008]

In a tournament with n players, each pair of players played exactly once and there were no ties. Let j, k be integers less than n such that $j < 1 + \frac{n\binom{(n-1)/2}{k}}{\binom{n}{k}}$. Show that there exist sets $\boldsymbol{A}$ and $\boldsymbol{B}$ of k players and j players respectively, such that each player in $\boldsymbol{A}$ beat each player in $\boldsymbol{B}$.

Answer

Count $(k + 1)$–tuples of the form $(P_1, P_2, \ldots, P_k, L)$ where L lost to each of the players $P_1, P_2, \ldots, P_k$. Let T be the total number of such tuples. If we fix L, we get $\binom{d_L}{k}$ tuples containing L, where d_L is the number of players L lost to. Summing over all n choices of L, $T \geq$

$\sum_{i=1}^{n}\binom{d_i}{k}$, where d_i is the number of losses of the ith player. Hence by Jensen's inequality, $T \geq n \times \binom{\sum_{i=1}^{n} d_i/n}{k} = n \times \binom{(n-1)/2}{k}$, since $\sum_{i=1}^{n} d_i = n\,(n-1)/2$.

Now assume to the contrary that there do not exist such sets ***A*** and ***B***. Then for any choice of $P_1, P_2, \ldots, P_k$, there are at most $(j-1)$ choices for L. Hence $T \leq \binom{n}{k}(j-1)$.

Combining these estimates gives $(j-1) \geq \frac{n\binom{(n-1)/2}{k}}{\binom{n}{k}}$, which contradicts the condition of the problem. Thus, our assumption in the second paragraph is false, and such sets ***A*** and ***B*** indeed exist.

Example 12 [IMO Shortlist 2010 C5]
$n \geq 4$ players participated in a tennis tournament. Any two players have played exactly one game, and there was no tie game. We call a company of four players *bad* if one player was defeated by the other three players, and these three players formed a cyclic triple (a set (A, B, C) such that A beat B, B beat C and C beat A). Suppose that there is no bad company in this tournament. Let w_i and l_i be respectively the number of wins and losses of the ith player. Prove that

$$\sum_{i=1}^{n}(w_i - l_i)^3 \geq 0.$$

Answer:
Note that

$$\sum_{i=1}^{n}(w_i - l_i)^3 = \sum_{i=1}^{n}(w_i^3 - l_i^3) + 3\sum_{i=1}^{n}(w_i l_i^2 - w_i^2 l_i)$$

We will show that

(i) $\sum_{i=1}^{n} w_i^3 \geq \sum_{i=1}^{n} l_i^3$

(ii) $\sum_{i=1}^{n} w_i l_i^2 \geq \sum_{i=1}^{n} w_i^2 l_i$

From now on, any summation without indices is assumed to be from 1 to n. Note that by using lemma 5, we can reduce **(i)** to the "more combinatorial" form

(iii) $\sum\binom{w_i}{3} \geq \sum\binom{l_i}{3}$

Let us define a "chained quadruple" as a set of 4 players with no cyclic triple amongst them. It is easy to see that a chained quadruple has

(a) A person who won against all the other three players, called the *winner*
(b) A person who lost against all the other three players, called the *loser*

The converse of (a) is not true, since the other three players may form a cyclic triple. However, the converse of (b) holds, since by assumption there is no bad quadruple. Let Q be the number of chained quadruples. If we count Q by picking the loser first, we get

$$Q = \sum\binom{l_i}{3}$$

If we count Q by picking the winner first, noting that the converse of (a) doesn't hold, then

$$Q \leq \sum\binom{w_i}{3}$$

Hence $\sum\binom{w_i}{3} \geq \sum\binom{l_i}{3}$, which proves **(iii)** and hence **(i)**.

To prove **(ii)**, subtract $\sum w_i l_i$ from both sides and divide by 2, to write it as:

(iv) $\sum w_i\binom{l_i}{2} \geq \sum l_i\binom{w_i}{2}$

Observe that the LHS of this expression counts pairs of the form (*quadruple*, *person*) such that the person won exactly one game against the other three in the quadruple. Similarly, the RHS counts pairs such that the person won exactly two games.

Now let us look at the types of quadruples we can have. If in a certain quadruple the number of games won by each person against the other three are a, b, c, d in non-increasing order, we say that this quadruple is of type (a, b, c, d). The only types we can have are:

(3, 1, 1, 1) – Note that this refers to a quadruple in which one person beat the other three, and the other three each won one game. This type of quadruple contributes 3 to the LHS of **(iv)** (three people won one game) and 0 to the RHS (no one won two games).
(2, 2, 1, 1) – This contributes 2 to both sides of **(iv)**
(3, 2, 1, 0) – This contributes 1 to both sides of **(iv)**
(2, 2, 2, 0) – This is not allowed: this is a bad company.

Thus we see that every allowed quadruple contributes at least as much to the LHS of **(iv)** as it does to the RHS. Hence **(iv)** indeed holds, which proves **(ii)**. Hence **(i)** and **(ii)** together give us the desired result and we are done.

Remark 1: This example shows the true power of "interpreting things combinatorially".

Remark 2: This problem was the first relatively hard (rated above C2 or C3) combinatorics problem I ever solved, and my solution was essentially the one above. The thought process behind this solution is fairly natural – keep expressing things "combinatorially", let these expressions guide what you choose to count, and exploit the fact that there is no "(2, 2, 2, 0)". Also note that it is not essential to prove **(i)** and **(ii)** separately: one can directly show that $\sum\binom{w_i}{3} + \sum w_i\binom{l_i}{2} \geq \sum\binom{l_i}{3} + \sum l_i\binom{w_i}{2}$ by comparing the contributions to each side by each type of quadruple.

Miscellaneous Applications

In this section we look at some unexpected applications of counting in two ways.

Example 13 [IMO 2001, Problem 4]
Let n be an odd integer greater than 1 and let $c_1, c_2, \ldots, c_n$ be integers. For each permutation $a = \{a_1, a_2, \ldots, a_n\}$ of $\{1, 2, \ldots, n\}$, define $S(a) = \sum_{i=1}^{n} a_i c_i$. Prove that there exist permutations $a \neq b$ of such that $n!$ is a divisor of $S(a) - S(b)$.

Answer:
Suppose to the contrary that all the $S(a)$'s are distinct modulo $(n!)$. Since there are $n!$ possibilities for a, this means that $S(a)$ takes each value in $\{1, 2, \ldots, n!\}$ modulo $n!$ Consider the sum of all the $S(a)$'s modulo $n!$. If the sum is S, then

$$S \equiv 1 + 2 + \ldots + n! \equiv n!\,(n! + 1)/2 \bmod n! \equiv n!/2 \bmod n!.$$

On the other hand, the coefficient of each c_i in S is

$$(n-1)!\,(1 + 2 + \ldots + n) = n!\,(n+1)/2 \equiv 0 \bmod n!,$$

since n is odd and 2 divides $n + 1$. Thus the coefficient of each c_i in S is divisible by $n!$, so $S \equiv 0 \bmod n!$. This is a contradiction to the result in the first paragraph. ■

Example 14 [IMO Shortlist 2003, C4]
Given n real numbers $x_1, x_2, \ldots, x_n$ and n further real numbers $y_1, y_2, \ldots, y_n$. The entries a_{ij} (with $1 \leq i, j \leq n$) of an $n \times n$ matrix A are defined as follows:

$$a_{ij} = \begin{cases} 1 & \text{if } x_i + y_j \geq 0; \\ 0 & \text{if } x_i + y_j < 0. \end{cases}$$

Further, let B be an $n \times n$ matrix whose elements are numbers from the set {0, 1} satisfying the following condition: The sum of all elements of each row of B equals the sum of all elements of the corresponding row of A; the sum of all elements of each column of B equals the sum of all elements of the corresponding column of A. Show that in this case $A = B$.

Answer:
Let b_{ij} denote the entry in the ith row and jth column of B. Define

$$S = \sum_{1 \le i, j \le n} (a_{ij} - b_{ij})(x_i + y_j)$$

On the one hand,

$$S = \sum_{i=1}^{n} x_i \left(\sum_{j=1}^{n} a_{ij} - \sum_{j=1}^{n} b_{ij}\right) + \sum_{j=1}^{n} y_j \left(\sum_{i=1}^{n} a_{ij} - \sum_{i=1}^{n} b_{ij}\right) = 0,$$

since $\sum_{j=1}^{n} a_{ij} = \sum_{j=1}^{n} b_{ij}$ and $\sum_{i=1}^{n} a_{ij} = \sum_{i=1}^{n} b_{ij}$ by the conditions of the problem.

On the other hand, note that if $x_i + y_j \ge 0$, then $a_{ij} = 1$ so $(a_{ij} - b_{ij}) \ge 0$. If $x_i + y_j < 0$, then $a_{ij} = 0$ so $a_{ij} - b_{ij} \le 0$. Thus in both cases, $(x_i + y_j)(a_{ij} - b_{ij}) \ge 0$. Hence each term in the summation is nonnegative, but the total sum is 0. Thus each term is 0. Hence whenever $(x_i + y_j) \ne 0$, we must have $a_{ij} = b_{ij}$. Whenever $(x_i + y_j) = 0$, then $a_{ij} = 1$. In these cases we must have $b_{ij} = 1$ since the sum of all the entries in both matrices is the same. Hence in all cases $a_{ij} = b_{ij}$, and we are done. ■

Remark: Where on earth does the expression

$$S = \sum_{1 \le i, j \le n} (a_{ij} - b_{ij})(x_i + y_j)$$

come from?!?! Note that one way of proving that several different real numbers are 0 is to show that their squares sum to 0, since no square is negative. Thus, a first approach to the problem may be to show that the sum

$$S' = \sum_{1\le i,j\le n}(a_{ij} - b_{ij})^2$$

is 0. This doesn't work as it doesn't utilize information about the *x*'s and *y*'s. Instead we try the following modification: we seek to weight each term $(a_{ij} - b_{ij})$ by some other quantity that still ensures that each term in the summation is nonnegative, and additionally enables us to use the information about the *x*'s and *y*'s to show that the entire sum is 0.

Example 15 [Indian TST 2010]
Let $A = (a_{jk})$ be a 10 x 10 array of positive real numbers such that the sum of numbers in each row as well as in each column is 1. Show that there exists $j < k$ and $l < m$ such that

$$a_{jl}a_{km} + a_{jm}a_{kl} \ge \frac{1}{50}$$

Answer:
To make things more intuitive, let us interpret the algebraic expression $a_{jl}a_{km} + a_{jm}a_{kl}$visually. The centers of the squares containing entries a_{jl}, a_{jm}, a_{km} and a_{kl} form a rectangle with sides parallel to grid lines. Define the *value* of this rectangle to be $a_{jl}a_{km} + a_{jm}a_{kl}$. Assuming to the contrary that the *value* of any such rectangle is strictly less than 1/50.

Observe that as *j*, *k*, *l*, *m* vary within the bounds $1 \le j < k \le 10$ and $1 \le l < m \le 10$, we obtain $\binom{10}{2}^2 = 45^2$ such rectangles. Let *S* be the sum of *values* of these 45^2 rectangles. By our earlier assumption, we obtain $S < 45^2/50 = 40.5$. We will now compute *S* in a different way to yield a contradiction.

Note that a_{jl} and a_{km} lie diagonally opposite and a_{jm} and a_{kl} lie diagonally opposite each other. Thus in each rectangle the diagonally opposite pairs of entries are multiplied. Hence, when the sum of *values* is taken over all rectangles, each entry a_{ij} occurs in products with every other entry in the array *except those in its own row or column*, since two entries in the same row or column can never be diagonally opposite in a rectangle. Therefore,

$$S = \frac{1}{2}\sum_{1\le i,j\le 10} a_{ij}\, S_{ij},$$

where S_{ij} is the sum of all entries except those in the ith row and jth column. Note that we have divided by two since if we simply sum the terms $a_{ij}S_{ij}$, we will be counting each product $a_{ij}a_{kl}$ twice.

Observe that $S_{ij} = (10 - 1 - 1 + a_{ij}) = (8 + a_{ij})$, since the sum of all entries is 10 and the sum in each row and column is 1. Note that the " $+ a_{ij}$" occurs since when we subtract all elements in row i and in column j, a_{ij} is subtracted twice. Thus the total sum is

$$S = \frac{1}{2}\sum_{1\le i,j\le 10} a_{ij}\, S_{ij} = \frac{1}{2}\sum_{1\le i,j\le 10} a_{ij}(a_{ij} + 8)$$
$$= 4\sum_{1\le i,j\le 10} a_{ij} + \frac{1}{2}\sum_{1\le i,j\le 10} a_{ij}^2$$

Now $\sum_{1\le i,j\le 10} a_{ij} = 10$ and $\sum_{1\le i,j\le 10} a_{ij}^2 \ge \frac{(\sum_{1\le i,j\le 10} a_{ij})^2}{100} = 1$, using Cauchy Schwarz. Thus

$S \ge 4\text{x}10 + 0.5 = 40.5$, a contradiction. ■

Remark: The visual interpretation as "diagonally opposite entries in rectangles" is by no means essential (and entails a little abuse of notation as well, for which I apologize). Simply taking a suitable double summation would lead to a significantly shorter but equivalent proof. However, I felt the basic intuition behind the problem may have been lost in a sea of symbols that would have mysteriously spat out the solution, so I chose to write the proof this way.

Exercises

1. **[Due to Grigni and Sipser]**
Consider an $m \times n$ table (m rows, n columns), in which each cell either contains a 0 or a 1. Suppose the entire table contains at least αmn 1s, where $0 < \alpha < 1$. Show that at least one of the following must be true:
(i) There exists a row containing at least $n\sqrt{\alpha}$ 1s
(ii) There exist at least $m\sqrt{\alpha}$ rows containing at least αn 1s.

2. **[Italy TST 2005, Problem 1]**
A class is attended by n students ($n > 3$). The day before the final exam, each group of three students conspire against another student to throw him/her out of the exam. Prove that there is a student against whom there are at least $\sqrt[3]{(n-1)(n-2)}$ conspirators.

3. **[Important Lemmas on incident matrices] (U*)**
Let A be an $r \times c$ matrix with row sums R_i (that is, the sum of the elements in the ith row is R_i) and column sums C_j. Suppose R_i and C_j are positive for all $1 \leq i \leq r$ and $1 \leq j \leq c$.

(i) Show that $\sum_{i,j} \frac{a_{ij}}{R_i} = r$ and $\sum_{i,j} \frac{a_{ij}}{C_j} = c$
(ii) Suppose $C_j \geq R_i$ whenever $a_{ij} = 1$. Using (i), show that $r \geq c$.
(iii) Suppose instead of the condition in (ii) we were given that $0 < R_i < c$ and $0 < C_j < r$ for each i and each j, and furthermore, $C_j \geq R_i$ whenever $a_{ij} = 0$. Prove that $r \geq c$.

4. **[IMO 1987, Problem 1]**
Let $p(n, k)$ denote the number of permutations of $\{1, 2, \ldots, n\}$ with exactly k fixed points. Show that $\sum_{k=1}^{n} kp(n,k) = n!$

5. [Corradi's Lemma] (U*)
Let $A_1, A_2, \ldots, A_n$ be r-element subsets of a set $\boldsymbol{X}$. Suppose that $A_i \cap A_j \leq k$ for all $1 \leq i < j \leq n$. Show that $|\boldsymbol{X}| \geq \frac{nr^2}{r+(n-1)k}$.

6. [Erdos-Ko-Rado] (U*)
Let $\boldsymbol{F}$ be a family of k-element subsets of $\{1, 2, \ldots, n\}$ such that every two sets in $\boldsymbol{F}$ intersect in at least one element. Show that $|\boldsymbol{F}| \leq \binom{n-1}{k-1}$.

7. [Indian Postal Coaching 2011]
In a lottery, a person must select six distinct numbers from $\{1, 2, \ldots, 36\}$ to put on a ticket. The lottery committee will then draw six distinct numbers randomly from $\{1, 2, \ldots, 36\}$. Any ticket not containing any of these 6 numbers is a winning ticket. Show that there exists a set of nine tickets such that at least one of them will certainly be a winning ticket, whereas this statement is false if 9 is replaced by 8.

8. [Hong Kong 2007]
In a school there are 2007 girls and 2007 boys. Each student joins at most 100 clubs in the school. It is known that any two students of opposite genders have joined at least one common club. Show that there is a club with at least 11 boys and 11 girls.

9. [IMO Shortlist 1995, C5]
At a meeting of $12k$ people, each person exchanges greetings with exactly $(3k + 6)$ others. For any two people, the number of people who exchange greetings with both of them is the same. How many people are at the meeting?

10. [Based on Furedi's result on maximal intersecting families]
Let n and k be positive integers with $n > 2k - 1$, and let $\boldsymbol{F}$ be a family of subsets of $\{1, 2, \ldots, n\}$ such that each set in $\boldsymbol{F}$ contains k elements, and every pair of sets in $\boldsymbol{F}$ has nonzero

intersection. Suppose further that for any k-element subset $\boldsymbol{X}$ of {1, 2, …, n} not in $\boldsymbol{F}$, there exists a set $\boldsymbol{Y}$ in $\boldsymbol{F}$ such that $\boldsymbol{X} \cap \boldsymbol{Y} = \emptyset$. Show that there are at least $\frac{\binom{n}{k}}{\binom{n-k}{k}+1}$ sets in $\boldsymbol{F}$.

11. [IMO Shortlist 2000, C3]
Let $n > 3$ be a fixed positive integer. Given a set $\boldsymbol{S}$ of n points P_1, P_2, …, P_n in the plane such that no three are collinear and no four concyclic, let a_t be the number of circles $P_i P_j P_k$ that contain P_t in their interior, and let $m(\boldsymbol{S}) = a_1 + a_2 + \ldots + a_n$. Prove that there exists a positive integer $f(n)$ depending only on n such that the points of $\boldsymbol{S}$ are the vertices of a convex polygon if and only if $m(\boldsymbol{S}) = f(n)$.

12. [Iran 2010]
There are n students in a school, and each student can take any number of classes. There are at least two students in each class. Furthermore, if two different classes have two or more students in common, then these classes have a different number of students. Show that the number of classes is at most $(n - 1)^2$.

13. [IMO Shortlist 2004, C4]
Consider a matrix of size $n \times n$ whose entries are real numbers of absolute value not exceeding 1. The sum of all entries of the matrix is 0. Let n be an even positive integer. Determine the least number C such that every such matrix necessarily has a row or a column with the sum of its entries not exceeding C in absolute value.

14. [Generalization of USAMO 2011, Problem 6]
Let $A_1, A_2, \ldots, A_n$ be sets such that $|A_i| = \binom{n-1}{2}$ for each $1 \le i \le n$ and $|A_i \cap A_j| = (n - 2)$ for each $1 \le i < j \le n$. Show that $|A_1 \cup A_2 \cup \ldots \cup A_n| \ge \binom{n}{3}$, and show that it is possible for equality to occur.

15. [Iran 1999]

Suppose that $C_1, C_2, \ldots, C_n$ ($n \geq 2$) are circles of radius one in the plane such that no two of them are tangent, and the subset of the plane formed by the union of these circles is connected.

Let $\boldsymbol{S}$ be the set of points that belong to at least two circles. Show that $|\boldsymbol{S}| \geq n$.

16. [IMO Shortlist 2000, C5]

Suppose n rectangles are drawn in the plane. Each rectangle has parallel sides and the sides of distinct rectangles lie on distinct lines. The rectangles divide the plane into a number of regions. For each region R let $v(R)$ be the number of vertices. Take the sum of $v(R)$ over all regions which have one or more vertices of the rectangles in their boundary. Show that this sum is less than $40n$.

17. [Indian TST 1998]

Let $\boldsymbol{X}$ be a set of $2k$ elements and $\boldsymbol{F}$ a family of subsets of $\boldsymbol{X}$ each of cardinality k such that each subset of $\boldsymbol{X}$ of cardinality $(k - 1)$ is contained in exactly one member of $\boldsymbol{F}$. Show that (k + 1) is a prime.

18. [IMO Shortlist 1988]

For what values of n does there exist an $n \times n$ array of entries -1, 0 or 1 such that the $2n$ sums obtained by summing the elements of the rows and the columns are all different?

19. [IMO 2001, Problem 3]

Twenty-one girls and twenty-one boys took part in a mathematical competition. It turned out that each contestant solved at most six problems, and for each pair of a girl and a boy, there was at least one problem that was solved by both the girl and the boy. Show that there is a problem that was solved by at least three girls and at least three boys.

20. [IMO 2005, Problem 6]
In a mathematical competition 6 problems were posed to the contestants. Each pair of problems was solved by more than 2/5 of the contestants. Nobody solved all 6 problems. Show that there were at least 2 contestants who each solved exactly 5 problems.

21. Let $\boldsymbol{A}$ be a set with n elements, and let $\boldsymbol{A}_1, \boldsymbol{A}_2, \ldots, \boldsymbol{A}_n$ be subsets of $\boldsymbol{A}$ such that $|\boldsymbol{A}_i| \geq 2$ for each $1 \leq i \leq n$. Suppose that for each 2-element subset $\boldsymbol{A}'$ of $\boldsymbol{A}$, there is a unique i such that $\boldsymbol{A'}$ is a (not necessarily proper) subset of $\boldsymbol{A}_i$. Show that for all pairs (i, j) such that $1 \leq i < j \leq n$, $\boldsymbol{A}_i \cap \boldsymbol{A}_j > 0$.

22. [USAMO 1999 proposal]
Let n, k and m be positive integers with $n > 2k$. Let $\boldsymbol{S}$ be a nonempty set of k-element subsets of $\{1, 2, \ldots, n\}$ such that every $(k + 1)$-element subset of $\{1, 2, \ldots, n\}$ contains exactly m elements of $\boldsymbol{S}$. Prove that $\boldsymbol{S}$ must contain every k-element subset of $\{1, 2, \ldots, n\}$.

23. [Based on Zarankeiwicz' problem]
At a math contest there were m contestants and n problems. It turned out that there were numbers $a < m$ and $b < n$ such that there did not exist a set of a contestants and b problems such that all a contestants solved all b problems. Define the *score* of each contestant to be the number of problems he solved, and let S denote the sum of the scores of all m contestants. Show that $S \leq (a - 1)^{1/b} n m^{1\text{-}1/b} + (b - 1)m$.

