Prof. Jim Loveys
MATH 352 : Problem Seminar

Rosalie Bélanger-Rioux
260235216

The Pigeonhole Principle

The pigeonhole principle, also known as Dirichlet's box or drawer principle, is a very straightforward principle which is stated as follows :

Given n boxes and m > n objects, at least one box must contain more than one object.

This was first stated in 1834 by Dirichlet. The proof is very easy : assume we are given n boxes and m > n objects. Then suppose, to the contrary, that no box contains more than one object, i. e., all n boxes contain either 0 or 1 object. This implies that the total number of objects, m, is smaller than or equal to n. This contradicts our hypothesis, and so the pigeonhole principle is true. Of course, this principle has many formulations and variations. For example, if n objects are put into n boxes, then at least one box is empty iff one box contains more than one object. Let us now give a more formal statement of Dirichlet's box principle :

There exists a one-one correspondence between two finite sets A and B, f : A -> B iff their number of elements is the same, i. E., |A| = |B|.

No matter how we state the pigeonhole principle, using it in a problem solving situation implies two things : finding your "boxes" and finding your "objects". The nice thing about this is that just about anything in mathematics can be made a "box" (a polygon, a volume in space, an edge, a color, a distance, an interval, a value…) or an "object" (a point, a polynomial, a vertex, a line…) to one's liking. Further, more often than not, noticing that the pigeonhole principle applies to a certain problem often reduces by a whole lot the mind-numbing calculations or analysis of different cases one has to perform in order to find the solution.

We shall start with three easy examples in which I hope it will be obvious to the reader that the pigeonhole principle applies, the "objects" being clearly stated in the problem statement. What we now have to find are the boxes. However, for the problems that will follow these examples, either the boxes are not very easy to find, or it takes a little imagination and understanding of what is going on to find the proper objects.

Example 1 (Putnam, 2002, A2)

We start with a simple problem, to illustrate our point.

Show that, given any 5 points on the surface of a sphere, there exists a closed hemisphere which contains at least four of the five points.

We know the pigeonhole principle must come into play here. We simply need to divide the sphere in such a way that it is clear n+1 points will have to be distributed in n parts of the sphere. Here, the clever trick is to divide the sphere into two hemispheres by choosing a great circle passing through any two points (such a great circle through any two points, of the same radius as the sphere, always exists). Clearly, no matter which hemisphere we choose, those two points will be in it, and we are left with three points in two hemispheres. At least two of them will be in either hemisphere, so we are done.

At least four of any five points given on the surface of a sphere are contained in a hemisphere. Here the pigeonhole principle renders trivial a Putnam problem!

Example 2

Again, here, "dividing" judiciously into boxes does a lot of the work.

We place 10 points on a disk of diameter 5.
Show at least two of these points are at a distance less than or equal to 2.

Simply draw a circle of radius 1 concentric with the disk, and divide the rest of the disk into 8 congruent wedges, as shown in the following diagram. We have 10 points and 9 boxes, so one of those will contain two points. But the small circle at the center clearly cannot contain two, else we are done. So any one of the 8 wedges contains two points. It is obvious that the distance between points A and D (or B and C) is 1,5, which is less than 2. So we will calculate the distance between A and B and A and C and show these are less than 2 also, taking $(1/\sqrt{2}) \approx 0{,}707$ and A(0; 2,5), B($2{,}5\sqrt{2}$; $2{,}5\sqrt{2}$), C($\sqrt{2}$; $\sqrt{2}$), D(0; 1) without loss of generality.

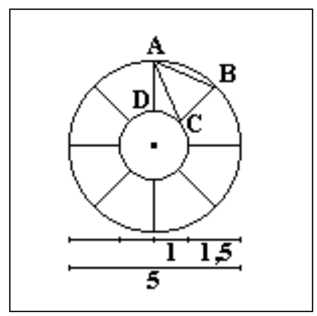

$AC = \sqrt{\left((1/\sqrt{2})^2 + (5/2 - (1/\sqrt{2}))^2 \right)}$;
$AC^2 = ½ + 25/4 + ½ - (5/\sqrt{2})$;
$AC^2 \approx 7 + ¼ - 3{,}535$;
$AC^2 \approx 3{,}714 < 4$.

