Sponsored by: UGA Math Department and UGA Math Club

Team Round / 1 hour / 210 points
November 8, 2014

WITH SOLUTIONS

Problem 1 (Lights out)**.** In the game of "Lights out" there is a collection of lights, some of which are on, and some off. If you touch any light, that light and all of the adjacent lights will change; those that are on will turn off, and those that are off will turn on. You win if you can touch a sequence of lights so that all of the lights are off. Depending on which lights are initially on, this may not be possible.

For this problem, there are 8 lights, located at the corners of a cube, so that each light is adjacent to 3 other lights. How many "winning positions" are there? In other words, for how many initial configurations of on/off lights is it possible to turn off all of the lights? Having all of the lights initially off counts as a winning position (you've already won!).

Answer. 16

Solution. For reference, we label the vertices of the cube as shown.

First, some observations:

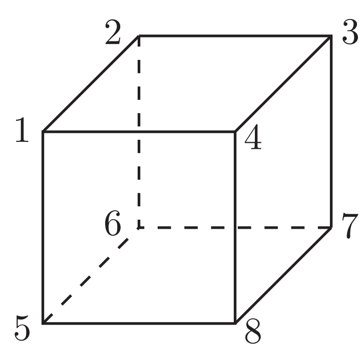

1. The status of the lights can be represented by the 8-dimensional column vector $\mathbf{a} = (a_1, a_2, \ldots, a_8)^T$ with $a_i = 1$ if the ith light is on and $a_i = 0$ if the ith light is off.

2. We can think of the components of this vector as belonging to the field $\mathbb{F}_2$ containing only the two elements 0 and 1; in this field, $1 + 1 = 0$.

3. Touching a light has the effect of adding a vector $\mathbf{v}_i$ to the current status vector. The vector to be added has 1s in the 4 positions corresponding to the vertex touched and its 3 adjacent vertices; the other 4 positions are 0; e.g., $\mathbf{v}_1 = (1, 1, 0, 1, 1, 0, 0, 0)^T$.

4. Since vector addition is commutative ($\mathbf{v}_i + \mathbf{v}_j = \mathbf{v}_j + \mathbf{v}_i$), it follows from observation (3) that the order in which the lights are touched is irrelevant. Moreover, since $\mathbf{v}_i + \mathbf{v}_i = \mathbf{0}$, no light need ever be touched twice.

5. From the last observation in (4), if $(a_1, a_2, \ldots, a_8)^T$ is a "winning position," then you can reach $(0, 0, \ldots, 0)^T$ by touching at most 8 lights; we can represent the lights touched by a column vector $(b_1, \ldots, b_8)^T$, where $b_i = 1$ if the light is touched and $b_i = 0$ otherwise.

6. If you start in position $(0, 0, \ldots, 0)^T$, then that same sequence of light touches represented by $(b_1, \ldots, b_8)^T$ will return you to $(a_1, \ldots, a_8)^T$; i.e., the winning positions are those that can be reached in eight or fewer light touches.

7. Thus, the winning positions are the linear combinations of the vectors $\mathbf{v}_i$ with coefficients from $\mathbb{F}_2$. In other words, they are the vectors that belong to the column space of

$$V = \begin{bmatrix} 1 & 1 & 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 \end{bmatrix}.$$

(Look carefully and you'll see this matrix hidden on this year's T-shirt!)

8. The dimension of the column space (over $\mathbb{F}_2$) is the rank of the matrix, which we claim is 4 (see below). So the number of winning positions is $2^4 = 16$.

To see that the rank is 4 you can simply row reduce the matrix. The 4×4 identity matrix in the lower left corner shows that the rank is at least 4; on the other hand, the row operations

$$\begin{aligned}
R_1 &\mapsto R_1 + (R_5 + R_6 + R_8) = \mathbf{0} \\
R_2 &\mapsto R_2 + (R_5 + R_6 + R_7) = \mathbf{0} \\
R_3 &\mapsto R_3 + (R_6 + R_7 + R_8) = \mathbf{0} \\
R_4 &\mapsto R_4 + (R_5 + R_7 + R_8) = \mathbf{0}
\end{aligned}$$

show that rank ≤ 4.

