Prelude

Even far into my math contest career, I had a really bad understanding of the basic principles, even though I could state each of them if you asked me to. Sometimes, when I fail to solve a problem and see this:

> *Let x be minimal. [...]*
>
> *[...]*
>
> *[...]. This violates the assumption that x is minimal, contradiction.*

...I would just tell myself that next time I needed to *remember* to try using the extremal principle. Or induction. Or monovariants.

The thing is, Olympiad Math was never meant to be a checklist of techniques[1] that you had to remember to try. If these ideas aren't *native* in your head, it will be difficult to realize when you should be using certain techniques.

My perspective started to change when I met the 2014/5 at IMO[2]. I was blown away by how much flexibility in thought that the problem demanded from the solver. Starting from then, the previously disparate ideas started to unify into a single framework.

Now, when I see "let x be minimal", I know what it actually means: from x, there is a way to "go to" a smaller "x". Whoever solved the problem did not simply "let x be minimal" and have the rest of the solution fall out of the sky; they saw a potential *process* where you could always produce something smaller, which is impossible.

Another name for a process is an *algorithm*, but strictly speaking this handout isn't about algorithms per se[3]. Rather, it's about the philosophy to consider algorithms *even when the problem does not ask for it.* Privately, I called this *dynamic reasoning*, because we stop thinking about objects as static. Perhaps you must walk around on a graph, or watch a number sequence evolve, or tie coins together. You must imagine yourself *in* the setting of the problem, doing all kinds of things to the objects.

Nonetheless, it was that one specific problem that finally knocked this idea into my head. I hope that at least one such problem in this handout will give you that same inspiration.

[1] Actually, if you feel that a checklist can have enough persuasive power over yourself, this is a perfectly valid strategy. Personally I didn't find checklists convincing.

[2] It was my second IMO, too! Just shows that sometimes you accidentally get very far even with big holes in your knowledge.

[3] which would focus on greedy, BFS, divide-and-conquer, recursion and the like

1 Algorithms for construction

A problem may ask you to show that something exists, or something is possible. Occasionally, the best way to do so is an algorithm, even if the problem itself mentions nothing about an algorithm.

1.1 "Just Do It"

Adapted from MOP 2016 notes.

The philosophy of the greedy algorithm is to build a required set/sequence up one element at a time. Sometimes, you find that it is never difficult to continue building.

Example 1 (folklore)

Show that any graph with degree d is $(d+1)$-colorable (i.e. the vertices can be colored in $(d+1)$ colors such that no two adjacent vertices have the same color).

Solution. Just color in the vertices one by one. This is okay because at each vertex, there can be at most d forbidden colors, so there's always an allowable color.

If you dial this up to eleven, we have the philosophy of "Just [Doing] It". Linus Hamilton writes:

> Do you need a complicated construction? Can't come up with anything that works? Maybe you should stop crying and "Just Do It".

Example 2 (MOP 2016)

Show that there is a sequence of positive integers $a_1, a_2, \dots$ containing every positive integer exactly once, such that first differences $|a_1 - a_2|, |a_2 - a_3|, \dots$ also contain every positive integer exactly once.

Solution. Suppose that you've partially built your sequence. Let $A = \{a_1, \dots, a_n\}$, $B = \{|a_1 - a_2|, \dots, |a_{n-1} - a_n|\}$, and let a, b be the smallest numbers not already in A, B respectively. We will attempt to add a few more terms so that A contains a and B contains b.

For some N really large, add $N, N+b, a$. In particular, N can be chosen large enough for these things not to have any clashes.

1.2 Constructing incrementally

Stupid ideas don't always work. But you can start somewhere and work your way to success. Maybe you've solved a small case, or a special case, or have a suboptimal construction. Then, find a way to make your construction better. (Related words: *smoothing*, *adjusting*, *local*, *perturbation*, *induction*.)

Example 3 (ISL 1998 C2)

A rectangular array of numbers is given. In each row and each column, the sum of all numbers is an integer. Prove that each nonintegral number x in the array can be changed into either $\lceil x \rceil$ or $\lfloor x \rfloor$ so that the row-sums and column-sums remain unchanged.

