Algorithms for Combinatorics - Atli FF:

Introduction:

This text will consider algorithms that are relevant to solving problems within combinatorics. This of course overlaps with the theory of algorithms, which belongs to computer science. We will focus on the combinatorics section of things, but plenty of optional exercises will dip into computer science territory.

An algorithm is a finite sequence of mathematically rigorous instructions. What exactly this means in perfectly precise mathematical terms (turing machines), is beyond the scope of this text. Despite this we will find plenty of use in defining and utilising algorithms for combinatorial proofs, most commonly for constructive proofs. A finite sequence of mathematically rigorous instructions for constructing an object for a constructive proof is just proof by algorithm. Similarly many proofs in game theory which just give a winning strategy are essentially producing an algorithm as a proof.

There's many algorithms used on the regular in mathematics which we could start with, but quite a few of them are a bit too trivial to be of much interest (though the process of multiplying numbers does hide quite a lot more complexity than one might expect). So we will look first at the Euclidean algorithm.

Definition (Euclidean Algorithm)**.** Start with two non-negative integers x, y. While neither is zero, subtract the smaller one from the larger one. The number you end up with, aside from the zero, is the greatest common divisor of the starting numbers.

Computer scientists would format this algorithm somewhat differently, for good reason, but we will let this suffice for now. At this point one could investigate various properties of this algorithm. For example, is it correct? How many steps does it take to get to an answer?

Greedy algorithms:

An algorithm that always makes locally optimal moves is called a greedy algorithm. For some kinds of problems this will also give a globally optimal solution as well. Let's consider a very simple example.

Definition (Cashier's Algorithm)**.** Start with an amount of money x to return to a customer. While $x > 0$ give the customer the largest denomination possible out of the coins available.

Now, the optimality of this algorithm will depend on the coin system. For example if the coins denominations are 5 and 2 and the cashier is to return 11, this won't even return a valid answer. And for a system with denominations $1, 8, 20$ we would return $20, 1, 1, 1, 1$ for 24 instead of the optimal $8, 8, 8$. But for all sensibly designed coin denominations this will be optimal.

But this highlights important properties of greedy algorithms. One might have to prove they even return a solution at all, and in the case that they do, that this solution is indeed globally optimal. Not much can be said in the general case of how to prove the solution will be valid. But for proving optimality there are two general methods that are often employed.

Suppose we have a set of lilypads at coordinates $x_1 < x_2 < \cdots < x_n$ along the x-axis. A frog can jump at most distance r. Supposing it starts at x_1 and can reach x_n, how should it reach x_n in the fewest jumps possible? Intuitively it feels obvious that it should just jump as far to the right as possible until we can reach x_n. But how do we prove it?

We have to be careful in proving it, as tweaking the problem just a bit will make the algorithm unoptimal. For example if the frog can jump any distance in $[\epsilon, r]$ instead of $[0, r]$ this will fail. To prove this one can use a method called "Greedy stays ahead", the name of which should be rather self-explanatory. Working through proving that the lilypad algorithm is optimal by showing it always stays ahead of any other solution is a good exercise. Just let J be the set of jumps done by the algorithm and J^* be some optimal set of jumps. Then by proving that the first i jumps of J go at least as far as the first i jumps of J^* by induction takes you there.

The other one is the exchange argument. For this we look at Prim's algorithm. Suppose we have a weighted graph G. We want to find a spanning tree in G of minimal weight. Prim's algorithm starts at some vertex v, then repeatedly picks the cheapest edge that leaves the connected component of v. Showing this produces a spanning tree is not so hard, but how do we know it's of minimal weight?

Let T be the set of edges chosen by Prim's algorithm and T^* be an optimal set of edges. Let $c(T)$ denote the total cost of a set of edges. We will prove that $c(T) = c(T^*)$. If $T = T^*$ the result is obvious, so we can assume $T \neq T^*$. Then let (u, v) be an edge in $T \setminus T^*$. Let S be the connected component of the algorithm's starting vertex when (u, v) was added to T and V be the set of all vertices. Then (u, v) is the cheapest esdge from S to $V \setminus S$. Since T^* is spanning it must contain some path from u to v. This path begins in S and ends in $V \setminus S$, so there is some edge (x, y) in the path such that $x \in S$ and $y \in V \setminus S$. Since (u, v) was the cheapest we must have $c(\{u, v\}) \leq c(\{x, y\})$.

