IMO 2026: Two Proofs of Every Problem

The GPT and Axiom Lean proof developments

July 2026

These notes give the six problem statements and the mathematical content of both formal solutions. Shared reductions are written once. When the developments diverge, both proofs are given. Routine Lean bookkeeping and repetitive algebra are omitted, but every load-bearing identity is displayed.

Problem	Relation	Main difference
1	mostly shared	two termination measures
2	same coordinates	kernel elimination versus certificate
3	partly different	$\mathbb{F}_3$ kernel versus graph coloring
4	substantially different	two constructive cutting strategies
5	essentially shared	one estimate versus two inequalities
6	substantially different	global good set versus stabilization

Sources. The GPT proof files are from junaid-hasan/2026_problems; AxiomProver's files are from AxiomMath/IMO2026.

1. Blackboard gcd and lcm process

Problem

There are 2026 integers greater than 1 on a blackboard. In one move, two entries $m, n > 1$ are replaced by

$$\gcd(m,n), \qquad \frac{\operatorname{lcm}(m,n)}{\gcd(m,n)}.$$

Prove that every play terminates with exactly one entry $M > 1$, and that M is independent of the choices.

For a board B and a prime p, define

$$g_p(B) = \gcd\{\nu_p(a) : a \in B\}.$$

If $x = \nu_p(m)$ and $y = \nu_p(n)$, then a move replaces (x, y) by

$$\big(\min(x,y),\ |x-y|\big).$$

Since $\gcd(\min(x,y), |x-y|) = \gcd(x,y)$, every $g_p(B)$ is invariant. Here zeros among the valuations do not affect the gcd, so $g_p(B) > 0$ exactly when p divides at least one board entry.

Thus, if a terminal board has unique entry $M > 1$, then

$$\nu_p(M) = g_p(B_0) \quad \text{for every prime } p, \qquad M = \prod_{p \mid \prod B_0} p^{g_p(B_0)}. \tag{1}$$

It remains only to prove termination and to rule out a terminal board of all ones.

The GPT proof

Set

$$\rho(B) = \left(\prod_{a \in B} a,\ \#\{a \in B : a > 1\} \right),$$

ordered lexicographically. Let $g = \gcd(m,n)$. The product of the two new entries is mn/g, so the board product is divided by g. If $g > 1$, the first coordinate of ρ decreases. If $g = 1$, the move is $\{m,n\} \mapsto \{1, mn\}$, so the product is unchanged and the second coordinate decreases. Hence ρ strictly decreases in the well-founded order $\mathbb{N}^2_{\text{lex}}$, and every play terminates.

A move cannot replace two numbers > 1 by two ones. If $\gcd(m,n) = 1$, the other output is $mn > 1$; if $g = \gcd(m,n) > 1$, the first output is already greater than 1. Therefore every reachable board contains an entry > 1. A terminal board contains at most one such entry, hence exactly one. Formula (1) proves that its value is determined by the initial board.

The Axiom proof

Let $\Omega(r)$ be the number of prime factors of r, counted with multiplicity, and put

$$\mu(B) = 2027 \sum_{a \in B} \Omega(a) + \#\{a \in B : a > 1\}.$$

For the chosen pair, the total Ω decreases by $\Omega(g)$, where $g = \gcd(m, n)$. If $g = 1$, the Ω-sum is unchanged and the number of entries > 1 decreases by one. If $g > 1$, the first term decreases by at least 2027, while the second term does not increase. Thus $\mu(B)$ strictly decreases in $\mathbb{N}$. Termination follows.

The same observation as above shows that a reachable board never consists entirely of ones. The invariant $g_p(B)$ then gives formula (1), so the terminal value is unique.

Comparison

> The number-theoretic invariant is identical. The only real difference is the descent: GPT uses the board product and a lexicographic tie-breaker; Axiom encodes the same two cases in one weighted integer potential.

2. A coordinate identity in a triangle

Problem

Let ABC be a triangle, and let M, N be the midpoints of AB, AC. Choose K inside $\triangle BMC$ and L inside $\triangle BNC$, with K strictly inside $\triangle LBA$ and L strictly inside $\triangle ACK$. Suppose

$$\angle KBA = \angle ACL, \qquad \angle LBK = \angle LNC, \qquad \angle LCK = \angle BMK.$$

If O is the circumcenter of $\triangle AKL$, prove that $OM = ON$.

