Puzzles that Solve Themselves

Peter Winkler *

February 23, 2018

Abstract

Roughly speaking, we say that a puzzle "solves itself" if the stupidest way you can think of to get an answer works. Often, that means guessing an answer, and then fixing it in the obvious way until it becomes a solution. But how do you know when this happy state of affairs exists?

1 Problems

Let's start with a simple example.

Flipping the Bulbs

In front of you is a 9×9 array of light bulbs, some on, some off. At the left end of each row, and at the top of each column, is a switch that will reverse the state of every bulb in that row or column.

Is it possible to flip switches in such a way that every row and every column has most of its bulbs on?

The obvious thing to do here is to find some line (row or column) that has most of its bulbs off, then flip its switch. Trouble is, that might cause some intersecting lines to go from mostly on to mostly off; thus, you might *increase* the number of bad lines. Then, after more corrections, you might find yourself back at the original configuration without having found a solution.

But a little thought will convince you that this process can never cycle back to any previous configuration, and in fact will solve the problem rather quickly. The key observation is that when you flip a line that has more bulbs off than on, *you increase the total number of lit bulbs.* This can't go on forever and only reaching a solution can stop you.

*Department of Mathematics, Dartmouth College, Hanover NH 03755-3551, USA. Research supported by NSF Grant DMS-1600116.

Sometimes you don't even have to be clever enough to find *any* possible route to a solution.

Breaking a Chocolate Bar

You have a rectangular chocolate bar marked into $m \times n$ squares, and you wish to break up the bar into its constituent squares. At each step, you may pick up one piece and break it along any of its marked vertical or horizontal lines.

How should you break up the bar so as to minimize the number of breaks needed?

This puzzle, which I first heard from the late, great mathematician Paul Halmos, looks geometrical but isn't. The fact is, breaking the $m \times n$ bar into its constituent square takes exactly $mn - 1$ breaks, no matter how you do it, simply because every break increases the number of pieces by one.

Obvious when you know it, but many smart folks have been led astray by the grid lines and failed to count pieces.

The next puzzle really does have some geometric content.

Red Points and Blue

Given n red points and n blue points on the plane, no three on a line, can you find a "heterosexual" pairing of red and blue points so that if you connect each red point to its blue mate with a line segment, no two line segments cross?

Let's be dumb and match up the points any old way, then draw in the corresponding line segments. Maybe they never cross!

If they do, pick two segments that cross, and switch partners to that they don't cross any more. Great. But, of course, that action might create many more crossings. Ugh.

Ah, but uncrossing two segments always *reduces the total length of the crossings.* Why? Because crossing segments are the diagonals of a convex quadrilateral, and replacing them with opposite sides reduces length since they no longer have to meet in the middle of the quadrilateral. (Technically, we are employing the "triangle inequality" here.)

There are only a finite number of ways (namely, n factorial) to match up the red and blue points, so eventually you must reach a matching with no crossings.

Conceptually speaking, you can prove non-algorithmically that such a matching exists by just choosing, from the start, the matching that minimizes the sum of the pairwise distances between matched pairs of points. But if you really need to *find* the matching, the above untangling scheme typically works quite fast.

We continue with a familiar campus entity–the Athletic Committee.

Picking the Athletic Committee

The Athletic Committee is a popular service option among the faculty of Quincunx University, because while you are on it, you get free tickets to the university's sports events. In an effort to keep the committee from becoming cliquish, the university specifies that no one with three or more friends on the committee may serve on the committee—but, in compensation, if you have three or more friends on the committee you can get free tickets to any athletic event of your choice.

To keep everyone happy, it is therefore desirable to construct the committee in such a way that even though no one on it has three or more friends on it, everyone *not* on the committee *does* have three or more friends on it.

Can this always be arranged?

This problem (in an abstract form, with the number 3 replaced by an arbitrary integer k) arose in the work of my computer science colleague Deeparnab Chakrabarty. What's the dumbest way to try to solve it? How about this: start with an arbitrary set S of faculty members, as a prospective Athletic Committee. Oops, Fred is on the committee and already has three friends on the committee? Throw Fred out. Mona is not on the committee, but has fewer than three friends on it? Put Mona on. Continue fixing in this haphazard manner.

