Solutions to the 69th William Lowell Putnam Mathematical Competition
Saturday, December 6, 2008

Kiran Kedlaya and Lenny Ng

A–1 The function $g(x) = f(x,0)$ works. Substituting $(x,y,z) = (0,0,0)$ into the given functional equation yields $f(0,0) = 0$, whence substituting $(x,y,z) = (x,0,0)$ yields $f(x,0) + f(0,x) = 0$. Finally, substituting $(x,y,z) = (x,y,0)$ yields $f(x,y) = -f(y,0) - f(0,x) = g(x) - g(y)$.

Remark: A similar argument shows that the possible functions g are precisely those of the form $f(x,0) + c$ for some c.

A–2 Barbara wins using one of the following strategies.

First solution: Pair each entry of the first row with the entry directly below it in the second row. If Alan ever writes a number in one of the first two rows, Barbara writes the same number in the other entry in the pair. If Alan writes a number anywhere other than the first two rows, Barbara does likewise. At the end, the resulting matrix will have two identical rows, so its determinant will be zero.

Second solution: (by Manjul Bhargava) Whenever Alan writes a number x in an entry in some row, Barbara writes $-x$ in some other entry in the same row. At the end, the resulting matrix will have all rows summing to zero, so it cannot have full rank.

A–3 We first prove that the process stops. Note first that the product $a_1 \cdots a_n$ remains constant, because $a_j a_k = \gcd(a_j, a_k)\operatorname{lcm}(a_j, a_k)$. Moreover, the last number in the sequence can never decrease, because it is always replaced by its least common multiple with another number. Since it is bounded above (by the product of all of the numbers), the last number must eventually reach its maximum value, after which it remains constant throughout. After this happens, the next-to-last number will never decrease, so it eventually becomes constant, and so on. After finitely many steps, all of the numbers will achieve their final values, so no more steps will be possible. This only happens when a_j divides a_k for all pairs $j < k$.

We next check that there is only one possible final sequence. For p a prime and m a nonnegative integer, we claim that the number of integers in the list divisible by p^m never changes. To see this, suppose we replace a_j, a_k by $\gcd(a_j, a_k), \operatorname{lcm}(a_j, a_k)$. If neither of a_j, a_k is divisible by p^m, then neither of $\gcd(a_j, a_k), \operatorname{lcm}(a_j, a_k)$ is either. If exactly one a_j, a_k is divisible by p^m, then $\operatorname{lcm}(a_j, a_k)$ is divisible by p^m but $\gcd(a_j, a_k)$ is not.

$\gcd(a_j, a_k), \operatorname{lcm}(a_j, a_k)$ are as well.

If we started out with exactly h numbers not divisible by p^m, then in the final sequence $a'_1, \ldots, a'_n$, the numbers $a'_{h+1}, \ldots, a'_n$ are divisible by p^m while the numbers $a'_1, \ldots, a'_h$ are not. Repeating this argument for each pair (p,m) such that p^m divides the initial product $a_1, \ldots, a_n$, we can determine the exact prime factorization of each of $a'_1, \ldots, a'_n$. This proves that the final sequence is unique.

Remark: (by David Savitt and Noam Elkies) Here are two other ways to prove the termination. One is to observe that $\prod_j a_j^j$ is *strictly* increasing at each step, and bounded above by $(a_1 \cdots a_n)^n$. The other is to notice that a_1 is nonincreasing but always positive, so eventually becomes constant; then a_2 is nonincreasing but always positive, and so on.

Reinterpretation: For each p, consider the sequence consisting of the exponents of p in the prime factorizations of $a_1, \ldots, a_n$. At each step, we pick two positions i and j such that the exponents of some prime p are in the wrong order at positions i and j. We then sort these two position into the correct order for every prime p simultaneously.

It is clear that this can only terminate with all sequences being sorted into the correct order. We must still check that the process terminates; however, since all but finitely many of the exponent sequences consist of all zeroes, and each step makes a nontrivial switch in at least one of the other exponent sequences, it is enough to check the case of a single exponent sequence. This can be done as in the first solution.

Remark: Abhinav Kumar suggests the following proof that the process always terminates in at most $\binom{n}{2}$ steps. (This is a variant of the worst-case analysis of the *bubble sort* algorithm.)

Consider the number of pairs (k,l) with $1 \le k < l \le n$ such that a_k does not divide a_l (call these *bad pairs*). At each step, we find one bad pair (i,j) and eliminate it, and we do not touch any pairs that do not involve either i or j. If $i < k < j$, then neither of the pairs (i,k) and (k,j) can become bad, because a_i is replaced by a divisor of itself, while a_j is replaced by a multiple of itself. If $k < i$, then (k,i) can only become a bad pair if a_k divided a_i but not a_j, in which case (k,j) stops being bad. Similarly, if $k > j$, then (i,k) and (j,k) either stay the same or switch status. Hence the number of bad pairs goes down by at least 1 each time; since it is at most $\binom{n}{2}$ to begin with, this is an upper bound for the number of steps.

