ALGEBRA

A1. Determine all functions f from the set of non-negative integers to itself such that

$$f(a+b) = f(a) + f(b) + f(c) + f(d),$$

whenever a, b, c, d, are non-negative integers satisfying $2ab = c^2 + d^2$.

Solution 1. (*Ilya Bogdanov*) The required functions are $f(n) = kn^2$, where k is a non-negative integer — these clearly satisfy the condition in the statement.

Conversely, let f be a function satisfying the condition in the statement. Setting $(a, b, c, d) = (n, n, n, n)$ in the functional relation yields $f(2n) = 4f(n)$ for all n. In particular, $f(0) = 0$ and $f(2) = 2k$, where $k = f(1)$.

Setting successively $(a, b, c, d) = (n^2, 1, n, n)$, $(a, b, c, d) = (n^2, 2, 2n, 0)$ and $(a, b, c, d) = (n^2 + 1, 1, n + 1, n - 1)$ in the functional relation yields

$$\begin{aligned} f(n^2+1) &= f(n^2) + k + 2f(n), \\ f(n^2+2) &= f(n^2) + 4k + f(2n) = f(n^2) + 4k + 4f(n), \\ f(n^2+2) &= f(n^2+1) + k + f(n+1) + f(n-1). \end{aligned}$$

Subtraction of the second relation above from the sum of the other two yields $f(n+1) = 2f(n) - f(n-1) + 2k$.

A straightforward induction on n now shows that $f(n) = kn^2$ for all non-negative n, and completes the proof.

Solution 2. As in the first solution, consider a function f satisfying the condition in the statement, and establish that $f(2n) = 4f(n)$ for all n; in particular, $f(0) = 0$.

We now show by induction on m that $f(mn) = m^2 f(n)$ for all n. By the preceding, this is clearly the case if $m = 0, 1, 2$. For the induction step, let $m > 2$.

Since $\sqrt{m} < m$, if m is a square, the conclusion follows by the induction hypothesis: $f(mn) = f(\sqrt{m}(\sqrt{m}\,n)) = mf(\sqrt{m}\,n) = m^2 f(n)$.

Otherwise, use Lagrange's four-square theorem to write $m = A + B$, where A and B are positive integers, each of which is a sum of two squares. Recall that if each factor of a product is a sum of two squares, then so is the product (use standard identities such as $(w^2+x^2)(y^2+z^2) = (wy+xz)^2+(wz-xy)^2$ repeatedly), to write $2AB = (1^2+1^2)AB = C^2+D^2$ for some non-negative integers C and D. Clearly, $A < m$ and $B < m$; $C \le \sqrt{2AB} \le (A+B)/\sqrt{2} = m/\sqrt{2} < m$, and similarly, $D < m$. Setting $a = An$, $b = Bn$, $c = Cn$ and $d = Dn$ in the functional relation and applying the induction hypothesis completes the induction step:

$$\begin{aligned} f(mn) &= f((A+B)n) = f(An) + f(Bn) + f(Cn) + f(Dn) \\ &= A^2 f(n) + B^2 f(n) + C^2 f(n) + D^2 f(n) = (A^2 + B^2 + C^2 + D^2) f(n) \\ &= (A^2 + B^2 + 2AB) f(n) = (A+B)^2 f(n) = m^2 f(n), \end{aligned}$$

for all non-negative integers n.

Consequently, $f(mn) = m^2 f(n)$ for all m and all n. Setting $n = 1$ and $k = f(1)$ concludes the proof.

A2. Let $p > 3$ be a prime number, and let $\mathbb{F}_p$ denote the (finite) set of residue classes modulo p. Let S_d denote the set of 2-variable polynomials $P(x, y)$ with coefficients in $\mathbb{F}_p$, total degree $\leq d$, and satisfying $P(x, y) = P(y, -x - y)$. Show that

$$|S_d| = p^{\lceil (d+1)(d+2)/6 \rceil}.$$

The total degree of a 2-variable polynomial $P(x, y)$ is the largest value of $i + j$ among monomials $x^i y^j$ appearing in P.

Solution. Let T_d denote the set of 3-variable polynomials $Q(x, y, z)$ with coefficients in $\mathbb{F}_p$, total degree $\leq d$, and satisfying $Q(x, y, z) = Q(y, z, x)$. By construction, if $Q(x, y, z) \in T_d$, then $P(x, y) := Q(x, y, -x - y) \in S_d$. We use this construction to relate the sizes of S_d and T_d.

