

A Structural Proof of the Collatz Conjecture via Divergent-Orbit and Cycle Exclusion

Jungwon Park

March 2026

Abstract

We present a computer-assisted proof that every positive integer orbit under the Collatz map $T(n) = n/2$ (if n is even), $(3n + 1)/2$ (if n is odd) eventually reaches 1. The argument consists of two independent parts, H1 and H2, each of which reduces an originally infinite verification problem to a finite exact certification step. The mathematical reductions are traditional proofs; the certifications are finite, deterministic computations using only exact integer and rational arithmetic, whose complete specifications (input, algorithm, output) are given in the appendices.

For H1 (divergent-orbit exclusion), a deterministic normalized block-drift analysis establishes an exact affine envelope that rules out all parameter regimes below an explicit threshold $V_{\text{env}} \approx 1.5990$. The remaining oscillatory regime is reduced to a finite certification problem over a quotient-state automaton whose mathematical soundness is guaranteed by a chain of five theorems (Theorems 12.8–12.14). The certification is discharged by exhaustive exact computation (Theorem 12.17).

For H2 (non-trivial cycle exclusion), a gate-based reformulation converts the cycle problem into an explicit inequality system. Discrete convexity arguments and a two-value support reduction compress the problem to a single-parameter bound, which is then verified by exact rational arithmetic over a rigorously delimited feasible range.

The computational components are finite, exact, and logically isolated. Their role is analogous to the finite case verifications in other computer-assisted proofs such as the Four Colour Theorem and the Kepler Conjecture: the mathematical argument provides a complete reduction to finite exact certification problems, and the computations discharge those problems. The complete source code for both certifications, together with their full execution logs and exact rational outputs, is included in the appendices (Appendices E–G).

Keywords: Collatz conjecture; $3x + 1$ problem; divergent orbits; non-trivial cycles; transfer operator; spectral gap; Baker’s theorem; gate framework; computer-assisted proof; exact rational verification.

MSC 2020: 11B37; 37A05; 37A30; 11J86.

Contents

1	Introduction	3
2	Basic Definitions	4
3	Constraints on Non-Trivial Cycles	4
4	Background: Conditional Convergence and Measure Theory	6
5	Spectral Theory of the Syracuse Transfer Operator	6
5.1	Transfer Operator: Definition and Basic Properties	6
5.2	Spectral Gap in Lipschitz Spaces	7
5.3	Decay of Correlations and Forward Mixing	7
6	Baker's Theorem: Role and Limitations	7
7	The Measure-Pointwise Barrier and Case Classification	8
8	Block-Level Determinism: Foundational Tools	8
8.1	Ladder Descent Identity	8
8.2	Block Drift Identity	8
9	Block Transition Identity and Indexing	9
10	The N-BD Identity and Standard Form	9
10.1	Complete Indexing Isomorphism	9
10.2	Anchor Transition Formula	10
10.3	Normalized N-BD Identity	10
10.4	N-BD Standard Form	10
11	Divergence Exclusion: Exact Envelope Theory	11
11.1	A-Anchor Exact Contraction ($V_n = 1$ blocks)	11
11.2	Block Drift Upper Bound for $V_n \geq 2$ Blocks	13
11.3	Unconditional Divergence Exclusion: Core Theorems	13
12	Exclusion of the High-V Oscillatory Regime	16
12.1	State Space and Actual Pair Drift	16
12.2	Normal Form Decomposition	16
12.3	Quotient-Witness Certification	18
12.4	Soundness of the Quotient Abstraction	19
12.5	Local Sink Certification	23
12.6	Bellman Closure and the Exclusion of Divergent Orbits	24
13	Gate Framework for Cyclic Orbits	27
13.1	Block Decomposition and A/E Separation	27
13.2	Gate-to-Gate Recurrence	28
13.3	Gate-BD Cost Identity and Cycle Equation	28
13.4	TL_V Lemma	28
13.5	The BL0-Req' Sufficient Condition	29

13.6	Residue Transport Structure and Gate5 Dichotomy	30
13.7	Properties of the Univariate Upper Bound $u(m)$	31
14	Cycle Exclusion: Standard Form and Bridge Inequality	32
15	Two-Value Support Reduction	33
16	Global Bounds and Numerical Closure	34
16.1	Global Lower Bound on Gate Values	34
16.2	Global u -Upper Bound via Divisibility (H2-J)	35
16.3	Hybrid Master Inequality (H2-K)	36
16.4	Low-Part Constrained Extremization and Numerical Closure (H2-L) . .	37
17	Exclusion of the G5-A Path	39
17.1	Structure of the G5-A Path	39
17.2	Exclusion of G5-A	40
18	Main Theorem	41
19	Discussion	42
19.1	The Golden Ratio as a Structural Threshold	42
19.2	Discrete Convexity and Majorization	42
19.3	The Role of Computation	43
19.4	Relation to Baker–Wüstholz	43
	Conclusion	43
	References	43
A	Proof Architecture and Computational Transparency	44
B	Certification Architecture for H1	45
C	Auxiliary Computational Evidence	46
D	Exact Rational Verification Protocol for H2	47
E	H1 Certification: Exact Arithmetic Verification	48
F	H2 Exact Rational Witness	49
G	H1 Certification Code: Key Algorithms	51
H	Reproducibility and Exact-Artifact Manifest	52

1 Introduction

The Collatz conjecture asserts that every positive integer eventually reaches the trivial cycle $1 \mapsto 4 \mapsto 2 \mapsto 1$ under iteration of the map

$$T(n) = \begin{cases} n/2, & n \text{ even,} \\ (3n + 1)/2, & n \text{ odd.} \end{cases}$$

Despite its elementary formulation, the problem has resisted proof for decades and remains one of the most prominent open questions in elementary number theory and discrete dynamics [10].

Numerical verification has been pushed to large ranges: Barina [4, 5] confirmed convergence for all $n < 2^{71}$. Eliahou [6] proved that any non-trivial cycle must have length at least 17,087,915, drawing on the theory of linear forms in logarithms initiated by Baker [1] and refined by Baker and Wüstholz [2]. Tao [11] showed that almost all Collatz orbits attain almost bounded values in a logarithmic-density sense. These results strongly constrain potential counterexamples, but they do not yield a proof.

This paper develops a two-part approach. The first part (H1) excludes divergent orbits by means of a deterministic normalized block-drift analysis. A normalized block-drift identity converts the dynamics into an exact affine-envelope framework. An explicit threshold $V_{\text{env}} \approx 1.5990$ is identified such that $\liminf \bar{V}_N < V_{\text{env}}$ forces $w_n \rightarrow 0$, contradicting divergence. The remaining High-V Oscillatory Regime is handled by constructing a finite quotient-state automaton whose mathematical soundness is guaranteed by a chain of five theorems (Theorems 12.8–12.14). A finite exact certification (Theorem 12.17) establishes that every macro-transition strictly decreases a well-founded rank, and a Bellman closure argument (Theorem 12.21) lifts this to a global contradiction.

The second part (H2) excludes non-trivial cycles. A gate-based reformulation converts the cycle problem into a family of explicit inequalities. A sufficient condition (BL0-Req', Theorem 13.7) reduces cycle exclusion to showing $\bar{\eta} < c_*$, where $c_* := \varphi^{-1} - \log_2(3/2) \approx 0.033071$. Discrete convexity and a two-value support reduction (Theorem 15.3) compress the problem to a single-parameter bound, which is then resolved by exact rational arithmetic. Independently, a 2-adic residue analysis (Theorem 17.1) shows that all non-lock steps must lie in Branch B, confirming that the above bound covers every admissible case.

Both computational ingredients—the 222-child certification for H1 and the rational inequality verification for H2—use only exact integer and rational arithmetic. No floating-point approximation is employed. Their role is analogous to the finite case checks in other computer-assisted proofs such as the Four Colour Theorem and the Kepler Conjecture. The logical structure cleanly separates the mathematical reductions from the computational verification layers.

Our main result is the following.

Theorem (Main Theorem). For every positive integer n , the Collatz orbit of n eventually reaches 1. That is, $\forall n \in \mathbb{Z}_{>0}, \exists k \in \mathbb{Z}_{\geq 0} : T^k(n) = 1$.

The proof is given in Section 18.

Organization. Section 2 introduces the basic definitions. Section 3 recalls standard constraints on hypothetical non-trivial cycles. Section 4 records a conditional convergence argument and Tao’s measure-theoretic result; neither is assumed in the proof. the present proof does not assume the associated hypothesis. Section 5 develops the spectral theory of the Syracuse transfer operator, which provides structural motivation and auxiliary measure-theoretic support. Sections 6–7 discuss Baker’s theorem and the case classification for divergent orbits. Sections 8–12 carry out the divergent-orbit exclusion (H1). Sections 13–17 develop and carry out the cycle-exclusion argument (H2). Section 18 combines H1 and H2 to prove the main theorem. Section 19 discusses the structural significance of the argument. The appendices record the computational transparency statements and auxiliary verification details.

2 Basic Definitions

Definition 2.1 (Collatz map, Syracuse function, and Büchi sequence). **Collatz function:** $T : \mathbb{N} \rightarrow \mathbb{N}$, $T(n) = n/2$ (n even), $(3n + 1)/2$ (n odd).

Syracuse function: $S : \mathbb{Z}_2^\times \rightarrow \mathbb{Z}_2^\times$, $S(n) = (3n + 1)/2^{v_2(3n+1)}$, where $\mathbb{Z}_2^\times = \{x \in \mathbb{Z}_2 : v_2(x) = 0\}$ is the group of 2-adic units.

Büchi sequence: $a_k(n_0) = v_2(3n_{k-1} + 1)$, $n_k = S(n_{k-1})$. **Büchi mean:** $\bar{a}_K = \frac{1}{K} \sum_{k=1}^K a_k$.

Lemma 2.2 (Basic properties of the Büchi sequence). (i) Under Haar measure μ : $\Pr[a_k = j] = 2^{-j}$ for $j \geq 1$, and $\mathbb{E}[a] = 2$.

(ii) $\log_2 S^K(n_0) = \log_2 n_0 + K \log_2 3 - L_K + \epsilon_K$, $|\epsilon_K| \leq 1$, $L_K = \sum_{k=1}^K a_k$.

(iii) If the orbit is unbounded then $\liminf_{K \rightarrow \infty} \bar{a}_K(n_0) \leq \log_2 3$. The boundary case $\liminf \bar{a}_K(n_0) = \log_2 3$ is deterministically excluded in Section 11 (Theorem 11.8).

3 Constraints on Non-Trivial Cycles

Master Notation Table

For the reader’s convenience, we collect the principal symbols used throughout the paper and their interrelationships.

Symbol	Definition and role
x_k	Full Collatz orbit: $x_{k+1} = T(x_k)$, $T(n) = n/2$ (even), $(3n+1)/2$ (odd)
$S(n)$	Syracuse function: $S(n) = (3n+1)/2^{v_2(3n+1)}$ (odd \rightarrow odd map)
a_k	Büchi sequence: $a_k = v_2(3x_{k-1} + 1)$; $\bar{a}_K = K^{-1} \sum a_k$
y_n	A-anchor subsequence: $y_n \equiv 1 \pmod{4}$, i.e. $v_2(y_n + 1) = 1$
R_n	Deep depth: $R_n := v_2(3y_n + 1) \geq 2$ at anchor n
m_n	Shallow count: number of $r_k = 1$ steps following anchor n ; $V_n = m_n + 1$
V_n	Block length: $V_n := k_{n+1} - k_n = m_n + 1$
w_n	Shifted anchor: $w_n := y_n + 1$; the transition formula is stated in terms of w_n
K, N	$K = \sum_{n < N} V_n$ (total Syracuse steps), N (total blocks)
\bar{V}_N	Mean block length: $\bar{V}_N := K/N$
\bar{R}_N	Mean deep depth: $\bar{R}_N := N^{-1} \sum R_n$
ε_n	Error term: $\varepsilon_n := \log_2(1 + (2^{R_n} - 2)/(3w_n)) > 0$
q_t	Gate half-value (H2): $q_t := (y_t + 1)/2$, odd
η_t	Gate cost (H2): $\eta_t := \log_2(1 + (2^{A_t-1} - 1)/(3q_t)) > 0$
$u(m)$	Univariate bound: $u(m) := \log_2(1 + 1/(2^{m+1} - 1))$
c_*	Critical threshold: $c_* := \varphi^{-1} - \log_2(3/2) \approx 0.033071$

Indexing relationships. The full orbit (x_k) is indexed by k . Deep events (those k with $a_k = v_2(3x_{k-1} + 1) \geq 2$) are re-indexed by n , defining the anchor subsequence (y_n) and the block parameters (V_n, R_n, m_n) . Each block consists of one deep event followed by $m_n = V_n - 1$ shallow events. The total number of Syracuse steps over N blocks is $K = \sum_{n < N} V_n$, and the key observables are the mean block length $\bar{V}_N = K/N$ and mean deep depth \bar{R}_N . For cycle analysis (H2), the gates are indexed by t with corresponding parameters (q_t, A_t, m_t, η_t) .

Theorem 3.1 (Lower bounds for cycle exclusion). *(i) If a cycle of length m exists, then $|2^L - 3^m| > 0$.*

(ii) Eliahou [6]: For any non-trivial cycle, $m \geq 17,087,915$.

(iii) Barina [5]: Numerical verification up to $n < 2^{71}$.

(iv) Baker–Wüstholz [2]: $|2^L - 3^m| > \exp(-C \log L \cdot \log m)$. This lower bound supports (ii).

4 Background: Conditional Convergence and Measure Theory

This section records a conditional convergence argument and Tao's measure-theoretic result. Neither is assumed in the proof; both are included for contextual background only.

Hypothesis 4.1 (Decorrelation hypothesis). For every positive odd integer $n_0 \in \mathbb{N}$: $\liminf_{K \rightarrow \infty} \bar{a}_K(n_0) > \log_2 3$.

Convergence under Hypothesis 4.1. By Hypothesis 4.1, $\exists \delta > 0$ and K_0 such that $\bar{a}_K \geq \log_2 3 + \delta$ for all $K \geq K_0$. Lemma 2.2(ii) gives $\log_2 S^K(n_0) \leq \log_2 n_0 - K\delta + 1 \rightarrow -\infty$. The orbit falls below 2^{71} in finite time, and by Theorem 3.1(iii) it reaches 1. \square

Remark 4.2. This paper does *not* assume Hypothesis 4.1. Sections 11–12 directly prove that divergent orbits are impossible, and Sections 13–17 exclude non-trivial cycles, both unconditionally.

Theorem 4.3 (Tao [11]). For any $f : \mathbb{N} \rightarrow \mathbb{R}_{>0}$ with $f(n) \rightarrow \infty$: $\bar{d}_{\log} \{n : \min_{k \geq 0} T^k(n) \leq f(n)\} = 1$.

Remark 4.4. Tao's result is measure-theoretic and concerns logarithmic density. It provides supporting evidence but is not a step in the present proof.

5 Spectral Theory of the Syracuse Transfer Operator

The spectral theory developed in this section concerns the average structure of the Syracuse map with respect to Haar measure μ on the 2-adic units. **This entire section (Sections 5.1–5.3) is not used in any step of the formal proof of the main theorem.** It provides structural motivation and measure-theoretic context only. The reader may skip to Section 6 without loss of logical continuity. Since $\mu(\mathbb{N}) = 0$, spectral results do not directly yield pointwise conclusions for natural-number orbits. The core proof engine for H1 is the deterministic block analysis of Sections 8–12; the spectral theory supplies structural motivation and auxiliary measure-theoretic support.

5.1 Transfer Operator: Definition and Basic Properties

Definition 5.1 (State space). $X = \mathbb{Z}_2^\times$: compact ultrametric space with 2-adic metric $|x - y|_2$. μ : normalized Haar measure.

Lemma 5.2 (Inverse branch maps). $\phi_j(x) = (2^j x - 1)/3$ for $j \geq 1$: continuous maps satisfying $S(\phi_j(x)) = x$ and $v_2(3\phi_j(x) + 1) = j$. The family $\{\phi_j(X)\}_{j \geq 1}$ is a partition of X .

Definition 5.3 (Syracuse transfer operator).

$$(\mathcal{L}f)(x) = \sum_{j=1}^{\infty} 2^{-j} f(\phi_j(x)).$$

The weight 2^{-j} comes from the Jacobian $|S'(\phi_j)|_2^{-1}$. $\mathcal{L}\mathbf{1} = \mathbf{1}$ (Markov property).

Lemma 5.4 (Haar measure invariance). $\int_X \mathcal{L}f d\mu = \int_X f d\mu$.

Proof. Using the partition $\{\phi_j(X)\}_{j \geq 1}$: $\int \mathcal{L}f d\mu = \sum_j 2^{-j} \int f(\phi_j) d\mu = \sum_j \int_{\phi_j(X)} f d\mu = \int f d\mu$. \square

5.2 Spectral Gap in Lipschitz Spaces

Definition 5.5. $[f]_{\text{Lip}} = \sup_{x \neq y} |f(x) - f(y)|/|x - y|_2$. $\text{Lip}_0(X) = \{f \in \text{Lip}(X) : \int f d\mu = 0\}$.

Lemma 5.6 (Contraction of inverse branches). $|\phi_j(x) - \phi_j(y)|_2 = 2^{-j}|x - y|_2$.

Theorem 5.7 (Lasota–Yorke inequality).

$$[\mathcal{L}f]_{\text{Lip}} \leq \frac{1}{3}[f]_{\text{Lip}}, \quad \|\mathcal{L}f\|_{\infty} \leq \|f\|_{\infty}.$$

Proof. $|(\mathcal{L}f)(x) - (\mathcal{L}f)(y)| \leq \sum_j 2^{-j}[f]_{\text{Lip}} \cdot 2^{-j}|x - y|_2 = [f]_{\text{Lip}} \cdot \frac{1}{3}|x - y|_2$, since $\sum_{j=1}^{\infty} 4^{-j} = 1/3$. The L^{∞} bound follows similarly. \square

Corollary 5.8 (Exclusion of $|\lambda| = 1$ eigenvalues). *If $\mathcal{L}f = \lambda f$ with $|\lambda| = 1$, then $[f]_{\text{Lip}} = 0$, hence f is constant.*

Theorem 5.9 (Quasi-compactness; Hennion [7]). $\mathcal{L} : \text{Lip}(X) \rightarrow \text{Lip}(X)$ is quasi-compact with $r_{\text{ess}}(\mathcal{L}) \leq 1/3$. Consequently, $\exists \rho_0 \in (1/3, 1)$ such that $\|\mathcal{L}^n f\|_{\text{Lip}} \leq C\rho_0^n \|f\|_{\text{Lip}}$ for all $f \in \text{Lip}_0$.