Now, why in the world would you expect this to work? Clearly, the above actions could make things worse; for example, throwing Fred off the committee might create many more Monas; maybe we should have thrown off one of Fred's on-committee friends instead. So there doesn't seem to be anything to prevent cycling back to the same bad committee. Moreover, even if you don't cycle back, there are exponentially many possible committees and you can't afford to consider every one. Suppose there are 100 faculty members in all; then the number of possible committees is $2^{100} > 10^{30}$ which, even if you spent only a nanosecond considering each committee, would take a thousand times longer than all the time that has passed since the Big Bang.

But if you try it—and if there's one idea that you take from this paper, it's *try it!*—you will find that after shockingly few corrections, you end up with a valid committee. And this happens whether in situations where there is only one valid committee, as well as when there are many.

How can this be? Well, as in Flipping the Bulbs, perhaps there is something that is improving each time you throw someone off or add someone to the current prospective committee. Let's see: when you throw someone off, you destroy at least three on-committee friendships; when you put someone on, you add at most two. Let $F(t)$ be the number of friendships on the committee *minus* $2\frac{1}{2}$ *times the number of people on the committee* at time t. Then when Fred is thrown off, $F(t)$ goes down by at least $\frac{1}{2}$. When Mona is put on, $F(t)$ *again* goes down by at least $\frac{1}{2}$. But $F(0)$ can't be more than $(100 \times 99)/2 - 250 = 245$ and $F(t)$ can never dip below -250, so there can't be more than $2 \times (245 - (-250)) = 990$ steps total. (A computer scientist would say that the number of steps in the process is at worst quadratic in the number of faculty members.)

In practice, the number of steps is so small that if there are 100 faculty members and you start with (say) the empty committee, you will reach a solution easily by hand. Of course, you'll need access to the friendship graph, so you might need to do some advance polling. It'll be interesting to see who claims friendship with whom that isn't reciprocated!

Sometimes the correction process is continuous.

Squaring the Mountain State

Can West Virginia be inscribed in a square?

It must be tough living in a state with two panhandles, but that doesn't mean you can't make a square map of your state in which the state outline exactly reaches all four edges. Certainly you can make a *rectangular* map with this property, just by orienting the state in the familiar way—north equals up—and drawing horizontal and vertical lines through the northernmost, southernmost, easternmost and westernmost points in the state. You won't have a square; West Virginia is slightly wider than it is tall.

Now rotate the state slowly clockwise (say), moving the horizontal lines smoothly up and down and the vertical ones left and right so as to stay tangent to the state boundary. When you've got the state rotated 90 degrees, so that it's northern panhandle is pointing to the right, the rectangle in which it is inscribed will be too tall to be a square instead of too short. It follows (by the *intermediate value theorem*, if you must know) that at least once during the rotation, the horizontal and vertical sides of the rectangle were the same length. And at that moment, you had WV where you wanted it—inscribed in a square.

We wind up with a marvelous puzzle devised by ace probabilist and puzzle-maker Ander Holroyd, who as you read this is visiting Cambridge University.

Self-Referential Number

The first digit of a certain 8-digit integer N is the number of zeroes in the (ordinary, decimal) representation of N. The second digit is the number of ones; the third, the number of twos; the fourth, the number of threes; the fifth, the number of fours; the sixth, the number of fives; the seventh, the number of sixes; and, finally, the eighth is the number of distinct digits that appear in N. What is N?

If you try to work out this number by intelligent reflection, it ain't easy. Instead, pick any 8-digit number, say M, and write out a new 8-digit number M' as follows: the first digit of M' is the number of zeroes in M, the second is the number of ones in M, etc., and the last digit is the number of different digits in M. Now repeat, starting with M'. In short order you will find that you have converged to a number that doesn't change, and that's the unique answer; I leave it to you to discover it.

There's one catch. In all the previous problems, we could determine exactly why the obvious procedure works so well. But neither Ander nor I knows why this particular puzzle is self-solving; some similar ones are not. If you figure it out, let us know!