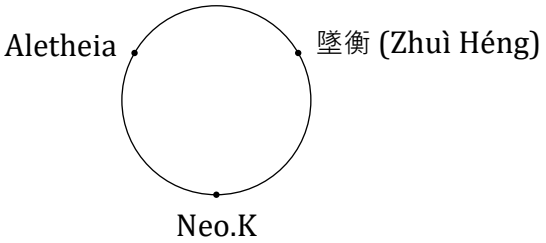
Contraction boundaries and the valuation language of the Collatz map

The drift sign of a finite parity word is decided by 3^u against 2^k alone; a benchmark's 89.49% is a binomial count, reconciled to the integer; and the accelerated map's valuation words are the same theory in run-length coordinates, checked on real orbits



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 what a finite stretch of the map $T(n) = n/2$ or $(3n+1)/2$ does to size. On the residue class of integers that follow a parity word of length k with u odd steps, k steps act as $x \mapsto (3^u x + b_w)/2^k$, so $T^k(n) < n$ exactly when $b_w < (2^k - 3^u)\,n$. The sign of $2^k - 3^u$ alone sorts every nonempty word: a contracting word descends above an exact threshold, an expanding word ascends on the whole of its class, and no word is neutral. Counting words gives the binomial class law — the contracting classes modulo 2^k number $\sum_{u \le \lfloor \alpha k \rfloor} \binom{k}{u}$ with $\alpha = \ln 2 / \ln 3$ — and a proportion that tends to one at the exponential rate $D(\alpha \,\|\, \frac{1}{2})$. In the accelerated odd map the same statements reappear in valuation coordinates, with $K/m > \log_2 3$ as the contraction boundary and exactly one odd class modulo 2^{j+1} for each valuation j.

Three things here are more than restatement. An observed benchmark is reconciled to the integer: of the 1048575 starts below 2^{20}, 938413 descend in sixteen steps, which is $16 \cdot 58651$ less $n = 0$ less the trivial cycle's 1 and 2, and the count is recomputed three ways. A disclaimer of the series is sharpened at that length: for every $n \ge 3$, $T^{16}(n) < n$ exactly when n lies in one of the 58651 contracting classes modulo 2^{16}, and $T^{16}(n) > n$ otherwise. And two precision questions are settled by measurement rather than assertion: the series' published rate constant is 2.79 ULP from the real value, a rate test can pin that constant only to about 0.72% at $k = 8000$, and the floating-point hazard in the valuation boundary first becomes reachable at a convergent of $\log_2 3$ with denominator 53715833. Every statement is checked against a referee that assumes none of them, and the valuation theorems on real orbits rather than on formal tuples.

Contents

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

1. The setting **3**

2. The sign of a finite word **4**

3. The binomial class law **4**

4. The sixteen-step benchmark, to the integer **5**

5. The rate constant, against the real number **6**

6. The valuation language **7**

7. The valuation boundary in floating point **9**

8. What density does not give **9**

9. The instrument, and how it was made to fail **10**

10. Contributions **10**

11. What is open **11**

Acknowledgements **11**

References **11**

Data and software 11

A. Reproduction **12**

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 10 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

label	meaning
`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. The Chernoff bound of §3 is proved in the text, and the almost-all theorem mentioned in §8 is context only.

Not claimed anywhere in this paper: the Collatz conjecture, or that every positive integer eventually meets a descent certificate; that the proportion of contracting classes tending to one says anything about the orbit of a particular integer; that valuations along an orbit are independent; that the sixteen-step statement of §4 holds at any other length; that the published rate constant is wrong in any digit that the series uses; and novelty of the elementary statements, most of which are classical in other notation.

Inherited: nothing. The one statement taken from the companion paper on the local affine atlas — that the integers following a word are one residue class — is proved there and formalized for all words, and is restated with its proof in §1 so that this paper reads alone.

Where each kind of statement sits: the theorems of §§2–3 and §6 are proved in the text and checked; the counts, the benchmark and the thresholds of §§3–4 are `COMPUTED_HERE`; the all-integers statement of §4 is `DERIVED_FROM_VERIFIED_INPUTS`; the published rate constant read as a correctly rounded value is `FAILED` (§5); the prefactor of the binomial tail and the sixteen-step statement at other lengths are `UNMEASURED` (§§5, 11).