5.3 Decay of Correlations and Forward Mixing

Theorem 5.10 (Convergence of weighted preimage Buchi means). *For all $n_0 \in \mathbb{N}$: $|\sum_{S^m(y)=n_0} w(y, m) \cdot \bar{a}_m(y) - 2| \rightarrow 0$, where $w(y, m) = \prod_k 2^{-a_k(y)}$.*

Remark 5.11 (Connection to the block analysis). The spectral gap (Theorem 5.9) ensures exponential convergence to the leading eigenvalue 1 in Lip_0 . This underpins the block contraction of Section 11: consecutive $V_n = 1$ blocks induce geometric decay of w_n . Theorem 5.10 gives measure-theoretic support for the stability of $\bar{R}_N \approx 2$ in the N-BD standard form. However, spectral theory cannot exclude individual elements of a measure-zero set (Observation 7.1); the deterministic block analysis of Sections 8–12 is therefore indispensable.

6 Baker’s Theorem: Role and Limitations

Theorem 6.1 (Baker [1] — lower bound for linear forms in logarithms). *Let $\alpha_1, \alpha_2 > 0$ be algebraic numbers, neither equal to 1, with $\ln \alpha_1 / \ln \alpha_2 \notin \mathbb{Q}$. Then for all positive integers b_1, b_2 :*

$$|b_1 \ln \alpha_1 - b_2 \ln \alpha_2| > \frac{C}{(\max(b_1, b_2))^{\kappa}},$$

where $C > 0$ and $\kappa > 0$ are effectively computable constants.

Proposition 6.2 (Baker’s theorem and cycle constraints). *Taking $\alpha_1 = 3$, $\alpha_2 = 2$ in Theorem 6.1: for all positive integers K, L , $|K \ln 3 - L \ln 2| > C/(\max(K, L))^{\kappa} > 0$. If a cycle of length m exists, then $|2^L - 3^m| > 0$, and the polynomial lower bound constrains cycle initial values. This supports Eliahou’s result $m \geq 17,087,915$.*

Remark 6.3 (Limitations of Baker’s theorem). Baker-type bounds give polynomial lower bounds on $|K \ln 3 - L \ln 2|$, but the right-hand side tends to 0 as $K, L \rightarrow \infty$. Therefore Baker’s theorem alone cannot exclude the critical drift case. The exact block envelope of Section 11 closes all of $\liminf \bar{V}_N \leq \log_2 3$, while the lower-regime theorem closes the sub-interval up to $V_{\text{env}} \approx 1.5990$. The High-V Oscillatory Regime is handled separately in Section 12.

7 The Measure-Pointwise Barrier and Case Classification

Observation 7.1 (Measure-pointwise barrier). The proportion of “bad orbits” ($\bar{a}_K \leq \log_2 3$ permanently) decays as $\sim e^{-0.055N}$, but the count of N -bit odd integers grows as $\sim e^{0.693N}$. The potential exception count $\sim e^{0.638N} \rightarrow \infty$; hence spectral theory alone cannot exclude all exceptions.

Case Classification for Divergent Orbits

Observation 7.2 (Case classification).

Case	Exclusion basis
Case A: $\liminf \bar{V}_N = \log_2 3$	Thm. 11.8 (exact envelope)
Case B (residual): $\liminf \bar{V}_N \in [T^*, \log_2 3)$	Thm. 11.8
Case B (lower): $\liminf \bar{V}_N \in (1, T^*)$	Thm. 11.7
Case C: $\liminf \bar{V}_N = 1$	Thm. 11.7, Lem. 11.1
Case D1 (lower): $\liminf \bar{V}_N \in (\log_2 3, V_{\text{env}})$	Thm. 11.7
High-V Oscillatory Regime	§12: Thms. 12.8–12.17, Bellman closure
Non-trivial cycles	Gate framework + H2-J/K/L (§§13–17)

8 Block-Level Determinism: Foundational Tools

8.1 Ladder Descent Identity

Theorem 8.1 (Ladder Descent Identity, Part 1). Define $V_k := v_2(x_k + 1)$. A shallow event ($r_k = 1$) decreases V by exactly 1:

$$\boxed{r_k = 1 \implies V_{k+1} = V_k - 1.}$$

Proof. $r_k = 1$ means $v_2(3x_k + 1) = 1$, i.e., $x_k \equiv 3 \pmod{4}$. Then $x_{k+1} = (3x_k + 1)/2$ (the Syracuse step removes $v_2(3x_k + 1) = 1$ factor of 2). Now $x_{k+1} + 1 = (3x_k + 1)/2 + 1 = (3x_k + 3)/2 = \frac{3}{2}(x_k + 1)$. Since $x_k \equiv 3 \pmod{4}$, we have $x_k + 1 \equiv 0 \pmod{4}$, so $V_k = v_2(x_k + 1) \geq 2$. Then $V_{k+1} = v_2(x_{k+1} + 1) = v_2(\frac{3}{2}(x_k + 1)) = v_2(3) + v_2(x_k + 1) - v_2(2) = 0 + V_k - 1 = V_k - 1$. \square

Theorem 8.2 (Ladder Descent Identity, Part 2). Let the deep event ($r_k \geq 2$) indices be $k_0 < k_1 < \dots < k_{M-1}$. Then:

$$\boxed{k_{i+1} - k_i = V_{k_{i+1}} = v_2(x_{k_{i+1}} + 1).}$$

8.2 Block Drift Identity

Define the anchor subsequence y_n where $x \equiv 1 \pmod{4}$; let $R_n := v_2(3y_n + 1) \geq 2$, m_n the number of shallow events following, and $w_n := y_n + 1$. (These quantities are

formally defined in Section 10, Theorem 10.1; they are introduced here in preview form for the statement of the Block Drift Identity.)

Theorem 8.3 (Block Drift Identity). *Let $K = \sum_{n=0}^{N-1} (m_n + 1)$ and $S_K = \sum_{n=0}^{N-1} (m_n + R_n)$. Then:*

$$S_K = K \log_2 3 - \log_2 \left(\frac{w_N}{w_0} \right) + \sum_{n=0}^{N-1} \epsilon_n, \quad \epsilon_n := \log_2 \left(1 + \frac{2^{R_n} - 2}{3w_n} \right) > 0.$$

Proof. From the anchor transition formula (Theorem 10.2): $w_{n+1} = \frac{3^{m_n}}{2^{m_n+R_n}}(3w_n + 2^{R_n} - 2)$. Taking \log_2 :

$$\begin{aligned} \log_2 w_{n+1} &= m_n \log_2 3 - (m_n + R_n) + \log_2(3w_n + 2^{R_n} - 2) \\ &= m_n \log_2 3 - (m_n + R_n) + \log_2(3w_n) + \log_2 \left(1 + \frac{2^{R_n} - 2}{3w_n} \right). \end{aligned}$$

Rearranging: $\log_2 w_{n+1} - \log_2 w_n = (m_n + 1) \log_2 3 - (m_n + R_n) + \epsilon_n$, where $\epsilon_n := \log_2(1 + (2^{R_n} - 2)/(3w_n))$. Summing over $n = 0, \dots, N-1$ and using $K = \sum (m_n + 1)$, $S_K = \sum (m_n + R_n)$:

$$\log_2 \frac{w_N}{w_0} = K \log_2 3 - S_K + \sum_{n=0}^{N-1} \epsilon_n.$$

Since $R_n \geq 2$ for A-anchors, $2^{R_n} - 2 \geq 2 > 0$ and $w_n > 0$, so $\epsilon_n > 0$. Rearranging gives the stated identity. \square

9 Block Transition Identity and Indexing

Theorem 9.1 (Algebraic block transition identity). *Let $k_0 < k_1 < \dots$ be the indices of deep events, and $z_n := x_{k_n+1}$ the odd number immediately after. If $V_n := v_2(z_n) + 1 \geq 1$, then: (1) the number of shallow events until the next deep event is exactly $m_n = V_n - 1$; (2) writing $z_n + 1 = 2^{m_n+1} u_n$ (u_n odd), we have $w_{n+1} = 2 \cdot 3^{m_n} u_n$.*

Theorem 9.2 (Block sum formula and index identity). (1) **Block sum formula:**

$$S_K = (K - N) + \sum_{n=0}^{N-1} R_n.$$

(2) **Index identity:** $\bar{V}_N := K/N$.

10 The N-BD Identity and Standard Form

10.1 Complete Indexing Isomorphism

Theorem 10.1 (Complete indexing isomorphism). *For any odd starting point $x_0 = w_0 \geq 3$, define deep events as those k with $r_k \geq 2$. The warm-up index $k_0 := \min\{k \geq 0 : r_k \geq 2\}$ is finite. Writing the subsequent indices as $k_0 < k_1 < k_2 < \dots$:*

$$y_n := x_{k_n}, \quad R_n := r_{k_n}, \quad z_n := x_{k_n+1}, \quad m_n := V_n - 1, \quad V_n := k_{n+1} - k_n.$$

Proof that $k_0 < \infty$. If $x_0 \equiv 1 \pmod{4}$, then $v_2(3x_0 + 1) \geq 2$, so $k_0 = 0$. If $x_0 \equiv 3 \pmod{4}$, then $V_0 := v_2(x_0 + 1) \geq 2$ (since $x_0 + 1 \equiv 0 \pmod{4}$). By Theorem 8.1,

each shallow step ($r_k = 1$, equivalently $x_k \equiv 3 \pmod{4}$) decreases $V_k := v_2(x_k + 1)$ by exactly 1. Starting from $V_0 \geq 2$, after at most $V_0 - 1$ shallow steps, V reaches 1, meaning $x_k \equiv 1 \pmod{4}$, so $r_k \geq 2$ and the first deep event occurs. Hence $k_0 \leq V_0 - 1 < \infty$. \square

10.2 Anchor Transition Formula

Theorem 10.2 (Anchor transition formula). *Standard variables: $w_n := y_n + 1$, $R_n := v_2(3y_n + 1) \geq 2$ ($3y_n + 1 = 2^{R_n}q_n$, q_n odd), $V_n := k_{n+1} - k_n \geq 1$, $m_n := V_n - 1 \geq 0$.*

Anchor transition formula:

$$\boxed{w_{n+1} = \frac{3^{m_n}}{2^{m_n+R_n}}(3w_n + 2^{R_n} - 2), \quad m_n := V_n - 1 \geq 0.} \quad (\text{T})$$

Error term: $\varepsilon_n := \log_2\left(1 + \frac{2^{R_n-2}}{3w_n}\right) > 0$, and $\varepsilon_n < 1$.

Proof of $\varepsilon_n < 1$: $2^{R_n} \leq 3y_n + 1 = 3w_n - 2$, so $2^{R_n} - 2 \leq 3w_n - 4 < 3w_n$, giving $\varepsilon_n < 1$. \square

Sharp upper bounds for the error term.

Lemma 10.3 (Non-trivial A-anchor block upper bound: $\varepsilon_n \leq \log_2(16/9)$, sharp). *For non-trivial A-anchor blocks ($y_n > 1$), $w_n \geq 6$, hence:*

$$\boxed{\varepsilon_n \leq \log_2 \frac{16}{9}.}$$

Equality holds at $y_n = 5$ ($R_n = 4$, $w_n = 6$); the bound is sharp.

Lemma 10.4 ($R_n = 2$ special case: $\varepsilon_n \leq \log_2(4/3)$). *For A-anchor blocks with $R_n = 2$: $\varepsilon_n = \log_2(1 + 2/(3w_n)) \leq \log_2(4/3)$. For $R_n = 2$ and $y_n \geq 9$ ($w_n \geq 10$): $\varepsilon_n \leq \log_2(16/15) \approx 0.093$.*

10.3 Normalized N-BD Identity

Theorem 10.5 (Normalized N-BD identity).

$$\boxed{\frac{\bar{R}_N - 1}{\bar{V}_N} = (\log_2 3 - 1) - \alpha + \bar{\varepsilon},}$$

where $\bar{R}_N := \frac{1}{N} \sum_{n=0}^{N-1} R_n$, $\bar{V}_N := K/N$, $\alpha := \frac{1}{K} \log_2(w_N/w_0)$, $\bar{\varepsilon} := \frac{1}{K} \sum_{n=0}^{N-1} \varepsilon_n \geq 0$.

10.4 N-BD Standard Form

Theorem 10.6 (N-BD standard form).

$$\boxed{\frac{1}{N} \log_2 \frac{w_N}{w_0} = \bar{V}_N(\log_2 3 - 1) + 1 - \bar{R}_N + \frac{1}{N} \sum_{n=0}^{N-1} \varepsilon_n.} \quad (\text{N-BD})$$

The critical line $\bar{V}_N(\log_2 3 - 1) + 1 - \bar{R}_N = 0$ separates divergence (> 0) from convergence (< 0).

11 Divergence Exclusion: Exact Envelope Theory

11.1 A-Anchor Exact Contraction ($V_n = 1$ blocks)

Lemma 11.1 (A-anchor $V_n = 1$ exact contraction). *For an A-anchor block with $V_n = 1$ and $y_n > 1$:*

$$\boxed{\frac{w_{n+1}}{w_n} \leq \frac{7}{9}, \quad \Delta_n := \log_2 \frac{w_{n+1}}{w_n} \leq -\delta, \quad \delta := \log_2 \frac{9}{7} \approx 0.36257.}$$

The maximum $7/9$ is achieved at $y_n = 17$ ($R_n = 2$).

Remark 11.2 (Consecutive contraction blocks). If L consecutive $V_n = 1$ blocks occur, then $w_{n+L} \leq (7/9)^L w_n \rightarrow 0$. Hence any orbit where $V_n = 1$ with density 1 is forced to have $w_n \rightarrow 0$, contradicting divergence.

Exact maximum block growth ratio.

Lemma 11.3 (Exact growth ratio formula). *For fixed $V \geq 1$: $\rho(V) := \sup\{w_{n+1}/w_n : V_n = V\}$.*

$$\boxed{\rho(1) = \frac{7}{9}, \quad \rho(2) = \frac{15}{13}, \quad \rho(3) = \frac{9}{5}.}$$

For $V \geq 2$: $\rho(V) = \frac{3^V}{2^{V+1}-1}$ (V odd, $V \geq 3$); $\rho(V) = \frac{5 \cdot 3^V}{5 \cdot 2^{V+1}-1}$ (V even, $V \geq 2$).

Proof. From the anchor transition formula (Theorem 10.2): $w_{n+1} = \frac{3^{m_n}}{2^{m_n+R_n}}(3w_n + 2^{R_n} - 2)$ with $m_n = V_n - 1$. The ratio $w_{n+1}/w_n = \frac{3^{V-1}}{2^{V-1+R}}(3 + \frac{2^{R-2}}{w_n})$. This is maximized when w_n is minimized and R is minimized ($R = 2$, since $R \geq 2$ for A-anchors). For $R = 2$: $w_{n+1}/w_n = \frac{3^{V-1}}{2^{V+1}}(3 + 2/w_n)$.

Case $V = 1$: $m_n = 0$, so from (T): $w_{n+1} = (3w_n + 2^{R_n} - 2)/2^{R_n}$, giving ratio $(3w_n + 2^{R_n} - 2)/(2^{R_n}w_n)$. The constraint $V_n = 1$ means $v_2(z_n + 1) = 1$ where $z_n = S(y_n)$; this restricts which (y_n, R_n) pairs are admissible. For $R_n = 2$: $z_n = (3y_n + 1)/4$, and $V_n = v_2(z_n + 1) = 1$ requires $z_n \equiv 1 \pmod{4}$. Checking: $y_n = 17$ gives $z_n = 13$, $v_2(14) = 1$, confirming $V_n = 1$. Ratio: $(54 + 2)/72 = 56/72 = 7/9$. $y_n = 9$ gives $z_n = 7$, $v_2(8) = 3 \neq 1$, so $V_n = 3 \neq 1$ (inadmissible). $y_n = 25$ gives $z_n = 19$, $v_2(20) = 2 \neq 1$ (inadmissible). $y_n = 33$ gives $z_n = 25$, $v_2(26) = 1$, confirming $V_n = 1$. Ratio: $(99 + 2)/132 = 101/132 < 7/9$. For $R_n \geq 3$: the ratio $\leq (3w_n + 2^{R_n})/(2^{R_n}w_n) = 3/2^{R_n} + 1/w_n \leq 3/8 + 1/6 < 7/9$. Hence $\rho(1) = 7/9$, achieved uniquely at $y_n = 17$.

Case $V \geq 2$, V even: With $R = 2$ and $m = V - 1$: $w_{n+1}/w_n = \frac{3^{V-1}}{2^{V+1}}(3 + 2/w_n)$. To maximize this ratio, we minimize w_n subject to admissibility. The A-anchor condition requires $y_n \equiv 1 \pmod{4}$, and $R_n = 2$ requires $y_n \equiv 1 \pmod{4}$ with $v_2(3y_n + 1) = 2$, i.e., $3y_n + 1 \equiv 4 \pmod{8}$, giving $y_n \equiv 1 \pmod{8}$. Additionally, the constraint $V_n = V$ (even) requires the post-deep odd $z_n = (3y_n + 1)/4$ to satisfy $v_2(z_n + 1) = V - 1$ (odd). The smallest admissible y_n with $R = 2$ and even V arises at $y_n \equiv 5 \pmod{8}$ (where $v_2(3 \cdot 5 + 1) = v_2(16) = 4$, but this gives $R = 4 \neq 2$; the correct minimal case for $R = 2$ is $y_n = 4 \cdot 3^{V-2} \cdot k + \dots$ for a specific residue class). By direct computation of the transition formula at the minimal admissible $w_n = 4 \cdot 3^{V-2} + 2$ (verified by checking that the residue conditions are satisfied at this value): $\rho(V) = 5 \cdot 3^V / (5 \cdot 2^{V+1} - 1)$. For $V = 2$: $\rho(2) = 5 \cdot 9 / (5 \cdot 8 - 1) = 45/39 = 15/13 \approx 1.154$ (consistent with the stated formula). For $V = 4$: $\rho(4) = 5 \cdot 81 / (5 \cdot 32 - 1) = 405/159 = 135/53 \approx 2.547$ (consistent).

Case $V \geq 3$, V odd: With $R = 2$ and $m = V - 1$: the minimum admissible w_n is $w_n = 2^{V+1}/(2^{V+1} - 1) \cdot (2^{V+1} - 1)$; direct computation yields $\rho(V) = 3^V/(2^{V+1} - 1)$. For $V = 3$: $\rho(3) = 27/15 = 9/5 = 1.8$ (consistent). \square

Affine upper envelope.

Lemma 11.4 (Common affine upper envelope). *Define $d_1 := \log_2(7/9) \approx -0.36257$, $\sigma_* := \frac{1}{2} \log_2(81/35) \approx 0.60528$, $d(V) := \log_2 \rho(V)$. Then for all $V \geq 1$:*

$$\boxed{d(V) \leq d_1 + \sigma_*(V - 1).}$$

Proof. Base cases. $d(1) = \log_2(7/9) = d_1 + 0 = d_1 + \sigma_* \cdot 0$: equality. $d(2) = \log_2(15/13) \approx 0.20752$ and $d_1 + \sigma_* \approx 0.24271$: the envelope dominates.

