Preface

This book is about taking familiar ideas (and perhaps some not-so-familiar ones!) and extending them to solve a broad variety of problems. The intended audience is the ambitious high school or college student who is seeking a substantive "big picture" view of several topics that often present themselves at the Olympiad level of mathematics competitions or in proofs and applications of classical results. This "big picture" view of a topic suggests the metaphor of a *landscape*. Landscapes are rich in detail and often contain subtleties that are embedded for the discerning eye to discover.

And so are the chapters and problems that follow. As we shall soon see, sometimes the simplest of ideas can hold very powerful abstractions, or can be used to solve a wide range of problems. It is our hope that this book will offer the reader an ample variety of both breadth and depth in its subject matter. The topics covered span the broad subject areas of algebra, geometry, number theory, and even a few elements of mathematical analysis while each chapter explores specific themes and ideas that illustrate the aforementioned subject areas.

Each chapter is composed of three parts: the theoretical discussion, proposed problems, and solutions to the proposed problems. In each chapter, the theoretical discussion sets the stage for a different landscape by introducing and motivating the themes of that chapter — often with a review of some definitions or classical results. The remainder of the theoretical part of each chapter is devoted to examining illustrative examples — that is, several problems are presented, each followed by at least one solution. It is assumed that the reader is intimately familiar with topics covered in the standard high school mathematics curriculum up through and including precalculus. For some chapters, it will be helpful for the reader to have had previous exposure to topics from discrete mathematics and number theory.

As you experience these "landscapes" in each chapter, the reader is encouraged to carefully seek out the finer points and subtleties. To summarize: *don't rush.* These landscapes provide a "view" into areas that are not typically encountered in great depth in standard coursework, but nonetheless have profound implications. It is also our hope that even if the reader happens not to be well-versed in proofs, he or she will appreciate and grasp the level of *rigor* and thinking that is required in formulating logical and convincing arguments. We hope the examples and the suggested solutions to the problems lend insight not only to the underlying mathematics, but also serve as examples of good mathematical exposition.

We want to express our thanks to our friends Gabriel Dospinescu and Chris Jeuell for their great help throughout our work to the manuscript; they surely had a big contribution to improving it and giving it a better final form.

We invite you to enjoy and appreciate our selection of these beautiful landscapes that abound in mathematics.

Titu Andreescu, Cristinel Mortici, Marian Tetiva

Contents

Chapter 1

The Ubiquitous Pigeonhole Principle

Oh, no, not again! — we hear you say, and you might be right — or you might not. The pigeonhole principle appears in numerous books and magazines and places on the Internet, and everyone is familiar with it, and everybody knows what is "all" about this principle with a bizarre name that.... Who needs another lesson on it? Well, this point of view may be valid — or it may be not. Just read and smile (since we start with the simplest examples), then enjoy (because we will move on to harder problems). Constantin Noica, a great Romanian philosopher, used to say that you can begin to learn philosophy from anywhere — start with any book you want, with any philosopher that you please. Apparently this doesn't work very well for mathematical training, but we say that a good point to start in mathematics is the *pigeonhole principle*. As humble as it may seem, the pigeonhole principle is an indispensable tool throughout all of mathematics. As said in the chapter title, you can find it anywhere in this beautiful realm of human knowledge to which we (and our readers) dedicated our lives, thus we decided to offer the opening chapter of our book to this marvelous (by its simplicity) idea. Last, but not least, remember that another name for the pigeonhole principle is *Dirichlet's drawer principle* (it seems that he was the first to use it explicitly); if it is related to such a great mathematician's name as Dirichlet, it must be worth studying, right?

We are sure that the reader is acquainted with the spectacularly simple statement of the pigeonhole principle (or, if he/she is not, let him/her be from now on). It says that if you try to put in a few boxes a few objects, the number of which exceeds the number of boxes, then, necessarily, there is (at least) one box that contains more than one object. (If you put five pigeons in four pigeonholes, two or more pigeons need to jostle each other into the same pigeonhole.) To put it more formally, we have:

The pigeonhole principle/Dirichlet's drawer principle. If we want to distribute $m > n$ objects into n boxes, then there will exist (at least) one box that contains (at least) two objects.

If we want to push this formalism further, we can state it as follows:

The pigeonhole principle/Dirichlet's drawer principle. If A and B are finite sets with $|A| > |B|$ (the number of elements of A is greater than the number of elements of B), then there is no injective function from A to B.

(Consider A as the set of pigeons and consider B as the set of pigeonholes; a function from A to B puts pigeons in the pigeonholes, and there must be more than one pigeon in some pigeonhole, thus the function cannot be injective.) Actually, this form won't be necessary for our tasks, but it might be interesting for further developments.