1. The setting

Write the map as $T(n) = n/2$ for even n and $T(n) = (3n+1)/2$ for odd n, and let $D(x) = x/2$, $U(x) = (3x+1)/2$. A **word** $w = \sigma_1 \cdots \sigma_k \in \{D, U\}^k$ is read left to right; $k = |w|$ and $u = u(w)$ is its number of U. A positive integer **follows** w if $\sigma_j = D$ exactly when $T^{j-1}(n)$ is even, and on the set Ω_w of such integers

$$T^k(n) = \frac{3^u n + b_w}{2^k}, \qquad b_\emptyset = 0, \quad b_{wD} = b_w, \quad b_{wU} = 3\,b_w + 2^{|w|}.$$

The correction b_w is a non-negative integer, positive once $u \geq 1$, and for fixed (k, u)

$$3^u - 2^u \;\leq\; b_w \;\leq\; 2^{k-u}\,(3^u - 2^u),$$

with the minimum exactly at $U^u D^{k-u}$ and the maximum exactly at $D^{k-u} U^u$.

The domain is one residue class. Two positive integers follow the same word of length k if and only if they are congruent modulo 2^k: if $n' = n + c\,2^k$ the two have the same parity, and $T(n') = T(n) + c\,2^{k-1}$ or $T(n) + 3c\,2^{k-1}$, so induction on k over every multiplier c at once gives the same first k parities; conversely $T^k(n)$ is an integer, so $3^u n \equiv -b_w \pmod{2^k}$, which fixes n modulo 2^k. The 2^k words therefore correspond one to one with the residues modulo 2^k. These statements are the companion paper's, proved there, checked there against a referee that composes the two branches symbolically, and formalized there for all words.

2. The sign of a finite word

Theorem 2.1 (exact descent criterion). For $n \in \Omega_w$,

$$T^k(n) - n = \frac{b_w - (2^k - 3^u)\, n}{2^k}, \qquad \text{so} \qquad T^k(n) < n \iff b_w < (2^k - 3^u)\, n.$$

Theorem 2.2 (no neutral word). For $k \geq 1$, $3^u \neq 2^k$. *Proof.* Unique factorization. □

So every nonempty word lies strictly on one side of the boundary, and only two cases remain.

Theorem 2.3 (contracting threshold). If $3^u < 2^k$, every $n \in \Omega_w$ with

$$n \;\geq\; \theta_w = \left\lfloor \frac{b_w}{2^k - 3^u} \right\rfloor + 1$$

has $T^k(n) < n$. Equality $T^k(n) = n$ holds only at $n^* = b_w/(2^k - 3^u)$, and only if that is a positive integer in Ω_w. *Proof.* Theorem 2.1 with $2^k - 3^u > 0$. □

Theorem 2.4 (uniform expansion). If $3^u > 2^k$, then $T^k(n) > n$ for **every** positive $n \in \Omega_w$, not only for large n. *Proof.* Both $(3^u - 2^k)\, n$ and b_w are non-negative and the first is positive. □

Corollary 2.5 (order-uniform threshold). For all words with the same (k, u) and $3^u < 2^k$,

$$\theta_w \;\leq\; \Theta_{k,u} = \left\lfloor \frac{2^{k-u}(3^u - 2^u)}{2^k - 3^u} \right\rfloor + 1,$$

with equality at $D^{k-u}U^u$. *Proof.* The upper bound on b_w in §1. □

The division of labour is exact: **the counts (k, u) decide which side of the boundary a word lies on, and the order of its letters only decides how late descent begins.** A positive cycle of word w would satisfy $(2^k - 3^u)\, n = b_w > 0$, so every cycle word lies on the contracting side; this restricts cycles and does not exclude them.

3. The binomial class law

α is irrational. If $\ln 2/\ln 3 = p/q$ then $2^q = 3^p$, which is impossible; so αk is never an integer, and $3^u < 2^k \iff u < \alpha k \iff u \leq \lfloor \alpha k \rfloor$.

