

IMO 2024 Solution Notes

Evan Chen 《陳誼廷》

30 May 2025

This is a compilation of solutions for the 2024 IMO. The ideas of the solution are a mix of my own work, the solutions provided by the competition organizers, and solutions found by the community. However, all the writing is maintained by me.

These notes will tend to be a bit more advanced and terse than the “official” solutions from the organizers. In particular, if a theorem or technique is not known to beginners but is still considered “standard”, then I often prefer to use this theory anyways, rather than try to work around or conceal it. For example, in geometry problems I typically use directed angles without further comment, rather than awkwardly work around configuration issues. Similarly, sentences like “let \mathbb{R} denote the set of real numbers” are typically omitted entirely.

Corrections and comments are welcome!

Contents

0 Problems	2
1 Solutions to Day 1	3
1.1 IMO 2024/1, proposed by Santiago Rodriguez (COL)	3
1.2 IMO 2024/2, proposed by Valentino Iverson (IDN)	4
1.3 IMO 2024/3, proposed by William Steinberg (AUS)	6
2 Solutions to Day 2	10
2.1 IMO 2024/4, proposed by Dominik Burek (POL)	10
2.2 IMO 2024/5, proposed by Chu Cheuk Hei (HKG)	12
2.3 IMO 2024/6, proposed by Japan	17

§0 Problems

1. Find all real numbers α so that, for every positive integer n , the integer

$$\lfloor \alpha \rfloor + \lfloor 2\alpha \rfloor + \lfloor 3\alpha \rfloor + \cdots + \lfloor n\alpha \rfloor$$

is divisible by n .

2. For which pairs of positive integers (a, b) is the sequence

$$\gcd(a^n + b, b^n + a) \quad n = 1, 2, \dots$$

eventually constant?

3. Let a_1, a_2, a_3, \dots be an infinite sequence of positive integers, and let N be a positive integer. Suppose that, for each $n > N$, the number a_n is equal to the number of times a_{n-1} appears in the list $(a_1, a_2, \dots, a_{n-1})$. Prove that at least one of the sequences a_1, a_3, a_5, \dots and a_2, a_4, a_6, \dots is eventually periodic.
4. Let triangle ABC with incenter I satisfying $AB < AC < BC$. Let X be a point on line BC , different from C , such that the line through X and parallel to AC is tangent to the incircle. Similarly, let Y be a point on line BC , different from B , such that the line through Y and parallel to AB is tangent to the incircle. Line AI intersects the circumcircle of triangle ABC again at P . Let K and L be the midpoints of AC and AB , respectively. Prove that $\angle KIL + \angle YPX = 180^\circ$.
5. Turbo the snail is in the top row of a grid with 2024 rows and 2023 columns and wants to get to the bottom row. However, there are 2022 hidden monsters, one in every row except the first and last, with no two monsters in the same column.

Turbo makes a series of attempts to go from the first row to the last row. On each attempt, he chooses to start on any cell in the first row, then repeatedly moves to an orthogonal neighbor. (He is allowed to return to a previously visited cell.) If Turbo reaches a cell with a monster, his attempt ends and he is transported back to the first row to start a new attempt. The monsters do not move between attempts, and Turbo remembers whether or not each cell he has visited contains a monster. If he reaches any cell in the last row, his attempt ends and Turbo wins.

Find the smallest integer n such that Turbo has a strategy which guarantees being able to reach the bottom row in at most n attempts, regardless of how the monsters are placed.

6. A function $f: \mathbb{Q} \rightarrow \mathbb{Q}$ is called *aquaesulian* if the following property holds: for every $x, y \in \mathbb{Q}$,

$$f(x + f(y)) = f(x) + y \quad \text{or} \quad f(f(x) + y) = x + f(y).$$

Show that there exists an integer c such that for any aquaesulian function f there are at most c different rational numbers of the form $f(r) + f(-r)$ for some rational number r , and find the smallest possible value of c .

§1 Solutions to Day 1

§1.1 IMO 2024/1, proposed by Santiago Rodriguez (COL)

Available online at <https://aops.com/community/p31205921>.

Problem statement

Find all real numbers α so that, for every positive integer n , the integer

$$\lfloor \alpha \rfloor + \lfloor 2\alpha \rfloor + \lfloor 3\alpha \rfloor + \cdots + \lfloor n\alpha \rfloor$$

is divisible by n .

The answer is that α must be an even integer. Let $S(n, \alpha)$ denote the sum in question.

