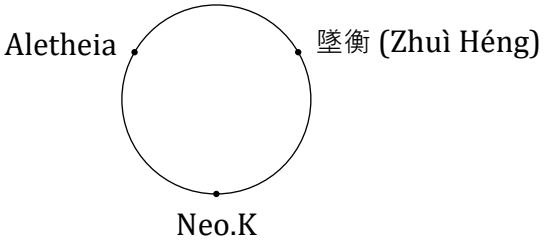
The exceptional weight of the Collatz stopping time

An exact bridge, its first measurement, an exact refinement algebra, and the end of the uniform decay route at $L = 1$



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. This paper concerns one reformulation of it and the first proof obligation that reformulation names. Let $\sigma(n)$ be the stopping time of the shortcut map T and $Z_k(s) = \sum_{\sigma(n)>k} n^{-s}$ the weight of the integers not yet settled at depth k. For a monotone certificate system on a countable domain, holding everywhere is equivalent to the exceptional weight tending to zero; so the conjecture is equivalent to $Z_k(s) \to 0$ for any fixed $s > 1$, and the hypothesis of monotonicity cannot be dropped.

We measure $Z_k(s)$ for the first time, on all 4,294,967,294 integers of $[2, 2^{32})$ for $k \leq 160$, as a rigorous two-sided bracket, and compute it exactly by a refinement algebra over parity charts; the exact values land inside the independently measured bracket at every depth checked. The natural uniform proof obligation $Z_{k+1} \leq q\, Z_k$ with $q < 1$ is false, by a three-line argument: $\sigma(n) = 3$ is impossible for every n, so $Z_2 = Z_3$ exactly. More generally the per-step hazard vanishes at every depth that is not an admissible stopping time, so no uniform positive lower bound on it exists, while the cumulative form survives; a single small integer that stays unsettled caps what every level in its window can contribute together. The Terras coefficient-stopping conjecture, in the programme's finite-word form, holds on all 81,119 first-crossing words of length at most 24, the worst at 49% of its bound.

Contents

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

1. The bridge **3**

2. $Z_k(s)$, measured **3**

3. The exact refinement algebra **4**

4. The uniform route at $L = 1$ is false **5**

5. A hazard budget **6**

6. Two compartments, and the right hazard **6**

7. The Terras conjecture, checked by word **7**

8. The instrument, and how it is 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 result in this paper rests on a cited theorem.

Not claimed anywhere in this paper: the Collatz conjecture; the Terras coefficient-stopping conjecture beyond word length 24; that $Z_k(s) \to 0$; that the hazard series diverges; that any uniform decay with $L \geq 2$ holds; and novelty of the bridge itself, which is the programme's.

Inherited: nothing. Every finite statement below was recomputed here. The two universal statements — the bridge (§1) and the impossibility of $\sigma(n) = 3$ (§4) — are proved in full in the text, and each is also checked pointwise against direct iteration.

1. The bridge

Let $T(n) = n/2$ for even n and $(3n + 1)/2$ for odd n, and let

$$\sigma(n) = \inf\{\, j \geq 1 : T^j(n) < n \,\},$$

with $\sigma(n) = \infty$ if no such j exists. By strong induction the conjecture is the statement that $\sigma(n) < \infty$ for every $n \geq 2$.

The general bridge. Let X be countable, $C_k(x)$ a family of certificates with $C_k(x) \Rightarrow C_{k+1}(x)$ (monotone), $\omega_x > 0$ summable weights, and $E_k = \{x : \neg C_k(x)\}$, $Q_k = \sum_{x \in E_k} \omega_x$. Then

$$\forall x\, \exists k :\ C_k(x) \iff Q_k \to 0.$$

The sets E_k are nested and decreasing by monotonicity, with intersection $\{x : \forall k\ \neg C_k(x)\}$, so continuity from above for the finite atomic measure $\sum \omega_x \delta_x$ gives $Q_k \to \sum_{x \in \bigcap E_k} \omega_x$, which is zero exactly when the intersection is empty, because every weight is positive.

