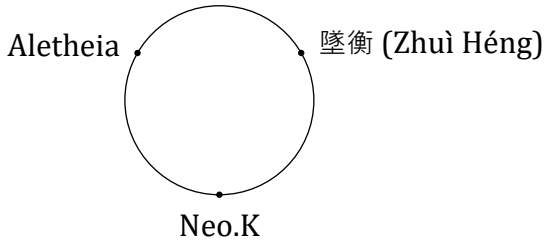
Sparse lifts cannot keep up: the packing barrier for subcritical Collatz orbits, and where it stops

Four rounds of a Collatz programme — an entropy constant below 3, a sharp packing constant, a second-order term from Stirling's prefactor, and a counting error that two implementations agreed on



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 orbit that never descended would have to stay *subcritical* forever: after m accelerated steps its total valuation K_m stays below $m \log_2 3$. Measure how far it lags the mechanical bound by the deficit $d_m = \lfloor m \log_2 3 \rfloor - K_m$. Four consecutive rounds of a Collatz programme prove that such an orbit cannot keep its deficit small: two equal blocks of valuations force the states before them into one residue class modulo a power of two, so repeated blocks force states far apart, while blocks of a thin-deficit word are too few to avoid repeating — their number grows like Λ_γ^r with $\Lambda_\gamma = (1+\gamma)^{1+\gamma}/\gamma^\gamma = 2.8395\ldots < 3$, $\gamma = \log_2 3 - 1$.

The consequence is a barrier on the record deficit D_N: $\limsup D_N / \log_2 N \geq 0.01$, improved by counting occurrences rather than repeats to the sharp $c_{\text{pack}} = 0.0358567600\ldots$, and by Stirling's prefactor to a second-order term $d_{\text{pack}} = 0.3689789787\ldots$; a further lever, the prefix constraint, provably buys nothing at first order, and the branch is closed. We recompute every constant from scratch by a different method to 80 digits or more, check the finite arithmetic the arguments rest on, and eliminate the excluded families on explicit instances. The barrier itself is about infinite orbits and is the programme's. We also record a counting error — paths counted where words were defined — that this line reproduced faithfully with two agreeing implementations and did not notice.

Contents

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

1. The setting **3**

2. Two finite facts the barrier rests on **3**

3. The squeeze: an entropy constant below 3 **4**

4. Counting occurrences: the sharp constant c_{pack} **4**

5. The second order, and a lever that bought nothing **5**

6. Paths counted where words were defined **5**

7. The instrument, and how it was made to fail **6**

8. Contributions **6**

9. What is open **7**

Acknowledgements **7**

References **7**

Data and software 7

A. Reproduction **7**

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 8 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 here rests on a cited theorem.

Not claimed anywhere in this paper: the Collatz conjecture; that a subcritical orbit exists or does not; that c_{pack} or d_{pack} is the true sharp constant — each is optimal only for its envelope; that the four levers the rounds list and do not pull are useless; and novelty of the elementary steps.

Inherited, and marked: the barriers themselves — the asymptotic statements of §§3–5 about infinite subcritical orbits — are the programme's proofs. This line checked the finite arithmetic each one rests on, recomputed every constant by an independent method, and checked the optimality claims as exhibited failures on either side; it did not re-prove the asymptotic counting, and it says so wherever a barrier is used. The saturation lower bound of the last round leans on standard local-limit and ballot estimates, which the round names as an analytic dependency rather than proving.

round	second-layer record	what this paper takes from it
A-U.2b	RUN-015	the repeated-block squeeze and $\Lambda_\gamma < 3$
A-U.2b.1	RUN-016	occurrence packing and c_{pack}
A-U.2b.2	RUN-017	prefix-constraint no-gain and d_{pack}
A-U.2b.3	RUN-018	the pointed/unpointed correction; the branch closed