Solution. We will incrementally modify all numbers to integers. Find a "cycle" of nonintegers (i.e. same row, same col, same row...), then $+\varepsilon$ and $-\varepsilon$ alternately along the cycle. Increase ε gradually, and eventually you will zero some number out. Find another cycle, and repeat. This stops because the number of non-integers strictly decreases after each step.

Example 4 (NewStar Math Problems Column)

There are n vectors on the plane which sum to zero. Show that for any $k \leq n$, there exists a subset of k vectors having a sum at most as long as the longest vector (among the original n vectors).

Solution. By rescaling, assume that the longest vector has length 1. Let the n vectors be $v_1, \ldots, v_n$, so we would like $c_i \in \{0, 1\}$ such that

$$\sum_{i=1}^{n} c_i v_i \leq 1$$
$$\sum_{i=1}^{n} c_i = k$$

The key is to imagine allowing *real* $c_i \in [0, 1]$ instead: a solution to

$$\sum_{i=1}^{n} c_i v_i = 0$$
$$\sum_{i=1}^{n} c_i = k$$

is simply $c_1 = \cdots = c_n = k/n$. Now we try to adjust c_i so that most of them are 0 or 1.

Consider $(v_i, 1)$ as a 3D-vector, then for every 4 indices i, j, k, ℓ, there exists reals a_i, a_j, a_k, a_ℓ not all zero such that

$$a_i(v_i, 1) + a_j(v_j, 1) + a_k(v_k, 1) + a_\ell(v_\ell, 1) = 0$$

Thus,

$$c_i v_i + c_j v_j + c_k v_k + c_\ell v_\ell = (c_i + \epsilon a_i)v_i + (c_j + \epsilon a_j)v_j + (c_k + \epsilon)v_k + (c_\ell + \epsilon)v_\ell$$
$$c_i + c_j + c_k + c_\ell = (c_i + \epsilon a_i) + (c_j + \epsilon a_j) + (c_k + \epsilon a_k) + (c_\ell + \epsilon a_\ell)$$

so we can adjust all but three of them to be 0 or 1. (WLOG, it's c_1, c_2, c_3 remaining.)

Now we show that the remaining adjustment shifts the sum of vectors by at most 1. If $c_1 + c_2 + c_3$ is 0 or 3 then we're trivially done, and if $c_1 + c_2 + c_3 = 2$, we can consider $(1 - c_1) + (1 - c_2) + (1 - c_3) = 1$ instead (since the shift in the vector sum is the same).

Thinking geometrically, if $v_i = \overrightarrow{OP_i}$, then $c_1 v_1 + c_2 v_2 + c_3 v_3$ is a point inside (or on the boundary) of $\triangle P_1 P_2 P_3$, and we want to know the minimal distance to any vertex. This is precisely the circumradius of $\triangle P_1 P_2 P_3$, which is at most 1, so we are done.

1.3 When to stop

For an algorithm to produce an object, it must stop at some point. By designing how you stop the algorithm, you can make your object have additional properties.

1.3.1 Unsustainable quantities

This is when you are forced to stop because a certain monovariant quantity cannot keep increasing/decreasing. Some ideas: (a) being "greedy"; (b) keep making something bigger/smaller; (c) flouting an "exchange rate".

Example 5

A 5×5 grid of light bulbs contains a light switch for each row and each column. The light switch toggles every light in the corresponding row or a column. Prove that given any initial state of the light bulbs, you can hit a finite sequence of light switches such that in each row and each column, there are more lights on than lights off.

Solution. If in some row (or column) there are more lights off than on, you can toggle the entire row (or column) to increase the number of light bulbs. So keep doing this; because you can't increase the number of light bulbs forever, you stop, and the state has the desired properties.

Example 6 (Mediterrenean 2015)

In a mathematical contest, some of the competitors are friends and friendship is mutual. Prove that there is a subset M of the competitors such that each element of M has at most three friends in M and such that each competitor who is not in M, has at least four friends in M.

Solution. It's reasonable to attempt a greedy algorithm, but it's quickly evident that it doesn't work: adding a new person with less than four friends in M could cause some existing member of M to have more than 4 friends in M.