Now we can do the exchange step, let $T' = T^* \cup \{(u, v)\} \setminus \{(x, y)\}$. Since T^* doesn't use (u, v) to connect S and $V \setminus S$ internally, T' connects those components as well. But then we see that T' is spanning as well. But $c(T') = c(T^*) + c(\{u, v\}) - c(\{x, y\}) leq c(T^*)$. Since T^* is optimal this means $c(T') \geq c(T^*)$, so they are equal. We also see that $|T \setminus T'| = |T \setminus T^*| - 1$ so if we repeat this argument enough times we will have converted T^* into T without changing the weight, so $c(T) = c(T^*)$.

Those are two ways to prove optimality, so let's look at examples involving showing the algorithm produces a solution. The following is from the IMO 2001 shortlist. A set of three nonnegative integers $\{x, y, z\}$ with $x < y < z$ satisfying $\{z - y, y - x\} = \{1776, 2001\}$ is called a historic set. Show that the set of all nonnegative integers can be written as a disjoint union of historic sets.

Let $a = 1776$ and $b = 2001$. The historic sets are of the form $\{x, x + a, x + a + b\}$ or $\{x, x + b, x + a + b\}$, which we will call small and big sets respectively.

We want to cover every number using these sets. Generally when we make greedy algorithms we want to make extremal choices, take the biggest, take the smallest, first, last, etc. Here we choose to try to cover the smallest number and work our way up. So at any given step let x be the smallest number not yet covered and use a small set if possible, otherwise use the big set.

So how could this fail? Let x_i be the value chosen in the i-th step and suppose we get stuck in the $(n+1)$-st step. By the definition of our algorithm x_{n+1} will be uncovered when we choose it. Furthermore $x_{n+1} > x_i$ for all $i < n+1$ by the definition of our algorithm. This means $x_{n+1} + a + b$ will be uncovered as we could never have chosen a value this large before. So the only option for how we get stuck is that moth $x_{n+1} + a$ and $x_{n+1} + b$ are covered already. We see that $x_{n+1} + b$ must have been the largest value in its set, so the smallest number in that set must have been $x_{n+1} - a$. But when choosing that set x_{n+1} was uncovered, so the algorithm would have used a small set, covering x_{n+1}. This is a contradiction, so the algorithm produces a solution.

Here one could have designed the solution but deferred the decision of which type of set to prioritise. Then once that becomes relevant in the proof as one writes it up, one could either see that choosing the small set is better, or just try writing out the proof both ways and seeing which one works.

Let us consider one more example, the sixth problem on IMO 2014. Prove that for all sufficiently large n, in any set of n lines in general position it is possible to colour at least $\sqrt{n}$ lines blue such that none of its finite regions has a completely blue boundary.

We solve it greedily, we just colour lines blue until we get stuck. What does it mean that we got stuck? Well, suppose we ended up with k blue lines and $n-k$ uncoloured lines. We have $\binom{k}{2}$ intersections of blue lines, call them blue vertices. We also note that every uncoloured line must be part of a boundary to a region where it is the only uncoloured line bounding that region. For each such line select one such region and take the next counterclockwise vertex. This vertex must be blue, and it's not too hard to show each blue vertex is chosen at most twice. Thus we get that $n - k \leq 2\binom{k}{2}$ which gives $n \leq k^2$ as desired.

Invariants and Monovariants:

Many problems will ask not just for existence but in what cases a solution exists. Then a constructive algorithm is useful for the cases where a solution does indeed exist, but we generally need some other tools to prove non-existence. A very common form of proof for results of this form involves invariants or monovariants. An invariant is some quantity in the problem that never changes, while a monovariant is some quantity that can only increase or only decrease.

A nearly canonical example of a problem solveable with this approach is the following. Two diagonally opposite corners of a 8×8 chessboard are removed. Is it possible to completely cover the remaining 62 squares with 31 2×1 dominoes?

It is in fact impossible, and the reason is the difference between the number of covered white and black squares. This difference can never change as each domino must cover one black square and one white square. But if we remove two opposite corners at the start we removed either two black squares or two white squares, so this difference is 2 to begin with. To completely cover everything we must reach a difference of zero, so therein lies the impossiblity.

Many invariants involve colourings such as this. But many invariants are also numeric. So let us take a look at a numerical invariant. We consider a problem from the 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 $a + b \leq 3$.