For instance, if you have a pile of red and blue balls, you have to take (at least) three of them in order to be sure that you get two balls of the same color. In this example the pigeonholes are the colors (red and blue) and their number is, of course, 2 — so you need have three or more pigeons (balls) to be sure that at least two of them fall into the same box (that is, they have the same color). Or, because everyone of us has on his or her head at most (say) a million hairs, you can gather together arbitrarily 1000001 people and bet them that there are two of them with the same number of hairs on their heads. Bet on one dollar to each of them that your guess is true and you will surely win, won't you? And you will be a millionaire only by properly applying the pigeonhole principle. (Of course, it will be hard to count all the hairs in their heads, so you can think of something more practical, like: gather together q people and bet with them that there exist among them two with the same birthday. What value of q should be chosen to assure winning? We are sure you already know. And don't be sad about losing a million: there will be less money, but this win is accomplishable and it is easy money, isn't it?) Let us move on now to "serious" problems.

Problem. From any n people we can always choose two with the same number of acquaintances among the others $n - 1$.

Solution. If $n = 1$, this is vacuously true, while if $n = 2$, the two persons can be acquainted to each other, or not: in both cases they have the same number of acquaintances among the other(s). (Of course we consider "acquaintanceship" to be a symmetric relation, as will it be in all such problems: A knows B if and only if B knows A. Also, we don't consider that A knows A — the relation is not reflexive.) If there are three people each of whom knows each other (or each of whom does not know each other) the conclusion is plain. Two cases remain: when two persons (say A and B) know the third (name it C) but don't know each other, and when A and B are acquainted, but neither of them is acquainted to C. In both cases A and B have the same number of acquaintances among the others.

However, this is clearly *not* a path to follow for solving this problem: things get more and more complicated when n increases. On the other hand, the pigeonhole principle solves it (almost) immediately. Namely, each of the n persons can have at most $n-1$ acquaintances among the others. There are n persons (pigeons) and only $n-1$ possibilities for the number of the acquaintances (pigeonholes), so there must be two persons with the same number of acquaintances. Done... or not?

Of course, we forgot that the number of persons that someone knows among the others could also be 0: we didn't say that there isn't a person knowing none of the others, so the above solution works *only* in the case when everybody knows at least one person in the group. However, this flaw is not hard to fix. If there are *at least two* such persons, the problem is clearly solved. Otherwise, each and every one of the other $n-1$ has a number of acquaintances from 1 to $n-2$ (that is, at most $n-1$ known people), and the above reasoning shows that there are two of them who know the same number of people among the others.

Solution. After analyzing the above facts, one can come with the following alternative (and much more simple) wording for solving the problem. Choose from the n people those k who know nobody in the entire group. Each of the $n-k$ remaining persons has at least one and at most $n-k-1$ acquaintances among the others, and the pigeonhole principle ensures us that there are two of them with the same number of acquaintances in the group. Simple, isn't? Actually, we see that the first case (when everybody knows at least somebody else) suffices to solve the problem anyway — the problem that can also be rephrased in terms of graph theory: in a (simple, undirected) graph on n vertices, there are always two vertices with the same degree (that is, two vertices from which the same number of edges come out).

Another well-known problem about people who know (or don't know) people is the following.

Problem. Prove that in any group of six people there are three who know each other, or there are three who don't know each other (mutually).

Solution. It is easier to speak if we use the language of graphs. Namely, let the six persons be represented by six points in the plane (the vertices of the graph) and draw a red (blue) line between two points if the persons represented by that line know each other (respectively if they do not know each other). We get the complete graph on 6 vertices with all edges (How many? It is not important for this problem, but think about it!) colored red or blue, and we want to prove that (for any such coloring) there exists either a red triangle, or a blue one. Of course, by a "red/blue triangle" we mean a triangle with all its edges colored red/blue. In order to solve this problem, we need the following more general, often used, form of the principle:

General pigeonhole principle. If more than nk objects are to be placed in n boxes, then there exists (at least) one box containing (at least) $k+1$ objects.

(For if each of the n boxes contained at most k objects, there would be at most nk objects, which is not the case. For $k=1$ we get the principle as stated above.) Now let us go back to our problem. Look at one (arbitrarily chosen) vertex A and the five edges incident with it. Because $5 = 2 \cdot 2 + 1$ and there are only two colors (the boxes, or the pigeonholes) there must be three edges from that vertex that have all the same color. Name those edges AB, AC, and AD, and suppose (without loss of generality) that their common color is red. We are now almost done. Indeed, if one of the edges BC, BD, and CD is red, we have a red a triangle: ABC, or ABD, or ACD. Otherwise BC, BD, and CD are all blue, and the triangle BCD is blue.