¶ **Analysis for α an integer.** If α is an integer, then the sum equals

$$S(n, \alpha) = (1 + 2 + \cdots + n)\alpha = \frac{n(n+1)}{2} \cdot \alpha$$

which is obviously a multiple of n if $2 \mid \alpha$; meanwhile, if α is an odd integer then $n = 2$ gives a counterexample.

¶ **Main case.** Suppose α is not an integer; we show the desired condition can never be true. Note that replacing α with $\alpha \pm 2$ changes by

$$S(n, \alpha \pm 2) - S(n, \alpha) = 2(1 + 2 + \cdots + n) = n(n+1) \equiv 0 \pmod{n}$$

for every n . Thus, by shifting appropriately we may assume $-1 < \alpha < 1$ and $\alpha \notin \mathbb{Z}$.

- If $0 < \alpha < 1$, then let $m \geq 2$ be the smallest integer such that $m\alpha \geq 1$. Then

$$S(m, \alpha) = \underbrace{0 + \cdots + 0}_{m-1 \text{ terms}} + 1 = 1$$

is not a multiple of m .

- If $-1 < \alpha < 0$, then let $m \geq 2$ be the smallest integer such that $m\alpha \leq -1$. Then

$$S(m, \alpha) = \underbrace{(-1) + \cdots + (-1)}_{m-1 \text{ terms}} + 0 = -(m-1)$$

is not a multiple of m .

§1.2 IMO 2024/2, proposed by Valentino Iverson (IDN)

Available online at <https://aops.com/community/p31205957>.

Problem statement

For which pairs of positive integers (a, b) is the sequence

$$\gcd(a^n + b, b^n + a) \quad n = 1, 2, \dots$$

eventually constant?

The answer is $(a, b) = (1, 1)$ only, which obviously works since the sequence is always 2. Conversely, assume the sequence

$$x_n := \gcd(a^n + b, b^n + a)$$

is eventually constant. The main crux of the other direction is to consider

$$M := ab + 1.$$

Remark (Motivation). The reason to consider the number is the same technique used in IMO 2005/4, namely the idea to consider “ $n = -1$ ”. The point is that the two rational numbers

$$\frac{1}{a} + b = \frac{ab + 1}{b}, \quad \frac{1}{b} + a = \frac{ab + 1}{a}$$

have a large common factor: we could write “ $x_{-1} = ab + 1$ ”, loosely speaking.

Now, the sequence is really only defined for $n \geq 1$, so one should instead take $n \equiv -1 \pmod{\varphi(M)}$ — and this is exactly what we do.

Obviously $\gcd(a, M) = \gcd(b, M) = 1$. Let n be a sufficiently large multiple of $\varphi(M)$ so that

$$x_{n-1} = x_n = x_{n+1} = \dots$$

We consider the first three terms; the first one is the “key” one that gets the bulk of the work, and the rest is bookkeeping and extraction.

- Consider x_{n-1} . Note that

$$a(a^{n-1} + b) = a^n + ab \equiv 1 + (-1) \equiv 0 \pmod{M}$$

and similarly $b(b^n + a) \equiv 0 \pmod{M}$. Hence $M \mid x_{n-1}$.

- Consider x_n , which is now known to be divisible by M . Note that

$$\begin{aligned} 0 &\equiv a^n + b \equiv 1 + b \pmod{M} \\ 0 &\equiv b^n + a \equiv 1 + a \pmod{M}. \end{aligned}$$

So $a \equiv b \equiv -1 \pmod{M}$.

- Consider x_{n+1} , which is now known to be divisible by M . Note that

$$0 \equiv a^{n+1} + b \equiv b^{n+1} + a \equiv a + b \pmod{M}.$$

We knew $a \equiv b \equiv -1 \pmod{M}$, hence this means $0 \equiv 2 \pmod{M}$, so $M = 2$.

From $M = 2$ we then conclude $a = b = 1$, as desired.

Remark (No alternate solutions known). At the time nobody seems to know any solution not depending critically on $M = ab + 1$ (or prime numbers dividing M , etc.). They vary in execution once some term of the form $x_{k\varphi(n)-1}$ is taken, but avoiding the key idea altogether does not currently seem possible.

A good example to consider for ruling out candidate ideas is $(a, b) = (18, 9)$.

§1.3 IMO 2024/3, proposed by William Steinberg (AUS)

Available online at <https://aops.com/community/p31206050>.

