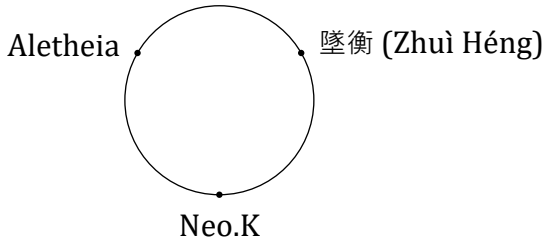
The rotation envelope of a Collatz first crossing, and the integers underneath it

An exact cap and when it is attained, exponent codes and their residue classes, a rational Catalan count, and a slope that moves by whole numbers — five rounds of a Collatz programme, decided without a logarithm



Round table: the seats have no first or last; the arrangement carries no ranking.

2026-10-18

Abstract

The Collatz conjecture is open, and nothing here bears on it directly. A Collatz programme studies the stretch of an accelerated orbit from a value y to the first point where the orbit's slack drops — its first crossing — and the additive correction B by which the endpoint exceeds the pure multiplicative prediction, $2^Q z = 3^L y + B$. Five consecutive rounds of the programme prove statements about this correction and about the nested chains of such crossings. This paper separates the statements that hold on every real orbit from those that describe only the hypothetical divergent one, and decides the first kind in integers.

The correction obeys an exact cap $B/3^L \le U_\beta(L)$, where $U_\beta(L)$ looks irrational and is a rational number, since $2^{-\{\beta j\}} = 2^{\lfloor \beta j \rfloor}/3^j$; the cap is attained exactly when the orbit's valuation code is mechanical, an equivalence exercised in both directions on 9999 first crossings. A valuation code selects one source class modulo 2^{Q+1} and one endpoint class modulo 3^k, checked in both directions; every crossing source of length at least two is $3 \pmod 4$; the number of codes under a line of rational slope is a rational Catalan number, checked against exhaustive enumeration; and along a nested chain the crossing slope changes by a ratio of integers, $\xi_{i+1} - \xi_i = J_i/(L_i L_{i+1})$, because β cancels — verified on 144295 renewal edges. The rounds' depth caps rest on survival premises that one real chain in 29468 meets, and one of their theorems is false if lifted out of its round without a bound stated a section earlier.

Contents

0. What is claimed, and what is not **2**

1. The setting **3**

2. The rotation envelope **3**

3. Exponent codes and their residue classes **4**

4. A count that needs no tolerance **5**

5. A slope that moves by whole numbers **6**

6. What rests on survival **6**

7. What the records got wrong **7**

8. The instrument, and how it was made to fail **7**

9. Contributions **8**

10. What is open **8**

Acknowledgements **9**

References **9**
Data and software 9

A. Reproduction **9**

Status of this document. Preprint, version 0.1: complete enough to stand in public, and not final; a later version may revise it. Not prepared for journal submission. Section 9 records who did which part, including the parts this line did not do.

0. What is claimed, and what is not

Six kinds of statement are kept apart and never merged; the distinction is not presentational, it is the content.

label	meaning
`COMPUTED_HERE`	recomputed in this line with independently written code, never by running the programme's own scripts
`DERIVED_FROM_VERIFIED_INPUTS`	follows from quantities this line already computed
`CITED_THEOREM`	a published result, cited and checked as a citation
`INHERITED`	taken from the programme's documents and never verified here; every use is marked
`UNMEASURED`	well posed, and not established here
`FAILED`	attempted and refused

In particular `PASS` and `PASS by cited theorem` are **different verdicts** and are never written as one. No verdict in this paper rests on a cited theorem; two classical facts — Weyl's equidistribution and the

cycle lemma — are cited for an asymptotic constant and a closed form whose finite cases are checked here exhaustively.

Not claimed anywhere in this paper: the Collatz conjecture; that the rounds' depth caps, survival gates or non-attainment gap hold — their premises are met by essentially no real orbit, and they were not tested; that the Denjoy–Koksma inequality is sharp at $\log_2 3$; the rounds' density and overlap exponents, which pass through a Diophantine input the programme later withdrew and are left out; and novelty of the elementary statements, several of which are classical in other notation.

