A Global Occupation Conjecture for the Accelerated $3x+1$ Map: Residue–Time Duality, Divergent-Orbit Sparsity, and the Pointwise Realizer Problem

Elias De Jesús
Independent Researcher
ORCID: 0009-0007-0190-9143
dejesuselias10@gmail.com

Revision 9, 2026

Abstract

We study the Effective Occupation Conjecture (EOC) for the accelerated $3x+1$ map $T(m) = (3m+1)/2^{a(m)}$, $a(m) = \nu_2(3m+1)$. Writing $R_n = S_n - n\log_2 3$ for the drift of the accumulated 2-adic valuation and $\mathcal{O}_c(m_0) = \#\{n \geq 0 : R_n \leq c\}$, EOC asserts $\mathcal{O}_c(m_0) = O(\log m_0)$ for every seed and every corridor width c.

This paper organizes the programme around four axes — occupation, residue, time, and divergence — and the exact bridges between them. Record chronology identifies the residue and time axes precisely: the least-realizer frontier $r_{\min}(N, c)$ is the inverse function of the single-window confinement lifetime $L_c(m)$, so every nondecreasing pointwise lifetime bound transfers to a realizer floor, on a graded scale that we tabulate. On the divergence side we retain the strongest unconditional input: via Garcia–Tal's collision-free window theorem and Curry's explicit quantitative form, every divergent orbit has summable reciprocals and admits no eventual logarithmic drift floor $R_n \geq -B\log_2 n + O(1)$ for any $B < 1/\beta_* \approx 1.0358567$. Curry's input plays a second, qualitative role: after the last global maximum of R, a hypothetical divergent orbit has an infinite zero-corridor tail, so *universal drift exit* — every positive seed eventually attains $R_n > 0$ — is equivalent to the exclusion of divergent orbits, while the drift-exit argument supplies no cycle exclusion: no implication from it to cycle-freeness is established here, and none is claimed in either direction (Rem. 6.17).

We then record what a sustained programme of structural audits has closed. Finite information cannot decide infinite realization: every exact-realizer cylinder contains arbitrarily large positive integers, so $\forall N\, \exists m_N$ is free and all content lies in $\exists m\, \forall N$. Nor do the *known divergent-orbit* valuation-word shadows decide it: we exhibit an explicitly indexed uncountable family of zero-confined words with linear negative drift, summable $\sum 2^{R_n}$ and $O(G)$ occupation shadow, of which all but countably many have no positive-integer realizer. Word-level criteria in general are not excluded, and one is exact (Rem. 9.6). Several apparently independent coordinates — normalized carry, lift digits, terminal zero gap, transport budgets — are invertible reparameterizations of the exact realizer, and the Archimedean and 2-adic readings of the natural divisibility object are complementary factors of one identity rather than independent constraints. Launch renormalization remains a valid bridge from finite words to actual positive-orbit dynamics, but its corridor-uniform closing hypothesis is, once the corridor width is computed rather than assumed, essentially an $O(\log m)$-scale pointwise lifetime conjecture — it bounds the *injective* initial segment, so on its own it yields entry into *a* cycle, not arrival at 1; and deep negative drift is self-financing, since the ordinary height it creates is exactly the resource that permits the recovery valuations repaying it.

Neither EOC nor the Collatz conjecture is proved here. The surviving frontier is actual positive-orbit arithmetic at three nested levels: qualitative exclusion of an infinite zero-corridor tail, quantitative control of single-window lifetime, and control of separated returns sufficient for full occupation.

Non-claims. This paper proves neither the Collatz conjecture nor any tier of EOC. It does not prove an exponential realizer floor for arbitrary confined words. It does not promote any computational finding to a theorem: every claim is labelled with its status (§B). Where a statement reorganizes or reclassifies earlier work rather than proving something new, this is said explicitly.

Contents

1 Introduction **3**
- 1.1 Four axes, not three layers . . . 3
- 1.2 What Revision 9 changes . . . 4
- 1.3 What Revision 8 changes . . . 5
- 1.4 What Revision 7 changes . . . 6
- 1.5 What Revision 6 added . . . 7

2 Preliminaries **8**

3 The occupation axis: EOC and its logical status **9**
- 3.1 Four tiers . . . 9
- 3.2 The Collatz reduction . . . 9
- 3.3 The entropy scale . . . 10
- 3.4 EOC as an occupation-time shadow of Rozier's hypothesis . . . 10

4 The divergence axis **10**
- 4.1 An elementary 8/9 bound and its terminality . . . 10
- 4.2 Garcia–Tal/Curry sparsity . . . 11
- 4.3 Curry's second, qualitative role . . . 11

5 The residue axis **12**
- 5.1 Proved sectors . . . 12
- 5.2 The finite-defect sector: an unproved claim inherited from Revision 5 . . . 12
- 5.3 What the height invariant actually separates . . . 13
- 5.4 Endpoint depth and the frontier defect . . . 14

6 The time axis **14**
- 6.1 Three temporal scales . . . 14
- 6.2 Residue–time duality . . . 15
- 6.3 The general transfer principle . . . 15
- 6.4 Drift exit and the future-minimum theorem . . . 16
- 6.5 Universal drift exit and divergence . . . 17

7 Launch renormalization: the bridge to actual orbits **18**
- 7.1 Reclassifying the corridor-uniform hypothesis . . . 18

8 Self-financing recovery **20**

9 Boundaries of the programme **21**
9.1 Finite cylinders do not encode infinite realization 21
9.2 Coordinate sufficiency is not independent information 22
9.3 The known divergent-orbit word shadows do not characterize positive realization . . 22
9.4 Population sparsity is not pointwise placement 23
9.5 Deep drift creates its own recovery capacity 24
9.6 Single-window lifetime is not the same functional as occupation 24

10 The programme hierarchy **24**

11 The programme map **25**
11.1 Symbolic freedom, finite and infinite 25

12 Closed routes **25**
12.1 External arithmetic 26

13 Open problems **26**

14 Formal verification status **29**

15 Computational calibration **31**

16 Relation to prior work **31**

17 Conclusion **32**

A A complete proof of the carry budget, with an explicit constant **33**

A Band rigidity **36**

B Claim-status conventions **37**

1 Introduction

The accelerated $3x + 1$ map sends an odd positive integer to the next odd term of its Collatz trajectory,

$$T(m) = \frac{3m+1}{2^{a(m)}}, \qquad a(m) = \nu_2(3m+1) \geq 1.$$

Writing S_n for the accumulated 2-adic valuation after n accelerated steps and $R_n = S_n - n\alpha$ with $\alpha = \log_2 3$ for its drift below the critical slope, the *Effective Occupation Conjecture* asserts that the time an orbit spends at or below any fixed drift level is logarithmic in the seed. This single pointwise statement forces every orbit to be eventually periodic and — *once effective constants are supplied*, which Conjecture 3.1 does not itself provide — reduces Collatz to a finite cycle search (Theorem 3.6).

1.1 Four axes, not three layers

Revision 5 of this work was organized in three layers: EOC and its logical status; unconditional drift-axis progress; and the residue-axis reformulation. That organization is no longer adequate, for a specific reason: the residue axis and the confinement-lifetime question are not two problems

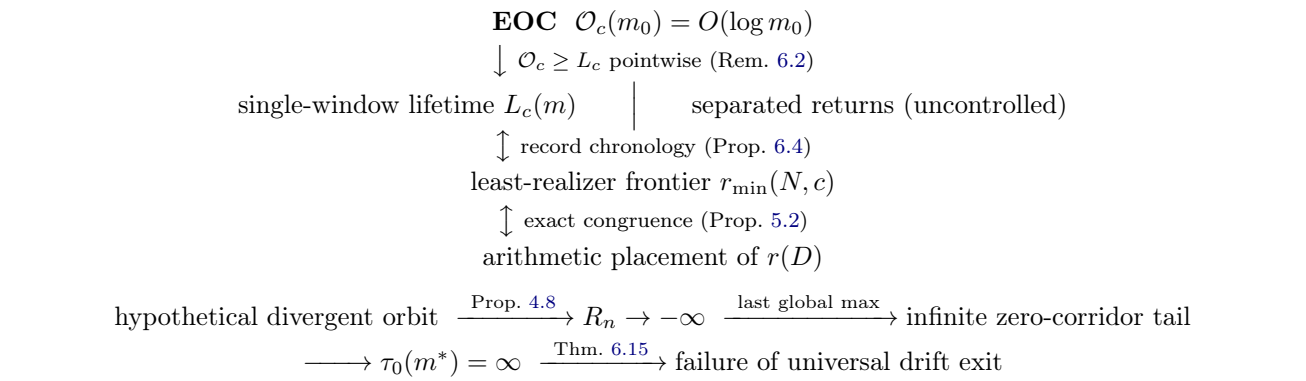


Figure 1: The programme map. Vertical double arrows are exact equivalences; single arrows are implications. The top arrow records the *pointwise* inequality $\mathcal{O}_c \geq L_c$; it does *not* assert that the two universal asymptotic bounds are inequivalent, which is open (Rem. 10.1). The lower chain is the qualitative divergence reduction. Every implication is referenced in §11.

but two coordinates on one problem, and a fourth, genuinely qualitative question — exclusion of divergent orbits — sits beside both. The paper is therefore organized around four axes.

Occupation axis (§3)
The functional $\mathcal{O}_c(m_0) = \#\{n \geq 0 : R_n \leq c\}$, counting *all* times in the corridor, including separated later returns. EOC lives here, and it is the strongest of the four.

Residue axis (§5)
The exact realizer congruence: which integers realize a given valuation word, and how small the least such integer, $r_{\min}(N, c)$, can be.

Time axis (§6)
The single-window lifetime $L_c(m)$ and the drift exit time $\tau_c(m)$. Record chronology makes the residue and time axes *exactly* dual.

Divergence axis (§4)
What is unconditionally known about a hypothetical divergent orbit: sparsity of its value set, summability of reciprocals, and exclusion of logarithmic drift floors.

The programme map in Figure 1 records the exact bridges. Every arrow in it is a theorem, and §11 gives the reference for each.

1.2 What Revision 9 changes

Revision 9 is an *erratum and literature sweep.* Nothing mathematical moves: no theorem, proof, conjecture, open problem, or epistemic tier changes, every statement labelled proved in Revision 8 is proved here by the same argument, every open problem stands open at the same tier, and the claim-status conventions of §B are untouched. What changes is attribution, verification scope, and the literature around four open problems.

A1. The drift exit time has a standard name. §16 claimed "$\exists n : S_n > \alpha n$" has none. It is Terras's *coefficient stopping time* κ [3, 4]; Proposition 6.11 is his $\kappa \leq \sigma$, and the $\tau_0 = \sigma$ diagnostic is the Coefficient Stopping Time Conjecture (Remark 6.12). *Verified:* Rozier–Terracol [5] read at

arXiv:2502.00948v5 — their Cor. 5.4 gives $t(n) = \tau(n)$ for all $2 \le n \le 2.8 \times 10^{19}$, superseding §15's scan, and their Thm. 3.2 is quoted as printed.

A2. The divergence chain is now formalized without hypotheses, in a different repository from the one §14 describes, and the scope of that is bounded: Props. 4.8, 4.11 and Thm. 6.15 unconditionally; Thm. 4.9 at the sharp $1/\beta_*$ *not* covered. *Verified:* every Lean name read at `e56950f`, each built with `#print axioms` reporting `propext`, `Classical.choice`, `Quot.sound`.

A3. Uniformity over orbits is the formalization's, not Curry's. Theorem 4.7 as quoted is uniform in the window start only. *Verified:* both integer certificates checked exactly.

A4. Proposition 4.8 is Curry's Prop. 3.1; the *[Revision 5]* label records entry into this paper, not priority. *Verified:* Curry's note retrieved and read.

B1. Rozier's LBH is Open Problem C stated uniformly over (j, q), and his Lemma 3.1 is this paper's Level 1 at $c = 0$ (§13). *Clarified: two critical ratios.* At the divergence-forced ones-ratio $1/\alpha$, LBH's exponent is $1 - H_2(1/\alpha) = 0.0500445$ per T-step, i.e. I_0 per accelerated step — the identity $I = \alpha(1 - H_2(1/\alpha))$ of §3. Rozier's $\gamma_* \approx 0.6090897$ is a different ratio, the Garcia–Tal/Curry sparsity exponent of Theorem 4.7, where $(1 - H(\gamma_*))/\gamma_* = 0.0568316$. *Verified at arXiv:1510.01610.* His $C \ge 0.980916600$ and the sharp form's $+\frac{3}{2}\log_2 N$ remain in different normalizations; no correspondence is asserted.

B2. Where a divergent word must live (Rem. 9.4): above the Dubickas complexity threshold $1.70951129\,n$ and at or above critical density, hence outside every class for which Φ-irrationality is proved. *Verified:* Dubickas Thm. 5 read first-hand previously and quoted verbatim from that record; Monks–Yazinski Thm. 2.7(b) read at source. The density half is in any case internal (Prop. 4.11).

B3. An external refinement reaches the endpoint Thm. 4.9 leaves open, without closing it (Rem. 4.10). *Verified:* Esm's theorem statements and its own strictness disclaimer read in its repository; the numerical value of γ was **not** found at source.

B4. A refereed three-term S-unit route is recorded in §12.1, consistent with the growing-term-count obstruction stated there. *Verified:* the full text read — the wave definition, $h = \nu_2(\mu + 1)$, the identity of Thm. 1.2(ii), and the bounded-height hypothesis of Thm. 1.2(iii).

C1. An AI-assistance section is added; it was missing from Revisions 5–8 and is now stated as in the companion notes.

Revision 8 is preserved unchanged; this revision supersedes it only in the items above.

1.3 What Revision 8 changes

Revision 8 is an *erratum and arithmetic sweep*. Nothing mathematical moves: no theorem, proof, conjecture, or epistemic tier changes, every statement labelled proved in Revision 7 is proved here by the same argument, every open problem stands open at the same tier, and the claim-status conventions of §B are untouched. What changes is one count, one convention, and some bookkeeping.

1. **An occupation count corrected** (Rem. 6.2). For $m_0 = 285175$ the paper printed $\mathcal{O}_1 = 97$; Definition 6.1 counts from $n = 0$ and $R_0 = 0$ lies in the corridor, so the value is $\mathcal{O}_1 = 98$. The example dates from Revision 5, so the misprint stood through Revisions 5–7. $L_1 = 14$ and

$\tau_1 = 15$ are unchanged, and so is the use made of the example. The remark now also records the episode decomposition $15 + 4 + 6 + 64 + 1 + 5 + 3 = 98$, which shows that 83 of the 98 occupation times are separated returns and sharpens the remark's point into Open Problem F.

2. **The post-arrival convention, stated** (Def. 6.1). L_c, $\mathcal{O}_c$ and τ_c are read along the full forward orbit, which continues at the fixed point because $a(1) = 2$ and $T(1) = 1$; R then rises by $2 - \alpha$ per step. The convention was implicit in every earlier revision and is now stated once, where the three functionals are defined.

3. **Every corridor number rechecked.** All remaining numerical values of $\mathcal{O}_c$, L_c, τ_c and $r_{\min}$ were recomputed in exact integer arithmetic and are confirmed, including $\tau_0(1) = 1$ (Prop. 6.11), $\mathcal{O}_0(1) = 1$ (Prop. 7.7) and the printed constants of §3.3 and Rem. 6.13. Two further bookkeeping errors are corrected. The self-financing margin $+1.7370$ (Rem. 8.2) is the minimum over $3 \le m < 200001$, not over all odd $m < 200001$, a range that also admits $m_0 = 1$. And the occupation-to-$\log_2$-seed ratio of §15 does *not* remain under 10. That claim is replaced by an exhaustive maximum over odd $3 \le m < 2^{33}$, computed in exact integer arithmetic: the ratio tops out at 11.3402 at $c = 8$, and already exceeds 10 at $c = 2$, while staying below $1/I(\alpha)$ throughout.