1. The setting

For odd n, $\mathrm{Syr}(n) = (3n+1)/2^{q(n)}$; the orbit $Y_0 = n, Y_{m+1} = \mathrm{Syr}(Y_m)$ has valuations q_m and $K_m = q_0 + \cdots + q_{m-1}$. Write $\beta = \log_2 3$ and $\gamma = \beta - 1$. The orbit is **subcritical** to step N if $K_m < \beta m$ for $1 \le m \le N$ — it has not yet come below its start in the multiplicative sense — and a never-descending orbit would be subcritical for all N. Then $K_m \le \lfloor \beta m \rfloor$, and the **deficit** $d_m = \lfloor \beta m \rfloor - K_m \ge 0$ measures how far it lags the **mechanical** word, whose valuations are $\lfloor \beta m \rfloor - \lfloor \beta(m-1) \rfloor$ and whose deficit is identically zero. $D_N = \max_{m \le N} d_m$ is the record deficit.

Over a block of r steps starting at i, the excess of the valuations over 1 is the ledger

$$E_{i,r} = \lfloor \gamma(i+r) \rfloor - \lfloor \gamma i \rfloor + d_i - d_{i+r},$$

since $\lfloor \beta m \rfloor = m + \lfloor \gamma m \rfloor$; with $d \le D$ it lies within $D+1$ of γr, since the difference of the floors is itself within 1 of γr. Write

$$H(z) = (1+z)\log_2(1+z) - z\log_2 z,$$

so that the number of ways to distribute an excess near zr over r steps grows like $2^{rH(z)}$, and $\Lambda_\gamma = 2^{H(\gamma)} = (1+\gamma)^{1+\gamma}/\gamma^\gamma$.

2. Two finite facts the barrier rests on

Proposition 2.1 (repeated blocks). If the same block of r valuations, with total Q, follows positions a and b of an orbit, then $Y_a \equiv Y_b \pmod{2^{Q+1}}$.

Proof. A block of valuations selects one residue class of its source modulo 2^{Q+1}: in the map $n \mapsto n/2, (3n+1)/2$ it is one parity word of length $Q+1$, and the integers following a parity word of length m form one class modulo 2^m. □

Proposition 2.2 (no repeated state). An orbit subcritical for all N never repeats a state, so two occurrences of one block lie at least 2^{Q+1} apart.

Proof. A repeat is a cycle of period p with valuation sum Q_{cyc}, and $(2^{Q_{\text{cyc}}} - 3^p)Y = B_{\text{cyc}} > 0$ forces $Q_{\text{cyc}} > \beta p$; going round it indefinitely contradicts $K_m < \beta m$. □

With $Q \geq \lfloor \beta r \rfloor - D$, the separation is $2^{Q+1} > 3^r/2^D$: it operates at base 3. And the states cannot be far apart, because $Y_m = 2^{\delta_m} A_m$ with $\delta_m = \beta m - K_m < 1 + D_N$ and the renormalized anchor $A_m \leq n + m/3$, so the peak satisfies $M_N < 2^{D_N+1}(n + N/3)$.

These were checked on real orbits: the repeated-block congruence on actual repeats, with a guard that repeats occur at all; the excursion bound in exact rationals; the ledger on every block, with the requirement that the block excess actually vary across the sample (it spans 8 values), since on a degenerate sample the band would hold for free; and the peak law $M_{p(r)+1} \geq 2^{r+1}$ relating factor complexity to height.

3. The squeeze: an entropy constant below 3

Thin-deficit words cannot supply enough distinct blocks of length r — their number grows like Λ_γ^r — to keep pace with the separation the integers force at base 3, so a block repeats too early, and the excursion bound says the states are too small to be that far apart (inherited: the counting argument is the round's). The whole argument turns on

$$\Lambda_\gamma = 2.83951373049775259640 \ldots < 3,$$

computed here to the digits shown and matching every digit the round states. Its explicit form needs $1/\beta < c < 1/\log_2 \Lambda_\gamma$ to be non-empty — it is $(0.630930, 0.664168)$, and the round's $c = 0.645$ lies inside — and, at $\varepsilon = 0.01$, two inequalities that clear by 1.008×10^{-4} and 5.986×10^{-4}: the constants were chosen to just pass, and each is drilled by a perturbation smaller than its own margin, so the check reads the value and not a rounding of it. The barrier is