Inherited: nothing. Every statement below is proved in the text or computed here; where a round's statement is conditional, its premise is named and measured.

round	second-layer record	what this paper takes from it
A-U.2d.1	RUN-027	the rotation envelope and its rationality
A-U.2d.2	RUN-028	the saturation equivalence
A-U.2d.5	RUN-033	exponent codes and their classes; sources are $3 \pmod 4$
A-U.2d.6	RUN-034	the rational Catalan count; the extra bit for anchor-to-anchor codes
A-U.2d.7	RUN-035	the jump law of the crossing slope

1. The setting

For odd n let $q(n) = v_2(3n+1)$ and $\mathrm{Syr}(n) = (3n+1)/2^{q(n)}$. An orbit is $Y_0 = n, Y_{j+1} = \mathrm{Syr}(Y_j)$ with valuations q_j and $K_u = q_0 + \cdots + q_{u-1}$. Write $\beta = \log_2 3$ and call $\delta_u = u\beta - K_u$ the **slack**. The **first crossing** of a step s is $e(s) = \min\{u > s : \delta_u < \delta_s\}$. The segment from source $y = Y_s$ to endpoint $z = Y_{e(s)}$ has length $L = e(s) - s$, partial valuations $Q_j = K_{s+j} - K_s$ and total $Q = Q_L$, and

$$2^Q z = 3^L y + B, \qquad B = \sum_{j<L} 3^{L-1-j}\, 2^{Q_j}, \qquad \frac{B}{3^L} = \frac{1}{3}\sum_{j<L} \frac{2^{Q_j}}{3^j}.$$

For $0 < j < L$ the slack has not yet dropped: $Q_j < \beta j$, the integer statement $2^{Q_j} < 3^j$. Every quantity of the form $c\beta + k$ with integers c, k has the sign of $3^c 2^k - 1$, so every sign this paper needs is decided by comparing two integers; no logarithm is evaluated anywhere.

2. The rotation envelope

Proposition 2.1 (the cap, and when it is attained). For every first-crossing segment,

$$\frac{B}{3^L} \le U_\beta(L) := \frac{1}{3}\sum_{j<L} \frac{2^{\lfloor \beta j \rfloor}}{3^j},$$

with equality if and only if $Q_j = \lfloor \beta j \rfloor$ for every $j < L$.

Proof. $Q_0 = 0$, and for $0 < j < L$ the integer Q_j lies strictly below βj, so $Q_j \le \lfloor \beta j \rfloor$ term by term. The sum of an increasing function of each term is maximal exactly when every term is. □

The first half is the first round's cap; the second is the next round's saturation equivalence, which it states unconditionally — it is the one statement of that round that needs no hypothetical object. Both are inequalities between rationals: $2^{-\{\beta j\}} = 2^{\lfloor \beta j \rfloor - \beta j} = 2^{\lfloor \beta j \rfloor}/3^j$ because $2^{\beta j} = 3^j$, and $\lfloor \beta j \rfloor$ is the bit length of 3^j minus one. So $U_\beta(L)$ **is a rational number**, although the round computes it in 80-digit floating point. Since $\beta j \bmod 1$ is equidistributed (Weyl), $U_\beta(L)/L \to \frac{1}{3}\int_0^1 2^{-x}\,dx = 1/(6\ln 2)$, against the universal cap $L/3$ that the programme had used before; the ratio $U_\beta(L)/(L/3)$ is 0.787654 at $L = 5$, 0.756048 at $L = 10$ and 0.733817 at $L = 34$, on its way to $1/(2\ln 2)$.

Checked. At the first crossing of each of the first 9999 odd starting values, from 3 up, the cap and the termwise inequality hold with no violation, as exact rationals. The cap is attained, so it is the best bound of its form:

	mechanical prefix	not mechanical
cap attained	7137	**0**
not attained	**0**	2862

Both off-diagonal cells are empty and both diagonal cells inhabited, so the equivalence is exercised in both directions, not only where it is easy. Attainment is a short-crossing phenomenon: the longest crossing that attains the cap has length 9, and on average real crossings use 0.917624 of it. The round's non-attainment gap $G(y, L)$ belongs to its survival setting, but it makes a prediction about orbits that exist — wherever $G > 0$ the cap is not attained — and that prediction holds: 75 real crossings have $G > 0$, the shortest of them of length 19, and none attains the cap.