4. **The realizer modulus, confirmed and attributed** (Props. 5.2 and 6.4). The exact realizer class was already stated modulo 2^{S_N+1} and needed no correction; the erratum establishing that modulus, and the record-chronology reformulation of the single-window realizer floor, are now cited [25].

5. **Two textual defects fixed.** The one-density $\rho_c = N/S$ (Rem. 5.10) was used but never defined in Revisions 5–7, and a cross-reference to a "record table" in §15 pointed at something that does not exist — that section records frontier figures, not a table — and is reworded. The parenthesization hazard of §3.3 was checked and is already correct.

6. **Bibliography.** The two companion notes [27] and [28] gain their Zenodo DOIs, and [28] its full subtitle.

7. **Acknowledgments and a software reference** added, for the independent reimplementation of the least-realizer and occupation computations that reproduced Proposition 5.2, found the corrected $\mathcal{O}_1$ and supplied its episode decomposition [33].

Revision 7 is preserved unchanged; this revision supersedes it only in the six items above.

1.4 What Revision 7 changes

Revision 7 is a *consolidation of corrections*. Stated plainly: the verification record improved, several claims were narrowed, and universal drift exit remains open. Nothing here proves divergence exclusion or any tier of EOC.

1. **Verification improved.** Reciprocal summability (Prop. 4.8) is no longer an unformalized input: it is now *derived in Lean* from the windowed sparsity count (Rem. 14.1). The remaining external input in that chain is Theorem 4.7 itself, and the open hypothesis (DE) remains open; both are retained explicitly. Curry's note was also audited from the source text, and its two local defects — a typographical slip in the light-case bound and a missing integrality step — were identified and repaired; neither affects Proposition 4.8.

2. **Separated returns, made precise** (Open Problem F). The entry lemma pins every re-entry threshold below $\alpha - 1$ (Prop. 13.1); a *uniformly bounded* episode count is refuted constructively (Prop. 13.2); and the coarse per-episode accounting is shown lossy by a factor $\Omega(\log m_0)$ on the realized prefix (Rem. 13.3). Finite-prefix statements are kept separate from total occupation throughout.

3. **The corridor-uniform hypothesis, clarified — and a gap in it closed** (Rems. 7.5, 7.8). Revision 6 read the hypothesis as giving an $O(\log m)$ bound on the total stopping time *to* 1. It does not: it bounds the *injective* initial segment, so the orbit enters *a* cycle within $O(\log m)$ steps, and arrival at 1 needs cycle exclusion, which the hypothesis does not visibly supply. The EOC consequence is restated with that assumption explicit, and what survives without it is stated separately. Also: the analytic absorption inside Corollary 7.4 is identified as the unformalized layer, and corridor-uniformity must not be conflated with a bound at one fixed corridor.

4. **Scope overstatements corrected.** The closed-route table no longer labels fixed-modulus *terminality* as a formal result; and the cardinality barrier is stated for the *known divergent-orbit* shadows, not for word-level criteria in general — one such criterion is exact (Rem. 9.6).

5. **Hierarchy strictness audited** (Rem. 10.1). "No converse holds in general" and "Level 4 is strictly stronger than Level 3" are withdrawn: a pointwise example with $\mathcal{O}_c > L_c$ shows the quantities differ, not that the universal bounds are inequivalent. For the corridor-uniform form the implication in fact reverses, *given cycle exclusion*.

6. **Bibliography.** Four entries recovered their DOIs, and Bařina's entry — which had merged two distinct papers under one title — is split into [21] and [22], each with its own metadata.

   Revision 6 is preserved unchanged; Revision 7 supersedes it only where stated above.

1.5 What Revision 6 added

Revision 6 is primarily a *synthesis*. Its main additions are:

1. A general transfer principle (§6.3) turning any nondecreasing pointwise lifetime bound into a realizer floor, with the exponents worked out: logarithmic, polylogarithmic, polynomial and linear lifetime bounds give exponential, stretched-exponential, polynomial and linear frontiers respectively. EOC thereby becomes a *graded* programme rather than an all-or-nothing one.

2. The qualitative time axis (§6.4): the drift exit time τ_0, its exact integral characterisation $R_n > 0 \iff 3^n < 2^{S_n}$, and the equivalence between universal drift exit and the exclusion of divergent orbits.

3. A boundaries section (§9) collecting the strongest no-go results: the finite-cylinder barrier, the coordinate-collapse ledger, and the word-shadow cardinality barrier.

4. A corrected reading of the launch-renormalization closing hypothesis (§7) and the self-financing obstruction to temporal impossibility arguments (§8).

5. A closed-route map (§12) recording, for each mechanism tried, what it achieved and where it stops.

We are explicit about provenance throughout. Several statements that read as new are in fact consequences of identities already in Revision 5 — the future-minimum theorem (Prop. 6.9) follows from the Eliahou–Rozier identity, and the positive cycle drift (Lemma 4.2) is Revision 5's Lemma on terminal and cyclic drift. Where this happens we say so and record only the new proof or the new role.

2 Preliminaries

Let m_0 be a positive odd integer, $m_{n+1} = T(m_n)$, $d_n = a(m_n)$, $S_n = \sum_{i<n} d_i$, $\alpha = \log_2 3$, $R_n = S_n - n\alpha$, so $R_0 = 0$ and each step changes R by $d_n - \alpha$.

Remark 2.1. If every accelerated orbit reaches 1 then Collatz holds: every positive integer's odd part has an accelerated orbit, and the classical trajectory interleaves accelerated iterates with halvings. We work exclusively with T. [standard]

For fixed $c \geq 0$, $\mathcal{O}_c(m_0) = \#\{n \geq 0 : R_n \leq c\} \in \mathbb{Z}_{\geq 1} \cup \{\infty\}$; the four conjectures below fix $c > 0$, and statements using $c = 0$ say so. This is a global sojourn count: it neither halts at first passage nor counts time far above the critical line.

Lemma 2.2 (Eliahou–Rozier identity)**.** *For every odd $m_0 \geq 1$ and every $n \geq 0$,*

$$m_n\, 2^{R_n} = m_0 \prod_{i<n} \Big(1 + \tfrac{1}{3m_i}\Big), \qquad \textit{equivalently} \qquad R_n = \log_2 \frac{m_0}{m_n} + E_n, \quad E_n = \sum_{i<n} \log_2\Big(1 + \tfrac{1}{3m_i}\Big) \geq 0.$$

[due to Eliahou [1] and Rozier [2]; no novelty claimed]

Remark 2.3 (Notation warning: two quantities called E). Throughout, $E_n = \log_2 Q_n$ denotes the *carry excess* of Lemma 2.2. The *endpoint depth* of a word, elsewhere written $E(D)$, is a different quantity and is written $\mathsf{E}(D) = S - \log_2 r(D)$ here (§5.4). The two are unrelated and were a genuine source of confusion in earlier drafts.

There is a second, purely integral form of the same information, used whenever real logarithms would be inconvenient.

Lemma 2.4 (Aggregate identity)**.** *With $C_n = \sum_{j<n} 3^{n-1-j} 2^{S_j}$ (so $C_0 = 0$, $C_{n+1} = 3C_n + 2^{S_n} > 0$ for $n \geq 0$),*

$$2^{S_n} m_n = 3^n m_0 + C_n \qquad \textit{for all } n \geq 0.$$

[elementary; formalized, §14]

Proof. Induction on n. The case $n = 0$ is $m_0 = m_0$. If $2^{S_n} m_n = 3^n m_0 + C_n$ then, multiplying $2^{d_n} m_{n+1} = 3m_n + 1$ by 2^{S_n}, $2^{S_{n+1}} m_{n+1} = 3 \cdot 2^{S_n} m_n + 2^{S_n} = 3^{n+1} m_0 + (3C_n + 2^{S_n}) = 3^{n+1} m_0 + C_{n+1}$. □

Lemma 2.5 (Integral corridor)**.** *Since $S_n \in \mathbb{Z}$ and $n\alpha = \log_2(3^n)$,*

$$R_n \leq 0 \iff 2^{S_n} \leq 3^n, \qquad R_n > 0 \iff 3^n < 2^{S_n}, \qquad R_n \leq 0 \iff S_n \leq \lfloor n\alpha \rfloor.$$

[elementary]

Lemma 2.6 (Wedge)**.** *If $m_i \geq m_0$ for all $0 \leq i \leq n$, then $R_i \leq E_i \leq i/(3m_0 \ln 2)$ for $0 \leq i \leq n$. In particular $R_i \leq c$ for $0 \leq i \leq \min\{n, \lfloor 3c \ln 2 \cdot m_0 \rfloor\}$.* *[Revision 5]*

3 The occupation axis: EOC and its logical status

3.1 Four tiers

Conjecture 3.1 (Existence EOC)**.** For every fixed $c > 0$ there exist constants K_c, C_c such that every positive odd m_0 satisfies $\mathcal{O}_c(m_0) \le K_c \log_2 m_0 + C_c$.

Conjecture 3.2 (Effective EOC at (c, K, C))**.** For stated $c > 0$ and stated numerical K, C, every positive odd m_0 satisfies $\mathcal{O}_c(m_0) \le K \log_2 m_0 + C$.

Conjecture 3.3 (Sharp-leading EOC)**.** For every fixed $c > 0$, $\mathcal{O}_c(m_0) \le K \log_2 m_0 + O_c(\log\log m_0)$ with $K = 1/I(\alpha)$, where $I(\mu) = \mu - \mu \log_2 \mu + (\mu - 1)\log_2(\mu - 1)$.

Conjecture 3.4 (Strong sharp EOC)**.** With $K = 1/I(\alpha)$ as in Conjecture 3.3: for every fixed $c > 0$ there is C_c with $\mathcal{O}_c(m_0) \le K \log_2 m_0 + C_c$ for every positive odd m_0. (Revision 6 left K unbound here.)

Conjecture 3.4⇒3.3⇒3.1, and Conjecture 3.2 at any constants implies Conjecture 3.1 at that c. Crucially, Conjecture 3.1 does *not* imply Conjecture 3.2 at any stated constants: it asserts existence without a bound. All four tiers are pointwise statements about every seed, not almost-every-seed statements.

3.2 The Collatz reduction

Definition 3.5. For $c > 0$ and constants K, C, let $m^* = m^*(K, C, c)$ be least such that $3c \ln 2 \cdot m > K \log_2 m + C + 1$ for all $m \ge m^*$; it exists and is computable by finitely many exact comparisons.

Theorem 3.6. *Fix $c > 0$ and suppose Conjecture 3.1 holds at (c, K, C); put $m^* = m^*(K, C, c)$. Then*

(i) no accelerated orbit is injective; every orbit is eventually periodic;

(ii) every odd $m_0 \ge m^*$ has an iterate $m_i < m_0$ with $i \le \lfloor 3c \ln 2 \cdot m_0 \rfloor$;

(iii) if some seed's orbit does not reach 1, there is a nontrivial accelerated cycle with minimum member below m^*;

(iv) if additionally no nontrivial cycle has minimum member below m^*, the Collatz conjecture holds.

[Revision 5, retained unchanged]

Corollary 3.7 (Explicit instance)**.** *Assume Conjecture 3.2 at $c = 1$, $K = 13$, $C = 100$. Then $m^*(13, 100, 1) = 90$; every odd $m < 90$ reaches 1 by direct inspection (longest: $m = 73$, 42 accelerated steps; $m = 27$ and $m = 55$ take 41); hence Collatz holds.* *[Revision 5; witness corrected in Revision 7]*

Remark 3.8 (Why effective constants are genuinely needed). m^* is computable from (K, C, c), but Conjecture 3.1 supplies no bound on (K, C) themselves. Minimal-counterexample reasoning does not repair this: Theorem 3.6(ii) localizes a counterexample below m^* but does not eliminate it. Since Bařina's verification [22] settles all seeds below 2^{71}, any nontrivial cycle's minimum member exceeds 2^{71}, and the wedge forces $\mathcal{O}_1(m_{\min}) \gtrsim 3 \ln 2 \cdot 2^{71} \approx 4.9 \times 10^{21}$ — a value Conjecture 3.1 can accommodate at unspecified constants but which is wildly larger than anything the frontier-record figures of §15 suggest. [Revision 5]

3.3 The entropy scale

The natural population constant is the per-odd-step entropy deficit

$$I_{\text{Collatz}} \;=\; \alpha\big(1 - H_2(1/\alpha)\big) \;=\; \frac{1 - H_2(\rho)}{\rho}, \qquad \rho = \frac{1}{\alpha} = 0.6309297\ldots,$$

with H_2 the binary entropy. Numerically $I_{\text{Collatz}} = 0.0793186127748554\ldots$, and $1/I(\alpha) = 12.6074\ldots$ is the predicted asymptotic slope in Conjecture 3.3. The parenthesization matters: the constant is $\big(1 - H_2(\rho)\big)/\rho$, *not* $1 - H_2(\rho)/\rho$. [arithmetic; a known typesetting hazard in earlier drafts]

This constant calibrates the population scale and supplies, by itself, no pointwise anti-concentration. That distinction is one of the recurring themes of §9.

3.4 EOC as an occupation-time shadow of Rozier's hypothesis

Theorem 3.9. *Assume Rozier's Lower Bound Hypothesis [2, Hyp. 2.3]. Then for every fixed $c > 0$, $\mathcal{O}_c(m_0) \le \frac{1}{I(\alpha)} \log_2 m_0 + O_c(\log\log m_0)$; that is, LBH implies Conjecture 3.3, hence Conjecture 3.1. [Revision 5]*

The proof applies Rozier's own manoeuvre [2, Thm. 4.1] at occupation times rather than at total stopping time, via the entropy identity $\mu(1 - H_2(1/\mu)) = I(\mu)$. We regard it as a corollary of his framework, recorded because the occupation functional had not previously been isolated. It does not advance the provability of EOC — LBH is itself open, and arguably harder — but positions EOC precisely as a weaker occupation-time shadow of LBH. Making the additive term explicit would supply Conjecture 3.2 and, with it, Collatz *provided* the accompanying cycle search below m^* is completed — at the constants of Corollary 3.7 it is, since $m^* = 90$; at other constants m^* may exceed the verified range (Open Problem B). This is Open Problem A.

4 The divergence axis

This section is unconditional throughout; nothing here assumes EOC.

Definition 4.1. An orbit (m_n) is *injective* if $m_0, m_1, \ldots$ are pairwise distinct.