We will attempt to find some "extremal" assumptions that automatically guarantee the properties of M. We go back to the greedy algorithm earlier. In an ideal world, we would be able to remove the existing member with more than 4 friends. However, the worry here is that the algorithm now is not guaranteed to terminate, since we could add a person and remove a person and continue in some kind of cycle. This is not a problem because there is another monovariance at work. But consider this perspective: at each step we add at most 3 friendships each step or take away at least 4 friendships. If we have an ideal picture of exchanging the number of elements of M for the number of friendships, we can see that in some sense we "losing" something gradually and eventually have to stop. This can be made concrete by considering a quantity like $|E| - (7/2)|V|$, where V, E refer to the vertices and edges within M. Since this quantity always decreases, the algorithm terminates.

1.3.2 Cyclic

Stop when you loop. This is good for finding loops. For this to work, you should have finitely many states.

Example 7 (ISL 2007 C2)

A rectangle D is partitioned in several (≥ 2) rectangles with sides parallel to those of D. Given that any line parallel to one of the sides of D, and having common points with the interior of D, also has common interior points with the interior of at least one rectangle of the partition; prove that there is at least one rectangle of the partition having no common points with D's boundary.

Solution. The official solution in the IMO shortlist booklet is basically case bash, which is rather unfortunate. Here's a different idea: imagine a particle moving along the boundaries of the diagram (consisting of all sides of sub-rectangles). One way to isolate a rectangle is to just keep turning anticlockwise at the earliest opportunity. This doesn't really help because we need an "interior" rectangle... so instead we will turn whichever way that doesn't lead out into the boundary (which makes perfect use of the condition that no line cuts straight through all the rectangles). Eventually, because the number of vertices are finite, we will bound some non-empty region with the traced path, and clearly that contains a rectangle.

1.3.3 I have no idea when to stop

... but maybe you know that along the way you must stop somewhere. Sometimes, this looks like the discrete Intermediate Value Theorem. (This has some overlap with *incremental constructions.*)

Example 8 (2D Ham-Sandwich)

There are $4n+2$ points on the plane, no three of which are collinear. Exactly half of them are colored red, and the other half is colored blue.

Show that we can draw a line ℓ connecting a red and a blue point such that there are n red points and n blue points on either side of the line ℓ.

Solution. Start with ℓ passing through some blue point such that there are n blue points on either side. (Either choose the point first, then rotate the line, or choose the line first then the point.) Now rotate ℓ while keeping this property. At some point, the number of red points on either side are equal!

Example 9 (Russia 4th Round 2004)

Obligatory context about a fruit seller (otherwise the spacing looks weird).

(a) There are 99 boxes, each with some apples and oranges. Show that you can pick 50 boxes that contains (at least) half of all the apples and half of all the oranges.

(b) There are 100 boxes, each with some apples, oranges and bananas. Show that you can pick 51 boxes that contains (at least) half of all the apples, half of all the oranges and half of all the bananas.

Solution.

1.4 Real Induction

As Jacob Tsimermann puts it:

Induction is awesome and should be used to its full potential!

What does this mean?

There's the usual vanilla *induction*: you reduce the case of n to the case of $n-1$.

There's *strong induction*: you reduce the case of n to the case of anything smaller than n.

And then there's *real induction*: you weaponize every possible case covered by induction and bombard the fuck out of the case of n.

Here's an example of what I mean:

Example 10 (Canada 2010/4)

There are n lamps in a room, with certain lamps connected by wires. Initially all lamps are off. You can press the on/off button on any lamp A, but this also switches the state of all the lamps connected to lamp A from on to off and vice versa. Prove that by pressing enough buttons you can make all the lamps on. (Connections are two-way.)

Solution. We use induction on n. The base case of $n=1$ is trivial (just turn the lamp on). Now assume the case of $n-1$ lamps. Now look at the set of n lamps and ignore some lamp A. Then by induction, we can turn the remaining lamps on by pressing the buttons on some subset of them. Now if at the end of doing this A is also on, we are done. So we can assume that at the end of doing this A is off.