Claim 1. *Every $P(x, y)$ in S_d arises from some $Q_0(x, y, z)$ in T_d as described above.*

Proof. Fix some P in S_d, and notice that $P(x, y) = P(y, -x - y) = P(-x - y, x)$. Recall that $p > 3$ and set $Q_0(x, y, z) = (P(x, y) + P(y, z) + P(z, x))/3$. This Q_0 clearly lies in T_d, and gives rise to P by the above identity for P. The claim follows.

We have seen that every P in S_d arises from some Q_0 in T_d. In fact, we can precisely determine which Q in T_d give rise to P.

Claim 2. *Let P be a polynomial in S_d and let Q_0 be some element of T_d giving rise to P. A polynomial Q in T_d gives rise to P if and only if $Q(x, y, z) = Q_0(x, y, z) + (x + y + z)Q_1(x, y, z)$ for some Q_1 in T_{d-1}.*

Proof. It is clear that any Q of this form gives rise to P (since $x + y + (-x - y) = 0$), so we just need to prove the converse direction. Thus assume Q gives rise to P. Since $Q(x, y, -x - y) = Q_0(x, y, -x - y)$, it follows that $Q - Q_0$ is divisible by[1] $x + y + z$, so we write $Q = Q_0 + (x + y + z)Q_1$. The degree of Q_1 is then $\leq d - 1$ and it inherits the symmetry condition from Q and Q_0, so Q_1 is indeed in T_{d-1}. This concludes the proof of the claim.

By the two preceding claims, every element of S_d arises from exactly $|T_{d-1}|$ elements of T_d, so $|S_d| = |T_d|/|T_{d-1}|$.

We now evaluate $|T_d|$. To this end, notice that a polynomial $Q(x, y, z) = \sum_{i+j+k \leq d} a_{ijk} x^i y^j z^k$ lies in T_d if and only if $a_{ijk} = a_{jki}$ for all i, j, k. Hence, among the $(r + 1)(r + 2)/2$ coefficients a_{ijk} with $i + j + k = r$, we may choose $\lceil (r + 1)(r + 2)/6 \rceil$ arbitrarily, and this determines the other coefficients uniquely. Consequently,

$$|T_d| = p^{\sum_{r=0}^{d} \lceil (r+1)(r+2)/6 \rceil},$$

and the conclusion follows.

Remark. (*Ilya Bogdanov*) Let us reformulate the solution above in more advanced terms.

Clearly, the set S_d is a linear space over $\mathbb{F}_p$. Let C_d be the space of all 3-variable polynomials $Q(x, y, z)$ of total degree at most d satisfying $Q(x, y, z) = Q(y, z, x) = Q(z, x, y)$, and define the linear maps $\phi \colon C_d \to S_d$ and $\psi \colon S_d \to C_d$ by

$$\phi Q(x, y) = Q(x, y, -x - y), \qquad \psi P(x, y, z) = (P(x, y) + P(y, z) + P(z, x))/3.$$

It is readily checked that they are well-defined, and $\phi\psi = \mathrm{id}_{S_d}$, so ϕ is surjective, and $\dim S_d = \dim \operatorname{Im} \varphi = \dim C_d - \dim \operatorname{Ker} \varphi$.

To find $\operatorname{Ker} \varphi$, detect all polynomials Q in C_d vanishing upon the substitution $z \mapsto -x - y$. Viewing Q as a polynomial in z over $\mathbb{F}_p(x, y)$, this is equivalent to Q being divisible by $z + x + y$. Thus $\operatorname{Ker} \varphi = (z + x + y) \cdot C_{d-1}$, and $\dim \operatorname{Ker} \varphi = \dim C_{d-1}$.

Consequently, $\dim S_d = \dim C_d - \dim C_{d-1}$, that is, the dimension of the space of all homogeneous cyclically-symmetric polynomials of total degree d. This can be evaluated in several ways.

[1] To prove this formally, use division-with-remainder to write $Q - Q_0 = (x + y + z)Q_1 + P_1(x, y)$ and observe that the condition forces $P_1 = 0$.

COMBINATORICS

C1. We start with any finite list of distinct positive integers. We may replace any pair n, $n+1$ (not necessarily adjacent in the list) by the single integer $n-2$, now allowing negatives and repeats in the list. We may also replace any pair n, $n+4$ by $n-1$. We may repeat these operations as many times as we wish. Either determine the most negative integer which can appear in a list, or prove that there is no such minimum.

