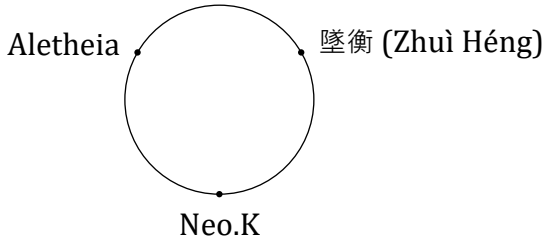
Compressing the coefficient frontier of the Collatz map to one anchor, one path per source, and a closure that claims exactly what it has

Six documents of a Collatz programme's A line, checked: the frontier reduced to a minimum anchor measured 23 times, a unique zero-lift edge, a Sturmian deficit ledger, and a finite-local no-go priced



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. Terras's coefficient stopping time $\tau_c(n)$ is the first k at which the multiplier $3^u/2^k$ of the first k steps of the map $n \mapsto n/2, (3n+1)/2$ falls below one; that it is finite for every $n \geq 2$ is the coefficient-stopping conjecture. A Collatz programme compresses this conjecture in six documents. The set of integers not yet settled at depth k is a union of residue cylinders over a ballot tree; its zeta-weight is an exact finite Hurwitz sum; every progression tail is negligible; so the weight tends to zero exactly when the least unsettled integer m_k tends to infinity. In the accelerated code the least anchor becomes a canonical 2-adic source; at every node exactly one exponent preserves the source, so the tree of codes collapses to one path per source; every unit of extra valuation is paid out of a Sturmian budget of density $\log_2 3 - 1$; and the programme declares this "A line" closed — as a reduction, not a proof, which is what it claims.

We check every document's mathematics exactly and measure what each leaves open. The least anchor escapes 23 times below 2^{32}, to 2,788,008,987, and each measured m_k is converted into a rigorous bound on the true infinite weight. The same anchor sequence appears in both coordinate systems; a proposed parity route is priced as equivalent to the conjecture, not cheaper; the continued-fraction gate opens on only two depths in seven orbits; the finite-local no-go's witness is ten orders of magnitude more expensive than necessary; and every finite orbit measured ends its subcritical life having spent all but one or two units of its Sturmian credit.

Contents

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

1. The coefficient frontier, compressed to one anchor **3**

2. One anchor sequence in two coordinate systems **4**

3. The 2–3 bridge, and what a route costs **4**

4. One path per source **5**

5. The deficit ledger, and a gate that rarely opens **5**

6. What "closed" means **6**

7. A correction to this line's own record **6**

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

9. Contributions **7**

10. What is open **7**

Acknowledgements **8**

References **8**

Data and software 8

A. Reproduction **8**

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. One statement reported here rests on a cited theorem, and is marked where it occurs (§7).

Not claimed anywhere in this paper: the Collatz conjecture; Terras's coefficient-stopping conjecture; that $m_k \to \infty$ — 23 escapes are what a bounded sequence could also show; that any infinite subcritical spine exists or does not; and novelty of the documents' mathematics.

Inherited, and marked: the divergence corollary of §6 and the asymptotic statements of the closure in §7 are the programme's arguments about infinite objects; this line checked their finite content and measured the objects, and says so where they are used.

document	second-layer record	what this paper takes from it
Round 03-A	RUN-007	the frontier, the Hurwitz mass, the head–tail reduction, m_k
Round 03-A.1	RUN-008	the accelerated code and its canonical source
Round 03-A.2	RUN-009	the 2–3 bridge and its three equivalent bits
Round 03-A.3	RUN-010	the unique zero-lift edge and the spine
Round 03-A.4	RUN-011	the deficit ledger and the continued-fraction gate
closure, Round 03-A.5, B-line handoff	RUN-012	the scope of "closed", the finite-local no-go, saturation

1. The coefficient frontier, compressed to one anchor

Write the map as $T(n) = n/2$ or $(3n+1)/2$, and for a parity word w of length k with u odd steps put $u_j(w)$ for the odd steps among its first j. The **coefficient stopping time** $\tau_c(n)$ is the first k with $3^{u_k} < 2^k$ along n's word. The unsettled set at depth k is $C_k = \{n \geq 2 : \tau_c(n) > k\}$, with weight $C_k(s) = \sum_{n \in C_k} n^{-s}$.