Since A was an arbitrary lamp, we can assume that by pressing a sequence of buttons we can flip the states of all lamps except one of our choosing. Now, taking A and B to be 2 different lamps and flipping the states first of all lamps different from A and then all lamps different from B, we see that we can flip the states of only A and B. So this means we can flip the states of any number of even lamps. Now we have 2 cases:

- n is even: Then we are already done, since we can flip the states of any number of even lamps.
- n is odd: In this case, there must be some lamp A connected to an even number of lamps (prove this!) so first press the button on lamp A. Now, including A, an odd number of lamps are on, so an even number of lamps that are off remain. Flip their states to finish the proof!

Note that instead of using the induction hypothesis once to finish off the solution, we used it once per vertex in order to get a huge amount of infromation (flipping the state of any even set of lamps, that's clearly huge) so that finishing off the proof was much more straightforward.

1.5 Losing nothing

Adapted from Evan's blog, and very related to the previous idea.

Evan Chen writes:

> *Usually, a good thing to do whenever you can is to make "safe moves" which are implied by the property P. [...] But often times such "safe moves" or not enough to solve the problem, and you have to eventually make "leap-of-faith moves".*
>
> *However, a strange type of circular reasoning can sometimes happen, in which a move that would otherwise be a leap-of-faith is actually known to be safe because you also know that the problem statement you are trying to prove is true.*

I call this a "free reduction" - as opposed to an assumption which incurs a cost of being potentially wrong, sometimes assumptions are guaranteed to be correct because they are just another case of the problem. This can also be thought of smoothing towards another case of the problem.

(Related terms: *induction*, *recursion*, *WLOG* but in an important way.)

Example 11

Prove that the area of a parallelogram inside a triangle is at most half the area of the triangle.

Solution. Perform the following steps:

1. WLOG that one vertex is along the side of the triangle (otherwise just extend along one side of the parallelogram)
2. WLOG that one side is along the side of the triangle (since moving the other side along the same direction doesn't change the area). Call this side a.
3. WLOG the other two vertices are also on the sides of the triangle (just expand along the direction of a)
4. WLOG one of the vertices is a vertex of the triangle (move a to one of the ends of the triangle side)
5. WLOG it's a right-angled triangle.
6. Now it's really easy!

This idea is even useful in some "real math" scenarios:

Example 12 (Cauchy-Davenport)

Given nonempty subsets $A, B \pmod p$, show that

$$|A + B| \geq \min\{|A| + |B| - 1, p\}$$

where $A + B = \{a + b \pmod p \mid a \in A, b \in B\}$.

Solution. (due to Terence Tao) The question is trivial if $|A| + |B| \geq p + 1$ (think about the sets A and $x - B$ for any element x).

Otherwise, we wish to show $|A+B| \geq |A|+|B|-1$. Note that if $A \cap B$ is non-empty then

$$(A \cap B)+(A \cup B)=A+B$$

$$|A \cap B|+|A \cup B|=|A|+|B|$$

so we can really just reduce to the case where $A \subseteq B$.

But then there's also the "translation" invariance: nothing really changes when shift all the elements in one of the sets by the same amount, so we can consider instead $A+x, B$. The x for which $A+x, B$ is not empty is really just $B-A$. Then for any such x, we can perform the translation, which has the effect of shrinking A.

If we can't shrink A any more, then for all $x \in B-A$, $A+x \subseteq B$. So $A+B-A \subseteq B$. However, this forces $|A|=1$, since otherwise a nonzero element of $A-A$ would imply that all elements are in B, and then we are clearly done. If $|A|=1$, we are also clearly done.

1.6 Finale

Here is an IMO problem that uses a combination of a few techniques:

Example 13 (IMO 2014/5)

For each positive integer n, the Bank of Cape Town issues coins of denomination $\frac{1}{n}$. Given a finite collection of such coins (of not necessarily different denominations) with total value at most $99+\frac{1}{2}$, prove that it is possible to split this collection into 100 or fewer groups, such that each group has total value at most 1.

Solution. If we see two coins of value $1/2k$, we can tie them together and pretend it was a one coin of value $1/k$. Thus, WLOG there is only one coin of value $1/2k$ for each k.

If we see k coins of value $1/k$, we can tie them together and pretend it was one coin of value 1. Thus, WLOG there is at most $(k-1)$ coins of value $1/k$ for each $k>1$.

If we see a coin of value 1, we can just put it into its own group. This suggests (but not wlog) generalizing the hypothesis to being able to put coins of total value $n-1/2$ into n groups.