Solution. The minimal value that can be achieved is -3, for instance by

$$1,2,3,4,5 \;\mapsto\; 0,1,4,5 \;\mapsto\; 0,1,2 \;\mapsto\; -1,0 \;\mapsto\; -3.$$

To show that no lesser value is achievable, let α be the root of the polynomial X^3+X^2-1 lying between 0 and 1, and note that α is also a root of the polynomial

$$X^5+X-1=(X^3+X^2-1)(X^2-X+1).$$

Assign the value of α^n to each integer n. By construction, $\alpha^{n+1}+\alpha^n=\alpha^{n-2}$ and $\alpha^{n+4}+\alpha^n=\alpha^{n-1}$, so the given operations do not change the total value of a list. Since the initial total value of the list is *strictly* less than $\alpha+\alpha^2+\cdots=\alpha/(1-\alpha)=\alpha^{-4}$, no integer less than or equal to -4 can ever appear on the list.

C2. A frog trainer places one frog at each vertex of an equilateral triangle ABC of unit side-length. The trainer can make one frog jump over another along the line joining the two, so that the total length of the jump is an even multiple of the distance between the two frogs just before the jump. Let M and N be two points on the rays AB and AC, respectively, emanating from A, such that $AM = AN = \ell$, where ℓ is a positive integer. After a finite number of jumps, the three frogs all lie in the triangle AMN (inside or on the boundary), and no more jumps are performed. Determine the number of final positions the three frogs may reach in the triangle AMN. (During the process, the frogs may leave the triangle AMN; only their final positions are to be in that triangle.)

Solution. The required number is $\binom{(\ell+3)/2}{2}^3$ if ℓ is odd, and $\binom{\ell/2+2}{2}\binom{\ell/2+1}{2}^2$ if ℓ is even.

Fix the origin at A, and consider the x-axis along AB and the y-axis along AC, so that $B = (1, 0)$ and $C = (0, 1)$.

Colour the integral lattice points (points whose coordinates are both integral) as follows: colour *red* all points whose coordinates are both even, so A is red; colour *green* all points whose x-coordinate is odd and y-coordinate is even, so B is green; colour *blue* all points whose x-coordinate is even and y-coordinate is odd, so C is blue; finally, let the other points (whose coordinates are both odd) be all *white*. A triangle with one red vertex, one green vertex and one blue vertex will be called an RGB-triangle. Notice that there are no degenerate RGB-triangles.

Since all vectors of jumps have even coordinates, jumps preserve colours, so at each stage the three frogs are located at the vertices of an RGB-triangle.

We prove that the three frogs can reach any RGB-triangle through a finite chain of successive jumps. In particular, the RGB-triangles in the triangle AMN are all reachable, so their number yields the required number. The triangle AMN contains exactly $\binom{(\ell+3)/2}{2}$ points from each of the first three classes if ℓ is odd; otherwise, it contains exactly $\binom{\ell/2+2}{2}$ red points and exactly $\binom{\ell/2+1}{2}$ points from each of the other three classes. Since there are no degenerate RGB-triangles, this yields a total of $\binom{(\ell+3)/2}{2}^3$ RGB-triangles if ℓ is odd, and $\binom{\ell/2+2}{2}\binom{\ell/2+1}{2}^2$ if ℓ is even.

We now show that any RGB-triangle can be traced back to ABC through a finite number of backward jumps. In what follows, a k-jump is a jump of one frog over another, whose total length is $2k$ times the distance between the two frogs just before the jump. Notice that the inverse of a 1-jump is itself a 1-jump. The backward procedure consists of the three steps described below.

Step 1. *There exists a finite sequence of* 1*-jumps, possibly followed by a single backward jump, from any RGB-triangle to one having a horizontal side of unit length.*

To prove this, label the vertices of an RGB-triangle $A_1 = (x_1, y_1)$, $A_2 = (x_2, y_2)$, $A_3 = (x_3, y_3)$, so that $y_1 \le y_2 \le y_3$, ans call the positive integer $y_3 - y_1$ the y-breadth of the triangle. Since lattice points on horizontal lines are dichromatic, at least one of the two inequalities is strict. If they are both strict, a 1-jump of A_1 across A_2 yields a triangle with a smaller y-breadth. A finite number of such 1-jumps produces eventually an RGB-triangle DEF with one horizontal edge, say DE.