Remark: This problem is closely related to the classification theorem for finite abelian groups. Namely, if $a_1, \ldots, a_n$ and $a'_1, \ldots, a'_n$ are the sequences obtained at two different steps in the process, then the abelian

groups $\mathbb{Z}/a_1\mathbb{Z} \times \cdots \times \mathbb{Z}/a_n\mathbb{Z}$ and $\mathbb{Z}/a'_1\mathbb{Z} \times \cdots \times \mathbb{Z}/a'_n\mathbb{Z}$ are isomorphic. The final sequence gives a canonical presentation of this group; the terms of this sequence are called the *elementary divisors* or *invariant factors* of the group.

Remark: (by Tom Belulovich) A *lattice* is a partially ordered set L in which for any two $x, y \in L$, there is a unique minimal element z with $z \geq x$ and $z \geq y$, called the *join* and denoted $x \wedge y$, and there is a unique maximal element z with $z \leq x$ and $z \leq y$, called the *meet* and denoted $x \vee y$. In terms of a lattice L, one can pose the following generalization of the given problem. Start with $a_1, \ldots, a_n \in L$. If $i < j$ but $a_i \not\leq a_j$, it is permitted to replace a_i, a_j by $a_i \vee a_j, a_i \wedge a_j$, respectively. The same argument as above shows that this always terminates in at most $\binom{n}{2}$ steps. The question is, under what conditions on the lattice L is the final sequence uniquely determined by the initial sequence?

It turns out that this holds if and only if L is *distributive*, i.e., for any $x, y, z \in L$,

$$x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z).$$

(This is equivalent to the same axiom with the operations interchanged.) For example, if L is a *Boolean algebra*, i.e., the set of subsets of a given set S under inclusion, then $\wedge$ is union, $\vee$ is intersection, and the distributive law holds. Conversely, any finite distributive lattice is contained in a Boolean algebra by a theorem of Birkhoff. The correspondence takes each $x \in L$ to the set of $y \in L$ such that $x \geq y$ and y cannot be written as a join of two elements of $L \setminus \{y\}$. (See for instance Birkhoff, *Lattice Theory*, Amer. Math. Soc., 1967.)

On one hand, if L is distributive, it can be shown that the j-th term of the final sequence is equal to the meet of $a_{i_1} \wedge \cdots \wedge a_{i_j}$ over all sequences $1 \leq i_1 < \cdots < i_j \leq n$. For instance, this can be checked by forming the smallest subset L' of L containing $a_1, \ldots, a_n$ and closed under meet and join, then embedding L' into a Boolean algebra using Birkhoff's theorem, then checking the claim for all Boolean algebras. It can also be checked directly (as suggested by Nghi Nguyen) by showing that for $j = 1, \ldots, n$, the meet of all joins of j-element subsets of $a_1, \ldots, a_n$ is invariant at each step.

On the other hand, a lattice fails to be distributive if and only if it contains five elements $a, b, c, 0, 1$ such that either the only relations among them are implied by

$$1 \geq a, b, c \geq 0$$

(this lattice is sometimes called the *diamond*), or the only relations among them are implied by

$$1 \geq a \geq b \geq 0, \qquad 1 \geq c \geq 0$$

(this lattice is sometimes called the *pentagon*). (For a proof, see the Birkhoff reference given above.) For each of these examples, the initial sequence a, b, c fails to determine the final sequence; for the diamond, we can end up with $0, *, 1$ for any of $* = a, b, c$, whereas for the pentagon we can end up with $0, *, 1$ for any of $* = a, b$.

Consequently, the final sequence is determined by the initial sequence if and only if L is distributive.

A–4 The sum diverges. From the definition, $f(x) = x$ on $[1, e]$, $x \ln x$ on $(e, e^e]$, $x \ln x \ln \ln x$ on $(e^e, e^{e^e}]$, and so forth. It follows that on $[1, \infty)$, f is positive, continuous, and increasing. Thus $\sum_{n=1}^{\infty} \frac{1}{f(n)}$, if it converges, is bounded below by $\int_1^{\infty} \frac{dx}{f(x)}$; it suffices to prove that the integral diverges.

Write $\ln^1 x = \ln x$ and $\ln^k x = \ln(\ln^{k-1} x)$ for $k \geq 2$; similarly write $\exp^1 x = e^x$ and $\exp^k x = e^{\exp^{k-1} x}$. If we write $y = \ln^k x$, then $x = \exp^k y$ and $dx = (\exp^k y)(\exp^{k-1} y) \cdots (\exp^1 y) dy = x(\ln^1 x) \cdots (\ln^{k-1} x) dy$. Now on $[\exp^{k-1} 1, \exp^k 1]$, we have $f(x) = x(\ln^1 x) \cdots (\ln^{k-1} x)$, and thus substituting $y = \ln^k x$ yields

$$\int_{\exp^{k-1} 1}^{\exp^k 1} \frac{dx}{f(x)} = \int_0^1 dy = 1.$$

It follows that $\int_1^{\infty} \frac{dx}{f(x)} = \sum_{k=1}^{\infty} \int_{\exp^{k-1} 1}^{\exp^k 1} \frac{dx}{f(x)}$ diverges, as desired.