Lemma 4.2 (Terminal and cyclic drift)**.** *If an orbit reaches 1, or enters a nontrivial positive cycle, then $R_N \to +\infty$. Consequently every non-injective orbit has $R_N \to +\infty$. In particular, closure of a cycle of length L with total valuation S forces $2^S > 3^L$: strictly positive drift per period. [Revision 5; the integral form $2^S > 3^L$ is immediate from Lemma 2.4 and is formalized, §14]*

4.1 An elementary $8/9$ bound and its terminality

Lemma 4.3 (Carry budget)**.** *Every accelerated iterate m_k, $k \ge 1$, satisfies $m_k \not\equiv 0 \pmod 3$. Hence for $k \ge 1$ the states lie in the density-$1/3$ set of integers coprime to 6, and for every injective orbit $E_N \le \frac{1}{9}\log_2 N + O(1)$, so $2^{E_N} \ll N^{1/9}$. [Revision 5, Lem. 4.3 — where it carries a proof* sketch *only. Theorem 4.4 below rests on it, so a complete elementary proof with an explicit constant is given in Appendix A]*

Theorem 4.4 (Logarithmic-floor exclusion, $8/9$)**.** *Let m_0 have an injective accelerated orbit. For every $B < 8/9$, $R_k \ge -B\log_2 k$ fails for infinitely many k. [Revision 5, Thm. 4.5, verbatim; fully proved there, but via Lemma 4.3, which Revision 5 only sketches]*

Corollary 4.5. *For every injective orbit and fixed $c > 0$, $\#\{0 \le j < N : R_j > c\} \ll_{m_0,c} N^{1/9}$; hence $\mathcal{O}_c(m_0) = \infty$ at density-one rate, and Conjecture 3.1 directly excludes injective orbits. This is what Theorem 3.6(i) uses.* *[Revision 5]*

Remark 4.6 (Terminality of the mod-3 mechanism). For a prime $p > 3$, both 2 and 3 are units mod p, so $3m + 1 \equiv 0 \pmod p$ is solvable and no annihilation analogous to the mod-3 case occurs. This one-step fixed-modulus packing mechanism supplies no further density reduction at any other fixed modulus; pushing past 8/9 requires a genuinely different mechanism. [Revision 5]

4.2 Garcia–Tal/Curry sparsity

Theorem 4.7 (Garcia–Tal [13]; explicit form, Curry [14])**.** *Let $\mathcal{O}$ be the value set of an infinite aperiodic orbit of the raw map. For every $\beta > \beta_*$ there is C_β with $\#(\mathcal{O} \cap [a, a+X)) \le C_\beta X^\beta \log(2X)$ uniformly in a, where $\beta_* = \gamma_* \log_2 3 \approx 0.9653844$ and $\gamma_* \approx 0.6090897$ is the unique root of $H(\gamma) = \gamma \log_2 3$.* *[external, cited]*

As quoted, the constant C_β is uniform in the window start a and the statement is silent about uniformity over orbits. That uniformity is supplied not by [14]'s text but by the formalization [11]: there the base $\vartheta = 3^{31/50} < 2$ is produced by an existential quantified *before* any orbit is named, from the two exact integer certificates $50^{50} < 2^{50} \cdot 31^{31} \cdot 19^{19}$ and $3^{31} < 2^{50}$, so it is absolute. Uniformity over orbits is therefore a property of the formalized proof. [verification note]

Proposition 4.8 (Reciprocal summability)**.** *Every divergent (equivalently, injective) orbit satisfies $\sum_{x \in \mathcal{O}} 1/x < \infty$, hence $\sum_{n \ge 0} 1/m_n < \infty$ for its accelerated iterates. Consequently $Q_n \uparrow Q_\infty < \infty$, $E_n \uparrow E_\infty < \infty$, $m_n \to \infty$, and $R_n \to -\infty$.* *[Revision 5, from Theorem 4.7]*

Proposition 4.8 is Curry's Proposition 3.1 [14]; the label *[Revision 5]* records when it entered this paper, not priority. Curry's note is the single external input on which the divergence axis rests, and it answered Open Problem 12.8 of an earlier revision of this paper. [attribution]

Theorem 4.9 (Occupation bound and the $1/\beta_*$ threshold)**.** *Fix $\beta \in (\beta_*, 1)$. There are constants c_1, c_2 depending on β and on the orbit with $\#\{n \ge 0 : R_n \ge -G\} \le c_1 2^{\beta G}(G + c_2)$ for all $G \ge 1$. Consequently no divergent orbit satisfies $R_n \ge -B \log_2 n + O(1)$ eventually, for any $B < 1/\beta_* \approx 1.0358567$. The endpoint $B = 1/\beta_*$ is not excluded.* *[Revision 5]*

Remark 4.10 (An external refinement at the endpoint). Esm [10] formalizes a logarithmic refinement that operates at exactly the endpoint Theorem 4.9 leaves open. Its `uniform_window_bound` replaces $C_\beta X^\beta \log(2X)$ by $400\, X^\beta/(\log_2 X)^\gamma$, and its `divergent_drift_limsup` excludes $R_n \ge -\big(\beta_*^{-1} \log_2 n + c \log_2 \log_2 n + B\big)$ for every $c < \gamma/\beta$. So the main term $\beta_*^{-1} \log_2 n$ — the endpoint itself — is reached, but not closed: its own documentation states that "the cutoff is strict in the logarithmic exponent c; no exclusion at $c = \gamma/\beta$ is claimed". [external, unrefereed, Lean-claimed, not rebuilt here]

4.3 Curry's second, qualitative role

Revision 5 treated Theorem 4.7 purely as quantitative divergence sparsity. It has a second role, which organizes the whole time axis.

Proposition 4.11 (Last-maximum reduction)**.** *Let m_0 have a divergent orbit. By Proposition 4.8, $R_n \to -\infty$, so R attains a global maximum at finitely many indices; let n_0 be the* last. *Put $m^* = m_{n_0}$, an actual positive odd integer, and $R'_k = R_{n_0+k} - R_{n_0}$ its relative drift. Then $R'_k < 0$ for every $k \ge 1$,*

hence $S'_k \le \lfloor k\alpha \rfloor$ for every $k \ge 1$: the restarted orbit has an infinite zero-corridor tail. [companion work. Formalized as `exists_zero_confined_seed_tendsto`, but under the hypothesis `ReciprocalSummable` — i.e. conditional on Theorem 4.7, which is external and unformalized. It is not *formalized from "divergent orbit" directly]*

This reduction is not new to Revision 6 — it is the zero-corridor tail theorem of the Curry foundation notes, where it is proved and formalized, and where it is also recorded that the step requires no irrationality of α. Its role here is to make the divergence axis and the time axis meet.

Remark 4.12 (Scope). Nothing in this section constrains hypothetical nontrivial cycles, proves any tier of EOC, or touches the occupation of convergent orbits. [Revision 5]

5 The residue axis

Definition 5.1. A valuation word $D = (d_0, \dots, d_{N-1})$ has $d_i \ge 1$; $s_0 = 0$, $s_j = \sum_{i<j} d_i$, $S_N(D) = s_N$, and carry sum $C_N(D) = \sum_{j=0}^{N-1} 3^{N-1-j} 2^{s_j}$. An odd m_0 *realizes* D if $a(m_i) = d_i$ for all $i < N$.

Proposition 5.2 (Exact realizer congruence)**.** *The exact realizers of D form a single residue class modulo $2^{S_N(D)+1}$, not $2^{S_N(D)}$. Writing $r(D)$ for the least positive representative,*

$$r(D) \equiv \left(2^{S_N(D)} - C_N(D) \cdot 3^{-N}\right) \pmod{2^{S_N(D)+1}}.$$

[Revision 5; the modulus is $2^{S_N(D)+1}$ and not the Terras–Everett modulus $2^{S_N(D)}$, as established in [25]]

Definition 5.3 (Coarse anchor versus exact lift)**.** For $S \le S_N(D)$ the *coarse anchor* at depth S is $a_S(D) := \left(-C_N(D)3^{-N}\right) \bmod 2^S$. Proposition 5.2 shows $r(D)$ reduces mod 2^{S_N} to $a_{S_N}(D)$ and is one of $\{a_{S_N}(D), a_{S_N}(D) + 2^{S_N}\}$, terminal parity selecting which. Unconditionally $r(D) \ge a_{S_N}(D)$; this inequality, not the equality, is what every floor argument below uses. [Revision 5]

5.1 Proved sectors

Theorem 5.4 (Periodic realizer floor)**.** *For a fixed block X put $\xi_X = -C(X)/\delta_X \in \mathbb{Z}_2$ with $\delta_X = 3^{q_X} - 2^{S_X}$ odd. If $\xi_X \notin \mathbb{Z}_{>0}$, there is a finite B_X with $r(X^n) \ge 2^{nS_X - B_X}$ for all large n, hence $\log_2 r(X^n)/(nS_X) \to 1$. [Revision 5; full proof in companion note [24]]*

Theorem 5.5 (Eventually periodic words)**.** *For a fixed prefix Y the same mechanism applies to YX^n with transformed anchor $\xi_{X,Y}$: if $\xi_{X,Y} \notin \mathbb{Z}_{>0} \cup \{0\}$ then $r(YX^n) \ge 2^{S_Y + nS_X - B_{X,Y}}$* for every n with $S_Y + nS_X \ge P_{X,Y} + B_{X,Y}$, *and $\log_2 r(YX^n)/(S_Y + nS_X) \to 1$. [Revision 5; proved unconditionally in [24], Thm. 9.1. The exceptional case is* not *literally the same as in Thm. 5.4: there $\xi_X \ne 0$ holds for every block, whereas $\xi_{X,Y} \ne 0$ is proved in [24] only when $b_X > 1$ (Cor. 10.2), the boundary case $b_X = 1$ being checked directly when it arises]*

5.2 The finite-defect sector: an unproved claim inherited from Revision 5

Nothing in this paper depends on the statement below. It is retained only because Revision 5 recorded it and a reader of that revision will look for it; it is **not** a result of this programme, and it is placed outside §5's proved sectors deliberately.

Proposition 5.6 (Bulk-defect anchor stability — **stated, not proved**)**.** *For $D = X^a E X^b$ (fixed X, arbitrary finite defect E, arbitrary b), the coarse anchor satisfies $a_{aS_X}(D) = \xi_X \bmod 2^{aS_X}$ exactly, independent of E and b. Hence, outside the exceptional case, $r(D) \ge 2^{aS_X - B_X}$* for a sufficiently large. *[Revision 5, Prop. 5.10.* ***No proof is on file*** *— see Rem. 5.7]*

Remark 5.7 (The finite-defect sector is open, and this proposition is not a proved sector). Revision 6 and the first draft of Revision 7 listed Proposition 5.6 under "proved sectors". An audit of the sources withdraws that classification.

(i) The companion note [24], which this subsection cites for the periodic sector, contains *no* statement about defect words. Its epistemic ledger lists "Finite-defect words — **Open, major**", and its limitations section says of exactly this class: "This note does not attempt to resolve this question." It further warns that it "does not establish $r(D) \geq 2^{S(D)-O(1)}$ for arbitrary confined words, and no such claim should be inferred from the results proved here."

(ii) Revision 5 states Proposition 5.10 with a one-paragraph mechanism and *no proof*, and carries the qualifier "for a sufficiently large", which Revision 7 had dropped and which is genuinely needed (the analogue in Thm. 5.4 holds only once $nS_X \geq P_X + B_X$).

(iii) The statement is not contradicted by (i): its content concerns the coarse anchor at depth aS_X, fixed by the X^a prefix alone, and its floor is exponential in a only, not in $|D| = aS_X + |E| + bS_X$. It is weaker than what the note calls open. But weaker-and-unproved is not proved.

We therefore retain the statement, as a recorded Revision 5 claim, and remove it from the proved sectors. Theorems 5.4 and 5.5 are unaffected: both are proved unconditionally in [24] (Thms. 6.2 and 9.1). [Revision 7, correction of status]

Remark 5.8 (Arbitrary irregular words: open). For a general confined word the coarse anchor $\xi_N(D) = -C_N(D)3^{-N}$ changes with N: there is no single fixed rational whose periodicity can be invoked. This is the moving-anchor sector, and it is where actual Collatz orbits live. [Revision 5]

5.3 What the height invariant actually separates

Revision 5 framed the split as periodic versus irregular. Later audits sharpened it. Every realizer condition can be written

$$2^{S_N} \mid A + \chi\Lambda$$

for integers A, Λ *generated by the word*: for a general word $(A, \Lambda) = (C_N, 3^N)$, while for a purely periodic word of period p the carry collapses to $C_{Np} \cdot \Delta = c_0(3^{Np} - 2^{N\sigma})$ with c_0, Δ independent of N. The only general tool converting divisibility into a size statement is that a nonzero multiple of 2^m has absolute value at least 2^m, which gives a *dichotomy*: either $A + \chi\Lambda = 0$, pinning χ to the single rational $-A/\Lambda$, or

$$|\chi| \;\geq\; 2^{S_N}/H - 1, \qquad H = \max(|A|, |\Lambda|).$$

Writing $\rho_N = \log_2 H / S_N$, the second branch is useful exactly when ρ_N is bounded away from 1.

Observation 5.9 (Height amortization, not complexity). Periodic and eventually periodic words have $\rho_N \to 0$, at rate $1/(\text{number of repetitions})$: a fixed height is amortized over unboundedly many periods. (The same amortization is what Proposition 5.6 would give for one-defect words, but that proposition is unproved — Rem. 5.7.) Arbitrary confined words supply only $(C_N, 3^N)$, with $\log_2 H = \alpha N \geq S_N$, so $\rho_N \geq 1$ and the second branch is vacuous. The separating invariant is therefore height amortization, not symbolic complexity: the Beatty/Sturmian zero-corridor word, of minimal aperiodic subword complexity $p(n) = n + 1$, already has $\rho_N > 1$ at every depth examined, while its periodic continued-fraction approximants keep $\rho_N \approx 0.26$. [structural interpretation, supported by exact identities and by computation; not a theorem about all words]

5.4 Endpoint depth and the frontier defect

Write $\mathsf{E}(D) = S - \log_2 r(D)$ for the endpoint depth (Remark 2.3) and $g(D) = S - \lfloor \log_2 r(D) \rfloor$ for the discrete terminal gap. Then

$$\mathsf{E}(D) \;\le\; g(D) \;<\; \mathsf{E}(D) + 1, \qquad \text{i.e. } g = \lceil \mathsf{E} \rceil,$$

the two agreeing only when $r(D)$ is a power of 2. [proved; the inequalities are formalized (`discreteGap_bracket`); the equality clause is not]

Remark 5.10 (Correction to Revision 5's Open Problem E). Revision 5 asked whether the frontier-defect target $\mathsf{E}(D) \le H_2(\rho_c)S + \Phi(N)$, $\Phi = o(N)$ — where $\rho_c = N/S$ is the one-density of the valuation word, a symbol left undefined in Revisions 5–7 — is genuinely easier than an exponential floor, and recorded that question as open. It is not easier in any useful sense: by the bracket above the statement is equivalent, up to an additive 1, to $g(D) \le H_2(\rho_c)S + \Phi(N)$, and under confinement $S \le \lfloor \alpha N \rfloor$ this is an exponential floor with $\varepsilon = I_{\text{Collatz}} + o(1)$. Open Problem E of Revision 5 is therefore the sharp-leading form of Open Problem C, not an independent weaker route, and is restated as such in §13. [correction to Revision 5]

6 The time axis

6.1 Three temporal scales

Definition 6.1. Fix $c \ge 0$ and an odd seed m.

$$\begin{aligned} L_c(m) &= \max\{N \ge 0 : R_j(m) \le c \text{ for all } 0 \le j \le N\} && \text{(single-window lifetime)},\\ \mathcal{O}_c(m) &= \#\{n \ge 0 : R_n(m) \le c\} && \text{(total occupation)},\\ \tau_c(m) &= \inf\{n \ge 1 : R_n(m) > c\} && \text{(drift exit time)}. \end{aligned}$$

All three are read along the full forward orbit, including after it reaches 1: since $a(1) = 2$ and $T(1) = 1$, the orbit continues at the fixed point, where R rises by $2 - \alpha > 0$ per step. This is the only convention consistent with Definition 6.1 as written — the orbit is defined for every $n \ge 0$ and nothing in the three definitions halts it — and it makes all three quantities finite for every seed that reaches 1, so no separate stopping rule is needed. [convention; fixed here for the first time]

These answer different questions. L_c measures the initial confined episode; $\mathcal{O}_c$ sums all episodes, including separated later returns; τ_c asks only whether the corridor is ever left. Always $L_c(m) \le \mathcal{O}_c(m)$ and $\tau_c(m) = L_c(m) + 1$ when finite.

Remark 6.2 (Single window does not control occupation). The inequality $\mathcal{O}_c \ge L_c$ is strict in general, and dramatically so: $m_0 = 285175$ has $L_1 = 14$ but $\mathcal{O}_1 = 98$. So the two quantities are genuinely different functionals, and a single-window theory does not *directly* deliver EOC.

Revisions 5–7 printed $\mathcal{O}_1 = 97$ here. That was an off-by-one: Definition 6.1 counts from $n = 0$, and $R_0 = 0 \le 1$, so the seed time lies in the corridor and is counted. The corrected value is $\mathcal{O}_1(285175) = 98$, recomputed in exact integer arithmetic; $L_1 = 14$ and $\tau_1 = 15$ are unchanged, and so is the point of the example.

The decomposition into episodes sharpens that point. At $c = 1$ the occupation set of $m_0 = 285175$ splits into seven maximal runs of consecutive corridor times, of lengths $15+4+6+64+1+5+3 = 98$: the first is the initial window itself ($L_1 + 1 = 15$ times, counting $n = 0$), so 83 of the 98 occupation times are separated returns, and the longest single return is a run of 64 — more than four times the initial window. An optimal bound on L_c would therefore still leave the bulk of $\mathcal{O}_c$ uncontrolled,

which is exactly the content of Open Problem F. Both the corrected count and this decomposition were computed independently [33].