We want the invariant to be some numeric quantity such that the move does not change it. A natural choice would be to associate each pebble to a numeric quantity based on its position, then sum all of those. After a bit of experimentation one might arrive at the value 2^{i+j} for position (i, j). Then clearly the sum is invariant as each pebble is always replaced with two pebbles of half the weight. But then the question is how this solves the problem? Well, we can sum all the pebble values outside $a+b \leq 3$ and see (calculations omitted) that they sum to $3/4$, but the initial sum is 1 and never changes. So we must always have some pebble with $a + b \leq 3$.

Invariants can be much more involved than this, but often relatively simple invariants can get strong results. Despite that, let us look at a much more far-fetched example.

We start with the three numbers $0, 1, \sqrt{2}$, to which the following operation is repeatedly applied: one of the numbers is chosen and an arbitrary rational multiple of the difference of the two others is added. Is it possible to obtain the triple $0, \sqrt{2} - 1, \sqrt{2} + 1$ after a number of applications of this operation?

To solve this we will construct a very clever invariant. Since $\sqrt{2}$ is irrational, all numbers that are obtained during this process must have the form $a + b\sqrt{2}$ for some rational a, b. Such a number can be represented by the point (a, b) in the plane. We consider the triangle that is formed by the three numbers. In the beginning its vertices are $(0, 0), (1, 0)$ and $(1, 0)$. The described operation amounts to a translation of one of the points along a line parallel to the opposite side of the triangle. This operation does not change the area of the triangle, so the area remains constant. In the beginning, the area is $1/2$. The triangle that is formed by the three points $(0, 0)$, $(-1, 1)$ and $(1, 1)$ has area 1 however. So it is impossible to reach the triple $0, \sqrt{2} - 1, \sqrt{2} + 1$.

Further reading:

For those willing to dip their toes a bit further into computer science topics, there are some more algorithms that have very nice pure math applications. Many of them reside within graph theory, with algorithms like Gale-Shepley and Ford-Fulkerson.

Winning strategies in combinatorial games are also nothing but algorithms, so for further applications of these techniques one could read more about combinatorial games. Common winning strategies such as mirroring, common invariants such as colourings and more.

Problems:

Problems rated on a scale of difficulty from 1 to 5. All ratings are my subjective opinion, based on how much trouble I had. Problems with several subproblems are rated by their hardest subproblem. These difficulty levels are only my subjective experience with solving the problems, but I hope they give a rough idea at least. Starred problems require results from outside this lecture that the reader might not be familiar, but is standard for the IMO. Double star means the same but it might not be standard IMO material. Problems marked with a dagger dip a bit more into computer science, and might be slightly less relevant to those only interested in the combinatorics side of things. The ones with a double dagger are just straight up computer science problems. The source of the problems is given after the difficulty. I tried to give the earliest source I could find in each case.

- (1, Icelandic final contest 2026) You start with a (finite) list L of numbers. You replace this list L with a new list that counts how often each number appeared in L. For example if a number appeared thrice in L the new list contains a three. If $L = [7, 1, 1, 1, 4, 5, 4]$ then the new list is $[1, 1, 2, 3]$. Then we replace the new list with a newer one in the same way, repeatedly. Prove that the list eventually converges to containing just a single copy of the number 1.

- (1, classic) Alice is thinking of a number from 1 to 1000. If you guess it wrong, she will tell you if the true value is higher or lower than your guess. Show that you can find the answer in a maximum of 10 guesses in all cases.

- (1, classic) Show that the vertices of a graph G with maximum degree Δ can be coloured using at most $\Delta + 1$ colours.

- (2, Icelandic Final Contest 2022) Atli and Fanney play a game on a rectangular board with an $n \times m$ grid. They start with every cell of the grid empty and then alternate taking turns. On their turn a player may colour one or more cells that are all in the same column or in the same row. Each cell may only be coloured once. The player who colours the last cell wins. Atli moves first. Who wins on a 4×6 board? But a 5×6 board?

- (2, University of Iceland programming contest 2020) You have $n+1$ spots, of which the first n of them have numbers on them. The one move you may make is moving an number from its spot to the one unoccupied spot. For a given permutation π on $1, 2, \ldots, n$, how many moves do you need to put the numbers in ascending order in the first n spots?

- (2, OTIS Excerpts) Suppose 4951 distinct points in the plane are given such that no four points are collinear. Show that it is possible to select 100 of the points for which no three points are collinear.