A–5 Form the polynomial $P(z) = f(z) + ig(z)$ with complex coefficients. It suffices to prove that P has degree at least $n-1$, as then one of f, g must have degree at least $n-1$.

By replacing $P(z)$ with $aP(z) + b$ for suitable $a, b \in \mathbb{C}$, we can force the regular n-gon to have vertices $\zeta_n, \zeta_n^2, \ldots, \zeta_n^n$ for $\zeta_n = \exp(2\pi i/n)$. It thus suffices to check that there cannot exist a polynomial $P(z)$ of degree at most $n-2$ such that $P(i) = \zeta_n^i$ for $i = 1, \ldots, n$.

We will prove more generally that for any complex number $t \notin \{0, 1\}$, and any integer $m \geq 1$, any polynomial $Q(z)$ for which $Q(i) = t^i$ for $i = 1, \ldots, m$ has degree at least $m-1$. There are several ways to do this.

First solution: If $Q(z)$ has degree d and leading coefficient c, then $R(z) = Q(z+1) - tQ(z)$ has degree d and leading coefficient $(1-t)c$. However, by hypothesis, $R(z)$ has the distinct roots $1, 2, \ldots, m-1$, so we must have $d \geq m-1$.

Second solution: We proceed by induction on m. For the base case $m = 1$, we have $Q(1) = t^1 \neq 0$, so Q must be nonzero, and so its degree is at least 0. Given the assertion for $m-1$, if $Q(i) = t^i$ for $i = 1, \ldots, m$, then the polynomial $R(z) = (t-1)^{-1}(Q(z+1) - Q(z))$ has degree one less than that of Q, and satisfies $R(i) = t^i$ for $i = 1, \ldots, m-1$. Since R must have degree at least $m-2$ by the induction hypothesis, Q must have degree at least $m-1$.

Third solution: We use the method of *finite differences* (as in the second solution) but without induction.

Namely, the $(m-1)$-st finite difference of P evaluated at 1 equals

$$\sum_{j=0}^{m-1}(-1)^j\binom{m-1}{j}Q(m-j)=t(1-t)^{m-1}\neq 0,$$

which is impossible if Q has degree less than $m-1$.

Remark: One can also establish the claim by computing a Vandermonde-type determinant, or by using the Lagrange interpolation formula to compute the leading coefficient of Q.

A–6 For notational convenience, we will interpret the problem as allowing the empty subsequence, whose product is the identity element of the group. To solve the problem in the interpretation where the empty subsequence is not allowed, simply append the identity element to the sequence given by one of the following solutions.

First solution: Put $n=|G|$. We will say that a sequence S *produces* an element $g\in G$ if g occurs as the product of some subsequence of S. Let H be the set of elements produced by the sequence S.

Start with S equal to the empty sequence. If at any point the set $H^{-1}H=\{h_1h_2:h_1^{-1},h_2\in H\}$ fails to be all of G, extend S by appending an element g of G not in $H^{-1}H$. Then $Hg\cap H$ must be empty, otherwise there would be an equation of the form $h_1g=h_2$ with $h_1,h_2\in G$, or $g=h_1^{-1}h_2$, a contradiction. Thus we can extend S by one element and double the size of H.

After $k\leq\log_2 n$ steps, we must obtain a sequence $S=a_1,\ldots,a_k$ for which $H^{-1}H=G$. Then the sequence $a_k^{-1},\ldots,a_1^{-1},a_1,\ldots,a_k$ produces all of G and has length at most $(2/\ln 2)\ln n$.

Second solution:

Put $m=|H|$. We will show that we can append one element g to S so that the resulting sequence of $k+1$ elements will produce at least $2m-m^2/n$ elements of G. To see this, we compute

$$\begin{aligned}\sum_{g\in G}|H\cup Hg| &= \sum_{g\in G}(|H|+|Hg|-|H\cap Hg|)\\ &= 2mn-\sum_{g\in G}|H\cap Hg|\\ &= 2mn-|\{(g,h)\in G^2:h\in H\cap Hg\}|\\ &= 2mn-\sum_{h\in H}|\{g\in G:h\in Hg\}|\\ &= 2mn-\sum_{h\in H}|H^{-1}h|\\ &= 2mn-m^2.\end{aligned}$$

By the pigeonhole principle, we have $|H\cup Hg|\geq 2m-m^2/n$ for some choice of g, as claimed.