Revision 6 drew the stronger conclusion that "a complete single-window theory — even an optimal one — does not yield EOC". That does not follow from a pointwise example, and is withdrawn; see Remark 10.1. For the corridor-uniform form of the hypothesis it is in fact false, *given cycle exclusion* (Rem. 7.8(i)). [Revision 5 example; conclusion corrected in Revision 7; occupation value corrected in Revision 8]

6.2 Residue–time duality

Definition 6.3. Fix $c > 0$. A word D of length N is *c-confined* if $R_j(D) \le c$ for $0 \le j \le N$. For the c-confined family at depth N, $r_{\min}(N, c) := \min_D r(D)$.

Proposition 6.4 (Record-chronology inversion)**.** $r_{\min}(N, c) = \min\{m \text{ odd} : L_c(m) \ge N\}$. *[Revision 5, Proposition 5.6; record-chronology reformulation of the single-window realizer floor in [25]]*

Proposition 6.5 (Single-window equivalence)**.** *Fix $c > 0$ and $\varepsilon \in (0, 2]$. The following are equivalent:*

(i) $r(D) \ge 2^{\varepsilon N}$ for every c-confined D of length N with $r(D) \ge 3$;

(ii) $L_c(m_0) \le \varepsilon^{-1} \log_2 m_0$ for every odd $m_0 \ge 3$.

[Revision 5, Prop. 5.7, verbatim. Revision 5 labels its argument a proof sketch*; it is complete in substance in both directions, and both directions are re-derived in the Revision 7 audit]*

Revision 5 presented Propositions 6.4 and 6.5 as technical observations inside the residue section. They deserve a more prominent role: together they say that *the residue and time axes are the same axis in two coordinates*. Every statement about how small a realizer can be is a statement about how long an orbit can stay confined, and conversely. Revision 6 takes this as an organizing principle rather than a lemma.

6.3 The general transfer principle

Proposition 6.5 matches one specific pair of scales. The inversion is general.

Proposition 6.6 (Lifetime–frontier transfer)**.** *Let $F : \mathbb{Z}_{>0} \to \mathbb{Z}_{>0}$ be nondecreasing and suppose $L_c(m) \le F(m)$ for every odd m. Then*

$$r_{\min}(N, c) \;\ge\; \min\{m : F(m) \ge N\} \;=\; F^{-1}(N).$$

[proved here; the finite form is formalized, §14]

Proof. By Proposition 6.4, $r_{\min}(N, c)$ is the least odd m with $L_c(m) \ge N$. If $m < F^{-1}(N)$ then $F(m) < N$ by monotonicity, so $L_c(m) \le F(m) < N$; hence no such m qualifies. □

Writing the hypothesis in terms of $x = \log_2 m$ gives the scales of interest.

Lifetime hypothesis	Frontier obtained	Character
$L_c(m) \le K \log_2 m$	$\log_2 r_{\min}(N, c) \ge N/K$	exponential (Open Problem C)
$L_c(m) \le K(\log_2 m)^p$, $p > 1$	$\log_2 r_{\min}(N, c) \ge (N/K)^{1/p}$	stretched exponential
$L_c(m) \le K m^\theta$, $\theta > 0$	$r_{\min}(N, c) \ge (N/K)^{1/\theta}$	polynomial
$L_c(m) \le K m$	$r_{\min}(N, c) \ge N/K$	linear

Remark 6.7 (A graded programme). This table is the most constructive consequence of the recent audits. EOC's residue axis has usually been posed all-or-nothing — prove an exponential floor or prove nothing. Transfer shows otherwise: *any* unconditional pointwise lifetime bound by a nondecreasing F, however weak, yields a corresponding unconditional frontier bound. A bound as weak as $L_c(m) \le Km$ would already give a linear frontier, and none is currently known. We record this as Open Problem D. [proved; the assessment of what is currently known is in §6.8]

Remark 6.8 (No unconditional lifetime bound can be easy). No unconditional bound $L_0(m) \le F(m)$ is currently known, and none can be proved without excluding divergent orbits: applied to the restarted seed m^* of Proposition 4.11, which has $L_0(m^*) = \infty$, any such bound is immediately contradicted. The transfer table therefore grades a programme whose every entry is open. [proved]

6.4 Drift exit and the future-minimum theorem

Proposition 6.9 (Future minimum). *For $n \ge 1$: if $R_n \le 0$ then $m_n > m_0$. Hence an orbit with $R_n \le 0$ for all $n \ge 1$ has its seed as a* strict minimum of its entire future. *Contrapositively, for $n \ge 1$, $m_n \le m_0$ already forces $R_n > 0$.*

The restriction $n \ge 1$ is necessary, not cosmetic: at $n = 0$ we have $R_0 = 0 \le 0$ while $m_0 > m_0$ fails. Both proofs below use it — the first through $E_n > 0$, the second through $C_n > 0$, and $E_0 = C_0 = 0$. [Revision 5; hypothesis $n \ge 1$ made explicit in Revision 7]

Proof. Two proofs. Analytically, Lemma 2.2 gives $R_n = \log_2(m_0/m_n) + E_n$ with $E_n > 0$ for $n \ge 1$, so $m_n \le m_0$ implies $R_n \ge E_n > 0$. Integrally, Lemma 2.4 and $C_n > 0$ give $2^{S_n} m_n = 3^n m_0 + C_n > 3^n m_0 \ge 2^{S_n} m_0$ when $2^{S_n} \le 3^n$, so $m_n > m_0$. □

Remark 6.10 (Provenance). This is *not* a new theorem. The analytic proof is an immediate consequence of the Eliahou–Rozier identity already in Revision 5 (Lemma 2.2), and is close in content to Revision 5's wedge (Lemma 2.6). What is new is only the second proof, which uses no logarithms and is therefore convenient to formalize, and the role the statement plays below. [reclassification]

Proposition 6.11 (Drift exit is weaker than ordinary descent). *Let $\sigma(m) = \inf\{n \ge 1 : m_n < m\}$ be the (Terras) stopping time. Then $m_n < m_0 \iff R_n > E_n$, whence $\tau_0(m) \le \sigma(m)$, and $\sigma(m) < \infty \Rightarrow \tau_0(m) < \infty$. The inequality is strict at cycle minima: for $m = 1$, $R_n = E_n = (2-\alpha)n$ exactly, so $\tau_0(1) = 1$ while $\sigma(1) = \infty$. [proved]*

Remark 6.12 (Terras's coefficient stopping time). Proposition 6.11 is not new: τ_0 is Terras's coefficient stopping time κ [3, 4] and the inequality is his $\kappa \le \sigma$. Equality $\tau_0 = \sigma$ is the Coefficient Stopping Time Conjecture (so named by Lagarias), verified by Rozier and Terracol [5] for all $2 \le n \le 2.8 \times 10^{19}$. Their Theorem 3.2 — an integer of infinite stopping time generates infinitely many paradoxical sequences, i.e. sequences with ones-ratio below $\log 2/\log 3$ which nonetheless never descend below their first term — is the refereed dual of Proposition 4.11: both convert a hypothetical non-descending seed into an infinite supply of extremal finite configurations. [literature audit; the attribution is corrected, the mathematics is unchanged]

Remark 6.13 (How much weaker, quantitatively). Off cycle minima the gap is small. At $n = \tau_0$, strictness needs $R_n \le E_n$; under confinement for $j < n$, Proposition 6.9 gives $m_j > m_0$, so $E_n \le (3\ln 2)^{-1} \sum_{j<n} 1/m_j \le 0.4809\, n/m_0$, while integrality of S_n with $R_n > 0$ forces $R_n \ge 1 - \{\alpha n\}$. Strictness at depth n therefore requires

$$r_{\min}(n, 0) \;\le\; m_0 \;\le\; 0.4809\, n/(1 - \{\alpha n\}),$$

whose right side is polynomial in n because α has finite irrationality measure. So *any* superpolynomial realizer floor forces $\tau_0 = \sigma$ away from cycle minima. An exhaustive search of every odd m below this size bound at every depth $n < 200$ where the bound is not already violated found no strict example. [the inequality is proved; the search is a computational diagnostic]

6.5 Universal drift exit and divergence

Definition 6.14 (Universal drift exit)**.**

$$\text{(DE)} \qquad \text{for every odd } m \geq 1 \text{ there is } n \geq 1 \text{ with } 3^n < 2^{S_n(m)}.$$

Equivalently, by Lemma 2.5, $\tau_0(m) < \infty$ for every m. The phrase "drift exit time" is terminology used here; the quantity itself is Terras's coefficient stopping time κ (Remark 6.12, §16).

Theorem 6.15 (Divergence equivalence)**.** (DE) *holds if and only if there is no divergent positive accelerated orbit.*

Proof. ($\Rightarrow$, contrapositive.) If a divergent orbit exists, Proposition 4.11 produces an actual positive odd integer m^* with $R'_k < 0$ for every $k \geq 1$, so $\tau_0(m^*) = \infty$ and (DE) fails.

($\Leftarrow$, contrapositive.) Suppose (DE) fails at m, so $2^{S_n} \leq 3^n$ for all $n \geq 1$. The orbit cannot repeat a state: a repeat makes it eventually periodic, and Lemma 4.2 then gives $R_N \to +\infty$, contradicting $R_n \leq 0$. So the orbit is injective, hence divergent. □

Corollary 6.16 (Cycle compatibility)**.** *By Lemma 4.2 every point of a nontrivial positive cycle, and every transient point entering one, has finite τ_0; so such a cycle is consistent with the defining condition of* (DE) *holding at every seed. Hence the drift-exit argument supplies no cycle exclusion, Collatz additionally requires excluding nontrivial cycles, and whether* (DE) *already does so is precisely the open cycle question — see Remark 6.17, which claims nothing either way.*

Remark 6.17 ("Strictly weaker" is itself open). Revisions 5–7 up to this point wrote that (DE) is *strictly weaker* than Collatz. That is **not established**, and the argument just given does not establish it. Since (DE) is equivalent to the absence of divergent orbits (Thm. 6.15) and Collatz is the conjunction of that with the absence of nontrivial cycles, Collatz $\Rightarrow$ (DE) always, while (DE) $\Rightarrow$ Collatz holds *if and only if* no nontrivial cycle exists. So (DE) is strictly weaker precisely when a nontrivial cycle exists — an open question whose expected answer is *no*, in which case the two are equivalent.

What is established, precisely, is this list and no more:

(i) Collatz $\Rightarrow$ (DE);

(ii) (DE) excludes divergence (Thm. 6.15);

(iii) (DE) together with the exclusion of nontrivial positive cycles $\Rightarrow$ Collatz;

(iv) every point of a cycle has finite τ_0 (Lem. 4.2), so the drift-exit argument itself supplies no cycle exclusion.

Accordingly: **no implication from** (DE) **to cycle-freeness is established here.** We claim neither that (DE) is strictly weaker than Collatz nor that (DE) fails to entail cycle-freeness — the latter is *also* unproved, and would be false if no nontrivial cycle exists, which is the expected case. Item (iv) is a statement about what this argument delivers, not about what is true. This is the same error of form corrected in Remark 10.1: a non-implication of *mechanism* is neither strictness nor non-entailment of *content*. [Revision 7, correction of scope]

Remark 6.18 (Provenance, emphatically). Neither direction is new mathematics. The forward direction is the zero-corridor tail theorem of the companion Curry foundation, already proved and formalized there; the reverse direction is Revision 5's Lemma 4.2. Moreover the *statement* (DE) is, in negated form, exactly the frontier already posed in that companion work: "determine whether an infinite zero-confined valuation word can be realized by a positive ordinary integer." Revision 6 contributes the reclassification — recognising this as the qualitative endpoint of the *time* axis rather than an isolated realizability question — together with the integral form and its exact position in the hierarchy of §10. We do not coin a new acronym for it; in particular it is distinct from the repository's ZCRE statement, which is an equivalence about bounded prefix realizers. [reclassification, not a new theorem]

7 Launch renormalization: the bridge to actual orbits

Finite word and residue analysis becomes powerful exactly when it produces an *actual* positive-orbit segment. Launch renormalization does this. We summarize only what the programme map needs; the full treatment is the companion note [27].

Theorem 7.1 (Launch-state collapse, [27])**.** *Let D be a word and j any index with $2^{S_j} > r(D)$. Then $m_j^* := (3^j r(D) + C_j(D))/2^{S_j}$ is a positive integer, in fact the genuine j-th accelerated iterate of $r(D)$, and the letters $(d_{j+1}, \dots, d_N)$ are exactly the accelerated valuation word of m_j^*.* [companion work]

Call D of length N *ε-dangerous* if it is c-confined, has $r(D) \geq 3$, and satisfies $\log_2 r(D) < \varepsilon N$ (with $c \geq 1$ and $\varepsilon \in (0,1)$). For such a word the pinning index t_1 and the launch data satisfy

$$t_1 \leq \varepsilon N + O(1), \qquad L \;\geq\; (1-\varepsilon)N - O(1), \qquad \log_2 m_{t_1}^* \leq \alpha\varepsilon N + O(\log N),$$

with inherited *relative* corridor $A \leq c + (\alpha - 1)\varepsilon N + O(1)$. The leading coefficient in the launch-size bound is $\alpha\varepsilon$, not ε.

Remark 7.2 (Hypotheses the launch-size bound carries). The third bound — the launch-size bound — holds in [27] (Prop. 6.3(d)) only *if the orbit segment $r(D) = m_0, \dots, m_{t_1}$ is injective.* That hypothesis is not free. The note closes the trivial-cycle branch unconditionally and sharply for $N > N_c$, but records, in a retraction worth quoting, that the other branch is open: an earlier draft's claim that cycle exclusions render a nontrivial-cycle tail incompatible with danger "was incorrect, with the asymptotics running in the wrong direction, and is retracted … the cycle branch is closed unconditionally … for the trivial cycle, and open for hypothetical nontrivial cycles at large N." The note then carries the hypothesis explicitly rather than, in its own words, assuming nonexistence silently.

Revision 6 and the first draft of Revision 7 stated the launch-size bound flatly. It is restated here with its hypothesis. The renormalization-stability sentence below inherits the same condition, together with [27]'s $N > N_c$, its period cap q_0, and the threshold $\varepsilon < (2-\alpha)/(2+\alpha) \approx 0.1158$ below which its terminal-repetition bound is useful. [Revision 7, restored hypotheses]

The residual irregular class is renormalization-stable: the forced suffix of a dangerous word in that class is again a small integer with a long, near-critically confined injective orbit segment, so finite structural enumeration cannot close the programme [27].

7.1 Reclassifying the corridor-uniform hypothesis

The companion note closes the route conditionally on a corridor-uniform lifetime hypothesis: for some $K < \infty$, every injective initial segment of m with $R_i \leq A$ for $0 \leq i \leq L$ has $L \leq K(\log_2 m + A)$. This reads as a mild statement about corridors. It is not.

Proposition 7.3 (The corridor is automatic)**.** *For every positive orbit and every n, $2^{S_n} \le 3^n m_0 + C_n$, i.e. $R_n \le \log_2 m_0 + E_n$; and equality holds as soon as the orbit reaches* 1*. Hence the least valid corridor width for a full orbit is exactly $\log_2 m_0 + E$, not merely bounded by it.* *[proved; formalized, §14]*

Proof. Immediate from Lemma 2.4 and $m_n \ge 1$, with equality when $m_n = 1$. □

Corollary 7.4 (Reclassification)**.** *Instantiating the corridor-uniform hypothesis at its own automatic corridor gives $L \le K(2\log_2 m + E)$. For an injective segment the iterates are distinct odd integers, so $E_L \le (3\ln 2)^{-1} \sum_{j<L} 1/m_j \le 0.2404 \ln L + O(1)$, which grows slower than L and therefore bounds L. Consequently the hypothesis implies that no accelerated orbit is injective forever — i.e. it implies the absence of divergent orbits — and, more precisely, that* every *orbit repeats a value, and so enters a cycle, within $O(\log m)$ steps. Conversely a bound $L \le K' \log_2 m$ returns the hypothesis for every $A \ge 0$.* *[proved]*

Remark 7.5 (What the injective-segment bound does *not* give). The quantity bounded is the length of an *injective* initial segment, which is the pre-period plus the period: the orbit repeats a value within $O(\log m)$ steps and is thereafter periodic. It does *not* follow that the orbit reaches 1: a hypothetical nontrivial positive cycle is entirely consistent with the bound, and the corridor-uniform hypothesis does not visibly exclude one. (Applied to a cycle point it bounds that cycle's *period* by $O(\log m^*)$, which is a strong constraint but not a contradiction.)