24. [IMO Shortlist 2007, C7]
Let $\alpha < \frac{3-\sqrt{5}}{2}$ be a positive real number. Prove that there exist positive integers n and $p > \alpha 2^n$ for which one can select $2p$ pairwise distinct subsets $S_1, S_2, \ldots, S_p, T_1, T_2, \ldots, T_p$ of the set $\{1, 2, \ldots, n\}$ such that $S_i \cap T_j \neq \emptyset$ for all $1 \leq i, j \leq p$.

Olympiad Combinatorics

Pranav A. Sriram

August 2014

Copyright notices

About the Author

Pranav Sriram graduated from high school at The International School Bangalore, India, and is currently pursuing undergraduate studies at Stanford University.

7. Extremal Combinatorics

Introduction

Extremal combinatorics is, in essence, the combinatorics of superlatives. Problems in this field typically revolve around finding or characterizing the maximum, minimum, best, worst, biggest or smallest object, number or set satisfying a certain set of constraints. This chapter and the next two will take us to the heart of combinatorics, and will represent a deep dive into the intersection of Olympiad mathematics, classical combinatorics, and modern research in the field. Extremal combinatorics is an actively researched area, with deep connections to fundamental problems in theoretical computer science, operations research and statistical learning theory. Entire books have been devoted to the subject (and rightfully so), so we will not be able to do complete justice to this field in a single chapter. However, the powerful arsenal of tools we have built up in the first six chapters has already done much of our work for us: indeed, pretty much every technique we have seen so far has a role to play in extremal combinatorics. This chapter, then, will develop specialized, niche methods for extremal combinatorics, as well as demonstrate how to effectively exploit combinatorial structure to apply classical techniques like induction effectively in the context of extremal problems.

Injections and Bijections

One simple way to compare the cardinalities of sets is to find mappings between them. For instance, suppose I have two sets S and T, and a function f mapping elements from S to elements of T. If f is *injective* – that is, $f(x) = f(y)$ if and only if $x = y$ – then it follows that T must have *at least* as many elements as S. Moreover, if f is also *surjective* - that is, all elements in T are mapped to by some element in S - then it follows that |S| = |T|. A function that is injective and surjective is called a *bijection*.

The basic idea in this section will be to construct functions between carefully chosen sets. The choice of these sets will enable us to exploit information given in the problem in order to conclude that our function is injective or bijective. This conclusion will give us quantitative results relating the sizes of the sets, which will hopefully reduce to the result we are trying to prove.

Example 1 [APMO 2008]
Students in a class form groups. Each group contains exactly three members and any two distinct groups have at most one member in common. Prove that if there are 46 students in the class, then there exists a set of at least 10 students in which no group is properly contained.

Answer:
Let T be the set of 46 students. Take the largest set S of students such that no group is properly contained in S. Now take any student X not in S. By the maximality of S, there exists a group containing X and two students of S (otherwise we could add X to S, contradicting maximality). This suggests the following mapping: if (A, B, X) is this group, define a mapping from $T \setminus S$ to pairs in S such that $f(X) = (A, B)$. This mapping is injective because if $f(Y) = f(Z)$ for some $Y \neq Z$, then both (Y, C, D) and (Z, C, D) are groups for

some (C, D), contradicting the fact that any two groups have at most one common student. The injectivity implies $| T \setminus S | \leq \binom{|S|}{2}$, or $(46-|S|) \leq \binom{|S|}{2}$. Simplifying gives $|S| \geq 10$. ■

Example 2 [IMO Shortlist 1988]
Let $N = \{1, 2, \ldots, n\}$, with $n \geq 2$. A collection $F = \{A_1, A_2, \ldots, A_t\}$ of subsets of N is said to be *separating*, if for every pair $\{x, y\}$ there is a set $A_i \in F$ so that $A_i \cap \{x, y\}$ contains just one element. F is said to be *covering*, if every element of N is contained in at least one set $A_i \in F$. What is the smallest value of t in terms of n so that there is a family $F = \{A_1, A_2, \ldots, A_t\}$ which is simultaneously separating and covering?

Answer:
Associate each element m of N with a binary string $x_1x_2\ldots x_t$, where $x_i = 1$ if m is in set A_i and 0 if m is not in A_i. The condition that F is separating simply means that distinct elements of N will be mapped to distinct binary strings. The condition that F is covering means that no element of N will be mapped to (0, 0, …, 0).

Thus we have n distinct binary strings of length t, none of which is the all 0 string. This implies $n \leq 2^t-1$. Conversely, if we indeed have $n \leq 2^t-1$, then a construction is easy by reversing the above process: first label each element with a different binary string and then place it into the appropriate sets. Thus $n \leq 2^t-1$ is necessary and sufficient, so $t = \lfloor \log_2 n \rfloor+1$ is the answer. ■

Remark: This idea of associating elements with binary strings is more than just a useful trick on Olympiads – in fact, it plays an important role in a whole branch of combinatorics known as algebraic combinatorics, where these binary “strings” are actually treated as vectors. Algebraic manipulations of these vectors (which often take place mod 2 or in some other field) can produce surprising combinatorial results.

Example 3 [IMO 2006-2]
A diagonal of a regular 2006-gon is called *odd* if its endpoints divide the boundary into two parts, each composed of an odd number of sides. Sides of the 2006-gon are also regarded as odd diagonals. Suppose the 2006-gon has been dissected into triangles by 2003 nonintersecting diagonals. Define a *good triangle* as an isosceles triangles with two odd sides. Find the maximum number of good triangles.

Answer:
Note a good triangle has two odd sides and an even side; hence the pair of equal sides must be the odd sides.

Experimentation with 2006 replaced by small even numbers hints that the general answer for a regular $2n$-gon is n. This is attainable by drawing all diagonals of the form $A_{2k}A_{2k+2}$, where A_1, A_2, …, A_{2n} are the vertices of the $2n$-gon. Now we show this is indeed the maximum.

Consider a $2n$-gon $\boldsymbol{P}$. To simplify notation, draw the circumcircle of $\boldsymbol{P}$. For a side AB in a triangle ABC, "arc AB" will denote the arc of the circumcircle not containing C. Arc AB is a "good arc" if AB is odd in a good triangle ABC.

Our basic idea is to construct a mapping f from sides of $\boldsymbol{P}$ to good triangles such that each good triangle is mapped to by at least 2 sides, and no side is mapped to more than one good triangle. This will immediately imply the result.

Consider a side XY of $\boldsymbol{P}$. Let AB denote the smallest good arc containing vertices X and Y, if it exists. (Note that {A, B} may be equal to {X, Y}.) Let C be the third vertex of the good triangle ABC. Then we will map XY to ABC: f(XY) = ABC.

All we need to show is that each good triangle is mapped to by at least two sides. In fact, for a good triangle DEF, with DE and EF odd, we will show that at least one side of ***P*** with vertices in arc DE is mapped to triangle DEF; the same argument will hold for EF and we will hence have two sides mapped to DEF.

Suppose to the contrary that no side with vertices in arc DE is mapped to DEF. Consider some side RS of ***P*** with vertices R and S in arc DE. Let *f*(RS) = D'E'F' for some D', E', F' lying on arc DE with D'E' and E'F' the odd sides of good triangle D'E'F'. Then by symmetry, the side of ***P*** that is the reflection of RS across the perpendicular bisector of D'F' will also be mapped to D'E'F'.

In this manner, sides in arc DE can be paired up, with each pair of sides being mapped to the same triangle. But there are an odd number of sides in arc DE, so they cannot all be paired up. Contradiction. ■

The Alternating Chains Technique

(Yes, I made that name up)

The basic idea that we will use in some form or the other for the next few problems is a simple consequence of the pigeonhole principle. Suppose you have n points on a line, and you are allowed to mark some of them such that no consecutive points are marked. Then the maximum number of points you can mark is $\lceil n/2 \rceil$, and this can be achieved by marking alternate points. If the n points were on a circle and not a line segment, then the maximum would be $\lfloor n/2 \rfloor$. These obvious statements can be cleverly applied in several combinatorial settings.

Example 4 [IMO Shortlist 1990]
Let $n \geq 3$ and consider a set E of $2n$-1 distinct points on a circle. Suppose that exactly k of these points are to be colored red. Such a coloring is *good* if there is at least one pair of red points such that

the interior of one of the arcs between them contains exactly n points from E. Find the smallest value of k so that every such coloring of k points of E is good.

Answer:
Let j be maximum number of colored points a *bad* coloring can have. Then $k = j+1$, so it suffices to find j.

Let the points be $A_1, A_2, \ldots, A_{2n-1}$. Join vertices A_i and A_{i+n+1} by an edge for each i (indices modulo $2n$-1). This decomposes E into disjoint cycles. The coloring is bad if and only if no two red vertices are joined by an edge. In other words, no two consecutive vertices in the cycle are both red. How many cycles are there? Using elementary number theory, it is easy to show that the number of cycles is equal to gcd $(n+1, 2n-1)$.

Since $2n-1 = 2(n+1)-3$, *gcd* $(n+1, 2n-1) = 3$ if $n+1$ is divisible by 3, and *gcd* $(n+1, 2n-1) = 1$ otherwise. If *gcd* $(n+1, 2n-1) = 1$, we get only one cycle containing all $2n$-1 points. Then $j = n-1$ by our earlier discussion. Hence $k = j+1 = n$.

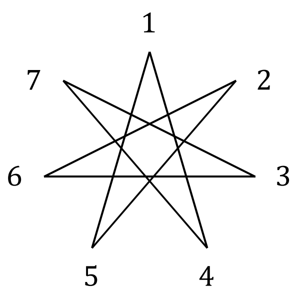

Figure 7.1.

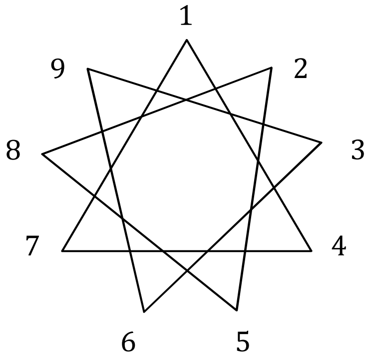

Figure 7.2.

If $gcd(n+1, 2n-1) = 3$, then we get 3 cycles, each containing $\frac{2n-1}{3}$ vertices. Each cycle then can have at most $\lfloor(2n-1)/6\rfloor$ red points in a bad coloring. Thus at most $3\lfloor(2n-1)/6\rfloor$ points in total can be colored in a bad coloring. Hence $k = 3\lfloor(2n-1)/6\rfloor+1$ if 3 divides $n+1$, and $k = n$ otherwise. ■

Example 5 [USAMO 2008, Problem 3]

Let n be a positive integer. Denote by S_n the set of points (x, y) with integer coordinates such that $|x| + \lfloor y+½\rfloor \leq n$. A path is a sequence of distinct points $(x_1, y_1), (x_2, y_2), \ldots, (x_k, y_k)$ in S_n such that, for $i = 1, 2, \ldots, k-1$, the distance between (x_i, y_i) and (x_{i+1}, y_{i+1}) is 1 (in other words, the points (x_i, y_i) and (x_{i+1}, y_{i+1}) are neighbors in the lattice of points with integer coordinates). Prove that the points in S_n cannot be partitioned into fewer than n paths (a partition of S_n into m paths is a set P of m nonempty paths such that each point in S_n appears in exactly one of the m paths in P).

Answer:

Color the points of each row of S_n alternately red and black, starting and ending with red. Any two neighboring points are of opposite color, unless they are from the middle two rows.

Consider a partition of S_n into m paths, for some m. For each of the m paths, split the path into two paths wherever there are consecutive red points. Now no path has consecutive red points. Further, since there are n pairs of consecutive red points in S_n (from the middle two rows), we have split paths at most n times. Thus we now have at most $m+n$ paths (each split increases the number of paths by one).

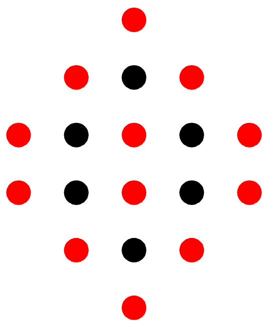

Figure 7.3. Example for n = 3

Now, there are $2n$ more red points than black points in S_n, but each of the $m+n$ paths contains at most one more red point than black point (since no path contains consecutive red points). Thus we obtain $m+n \geq 2n$, or $m \geq n$, proving the result. ■

In the next example, in order to use the chain decomposition idea we need some starting point, that is, some information or assumption that lends enough structure to the problem for us to exploit. To this end, we use the extremal principal by considering the smallest element in a certain set.

Example 6 [USAMO 2009-2]

Let n be a positive integer. Determine the size of the largest subset S of $N = \{-n, -n+1, \ldots, n-1, n\}$ which does not contain three elements

(a, b, c) (not necessarily distinct) satisfying $a + b + c = 0$.

Answer:
Obviously 0 is not in S since $0+0+0 = 0$. We claim the answer is n if n is even and $(n+1)$ if n is odd. These bounds can clearly be achieved by taking all the odd numbers in N, since the sum of three odd numbers can never be 0. To show this is the maximum, let j be the element of smallest absolute value in S (if both j and $-j$ are present, consider the positive one). Assume WLOG that $j > 0$, and let T denote the set of elements with absolute value at least j. Note that all elements of S are in T.

Case 1: $(-j)$ is not in S. Consider the pairs $(j, -2j)$, $(j+1, -(2j+1))$, $(j+2, -(2j+2))$, …., $(n-j, -n)$. In each of these pairs the sum of the numbers is $(-j)$, so at most one of the two elements is in S (otherwise the sum of the two elements plus j would be 0). There are exactly $(n-2j+1)$ pairs, so at most $(n-2j+1)$ of the paired numbers are in S.

Furthermore, there are exactly $2j-1$ unpaired numbers in $T \setminus \{-j\}$: j positive unpaired numbers (namely n, $n-1$, …, $n-j+1$), and $(j-1)$ negative unpaired numbers (namely $-(j+1)$, $-(j+2)$, …, $-(2j-1)$). Thus the maximum number of elements in S is $(2j-1) + (n-2j+1) = n$.

Case 2: $(-j)$ is also in S. Now we use the chain decomposition idea. If a and b are elements in T, then a is joined to b by an edge if and only if $a+b = j$ or $a+b = -j$. This ensures that no two elements joined by an edge can both be in S. Each element x in T is joined to at least one and at most two other elements of T ($(j-x)$ and $(-j-x)$). Hence the elements of T have been partitioned into disjoint chains.

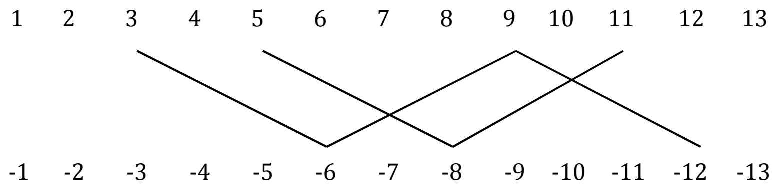

Figure 7.4: n = 13, j = 3. There are 6 chains, of which 2 are shown.

Now the rest is just counting. There are exactly $2j$ chains, namely the chains starting with $\pm j$, $\pm(j+1)$, ..., $\pm(2j-1)$. Let the lengths of these chains be l_1, l_2, ..., l_{2j}. From a chain of length l_i, at most $\left\lceil \frac{l_i}{2} \right\rceil$ elements can be in S. Thus the total number of elements we can take is at most

$$\sum_{i=1}^{2j} \left\lceil \frac{l_i}{2} \right\rceil \leq \sum_{i=1}^{2j} \frac{l_i+1}{2} = j + \frac{\sum_{i=1}^{2j} l_i}{2} = j + \frac{2n-2(j-1)}{2} = n+1.$$

Here we used $\sum_{i=1}^{2j} l_i = 2n - 2(j-1)$ since both sides are equal to the number of elements of absolute value at least j.

Now we are done if n is odd. If n is even, we need to tighten our bound to n. For this we note that the inequality $\sum_{i=1}^{2j} \left\lceil \frac{l_i}{2} \right\rceil \leq \sum_{i=1}^{2j} \frac{l_i+1}{2}$ is strict if there is a chain of even length, since $\left\lceil \frac{x}{2} \right\rceil$ is strictly less than $\frac{x+1}{2}$ for even x. It now suffices to prove the existence of a chain of even length if n is even; this is pretty simple and is left to the reader. ■

Two problems on boards

We now look at two problems based on $n \times n$ boards that initially look quite similar but are actually very different. Through these two problems we will demonstrate two important ways of thinking about and exploiting the structure of boards.

The following example uses the idea of examining individual objects' contributions toward some total. We saw a similar idea in examples 5 and 6 from chapter 3.

Example 7 [USAMO 1999-1]
Some checkers placed on an $n \times n$ checkerboard satisfy the following conditions:

a) Every square that does not contain a checker shares a side with one that does;
b) Given any pair of squares that contain checkers, there is a sequence of squares containing checkers, starting and ending with the given squares, such that every two consecutive squares of the sequence share a side.

Prove that at least $\frac{n^2-2}{3}$ checkers have been placed on the board.

Answer:
Suppose we have an empty board, and we want to create an arrangement of k checkers satisfying (a) and (b). Call a square *good* if it contains a checker or shares a side with a square containing a checker. By (a), every square must eventually be *good*. Let us place the checkers on the board as follows: place one checker on the board to start, and then in each step place one checker adjacent to one that has already been placed. Since any arrangement of checkers that satisfies the problem must be

connected by (b), we can form any arrangement of checkers in this manner.

In the first step at most 5 squares become *good* (the square we placed the checker on and its neighbors). In each subsequent step, at most 3 squares that are **not** already *good* become *good*: the square we just put a checker on and the square next to it are already good, leaving 3 neighbors that could become *good*. Thus, at the end of placing k checkers, at most $5+3(k-2) = 3k+2$ squares are good. But we know all n^2 squares are *good* at the end, so $n^2 \leq 3k+2$, proving the result. ■

Remark 1: Initially the problem appears difficult due to the fact that a given square may be *good* due to more than one checker. This makes it hard to calculate "individual contributions", that is, the number of squares that are good because of a certain checker. We get around this problem by imagining the k checkers being *added sequentially*, rather than simply "being there." This allowed us to measure the "true contribution" of a checker by not counting its neighbors that are already *good*. This was just a simple example, but the ideas of introducing an element of time and adopting a dynamic view of a static problem have powerful applications in combinatorics and algorithms.

Remark 2: The only property of the board we are using is that any square has at most 4 neighboring squares. Thus we can actually think of the board as a graph with vertices representing squares and two vertices being connected if and only if they correspond to adjacent squares on the board. This interpretation of problems involving $n \times n$ boards is often very useful, and we have already used this idea in example 3 of chapter 3. With the present problem, we can easily generalize to graphs with maximum degree Δ.

At first glance, the next example looks very similar to the previous one. However, it is significantly more difficult, and the solution uses a clever coloring. Coloring is another extremely important way of exploiting the structure of boards.

Example 8 [IMO 1999, Problem 3]
Let n be an even positive integer. We say that two different cells of an $n \times n$ board are **neighboring** if they have a common side. Find the minimal number of cells on the $n \times n$ board that must be marked so that any cell (marked or not marked) has a marked neighboring cell.

Answer:
Let $n = 2k$. Color the board black and white in layers as shown in figure. Note that any square (black or white) neighbors exactly two black squares. Hence, since the number of black squares is $2k(k+1)$, we must mark at least $k(k+1)$ squares. On the other hand, this bound can be achieved by marking alternate black squares in each layer, in such a way that each white cell neighbors exactly one marked black square. ■

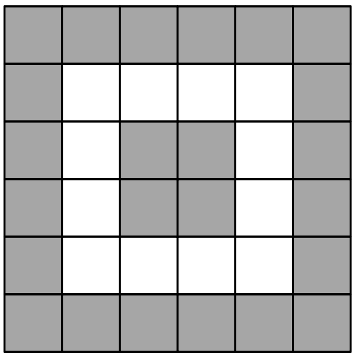

Figure 7.5: Example for n = 6.

The Classification Method

Suppose we have a set S of objects, and we want to show that there exists a large subset S' of these objects such that S' satisfies a particular condition. The idea behind the classification method is to partition (split) S into sets $S_1, S_2, \ldots, S_k$ such that $S_1 \cup S_2 \cup \ldots \cup S_k = S$, and furthermore, each of the sets $S_1, S_2, \ldots, S_k$ satisfies the given condition. Then, by the pigeonhole principle, at least one of these sets will have size at least $|S|/k$, thereby proving the existence of a subset of S of size at least $|S|/k$ satisfying the given condition.

Example 9 [IMO Shortlist 2001, C6]
For a positive integer n define a sequence of zeros and ones to be *balanced* if it contains n zeros and n ones. Two balanced sequences a and b are *neighbors* if you can move one of the $2n$ symbols of a to another position to form b. For instance, when $n = 4$, the balanced sequences 01101001 and 00110101 are neighbors because the third (or fourth) zero in the first sequence can be moved to the first or second position to form the second sequence. Prove that there is a set S of at most $\frac{1}{n+1}\binom{2n}{n}$ balanced sequences such that every balanced sequence is equal to or is a neighbor of at least one sequence in S.

Answer:
Call such a set S a *dominating set*. Our idea is to partition the set of $\binom{2n}{n}$ balanced sequences into (n+1) classes, so that the set of sequences in any class form a dominating set. Then we will be done by the pigeonhole principle, since some class will have at most $\frac{1}{n+1}\binom{2n}{n}$ balanced sequences.

To construct such a partition, for any balanced sequence A let $f(A)$ denote the sum of the positions of the ones in A (mod $(n+1)$). For example, $f(100101) \equiv 1+4+6 \pmod 4 \equiv 3 \bmod 4$. A sequence is in class i if and only if $f(A) \equiv i \bmod (n+1)$. It just remains to show that every class is indeed a dominating set, that is, for any class C_i and any balanced sequence A not in C_i, A has a neighbor in C_i.

This isn't difficult: if A begins with a one, observe that moving this one immediately to the right of the kth zero gives a sequence B satisfying $f(B) \equiv f(A) + k \bmod (n+1)$. Hence simply choose $k \equiv I - f(A) \bmod (n+1)$, and then by shifting the first one to the right of the kth zero we end up with a sequence B satisfying $f(B) \equiv i \bmod (n+1)$. Hence B is a sequence in C_i. The case when A begins with a zero is similar. Thus each class is indeed a dominating set and we are done by the first paragraph. ■

Remark: It is worth mentioning that the reason for naming S a "dominating set" is that this problem has a very nice graph theoretic interpretation. Dominating sets are graph structures that we will encounter in chapters 8 and 9.