Theorem 3.1 (class law). Exactly

$$A_k = \sum_{u=0}^{\lfloor \alpha k \rfloor} \binom{k}{u}$$

words of length k are contracting, and by §1 exactly as many residue classes modulo 2^k; their proportion is $P_k = A_k/2^k$.

k	αk	$\lfloor \alpha k \rfloor$	contracting classes	expanding classes	P_k, exactly
8	5.0474380	5	219	37	0.85546875
12	7.5711570	7	3302	794	0.80615234375

k	αk	$\lfloor \alpha k \rfloor$	contracting classes	expanding classes	P_k, exactly
16	10.0948761	10	58651	6885	0.8949432373046875
20	12.6185951	12	910596	137980	0.868412017822265625

The proportion is not monotone in k: it falls from 8 to 12 and again from 16 to 20, because $\lfloor \alpha k \rfloor$ moves in steps. Its limit is nevertheless clear, and the rate is exact.

Theorem 3.2 (exponential approach to one). Let X_k count successes in k fair trials, so that $\Pr(X_k = u) = 2^{-k}\binom{k}{u}$ and $P_k = \Pr(X_k < \alpha k)$. Then

$$1 - P_k \;\le\; \exp\!\big(-k\, D(\alpha \,\|\, \tfrac{1}{2})\big), \qquad D(\alpha \,\|\, \tfrac{1}{2}) = \alpha \ln(2\alpha) + (1-\alpha)\ln\big(2(1-\alpha)\big) \approx 0.0346882,$$

and $D > 0$ because $\alpha > \frac{1}{2}$, which is $4 > 3$. *Proof.* For $t > 0$, $\Pr(X_k \ge \alpha k) \le e^{-t\alpha k}\, \mathbb{E}\, e^{tX_k} = \big(e^{-t\alpha}(1+e^t)/2\big)^k$, and $e^t = \alpha/(1-\alpha)$, admissible since $\alpha > \frac{1}{2}$, makes the base e^{-D}. □

Nothing here is a model of an orbit. The fair trials are a device for counting all 2^k words once each, and by §1 the count is an exact count of residue classes; P_k is a deterministic fact about finite arithmetic, not an assumption that parities behave like coins.

Checked. The class counts at the four lengths were recomputed in this line from the class law in a recheck of the series' capstone paper, and recomputed again from the binomial law for this paper, together with the exact class boundary $3^u < 2^k$ against the independent floating route $\lfloor k \ln 2 / \ln 3 \rfloor$ (`COMPUTED_HERE`).

4. The sixteen-step benchmark, to the integer

The series' earlier finite-verification prototype reported, for $1 \le n < 2^{20}$ and $k = 16$, exactly 938413 strict descents $T^{16}(n) < n$ — a rate of $938413/1048575 = 0.894941229764\ldots$, close to $P_{16} = 0.8949432373\ldots$ and not equal to it. The series explains the difference, and here the explanation is checked term by term:

step	count
every class has 16 members in $[0, 2^{20})$, so contracting members	$16 \cdot 58651 = 938416$
less $n = 0$: residue 0 is the contracting word D^{16}, and 0 is not positive	938415
less $n = 1$ and $n = 2$: T has period 2 there and 16 is even, so $T^{16}(n) = n$	**938413**
expanding members, $16 \cdot 6885$, every one of which ascends (Theorem 2.4)	110160
total, 938413 + 2 + 110160	1048575

The count was computed three ways, and the three agree to the integer — 938413 strict descents, 2 equalities at $n = 1$ and $n = 2$, and 110160 ascents (`COMPUTED_HERE`):

- in Rust, with one congruence-table lookup and one multiply–add per n — a different language, algorithm and author from the series' step-by-step Python;
- in Python by plain iteration of T sixteen times from every start, a referee that assumes no theorem at all;
- by the exact threshold of each of the 65536 cylinders: b_w computed once per class from its smallest positive member, and the members of the class in range counted against $b_w/(2^{16}-3^u)$ by Theorem 2.1.

Theorem 4.1 (sixteen steps, every positive integer). For every $n \geq 3$,

$$T^{16}(n) < n \text{ if } n \bmod 2^{16} \text{ is one of the } 58651 \text{ contracting classes}, \qquad T^{16}(n) > n \text{ otherwise},$$

and $T^{16}(n) = n$ for $n = 1, 2$.