Now, imagine picking up the coins one by one, starting from the heaviest. After picking up each coin, check if you can still do the splitting. If the newest coin of weight $1/m$ makes the splitting impossible, then any splitting just before picking up this coin has each group having value in $(1-1/m, 1]$, so $n(1-1/m)+m \leq n-1/2$, or $m<2n-2$.

So we may assume WLOG that there is at most one coin of weight $1/(2k)$ and $(2k-2)$ coins of weight $1/(2k-1)$. But the total value now is at most

$$\sum_{k=1}^{n-1}\left(\frac{1}{2k}+\frac{2k-2}{2k-1}\right) \leq 1/2+(n-2) \cdot 1=n-3/2$$

so this can in fact fit into $(n-1)$ groups, by a suitable induction on n.

1.7 Fun/misc examples

Sometimes these idea shows up in non-combinatorial problems.

Example 14

Show that for any positive $a_1, \ldots, a_n$,

$$\sum_{i=1}^{n} \frac{a_i}{a_{i+1}+a_{i+2}}>\frac{n}{4}$$

where the indices are taken mod n.

Solution. Consider $a_i \to \max\{a_{i+1}, a_{i+2}\}$, starting from the maximal a_i. It must loop back. This chain contains at least half of all the terms, and we can just AMGM to show that it's at least $\frac{n}{4}$.

Also, the laser technique:

Example 15 (MEMO 2015/T-4)

Let N be a positive integer. In each of the N^2 unit squares of an $N \times N$ board, one of the two diagonals is drawn. Show that the drawn diagonals divide the $N \times N$ board into at least $2N$ regions.

Solution. Treat the diagonals as mirrors. There are two kinds of regions: ones that touch the edge and have a horizontal or vertical side, and ones that don't. For the ones that do touch the edge, if we shine a laser straight through that edge and let it bounce off the diagonals, it always eventually exits out of another edge after traversing the entire region. Therefore, exactly two unit segments of the board's boundary are part of each edge-touching region. This means the number of edge-touching regions is always $2N$.

This last problem is self-explanatory:

Example 16

There are 4 grasshoppers on a unit square. A grasshopper at point A may jump over another grasshopper at point B to point C - the reflection of A about B. Show that the grasshoppers cannot form a larger square.

Solution. Suppose it does form a larger square of side length $d > 1$. By reversing this process, the grasshopper can form a square of length $1/d^n$, but the grasshoppers always remain on a unit grid, where the minimal area of a non-degenerate triangle is $1/2$.

1.8 Anti-problems

Some failure modes of the algorithmic approach:

- There is a secret (global) invariance being preserved.
- There are end-states of being "stuck" that must be avoided.
- The existence is actually non-constructive (think: pigeonhole, double counting etc.)

Example 17 (USA TSTST 2019 1/3)

On an infinite square grid we place finitely many cars, which each occupy a single cell and face in one of the four cardinal directions. Cars may never occupy the same cell. It is given that the cell immediately in front of each car is empty, and moreover no two cars face towards each other (no right-facing car is to the left of a left-facing car within a row, etc.). In a move, one chooses a car and shifts it one cell forward to a vacant cell. Prove that there exists an infinite sequence of valid moves using each car infinitely many times.

Solution. Color the rows alternately white and black. Move all the vertically faced cars in black rows into white rows. Now clear all the horizontal cars in black rows by moving them very far to the right or left. Now, the black rows are empty.

Move the vertically faced cars into black rows. Now clear the horizontal cars in white rows by moving them very far out. At this point, we have essentially moved all horizontal cars very far out, which means that we can freely clear the vertically faced cars, and we're done.

2 Problems

1. A magical castle has n identical rooms, each of which contains k doors arranged in a line. In room $i, 1 \leq i \leq n-1$ there is one door that will take you to room $i+1$, and in room n there is one door that takes you out of the castle. All other doors take you back to room 1. When you go through a door and enter a room, you are unable to tell what room you are entering and you are unable to see which doors you have gone through before. You begin by standing in room 1 and know the values of n and k.

   Determine for which values of n and k there exists a strategy that is guaranteed to get you out of the castle and explain the strategy. For such values of n and k, exhibit such a strategy and prove that it will work.