Two Graph Theory Problems

Graphs are a rich source of extremal problems. We present two here, and you will see several more in the next chapter. The main thing to keep in mind when dealing with such problems is dependencies – how are all the quantities in question related? How does the size of one impact the size of another? Indeed, the entire purpose of a graph is to model connections and dependencies, so the structure of a graph invariably proves useful for answering these types of extremal questions.

Example 10 [Swell coloring]
Let K_n denote the complete graph on n vertices, that is, the graph with n vertices such that every pair of vertices is connected by an edge. A *swell coloring* of K_n is an assignment of a color to each of

the edges such that the edges of any triangle are either all of distinct colors or all the same color. Further, more than one color must be used in total (otherwise trivially if all edges are the same color we would have a swell coloring). Show that if K_n can be swell colored with k colors, then $k \geq \sqrt{n} + 1$.

Answer:
Let $N(x, c)$ denote the number of edges of color c incident to a vertex x. Fix x_0, c_0 such that $N(x_0, c_0)$ is maximum and denote this maximum value by N. There are $n-1$ edges incident to x_0, colored in at most k colors, with no color appearing more than N times. Hence $Nk \geq n-1$.

Now consider vertex y such that edge yx_0 is **not** of color c_0. Also let $x_1, x_2, \ldots, x_N$ be the N vertices joined to x_0 by color c_0. Note that for any i and j in $\{1, 2, \ldots, N\}$, x_ix_j is of color c_0 since x_0x_i and x_0x_j are of color c_0.

Suppose yx_i is of color c_0 for some i in $\{1, 2, \ldots, N\}$. Then triangle yx_ix_0 contradicts the swell coloring condition, since two sides (yx_i and xx_i) are the same color c_0 but the third side isn't. Hence the color of yx_i is not c_0 for $i = 1, 2, \ldots, N$.

Now suppose yx_i and yx_j are the same color for some distinct i and j in $\{0, 1, 2, \ldots, N\}$. Then x_ix_j also must be this color. But x_ix_j is of color c_0, which implies yx_i and yx_j are also of color c_0, contradicting our earlier observation.

It follows that $yx_0, yx_1, yx_2, \ldots, yx_N$ are all different colors, and none of them is c_0. This implies that there are at least $N+2$ distinct colors, so $k \geq N+2$. Since we already showed $Nk \geq n-1$, it follows that $k(k-2) \geq n-1$, from which the desired bound follows. ■

Remark: Basically, one can think about the above proof as follows: either there is a big clique of one color, or there isn't. If there isn't, then we need many colors to avoid big monochromatic cliques. If there is, then anything outside this clique needs many different colors to be connected to the clique.

Example 11 [Belarus 2001]
Given n people, any two are either friends or enemies, and friendship and enmity are mutual. I want to distribute hats to them, in such a way that any two friends possess a hat of the same color but no two enemies possess a hat of the same color. Each person can receive multiple hats. What is the minimum number of colors required to always guarantee that I can do this?

Answer:
Set up a graph in the usual way, with vertices standing for people and edges between two people if and only if they are friends. Note that if we have a complete bipartite graph with $\lfloor n/2 \rfloor$ vertices on one side and $\lceil n/2 \rceil$ vertices on the other, then we need at least $\lfloor n/2 \rfloor \lceil n/2 \rceil$ colors. This is because we need one for each pair of friends and no color could belong to more than two people (otherwise some two people on the same side of the bipartition would have the same color, which is not possible since they are enemies). We claim this is the worst case, that is, given $\lfloor n/2 \rfloor \lceil n/2 \rceil$ colors we can always satisfy the given conditions. We will use strong induction, the base cases n = 1, 2, 3 being easy to check.

Obviously if the graph has fewer than $\lfloor n/2 \rfloor \lceil n/2 \rceil$ edges we are done, since we can assign a separate color for each pair of friends. Now if the graph has more than $\lfloor n/2 \rfloor \lceil n/2 \rceil$ edges, then by the contrapositive of corollary 1 in chapter 6 (example 8), the graph contains a triangle. Use one color for the triangle (that is, give each member of the triangle a hat of that color). Using at most n–3 colors, we can ensure that each person not in the triangle who is friends with some member(s) of the triangle has a common color with them. Now among the remaining n–3 people, we need at most $\lfloor (n-3)/2 \rfloor \lceil (n-3)/2 \rceil$ more colors by the induction hypothesis. Hence in total we use at most

$(1+(n-3)+\lfloor (n-3)/2 \rfloor \lceil (n-3)/2 \rceil) \leq \lfloor n/2 \rfloor \lceil n/2 \rceil$ colors. ■

Remark: The important part of the previous example is guessing the worst case scenario. Intuitively, when there are too many or too few edges, we don't need many colors, because we would then have either very few enemies or very few friends. This leads us to guess that the worst case is "somewhere in the middle". In these cases, bipartite graphs should be your first suspects (followed by multipartite graphs). This gives us the intuition needed to complete the solution- the simplest structure we can exploit in a graph with "too many edges" is that it will have triangles.

Induction and Combinatorics of Sets

In this section we will use induction to solve extremal problems on sets. We first establish two simple lemmas.

Lemma 7.1: Let S be a set with n elements, and let F be a family of subsets of S such that for any pair A, B in F, $A \cap B \neq \emptyset$. Then $|F| \leq 2^{n-1}$.

Proof: For any set A in F, the complement of A, that is, $S \backslash A$, cannot be in F. So at most ½ of the total number of subsets of S can be in F. Equality is achieved by taking all subsets of S containing a fixed element x in S.

Lemma 7.2: Let S be a set with n elements, and let F be a family of subsets of S such that for any pair A, B in F, S is not contained by $A \cup B$. Then $|F| \leq 2^{n-1}$.

Proof: The proof is identical to that of lemma 7.1. Equality holds by taking all subsets of S excluding a fixed element x of S.

Example 12 [Iran TST 2008]

Let $S = \{1, 2, \ldots, n\}$, and let F be a family of 2^{n-1} subsets of S.

Suppose for all A, B and C in S, $A \cap B \cap C \neq \emptyset$. Show that there is an element in S belonging to all sets in F.

Answer

Suppose there are sets X and Y in F such that $|X \cap Y| = 1$. Then we are trivially done since all sets in F must contain this element by the intersection condition of the problem. Now assume that F is 2-intersecting; that is, $|X \cap Y| \geq 2$ for each X and Y in F. We prove by induction on n that in this case, $|F| < 2^{n-1}$, showing that this case cannot occur.

Base cases $n = 1$ and 2 are trivial, so assume the result holds for $n-1$ and that we are trying to prove it for n. We can write F as $F = F_n \cup F_{n-1}$, where F_n consists of all sets in F containing n and F_{n-1} consists of all sets in F **not** containing n.

By the induction hypothesis, $|F_{n-1}| < 2^{n-2}$. Now define $F_n' = \{S \backslash n \mid S \in F_n\}$. In other words, define F_n' is obtained from F_n by deleting n from all sets in F_n. Since F_n is 2-intersecting, F_n' is still intersecting, so by lemma 7.1, $|F_n'| \leq 2^{n-2}$. Note that $|F_n'| = |F_n|$, so we get $|F| = |F_n| + |F_{n-1}| = |F_n'| + |F_{n-1}| < 2^{n-2} + 2^{n-2} < 2^{n-1}$, as desired. ■

Example 13 [Kleitman's lemma] (U*)

A set family F is said to be *downwards closed* if the following holds: if X is a set in F, then all subsets of X are also sets in F. Similarly, F is said to be *upwards closed* if whenever X is a set in F, all sets containing X are also sets in F. Let F_1 and F_2 be downwards closed families of subsets of $S = \{1, 2, ..., n\}$, and let F_3 be an upwards closed family of subsets of S.

(a) Show that $|F_1 \cap F_2| \geq \frac{|F_1||F_2|}{2^n}$.

(b) Show that $|F_1 \cap F_3| \leq \frac{|F_1||F_3|}{2^n}$. .

Answer

a) Induct on n. The base case $n = 1$ is trivial, so assume the result holds for $(n-1)$ and that we are trying to prove it for n. Let X_1 be the family of sets in F_1 containing n and Y_1 be the family of sets not containing n. Delete n from each set in X_1 to obtain a new family X_1'. Note that $|X_1'| = |X_1|$ and $|X_1'|+|Y_1| = |F_1|$. Also note that X_1 and Y_1 are still downwards closed. Analogously define X_2, Y_2 and X_2' for the family F_2.

Observe that $|F_1 \cap F_2| = |Y_1 \cap Y_2| + |X_1' \cap X_2'| \geq \frac{|Y_1||Y_2|}{2^{n-1}} + \frac{|X_1'||X_2'|}{2^{n-1}}$ (applying the induction hypothesis). Since F_1 and F_2 are downwards closed, X_1' is a subset of Y_1, and similarly X_2' is a subset of Y_2. Hence $|Y_1| \geq |X_1'|$, $|Y_2| \geq |X_2'|$, so Chebyshev's inequality (or just basic algebra) yields $\frac{|Y_1||Y_2|}{2^{n-1}} + \frac{|X_1'||X_2'|}{2^{n-1}} \geq (|Y_1| + |X_1'|)(|Y_2| + |X_2'|)/2^n = \frac{|F_1||F_2|}{2^n}$.

b) The proof is similar to that in part (a), but with inequality signs reversed. ■

Lemmas (i) and (ii) in this section are fairly straightforward. Note that in both cases, the imposition of a certain constraint decreases the number of sets we can have by a factor of ½ (we can include 2^{n-1} sets out of 2^n total possibilities). A natural question is what happens if we impose both conditions simultaneously – will the number of possible sets in the family be decreased by a factor of 4? Interestingly, Kleitman's lemma answers this question in the affirmative.

Example 14 (U*)

Let $F = \{A_1, A_2, \ldots, A_k\}$ be a family of subsets of $S = \{1, 2, \ldots, n\}$ ($n > 2$), such that for any distinct subsets A_i and A_j, $A_i \cap A_j \neq \emptyset$ and $A_i \cup A_j \neq S$. Show that $k \leq 2^{n-2}$.

Answer

F can be extended to a downward closed system D by adding all subsets of the sets in F. Similarly, F can be extended to an upward closed system U by adding all subsets of S that contain some set in

F. Note that $F = U \cap D$. Since F is intersecting, so is U (since in creating U we only added "big" sets). Hence by lemma 7.1, $|U| \le 2^{n-1}$. Similarly, since the union of no two sets in F covers S, the same holds for D. Hence by lemma 7.2, $|D| \le 2^{n-1}$. Then by part (b) of the previous problem,
$k = |F| = |U \cap D| \le |U||D|/2^n \le 2^{n-1} \times 2^{n-1}/2^n = 2^{n-2}$. ■

Example 15 [stronger version of USA TST 2011]
Let $n \ge 1$ be an integer, and let S be a set of integer pairs (a, b) with $1 \le a < b \le 2^n$. Assume $|S| > n2^n$. Prove that there exist four integers $a < b < c < d$ such that S contains all three pairs (a, c), (b, d) and (a, d).

Answer:
We induct on n. The base cases being trivial, suppose the result holds for $(n-1)$. Let S' be the set of pairs (a, b) in S with $a < b \le 2^{n-1}$. If $|S'| \ge (n-1)2^{n-1}$, we would be done by applying the induction hypothesis to S'.

Similarly, let S'' be the set of pairs (a, b) with $2^{n-1} < a < b$. If $|S''| \ge (n-1)2^{n-1}$ we would again be done by applying the induction hypothesis to S'', treating the pair (a, b) as if it were the pair $(a-2^{n-1}, b-2^{n-1})$. (Take a moment to fully understand this.)

Now suppose neither of these cases arises. Then more than $(n2^n-2(n-1)2^{n-1}) = 2^n$ pairs (a, b) would have to satisfy $a \le 2^{n-1} < b$.

We call a pair (a, b) in S a *B-champion* if $a \le 2^{n-1} < b$, and b is the smallest number greater than 2^{n-1} with which a occurs in a pair. Note that there is at most 1 *B-champion* for fixed a, and at most 2^{n-1} choices for a. Thus there are at most 2^{n-1} *B-champions*. Similarly, define an *A-champion* to be a pair (a, b) in S if $a \le 2^{n-1} < b$

such that a is the largest number less than or equal to 2^{n-1} with which B is paired. The same argument shows that there are at most 2^{n-1} *A-champions*.

Since there are more than 2^n pairs (a, b) with $a \leq 2^{n-1} < b$, at least one of these pairs, say (x, y) is neither an A-champion nor a B-champion. Then there must exist z such that $2^{n-1} < z < y$ and (x, z) is in S (since (x, y) is not a B- champion). Similarly, there exists w such that $x < w \leq 2^{n-1}$ and (y, w) is in S. Hence $x < w < z < y$, and (x, z), (w, y) and (x, y) are all in S, proving the statement of the problem. The induction step, and hence the proof, is complete. ■

Remark 1: Observe the structure of this solution- we first tried to find a suitable subset of S to which we could apply the induction hypothesis; that is, we tried to break the problem down. We then solved the problem for the cases in which this didn't work. By dealing with an easy case of a problem first, we do more than just get the easy case out of the way. We actually learn some important conditions a case must satisfy to *not* be easy- and this information is crucial for handling the hard case. This subtly illustrates the following problem solving tenet - *the first step in solving a hard problem often lies in identifying what makes the problem hard.*

Remark 2: The definition of *A-champions* and *B-champions* in the above solution initially appears to come out of nowhere. However, after carefully reading the whole solution, the purpose behind it becomes clear: we are trying to find a pair (x, y) with w and z "squished between" x and y such that (x, z) and (w, y) are pairs in S. The only way this can happen is if (x, y) is neither an *A-champion* nor a *B-champion*.

Example 16 [The Sunflower Lemma]

A sunflower with k petals and a core X is a family of sets $S_1, S_2, \ldots, S_k$ such that $S_i \cap S_j = X$ for each $i \neq j$. (The reason for the name is that the Venn diagram representation for such a family resembles a sunflower.) The sets $S_i \setminus X$ are known as petals and must be nonempty, though X can be empty. Show that if F is a family of sets of cardinality s, and $|F| > s!(k\text{-}1)^s$, then F contains a sunflower with k petals.

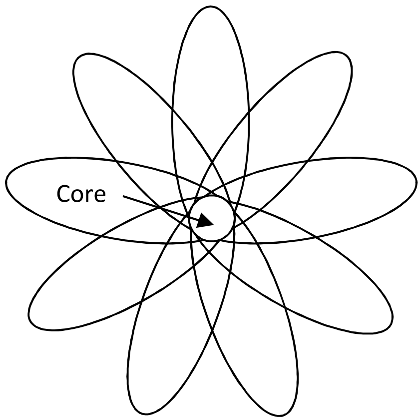

Figure 7.6: A Sunflower.

Answer:
Induct on s. The result is trivial for $s = 1$, since then k singleton sets will form the sunflower (its core will be empty but that's okay, no one said the sunflower has to be pretty). Passing to the inductive step, let $s \geq 2$ and take a maximal family $A = \{A_1, A_2, \ldots, A_t\}$ of pairwise disjoint sets in F. If $t \geq k$ we are done, since this family will form our required sunflower (with empty core). Now suppose $t \leq k - 1$ and let $B = A_1 \cup A_2 \cup \ldots \cup A_t$. Note that $|B| = st \leq s(k-1)$. Also, by the maximality of A, it follows that B intersects all sets in F (otherwise we could add more sets to A). Hence by the pigeonhole principle, some element x of B must be contained in at least

$$\frac{|F|}{|B|} > \frac{s!(k-1)^s}{s(k-1)} = (s-1)!(k-1)^{s\text{-}1}$$

sets in F. Deleting x from these sets and applying the induction hypothesis to these sets (which now contain s–1 elements each), we see that there exists a sunflower with k petals. Adding x back to all these sets doesn't destroy the sunflower since it just goes into the core, so we get the desired sunflower. ■

Example 17 [IMO Shortlist 1998, C4]
Let $\boldsymbol{U} = \{1, 2, \ldots, n\}$, where $n \geq 3$. A subset $\boldsymbol{S}$ of $\boldsymbol{U}$ is said to be *split* by an arrangement of the elements of $\boldsymbol{U}$ if an element not in $\boldsymbol{S}$ occurs in the arrangement somewhere between two elements of $\boldsymbol{S}$. For example, 13542 splits {1, 2, 3} but not {3, 4, 5}. Prove that for any family $\boldsymbol{F}$ of $(n-2)$ subsets of $\boldsymbol{U}$, each containing at least 2 and at most $n-1$ elements, there is an arrangement of the elements of $\boldsymbol{U}$ which splits all of them.

Answer:
We induct on n. As always, we unceremoniously dismiss the base case as trivial and pass to the induction step, assuming the result for $n-1$ and prove it for n. We first prove a claim.

Claim: There exists an element a in $\boldsymbol{U}$ that is contained in all subsets of $\boldsymbol{F}$ containing $n - 1$ elements, but in at most one of the 2-element subsets.
Proof: A simple counting argument suffices. Let $\boldsymbol{F}$ contain k $(n\text{-}1)$-element subsets and m 2-element subsets. Note that $k+m$ is at most the total number of subsets in $\boldsymbol{F}$, which is $n-2$. Hence $(k+m) \leq (n-2)$. The intersection of the k $(n-1)$-element subsets contains exactly $(n-k)$ elements. This is because for each of these subsets there is exactly one element it doesn't contain.

But $n-k \geq m+2$ and at most m elements can be in more than one of the two-element sets. Thus one of these elements that is in the intersection of all the $(n-1)$-element subsets is in at most one of the 2-element sets, proving the claim.

Now let $\boldsymbol{A}$ be the 2-element subset that contains a if it exists; otherwise let it be an arbitrary subset of $\boldsymbol{F}$ containing a. Now exclude a from all subsets in $\boldsymbol{F} \setminus \boldsymbol{A}$. We get at most $n-3$ subsets of $\boldsymbol{U} \setminus \{a\}$ containing at least 2 and at most $n-2$ elements. Applying the inductive hypothesis, we can arrange the elements of $\boldsymbol{U} \setminus a$ so as to split all subsets of $\boldsymbol{F} \setminus \boldsymbol{A}$. Replace a anywhere away from $\boldsymbol{A}$ and we are done. ■

Exercises

1. **[Generalization of USAMO 1999, Problem 4]**
   Find the smallest positive integer m such that if m squares of an $n \times n$ board are colored, then there will exist 3 colored squares whose centers form a right triangle with sides parallel to the edges of the board.

2. **[Erdos-Szekeres Theorem] (U*)**
   Show that any sequence of n^2 distinct real numbers contains a subsequence of length n that is either monotonically increasing or monotonically decreasing.

3. **[USA TST 2009, Problem 1]**
   Let m and n be positive integers. Mr. Fat has a set $\boldsymbol{S}$ containing every rectangular tile with integer side lengths and area of a power of 2. Mr. Fat also has a rectangle R with dimensions $2^m \times 2^n$ but with a 1×1 square removed from one of the corners. Mr. Fat wants to choose $(m+n)$ rectangles from $\boldsymbol{S}$, with respective areas 2^0, 2^1, …, 2^{m+n-1} and then tile R with the chosen rectangles. Prove that this can be done in at most $(m+n)!$ ways.

4. **[Generalization of APMO 2012, Problem 2]**
   Real numbers in [0, 1] are written in the cells of an $n \times n$ board. Each gridline splits the board into two rectangular parts. Suppose that for any such division of the board into two parts along a gridline, at least one of the parts has weight at most 1, where the weight of a part is the sum of all numbers written in

cells belonging to that part. Determine the maximum possible sum of all the numbers written on the board.
[Challenge: generalize to k dimensional boards.]

5. **[Indian postal coaching 2011]**
Consider 2011^2 points arranged in the form of a 2011 x 2011 grid. What is the maximum number of points that can be chosen among them so that no four of them form the vertices of either an isosceles trapezium or a rectangle whose parallel sides are parallel to the grid lines?

6. **[China Girls Math Olympiad 2004]**
When the unit squares at the 4 corners are removed from a 3×3 square, the resulting shape is called a cross. Determine the maximum number of non-overlapping crosses that can be placed within the boundary of a 10×11 board.

7. **[IMO Shortlist 2010, C2]**
Let $n > 3$ be a positive integer. A set of n distinct binary strings of length n is said to be *diverse* if there exists an n×n array whose rows are these n binary strings in some order, and all entries along the main diagonal of this array are equal. Find the smallest integer m, such that among any m binary strings of length n, there exist n strings forming a diverse set.

8. **[Iran TST 2007]**
Let $\boldsymbol{A}$ be the largest subset of $\{1, 2, \ldots, n\}$ such that for each $x \in \boldsymbol{A}$, x divides at most one other element in $\boldsymbol{A}$. Show that $2n/3 \leq |\boldsymbol{A}| \leq \left\lceil \frac{3n}{4} \right\rceil$.

9. **[IMO 2014, Problem 2]**
Let n be a positive integer, and consider an n×n board. Suppose some rooks are placed on this board such that each row contains exactly one rook and each column contains exactly one rook. Find the largest integer k such that for any such configuration as described above, there necessarily

exists a $k \times k$ square which does not contain a rook on any of its k^2 squares.

10. [IMO 2013, Problem 2]

A configuration of 4027 points in the plane is called Colombian if it consists of 2013 red points and 2014 blue points, and no three points of the configuration are collinear. By drawing some lines, the plane is divided into several regions. An arrangement of lines is *good* for a Colombian configuration if the following conditions are satisfied:

(i) No line passes through any point of the configuration

(ii) No region contains points of both colors

Find the least value of k such that for any Colombian configuration of 4027 points, there is a good arrangement of k lines.

11. [IMO Shortlist 2005, C3]

Consider an $m \times n$ rectangular board consisting of mn unit squares. Two of its unit squares are called adjacent if they have a common edge, and a path is a sequence of unit squares in which any two consecutive squares are adjacent. Two paths are called non-intersecting if they don't share any common squares. Suppose each unit square of the rectangular board is colored either black or white.