The point E separates exactly one of its lattice neighbours on the line DE, say G, from D. Since colours alternate on the line DE, and D and E have different colours, it follows that D and G have like colours, and the segment DG has an even length, say $2k$. Consequently, D is the image of G under a k-jump of G across E, and the triangle EFG satisfies the desired conditions.

Step 2. *There exists a finite sequence of* 1*-jumps, possibly followed by a single backward jump, transforming any RGB-triangle with one horizontal edge of unit length into another such with (at least) one more edge of unit length.*

To prove this, refer to the triangle EFG, the outcome of step 1. Finitely many alternate 1-jumps of E and G over one another bring the two at some positions H and I, so that H and F have the same x-coordinate. During the process, the horizontal edge of unit length moves rigidly from EG to HI.

At this stage, a backward jump of F across H, as the one described at step 1, brings the former at one of the lattice neighbours of H on the line FH, say J, and the triangle HIJ satisfies the desired conditions.

The triangle HIJ either is equilateral of unit side-length or it can be transformed into one under a further 1-jump of I (or J) across H; this equilateral triangle of unit side-length is, of course, an RGB-triangle.

Step 3. *There exists a finite sequence of* 1*-jumps moving any equilateral RGB-triangle of unit side-length to* ABC*.*

To prove this final step, consider such a triangle, and locate its red vertex at $(2m, 2n)$. It is readily checked that there exist finite sequences of 1-jumps moving this triangle to each of the four translates of ABC by $(2m \pm 2, 2n)$ and $(2m, 2n \pm 2)$, and thence inductively all the way back to ABC through a finite chain of successive 1-jumps. This completes the proof.

Remark. To avoid case consideration one may fix the parity of ℓ.

C3. A set $S = \{s_1, \ldots, s_k\}$ of positive real numbers is *polygonal* if $k \geq 3$ and there is a non-degenerate planar k-gon whose side lengths are exactly s_1, …, s_k; the set S is *multipolygonal* if in every partition of S into two subsets, each of which has at least three elements, exactly one of these two subsets in polygonal. Fix an integer $n \geq 7$.

(a) Does there exist an n-element multipolygonal set, removal of whose maximal element leaves a multipolygonal set?

(b) Is it possible that every $(n-1)$-element subset of an n-element set of positive real numbers be multipolygonal?

Solution. Recall that a necessary and sufficient condition for $k \geq 3$ positive real numbers s_1, …, s_k to be the side lengths of a non-degenerate planar k-gon is that a maximal s_i be less than the sum of the other s_j.

(a) The answer is in the affirmative. Given pairwise distinct positive real numbers ϵ_1, ϵ_2, ϵ_3 less than $1/2$, we show that the sets $S = \{1, 2, 4, \ldots, 2^{n-5}, 2^{n-4}+\epsilon_1, 2^{n-4}+\epsilon_2, 2^{n-4}+\epsilon_3, 2^{n-3}-1/2\}$ and $S \smallsetminus \{2^{n-3} - 1/2\}$ are both multipolygonal.

Split any of the two sets into two subsets each of which has at least three elements, let A be the part containing at least two of the $2^{n-4} + \epsilon_i$, and let B be the other part.

The set A is polygonal since its maximal element is either one of the $2^{n-4} + \epsilon_i$ or $2^{n-3} - 1/2$, each of which is smaller than the sum of other elements in A.

To prove that B is not polygonal, notice that its maximal element is either $2^{n-3} - 1/2$, or one of the $2^{n-4} + \epsilon_i$, or some 2^k, $k \leq n - 5$. In the first case, the sum of all other elements in B is less than $1 + 2 + \cdots + 2^{n-5} + 2^{n-4} + \epsilon_i = 2^{n-3} - 1 + \epsilon_i < 2^{n-3} - 1/2$; in the second case, this sum does not exceed $1 + 2 + \cdots + 2^{n-5} = 2^{n-4} - 1 < 2^{n-4} + \epsilon_i$; and in the third case, this sum is at most $1 + 2 + \cdots + 2^{k-1} = 2^k - 1 < 2^k$. Consequently, B is not polygonal.

(b) The answer is in the negative. Suppose, if possible, that S is an n-element set of positive real numbers, each $(n-1)$-element subset of which is multipolygonal.

Consider an integer k in the range 3, …, $n-4$. The key fact is that if some k-element subset of S is polygonal, then so are all k-element subsets of S, and hence so are all subsets of S of cardinality at least k.