Revision 6 wrote "and, with that, an $O(\log m)$ bound on accelerated total stopping time", where the total stopping time is the time to reach 1. That step is **withdrawn**: it silently assumed cycle exclusion. Everything downstream that used arrival at 1 is restated in Remark 7.8 with the assumption made explicit. [Revision 7, correction]

Proposition 7.6 (Corridor times are large-value times)**.** *For every positive odd seed and every n, if $R_n \le c$ then $m_n \ge m_0 2^{-c}$.*

Proof. *By Lemma 2.4, $2^{S_n} m_n = 3^n m_0 + C_n \ge 3^n m_0$, while $R_n \le c$ reads $2^{S_n} \le 2^c 3^n$. Hence $3^n m_0 \le 2^c 3^n m_n$.* □

No cycle hypothesis is used, and none is available: this is the only statement of its kind we can make unconditionally. *[Revision 7; formalized, §14]*

Proposition 7.7 (Occupation after the orbit becomes periodic)**.** *Suppose the orbit of m_0 is eventually periodic, entering a cycle at time T with value set $\mathcal{C}$.*

(i) If $\max \mathcal{C} < m_0 2^{-c}$ then $\mathcal{O}_c(m_0) \le T$.

(ii) In particular, if the orbit reaches 1 *at time $n^* \ge 1$ and $m_0 > 2^c$, then $\mathcal{O}_c(m_0) \le n^*$.*

Proof. *By Proposition 7.6 every corridor time n satisfies $m_n \ge m_0 2^{-c}$; under (i) no $n \ge T$ does, so all corridor times precede T. For (ii), $\mathcal{C} = \{1\}$ and $1 < m_0 2^{-c}$ exactly when $m_0 > 2^c$.* □

Remark (The value criterion is vacuous in the remaining case). Revision 7 first stated this as $\mathcal{O}_c(m_0) \le T + \#\{n \ge T : m_n \ge m_0 2^{-c}\}$. That inequality is true but *useless* outside case (i): if any cycle value meets the threshold, the cycle repeats it forever, so the counting term is **infinite**, while the actual occupation is finite. Occupation is not governed by the value criterion alone — corridor membership is a condition on R_n, and by Lemma 4.2 R gains a fixed positive amount per period, so the corridor is left permanently after finitely many periods. That finite bound depends on the cycle — its period, its per-period drift gain, and the drift at entry — and we have no bound on it uniform over hypothetical nontrivial cycles. The display has been removed rather than retained with a caveat. [Revision 7, correction]

Remark. The hypotheses $n^* \geq 1$ and $m_0 > 2^c$ (rather than $m_0 \geq 2^c$) are both needed, and the boundary case shows why: for $m_0 = 1$, $c = 0$ we have $n^* = 0$, yet $R_0 = 0 \leq c$, so time 0 lies in the corridor and $\mathcal{O}_0(1) = 1 > n^*$. Proposition 7.6 diagnoses this exactly — at $m_0 = 1, c = 0$ the criterion $m_n \geq m_0 2^{-c}$ is satisfied by the value 1 itself. Revision 6's formulation omitted these hypotheses.

What the Lean statement carries is $1 \leq n^*$ together with the *non-strict* $2^c \leq m_0$; that pair suffices, and it is not the same pair as stated above. Revision 7 first claimed the Lean hypothesis was $m_0 > 2^c$; it is not.

[Revision 7. The reaches-1 corollary is formalized (`occupation_le_of_reaches_one`); the general eventually-periodic display is not *formalized]*

Remark 7.8 (What Corollary 7.4 does and does not settle). Three points, kept apart.

(i) The hypothesis implies EOC — given cycle exclusion. Corollary 7.4 bounds the injective initial segment, i.e. the time T at which the orbit first repeats a value, by $O(\log m_0)$. If the only positive cycle is $\{1\}$ — that is, *given cycle exclusion* — then every orbit reaches 1 by time T, and Proposition 7.7(ii) gives $\mathcal{O}_c(m_0) \leq T = O(\log m_0)$ for every $m_0 > 2^c$: the corridor-uniform hypothesis implies Existence EOC at every c. Under the same additional assumption it also implies the Collatz conjecture, so it relocates the difficulty rather than reducing it.

Without cycle exclusion the conclusion is weaker and should be stated as such: occupation is bounded by $O(\log m_0)$ plus the corridor time spent inside whatever cycle the orbit enters. That residual is finite for each fixed cycle — a cycle has strictly positive period drift (Lem. 4.2), so its drift exceeds c after finitely many periods — but we have no bound on it that is uniform over hypothetical nontrivial cycles, whose period and values are unconstrained above. Supplying such a uniform bound, or excluding the cycles, is a genuine additional obligation. Revision 6 stated (i) without it. [Revision 7, correction]

(ii) One layer is not formalized. Proposition 7.7 and the aggregate identity it rests on are machine-checked. The passage from the corridor-uniform hypothesis to the stopping-time bound — Corollary 7.4 — is not: it absorbs the term $E_L = O(\log L)$ into a bound on L, an elementary but genuinely real-analytic step. The chain is therefore a formalized bridge sitting on an unformalized analytic layer, and should be described that way.

(iii) Do not conflate corridor-uniformity with a fixed corridor. The hypothesis quantifies over *every* corridor width A simultaneously; it is that uniformity which permits instantiation at the automatic corridor $A = \log_2 m_0 + E$ and hence the collapse in (i). A bound on L_c at one fixed c carries none of this: it is implied by the corridor-uniform form, and the converse is not known. (We avoid "strictly weaker" here for the reason given in Rem. 6.17.) Remarks elsewhere in this paper about "single-window" bounds refer to the fixed-corridor form. [Revision 7]

So the corridor-uniform hypothesis is *noncircular* with respect to Open Problem C — it is not a repackaging of the realizer floor — but it is a substantially stronger *quantitative* hypothesis than is required merely to exclude divergence. In particular, once the automatic corridor is inserted it is essentially an $O(\log m)$-scale pointwise stopping/lifetime hypothesis, and it should not be presented as a plausibly easier closing hypothesis. The companion note's weakened variant, which exempts segments majority-covered by a bounded-period terminal repetition, does not change this: over all odd $m < 60001$ the exemption applies to none of the 29 999 orbits tested. [computational diagnostic only]

8 Self-financing recovery

A natural strategy for the time axis is to argue that a deep negative drift excursion cannot be repaid, because the recovery would demand an impossibly large valuation. This fails at its source.

Proposition 8.1 (Self-financing)**.** *A single step of valuation d from state m_n achieves drift exit as soon as $2^d m_0 > 3\,m_n$. The available valuation is capped by $2^{d_n} \le 3m_n + 1$, since $2^{d_n} \mid 3m_n + 1$.* *[proved; formalized, §14]*

Proof. For the cap, $2^{d_n} m_{n+1} = 3m_n + 1$ with $m_{n+1} \ge 1$. For the sufficient condition, multiply the target $3^{n+1} < 2^{S_{n+1}}$ by m_n: by Lemma 2.4, $2^{S_{n+1}} m_n = 2^d\, 2^{S_n} m_n \ge 2^d 3^n m_0 > 3 \cdot 3^n m_n = 3^{n+1} m_n$ whenever $2^d m_0 > 3m_n$. □

The cap therefore exceeds the sufficient threshold $3m_n/m_0$ by a factor of about m_0. The underlying reason is the identity $\log_2 m_n = \log_2 m_0 - R_n + E_n$ of Lemma 2.2: sustained negative drift makes the ordinary iterate *larger*, and the ordinary height so created is exactly the resource that permits the large valuations needed to repay the deficit.

Remark 8.2. What this rules out is a class of arguments, not a theorem: any argument of the form "a sufficiently deep drift excursion cannot recover, because the required valuation exceeds what the current iterate can supply" is blocked before it starts. Over all odd m with $3 \le m < 200001$, the minimum over all steps of $\log_2(3m_n + 1) - (\alpha - R_n)$ — available valuation minus what a one-step return to $R = 0$ requires — is $+1.7370$, attained at the seed step of $m_0 = 3$, and is never negative. (Revision 6 reported $+1.8301$; that is the minimum over steps $n \ge 1$ only, and silently omitted the seed step. Revisions 5–7 wrote the range as all odd $m < 200001$, which also admits $m_0 = 1$, where the margin is $+0.4150$; only the quoted minimum was affected, not the sign. The qualitative conclusion is unaffected.) [the inequality is proved; the margin is a computational diagnostic]

9 Boundaries of the programme

This section collects the strongest no-go results. Their scope is deliberately narrow: each says that a specific class of information cannot distinguish positive-integer realization. None says that Collatz or EOC is unprovable.

9.1 Finite cylinders do not encode infinite realization

Proposition 9.1. *For $k \ge 1$, every odd residue class mod 2^k contains arbitrarily large positive odd integers. Consequently every exact-realizer cylinder is nonempty and unbounded, and no property of integers determined by finitely many 2-adic digits can be equivalent to positivity.* *[proved; formalized, §14]*

Proof. $r + 2^k(B+1)$ lies in the class and exceeds B; oddness is preserved since $2 \mid 2^k$. For the second claim, 1 and $1 - 2^k$ share a class mod 2^k and differ in sign. □

The consequence is a quantifier distinction that should be displayed:

$\forall N\ \exists m_N$ (free)	versus	$\exists m\ \forall N$ (the whole problem)

A valuation prefix of length N pins exactly one odd class mod 2^{S_N+1}, which contains unboundedly many positive integers; so finite-prefix realizability never obstructs. The content lies entirely in whether *one* integer works for all N, which is the bounded-prefix-realizer equivalence of the companion work.

Remark 9.2 (Scope). This does *not* say finite information is useless. Finite data can still prove quantitative lower bounds — Open Problem C is precisely such a statement at each finite N. The claim is only about *deciding* the infinite realization property. [scope statement]

9.2 Coordinate sufficiency is not independent information

Several coordinates introduced in the programme turn out to be invertible reparameterizations of the exact realizer. The ledger below records each, with the reference where the collapse is proved.

Coordinate	Status
Normalized carry Z_N	Same moving-anchor information as $-C_N 3^{-N}$; no independent obstruction.
Lift digits t_N	Exactly the binary digits of a fixed realizer, read in blocks cut at the valuation sums.
Eventually-zero lift tail	Equivalent to eventual constancy of the prefix realizers, hence to positive-integer realization.
Terminal zero gap $g(D)$	$\lceil \mathsf{E}(D) \rceil$; a restatement of the target, not a mechanism (§5.4).
Transport depths F, M, H	Exact bookkeeping. On a genuine orbit the matching precision M_j is the length of the zero-run of the *seed*'s binary expansion beginning at bit S_j.
Cumulative excess regeneration	Bit intervals at distinct failures are disjoint, so the total is at most $\lfloor \log_2 m_0 \rfloor + 1$.
$B_N = C_N + 3^N \chi$	Equals $2^{S_N} m_N$ exactly for a realizing χ; hence $v_2(B_N) = S_N$ exactly and $\log_2 \lvert B_N \rvert - v_2(B_N) = \log_2 m_N$.

The last row deserves emphasis. The Archimedean size and the 2-adic valuation of B_N are *complementary factors of one exact identity*, not two independent constraints; knowing the word and either one determines the other. This is why attempts to intersect a p-adic condition with an Archimedean one at this object produced nothing: there was only ever one condition. It also explains why the product formula yields exactly the trivial bound $v_2(B) \le \log_2 |B|$ here, and why the odd-prime support of B_N — being that of the arbitrary odd integer m_N — is unconstrained. [proved in companion audits; identities verified computationally without exception]

9.3 The known divergent-orbit word shadows do not characterize positive realization

Theorem 9.3 (Word-shadow cardinality barrier)**.** *For $b \subseteq \mathbb{N}$ define the valuation word (digits indexed from 0)*

$$w_b(i) = \begin{cases} 2 & \text{if } i \equiv 2 \pmod 3 \text{ and } \lfloor i/3 \rfloor \in b, \\ 1 & \text{otherwise.} \end{cases}$$

Then every w_b is an admissible valuation word satisfying

(a) $\sum_{i<n} w_b(i) \le n + \lfloor n/3 \rfloor$, *hence* $2^{S_n} \le 3^n$ *for every n: zero confinement, in the integral form of Lemma 2.5;*

(b) $R_n \le (4/3 - \alpha)n = -0.2516\ldots n \to -\infty$ *linearly;*

(c) $\sum_n 2^{R_n} < \infty$;

(d) $\#\{n : R_n \ge -G\} = O(G)$.

The map $b \mapsto w_b$ is injective, so the family has cardinality $2^{\aleph_0}$. Since one positive integer realizes at most one infinite valuation word, all but countably many members have no *positive-integer realizer.* [proved. Formalized: item (a) (as `famWord_sum_le` with `famWord_confined`), injectivity, and the abstract countability argument. Items (b), (c), (d) are proved here but not formalized, and the final sentence is not instantiated at w_b in Lean]

Proof. (a) At most one digit in three is a 2, so $S_n \le n + \lfloor n/3 \rfloor$; and $2^{n+\lfloor n/3 \rfloor} \le 3^n$ follows by cubing both sides, since $3(n + \lfloor n/3 \rfloor) \le 4n$ and $16 \le 27$. (b) $4/3 < \alpha$. (c) The bound in (b) makes 2^{R_n} dominated by a geometric series. (d) $R_n \ge -G$ forces $n \le G/0.2516$. Injectivity: evaluate w_b at $i = 3k + 2$. The last sentence is Cantor's theorem: a realizing assignment would give an injection $\mathcal{P}(\mathbb{N}) \to \mathbb{N}$. □

Remark 9.4 (Where a divergent word must actually live). Theorem 9.3 says the *known* shadows are too coarse. Two further facts say where the shadow of a genuine divergent orbit must sit, and they sharpen the scope recorded in Remark 9.5 without changing it. Dubickas [6], Theorem 5, shows that the parity word of a divergent orbit of the shortcut map has factor complexity $P(X, n) > 1.70951129\, n$ for all large n, the constant being $1/\log_2(3/2)$; the mechanism is the growth bound $x_n < (3/2)^n(x_0 + 1)$ against mod-2^L injectivity of length-L parity factors (Bernstein–Lagarias). The density half needs nothing external: by Proposition 4.11 the restarted seed satisfies $S'_k \le \lfloor k\alpha \rfloor$ for every $k \ge 1$, and a block of k accelerated steps occupies S'_k ordinary steps of which k are odd, so the ones-ratio is $k/S'_k \ge 1/\alpha$ from the last drift maximum onwards. Monks and Yazinski [7], Theorem 2.7(b), give the generalization to rational seeds: an $x \in \mathbb{Q}_{\text{odd}}$ with infinite orbit has $\liminf_n \kappa_n(x)/n \ge \log 2/\log 3 = 1/\alpha$, κ_n counting the odd iterates among the first n. Either way the ones-density of a divergent orbit is at or above the critical ratio. A divergent-orbit word must therefore lie *above* the Dubickas complexity threshold and *at or above* critical density — outside every word class for which Φ-irrationality is currently proved. Both halves of Dubickas's mechanism are now formalized in [11] (`Divergence/GrowthBound.lean`; `modEq_of_parityPrefix_eq` in `Divergence/RawMap.lean`), with the pigeonhole step joining them left unformalized. [external; both read first-hand]

Remark 9.5 (What this does and does not establish). The shared global word-level shadows established in Theorem 9.3 — zero confinement, linear negative drift, summability of $\sum 2^{R_n}$, and the stated occupation bound — therefore do not characterize positive-integer realization. More generally, any proposed criterion depending only on properties shared by this family cannot do so. This is a statement about characterization, not about quantitative bounds: it does not preclude finite or word-level information yielding realizer lower bounds. [scope statement]

Remark 9.6 (The barrier is sharply delimited). It is worth saying how far Theorem 9.3 reaches, because a broader reading has circulated. Since the least prefix realizer $r(D)$ is a function of the valuation word alone, *bounded prefix realizers* is itself a global word-level property, and it characterizes positive-integer realizability *exactly*: a word has bounded prefix realizers if and only if some positive odd integer realizes it (§9.1). It escapes the theorem because it is not shared by the family w_b — it holds for the countably many realizable members and fails for the rest.