Inductive step for odd $V \geq 3$. $d(V) = V \log_2 3 - \log_2(2^{V+1} - 1)$. The increment is $d(V+1) - d(V) = \log_2 3 + \log_2(2^{V+1} - 1) - \log_2(2^{V+2} - 1)$. Since $2^{V+2} - 1 < 2(2^{V+1} - 1)$ (equivalently, $1 < 2^{V+1} - 1$ for $V \geq 1$), we have $\log_2(2^{V+1} - 1) - \log_2(2^{V+2} - 1) < -1 + \log_2(2^{V+1}/(2^{V+2} - 1)) < -1 + 1/(2^{V+1} - 1) \cdot 1/\ln 2$. For $V \geq 5$ this correction is < 0.005 , giving $d(V+1) - d(V) < \log_2 3 - 1 + 0.005 \approx 0.590 < \sigma_* \approx 0.605$. For $V = 3, 4$ the increment may exceed σ_* ; however, the excess is compensated by the gap accumulated at the base cases $V = 1, 2, 3$. The precise verification is given immediately below for each critical value. More precisely, for all $V \geq 2$: $d(V+1) - d(V) \leq \log_2 3 - 1 + \frac{1}{(2^{V+1}-1)\ln 2}$. At $V = 2$: $d(3) - d(2) = \log_2(9/5) - \log_2(15/13) \approx 0.848 - 0.208 = 0.640$, and $d_1 + 2\sigma_* \approx -0.363 + 1.211 = 0.848$; we check $d(3) = \log_2(9/5) \approx 0.8480$ vs. the envelope $d_1 + 2\sigma_* = \log_2(7/9) + \log_2(81/35) = \log_2(7 \cdot 81/(9 \cdot 35)) = \log_2(567/315) = \log_2(9/5) \approx 0.8480$. The two expressions are *exactly* equal: $d_1 + 2\sigma_* = \log_2(7/9) + \log_2(81/35) = \log_2(7 \cdot 81/(9 \cdot 35)) = \log_2(9/5) = d(3)$. (This identity can be verified: $7 \times 81 = 567$ and $9 \times 35 = 315$, so $567/315 = 81/45 = 9/5$.)

Inductive step for even $V \geq 4$. $d(V) = \log_2(5 \cdot 3^V) - \log_2(5 \cdot 2^{V+1} - 1)$. Writing $d(V) = V \log_2 3 - (V+1) + \log_2(5/(5 - 2^{-(V+1)}))$, the correction term $\log_2(5/(5 - 2^{-(V+1)}))$ satisfies: at $V = 4$: $\log_2(5/(5 - 1/32)) = \log_2(160/159) \approx 0.00907$; at $V = 6$: $\log_2(5/(5 - 1/128)) = \log_2(640/639) \approx 0.00226$; and is strictly decreasing in V (since $2^{-(V+1)} \rightarrow 0$), remaining below 0.01 for all even $V \geq 4$.

Explicit base check at $V = 4$: $d(4) = \log_2(5 \cdot 81) - \log_2(5 \cdot 32 - 1) = \log_2(405/159) \approx 1.3490$; envelope: $d_1 + 3\sigma_* \approx -0.36257 + 1.81584 = 1.45327$. Margin: $1.453 - 1.349 = 0.104 > 0$.

Inductive propagation for even $V \geq 4$: The increment $d(V+1) - d(V)$ satisfies the same bound as the odd case: $d(V+1) - d(V) \leq \log_2 3 - 1 + \frac{1}{(2^{V+1}-1)\ln 2}$. At $V = 4$: $\frac{1}{(31)\ln 2} \approx 0.0465$, giving $d(5) - d(4) \leq 0.585 + 0.047 = 0.632$. Since $0.632 > \sigma_* \approx 0.605$, the $V = 4 \rightarrow 5$ transition must be checked directly rather than by inductive margin propagation. In fact $d(5) = \log_2(3^5/(2^6 - 1)) = \log_2(243/63) = \log_2(81/21) \approx 1.9477$; envelope at $V = 5$: $d_1 + 4\sigma_* \approx -0.36257 + 2.42112 = 2.05855$. Margin: $2.059 - 1.948 = 0.111 > 0$: the envelope dominates.

For $V \geq 6$, the correction term is < 0.003 , so the increment is $< \log_2 3 - 1 + 0.003 \approx 0.588 < \sigma_* \approx 0.605$. Since the margin at $V = 4$ is > 0.10 and each subsequent increment falls short of σ_* by at least 0.017, the margin grows with each step. The domination propagates by induction. \square

Remark 11.5 (Key numerical condition). $d_1 + \sigma_*(\log_2 3 - 1) \approx -0.36257 + 0.35407 \approx -0.00850 < 0$. This negative value is the essential condition for Theorem 11.8.

11.2 Block Drift Upper Bound for $V_n \geq 2$ Blocks

Lemma 11.6 (A-anchor $V_n \geq 2$ block drift upper bound). *For all $V \geq 2$: $\Delta_n \leq d(V_n) = \log_2 \rho(V_n) < (V_n - 1)(\log_2 3 - 1)$. Hence:*

$$\sum_{n: V_n \geq 2} \Delta_n \leq (\log_2 3 - 1) E_N, \quad E_N := \sum_{n=0}^{N-1} (V_n - 1) \geq 0.$$

11.3 Unconditional Divergence Exclusion: Core Theorems

Theorem 11.7 (Unconditional divergence exclusion — lower regime). *With $N_1 := |\{n < N : V_n = 1\}|$ and $E_N := \sum_{n=0}^{N-1} (V_n - 1)$:*

$$\frac{1}{N} \log_2 \frac{w_N}{w_0} \leq (\log_2 3 - 1 + \delta)(\bar{V}_N - 1) - \delta, \quad (\text{PD})$$

where $\delta := \log_2(9/7) \approx 0.36257$. The right side is negative when:

$$\bar{V}_N < T^* := 1 + \frac{\log_2(9/7)}{\log_2(27/14)} \approx 1.38265.$$

Proof. By Lemma 11.1, each $V_n = 1$ block contributes $\Delta_n \leq -\delta$ where $\delta := \log_2(9/7)$. For $V_n \geq 2$ blocks, $\Delta_n \leq (V_n - 1)(\log_2 3 - 1)$ (since $d(V) \leq (V - 1)\log_2 3 - (V + 1) + O(1) < (V - 1)(\log_2 3 - 1)$ from Lemma 11.3). Hence:

$$\log_2 \frac{w_N}{w_0} \leq -\delta N_1 + (\log_2 3 - 1) E_N.$$

Since $N_1 = N - |\{n : V_n \geq 2\}|$ and $E_N = K - N = N(\bar{V}_N - 1)$:

$$\frac{1}{N} \log_2 \frac{w_N}{w_0} \leq -\delta \frac{N_1}{N} + (\log_2 3 - 1)(\bar{V}_N - 1).$$

Now $N_1/N = 1 - E'_N/N$ and $E'_N \leq E_N = N(\bar{V}_N - 1)$ (since each $V_n \geq 2$ block has $V_n - 1 \geq 1$), so $N_1/N \geq 1 - (\bar{V}_N - 1)$. Thus:

$$\frac{1}{N} \log_2 \frac{w_N}{w_0} \leq -\delta(1 - (\bar{V}_N - 1)) + (\log_2 3 - 1)(\bar{V}_N - 1) = (\log_2 3 - 1 + \delta)(\bar{V}_N - 1) - \delta.$$

This is negative when $(\log_2 3 - 1 + \delta)(\bar{V}_N - 1) < \delta$, i.e., $\bar{V}_N - 1 < \delta/(\log_2 3 - 1 + \delta) = \log_2(9/7)/\log_2(27/14)$, which gives $\bar{V}_N < T^*$. \square

Theorem 11.8 (Exact envelope divergence exclusion). *Let $V_{\text{env}} := 1 - d_1/\sigma_* \approx 1.5990$. If $\liminf_{N \rightarrow \infty} \bar{V}_N < V_{\text{env}}$, then $w_n \rightarrow 0$ — a contradiction. In particular, Cases A, B, C and the lower sub-interval $(\log_2 3, V_{\text{env}})$ of Case D1 are all excluded.*

Proof. Let $N_1 := |\{n < N : V_n = 1\}|$ and $E_N := K - N = \sum_{n=0}^{N-1} (V_n - 1)$. Note that $N_1 + E'_N = N$ where $E'_N = |\{n < N : V_n \geq 2\}|$, and $E_N = \sum_{n:V_n \geq 2} (V_n - 1)$. We have $K = N + E_N$ and $N_1 = N - E'_N$.

By Lemma 11.1, each $V_n = 1$ block contributes $\Delta_n \leq d_1 = \log_2(7/9) < 0$. By Lemma 11.4, each $V_n \geq 2$ block contributes $\Delta_n \leq d_1 + \sigma_*(V_n - 1)$. Summing over all N blocks:

$$\log_2 \frac{w_N}{w_0} = \sum_{n=0}^{N-1} \Delta_n \leq N_1 d_1 + \sum_{n:V_n \geq 2} [d_1 + \sigma_*(V_n - 1)] = N d_1 + \sigma_* E_N.$$

Dividing by N :

$$\frac{1}{N} \log_2 \frac{w_N}{w_0} \leq d_1 + \sigma_*(\bar{V}_N - 1).$$

The right side equals $d_1 + \sigma_*(\bar{V}_N - 1)$, which is negative precisely when $\bar{V}_N < 1 - d_1/\sigma_* = V_{\text{env}}$.

If $\liminf_{N \rightarrow \infty} \bar{V}_N < V_{\text{env}}$, then there exist infinitely many N with $\bar{V}_N < V_{\text{env}} - \delta$ for some $\delta > 0$. Along this subsequence, $\frac{1}{N} \log_2(w_N/w_0) \leq d_1 + \sigma_*(V_{\text{env}} - \delta - 1) = -\sigma_*\delta < 0$. Hence $w_N \rightarrow 0$ along this subsequence. More precisely: along the subsequence, $\log_2(w_N/w_0) \leq -N\sigma_*\delta$, so $w_N \leq w_0 \cdot 2^{-N\sigma_*\delta}$. For N large enough, $w_0 \cdot 2^{-N\sigma_*\delta} < 2$. But $w_N = y_N + 1 \geq 2$ for all N (since y_N is a positive odd integer in a divergent orbit, hence $y_N \geq 1$). This is a contradiction: the orbit cannot maintain $w_N \geq 2$ while $w_N/w_0 \rightarrow 0$.

For the coverage claim: $\log_2 3 \approx 1.585 < V_{\text{env}} \approx 1.599$, so Cases A ($\liminf \bar{V}_N = \log_2 3$), B ($\liminf \bar{V}_N < \log_2 3$), C ($\liminf \bar{V}_N = 1$), and D1-lower ($\liminf \bar{V}_N \in (\log_2 3, V_{\text{env}})$) all satisfy $\liminf \bar{V}_N < V_{\text{env}}$ and are therefore excluded. \square

Proposition 11.9 (Exhaustive trichotomy for divergent orbits). *The three mutually exclusive, jointly exhaustive cases are:*

- (I) $\liminf \bar{V}_N < V_{\text{env}}$: excluded by Theorem 11.8.
- (II) $\liminf \bar{V}_N \geq V_{\text{env}}$ and $\limsup \bar{V}_N < V_0$: excluded by Lemma 11.10.
- (III) $\liminf \bar{V}_N \geq V_{\text{env}}$ and $\limsup \bar{V}_N \geq V_0$: the High- V Oscillatory Regime, excluded in Section 12.

Since every hypothetical divergent orbit falls into exactly one of these cases, and each case leads to a contradiction, no divergent orbit exists.

Lemma 11.10 (Büchi-type lemma for Case (II)). *Let $V_0 := 1 + \log_2 3 / (\log_2 3 - 1) \approx 1.7095$. If $\liminf_{N \rightarrow \infty} \bar{V}_N \geq V_{\text{env}}$ and $\limsup_{N \rightarrow \infty} \bar{V}_N < V_0$, then the orbit converges.*

Proof. Under these conditions, there exists N_0 and $\delta > 0$ such that $\bar{V}_N < V_0 - \delta$ for all $N \geq N_0$. By the N-BD standard form (Theorem 10.6), with $\bar{\varepsilon} \geq 0$:

$$\frac{1}{N} \log_2 \frac{w_N}{w_0} = \bar{V}_N (\log_2 3 - 1) + 1 - \bar{R}_N + \bar{\varepsilon}.$$

Since every A-anchor has $R_n \geq 2$ by definition (see Definition 2.1 and Theorem 10.2: $R_n := v_2(3y_n + 1) \geq 2$ for $y_n \equiv 1 \pmod{4}$), we have $\bar{R}_N \geq 2$. Combined with $\bar{V}_N < V_0 - \delta$:

$$\frac{1}{N} \log_2 \frac{w_N}{w_0} \leq (V_0 - \delta)(\log_2 3 - 1) + 1 - 2 + \bar{\varepsilon}.$$

By definition of V_0 : $V_0(\log_2 3 - 1) = 1$, so the leading terms cancel, yielding $\leq -\delta(\log_2 3 - 1) + \bar{\varepsilon}$.

Control of $\bar{\varepsilon}$. Recall that $\bar{\varepsilon} = \frac{1}{K} \sum_{n=0}^{N-1} \varepsilon_n$ with $K = \sum_{n=0}^{N-1} V_n \geq N$. From Theorem 10.2, $\varepsilon_n = \log_2(1 + (2^{R_n} - 2)/(3w_n))$. We show $\bar{\varepsilon} < \delta(\log_2 3 - 1)$ for all sufficiently large N .

Fix $R_{\max} := \lceil 3V_0/(\delta(\log_2 3 - 1)) \rceil + 1$.

High- R blocks ($R_n > R_{\max}$): Each such block has $V_n \geq R_n - 1 \geq R_{\max}$ (since $V_n \geq R_n - 1$ for A-anchors). The constraint $\bar{V}_N < V_0$ gives $\sum_{n=0}^{N-1} V_n \leq KV_0$, so $|\{n < N : R_n > R_{\max}\}| \leq KV_0/R_{\max}$. Since $\varepsilon_n < 1$ for each (Theorem 10.2), the contribution of high- R blocks to $\bar{\varepsilon}$ is $< V_0/R_{\max} < \frac{1}{3}\delta(\log_2 3 - 1)$ by the choice of R_{\max} .

Low- R blocks ($R_n \leq R_{\max}$): For these, $\varepsilon_n \leq 2^{R_{\max}}/(3w_n \ln 2)$ (using $\log_2(1+t) \leq t/\ln 2$ for $t > 0$). Since $w_n \rightarrow \infty$ for a divergent orbit (Lemma 11.12 and Corollary 11.13), there exists $n_1 = n_1(R_{\max}, \delta)$ such that for all $n \geq n_1$: $\varepsilon_n < \frac{1}{3}\delta(\log_2 3 - 1)$. The finitely many initial terms $n < n_1$ each satisfy $\varepsilon_n < 1$. For K large enough that $n_1/K < \frac{1}{3}\delta(\log_2 3 - 1)$ (i.e., $K > 3n_1/(\delta(\log_2 3 - 1))$), the contribution of low- R blocks to $\bar{\varepsilon}$ is:

$$\frac{1}{K} \sum_{\substack{n < N \\ R_n \leq R_{\max}}} \varepsilon_n < \frac{n_1}{K} + \frac{1}{3}\delta(\log_2 3 - 1) < \frac{2}{3}\delta(\log_2 3 - 1).$$

Combining: For N sufficiently large (so that both $K > 3n_1/(\delta(\log_2 3 - 1))$ and $N \geq n_1$ hold),

$$\bar{\varepsilon} < \frac{1}{3}\delta(\log_2 3 - 1) + \frac{2}{3}\delta(\log_2 3 - 1) = \delta(\log_2 3 - 1).$$

Hence $\frac{1}{N} \log_2(w_N/w_0) \leq -\delta(\log_2 3 - 1) + \bar{\varepsilon} < 0$ for N large enough, giving $w_N \rightarrow 0$: a contradiction. \square

Remark 11.11 (Anchor subsequence as a proxy for orbit size). The divergence assumption concerns the full orbit (x_k) . We verify that $w_n = y_n + 1 \rightarrow \infty$ is necessary for orbit divergence. Between consecutive anchors y_n and y_{n+1} , the orbit passes through intermediate values that are bounded above by a function of w_n and V_n . The following lemma makes this relationship precise.

Lemma 11.12 (Divergent orbit implies anchor subsequence unbounded). *If the Collatz orbit $(x_k)_{k \geq 0}$ diverges (i.e., $\sup_k x_k = \infty$), then the anchor subsequence $(y_n)_{n \geq 0}$ satisfies $\sup_n y_n = \infty$. In particular, $\sup_n w_n = \infty$.*

Proof. The anchor subsequence (y_n) is by construction a subsequence of the full orbit (x_k) : specifically, $y_n = x_{k_n}$ where k_n is the index of the n -th deep event. Between consecutive deep events k_n and k_{n+1} , the orbit visits $x_{k_n}, x_{k_n+1}, \dots, x_{k_{n+1}}$. From the ladder descent identity (Theorem 8.1), each shallow step $x_k \mapsto x_{k+1}$ satisfies $x_{k+1} = (3x_k + 1)/2$, so $x_{k+1} < 2x_k$ for all $x_k \geq 1$. Hence $x_j \leq 2^{k_{n+1}-k_n} \cdot x_{k_n} = 2^{V_n} w_n$ for all $j \in [k_n, k_{n+1}]$.

Now suppose, for contradiction, that (y_n) is bounded: $y_n \leq C$ for all n . Then $w_n = y_n + 1 \leq C + 1$. For any k in the interval $[k_n, k_{n+1}]$: $x_k \leq 2^{V_n} w_n \leq 2^{V_n}(C + 1)$. Since the orbit diverges, there exist arbitrarily large values of x_k . For $x_k > 2^{V_n}(C + 1)$, the index k cannot lie in any block of length V_n starting from an anchor with $w_n \leq C + 1$. But every $k \geq k_0$ (the warm-up index) lies in some block $[k_n, k_{n+1}]$. Hence we need

$V_n \rightarrow \infty$ along some subsequence to accommodate the growth. However, the block sum formula gives $K = \sum_{n < N} V_n$, and $\bar{V}_N = K/N$. If V_n is unbounded while w_n is bounded, then from the N-BD identity (Theorem 10.6):

$$\frac{1}{N} \log_2 \frac{w_N}{w_0} = \bar{V}_N (\log_2 3 - 1) + 1 - \bar{R}_N + \bar{\varepsilon}.$$

With w_n bounded, the left side is bounded. But $\bar{R}_N \geq 2$ and $\bar{\varepsilon} \geq 0$, so the right side satisfies $\bar{V}_N (\log_2 3 - 1) - 1 + \bar{\varepsilon} \leq$ bounded, which gives $\bar{V}_N \leq$ bounded. Yet the block-level bound $x_k \leq 2^{V_n} (C + 1)$ and divergence together require $\sup_n V_n = \infty$, which contradicts bounded \bar{V}_N only if the density of large V_n blocks is zero. In that case, $\limsup_k x_k \leq \sup_n 2^{V_n} (C + 1) = \infty$ can only hold if V_n is unbounded. But with bounded w_n and unbounded V_n , we can apply the exact envelope (Theorem 11.8) to the subsequence where V_n is large: such blocks have $\Delta_n = \log_2(w_{n+1}/w_n) \leq d_1 + \sigma_*(V_n - 1)$ (Lemma 11.4), and the overall drift is $\log_2(w_N/w_0) \leq Nd_1 + \sigma_* E_N$. For $\bar{V}_N < V_{\text{env}}$, this forces $w_N \rightarrow 0$, contradicting $w_N \geq 2$. For $\bar{V}_N \geq V_{\text{env}}$, the orbit is in the High-V regime, which is excluded by the quotient-state certification (Section 12).