To prove this, let A be a k-element polygonal subset of S and let A' be any other k-element subset of S. Since A' is obtained from A by a finite number of one-element exchanges — i.e., A and A' are joined by a finite chain of k-element subsets $A = A_0$, A_1, …, $A_r = A'$ such that $|A_i \cap A_{i+1}| = k-1$ for all indices i, — we may and will further assume that $|A \cap A'| = k-1$. In this case, let $B = S \smallsetminus (A \cup A')$ and notice that A, B and A', B are both partitions of $(n-1)$-element subsets of S. Since A is polygonal, B is not, so A' is polygonal.

By the preceding, if $3 \leq k \leq (n-1)/2$, then no k-element subset of S is polygonal — otherwise, some $(n-1)$-element subset of S would split into two polygonal sets, one of cardinality k, and the other of cardinality $n-k-1$. In particular, n must be even — otherwise, any $(n-1)$-element subset of S would split into non-polygonal halves, — and the $n/2$-element subsets of S are all polygonal. Hence, no $(n/2-1)$-element subset of S is polygonal.

Finally, label the elements of S decreasingly $s_1 > s_2 > \cdots > s_n$. By the preceding paragraph, neither s_1, s_2, s_3, nor s_3, s_4, …, $s_{n/2+1}$ form a polygonal set, so $s_1 \geq s_2 + s_3$ and $s_3 \geq s_4 + \cdots + s_{n/2+1}$. Consequently, $s_1 \geq s_2 + s_4 + \cdots + s_{n/2+1}$, implying that s_1, s_2, s_4, …, $s_{n/2+1}$ form an $n/2$-element subset of S which is not polygonal — a contradiction.

C4. Prove that a 46-element set of integers contains two distinct doubletons $\{u, v\}$ and $\{x, y\}$ such that $u + v \equiv x + y \pmod{2016}$.

Solution. Suppose that $S \subseteq \mathbb{Z}/2016\mathbb{Z}$ has the property that its pairwise sums of unordered pairs are distinct, and consider the pairwise differences of *ordered* pairs.

The crucial observation is that if a difference $d \neq 0$ occurs twice in S, then these two occurrences must be adjacent in a 3-term arithmetic progression in S (with difference d). Indeed, if $a, a + d, a', a' + d$ are all in S with $a \neq a'$, then $a + (a' + d) = a' + (a + d)$ so (by assumption on S) we must have $a' = a \pm d$, i.e. the ordered pairs $(a, a + d)$ and $(a', a' + d)$ are adjacent in a 3-term arithmetic progression.

We also observe that, excepting the arithmetic progression of common difference 1008, no two 3-term arithmetic progressions can have the same central term, since if $a, a \pm d, a \pm d'$ are all in S with $d, d' \neq 0, 1008$ distinct, then $(a + d') + (a - d') = (a + d) + (a - d)$ would violate our assumption on S.

It follows that the number of 3-term arithmetic progressions in S is at most $|S| + 2$. Now any non-zero difference not occurring in a 3-term arithmetic progression can appear at most once in S, and those which do appear in 3-term arithmetic progressions can appear at most twice, except $\pm 672 = \pm 2016/3$, which can appear three times.

Consequently, the total number of non-zero differences appearing in S is at least

$$|S|(|S| - 1) - (|S| + 2) - 2 = |S|^2 - 2|S| - 4.$$

If $|S| = 46$, this quantity is greater than 2015 — a contradiction.

Remark. Given a positive integer n, the argument in the solution shows that for no set of integers modulo n, whose size exceeds $1 + \sqrt{n + 4}$, is it possible that the pairwise sums be all distinct. This agrees with the Erdős-Turán estimate for Sidon sequences (or Sidon sets). A *Sidon sequence* is a sequence a_0, a_1, a_2, ... of natural numbers all of whose pairwise sums $a_i + a_j$, $i \leq j$, are different. Erdős and Turán showed that, for every positive real number x, the number of elements smaller than x in a Sidon sequence is at most $\sqrt{x} + O(\sqrt[4]{x})$.

GEOMETRY

G1. Two circles, ω_1 and ω_2, centred at O_1 and O_2, respectively, meet at points A and B. A line through B meets ω_1 again at C, and ω_2 again at D. The tangents to ω_1 and ω_2 at C and D, respectively, meet at E, and the line AE meets the circle ω through A, O_1, O_2 again at F. Prove that the length of the segment EF is equal to the diameter of ω.