So the barrier does not say that word-level information is insufficient in principle. It says that the *divergent-orbit shadows* — the properties a hypothetical Type-II orbit's word is known to have — are too coarse, while the one word-level invariant sharp enough to decide the question is a restatement of realizability rather than an independent handle. [Revision 7, scope]

Remark 9.7 (Explicit small witnesses). The barrier is not an artefact of uncountability. The all-ones word (realizer -1) and the word $1, 2, 1, 2, \ldots$ (realizer -5) are both zero-confined with $R_n \to -\infty$ linearly, and neither has a positive-integer realizer. [computational verification of an exact statement]

9.4 Population sparsity is not pointwise placement

Theorem 4.7 and the entropy constant of §3.3 are population statements: they count, respectively, orbit values in a window and words in a shell. Open Problem C is a statement about the *minimum* over a family. A density-zero exceptional set may contain every least-realizer champion, and no

argument currently rules this out. The same gap defeats density-one results in the literature (§16) as applied to launch integers, which by construction form an exceptional set. [structural interpretation]

9.5 Deep drift creates its own recovery capacity

§8. [proved]

9.6 Single-window lifetime is not the same functional as occupation

Remark 6.2. This concerns the two *functionals*, witnessed by one seed; it is *not* the claim that the universal bounds are inequivalent, which is open (Rem. 10.1). [Revision 5, computational example]

Remark 9.8 (Net effect). These boundaries narrow the remaining search to actual positive-orbit arithmetic and to genuinely pointwise or cross-scale mechanisms. They do not constitute an argument that the programme cannot succeed.

10 The programme hierarchy

Level	Statement	Status
0	Finite word realizability	Solved exactly (Prop. 5.2); every cylinder is nonempty and unbounded (Prop. 9.1).
1	Universal drift exit (DE)	Equivalent, given the external Thm. 4.7, to excluding divergent orbits (Thm. 6.15); no implication to cycle-freeness established (Rem. 6.17); *open*.
2	Quantitative single-window: bounds on $L_c(m)$	Transfers to frontier bounds (Prop. 6.6); no unconditional bound known (Rem. 6.8).
3	Exponential single-window floor	Open Problem C; equivalent to $L_c(m) = O(\log m)$ for odd $m \geq 3$, words taken with $r(D) \geq 3$ (Prop. 6.5).
4	Total occupation $\mathcal{O}_c(m) = O(\log m)$	Full EOC (Conjecture 3.1); implies Level 3, and implied by its corridor-uniform form *given cycle exclusion*; strictness is *open* (Rem. 10.1).
5	Effective EOC + finite cycle verification	Collatz (Thm. 3.6, Cor. 3.7).

Levels $5 \Rightarrow 4 \Rightarrow 3 \Rightarrow 2$ and $3 \Rightarrow 1$ hold; $2 \Rightarrow 1$ holds by Remark 6.8. *No converse is claimed*, and the reader should note what that does and does not mean.

Remark 10.1 (Strictness is not established by the pointwise gap). Revision 6 asserted that "no converse holds in general" and that Level 4 is *strictly* stronger than Level 3. Those assertions are withdrawn as stated. The evidence offered for them was the pointwise gap $\mathcal{O}_c > L_c$ (Rem. 6.2), and that establishes only that the two *quantities* differ; it does not show that the two *universal asymptotic bounds* "$L_c(m) = O(\log m)$ for all m" and "$\mathcal{O}_c(m) = O(\log m)$ for all m" are logically inequivalent. Proving inequivalence would require exhibiting a consistent state of affairs in which every single-window lifetime is logarithmic while some total occupation is not, and no such argument is given here or known to us.

Moreover the evidence points the other way for one natural form of Level 3. If the lifetime bound is assumed *corridor-uniformly* — for every $c' \geq 0$ simultaneously, not at one fixed corridor — then Corollary 7.4 bounds the injective initial segment, and Proposition 7.7 converts that into Level 4 *given cycle exclusion* (Rem. 7.8(i)). So corridor-uniform Level 3 implies Level 4 under that additional assumption, and for that form strictness is at best unproved rather than established.

What survives: Level 4 implies Level 3 (immediately, since $L_c \leq \mathcal{O}_c$); no implication from Level 1 to cycle-freeness is established here, so the drift-exit route alone does not reach Collatz (Cor. 6.16, Rem. 6.17) — while *neither* strictness *nor* non-entailment is claimed; and for a *fixed* corridor the converse Level 3 $\Rightarrow$ Level 4 is open. [Revision 7, correction of scope]

11 The programme map

Axis	Object	Exact bridge	Strongest proved	Open target	Closed routes
Occupation	$\mathcal{O}_c(m)$	Thm. 3.6 to Collatz	Reduction to a finite cycle search, given effective constants	Separated-return control	—
Time	$L_c(m)$, $\tau_0(m)$	Record chronology (Prop. 6.4); transfer (Prop. 6.6)	(DE) $\iff$ no divergence (Thm. 6.15)	Any unconditional bound on L_c	Self-financing recovery blocks excursion arguments (§8)
Residue	$r_{\min}(N, c)$, $r(D)$	Exact congruence (Prop. 5.2) + inversion	Fixed-anchor sectors (Thms. 5.4–5.5)	Finite-defect (§5.2) and arbitrary irregular floors	Coordinate collapse (§9.2)
Divergence	R_n, Q_n, value set	Curry sparsity; last-maximum reduction (Prop. 4.11)	Reciprocal summability; $B < 1/\beta_*$ (Thm. 4.9)	Exclude divergence	The *known* shadows are insufficient (Thm. 9.3)
Symbolic	Valuation words; return labels	Realizer correspondence	Finite and infinite symbolic freedom	Positive realization	Cardinality barrier (Thm. 9.3)

11.1 Symbolic freedom, finite and infinite

Revision 5 established that no *finite* forbidden-pattern mechanism can settle pointwise balance for a natural mod-32 return observable: every finite binary history is realizable by a genuine orbit, with prescribed events consecutive, and this was formally verified. The symbolic conclusion is now stronger. A machine-checked construction produces an infinite zero-confined valuation sequence carrying arbitrary prescribed binary labels at designated event positions: a *full shift* at the word level, inside the zero corridor.

The distinction that matters is therefore sharper than "no finite forbidden patterns": even substantial *infinite* symbolic freedom coexists with zero confinement. Positive-integer realization of such words remains open and, by Theorem 9.3, is not decided by the global word-level shadows established there. The missing restriction is not local symbolic dynamics. [finite universality: Revision 5, formalized; infinite full shift: companion formalization]

12 Closed routes

For each mechanism, what it achieved and where it stops. Status codes: T formal result, C computational diagnostic, I interpretive closure.

	Mechanism	Where it stops	
1	Fixed-modulus drift packing	Bound 8/9 proved; *terminality* of the method rests on a structural argument about primes $p > 3$ (Rem. 4.6), not a theorem about all moduli; surpassed by a different mechanism	T/C
2	Bulk entropy / I_{Collatz}	Population scale only; no pointwise anti-concentration	I
3	Shell injectivity	Counts, does not place	I
4	Moving-anchor collision law	Symbolic prefix geometry only	I
5	Chang finite/local exclusions	Finite *and* infinite symbolic freedom (§11.1)	T
6	Lift digits	Binary blocks of the realizer (§9.2)	T
7	Terminal zero gap	$\lceil \mathsf{E} \rceil$: target, not mechanism (§5.4)	T
8	Normalized carry	Same anchor information	T
9	Transport survival budget	Exact bookkeeping; on genuine orbits matching precision collapses to seed bits, cumulative excess $O(\log m_0)$	T
10	External p-adic arithmetic	No applicable uniform bound found; the dimension grows with N (§12.1)	I
11	The *known divergent-orbit* word shadows	Cardinality barrier (Thm. 9.3); word-level criteria in general are *not* excluded (Rem. 9.6)	T
12	Corridor-uniform lifetime hypothesis	Bounds the injective initial segment on an $O(\log m)$ scale (Cor. 7.4); reaching 1 needs cycle exclusion (Rem. 7.5)	T
13	Deep-drift recovery impossibility	Self-financing (§8)	T

12.1 External arithmetic

We record only the durable conclusion. The natural cleared-denominator object attached to a confined word is $B_N(D, \chi) = C_N + 3^N\chi$, an $(N{+}1)$-term sum of $\{2, 3\}$-units whose 2-adic valuation one would like to bound. Three families of deep results were examined against it. p-adic lower bounds for linear forms in logarithms bound $\alpha_1^{b_1} \cdots \alpha_n^{b_n} - 1$, a *product* minus one, never a sum of three or more terms; their constants are at best geometric in the number of logarithms; and at $p = 2$ with 2 among the bases the hypotheses fail and the quantity is identically zero. Quantitative results on S-unit equations bound the *number* of non-degenerate solutions, not their size, and are ineffective beyond two variables; moreover explicit lower bounds for that count forbid any bound uniform in the number of variables. The p-adic Subspace Theorem does apply formally and yields a bound of the right shape, but it is ineffective, its subspace count and height threshold depend on the dimension, and here the dimension is $N{+}1$. In every case the obstruction is the same: the number of additive terms grows with N, while the available theorems are quantified "for each fixed dimension" with constants that provably blow up. [applicability audit; external theorems cited but not reproved]

Verger-Gaugry [9] reaches the same arithmetic from the opposite direction and is consistent with this obstruction. Assuming a divergent trajectory exists, he shows a three-term Diophantine exponential equation $x + y = w2^n$ governs it, and applies the Bennett–Billerey finiteness theorem on S-units to the primes generated by a collection of local minima ("wave-minima"). A *wave* of height $h \geq 2$ is a local minimum μ followed by $h - 1$ odd iterates and then an even one; his Theorem 1.2(ii) gives $h = \nu_2(\mu + 1)$ and the identity $3^h(\mu + 1) - 2^h = 2^h T^h(\mu)$, and his Theorem 1.2(iii) concludes that the set of primes generated by the wave-minima is infinite *conditional on the wave heights being bounded* ($h \leq M$ for some $M \geq 2$). On that basis he defines a density $\delta_{\mathcal{T}}$ of a trajectory and states conjectures 3.1 and 3.2 about it. The contrast with the paragraph above is instructive: his equation has three terms because it is attached to one wave rather than to the whole prefix, which is exactly the dimension restriction the available theorems need, and exactly what $B_N(D, \chi)$ does not satisfy. [external, refereed, cited but not reproved]

13 Open problems

Open Problem A (Effective EOC constants)**.** Establish Conjecture 3.2 at any explicit (c, K, C), or any explicit upper bound on admissible constants in Conjecture 3.1. With the cycle search below

the resulting m^* also completed, this closes the Collatz implication — Corollary 3.7 does complete it at $m^* = 90$; in general that search is Open Problem B.

Open Problem B (Cycle exclusion below threshold)**.** Show no nontrivial accelerated cycle has minimum member below $m^*(K, C, c)$. Settled for $m^* \le 2^{71}$ by [22]; open beyond.

Open Problem C (Exponential realizer floor / logarithmic lifetime)**.** Find $\varepsilon > 0$ with, equivalently (Proposition 6.5),

$$\text{(residue) } r_{\min}(N, c) \;\ge\; 2^{\varepsilon N - O(1)} \qquad \Longleftrightarrow \qquad \text{(time) } L_c(m) \;\le\; \varepsilon^{-1} \log_2 m + O(1),$$

for all relevant c-confined words. This is the residue–time core of EOC. Revision 5's separate Open Problems C and D are the same question at two levels of precision and are merged here; Revision 5's Open Problem E (frontier defect) is its sharp-leading form, not an independent weaker target (Remark 5.10).

Rozier's Lower Bound Hypothesis [2, Hyp. 2.3] is this problem stated uniformly over all classes (j, q): there is $C \ge 0$ with $n \ge j^{-C} 2^{(1-H(q/j))j}$ for all j and all n, q being the number of odd terms among the first j iterates. His Lemma 3.1 — LBH implies every positive integer reaches the trivial cycle — is this paper's Level 1 implication at $c = 0$ (§10). Two remarks. First, the entropy points agree where it matters, and two distinct critical ratios should not be confused. At the ones-ratio $1/\alpha \approx 0.6309298$ — the ratio forced on a divergent orbit, and the one Rozier's Lemma 3.1 uses — LBH's exponent is $1 - H_2(1/\alpha) = 0.0500445\ldots$ per T-step, hence $\alpha(1 - H_2(1/\alpha)) = I_0 = 0.0793186\ldots$ per accelerated step, which is exactly the identity $I = \alpha(1 - H_2(1/\alpha))$ of §3. Rozier's $\gamma_* \approx 0.6090897$ is a *different* ratio: it is the Garcia–Tal/Curry sparsity exponent of Theorem 4.7, the root of $H(\gamma) = \gamma \log_2 3$, and there $(1 - H(\gamma_*))/\gamma_* = 0.0568316\ldots$. Second, Rozier's observed polynomial loss ($C \ge 0.980916600\ldots$ from all $n \le 10^9$) and the $+\frac{3}{2}\log_2 N$ of the sharp form here are stated in different normalizations — his j counts ordinary steps, this paper's N counts accelerated steps — and we do not assert a correspondence between them.

Open Problem D (Intermediate time-axis bounds)**.** Prove *any* unconditional pointwise upper bound $L_c(m) \le F(m)$ with F nondecreasing. By Proposition 6.6 even $L_c(m) \le Km$ yields a linear frontier, and $L_c(m) \le K(\log_2 m)^2$ yields a stretched-exponential one. Nothing of this kind is known (Remark 6.8), and any such bound would be new.

Open Problem E (Universal drift exit)**.** Show that for every positive odd m there is $n \ge 1$ with $3^n < 2^{S_n(m)}$. By Theorem 6.15 this is equivalent to excluding divergent orbits; no implication from it to cycle-freeness is established here (Rem. 6.17). This is the previously posed zero-corridor realizability frontier, reclassified as the qualitative endpoint of the time axis (Remark 6.18); it is not newly invented here.

Open Problem F (Separated returns)**.** Control the *weighted* total $\mathcal{O}_c = \sum_i \ell_i$ of the confined episodes. Revision 6 asked for control of the episode count; the results below fix what such control can and cannot be.

Proposition 13.1 (Entry lemma)**.** *Let $a > 0$ be a re-entry time, so $R_{a-1} > c \ge R_a$. Then $d_{a-1} = 1$, and consequently*

$$0 \;\le\; c - R_a \;<\; \alpha - 1 = \log_2(3/2) = 0.58496\ldots$$

Proof. *$R_a = R_{a-1} + d_{a-1} - \alpha$ with $1 < \alpha < 2$ and d_{a-1} a positive integer; a digit ≥ 2 would give $R_a > R_{a-1} > c$.* □

So only the first *episode has threshold c: every later episode begins in a corridor of width less than* 0.585*, a bound independent of c, of m_0, and of how deep the drift fell during the excursion. Drift may of course go far below c* inside *an episode.* *[Revision 7; formalized, §14]*

Proposition 13.2 (Episode counts are unbounded)**.** *Let $e_0 = e_0(c)$ be the first exit time of the oscillating word below. For every $c \geq 0$, every B, and every prescribed bound M, there are N and a positive odd $m_0 > M$ whose genuine accelerated orbit follows that word for N steps and satisfies*

$$\#\{\textit{corridor re-entries before } N\} \;\geq\; B, \qquad \#\{n < N :\; R_n \leq c\} \;\leq\; e_0 + 2B.$$

Both conclusions concern the realized prefix $[0, N)$ only.