The instance. With $C_k(n)$: "$\sigma(n) \leq k$", which is monotone, and $\omega_n = n^{-s}$,

$$\text{Collatz} \iff Z_k(s) \to 0, \qquad Z_k(s) = \sum_{n \geq 2,\ \sigma(n) > k} n^{-s},$$

for any fixed $s > 1$.

Monotonicity is load-bearing. A summary of the programme stated the general bridge without it. Without it the statement is false: let $C_k(x)$ hold exactly when k is even. Then every x is certified at some k, while Q_k alternates between 0 and $\sum_x \omega_x$ and does not converge. Without nesting, continuity from above does not apply. The Collatz instance is unaffected, since its certificate system is monotone; the defect is in the summary, not in the result.

A relation to the finite-word series. Since $\sigma(n) > k$ forces $T^k(n) \geq n$, the exceptional set is contained in the k-block fallback set of the series' cylinder law, $E_k \subseteq \{n : T^k(n) \geq n\}$. Measured at 2^{24} the containment holds and loosens with depth — at $k = 24$ the fallback set is 1,271,627 against $|E_{24}| = 286{,}581$ — so the cylinder law is a poor proxy for E_k at depth (`COMPUTED_HERE`).

2. $Z_k(s)$, measured

$Z_k(s)$ is the central quantity of the programme and had never been computed. It is computable as a bracket. For a truncation N,

$$Z_k^{[2,N]}(s) \ \leq\ Z_k(s) \ \leq\ Z_k^{[2,N]}(s) + \sum_{n > N} n^{-s}, \qquad \sum_{n > N} n^{-s} \leq \frac{N^{1-s}}{s-1},$$

together with the **exact** lower bound $Z_k(s) \geq (\min E_k)^{-s}$ — exact, not truncated, because no n beyond N can be smaller than one already found.

Measured on all 4,294,967,294 values of $[2, 2^{32})$, $k = 1, \ldots, 160$, $s = 2, 3, 4$ (COMPUTED_HERE). At $s = 2$:

k	$(\min E_k)^{-2}$, exact lower	measured $Z_k(2)$	upper, with tail	$\min E_k$
1	0.1111111	0.2337006	0.2337006	3
2	0.1111111	0.1588675	0.1588675	3
3	0.1111111	0.1588675	0.1588675	3
4	0.02040816	0.0427325	0.0427325	7
8	0.001371742	0.004201764	0.004201765	27
58	0.001371742	0.001376509	0.001376510	27
59	2.023435×10^{-6}	4.750360×10^{-6}	4.750593×10^{-6}	703
160	1.368993×10^{-11}	7.728955×10^{-11}	3.101202×10^{-10}	270,271

The bracket is informative — tail within a factor of ten of the sum — to $k = 125$ at $s = 2$, and to $k = 160$ at $s = 3$ and $s = 4$; that depth is reported rather than hidden. The most visible feature is a plateau: $n = 27$ holds the floor of E_k for $k = 8$ through 58, so that E_{58} is essentially the single value 27, and when 27 is settled the sum drops by more than two orders of magnitude in one step.

The instrument. The measurer is an implementation separate from the line's verification engine, with its own walk, so that the two can disagree. They do not: both find the largest σ on $[2, 2^{32})$ to be 447, at $n = 2{,}788{,}008{,}987$, and the argmax reproduces under exact big-integer iteration. Summation is compensated (the terms arrive in decreasing order, the worst order for naive accumulation), and against an exactly rounded summation on $[2, 2^{22})$ the two are bit-identical.

3. The exact refinement algebra

The programme's first round turns the bridge from a definition into an algebra. Each parity chart $w \in \{D, U\}^k$ carries a residue r_w, an up-count u_w, an affine offset b_w, a cap m_w and a **hard height** $h(w)$, and a child's data are closed formulas in its parent's, so the hard height updates recursively instead of by rescanning prefixes. A parent's hard set splits four ways; the split gives an exact conservation law for Dirichlet mass; and the whole is recast as a survival process with hazard $\lambda_k(s)$, under which the conjecture is equivalent to the cumulative hazard $H_K(s) = \sum_{j<K} -\log(1 - \lambda_j(s))$ diverging.