Solution. Begin by noticing that the lines CO_1 and DO_2 meet at a point P on ω, since $\angle(PO_1, PO_2) = \angle(O_1C, CB) + \angle(BD, DO_2) = \angle(CB, BO_1) + \angle(O_2B, BD) = \angle(O_2B, BO_1) = \angle(O_1A, AO_2)$. In what follows, we consider the case when O_1 and O_2 lie on the segments CP and DP, respectively; other cases are similar.

Since the angles PCE and PDE are both right, and $2\angle ACP = \angle AO_1P = \angle AO_2P = 2\angle ADP$ (the equality in the middle holds on account of P lying on ω), the points A, C, D, E, P all lie on the circle on diameter EP, so FP is a diameter of ω, and it is therefore sufficient to show that $EF = FP$. Finally, since $\angle AFP = \angle AO_1P = 2\angle ACP = 2\angle AEP$ (the first, respectively third, equality holds on account of $APFO_1$, respectively $ACEP$, being inscribed), it follows that the triangle EFP is isosceles with apex at F.

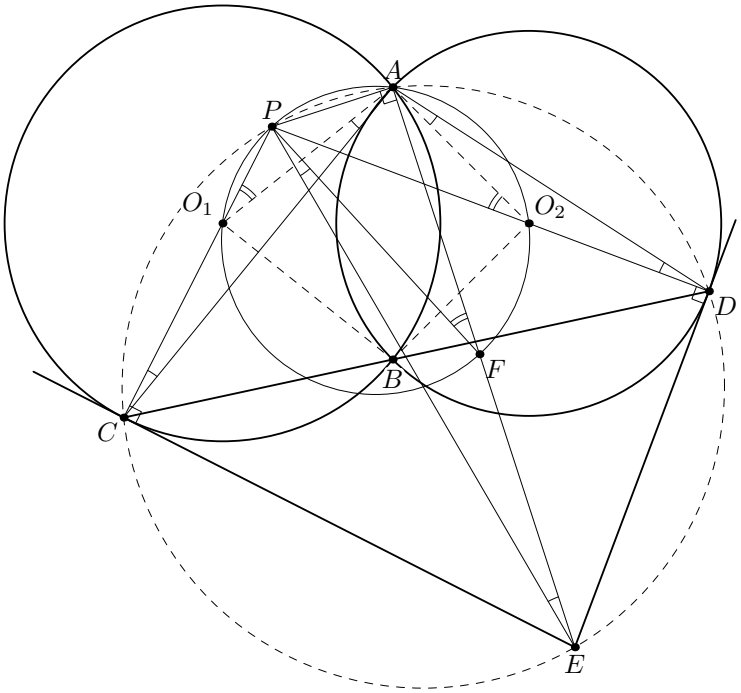

Fig. 1

NUMBER THEORY

N1. Determine all integers $n \geq 3$ whose decimal expansion has less than 20 digits, such that every quadratic non-residue modulo n is a primitive root modulo n.

An integer a is a *quadratic non-residue modulo* n, if there is no integer b such that $a - b^2$ is divisible by n. An integer a is a *primitive root modulo* n, if for every integer b relatively prime to n there is a positive integer k such that $a^k - b$ is divisible by n.

Solution 1. The required numbers are the first five Fermat numbers: $n = 2^{2^k} + 1$, $k = 0, 1, \ldots, 4$. The main thrust is that an integer $n \geq 3$ satisfies the second condition in the statement if and only if n is a Fermat prime. The first five Fermat numbers are indeed prime (and have at most five decimal digits), the sixth is not (Euler showed it divisible by 641), and the higher rank Fermat numbers all have at least 20 decimal digits: if $k \geq 6$, then $F_k = 2^{2^k} + 1 \geq 2^{2^6} + 1 = 2^4 \cdot (2^{10})^6 + 1 > 2^4 \cdot (10^3)^6 + 1 = 16 \cdot 10^{18} + 1 > 10^{19}$.

Since a primitive root modulo n is necessarily a quadratic non-residue modulo n (otherwise, the multiplicative order of the former would be at most $\phi(n)/2$), and for a Fermat prime the number of primitive roots and the number of quadratic non-residues are equal (for a Fermat prime F_k both are 2^k), all Fermat primes satisfy the second condition in the statement.