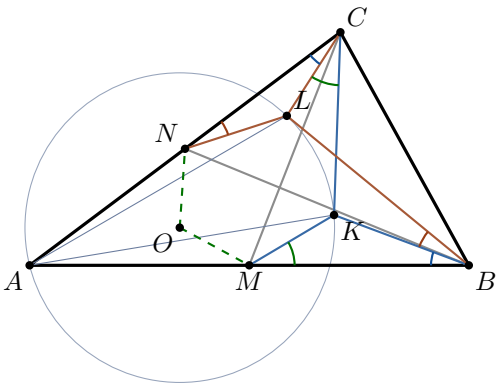

Figure 1: All segments used in the hypotheses are drawn. In particular, CM and BN are visible, and the dashed segments are OM and ON.

Place A at the origin. Write

$$B = b, \quad C = c, \quad a = \langle b, b \rangle, \quad \beta = \langle c, c \rangle, \quad g = \langle b, c \rangle,$$

$$K = xb + yc, \qquad L = zb + wc.$$

The interior conditions imply

$$x, y, z, w > 0, \quad x + y < 1, \quad z + w < 1,$$

and

$$E = w(1 - x) - y(1 - z) > 0, \qquad F = x(1 - w) - z(1 - y) > 0. \tag{2}$$

The last two inequalities follow from the assumptions that K lies inside $\angle LBA$ and L lies inside $\angle ACK$.

Using

$$\cot \angle(u, v) = \frac{\langle u, v \rangle}{|[u, v]|},$$

the three angle equalities become three homogeneous linear equations in (a, β, g). Set

$$P = a(1-x)(1-z) + g\big(w(x-1) + y(z-1)\big) + \beta wy,$$
$$Q = axz + g\big(x(w-1) + z(y-1)\big) + \beta(1-w)(1-y).$$

Then

$$c_1 := az(1-x) - \beta y(1-w) = 0, \tag{3}$$
$$c_2 := 2zP - E\big((2w-1)\beta + 2zg\big) = 0, \tag{4}$$
$$c_3 := 2yQ - F\big((2x-1)a + 2yg\big) = 0. \tag{5}$$

For example, the first angle equality gives

$$z\big((1-x)a - yg\big) = y\big((1-w)\beta - zg\big),$$

and the mixed terms cancel, giving (3). The other two calculations are the same cotangent identity applied at B, N and at C, M.

Put $\Delta = xw - yz$, and let

$$U = ax^2 + 2gxy + \beta y^2 = |K|^2, \qquad V = az^2 + 2gzw + \beta w^2 = |L|^2.$$

Since O is the circumcenter of AKL, $2\langle O, K\rangle = U$ and $2\langle O, L\rangle = V$. If $\Delta \neq 0$, solving $c - b = \lambda K + \mu L$ gives

$$\lambda = -\frac{w+z}{\Delta}, \qquad \mu = \frac{x+y}{\Delta}.$$

A short calculation now shows that $OM = ON$ is equivalent to

$$\mathcal{T} := 2\big(-(w+z)U + (x+y)V\big) - \Delta(\beta - a) = 0. \tag{6}$$

The GPT proof

Let $\mathcal{G}$ be the 3×3 coefficient matrix of (3)–(5) in the variables (a, β, g). Direct expansion gives

$$\det \mathcal{G} = 4yz(1-w)(1-x)(xE - wF), \tag{7}$$

Since $(a, \beta, g) \neq 0$ belongs to $\ker \mathcal{G}$, the determinant vanishes. Hence

$$xE = wF. \tag{8}$$

We also have $\Delta \neq 0$. Indeed, if $\Delta = 0$, then the identity

$$xE - wF = (wz - xy) + (w - x)\Delta$$

and (8) give $wz = xy$. Together with $xw = yz$, positivity yields $w = y$ and then $x = z$. This makes $E = 0$, contradicting (2).

The minor of rows $1, 2$ in columns a, g is $-4wz^2(x-1)^2 \neq 0$, which licenses the elimination below. Solve (3) for β, substitute into (4) to solve for g, and insert both into (6). After cancellation,

$$\mathcal{T} = \frac{a(z+w)}{w(1-w)}(xE - wF) = 0$$

by (8). Therefore $OM = ON$.