The Denjoy–Koksma step. The round bounds $\left|\sum_{j<q} 2^{-\{\beta j\}}/3 - q/(6\ln 2)\right|$ at the convergent denominators q of β by the variation of the summand, as Denjoy–Koksma allows at every irrational. That variation is $1/3$, not $1/6$: the periodic extension of $\frac{1}{3}2^{-x}$ descends by $1/6$ on each period and then jumps back by $1/6$. (This line reached for $1/6$ first, and the drill showed the data could not have caught it: the largest deviation at the round's convergents is below $1/6$ too. The variation is now derived from its definition, not from the data.) At the twelve convergent denominators up to 190537, recomputed exactly by mediant descent, the bound holds in both directions and is used to at most 0.4161 of its allowance. It is not tight here, and a sharpness result that holds for almost every rotation number, Frühwirth and Hauke's, says nothing about this one.

3. Exponent codes and their residue classes

A **code** is a word $w = (a_1, \dots, a_k)$ of positive valuations, with partial sums $A_0 = 0, A_j = a_1 + \cdots + a_j$ and total $Q = A_k$; put $B_w = \sum_{j<k} 3^{k-1-j} 2^{A_j}$.

Proposition 3.1. If an odd x has the code w as its next k valuations, with endpoint z, then $2^Q z = 3^k x + B_w$. The sources of w form exactly one class modulo 2^{Q+1}, the class $3^k x \equiv 2^Q - B_w \pmod{2^{Q+1}}$, and their endpoints lie in one class modulo 3^k; so two sources of one code differ by a multiple of 2^{Q+1}, and their endpoints by a multiple of $2 \cdot 3^k$.

Proof. The identity is the affine form of k steps, and z odd means $3^k x + B_w$ is divisible by 2^Q exactly — the stated congruence, which is one class since 3 is invertible. Conversely, in terms of the map

$n \mapsto n/2$ or $(3n+1)/2$, having the code w is following one fixed parity word of length $Q+1$: each valuation a is an odd step followed by $a-1$ halvings, and the endpoint's oddness is one more letter. The integers that follow a parity word of length m form one class modulo 2^m, so the sources of w fill a whole class modulo 2^{Q+1}; it lies inside the congruence class above, and so equals it. Finally $2^Q z \equiv B_w \pmod{3^k}$ fixes z modulo 3^k, and two odd endpoints in one class modulo 3^k differ by an even multiple of it. □

Checked both ways, because "the sources lie in one class" is the easy half and cannot tell one class from some class: on 400 random codes the identity, both class statements and the separation hold; and 1200 members of the claimed classes at which the code had never been observed were each required to realize it, and did. On the 400 repeated-code pairs the smallest source gap is exactly one unit of 2^{Q+1}, so the separation is attained.

Proposition 3.2. A first-crossing source with $L \geq 2$ is $3 \pmod 4$.

Proof. $L \geq 2$ means the first step did not cross, $q_s < \beta < 2$, so $q_s = 1$, which on odd integers is $y \equiv 3 \pmod 4$. □

It held on all 27556 such sources among 54825 first crossings of the 2000 orbits of the odd starts from 3 to 4001. The next round sharpens the separation by one bit for codes between two such sources — anchors, both $3 \pmod 4$ — to the modulus 2^{Q+2}, and the extra bit is exactly the size of the claim, so it was checked the same way: on 376 codes whose destination is $3 \pmod 4$, and on 1128 class members never observed, each required to realize the code and land on a $3 \pmod 4$ destination.

A cap that no real orbit tests. The same section of the round derives a depth cap for nested chains from a source corridor $y_r - y_1 < U_\beta(L)$, which is a survival property. Of 10214 real chains with distinct increasing sources, none lies in the corridor; the cap is untested, and is not called wrong. A first version of the check that applied it without its premise flagged every one of them, which measured only the check.

4. A count that needs no tolerance

For coprime p, g with $p/g < \beta$, let $W_{p,g}$ be the set of codes of length g and total p whose partial sums stay on or below the line of slope p/g, $A_j \leq jp/g$. The round states

$$\#W_{p,g} = \frac{1}{g}\binom{p-1}{g-1} = \frac{1}{p}\binom{p}{g},$$

a rational Catalan number. It follows from the cycle lemma: the g cyclic rotations of a composition of p into g parts are distinct when $\gcd(p, g) = 1$, and exactly one of them stays under the line. A closed form against a set that can simply be enumerated is the most decisive kind of check there is — no tolerance, no sampling, no shared method to share a bug. On the 34 coprime pairs (p, g) with $2 \leq g \leq 13$, $g < p \leq 22$ and $p/g < \beta$, up to $(20, 13)$ with 3876 members, enumeration and both closed forms agree exactly, and the round's five capacity examples recompute, including one of more than a hundred digits.