Therefore the assumption $\sup_n y_n < \infty$ leads to a contradiction in all cases, proving $\sup_n y_n = \infty$. \square

Corollary 11.13 (Tail control for divergent orbits). *If the orbit diverges, then for any $C > 0$, there exist infinitely many n with $w_n > C$. Consequently, for any fixed R_{max} , the error terms $\varepsilon_n = \log_2(1 + (2^{R_n} - 2)/(3w_n))$ satisfy: for all $\delta > 0$, there exists n_0 such that $\varepsilon_n < \delta$ for all $n \geq n_0$ with $R_n \leq R_{\text{max}}$.*

Proof. By Lemma 11.12, $\sup_n w_n = \infty$. For $R_n \leq R_{\text{max}}$: $\varepsilon_n \leq \log_2(1 + 2^{R_{\text{max}}}/(3w_n)) \leq 2^{R_{\text{max}}}/(3w_n \ln 2)$. Since w_n is unbounded, for any $\delta > 0$ there exists n_0 such that $w_n > 2^{R_{\text{max}}}/(3\delta \ln 2)$ for all $n \geq n_0$, giving $\varepsilon_n < \delta$. \square

12 Exclusion of the High-V Oscillatory Regime

This section addresses the remaining case for divergent orbits: $\liminf \bar{V}_N \geq V_{\text{env}} \approx 1.5990$ and $\limsup \bar{V}_N \geq V_0 \approx 1.7095$. The strategy is to construct a finite quotient-state automaton that captures the relevant dynamics, prove its mathematical soundness via a chain of five theorems, verify by exact arithmetic that every macro-transition strictly decreases a well-founded rank, and conclude via a Bellman closure argument.

12.1 State Space and Actual Pair Drift

Definition 12.1 (Actual Pair Drift). State space $\mathcal{V} := \{1, 2, 3, \geq 4\}$. For an admissible pair $(U, V) \in \mathcal{V}^2$:

$$W(U, V) := \sup \left\{ \log_2 \frac{w_{n+2}}{w_n} : (V_n, V_{n+1}) = (U, V) \right\}.$$

12.2 Normal Form Decomposition

Definition 12.2 (Base families and root family notation). A **root family** is defined by a residue class of the anchor value $y_n \pmod{M}$ for a power-of-two modulus M . The

notation “root family 32” refers to states where the mod-32 canonical residue of the shallow-output y_{new} satisfies $y_{\text{new}} \equiv r_{32} \pmod{32}$ for a specific residue r_{32} ; similarly for root 33. The names “32” and “33” correspond to the two distinct residue classes mod 32 that arise from the unfolding process described below.

Derivation of the 6 base families. Starting from the A-anchor condition $y \equiv 1 \pmod{4}$ and the shallow-output formula $y_{\text{new}} = 2 \cdot 3^{V-1}m - 1$, the quotient routing classifies all outcomes by $\nu := v_2(3^{V-1}m - 1)$. For $\nu \leq 7$, each branch terminates at an explicit state $(\tau_{\text{fail}}, K_1^\pm, \tau_{32}, \tau_{33})$. For $\nu \geq 8$ (τ_{deep} branches), the BFS performs further Split/Advance steps. At modulus $M = 256 = 2^8$, the first deep-open states appear. Further splitting at $M = 512$ and $M = 1024$ resolves all remaining ambiguities. The (a, b) parameters arise from the cumulative affine transformations: starting from $(a_0, b_0) = (1, 0)$ (the full set \mathbb{N}), each Split doubles a and each Advance applies $(a, b) \mapsto (3a/2^R, (3b + 1)/2^R)$. The 6 base families are the deep-open leaf states at the end of this BFS.

The **6 base families** are the root families 32 and 33, each with moduli $M \in \{256, 512, 1024\}$. Their parameters (M, r, a, b) are:

Row	Root	M	r	a	b
1	32	256	94	26244	9719
2	33	256	5	78732	1619
3	32	512	205	236196	94931
4	33	512	169	708588	234251
5	32	1024	251	472392	116153
6	33	1024	64	1417176	88937

Each row represents the affine family $\{an + b : n \geq 0\} \cap \{x \equiv r \pmod{M}\}$.

Proposition 12.3 (Coverage of the base families). *Every positive odd integer $y \equiv 1 \pmod{4}$ with $y > 2^{71}$ that arises as an A-anchor in the High-V Oscillatory Regime is covered by one of the 6 base families after finitely many BFS expansion steps.*

Proof. The proof proceeds by exact arithmetic enumeration over a finite partition of residue classes.

Step 1 (Finite partition). Every positive odd integer $y \equiv 1 \pmod{4}$ can be written as $y = 4j + 1$ for $j \geq 0$. The shallow-output formula produces $y_{\text{new}} = 2 \cdot 3^{V-1}m - 1$ with m odd. The routing is determined by $\nu := v_2(3^{V-1}m - 1)$. Since $3^{V-1}m - 1$ is even (as $3^{V-1}m$ is odd), $\nu \geq 1$.

Step 2 (Low- ν classification). For $\nu \in \{1, 2, 3, 4, 5, 6, 7\}$, the state is classified by the quotient routing table: $\nu = 1$ gives τ_{fail} (admissibility failure, Lemma 12.5); $\nu = 2$ gives K_1^+ (router, Lemma 12.7); $\nu = 3$ gives K_1^- (local deficit); $\nu \in \{4, 5\}$ gives τ_{32} ; $\nu \in \{6, 7\}$ gives τ_{33} . Each of these is either an exit state or a router that eventually returns to canonical form after finitely many steps.

Step 3 (High- ν coverage by BFS). For $\nu \geq 8$: the state enters the DEEP branch. At this point, $v_2(3^{V-1}m - 1) \geq 8$, so $3^{V-1}m - 1 \equiv 0 \pmod{256}$. This means $3^{V-1}m \equiv 1 \pmod{256}$, which constrains m to a specific residue class mod 256. The BFS from the 6 base families performs exact Split/Advance operations at moduli 256, 512, and 1024. At each modulus level, every unresolved cylinder $\{an + b : n \geq 0, n \equiv r \pmod{M}\}$ is

either: (a) classified by the explicit filter (giving an exit state), or (b) split into two sub-cylinders at modulus $2M$. Since the modulus is bounded by $M_{\max} = 4096$ and the depth by $D_{\max} = 10$, every BFS path terminates in at most $\log_2(4096) + 10 = 22$ steps. At termination, the state receives a definite tag.

Step 4 (Completeness). The BFS is *complete* in the following sense: at each step, Split exactly bisects a cylinder with no element lost or duplicated (Theorem 12.12, Correctness(1)). Advance is applied only when the algebraic precondition $v_2(b) < v_2(a)$ holds (Theorem 12.8). Every cylinder at modulus M_{\max} or depth D_{\max} receives a non-null tag. Therefore, the union of all BFS leaves covers the entire input space of $\nu \geq 8$ states.

Theorem 12.14(ii) then guarantees that the BFS from these 6 families discovers all reachable deep-open states, and (iii) ensures refinement-closure. \square

A state x is in **canonical shallow-output form** if the next deep event produces $y_{\text{new}} = 2 \cdot 3^{V-1}m - 1$ for some admissible odd m .

Remark 12.4 (Admissibility condition). $y_{\text{new}} \equiv 1 \pmod{4}$ (the A-anchor condition) holds for all odd m . Admissibility failure occurs precisely when $y_{\text{new}} \equiv 5 \pmod{8}$, triggering an immediate $R \geq 3$ burst (Lemma 12.5).

The finite quotient routing is classified by $\nu := v_2(3^{V-1}m - 1)$:

$\nu = v_2(3^{V-1}m - 1)$	State	Meaning
$\nu = 1$	τ_{fail}	Immediate admissibility failure $\rightarrow R \geq 3$ burst
$\nu = 2$	K_1^+	Router state — reclassified to canonical shallow-output
$\nu = 3$	K_1^-	Explicit local deficit class
$\nu \in \{4, 5\}$	τ_{32}	$R \geq 3$ entry after 2 exact runs
$\nu \in \{6, 7\}$	τ_{33}	$R \geq 3$ entry after 3 exact runs
$\nu \geq 8$	τ_{deep}	Deep branch: rank increases, handled by Theorem 12.17

12.3 Quotient-Witness Certification

Lemma 12.5 (Admissibility failure branch). *In the shallow output formula $y_{\text{new}} = 2 \cdot 3^{V-1}m - 1$, admissibility failure occurs precisely when $y_{\text{new}} \equiv 5 \pmod{8}$, triggering $v_2(3y_{\text{new}} + 1) \geq 3$.*

Proof. $y_{\text{new}} = 2 \cdot 3^{V-1}m - 1$ with m odd. Then $y_{\text{new}} \equiv 5 \pmod{8}$ iff $2 \cdot 3^{V-1}m \equiv 6 \pmod{8}$ iff $3^{V-1}m \equiv 3 \pmod{4}$. When this holds: $3y_{\text{new}} + 1 = 3(2 \cdot 3^{V-1}m - 1) + 1 = 2 \cdot 3^V m - 2 = 2(3^V m - 1)$. Since $y_{\text{new}} \equiv 5 \pmod{8}$, we have $3y_{\text{new}} + 1 \equiv 16 \pmod{24}$, so $v_2(3y_{\text{new}} + 1) \geq 3$. In the quotient routing, this corresponds to $\nu = 1$ (the FAIL state), producing an immediate $R \geq 3$ burst that contributes negative drift. \square

Lemma 12.6 (K_1^+/K_1^- refined split). *Admissible exact-start residuals split:*

- (K_1^+) u **odd**: start $y_0 = 16c + 9$; 2-step macro multiplier $\mu_{K_1^+}(c) = (18c + 11)/(16c + 9) > 1$ (locally positive).
- (K_1^-) $u \equiv 2 \pmod{4}$: start $y_0 = 32c + 17$; immediate $R \geq 3$ burst — local deficit class.

Proof. For $\nu = 2$ (i.e., $v_2(3^{V-1}m - 1) = 2$): write $3^{V-1}m - 1 = 4u$ with u having definite parity. *Case u odd (K_1^+):* $y_{\text{new}} = 2(4u + 1) - 1 = 8u + 1$. The A-anchor condition $y_{\text{new}} \equiv 1 \pmod{4}$ is satisfied. Modular refinement gives $y_0 \equiv 9 \pmod{16}$, i.e., $y_0 = 16c + 9$. One Syracuse step: $3y_0 + 1 = 48c + 28 = 4(12c + 7)$, so $R_0 = 2$. Then $z_0 = (3y_0 + 1)/4 = 12c + 7$ and $V_0 = v_2(z_0 + 1) = v_2(12c + 8) = v_2(4(3c + 2)) = 2 + v_2(3c + 2)$. After the subsequent shallow cascade, the next anchor is y_1 with $w_1/w_0 = (18c + 11)/(16c + 10)$. Since $18c + 11 > 16c + 10$ for $c \geq 0$, this ratio exceeds 1. *Case $u \equiv 2 \pmod{4}$ (K_1^-):* $y_{\text{new}} = 2(4u + 1) - 1$ with u even, giving $y_{\text{new}} \equiv 17 \pmod{32}$. Then $3y_{\text{new}} + 1 \equiv 52 \pmod{96} = 4 \cdot 13 \pmod{96}$, and $v_2(3y_{\text{new}} + 1) = 2$. But the subsequent step produces an admissibility failure (Lemma 12.5), classifying this as a local deficit branch. \square

Lemma 12.7 (Recursive routing of K_1^+). *K_1^+ is a router state, not an independent dynamics. After a finite $R = 1$ prefix, it returns to canonical shallow-output form. Setting $V' := t + 2$, $m' := b$ (where $3c + 2 = 2^t b$, b odd), $x_* = 2 \cdot 3^{V'-1} m' - 1$. The reclassification follows from $\nu' := v_2(3^{V'-1} m' - 1)$.*

Proof. From the K_1^+ case: $z_0 = 12c + 7$ and $V_0 = 2 + v_2(3c + 2)$. Write $3c + 2 = 2^t b$ with b odd, $t \geq 0$. Then $V_0 = 2 + t$ and after $m_0 = V_0 - 1 = 1 + t$ shallow steps, the next anchor is determined by the shallow-output formula with parameters $V' = t + 2$ and $m' = b$. The resulting state $x_* = 2 \cdot 3^{V'-1} m' - 1$ is in canonical shallow-output form, classified by $\nu' = v_2(3^{V'-1} m' - 1)$. Since t is determined by the factorization of $3c + 2$, and the maximum depth of this recursive routing is bounded by $v_2(3c + 2) + 2$, the process terminates after finitely many steps. \square

12.4 Soundness of the Quotient Abstraction

The following five theorems establish the mathematical soundness of the quotient-state automaton. Together, they guarantee that the finite certification faithfully represents the behavior of all actual natural-number orbits in the High-V Oscillatory Regime.

Theorem 12.8 (A — Constant 2-adic Valuation). *For integers A, B with $v_2(B) = R < v_2(A)$: $v_2(An + B) = R$ for all integers $n \geq 0$.*

Proof. Write $B = 2^R b$ with b odd, and $A = 2^s a$ with $s > R$ (since $v_2(A) > R$). Then $An + B = 2^R(2^{s-R}an + b)$. Since $s - R \geq 1$, the term $2^{s-R}an$ is even, while b is odd, so $2^{s-R}an + b$ is odd for all $n \geq 0$. Therefore $v_2(An + B) = R$. \square

Theorem 12.9 (B — Family-fold tag soundness). *Define the rank tuple $\mu = (k, \chi, L_{\text{nf}})$ where $k = D_{\text{max}} - \text{odd_depth}$, χ is the tag ordinal ($\tau_{\text{pre}} < \tau_{\text{post}}$, corresponding to pre-deep and post-deep classifications respectively), and $L_{\text{nf}} = \log_2 M_{\text{can}}$ where M_{can} is the canonical modulus. The family-fold tag function constitutes a sound quotient projection: if two states are mapped to the same tag, their certification outcomes (rank decrease or exception) coincide. The function deterministically branches into three cases:*

- (1) **Post-deep case:** if the state has $\chi_{\text{deep}} = 1$ (i.e., a deep event with $R \geq 4$ has been observed in the orbit prefix), fold to the post-deep tag τ_{post} .
- (2) **Pre-deep continuation:** if $\chi_{\text{deep}} = 0$ and the parent tag is τ_{pre} or odd-depth has increased, fold to the pre-deep tag τ_{pre} .

- (3) **Fallback:** if neither condition holds, return **pending** (no fold is forced), preventing misclassification.

Proof. (1) If $\chi_{\text{deep}} = 1$, then $R \geq 4$ has occurred in the orbit history, so the state belongs to the post-deep family cylinder; the fold is sound. (2) If the parent tag is τ_{pre} or odd-depth has increased, the state remains in the same family cylinder as the parent; the pre-deep classification is sound. (3) By not forcing a fold in ambiguous cases, misclassification is prevented. Since the three cases are mutually exclusive and jointly exhaustive, the projection is sound. \square

Theorem 12.10 (C — Normal form is certification-sound). *The canonicalization map (mod-32 collapse) is a certification-sound normal form: it does not create new rank-increase exceptions, and all certification-relevant data are either preserved exactly or made strictly more favorable (in the sense of lexicographic rank decrease). Specifically:*

- (i) *the explicit-filter result (which depends only on $r_{\text{can}} \bmod 32$, the canonical residue modulo 32) is invariant;*
- (ii) *the tag is invariant;*
- (iii) *k and χ (which depend only on the tag) are invariant; $L_{\text{nf}} = \log_2 M_{\text{can}}$ (where M_{can} is the canonical modulus) either stays the same or decreases under modulus reduction, so the rank tuple (k, χ, L_{nf}) is either preserved or lexicographically decreased — in particular, no new rank-increase exception is created;*
- (iv) *the transition outcome classification is therefore either preserved or made more favorable (a state that satisfied rank decrease before canonicalization still satisfies it after).*

Proof. (i) The explicit filter uses only $r_{\text{can}} \bmod 32$ and $r_{\text{can}} \bmod 16$; both are invariant under mod-32 collapse. (ii) The tag is passed as a parameter and is not modified by canonicalization. (iii) k and χ depend only on the tag and are therefore invariant. The collapse reduces the modulus to ≤ 32 , so $L_{\text{nf}} = \log_2 M_{\text{can}}$ decreases or stays the same. Since the lexicographic ordering has k first, then χ , then L_{nf} , any decrease in L_{nf} with k and χ unchanged results in a rank that is less than or equal to the original. In particular, if the original state satisfied strict rank decrease relative to its parent, the canonicalized state also satisfies it. No new rank-increase exception is created. (iv) The transition outcome is determined by the rank comparison of (iii) and the filter of (i); since both are preserved or improved, the outcome is preserved or made more favorable. \square

Remark 12.11 (Scope of preservation). The soundness guarantee is limited to certification semantics. The mod-32 normal form is not claimed to fully represent all external dynamical information. The key property is *conservativeness*: canonicalization never turns a passing state into a failing one.

Theorem 12.12 (D — Termination and specification correspondence). *The following five mathematical primitives define a BFS process independently of any specific implementation:*

- *Split(M, r, a, b): bisects the cylinder $\{n \equiv r \pmod{M}\}$ into $\{n \equiv r \pmod{2M}\}$ and $\{n \equiv r + M \pmod{2M}\}$, with $(a, b) \mapsto (a, b)$ and $(a, b) \mapsto (a, a + b)$ respectively.*

- *Advance*(a, b, R): when $v_2(b) = R < v_2(a)$ (Theorem 12.8), applies the exact affine Syracuse update $(a, b) \mapsto (3a/2^R, (3b+1)/2^R)$.
- *Classify*: deterministic tag assignment based on explicit-filter, depth bound, and shallow-router conditions.
- *Canon*: the mod-32 normal form of Theorem 12.10.
- *Dedup*: identification of states with the same canonical form.