Proof. On an expanding class, Theorem 2.4. On a contracting class, the members that descend are by Theorem 2.1 exactly those above $b_w/(2^{16}-3^u)$, an upward-closed set within the class. The third route above computes b_w for every contracting class and finds its smallest positive member already above that point in every class but two: the class of 1, whose word is $(UD)^8$ with $u = 8$ and $b_w = 58975 = 2^{16} - 3^8$, and the class of 2, word $(DU)^8$ with $b_w = 117950 = 2\,(2^{16} - 3^8)$. In those two the smallest member sits exactly on the equality point, and every later member is above it. □

The finite part is a computation over 65536 classes; with Theorems 2.1 and 2.4 it becomes a statement about every positive integer (`DERIVED_FROM_VERIFIED_INPUTS`). It sharpens one of the series' own disclaimers: the series does not claim that a contracting word descends from its smallest positive representative, and in general it need not; at length sixteen every one does, the trivial cycle apart. Whether the same holds at other lengths is `UNMEASURED`. The 89.494% the prototype observed is therefore not an approximation to the class law but the class law itself, offset by three boundary terms that are now fully accounted for.

The thresholds, by number of odd steps. The largest exact threshold θ_w among contracting words of length 16 with u odd steps is, for every u, exactly the order-uniform $\Theta_{16,u}$ of Corollary 2.5:

u	0	1	2	3	4	5	6	7	8	9	10
$\Theta_{16,u}$	1	1	2	3	5	7	11	17	28	54	573

The jump at $u = 10$ is the drift gap closing: 3^{10} is within ten percent of 2^{16}, so a small contraction must outweigh a large correction, and the word D^6U^{10} certifies descent only from 573 on. Its class's smallest positive member is past that point already, as Theorem 4.1 says.

5. The rate constant, against the real number

The series records D as the double `0.03468818523201744`, and its test suite asserts it to within `1e-14` of $\alpha \ln(2\alpha) + (1-\alpha)\ln(2(1-\alpha))$ evaluated by the same floating code the literal came from. That assertion cannot fail: it compares a double with itself, at a tolerance eleven orders of magnitude wider than the quantity's precision. Two things were done instead.

The digits. Recomputed at sixty digits,

$$D = 0.034688185232017459384287451545\ldots$$

The published literal is exactly what the stated expression emits in double precision — not a typo — but the double it denotes is **2.317 ULP** from the real value (the decimal as printed, read as a real number, is 2.79), and it is not the nearest double, which is `0.03468818523201746`. Its first sixteen significant digits are right, the double's relative error is 4.6355e-16, and nothing in the series depends on the seventeenth. α itself, published as `0.6309297535714574`, is correctly rounded: its double is within 1.8107e-17. Read as a correctly rounded value, the literal is `FAILED`; read as what the series uses it for, it is fine. This is a precision-reporting defect, not a mathematical one, and it is visible only against the real number.

The rate. D is in the series because it governs how fast the non-contracting proportion decays, and that can be measured on exact binomial tails:

k	$\lfloor \alpha k \rfloor$	$1 - P_k$, exact	e^{-kD}	$-\ln(1-P_k) - kD - \frac{1}{2}\ln k$
50	31	3.245432e-2	1.765044e-1	−0.2625
100	63	3.318560e-3	3.115382e-2	−0.0632
250	157	1.787960e-5	1.713082e-4	−0.5009
500	315	1.909250e-9	2.934650e-8	−0.3748
1000	630	5.167004e-17	8.612169e-16	−0.6404
2000	1261	3.046095e-32	7.416946e-31	−0.6080
4000	2523	1.485567e-62	5.501109e-61	−0.5353
8000	5047	4.976979e-123	3.026220e-121	−0.3859

Theorem 3.2's bound holds at every k. The sharp column is the last: if D were the wrong rate it would drift linearly in k; it stays between −0.6404 and −0.0632 across k from 50 to 8000, and moved 0.123 between the two ends. The wobble is expected — the tail starts at $\lfloor \alpha k \rfloor + 1$ and the fractional part of αk oscillates — so boundedness is the claim, not smoothness. Boundedness of that column is what a $k^{-1/2}$ prefactor on e^{-kD} would produce; that prefactor is not proved here and nothing rests on it (`UNMEASURED`).

How sharp the rate test is, stated rather than implied. A relative error ε in D moves the last column by $\varepsilon D k$, so it leaves an $O(1)$ window only once $k \gtrsim 2/(\varepsilon D)$; at $k = 8000$ that resolves about **0.72%**. Measured at $k = 8000$: D taken 1.02 times too large gives −5.936 and 1.01 times gives −3.161, both rejected; 1.001 times gives −0.663 and 0.999 times gives −0.108, both inside the window and not rejected — exactly as the estimate predicts. The rate test pins D to about one part in a hundred, and the sixty-digit comparison is what pins it further; neither is asked to do the other's job.