So what are the winning positions? First notice that

$$\begin{aligned}
\mathbf{v}_5 &= \mathbf{v}_1 + \mathbf{v}_2 + \mathbf{v}_4 \\
\mathbf{v}_6 &= \mathbf{v}_1 + \mathbf{v}_2 + \mathbf{v}_3 \\
\mathbf{v}_7 &= \mathbf{v}_2 + \mathbf{v}_3 + \mathbf{v}_4 \\
\mathbf{v}_8 &= \mathbf{v}_1 + \mathbf{v}_3 + \mathbf{v}_4
\end{aligned}$$

so any position can be reached using only lights $1, 2, 3$, and 4. The 2^4 combinations of these give the different winning positions.

Problem 2 (Counting lattice points)**.** Partition the plane into 1×1 squares using the lines $x = n$ and $y = m$ for all integers m and n. Then draw a circle of radius 100 centered at $(0, 0)$. How many of the 1×1 squares does the circle pass through the interior of? Notice that the circle passes through the interior of the square whose lower left corner is $(0, 99)$, but it does not pass through the interior of the square whose lower left corner is $(0, 100)$.

Answer. 780

Solution. We begin by counting the squares in the first quadrant between the lines $y = x$ and the positive y-axis. In this half of the quadrant, the tangent lines to the circle have slopes in $[-1, 0]$, and so the circle can pass through a square in these three ways:

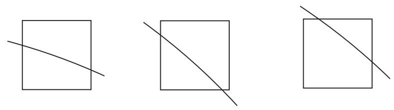

but not this

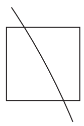

Also note that this

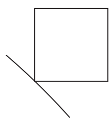

is not passing through the *interior* of the square.

This shows that this piece of the circle intersects every column of boxes once, except for the columns in which the circle crosses a horizontal line. In that case, the circle passes through 2 boxes unless the circle crosses the line at a corner, in which case there is again 1 box in that column.

Notice that the line $y = x$ intersects the circle $x^2 + y^2 = 100^2$ when $x = \sqrt{5000}$, so $70 < x < 71$; i.e., the circle passes through the interior of the square whose lower left corner is $(70, 70)$.

To get from $(0, 100)$ to $(70, 70)$, the circle passes through 70 columns (those whose left sides range from $x = 0$ to $x = 69$) which accounts for 70 squares. The circle also crosses 29 horizontal lines ($y = 99 \to y = 71$). Allowing for two Pythagorean triples $(60, 80, 100)$ and $(28, 96, 100)$ reduces this to 27 new squares. So we have a total of 97 squares in this half-quadrant; adding the 8 half-quadrants and the 4 squares lying on $y = \pm x$, we get $8 \times 97 + 4 = 780$ squares.

Problem 3 (Primitive vertices)**.** Let P_n denote the regular n-gon centered at the origin and having one vertex at $(1, 0)$. We adopt the convention that P_1 consists of the single point $(1, 0)$ and that P_2 consists of the line segment connecting $(-1, 0)$ and $(1, 0)$. A vertex of P_n is called *primitive* if it is not a vertex of P_m for any $m < n$. For example, each vertex of P_3 is primitive except $(1, 0)$. Let C_n denote the center of mass of the primitive vertices of P_n. For how many $n \leq 100$ is C_n located at $(0, 0)$?

Answer. 39

Solution. By de Moivre's theorem, the vertices (x, y) of P_n correspond to the complex numbers $\zeta = x+yi$ satisfying $\zeta^n = 1$. A vertex of P_n is primitive if the corresponding value of ζ does not satisfy $\zeta^m = 1$ for any $m < n$. Such a ζ is called a *primitive* nth root of unity.

For each natural number n, let $U_n = \{\zeta \in \mathbb{C} : \zeta^n = 1\}$, and let $V_n = \{\zeta \in \mathbb{C} : \zeta \text{ is a primitive } n\text{th root of unity}\}$. We are asked for the number of $n \leq 100$ for which $S(n) = 0$, where the function S is defined by

$$S(n) = \sum_{\zeta \in V_n} \zeta.$$

If $\zeta \in U_n$, then $\zeta \in V_m$ for a unique positive integer m dividing n. Conversely, if m divides n and $\zeta \in V_m$, then $\zeta \in U_n$. Thus, U_n is the disjoint union of the sets V_m for m dividing n. Note that as complex polynomials,

$$x^n - 1 = \prod_{\zeta \in U_n} (x - \zeta).$$

If $n > 1$, comparing the coefficients of x^{n-1} yields $\sum_{\zeta \in U_n} \zeta = 0$. Thus,

$$S(1) = 1, \qquad \text{and for all } n > 1, \quad \sum_{m|n} S(m) = 0. \tag{*}$$

Here the sum is over all positive integers m dividing n.