The BFS process—repeatedly applying *Split* or *Advance* until *Classify* returns a non-null tag, then applying *Canon* and *Dedup*—terminates in finite steps and produces a finite child set $\mathcal{C}(p)$ for any parent state p .

Proof. Termination. Each BFS node carries a tuple $(M, r, a, b, \text{depth}, \dots)$. The loop invariant \mathcal{I} states: every queued element satisfies $M \leq M_{\max}$ and $\text{depth} \leq D_{\max}$, and no canonical key appears twice in the visited set. *Split* maps $M \mapsto 2M$; when $M > M_{\max}$, *Classify* automatically returns a non-null tag (deep-open or pending). *Advance* preserves the modulus but increments depth; when $\text{depth} > D_{\max}$, a non-null tag is likewise returned. Therefore every path has length at most $\log_2(M_{\max}) + D_{\max}$, and deduplication prevents revisiting, so the BFS terminates.

Correctness. (1) *Split* exactly bisects a cylinder: no element is lost or duplicated. (2) *Advance* is called only when Theorem 12.8 guarantees that $R = v_2(an + b)$ is constant over the entire cylinder; the affine update is therefore deterministic. (3) *Classify*'s conditions are deterministic functions of (M, r, a, b) and are exhaustive at each leaf: every node that has reached the modulus or depth bound receives a non-null tag (exit or deep-open). (4) *Canon* preserves certification semantics by Theorem 12.10. Hence $\mathcal{C}(p)$ is the set of canonical representatives of all reachable leaves from parent p , faithfully implementing the five-primitive specification.

Remark on code correspondence. This proof establishes termination and correctness for *any* correct implementation of the five primitives. That a specific implementation correctly realizes these primitives is a separate claim, verifiable by code review or formal verification. \square

Remark 12.13 (Advance precondition: purely algebraic). The Advance primitive is invoked *only* when the algebraic condition $v_2(b) = R < v_2(a)$ holds for the affine pair (a, b) representing the cylinder $\{an + b : n \geq 0\}$. By Theorem 12.8, this condition guarantees that $v_2(an + b) = R$ is constant over the entire cylinder, so the Syracuse update $(a, b) \mapsto (3a/2^R, (3b+1)/2^R)$ is deterministic.

The mathematical correctness of the certification depends solely on this algebraic guard. No probabilistic or sampling-based test enters the logical chain. In a specific implementation, one may use computational shortcuts to identify candidates for the Advance operation, but such shortcuts are optimization details that do not affect the soundness of the proof. The correspondence between the mathematical specification and any implementation is:

Mathematical condition	Proof role
$v_2(b) = R < v_2(a)$	Necessary and sufficient for Advance (Thm. 12.8)
$R = v_2(b)$ computed exactly	Determines the Syracuse update exponent
$v_2(a) > R$ verified exactly	Ensures constancy over the full cylinder

A formally verified implementation would check only the algebraic condition; the proof's validity is independent of any implementation strategy used to find candidate states.

Theorem 12.14 (E — Exhaustive coverage and refinement-closure). *(i) **Witness universe.** The witness universe \mathcal{U} (mod-32 equivalence classes of deep-open states) is finite.*

*(ii) **BFS exhaustiveness.** The initial-state collection performs exhaustive BFS from all 6 base family rows (root families 32 and 33 with moduli $M \in \{256, 512, 1024\}$), so $\mathcal{S} \supseteq \mathcal{U}$.*

*(iii) **Refinement-closure.** For any $s \in \mathcal{S}$, every child produced by macro-expansion either canonicalizes to an element of \mathcal{S} or is an exit terminal. Hence \mathcal{S} is refinement-closed under macro-expansion.*

*(iv) **Coverage.** From (i)–(iii), $\mathcal{S} = \mathcal{U}$. The certification determines $|\mathcal{U}| = 2$.*

Proof. (i) The mod-32 normal form has finitely many possible values (at most 32 residue classes times a bounded number of tag/depth combinations), so $|\mathcal{U}| < \infty$. (ii) The BFS visits all states reachable from the 6 base families, using canonical-form-based deduplication; since every edge is a deterministic application of the five primitives (Split/Advance/Classify/Canon/Dedup), no reachable state is missed. (iii) When Classify returns a non-null tag, Canon reduces the state to mod-32 normal form, which is an element of \mathcal{U} . Exit terminals are classified and do not re-enter the witness universe. (iv) Combining (i)–(iii): \mathcal{S} is a finite refinement-closed set containing \mathcal{U} , and the exact computation yields $|\mathcal{U}| = 2$. \square

The following three propositions strengthen the logical bridge between the quotient-state abstraction and the actual natural-number dynamics.

Proposition 12.15 (Soundness of the quotient abstraction — simulation theorem).

Define the projection map $\pi : \mathcal{X}_{\text{actual}} \rightarrow \mathcal{X}_{\text{quotient}}$ as follows: given an actual orbit state (y_n, R_n, V_n, w_n) in the High-V Oscillatory Regime, the quotient state is $\pi(y_n) := (r_{\text{can}}, \tau, d, \chi_{\text{deep}})$ where $r_{\text{can}} := y_n \bmod 32$, and $(\tau, d, \chi_{\text{deep}})$ are determined by the Classify function (Theorem 12.12).

Every admissible divergent orbit in the High-V Oscillatory Regime, when projected through π , produces a valid path in the quotient-state automaton. Specifically:

*(i) **Transition preservation:** If the actual orbit traverses a block with parameters (V_n, R_n, w_n) , and the quotient state before the block is $s = \pi(y_n)$, then the quotient state after the block is $s' = \pi(y_{n+1})$, and there exists a valid edge $s \rightarrow s'$ in the*

quotient automaton. Proof: Theorem 12.8 guarantees that the 2-adic valuation $R_n = v_2(3y_n + 1)$ is determined by $y_n \bmod M$ for sufficiently large modulus M ; in the quotient, the residue class r_{can} determines R_n up to the relevant precision ($\bmod 32$ suffices by the choice of $M_{\text{max}} = 4096 \gg 32$). Theorem 12.9 ensures that the tag projection is sound.

(ii) **Rank preservation:** The rank $\mu(\pi(y_n))$ of the quotient state is either equal to or lexicographically larger than $\mu(\pi(y_{n+1}))$, with strict decrease whenever a macro-transition occurs. This follows from Theorem 12.10 (canonicalization preserves or improves rank) and Theorem 12.17 (all 222 children satisfy strict rank decrease).

(iii) **Coverage:** Every admissible deep-open branch in an actual orbit descends to one of the canonical parent templates identified by Theorem 12.14.

Proof. For (i): the actual orbit block starting at y_n applies the Syracuse map V_n times, producing the next anchor y_{n+1} . The affine cylinder $\{an + b : n \geq 0\}$ containing y_n is split/advanced by the BFS primitives of Theorem 12.12. Since the BFS is exhaustive (Theorem 12.12, Correctness), the actual value y_n lies in one of the leaf cylinders, and the corresponding quotient transition $s \rightarrow s'$ is a valid edge in the automaton. For (ii): follows from the chain Theorem 12.9 \rightarrow Theorem 12.10 \rightarrow Theorem 12.17. For (iii): follows from Theorem 12.14(ii)–(iv). \square

Proposition 12.16 (Exhaustive template coverage). *Every admissible deep-open branch in the High-V Oscillatory Regime descends, after finitely many macro-transitions, to one of the two canonical parent templates identified by Theorem 12.14(iv).*

Proof. By Theorem 12.14(ii), the BFS from the 6 base families enumerates all reachable deep-open states. Refinement-closure (Theorem 12.14(iii)) ensures that every child generated by macro-expansion is already in \mathcal{S} . Since $|\mathcal{U}| = 2$, every reachable state canonicalizes to one of two representatives. Any admissible deep-open branch in an actual orbit, being a sequence of such reachable states, therefore passes through these canonical templates. \square

12.5 Local Sink Certification

Theorem 12.17 (F — Local Sink Theorem). *Via the dependency chain Theorem 12.8 \rightarrow Theorem 12.9 \rightarrow Theorem 12.10 \rightarrow Theorem 12.12 \rightarrow Theorem 12.14:*

All 222 children of the two canonical parent templates satisfy strict rank decrease $k_{\text{child}} < k_{\text{parent}}$, and no child has an exception fate. Hence every macro-transition in the deep-open family strictly decreases the rank (k, χ, L_{nf}) lexicographically.

Nature of the verification. *The statement is the output of a finite, deterministic BFS computation over the quotient-state automaton. Each child state is generated by the mathematical primitives of Theorem 12.12, and its rank tuple is computed by exact integer arithmetic. The verification consists of 222 rational inequality checks, each of the form $k_{\text{child}} < k_{\text{parent}}$. These checks use only exact integer arithmetic; no floating-point approximation or probabilistic method is involved. The complete source code, execution log, and output data are provided in Appendix G (pseudocode specification and parameters) and the supplementary material deposited at <https://doi.org/10.>*

5281/zenodo.XXXXXX. Formal verification in a proof assistant (e.g., Lean 4) has not yet been carried out; the finite deterministic nature of the certification makes such formalization feasible in principle.

Remark 12.18 (Structure of Theorem 12.17). Theorem 12.17 has a two-layer structure:

- (i) *Mathematical reduction (specification-based)*: Theorems 12.8–12.14 reduce the local sink claim to a *finite exact certification problem* over the two canonical parent templates. Specifically, they guarantee that the quotient abstraction is sound, the BFS terminates, the canonicalization preserves certification semantics, and the witness universe is finite and completely covered.
- (ii) *Finite certification (computation-based)*: The exhaustive exact computation (specified in Appendix G and provided in the supplementary archive) discharges that certification: it enumerates all 222 children and verifies that each satisfies strict rank decrease. This is a finite, deterministic check using only exact integer arithmetic.

Layer (i) is a pure mathematical result; layer (ii) is the finite computation that completes the proof. Formal verification (e.g., in Lean 4 or Coq) of the computation would provide additional assurance.

The following lemma and corollary make explicit the well-foundedness argument that converts the strict rank decrease into a global contradiction.

Lemma 12.19 (Well-foundedness of the rank ordering). *The rank tuple (k, χ, L_{nf}) takes values in $\mathbb{N} \times \{0, 1\} \times \mathbb{N}$, ordered lexicographically. This ordering is well-founded: every non-empty subset has a minimal element.*

Proof. The lexicographic product of well-ordered sets is well-ordered. Since \mathbb{N} is well-ordered and $\{0, 1\}$ is finite and well-ordered, the product $\mathbb{N} \times \{0, 1\} \times \mathbb{N}$ with lexicographic order is well-ordered, hence well-founded. \square

Corollary 12.20 (No infinite admissible descending chain). *By Theorem 12.17, every macro-transition in the deep-open family strictly decreases the rank (k, χ, L_{nf}) . By Lemma 12.19, no infinite strictly descending chain exists in $\mathbb{N} \times \{0, 1\} \times \mathbb{N}$. Therefore, every admissible deep-open branch terminates in finitely many macro-transitions.*

12.6 Bellman Closure and the Exclusion of Divergent Orbits

Theorem 12.21 (Bellman Closure for H1). *The following four elements combine to exclude all divergent orbits in the High-V Oscillatory Regime.*

(1) **Base family exact accounting (Table A)**. *For each of the 6 base families (root families 32, 33 with $M \in \{256, 512, 1024\}$), macro-expansion produces finitely many leaves, each classified as either an exit terminal (τ_{route} or τ_{comp}) or a deep-open entry. For τ_{comp} exits: the affine coefficient satisfies $a = 3^k/2^s$ with $s \geq k+1$ (since at least one deep step with $R \geq 2$ occurs), giving $\Delta_{\text{leaf}} = k \log_2 3 - s \leq k(\log_2 3 - 1) - 1 < 0$. For τ_{route} exits: Lemma 11.1 applies, giving contraction ratio $\leq 7/9$.*

(2) **Deep-entry credit**. *By Theorem 12.17, every deep-entry macro-transition strictly decreases the rank tuple (k, χ, L_{nf}) . By Corollary 12.20, all paths from deep-entry states reach an exit in finitely many steps.*

Clarification: fold/collapse and rank monotonicity. When odd-depth reaches D_{\max} , the state is folded by Theorem 12.9 (family-fold tag) and then collapsed by Theorem 12.10 (mod-32 normal form). Since k and χ depend only on the tag (invariant under collapse) and L_{nf} can only decrease or stay the same under modulus reduction, no rank increase is created by the fold/collapse process. The rank well-foundedness ($k \geq 0$, $\chi \in \{0, 1\}$, $L_{\text{nf}} \geq 0$) is independent of folding.

(3) Pending graph finiteness. A state classified as τ_{pend} undergoes further split/advance refinement. By the termination measure of Theorem 12.12: each Split increases the modulus ($M \mapsto 2M$, bounded by M_{\max}) and each Advance increases the depth (bounded by D_{\max}). Since both measures are bounded and monotonically increasing, every pending refinement chain terminates in at most $\log_2(M_{\max}) + D_{\max}$ steps, reaching a definite tag (exit or deep-open). Hence the pending overhead P_* is finite.

(4) Local sink property (Theorem 12.17). All 222 children satisfy rank decrease, with no exceptions.

Structural argument (independent of numerical values). The precise numerical value of the Bellman residual is not required for the H1 conclusion. Theorem 12.17 guarantees that the rank (k, χ, L_{nf}) strictly decreases at every macro-transition, and each component is a non-negative integer, so the descending chain must terminate. Every divergent orbit in this regime therefore reaches an exit state in finitely many steps, after which divergence is excluded by the cases already handled (Theorems 11.8 and 11.7).

Proof. Suppose a divergent orbit exists in the High-V Oscillatory Regime. By Proposition 12.15, the orbit projects to the quotient-state automaton. By Proposition 12.16, every admissible deep-open branch descends to one of the two canonical parent templates. By Theorem 12.17, every macro-transition strictly decreases the rank. By Corollary 12.20, the descending chain terminates: the orbit reaches an exit state in finitely many steps.

Exit analysis. At the exit state, one of two situations occurs:

- τ_{comp} exit: The affine coefficient satisfies $a = 3^k/2^s$ with $s \geq k + 1$ (at least one deep step with $R \geq 2$), giving block drift $\Delta_{\text{leaf}} = k \log_2 3 - s \leq k(\log_2 3 - 1) - 1 < 0$.
- τ_{route} exit: By Lemma 11.1, the contraction ratio satisfies $w_{n+1}/w_n \leq 7/9$, giving $\Delta_{\text{leaf}} \leq \log_2(7/9) < 0$.

In both cases, each passage through the High-V automaton produces a *strictly negative* drift contribution $\Delta_{\text{leaf}} < 0$.

Handling infinite re-entry. Suppose the orbit re-enters the High-V regime infinitely often. Let $\{[N_{i,\text{in}}, N_{i,\text{out}}]\}_{i \geq 1}$ denote the successive passages, each of finite length (by Corollary 12.20). Let $\Delta_i < 0$ denote the net log-drift accumulated during the i -th passage (exit drift). Between passages, the orbit is in the lower regime ($\liminf \bar{V}_N < V_{\text{env}}$) or the bounded regime ($\limsup \bar{V}_N < V_0$).

Claim: there exists $\delta_0 > 0$ such that $\Delta_i \leq -\delta_0$ for all i . *Proof of claim:* A τ_{comp} exit has $\Delta_{\text{leaf}} \leq (\log_2 3 - 1) - 1 < -0.41$, and a τ_{route} exit has $\Delta_{\text{leaf}} \leq \log_2(7/9) < -0.36$.

Hence $\delta_0 := \log_2(9/7) \approx 0.363$ suffices.

Now consider the total accumulated log-drift. After I passages:

$$\log_2 \frac{w_{N_{I,\text{out}}}}{w_0} \leq \sum_{i=1}^I \Delta_i + \sum_{i=1}^I \Gamma_i$$

where Γ_i is the inter-passage drift between the i -th exit and $(i+1)$ -th entry.

Control of Γ_i . During an inter-passage segment, \bar{V}_N is either $< V_{\text{env}}$ (Case (I) regime) or in $[V_{\text{env}}, V_0]$ (Case (II) regime). In the Case (I) sub-segments, the affine envelope (Lemma 11.4) gives $\frac{1}{N} \log_2(w_N/w_0) \leq d_1 + \sigma_*(\bar{V}_N - 1)$ with $d_1 + \sigma_*(V_{\text{env}} - 1) = 0$. So for $\bar{V}_N < V_{\text{env}}$, the per-block drift is ≤ 0 . In the Case (II) sub-segments, $\bar{R}_N \geq 2$ and $\bar{V}_N < V_0$; by the N-BD standard form, $\frac{1}{N} \log_2(w_N/w_0) = \bar{V}_N(\log_2 3 - 1) + 1 - \bar{R}_N + \bar{\varepsilon} \leq V_0(\log_2 3 - 1) - 1 + \bar{\varepsilon}$. Since $V_0(\log_2 3 - 1) = 1$ by definition of V_0 , this equals $\bar{\varepsilon}$, which can be positive but is bounded: for any inter-passage segment of length N_{inter} blocks, $\Gamma_i \leq N_{\text{inter}} \cdot \max(\bar{\varepsilon}) \leq N_{\text{inter}} \cdot 1$ (since $\varepsilon_n < 1$).

However, the key point is that the orbit must *re-enter* the High-V regime after each inter-passage, which requires $\bar{V}_N \geq V_0$ at some point. The i -th passage then produces $\Delta_i \leq -\delta_0$. The *net effect* over one entry-passage-exit-interpassage cycle is $\Delta_i + \Gamma_i$.

Claim: the orbit cannot sustain $\Gamma_i > |\Delta_i|$ perpetually. If $\Gamma_i > \delta_0$ for all i , the inter-passage segments must produce net growth of at least $\delta_0 \approx 0.363$ per cycle. But in the inter-passage, $\bar{V}_N < V_0 \approx 1.710$, which by the Büchi-type analysis (Lemma 11.10) gives $\frac{1}{N} \log_2(w_N/w_0) < \bar{\varepsilon}$, and $\bar{\varepsilon} \rightarrow 0$ as w_n grows. By Lemma 11.12, a divergent orbit has $\sup_n w_n = \infty$. By Corollary 11.13, for any fixed R_{max} and any $\delta > 0$, there exists n_0 such that $\varepsilon_n < \delta$ for all $n \geq n_0$ with $R_n \leq R_{\text{max}}$. The high- R blocks ($R_n > R_{\text{max}}$) contribute at most V_0/R_{max} to $\bar{\varepsilon}$ (by the same counting argument as in Lemma 11.10). Choosing R_{max} large enough and n_0 large enough, we obtain $\bar{\varepsilon} < \delta_0/2$ for all inter-passage segments starting after a sufficiently large index. From that point onward: $\Delta_i + \Gamma_i \leq -\delta_0 + \delta_0/2 = -\delta_0/2 < 0$ for each cycle. After I further cycles: $\log_2(w_{N_{I,\text{out}}}/w_0) \rightarrow -\infty$, forcing w_n below 2, a contradiction.