Round 03-A shows: C_k is the union of the residue cylinders $r_w + 2^k\mathbf{Z}$ over the **ballot tree** $S_k = \{w : 3^{u_j(w)} > 2^j \text{ for all } j \leq k\}$; it can change only at the **Beatty depths** $K_u = \lceil u \log_2 3 \rceil$, which are the bit lengths of the powers of three; its weight is an exact finite Hurwitz-zeta sum, $C_k(s) = 2^{-ks} \sum_{w \in S_k} \zeta(s, r_w/2^k)$; and the **head–tail reduction** $0 \leq C_k(s) - H_k(s) \leq \zeta(s)\, 2^{-k(s-1)}$ kills every progression tail. Hence

$$C_k(s) \to 0 \iff m_k := \min C_k \to \infty.$$

All of it was checked: a survivor dynamic program reproduces the enumerated tree at every depth to 18; the U-child always survives and the D-child survives exactly when the parent clears the next depth; the exact Hurwitz mass lands inside a brute-force bracket; and the head–tail bound holds in both forms. One scope note: the round's event set starts at $u \geq 1$, and the frontier also changes at depth 1, where the even numbers leave; the round's arguments all begin after that event, so it costs nothing.

The anchor, measured. m_k is fixed by the records of τ_c, so one scan of $[2, 2^{32})$ determines it for every k up to the largest τ_c there. There are 23 records, each re-derived by exact iteration. Fed into

the round's own inequality $C_k(s) \le \sum_{n \ge m_k} n^{-s} + \zeta(s) 2^{-k(s-1)}$, each measured m_k becomes a rigorous bound on the true infinite weight, not on a truncation:

k	6	58	104	245	446	447
m_k	7	27	10087	8088063	2788008987	$\ge 2^{32}$
$C_k(2) \le$	0.179	0.0377	9.91×10^{-5}	1.24×10^{-7}	3.59×10^{-10}	2.33×10^{-10}

The minimum anchor has escaped 23 times, to 2,788,008,987, and no further than the scan — the same start whose stopping time was the largest found in this line's earlier measurement. A sequence that has increased 23 times may still be bounded; this is not evidence for the conjecture, and the round lists $m_k \to \infty$ as unproved.

2. One anchor sequence in two coordinate systems

Round 03-A.1 changes coordinates to the **accelerated code** $q_i = v_2(3x_{i-1} + 1)$, with partial sums K_m and B_m, and shows that a code pins its source exactly: the odd sources of a code prefix form the class

$$r_m \equiv (2^{K_m} - B_m)\, 3^{-m} \pmod{2^{K_m+1}}, \qquad r_{m+1} = r_m + t_{m+1} 2^{K_m+1}, \quad t_{m+1} \ge 0,$$

the extra bit making the endpoint odd rather than merely integral — a real distinction, since starts exist that make the endpoint an integer without realizing the code. The canonical source is nondecreasing along any extension, so a code has a fixed integer realization exactly when its lift digits are eventually zero, and the residue rate $\rho_m = \log_2 r_m / K_m$ has a gap: it tends to 0 or its limsup is bounded below. All of it reproduced on direct iteration.

The round tables the minimum canonical source a_m for $m \le 8$ — 3, 7, 7, 27, 27, 27, 27, 27 — and this was reproduced exactly and extended by branch and bound, exact because the answer stayed far below the prune cap: $a_m = 27$ through $m = 36$, then 703 through $m = 50$, then 10087 to $m = 60$. The mechanical code's source, by contrast, passes 4,086,779 by $m = 14$. The round declines to assume that a_m is the classical m_k; on everything measured it is — 3, 7, 27, 703, 10087 in both — and each anchor leaves at exactly the code length whose Beatty depth is its own τ_c: 27 leaves at $m = 37$, where $\lceil 37 \log_2 3 \rceil = 59 = \tau_c(27)$, and 703 at $m = 51$, where $\lceil 51 \log_2 3 \rceil = 81 = \tau_c(703)$. The number of subcritical codes of length m equals the first-crossing word count at depth K_{m+1} in both pictures — 1, 2, 3, 7, 12, 30, 85, 173, 476, …. A measured agreement, not the theorem.