The relation (*) inductively determines all of the values of $S(m)$. Moreover, we can compute enough examples to guess a formula for $S(m)$. To start off, notice that if p is prime, then (*) shows $S(1) + S(p) = 1 + S(p) = 0$, so $S(p) = -1$. On the other hand, if $n = pq$ is a product of two different primes, (*) gives $0 = S(1)+S(p)+S(q)+S(pq) = 1-1-1+S(pq)$, so that $S(pq) = 1$. If $n = p^2$ is the square of a prime, then $0 = S(1)+S(p)+S(p^2) = 1-1+S(p^2)$, and so $S(p^2) = 0$. Continuing to experiment in this way, we arrive at the following guess:

$S(n) = (-1)^k$ if n is a product of k distinct primes, and $S(n) = 0$ otherwise, that is, whenever n has a repeated prime factor.

This guess is correct and we prove it below. Taking its correctness as given, we can complete our analysis of how often $S(n) = 0$. This occurs exactly when n is divisible by p^2 for some prime p. Since $n \leq 100$, the only possibilities for p are $2, 3,$ 5, or 7. Moreover, the only products of two primes whose squares can divide a number $n \leq 100$ are 6 and 10, and no $n \leq 100$ can be divisible by the square of three different primes. So by inclusion-exclusion, the desired count of n is

$$\lfloor 100/2^2 \rfloor + \lfloor 100/3^2 \rfloor + \lfloor 100/5^2 \rfloor + \lfloor 100/7^2 \rfloor - \lfloor 100/6^2 \rfloor - \lfloor 100/10^2 \rfloor$$
$$= 25 + 11 + 4 + 2 - 2 - 1 = 39.$$

Proof of our guess. We use the method of generating functions. Given an infinite sequence of complex numbers $\{a_n\}_{n=1}^{\infty}$, we define the asociated *formal Dirichlet series* by the expression $\sum_{n=1}^{\infty} a_n n^{-s}$. We multiply two formal Dirichlet series by the rule

$$\left(\sum_{n=1}^{\infty} a_n n^{-s}\right)\left(\sum_{n=1}^{\infty} b_n n^{-s}\right) = \sum_{n=1}^{\infty}\Big(\sum_{\substack{d\geq 1,\ e\geq 1\\ de=n}} a_d \cdot b_e\Big)n^{-s}$$
$$= \sum_{n=1}^{\infty}\Big(\sum_{e|n} a_{n/e} \cdot b_e\Big)n^{-s}.$$

In the language of formal Dirichlet series, (*) says that

$$\left(\sum_{n=1}^{\infty} n^{-s}\right)\left(\sum_{n=1}^{\infty} S(n) n^{-s}\right) = 1,$$

and thus

$$\sum_{n=1}^{\infty} S(n) n^{-s} = \frac{1}{\sum_{n=1}^{\infty} n^{-s}}.$$

Because each positive integer can be factored uniquely as a product of primes,

$$\sum_{n=1}^{\infty} n^{-s} = (1 + 2^{-s} + 2^{-2s} + \dots)(1 + 3^{-s} + 3^{-2s} + \dots)(1 + 5^{-s} + 5^{-2s} + \dots)$$
$$= \prod_{p \text{ prime}} \frac{1}{1 - p^{-s}}.$$

(This famous identity due to Leonhard Euler plays an important role in modern number theory.) Taking reciprocals, we find

$$\sum_{n=1}^{\infty} S(n)n^{-s} = \frac{1}{\sum_{n=1}^{\infty} n^{-s}} = \prod_{p \text{ prime}} (1 - p^{-s})$$
$$= \sum_{n=1}^{\infty} \mu(n)n^{-s},$$

where $\mu(n) = (-1)^k$ if n is a product of k distinct primes, and $\mu(n) = 0$ otherwise. Comparing the coefficients of n^{-s} on both sides, we see $S(n) = \mu(n)$ for all values of n. This confirms the guess we made above. □

Authors. Problems and solutions were written by Mo Hendon and Paul Pollack.

Sources. Problem #2 is adapted from a problem in *Math Girls Talk about Trigonometry*, to be released by Bento Books in late 2014.