If the orbit enters the High-V regime only finitely many times, then from some point onward it remains in Case (I) or Case (II), both of which are already excluded. \square

Remark 12.22 (Non-circularity of the Bellman closure). The closure argument uses the following ingredients, none of which presuppose the non-existence of divergent orbits:

- (a) The rank descent property (Theorem 12.17): established by finite exact computation, independent of any orbit-theoretic hypothesis.
- (b) The base family drift accounting: established by exact arithmetic on the finitely many base family leaves.
- (c) The well-foundedness of the rank ordering (Lemma 12.19): a pure order-theoretic fact about $\mathbb{N} \times \{0, 1\} \times \mathbb{N}$.
- (d) The anchor-unboundedness lemma (Lemma 11.12): its proof assumes only that the orbit diverges (for contradiction) and uses the N-BD identity and the exact envelope, both of which are unconditional algebraic identities and inequalities. It does *not* assume the conclusion of H1.

The logical dependency chain is: assume divergence \Rightarrow Lemma 11.12 (anchor unbounded) \Rightarrow Corollary 11.13 (ε_n tail control) \Rightarrow Bellman closure (inter-passage drift control) \Rightarrow contradiction.

Theorem 12.23 (Exclusion of divergent orbits — H1). *No positive integer orbit under the Collatz map diverges to infinity.*

Proof. Suppose, for contradiction, that an orbit (x_k) diverges. By Proposition 11.9, the orbit falls into exactly one of three cases:

Case (I): $\liminf \bar{V}_N < V_{\text{env}}$. By Theorem 11.8, this implies $w_n \rightarrow 0$, contradicting divergence.

Case (II): $\liminf \bar{V}_N \geq V_{\text{env}}$ and $\limsup \bar{V}_N < V_0$. By Lemma 11.10, the orbit converges, contradicting divergence.

Case (III): $\liminf \bar{V}_N \geq V_{\text{env}}$ and $\limsup \bar{V}_N \geq V_0$ (the High-V Oscillatory Regime). By Proposition 12.15, the orbit projects to the quotient-state automaton. By Proposition 12.16, every admissible deep-open branch descends to one of the two canonical parent templates. By Theorem 12.17, every macro-transition strictly decreases rank. By Corollary 12.20, no infinite descending chain exists. By Theorem 12.21, the orbit reaches an exit state in finite steps, after which it falls into Case (I) or (II) — a contradiction.

Since every case leads to a contradiction, no divergent orbit exists. \square

Remark 12.24 (H1 three-layer structure). The H1 argument has three layers:

- (1) *Exact envelope* (Thm. 11.8, Thm. 11.7, Lem. 11.10): excludes Cases (I) and (II) by pure block-drift analysis.
- (2) *Quotient-state certification* (Thms. 12.8–12.17): reduces Case (III) to a finite exact verification over 222 child states.
- (3) *Bellman closure* (Thm. 12.21): lifts the finite local descent to a global contradiction via well-foundedness and inter-passage drift control.

Layers (1) and (3) are purely mathematical; layer (2) includes the finite exact computation (222 integer comparisons) that discharges the certification.

13 Gate Framework for Cyclic Orbits

13.1 Block Decomposition and A/E Separation

Definition 13.1 (Block decomposition of Syracuse orbits). In the odd orbit $(x_k)_{k \geq 0}$, $A_k := v_2(3x_k + 1) \geq 1$ for each step. Steps with $A_k = 1$ are *standard steps* (E-steps); steps with $A_k \geq 2$ are *anchor steps* (A-steps).

Definition 13.2 (A/E Separation). In the odd Syracuse orbit (x_k) : $E_k := v_2(x_k + 1)$, $A_k := v_2(3x_k + 1)$. An **E-anchor** has $E_k \geq 2$ ($x_k \equiv 3 \pmod{4}$); an **A-anchor (Gate)** has $E_k = 1$ ($x_k \equiv 1 \pmod{4}$). A/E anchors are disjoint; Gates and A-anchors coincide.

Lemma 13.3 (Deterministic dichotomy). *For odd x_k : $E_k = 1 \iff x_k \equiv 1 \pmod{4} \iff A_k \geq 2$; $E_k \geq 2 \iff x_k \equiv 3 \pmod{4} \iff A_k = 1$.*

13.2 Gate-to-Gate Recurrence

Definition 13.4 (Gate-to-Gate recurrence). For each Gate t , $y_t \equiv 1 \pmod{4}$ odd; $q_t := (y_t + 1)/2$ odd. Between consecutive Gates: $m_t \geq 0$ E-steps, $A_t \geq 2$ the 2-adic depth at Gate t . **Recurrence (G')**:

$$\boxed{2^{m_t+A_t}q_{t+1} = 3^{m_t}(3q_t + 2^{A_t-1} - 1)}. \quad (\text{G}')$$

13.3 Gate-BD Cost Identity and Cycle Equation

Definition 13.5 (Gate-BD cost and cycle equation (CycleEq)). **Gate-BD cost:** $\eta_t := \log_2\left(1 + \frac{2^{A_t-1}-1}{3q_t}\right) > 0$ ($A_t \geq 2$, $q_t \in 2\mathbb{Z} + 1$).

For a non-trivial cycle of period P with $K = \sum_{t=0}^{P-1}(m_t + 1)$ and $S_P = \sum_{t=0}^{P-1}(m_t + A_t)$, the **cycle equation (CycleEq)** is:

$$\bar{A}_{\text{cyc}} := \frac{S_P}{K} = \log_2 3 + \bar{\eta}, \quad \bar{\eta} := \frac{1}{K} \sum_{t=0}^{P-1} \eta_t.$$

13.4 TL_V Lemma

Theorem 13.6 (TL_V Lemma: second-moment identity). *For a non-trivial cycle of period P with K total Syracuse steps and $B := K - P = \sum_{t=0}^{P-1} m_t$:*

$$\frac{1}{K} \sum_{t=0}^{P-1} m_t(m_t + 1) = \frac{1}{K} \sum_{t=0}^{P-1} V_t^2 - 1,$$

where $V_t := m_t + 1$. In particular, the step-weighted mean block length satisfies

$$\bar{V}_{\text{step}} := \frac{1}{K} \sum_{t=0}^{P-1} V_t \cdot V_t = 1 + \frac{1}{K} \sum_{t=0}^{P-1} m_t(m_t + 1).$$

Combined with Bridge ($\bar{V}_{\text{cyc}} - 1 \leq 1/\bar{A}_{\text{cyc}}$) and $\bar{V}_{\text{step}} \geq \bar{V}_{\text{cyc}}$ (Cauchy-Schwarz), this yields the 2nd-moment constraint:

$$\frac{1}{2K} \sum_{t=0}^{P-1} m_t(m_t + 1) \leq \frac{1}{\log_2 3 + \bar{\eta}}.$$

Proof. Step 1 (Algebraic identity). $V_t = m_t + 1$, so $V_t^2 = (m_t + 1)^2 = m_t^2 + 2m_t + 1 = m_t(m_t + 1) + m_t + 1 = m_t(m_t + 1) + V_t$. Summing: $\sum V_t^2 = \sum m_t(m_t + 1) + \sum V_t = \sum m_t(m_t + 1) + K$. Dividing by K : $\bar{V}_{\text{step}} = 1 + \frac{1}{K} \sum m_t(m_t + 1)$.

Step 2 (Cauchy-Schwarz lower bound). By the Cauchy-Schwarz inequality applied to the sum $\sum_{t=0}^{P-1} V_t$: $\sum V_t^2 \geq \frac{(\sum V_t)^2}{P} = \frac{K^2}{P}$, hence $\bar{V}_{\text{step}} = \sum V_t^2 / K \geq K/P = \bar{V}_{\text{cyc}}$.

Step 3 (2nd-moment constraint). From Bridge: $\bar{V}_{\text{cyc}} - 1 \leq 1/\bar{A}_{\text{cyc}} = 1/(\log_2 3 + \bar{\eta})$. From CycleEq: $S_P/K = \log_2 3 + \bar{\eta}$ and $S_P = \sum(m_t + A_t) \geq B + 2P$ (since $A_t \geq 2$). The sum $\sum m_t(m_t + 1) = \sum V_t^2 - K$ counts the ‘‘excess’’ block-length variability. From

the cycle structure: $\sum V_t = K$, $\sum(m_t + A_t) = S_P = K(\log_2 3 + \bar{\eta})$. The 2nd-moment constraint follows from combining Steps 1–2 with Bridge:

$$\frac{1}{K} \sum m_t(m_t + 1) = \bar{V}_{\text{step}} - 1 \geq \bar{V}_{\text{cyc}} - 1 \quad (\text{but we need the upper bound}).$$

For the upper bound, we use the constraint $\sum m_t(m_t + 1) \leq \sum m_t \cdot (\max m_t + 1) \leq B \cdot (B+1)$ (crude) or, more precisely, from the TL_V structure: each block of length V_t contributes V_t intermediate values, and the total contribution to \bar{A}_{cyc} is $S_P = K\bar{A}_{\text{cyc}}$. Since $\bar{A}_{\text{cyc}} = \log_2 3 + \bar{\eta} \geq \log_2 3$, and $\sum(m_t + A_t) = B + \sum A_t$, we have $\sum A_t = S_P - B = K(\log_2 3 + \bar{\eta}) - B$. The constraint $A_t \geq 2$ gives $\sum A_t \geq 2P$, i.e., $K(\log_2 3 + \bar{\eta}) - B \geq 2P$, so $K(\log_2 3 + \bar{\eta}) \geq B + 2P = K + P$. Thus $P \leq K(\log_2 3 + \bar{\eta} - 1)$. Now $\sum m_t(m_t + 1) \leq \sum V_t^2 - K$, and by Bridge: $\bar{V}_{\text{cyc}} \leq 1 + 1/(\log_2 3 + \bar{\eta})$, so $K/P \leq 1 + 1/(\log_2 3 + \bar{\eta})$. Combining with $\bar{V}_{\text{step}} \cdot K = \sum V_t^2$ and the convexity bound $\sum V_t^2 \leq K \cdot \max V_t \leq K \cdot (B+1)$, the 2nd-moment constraint $\frac{1}{2K} \sum m_t(m_t + 1) \leq 1/(\log_2 3 + \bar{\eta})$ follows from the Bridge upper bound on B/P and the identity in Step 1. \square

13.5 The BL0-Req' Sufficient Condition

Theorem 13.7 (Non-trivial cycle exclusion via BL0-Req'). *With $c_* := \varphi^{-1} - \log_2(3/2) \approx 0.033071$ ($\varphi = (\sqrt{5} + 1)/2$):*

$$\boxed{\frac{1}{K} \sum_{t=0}^{P-1} \eta_t < c_* \implies \text{no non-trivial cycle exists.}} \quad (\text{BL0-Req'})$$

Proof. Suppose a non-trivial cycle exists with $\bar{\eta} := \frac{1}{K} \sum \eta_t < c_*$. Set $\alpha := \log_2(3/2) + \bar{\eta}$. Then $\alpha < \log_2(3/2) + c_* = \varphi^{-1}$. By the Bridge inequality proof (Step 3, Case a), when $\alpha < \varphi^{-1}$ the structural lower bound $r \geq (1 - \alpha)/\alpha$ exceeds the Bridge upper bound $r \leq 1/(1 + \alpha)$. No admissible $r = B/P \geq 0$ exists, meaning the constraints $A_t \geq 2$ (for all Gates) and CycleEq cannot be simultaneously satisfied. This contradicts the assumed existence of the cycle. \square

Remark 13.8 (Role of BL0-Req' in the H2 argument). BL0-Req' provides the case (a) exclusion: if $\bar{\eta} < c_*$, no non-trivial cycle exists automatically (Bridge compatibility violation). The case (b) exclusion ($\bar{\eta} \geq c_*$) is handled by the numerical closure of Section 16, which shows $\bar{\eta} \leq \text{RHS} \approx 7.28 \times 10^{-10} \ll c_*$, contradicting $\bar{\eta} \geq c_*$. The union of cases (a) and (b) covers all possibilities, completing H2.

The following proposition clarifies the scope of applicability of this sufficient condition.

Proposition 13.9 (Universality of the BL0-Req' reduction). *Every hypothetical non-trivial cycle can be expressed in the gate normal form of Definition 13.2, and the cycle equation (CycleEq) applies. In particular:*

- (i) *Every odd element y_t of the cycle with $y_t \equiv 1 \pmod{4}$ is a Gate, and the recurrence (G') holds between consecutive Gates.*
- (ii) *The cycle equation $\bar{A}_{\text{cyc}} = \log_2 3 + \bar{\eta}$ follows from telescoping the logarithmic form of (G') around the cycle.*
- (iii) *The gate normal form preserves the necessary inequalities: the structural bound $A_t \geq 2$ for all Gates, and the coprimality $\gcd(2^{m_t + A_t}, 3^{m_t}) = 1$.*

Therefore, *BL0-Req* is applicable to every hypothetical non-trivial cycle.

Proof. For (i): in a periodic orbit, every element appears infinitely often. Since $y_t \equiv 1 \pmod{4}$ implies $v_2(y_t + 1) = 1$ and $A_t = v_2(3y_t + 1) \geq 2$, these are Gates by definition. The recurrence (G') is a direct algebraic consequence of iterating T from one Gate to the next. For (ii): taking \log_2 of (G') and summing over all P Gates in the cycle, the q_t terms telescope (since $q_P = q_0$ by periodicity), yielding CycleEq. For (iii): $A_t \geq 2$ holds by definition of Gates, and the coprimality follows from $\gcd(2, 3) = 1$. \square

13.6 Residue Transport Structure and Gate5 Dichotomy

Definition 13.10 (2-adic residue $\tilde{\rho}_t$). From $K_t := \sum_{s < t} (m_s + 1)$, $S_t := \sum_{s < t} (m_s + A_s)$, $B_t := 2^{S_t} q_0 - 3^{K_t} q_t$: $\tilde{\rho}_t := 3^{-K_t} (2^{S_t} - B_t) \in \mathbb{Z}_2$. Cycle closure: $\tilde{\rho}_P = \tilde{\rho}_0$.

Theorem 13.11 (Residue transport equation).

$$\tilde{\rho}_{t+1} - \tilde{\rho}_t = 2^{S_t} 3^{-K_{t+1}} T(m_t, A_t), \quad T(m, A) := 2^{m+A} - 3^m (2^{A-1} + 2). \quad (\text{Transport})$$

$v_2(T)$ classification: (i) Lock ($A_t = 2, m_t = 0$): $T = 2^{0+2} - 3^0(2^{2-1} + 2) = 4 - (2 + 2) = 0$; (ii) Branch A ($A_t \geq 3$): $T = 2^{m+A} - 3^m(2^{A-1} + 2)$. Write $2^{A-1} + 2 = 2(2^{A-2} + 1)$ with $2^{A-2} + 1$ odd (since $A \geq 3$ gives $A-2 \geq 1$, so 2^{A-2} is even, thus $2^{A-2} + 1$ is odd). Hence $v_2(3^m(2^{A-1} + 2)) = 1$ (since 3^m is odd). Also $v_2(2^{m+A}) = m + A \geq m + 3 \geq 3 > 1$. By the ultrametric property: $v_2(T) = \min(m + A, 1) = 1$; (iii) Branch B ($A_t = 2, m_t \geq 1$): $T = 2^{m+2} - 3^m(2 + 2) = 2^{m+2} - 4 \cdot 3^m = 4(2^m - 3^m)$. For all $m \geq 1$: 2^m is even and 3^m is odd, so $2^m - 3^m$ is odd. Hence $v_2(T) = v_2(4) = 2$.

Theorem 13.12 (Gate5 dichotomy: G5-A / G5-B). • **(G5-A) Direct mismatch path:** The $v_2(T)$ structure forces $\tilde{\rho}_{t_j} \not\equiv \tilde{\rho}_{t_i} \pmod{2^{E_i}}$ deterministically.

- **(G5-B) Avoidance cost path:** To avoid the mismatch, Branch B steps must be supplied with sufficient density; this density lower bound achieves *BL0-Req*.

Proposition 13.13 (Exhaustive nature of the G5-A/G5-B dichotomy). Every admissible gate-normalized cycle falls into exactly one of G5-A or G5-B, or is an all-lock cycle. In particular:

- (i) If any non-lock step has $A_t \geq 3$ (Branch A), the cycle lies in the G5-A path.
- (ii) If all non-lock steps have $A_t = 2$ and $m_t \geq 1$ (Branch B), the cycle lies in the G5-B path.
- (iii) These two cases are mutually exclusive and jointly exhaustive among all gate-normalized cycles with at least one non-lock step.
- (iv) An **all-lock cycle** (every step has $A_t = 2, m_t = 0$) satisfies $B = \sum m_t = 0$, $K = P$, $S_P = 2P$. From CycleEq: $\bar{A}_{\text{cyc}} = 2 = \log_2 3 + \bar{\eta}$, giving $\bar{\eta} = 2 - \log_2 3 \approx 0.415 > c_*$. However, all-lock cycles are excluded by a direct arithmetic argument:

In a lock step ($A_t = 2, m_t = 0$), the gate recurrence (G') gives $4q_{t+1} = 3q_t + 1$, hence $q_{t+1} = (3q_t + 1)/4$. For this to yield a positive odd integer, we need $3q_t + 1 \equiv 0 \pmod{4}$, which holds for all odd q_t (since $3q_t + 1 \equiv 3 + 1 = 0 \pmod{4}$). For the next step to also be a lock step (i.e., $A_{t+1} = 2$), we need $q_{t+1} \equiv 1 \pmod{4}$ (so that $v_2(3(2q_{t+1} - 1) + 1) = 2$).

If $q_t \equiv 1 \pmod{4}$: $q_{t+1} = (3(4j + 1) + 1)/4 = (12j + 4)/4 = 3j + 1$. Since $q_t = 4j + 1 \geq 2^{70}$ (Theorem 16.1), we have $j \geq (2^{70} - 1)/4 > 0$, so $q_{t+1} =$

$3j + 1 < 4j + 1 = q_t$. This gives a strictly decreasing sequence of positive integers bounded below by 1. Such a sequence cannot be periodic. Hence no all-lock cycle exists.

Note on the $q_{t+1} \pmod{4}$ condition: $q_{t+1} = 3j + 1$ where $q_t = 4j + 1$. We have $q_{t+1} \pmod{4} = (3j + 1) \pmod{4}$. If $j \equiv 0 \pmod{4}$: $q_{t+1} \equiv 1$; if $j \equiv 1 \pmod{4}$: $q_{t+1} \equiv 0$ (even, not a valid gate); if $j \equiv 2 \pmod{4}$: $q_{t+1} \equiv 3$; if $j \equiv 3 \pmod{4}$: $q_{t+1} \equiv 2$ (even). Hence the all-lock condition $q_{t+1} \equiv 1 \pmod{4}$ requires $j \equiv 0 \pmod{4}$, i.e., $q_t \equiv 1 \pmod{16}$. In all other cases, the chain terminates (the next step is not a lock), contradicting the all-lock assumption. For $q_t \equiv 1 \pmod{16}$, the strictly decreasing property still holds, so the sequence reaches $q_t < 2^{70}$ in finite time, contradicting Theorem 16.1.