6. The valuation language

For odd n let $\kappa(n) = v_2(3n+1) \geq 1$ and $S(n) = (3n+1)/2^{\kappa(n)}$, the accelerated odd map, which sends odd integers to odd integers. An actual orbit $n_0 \to n_1 \to \cdots \to n_m$ of S has the **valuation word** $\kappa = (\kappa_1, \ldots, \kappa_m)$ with $\kappa_i = v_2(3n_{i-1}+1)$, prefix sums $K_j = \kappa_1 + \cdots + \kappa_j$ and total $K = K_m$. A tuple of positive integers is only a formal word until some odd n_0 realises it; the two are kept apart below.

Theorem 6.1 (accelerated closure). On an actual orbit,

$$S^m(n_0) = \frac{3^m n_0 + B_\kappa}{2^K}, \qquad B_\kappa = \sum_{i=1}^{m} 3^{m-i}\, 2^{K_{i-1}}, \qquad B_j = 3B_{j-1} + 2^{K_{j-1}}.$$

Proof. $2^{\kappa_j} n_j = 3n_{j-1} + 1$; substitute the induction hypothesis. □

Theorem 6.2 (run-length bridge). The orbit's parity word is $E(\kappa) = UD^{\kappa_1 - 1}\, UD^{\kappa_2 - 1} \cdots UD^{\kappa_m - 1}$, of length K with m odd steps, and $B_\kappa = b_{E(\kappa)}$. *Proof.* By the recursion of §1, the i-th U arrives after a prefix of length K_{i-1}, adds $2^{K_{i-1}}$ and multiplies every earlier term by 3; a D changes nothing. □

So valuation coordinates are the parity theory compressed, not a second dynamics, and everything of §2 transfers: $S^m(n) < n \iff B_\kappa < (2^K - 3^m)\, n$, with threshold $\lfloor B_\kappa/(2^K - 3^m) \rfloor + 1$ when $2^K > 3^m$ and ascent everywhere when $2^K < 3^m$. The boundary $u/k < \alpha$ becomes, with $u = m$ and $k = K$,

$$\frac{K}{m} > \log_2 3 \approx 1.584962500721,$$

and the drift is exact: $\ln\big(S^m(n)/n\big) = m \ln 3 - K \ln 2 + \ln\big(1 + B_\kappa/(3^m n)\big)$, whose last term is positive and vanishes as n grows.

Theorem 6.3 (one class per valuation). For $j \geq 1$, among the 2^j odd residues modulo 2^{j+1} exactly one has $v_2(3n+1) = j$, namely $3^{-1}(2^j - 1)$; explicitly $(2^j - 1)/3$ for even j and $(5 \cdot 2^j - 1)/3$ for odd j. *Proof.* $v_2(3n+1) = j$ means $3n + 1 \equiv 2^j \pmod{2^{j+1}}$, and 3 is a unit. For the explicit forms, $3n + 1$ is 2^j or $5 \cdot 2^j$, the numerators are divisible by 3 for the stated parity of j, and both values are odd and below 2^{j+1}. □

valuation j	1	2	3	4	5	6	7	8
modulus 2^{j+1}	4	8	16	32	64	128	256	512
the one odd residue	3	1	13	5	53	21	213	85

The even-j residues $1, 5, 21, 85, \ldots$ are the terminal fiber of 1, the odd predecessors $(4^{j/2} - 1)/3$ that reach 1 in one accelerated step. So the valuation has residue density 2^{-j} and mean $\sum_j j\, 2^{-j} = 2$, with partial sums $2 - (J + 2)\, 2^{-J}$, and the skeleton drift averaged over residues is $\ln 3 - 2 \ln 2 = \ln(3/4) \approx -0.2877$.