In other words, by extending the sequence by one element, we can replace the ratio $s=1-m/n$ (i.e., the fraction of elements of G not generated by S) by a quantity no greater than

$$1-(2m-m^2/n)/n=s^2.$$

We start out with $k=0$ and $s=1-1/n$; after k steps, we have $s\leq(1-1/n)^{2^k}$. It is enough to prove that for some $c>0$, we can always find an integer $k\leq c\ln n$ such that

$$\left(1-\frac{1}{n}\right)^{2^k}<\frac{1}{n},$$

as then we have $n-m<1$ and hence $H=G$.

To obtain this last inequality, put

$$k=\lfloor 2\log_2 n\rfloor<(2/\ln 2)\ln n,$$

so that $2^{k+1}\geq n^2$. From the facts that $\ln n\leq\ln 2+(n-2)/2\leq n/2$ and $\ln(1-1/n)<-1/n$ for all $n\geq 2$, we have

$$2^k\ln\left(1-\frac{1}{n}\right)<-\frac{n^2}{2n}=-\frac{n}{2}<-\ln n,$$

yielding the desired inequality.

Remark: An alternate approach in the second solution is to distinguish betwen the cases of H small (i.e., $m<n^{1/2}$, in which case m can be replaced by a value no less than $2m-1$) and H large. This strategy is used in a number of recent results of Bourgain, Tao, Helfgott, and others on *small doubling* or *small tripling* of subsets of finite groups.

In the second solution, if we avoid the rather weak inequality $\ln n\leq n/2$, we instead get sequences of length $\log_2(n\ln n)=\log_2(n)+\log_2(\ln n)$. This is close to optimal: one cannot use fewer than $\log_2 n$ terms because the number of subsequences must be at least n.

B–1 There are at most two such points. For example, the points $(0,0)$ and $(1,0)$ lie on a circle with center $(1/2,x)$ for any real number x, not necessarily rational.

On the other hand, suppose $P=(a,b),Q=(c,d),R=(e,f)$ are three rational points that lie on a circle. The midpoint M of the side PQ is $((a+c)/2,(b+d)/2)$, which is again rational. Moreover, the slope of the line PQ is $(d-b)/(c-a)$, so the slope of the line through M perpendicular to PQ is $(a-c)/(b-d)$, which is rational or infinite.

Similarly, if N is the midpoint of QR, then N is a rational point and the line through N perpendicular to QR has rational slope. The center of the circle lies on both of these lines, so its coordinates (g,h) satisfy two linear equations with rational coefficients, say $Ag+Bh=C$ and $Dg+Eh=F$. Moreover, these equations have a unique solution. That solution must then be

$$\begin{aligned}g&=(CE-BD)/(AE-BD)\\ h&=(AF-BC)/(AE-BD)\end{aligned}$$

(by elementary algebra, or Cramer's rule), so the center of the circle is rational. This proves the desired result.

Remark: The above solution is deliberately more verbose than is really necessary. A shorter way to say this is that any two distinct rational points determine a *rational line* (a line of the form $ax+by+c=0$ with a,b,c rational), while any two nonparallel rational lines intersect at a rational point. A similar statement holds with the rational numbers replaced by any field.

Remark: A more explicit argument is to show that the equation of the circle through the rational points $(x_1,y_1),(x_2,y_2),(x_3,y_3)$ is

$$0 = \det \begin{pmatrix} x_1^2+y_1^2 & x_1 & y_1 & 1 \\ x_2^2+y_2^2 & x_2 & y_2 & 1 \\ x_3^2+y_3^2 & x_3 & y_3 & 1 \\ x^2+y^2 & x & y & 1 \end{pmatrix}$$

which has the form $a(x^2+y^2)+dx+ey+f=0$ for a,d,e,f rational. The center of this circle is $(-d/(2a),-e/(2a))$, which is again a rational point.

B–2 We claim that $F_n(x) = (\ln x - a_n)x^n/n!$, where $a_n = \sum_{k=1}^n 1/k$. Indeed, temporarily write $G_n(x) = (\ln x - a_n)x^n/n!$ for $x>0$ and $n \geq 1$; then $\lim_{x\to 0} G_n(x) = 0$ and $G_n'(x) = (\ln x - a_n + 1/n)x^{n-1}/(n-1)! = G_{n-1}(x)$, and the claim follows by the Fundamental Theorem of Calculus and induction on n.

Given the claim, we have $F_n(1) = -a_n/n!$ and so we need to evaluate $-\lim_{n\to\infty} \frac{a_n}{\ln n}$. But since the function $1/x$ is strictly decreasing for x positive, $\sum_{k=2}^n 1/k = a_n - 1$ is bounded below by $\int_2^n dx/x = \ln n - \ln 2$ and above by $\int_1^n dx/x = \ln n$. It follows that $\lim_{n\to\infty} \frac{a_n}{\ln n} = 1$, and the desired limit is -1.

B–3 The largest possible radius is $\frac{\sqrt{2}}{2}$. It will be convenient to solve the problem for a hypercube of side length 2 instead, in which case we are trying to show that the largest radius is $\sqrt{2}$.