To prove the converse, we deliberately ignore the characterisation of moduli admitting primitive roots, and first show that if all quadratic non-residues modulo n are relatively prime to n, then n is prime. To this end, let U be the set of integers in the range 0, 1, ..., $n-1$ that are relatively prime to n, so $|U| = \phi(n)$, and let $S = \{0, 1, \ldots, n\} \setminus U$. By hypothesis, all quadratic non-residues modulo n are in U, so all of S are quadratic residues modulo n. Therefore, the function $f\colon S \to S$, $f(k) \equiv k^2 \pmod{n}$ is surjective, and hence injective. Since $f(k) = f(n-k)$ for all k, it follows that $2k \equiv 0 \pmod{n}$ for all k in S, so either $k = 0$ or the additive order of k modulo n is 2. Since $(\mathbb{Z}_n, +)$ has at most two elements of order at most 2, it follows that $n - \phi(n) = n - |U| = |S| \leq 2$. Recall that $n \geq 3$, so $\phi(n)$ is even. If $n - \phi(n) = 2$, then n is even, so $\phi(n) \leq n/2$, in which case $n = 4$ and f is not injective. Consequently, $\phi(n) = n - 1$, and therefore n is prime. Hence the number of quadratic non-residues is $(n-1)/2$, and the number of primitive roots is $\phi(n-1)$.

If, in addition, every quadratic non-residue is a primitive root, then $\phi(n-1) \geq (n-1)/2$. Since n is odd, it follows that $\phi(n-1) \leq (n-1)/2$, so $\phi(n-1) = (n-1)/2$, in which case $n-1$ is a power of 2; that is, n is a Fermat prime: $n = 2^{2^k} + 1$, where k is a non-negative integer.

Solution 2. We prove that an integer $n \geq 3$ satisfying the second condition in the statement must be a Fermat prime. To this end, we show that $\phi(\phi(n)) = \phi(n)/2$, so $\phi(n)$ is a power of 2, and then use the fact that the only moduli admitting primitive roots are 2, 4, p^α and $2p^\alpha$, where p is an odd prime, and α is a positive integer.

To show that $\phi(\phi(n)) = \phi(n)/2$, begin by noticing that at least half of the elements of U (see the notation in Solution 1), are quadratic non-residues — this is because every square $k^2 \pmod{n}$ in U comes from at least two distinct elements of U, namely, k and $n - k$. Consequently, there are at least $|U|/2 = \phi(n)/2$ quadratic non-residues, so, by hypothesis, at least as many primitive roots.

Now, a primitive root is necessarily a quadratic non-residue (otherwise, the multiplicative order of the former would be at most $\phi(n)/2$). Since there are at most $\phi(|U|) = \phi(\phi(n))$ primitive roots, $\phi(\phi(n)) \geq \phi(n)/2$, by the preceding. But $\phi(n)$ is even if $n \geq 3$, so $\phi(\phi(n)) \leq \phi(n)/2$, and consequently $\phi(\phi(n)) = \phi(n)/2$, forcing in turn $\phi(n)$ to be a power of 2.

By the preceding, inspection of moduli admitting primitive roots rules out the first two cases and forces $\alpha = 1$ in the other two, so either $n = p$, in which case n is a Fermat prime, or $n = 2p$. To rule out the latter case, it is sufficient to produce an even quadratic non-residue modulo $2p$. To this end, recall that there are exactly $(p-1)/2$ quadratic residues modulo p, and exactly as many quadratic non-residues modulo p. Since 1 is clearly a quadratic residue modulo p, some quadratic non-residue modulo p must be even. The latter is also a quadratic non-residue modulo $2p$, and the conclusion follows.

Remarks. (1) To avoid Euler's result, we may require n to have less than 10 decimal digits.

(2) (*Ilya Bogdanov*) The integers $n \geq 3$ such that every quadratic non-residue modulo n relatively prime to n is a primitive root modulo n are 4, the Fermat primes, and the doubles of the latter.

In what follows, the obvious case $n = 4$ is left aside.

Consider such an n and let U be the set of integers in the range 0, 1, ..., $n - 1$ that are relatively prime to n. Since $k^2 \equiv (n-k)^2 \pmod{n}$ for all k in U, and $n/2$ is not in U for $n \geq 3$, it follows that U contains at most $|U|/2 = \phi(n)/2$ quadratic residues, and hence at least as many quadratic non-residues. Therefore, there is at least one primitive root a. With reference to the characterisation of moduli admitting primitive roots, $n = p^\alpha$ or $n = 2p^\alpha$, where p is an odd prime and α is a positive integer.