Note also that the general form of the principle can be extended to the following one, referring to infinitely many objects:

The pigeonhole principle — the infinite case. Consider infinitely many objects to be placed into n (thus finitely many) boxes. Then there exists a box that contains infinitely many objects.

We use this version to prove the following theorem of Schur.

Theorem. *Let the set of positive integers be partitioned into finitely many classes. Then there exist three numbers x, y, and z in the same class of the partition such that $x+y=z$ (x and y need not necessarily be distinct).*

Proof. Let us consider that $\mathbb{N}^*$ is partitioned into three classes A, B, and C, and let $a_1 < a_2 < \cdots$ be the elements of A (of course, at least one of A, B, and C must be infinite, and we can assume, without loss of generality, that A is infinite). If any of the positive integers $a_2 - a_1, a_3 - a_1, \ldots$ belongs to A, we are done (if, say, $a_i - a_1 = a_j$, we also have $a_1 + a_j = a_i$, hence a solution for $x+y=z$ with all x, y, and z belonging to A). Otherwise, since any positive integer belongs to one (and only) of A, B, and C, one of the sets B and C contains infinitely many of the above differences (by the infinite case of the pigeonhole principle). Suppose (again without loss of generality) that the numbers $b_1 = a_{i_1} - a_1$, $b_2 = a_{i_2} - a_1, \ldots$ belong to B. If any of the differences $b_2 - b_1,\ b_3 - b_1, \ldots$ belongs to B, we clearly finished our proof (we get $x, y, z \in B$ such that $x+y=z$). Yet, if any of these numbers belongs to A, we have also completed our proof; indeed, if $b_m - b_1 = a_{i_m} - a_{i_1}$ belongs to A, this yields a solution to $x+y=z$ with all components in A. Now if $c_1 = b_2 - b_1$, $c_2 = b_3 - b_1, \ldots$ all belong to C, the differences $c_n - c_1 = b_{n+1} - b_2 = a_{i_{n+1}} - a_{i_2}$ have to be in one (and only one) of the sets A, B, and C, but (due to the threefold aspect of these differences) in any case a solution to $x+y=z$ appears in one of the three classes of the partition.

The reader is invited to write the proof in the general case, with an arbitrary (but finite) number of classes. Note that, actually, this is just a weak form of Schur's theorem, whose statement is as follows:

Theorem. *For every positive integer k, there exists a positive integer $S = S(k)$ such that, for any $N \geq S$, and for every k-class partition of $\{1, 2, \ldots, N\}$, one of the classes of the partition contains three numbers x, y and z such that $x + y = z$.*

The numbers $S(k)$ are named Schur's numbers, and their only known values are $S(1) = 2$, $S(2) = 5$, $S(3) = 14$, and $S(4) = 45$. The reader can try to prove that these values are correct, however we will move to other topics, since these problems are beyond the scope of this book.

Everybody knows that from any n integers, one can choose a few (at least one) whose sum is divisible by n. It is a classic result using Dirichlet's drawer principle. The other important idea for solving this problem is that two integers leave the same remainder when divided by n if and only if their difference is divisible by n. Based on the same ideas, one can prove the following result.

Problem. For any three integers a, b, and c, the number

$$N = abc(a-b)(a-c)(b-c)$$

is divisible by 3.

Solution. If any of the three numbers is divisible by 3, clearly N has the same property, too. Otherwise, a, b, and c give, when divided by 3, one of the remainders 1 or 2. So, we have three numbers (pigeons) that must fall into two residue classes modulo 3 (only 2 pigeonholes). The principle says that there exist two of the numbers in the same residue class — and then their difference is divisible by 3, therefore N (which has all three differences as factors) is divisible by 3, too. The reader can immediately generalize this to n numbers: if $a_1, \ldots, a_n$ are integers, then

$$N = a_1 \cdots a_n \prod_{1 \leq i < j \leq n} (a_i - a_j)$$

is divisible by n. Of course, this is not a very powerful statement (as long as one knows that $\prod_{1 \leq i < j \leq n} (a_i - a_j)$ is always divisible by $\prod_{1 \leq i < j \leq n} (i - j)$), but it is a good exercise for practicing (as a beginner) with the pigeonhole principle. Or, the reader may wish to prove that for all integers a, b, and c, the number

$$abc(a-1)(b-1)(c-1)(a-b)(a-c)(b-c)$$

is divisible by 4.