Choose coordinates so that the interior of the hypercube is the set $H = [-1,1]^4$ in $\mathbb{R}^4$. Let C be a circle centered at the point P. Then C is contained both in H and its reflection across P; these intersect in a rectangular paralellepiped each of whose pairs of opposite faces are at most 2 unit apart. Consequently, if we translate C so that its center moves to the point $O=(0,0,0,0)$ at the center of H, then it remains entirely inside H.

This means that the answer we seek equals the largest possible radius of a circle C contained in H *and centered at* O. Let $v_1 = (v_{11},\dots,v_{14})$ and $v_2 = (v_{21},\dots,v_{24})$ be two points on C lying on radii perpendicular to each other. Then the points of the circle can be expressed as $v_1\cos\theta + v_2\sin\theta$ for $0 \leq \theta < 2\pi$. Then C lies in H if and only if for each i, we have

$$|v_{1i}\cos\theta + v_{2i}\sin\theta| \leq 1 \qquad (0 \leq \theta < 2\pi).$$

In geometric terms, the vector (v_{1i},v_{2i}) in $\mathbb{R}^2$ has dot product at most 1 with every unit vector. Since this holds for the unit vector in the same direction as (v_{1i},v_{2i}), we must have

$$v_{1i}^2 + v_{2i}^2 \leq 1 \qquad (i=1,\dots,4).$$

Conversely, if this holds, then the Cauchy-Schwarz inequality and the above analysis imply that C lies in H.

If r is the radius of C, then

$$\begin{aligned} 2r^2 &= \sum_{i=1}^4 v_{1i}^2 + \sum_{i=1}^4 v_{2i}^2 \\ &= \sum_{i=1}^4 (v_{1i}^2 + v_{2i}^2) \\ &\leq 4, \end{aligned}$$

so $r \leq \sqrt{2}$. Since this is achieved by the circle through $(1,1,0,0)$ and $(0,0,1,1)$, it is the desired maximum.

Remark: One may similarly ask for the radius of the largest k-dimensional ball inside an n-dimensional unit hypercube; the given problem is the case $(n,k)=(4,2)$. Daniel Kane gives the following argument to show that the maximum radius in this case is $\frac{1}{2}\sqrt{\frac{n}{k}}$. (Thanks for Noam Elkies for passing this along.)

We again scale up by a factor of 2, so that we are trying to show that the maximum radius r of a k-dimensional ball contained in the hypercube $[-1,1]^n$ is $\sqrt{\frac{n}{k}}$. Again, there is no loss of generality in centering the ball at the origin. Let $T:\mathbb{R}^k \to \mathbb{R}^n$ be a similitude carrying the unit ball to this embedded k-ball. Then there exists a vector $v_i \in \mathbb{R}^k$ such that for $e_1,\dots,e_n$ the standard basis of $\mathbb{R}^n$, $x\cdot v_i = T(x)\cdot e_i$ for all $x \in \mathbb{R}^k$. The condition of the problem is equivalent to requiring $|v_i| \leq 1$ for all i, while the radius r of the embedded ball is determined by the fact that for all $x \in \mathbb{R}^k$,

$$r^2(x\cdot x) = T(x)\cdot T(x) = \sum_{i=1}^n x\cdot v_i.$$

Let M be the matrix with columns $v_1,\dots,v_k$; then $MM^T = r^2 I_k$, for I_k the $k\times k$ identity matrix. We then have

$$\begin{aligned} kr^2 &= \mathrm{Trace}(r^2 I_k) = \mathrm{Trace}(MM^T) \\ &= \mathrm{Trace}(M^T M) = \sum_{i=1}^n |v_i|^2 \\ &\leq n, \end{aligned}$$

yielding the upper bound $r \leq \sqrt{\frac{n}{k}}$.

To show that this bound is optimal, it is enough to show that one can find an orthogonal projection of $\mathbb{R}^n$ onto $\mathbb{R}^k$ so that the projections of the e_i all have the same norm (one can then rescale to get the desired configuration of $v_1,\dots,v_n$). We construct such a configuration by a "smoothing" argument. Startw with any projection. Let

$w_1, \ldots, w_n$ be the projections of $e_1, \ldots, e_n$. If the desired condition is not achieved, we can choose i, j such that

$$|w_i|^2 < \frac{1}{n}(|w_1|^2 + \cdots + |w_n|^2) < |w_j|^2.$$

By precomposing with a suitable rotation that fixes e_h for $h \neq i, j$, we can vary $|w_i|, |w_j|$ without varying $|w_i|^2 + |w_j|^2$ or $|w_h|$ for $h \neq i, j$. We can thus choose such a rotation to force one of $|w_i|^2, |w_j|^2$ to become equal to $\frac{1}{n}(|w_1|^2 + \cdots + |w_n|^2)$. Repeating at most $n-1$ times gives the desired configuration.