- (2, Icelandic high school programming contest 2026) You have n boys, m girls and k non-binary kids to arrange into a line with $n+m+k$ seats. You don't want to place any two boys adjacent to each other, nor two girls or two non-binary kids. For which n, m, k is this possible?

- (2, OTIS Excerpts) Let G be a finite simple graph. Show that one can partition the vertices into two groups such that for each vertex, at least half the neighbors are in the other group.

- ($2^{*\dagger}$, University of Iceland programming contest 2022) You are given $n > 0$ variables, all of which must be set to `TRUE` or `FALSE`. You are then given $m > 0$ clauses of the form x_i `OR` x_j `OR` x_k for some indices i, j, k. Sho w that you can choose values for the n variables to satisfy at least $\lfloor 7m/8 \rfloor$ clauses.

- (2, Rutgers Spring Contest) For $n \geq 4$ show there exists a permutation $\pi_1, \ldots, \pi_n$ of the numbers $1, \ldots, n$ such that $|\pi_{i+1} - \pi_i|$ is prime for each $i = 1, \ldots, n-1$.

- ($2^{\dagger}$, BAPC preliminaries 2022) You have two fragile books, one full of text and one blank, both with n pages. You plan to copy the text over, but reordering the pages. What is on page i in the first book should be on page π_i in the second book, for some permutation π. To copy something from page i in the first book and onto page j in the second book they have to be open to those pages (we leave the back of pages blank). They start open on page 1 and you can only turn one page at a time. Due to their fragility you want to turn pages as few times as possible. Show you can copy everything over with at most $\lceil 2n\sqrt{n} \rceil$ page flips.

- (2, Icelandic high school programming contest 2025) Alice and Bob play a variant of Tic-tac-toe. Alice starts with 4 `x` and 2 `X` while Bob starts with 4 `o` and 2 `O`. They take turns placing letters in a 3×3 grid and anyone who gets three of their letter in a row, horizontically, vertically or diagonally, wins. On their turn a player may place a lower case letter in an empty cell, expending it. Alternatively they may place an upper case letter on a cell with a lower case letter, expending their upper case letter and removing the lower case letter. Lastly a player may move one of their upper case letters that is on the board onto a cell with a lower case letter to remove that lower case letter. If a player can not move they lose. Who has a winning strategy if Alice goes first?

- (2, India RMO 2013b) For a natural number n, let $T(n)$ denote the number of ways we can place n objects of weights $1, 2, \ldots, n$ on a balance such that the sum of the weights in each pan is the same. Prove that $T(100) > T(99)$.

- (2, Columbia National Olympiad 2025) The numbers $1, 2, \ldots, 2025$ are written around a circle in some order. On each turn, Celeste chooses an integer $1 \le n \le 2025$. Then she selects the number n and the next $n-1$ numbers following it in the clockwise direction and inverts their order on the circle. Prove that, after a finite amount of turns, Celeste can put the numbers on the circle in the order $1, 2, 3, \ldots, 2025$ in the clockwise direction, regardless of their original arrangement on the circle.
- (2, GCPC 2020) A number n is called adorable if there exists positive integers a, b, c such that an equilateral triangle with side length c that can be tiled with n smaller equilateral triangles, each having side length a or b. What numbers are adorable?
- (2, 102 problems USA IMO) Determine if it is possible to partition the set of positive integers into sets A and B such that A does not contain any 3-element arithmetic sequence and B does not contain any infinite arithmetic sequence.
- (2, ICPC world finals 2022/3) You know axexes, basilisks and centaurs existed, but not how many legs each creature had (only that it had a nonnegative integer number of them). But thankfully you can ask a sphinx five questions, each time asking how many total legs a axexes, b basilisks and c centaurs have for $0 \le a, b, c \le 10$. Sphinxes are tricky creatures however, and the sphinx may lie in one of its answers. How can you make sure you figure out all leg counts in only five questions?
- (2, CodeSprint LA 2021) Alice and Bob are each thinking of a number from 1 to 500. When you guess a number they will both independently tell you whether your guess is correct, too high or too low. Show you can figure out both of their numbers in at most 10 guesses.
- (2^*, NCPC 2012) Which permutatations can be put into sorted order by repeatedly cyclically rotating three adjacent elements? This means if we have elements a, b, c next to each other they can be rearranged as b, c, a or c, a, b.
- (2, Pan-American Girls MO 2021) There are $n \ge 2$ coins numbered from 1 to n. These coins are placed around a circle, not necessarily in order. On each turn, if we are on the coin numbered i, we will jump to the coin i places forward in clockwise order. We start on coin 1. Find all values of n for which there exists an arrangement of the coins in which every coin will be visited.
- (2^*, 2018 ICPC Asia Hanoi Regional Contest) Alice and Bob play a game on N bipartite graphs, the i-th of which has halves with a_i and b_i vertices. Alice and Bob take turns deleting one edge or one vertex from any of the graphs. If a vertex is deleted all of its incident edges are deleted as well. A player who can't move loses, and Bob goes first. In how many ways can the edges of the starting graphs be chosen such that Alice wins?