$AB^2 = (5/(2\sqrt{2}))^2 + (5/2)(1-(1/\sqrt{2}))^2$;
$AB^2 = 25/8 + (25/4)(3/2 - \sqrt{2})$;
$AB^2 \approx 25/8 + (25/4)*0{,}085$;
$AB^2 \approx 3{,}661 < 4$.

So we have shown, using the pigeonhole principle, that for any 10 points taken on a disk of diameter 5, at least two of those are at a distance less than or equal to 2.

Example 3

Our last example is also a geometric one, but here the "divisions", i. e. the boxes, are not whole parts of the given geometric figure, indeed, they are points in it. The difficulty is realizing that it is these precise points that will help (or having a hunch of what those points are by making a few drawings).

We are given a square and 9 lines cutting through it. Each line cuts the square into two quadrilaterals, and the ration of the areas of those quadrilaterals is 2:3. Prove that at least three of these 9 lines pass through the same point.

Since each line divides the square in two quadrilaterals, each line cuts the square at two non consecutive edges. The figures obtained are trapezoids with two right angles, and their area is equal to their base (here the side of the square) multiplied by the height of their "midline", the line perpendicular to their two parallel edges and at equal distance from those two edges. Since for any of those trapezoids their base is the same, the ration of their midlines must be 2:3. But there are only two possible midlines for the trapezoids, so we get four different points, call them midpoints, two on each midline, where the midline of two trapezoids intersects with one of the 9 given lines. Every one of those lines must pass through one of these midpoints. We have 4 points and 9 lines, so there is at least one point through which pass three lines.

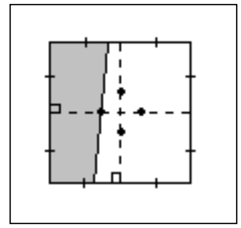

We have shown that, given 9 lines, each cutting a square into two trapezoids whose areas are in the ratio 2:3, at least three of the lines pass through the same point.

We now move on to more difficult problems, which generally necessitate a little more twist, or work, before applying the pigeonhole principle.

Problem 1 (Putnam, 1994, A3)

Here is a coloring problem where the pigeonhole principle comes in handy; it simplifies and clarifies the proof.

We 4-color a right isosceles triangle, its two equal sides being of length one. Show there are always 2 points of the same color that are at least at a distance of 2-√2 apart.

For this we consider the following 9 points on the triangle.

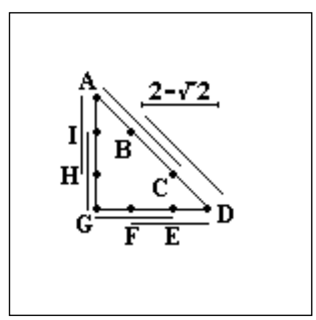

A(0 ; 1) D(1 ; 0) G(0 ; 0)
B(√2 – 1 ; 2 - √2) E(2 - √2 ; 0) H(0 ; √2 – 1)
C(2 - √2 ; √2 – 1) F(√2 – 1 ; 0) I (0 ; 2 - √2)

Each of these points will be colored, just like the rest of the triangle. We shall first calculate the distances between each pair of points, then draw a graph where each edge means two points can be of the same color, i. e. they are closer than 2 - √2. The pigeonhole principle will then help solve the problem. Let us note that the points were chosen such that the following distances are exactly 2 - √2 : AB and CD. But it also happens that the following distances are also exactly 2 - √2 : AH, BF, CH, DF, EG, GI. Obviously, looking at the diagram, the following distances are less than 2 -√2 : AI, BC, BH, BI, CE, CF, DE, EF, FG, GH, HI. We shall now calculate FH, all the other distances being greater that 2 -√2 (or previously listed).

$FH = \sqrt{(2(\sqrt{2} - 1)^2)}$; $FH = \sqrt{2}(\sqrt{2} - 1)$; $FH = 2 - \sqrt{2}$.