B–4 We use the identity given by Taylor's theorem:

$$h(x+y) = \sum_{i=0}^{\deg(h)} \frac{h^{(i)}(x)}{i!} y^i.$$

In this expression, $h^{(i)}(x)/i!$ is a polynomial in x with integer coefficients, so its value at an integer x is an integer.

For $x = 0, \ldots, p-1$, we deduce that

$$h(x+p) \equiv h(x) + ph'(x) \pmod{p^2}.$$

(This can also be deduced more directly using the binomial theorem.) Since we assumed $h(x)$ and $h(x+p)$ are distinct modulo p^2, we conclude that $h'(x) \not\equiv 0 \pmod{p}$. Since h' is a polynomial with integer coefficients, we have $h'(x) \equiv h'(x+mp) \pmod{p}$ for any integer m, and so $h'(x) \not\equiv 0 \pmod{p}$ for *all* integers x.

Now for $x = 0, \ldots, p^2-1$ and $y = 0, \ldots, p-1$, we write

$$h(x+yp^2) \equiv h(x) + p^2 y h'(x) \pmod{p^3}.$$

Thus $h(x), h(x+p^2), \ldots, h(x+(p-1)p^2)$ run over all of the residue classes modulo p^3 congruent to $h(x)$ modulo p^2. Since the $h(x)$ themselves cover all the residue classes modulo p^2, this proves that $h(0), \ldots, h(p^3-1)$ are distinct modulo p^3.

Remark: More generally, the same proof shows that for any integers $d, e > 1$, h permutes the residue classes modulo p^d if and only if it permutes the residue classes modulo p^e. The argument used in the proof is related to a general result in number theory known as *Hensel's lemma*.

B–5 The functions $f(x) = x+n$ and $f(x) = -x+n$ for any integer n clearly satisfy the condition of the problem; we claim that these are the only possible f.

Let $q = a/b$ be any rational number with $\gcd(a,b) = 1$ and $b > 0$. For n any positive integer, we have

$$\frac{f(\frac{an+1}{bn}) - f(\frac{a}{b})}{\frac{1}{bn}} = bnf\left(\frac{an+1}{bn}\right) - nbf\left(\frac{a}{b}\right)$$

is an integer by the property of f. Since f is differentiable at a/b, the left hand side has a limit. It follows that for sufficiently large n, both sides must be equal to some integer $c = f'(\frac{a}{b})$: $f(\frac{an+1}{bn}) = f(\frac{a}{b}) + \frac{c}{bn}$. Now c cannot be 0, since otherwise $f(\frac{an+1}{bn}) = f(\frac{a}{b})$ for sufficiently large n has denominator b rather than bn. Similarly, $|c|$ cannot be greater than 1: otherwise if we take $n = k|c|$ for k a sufficiently large positive integer, then $f(\frac{a}{b}) + \frac{c}{bn}$ has denominator bk, contradicting the fact that $f(\frac{an+1}{bn})$ has denominator bn. It follows that $c = f'(\frac{a}{b}) = \pm 1$.

Thus the derivative of f at any rational number is ± 1. Since f is continuously differentiable, we conclude that $f'(x) = 1$ for all real x or $f'(x) = -1$ for all real x. Since $f(0)$ must be an integer (a rational number with denominator 1), $f(x) = x+n$ or $f(x) = -x+n$ for some integer n.

Remark: After showing that $f'(q)$ is an integer for each q, one can instead argue that f' is a continuous function from the rationals to the integers, so must be constant. One can then write $f(x) = ax+b$ and check that $b \in \mathbb{Z}$ by evaluation at $a = 0$, and that $a = \pm 1$ by evaluation at $x = 1/a$.

B–6 In all solutions, let $F_{n,k}$ be the number of k-limited permutations of $\{1, \ldots, n\}$.

First solution: (by Jacob Tsimerman) Note that any permutation is k-limited if and only if its inverse is k-limited. Consequently, the number of k-limited permutations of $\{1, \ldots, n\}$ is the same as the number of k-limited involutions (permutations equal to their inverses) of $\{1, \ldots, n\}$.

We use the following fact several times: the number of involutions of $\{1, \ldots, n\}$ is odd if $n = 0, 1$ and even otherwise. This follows from the fact that non-involutions come in pairs, so the number of involutions has the same parity as the number of permutations, namely $n!$.

For $n \leq k+1$, all involutions are k-limited. By the previous paragraph, $F_{n,k}$ is odd for $n = 0, 1$ and even for $n = 2, \ldots, k+1$.

For $n > k+1$, group the k-limited involutions into classes based on their actions on $k+2, \ldots, n$. Note that for C a class and $\sigma \in C$, the set of elements of $A = \{1, \ldots, k+1\}$ which map into A under σ depends only on C, not on σ. Call this set $S(C)$; then the size of C is exactly the number of involutions of $S(C)$. Consequently, $|C|$ is even unless $S(C)$ has at most one element. However, the element 1 cannot map out of A because we are looking at k-limited involutions. Hence if $S(C)$ has one element and $\sigma \in C$, we must have $\sigma(1) = 1$. Since σ is k-limited and $\sigma(2)$ cannot belong to A, we must have $\sigma(2) = k+2$. By induction, for $i = 3, \ldots, k+1$, we must have $\sigma(i) = k+i$.