Let N be the number of colorings of the board such that there exists at least one black path from the left edge of the board to its right edge. Let M be the number of colorings of the board for which there exist at least two non-intersecting black paths from the left edge of the board to its right edge.

Prove that $N^2 \geq M \times 2^{mn}$.

12. [Balkan Math Olympiad 1994]

Find the smallest number $n \geq 5$ for which there can exist a set of n people, such that any two people who are friends have no common friends, and any two people who are not friends have exactly two common acquaintances.

13. Given a set ***S*** of n points in 3-D space, no three on a line, show that there exists a subset ***S'*** of ***S*** containing at least $n^{1/4}$ points, such that no subset of the points in ***S'*** form a regular polygon.

14. [IMO Shortlist 2000, C4]
Let n and k be positive integers such that $n/2 < k \le 2n/3$. Find the smallest number m for which it is possible to mark m squares on an $n \times n$ board such that no row or column contains a block of k adjacent unoccupied squares.

15. [IMO Shortlist 1988]
The code for a safe is a three digit number with digits in {1, 2, …, 8}. Due to a defect in the safe, it will open even if the number we enter matches the correct code in two positions. (For example, if the correct code is 245 and we enter 285, it will open.) Determine the smallest number of combinations that must be tried in order to guarantee opening the safe.

16. [Bulgaria 1998]
Let n be a given positive integer. Determine the smallest positive integer k such that there exist k binary (0-1) sequences of length $2n+2$, such that any other binary sequence of length $2n+2$ matches one of the k binary sequences in at least $n+2$ positions.

17. [IMO Shortlist 1996, C3]
Let k, m, n be integers satisfying $1 < n \le m-1 \le k$. Determine the maximum size of a subset ***S*** of the set {1, 2, …, k} such that no n distinct elements of ***S*** add up to m.

18. [IMO Shortlist 1988]
49 students took part in a math contest with three problems. Each problem was worth 7 points, and scores on each problem were integers from 0 to 7. Show that there exist two

students A and B such that A scored at least as many points as B on each of the three problems.

19. [USAMO 2007, Problem 3]

Let $\boldsymbol{S}$ be a set containing (n^2+n-1) elements, for some positive integer n. Suppose that the n-element subsets of $\boldsymbol{S}$ are partitioned into two classes. Prove that there are at least n pairwise disjoint sets in the same class.

20. [IMO Shortlist 2009, C6]

On a 999 × 999 board a limp rook can move in the following way: from any square it can move to any of its adjacent squares, that is, a square having a common side with it, and every move must be a turn, that is, the directions of any two consecutive moves must be perpendicular. A non-intersecting route of the limp rook consists of a sequence of pairwise different squares that the limp rook can visit in that order by an admissible sequence of moves. Such a non-intersecting route is called cyclic, if the limp rook can, after reaching the last square of the route, move directly to the first square of the route and start over.

How many squares does the longest possible cyclic, non-intersecting route of a limp rook visit?

21. [USAMO 2002, Problem 6]

Some trominoes (3×1 tiles) are to be placed on an $n \times n$ board without overlaps or trominoes sticking out of the board. Let $b(n)$ denote the minimum number of trominoes that must be placed so that no more can be placed according to the above rules. Show that there exist constants c and d such that $n^2/7 - cn \le b(n) \le n^2/5 + dn$.

22. [IberoAmerican Math Olympiad 2009, Problem 6]

6000 points on the circumference of a circle are marked and colored with 10 colors such that every group of 100 consecutive points contains all ten colors. Determine the smallest positive integer k such that there necessarily exists a group of k consecutive points containing all ten colors.

23. [IMO Shortlist 2011, C6]
Let n be a positive integer, and let $\boldsymbol{W} = \ldots x_{-1}\ x_0\ x_1 \ldots$ be an infinite periodic word, consisting of just letters a and/or b. Suppose that the minimal period N of $\boldsymbol{W}$ is greater than 2^n.
A finite nonempty word $\boldsymbol{U}$ is said to appear in $\boldsymbol{W}$ if there exist indices $k \le l$ such that $\boldsymbol{U} = x_k \ldots x_l$. A finite word U is called ubiquitous if the four words $\boldsymbol{Ua}$, $\boldsymbol{Ub}$, $\boldsymbol{aU}$ and $\boldsymbol{bU}$ all appear in $\boldsymbol{W}$. Prove that there are at least n ubiquitous finite nonempty words.

24. [IMO Shortlist 2007, C8]
Consider a convex polygon $\boldsymbol{P}$ with n vertices. A triangle whose vertices lie on vertices of $\boldsymbol{P}$ is called good if all its sides have equal length. Prove that there are at most $2n/3$ good triangles.

25. [Stronger version of IMO Shortlist 2008, C6]
For $n > 2$, let $\boldsymbol{S}_1$, $\boldsymbol{S}_2, \ldots, \boldsymbol{S}_{2^n}$ be 2^n subsets of $\boldsymbol{A} = \{1, 2, \ldots, 2^{n+1}\}$ that satisfy the following property: There do not exist indices a and b with $a < b$ and elements $x, y, z \in \boldsymbol{A}$ with $x < y < z$ and $y, z \in \boldsymbol{S}_a$ and $x, z \in \boldsymbol{S}_b$. Prove that at least one of the sets S_1 , $S_2, \ldots, S_{2^n}$ contains at most $2n+1$ elements. (Note: the original problem had a weaker bound of $4n$ instead of $2n+1$.)

26. [IMO Shortlist 2005, C8]
In a certain n-gon, some $(n-3)$ diagonals are colored black and some other $(n-3)$ diagonals are colored red, so that no two diagonals of the same color intersect strictly inside the polygon, although they can share a vertex. (Note: a side is not a diagonal.) Find the maximum number of intersection points between diagonals colored differently strictly inside the polygon, in terms of n.

27. [IMO Shortlist 2011, C7]

On a 2011×2011 square table we place a finite number of napkins that each cover a square of 52 by 52 cells. Napkins can overlap, and in each cell we write the number of napkins covering it, and record the maximal number k of cells that all contain the same nonzero number. Considering all possible napkin configurations, what is the largest value of k?

Olympiad Combinatorics

Pranav A. Sriram

August 2014

Copyright notices

About the Author

Pranav Sriram graduated from high school at The International School Bangalore, India, and is currently pursuing undergraduate studies at Stanford University.

8. GRAPH THEORY

Introduction

Graphs rule our lives: from Google Search to molecular sequencing, flight scheduling to Artificial Intelligence, graphs are the underlying mathematical abstraction fueling the world's most advanced technology. Graphs are also pervasive in models of speech, group dynamics, disease outbreaks and even the human brain, and as such play a crucial role in the natural and social sciences. Where does this versatility come from? Problems in several of the above-mentioned fields involve, at their core, entities existing in complex relationships with each other: hyperlinks between web pages, subject-object relationships between words, flights between cities and synapses between neurons. The power of graph theory stems from the simplicity and elegance with which graphs can model such relationships. Once you've cast a problem as a problem on graphs, you have at your disposal powerful machinery developed by mathematicians over the centuries. This is the power of abstraction.

This chapter is by no means an exhaustive reference on the subject – graph theory deserves its own book. However, we will see several powerful lemmas and techniques that underlie a vast majority of Olympiad and classical graph theory problems, and hopefully build plenty of graph theoretic intuition along the way.

In the final section of this chapter, we will leverage the power of graphs mentioned in the first paragraph to solve Olympiad problems that initially appear to have nothing to do with graphs.

We now recall some results we have proven in earlier chapters, that prove to be extremely useful both in Olympiad problems and in classical graph theory problems. We also advise the reader to go over their proofs again, because the proof techniques for these results are as important as the results themselves.

Some Useful Results

(i) In a graph ***G*** with n vertices, suppose no vertex has degree greater than Δ. Then one can color the vertices using at most Δ+1 colors, such that no two neighboring vertices are the same color. [Chapter 1, example 1]

(ii) In a graph ***G*** with V vertices and E edges, there exists an induced subgraph ***H*** with each vertex having degree at least E/V. (In other words, a graph with average degree d has an induced subgraph with minimum degree at least $d/2$) [Chapter 1, example 3]

(iii) Given a graph ***G*** in which each vertex has degree at least $(n-1)$, and a tree ***T*** with n vertices, there is a subgraph of ***G*** isomorphic to ***T***. [Chapter 2, example 3]

(iv) In a graph ***G***, if all vertices have degree at least δ, then there exists a path of length at least $\delta+1$. [Chapter 4, example 6]

(v) The vertex set ***V*** of a graph ***G*** on n vertices can be partitioned into two sets $\boldsymbol{V_1}$ and $\boldsymbol{V_2}$ such that any vertex in $\boldsymbol{V_1}$ has at least as many neighbors in $\boldsymbol{V_2}$ as in $\boldsymbol{V_1}$ and vice versa. [Chapter 4, example 8]

(vi) A *tournament* on n vertices is a directed graph such that for any two vertices u and v, there is either a directed edge from u to v or from v to u. A *Hamiltonian path* is a path passing through all the vertices. Every tournament has a Hamiltonian path. [Chapter 4, example 10]

More Useful Results and Applications

Dominating Sets

In a graph $\boldsymbol{G}$ with vertex set $\boldsymbol{V}$, a subset $\boldsymbol{D}$ of $\boldsymbol{V}$ is set to be a *dominating set* if every vertex v is either in $\boldsymbol{D}$ or has a neighbor in $\boldsymbol{D}$. The next lemma tells us that under certain simple conditions, there exists a fairly small dominating set.

Lemma 8.1: If $\boldsymbol{G}$ has no isolated vertices, then it has a dominating set of size at most $\frac{|V|}{2}$.

Proof: By **(v)**, there exists a bipartition $\boldsymbol{V} = \boldsymbol{V}_1 \cup \boldsymbol{V}_2$ so that every vertex in $\boldsymbol{V}_1$ has at least as many neighbors in $\boldsymbol{V}_2$ and vice versa. Since each vertex has degree at least 1, this implies that every vertex in $\boldsymbol{V}_1$ has at least one neighbor in $\boldsymbol{V}_2$ and vice versa. Thus both $\boldsymbol{V}_1$ and $\boldsymbol{V}_2$ are dominating sets. One of them has at most $\frac{|V|}{2}$ vertices and we are done.

Remark: In the next chapter we will show that if the minimum degree in an n-vertex graph $\boldsymbol{G}$ is $d > 1$, then $\boldsymbol{G}$ has a dominating set containing at most $n\,\frac{1+\ln(d+1)}{d+1}$ vertices.

Spanning Trees

Recall that a *spanning subgraph* of a graph $\boldsymbol{G}$ is a subgraph of $\boldsymbol{G}$ containing all of the vertices of $\boldsymbol{G}$. A *spanning tree* in $\boldsymbol{G}$ is a spanning subgraph that is a tree (that is, it is acyclic). Note that if

G is **not** connected, it cannot have a spanning tree (because otherwise there would be a path between every pair of vertices along edges in this tree, contradicting disconnectedness). Do all connected graphs have spanning trees?

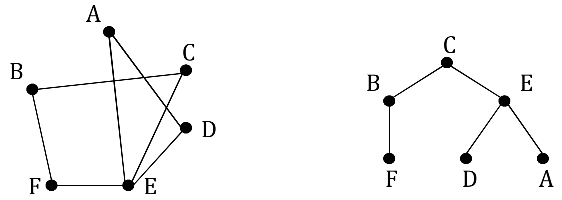

*Figure 8.1. A graph **G** and a spanning tree of **G***

Lemma 8.2: Every (finite) connected graph ***G*** = (***V***, ***E***) has a spanning tree.

Proof: Delete edges from ***G*** as follows. As long as there is at least one cycle present, take one cycle and delete one edge in that cycle. Notice that this procedure cannot destroy connectivity, so the graph obtained at each stage is connected. This process cannot continue indefinitely (we are dealing with finite graphs), so eventually we get a connected graph with no cycles. This is the required spanning tree (note that all the vertices of ***V*** are still present since we only deleted edges).

Spanning trees arise very often in the study of graphs, especially in optimization problems. I like to think of them as the "skeleton" of the graph, since they are in a sense the minimal structure that is still connected on its own. Their main usage on Olympiad problems is that instead of focusing on general graphs ***G*** which may have a complicated structure, we can sometimes find what we are looking for just by taking a spanning tree. For our purposes, all you really need to know about spanning trees is

a) They are trees
b) They exist (unless ***G*** isn't connected).

We'll now use our arsenal of lemmas to reduce some rather challenging Olympiad problems to just a few lines.

Example 1 [Based on ELMO Shortlist 2011, C7]
Let $\boldsymbol{T}$ be a tree with t vertices, and let $\boldsymbol{G}$ be a graph with n vertices. Show that if $\boldsymbol{G}$ has at least $(t-1)n$ edges, then $\boldsymbol{G}$ has a subgraph isomorphic to $\boldsymbol{T}$.

Answer:
By (ii), $\boldsymbol{G}$ has a subgraph $\boldsymbol{H}$ such that all vertices in $\boldsymbol{H}$ have degree at least $t-1$ (in $\boldsymbol{H}$). Applying (iii), $\boldsymbol{H}$ has a subgraph isomorphic to $\boldsymbol{T}$. ■
Remark: This is probably the shortest solution to a problem in this book. Note that it would actually be quite a difficult problem if you didn't know the super useful lemmas listed above!

Example 2 [Based on ELMO Shortlist 2011, C2]
Let $\boldsymbol{G}$ be a directed graph with n vertices such that each vertex has indegree and outdegree equal to 2. Show that we can partition the vertices of $\boldsymbol{G}$ into three sets such no vertex is in the same set as both the vertices it points to.

Answer:
Take a partition that maximizes the number of "crossing edges", that is, edges between distinct sets. If some v belongs to the same set as both of its out-neighbors, moving v to one of the other two sets (whichever has fewer in-neighbors of v) will add 2 crossing edges but destroy at most 1 old one. Then we get a partition with even more crossing edges, contradiction. Thus the original partition indeed works. ■
Remark: This is essentially the same idea used to prove (v).

Example 3 [Russia 2001]
A company with $2n+1$ people has the following property: For each group of n people, there exists a person amongst the remaining $n+1$ people who knows everyone in this group. Show that there exists a person who knows all the people in the company. (As

usual, knowing is mutual: A knows B if and only if B knows A).

Answer:
Assume to the contrary that no one knows everyone else. Construct a graph ***G*** with $2n+1$ vertices representing the people, and an edge between two vertices if and only if those two people **do not** know each other. Our assumption implies that every vertex has degree at least 1. Now applying lemma 8.1, there exists a dominating set of ***G*** containing n vertices. This means that each of the other $n+1$ vertices has a neighbor in this set of n vertices. In other words, no person outside this set of n people knows everyone in this set, contradicting the problem statement. This contradiction establishes the result. ■

The Extremal Principle

We've already encountered the extremal principle several times in various forms. The true power of this technique lies in its ubiquitous use in graph theory. In each of the next five examples, the step in which we use the extremal principle is marked in bold letters.

Example 4 [IMO Shortlist 2004, C3]
The following operation is allowed on a finite graph: choose any cycle of length 4 (if one exists), choose an arbitrary edge in that cycle, and delete this edge from the graph. For a fixed integer $n \geq 4$, find the least number of edges of a graph that can be obtained by repeated applications of this operation from the complete graph on n vertices (where each pair of vertices is joined by an edge).

Answer:
Clearly the answer cannot be less than $n-1$, since the graph obtained at each stage will always be connected. We claim that the

graph obtained at each stage is also non bipartite. This will imply that the answer is at least n (since a graph with n-1 vertices is a tree which is bipartite).

K_n is non-bipartite (it has a triangle), so suppose to the contrary that at some stage the deletion of an edge makes the graph bipartite. Consider the **first** time this happens. Let edge AB from the 4 cycle ABCD be the deleted edge. Since the graph is non bipartite before deleting AB but bipartite afterwards, it follows that A and B must lie on the same side of the partition. But since BC, CD, DA are edges in the now-bipartite graph, it follows that C and A are on one side and B and D are on the other side of the bipartition. Contradiction.

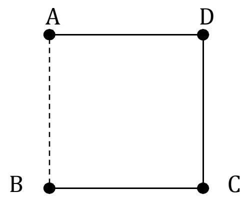

Figure 8.2.

To show n can be achieved, let the vertices be V_1, V_2, ..., V_n. Remove every edge V_iV_j with $3 \le i < j < n$ from the cycle $V_2\,V_i\,V_j\,V_n$. Then for $3 \le i < n$ delete edges $V_2\,V_i$ and $V_i\,V_n$ from cycles $V_1\,V_2\,V_i\,V_n$ and $V_1\,V_i\,V_n\,V_2$ respectively. This leaves us with only n edges: $V_1\,V_i$ for $2 \le i \le n$ and $V_2\,V_n$. ■

Remark: You may have been tempted to guess that the answer is (n−1), since the problem looks like the algorithm for obtaining a spanning tree. While guessing and conjecturing is an important part of solving problems, it is important to verify these guesses by experimenting a bit before trying to prove the guess. The process of realizing your guess was wrong may give you a clue as to how to proceed with the proof. In this example, you may have noticed that the graph you end up with always had an odd cycle, which would lead to the correct claim that the graph obtained is never bipartite.

Example 5 [Croatian TST 2011]
There are *n* people at a party among whom some are friends. Among any 4 of them there are either 3 who are all friends with each other or 3 who aren't friends with each other. Prove that the people can be separated into two groups ***A*** and ***B*** such that ***A*** is a clique (that is, everyone in ***A*** is friends with each other) and ***B*** is an independent set (no one in ***B*** knows anyone else in ***B***). (Friendship is a mutual relation).

Answer
Construct a graph ***G*** with vertices representing people and edges between two people if they are friends. The natural idea is to let ***A*** be the **largest clique in G**, and the remaining people as ***B***. We prove that this works.

If ***A***=***G*** or |***A***|=1 we are trivially done, so assume that $n > |A| \geq 2$. We only need to show that ***B*** is independent, that is, ***G-A*** is independent. Assume to the contrary v_1, v_2 belong to ***G*−*A*** and v_1v_2 is an edge. Since ***A*** is the largest clique, there exists u_1 in ***A*** such that v_1u_1 is not an edge (otherwise we could add v_1 to ***A***, forming a bigger clique).

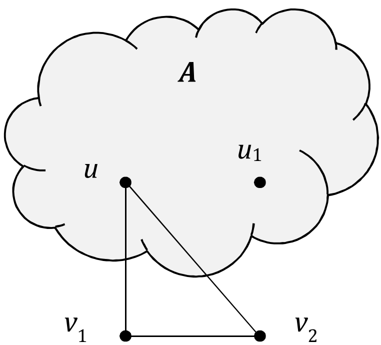

Figure 8.3.

If u_1v_2 is not an edge, then let u be any other vertex in $\boldsymbol{A}$. Since v_1v_2 and uu_1 are edges, by the condition of the problem there must be a triangle amongst these four vertices. The only possibility is uv_1v_2 since u_1v_2 and u_1v_1 are not edges. Then uv_1 and uv_2 are edges for all u in $\boldsymbol{A}$, so $\boldsymbol{A}$ U $\{v_1, v_2\}\setminus \{u_1\}$ is a larger clique, contradiction.

Similarly, if v_2u_1 is an edge, then for all u in $\boldsymbol{A}$ either v_2u_1u or v_2v_1u must be a triangle. In either case v_2u is an edge for all u in $\boldsymbol{A}$. Thus $\boldsymbol{A}$ U $\{v_2\}$ is a larger clique, contradiction. ■

Example 6 [Degree vectors]
A vector $\boldsymbol{v} = [d_1\ d_2\ \ldots\ d_k]$ with $d_1 \geq d_2 \geq \ldots \geq d_k$ is said to be a *graphical* vector if there exists a graph $\boldsymbol{G}$ with k vertices $x_1, x_2, \ldots, x_k$ having degrees d_1, d_2, …, d_k respectively. Note that there could be multiple graphs $\boldsymbol{G}$ with degree vector $\boldsymbol{v}$. Let $\boldsymbol{v}'$ be the vector obtained from $\boldsymbol{v}$ by deleting d_1 and subtracting 1 from the next d_1 components of $\boldsymbol{v}$. Let $\boldsymbol{v}_1$ be the non-increasing vector obtained from $\boldsymbol{v}'$ by rearranging components if necessary. (For example, if $\boldsymbol{v}$ = [4 3 3 2 2 2 1 1] then $\boldsymbol{v}'$ = [**2 2 1 1** 2 1 1] and $\boldsymbol{v}_1$ = [2 2 2 1 1 1 1].) Show that $\boldsymbol{v}_1$ is also a graphical vector.

Answer:
Let S be the sum of the indices of the neighbors of x_1 (for instance, if x_1 is adjacent to x_3, x_4 and x_8, then $S = 15$). Take the graph $\boldsymbol{G}$ with degree vector v such that the **S is as small as possible**.

Now we claim that there do not exist indices $i < j$ such that x_1x_i is **not** an edge and x_1x_j is an edge in $\boldsymbol{G}$. Suppose the contrary. Since $d_i \geq d_j$, there must be some vertex x_t such that x_ix_t is an edge but x_jx_t is not an edge. Now in $\boldsymbol{G}$, delete edges x_1x_j and x_ix_t and replace them with edges x_1x_i and x_jx_t. Note that all degrees remain unchanged, but the sum of indices of neighbors of $\boldsymbol{v}_1$ has decreased by $(j-i)$, contradicting our assumption on $\boldsymbol{G}$. This proves our claim.

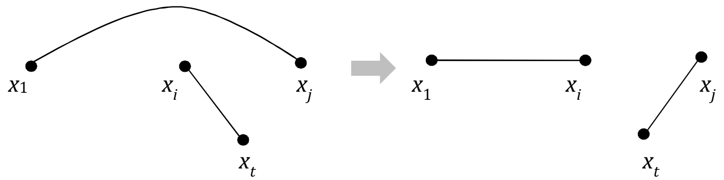

Figure 8.4. Illustration of a swap that decreases S

Our claim implies that x_1 is adjacent to the next d_1 vertices, namely x_2, x_3, …, x_{d_1+1}. Hence $\boldsymbol{v}_1$ is nothing but the graphical vector of the graph obtained from $\boldsymbol{G}$ by deleting x_1, since then the degrees of its neighbors all reduce by one. Hence $\boldsymbol{v}_1$ is graphical. ■

Remark 1: How did we come up with our rather strange extremal condition in the first paragraph? The problem provides a hint: it says we delete d_1 and subtract 1 from the next d_1 components to obtain $\boldsymbol{v'}$. Hence we wanted a graph such that x_1 is connected to the next d_1 vertices, since in this case simply removing vertex x_1 would have precisely this effect on the degree vector (this is our reasoning in the last paragraph of the proof). Now to prove such a graph exists, we needed a simple extremal property satisfied by such a graph. This naturally leads to our definition of S and the extremal condition that $\boldsymbol{G}$ minimizes S.