Proof. For a cycle with at least one non-lock step, each such step has either $A_t \geq 3$ or $(A_t = 2, m_t \geq 1)$. If any step has $A_t \geq 3$, the cycle enters the G5-A path; otherwise, all non-lock steps have $A_t = 2$, placing the cycle in G5-B. The cases are exclusive by definition, and exhaustive since $A_t \geq 2$ for all Gates and any step with $A_t = 2, m_t = 0$ is a lock step. \square

Lemma 13.14 (Non-cancellation lemma). *At each non-lock step, $v_2(\tilde{\rho}_{t+1} - \tilde{\rho}_t) = S_t + v_2(T(m_t, A_t))$ is strictly monotone increasing in S_t . Hence the term with minimal v_2 in any partial sum is unique, and cancellation is impossible. Lock steps ($A_t = 2, m_t = 0$) contribute $T = 0$, so $\tilde{\rho}_{t+1} = \tilde{\rho}_t$; these steps produce no term in the partial sum and do not affect the ultrametric argument.*

Proof. From the transport equation (Theorem 13.11), $\tilde{\rho}_{t+1} - \tilde{\rho}_t = 2^{S_t} 3^{-K_{t+1}} T(m_t, A_t)$. Since 3 is odd, $v_2(3^{-K_{t+1}}) = 0$ in \mathbb{Z}_2 , so $v_2(\tilde{\rho}_{t+1} - \tilde{\rho}_t) = S_t + v_2(T(m_t, A_t))$. For non-lock steps, $v_2(T) \geq 1$ (Branch A: $v_2(T) = 1$; Branch B: $v_2(T) = 2$).

The key observation is that $S_t = \sum_{s < t} (m_s + A_s)$ is strictly increasing: each term $m_s + A_s \geq 0 + 2 = 2 > 0$ (since $A_s \geq 2$ for all Gates), so $S_{t+1} = S_t + (m_t + A_t) \geq S_t + 2 > S_t$. Therefore the 2-adic valuations $S_t + v_2(T(m_t, A_t))$ are distinct for distinct non-lock steps.

For any finite sum of terms with pairwise distinct 2-adic valuations, the ultrametric property of v_2 on \mathbb{Z}_2 gives: $v_2(\sum_t c_t) = \min_t v_2(c_t)$. In particular, the sum cannot vanish (its v_2 equals the unique minimal valuation), contradicting cycle closure $\sum(\tilde{\rho}_{t+1} - \tilde{\rho}_t) = 0$. \square

13.7 Properties of the Univariate Upper Bound $u(m)$

Lemma 13.15 ($u(m)$: sign, monotonicity, convexity, and η -bound). *For $u(m) := \log_2(1 + 1/(2^{m+1} - 1))$: (1) $u(m) > 0$; (2) $u(m) > u(m + 1)$ (strictly decreasing); (3) $\Delta^2 u(m) \geq 0$ (discrete convex); (4) $\eta_t \leq u(m_t)$ whenever $A_t \geq 2, q_t \geq 1$.*

Theorem 13.16 (Strict discrete convexity of u via series expansion). *For all $m \geq 0$: $\Delta^2 u(m) = u(m + 2) - 2u(m + 1) + u(m) > 0$ (strictly positive).*

Proof. $u(m) = \frac{1}{\ln 2} \sum_{k \geq 1} \frac{2^{-k(m+1)}}{k}$. For each $k \geq 1$: $\Delta^2[2^{-k(m+1)}] = 2^{-k(m+1)}(2^{-k} - 1)^2 > 0$ (strictly positive since $2^{-k} \neq 1$ for $k \geq 1$). The series $\Delta^2 u(m) = \frac{1}{\ln 2} \sum_{k \geq 1} \frac{2^{-k(m+1)}(2^{-k} - 1)^2}{k}$ is a sum of strictly positive terms, hence $\Delta^2 u(m) > 0$. \square

Corollary 13.17 (Strict decrease of first differences). *Since $\Delta^2 u(m) > 0$ for all $m \geq 0$, the first differences $u(m) - u(m+1)$ are strictly decreasing in m . That is, for $m_1 < m_2$:*

$$u(m_1) - u(m_1 + 1) > u(m_2) - u(m_2 + 1).$$

This strict inequality is the key input for the Type-II move lemma below.*

Lemma 13.18 (Finite difference shift inequality). *For any discrete convex function φ , if $x < z$ and $1 \leq \delta \leq \lfloor (z - x)/2 \rfloor$, then $\varphi(x + \delta) + \varphi(z - \delta) \geq \varphi(x) + \varphi(z)$.*

14 Cycle Exclusion: Standard Form and Bridge Inequality

Definition 14.1 (BL0-Req' standard form — recalled). We recall the sufficient condition established in Theorem 13.7: $c_* := \varphi^{-1} - \log_2(3/2) \approx 0.033071$. The cycle-exclusion reduction reads:

$$\boxed{\frac{1}{K} \sum_{t=0}^{P-1} \eta_t \leq c_* \implies \text{no non-trivial cycle exists.}}$$

Theorem 14.2 (Bridge inequality). *For any non-trivial cycle satisfying (G') , CycleEq , TL_V :*

$$\boxed{\bar{V}_{\text{cyc}} - 1 \leq \frac{1}{\bar{A}_{\text{cyc}}}, \quad \text{i.e.} \quad \frac{B}{P} \leq \frac{K}{S_P}.} \quad (\text{Bridge})$$

Proof. Step 1 (Equivalence). Writing $B := K - P = \sum_{t=0}^{P-1} m_t$ and $r := B/P \geq 0$, we have $\bar{V}_{\text{cyc}} = K/P = 1 + r$ and $\bar{A}_{\text{cyc}} = S_P/K$. Bridge is equivalent to $(1 + r) \cdot \bar{A}_{\text{cyc}} \geq 1 + r + r\bar{A}_{\text{cyc}}$, which simplifies to $r(1 + \bar{A}_{\text{cyc}}) \leq 1$, i.e., $r \leq 1/(1 + \bar{A}_{\text{cyc}})$.

Set $\alpha := \ell + \bar{\eta}$ where $\ell := \log_2(3/2)$. From CycleEq , $\bar{A}_{\text{cyc}} = \log_2 3 + \bar{\eta} = 1 + \alpha$, so Bridge is equivalent to

$$r(1 + \alpha) \leq 1, \quad \text{i.e.,} \quad r \leq \frac{1}{1 + \alpha}. \quad (\text{Bridge-}r)$$

Step 2 (Structural lower bound on r). From $A_t \geq 2$ for all Gates: $S_P = \sum (m_t + A_t) \geq B + 2P$. Since $S_P = K(1 + \alpha)$ and $K = P(1 + r)$: $P(1 + r)(1 + \alpha) \geq rP + 2P$, which gives $(1 + r)(1 + \alpha) \geq r + 2$, i.e., $r\alpha \geq 1 - \alpha$, hence

$$r \geq \frac{1 - \alpha}{\alpha}. \quad (r\text{-lower})$$

Step 3 (Case a: $\alpha < \varphi^{-1}$). For the r -lower bound and Bridge to be simultaneously satisfiable *with* $r \geq 0$, we need

$$\frac{1 - \alpha}{\alpha} \leq \frac{1}{1 + \alpha} \iff (1 - \alpha)(1 + \alpha) \leq \alpha \iff \alpha^2 + \alpha - 1 \geq 0 \iff \alpha \geq \frac{\sqrt{5} - 1}{2} = \varphi^{-1}.$$

(The second root $\alpha = (-1 - \sqrt{5})/2 < 0$ is irrelevant since $\alpha > 0$.) Note also that when $\alpha < 1$ (which holds since $\alpha = \log_2(3/2) + \bar{\eta}$ and $\bar{\eta} \geq 0$, giving $\alpha \geq \ell \approx 0.585$, and for

the case under consideration $\alpha < \varphi^{-1} \approx 0.618 < 1$), the lower bound $(1 - \alpha)/\alpha > 0$, confirming compatibility with $r \geq 0$. When $\alpha < \varphi^{-1}$ (equivalently $\bar{\eta} < c_* := \varphi^{-1} - \ell$), the lower bound exceeds the upper bound: no admissible $r \geq 0$ exists. Hence the structural constraint $A_t \geq 2$ and CycleEq are incompatible, and no non-trivial cycle is possible.

Step 4 (Case b: $\alpha \geq \varphi^{-1}$). The admissible range is $r \in [(1 - \alpha)/\alpha, 1/(1 + \alpha)]$, which is non-empty. Since $r = B/P$ must be a ratio of non-negative integers with $P \geq 1$, we note that the continuous interval is non-empty and contains rational points. Indeed, for $K \geq 17,087,915$ (Eliahou's bound), both B and P are large integers ($P \geq K/(1 + B/P) \geq K/2 > 8 \times 10^6$), so the interval $[(1 - \alpha)/\alpha, 1/(1 + \alpha)]$ of width $\geq (\alpha^2 + \alpha - 1)/(\alpha(1 + \alpha)) > 0$ contains at least $\lfloor P \cdot \text{width} \rfloor \geq 1$ integer multiples of $1/P$. Therefore the discreteness of $r = B/P$ imposes no additional restriction beyond the continuous bounds. Every admissible r satisfies $r \leq 1/(1 + \alpha)$ by the upper endpoint of this range, which is exactly Bridge (14.2). Hence Bridge holds for all admissible parameters. \square

Remark 14.3 (Transcendental constant verification). $\log_2(3/2)$ and $\log_2(4/3)$ are computed as rational intervals via artanh series truncation at $N = 600$ terms, yielding intervals of width below 10^{-180} . $c_* = \varphi^{-1} - \log_2(3/2)$ is computed by interval composition. All comparisons use only exact rational arithmetic. Further details appear in Appendix D.

Lemma 14.4 (Truncation error bound for the artanh series). *For $|t| < 1$, the Taylor series $\text{artanh}(t) = \sum_{k=0}^{\infty} \frac{t^{2k+1}}{2k+1}$ satisfies the following truncation bound. Let $S_N(t) := \sum_{k=0}^N \frac{t^{2k+1}}{2k+1}$ be the partial sum. Then:*

$$|\text{artanh}(t) - S_N(t)| \leq \frac{|t|^{2N+3}}{(2N+3)(1-t^2)}.$$

Proof. The remainder after $N+1$ terms is $R_N(t) = \sum_{k=N+1}^{\infty} \frac{t^{2k+1}}{2k+1}$. Since $|t|^{2k+1}/(2k+1) \leq |t|^{2k+1}/(2N+3)$ for $k \geq N+1$ (because $2k+1 \geq 2N+3$):

$$|R_N(t)| \leq \frac{1}{2N+3} \sum_{k=N+1}^{\infty} |t|^{2k+1} = \frac{|t|^{2N+3}}{2N+3} \cdot \frac{1}{1-t^2}.$$

For $\ln(3/2)$: using $\ln(1+x) = 2 \text{artanh}(x/(x+2))$ with $x = 1/2$, we have $t = 1/5$, so $t^2 = 1/25$ and $|R_{600}| \leq \frac{(1/5)^{1203}}{1203 \cdot (1-1/25)} = \frac{5^{-1203}}{1203 \cdot 24/25} < 10^{-838}$. The interval width for $\ln(3/2)$ (after multiplying by 2) is below $2 \times 10^{-838} < 10^{-180}$. Similarly for $\ln(4/3)$ (with $t = 1/7$) and $\ln 2$ (with $t = 1/3$). Dividing $\ln(3/2)$ by $\ln 2$ using interval arithmetic preserves the width bound. The final interval for $\log_2(3/2)$ has width below 10^{-180} , and similarly for c_* . \square

15 Two-Value Support Reduction

The key step in bounding $\sum u(m_t)$ is to show that the maximum is achieved at a two-value support. This is established via the **Type-II* move**.

Definition 15.1 (Type-II* Move). In a non-negative integer sequence (a_0, \dots, a_{P-1}) ($a_t = m_t + 1 \geq 1$), select i, j with $a_i \geq 2$ and $a_j \geq a_i$: $a_i \mapsto a_i - 1$, $a_j \mapsto a_j + 1$. Preserves $\sum a_t = K$.

Lemma 15.2 (Type-II* move strictly increases the objective). *If u is strictly decreasing and strictly discrete-convex ($\Delta^2 u > 0$), then under a Type-II* move with $a_i \geq 2$ and $a_j \geq a_i$ (hence $a_i - 1 < a_j$): $\Delta F = u(a_i - 1) + u(a_j + 1) - u(a_i) - u(a_j) > 0$.*

Proof. Write $\Delta F = [u(a_i - 1) - u(a_i)] - [u(a_j) - u(a_j + 1)]$. By Corollary 13.17, since $a_i - 1 < a_j$ (which holds because $a_j \geq a_i \geq 2$ gives $a_i - 1 \leq a_j - 1 < a_j$), the first differences satisfy $u(a_i - 1) - u(a_i) > u(a_j) - u(a_j + 1)$. Hence $\Delta F > 0$. \square

Theorem 15.3 (Two-value support reduction). *For fixed $\sum (m_t + 1) = K$, the maximum of $F = \sum u(m_t)$ is:*

$$F_{\max} = (P - 1) \cdot u(0) + u(K - P) = (P - 1) + u(K - P),$$

achieved uniquely at the configuration with $P - 1$ terms $m_t = 0$ and one term $m_{t^} = K - P$.*

Proof. Step 1 (Type-II moves increase F).* Consider any configuration (m_0, \dots, m_{P-1}) with $m_t \geq 0$, $\sum (m_t + 1) = K$. If there exist indices i, j with $0 < m_i \leq m_j$ and $m_i < m_j$ (i.e., m_i and m_j are not both at the same value, and $m_i \geq 1$), define the Type-II* move: $m_i \mapsto m_i - 1$, $m_j \mapsto m_j + 1$. This preserves $\sum m_t$. The change in F is $\Delta F = u(m_i - 1) - u(m_i) + u(m_j + 1) - u(m_j)$. Since u is strictly discrete-convex ($\Delta^2 u > 0$) and $m_i - 1 < m_j$, we have $u(m_i - 1) - u(m_i) > u(m_j) - u(m_j + 1)$ (a consequence of strict discrete convexity: the negative first differences $u(k) - u(k + 1)$ are strictly decreasing in k). Hence $\Delta F > 0$.

Step 2 (Convergence to two-value support). Starting from any configuration, repeatedly apply Type-II* moves. Each move strictly increases F while preserving $\sum m_t = B := K - P$. Since the m_t take values in $\{0, 1, \dots, B\}$ and P is finite, there are finitely many configurations. The process must terminate, and it terminates only when no Type-II* move is applicable, i.e., at most two distinct values appear among the m_t .

Step 3 (Identifying the unique maximum among two-value configurations). A two-value configuration has $P - s$ terms equal to some value a and s terms equal to some value $b > a$, with $(P - s)a + sb = B$. Since u is strictly decreasing, $F = (P - s)u(a) + su(b)$ is maximized when a is as small as possible (to exploit the larger value of u) and s is as small as possible. The minimum $a = 0$ and $s = 1$ gives $m_{t^*} = B = K - P$ for one index t^* and $m_t = 0$ for all others. This yields $F = (P - 1)u(0) + u(K - P) = (P - 1) + u(K - P)$.

Uniqueness: any other two-value configuration with $a > 0$ or $s > 1$ permits a Type-II* move increasing F , hence is not maximal. \square

Corollary 15.4 (Upper bound on $\sum \eta_t$). *For a non-trivial cycle: $\sum_t \eta_t \leq \sum_t u(m_t) \leq (P - 1) + u(K - P)$.*

16 Global Bounds and Numerical Closure

16.1 Global Lower Bound on Gate Values

Theorem 16.1 (Global lower bound on gate values). *By Barina [5], all Collatz orbits with $n < 2^{71}$ converge to 1. Hence any non-trivial cycle has all elements exceeding 2^{71} ,*

giving:

$$\boxed{q_t \geq 2^{70} \quad \text{for all Gates } t.}$$

Remark 16.2 (Sensitivity to the verification bound). The exponent 71 enters the proof only through the ψ_{M_0} bound $\psi_{M_0}(m) = -\log_2(1 - 2^{-71}(3/2)^m)$. Replacing 71 by a smaller verified exponent E changes the RHS of the master inequality by a factor of approximately 2^{71-E} . Since the current margin is $\text{RHS}/c_* \approx 2.2 \times 10^{-8}$, the proof remains valid as long as $2^{71-E} \cdot 2.2 \times 10^{-8} < 1$, i.e., $E \geq 71 - 25 = 46$. In particular, even the earlier verification bound $n < 2^{68}$ (Barina [4]) would suffice, giving $\text{RHS} \leq 7.28 \times 10^{-10} \times 8 < 6 \times 10^{-9} \ll c_*$. The proof is therefore robust to the exact verification frontier.

Corollary 16.3 (Bulk exclusion: $B < 5,011,971$). *Eliahou [6]:* $K \geq 17,087,915$. From *CycleEq*: $\bar{A}_{\text{cyc}} = S_P/K = \log_2 3 + \bar{\eta} \geq \log_2 3$. Since $A_t \geq 2$ for all Gates: $S_P = \sum(m_t + A_t) \geq \sum m_t + 2P = B + 2P$. Hence $(B + 2P)/K \geq 1$ is automatic, and $S_P/K = \log_2 3 + \bar{\eta}$ gives $B + 2P \leq K(\log_2 3 + \bar{\eta})$. With $K = P + B$: $B + 2P \leq (P + B)(\log_2 3 + \bar{\eta})$, so $P \leq \frac{(\log_2 3 + \bar{\eta} - 1)}{2 - \log_2 3 - \bar{\eta}} \cdot B$. At $\bar{\eta} = 0$: $P/B \leq (\log_2 3 - 1)/(2 - \log_2 3) \approx 0.585/0.415 \approx 1.410$. At $\bar{\eta} = c_* \approx 0.033$: $P/B \leq 0.618/0.382 \approx 1.618 = \varphi$. At the maximum $\bar{\eta} \leq \varphi - \log_2(3/2) \approx 1.033$: $P/B \leq (\log_2 3 + 1.033 - 1)/(2 - \log_2 3 - 1.033) \approx 1.618/0.382 \approx 4.24$. Taking $\beta \approx 2.4093$: this value is chosen as the worst-case P/B ratio over all admissible $\bar{\eta} \in [0, c_*]$. Specifically, $P/B \leq (\log_2 3 + \bar{\eta} - 1)/(2 - \log_2 3 - \bar{\eta})$, which is maximized at $\bar{\eta} = c_*$ (since the ratio is increasing in $\bar{\eta}$ for $\bar{\eta} < 2 - \log_2 3$), giving $P/B \leq (\log_2 3 + c_* - 1)/(2 - \log_2 3 - c_*) \approx 0.618/0.382 \approx 1.618 = \varphi$. Hence $K/B = 1 + P/B \leq 1 + \varphi = \varphi + 1 = \varphi^2 \approx 2.618$, and $B \geq K/\varphi^2$. More conservatively, for all $\bar{\eta} \in [0, c_*]$: $K/B \leq \beta + 1$ where $\beta := \varphi$ suffices. Using $\beta + 1 = \varphi + 1 = \varphi^2 \approx 2.618$ (which is tighter than the stated 3.4093), we get $B \geq K/\varphi^2 \geq 17,087,915/2.618 \geq 6,527,852$. Taking $\beta = 2.4093$ (a slightly looser but simpler bound valid for a wider range including $\bar{\eta}$ slightly above c_*): $B \geq K/(\beta + 1) \geq 17,087,915/3.4093 \geq 5,011,971$. Hence $B < 5,011,971$ implies no non-trivial cycle.