$$\limsup_{N \to \infty} \frac{D_N}{\log_2 N} \geq 0.01,$$

and it eliminates the mechanical word, bounded-deficit and ultimately periodic subcritical codes, and thin low-complexity words. On instances: the mechanical word has $d_m = 0$ for every m up to 400 and factor complexity exactly $r + 1$ for $r \leq 24$ — a second, independent exclusion through complexity; an ultimately periodic tail's source $B_{\text{per}}/(2^Q - 3^p)$ is negative for every subcritical period tested and positive exactly for the supercritical ones, both outcomes present among 12 periods.

How far this scheme goes. The two constraints the argument needs — block count $o(N)$, and separation beating the peak — can be optimized jointly in c. The published 0.01 is a safe round number: the same argument supports about 0.01502, at $c = 0.650$, and nothing past it, because there the constraints meet.

4. Counting occurrences: the sharp constant c_{pack}

The next round replaces "a block must recur" by "a block can recur only so often". All occurrences of one block are positive, distinct, congruent modulo 2^{Q+1} and at most M_N, so

$$\text{occ}_N(v) \leq 1 + M_N/2^{Q+1},$$

and summing over blocks gives the packing inequality $N-r+1 \le A(r, D_N)+(M_N/2^{r+1})\,B(r, D_N)$, with A, B the weak-composition sums of the thin-deficit envelope (inherited: the asymptotic conclusion). It reaches

$$c_{\text{pack}} = x^*/\beta = 0.03585676003404867\ldots, \qquad H(\gamma + x^*) = \beta,$$

2.388 times the previous scheme's ceiling — new arithmetic, where tuning could not get there. And this constant is the supremum of its envelope, not a safe value below it: $H'(z) = \log_2(1 + 1/z)$ and $H(z) - zH'(z) = \log_2(1 + z)$ make $F(x) = x/H(\gamma + x)$ strictly increasing, so $\sup F = F(x^*) = x^*/\beta$; past x^* the second packing term stops being $o(N)$, and below it the first does — both failures exhibited numerically.

Recomputed by a different method. The round ships a script and an 80-digit constants file; re-running the script would only test that it is deterministic. The constants were recomputed here by bisection in exact decimal arithmetic against the round's root-finder, and agree to 80 to 83 digits: β, γ, $z^* = \gamma + x^*$, x^* and c_{pack}. A drill defect perturbs a single digit of c_{pack} far past anything a floating comparison would notice, and the check catches it.

What a finite orbit can see. The packing inequality holds in exact rationals on every window tested, but its right side exceeds its left by a factor of at least 44.8 on all 28 windows: the multi-occurrence refinement matters as $N \to \infty$, and the longest orbits computable are short. The check therefore requires the first term alone to be insufficient somewhere, so that the second is exercised at all.

5. The second order, and a lever that bought nothing

The next round pulls the first of five listed levers — the **prefix constraint**, that the deficit stay in its corridor at every prefix — on the reasonable hypothesis that it must cost the block count something. Its own theorem says it does not: the all-prefix language has the same first-order entropy, and c_{pack} stands. The gain came instead from the **Stirling prefactor** $r^{-1/2}$ the first-order argument discarded:

$$h^* = H'(z^*) = \log_2(1 + 1/z^*) = 1.3550907472\ldots, \qquad d_{\text{pack}} = \frac{1}{2h^*} = 0.3689789787331466\ldots,$$

$$\limsup_{N\to\infty} \frac{D_N - c_{\text{pack}} \log_2 N}{\log_2 \log_2 N} \ge d_{\text{pack}} \quad \text{(inherited)}.$$

Both constants recompute independently, h^* to 80 digits and d_{pack} to 82. Generalizing the block scale by a parameter s gives two exponents that must both be negative; $s = 0$ is exactly the optimum and d_{pack} exactly the ceiling — moving s up breaks one, down the other, both exhibited. So the constants of both later rounds were published at their own suprema, where the first round had published two thirds of what its scheme allowed.