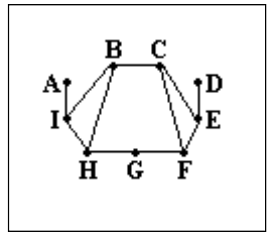

Finally, we can draw the graph of the 9 points, with edges between points that are strictly less than 2 - √2 apart. If we look at the graph, we see there are only two complete graphs, both with three vertices. So color those triangles with one color each; we can see two colors are not enough to color the three remaining vertices, and so we will have two points of the same color that will be apart by more than 2 - √2. So color only one of the triangles, say BHI with color 1, and let us color vertex A with color 2 (without loss of generality, by symmetry). It is clear that the remaining 5 vertices (C, D, E, F, G) cannot be colored using only colors 3 and 4, else we get two points that are at least 2 - √2 units apart. But the pigeonhole principle tells us that for 9 vertices and 4 colors, at least 3 vertices will have the same color, which is not possible as we have just seen. So we do not need to check other possibilities.

We have proved that, given an isosceles right triangle of equal sides 1 and 4 colors to color all its points, there at least two points which are at least a distance of 2 - √2 apart. Of course, in this problem, it takes some insight to decide to consider only a limited number of points instead of the whole triangle, to pick the right points, and the right number of points (5 points would not be enough for the proof, we need 9).

Problem 2 (Putnam, 1953, A2)

Of course, this result is well-known now, but in the 1950's this topic was only beginning to surface.

We color the edges of the complete graph with 6 vertices (and so 15 edges) with two colors. Show we can always find three points such that the edges joining them are of the same color.

The proof is quite simple. Take any vertex, name it a. It has 5 edges, so by the pigeonhole principle at least three of those are of the same color. Say ab, ac and ad are of the same color, blue (without loss of generality). Then if either bc, bd or cd are blue, we are done. But if none is, then bcd is a red triangle. Either way, we have found three points in the complete graph of three vertices such that the three edges joining them are of the same color.

Remarks. A very similar problem in graph theory, also easily solved using the pigeonhole principle, is the following : suppose there are n people at a party, then show at least two of them have the same number of acquaintances, where a person is not acquainted with himself or herself. This is easy, for the maximum number of acquaintances for one person is n-1. For everyone to have a different number of friends would mean someone at the party has no friends at all, a contradiction to the fact that someone is friends with everyone (let alone that this person would not be at the party if he or she did not know anyone)! Also, consider the complete graph of six vertices, assume all edges have different length. Some edges form triangles between themselves. Prove there is at least one edge that is both the shortest edge of a triangle and the longest edge of another triangle. To prove this, let us start by coloring blue all edges that are the shortest edge of at least one triangle, and coloring the rest, if any, red. Then, we know there is at least one triangle such that its edges are all blue (we cannot have a triangle all red since its shortest edge should be blue). Take the longest edge of that triangle; clearly it is blue because it is the shortest edge of some other triangle.

Problem 3 (Putnam)

Here, it is more a matter of counting the elements of a set in two different ways, one of them using the drawer principle.

Consider the set V of all the vertices of a hypercube in n space, $V = \{(\varepsilon_1, \varepsilon_2, ..., \varepsilon_n) : \varepsilon_i = \pm 1\}$.
Prove that, for any subset S of V such that $|S| > (2^{n+1})/n$,
S contains at least one equilateral triangle, that is
three vertices that together form an equilateral triangle.

Let us consider the set X, a subset of V × V,

$X = \{(a, b) : a = (a_1, a_2, ..., a_n) \in S, b = (a_1, a_2, \ldots, -a_j, ..., a_n)$ for some $0 \le j \le n\}$,

or X the set that contains pairs of vertices (a, b) such that a is in S, b is adjacent to a. Since for any vertex a, a has n adjacent vertices, $|X| = |S| \cdot n > (2^{n+1})$ by hypothesis. Now, if there exists a vertex k, not necessarily in S, such that (c, k), (d, k) and (f, k) are in X, so c, d and f all different in S, then we have an equilateral triangle : c, d and f each have, pairwise, precisely two different coordinates – so the distance between each pair of them is equal, and so they define an equilateral triangle.