- ($3^\dagger$, classic) Reformulate the Euclidean algorithm so that instead of replacing x, y by $x - y, y$ you replace it by $x \pmod{y}, y$. For a given upper bound N, find the numbers $x, y \leq N$ such that the Euclidean algorithm will take the most steps to terminate.
- (3, Icelandic final contest 2019) You stand before a 31 floor building with two identical glass spheres. You know that if you drop a sphere from the 31st floor down to the ground it will break. There is some floor n such that if you drop them from floor $\geq n$ the sphere will break, but if you drop it from a floor $< n$ the sphere will sustain no damage. Describe a method for determining n is as few drops as possible. You may drop each sphere as often as you like until it breaks.
- (3, GCPC 2024) For which h, w can an $h \times w$ grid of cells be covered by pentominoes without any two identical pentominoes touching each other orthogonally?
- (3, Icelandic high school programming contest 2025) You are given a set of line segments in the plane representing roads. No intersection point is the intersection of more than two line segments. Each intersection point represents a road intersection, and as such has traffic lights. A traffic light can either show green along one road and red along the other, or yellow along both. Show that the colours of the traffic lights can be chosen such that when traveling along any given road one never sees the same colour twice in a row.
- (3, Putnam 1979) Given n red points and n blue points in the plane, no three collinear, prove that we can draw n segments, each joining a red point to a blue point, such that no segments intersect.
- (3, KTH Challenge 2025) For a name to be valid it has to consist of 4 English lowercase letters. Each name must have two vowels (a, e, i, o, u) and two consonants (the rest). No two names may be too similar either, so no two names may differ in only one position. For example nail and mail are too similar, but seal and sale are not. Show there exists a set of 3150 valid names.
- (3, USAJMO 2019) There are $a + b$ bowls arranged in a row, numbered 1 through $a + b$ where a and b are given positive integers. Initirally, each of the first a bowls contains an apple and each of the last b bowls contains a pear. A legal move consists of moving an apple from bowl i to bowl $i + 1$ and a pear from bowl j to bowl $j - 1$, provided that the difference $i - j$ is even. We permit multiple fruits in the same bowl at the same time. The goal is to end up with the first b bowls containing a pear and the last a bowls each containing an apple. Show that this is possible if and only if the product ab is even.
- (3, IMO 1989 shortlist) A natural number is written in each square of an $m \times n$ chessboard. The allowed move is to add an integer k to each of two

orthogonally adjacent numbers in such a way that nonnegative numbers are obtained. Find a necessary and sufficient condition for it to be possible for all the numbers to be zero after finitely many operations.

- (3, Pan-American Girls MO 2021) Celeste has an unlimited amount of each type of n types of candy, numbered type 1, type 2, ..., type n. Initially she takes $m > 0$ candy pieces and places them in a row on a table. Then she choose one of the following operations (if available) and executes it: a) She eats a candy of type k and in its position in the row she places one candy of type $k-1$ followed by one candy of type $k+1$, indices modulo n. b) She chooses two consecutive candies which are the same type and eats them. Find all positive integers n for which Celese can leave the table empty for any value of m and any configuration of candies on the table.

- (3, 2023 ICPC North America Championship) You want to figure out which one of n compounds you are allergic to. On each of the d days you have to do this you perform a test on your arm, on which you have marked testing sites. Each day you apply each compound to some (possibly empty) subset of your arm. You can apply more than one compound to the same site. Then you wait and see, and for each site see if you had a reaction to at least one of the compounds applied, or not. What is the minimum number of sites you need to guarantee being able to determine the offending compound in the timeframe you have?