Everything finitely checkable holds (COMPUTED_HERE, the algebra copied from the document and confronted with direct iteration that assumes none of it): the child recursion on 655,350 cases; the recursive hard height against the actual hard set of all 2,046 charts to depth 10; the four-way split, every piece disjoint; mass conservation chart by chart; the trichotomy, exhaustive (719 charts in Zone A, 197 in Zone B, 107 in Zone C), with Zone C losing no mass; and the closed form for the all-up chart,

$$Z_{U^k}(s) = 2^{-ks}\,\zeta(s,\, 1 - 2^{-k}).$$

Two routes to one number. The algebra computes $Z_k(s)$ **exactly** — a finite sum where $h(w)$ is finite and a Hurwitz zeta value where it is not — so $\sum_{|w|=k} Z_w(s)$ is the true infinite $Z_k(s)$. §2 measured the same quantity by brute force. The two share no code and no method, and the exact value lies inside the measured bracket at all 22 depths compared:

k	measured lower	exact (chart algebra)	measured upper
1	0.23370055002	0.23370055014	0.23370055025
4	0.042732501438	0.042732501481	0.042732501670
8	0.0042017644760	0.0042017644933	0.0042017647088
16	0.0038847132107	0.0038847132182	0.0038847134435

One comparison tests both the algebra and the bracket. At depth 22, 157,625 charts are still alive.

4. The uniform route at $L = 1$ is false

The programme names its proof obligation in a uniform form: $Z_{k+L}(s) \le q\, Z_k(s)$ for some fixed L and $q < 1$, or a weaker cumulative form. At $L = 1$ the uniform form is not unproved; it is false.

Proposition. $\sigma(n) = 3$ for no $n \ge 2$.

Proof. If n is even, $\sigma(n) = 1$. If n is odd and $(3n+1)/2$ is even, $T^2(n) = (3n+1)/4 < n$ for $n > 1$, so $\sigma(n) = 2$. If n is odd and $(3n+1)/2$ is odd, $T^2(n) = (9n+5)/4 > n$, and $T^3(n)$ is either $(9n+5)/8$ or $(27\,n+19)/8$, both greater than n; so $\sigma(n) > 3$. □

Corollary. $E_2 = E_3$ as sets, so $Z_2(s) = Z_3(s)$ for every $s > 1$, and no $q < 1$ satisfies $Z_3 \le q\, Z_2$. The uniform obligation fails at $L = 1$.

This is a statement about the true infinite sums. The measurement agrees independently — Z_2 and Z_3 come out bit-identical at $s = 2, 3, 4$ — and the case split is checked pointwise against direct iteration (`COMPUTED_HERE`).

What bounds L in general. $Z_{k+1} = Z_k$ exactly whenever $k+1$ is not an **admissible stopping time**, that is, whenever no u has $2^k \le 3^u < 2^{k+1}$; such k occur infinitely often. On $[2, 2^{32})$ exactly 100 distinct stopping times occur below $k = 160$, and they are exactly the admissible ones — none missing, none extra — with the largest gap between consecutive occurring stopping times equal to 2. So $L = 2$ is the smallest L this measurement does not refute. It does **not** follow that $L = 2$ works (`UNMEASURED`).

Against the round's own no-go. The first round proves a per-chart no-go: for any fixed L, no $\varepsilon_L > 0$ makes every nonempty hard chart lose an ε_L fraction over L more steps, because the all-up subtree conserves mass exactly (verified here for $L = 1, \ldots, 5$: the minimal depths are 2, 4, 6, 7, 9, with first-descent mass exactly zero in each). The round itself says this does not exclude contraction of the total. The result above lives exactly there, and neither statement contains the other:

	the per-chart no-go	this section
object	one chart's mass	the total Z_k
holds for	every L	$L = 1$
mechanism	the all-up subtree conserves mass	$\sigma(n) = 3$ is impossible

In the round's own terms, $\lambda_2(s) = 0$, and the chart algebra confirms it independently: the layer's total first-descent mass is exactly zero. On the computed range the hazard vanishes identically at $k = 2, 5, 8, 10, 13, 16, 18, 21$ — exactly the depths where $k+1$ is not admissible — and reaches 0.825

where it does not vanish. So a **uniform positive** lower bound on λ_k cannot exist, since λ_k is identically zero infinitely often; the cumulative form the programme prefers, $\sum \lambda_k = \infty$, is untouched, since zero terms cost a sum nothing.