The Axiom proof

First, $\Delta \neq 0$ follows directly from the sign conditions. If $xw = yz$, then

$$E = w - y > 0, \qquad F = x - z > 0.$$

Thus $w > y$ and $x > z$, so $xw > yz$, a contradiction.

Axiom finishes with the following polynomial certificate. Let

$$R = 3wx + 2wy - w + 2xz - x + yz - y - z.$$

Expanding both sides gives the identity

$$(x-1)(w-1)\mathcal{T} = -Rc_1 + (x+y)(w-1)c_2 - (w+z)(x-1)c_3. \tag{9}$$

Equations (3)–(5) make the right side zero. Since $x < 1$ and $w < 1$, the factor $(x-1)(w-1)$ is nonzero, so $\mathcal{T} = 0$. By (6), $OM = ON$.

Comparison

> Both are coordinate eliminations. GPT extracts the extra relation $xE = wF$ from the kernel of a 3×3 matrix and then substitutes. Axiom writes the target directly as a linear combination of the three angle equations. The latter is shorter; the former exhibits an additional relation among the positions of K and L.

3. The stick-cutting game

Problem

Liu Bang marks at most n interior points of a unit stick. Xiang Yu then marks at most n further interior points. The stick is cut at all marks. They alternately claim an unclaimed piece, Liu Bang first, each maximizing his own total length. Determine the largest amount c_n that Liu Bang can guarantee.

The answer is

$$c_n = \frac{2^n}{2^{n+1} - 1}.$$

Put $D = 2^{n+1} - 1$. If the final piece lengths in decreasing order are $q_1 \geq \cdots \geq q_m$, optimal play takes them in that order. Therefore Liu Bang receives

$$\frac{1+G}{2}, \qquad G = q_1 - q_2 + q_3 - q_4 + \cdots . \tag{10}$$

We use one elementary fact. If the multiset of pieces is a union of equal pairs together with a residual multiset of total length r, then

$$0 \leq G \leq r. \tag{11}$$

Indeed, equal maximal pairs may be deleted without changing G; if a maximal residual piece is deleted, the remaining alternating sum changes sign, so the new value is at most that piece plus the remaining residual total. This proves (11) by induction.

The common upper bound

Suppose Liu Bang creates $k+1$ initial segments. If $k < n$, Xiang Yu bisects every segment, using at most n marks. All pieces occur in equal pairs, so $G = 0$.

Now let $k = n$. The 2^{n+1} subset sums of the initial lengths lie in $[0, 1]$. Partition $[0, 1]$ into D intervals of length $1/D$. Two subset sums differ by at most $1/D$. After removing their intersection and swapping them if needed, we obtain disjoint sets S, T such that

$$0 \leq \delta = \sum_{i \in S} a_i - \sum_{i \in T} a_i \leq \frac{1}{D}. \tag{12}$$

Concatenate the segments of S into one bar and those of T into another, aligned at the left. Transfer the division points of each bar to the other. Their common initial parts then split into equal pieces, while the longer bar has an overhang of total length δ. Bisect every segment outside $S \cup T$. The number of new marks is at most

$$|T| + (|S| - 1) + (n + 1 - |S| - |T|) = n.$$

Thus all pieces except a residual of total δ occur in equal pairs. By (11), $G \leq \delta \leq 1/D$.

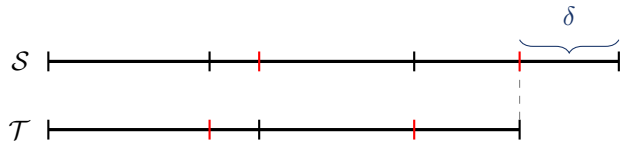

Figure 2: Overlay the two partitions. Everything before the dashed line is paired; the overhang has total length δ.

The common dyadic opening

Liu Bang cuts the stick into lengths

$$a_i = \frac{2^i}{D}, \qquad 0 \le i \le n. \tag{13}$$

After Xiang Yu's cuts, label each final piece by the initial segment containing it. Pair the sorted pieces as $(q_1, q_2), (q_3, q_4), \ldots$. There are at most n pairs because there are at most $2n+1$ pieces.