If an odd prime q divided $\phi(n)$, then a^q would be a quadratic non-residue that is not a primitive root — a contradiction.

Consequently, $\phi(n)$ is a power of 2 and, by the preceding, n is a Fermat prime or twice a Fermat prime.

Conversely, if n is a Fermat prime or twice a Fermat prime, and a is a primitive root modulo n, then the multiplicative order of a is a power of 2, so a^k is a primitive root if and only if k is odd, and this is the case if and only if a^k is a quadratic non-residue.

N2. Given a prime p, prove that the sum $\sum_{k=1}^{\lfloor q/p \rfloor} k^{p-1}$ is not divisible by q for all but finitely many primes q.

Solution 1. In what follows, k, m, n are non-negative integers. A straightforward induction on n shows that $X^n = \sum_{k=0}^{n} a_k \binom{X}{k}$ in $\mathbb{Q}[X]$, where the a_k are all integral, $a_n = n!$, and $\binom{X}{k} = X(X-1)\cdots(X-k+1)/k!$. Since $\sum_{j=0}^{m} \binom{j}{k} = \binom{m+1}{k+1}$, it follows that $\sum_{j=0}^{m} j^n = \sum_{k=0}^{n} a_k \binom{m+1}{k+1}$.

Consider the particular case $n = p-1$ and the degree p polynomial with rational coefficients

$$f = \sum_{k=0}^{p-1} a_k \binom{X+1}{k+1} = \frac{1}{p}(X+1)X(X-1)\cdots(X-p+2) + \cdots .$$

Fix r in the range $1, \ldots, p-1$, and write $f(-r/p) = a/b$, where a and b are integers, and $b > 0$ is minimal. Notice that b is divisible by no prime $q > p$.

We now consider a prime $q = mp + r > |a|$, m a positive integer, and show that $\sum_{k=1}^{m} k^{p-1} = f(m)$ is not divisible by q. Notice that $bf(m) - a = b(f(m) - f(-r/p))$ is divisible by q, since $m - (-r/p) = q/p$. Since b is not divisible by q and $|a| < q$, it is sufficient to show that $a \neq 0$. To this end, notice that $p^p \cdot p! f(-r/p) = (p-1)! \prod_{i=0}^{p-1}((i-1)p - r) + p^2 s$ for some integer s. By Wilson's and Fermat's theorems, the right-hand member is an integer congruent to r modulo p, so $f(-r/p)$ is not integral. In particular, $a \neq 0$, and the conclusion follows.

Solution 2. (*Ilya Bogdanov*) Although conceptually the same, this approach seems more down-to-earth.

We first prove the existence of a degree p monic polynomial with rational coefficients $f = X^p + a_{p-1}X^{p-1} + \cdots + a_1 X$ such that p divides the denominator of no a_k (in lowest terms), and

$$f(k) = p(1^{p-1} + 2^{p-1} + \cdots + k^{p-1}) \qquad (*)$$

for all positive integers k. Since the polynomials of the desired form all have a root at 0, $(*)$ is equivalent to $f(k) - f(k-1) = pk^{p-1}$ for all positive integers k. Alternatively, but equivalently, $f(X) - f(X-1) = pX^{p-1}$. Identification of coefficients (or, which is the same, successive formal differentiation) on both sides yields a triangular $(p-1)$-by-$(p-1)$ system of linear equations in the a_k:

$$ka_k + \phi_k(a_{k+1}, \ldots, a_{p-1}) = 0, \quad k = 1, \ldots, p-1,$$

where ϕ_k is linear and has integral coefficients; of course, ϕ_{p-1} is an integral constant. An obvious backwards recursion shows the a_k rational and no denominator divisible by p.

Next, consider an r in the range $1, \ldots, p-1$, and notice that $f(-r/p) = (-r/p)^p + a$, where a is a rational number whose denominator (in lowest terms) is not divisible by p^p. Consequently, $f(-r/p) \neq 0$.

The polynomials f and $pX + r$ are therefore coprime in $\mathbb{Q}[X]$, so there exist a positive integer N and two polynomials ϕ and ψ with integral coefficients such that $f\phi + (pX + r)\psi = N$.

Finally, if $q = mp + r$ is prime, and $f(m)$ is divisible by q, then so is $N = f(m)\phi(m) + (pm + r)\psi(m)$, and the conclusion follows.

Remark. The original version of the problem dealt only with primes q of the form $q = pk + 1$.