Proof sketch. *The word playing digit 2 inside the corridor and digit 1 outside has drift confined to a fixed two-sided window, re-enters immediately after each exit, and exits again within three steps, so exits recur at gaps of at most 3 and $N \leq e_0 + 3B + 1$ suffices for B re-entries; the prefix occupation is then N minus the number of exits, which is at most $e_0 + 2B$. Every finite prefix is realized by a positive odd integer via Proposition 5.2, and since the realizers of a prefix form a full residue class, one may take $m_0 > M$ without changing the word or either count.* □*[Revision 7; formalized, §14]*

Remark 13.3 (What this rules out, and how far). A *uniformly bounded* episode count is therefore false, and cannot be assumed as a closing hypothesis. Two scope points must travel with this.

Finite prefix versus total occupation. Propositions 13.1 and 13.2 and everything in the next paragraph concern the realized prefix $[0, N)$. Nothing is claimed about the *total* occupation of the constructed seeds, which depends on the orbit beyond N.

One particular coarse accounting is lossy, by a stated factor. Fix K and suppose the corridor-uniform single-window bound charges each episode $\ell_i \leq 1 + K(\log_2 m_{a_i} + c_i)$. Define the *surrogate total* for the realized prefix as the sum of those charges,

$$Q \;:=\; \sum_{i=1}^{P} \Big(1 + K\big(\log_2 m_{a_i} + c_i\big)\Big), \qquad P = \#\{\text{episodes in } [0, N)\},$$

so that $\mathcal{O}_c^{\text{prefix}} \leq Q$ is precisely what this route would use. Since $R_{a_i} \leq c$ and $E_{a_i} \geq 0$ give $\log_2 m_{a_i} + c_i \geq \log_2 m_0 - c$, every charge is at least $\log_2 m_0 - c$, whence $Q \geq P(\log_2 m_0 - c)$.

Now write e_0 for the first exit time of the oscillating word (a constant depending on c alone). Proposition 13.2 supplies, for each B and each prescribed bound, an odd realizing seed *exceeding that bound* whose prefix has $P \geq B$ episodes and prefix occupation $\mathcal{O}_c^{\text{prefix}} \leq e_0 + 2B$. Since $(2 + e_0)(B + 1) \geq e_0 + 2B$,

$$\frac{Q}{\mathcal{O}_c^{\text{prefix}}} \;\geq\; \frac{B(\log_2 m_0 - c)}{e_0 + 2B} \;\geq\; \frac{\log_2 m_0 - c}{2 + e_0},$$

in which B cancels. Along a sequence of such seeds with $m_0 \to \infty$ the ratio is therefore unbounded: *this* accounting overestimates the prefix occupation it bounds by a factor $\Omega(\log m_0)$, and no choice of K repairs it.

Two things this does *not* say. It is not a statement about per-episode methods in general — only about the single-window charge just displayed; a different per-episode estimate is not excluded. And it is not an impossibility theorem about EOC.

A distinction the argument depends on. The seeds above are obtained by moving *within the residue class* of a fixed prefix (Prop. 5.2), which makes m_0 as large as desired while leaving the word, and hence P and $\mathcal{O}_c^{\text{prefix}}$, unchanged. That is not the same as showing the *least* realizer grows, and nothing here asserts that it does.

What the problem needs, then, is not a bound on the episode count but a bound on $\sum_i \ell_i$ that does not factor through per-episode single-window estimates. [Revision 7]

Open Problem G (Density-one EOC)**.** Can a density-one form of EOC be proved by the methods of Tao [17]? This remains meaningful: divergent-orbit sparsity says nothing about the occupation of orbits that do converge.

Open Problem H (Pointwise return balance)**.** For an infinite aperiodic orbit, does the natural mod-32 return label equidistribute along the orbit? By §11.1 no finite symbolic restriction settles this, and neither do the *known divergent-orbit* word shadows (Thm. 9.3); word-level criteria in general are *not* excluded (Rem. 9.6). Demoted here from a top-level EOC problem to an auxiliary one.

14 Formal verification status

The accompanying Lean 4 development (Mathlib) formalizes the elementary bridges of this paper. The following are machine-checked, with no `sorry`, `admit`, added axiom, or `opaque` escape; each depends only on the standard axioms `propext`, `Classical.choice`, `Quot.sound` (several on a strict subset).

Two procedural notes, added in Revision 7 after an audit of this table. First, every module named below is now imported by the library root, so the default build target checks them; previously all but one were reachable only by building them individually by name, and "machine-checked" was not evidenced by the repository's default target. Second, the axiom claim is established by `#print axioms` run on each theorem; those commands are not themselves committed, so the claim is reproducible but not self-evidencing from the source tree. Descriptions below name the hypotheses each formalization actually carries.

This paragraph is corrected in Revision 9. Revisions 7–8 recorded that several entries are conditional on `ReciprocalSummable`, Theorem 4.7's conclusion, which was then "not formalized". That remains true of the development described in this table (`eoc-lean-verification`, branch `rev7-finalization`), but it is no longer true of the programme: the separate repository `eoc-divergence` [11], `main` at `e56950f`, proves the chain with no hypotheses — `exists_window_sparsity` (the window bound with absolute base $\vartheta = 3^{31/50}$), `summable_inv_orbit`, `uniform_partial_sums` (absolute constant $\Sigma_* = 4/(1-\vartheta/2)^2 \approx 28077.6$), `exists_zeroConfined_of_divergent` and `divergent_iff_zeroConfined`, each with axioms `propext`, `Classical.choice`, `Quot.sound` and continuous integration passing. The scope of that is bounded, and the bound matters:

statement	status in [11] at `e56950f`
Prop. 4.8 (reciprocal summability)	formalized unconditionally
Prop. 4.11 (last-maximum reduction)	formalized unconditionally
Thm. 6.15 (divergence equivalence)	formalized unconditionally
Thm. 4.9 at the sharp $1/\beta_* \approx 1.0358567$	**not covered**: the formalized exponent is $\log_2 \vartheta = 0.98268\ldots$, which reaches only $B < 1/0.98268 \approx 1.0176$; the sharp constant remains conditional on Theorem 4.7 as stated in [14], which is unformalized

Module	Paper reference	Content
`CurryFoundation`	Prop. 4.11	Curry identities; last-global-maximum existence (from $R_n \to -\infty$ alone); zero-corridor tail and deficit results. The hypothesis `ReciprocalSummable` these carried is discharged in [11] by `summable_inv_orbit`.
`ChangHistory`, `ChangFullShift`	§11.1	Finite return-history universality; infinite zero-corridor full shift.
`ZCRERealizerGrowth`	§9.1	Bounded prefix realizers $\iff$ positive-integer realizability.
`RealizerLiftDigit`, `TerminalZeroGap`	§9.2, §5.4	Lift digits are binary blocks; the discrete gap. (The bracket $g = \lceil \mathsf{E} \rceil$ is in `TransportRegeneration`, not here.)
`TransportRegeneration`, `SuperBudgetCancellation`	§9.2	Transport bookkeeping; the gap bracket; disjoint-run bound (disjointness is a hypothesis, not a conclusion). The forced/excess split is *not* formalized.
`ConstraintHeight`	§5.3	Height inequality and the separation criterion.
`GlobalSeparation`	Prop. 9.1, Thm. 9.3	Cylinder positivity; finite-precision sign impossibility; the family w_b, its confinement and injectivity; the *abstract* countability argument (not instantiated at w_b).
`OrbitLifetime`	Lem. 2.4, Prop. 7.3	Aggregate identity; growth envelope; automatic corridor.
`DriftExit`	Props. 6.9, 8.1, 6.6	Future-minimum theorem; cycle drift; self-financing inequality; frontier inversion.
`SeparatedReturns`	Props. 13.1, 7.6, 13.2	Restart relation; entry lemma; the corridor value criterion (`corridor_value_lower`); occupation bounded by the arrival time n^* *under* $1 \le n^*$, arrival at 1 from n^* onwards, and $2^c \le m_0$; the oscillating word and unbounded realizable episode counts.
`ZeroConfinedSeed`	Prop. 4.11	The restart at the last drift maximum, as an *actual positive odd seed*: zero-confined at every horizon, with drift $\to -\infty$. Carried here under `ReciprocalSummable`; discharged unconditionally in [11] as `exists_zeroConfined_of_divergent`.
`CurryInterface`	§4, Thm. 6.15	Injective $\iff$ divergent; the external inputs as explicit hypotheses; divergence $\Rightarrow$ a zero-confined positive seed; universal drift exit $\Rightarrow$ no divergent orbit. Those explicit hypotheses are discharged in [11] by `divergent_iff_zeroConfined`, which carries none.
`PowerSavingSummable`	Prop. 4.8	The dyadic summation: a windowed power-saving value count, with injectivity, yields reciprocal summability.

Remark 14.1 (The divergence chain, layer by layer). Revision 6 listed reciprocal summability (Prop. 4.8) as an unformalized external input. It no longer is: `PowerSavingSummable` *derives* it in Lean from the windowed sparsity count, by the dyadic summation, with injectivity carrying the value-set sum to the orbit-indexed one. The ledger is now

$$\underbrace{\text{windowed sparsity}}_{\text{at base } \vartheta\text{: formalized [11]; at Curry's sharp } \beta_*\text{: external}} + \underbrace{(\text{DE})}_{\text{open hypothesis — the target}} \implies \text{no divergent orbits},$$

and both must be retained in any summary. The two are of different kinds: formalizing Theorem 4.7 would remove the *external* dependency, and would neither discharge (DE) nor prove divergence exclusion. Revision 9 narrows the left-hand label: the window bound *at the absolute base* $\vartheta = 3^{31/50}$

is formalized in [11], which suffices for Proposition 4.8 and hence for this ledger; Curry's sharp exponent β_*, which Theorem 4.9 needs, remains external and unformalized. [Revision 7; scope narrowed in Revision 9]

Statements in this paper that are *not* formalized include Theorem 4.7 *at Curry's sharp exponent* β_* (external input; audited from the source text — the weaker form at base ϑ *is* formalized in [11], §14), Corollary 7.4's analytic absorption step (Rem. 7.8(ii)), the proved-sector floors of §5 (companion note), and every conjecture.

15 Computational calibration

We retain only computations that materially support interpretation, and label each as a diagnostic.

- *Frontier records.* A direct seed scan, justified by Proposition 6.4, produces the least realizers of confined words at each depth. The ratio of occupation to $\log_2$ seed passes 10 once the corridor is widened. Exhaustively over every odd $3 \le m < 2^{33}$, in exact integer arithmetic, the maximum of $\mathcal{O}_c(m)/\log_2 m$ is 9.7588, 9.9902, 10.1059, 10.4078 and 11.3402 for $c = 0, 1, 2, 4, 8$ respectively; all but the $c = 4$ maximum are attained at $m = 63728127$, and that one at $m = 4890328815$. Every value lies below the sharp-leading constant $1/I(\alpha) = 12.6074\ldots$, and no trend is inferred beyond the scanned range. Revisions 5–7 reported this ratio as remaining under 10: that holds for $c \le 1$, but fails from $c = 2$ on.

- *Identity verification.* The identities of Lemmas 2.2, 2.4 and 2.5, the future-minimum theorem, the descent criterion $m_n < m_0 \iff R_n > E_n$, and the factorization $B_N = 2^{S_N} m_N$ were each verified without exception over the ranges scanned.

- *Self-financing margin.* §8, Remark 8.2.

- *Drift-exit diagnostics.* $\tau_0 = \sigma$ was *observed* for every odd seed tested away from cycle minima ($3 \le m < 60001$; 0 exceptions) — superseded, under the identification of §16, by Rozier–Terracol [5, Cor. 5.4], which gives the same equality for every $2 \le n \le 2.8 \times 10^{19}$ — consistent with Remark 6.13. What is *proved* there is only $\tau_0 \le \sigma$ in general, together with the polynomial window to which any strict example is confined; the equality is not proved for all seeds.

Remark 15.1 (Discipline). We make no empirical saturation claims. In particular we do not assert that any normalized stopping-time or lifetime ratio is bounded: a comparison of a limited scan against published record tables corrected an earlier such claim in the underlying audits. Where a ratio is quoted, the scanned range is stated and no trend is inferred beyond it. Units also require care: the classical total stopping time σ_∞ counts *even* steps — equivalently the valuation depth S_L — whereas L counts *odd* accelerated steps; the two must not be compared without conversion. [methodological]

16 Relation to prior work

The occupation functional and its four tiers are, to our knowledge, not isolated elsewhere; the Collatz reduction of Theorem 3.6 is ours. Lemma 2.2 is due to Eliahou [1] and Rozier [2]. Theorem 4.7 is Garcia–Tal [13] with Curry's explicit form [14]. The 2-adic conjugacy underlying the realizer correspondence is Bernstein [29] and Bernstein–Lagarias [30]; cycle-structure bounds are Steiner [18], Simons–de Weger [19] and Hercher [20]; verification is Bařina [21, 22]; stochastic models are

Lagarias–Weiss [15] and Kontorovich–Lagarias [16]. Sturmian words are treated in Lothaire [23]. The return observable and its structural reduction are Chang's [31, 32].

One qualification on the divergent-orbit word (Remark 9.4). "At or above" the critical density is a genuine disjunction, and the two branches behave differently: density exactly $1/\alpha$ is the subexponential case, in which m_n grows slower than any exponential, while density strictly above $1/\alpha$ forces m_n to grow exponentially. What the refereed literature supplies is only the lower bound — $\liminf \geq 1/\alpha$, by Monks and Yazinski [7, Thm. 2.7(b)] for rational seeds, and by Proposition 4.11 here for integer ones — so the exponential branch is *open*. López and Stoll [8, Thm. 1] claim to exclude it, asserting density exactly $1/\alpha$ for a divergent rational seed; that theorem has two halves, of which the half below $1/\alpha$ is quoted there from [7], while the half above $1/\alpha$ rests on a step from the real value to the 2-adic one that has been questioned. Nothing in this paper depends on it.

On the time axis specifically: the condition "$\exists n : S_n > \alpha n$" *does* have a standard name, which Revisions 5–8 missed. It is Terras's *coefficient stopping time* κ [3, 4]; the quantity written τ_0 here is κ, and Proposition 6.11 ($\tau_0 \leq \sigma$) is Terras's $\kappa \leq \sigma$. It is not Terras's stopping time σ [3], which is the condition $m_n < m_0$, i.e. $R_n > E_n$. The diagnostic $\tau_0 = \sigma$ of Remark 6.13 and §15 is the *Coefficient Stopping Time Conjecture*, named by Lagarias; Rozier and Terracol [5] verify it for every n with $2 \leq n \leq 2.8 \times 10^{19}$ (their Cor. 5.4), which supersedes the scan reported in §15. Their "paradoxical" sequences — ones-ratio below $\log 2/\log 3$ yet $T^j(n) \geq n$ — are exactly drift exit without descent, the gap Proposition 6.11 measures. The classical results on stopping time — Terras [3], Everett [12], and in a different framework Tao [17] — are density statements and carry no pointwise content; none bounds the number of steps an orbit may spend at or above its starting value in terms of that value, and such a bound would prove Collatz by strong induction. [literature audit]

17 Conclusion

EOC is best viewed as a programme with an exact residue–time duality and a global occupation layer above it — one that implies the single-window layer, and whose converse at a fixed corridor is open rather than refuted (Rem. 10.1). The residue axis is solved for finite realizability and for several fixed-anchor sectors, and remains open at the arbitrary-word frontier. The time axis separates into a qualitative drift-exit problem — equivalent, under the current divergence theory, to excluding divergent orbits, with no implication to cycle-freeness established — and quantitative lifetime bounds that invert directly to realizer floors on a graded scale. Garcia–Tal/Curry supplies strong unconditional divergence-side sparsity and the last-maximum reduction, but not drift exit.