Both proofs produce a nonzero vector $\varepsilon = (\varepsilon_0, \ldots, \varepsilon_n) \in \{-1, 0, 1\}^{n+1}$ such that the labels u, v in each adjacent pair satisfy $\varepsilon_u + \varepsilon_v = 0$. Then

$$\left| \sum_{i=0}^{n} \varepsilon_i a_i \right| \le G. \tag{14}$$

On the other hand, if j is the largest index with $\varepsilon_j \neq 0$, then

$$\left| \sum_{i=0}^{n} \varepsilon_i 2^i \right| \ge 2^j - (2^j - 1) = 1.$$

Using (13) and (14), $G \ge 1/D$. Together with the upper bound and (10), this gives $c_n = 2^n/D$.

The GPT proof

Work over $\mathbb{F}_3$. For every adjacent pair with labels u, v, impose the linear equation

$$X_u + X_v = 0.$$

There are at most n equations in $n+1$ unknowns, so the homogeneous system has a nonzero solution. Identify the elements of $\mathbb{F}_3$ with $0, 1, -1$. The resulting vector is exactly the required ε, and the common calculation above finishes the proof.

The Axiom proof

Make a multigraph on vertices $0, 1, \ldots, n$, with one edge for every adjacent pair of labels. It has at most n edges and $n+1$ vertices, so some connected component has fewer edges than vertices. That component is a tree. Give its bipartition the values 1 and -1, and give every vertex outside it the value 0. Every edge has opposite endpoint values or two zeros. This supplies the required ε, and the common calculation finishes the proof.

Comparison

> The extremal dyadic division and the subset-sum upper bound are the same. The lower-bound certificate is genuinely different: GPT uses dimension over $\mathbb{F}_3$; Axiom uses a tree component and bipartite coloring.

4. The triangle-cutting game

Problem

Fix $0 < \theta < 180°$. Shan-Yu chooses a triangle. If it has an angle θ, Mulan wins. Otherwise Mulan cuts from a nonvertex boundary point to the opposite vertex, and Shan-Yu keeps one of the two triangles. Determine the values of θ for which Mulan can force a win in finitely many cuts.

The answer is

$$\theta = \frac{180°}{n} \qquad (n \ge 2).$$

We omit degree signs below. If a triangle has angles α, β, γ and Mulan cuts from a point on the side between the β and γ vertices to the opposite vertex, the children have angle triples

$$\{\beta, x, 180 - \beta - x\}, \qquad \{\gamma, 180 - x, x - \gamma\}, \tag{15}$$

where $\gamma < x < 180 - \beta$.

We first record a lemma used by both solutions. If a triangle contains $k\theta$ for a positive integer k, Mulan wins. Induct on k. For $k = 1$ she has already won. For $k \ge 2$, use the angle $k\theta$ as the apex and choose $x = \gamma + \theta$ in (15). One child contains θ and the other contains $(k-1)\theta$.

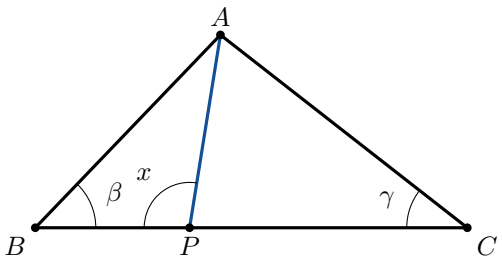

Figure 3: The angle parameter $x = \angle APB$ gives the two triples in (15).

The GPT proof

Assume first that $\theta = 180/n$, $n \geq 2$. Then $\theta \leq 90$.

If a triangle has an angle $a < \theta$, let its other angles be $\beta \leq \beta'$. Use β' as the apex and cut with $x = \theta$. This is legal because $a < \theta < 180 - \beta$. One child contains θ and the other contains $180 - \theta = (n-1)\theta$, so both are winning.

Now take any triangle. If it has an angle below θ, we are done. Otherwise choose an angle $\alpha \geq \theta$ and write

$$k = \left\lfloor \frac{\alpha}{\theta} \right\rfloor, \qquad \delta = k\theta.$$

If $\alpha = \delta$, the common lemma applies. Otherwise $0 < \alpha - \delta < \theta$. Cut with α as apex and $x = \gamma + \delta$. One child contains $\delta = k\theta$ and the other contains $\alpha - \delta < \theta$. Both are winning by the preceding arguments.