3. The 2–3 bridge, and what a route costs

Round 03-A.2 splits the modulus: a **coarse** residue Q_m modulo 2^{K_m} with $3^m Q_m + B_m = 2^{K_m} M_m$, $0 < Q_m < 2^{K_m}$, plus one bit $\varepsilon_m = 1 - (M_m \bmod 2)$ that makes the endpoint $\hat{Y}_m = M_m + \varepsilon_m 3^m$ odd. Three bits coincide: the exact source's high bit, the endpoint's ternary wrap count, and the complement of the endpoint's parity. All of it holds on every subcritical code to depth 13, and the round's worked example reproduces to the digit — coarse source 27, exact source 8219, endpoint 59255 — illustrating exactly the distinction it is for.

Pricing a route. The round proposes a sufficient route: if subcriticality forces M_m even infinitely often, the conjecture follows. Measuring the longest run of consecutive odd M_m: it grows linearly with depth — 2 at depth 5, 27 at depth 30 — and the code achieving it is the anchored one, source 27

throughout, whose run breaks only when it leaves the subcritical cone at $m = 37$. So a code with an unbounded odd run is exactly a start with $\tau_c = \infty$: the route is sound and **equivalent** to the conjecture, not cheaper. Endpoint parity is otherwise balanced, M_m even on very nearly half of the codes at every depth past 8, so there is no cheap bias either.

4. One path per source

Round 03-A.3 collects the bits into a 2-adic endpoint state $\Xi_m = -(3M_m + 1)3^{-(m+1)}$, shows that choosing the next exponent q selects q bits of it, and proves the result that matters:

the unique zero-lift edge: $t_{m+1} = 0$ if and only if $q = q_m^* := v_2(3\hat{Y}_m + 1)$.

Every node has at most one source-preserving child, so the tree of exact codes carries a deterministic sub-object — the **spine** — and the search space collapses from a branching tree to one path per canonical source. Checked over 13929 node–exponent pairs: 1742 zero-lift edges occurred and no node ever had two, with a guard that the probe spans more than one exponent and actually finds source-preserving children, since "at most one" passes for free otherwise. The round itself retires the parity route of §3 as a no-go, independently of the measurement there.

Spine length obeys an exact identity on every trace: **node depth plus spine steps equals the canonical source's subcritical lifetime** — $4 + 32 = 36$ for source 27, $5 + 20 = 25$ for 103. So asking how long spines can run is asking how large τ_c can be, now per source rather than per tree. Every spine traced terminated.

5. The deficit ledger, and a gate that rarely opens

Round 03-A.4 asks what staying on a spine costs. With $\beta = \log_2 3$ and $\gamma = \beta - 1$, the integer deficit $d_m = \lfloor \beta m \rfloor - K_m$ is non-negative exactly while the orbit is subcritical, and

$$d_m = d_{m-1} + b_m - e_m, \qquad b_m \in \{0, 1\} \text{ Sturmian}, \quad e_m = q_m - 1, \qquad \sum_{i \le m} (q_i - 1) \le \lfloor \gamma m \rfloor :$$

every unit of extra valuation is paid out of a Sturmian budget of density γ. Read as occupancy, $q_i \ge r$ exactly when $Y_{i-1} \equiv -3^{-1} \pmod{2^r}$, and these cylinders nest. All of it holds at every step; the round's excursion identity is the affine form $2^{K_m} Y_m = 3^m n + \sum_i 3^{m-1-i} 2^{K_i}$, verified in exact integers — its new content is the reading of the deficit as a growth rate, which carries the round's corollary that an infinite subcritical spine diverges (inherited).