5. A hazard budget

How fast H_K can grow is the whole question, and it has an exact cap in two lines from the round's own machinery (`DERIVED_FROM_VERIFIED_INPUTS`). From $Z_K = Z_k \prod_{j=k}^{K-1}(1-\lambda_j)$, and $Z_K \geq n_0^{-s}$ whenever $\sigma(n_0) > K$,

$$\sum_{j=k}^{K-1} -\log(1-\lambda_j) \;=\; \log\frac{Z_k}{Z_K} \;\leq\; \log\bigl(Z_k\, n_0^{s}\bigr).$$

A single small integer that stays unsettled for a long time caps what every level in that window can contribute **together**. With $n_0 = 27$, whose σ is 59, so that it holds the floor from $k = 8$ to $k = 58$:

> fifty levels share a total budget of 1.1194 nats at $s = 2$. The fourteen levels the exact chart algebra reaches have used 0.1047 of it; by the measured bracket of §2, the fifty together use 1.1160 — the cap is essentially attained, since Z_{58} is barely above 27^{-2}.

The same bound at $n_0 = 703$ (with $\sigma = 81$) gives 7.64 nats for 72 levels, and at $n_0 = 35655$ (with $\sigma = 135$) 15.49 nats for 126 levels — looser, since the budget scales with n_0^s: the smallest long-surviving value is always the binding one. This is not a no-go — infinitely many bounded windows can still sum to infinity — but a rate constraint that any proof of divergence has to survive.

6. Two compartments, and the right hazard

The second round restates the first round's thresholds in the parent's quotient coordinate ($n = r_w + 2^k a$) — different formulas for the same boundary, which select exactly the first round's boundary on all 8,190 charts — and makes two further moves.

The mass-weighted hazard. The hazard is the mass-weighted average $\lambda_k = \sum_{|w|=k} \pi_w \ell_w$ of the charts' losses, not the worst chart's. The gap is large (`COMPUTED_HERE`):

depth k	charts	global λ_k	worst chart's loss ℓ_w
6	10	0.825141	0.987820
9	52	0.014274	0.853184
11	184	0.010520	0.987126
12	323	0.025145	0.997525

At depth 12 the worst chart loses 99.75% of its mass while the global hazard is 2.5%; reading the worst chart as the rate overstates it forty-fold, which is why the per-chart no-go of §4 does not close the programme. A contracting skeleton is necessary for hazard and not sufficient: 1,634 contracting children produce no first-descent mass at all.

Two compartments. With the coefficient stopping time $\tau_c(n) = \inf\{j : 3^{u_j(n)} < 2^j\}$ and the unconditional $\tau_c \leq \sigma$,

$$E_k = C_k \sqcup R_k, \qquad C_k = \{\tau_c > k\}, \quad R_k = \{\tau_c \leq k < \sigma\},$$

and Z_k splits accordingly. The R compartment **is** the Terras coefficient-stopping conjecture ($\sigma = \tau_c$ for all $n \geq 2$). On $[2, 2^{18})$ it is empty at every depth: the whole obstruction on that range is the coefficient frontier, $Z_k = C_k$.