2. Given $a > 1$, show that there exists a bounded sequence of reals $x_1, x_2, ...$ such that

$$|x_i - x_j| \geq \frac{1}{|i-j|^a}$$

3. Does there exist a triangle-free graph with chromatic number at least 2020?

4. A table with 5 rows and 5 columns is filled with nonnegative integers. With an add move we can select one of the 5 rows or one of the 5 columns and add 1 to each number in that row or column. With a subtract move we can select one of the 5 rows or one of the 5 columns and subtract 1 from each number in that row or column.

   Suppose that we can find a sequence of add moves and subtract moves that will reduce our table of nonnegative integers to a table filled with zeros only. Prove that in that case the original table of nonnegative integers can also be reduced to all zeros using only subtract moves.

5. Given a graph G with vertex set V, there exists an independent set (i.e. no two vertices are adjacent) of size at least

$$\sum_{v \in V} \frac{1}{\deg v + 1}.$$

6. Given a graph with spanning trees with k and m leaves (terminal vertices, of degree 1 within the tree) respectively, show for all $k < l < m$ there exists a spanning tree with l leaves.

7. Let S be a finite set of positive integers. Suppose for any positive integer n, (n, s) is either 1 or s. Show that there exists $s, t \in S$ such that (s, t) is prime.

3 Meme Problems

1. Rothschild the benefactor has a certain number of coins. A man comes, and Rothschild wants to share his coins with him. If he has an even number of coins, he gives half of them to the man and goes away. If he has an odd number of coins, he donates one coin to charity so he can have an even number of coins, but meanwhile another man comes. So now he has to share his coins with two other people. If it is possible to do so evenly, he does so and goes away. Otherwise, he again donates a few coins to charity (no more than 3). Meanwhile, yet another man comes. This goes on until Rothschild is able to divide his coins evenly or until he runs out of money.

   Does there exist a natural number N such that if Rothschild has at least N coins in the beginning, he will end with at least one coin?

2. Do there exist positive integers $a_1, \dots, a_{1000}$ such that $d(a_1 + \cdots + a_k) = a_k$ for all $1 \leq k \leq 1000$?

3. A rectangle $\mathcal{R}$ with odd integer side lengths is divided into small rectangles with integer side lengths. Prove that there is at least one among the small rectangles whose distances from the four sides of $\mathcal{R}$ are either all odd or all even.

4. Show that a convex $2n$-gon whose vertices have integer coordinates has area at least $\binom{n}{2}$. (Bonus: does it have area at least cn^3 for some constant $c > 0$?)

5. Suppose the function $f : \mathbb{N} \to \mathbb{N}$ satisfies $f^n(n) = 2n$ for all natural numbers n. Must $f(n) = n + 1$?

6. Draw a $2^n \times n$ array (n columns), and along the rows write down every possible string of -1 and 1. Replace some cells with 0. Show that there exist some rows that add up to be 0 along every column.

4 All examples

1. Show that any graph with degree d is $(d+1)$-colorable (i.e. the vertices can be colored in $(d+1)$ colors such that no two adjacent vertices have the same color).

2. Show that there is a sequence of positive integers $a_1, a_2, \ldots$ containing every positive integer exactly once, such that first differences $|a_1 - a_2|, |a_2 - a_3|, \ldots$ also contain every positive integer exactly once.

3. A rectangular array of numbers is given. In each row and each column, the sum of all numbers is an integer. Prove that each nonintegral number x in the array can be changed into either $\lceil x \rceil$ or $\lfloor x \rfloor$ so that the row-sums and column-sums remain unchanged.

4. There are n vectors on the plane which sum to zero. Show that for any $k \le n$, there exists a subset of k vectors having a sum at most as long as the longest vector (among the original n vectors).

5. A 5×5 grid of light bulbs contains a light switch for each row and each column. The light switch toggles every light in the corresponding row or a column. Prove that given any initial state of the light bulbs, you can hit a finite sequence of light switches such that in each row and each column, there are more lights on than lights off.

6. In a mathematical contest, some of the competitors are friends and friendship is mutual. Prove that there is a subset M of the competitors such that each element of M has at most three friends in M and such that each competitor who is not in M, has at least four friends in M.