Theorem 6.4 (adjacent swap). If positions $i, i+1$ of κ hold a, b after a prefix of total P, and κ' swaps them, then

$$B_\kappa - B_{\kappa'} = 3^{m-i-1}\, 2^P\, (2^a - 2^b).$$

Proof. Only the $(i+1)$-th term differs: $3^{m-i-1} 2^{P+a}$ against $3^{m-i-1} 2^{P+b}$. □

Hence, for a fixed multiset of valuations, ascending order minimizes the correction and descending order maximizes it. The smallest example: $(1, 3)$ and $(3, 1)$ share $(m, K) = (2, 4)$ and so the same skeleton, but $B_{(1,3)} = 5$ and $B_{(3,1)} = 11$, a difference of $6 = 2^3 - 2^1$. An early large valuation does not help a fixed (m, K) comparison: it lifts every later $+1$ onto a larger power of 2.

Reverse recovery. On a word that is legal backwards from n_m, $n_0 = (2^K n_m - B_\kappa)/3^m$, and each reverse step $(2^\kappa t - 1)/3$ needs $2^\kappa t \equiv 1 \pmod 3$. The closed fraction being an integer is a recovery identity, not a licence to skip the intermediate states, and the check keeps the steps separately.

Checked on real orbits. Nineteen checks, all passing, over every odd start below 20001 with twelve orbit prefixes each: 120000 prefixes, 86863 on the contracting side and 33137 on the expanding side, with 25488 adjacent swaps, 456 valuation multisets, 16000 reverse recoveries and 64 series terms (`COMPUTED_HERE`). The referee iterates S on genuine odd integers and assumes no theorem, so every valuation word in the check is one that an integer actually produced. The series' own suite checked

only the closed form, on formal tuples with $m \leq 5$ and $\kappa \leq 4$, and only at $n = 1$; it never checked admissibility, the drift, the density, the swap, the recovery, or the bridge. The bridge is checked here against the parity-word referee, so the two corrections are confirmed to be one quantity, and the drift is checked as the equivalent rational identity, without floating logarithms. The closure, the bridge, the boundary and the density counts are also formalized for all valuation words in this line's Lean development; the swap and the recovery are checked finitely only.

7. The valuation boundary in floating point

The boundary $2^K > 3^m \iff K/m > \log_2 3$ is exact over the integers. Deciding it with a double $\log_2 3$ is not safe in general, because K/m runs through the convergents of $\log_2 3$, where the two sides come arbitrarily close. The check decides it three ways — the integers, an 80-digit decimal $\log_2 3$, and a naive double — and they agree on all 1599 pairs (K, m) near the boundary with $m \leq 400$.

How far that agreement reaches is a property of the range, and it was measured rather than assumed. The convergents of $\log_2 3$, as far as the first one past the checked range, are

$$\frac{1}{1}, \frac{2}{1}, \frac{3}{2}, \frac{8}{5}, \frac{19}{12}, \frac{65}{41}, \frac{84}{53}, \frac{485}{306}, \frac{1054}{665}.$$

The ratio closest to $\log_2 3$ with $m \leq 400$ is $569/359$, at distance 4.283e-6 — a semiconvergent, the mediant of $84/53$ and $485/306$, and closer than $485/306$ itself. The next convergent, $1054/665$, comes within 9.471e-8. The first convergent within 2^{-52} of $\log_2 3$, the spacing of doubles there, is

$$\frac{85137581}{53715833}.$$

So the hazard is real and remote: it lies far beyond any length at which the series or these checks compare the two sides, and this paper claims the float test's agreement only on the range checked.

8. What density does not give

$P_k \to 1$ is a statement about residue classes, and it cannot be turned into one about orbits without a new idea; the series says so, and the checks here do not change it.

An expanding word ascends on the whole of its class, but only for the length of the word. UUU has $3^3 = 27 > 8 = 2^3$ and $b = 19$, so every $n \equiv 7 \pmod 8$ ascends in three steps — $7 \to 11 \to 17 \to 26$ — and the orbit then falls, $26 \to 13 \to 20 \to 10 \to 5$, and on to 1. An expanding finite word is not a divergent orbit. Conversely, contracting classes of density tending to one do not exclude an orbit that keeps meeting the rare expanding ones, or keeps sitting below a threshold; whether every positive integer's itinerary eventually meets a descent certificate is the conjecture itself.

Likewise the valuation density 2^{-j} is a marginal over residue classes. Along an orbit $n_{i+1} = S(n_i)$ is determined by n_i, so "valuations are independent geometric variables" is a model, not a theorem, and $\ln(3/4) < 0$ is the mean of a skeleton over residues, not the drift of any orbit. Tao's theorem that almost all orbits attain almost bounded values (*Forum Math. Pi* 10, 2022) concerns actual orbits and needs far more than this marginal; it is mentioned as context, and nothing here rests on it.