Remark 2: This is quite a useful lemma for testing whether a given degree sequence is graphical (see exercise 27).

Example 7 [MOP 2008]
Prove that if the edges of K_n, the complete graph on n vertices, are colored such that no color is assigned to more than $n-2$ edges, there exists a triangle in which each edge is a distinct color.

Answer
Assume to the contrary that there exists no such triangle. Define a *C-connected component* to be a set of vertices such that for any

two vertices in that set, there exists a path between them, all of whose edges are of color C. Now let $\boldsymbol{X}$ be the **largest** C-connected component of the graph for any color C, and say the color of $\boldsymbol{X}$ is red.

Suppose there is a vertex v not in $\boldsymbol{X}$. Consider two vertices u_1 and u_2 that are joined by a red edge. Neither edge vu_1 nor edge vu_2 can be red (otherwise v could be added to $\boldsymbol{X}$). So vu_1 and vu_2 are the same color (otherwise vu_1u_2 would have three distinct colors, contradicting our assumption). It follows that v is joined to all elements of $\boldsymbol{X}$ by the same color edge, say blue. But then $\boldsymbol{X} \cup \{v\}$ is a larger connected component (of color blue), contradiction.

It follows that there cannot be any vertex v outside $\boldsymbol{X}$, so all n vertices are in $\boldsymbol{X}$. Now, since $\boldsymbol{X}$ is red-connected and has n vertices, there must be at least $n-1$ red edges, contradiction. ■

Example 8 [Generalization of USAMO 2007-4]
Given a connected graph $\boldsymbol{G}$ with V vertices, each having degree at most D, show that $\boldsymbol{G}$ can be partitioned into two connected subgraphs, each containing at least $\frac{V-1}{D}$ vertices.

Answer:
We induct on E, the number of edges of $\boldsymbol{G}$. When $E = 1$, there are only two vertices and the partition consists of two isolated vertices. For the induction step, note that if we can delete an edge and $\boldsymbol{G}$ remains connected, we are done by induction. Hence we only consider the case when $\boldsymbol{G}$ is a tree.

Pick the root x such that the size of the **largest subtree T is minimized**. Clearly $|\boldsymbol{T}| \geq \frac{V-1}{D}$, since there are at most D subtrees and V−1 vertices amongst them. Also, $|\boldsymbol{T}| \leq \frac{V}{2}$. This is because if $|\boldsymbol{T}| > \frac{V}{2}$, then instead of rooting the tree at x we could root the tree at the first vertex y of $\boldsymbol{T}$. This would decrease $|\boldsymbol{T}|$ by 1, but T\{y} would still be the largest subtree since it would still have at least

$\frac{V-1}{2}$ vertices. This would contradict our assumption on x, as we would have a smaller largest subtree.

Thus $\frac{V}{2} > |\boldsymbol{T}| \geq \frac{V-1}{D}$. Take $\boldsymbol{T}$ to be one subgraph and $\boldsymbol{G} - \boldsymbol{T}$ to be the other. Both have size at least $\frac{V-1}{D}$ by the bounds on $\boldsymbol{T}$ and are connected, so we have found a valid partition. ■

Remark: The "induction" in the first paragraph is really just a formal way of saying "Well, the *worst* case for us is when $\boldsymbol{G}$ is a tree, so let's just forget about general graphs and prove the result for trees: if it's true for trees, it's true for everyone". Another way of reducing the focus to just trees is to take a spanning tree of $\boldsymbol{G}$.

Hall's Marriage theorem

Given N sets (not necessarily distinct), we say that the family of N sets has a ***system of distinct representatives*** (SDR) if it is possible to choose exactly one element from each set such that all the chosen elements are distinct. For example, if we have 4 sets {1, 2, 3}, {2, 4}, {2, 3, 4} and {1, 3} then this family has (1, 2, 4, 3) as a system of distinct representatives. Under what conditions does a family have a system of distinct representatives? One obvious necessary condition is that for any subfamily of k sets, the union of these k sets must have at least k elements. It turns out that this condition, known as the marriage condition, is also sufficient.

Example 9 [Hall's Marriage Theorem]
Show that the marriage condition is sufficient for the existence of an SDR.

Proof: Let the marriage condition hold for the family $A_1, A_2, \ldots, A_n$. Keep deleting elements from these sets until a family $F' = A_1', A_2', \ldots, A_n'$ is reached such that the deletion of any element will cause the marriage condition to be violated. We claim that at this stage each set contains exactly one element. This would imply the result, since these elements would be distinct by the marriage condition, and would hence form the required SDR.

Suppose our claim is false. Then some set contains at least 2 elements. WLOG this set is A_1 and let x and y be elements of A_1. Deleting x or y would violate the marriage condition by the definition of F'. Thus there exists subsets $\boldsymbol{P}$ and $\boldsymbol{Q}$ of $\{2, 3, 4, \ldots, n\}$ such that $\boldsymbol{X}=(A_1'-x) \cup (\bigcup_{i\epsilon P} A_i')$ and $\boldsymbol{Y}=(A_1'-y) \cup (\bigcup_{i\epsilon Q} A_i')$ satisfy $|\boldsymbol{X}| \leq |\boldsymbol{P}|$ and $|\boldsymbol{Y}| \leq |\boldsymbol{Q}|$. Adding gives

$$|\boldsymbol{X}| + |\boldsymbol{Y}| = |\boldsymbol{X} \cap \boldsymbol{Y}| + |\boldsymbol{X} \cup \boldsymbol{Y}| \leq |\boldsymbol{P}| + |\boldsymbol{Q}|.$$

Now, $\boldsymbol{X} \cup \boldsymbol{Y} = \boldsymbol{A}_1' \cup (\bigcup_{i\epsilon P \cup Q} \boldsymbol{A}_i')$ and $\boldsymbol{X} \cap \boldsymbol{Y} = \bigcup_{i\epsilon P \cap Q} \boldsymbol{A}_i'$.

Thus the marriage condition implies that

$$|\boldsymbol{X} \cup \boldsymbol{Y}| \geq 1 + |\boldsymbol{P} \cup \boldsymbol{Q}|, \text{ and } |\boldsymbol{X} \cap \boldsymbol{Y}| \geq |\boldsymbol{P} \cap \boldsymbol{Q}|.$$

Adding gives

$$|\boldsymbol{X} \cap \boldsymbol{Y}| + |\boldsymbol{X} \cup \boldsymbol{Y}| \geq 1 + |\boldsymbol{P} \cup \boldsymbol{Q}| + |\boldsymbol{P} \cap \boldsymbol{Q}| = |\boldsymbol{P}| + |\boldsymbol{Q}| + 1,$$

contradicting our earlier bound. ■

Remark: The key idea in this proof was the fact that the marriage condition holds for the sets $\boldsymbol{A}_1', \boldsymbol{A}_2', \ldots \boldsymbol{A}_n'$ but **not** for $\boldsymbol{A}_1 \setminus \{x\}, \boldsymbol{A}_2', \ldots, \boldsymbol{A}_n'$ and $\boldsymbol{A}_1 \setminus \{y\}, \boldsymbol{A}_2', \ldots, \boldsymbol{A}_n'$. This proof illustrates an important idea: it's useful to exploit conditions given to us, but it's even more useful to exploit situations when the conditions **don't** hold.

Hall's marriage theorem was phrased above in the language of set theory, but we can also interpret it in graph theoretical terms. Consider a bipartite graph $\boldsymbol{G}$ with vertex set $\boldsymbol{V} = \boldsymbol{V}_1 \cup \boldsymbol{V}_2$. A *complete matching* of the vertices of $\boldsymbol{V}_1$ is a subset of the edges of $\boldsymbol{G}$

such that:

(i) Every vertex of $\boldsymbol{V}_1$ is incident on exactly one edge

(ii) Each vertex of $\boldsymbol{V}_2$ is incident on at most one edge

In other words, it is a pairing such that every vertex in $\boldsymbol{V}_1$ is paired with a vertex in $\boldsymbol{V}_2$ and no two vertices in $\boldsymbol{V}_1$ are paired with the same vertex of $\boldsymbol{V}_2$. The vertices in a pair are joined by an edge.

If the vertices of $\boldsymbol{V}_1$ represent the sets $\boldsymbol{A}_1$, $\boldsymbol{A}_2$, …, $\boldsymbol{A}_n$ and the vertices of $\boldsymbol{V}_2$ represent elements in $\bigcup_{i=1}^{n} \boldsymbol{A}_i$, then a complete matching of $\boldsymbol{V}_1$ gives us a system of distinct representatives (namely the vertices of $\boldsymbol{V}_2$ to which the vertices of $\boldsymbol{V}_1$ are matched).

Example 10 [Canada 2006-3]

In a rectangular array of nonnegative reals with m rows and n columns, each row and each column contains at least one positive element. Moreover, if a row and a column intersect in a positive element, then the sums of their elements are the same. Prove that $m = n$.

Answer:

Create a bipartite graph, with the left side representing rows and the right side for columns. Place an edge between two vertices if and only if the corresponding row and column intersect in a positive element. The idea is to show that there is a matching from rows to columns, so $n \geq m$. By symmetry the same argument will give $m \geq n$, implying $m = n$.

Assume to the contrary there is no such matching. Then the marriage condition must be violated, so there exists some set $\boldsymbol{S}$ of rows having a total set $\boldsymbol{T}$ with $|\boldsymbol{T}| < |\boldsymbol{S}|$ of columns in which positive entries appear. Let the sums of the $|\boldsymbol{S}|$ rows be s_1, …, s_k. By the property, each of the $|\boldsymbol{T}|$ columns has sum equal to one of the s_i. So the total sum of the elements in the $\boldsymbol{S}$ rows, when calculated from the column point of view (since entries outside are all nonnegative) is at most a sum of a subset of the s_i. Yet from the row point of view, it is the full sum. As all $s_i > 0$, this is a contradiction. ■

Remark: This duality between arrays of numbers (i.e. matrices) and graphs (especially bipartite graphs) comes up very often. Keep an eye out for this trick, since it can prove very useful. In fact, analyzing and algebraically manipulating these matrices allows graph theory to be studied from an algebraic viewpoint, and fast algorithms for multiplying matrices form the basis of a class of algorithms for large graphs known as algebraic graph algorithms.

Unexpected Applications of Graph Theory

Most problems in previous sections suggested a natural graph theoretic interpretation. In this section, we will leverage the power of graphs to model complex relationships in nonobvious ways. Carefully constructed graphs can reduce unfamiliar, complex problems to familiar graph theoretic ones.

Example 11 [Taiwan 2001]
Let $n \geq 3$ be an integer and let $\boldsymbol{A}_1, \boldsymbol{A}_2, \ldots, \boldsymbol{A}_n$ be n distinct subsets of $\boldsymbol{S} = \{1, 2, \ldots, n\}$. Show that there exists an element $x \in \boldsymbol{S}$ such that the subsets $\boldsymbol{A}_1 \setminus \{x\}, \boldsymbol{A}_2 \setminus \{x\}, \ldots, \boldsymbol{A}_n \setminus \{x\}$ are also distinct.

Answer:
We construct a graph $\boldsymbol{G}$ with vertices $A_1, A_2, \ldots, A_n$. For each element y, if there exist distinct sets $\boldsymbol{A}_i$ and $\boldsymbol{A}_j$ such that $\boldsymbol{A}_i \setminus \{y\} = \boldsymbol{A}_j \setminus \{y\}$, we select **exactly one** such pair $(\boldsymbol{A}_i, \boldsymbol{A}_j)$ and join them by an edge (even if there are multiple such pairs, we select only one for each y). Suppose to the contrary there doesn't exist x as stated in the problem. Then all elements of $\boldsymbol{S}$ contribute at least one edge to the graph. Moreover, it is impossible for two different elements

to contribute the same edge since if $\boldsymbol{A}_i \setminus y_1 = \boldsymbol{A}_j \setminus y_1$ and $\boldsymbol{A}_i \setminus y_2 = \boldsymbol{A}_j \setminus y_2$ for distinct y_1 and y_2, this would force $\boldsymbol{A}_i = \boldsymbol{A}_j$.

Thus $\boldsymbol{G}$ has at least n edges, and hence has a cycle, WLOG $A_1A_2 \ldots A_kA_1$ for some $k \geq 3$. Then there exists some distinct $x_1, x_2, \ldots, x_k$ such that $\boldsymbol{A}_1 \setminus \{x_1\} = \boldsymbol{A}_2 \setminus \{x_1\}$; $\boldsymbol{A}_2 \setminus \{x_2\} = \boldsymbol{A}_3 \setminus \{x_2\}, \ldots, \boldsymbol{A}_k \setminus \{x_k\} = \boldsymbol{A}_1 \setminus x_k$. Now x_1 is in exactly one of $\boldsymbol{A}_1$ and $\boldsymbol{A}_2$ (otherwise $\boldsymbol{A}_1 = \boldsymbol{A}_2$). WLOG it is in $\boldsymbol{A}_2$ but not in $\boldsymbol{A}_1$. But then x_1 must also be in $\boldsymbol{A}_3$ since $\boldsymbol{A}_2 \setminus \{x_2\} = \boldsymbol{A}_3 \setminus \{x_2\}$, and similarly must be in $\boldsymbol{A}_4$, and so on. We finally get that $x_1 \in \boldsymbol{A}_1$, a contradiction. ■

The next problem, like several others in this book, underscores the usefulness of induction. Problems with around "2^n" objects practically beg you to induct: all you need to do is find an appropriate way to split the set of objects into two parts, and apply the induction hypothesis to the larger or smaller part as applicable. But how does this connect to graph theory?

Example 12 [USA TST 2002]

Let n be a positive integer and let $\boldsymbol{S}$ be a set of (2^n+1) elements. Let f be a function from the set of two-element subsets of $\boldsymbol{S}$ to $\{0, 1, \ldots, (2^{n-1}-1)\}$. Assume that for any elements (x, y, z) of $\boldsymbol{S}$, one of $f(\{x, y\})$, $f(\{y, z\})$ and $f(\{z, x\})$ is equal to the sum of the other two. Show that there exist a, b, c in $\boldsymbol{S}$ such that $f\{a, b\}$, $f\{b, c)$ and $f\{c, a\}$ are all equal to 0.

Answer

Step 1: The basic strategy

Our idea is to find a subset $\boldsymbol{S}'$ of $\boldsymbol{S}$ such that $|\boldsymbol{S}'| \geq 2^{n-1} + 1$ and for all x, y in $\boldsymbol{S}'$ **$f(\{x, y\})$ is even**. Then if we let $g(\{x, y\}) = \frac{f(\{x,y\})}{2}$ for all x, y in $\boldsymbol{S}'$, we would have a function from pairs in $\boldsymbol{S}'$ to $\{0, 1, \ldots, 2^{n-2}-1\}$ satisfying the same conditions as f and we could apply the induction hypothesis to get the result. It remains to show that such a set $\boldsymbol{S}'$ exists.

Step 2: Constructing the graph and a new goal

Construct a graph ***G*** with 2^n+1 vertices representing elements in ***S*** as follows: there is an edge between a and b if and only if $f(\{a, b\})$ is odd. We now need to find an independent set in ***G*** of size at least $2^{n-1}+1$. Our hope is that ***G*** is bipartite: then we can just take the larger side of the bipartition, which will have size at least $\lceil (2^n + 1)/2 \rceil = 2^{n-1}+1$. Some experimentation confirms our hope – but of course we still need a proof.

Step 3: A key observation

Note that for any 3 vertices a, b, c there must be 0 or 2 edges amongst them, since $f(\{a, b\}) + f(\{b, c\}) + f(\{c, a\})$ is even (since one of these terms is the sum of the other two).

Step 4: Proving G is bipartite

If ***G*** is not bipartite, it has an odd cycle, so consider its *smallest* odd cycle $v_1v_2\ldots v_{2k+1}$. Consider vertices v_1, v_3, v_4. There must be an even number of edges amongst them. As v_3v_4 is an edge, v_1v_3 or v_1v_4 must be an edge. v_1v_3 is not an edge since otherwise amongst vertices v_1, v_2 and v_3 there would be three edges, contradicting our earlier observation. Hence v_1v_4 is an edge. But then $v_1v_4\ldots v_{2k+1}$ is a smaller odd cycle, contradicting our assumption. ■

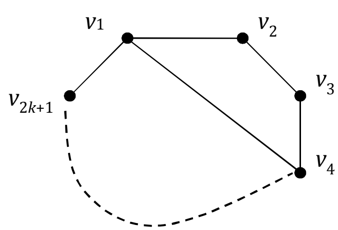

Figure 8.5. The edge v_1v_4 creates a smaller odd cycle

Remark 1: In step 4 we essentially proved the following result: If ***G*** is a graph in which for any three vertices there are either 0 or 2 edges between them, ***G*** is bipartite. This is a very handy lemma to keep in mind, especially since so many Olympiad problems boil down to proving that a certain graph is bipartite.

Remark 2: The idea of taking a shortest cycle arises very often.

Example 13 [IMO Shortlist 2002, C6]
Let n be an even positive integer. Show that there is a permutation $(x_1, x_2, ..., x_n)$ of $(1, 2, ..., n)$ such that for every $i \in (1, 2, ..., n)$, the number x_{i+1} is one of the numbers $2x_i$, $2x_i - 1$, $2x_i - n$, $2x_i - n - 1$. Here we use the cyclic subscript convention, so that x_{n+1} means x_1.

Answer:
Let $n = 2m$. We define a directed graph with vertices 1, 2, ..., m and edges numbered 1, 2,...,$2m$ as follows. For each i ≤ m, vertex i has two outgoing edges numbered $2i-1$ and $2i$, and two incoming edges labeled i and $i+m$. All we need is an Eulerian circuit, because then successive edges will be of one of the forms $(i, 2i-1)$, $(i, 2i)$, $(i+m, 2i)$ or $(i+m, 2i-1)$. Then we can let $x_1, x_2, ..., x_n$ be the successive edges encountered in the Eulerian circuit and they will satisfy the problem's conditions.

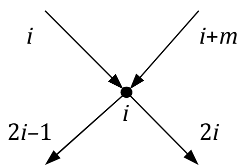

Figure 8.6.

Now, each vertex's indegree is equal to its outdegree, so we just need to show weak connectivity to establish that there is an Eulerian circuit. We do this by strong induction. There is a path from 1 to k: since there is a path from 1 to j where $2j = k$ or $2j-1 = k$, and an edge from j to k, there is a path from 1 to k. Thus ***G*** is weakly connected and hence has an Eulerian circuit. ■

Exercises

Hall's Theorem and Related Problems

1. **[Konig's marriage theorem]**
   Show that a *k-regular bipartite graph* $\boldsymbol{G}$ (all vertices have degree k) has a perfect matching (a matching covering all its vertices).

2. **[Konig's theorem]**
   A *matching* in a graph $\boldsymbol{G}$ is a set $\boldsymbol{M}$ of edges such that each vertex in $\boldsymbol{G}$ is incident to at most one edge in $\boldsymbol{M}$. A *vertex cover* in $\boldsymbol{G}$ is a set of vertices $\boldsymbol{C}$ such that each edge is incident to **at least** one vertex in $\boldsymbol{C}$. Using Hall's marriage theorem, show that in a bipartite graph $\boldsymbol{G}$, the maximum possible size of a matching is equal to the minimum possible size of a vertex cover.

3. **[Vietnam TST 2001]**
   A club has 42 members. Suppose that for any 31 members in this club, there exists a boy and a girl among these 31 members who know each other. Show that we can form 12 disjoint pairs of people, each pair having one boy and one girl, such that the people in each pair know each other.

4. **[IMO Shortlist 2006, C6]**
   Consider an upward equilateral triangle of side length n, consisting of n^2 unit triangles (by upward we mean with vertex on top and base at the bottom). Suppose we cut out n upward unit triangles from this figure, creating n triangular holes. Call the resulting figure a *holey triangle*. A *diamond* is a 60-120 unit rhombus. Show that a holey triangle T can be tiled by diamonds, with no diamonds overlapping, covering a hole or sticking out of T if and only if the following holds: every

upward equilateral of side length k in T has at most k holes, for all $1 \le k \le n$.

5. **[Dilworth's theorem]**
A *directed acyclic graph* (DAG) is a directed graph with no directed cycles. An *antichain* in a DAG (with some abuse of notation) is a set of vertices such that no two vertices in this set have a directed path between them. Show that the size of the maximum antichain of the DAG is equal to the minimum number of disjoint paths into which the DAG can be decomposed.

6. **[Romanian TST 2005]**
Let $\boldsymbol{S}$ be a set of n^2+1 positive integers such that in any subset $\boldsymbol{X}$ of $\boldsymbol{S}$ with $n+1$ integers, there exist integers $x \neq y$ such that $x|y$. Show that there exists a subset $\boldsymbol{S}'$ of $\boldsymbol{S}$ with $\boldsymbol{S}' = \{x_1, x_2, \dots, x_{n+1}\}$ such that $x_i|x_{i+1}$ for each $1 \le i \le n$.

Coloring Problems

7. **[Welsh-Powell theorem] (U*)**
A *proper coloring* of the vertices of a graph $\boldsymbol{G}$ is an assignment of one color to each vertex of a graph $\boldsymbol{G}$ such that no two adjacent vertices in $\boldsymbol{G}$ have the same color. Let $\boldsymbol{G}$ be a graph with vertices having degrees $d_1 \ge d_2 \ge \dots \ge d_n$. Show that there exists a proper coloring of $\boldsymbol{G}$ using at most $\max_i (\min\{i, d_i+1\}$ colors.

8. **[Dominating sets and coloring] (U*)**
If a graph on n vertices has no dominating set of size less than k, then show that its vertices can be properly colored in $n-k$ colors.