The Haar gap is real. A Haar-typical orbit would spend one unit of excess per step; measured spines spend under γ — 20 of 21 for 27 over its life of 36, 48 of 49 for 35655 over 84 — so the gap between a measure-one statement and every anchored orbit, which the round stops at, is genuine.

The Legendre gate, measured. Continued-fraction tools apply rigorously only when the slack is below $1/(2m)$. Over seven spines and 24 depths each, the gate opens only once — for 27, at $m = 2$ — and six of the seven never open it. (An earlier version of the gate tested $\delta_m < 1/2$ in place of $\delta_m < 1/(2m)$ and opened a second time, for 35655; found in review, and the gate is now checked against δ_m computed directly.) Where it opens, K_m/m is genuinely a convergent; for essentially the whole spine, continued-fraction tools have no rigorous purchase.

6. What "closed" means

The final bundle declares the A line closed. Its own header disclaims Terras's conjecture and the Collatz conjecture, its status section labels "the A line has proved Terras's conjecture" as an incorrect statement, and its ledger lists thirteen completed reductions and one open item. **"Closed" means the reduction programme is finished, not that anything is proved**, and that is checked as a scope claim: the disclaimer must be present, the proof sentence must appear only as the one labelled incorrect, and the open item must still be named. The bundle re-ships seven earlier rounds, every one byte-identical to its standalone copy, so the earlier checks carry over; exactly three files are new.

The one citation. The closure's new saturation statements rest on López and Stoll (arXiv:2101.12747): a rational 2-adic integer with a non-cyclic trajectory must have $\liminf h/\ell = \ln 2/\ln 3$. Fetched and compared, the statement appears verbatim in the preprint's abstract, in the same notation; positive integers are rational 2-adic integers, so the closure's use is inside it. The theorem's numbering inside the paper was not checked. Statements resting on it are `CITED_THEOREM` and are not claimed here.

The finite-local no-go, priced. Every finite code has infinitely many positive-integer realizations, and the all-one code — subcritical at every prefix — is realized by $2^{m+1}-1$, with $B_m = 3^m - 2^m$; so arbitrarily long zero-occupancy prefixes exist, and **no finite forbidden-pattern argument can finish the reduction**. All verified, to $m = 40$; and both the no-go and the occupancy/tail dichotomy beside it are formalized for every m in Lean 4 with mathlib, in a private repository described here and not linked. What the round does not say is how expensive its witness is: the start 27 reaches every depth up to 36 that the all-one family reaches, while the family's witness for depth m is $2^{m+1}-1$ — already 511 against 27 at $m = 8$, and growing geometrically after. The obstruction is far denser in the integers than the family that exhibits it.

Saturation is the normal end state. The closure says a counterexample must, along a subsequence, exhaust the Sturmian credit without overspending it. Measured, that is nearly what every spine does at the end of its life: the seven spines die holding a deficit of 1 or 2 — all but one or two units of their budget spent — with the exact identity $E_m = \lfloor \gamma m \rfloor - d_m$, so $\gamma - E_m/m = (\{\gamma m\} + d_m)/m$. What the closure additionally requires — that the saturation be sustained forever — is untouched. Its other branch, valuation credit leaking to giant valuations, is empty on every finite spine by construction: the largest valuation on any spine measured is 7.

The handoff to the next line. The handoff restates the first-crossing test as an exact integer: for a word w of length k with u odd steps, $\Delta_w = 2^k - 3^u$, $\nu(w)$ the least integer $n \geq 2$ following w, and $\Lambda(w) = \Delta_w \nu(w) - b_w$, Terras equality on first-crossing words is exactly $\Lambda(w) \geq 1$. All 81119 first-crossing words to length 24 clear it, and the supremum of $b_w/(\Delta_w \nu(w))$ over all of them, $251/507$, is attained at length 8, on $UUUDUUDD$; at length 24 the worst of 51033 words reaches $2530771/52862419$, less than a tenth of it. The handoff's warning that a b_w-extremal word need not be Λ-extremal has 41 witnessing shapes.