Here are a few problems involving divisibility.

Problem. If $n + 1$ numbers are arbitrarily chosen from the set $\{1, 2, \ldots, 2n\}$, prove that there exist two among them that are relatively prime.

Solution. The pigeonholes are the sets $\{1,2\},\{3,4\},\dots,\{2n-1,2n\}$, and the pigeons are our $n+1$ given numbers. Two pigeons must fall into the same pigeonhole, that is, two numbers are consecutive, therefore they are also relatively prime. Simple, isn't it? However read the following story before you decide that (if you didn't try to solve the problem first).

It seems that Erdős gave this problem (and others) to the twelve year old boy Louis Pósa, when he first met him, and they were having lunch. As Erdős himself (cited by Ross Honsberger) related, Pósa solved the problems before finishing his soup — and solved this particular problem in less than one minute. Which, by the way, Erdős didn't do at his time: he needed about ten minutes to find the above (simple and natural) solution.

Also note that n is the maximum number of numbers from $\{1,2,\dots,2n\}$ such that any two share a common factor greater than 1, for we can choose $2,4,\dots,2n$ (the even numbers) that have this property (actually they all share the same common factor, but this is not important for our problem).

Problem. Find the maximum number of numbers from $\{1,2,\dots,2n\}$ such that no number divides another.

Solution. We can find n numbers with this property, namely $n+1$, $n+2,\dots,2n$: no number divides another because the smallest multiple of the smallest number exceeds the greatest number. It remains to show that if we choose $n+1$ numbers from $\{1,2,\dots,2n\}$, then there are two (distinct, of course) such that one divides the other. In order to do that, let's note that any positive integer x can be written in a unique manner as $x=2^y(2z+1)$ with nonnegative integers y and z. Although 2^y could also be odd (when $y=0$) we will call $2z+1$ the odd part of x thus expressed. Now our $n+1$ numbers have their odd parts at most equal to $2n-1$ (because they themselves are at most equal to $2n$), hence there are at most n possible odd parts for $n+1$ numbers. The conclusion according to the pigeonhole principle is that there must be two of the numbers that have equal odd parts — of course, in this case, the smaller number is a divisor of the larger number. Observe that a similar (basically identical) problem can be formulated with the set $\{1,2,\dots,2n-1\}$ instead of $\{1,2,\dots,2n\}$.

Problem. Let $m\geq 2$ and a be relatively prime integers. Prove that there are positive integers x and y, both less than or equal to $\sqrt{m}$, and such that one of $ax\pm y$ is divisible by m.

Solution. We consider all expressions $au+v$, with u and v running through the set $\{0,1,\dots,[\sqrt{m}]\}$ (the square brackets indicate the integer part). Their number is $([\sqrt{m}]+1)^2>(\sqrt{m})^2=m$, therefore two of them have to yield the same remainder when divided by m. But if we have $au_1+v_1\equiv au_2+v_2 \pmod{m}$ (with either $u_1\neq u_2$, or $v_1\neq v_2$), then $ax+y\equiv 0 \pmod{p}$ or $ax-y\equiv 0 \pmod{p}$, for $x=|u_1-u_2|$ and $y=|v_1-v_2|$. If we content ourselves with the fact that the requested x and y are nonnegative and not both zero, we can finish the proof here. But we need to show that x and y are strictly positive, so we continue. If $u_1 = u_2$, the above congruence leads to $v_1 \equiv v_2 \pmod{m}$, which, for v_1 and v_2 nonnegative and at most equal to $[\sqrt{m}]$ (thus less than or equal to $\sqrt{m}$) is possible only if we had $v_1 = v_2$. But in this case the pairs (u_1, v_1) and (u_2, v_2) would be equal, which is false. Yet, $v_1 = v_2$ implies $au_1 \equiv au_2 \pmod{m}$, therefore $u_1 \equiv u_2 \pmod{m}$ (because a is prime to m, hence invertible modulo m), then, as before, $u_1 = u_2$ and $(u_1, v_1) = (u_2, v_2)$, a contradiction. Thus $u_1 \neq u_2$ and $v_1 \neq v_2$, and x and y are strictly positive, as desired.

As simple as it may seem, this result (known as Thue's theorem) is important in proving great theorems. We give here only the example of Fermat's theorem about the representation of integers as sums of two squares (actually just one of the important steps for the proof of the general result).