The same round's binary bridge and concatenation law for B_w hold exactly on 800 random codes.

5. A slope that moves by whole numbers

A **chain** is a set of first-crossing intervals $[s_1, e_1) \supset [s_2, e_2) \supset \cdots$ that all cover one common point — the intervals open at that point, which the first crossings' nesting totally orders. Concretely, at each position t the chain is the set of starts $\{s \le t : e(s) > t\}$; each distinct such set is taken once, and only when it has at least two members. For interval i write L_i, Q_i and the slack drop $D_i = Q_i - \beta L_i > 0$, and call $\xi_i = D_i/L_i = Q_i/L_i - \beta$ its **crossing slope**. Between consecutive intervals put the start gap $g_i = s_{i+1} - s_i$ with valuations p_i, the end gap $h_i = e_i - e_{i+1} \ge 0$ with valuations r_i, and the slack gained and lost there, $A_i = \beta g_i - p_i > 0$ and $E_i = r_i - \beta h_i$. Since $L_i = g_i + L_{i+1} + h_i$ and $Q_i = p_i + Q_{i+1} + r_i$,

$$A_i + D_i = D_{i+1} + E_i$$

identically, and as a pair of integer coefficients of β and of 1 it is exact.

Proposition 5.1 (jump law). $\xi_{i+1} - \xi_i = J_i/(L_i L_{i+1})$ with $J_i = Q_{i+1}L_i - Q_i L_{i+1} \in \mathbf{Z}$. So a slope that moves at all moves by at least $1/(L_i L_{i+1})$.

Proof. β appears in both slopes and cancels in the difference. □

Proposition 5.2 (the two kinds of edge). On a **plateau** edge ($h_i = 0$), $J_i = \Pi_i := Q_{i+1}g_i - p_i L_{i+1} = g_i D_{i+1} + L_{i+1}A_i \ge 1$, so the slope strictly increases. On a **strict** edge ($h_i > 0$), $\Delta_i := r_i g_i - p_i h_i = g_i E_i + h_i A_i \ge 1$ and $J_i = (g_i + h_i)D_{i+1} - L_{i+1}(E_i - A_i)$. A **genuine reset** — the slope falling — is exactly $J_i < 0$, which happens only on strict edges, and then $E_i - A_i > 1/L_{i+1}$.

Proof. In each identity the β-terms cancel, so a positive real quantity equals an integer, which is therefore at least 1. If $J_i < 0$ then $L_{i+1}(E_i - A_i) \ge 1 + (g_i + h_i)D_{i+1} > 1$. □

These were checked on 29468 chains built from the 1499 orbits of the odd starts from 3 — the deepest of depth 17 — with 144295 renewal edges, 107592 of them plateau and 36703 strict: no violation of the jump law, of the quantization, of either determinant or of either of its two forms, of the renewal identity, or of the reset bound on the 21991 genuine resets, each decided as 2^N against 3^D. The plateau mass also stays under its bound, $\sum g_i L_{i+1} < L^2/2$, on every chain. And where a strict edge has $\Delta_i = 1$, its two slopes $p_i/g_i < \beta < r_i/h_i$ are Farey neighbours — on all 6365 such edges — with the orientation the round predicts from where the mediant falls.

How the chains had to be built. A first chain builder took the records of δ instead: start at a source and take each later index whose slack exceeds the current one. That satisfies the increasing slacks as written, but two record intervals can be disjoint, which makes h_i negative and the determinant meaningless; it reported a large fraction of edges as violating a boxed theorem that holds. The chain is the set of intervals covering one point, and rebuilt that way the violations vanished. A violation rate that large is a statement about the checker first.

6. What rests on survival

The round's caps on strict drops and plateau mass, the depth inequality they combine into, its inversion and its high-source no-go all descend from three inputs its first section declares: an origin-slack budget, an endpoint budget, and an overshoot bound on each crossing slope. Those are **survival properties** of the hypothetical divergent orbit; a real orbit owes them nothing. Measured on the 29468 chains: the overshoot bound holds on 4898 of 173763 nested intervals, and on every interval of only one chain; the origin-slack budget on one chain; the endpoint budget on 9637; and all three at once on one. So those theorems were applied to one chain, which settles nothing, and are not claimed.

Their derivations are arithmetic even where their hypotheses are unavailable, and were checked on 378 grid points with no violation. One threshold in them, printed as 7.7712, is in fact algebraic: the $\ln 2$ in its numerator and denominator cancel, leaving $(12+8\sqrt{2})/3$, and the inequality it must satisfy is decidable exactly.