If $n < 2k+1$, this shows that no class C of odd cardinality can exist, so $F_{n,k}$ must be even. If $n \geq 2k+1$, the classes of odd cardinality are in bijection with k-limited involutions of $\{2k+2, \ldots, n\}$, so $F_{n,k}$ has the same parity as $F_{n-2k-1,k}$. By induction on n, we deduce the desired result.

Second solution: (by Yufei Zhao) Let $M_{n,k}$ be the $n \times n$ matrix with

$$(M_{n,k})_{ij} = \begin{cases} 1 & |i-j| \le k \\ 0 & \text{otherwise.} \end{cases}$$

Write $\det(M_{n,k})$ as the sum over permutations σ of $\{1,\dots,n\}$ of $(M_{n,k})_{1\sigma(1)} \cdots (M_{n,k})_{n\sigma(n)}$ times the signature of σ. Then σ contributes ± 1 to $\det(M_{n,k})$ if σ is k-limited and 0 otherwise. We conclude that

$$\det(M_{n,k}) \equiv F_{n,k} \pmod 2.$$

For the rest of the solution, we interpret $M_{n,k}$ as a matrix over the field of two elements. We compute its determinant using linear algebra modulo 2.

We first show that for $n \ge 2k+1$,

$$F_{n,k} \equiv F_{n-2k-1,k} \pmod 2,$$

provided that we interpret $F_{0,k} = 1$. We do this by computing $\det(M_{n,k})$ using row and column operations. We will verbally describe these operations for general k, while illustrating with the example $k = 3$.

To begin with, $M_{n,k}$ has the following form.

$$\left(\begin{array}{ccccccc|c} 1&1&1&1&0&0&0&\emptyset \\ 1&1&1&1&1&0&0&\emptyset \\ 1&1&1&1&1&1&0&\emptyset \\ 1&1&1&1&1&1&1&\emptyset \\ 0&1&1&1&1&1&1&? \\ 0&0&1&1&1&1&1&? \\ 0&0&0&1&1&1&1&? \\ \hline \emptyset&\emptyset&\emptyset&\emptyset&?&?&?&* \end{array}\right)$$

In this presentation, the first $2k+1$ rows and columns are shown explicitly; the remaining rows and columns are shown in a compressed format. The symbol $\emptyset$ indicates that the unseen entries are all zeroes, while the symbol ? indicates that they are not. The symbol $*$ in the lower right corner represents the matrix $F_{n-2k-1,k}$. We will preserve the unseen structure of the matrix by only adding the first $k+1$ rows or columns to any of the others.

We first add row 1 to each of rows $2,\dots,k+1$.

$$\left(\begin{array}{ccccccc|c} 1&1&1&1&0&0&0&\emptyset \\ 0&0&0&0&1&0&0&\emptyset \\ 0&0&0&0&1&1&0&\emptyset \\ 0&0&0&0&1&1&1&\emptyset \\ 0&1&1&1&1&1&1&? \\ 0&0&1&1&1&1&1&? \\ 0&0&0&1&1&1&1&? \\ \hline \emptyset&\emptyset&\emptyset&\emptyset&?&?&?&* \end{array}\right)$$

We next add column 1 to each of columns $2,\dots,k+1$.

$$\left(\begin{array}{ccccccc|c} 1&0&0&0&0&0&0&\emptyset \\ 0&0&0&0&1&0&0&\emptyset \\ 0&0&0&0&1&1&0&\emptyset \\ 0&0&0&0&1&1&1&\emptyset \\ 0&1&1&1&1&1&1&? \\ 0&0&1&1&1&1&1&? \\ 0&0&0&1&1&1&1&? \\ \hline \emptyset&\emptyset&\emptyset&\emptyset&?&?&?&* \end{array}\right)$$

For $i = 2$, for each of $j = i+1,\dots,2k+1$ for which the $(j,k+i)$-entry is nonzero, add row i to row j.

$$\left(\begin{array}{ccccccc|c} 1&0&0&0&0&0&0&\emptyset \\ 0&0&0&0&1&0&0&\emptyset \\ 0&0&0&0&0&1&0&\emptyset \\ 0&0&0&0&0&1&1&\emptyset \\ 0&1&1&1&0&1&1&? \\ 0&0&1&1&0&1&1&? \\ 0&0&0&1&0&1&1&? \\ \hline \emptyset&\emptyset&\emptyset&\emptyset&\emptyset&?&?&* \end{array}\right)$$

Repeat the previous step for $i = 3,\dots,k+1$ in succession.