7. A rectangle D is partitioned in several (≥ 2) rectangles with sides parallel to those of D. Given that any line parallel to one of the sides of D, and having common points with the interior of D, also has common interior points with the interior of at least one rectangle of the partition; prove that there is at least one rectangle of the partition having no common points with D's boundary.

8. There are $4n + 2$ points on the plane, no three of which are collinear. Exactly half of them are colored red, and the other half is colored blue.

   Show that we can draw a line ℓ connecting a red and a blue point such that there are n red points and n blue points on either side of the line ℓ.

9. Obligatory context about a fruit seller (otherwise the spacing looks weird).

   (a) There are 99 boxes, each with some apples and oranges. Show that you can pick 50 boxes that contains (at least) half of all the apples and half of all the oranges.

   (b) There are 100 boxes, each with some apples, oranges and bananas. Show that you can pick 51 boxes that contains (at least) half of all the apples, half of all the oranges and half of all the bananas.

10. There are n lamps in a room, with certain lamps connected by wires. Initially all lamps are off. You can press the on/off button on any lamp A, but this also switches the state of all the lamps connected to lamp A from on to off and vice versa. Prove that by pressing enough buttons you can make all the lamps on. (Connections are two-way.)

11. Prove that the area of a parallelogram inside a triangle is at most half the area of the triangle.

12. Given nonempty subsets $A, B \pmod p$, show that

$$|A + B| \ge \min\{|A| + |B| - 1, p\}$$

    where $A + B = \{a + b \pmod p \mid a \in A, b \in B\}$.

13. For each positive integer n, the Bank of Cape Town issues coins of denomination $\frac{1}{n}$. Given a finite collection of such coins (of not necessarily different denominations) with total value at most $99 + \frac{1}{2}$, prove that it is possible to split this collection into 100 or fewer groups, such that each group has total value at most 1.

14. Show that for any positive $a_1, \ldots, a_n$,

$$\sum_{i=1}^{n} \frac{a_i}{a_{i+1} + a_{i+2}} > \frac{n}{4}$$

where the indices are taken mod n.

15. Let N be a positive integer. In each of the N^2 unit squares of an $N \times N$ board, one of the two diagonals is drawn. Show that the drawn diagonals divide the $N \times N$ board into at least $2N$ regions.

16. There are 4 grasshoppers on a unit square. A grasshopper at point A may jump over another grasshopper at point B to point C - the reflection of A about B. Show that the grasshoppers cannot form a larger square.

17. On an infinite square grid we place finitely many cars, which each occupy a single cell and face in one of the four cardinal directions. Cars may never occupy the same cell. It is given that the cell immediately in front of each car is empty, and moreover no two cars face towards each other (no right-facing car is to the left of a left-facing car within a row, etc.). In a move, one chooses a car and shifts it one cell forward to a vacant cell. Prove that there exists an infinite sequence of valid moves using each car infinitely many times.

5 Hints

Examples.

2. Just do it.

Problems.

1. Just do it! Start by enumerating every possible door configuration.
2. Pick $x_1, x_2, \ldots$ sequentially. At each step, what is the total length of all the forbidden intervals?
3. A bit of incremental construction mixed with "just do it". How might you (abusively) force a single vertex to be of a new color, given many copies of a graph G?
5. If you wanted to pick the first vertex of an independent set, which vertex should you pick?
6. Given spanning trees T_1, T_2 with k, m leaves respectively, how might you transform T_1 into T_2 incrementally?
7. Try $|S| = 1$. It is obvious that if $(n, s) = 1$ or s always, then s itself has to be prime! But exactly which n tells us that?

Meme Problems.

1. Try to describe how the number of coins changes. Can you engineer a construction?
2. Consider $s_n = a_1 + a_2 + \ldots a_n$.
4. Can you assume that the convex polygon is centrally symmetric?
5. The answer is no. Imagine filling all natural numbers into chains of f. We can avoid clashes by using fast growing sequences, and we can fill in gaps by creating new chains.
6. Given $v \in \{0, 1\}^n$, there is always some row you can add to v such that it stays in $\{0, 1\}^n$. What happens when this cycles?