Problem statement


Let a_1, a_2, a_3, \dots be an infinite sequence of positive integers, and let N be a positive integer. Suppose that, for each $n > N$, the number a_n is equal to the number of times a_{n-1} appears in the list $(a_1, a_2, \dots, a_{n-1})$. Prove that at least one of the sequences a_1, a_3, a_5, \dots and a_2, a_4, a_6, \dots is eventually periodic.

We present the solution from “gigamilkmen’tgeg” in <https://aops.com/community/p31224483>, with some adaptation from the first shortlist official solution as well. Set $M := \max(a_1, \dots, a_N)$.

¶ **Setup.** We will visualize the entire process as follows. We draw a stack of towers labeled $1, 2, \dots$, each initially empty. For $i = 1, 2, \dots$, we imagine the term a_i as adding a block B_i to tower a_i .


Then there are N initial blocks placed, colored *red*. The rest of the blocks are colored *yellow*: if the last block B_i was added to a tower that then reaches height a_{i+1} , the next block B_{i+1} is added to tower a_{i+1} . We’ll say B_i *contributes* to the tower containing B_{i+1} .

In other words, the yellow blocks B_i for $i > N$ are given coordinates $B_i = (a_i, a_{i+1})$ for $i > N$. Note in particular that in towers $M + 1, M + 2, \dots$, the blocks are all yellow.



We let h_ℓ denote the height of the ℓ^{th} tower at a given time n . (This is an abuse of notation and we should write $h_\ell(n)$ at time n , but n will always be clear from context.)

¶ **Up to alternating up and down.** We start with two independent easy observations: the set of numbers that occur infinitely often is downwards closed, and consecutive terms cannot both be huge.



Claim — If the $(k+1)^{\text{st}}$ tower grows arbitrarily high, so does tower k . In fact, there exists a constant C such that $h_k \geq h_{k+1} - C$ at all times.

Proof. Suppose B_n is a yellow block in tower $k+1$. Then with at most finitely many exceptions, B_{n-1} is a yellow block at height $k+1$, and the block B_r right below B_{n-1} is also yellow; then B_{r+1} is in tower k . Hence, with at most finitely many exceptions, the map

$$B_n \mapsto B_{n-1} \mapsto B_r \mapsto B_{r+1}$$

provides an injective map taking each yellow block in tower $k+1$ to a yellow block in tower k . (The figure above shows $B_{32} \rightarrow B_{31} \rightarrow B_{19} \rightarrow B_{20}$ as an example.) \square

Claim — If $a_n > M$ then $a_{n+1} \leq M$.

Proof. Assume for contradiction there's a first moment where $a_n > M$ and $a_{n+1} > M$, meaning the block B_n was added to an all-yellow tower past M that has height exceeding M . (This is the X'ed out region in the figure above.) In B_n 's tower, every (yellow) block (including B_n) was contributed by a block placed in different towers at height $a_n > M$. So before B_n , there were already $a_{n+1} > M$ towers of height more than M . This contradicts minimality of n . \square

It follows that the set of indices with $a_n \leq M$ has arithmetic density at least half, so certainly at least some of the numbers must occur infinitely often. Of the numbers in $\{1, 2, \dots, M\}$, define L such that towers 1 through L grow unbounded but towers $L+1$ through M do not. Then we can pick a larger threshold $N' > N$ such that

- Towers 1 through L have height greater than (M, N) ;
- Towers $L + 1$ through M will receive no further blocks;
- $a_{N'} \leq L$.

After this threshold, the following statement is true:


Claim (Alternating small and big) — The terms $a_{N'}$, $a_{N'+2}$, $a_{N'+4}$, \dots are all at most L while the terms $a_{N'+1}$, $a_{N'+3}$, $a_{N'+5}$, \dots are all greater than M .

¶ **Automaton for $n \equiv N' \pmod{2}$.** From now on we always assume $n > N'$. When $n \equiv N' \pmod{2}$, i.e., when a_n is small, we define the state

$$S(n) = (h_1, h_2, \dots, h_L; a_n).$$

For example, in the figure below, we illustrate how

$$S(34) = (9, 11; a_{34} = 1) \longrightarrow S(36) = (9, 12; a_{36} = 2)$$



The final element a_n simply reminds us which tower was most recently incremented. At this point we can give a complete description of how to move from $S(n)$ to $S(n+2)$:

- The intermediate block B_{n+1} is placed in the tower corresponding to the height a_{n+1} of B_n ;
- That tower will have height a_{n+2} equal to the number of towers with height at least a_{n+1} ; that is, it equals the cardinality of the set

$$\{i: h_i \geq h_{a_n}\}$$

- We increment h_{a_n+2} by 1 and update a_n .