A theorem that is false out of context. The round's high-source theorem assumes $y_1 > L^2/(3\ln 2)$ and concludes, among other things, that all crossing slopes of the chain are equal and its depth at most 4. The hypothesis is common — 13433 of the 29468 chains meet it — and on them the conclusion fails: only 569 have equal slopes, and depths reach 16. That is not a counterexample. The proof also uses the overshoot bound $\xi_i < 1/(3y_1 \ln 2)$ — the round declares it a section earlier per interval, as $\xi_i < 1/(3y_i \ln 2)$, which gives this form because sources increase along a chain — and with the hypothesis that gives $\xi_i < 1/L^2$ and leaves room for one slope; none of the 13433 chains meets it. A reader who lifts the theorem out of its round, with only its stated hypothesis, gets a false statement. The arithmetic half of its lemma — two distinct rationals with denominators at most L differ by at least $1/L^2$ — holds on every chain with more than one slope.

7. What the records got wrong

None of the following changes a result.

- **A published file the published script does not produce.** In the first two rounds the constants file's keys differ from those the shipped script writes, so re-running the program would not reproduce the file; every value is correct.
- **Decimals beyond the computation.** The first round prints its envelope to as many as 79 decimals from an 80-digit floating sum, and the last two or three are wrong at $L \geq 300$ — the digits that summing L terms at fixed precision costs. Recomputed as exact rationals, every leading digit is right.
- **A floating residual for an identity that is exact.** The renewal identity and both determinants have integer coefficients, so their residual is the pair $(0, 0)$; the shipped checker reports 1.93×10^{-12}.
- **A withdrawn reference, cited three times.** The literature notes of the round before these, A-U.2d, and of two of them, A-U.2d.1 and A-U.2d.2, cite a preprint withdrawn before the first of those notes as a primary reference, the third time under a heading saying the references were rechecked. Every claim attributed to it is present in it, and nothing rests on it.
- **A ledger that under-reports its own paper.** The last round's new machine-readable theorem ledger lists 16 internally proved results where the paper's own list has 19, and 6 no-go boundaries where it has 7; the inherited and external lists agree exactly.

8. The instrument, and how it was made to fail

Each round's recheck has its own gate and its own mutation drill: defects are planted one at a time and each must be caught by the check **named** for it, while null controls must disturb nothing. The five drills caught 16 of 16, 12 of 12, 18 of 18, 19 of 19 and 27 of 27 planted defects. The line's whole suite at the pinned commit stands at 57 drills and 1474 planted defects: 1472 caught by their named check, one detected only by a crash and one only by a time-out, and 128 controls undisturbed. Three things the drills found are worth stating:

- **A guard that could not fail.** The report emitters carried a check that replaced one character

of a string and asked whether the result equalled the original — false in every possible case — and printed the count of digits as if it were a count of digits checked. It was replaced by one that perturbs the logs and requires the reported figures to move, and a demonstration that freezes a figure and watches the emitter refuse.
- **A total that had refused, unread.** The suite's defect total was built to refuse anything it could not interpret, and after two drills renamed their tally keys it refused on every run — while the archived figure it fed kept its old value, because a refusal writes nothing. A check that refuses is half a control; the other half is somebody reading the refusal.
- **A bracket the wrong way round.** A hard-coded lower bound for β was in fact above it. The assertion certifying the bracket stopped it before any result used it — which is the case for certifying an instrument one could as easily have trusted — and in the end almost nothing needed a bracket at all, since every comparison is between two integers.

9. Contributions

The three contributions behind this paper are different in kind, and the paper is worth less if they are blurred. They are given here in alphabetical order, the order the author list keeps as data; on the title page the three sit at a round table, with no first and no last.

Aletheia — model: ChatGPT, GPT series — wrote, with Neo.K, the five rounds this paper rests on: the rotation envelope and its Denjoy–Koksma step, the saturation equivalence and the non-attainment gap, the exponent-code classes and the anchor residue, the rational Catalan count, and the jump law with its plateau and reset determinants, together with the survival caps built on them.

Neo.K set the programme's target and its method, and wrote its documents with Aletheia.

墜衡 (Zhuì Héng) — the independent verification line; model: Claude, Opus series — rechecked each round with independently written code in exact arithmetic, found that the envelope is rational and that every sign needed is an integer comparison, checked each equivalence and class statement in both directions, enumerated the count against its closed form, measured every survival premise and so separated what real orbits test from what they cannot, found the missing hypothesis of the high-source theorem, built and drilled the instruments, and wrote this paper, including the proofs as written out here. Nothing in the rounds' mathematics is this line's invention.