A sustained programme of structural audits has narrowed what can supply the missing information. Finite cylinders cannot encode infinite realization; several apparently independent coordinates are invertible reparameterizations of the exact realizer, and the Archimedean and 2-adic readings of the natural divisibility object are complementary factors of one identity; the *known divergent-orbit* valuation-word shadows cannot characterize positive realization, by an explicit uncountable family — though word-level criteria in general are not excluded, and one such criterion is exact (Rem. 9.6); launch renormalization validly reaches actual positive-orbit dynamics, but its original corridor-uniform closure is essentially an $O(\log m)$-scale pointwise lifetime conjecture, bounding the injective initial segment rather than the time to reach 1; and deep negative drift is self-financing with respect to the recovery valuations it requires.

The surviving frontier is therefore actual positive-orbit arithmetic, at three nested levels: qualitative exclusion of an infinite zero-corridor tail; quantitative control of single-window lifetime, at any scale; and, beyond that, recurrence control sufficient for total occupation. Neither EOC nor the Collatz conjecture is proved here.

A A complete proof of the carry budget, with an explicit constant

Lemma 4.3 carries only a proof sketch in Revision 5, and Theorem 4.4 rests on it, which in turn feeds Corollary 4.5 and Theorem 3.6(i). Since that is a load-bearing chain, we give the argument in full. Nothing here is new mathematics; it is the sketch completed, with the constant made explicit so a reader can check it.

Proposition A.1. *Let m_0 be a positive odd integer whose accelerated orbit is injective. Then for every $N \geq 2$,*

$$E_N \;\leq\; \tfrac{1}{9}\log_2(3N-5) + \frac{2}{3\ln 2} \;\leq\; \tfrac{1}{9}\log_2 N + 1.138.$$

The restriction $N \geq 2$ is needed for the first expression to be defined: $3N-5$ is 1 at $N=2$ and negative at $N=1$. The two remaining cases are immediate and require no hypothesis: $E_0 = 0$ (empty sum), and $E_1 = \log_2\bigl(1+\frac{1}{3m_0}\bigr) \leq \log_2 \frac{4}{3} < 0.416$.

Proof. Step 1: the iterates avoid 3. For $k \geq 1$, $2^{d_{k-1}} m_k = 3m_{k-1}+1 \equiv 1 \pmod 3$, and 2 is a unit mod 3, so $m_k \not\equiv 0 \pmod 3$. Each m_k is odd, so $\gcd(m_k, 6) = 1$ for every $k \geq 1$. (The seed m_0 is excluded: it may be divisible by 3.)

Step 2: linearise the logarithm. From $\ln(1+x) \leq x$ for $x \geq 0$,

$$E_N = \sum_{k<N} \log_2\Bigl(1+\frac{1}{3m_k}\Bigr) \;\leq\; \frac{1}{3\ln 2}\sum_{k<N}\frac{1}{m_k}.$$

Step 3: injectivity turns the sum into a reciprocal sum over distinct admissible values. Let $a_1 < a_2 < \cdots$ enumerate the positive integers coprime to 6, and put $A_M = \sum_{j\leq M} 1/a_j$. By Step 1 and injectivity, $m_1, \ldots, m_{N-1}$ are $N-1$ *distinct* positive integers coprime to 6, so $\sum_{1\leq k<N} 1/m_k \leq A_{N-1}$, whence $\sum_{k<N} 1/m_k \leq 1 + A_{N-1}$ using $m_0 \geq 1$.

Step 4: $a_j \geq 3j-2$. Equivalently $\#\{n \leq 3j-3 : \gcd(n,6) = 1\} \leq j-1$. Each block of six consecutive integers contains exactly two integers coprime to 6 (residues 1 and 5). Write $3j-3 = 6q+s$ with $0 \leq s < 6$. Since $3j \equiv 0$ or $3 \pmod 6$, we have $s \in \{0,3\}$ — the cases $s \in \{1,2,4,5\}$ do not occur. If $s=0$ the count is $2q = (3j-3)/3 = j-1$; if $s=3$ the count is $2q+1 = (3j-3)/3 = j-1$. In both cases the bound holds.

Step 5: sum. Since $t \mapsto 1/(3t-2)$ is decreasing on $t > 2/3$,

$$A_M \;=\; 1 + \sum_{j=2}^{M}\frac{1}{a_j} \;\leq\; 1 + \sum_{j=2}^{M}\frac{1}{3j-2} \;\leq\; 1 + \int_1^M \frac{dt}{3t-2} \;=\; 1 + \tfrac{1}{3}\ln(3M-2).$$

Step 6: combine. With $M = N-1$, Steps 2, 3 and 5 give

$$E_N \;\leq\; \frac{1}{3\ln 2}\Bigl(2 + \tfrac{1}{3}\ln(3N-5)\Bigr) \;=\; \tfrac{1}{9}\log_2(3N-5) + \frac{2}{3\ln 2},$$

and $\frac{1}{9}\log_2(3N-5) \leq \frac{1}{9}\log_2 N + \frac{1}{9}\log_2 3$ with $\frac{1}{9}\log_2 3 + \frac{2}{3\ln 2} = 0.1761\ldots + 0.9618\ldots < 1.138$. □

Remark A.2 (Sharpness, as a computation only). The constant is not optimal, and Step 5's integral comparison is where most of the slack is: it loses $\frac{1}{3}\ln 3 = 0.366\ldots$. The exponent $\frac{1}{9}$ — the only part Theorem 4.4 uses — is unaffected by any of this.

Computed values of $A_M - \frac{1}{3}\ln M$: it equals 1 at $M=1$, drops to 0.9690 at $M=2$, and from there *increases*, approaching $0.9728\ldots$; so the sequence is not monotone from $M=1$, and its maximum over the computed range is the first term. (An earlier draft said it "increases to

0.9728, with maximum 1 at $M = 1$", which is self-contradictory as stated.) For $M \le 3 \cdot 10^5$ we find $A_M \le \frac{1}{3} \ln M + 1$ without exception, which would give an additive constant below 0.97 — but this is a finite check, not a proof, and no such inequality is used above. Over odd $m_0 < 60001$ the largest observed value of $E_N - (\frac{1}{9} \log_2 N + 1.138)$ is -1.004, so the proved bound is comfortable but not tight. [Proposition A.1 is proved; every numerical statement in this remark is a finite computation over the stated range and is used nowhere]

Acknowledgments

We thank Franklin Loeb Jr. for an independent reimplementation of the least-realizer and occupation computations [33], which reproduced the closed form of Proposition 5.2, identified the value $\mathcal{O}_1 = 98$ in Remark 6.2 (corrected in Revision 8), and supplied its episode decomposition.

AI assistance

Large language models — Claude (Anthropic) and ChatGPT (OpenAI), the latter including its Codex coding models — assisted with source auditing, derivation, the exact checkers, the Lean formalization drafts and drafting, and were used to cross-check one another, under the author's direction. Conceptual direction and all mathematical claims are the author's, who is responsible for any errors. AI output is not proof evidence: the content rests on the proofs, on the reproducible checks of §15, and on the machine-checked statements listed in §14, which were built and axiom-audited (`propext`, `Classical.choice`, `Quot.sound`) in the author's own sessions. No statement in this paper labelled [proved] depends on an unverified AI-generated step. The Revision 9 literature items (§1.2) were each verified at source as recorded there; where a source could not be retrieved, the entry says so. No external peer review of this paper has occurred.

References

[1] S. Eliahou, *The $3x + 1$ problem: new lower bounds on nontrivial cycle lengths*, Discrete Mathematics **118** (1993), 45–56.

[2] O. Rozier, *The $3x + 1$ problem: a lower bound hypothesis*, Functiones et Approximatio Commentarii Mathematici **56** (2017), no. 1, 7–23.

[3] R. Terras, *A stopping time problem on the positive integers*, Acta Arithmetica **30** (1976), 241–252.

[4] R. Terras, *On the existence of a density*, Acta Arithmetica **35** (1979), 101–102.

[5] O. Rozier and C. Terracol, *Paradoxical behavior in Collatz sequences*, Discrete Mathematics **349** (2026), 115167. doi:10.1016/j.disc.2026.115167 (arXiv:2502.00948v5)

[6] A. Dubickas, *On integer sequences generated by linear maps*, Glasgow Mathematical Journal **51** (2009), no. 2, 243–252. doi:10.1017/S0017089508004655

[7] K. G. Monks and J. Yazinski, *The autoconjugacy of the $3x + 1$ function*, Discrete Mathematics **275** (2004), 219–236.

[8] J. López and P. Stoll, *The $3x + 1$ Periodicity Conjecture in $\mathbb{R}$*, arXiv:2101.12747 (2021). Unrefereed preprint.

[9] J.-L. Verger-Gaugry, *Divergent trajectories in the* $3x+1$ *problem and an* S*-units equation*, Uniform Distribution Theory **21** (2026), no. 1, 1–9. doi:10.66833/UDT-2026-0001

[10] Esm, *A Logarithmic Endpoint Refinement of Collatz Orbit Sparsity: Lean 4 Formalization*, Zenodo, 12 October 2026. doi:10.5281/zenodo.23307433 Lean repository `github.com/Esythium/collatz-endpoint-lean`. Unrefereed; not rebuilt here.

[11] E. De Jesus, `eoc-divergence` (Lean 4), commit `e56950f73121054f12f6e0af9984f63b87372473`, `github.com/innerlightr-wq/eoc-divergence`.

[12] C. J. Everett, *Iteration of the number-theoretic function* $f(2n) = n$, $f(2n+1) = 3n+2$, Advances in Mathematics **25** (1977), 42–45.

[13] M. V. P. Garcia and F. A. Tal, *A note on the generalized* $3n+1$ *problem*, Acta Arithmetica **90** (1999), no. 3, 245–250.

[14] M. J. Curry, *An Explicit Windowed Sparsity Bound for Divergent* $3x+1$ *Orbits, with Logarithmic-Floor Exclusion beyond the Harmonic Barrier*, Zenodo, 24 August 2026. doi:10.5281/zenodo.22087163

[15] J. C. Lagarias and A. Weiss, *The* $3x+1$ *problem: two stochastic models*, Annals of Applied Probability **2** (1992), 229–261.

[16] A. V. Kontorovich and J. C. Lagarias, *Stochastic models for the* $3x+1$ *and* $5x+1$ *problems and related problems*, in: The Ultimate Challenge: The $3x+1$ Problem, Amer. Math. Soc., 2010, 131–188.

[17] T. Tao, *Almost all orbits of the Collatz map attain almost bounded values*, Forum of Mathematics, Pi **10** (2022), e12.

[18] R. P. Steiner, *A theorem on the Syracuse problem*, Proc. Seventh Manitoba Conf. Numer. Math. and Computing, 1977, 553–559.

[19] J. Simons and B. de Weger, *Theoretical and computational bounds for* m*-cycles of the* $3n+1$ *problem*, Acta Arithmetica **117** (2005), 51–70.

[20] C. Hercher, *There are no Collatz* m*-cycles with* $m \leq 91$, Journal of Integer Sequences **26** (2023), Article 23.3.5.

[21] D. Bařina, *Convergence verification of the Collatz problem*, The Journal of Supercomputing **77** (2021), 2681–2688. doi:10.1007/s11227-020-03368-x

[22] D. Bařina, *Improved verification limit for the convergence of the Collatz conjecture*, The Journal of Supercomputing **81** (2025), article 810. doi:10.1007/s11227-025-07337-0

[23] M. Lothaire, *Algebraic Combinatorics on Words*, Cambridge University Press, 2002. (Sturmian words: Chapter 2, by J. Berstel and P. Séébold.)

[24] E. De Jesús, *Periodic and Eventually Periodic Realizer Floors in the Accelerated Collatz Carry Equation: Fixed* 2*-Adic Anchors and Binary Periodicity*, Zenodo, 2026. doi:10.5281/zenodo.21889748

[25] E. De Jesús, *The Least-Realizer Minimum for Confined Words: An Erratum to the Terras–Everett Residue Class in* A Global Occupation Conjecture for the Accelerated $3x+1$ Map*, and the Record-Chronology Reformulation of the Single-Window Realizer Floor*, Zenodo, 2026. doi:10.5281/zenodo.21889704

[26] E. De Jesús, *The Entropy-Deficit Identity in the Accelerated Collatz Word Model: First-Moment Closure, Occupation-Tail Rates, the Geometric Codimension, and the Constant* $0.0793186127\ldots$, Zenodo, 2026. doi:10.5281/zenodo.21889866

[27] E. De Jesús, *Conditional Obstructions and Launch Renormalization in Accelerated Collatz Dynamics*, Zenodo, 2026. doi:10.5281/zenodo.21935927

[28] E. De Jesús, *Transport Deficits and* 2*-Adic Survival Budgets in Accelerated Collatz Words: Equality Cylinders, Past–Future Matching, and First-Failure Structure*, Zenodo, 2026. doi:10.5281/zenodo.21905991

[29] D. J. Bernstein, *A non-iterative* 2*-adic statement of the* $3N+1$ *conjecture*, Proc. Amer. Math. Soc. **121** (1994), 405–408.

[30] D. J. Bernstein and J. C. Lagarias, *The* $3x+1$ *conjugacy map*, Canadian Journal of Mathematics **48** (1996), 1154–1169.

[31] E. Y. Chang, *A Structural Reduction of the Collatz Conjecture to One-Bit Orbit Mixing*, preprint.

[32] E. Y. Chang, *Exploring Collatz Dynamics with Human–LLM Collaboration*, preprint.

[33] F. Loeb, *Collatz Box Universes Explorer Suite: EOC realizer comparison* (version v1.0-eoc-realizers) [computer software], Zenodo, 2026. doi:10.5281/zenodo.23144148

A Band rigidity

We retain the following structural observation from Revision 5, which constrains how a drift path may sit inside a unit band.

Theorem A.1 (Band rigidity, Revision 5 Thm. A.1)**.** *Let* $\theta = \alpha - 1$*. If a word's drift path satisfies* $\beta \le R_j < \beta + 1$ *for all* $j \le \ell$ *and some* $\beta \in \mathbb{R}$*, then every letter lies in* $\{1, 2\}$ *and, writing* $b_i = d_i - 1$*, the word is a length-*ℓ *factor of the Sturmian system of slope* θ*; conversely every such factor arises, and there are exactly* $\ell + 1$ *admissible words for each* ℓ*.* *[Revision 5, Thm. A.1, restated verbatim. Revision 5 gives a proof* sketch*, with the factor characterisation and the count* $\ell+1$ *cited to Morse–Hedlund (see [23])]*

Remark A.2 (Why the substitution $b_i = d_i - 1$ and the slope θ matter). The reduction is exact: with $d_i \in \{1, 2\}$ and $b_i = d_i - 1 \in \{0, 1\}$,

$$R_j = \textstyle\sum_{i<j} d_i - j\alpha = \sum_{i<j}(b_i + 1) - j\alpha = \sum_{i<j} b_i - j(\alpha - 1) = \sum_{i<j} b_i - j\theta,$$

so the band condition is the Beatty condition *at slope* θ, on the binary word b. Revision 7 initially paraphrased this as "the placement follows the Beatty gap sequence of α", dropping the substitution and naming the wrong slope, and dropped both the converse and the exact count while labelling the result "retained". It was a paraphrase, not a retention; the statement above is Revision 5's. [Revision 7, correction of attribution]

The relevance to this revision is limited but worth recording: band rigidity says that a *narrow* corridor forces exactly Sturmian structure, which is exactly the regime where Observation 5.9 shows height amortization fails. Narrow corridors are therefore not automatically easier.

B Claim-status conventions

Every substantive statement carries a bracketed status: [Revision 5] retained unchanged from the previous revision; [proved] proved in this paper; [companion work] proved in a cited companion note; [external, cited] a cited theorem of others, used but not reproved; [formalized] machine-checked, see §14; [computational diagnostic] supported by computation only, never promoted to a theorem; [structural interpretation] an organizing reading of proved facts; [reclassification] a previously known statement placed in a new role; [correction to Revision 5] a superseded claim, corrected here.