Theorem. *Any prime $p \equiv 1 \pmod{4}$ can be expressed as the sum of two squares.*

Proof. By multiplying side by side all congruences $j \equiv -(p-j) \pmod{p}$, for $j = 1, 2, \ldots, (p-1)/2$, we get

$$1 \cdot 2 \cdots \left(\frac{p-1}{2}\right) \equiv (-1)^{(p-1)/2} \left(\frac{p+1}{2}\right)\left(\frac{p+3}{2}\right) \cdots (p-1) \pmod{p}.$$

But $p \equiv 1 \pmod{4}$ implies that $(p-1)/2$ is even, and if we multiply again this congruence by $((p-1)/2)!$, we get

$$\left(\left(\frac{p-1}{2}\right)!\right)^2 \equiv (p-1)! \pmod{p}.$$

Now, Wilson's theorem tells us that $(p-1)! \equiv -1 \pmod{p}$, thus we have

$$\left(\left(\frac{p-1}{2}\right)!\right)^2 \equiv -1 \pmod{p}.$$

Or, we can say $a^2 + 1$ is divisible by p, for $a = ((p-1)/2)!$ — and we are done with the first step of the proof. In its classical variant, after obtaining a multiple of p that can be written as the sum of two squares, it is proven that the smallest such multiple is exactly p, thus getting the desired conclusion. However, with Thue's theorem at our disposal, we can go on like this. The number a with property that $a^2 + 1 \equiv 0 \pmod{p}$ is definitely relatively prime to p, hence, according to Thue's theorem, we can find positive integers x and y, each less than $\sqrt{p}$, such that either $ax - y \equiv 0 \pmod{p}$, or $ax + y \equiv 0 \pmod{p}$. In both cases we have $a^2x^2 - y^2 \equiv 0 \pmod{p}$. This congruence leads (by using $a^2 \equiv -1 \pmod{p}$) to $x^2 + y^2 \equiv 0 \pmod{p}$, that is, to the fact that $x^2 + y^2$ is a multiple of p. Yes, but we also have $x < \sqrt{p}$, and $y < \sqrt{p}$, hence

$x^2 + y^2 < p + p = 2p$. In conclusion, $x^2 + y^2$ is a nonzero multiple of p, which is also less than $2p$, so only the possibility $x^2 + y^2 = p$ remains, finishing the proof.

We end this section with a few problems with geometric flavor. For instance, everybody knows that among any five points within a unit square there exist two at distance at most $\sqrt{2}/2$ apart. (Here and further, "within" a geometric figure means in the interior of that figure, or on its border.) We prefer, however, the following similar exercise.

Problem. Let ten points be arbitrarily placed within an equilateral triangle of side length 3. Prove that there exist two of the ten points that are situated at a distance of at most 1 apart.

Solution. Divide each side of the triangle into three equal parts, then from each point of division draw parallels to the sides of the triangle. Thus we partition the triangle into nine smaller equilateral triangles, each of side 1. We have ten points, thus there are two of them in the same smaller triangle. (If a point belongs to a side common to two smaller equilateral triangles, we consider it in one of those triangles, freely chosen.) It only remains to see that two points within an equilateral triangle are at a distance at most equal to the side of the triangle apart.

Note that the last claim, as obvious as it may seem, needs a proof. Prove it!

Problem. Six points are within a (closed) disk of radius 1. Prove that we can find two of them at a distance of at most 1 apart.

Solution. If there were seven points, the problem's request would follow easily. It would be enough to consider a regular hexagon inscribed in the given circle (the border of the disk); the radii through the vertices of the hexagon partition the disk into six sectors (each with angle of 60°), and, if seven points are given, two of them must be in the same sector, therefore the distance between them is at most 1. (Prove the last statement!) However, we can use this partition of the disk even in our problem's case, when only six points are given. Choose a partition such that one of the six points (call it A) is on one of the radii that realize the partition (and that bound the sectors). If in one of the 60° sectors delimited by that radius there is another point (say B) then A and B are in the same (closed) sector and the distance between A and B is at most 1. Otherwise in those two sectors (that share the radius passing through A) there is none of the five remaining points (other than A), thus these five points are placed in only four sectors. By the pigeonhole principle, two of them share the same sector, and their mutual distance is less than or equal to 1.

A different approach is the following. Draw the radii of the given circle through its center O and each of the given points. There are two of the radii forming an angle of measure at most 60° (since the sum of the six adjacent angles formed by consecutive radii is 360°). Assume that A and B are two of the six points such that the measure of $\angle AOB$ is at most $60°$. The triangle AOB, however, also has an angle whose measure is at least $60°$ (because the measures of the three angles of the triangle sum to $180°$, the greatest angle has this property). Suppose (without loss of generality) that $\angle ABO$ has measure greater than or equal to $60°$, therefore it also has its measure greater than or equal to the measure of $\angle AOB$. But in a triangle a greater angle is opposite a greater side, thus we finally get $AB \leq OA \leq 1$, as we intended to show.