For necessity, write $\theta\mathbb{Z}$ for the integer multiples of θ. Induct on a winning strategy tree and prove:

$$\text{if a triangle is winning, then } 180 \in \theta\mathbb{Z} \text{ or one of its angles lies in } \theta\mathbb{Z}. \tag{16}$$

The base case is immediate. For a cut as in (15), assume $180 \notin \theta\mathbb{Z}$ and neither base angle β, γ is a multiple. By induction, the left child contributes x or $180 - \beta - x$, while the right contributes $180 - x$ or $x - \gamma$. The pair $x, 180 - x$ would make 180 a multiple. The other three combinations make, respectively, γ, β, or the apex angle a multiple. This proves (16).

Apply (16) to the equilateral triangle. Either 180 is a multiple of θ, or 60 is; in the latter case $180 = 3 \cdot 60$ is also a multiple. Hence $180 = n\theta$ with $n \geq 2$.

The Axiom proof

Again assume $\theta = 180/n$. Divide all three angles by θ, obtaining positive reals $r_1 + r_2 + r_3 = n$. If one is an integer, the common lemma applies. Suppose none is. After relabeling, there is an integer m with

$$\gamma < m < \gamma + \alpha, \qquad 1 \leq m \leq n - 1. \tag{17}$$

Indeed, if some normalized angle $\alpha > 1$, take $m = \lfloor \gamma \rfloor + 1$. If all three are below 1, then $n < 3$, so $n = 2$ and one may take $m = 1$.

Cut with parameter $x = m\theta$. By (17) the cut is legal. The two children contain $m\theta$ and $(n-m)\theta$, so both are winning by the common lemma.

For necessity, assume $180 \notin \theta\mathbb{Z}$ and call a triangle *safe* if none of its angles lies in $\theta\mathbb{Z}$. Every cut of a safe triangle has a safe child. To see this, use (15). If the left child is not safe, then either $x \in \theta\mathbb{Z}$ or $u = 180 - \beta - x \in \theta\mathbb{Z}$. In the first case neither $180 - x$ nor $x - \gamma$ is a multiple, since that would make 180 or γ a multiple. In the second case neither $180 - x = \beta + u$ nor $x - \gamma = \alpha - u$ is a multiple. Thus the right child is safe.

Shan-Yu can therefore keep a safe child forever, so Mulan cannot have a finite winning strategy from a safe triangle. The equilateral triangle is safe unless 60 is a multiple of θ, which would force 180 to be a multiple. Hence Mulan can guarantee a win only when $180 = n\theta$.

Comparison

> The sufficiency proofs are different. GPT cuts an angle into an integral part and a remainder below θ. Axiom chooses one integer m that places a positive multiple of θ in each child. The necessity arguments are positive and negative forms of the same invariant.

5. A functional inequality

Problem

Find all functions $f : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ such that

$$\sqrt{\frac{x^2 + f(y)^2}{2}} \geq \frac{f(x) + y}{2} \geq \sqrt{xf(y)}$$

for all $x, y > 0$.

The solutions are

$$f(x) = x + c \qquad (c \geq 0).$$

These functions work by QM–AM–GM applied to x and $y + c$.

For the converse, square the inequalities:

$$4xf(y) \leq (f(x) + y)^2 \leq 2(x^2 + f(y)^2). \tag{18}$$

Put $d(t) = f(t) - t$. Setting $x = f(y)$ makes the two outer expressions in (18) equal, so

$$f(f(y)) = 2f(y) - y, \qquad d(f(y)) = d(y). \tag{19}$$

Hence

$$f^k(t) = t + kd(t). \tag{20}$$

Since every iterate is positive, $d(t) \geq 0$.

The GPT proof

From (18) and the identities

$$(x + f(y))^2 - 4xf(y) = (x - f(y))^2,$$
$$2(x^2 + f(y)^2) - (x + f(y))^2 = (x - f(y))^2,$$

we obtain

$$|d(x) - d(y)|\big(2(x + y) + d(x) + d(y)\big) \leq (x - y - d(y))^2. \tag{21}$$

Suppose $d(x_0) = a > 0$ and $d(y_0) = b > 0$. By (20), both forward orbits are arithmetic progressions with steps a and b. Take a far point $Y = y_0 + mb$, and let $X = x_0 + na$ be the first point of the other orbit with $X \geq Y$. Then $0 \leq D = X - Y < a$. Applying (21) to X, Y gives

$$2X|a - b| \leq (D - b)^2 \leq (a + b)^2.$$

Letting $m \to \infty$ forces $a = b$. Therefore d takes values in $\{0, c\}$ for some $c \geq 0$.