- (3, MSOP 1999) Let X be a finite set of positive integers and A a subset of X. Prove that there exists a subset B of X such that A equals the set of elements of X which divide an odd number of elements of B.

- (3, CodeSprint LA 2021) Let s be a string of n bits. You may flip two adjacent bits if they are equal, so replace 00 with 11 or vice versa. What starting strings s can reach the string of n ones?

- (3, NUS Competitive Programming) Alice and Bob play a game based on a string s of n characters, all of which are either A or B. The game itself is played on a $(n+1) \times (n+1)$ grid, a token is initially placed on cell $(0,0)$ of the matrix. On each turn each player moves the token from (i,j) to either $(i+1,j)$ or $(i,j+1)$. The game ends once the token reaches the antidiagonal of the grid, so a cell (i,j) with $(i+j) = n$. If the token ends on cell $(i, N-i)$ and the i-th letter of s is A, then Alice wins. If the i-th letter is B, Bob wins. Among the 2^{n+1} possible strings s, for how many of them does Alice have a winning strategy?

- (3, Icelandic high school programming contest 2025) You have $2n$ batteries, of which you know n are dead and n are good. You need to power a flashlight which takes two batteries. If you put in two good batteries it turns on, otherwise nothing happens. What's the minimum number of attempts of putting batteries into the flashlight you need before it turns on in the worst case?

- (4, IMO 2017) For each integer $a_0 > 1$, define the sequence $a_0, a_1, a_2, \dots,$ by $a_{n+1} = \sqrt{a_n}$ if a_n is an integer and $a_{n+1} = a_n + 3$ otherwise for each $n \geq 0$. Determine all values of a_0 for which there is a number A such that $a_n = A$ for infinitely many values of n.

- (4, NCPC 2020) Let M be a $n \times m$ matrix (grid of numbers), all of which are 1 or 2. For which can you write M as a sum of three $n \times m$ matrices A, B, C such that their entries are all $0, 1$ and the 1s are all orthogonally connected in each matrix?

- (4, USAMO 2017) Let $P_1, P_2, \dots, P_{2n}$ be $2n$ distinct points on the unit circle $x^2 + y^2 = 1$, other than $(1, 0)$. Each point is colored either red or blue, with exactly n red points and n blue points. Let $R_1, R_2, \dots, R_n$ be any ordering of the red points. Let B_1 be the nearest blue point to R_1 traveling counterclockwise around the circle starting from R_1. Then let B_2 be the nearest of the remaining blue points to R_2 traveling counterclockwise around the circle from R_2, and so on, until we have labeled all of the blue points $B_1, \dots, B_n$. Show that the number of counterclockwise arcs of the form $R_i \to B_i$ that contain the point $(1, 0)$ is independent of the way we chose the ordering $R_1, \dots, R_n$ of the red points.

- (4, INMO 2025) Greedy goblin Griphook has a regular 2000-gon, whose every vertex has a single coin. In a move, he chooses a vertex, removes one coin each from the two adjacent vertices, and adds one coin to the chosen vertex, keeping the remaining coin for himself. He can only make such a move if both adjacent vertices have at least one coin. Griphook stops only when he cannot make any more moves. What is the maximum and minimum number of coins that he could have collected?

- ($4^{*\dagger}$, University of Iceland programming contest 2021) Fix some real $a > 0$. For what $b > 0$ does there exist some n and real numbers $x_1, \dots, x_n > 0$ such that $\sum_{i=1}^{n} x_i = a$ and $\prod_{i=1}^{n}(1 + x_i) = b$? Bonus: Find an mathematically exact method to compute these numbers given a, b (so not a numerical approximation algorithm).

- (4, USAMO 1997 Submission) A finite set of (distinct) positive integers is called a DS-set if eaech of the integers divides the sum of them all. Prove that every finite set of positive integers is a subset of some DS-set.

- (4, IMO shortlist 2005) 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.

- (4, SANDVIK Challenge 2020) You have a plate that can fit R rows and C columns of muffins. You have three types of muffins, having a, b, c of each type. They fit perfectly on the plate, so $a+b+c=R\cdot C$. For what values a,b,c can you arrange them on the plate so no two muffins of the same type are orthogonally adjacent?

- (4, IMO 2003) Let A be a 101-element subset of $S=\{1,2,\ldots,10^6\}$. Prove that there exist numbers $t_1,t_2,\ldots,t_{100}$ in S such that the sets $A_j=\{x+t_j|x\in A\}$ are pairwise disjoint for $j=1,2,\ldots,100$. Bonus: Do it for 198 numbers.