9. **[IMO 1992, Problem 3]**
Let $\boldsymbol{G}$ be the complete graph on 9 vertices. Each edge is either colored blue or red or left uncolored. Find the smallest value

of n such that if n edges are colored, there necessarily exists a monochromatic triangle.

10. [IMO Shortlist 1990]
The edges of a K_{10} are colored red and blue. Show that there exist two disjoint monochromatic odd cycles, both of the same color.

11. [Szekeres-Wilf theorem] (U*)
Show that any graph $\boldsymbol{G}$ can be properly colored using at most 1+max $\Delta(\boldsymbol{G}')$ colors, where the maximum is taken over all induced subgraphs $\boldsymbol{G}'$ of $\boldsymbol{G}$ and $\Delta(\boldsymbol{G}')$ refers to the maximum degree of a vertex in the induced subgraph $\boldsymbol{G}'$.

12. [Generalization of USA TST 2001]
Let $\boldsymbol{G}$ be a directed graph on n vertices, such that no vertex has out-degree greater than k. Show that the vertices of $\boldsymbol{G}$ can be colored in $2k$+1 colors such that no two vertices of the same color have a directed edge between them.

13. [Brook's theorem] (U*)
We know from (i) that Δ+1 colors are sufficient to properly color the vertices of a graph $\boldsymbol{G}$, where Δ is the maximum degree of any vertex in $\boldsymbol{G}$. Show that if $\boldsymbol{G}$ is connected and is neither a complete graph nor an odd cycle, then actually Δ colors suffice.
[Easy version: prove the result above with the added condition that not all vertices have the same degree.]

Turan's theorem and applications

14. [Turan's theorem] (U*)
The *Turan graph* $\boldsymbol{T}(n, r)$ is the graph on n vertices formed as follows: partition the set of n vertices into r equal or almost equal (differing by 1) parts, and join two vertices by an edge if and only if they are in different parts. Note that $\boldsymbol{T}(n, r)$ has no (r+1)-clique. Show that amongst all graphs having no (r+1)-

clique, the Turan graph has the most edges. Hence, deduce that the maximum number of edges in a K_{r+1}-free graph is $\frac{(r-1)n^2}{2r}$. This generalizes the bound of $n^2/4$ edges in triangle-free graphs that we proved in chapter 6.

[Hint: first prove the following claim: there do not exist three vertices u, v, w such that uv is an edge in $\boldsymbol{G}$ but uw and vw are not. Show this by assuming the contrary and then making some adjustments to $\boldsymbol{G}$ to obtain a graph with more edges, contradicting the fact that $\boldsymbol{G}$ has the maximum possible number of edges amongst all K_{r+1}-free graphs. The claim establishes that if two vertices u and v have a common non-neighbor, then u and v themselves are non-neighbors. This shows that $\boldsymbol{G}$ is k-partite for some k. Now show that the maximum number of edges will occur when $k = r$ and the parts are equal or differ by 1.]

Remark: This method of proving that $\boldsymbol{G}$ is multipartite is extremely important, and is sometimes called Zykov symmetrization.

15. [Poland 1997]

There are n points on a unit circle. Show that at most $n^2/3$ pairs of these points are at distance greater than $\sqrt{2}$.

16. [IMO Shortlist 1989]

155 birds sit on the circumference of a circle. It is possible for there to be more than one bird at the same point. Birds at points P and Q are mutually visible if and only if angle POQ ≤ 10°, where O is the center of the circle. Determine the minimum possible number of pairs of mutually visible birds.

17. [USA TST 2008]

Given two points (x_1, y_1) and (x_2, y_2) in the coordinate plane, their *Manhattan distance* is defined as $|x_1-x_2| + |y_1-y_2|$. Call a pair of points (A, B) in the plane *harmonic* if $1 < d(A, B) \leq 2$.

Given 100 points in the plane, determine the maximum number of harmonic pairs among them.

More Extremal Graph problems

18. [China TST 2012]

Let n and k be positive integers such that $n > 2$ and $n/2 < k < n$. Let $\boldsymbol{G}$ be a graph on n vertices such that $\boldsymbol{G}$ contains no $(k+1)$-clique but the addition of any new edge to $\boldsymbol{G}$ would create a $(k+1)$-clique. Call a vertex in $\boldsymbol{G}$ central if it is connected to all $(n-1)$ other vertices. Determine the least possible number of central vertices in $\boldsymbol{G}$.

19. [China TST 2011]

Let $\boldsymbol{G}$ be a graph on $3n^2$ vertices ($n > 1$), with no vertex having degree greater than $4n$. Suppose further that there exists a vertex of degree one and that for any two points, there exists a path of length at most 3 between them. Show that G has at least $(7n^2-3n)/2$ edges.

20. [Generalization of IMO Shortlist 2013, C6]

In a graph $\boldsymbol{G}$, for any vertex v, there are at most $2k$ vertices at distance 3 from it. Show that for any vertex u, there are at most $k(k+1)$ vertices at distance 4 from it.

21. [IMO Shortlist 2004, C8]

For a finite graph $\boldsymbol{G}$, let $f(\boldsymbol{G})$ denote the number of triangles in $\boldsymbol{G}$ and $g(\boldsymbol{G})$ the number of tetrahedra (K_4s). Determine the smallest constant c such that $g(\boldsymbol{G})^3 \leq c\, f(\boldsymbol{G})^4$ for all graphs $\boldsymbol{G}$.

22. [IMO Shortlist 2002, C7]

In a group of 120 people, some pairs are friends. A weak quartet is a group of 4 people containing exactly one pair of friends. What is the maximum possible number of weak quartets?

Tournaments

23. Show that if a tournament has a directed cycle, then it has a directed triangle.

24. [Landau's theorem]
Call a vertex v in a tournament $\boldsymbol{T}$ a champion if for every vertex u in $\boldsymbol{T}$, there is a directed path from v to u in $\boldsymbol{T}$ of length at most 2. Show that every tournament has a champion.

25. [Moon-Moser theorem]
A directed graph $\boldsymbol{G}$ is called *strongly connected* if there is a directed path from each vertex in $\boldsymbol{G}$ to every other vertex in $\boldsymbol{G}$. A directed graph with n vertices is called *vertex-pancyclic* if every vertex is contained in a cycle of length p, for each $3 \leq p \leq n$. Show that a strongly connected tournament is vertex pancyclic.

26. [Based on USA TST 2009]
Let $n>m>1$ be integers and let $\boldsymbol{G}$ be a tournament on n vertices with no $(m+1)$-cycles. Show that the vertices can be labeled 1, 2, …, N such that if $a \geq b+m-1$, there is a directed edge from b to a.

Miscellaneous

27. [Generalization of Saint Petersburg 2001]
For each positive integer n, show that there exists a graph on $4n$ vertices with exactly two vertices having degree d, for each $1 \leq d \leq 2n$.
Remark: The degree vector lemma kills this otherwise difficult problem.

28. [Iran 2001]
In an $n \times n$ matrix, a *generalized diagonal* refers to a set of n

entries with one in each row and one in each column. Let $\boldsymbol{M}$ be an $n \times n$ 0-1 matrix, and suppose $\boldsymbol{M}$ has exactly one generalized diagonal containing all 1s. Show that it is possible to permute the rows and columns of $\boldsymbol{M}$ to obtain a matrix $\boldsymbol{M}'$ such that (i, j) entry in $\boldsymbol{M}'$ is 0 for all $1 \leq j < i \leq n$.

29. [Generalization of Russia 2001]

Let $\boldsymbol{G}$ be a tree with exactly $2n$ leaves (vertices with degree 1). Show that we can add n edges to $\boldsymbol{G}$ such that $\boldsymbol{G}$ becomes 2-connected, that is, the destruction of any edge at this point would still leave $\boldsymbol{G}$ connected.

30. [Russia 1997]

Let m and n be odd integers. An $m \times n$ board is tiled with dominoes such that exactly one square is left uncovered. One is allowed to slide vertical dominoes vertically and horizontal dominoes horizontally so as to occupy the empty square (thereby changing the position of the empty square). Suppose the empty square is initially at the bottom left corner of the board. Show that by a sequence of moves we can move the empty square to any of the other corners.

31. [Generalization of Japan 1997]

Let $\boldsymbol{G}$ be a graph on n vertices, where $n \geq 9$. Suppose that for any 5 vertices in $\boldsymbol{G}$, there exist at least two edges with endpoints amongst these 5 vertices. Show that $\boldsymbol{G}$ has at least $n(n-1)/8$ edges. Determine all n for which equality can occur.

32. [Japan 1997]

Let n be a positive integer. Each vertex of a 2^n-gon is labeled 0 or 1. There are 2^n sequences obtained by starting at some vertex and reading the first n labels encountered clockwise. Show that there exists a labelling such that these 2^n sequences are all distinct.

33. [Based on USA TST 2011]

In an undirected graph $\boldsymbol{G}$, all edges have weight either 1 or 2. For each vertex, the sum of the weights of edges incident to it is odd. Show that it is possible to orient the edges of $\boldsymbol{G}$ such that for each vertex, the absolute value of the difference between its in-weight and out-weight is 1, where in-weight refers to the sum of weights of incoming edges and out-weight refers to the sum of weights of outgoing edges.

34. [IMO Shortlist 1990]

Consider the rectangle in the coordinate plane with vertices (0, 0), (0, m), (n, 0) and (m, n), where m and n are odd positive integers. This rectangle is partitioned into triangles such that each triangle in the partition has at least one side parallel to one of the coordinate axes, and the altitude on any such side has length 1. Furthermore, any side that is not parallel to a coordinate axis is common to two triangles in the partition. Show that there exist two triangles in the partition each having one side parallel to the x axis and one side parallel to the y axis.

Olympiad Combinatorics

Pranav A. Sriram

August 2014

Copyright notices

About the Author

Pranav Sriram graduated from high school at The International School Bangalore, India, and is currently pursuing undergraduate studies at Stanford University.

9. THE PROBABILISTIC METHOD

Introduction

Our final chapter will focus on an idea that has had a tremendous impact on combinatorics over the past sixty years, and that is playing a critical role in the "big data" driven applications of today's digitized world. The *probabilistic method* is a technique that broadly refers to using arguments based on probability to prove results in fields not necessarily directly related to probability. Probabilistic combinatorics is an extremely active area of research today, and several, if not most, recent developments in graph theory, extremal combinatorics, computational geometry, and combinatorial algorithms have relied extensively on probabilistic arguments. In addition, the probabilistic method has led to breakthroughs in number theory and additive group theory, and has surprisingly elegant connections to information theory and the theory of error correcting codes. Modern paradigms for efficiently monitoring computer network traffic, finding hidden patterns in large datasets, and solving extremely large scale mathematical problems have leveraged the power of the probabilistic method in advanced ways.

Each section in this chapter illustrates one general method of solving combinatorial problems via probabilistic arguments. We will assume knowledge of basic concepts from high school probability, such as the notion of random variables, expected value, variance, independent events, and mutually exclusive events. No formal background in Lebesgue integration or real analysis is required.

Linearity of Expectation

The linearity of expectation refers to the following (fairly intuitive) principle: if X is the sum of n random variables $X_1, X_2, \ldots, X_n$, then the expected value of X is equal to the sum of expected values of $X_1, X_2, \ldots, X_n$. In symbols,

$$E[X] = E[X_1 + X_2 + \ldots + X_n] = E[X_1] + E[X_2] \ldots + E[X_n].$$

This holds even when the variables $X_1, X_2, \ldots, X_n$ are not independent from each other. Another intuitive principle we will use, sometimes referred to as the "pigeonhole property" (PHP) of expectation, is that a random variable cannot always be less than its average. Similarly, it cannot always be more than its average. For example, if the average age at a party is 24 years, there must be at least one person at most 24 years old, and one person who is at least 24 years old.

How can we use this to solve combinatorial problems?

Example 1 [Szele's Theorem]
For each n, show that there exists a tournament on n vertices having at least $\frac{n!}{2^{n-1}}$ Hamiltonian paths.

Answer:
The basic idea is as follows: instead of explicitly constructing a tournament with the required properties, we use a *randomized procedure* for creating some tournament T. We will show that *on average*, the tournament we (randomly) construct contains $\frac{n!}{2^{n-1}}$ Hamiltonian paths. Then by the pigeonhole principle, *some* tournament must contain at least $\frac{n!}{2^{n-1}}$ Hamiltonian paths.

The randomized procedure is to construct a tournament where each of the matches is decided independently by a flip of a fair coin. Thus for two vertices u and v, the probability that u lost to v is 0.5 and vice versa. To compute the expected total number of Hamiltonian paths, observe that each Hamiltonian path corresponds to a permutation of the vertices of the graph.

Consider a random permutation of the vertices $(v_1, v_2, v_3, \ldots, v_n)$. For this sequence (in order) to form a Hamiltonian path, we need to have v_1 lost to v_2, v_2 lost to v_3, …, v_{n-1} lost to v_n. Each of these happens with probability 0.5, so the probability that they all occur is $(0.5)^{n-1}$. Thus, each permutation of the vertices has probability $(0.5)^{n-1}$ of being a Hamiltonian path, and there are $n!$ permutations that can give rise to Hamiltonian paths. Hence the **expected** total number of Hamilton paths, by the linearity of expectation, is $n!/2^{n-1.}$ It follows that *some* tournament must have at least the expected number of Hamiltonian paths. ■

Example 2 [Generalization of MOP 2010]
Let $\boldsymbol{G}$ be a graph with E edges and n vertices with degrees $d_1, d_2, \ldots, d_n$. Let k be an integer with $k \le \frac{2E}{n}$. Show that $\boldsymbol{G}$ contains an induced subgraph $\boldsymbol{H}$ such that $\boldsymbol{H}$ contains no $\boldsymbol{K}_{k+1}$ and $\boldsymbol{H}$ has at least $\frac{kn^2}{2E+n}$ vertices. (Recall that a $\boldsymbol{K}_{k+1}$ is a complete graph on k+1 vertices.)

Answer:

The basic idea

Note that it if all the vertices in ***H*** have degree at most k-1 (in ***H***), then ***H*** clearly cannot contain a $\boldsymbol{K_{k+1}}$. This gives us the following idea for a greedy construction.

Take a random permutation of the vertices $v_1v_2...v_n$. Starting from v_1 and continuing, we select vertex v_i if and only if amongst v_i and its neighbors, v_i is one of the first k appearing in the sequence. This ensures that every chosen vertex has at most k-1 chosen neighbors, so the final set of chosen vertices will not contain an induced $\boldsymbol{K_{k+1}}$. Then we just need to show that the expected number of chosen vertices is at least $\frac{kn^2}{2E+n}$.

Computing the expectation

Now, each vertex v_i will be chosen with probability $\frac{k}{d_i+1}$, where d_i is its degree. This is because v_i is selected if and only if amongst v_i and its d_i neighbors, v_i is amongst the first k in the permutation. The expected number of chosen vertices is hence $\sum_{i=1}^{n} \frac{k}{d_i+1}$.

And the rest is just algebra…

The function $1/(x+1)$ is convex (check this). Hence by Jensen's inequality, $\sum_{i=1}^{n} \frac{1}{d_i+1} \geq n\frac{1}{d+1}$, where d is the average degree ($\frac{\sum_{i=1}^{n} d_i}{n}$). Finally using $d = \frac{2E}{n}$ and multiplying by k on both sides we get the result. ■

Remark 1: Note that "$\frac{kn^2}{2E+n}$" in the problem can be replaced by "$\sum_{i=1}^{n} \frac{k}{d_i+1}$". Also note that $\frac{kn^2}{2E+n}$ is equal to $\frac{kn}{d+1}$, where d is the average degree of a vertex, that is, d = 2E/n. **Both results are extremely useful.** The summation form is obviously more useful in situations when we have (or can find) some information about the degrees of individual vertices in the graph we are analyzing (see example 5), whereas the other two forms are useful when we have some information about the total number of edges in a graph.

Corollary 1: Any graph $\boldsymbol{G}$ with n vertices and E edges contains an independent set containing at least $\frac{n^2}{2E+n}$ vertices. (Recall that an independent set is a set of vertices such that no two vertices in that set have an edge between them.)

Proof: Note that an independent set can be interpreted as a $\boldsymbol{K}_2$-free graph. Apply the above result for k = 1, and the corollary follows. ■

Corollary 2: A tree $\boldsymbol{T}$ on n vertices has an independent set of size greater than $n/3$.

Proof: Apply corollary 1, taking $E = n - 1$. ■
(Exercise: what is the best constant c such that a tree always has an independent set of size at least n/c?)

Remark 2: The above problem is representative of a typical class of problems in extremal graph theory. These problems ask us to show that in graphs under certain conditions there exist sufficiently large or small subgraphs with certain properties.

Example 3 [dominating sets]
Show that if the minimum degree in an n vertex graph $\boldsymbol{G}$ is $\delta > 1$, then $\boldsymbol{G}$ has a *dominating set* containing at most $n\,\frac{1+\ln(\delta+1)}{\delta+1}$ vertices.

Answer:
Form a subset $\boldsymbol{S}$ of the vertices of $\boldsymbol{G}$ by choosing each vertex with probability p, where p is a parameter in (0, 1) we will specify later. Let $\boldsymbol{T}$ be the set of vertices that have no neighbor in $\boldsymbol{S}$. Then $\boldsymbol{S} \cup \boldsymbol{T}$ is a dominating set. We now estimate its expected size.

Clearly $E[|S|] = np$. The probability that a vertex v is in $\boldsymbol{T}$ is the probability that neither v nor any of its neighbors is in $\boldsymbol{S}$. This is at most $(1-p)^{\delta+1}$, since v has at least d neighbors. Thus $E[|T|] \leq n(1-p)^{\delta+1} \leq ne^{-p(\delta+1)}$.

Finally, $E[|\boldsymbol{S} \cup \boldsymbol{T}|] \leq np + ne^{-p(\delta+1)}$. To minimize this expression, we choose $p = \frac{\ln(d+1)}{d+1}$, and we get $E[|\boldsymbol{S} \cup \boldsymbol{T}|] \leq n\,\frac{1+\ln(\delta+1)}{\delta+1}$. There exists some $\boldsymbol{S}$ and $\boldsymbol{T}$ such that $|\boldsymbol{S} \cup \boldsymbol{T}|$ is at most the expected value, and we are done. ■

Remark 1: The value of p minimizing the expression $np + ne^{-p(d+1)}$ is found by choosing p so that the derivative with respect to p is 0. Remember that p is a parameter we are free to choose, unlike n and d. We could have also directly chosen p to be the value minimizing $np + n(1-\mathrm{p})^{d+1}$, and we could have skipped the step where we estimate $n(1-p)^{d+1} \leq ne^{-p(d+1)}$. Although this would give a slightly sharper bound, it would have been a lot uglier. Anyway both bounds asymptotically approach $\frac{n \ln d}{d}$.

Remark 2: The basic idea in this proof is quite typical. Pick something randomly, and then make adjustments for whatever got missed.

Example 4 [USAMO 2010-6]

A blackboard contains 68 pairs of nonzero integers. Suppose that for no positive integer k do both the pairs (k, k) and $(-k, -k)$ appear on the blackboard. A student erases some of the 136 integers, subject to the condition that no two erased integers may add to 0. The student then scores one point for each of the 68 pairs in which at least one integer is erased. Determine, with proof, the largest number N of points that the student can guarantee to score regardless of which 68 pairs have been written on the board.

Note: The 68 pairs need not all be distinct; some may be repeated.

Answer:
Note that if (j, j) occurs then $(-j, -j)$ does not, so we can WLOG assume that if (j, j) occurs then $j > 0$ (by replacing j by $(-j)$).

Now for each integer $k > 0$, we can either delete all appearances of k on the board (if any) OR all appearances of $(-k)$ (if any), but not both. So, for each $k > 0$, we delete all k's with probability p, and otherwise delete all $(-k)$'s, where p is a parameter in $(1/2, 1)$ to be specified later. We now consider 3 types of pairs of numbers that can occur, and in each case bound the probability that we score a point.

(i) A pair of the form (k, k)
We score a point for this pair with probability p, since k is deleted with probability p.

(ii) A pair of the form $(k, -k)$
We score a point with probability 1, since we delete either k or $(-k)$

(iii) A pair of the form (a, b), where $b \neq \pm a$
We score a point with probability $(1 - P[a \text{ not deleted}] \times P[b \text{ not deleted}]) \geq (1 - p^2)$. Note that we are using $p > (1-p)$.

In all the above cases, the probability of scoring a point is **at least** $\min\{p, (1-p^2)\}$. Thus the expected number of points we score totally is at least $68 \times \min\{p, (1-p^2)\}$. This quantity is maximized by setting $p = \frac{\sqrt{5}-1}{2}$, and at this point the expectation is at least $68 \times p = 68 \times \frac{\sqrt{5}-1}{2}$, which is greater than 42. Therefore, it is always possible to score **at least 43 points**.

We leave it to the reader to construct an example demonstrating that it is not always possible to score 44 points.

This will show that the bound of 43 is tight, and hence the answer is 43. ■

Example 5 [IMO Shortlist 2012, C7]

There are 2^{500} points on a circle labeled 1, 2, …, 2^{500} in some order. Define the *value* of a chord joining two of these points as the sum of numbers at its ends. Prove that there exist 100 pairwise disjoint (nonintersecting) chords with equal values.

Answer:
Step 1: Basic observations
Let $n = 2^{499}$. There are $\binom{2n}{2}$ chords joining pairs of labeled points, and all chord *values* clearly belong to {3, 4, …, $4n$–1}. Furthermore, note that chords with a common endpoint have different values.

Step 2: Interpretation and reduction to familiar terms
Let $\boldsymbol{G_c}$ denote the graph with vertices representing chords with value c. Two vertices in $\boldsymbol{G_c}$ are neighbors if the corresponding chords intersect. (Note: don't confuse *vertices* in our graphs with *points* on the circle.) So we're basically just looking for a **large independent set** in *some* $\boldsymbol{G_c}$. We already know that by example 2, each $\boldsymbol{G_c}$ has an independent set of size at least $I(\boldsymbol{G_c}) = \sum_{v \in \boldsymbol{G_c}} \frac{1}{d_v+1}$.
By the pigeonhole principle, it now suffices to show that the average of $I(\boldsymbol{G_c})$ over all values c is at least 100.