If $c > 0$, (21) with $d(u) = c$ and $d(v) = 0$ gives

$$c\big(2(u + v) + c\big) \leq (u - v)^2,$$

so $|u - v| > c$. Thus points at distance at most c have the same d-value. The function d is locally constant on the connected interval $(0, \infty)$, hence constant. Therefore $f(x) = x + c$.

The Axiom proof

Let $d(X) = a$, $d(Y) = b$, and $D = X - Y$. Expanding the two sides of (18) gives

$$(a - D)^2 + 4X(a - b) \geq 0, \tag{22}$$
$$2(b - D)^2 - (a - D)^2 + 4X(b - a) \geq 0. \tag{23}$$

Take far points on the two arithmetic orbits and arrange $0 \leq D < b$. If $a \leq b$, (22) gives

$$4X(b - a) \leq (a - D)^2 \leq (a + b)^2.$$

If $a \geq b$, (23) gives

$$4X(a - b) \leq 2(b - D)^2 \leq 2b^2.$$

Letting $X \to \infty$ shows $a = b$. Again d takes only the values 0 and one positive constant c.

If $d(p) = 0$ and $d(q) = c$, the left inequality in (18) gives

$$4p(q + c) \le (p + q)^2, \qquad 4pc \le (p - q)^2. \tag{24}$$

The set $U = \{d = 0\}$ is open by (24). The set $V = \{d = c\}$ is also open: if $d(q) = c > 0$, then for p sufficiently close to q we have $(p - q)^2 < 4pc$, so (24) rules out $d(p) = 0$ and forces $d(p) = c$. Since $U \cup V = (0, \infty)$ and the interval is connected, one set is empty. Thus d is constant.

Comparison

> The proof architecture is the same: rigidity, arithmetic orbits, equality of positive steps, then connectedness. GPT packages both inequalities into (21); Axiom keeps the two polynomial estimates (22) and (23) separate.

6. A greedy gcd sequence

Problem

Let $a_1, a_2, \ldots > 1$. For each n, let a_{n+1} be the least integer greater than a_n satisfying

$$\gcd(a_{n+1}, a_i) > 1 \qquad (1 \le i \le n).$$

Prove that there are positive integers T, L such that

$$a_{n+T} = a_n + L$$

for every $n \ge 1$.

Call $b \ge 0$ *good* if $\gcd(b, a_i) > 1$ for every i. Every term a_n is good. Conversely, every good $b > a_1$ occurs in the sequence: choose the least N with $a_N \ge b$; if $a_N > b$, then b would have been an admissible smaller choice for a_N. Thus the sequence is the increasing enumeration of the good integers above a_1.

The GPT proof

First observe that translation-periodicity of the good set gives the theorem. Suppose

$$b \text{ is good} \iff b + L \text{ is good} \qquad (b \ge 0). \tag{25}$$

Choose T with $a_{1+T} = a_1 + L$. If $a_{n+T} = a_n + L$, then $a_{n+1} + L$ is a good integer above a_{n+T}, so $a_{n+T+1} \le a_{n+1} + L$. A strict inequality would make $a_{n+T+1} - L$ a good integer strictly between a_n and a_{n+1}, impossible. Hence $a_{n+T+1} = a_{n+1} + L$.

It remains to prove (25). We need four short facts.

1. Any two good integers $b, c > 1$ share a prime. A large power $b^r > a_1$ is good and therefore is a term; since c is good, it has a common prime with b^r, hence with b.
2. If $d > a_1$ is not good, some earlier term $a_i < d$ is coprime to d. Indeed, place d between two consecutive terms and use the minimality of the upper one. For the finitely many non-good $d \le a_1$, choose such witnesses and let C exceed all of them.
3. Every good $x > 1$ has a squarefree good divisor m minimal under divisibility. Take a minimal good divisor of the radical of x.
4. Every prime divisor of such an m is at most C.

For the last claim, use strong induction on m. Let $p \mid m$ and $d = m/p$. Minimality makes d non-good. If $d \le a_1$, choose a good term $a_i \le C$ coprime to d. Since m and a_i are good, they share a prime; it cannot divide d, so it is p. Thus $p \le C$.

