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 $0 < a_{j+1} - a_j \leq m$ $(1 \leq j \leq k-1)$.

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 $s \leq r$, $A_s \cup A_{s+1} \cup \ldots A_r$ 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 ϵ A, y ϵ B, z ϵ 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 ∈ 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 ϵ 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.

$$\begin{array}{cccccccc} \boldsymbol{X} & \boldsymbol{c_i} & \boldsymbol{c_i} & \boldsymbol{c_i} & \boldsymbol{c_i} & \boldsymbol{c_i} & \boldsymbol{c_i} & \boldsymbol{c_i} & \boldsymbol{Y} \\ & \boldsymbol{c_i} & c_5 & \boldsymbol{c_4} & \boldsymbol{c_4} & \boldsymbol{c_4} & c_3 & \boldsymbol{c_i} & \end{array}$$

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 S satisfying the required condition, and then **use the assumption** that S is as large as possible to prove that $|S|$ 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 X be an n-element set and let $A_1, A_2, \ldots, A_m$ be subsets of 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 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, …, 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	**2**	3	1	**2**	3	1
1	2	3	1	2	3	1	2
2	3	1	2	3	1	2	3
3	1	**2**	3	1	**2**	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:

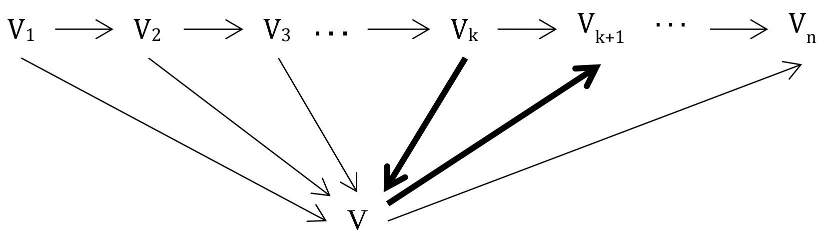

Figure 4.2: Illustration of how to extend the path to include V

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 $S_{j+1} \cap T_j = \emptyset$. ■

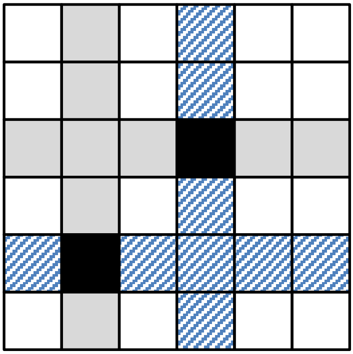

Figure 4.3: The black squares illustrate a contradiction as they cannot simultaneously belong to T_j and S_{j+1}.

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?

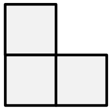

Figure 4.4: An L-tromino

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 $f(n) = b_n - (n\text{-}1) \leq 0$, 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. ■

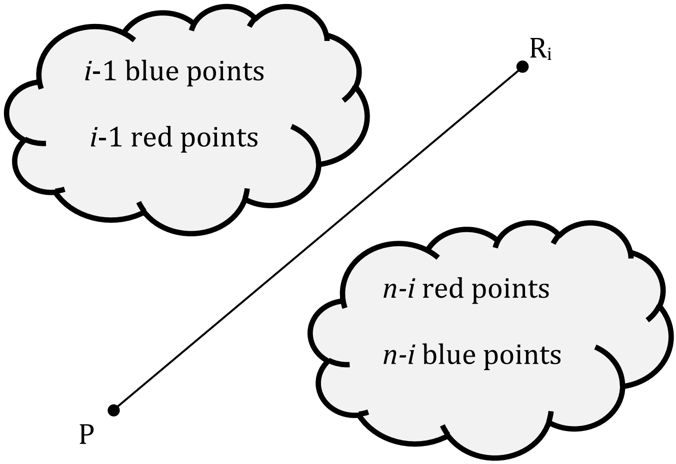

Figure 4.5: A balancing line for f(i) = 0

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 $x_0, x_1, x_2, \ldots, x_n$ 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 $x_1 + x_2 \ldots + x_n$ counts the total number of nonzero terms in the sequence; hence $x_1 + \ldots + x_n = m$. One of the nonzero terms is x_0, so there are exactly m-1 nonzero terms among $x_1, x_2, \ldots, x_n$. 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 $x_3, x_4, \ldots$ 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 $x_1, x_2, \ldots x_n$. 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 $1 \leq k \leq n$, 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?

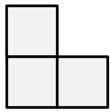

An L-tromino

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 $0, 1, \ldots, n\text{-}1$ 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.

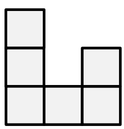

A hook