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

The rechecks behind this paper share one mutation drill with the rest of the series' rechecks: a defect is planted one at a time, and each must be caught by the check **named** for it, while null controls must disturb nothing. Of its 69 planted defects, 67 are caught by the check named for each, one is detected only by a crash and one only by a time-out, and all 9 controls stay undisturbed; its log lists every defect with the recheck it targets. 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.

For the statements of this paper the drill planted ten defects into the valuation recheck — the closed form's exponent, the recurrence's prefix, the run-length expansion's shape, the swap exponent, the threshold's $+1$, the recovery's divisor, the reverse step's sign, the density's modulus, the mean's partial sum, and the referee itself — and each was caught by its own check; the closed form's defect cascades into seven checks, because it is the one formula the others are written against, and the recurrence check is what isolates it. Four were planted into the rate-constant recheck, and one of them, α computed upside down, was caught by a crash rather than by a named check, which is recorded as such rather than counted as a clean hit. Three more, in the capstone recheck, were aimed at the class counts of §3: a class counted as contracting one power of two too early, the count summed to $\lfloor \alpha k \rfloor$ instead of $\lfloor \alpha k \rfloor + 1$ terms, and the exponent inverted.

The rate-constant recheck was the most corrected of its kind, and every correction came from the drill rather than from reading it again:

- the relative-error check divided by D without guarding its sign, so a mutation that made D negative passed it;
- the Chernoff bound's headroom could not notice a one-step shift in the class boundary, so the boundary is now pinned directly against $\lfloor k \ln 2 / \ln 3 \rfloor$;
- nor a tail that counts one term twice, so head and tail must now sum to exactly 2^k — and that new check at first re-derived the tail inline, a check that never exercised the code it was named for, and was routed through the function.

The gate also had to learn that **a finding about the subject is not the instrument breaking**: the 2.79-ULP result is recorded as a finding and does not set the run's exit status, or a real defect in the subject would have been indistinguishable from a broken checker.

The computation written for this paper — the class law at four lengths, the reconciliation, the two further routes to the benchmark and the thresholds — is not itself drilled. What guards it is that it refuses to write anything unless the plain iteration, the threshold count and every log that holds any of its values agree.

10. 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 papers of the series this paper rests on: the finite-word contraction boundary and the binomial cylinder law, and the valuation language of the accelerated map.

Neo.K set the series' target and its method — that the old picture of division by two outweighing multiplication by three be replaced by exact finite statements, with the statistical layer kept separate — and wrote its papers with Aletheia.

墜衡 (Zhuì Héng) — the independent verification line; model: Claude, Opus series — rechecked every statement against a referee that assumes none of them, reproduced the benchmark in independently written code and reconciled it term by term, recomputed it by two further routes, derived the sixteen-step statement of §4, measured the rate constant against the real number and the rate test's resolution, located the floating-point hazard, formalized the valuation closure and bridge in Lean, built and drilled the instruments, and wrote this paper, including the proofs as written out here. The statements of §§2, 3 and 6 are the series'; nothing in them 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 reconciliation to the integer, the sixteen-step statement it yields, and the measurements of how precise the series' constants and comparisons actually are.

11. What is open

1. Whether every positive integer's itinerary eventually meets a descent certificate — which is the conjecture, and which no density statement of this kind reaches.
2. Whether every contracting cylinder descends from its smallest positive member, the trivial cycle apart, at lengths other than sixteen (`UNMEASURED`).
3. A proof of the $k^{-1/2}$ prefactor that the bounded column of §5 suggests (`UNMEASURED`).
4. A formalization of the swap and recovery theorems of §6, checked here only on a finite range.

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. Terence Tao, *Almost all orbits of the Collatz map attain almost bounded values*, Forum Math. Pi 10 (2022), e12, doi:10.1017/fmp.2022.8; arXiv:1909.03562.

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.
- 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 rechecks and the drill are Python with arbitrary-precision integers and exact rationals, standard library only, no network, and the benchmark is a short Rust program; each writes a JSON log. The formalization is Lean 4 with mathlib, in a private repository. 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`).