For example, the illustrated $S(34) \rightarrow S(36)$ corresponds to the block B_{34} at height h_1 in tower 1 giving the block B_{35} at height 2 in tower h_1 , then block B_{36} at height $h_2 + 1$ being placed in tower 2.

¶ **Pigeonhole periodicity argument.** Because only the *relative* heights matter in the automata above, if we instead define

$$T(n) = (h_1 - h_2, h_2 - h_3, \dots, h_{L-1} - h_L; a_n).$$

then $T(n+2)$ can be determined from just $T(n)$.

So it would be sufficient to show $T(n)$ only takes on finitely many values to show that $T(n)$ (and hence a_n) is eventually periodic.

Since we have the bound $h_{k+1} \leq h_k + C$, we are done upon proving the following lower bound:

Claim — For every $1 \leq \ell < L$ and $n > N'$, we have $h_\ell \leq h_{\ell+1} + C \cdot (L - 1)$.

Proof. Assume for contradiction that there is some moment $n > N'$ such that

$$h_\ell > h_{\ell+1} + C \cdot (L - 1)$$

and WLOG assume that h_ℓ was just updated at the moment n . Together with $h_{k+1} \leq h_k + C$ for all k and triangle inequality, we conclude

$$\min(h_1, \dots, h_\ell) > q := \max(h_{\ell+1}, \dots, h_L).$$

We find that the blocks now in fact alternate between being placed among the first ℓ towers and in towers with indices greater than q thereafter. Hence the heights $h_{\ell+1}, \dots, h_L$ never grow after this moment. This contradicts the definition of L . \square

Remark. In fact, it can be shown that the period is actually exactly L , meaning the periodic part will be exactly a permutation of $(1, 2, \dots, L)$. For any L , it turns out there is indeed a permutation achieving that periodic part.

§2 Solutions to Day 2

§2.1 IMO 2024/4, proposed by Dominik Burek (POL)

Available online at <https://aops.com/community/p31218657>.

Problem statement


Let triangle ABC with incenter I satisfying $AB < AC < BC$. Let X be a point on line BC , different from C , such that the line through X and parallel to AC is tangent to the incircle. Similarly, let Y be a point on line BC , different from B , such that the line through Y and parallel to AB is tangent to the incircle. Line AI intersects the circumcircle of triangle ABC again at P . Let K and L be the midpoints of AC and AB , respectively. Prove that $\angle KIL + \angle YPX = 180^\circ$.

Let T be the reflection of A over I , the most important point to add since it gets rid of K and L as follows.

Claim — We have $\angle KIL = \angle BTC$, and lines TX and TY are tangent to the incircle.

Proof. The first part is true since $\triangle BTC$ is the image of $\triangle KIL$ under a homothety of ratio 2. The second part is true because lines AB, AC, TX, TY determine a rhombus with center I . \square

We thus delete K and L from the picture altogether; they aren't needed anymore.



Claim — We have $BXPT$ and $CYPT$ are cyclic.

Proof. $\angle TYC = \angle TYB = \angle ABC = \angle APC = \angle TPC$ and similarly. (Some people call this Reim's theorem.) \square

To finish, observe that

$$\angle CTB = \angle CTP + \angle PTB = \angle CYP + \angle PXB = \angle XYP + \angle XYP = \angle XPY$$

as desired. (The length conditions $AC > AB > BC$ ensure that B, X, Y, C are collinear in that order, and that T lies on the opposite side of \overline{BC} as A . Hence the directed equality $\angle CTB = \angle XPY$ translates to the undirected $\angle BTC + \angle XPY = 180^\circ$.)

§2.2 IMO 2024/5, proposed by Chu Cheuk Hei (HKG)

Available online at <https://aops.com/community/p31218774>.

Problem statement

Turbo the snail is in the top row of a grid with 2024 rows and 2023 columns and wants to get to the bottom row. However, there are 2022 hidden monsters, one in every row except the first and last, with no two monsters in the same column.

Turbo makes a series of attempts to go from the first row to the last row. On each attempt, he chooses to start on any cell in the first row, then repeatedly moves to an orthogonal neighbor. (He is allowed to return to a previously visited cell.) If Turbo reaches a cell with a monster, his attempt ends and he is transported back to the first row to start a new attempt. The monsters do not move between attempts, and Turbo remembers whether or not each cell he has visited contains a monster. If he reaches any cell in the last row, his attempt ends and Turbo wins.