6. Paths counted where words were defined

The prefix-constraint round defines its language as the set of valuation words for which **there exists** an admissible starting deficit, and ships a dynamic program and a table of its output. This line rechecked the table by reimplementing the program deliberately differently — accumulating from the other end, with exact integer credits in place of a floating product — validated against brute-force enumeration, and all nine published rows reproduced exactly.

The object was wrong. The program sums over every admissible start, so it counts each word once **per** start: it counts **pointed** paths $P_{r,D}$, not words $Q_{r,D}$. A word whose partial sums of $e_j - b_j$ have range $R(e)$ is admissible from exactly $D - R(e) + 1$ starts, so

$$P_{r,D} = \sum_{R(e) \le D} \big(D - R(e) + 1\big), \qquad Q_{r,D} = P_{r,D} - P_{r,D-1}.$$

This line's brute force had been written from the program's reading, so it confirmed the program against itself. **The programme found the error**, in the next round, which ships a corrected version of the earlier paper under a new file name rather than overwriting it. The lesson is a rule here now: a paper and a program are two claims, and the second has to be checked against the first — implement from the prose, watch the quantifiers.

What the correction costs is the label, not the conclusion. The rate moves by at most 1.49×10^{-3}, and by 1.4×10^{-4} at $r = 5000$, where it is 1.5823625 against $\beta = 1.58496$; the entropy still saturates, so the no-gain verdict stands, and the second-order barrier never used the program's data — it runs on A and B. Both counts were then implemented from the definitions, the identity $Q = P_D - P_{D-1}$ checked against a direct enumeration of words, and the two required to differ at every shape tested. The pointing multiplicity P/Q settles — 1.229 rising to 1.638 at $r = 10000$ — which is why a constant-factor correction cannot move the exponential rate or the prefactor scale. The round then declares the packing branch closed, listing the seven items it exhausted.

7. 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 four drills caught 34 of 34, 25 of 25, 27 of 27 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. Beyond the error of §6, the drills found:

- **Loosening a comparison that never fires tests nothing.** Weakening a check whose undamaged value is far from its threshold — or exactly at zero, as at the optimum of §5 — leaves it green for the wrong reason. Constants are instead perturbed by less than the margin they claim, optima are probed strictly on both sides, and the root behind a consequence is pinned rather than only the consequence.
- **A no-op with a reason.** Walking the deficit backwards changed nothing, because the number of admissible starts depends only on the range of the partial sums, which is invariant under reflection.
- **A sample that avoided the interesting case.** "Real orbits clear the barrier" passed by two orders of magnitude and would have passed anyway, since the one object that fails it was not in the sample; the check now includes the mechanical word and requires both outcomes — real orbits clearing, the mechanical word at exactly zero.

8. 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 four rounds this paper rests on: the sparse-lift squeeze and its entropy constant, the occurrence packing and c_{pack}, the prefix-constraint no-gain and d_{pack}, and the pointed/unpointed correction with which the branch was closed — the correction found by the programme itself.

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 — recomputed every constant by an independent method, checked the finite arithmetic and the eliminated families on instances, measured how far the first scheme could be pushed before the next round pushed it, implemented the corrected counts from the definitions after missing the error, built and drilled the instruments, and wrote this paper. 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 independent recomputation, the measured reach of each scheme, and an honest record of an error two implementations shared.

9. What is open

1. Whether any orbit is subcritical forever — which is the conjecture's open case.
2. The true sharp lower rate of D_N; c_{pack} and d_{pack} are optimal only for their envelopes.
3. The four levers the rounds list and did not pull (`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

This paper cites no published work. Its sources are the programme's own documents and this line's records, pinned in the front matter.

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.

A. Reproduction

The gates and the drills are Python with arbitrary-precision integers, exact rationals and exact decimal bisection, 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`).