If $d > a_1$, choose a term $a_i < d$ coprime to d, and let m' be a minimal good divisor of a_i. Then $m' < m$ and $\gcd(d, m') = 1$. Since m, m' are good, their common prime must be p, so $p \mid m'$. The induction hypothesis gives $p \le C$.

Let

$$L = \prod_{p \le C,\ p \text{ prime}} p.$$

Every good $b > 1$ has a minimal good divisor m with $m \mid b$ and $m \mid L$. Hence $m \mid b + L$, so $b + L$ is good. Conversely, if $b + L$ is good, a minimal good divisor m of $b + L$ divides both $b + L$ and L, hence divides b. The case $b = 1$ is therefore impossible because $m > 1$; in all remaining cases, b is a multiple of the good integer m and is good. The case $b = 0$ also follows directly by taking a minimal good divisor of a_1, which divides L. Thus (25) holds, and the first paragraph completes the proof for every n.

The Axiom proof

For an integer r, let $\mathcal{P}(r)$ be its set of prime divisors. Axiom first proves:

Lemma 6.1 (support reduction)**.** *For every i, there is $j \le i$ such that*

$$\mathcal{P}(a_j) \subseteq \mathcal{P}(a_i)$$

and every prime in $\mathcal{P}(a_j)$ is at most a_1.

Proof. Use strong induction on i. If a_i has no prime above a_1, take $j = i$. Otherwise choose $P > a_1$ dividing a_i, and let b_0 be the P-free part of a_i. A prime $q \mid \gcd(a_i, a_1)$ satisfies $q \le a_1 < P$, so $q \mid b_0$.

For each $\ell < i$, apply induction to a_ℓ and obtain a small-support term whose prime support lies inside $\mathcal{P}(a_\ell)$. It shares a prime with a_i; that prime is at most a_1, so it is not P and therefore divides b_0. Hence every multiple of b_0 shares a prime with every earlier term.

Choose t minimal with $b_0 q^t \ge a_i$, and put $c = b_0 q^{t-1}$. Then $c < a_i$. Also $cq \ge a_i \ge P b_0 \ge Pq$, so $c \ge P > a_1$. If $c > a_{i-1}$, it would be an admissible choice smaller than a_i, contradiction. Choose the least r with $a_r \ge c$. Since c shares a prime with every earlier term, minimality forces $a_r = c$; hence c is an earlier term. Its prime support is $\mathcal{P}(b_0) = \mathcal{P}(a_i) \setminus \{P\}$. Apply the induction hypothesis once more to this earlier term. □

Only finitely many subsets of the primes at most a_1 exist. Choose N so that every small support that ever occurs already occurs among $a_1, \dots, a_N$. By the support reduction lemma,

$$b \text{ is good} \iff \gcd(b, a_i) > 1 \text{ for } 1 \le i \le N. \tag{26}$$

Let $M = \prod_{i=1}^{N} a_i$. Since each a_i divides M, condition (26) depends only on $b \pmod M$. Thus goodness is periodic modulo M.

Among $a_N, a_{N+1}, \dots, a_{N+M}$, two terms have the same residue modulo M. Write them as a_u, a_v, $u < v$, and put

$$T = v - u, \qquad L = a_v - a_u.$$

Then $M \mid L$. The same greedy argument used in the GPT proof gives

$$a_{n+T} = a_n + L \qquad (n \ge u). \tag{27}$$

Finally extend (27) downwards. Suppose it holds at $n + 1$. The good integer $a_n + L$ is a term and lies below $a_{n+1} + L = a_{n+1+T}$, hence $a_n + L \le a_{n+T}$. Therefore $e = a_{n+T} - L$ is good and satisfies

$$a_n \le e < a_{n+1}.$$

The minimality of a_{n+1} forces $e = a_n$. This proves (27) for every $n \ge 1$.

Comparison

> GPT proves that all minimal good divisors use primes from one finite set, then obtains periodicity of the entire good set and therefore pure periodicity at once. Axiom proves a sharper support reduction to primes at most a_1, stabilizes the finite constraints, and first obtains eventual periodicity before propagating it backwards.

Overall assessment

Problems 1, 2, and 5 have the same mathematical backbone in both repositories. Problem 3 has a genuinely different certificate. Problems 4 and 6 show the clearest independent proof design. The concise versions above preserve those distinctions without reproducing formal bookkeeping that is useful in Lean but not in an olympiad solution note.