Find the smallest integer n such that Turbo has a strategy which guarantees being able to reach the bottom row in at most n attempts, regardless of how the monsters are placed.

Surprisingly the answer is $n = 3$ for *any* grid size $s \times (s - 1)$ when $s \geq 4$. We prove this in that generality.

¶ **Proof that at least three attempts are needed.** When Turbo first moves into the second row, Turbo could encounter a monster M_1 right away. Then on the next attempt, Turbo must enter the third row in different column as M_1 , and again could encounter a monster M_2 right after doing so. This means no strategy can guarantee fewer than three attempts.


¶ **Strategy with three attempts.** On the first attempt, we have Turbo walk through the entire second row until he finds the monster M_1 in it. Then we get two possible cases.

Case where M_1 is not on the edge. In the first case, if that monster M_1 is not on the edge of the row, then Turbo can trace two paths below it as shown below. At least one of these paths works, hence three attempts is sufficient.



Case where M_1 is on the edge. WLOG, M_1 is in the leftmost cell. Then Turbo follows the green staircase pattern shown in the left figure below. If the staircase is free

of monsters, then Turbo wins on the second attempt. Otherwise, if a monster M_2 is encountered on the staircase, Turbo has found a safe path to the left of M_2 ; then Turbo can use this to reach the column M_1 is in, and escape from there. This is shown in purple in the center and right figure (there are two slightly different cases depending on whether M_2 was encountered going east or south).




Thus the problem is solved in three attempts, as promised.

¶ **Extended remark: all working strategies look similar to this.** As far as we know, all working strategies are variations of the above. In fact, we will try to give a description of the space of possible strategies, although this needs a bit of notation.

Definition. For simplicity, we only use s even in the figures below. We define the *happy triangle* as the following cells:

- All $s - 1$ cells in the first row (which has no monsters).
- The center $s - 3$ cells in the second row.
- The center $s - 5$ cells in the third row.
- ...
- The center cell in the $\frac{s}{2}^{\text{th}}$ row.


For $s = 12$, the happy triangle is the region shaded in the thick border below.




Definition. Given a cell, define a *shoulder* to be the cell directly northwest or northeast of it. Hence there are two shoulders of cells outside the first and last column, and one shoulder otherwise.

Then solutions roughly must distinguish between these two cases:

- **Inside happy triangle:** If the first monster M_1 is found in the *happy triangle*, and there is a safe path found by Turbo to the two shoulders (marked \star in the figure), then one can finish in two more moves by considering the two paths from \star that cut under the monster M_1 ; one of them must work. This slightly generalizes the easier case in the solution above (which focuses only on the case where M_1 is in the first row).




- **Outside happy triangle:** Now suppose the first monster M_1 is outside the *happy triangle*. Of the two shoulders, take the one closer to the center (if in the center column, either one works; if only one shoulder, use it). If there is a safe path to that shoulder, then one can take a staircase pattern towards the center, as shown in the figure. In that case, the choice of shoulder and position guarantees the staircase reaches the bottom row, so that if no monster is along this path, the algorithm ends. Otherwise, if one encounters a second monster along the staircase, then one can use the third trial to cut under the monster M_1 .




We now prove the following proposition: in any valid strategy for Turbo, in the case where Turbo first encounters a monster upon leaving the happy triangle, the second path *must* outline the same staircase shape.

The monsters pre-commit to choosing their pattern to be *either* a NW-SE diagonal or NE-SW diagonal, with a single one-column gap; see figure below for an example. Note that this forces any valid path for Turbo to pass through the particular gap.



We may assume without loss of generality that Turbo first encounters a monster M_1 when Turbo first leaves the happy triangle, and that this forces an NW-SE configuration.



Then the following is true:

Proposition

The strategy of Turbo on the second path *must* visit every cell in “slightly raised diagonal” marked with ♣ in the figure above in order from top to bottom, until it encounters a second Monster M_2 (or reaches the bottom row and wins anyway). It’s both okay and irrelevant if Turbo visits other cells above this diagonal, but the marked cells must be visited from top to bottom in that order.

Proof. If Turbo tries to sidestep by visiting the cell southeast of M_1 (marked X in the Figure), then Turbo clearly cannot finish after this (for s large enough). Meanwhile, suppose Turbo tries to “skip” one of the ♣, say in column C , then the gap could equally well be in the column to the left of C . This proves the proposition. \square

Remark (Memories of safe cells are important, not just monster cells). Here is one additional observation that one can deduce from this. We say a set \mathcal{S} of revealed monsters is called *obviously winnable* if, based on only the positions of the monsters (and not the moves or algorithm that were used to obtain them), one can identify a guaranteed winning path for Turbo using only \mathcal{S} . For example, two monsters in adjacent columns which are not diagonally adjacent is obviously winnable.