So suppose X does not contain three such vertices c, d and f such that (c, k), (d, k) and (f, k) are in X, c, d and f all different. Then, for all vertices v in V, there are at most two other vertices u and w such that (u, v) and (w, v) are in X. What would then be the size of X? $|X| \le |V| \cdot 2 = 2^{n+1}$. This contradicts the previous fact that, by the hypothesis that $|S| > (2^{n+1})/n$, $|X| > (2^{n+1})$. So supposing there are no equilateral triangles formed by vertices in S arrives at a contradiction.

We have proved that, given a subset S of V, the set of vertices of a hypercube in n space, $|S| > (2^{n+1})/n$ implies that some vertices of S form at least one equilateral triangle. Of course, the pigeonhole principle is barely used here since by intuition we know that, in the case mentioned above, $|X| \le |V| \cdot 2$, but it simply confirms that this expression is true.

Problem 4 (Putnam, 1989, A4)

Okay, we have now worked on many problems where the boxes and objects could actually be "seen", or at least imagined quite easily. In this Putnam problem, it is a bit more difficult to see at first the concept of "boxes", but just as useful! In fact, we could call this a "continuous" version of the pigeonhole principle.

Given any point P inside a regular polygon of 2n + 1 sides inscribed in a circle of radius 1, show we can find a k > 0 and two vertices of the polygon such that their distances from point P differ by less than $1/n – k/n^3$.

Here, we shall use the pigeonhole principle in the following way : first of all, we need to find the supremum D of all distances between any point P inside the polygon and any vertex of the polygon. Next, we shall divide this distance by 2n, so that we get 2n boxes, which are the following intervals :

$$(0; D/2n], (D/2n; 2D/2n], \dots , ((2n-2)D/2n; (2n-1)D/2n], ((2n-1)D/2n; D].$$

Now, take any point P and find the distances d_i from P to each vertex of the polygon; there are 2n+1 such d_i's, with $0 < d_i \leq D$ for all i, so that by the pigeonhole principle at least one interval contains at least two of those distances. If we find a k positive and show that D/2n, the length of our intervals, is smaller than $1/n – k/n^3$, whatever n and the point P, we are done!

So, the maximal distance between a vertex and P is certainly less than 2. To determine this distance more accurately, let us consider the following : take a polygon with an odd number of vertices, then certainly the supremum of the distance between any point P in the polygon and a vertex of the polygon is the distance between two (almost) opposite vertices, say vertex 1 and vertex n+1, or n+2, if we label the vertices of the polygon from 1 to 2n+1 starting at the top and going clockwise around the polygon.

So $D = 2\cos\theta$, where $\theta = (\pi – (2\pi n / (2n+1)))/2$, or $D = 2 \cos(\pi (1/2 – n/(2n+1)))$, $D = 2 \cos(\pi/(4n+2))$, the maximal distance between a vertex of the polygon and P. So now, let us "divide" that maximal distance into 2n boxes, such that at least one of those boxes contains at least two of our 2n+1 vertices. If the "length" of our "boxes" is less than $1/n – k/n^3$ for some k positive, we are done! So we need to find a positive k and show that, for this k, $D/2n < 1/n – k/n^3$, or $\cos(\pi/(4n+2)) < 1- k/n^2$ for all n. This will be provided by the following lemma.

Lemma : for all k such that $0 < k < \pi^2/144$, we have $\cos(\pi/(4n+2)) < 1- k/n^2$.

Proof (of lemma): first, note that $n \geq 1$, so $0 < \pi/(4n+2) \leq \pi/6$ for all n. In this interval, we see that $\cos x < 1 – x^2/4$ in $(0, \pi/6]$. Indeed, the graphs of these two functions do not intersect in that interval, their value is the same at 0 but $1 – \pi^2/36\cdot 4 \approx 15/16 > 0,9$ is greater than $\cos(\pi/6) = \sqrt{3}/2 \approx 0,866$, and they both are continuous functions. So we have $\cos(\pi/(4n+2)) < 1- (\pi/(4n+2))^2/4$, and we claim that, for all $n \geq 1$,