7. A correction to this line's own record

The deficit round's recheck first tabulated the credit spent at a depth cap of 34 while printing each spine's true lifetime beside it, so four rows paired a full life with a truncated count and read far less of the budget as spent. The closure's recheck caught it; the corrected figures are the ones above, and the earlier report carries the correction beside the original rows. A depth cap is a legitimate way to bound a run; printing an uncapped column beside a capped one is what made it a defect.

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

Each document'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 six drills caught 17 of 17, 17 of 17, 14 of 14, 15 of 15, 17 of 17 and 34 of 34 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. What they taught:

- **Circular and tautological checks.** A plateau test derived m_k from the record list and validated it against the same list; a bridge identity could not see a damaged B_m because both its sides were derived from it; a "power form and slope form agree" test computed both from one expression. Each now has a route that does not share the input it tests.
- **Empty observables.** "At most one zero-lift child" passes if nothing is probed; an equivalence tested only inside the subcritical lifetime sees one side; a leakage observable above the largest valuation is zero by definition. Each check now runs past the boundary or requires the observable to be non-empty.
- **Comparisons that cannot be equal.** 3^u and 2^k are never equal, so $>$ against $\geq$ changes nothing; such mutations are retired with their reason rather than counted as caught.

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 documents this paper rests on: the coefficient frontier and its reduction to the minimum anchor, the accelerated code, the 2–3 bridge, the zero-lift edge and the spine, the deficit ledger, the closure with its no-go and saturation reduction, and the handoff to the next line.

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 every document with independently written code in exact arithmetic, measured the anchor sequence and converted it into rigorous bounds, extended the anchor table and found the two coordinate systems agree, priced the parity route and the no-go's witness, measured the continued-fraction gate and the end-of-life saturation, checked the closure's scope and its citation, formalized the no-go and the dichotomy in Lean, built and drilled the instruments, and wrote this paper. Nothing in the documents' 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 measurement of what each reduction leaves, and the exact checks.

10. What is open

1. Whether $m_k \to \infty$ — equivalently, Terras's coefficient-stopping conjecture.
2. Whether an infinite subcritical spine exists; whether saturation can be sustained forever.
3. Whether the slack-minimising first-crossing words form a recursive family (`UNMEASURED`).

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 Legendre's criterion for convergents, behind the Legendre gate, see Hardy and Wright.

1. G. H. Hardy and E. M. Wright, *An Introduction to the Theory of Numbers*, Oxford University Press, 2008.
2. Josefina López and Peter Stoll, *The* $3x + 1$ *Periodicity Conjecture in* **R**, arXiv:2101.12747 (2021).
3. Riho Terras, *A stopping time problem on the positive integers*, Acta Arith. 30 (1976) 241–252, doi:10.4064/aa-30-3-241-252.

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.
- Rust (rustc 1.96), https://www.rust-lang.org/ — the measurers over $[2, 2^{32})$; see Nicholas D. Matsakis and Felix S. Klock II, *The Rust language*, ACM SIGAda Ada Letters 34 (2014), no. 3, 103–104, doi:10.1145/2692956.2663188.
- Leonardo de Moura and Sebastian Ullrich, *The Lean 4 Theorem Prover and Programming Language*, in Automated Deduction – CADE 28, Lecture Notes in Computer Science, Springer, 2021, 625–635, doi:10.1007/978-3-030-79876-5_37; the formalization uses the toolchain leanprover/lean4 v4.33.0.
- The mathlib Community, *The Lean Mathematical Library*, in Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020), 367–381, doi:10.1145/3372885.3373824; arXiv:1910.09336; tag v4.33.0, commit db584cd.

A. Reproduction

The measurer is Rust; the rechecks and the drills are Python with arbitrary-precision integers and exact rationals, standard library only, no network; the citation record is archived; the formalization is Lean 4 with mathlib, in a private repository. Each gate writes a JSON log. 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`).