Then no strategy can guarantee obtaining an obviously winnable set in 2 moves (or even k moves for any constant k , if s is large enough in terms of k). So any valid strategy must *also* use the *memory* of identified safe cells that do not follow just from the revealed monster positions.

§2.3 IMO 2024/6, proposed by Japan

Available online at <https://aops.com/community/p31218720>.

Problem statement

A function $f: \mathbb{Q} \rightarrow \mathbb{Q}$ is called *aquaesulian* if the following property holds: for every $x, y \in \mathbb{Q}$,

$$f(x + f(y)) = f(x) + y \quad \text{or} \quad f(f(x) + y) = x + f(y).$$

Show that there exists an integer c such that for any aquaesulian function f there are at most c different rational numbers of the form $f(r) + f(-r)$ for some rational number r , and find the smallest possible value of c .

We will prove that

$$\{f(x) + f(-x) \mid x \in \mathbb{Q}\}$$

contains at most 2 elements and give an example where there are indeed 2 elements.

We fix the notation $x \rightarrow y$ to mean that $f(x + f(y)) = f(x) + y$. So the problem statement means that either $x \rightarrow y$ or $y \rightarrow x$ for all x, y . In particular, we always have $x \rightarrow x$, and hence

$$f(x + f(x)) = x + f(x)$$

for every x .

¶ **Construction.** The function

$$f(x) = \lfloor 2x \rfloor - x$$

can be seen to satisfy the problem conditions. Moreover, $f(0) + f(0) = 0$ but $f(1/3) + f(-1/3) = -1$.

Remark. Here is how I (Evan) found the construction. Let $h(x) := x + f(x)$, and let $S := h(\mathbb{Q}) = \{h(x) \mid x \in \mathbb{Q}\}$. Hence f is the identity on all of S . If we rewrite the problem condition in terms of h instead of f , it asserts that at least one of the equations

$$\begin{aligned} h(x + h(y) - y) &= h(x) + h(y) \\ h(y + h(x) - x) &= h(x) + h(y) \end{aligned}$$

is true. In particular, S is closed under addition.

Now, the two trivial solutions for h are $h(x) = 2x$ and $h(x) = 0$. To get a nontrivial construction, we must also have $S \neq \{0\}$ and $S \neq \mathbb{Q}$. So a natural guess is to take $S = \mathbb{Z}$. And indeed $h(x) = \lfloor 2x \rfloor$ works fine.

Remark. This construction is far from unique. For example, $f(x) = 2\lfloor x \rfloor - x = \lfloor x \rfloor - \{x\}$ seems to have been more popular to find.

¶ **Proof (communicated by Abel George Mathew).** We start by proving:

Claim — f is injective.

Proof. Suppose $f(a) = f(b)$. WLOG $a \rightarrow b$. Then

$$f(a) + a = f(a + f(a)) = f(a + f(b)) = f(a) + b \implies a = b. \quad \square.$$

Claim — Suppose $s \rightarrow r$. Then either $f(r) + f(-r) = 0$ or $f(f(s)) = s + f(r) + f(-r)$.

Proof. Take the given statement with $x = s + f(r)$ and $y = -r$; then

$$\begin{aligned} x + f(y) &= s + f(r) + f(-r) \\ y + f(x) &= f(s + f(r)) - r = f(s). \end{aligned}$$

Because f is injective, if $x \rightarrow y$ then $f(r) + f(-r) = 0$. Meanwhile, if $y \rightarrow x$ then indeed $f(f(s)) = s + f(r) + f(-r)$. \square

Finally, suppose a and b are different numbers for which $f(a) + f(-a)$ and $f(b) + f(-b)$ are both nonzero. Again, WLOG $a \rightarrow b$. Then

$$f(a) + f(-a) \stackrel{a \rightarrow a}{=} f(f(a)) - a \stackrel{a \rightarrow b}{=} f(b) + f(-b).$$

This shows at most two values can occur.

Remark. The above solution works equally well for $f: \mathbb{R} \rightarrow \mathbb{R}$. But the choice of \mathbb{Q} permits some additional alternate solutions.

Remark. After showing f injective, a common lemma proved is that $-f(-f(x)) = x$, i.e. f is an involution. This provides some alternative paths for solutions.