$$1 - (\pi/(4n+2))^2/4 < 1 - k/n^2 \text{ for all } 0<k<\pi^2/144,$$

or that $k/n^2 < (\pi)^2/4(4n+2)^2$. However, we have $k/n^2 < \pi^2/144n^2$ for all n, so it is sufficient to prove $\pi^2/144n^2 \leq \pi^2/4(16n^2+16n+4)$ for all n, or that $4(16n^2+16n+4) \leq 144n^2$. So if we prove $16n^2+16n+4 \leq 36n^2$, for all n, we have proved $1 - (\pi/(4n+2))^2/4 < 1 – k/n^2$ for all n and for $0 < k < \pi^2/144$, and so we have proved the lemma. But certainly $16n+4 \leq 20n^2$ for all n greater than or equal to 1, and so $16n^2+16n+4 \leq 36n^2$. This proves that, for all k such that $0 < k < \pi^2/144$, and for all n, we have $\cos(\pi/(4n+2)) < 1- k/n^2$.

Finally, we have proved that, given any point P inside a polygon of 2n+1 sides, there exists a $k > 0$ and at least two vertices of that polygon such that their distances to P differ by less than $1/n - k/n^3$, namely with $0 < k < \pi^2/144$. The trick here is to see that, having 2n + 1 objects, you want to create 2n boxes… by dividing the maximal distance between P and a vertex by 2n.

Remark. In fact, this can be viewed also as an application of the Fubini Principle : if the average number of objects per box is x, then at least one box has at least x objects and at least one contains at most x objects.

Problem 5

Given a subset S of size 10 of the first 100 natural numbers, show it is always possible to find two non-empty, disjoint subsets of S, S_1 and S_2, such that the sum of the elements of S_1 equals the sum of the elements of S_2.

Let us first count the number of subsets of S. Since each element of S can be either chosen or not, we have $2^{10} - 2 = 1022$ different possible proper non-empty subsets of S. Now, let us count all the different possible sums for any subset of S. Since S_1 (or S_2) must have at least one element each, and at most 9, the minimal sum of all the elements of either subset of S is 1, and the maximal sum is $864 = 92 + 93 + \ldots + 99 + 100$. Consequently, by Dirichlet's drawer principle, having more possible subsets than possible sums, at least two possible different subsets of S must share the same sum. But these subsets must be disjoint. However, it is clear that, since they share the same sum, neither one contains the other, and so if we call T the intersection of S_1 and S_2 , then $S_1 - T$ and $S_2 - T$ are disjoint subsets of S such that the sum of their respective elements is the same. Here are three other results in number theory that are quite well known and can be easily proved using Dirichlet's principle. We shall only discuss them briefly.

Result 1

Choose any n+1 elements of the set $S = \{1, 2, 3, \ldots, 2n\}$, then there is always one of those you picked which divides another one you picked. Proof: let us factor out any factor 2 that a positive integer m contains, and write $m = 2^r q$, with q an odd integer. But for all elements of S, there are only n odd numbers which can be used to express them in their factored form, i. e., 1, 3, 2n-3, 2n-1. By the box principle, at least two numbers in the n+1 picked will have the same q, so the smaller one will divide the bigger one.

Result 2

In any sequence of distinct mn+1 real numbers, there exist either one increasing subsequence of at least m+1 numbers, or one decreasing subsequence of at least n+1 numbers, when read from left to right. Proof: assign "coordinates" (d, i) to every number in the sequence, where for any number k, d is the length of the longest decreasing subsequence that starts with k, and i is the length of the longest increasing subsequence that starts with k. So we will obtain mn+1 pairs (d, i). Now, if any d is greater than or equal to n+1, or any i greater than or equal to m+1, we are done. But if every d and i are lower than, respectively, m+1 and n+1, then we have at most mn different possible coordinates, and so two distinct numbers in our sequence have the same coordinates, a contradiction.