Step 3: From graphs to individual vertices
Note that the average value of $I(\boldsymbol{G_c})$ is

$$\frac{1}{4n-3} \sum_{c=3}^{4n-1} I(\boldsymbol{G_c}) = \frac{1}{4n-3} \sum_{c=3}^{4n-1} \left(\sum_{v \in \boldsymbol{G_c}} \frac{1}{d_v+1}\right)$$

The double summation on the right is nothing but the sum of $\frac{1}{d_v+1}$ over **all 2*n* vertices**, since each vertex belongs to exactly one of the $4n$−3 graphs.

Step 4: Estimating the degrees
For a given chord L, let $m(L)$ denote the number of points contained by its minor arc. $m(L) = 0$ if the chord joins consecutive points and $m(L) = n-1$ if the chord is a diameter. Clearly, a chord L can intersect at most $m(L)$ other chords.

Now for each $i \in \{0, 1, 2, \ldots, n-1\}$, there are exactly $2n$ chords with $m(L) = i$. Thus for each $i \in \{0, 1, 2, \ldots, n-1\}$, there are at least $2n$ vertices with degree at most i. Therefore, the sum of $\frac{1}{d_v+1}$ over all vertices is at least $\sum_{i=1}^{n} \frac{2n}{i}$. Finally,

$$\text{Average of } I(\boldsymbol{G}_c) \geq \frac{2n}{4n-3} \sum_{i=1}^{n} \frac{1}{i} > \frac{1}{2} \ln(n+1)$$

$$> \frac{1}{2} \ln(2^{499}) = \frac{499 \ln(2)}{2} > 249.5 \times 0.69 > 100.$$

This proves the desired result. ■

Remark: In fact, $\frac{499 \ln(2)}{2} > 172$. Even with rather loose estimates throughout, this proof improves the bound significantly from 100 to 172. More generally, 2^{500} and 100 can be replaced by $2n$ and $\frac{2n \ln(n+1)}{4n-3}$.

Example 6 [Biclique covering]
In a certain school with n students, amongst any $(\alpha+1)$ students, there exists at least one pair who are friends (friendship is a symmetric relation here). The principal wishes to organize a series of basketball matches $M_1, M_2, \ldots, M_j$ under certain constraints. Each match M_i is played by two teams A_i and B_i, and each team consists of 1 or more students. Furthermore, each person in A_i must be friends with each person in B_i. Finally, for any pair of friends u and v, they must appear on opposite teams in at least one match. Note that a particular student can participate in

any number of matches, and different matches need not be played by a disjoint set of people. Let $|M_i|$ denote the number of students playing match M_i, and let S denote the sum of $|M_i|$ over all matches. Show that $S \geq n\log_2(n/\alpha)$.

Answer:
For each match M_i, randomly fix the winner by flipping a fair coin (note that we are free to do this as the problem mentions nothing about the winners of matches). Hence for each i, A_i wins with probability 0.5 and B_i wins with probability 0.5. Call a student a *champion* if he is on the winning team of each match that he plays.

Key observation: the total number of champions is at most α.
Proof: If there were α+1 champions, some pair must be friends (see first statement of the problem), and hence they must have played against each other in some match. But then one of them would have lost that match and would hence not be a champion; contradiction.

On the other hand, call the students $s_1, s_2, \ldots, s_n$. The probability that a student s_k becomes a champion is 2^{-m_k}, where m_k is the number of matches played by s_k. Hence the expected number of champions is $\sum_{k=1}^{n} 2^{-m_k}$. Applying the AM-GM inequality, $\sum_{k=1}^{n} 2^{-m_k} \geq n \times 2^{-(\sum_{k=1}^{n} m_k)/n} = n \times 2^{-S/n}$. But the number of champions is bounded by α, so the expectation obviously cannot exceed α. Hence $n \times 2^{-S/n} \leq \alpha$. This implies $S/n \geq \log_2(n/\alpha)$. Hence proved. ■

Remark 1: As a corollary, note that the number of matches must be at least $\log_2(n/\alpha)$, since each match contributes at most n people to the sum.

Remark 2: Interpreted graph theoretically, this problem is about covering the edges of a graph by bipartite cliques (complete bipartite graphs). S is the sum of the number of vertices in each bipartite clique and α is the independence number of the graph.

Mutually Exclusive Events

Two events A and B are said to be *mutually exclusive* if they cannot both occur. For instance, if an integer is selected at random, the events A: the integer is divisible by 5, and B: the last digit of the integer is 3 are mutually exclusive. More formally, $P(A|B) = P(B|A) = 0$. If n events are pairwise mutually exclusive, then no two of them can simultaneously occur. A useful property of mutually exclusive events is that if $E_1, E_2, \ldots, E_n$ are mutually exclusive, then the probability that some E_i occurs is equal to the sum of the individual probabilities: $P[\bigcup_{i=1}^{n} E_i] = P[E_1] + P[E_2] + \ldots + P[E_n]$. On the other hand, all probabilities are bounded by 1, so $P[\bigcup_{i=1}^{n} E_i] \leq 1$, which implies $P[E_1] + P[E_2] + \ldots + P[E_n] \leq 1$. Thus we have the following simple lemma:

Lemma: If $E_1, E_2, \ldots, E_n$ are mutually exclusive events, then $P[E_1] + P[E_2] + \ldots + P[E_n] \leq 1$.

Like so many other simple facts we have encountered in this book, this lemma can be exploited in non-trivial ways to give elegant proofs for several combinatorial results. The next five examples demonstrate how.

Example 7 [Lubell-Yamamoto-Meshalkin Inequality]
Let $A_1, A_2, \ldots, A_s$ be subsets of $\{1, 2, \ldots, n\}$ such that A_i is not a subset of A_j for any i and j. Let a_i denote $|A_i|$ for each i. Show that

$$\sum_{i=1}^{s} \frac{1}{\binom{n}{a_i}} \leq 1.$$

Answer
Take a random permutation of $\{1, 2, \ldots, n\}$. Let E_i denote the event that A_i appears as an initial segment of the permutation. (For example, if $n = 5$, the permutation is 3, 4, 2, 5, 1 and $A_2 = \{2, 3, 4\}$, then event E_2 occurs since the elements of A_2 match the first three

elements of the permutation.) The key observation is that the events E_i are **mutually exclusive**: if two different sets matched initial segments of the permutation, one set would contain the other. Also note that $P[E_i] = \frac{1}{\binom{n}{a_i}}$, as there are $\binom{n}{a_i}$ different choices for the first a_i elements. Therefore, the probability of *some* event occurring is $P[E_1] + P[E_2] + \ldots + P[E_s] = \sum_{i=1}^{s} \frac{1}{\binom{n}{a_i}}$. But probability is always at most 1, so $\sum_{i=1}^{s} \frac{1}{\binom{n}{a_i}} \leq 1$. ■

Corollary [Sperner's theorem]
Let $A_1, A_2, \ldots, A_s$ be subsets of {1, 2,..., n} such that A_i is not a subset of A_j for any i and j. Such a family of sets is known as an *antichain.* Show that $s \leq \binom{n}{\lfloor \frac{n}{2} \rfloor}$. In other words, an antichain over the power set P^n has cardinality at most $\binom{n}{\lfloor \frac{n}{2} \rfloor}$. Note that equality is achieved by taking all sets of size [n/2], since two sets of the same size cannot contain each other.

Proof:
Since $\binom{n}{\lfloor n/2 \rfloor}$ is the largest binomial coefficient, $\binom{n}{\lfloor n/2 \rfloor} \geq \binom{n}{a_i}$ for each a_i. Therefore $\sum_{i=1}^{s} \frac{1}{\binom{n}{a_i}} \geq \sum_{i=1}^{s} \frac{1}{\binom{n}{\lfloor n/2 \rfloor}} = \frac{s}{\binom{n}{\lfloor n/2 \rfloor}}$. Combining this with the Lubell-Yamamoto-Meshalkin inequality gives $\frac{s}{\binom{n}{\lfloor n/2 \rfloor}} \leq 1$, and the result follows. ■

Example 8 [Bollobas' Theorem]
Let $A_1, A_2, \ldots, A_n$ be sets of cardinality s, and let $B_1, \ldots, B_n$ be sets of cardinality t. Further suppose that $A_i \cap B_j = \emptyset$ if and only if $i = j$. Show that $n \leq \binom{t+s}{s}$.

Answer:
We use a visual trick similar to that used in the previous problem.

Take a random permutation of all elements in $\boldsymbol{S}$, where $\boldsymbol{S}$ is the union of all $2n$ sets. Let E_i be the event that all elements of A_i precede all elements of B_i in the permutation (in other words, the first element of B_i appears only after the last element of A_i). The key observation is that no two events E_i and E_j can occur simultaneously: if E_i and E_j occurred simultaneously, then $A_i \cap B_j$ or $A_j \cap B_i$ would be $\emptyset$, as these sets would be too "far apart" in the permutation. Hence, as in the previous problem, probability of *some* event occurring is $P[E_1] + \ldots + P[E_n] = \frac{n}{\binom{t+s}{s}} \leq 1$, implying the result. ■

Example 9 [Tuza's Theorem (Variant of Bollobas' Theorem)]
Let $A_1, A_2, \ldots, A_n$ and $B_1, B_2, \ldots, B_n$ be sets such that $A_i \cap B_i = 0$ for all i and for $i \neq j$ $(A_i \cap B_j) \cup (A_j \cap B_i) \neq \emptyset$. In other words, either $A_i \cap B_j \neq \emptyset$ or $A_j \cap B_i \neq \emptyset$ (or both). Then show that for any positive real number $p < 1$, $\sum_{i=1}^{n} p^{|A_i|} (1-p)^{|B_i|} \leq 1$.

Answer:
The form of the result we need, namely "$\sum_{i=1}^{n} p^{|A_i|} (1-p)^{|B_i|} \leq 1$", gives us a good clue as to what our random process should be. Let $\boldsymbol{U}$ be the universe of all elements in some A_i or B_i. Randomly select elements from $\boldsymbol{U}$ with probability p. Then $p^{|A_i|}(1-p)^{|B_i|}$ clearly represents the probability that all elements from A_i are selected and no elements from B_i are selected. Call this event E_i. For any pair $i \neq j$, we claim that E_i and E_j are independent. Suppose to the contrary that E_i and E_j both occur for some $i \neq j$. Then that means that all elements of $A_i \cup A_j$ have been selected, but no elements of $B_i \cup B_j$ have been selected. This implies that the sets $(A_i \cup A_j)$ and $(B_i \cup B_j)$ are disjoint, so $A_i \cap B_j = \emptyset$ and $A_j \cap B_i \neq \emptyset$. Contradiction. Thus the n events are pairwise mutually exclusive, so the probability that *some* event occurs is $P[E_1] + \ldots + P[E_n] \leq 1$, which implies $\sum_{i=1}^{n} p^{|A_i|} (1-p)^{|B_i|} \leq 1$. ■

Corollary:
Let $A_1, A_2, \ldots, A_n$ be sets of cardinality a and let $B_1, B_2, \ldots, B_n$ be sets

of cardinality b, such that $A_i \cap B_i = 0$ for all i and for $i \neq j$ $(A_i \cap B_j) \cup (A_j \cap B_i) \neq \emptyset$. Show that $n \leq \frac{(a+b)^{a+b}}{a^a b^b}$.

Proof:
Simply plug in $p = \frac{a}{a+b}$ in Tuza's theorem. ■

Example 10 [Original Problem, Inspired by China TST 1997]
There are n ice cream flavors offered at an ice cream parlor, and k different cone sizes. A group of kids buys ice creams such that no student buys more than one ice cream of a particular flavor. Suppose that for any two flavors, there exists some kid who buys both flavors but in different cone sizes. Show that the total number of ice creams bought by this group is at least $n\log_k n$.

Answer:
How do we go from innocuous ice cream to those nasty logarithms?! Answer: The Probabilistic Method.

Label the k sizes $S_1, S_2, \ldots, S_k$ and the n flavors $F_1, F_2, \ldots, F_n$. Now randomly label each student with one integer in $\{1, 2, \ldots, k\}$. For each flavor F_i, let E_i denote the event that "all students eating an ice cream of flavor F_i are eating the size corresponding to their label". For example, this would mean that a student with label 2 eating flavor F_i must be eating it in cone size S_2.

The probability of E_i is $k^{-|F_i|}$, where $|F_i|$ is the number of kids eating flavor F_i. This is because each of these $|F_i|$ kids has probability $1/k$ of eating flavor F_i from the "correct" size cone, as there are k sizes.

The key observation is that the E_i's are mutually exclusive: this follows from the fact that for any two flavors, there exists some kid who buys both flavors but in different cone sizes, meaning that he cannot be eating from the size corresponding to his label for both flavors.

It follows that $1 \geq P(\bigcup E_i) = \sum_{i=1}^{n} k^{-|F_i|} \geq n \times k^{-\sum_{i=1}^{n} |F_i|/n} =$

$nk^{-T/n}$, where $T = \sum_{i=1}^{n} |F_i|$ is the total number of ice creams bought. Here we have used the AM-GM inequality. Hence $k^{T/n} \geq n$, which implies that $T \geq n\log_k n$, as desired. ■

Example 11 [IMO Shortlist 2009, C4]

Consider a $2^m \times 2^m$ chessboard, where m is a positive integer. This chessboard is partitioned into rectangles along squares of the board. Further, each of the 2^m squares along the main diagonal is covered by a separate unit square. Determine the minimum possible value of the sum of perimeters of all the rectangles in the partition.

Answer:

Squares with position (i, i) have already been covered, and they symmetrically divide the board into two parts. By symmetry, we can restrict our attention to the bottom portion for now.

Suppose the portion of the board below the diagonal has been partitioned. Let R_i denote the set of rectangles covering at least one square in the ith row, and C_i the set of rectangles covering at least one square in the ith column. The total perimeter is $P = 2\sum(|R_i| + |C_i|)$, because each rectangle in R_i traces two vertical segments in row i and each rectangle in C_i traces two horizontal segments in column i.

Now randomly form a set $\boldsymbol{S}$ of rectangles by selecting each rectangle in the partition with probability ½. Let E_i denote the event that **all** rectangles from R_i are chosen and **no** rectangle from C_i is chosen. Clearly, $P[E_i] = 2^{-(|R_i|+|C_i|)}$. But for distinct i and j, E_i implies that the rectangle covering (i, j) **is** in $\boldsymbol{S}$, whereas E_j implies the rectangle covering (i, j) is **not** in $\boldsymbol{S}$. Hence **the $\mathbf{E_i}$'s are independent**.

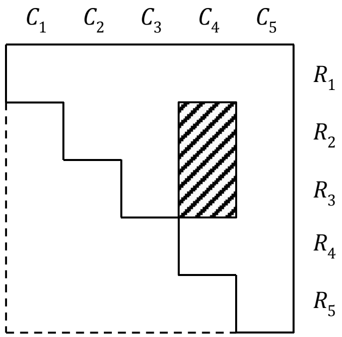

Figure 9.1. The shaded rectangle belongs to R2, R3 and C4. Its perimeter is 6, twice the number of sets it belongs to.

Therefore, $\sum 2^{-(|R_i|+|C_i|)} \leq 1$. But by the AM-GM inequality, $\sum_{i=1}^{2^m} 2^{-(|R_i|+|C_i|)} \geq 2^m \times 2^{-\sum(|R_i|+|C_i|)/2^m} = 2^m \times 2^{-P/2^{m+1}}$. Hence $2^{P/2^{m+1}} \geq 2^{\mathrm{m}}$, so $P \geq m2^{m+1}$.

The same bound holds for the upper portion, so the overall perimeter is at least $2P$ + sum of perimeters of unit squares on diagonal $\geq 2(m2^{m+1}) + 4(2^m) = (m+1)2^{m+2}$. To show that this can be achieved, use a recursive construction: split the board into 4 squares of sizes 2^{m-1} x 2^{m-1}. Two of them have no squares along the main diagonal; leave these squares as they are. Two of them already have unit squares along their diagonals; recursively partition these. ■

Bounding Violated Conditions

In practically all existence problems, we essentially need to prove the existence of an object that satisfies certain conditions. A useful idea is as follows: construct the object randomly, and then show that the expected number of conditions violated is less than 1. In other words, the expected number of "failures" - conditions that fail to hold - should be less than 1. Then *some* object must violate 0 conditions; this is the required object. The next four examples illustrate this idea.

Example 12 [Erdos' theorem on tournaments]
Say that a tournament has the property P_k if for every set of k players there is one who beats them all. If ${}^nC_k \times (1-2^{-k})^{n-k} < 1$, then show that there exists a tournament of n players that has the property P_k.

Answer:
Consider a random tournament of n players, i.e., the outcome of every game is determined by the flip of fair coin. For a set $\boldsymbol{S}$ of k players, let A_S be the event that no y not in $\boldsymbol{S}$ beats all of $\boldsymbol{S}$. Each y not in $\boldsymbol{S}$ has probability 2^{-k} of beating all of $\boldsymbol{S}$ and there are $n-k$ such possible y, all of whose chances are independent of each other. Hence $Pr[A_S] = (1-2^{-k})^{n-k}$. Thus the expected number of sets $\boldsymbol{S}$ such that event A_S occurs is at most ${}^nC_k \times (1-2^{-k})^{n-k}$, which is strictly smaller than 1. Thus, for some tournament $\boldsymbol{T}$ on n vertices, no event A_S occurs and we are done. ■

Example 13 [Taiwan 1997]
For $n \geq k \geq 3$, let $\boldsymbol{X} = \{1, 2, \ldots, n\}$and let $\boldsymbol{F}_k$ be a family of k-element subsets of $\boldsymbol{X}$ such that any two subsets in $\boldsymbol{F}_k$ have at most $k-2$ elements in common. Show that there exists a subset $\boldsymbol{M}$ of $\boldsymbol{X}$ with at least $\lfloor log_2 n \rfloor + 1$ elements and not containing any subset in $\boldsymbol{F}_k$.

Answer:
The basic plan
If $k \geq \log_2 n$ there is nothing to prove, so assume $k < \log_2 n$. Let $m = \lfloor log_2 n \rfloor + 1$. Our idea is to show that the expected number of sets in $\boldsymbol{F}_k$ that a randomly chosen m-element subset would contain is strictly less than 1.

How large is $|F_k|$?

Easy double counting provides this key detail. Since each $(k-1)$ element subset of $\boldsymbol{X}$ lies in at most one set in $\boldsymbol{F_k}$, and each set in $\boldsymbol{F_k}$ contains exactly k $(k-1)$-element subsets of $\boldsymbol{X}$, $|\boldsymbol{F_k}| \leq \frac{1}{k}\binom{n}{k-1} = \frac{1}{n-k+1}\binom{n}{k}$.

Finding the expectation

Now take a randomly chosen m-element subset $\boldsymbol{S}$ of $\boldsymbol{X}$. $\boldsymbol{S}$ contains $\binom{m}{k}$ k-element subsets of $\boldsymbol{X}$. Thus the expected number of elements of $\boldsymbol{F_k}$ it contains is

$$\frac{\text{number of } k \text{ element subsets of } X \text{ contained}}{\text{total number of } k \text{ element subsets of } X} \times (\text{number of subsets in } \boldsymbol{F_k})$$

$$= \frac{\binom{m}{k}|F_k|}{\binom{n}{k}} \leq \frac{\binom{m}{k}}{n-k+1}, \text{ using the bound on } |\boldsymbol{F_k}|.$$

The final blow

To prove that this expression is less than 1, it is enough to show that $\binom{m}{k} \leq \frac{3n}{4}$, since $\frac{3n}{4(n-k+1)} < 1$ for $n \geq 3$ (Check this: remember $k < \log_2 n$). Replacing $n \geq 2^{m-1}$, we just need $\binom{m}{k} \leq 3 \times 2^{m-3}$ for all $m \geq k$ and $m \geq 3$, which can be shown by induction on m: The base case holds, and observe that

$$\binom{m+1}{k} = \binom{m}{k} + \binom{m}{k+1} \leq 2 \max\{\binom{m}{k}, \binom{m}{k+1}\}$$

$$\leq 2 \times 3 \times 2^{m-3} = 2 \times 2^{m-2} \text{ (using the induction hypothesis).}$$

Thus the expectation is less than 1 and we are done. ■

Example 14 [Original problem, inspired by IMO Shortlist '87]

Let n, k, l and j be natural numbers with $k > l$, and $j > l \geq 1$. Let $\boldsymbol{S} = \{1, 2, 3, \ldots, n\}$. Each element of $\boldsymbol{S}$ is colored in one of the colors

c_1, c_2, …, c_j. Call a set of k terms in $\boldsymbol{S}$ in arithmetic progression *boring* if amongst them there are only l or fewer colors. (For instance, if $k = 4$, $l = 2$ and the numbers 1, 4, 7 and 10 are colored c_3, c_5, c_5 and c_3 respectively, then {1, 4, 7, 10} is boring as amongst these 4 numbers in AP there are just 2 colors.) Show that if $n < [(\frac{2k}{e^l})\,(\frac{j}{l})^{(k-l)}]^{1/2}$, there exists a coloring of $\boldsymbol{S}$ such that there are no boring sets. [Note: Here $e \approx 2.718$ is Euler's constant]

Answer:
Looks complicated? The solution is actually surprisingly straightforward. Take a random coloring (each number receives a particular color with probability $1/j$). Our goal is to show that the expected number of boring sets is less than 1. Then *some* outcome of this random coloring produces 0 boring sets. Now note that the expected number of boring sets is (number of k-term APs in $\boldsymbol{S}$) × (probability that a given k-term AP is boring). Denote the later quantity by p and the former by A. We just need to show that $pA < 1$.

Bounding A

Any k-term arithmetic progression $\{x, x+d, \ldots, x+(k-1)d\}$ is uniquely defined by its starting term x and its common difference, d where $x, d > 0$. Since we need all terms to be in $\boldsymbol{S}$, $x+(k-1)d \le n$, or $x \le n - (k-1)d$. Hence, for each $d \le n/(k-1)$, there are exactly $n - (k-1)d$ possible values x can take, and hence there are $(n - (k-1)d)$ k-term arithmetic progressions for each d. Therefore,

$$\mathrm{A} = \sum_{d=1}^{\left\lfloor \frac{n}{k-1} \right\rfloor} [n - (k-1)d]$$

$$= n\left\lfloor \frac{n}{k-1} \right\rfloor - \frac{1}{2}(k-1)\left\lfloor \frac{n}{k-1} \right\rfloor \left(\left\lfloor \frac{n}{k-1} \right\rfloor + 1\right) < \frac{n^2}{2k}$$

where the last step is simple (boring) algebra (do it yourself!). So $\mathrm{A} < \frac{n^2}{2k}$.