Note that five points can be chosen within a unit disk such that all their pairwise distances are greater than 1. (How?)

Problem. Show that among any seven points within a triangle of area 1, one can choose three of them such that the area of the triangle they determine is at most 1/4.

Solution. It would be nice if we had nine points. In that case, due to the generalized pigeonhole principle, three of them would lie in one and the same of the four triangles (each of area 1/4) determined by the mid-lines of the triangle and its sides, consequently, the triangle they form would be the one we are seeking. Unfortunately, we don't have nine, but only seven points. Fortunately, the above idea works for seven points, too — but we have to be a little more careful. Namely, let ABC be the given triangle, and let M, N, and P be the midpoints of the line segments BC, AC, and AB respectively. We partition the triangle into the triangles BMP, MCN, and the parallelogram $MNAP$. Three of the seven points must be in one of these three regions. If they are within one of the triangles, clearly they determine a triangle with area at most the area of BMP, or MCN, that is, at most 1/4. But the conclusion is the same if they are within the parallelogram. In this case, we use the following (non-trivial!) statement: the area of a triangle determined by three points within a parallelogram is less than or equal to half the area of the parallelogram. (Prove this claim!) In our case, the area of $MNAP$ is 1/2, hence three points inside it (or belonging to its sides) determine a triangle with area at most 1/4, and the solution ends here.

Proposed Problems

1. Prove the assertions that remained unproved in our text, namely:

   a) The distance between any two points within an equilateral triangle is at most equal to the length of the side of the triangle.

   b) The distance between two points within a sector of a disk with angle of $60°$ is at most equal to the radius of the disk.

   c) The distance between two points within a rectangle is less than or equal to the diagonal of the rectangle. (Although this part has not been proposed in the text, we'll see that a connection exists.)

d) The area of a triangle determined by three points within a parallelogram is less than or equal to half the area of the parallelogram.

2. Prove that among any nine points within a unit square one can find three such that the area of the triangle determined by them is at most 1/8.

3. Let $2n$ ($n \geq 2$) points be given in the plane and let some line segments joining them be drawn. Prove that if there are at least $n^2 + 1$ segments, then at least one triangle is drawn.

4. Prove that for any integers a, b, c, and d, the number

$$N = abcd(a^2 - b^2)(a^2 - c^2)(a^2 - d^2)(b^2 - c^2)(b^2 - d^2)(c^2 - d^2)$$

is divisible by 7.

5. Prove that from any three integers, one can choose two, say a and b, such that $ab(a - b)(a + b)$ is divisible by 10.

6. Prove that every polyhedron has two faces with the same number of edges.

7. What is the maximum number of bishops that can be placed on a chessboard and that do not attack each other? (Two bishops attack each other if they are on the same diagonal.)

8. On a chessboard 17 rooks are placed. Prove that three of them can be chosen such that they do not attack each other. (Two rooks attack each other if they are in the same row or in the same column of the board.)

9. Let $(G, \cdot)$ be a group with n elements and with identity element e, and let $a_1, a_2, \ldots, a_n$ be elements of G (not necessarily distinct). Prove that the product of some (i.e., one or more) of them is e. (And we don't mean an empty product, automatically considered to be e. However, a product with only one factor is allowed.)

10. Let $(G, \cdot)$ be a finite group and let A and B be subsets of G, such that $|A| + |B| > |G|$ (where $|X|$ denotes the cardinality — the number of elements — of X). Prove that $AB = G$; that is, for every $g \in G$, there exists $a \in A$ and $b \in B$ such that $g = ab$.

11. Let $(G, \cdot)$ be a finite group and let H be a nonempty subset of G closed under multiplication. Prove that H is a subgroup of G.

12. Let $(S, \cdot)$ be a semigroup (that is, the operation "$\cdot$" on S is associative and has an identity element). Prove that any element $s \in S$ has an idempotent power. (That is, for each $s \in S$, there exists a positive integer n such that $(s^n)^2 = s^n$.)

13. Write the complete proof of the (weak) theorem of Schur stating that if the positive integers are partitioned arbitrarily in a finite number of sets, then the equation $x + y = z$ has a solution with all components x, y, and z in the same class of the partition.