Result 3

This last result can be used to play card tricks, such as “teaching the cards to order themselves” (see *The Last Recreations*, by Martin Gardner, for an explanation and some further references). Quite simply, arrange nm numbers (distinct or not) in a rectangular array of m rows by n columns. Rearrange every row independently, one by one, so that the numbers in each row are in non-decreasing order, from left to right. For example, take the rows “1 6 2 6 9” and “13 12 2 2 1”, which when rearranged become “1 2 6 6 9” and “1 2 2 12 13”. When done, rearrange the columns similarily, one by one, so that the numbers are in non-decreasing order, from front to back. Then look at the rows again : even though permuting the numbers in each column has altered the arrangements in the rows, the numbers in the rows are still in non-decreasing order from front to back. Proof: assume not, then there is a smaller number a placed to the right, and maybe to the back also, of a larger number b after rearranging the columns of array A. But the columns have all been ordered, so all numbers to the back of b, in the same column, are no smaller than b, and bigger than a (since b>a), and all of those in front of a, in the same column, are no bigger than a, so smaller than b. Consider now the moment just before the columns where rearranged, after the rows had been ordered. We want to show that there is no prior arrangement of numbers into an n by m array, such that the numbers in each row of this array have been sorted in non-decreasing order, that can lead to having at least one row that is not ordered anymore after rearranging the columns. So we want to “return” all numbers in A to their position prior to rearranging the columns. We will start by replacing all numbers behind b in b's column, including b, and all numbers in front of a in a's column, including a. Since b is to the left, and possibly in front of, a then we have at least m+1 numbers to replace in m rows. But for two numbers to be replaced in the same row, they must be in different columns, so that a number greater than or equal to b will be in the same row, to the left, of a number smaller than or equal to a, with a<b. This is a contradiction, since we are returning to the moment before rearranging the columns, that is, after rearranging the rows. So no n by m array of numbers can lead to disordered rows when the numbers in the rows, and then in the columns, have been ordered in non-decreasing fashion.

Problem 6 (Putnam, 1994, A4)

In this problem, using Dirichlet's box principle saves us a long and cumbersome look of all the possible cases, while it is certainly not clear at first where and how this principle might be applied.

We are given two 2 by 2 matrices A and B with integral values such that A, A + B, A + 2B, A + 3B, and A + 4B all have inverses with integral values. Show the inverse of A + 5B also has integral values.

First of all, take any matrix $C = (c_{ij})$. It is invertible iff $\det(C) = c = c_{11}c_{22} - c_{12}c_{21} \neq 0$. Its inverse is $D = C^{-1} = (d_{ij})$ such that $d_{11} = c_{22}/c$, $d_{12} = -c_{21}/c$, $d_{21} = -c_{12}/c$, $d_{22} = c_{11}/c$. Note that D has integer values if and only if $c \mid c_{ij}$. Now, look at $\det(D) = d = (d_{11}d_{22} - d_{12}d_{21})$, $d = (c_{11}c_{22} - c_{12}c_{21})/c^2 = c/c^2 = 1/c$ (since c is not zero). However, in our situation, we require C and D to have integral values, and so the determinants c and d are integers. We get $dc = 1$, which implies $c = \pm 1$.

Now, let us find $\det(A + nB)$ for any n. We have :

$A = (a_{ij})$, $B = (b_{ij})$, $A + nB = (a_{ij} + nb_{ij})$, and $\det(A) = a$, $\det(B) = b$.

So $\det(a + nB) = (a_{11}a_{22} + na_{22}b_{11} + na_{11}b_{22} + n^2b_{11}b_{22}) - (a_{12}a_{21} + na_{12}b_{21} + na_{21}b_{12} + n^2b_{12}b_{21}) = a + kn + bn^2 = \pm 1$ for $n = 0, 1, 2, 3, 4$, with $k = (a_{22}b_{11} + a_{11}b_{22} - a_{12}b_{21} - a_{21}b_{12})$. So we have a polynomial that can take only two possible values, that is 1 or –1, for five different values of n = 0, 1, 2, 3 and 4. By the pigeonhole principle this implies that either 1 or –1 is taken at least three times by det(A + nB). This is impossible for a polynomial of degree 2 unless it is constant, so det(A + nB) is constant, and det(A + 5B) = ±1. The inverse of A + 5B has integral entries.