Bounding p

For a particular k-term AP, the total number of ways in which it can be colored is j^k. A coloring making it boring can be created by choosing l colors first and then coloring using only these l colors. There are $\binom{j}{l}l^k$ ways of doing this. (In fact, this may be higher than the total number of bad events as there is double counting - irrelevant here but worth noting). Therefore,

$$p \leq \binom{j}{l}l^k / j^k \leq \left(\frac{ej}{l}\right)^l \left(\frac{l}{j}\right)^k = e^l \left(\frac{j}{l}\right)^{l-k}.$$

And finally…

$pA < \frac{n^2}{2k} \times e^l\left(\frac{j}{l}\right)^{l-k} < 1$, using the bound on n given in the problem. ■

Example 15 [Based on a Result of Kleitman and Spencer]

Let n, m and k be positive integers with $m \geq k\, 2^k \ln(n)$ and $n > k > 3$. Define $\boldsymbol{S} = \{1, 2, \ldots, n\}$. Show that there exist m subsets of $\boldsymbol{S}$, $\boldsymbol{A}_1, \boldsymbol{A}_2, \ldots, \boldsymbol{A}_m$, such that for **any** k-element subset $\boldsymbol{T}$ of $\boldsymbol{S}$, the m intersections $\boldsymbol{T} \cap \boldsymbol{A}_i$ for $1 \leq i \leq m$ range over all 2^k possible subsets of $\boldsymbol{T}$.

Answer:

Take a few minutes to completely understand the question. Our basic strategy will be to choose the $\boldsymbol{A}_i$'s randomly and then show that the probability of "failure" is less than one.

Randomly create each set $\boldsymbol{A}_i$ by including each element with probability 0.5. Fix a pair $(\boldsymbol{T}, \boldsymbol{T}_1)$ where $\boldsymbol{T}$ is a k-element subset of $\boldsymbol{S}$ and $\boldsymbol{T}_1$ is a subset of $\boldsymbol{T}$. For each i, the probability that $\boldsymbol{A}_i \cap \boldsymbol{T} = \boldsymbol{T}_1$ is 2^{-k}, because we need each element in $\boldsymbol{T}_1$ to occur in $\boldsymbol{A}_i$ (0.5 probability each) and each element in $\boldsymbol{T} \setminus \boldsymbol{T}_1$ to **not** occur in $\boldsymbol{A}_i$ (0.5 probability each again). Therefore, the probability that **no** set $\boldsymbol{A}_i$ satisfies $\boldsymbol{A}_i \cap \boldsymbol{T} = \boldsymbol{T}_1$ is $(1 - 2^{-k})^m$.

There are $\binom{n}{k}$ choices for $\boldsymbol{T}$ and 2^k choices for $\boldsymbol{T}_1$ once we have chosen $\boldsymbol{T}$. Therefore, the overall probability of failure is at most

$$\binom{n}{k}2^k(1 - 2^{-k})^m \le \frac{n^k}{k!}2^k e^{-m/2^k}$$

$$\le \frac{n^k}{k!}2^k e^{-k2^k \ln(n)/2^k} = \frac{2^k}{k!} < 1, \text{ since } k > 3.$$

Thus, with positive probability, no condition of the problem is violated, and hence there exist m subsets of $\boldsymbol{S}$ satisfying the problem's requirements. ■

Useful Concentration Bounds

In many situations we want the outcome of a random process we design to closely match the expectation of that outcome. Concentration inequalities allow us to achieve this by bounding the probability that a random variable is far from its mean. The following three lemmas represent three of the most basic, widely used concentration inequalities. We omit rigorous measure theoretic proofs, but will provide some intuition as to why they hold.

1. **Markov's Bound**:
   If x is a nonnegative random variable and a is a positive number,
   $$P[x \ge a] \le E[x]/a.$$

   Equivalently written as
   $$P[x \ge aE[x]] \le 1/a.$$

   Intuition: It's not possible for 20% of a country's population to each earn more than 5 times the national average.

2. **Chebychev's Inequality:**
Chebychev's inequality allows us to bound the probability that a random variable is far from its mean in terms of its variance. It is a direct consequence of Markov's bound.

For a random variable x, applying Markov's bound to the positive random variable $(x-E[x])^2$, and applying the definition of variance,

$$P[|x-E[x]| \geq a] = P[(x-E[x])^2 \geq a^2] \leq \frac{E[(x-E[x])^2]}{a^2} = \frac{Var\{x\}}{a^2}.$$

This result is intuitive: The "spread" of the random variable is directly related to its variance.

3. **Chernoff's Bound**, simple version:
Chernoff's bound provides another bound on the probability that a random variable is far from its mean. It is particularly useful when we want relative guarantees rather than absolute, e.g. when we are not directly concerned about the difference between a random variable and its mean but rather want to estimate the probability that it is ten percent higher than its mean. In addition, we do not need information about the variance, unlike in Chebychev's inequality.

Let X be the sum of n independent 0-1 random variables X_1, X_2, …, X_n, where $E[X_i] = p_i$. Denote $E[X] = \sum_{i=1}^{n} p_i$ by μ. Then for any $0 < \delta < 1$,
(i) $P[X > (1+\delta)\mu] \leq e^{-\delta^2\mu/3}$
(ii) $P[X < (1-\delta)\mu] \leq e^{-\delta^2\mu/2}$

In some cases, it may be clear that you need to use a concentration inequality but you may not be sure which one to use. In such cases, there's no harm in trying all three of these (as long as the variance is not difficult to compute). These three bounds are the only ones you will need for the exercises in this book. However, there exist a wealth of other powerful, more specialized concentration inequalities used in probabilistic combinatorics, and interested readers are encouraged to refer to online resources devoted to these results.

The next example represents a very typical and intuitive use of concentration inequalities. Essentially, we want something to be roughly balanced, so we designed a randomized procedure that would, in expectation, produce something perfectly balanced. Then, we used a concentration inequality to show that with nonzero probability, we would get something "not too unbalanced."

Example 16 [Hypergraph discrepancy]
Let ***F*** be a family of m subsets of {1, 2, …, n} such that each set contains at most s elements. Each element of {1, 2, …, n} is to be colored either red or blue. Define the *discrepancy* of a set ***S*** in ***F***, denoted *disc*(***S***), to be the absolute value of the difference between the number of red elements in ***S*** and blue elements in ***S***. Define the discrepancy of ***F*** to be the maximum discrepancy of any set in ***F***. Show that we can color the elements such that the discrepancy of ***F*** is at most $\lceil 2\sqrt{s\ ln(2m)} \rceil$.

Answer:
Color each number in {1, 2, …, n} red with probability 0.5 and blue otherwise. We just need to show that for a particular ***S*** in ***F***, the probability that ***S*** has discrepancy higher than our desired threshold is less than $1/m$, for then the result will follow from the union bound.

Let ***S*** be a set with t elements, and denote by b the number of blue elements and by r the number of red elements of ***S***. Note that for a fixed number $d < t$, *disc*(***S***) > d if and only if either $b < t/2 - d/2$ or $r < t/2 - d/2$. Hence the probability that *disc*(***S***) > d is equal to

$$P[disc(\boldsymbol{S}) > d] = P[b < t/2 - d/2] + P[r < t/2 - d/2]$$

$$= 2P[b < t/2 - d/2] \text{ by symmetry.}$$

Note that b is the sum of t independent random 0-1 variables, and $E[b] = t/2$. Write $\mu = t/2$, $\delta = d/t$. Then applying the Chernoff bound and noting that $t \le s$,

$$2P[b < t/2 - d/2] = 2P[b < (1 - \delta)\,\mu] \le 2e^{-\delta^2\mu/2}$$

$$= 2\exp(-d^2/4t) \le 2\exp(-d^2/4s).$$

For the last quantity to be less than $1/m$, we just need $d > 2\sqrt{s\ ln(2m)}$, as desired. ■

Remark: Note that using the Chebychev inequality for this problem would prove a weaker bound. It is instructive for readers to try this, as it is a good exercise in computing the variance of a random variable.

The Lovasz Local Lemma

In the section on bounding "violated conditions," we used the fact that if the expected number of violated conditions is less than 1, then there exists an object satisfying all conditions. Unfortunately, we aren't always this lucky: in many cases, the expected number of violated conditions is large. The *Lovasz local lemma* provides a much sharper tool for dealing with these situations, especially when the degree of dependence (defined below) between violated conditions is low.

Definition: Let E_1, E_2, …, E_n be events. Consider a graph $\boldsymbol{G}$ with vertices v_1, …, v_n, such that there is an edge between v_i and v_j if and only if events E_i and E_j are **not** independent of each other. Let d be the maximum degree of any vertex in ***G***. Then d is known as the ***degree of independence*** between the events $E_1, E_2, \ldots, E_n$.

Theorem [The Lovasz Local Lemma]:
Let $E_1, E_2, \ldots, E_n$ be events with degree of dependence at most d and $P[E_i] \leq p$ for each $1 \leq i \leq n$, where $0 < p < 1$. If $ep(d+1) \leq 1$, then $P[E_1' \cup E_2' \cup \ldots \cup E_n'] > 0$, where E_i' denotes the complement of event E_i. In other words, if $ep(d+1) \leq 1$, then the probability of **none of the *n* events occurring** is strictly greater than 0. (Here e is Euler's constant.)

The Lovasz local lemma hence gives us a new way to solve problems in which we need to show that there exists an object violating 0 conditions. Given n conditions that need to hold, define E_i to be the event that the ith condition is violated. All we need to do is show compute p and d for these events, and show that $ep(d+1) \leq 1$.

Example 17 [Generalization of Russia 2006]
At a certain party each person has at least δ friends and at most Δ friends (amongst the people at the party), where $\Delta > \delta > 1$. Let k be an integer with $k < \delta$. The host wants to show off his extensive wine collection by giving each person one type of wine, such that each person has a set of at least $(k+1)$ friends who all receive different types of wine. The host has W types of wine (and unlimited supplies of each type).

Show that if $W > k\,[e^{k+1}(\Delta^2 - \Delta + 1)]^{1/(\delta-k)}$, the host can achieve this goal.

Answer:
Interpret the problem in terms of a friendship graph (as usual), and interpret "types of wine" as coloring vertices. Note the presence of *local* rather than *global* information- the problem revolves around the neighbors and degrees of individual vertices,

but doesn't even tell us the total number of vertices! This suggests using the Lovasz Local Lemma. We will show that $ep(d+1) \leq 1$, where p is the probability of a *local bad event* E_v (defined below) and d is the degree of dependency between these events.

Defining E_v and bounding $p = P[E_v]$

Randomly color vertices using the W colors (each vertex receives a particular color with probability $1/W$). Let E_v denote the event that vertex v does **not** have $(k+1)$ neighbors each receiving a different color. Hence,

$$P[E_v] \leq \binom{W}{k} \times \left(\frac{k}{W}\right)^{d_v} \leq \binom{W}{k} \times \left(\frac{k}{W}\right)^{\delta} \leq \left(\frac{k}{W}\right)^{\delta}\left(\frac{eW}{k}\right)^{k} = \mathrm{e}^{k}\left(\frac{W}{k}\right)^{k-\delta}.$$

Here we used $d_v \geq \delta$.

Bounding d

E_v and E_u will be non-independent if and only if v and u have a common neighbor - in other words, u must be a neighbor of one of v's neighbors. v has at most Δ neighbors, and each of these has at most $(\Delta-1)$ neighbors apart from v. Hence for fixed v, there are at most $(\Delta^2-\Delta)$ choices for u such that E_v and E_u are dependent. It follows that $d \leq (\Delta^2-\Delta)$.

And finally…

$ep(d+1) \leq e \times e^{k}\left(\frac{W}{k}\right)^{k-\delta} \times (\Delta^2-\Delta+1)$. Rearranging the condition in the problem (namely $W \geq k\ [e^{k+1}(\Delta^2-\Delta+1)]^{1/(\delta-k)}$) shows that this expression is at most 1. Hence proved. ■

Exercises

1. **[Max cut revisited]**
Give a probabilistic argument showing that the vertex set $\boldsymbol{V}$ of a graph $\boldsymbol{G}$ with edge set $\boldsymbol{E}$ can be partitioned into two sets $\boldsymbol{V}_1$ and $\boldsymbol{V}_2$ such that at least $|\boldsymbol{E}|/2$ edges have one endpoint in $\boldsymbol{V}_1$ and one endpoint in $\boldsymbol{V}_2$.

2. **[Hypergraph coloring]**
Let $\boldsymbol{F}$ be a family of sets such that each set contains n elements and $|\boldsymbol{F}| < 2^{n-1}$. Show that each element can be assigned either the color red or the color blue, such that no set in $\boldsymbol{F}$ is monochromatic (contains elements of only one color).

3. **[A Bound on Ramsey Numbers]**
The *Ramsey number* $R(k, j)$ is defined as the least number n such that for any graph $\boldsymbol{G}$ on n vertices, $\boldsymbol{G}$ either contains a clique of size k or an independent set of size j. Show that $R(k, k) \geq 2^{k/2}$ for each k.

4. **[Stronger version of IMO Shortlist 1999, C4]**
Let $\boldsymbol{A}$ be a set of n different residues mod n^2. Then prove that there exists a set $\boldsymbol{B}$ of n residues mod n^2 such that the set $\boldsymbol{A}+\boldsymbol{B} = \{a+b \mid a \in \boldsymbol{A}, b \in \boldsymbol{B}\}$ contains at least $(e-1)/e$ of the residues mod n^2.

5. **[Another hypergraph coloring problem]**
Let $\boldsymbol{F}$ be a family of sets such that each set contains n elements. Suppose each set in $\boldsymbol{F}$ intersects at most 2^{n-3} other sets in $\boldsymbol{F}$. Show that each element can be assigned either the color red or the color blue, such that no set in $\boldsymbol{F}$ is monochromatic. Improve this bound (that is, replace 2^{n-3} by a

larger number and prove the result for this number).

6. [Based on Russia 1999]
Let $\boldsymbol{G}$ be a bipartite graph with vertex set $\boldsymbol{V} = \boldsymbol{V}_1 \cup \boldsymbol{V}_2$. Show that $\boldsymbol{G}$ has an induced subgraph $\boldsymbol{H}$ containing at least $|\boldsymbol{V}|/2$ vertices, such that each vertex in $\boldsymbol{V}_1 \cap \boldsymbol{H}$ has odd degree in $\boldsymbol{H}$.

7. [USAMO 2012, Problem 2]
A circle is divided into 432 congruent arcs by 432 points. The points are colored in four colors such that some 108 points are colored Red, some 108 points are colored Green, some 108 points are colored Blue, and the remaining 108 points are colored Yellow. Prove that one can choose three points of each color in such a way that the four triangles formed by the chosen points of the same color are congruent.

8. [List coloring]
Each vertex of an n-vertex bipartite graph $\boldsymbol{G}$ is assigned a list containing more than $\log_2 n$ distinct colors. Prove that $\boldsymbol{G}$ has a proper coloring such that each vertex is colored with a color from its own list.

9. Let $\boldsymbol{A}$ be an $n \times n$ matrix with distinct entries. Prove that there exists a constant $c > 0$, independent of n, with the following property: it is possible to permute the rows of $\boldsymbol{A}$ such that no column in the permuted matrix contains an increasing subsequence of length $c\sqrt{n}$. (Note that consecutive terms in a subsequence need not be adjacent in the column; "subsequence" is different from "substring.")

10. [IMO Shortlist 2006, C3]
Let $\boldsymbol{S}$ be a set of points in the plane in general position (no three lie on a line). For a convex polygon P whose vertices are in $\boldsymbol{S}$, define $a(P)$ as the number of vertices of P and $b(P)$ as the number of points of $\boldsymbol{S}$ that are outside P. (Note: an empty set , point and line segment are considered convex polygons with 0, 1 and 2 vertices respectively.) Show that for each real

number x,

$$\sum x^{a(P)}(1-x)^{b(P)} = 1,$$

where the sum is taken over all convex polygons P.

11. [Bipartite Expanders]

A bipartite graph $\boldsymbol{G}$ with vertex set $\boldsymbol{V} = \boldsymbol{V}_1 \cup \boldsymbol{V}_2$ is said to be an (n, m, d, β) *bipartite expander* if the following holds:

(i) $|\boldsymbol{V}_1| = n$ and $|\boldsymbol{V}_2| = m$
(ii) Each vertex in $\boldsymbol{V}_1$ has degree d
(iii) For any subset $\boldsymbol{S}$ of $\boldsymbol{V}$ with $|\boldsymbol{S}| \leq n/d$, there are at least $\beta|\boldsymbol{S}|$ vertices in $\boldsymbol{V}_2$ that have a neighbor in $\boldsymbol{S}$

Show that for any integers $n > d \geq 4$, there exists an $(n, n, d, d/4)$ bipartite expander.

12. [Sphere packing]

Let n be a given positive integer. Call a set $\boldsymbol{S}$ of binary strings of length n *α-good* if for each pair of strings in $\boldsymbol{S}$, there exist at least $n\alpha$ positions in which the two differ. For instance, if $n = 100$ and $\alpha = 0.1$, then any pair of strings in $\boldsymbol{S}$ must differ in at least 10 positions. Show that for each integer n and real number $0 < \alpha < 0.5$, there exists an α-good set $\boldsymbol{S}$ of cardinality at least $\left\lfloor \sqrt{2}e^{n(0.5-\alpha)^2} \right\rfloor$.

As a (rather remarkable) corollary, deduce that the unit sphere in n dimensions contains a set of $\left\lfloor \sqrt{2}e^{n/16} \right\rfloor$ points such that no two of these points are closer than one unit (Euclidean) distance from each other.

Remark: The problem of finding large sets of binary strings of a given length such that any two differ in sufficiently many positions is one of the central problems of coding theory. Coding theory is a remarkable field lying in the intersection of mathematics and computer science, and uses techniques from combinatorics, graph theory, field theory, probability and linear algebra. Its applications range from file compression to mining large scale web data. Coding theory is also the reason your CDs often work even if they get scratched. While this exercise asks for an existence proof, finding constructive solutions to several similar problems remains an active research area with high practical relevance.

13. Let ***F*** be a collection of k-element subsets of {1, 2, ..., n} and let $x = |\boldsymbol{F}|/n$. Then there is always a set ***S*** of size at least $\frac{n}{4x^{1/(k-1)}}$ which does not completely contain any member of ***F***.

14. [Another list coloring theorem]
Each vertex of an n-vertex graph ***G*** is assigned a list containing at least k different colors. Furthermore, for any color c appearing on the list of vertex v, c appears on the lists of at most $k/2e$ neighbors of v. Show that there exists a proper coloring of the vertices of ***G*** such that each vertex is colored with a color from its list.

15. [Due to Furedi and Khan]
Let ***F*** be a family of sets such that each set in ***F*** contains at most n elements and each element belongs to at most m sets in ***F***. Show that it is possible to color the elements using at most $1+(n-1)m$ colors such that no set in ***F*** is monochromatic.

16. [Due to Paul Erdos]
A set ***S*** of distinct integers is called *sum-free* if there does not exist a triple $\{x, y, z\}$ of integers in ***S*** such that $x + y = z$. Show that for any set ***X*** of distinct integers, ***X*** has a sum-free subset ***Y*** such that $|\boldsymbol{Y}| > |\boldsymbol{X}|/3$.

17. [IMO 2012, Problem 3]
The *liar's guessing game* is a game played between two players A and B. The rules of the game depend on two positive

integers k and n which are known to both players.

A begins by choosing integers x and N with $1 \leq x \leq N$. Player A keeps x secret, and truthfully tells N to player B. Player B now tries to obtain information about x by asking player A questions as follows: each question consists of B specifying an arbitrary set $\boldsymbol{S}$ of positive integers (possibly one specified in some previous question), and asking A whether x belongs to $\boldsymbol{S}$. Player B may ask as many questions as he wishes. After each question, player A must immediately answer it with *yes* or *no*, but is allowed to lie as many times as she wants; the only restriction is that, among any (k+1) consecutive answers, at least one answer must be truthful.

After B has asked as many questions as he wants, he must specify a set $\boldsymbol{X}$ of at most n positive integers. If x belongs to $\boldsymbol{X}$, then B wins; otherwise, he loses. Prove that:

1. If $n \geq 2^k$, then B can guarantee a win.
2. For all sufficiently large k, there exists an integer $n \geq 1.99^k$ such that B cannot guarantee a win.

18. [Johnson-Lindenstrauss Lemma]

(i) Let $X_1, X_2, \ldots, X_d$ be d independent Gaussian random variables. Let $\boldsymbol{X} = (X_1, X_2, \ldots, X_d)$ be the d-dimensional vector whose coordinates are $X_1, X_2, \ldots, X_d$. Let $k < d$ and let $\boldsymbol{X}'$ be the k-dimensional vector $(X_1, X_2, \ldots, X_k)$. Show that:

a) If α is a constant greater than 1,

$$\text{Prob}[\ ||X_k|| \geq \alpha\, k/d\ ] \leq \exp\left(\frac{k(1-\alpha+\ln\alpha)}{2}\right)$$

b) If α is a constant smaller than than 1,

$$\text{Prob}[\ ||X_k|| \leq \alpha\, k/d\ ] \leq \exp\left(\frac{k(1-\alpha+\ln\alpha)}{2}\right)$$

(ii) Part (i) says that the norm of a vector of independent randomly distributed Gaussian variables projected onto its first k coordinates is sharply concentrated around its expectation. Equivalently, the norm of a fixed vector projected onto a random k-dimensional subspace is sharply concentrated around its expectation. Use this result to show that given a set $\boldsymbol{S}$ of n points in a d-dimensional space, there is a mapping f to a subspace of dimension $O(n \log n/\varepsilon^2)$ such that for all points u, v in $\boldsymbol{S}$,

$$(1-\varepsilon)\ \mathrm{dist}(u, v) \leq \mathrm{dist}(f(u), f(v)) \leq (1+\varepsilon)\ \mathrm{dist}(u, v).$$