$$\left(\begin{array}{ccccccc|c} 1&0&0&0&0&0&0&\emptyset \\ 0&0&0&0&1&0&0&\emptyset \\ 0&0&0&0&0&1&0&\emptyset \\ 0&0&0&0&0&0&1&\emptyset \\ 0&1&1&1&0&0&0&? \\ 0&0&1&1&0&0&0&? \\ 0&0&0&1&0&0&0&? \\ \hline \emptyset&\emptyset&\emptyset&\emptyset&\emptyset&\emptyset&\emptyset&* \end{array}\right)$$

Repeat the two previous steps with the roles of the rows and columns reversed. That is, for $i = 2,\dots,k+1$, for each of $j = i+1,\dots,2k+1$ for which the $(j,k+i)$-entry is nonzero, add row i to row j.

$$\left(\begin{array}{ccccccc|c} 1&0&0&0&0&0&0&\emptyset \\ 0&0&0&0&1&0&0&\emptyset \\ 0&0&0&0&0&1&0&\emptyset \\ 0&0&0&0&0&0&1&\emptyset \\ 0&1&0&0&0&0&0&\emptyset \\ 0&0&1&0&0&0&0&\emptyset \\ 0&0&0&1&0&0&0&\emptyset \\ \hline \emptyset&\emptyset&\emptyset&\emptyset&\emptyset&\emptyset&\emptyset&* \end{array}\right)$$

We now have a block diagonal matrix in which the top left block is a $(2k+1) \times (2k+1)$ matrix with nonzero determinant (it results from reordering the rows of the

identity matrix), the bottom right block is $M_{n-2k-1,k}$, and the other two blocks are zero. We conclude that

$$\det(M_{n,k}) \equiv \det(M_{n-2k-1,k}) \pmod 2,$$

proving the desired congruence.

To prove the desired result, we must now check that $F_{0,k}, F_{1,k}$ are odd and $F_{2,k}, \ldots, F_{2k,k}$ are even. For $n = 0, \ldots, k+1$, the matrix $M_{n,k}$ consists of all ones, so its determinant is 1 if $n = 0, 1$ and 0 otherwise. (Alternatively, we have $F_{n,k} = n!$ for $n = 0, \ldots, k+1$, since every permutation of $\{1, \ldots, n\}$ is k-limited.) For $n = k+2, \ldots, 2k$, observe that rows k and $k+1$ of $M_{n,k}$ both consist of all ones, so $\det(M_{n,k}) = 0$ as desired.

Third solution: (by Tom Belulovich) Define $M_{n,k}$ as in the second solution. We prove $\det(M_{n,k})$ is odd for $n \equiv 0, 1 \pmod{2k+1}$ and even otherwise, by directly determining whether or not $M_{n,k}$ is invertible as a matrix over the field of two elements.

Let r_i denote row i of $M_{n,k}$. We first check that if $n \equiv 2, \ldots, 2k \pmod{2k+1}$, then $M_{n,k}$ is not invertible. In this case, we can find integers $0 \leq a < b \leq k$ such that $n + a + b \equiv 0 \pmod{2k+1}$. Put $j = (n+a+b)/(2k+1)$. We can then write the all-ones vector both as

$$\sum_{i=0}^{j-1} r_{k+1-a+(2k+1)i}$$

and as

$$\sum_{i=0}^{j-1} r_{k+1-b+(2k+1)i}.$$

Hence $M_{n,k}$ is not invertible.

We next check that if $n \equiv 0, 1 \pmod{2k+1}$, then $M_{n,k}$ is invertible. Suppose that $a_1, \ldots, a_n$ are scalars such that $a_1 r_1 + \cdots + a_n r_n$ is the zero vector. The m-th coordinate of this vector equals $a_{m-k} + \cdots + a_{m+k}$, where we regard a_i as zero if $i \notin \{1, \ldots, n\}$. By comparing consecutive coordinates, we obtain

$$a_{m-k} = a_{m+k+1} \qquad (1 \leq m < n).$$

In particular, the a_i repeat with period $2k+1$. Taking $m = 1, \ldots, k$ further yields that

$$a_{k+2} = \cdots = a_{2k+1} = 0$$

while taking $m = n-k, \ldots, n-1$ yields

$$a_{n-2k} = \cdots = a_{n-1-k} = 0.$$

For $n \equiv 0 \pmod{2k+1}$, the latter can be rewritten as

$$a_1 = \cdots = a_k = 0$$

whereas for $n \equiv 1 \pmod{2k+1}$, it can be rewritten as

$$a_2 = \cdots = a_{k+1} = 0.$$

In either case, since we also have

$$a_1 + \cdots + a_{2k+1} = 0$$

from the $(k+1)$-st coordinate, we deduce that all of the a_i must be zero, and so $M_{n,k}$ must be invertible.

Remark: The matrices $M_{n,k}$ are examples of *banded matrices*, which occur frequently in numerical applications of linear algebra. They are also examples of *Toeplitz matrices*.