We have proved that, given A and B, 2 by 2 matrices with integral values such that the inverses of A, A + B, A +2B, A + 3B, and A + 4B have integral values, A + 5B also has an inverse that has integral entries. In fact, we have showed this is true for all n!

Problem 7 (Putnam 1993, A4)

The numbers 19 and 93 in the following problem were chosen because of the year of the contest, but this can be readily generalized.

Choose a sequence X of 19 positive integers no greater than 93, and a sequence Y of 93 positive integers no greater than 19, all integers not necessarily distinct. Prove we can find two subsequences, one in each sequence X and Y, such that the sum of their elements is equal.

So first, consider the sum a_i of a subsequence (x_i) of X, $a_i = x_1 + x_2 + \ldots + x_i$, and the sum b_j of a subsequence (y_i) of Y, $b_j = y_1 + y_2 + \ldots + y_j$. Here we take :

$0 \le x_1 \le x_2 \le \ldots \le x_i \le \ldots \le x_{19} \le 93$, $0 \le y_1 \le y_2 \le \ldots \le y_i \le \ldots \le y_{93} \le 19 \le 93$.

It is difficult to see where to start, but let us only consider for the proof those subsequences which start with the smallest element of X (or Y) and are non-decreasing.

If we can prove the result with only the help of those particular subsequences, certainly the result holds. Now, let us look at the two subsequences (x_{19}), (y_{93}), or the whole sequences X and Y, in non-decreasing order. Assume $a_{19} \leq b_{93}$, else in the following proof reverse the roles of x and y, a and b.

Now, since $a_i \leq a_{19} \leq b_{93}$, for all i smaller than or equal to 19 we shall define f(i) = j as the smallest index j of a sum b such that $a_i \leq b_j$. We can then form the set S :

$$S = \{b_{f(1)} - a_1, b_{f(2)} - a_2, \dots , b_{f(19)} - a_{19}\},$$

which contains 19 elements, our "objects". Should any element of S be zero, we are done, since we have found subsequences $(y_{f(k)})$ and (x_k) such that the difference of their sums is zero, so their respective sums are equal. So assume any element of S is greater than or equal to 0. Further, no element of S can be greater than or equal to 19, because we would have $b_{f(k)} \geq a_k + 19$, but since $b_i - b_j \leq 19$ for all i>j, this would imply $b_{f(k)} > b_{f(k)-1} \geq a_k$, which contradicts the way we constructed our function f (here we have to be careful not to confuse f(k)-1 and f(k-1), which are not necessarily the same). So elements of S can take the values from 1 to 18 only, those are our "boxes", so at least two elements in S will have the same value, say $b_{f(s)} - a_s = b_{f(t)} - a_t$, with say s<t. Then we are done, because this implies $a_t - a_s = b_{f(t)} - b_{f(s)}$, or $x_{s+1} + x_{s+2} + \dots x_{t-1} + x_t = y_{f(s)+1} + y_{f(s)+2} + \dots + y_{f(t)-1} + y_{f(t)}$. We have found a required subsequence.

Hence we have proved that, by choosing a sequence X of 19 positive integers no more than 93, and a sequence Y of 93 positive integers no more than 19, we can find a subsequence of X and one of Y such that the sum of their elements is equal. Of course, not knowing that this can be proved using the pigeonhole principle makes the problem that much harder. Even so, finding the right objects and the right boxes takes a lot of fiddling around with different types of subsequences and sums. This only confirms the fact that no amount of knowledge in mathematics can replace intuition.

REFERENCES

Internet

CUT-THE-KNOT – INTERACTIVE MATHEMATICS, "Party Acquaintances", http://www.cut-the-knot.org/Curriculum/Combinatorics/ThreeOrThree.shtml, Alexander Bogomolny, November 29, 2006.

PUTNAM PROBLEMS, "Archive", http://www.kalva.demon.co.uk/putnam.html, John Scholes, November 25, 2006.

MATHWORLD, "Fubini Principle", http://mathworld.wolfram.com/FubiniPrinciple.html Eric W. Weisstein, November 30, 2006.

Books

GARDNER, Martin. *The Last Recreations*, Copernicus, New-York, 1997, 392 p.