14. Prove that $14 \leq S(3) \leq 16$. That is, you must find a partition into three classes of $\{1, \ldots, 13\}$ such that each class of the partition does not contain three numbers x, y, and z such that $x + y = z$ (it is said about such sets that they are *sum-free*), and then prove that in any partition with three classes of $\{1, \ldots, 16\}$ one of the classes is not sum-free.

15. Prove Lagrange's four-square theorem, stating that any natural number can be expressed as the sum of (at most) four perfect squares (numbers like $3 = 1^2 + 1^2 + 1^2$ can be written as the sum of four squares only using a zero).

    a) Prove Euler's four-square identity:

$$(a_1^2 + a_2^2 + a_3^2 + a_4^2)(b_1^2 + b_2^2 + b_3^2 + b_4^2)$$
$$= (a_1b_1 + a_2b_2 + a_3b_3 + a_4b_4)^2 + (a_1b_2 - a_2b_1 + a_3b_4 - a_4b_3)^2$$
$$+(a_1b_3 - a_2b_4 - a_3b_1 + a_4b_2)^2 + (a_1b_4 + a_2b_3 - a_3b_2 - a_4b_1)^2,$$

    thus showing that the product of two (or more) sums of four squares is a sum of four squares, too. Infer that if Lagrange's theorem is true for the primes and for numbers 0 and 1, then it is true for every nonnegative integer.

    Further, let p be an odd prime.

    b) Prove that one can find integers a and b such that

$$a^2 + b^2 + 1 \equiv 0 \pmod{p},$$

    therefore there exists a (nonzero) multiple of p that can be expressed as the sum of four squares.

    c) Show that, for any integers a and b, the system of congruences

$$x \equiv az + bt \pmod{p}, \ \ y \equiv bz - at \pmod{p}$$

    has a solution (x, y, z, t) in nonnegative integers, not all 0, and all less than $\sqrt{p}$ in absolute value.

    d) Infer that p itself can be written as the sum of four squares.

Solutions

1. a) Let P and Q be any points situated in the interior of the equilateral triangle ABC, or on its border. We can draw the line through P and Q, consider its intersections P' and Q' with the triangle, and we have $PQ \leq P'Q'$, which shows that it suffices to prove the statement only for points on the border of the triangular surface. Thus we can consider, without loss of generality, P and Q to be on the sides of the triangle. Then note that there exists a vertex of the triangle, say A such that the parallel to PQ through A intersects the opposite side (in this case BC) in an interior point of it, say in M, in such a way that $PQ \leq AM$. But $AM \leq AB$ is clear since one of B and C is more remote from the midpoint of BC than M (and the oblique which is further from the perpendicular is longer), hence we have $PQ \leq AM \leq AB$, as desired.

   Actually, based on similar reasoning, one can prove that the maximum distance between two points within an arbitrary triangle is the length of the triangle's longest side.

   b) We consider the sector in a coordinate system with origin in its center (actually the center of the corresponding disk), the x-axis along one of the radii that delimit the sector, and the other one of these radii forming a positive angle of $60°$ with the x-axis. Let M and N be two points within the sector, whose position can be expressed in polar coordinates, or by their complex affixes. Thus we have $z_M = r_M(\cos t_M + i \sin t_M)$, and $z_N = r_N(\cos t_N + i \sin t_N)$ with $r_M = OM$ and $r_N = ON$ in the interval $[0, R]$ (R being the radius of the sector), and with t_M and t_N from the interval $[0°, 60°]$. A simple calculation (or an application of the cosine law) shows that

$$MN = |z_M - z_N| = \sqrt{r_M^2 + r_N^2 - 2r_M r_N \cos(t_M - t_N)}.$$

   The difference $t_M - t_N$ is in the interval $[-60°, 60°]$, therefore

$$\cos(t_M - t_N) \geq 1/2,$$

   yielding

$$r_M^2 + r_N^2 - 2r_M r_N \cos(t_M - t_N) \leq r_M^2 + r_N^2 - r_M r_N.$$

   Now if, for example, $r_M \leq r_N$, we can write

$$r_M^2 + r_N^2 - r_M r_N = r_N^2 + r_M(r_M - r_N) \leq r_N^2 \leq R^2,$$

and similarly, we prove that the expression under the radical is at most R^2 when $r_N \le r_M$. Thus, either way,

$$MN = \sqrt{r_M^2 + r_N^2 - 2r_M r_N \cos(t_M - t_N)} \le \sqrt{r_M^2 + r_N^2 - r_M r_N} \le R,$$

as we intended to prove.