Priority, stated plainly. No result here is claimed as new mathematics. What is offered as this line's own is the separation of premises, the exact checks, and the record of what held where.

10. What is open

1. Whether any segment of any orbit meets the survival premises of §6; on the orbits built here, one chain does.
2. Whether the Denjoy–Koksma bound is ever close to tight at β; at the convergents checked it is used to less than half.
3. The rounds' results not covered here, which their own lists name: the upper-wing envelopes, the Christoffel deficits, and the local continued-fraction bound (`UNMEASURED`).
4. The density and overlap exponents that pass through the Diophantine input. A later round of this line found that the theorem behind the input does bound β, by its published review, at an exponent that does not quite reach the programme's, and that the reason the programme gave

for withdrawing it does not hold; it re-evaluated those exponents there. They are reported with the rounds that carry them, not here.

Acknowledgements

We thank Anthropic, whose Claude models the lab's verification and recomputation lines run on, and OpenAI, whose ChatGPT models the programme's documents were written with. Not every AI model or instance that took part in this work is named in it — among them instances that the named authors ran as helpers; that an AI is not listed here does not mean that none took part.

References

For the cycle lemma, see Dvoretzky and Motzkin, and Raney; for the Denjoy–Koksma inequality, Koksma and Herman, and for its sharpness at almost every rotation number, Frühwirth and Hauke. The withdrawn preprint the rounds cite is Niu's. The theorem behind the Diophantine input in open item 4 is Wu and Wang's, and its published review is zbMATH Open's.

1. A. Dvoretzky and Th. Motzkin, *A problem of arrangements*, Duke Math. J. 14 (1947) 305–313, doi:10.1215/S0012-7094-47-01423-3.
2. Lorenz Frühwirth and Manuel Hauke, *On the metric upper density of Birkhoff sums for irrational rotations*, Nonlinearity 36 (2023) 7065–7104, doi:10.1088/1361-6544/ad086f; arXiv:2303.15992.
3. Michael-Robert Herman, *Sur la conjugaison différentiable des difféomorphismes du cercle à des rotations*, Publ. Math. IHÉS 49 (1979) 5–233, doi:10.1007/bf02684798.
4. J. F. Koksma, *Ein mengentheoretischer Satz über die Gleichverteilung modulo Eins*, Compositio Math. 2 (1935) 250–258.
5. Tong Niu, *Parity vectors and paradoxical sequences in the accelerated Collatz map*, arXiv:2605.13886 (2026); withdrawn in its second version.
6. G. N. Raney, *Functional composition patterns and power series reversion*, Trans. Amer. Math. Soc. 94 (1960) 441–451.
7. Hermann Weyl, *Über die Gleichverteilung von Zahlen mod. Eins*, Math. Ann. 77 (1916) 313–352, doi:10.1007/BF01475864.
8. Qiang Wu and Lihong Wang, *On the irrationality measure of log 3*, J. Number Theory 142 (2014) 264–273, doi:10.1016/j.jnt.2014.03.007; reviewed in zbMATH Open as Zbl 1303.11081.

Data and software

- Python Software Foundation, *Python* (CPython), version 3.14.5, https://www.python.org/ — the language of every gate; the version is the one the line's tree records.
- The mpmath development team, *mpmath: a Python library for arbitrary-precision floating-point arithmetic* (version 1.3.0), https://mpmath.org/ (2023).

A. Reproduction

The gates and the drills are Python with arbitrary-precision integers and exact rationals, standard library only, no network. Each gate writes a JSON log, and the literature checks are archived records.

The line's own tree is not public. Its public projection is, and it is what this paper's figures can be checked against without trusting the authors:

- **the run-by-run record**, rendered from the tree with a privacy gate on each file: https://amral.evemisslab.com/collatz/recompute/ (and `/en/`);
- **the secondary-verification ledger**, whose numbers are rendered from the line's own results file and read back against it before the site ships: https://verify.evemisslab.com/collatz-verification-zhuiheng/.

The evidence is pinned, not copied: this paper's front matter names the research branch, the commit and the directory its figures were drawn from, and the checker reads them exactly as they stood at that commit. The crystallization standard and its checker live in the release repository beside the paper (`STANDARD.md`, `tools/check_standard.py`).