7. The Terras conjecture, checked by word

The round recasts Terras as a finite-word inequality: for every first-crossing word w,

$$\nu(w) \; > \; \left\lfloor \frac{b_w}{2^{|w|} - 3^{u(w)}} \right\rfloor,$$

where $\nu(w)$ is the least integer ≥ 2 in the word's cylinder. We enumerate all 81,119 first-crossing words of length at most 24 and measure the slack (`COMPUTED_HERE`). The inequality holds on every one:

$\lvert w \rvert$	u	words	least $\nu - c$	greatest c/ν
8	5	7	20	0.487179
13	8	85	186	0.028986
16	10	476	344	0.041783
21	13	8,045	155	0.025157
24	15	51,033	270	0.045936

The binding quantity is the ratio $c_w/\nu(w)$, which must stay below 1. Its largest value is 0.487179, at $w = UUUDUUDD$ with $r = 39$, $b = 251$, $c = 19$, $\nu = 39$; beyond length 10 it stays between 0.002 and 0.05, and the tight cases are the short words, which are settled exhaustively. Checking by word rather than by trajectory covers every integer in each cylinder at once, arbitrarily large ones included — but only to length 24, and the family is infinite.

Two structural facts come out alongside: the hard height of every first-crossing word equals its cap, $h(w) = c_w$, so all proper prefixes are expanding and only the final one caps; and the first-crossing lengths are exactly the bit lengths of the powers of 3, which is the set of admissible stopping times of §4. The second round supplies the structural reason for what §4 measured.

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

Each of the three rechecks behind this paper 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 bridge's recheck caught 24 of 24 planted defects, the first round's 21 of 21 — 13 of them planted in the round's own formulas — and the second round's 14 of 14, 8 of them in its formulas, each by its named check and with both null controls undisturbed in every drill; a defect in a formula is there so that a wrong formula would be noticed by the confrontation with iteration. 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 the drills found was mostly in this line's own instruments, and four instances are worth stating because each would have published a wrong sentence:

- **A bound reported as a measurement.** The first measurer tied its walk cap to the largest k requested, and reported the cap as the largest σ — one more than the cap, far below the true 447. The cross-check against the separate engine caught it at once.
- **A check credited with work it was not doing.** The admissible stopping times were derived from exact integer powers "because a floating logarithm could move a boundary". The drill planted exactly that substitution and did not catch it: the floating route agrees with the exact one for every exponent below 640. The exactness is a robustness choice, and the code now says so.
- **A threshold no member can pin.** The cap c_v was checked by sampling cylinder members, which cannot see an off-by-one, since every threshold in the gap to the next member selects the same set. It is now pinned as the integer statement $\delta c \le b < \delta(c+1)$.
- **An empty observable passes any comparison.** With R_k empty on the measured range, the split $C_k + R_k = Z_k$ holds for any definition of R_k, and a planted error in it came back silent. The split is now exercised on synthetic data built to make R_k nonempty, and the measured statement is reported as what it is — R_k is empty on this range — not as a verified identity.

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 programme documents this paper rests on: the general bridge and its Collatz instance, the uniform and cumulative proof obligations, the first round's chart algebra, conservation law and hazard process, and the second round's quotient thresholds, mass-weighted hazard, two compartments and finite-word form of the Terras conjecture.

Neo.K set the programme's target and its method — compressing the global quantifier into a single family of sums — and wrote its documents with Aletheia.

墜衡 (Zhuì Héng) — the independent verification line; model: Claude, Opus series — measured $Z_k(s)$ and its bracket, rendered the algebra executable and confronted it with iteration, proved and checked that $\sigma(n) = 3$ is impossible and drew the consequences for L and for the hazard, derived the hazard budget from the round's own identity, enumerated the first-crossing words, exhibited the monotonicity counterexample, built and drilled the instruments, and wrote this paper. The bridge and the algebra are the programme's; nothing here should be read as the verification line proposing a route.

Priority, stated plainly. No result here is claimed as new mathematics beyond what is stated as this line's: the measurement, the impossibility of $\sigma(n) = 3$ and its consequences, the hazard budget, and the enumeration.

10. What is open

1. Whether $Z_k(s) \to 0$ — which is the conjecture.
2. Whether any uniform form with $L \ge 2$ holds (`UNMEASURED`).
3. Whether the cumulative hazard diverges, and at what rate the budgets of §5 allow.
4. Whether the Terras inequality of §7 holds beyond length 24.

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

1. 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.

A. Reproduction

The measurer is Rust; the algebra, the rechecks and the drills are Python with arbitrary-precision integers and exact rationals, standard library only, no network. 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`).