c) If M and N are points within the rectangle $ABCD$ and P and Q, respectively R and S are the projections of M and N on AB, respectively AD, one readily sees that $MN^2 = PQ^2 + RS^2$, and $PQ \le AB$, $RS \le AD$. Thus

$$MN = \sqrt{PQ^2 + RS^2} \le \sqrt{AB^2 + AD^2} = AC.$$

Clearly, we can have equality only when $\{M, N\} = \{A, C\}$, or $\{M, N\} = \{B, D\}$. In particular, the maximum distance between two points in a square is achieved when the points are opposite corners of the square, and is, of course, the length of the square's diagonal. Actually this does not appear in our text above, but it is useful to prove, for example, that from five points within a unit square there will always be two at distance at most $\sqrt{2}/2$ apart (and this was mentioned in the text).

d) Let X, Y, and Z be points within the parallelogram $ABCD$. The reader can figure out that there is a vertex of the triangle XYZ such that one of the parallels through that vertex to either AB or AD cuts the triangle in two smaller triangles (possibly the parallel coincides with one of the sides of the triangle, and one of the smaller triangles is degenerate, while the other is the initial triangle XYZ). Assume for convenience that the parallel through X to AB intersects the opposite side YZ at the point T, cutting the triangle into two triangular pieces XYT and XZT. Let $d = d(AB, CD)$ be the distance between the parallel lines AB and CD (the height of the parallelogram corresponding to its sides AB and CD). We then have

$$A_{XYZ} = A_{XYT} + A_{XZT} = \frac{1}{2} \cdot XT \cdot d(Y, XT) + \frac{1}{2} \cdot XT \cdot d(Z, XT),$$

where A_{XYZ} is the area of XYZ and $d(Y, XT)$ is the perpendicular distance from Y to XT. Now we have $XT \parallel AB \parallel CD$ and X and T are within the parallelogram, hence $XT \le AB$ and, also

$$d(Y, XT) + d(Z, XT) \le d,$$

therefore

$$A_{XYZ} = \frac{1}{2} \cdot XT \cdot (d(Y, XT) + d(Z, XT)) \le \frac{1}{2} \cdot AB \cdot d = \frac{1}{2} \cdot A_{ABCD},$$

finishing the proof.

2. Partition the square into four little squares with side length $1/2$. We have $9 = 4 \cdot 2 + 1$ points placed in these four little squares, hence three of them are within the same square. By the last part of the previous problem, the area of the triangle determined by them is at most half the area of the little square, that is, at most $1/8$. Of course, the square can be replaced with any parallelogram with area 1, and the result remains true, with the same proof.

3. In other words, a graph with $2n$ vertices and at least $n^2 + 1$ edges must contain a triangle (a complete subgraph with three vertices). We prove this by induction on n. For $n = 2$, let the points A, B, C, and D be given with five of the line segments joining them. Let AB be drawn, and note that either both AC and BC, or both AD and BD are drawn, otherwise, we can only have at most four segments. So, either triangle ABC, or triangle ABD appears. (Observe that the result is also vacuously true for $n = 1$.) Basically the same idea works for the induction step as well.

   Assume the result is true for n, and let $2n + 2$ points be given together with at least $(n + 1)^2 + 1$ segments joining some two of them. Consider again AB to be a drawn segment and look at the set of the other $2n$ points, different from A and B (denote this set by X). If among these points at least $n^2 + 1$ segments are drawn, some three of them form a triangle, by the inductive hypothesis. Otherwise, among them there are at most n^2 segments, thus there are at least $2n+1$ segments joining either A or B with one of the $2n$ points from X. Let $A_1, \ldots, A_k$ be the points from X joined with A, and let $B_1, \ldots, B_l$ be the points from X joined to B. The total number of these points is $k + l = 2n + 1 > 2n$, thus greater than the number of elements of X, to which they all belong, hence there must be two of them that coincide. But the A_is are all distinct, and so are the B_js. The only possibility that remains is that $A_i = B_j$ for some $1 \le i \le k$ and some $1 \le j \le l$. Now, if we denote $C = A_i = B_j$, we have the triangle ABC drawn.

   Note that n^2 segments can be drawn with no triangle appearing; just split arbitrarily the $2n$ points into two sets with n elements each and draw all segments that join a point from one set with a point from the other (thus obtaining a complete bipartite graph on $2n$ vertices). This is a case of Mantel's theorem asserting that the maximum number of edges of a graph on m vertices that does not contain triangles is $[m^2/4]$ (and seems to be another problem that Erdős gave to Pósa in his early years, and Pósa solved it very quickly).

4. If one of the four numbers a, b, c, d is divisible by 7, the conclusion is straightforward. Otherwise, they are congruent to ± 1, ± 2, or ± 3 modulo