- (4, IMO shortlist 1998) 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.

- (4, Icelandic final contest 2026) A hacker is hiding in one of $2N$ servers, each of which is specified by N bits. For example if $N=5$ one server is specified with 01101. To catch him, K policemen connect to K servers. Then the hacker and police alternate taking moves, with the police going first. They may choose to stay in place or move to a server with one bit different, for example one can move from 01101 to 11101. What does K need to be at a minimum to ensure a policeman can be at the same server as the hacker after a finite amount of time? You may assume everyone employs the best strategy possible. Everyone knows where everyone is. All police move at once, but each policeman can choose individually whether they move or not. The police move first.

- (4, Havel-Hakimi) Let G_1,G_2 be two simple graphs with the same vertex set V. For different vertices a,b,c,d replacing the edges $(a,b),(c,d)$ with $(a,c),(b,d)$ in a graph (making sure this does not introduce duplicate edges) is called a two-swap. Show that if the vertices have the same degrees in G_1 and G_2 we can transform G_1 into G_2 by a finite sequence of two-swaps.

- (4, IMO 2007) 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.

- (4, MAPS 2022) Three rival gangs have a,b and c units of resources, all positive integers. If one gang has x units of resources they can mount an attack to steal x units of resources from another gang, and will only do so if that other gang has at least x resources. So if one gang has x resources and another $y \geq x$, after the raid they will have $2x$ and $y - x$ units respectively. Show that if you can manipulate the gangs into doing the raids you ask them to do, you can bring one gang down to zero resources.

- (4, Vizing) Show that a finite simple graph with maximum degree Δ can be edge-coloured using at most $\Delta + 1$ colours.
- (4, NUS Competitive Programming) You have a list a of n integers $a_1, \ldots, a_n$ and an integer $k < n$. The operation cat-split consists of doing the following in order: Choose a subseequence of a with k elements, remove them from a, then prepend the values of the subsequence to the beginning of the list in reverse. For example if a is $[5, 1, 4, 2, 3]$ and k is 3, you could remove $1, 4, 3$. Then the result would be $[3, 4, 1, 5, 2]$. Show that you can sort a into ascending order with a finite number of cat-splits.
- (5, Sylvester-Gallai) For a finite set of points in the Euclidean plane, not all on the same line, show there exists a line containing exactly two of them.
- (5, 102 problems USA IMO) Given an initial sequence $a_1, a_2, \ldots, a_n$ of real numbers we perform a series of steps. At each step, we replace the current sequence $x_1, x_2, \ldots, x_n$ with $|x_1 - a|, |x_2 - a|, \ldots, |x_n - a|$ for some a. For each step, the value of a can be different. a) Prove that it is always possible to obtain the null sequence consisting of all 0s. b) Determine with proof the minimum number of steps required, regardless of initial sequence, to obtain the null sequence.
- (5, USA TST 2018) At a university dinner, there are 2017 mathematicians who each order two distinct entrées, with no two mathematicians ordering the same pair of entrées. The cost of each entrée is equal to the number of mathematicians who ordered it, and the university pays for each mathematician's less expensive entrée (ties broken arbitrarily). Over all possible sets of orders, what is the maximum total amount the university could have paid?
- (5, Atli F & Robert M) You play a game with the numbers $1, 2, \ldots, n$. You want to maximise the sum of the numbers you pick. Whenever you pick a number, the taxman takes all of its proper divisors so they can no longer be picked. Taxes are unavoidable, so if a number has no proper divisors left, it can't be picked. Show that for $n \neq 1, 3$ you can always get more than half the total sum.
- (5, Marshall Hall Jr) You are given a sequence of non-negative numbers $a_1, \ldots, a_n$ with the property that their average is an integer. Prove that you can permute the sequence to obtain a new sequence $b_1, \ldots, b_n$ with the property that all the values $b_i + i \pmod{n}$ are unique for $i = 1, \ldots, n$.
- (5, IMO 2009) 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 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 M.

References:

- Mathematics Olympiad Problem Solving Sessions, https://jpsaha.github.io/
- Olympiad Combinatorics by Pranav A. Sriram
- OTIS Excerpts by Evan Chen
- 102 Combinatorial problems from the training of the USA IMO team, Titu Andreescu and Zuming Feng
- Combinatorics, Stephan Wagner