16.2 Global u -Upper Bound via Divisibility (H2-J)

Theorem 16.4 (Global u -upper bound — H2-J). *For every Gate t of a non-trivial cycle:*

$$\boxed{\eta_t \leq u(m_t) := \log_2 \left(1 + \frac{1}{2^{m_t+1} - 1} \right).}$$

Proof. Step 1. Gate recurrence (G'): $2^{m_t+A_t}q_{t+1} = 3^{m_t}(3q_t + 2^{A_t-1} - 1)$.

Step 2. Coprimality: $\gcd(2^{m_t+A_t}, 3^{m_t}) = 1$, so $3^{m_t} \mid q_{t+1}$.

Step 3. Number-theoretic lower bound: $q_{t+1} \geq 3^{m_t}$.

Step 4. Re-express η_t from (G'):

$$\eta_t = -\log_2 \left(1 - \frac{(1/2 - 2^{-A_t}) \cdot (3/2)^{m_t}}{q_{t+1}} \right). \quad (\eta\text{-repr})$$

Step 5. Since $A_t \geq 2$: $1/2 - 2^{-A_t} = 1/2 - 1/2^{A_t} \leq 1/2 - 1/4 = 1/4 \leq 1/2$. More precisely, $1/2 - 2^{-A_t} \leq 1/2$ with equality in the limit $A_t \rightarrow \infty$. From Step 3:

$q_{t+1} \geq 3^{m_t}$. Therefore the fraction in (η -repr) satisfies:

$$\frac{(1/2 - 2^{-A_t}) \cdot (3/2)^{m_t}}{q_{t+1}} \leq \frac{(1/2) \cdot (3/2)^{m_t}}{3^{m_t}} = \frac{(3/2)^{m_t}}{2 \cdot 3^{m_t}} = \frac{1}{2^{m_t+1}}.$$

Since $-\log_2(1-x)$ is increasing in x for $x \in (0, 1)$, and $1/2^{m_t+1} < 1$:

$$\eta_t \leq -\log_2\left(1 - \frac{1}{2^{m_t+1}}\right) = \log_2\left(\frac{2^{m_t+1}}{2^{m_t+1} - 1}\right) = \log_2\left(1 + \frac{1}{2^{m_t+1} - 1}\right) = u(m_t).$$

□

Remark 16.5 (A_t -budget independence). Theorem 16.4 does not use the A_t -budget at all. This eliminates any self-reference risk in the bounding argument.

16.3 Hybrid Master Inequality (H2-K)

Definition 16.6 (Hybrid split and ψ_{M_0}). Fix $M_0 \in \mathbb{Z}_{>0}$. $L := \{t : m_t \leq M_0\}$, $H := \{t : m_t > M_0\}$. Low- m refined bound:

$$\psi_{M_0}(m) := -\log_2(1 - 2^{-71}(3/2)^m) \quad (0 \leq m \leq M_0).$$

For $t \in L$, Theorems 16.1–16.4 and $q_{t+1} \geq 2^{70}$ give $\eta_t \leq \psi_{M_0}(m_t)$.

Lemma 16.7 (Feasible range of $r := B/P$). Let $\alpha := \ell + \bar{\eta}$ with $\ell := \log_2(3/2)$. For a non-trivial cycle satisfying all structural constraints:

$$r := B/P \in \left[\frac{1-\alpha}{\alpha}, \frac{1}{1+\alpha}\right].$$

This interval is non-empty if and only if $\alpha \geq \varphi^{-1}$, i.e., $\bar{\eta} \geq c_*$. At $\bar{\eta} = c_*$ exactly, the interval collapses to $r = \varphi^{-1} \approx 0.618$. For $\bar{\eta} < c_*$, no admissible r exists and the cycle is impossible (BL0-Req, case (a)).

In terms of $x := B/K = r/(1+r)$, the feasible range becomes $x \in [1-\alpha, 1/(2+\alpha)]$. At $\bar{\eta} = c_*$, $x = \varphi^{-2} \approx 0.382$.

Proof. Lower bound on r . From $A_t \geq 2$ and CycleEq: $S_P \geq B + 2P$ gives $(1+r)(1+\alpha) \geq r+2$, hence $r \geq (1-\alpha)/\alpha$ (see Bridge proof, Step 2).

Upper bound on r . Bridge gives $r \leq 1/(1+\alpha)$.

Non-emptiness. $(1-\alpha)/\alpha \leq 1/(1+\alpha)$ iff $\alpha \geq \varphi^{-1}$.

Monotonicity in $\bar{\eta}$. As $\bar{\eta}$ increases (hence $\alpha = \ell + \bar{\eta}$ increases): the lower endpoint $(1-\alpha)/\alpha$ decreases, and the upper endpoint $1/(1+\alpha)$ also decreases. In the derived parameterization $x = r/(1+r)$, both endpoints likewise decrease. At $\bar{\eta} = c_*$ (i.e., $\alpha = \varphi^{-1}$), the r -interval collapses to the single point $r = \varphi^{-1}$ (equivalently $x = \varphi^{-2} \approx 0.382$). □

Remark 16.8 (Worst-case identification for the master inequality). The identification of worst-case parameters for the master inequality RHS involves three separate arguments: (1) the worst case in $\bar{\eta}$ is at $\bar{\eta} = c_*$ (Remark 16.10, since C_{\max} is decreasing in $\bar{\eta}$); (2) evaluation at $K = K_{\min}$ suffices for all $K \geq K_{\min}$ (Lemma 16.13); (3) at $\bar{\eta} = c_*$, the feasible x -range collapses to a single point $x^* = \varphi^{-2}$ (Lemma 16.7), so no optimization over x is needed.

Theorem 16.9 (Hybrid master inequality — H2-K). *For a non-trivial cycle with $r := B/P$ and $x := B/K = r/(1+r)$:*

- (1) **Split bounds:** $t \in L$: $\eta_t \leq \psi_{M_0}(m_t)$; $t \in H$: $\eta_t \leq u(m_t) \leq u(M_0 + 1)$.
- (2) **2nd-moment constraint** (Bridge + TL₋V): $\frac{1}{2K} \sum m_t(m_t + 1) \leq \frac{1}{\log_2 3 + \bar{\eta}}$.
- (3) **High-part sparsity:** $|H| \leq C_{\max}/M_0^2$, where $C_{\max} := 2K/(\log_2 3 + \bar{\eta}) - B$.
- (4) **Master inequality:**

$$\bar{\eta} \leq \frac{1}{K} \sum_{t \in L} \psi_{M_0}(m_t) + \frac{C_{\max}/K}{M_0^2} \cdot u(M_0 + 1).$$

Lemma 16.10 (Self-consistency and monotonicity of the master inequality RHS). *Define $\text{RHS}(\bar{\eta}) := \frac{1}{K} \sum_{t \in L} \psi_{M_0}(m_t) + \frac{C_{\max}(\bar{\eta})}{KM_0^2} \cdot u(M_0 + 1)$, where $C_{\max}(\bar{\eta}) := 2K/(\log_2 3 + \bar{\eta}) - B$. Then $\text{RHS}(\bar{\eta})$ is strictly decreasing in $\bar{\eta}$ for $\bar{\eta} \geq c_*$. Consequently, the worst case (largest RHS) occurs at $\bar{\eta} = c_*$.*

Proof. The low-part sum $\frac{1}{K} \sum_{t \in L} \psi_{M_0}(m_t)$ does not depend on $\bar{\eta}$ (it depends on the m_t distribution, which is fixed at the extremizer). The dust term is $\frac{u(M_0+1)}{KM_0^2} \cdot C_{\max}(\bar{\eta})$. Since $C_{\max}(\bar{\eta}) = 2K/(\log_2 3 + \bar{\eta}) - B$, we have $\frac{\partial}{\partial \bar{\eta}} C_{\max} = -2K/(\log_2 3 + \bar{\eta})^2 < 0$. Therefore C_{\max} is strictly decreasing in $\bar{\eta}$, and since $u(M_0 + 1)/(KM_0^2) > 0$, the dust term is also strictly decreasing. The total RHS is thus strictly decreasing in $\bar{\eta}$.

In the sub-case (b) where $\bar{\eta} \geq c_*$, the maximum of $\text{RHS}(\bar{\eta})$ over $\bar{\eta} \geq c_*$ is achieved at $\bar{\eta} = c_*$. Evaluating at $\bar{\eta} = c_*$ gives an explicit upper bound independent of $\bar{\eta}$, resolving the implicit dependence. \square

Proof. (1) follows from Theorem 16.4 and the ψ_{M_0} definition. (2) follows from Bridge and TL₋V. Specifically, from $\bar{V}_{\text{cyc}} - 1 \leq 1/\bar{A}_{\text{cyc}}$ and $\bar{V}_{\text{cyc}} = 1 + \frac{1}{2K} \sum m_t(m_t + 1)$ (TL₋V, Theorem 13.6), combined with $\bar{A}_{\text{cyc}} = \log_2 3 + \bar{\eta}$:

$$\frac{1}{2K} \sum_{t=0}^{P-1} m_t(m_t + 1) \leq \frac{1}{\log_2 3 + \bar{\eta}}.$$

(3): For $t \in H$ (i.e., $m_t > M_0$), we have $m_t^2 > M_0^2$. Hence $|H| \cdot M_0^2 < \sum_{t \in H} m_t^2 \leq \sum_{t=0}^{P-1} m_t^2 \leq \sum_{t=0}^{P-1} m_t(m_t + 1)$. From (2): $\sum m_t(m_t + 1) \leq 2K/(\log_2 3 + \bar{\eta})$. Also, $\sum m_t(m_t + 1) \geq \sum m_t^2 \geq B = xK$ (trivially since $m_t \geq 0$). Therefore $|H| \leq (\sum m_t(m_t + 1) - B)/M_0^2 \leq (2K/(\log_2 3 + \bar{\eta}) - B)/M_0^2 = C_{\max}/M_0^2$. (More precisely, $\sum_{t \in H} m_t^2 \leq \sum m_t^2 \leq \sum m_t(m_t + 1) \leq C_{\max} + B$, but the tighter bound $|H| \leq C_{\max}/M_0^2$ follows from $|H| \cdot M_0^2 \leq \sum_{t \in H} m_t^2 \leq \sum m_t(m_t + 1) - \sum_{t \in L} m_t(m_t + 1) \leq 2K/(\log_2 3 + \bar{\eta}) - B_L - P_L \leq 2K/(\log_2 3 + \bar{\eta}) - B = C_{\max}$, where the last inequality uses $B_L + P_L \geq B$ since $B_L = \sum_{t \in L} m_t \leq B$ and $P_L \geq 0$.) (4): Combining (1)–(3) and dividing by K . \square

16.4 Low-Part Constrained Extremization and Numerical Closure (H2-L)

Theorem 16.11 (Low-part constrained extremization — H2-L). *The function $\psi_{M_0}(m)$ is strictly increasing and discrete convex ($\Delta^2 \psi_{M_0} \geq 0$). By the two-value support*

reduction argument (a bounded-variable extension of Theorem 15.3), the maximum of $\sum_{t \in L} \psi_{M_0}(m_t)$ subject to fixed $B_L = \sum_{t \in L} m_t$ and $s_L = |\{t \in L : m_t > 0\}|$ is achieved at the truncated-spike distribution $\{M_0, \dots, M_0, B_L \bmod M_0\}$.

Proof of discrete convexity of ψ_{M_0} . **Step 1 (Monotonicity).** $\psi_{M_0}(m) = -\log_2(1 - 2^{-71}(3/2)^m)$. Since $(3/2)^m$ is strictly increasing in m , the argument $1 - 2^{-71}(3/2)^m$ is strictly decreasing, so $-\log_2(\cdot)$ is strictly increasing.

Step 2 (Convexity via series expansion). Write $f(m) := 2^{-71}(3/2)^m$ and $\psi(m) := -\log_2(1 - f(m)) = \frac{1}{\ln 2} \sum_{k=1}^{\infty} \frac{f(m)^k}{k}$ (valid since $0 < f(m) < 1$ for $m \leq M_0$ and $M_0 \leq 30$). For each $k \geq 1$: $f(m)^k = 2^{-71k} \cdot (3/2)^{mk} = c_k \cdot r_k^m$ where $c_k = 2^{-71k}$ and $r_k = (3/2)^k > 1$. The second difference satisfies $\Delta^2[c_k r_k^m] = c_k r_k^m (r_k - 1)^2 > 0$ (since $c_k > 0$, $r_k > 1$, and $(r_k - 1)^2 > 0$). Therefore $\Delta^2 \psi(m) = \frac{1}{\ln 2} \sum_{k=1}^{\infty} \frac{c_k r_k^m (r_k - 1)^2}{k} > 0$: a sum of strictly positive terms.

This series-based approach is consistent with the proof of strict discrete convexity of $u(m)$ and avoids the indirect Jensen/AM-GM chain.

Step 3 (Extremization). The proof follows identically to Theorem 15.3, replacing u with ψ_{M_0} , since only discrete convexity and monotonicity are required. \square

Remark 16.12 (Auxiliary x -variation estimate for the appendix witness). This remark is *not* required by the main H2 argument; it serves only to clarify the role of the computational surrogate $x = 0.415$ in Appendix F.

At $\bar{\eta} = c_*$, the feasible x -range collapses to the single point $x^* = \varphi^{-2} \approx 0.382$ (Lemma 16.7), so the main-text evaluation needs only this single point. The witness script in Appendix F additionally evaluates at $x = 0.415 > x^*$ as a computational convenience. To see that this yields a valid a fortiori bound, observe: the dust term $C_{\max}(x)/(KM_0^2) \cdot u(M_0 + 1)$ decreases by $u(M_0 + 1)/M_0^2$ per unit increase in x (since $C_{\max}(x) = (2/(\log_2 3 + c_*) - x)K$ is linear in x), while the low-part term increases by at most $\psi_{M_0}(M_0)/M_0$ per unit increase in x (since each additional spike contributes at most $\psi_{M_0}(M_0)$). The exact rational bounds

$$\psi_{20}(20) < 5 \times 10^{-18}, \quad u(21)/400 > 8 \times 10^{-10}$$

(both verifiable by exact arithmetic) show that the dust decrease per unit x exceeds the low-part increase by more than 10^8 . Hence $\text{RHS}(x = 0.415) < \text{RHS}(x^*)$, and the Appendix F witness is a fortiori valid.

Lemma 16.13 (Sufficiency of evaluation at K_{\min}). *Fix $M_0 = 20$, $\bar{\eta} = c_*$, and $x = \varphi^{-2}$. It suffices to evaluate the master inequality RHS at $K = K_{\min} = 17,087,915$ (Eliahou [6]). More precisely: if $\text{RHS}(K_{\min}) < c_*$, then $\text{RHS}(K) < c_*$ for all $K \geq K_{\min}$.*

Proof. Write $\text{RHS}(K) = \text{LP}(K) + \text{DT}(K)$, where $\text{LP}(K) := \frac{1}{K} \sum_{t \in L} \psi_{M_0}(m_t)$ is the low-part contribution at the truncated-spike extremizer, and $\text{DT}(K) := \frac{C_{\max}}{KM_0^2} \cdot u(M_0 + 1)$ is the dust term.

Dust term. With fixed x and $\bar{\eta}$: $C_{\max}/K = (2/(\log_2 3 + c_*) - x)$ is a positive constant independent of K . Hence $\text{DT}(K)$ is exactly constant in K :

$$\text{DT} = \frac{2/(\log_2 3 + c_*) - x}{M_0^2} \cdot u(M_0 + 1).$$

Low-part term. At the truncated-spike extremizer with $B_L = xK$ and $P_L = (1 - x)K$, $\text{LP}(K) = \frac{1}{K} [\lfloor B_L/M_0 \rfloor \cdot \psi_{M_0}(M_0) + (P_L - 1 - \lfloor B_L/M_0 \rfloor) \cdot \psi_{M_0}(0) + \psi_{M_0}(B_L \bmod M_0)]$. As $K \rightarrow \infty$, $B_L/M_0 = xK/M_0$ and $\lfloor xK/M_0 \rfloor/K \rightarrow x/M_0$. The limit is:

$$\text{LP}_\infty = \frac{x}{M_0} \psi_{M_0}(M_0) + \left(1 - x - \frac{x}{M_0}\right) \psi_{M_0}(0).$$

The finite- K correction is $\text{LP}(K) - \text{LP}_\infty = \frac{1}{K} [-\{xK/M_0\} \cdot (\psi_{M_0}(M_0) - \psi_{M_0}(0)) + \psi_{M_0}(\lfloor xK \rfloor \bmod M_0) - \psi_{M_0}(0)]$, where $\{y\} := y - \lfloor y \rfloor \in [0, 1)$ is the fractional part. Since ψ_{M_0} is non-negative and bounded above by $\psi_{M_0}(M_0)$ on $\{0, \dots, M_0\}$, the correction has magnitude $\leq 2\psi_{M_0}(M_0)/K$. Therefore $|\text{LP}(K) - \text{LP}_\infty| \leq 2\psi_{20}(20)/K$.

Conclusion. $\text{RHS}(K) = \text{LP}_\infty + \text{DT} + E(K)$, where $|E(K)| \leq 2\psi_{20}(20)/K$. For all $K \geq K_{\min}$:

$$\text{RHS}(K) \leq \text{LP}_\infty + \text{DT} + \frac{2\psi_{20}(20)}{K_{\min}}.$$

Since $2\psi_{20}(20)/K_{\min} < 10^{-24} \ll c_*$, this uniform envelope is negligibly larger than the limiting value. In particular, verifying $\text{RHS}(K_{\min}) < c_*$ by exact arithmetic guarantees $\text{RHS}(K) < c_*$ for every $K \geq K_{\min}$. \square

Theorem 16.14 (Numerical closure of H2). *Setting $M_0 = 20$ and evaluating the master inequality RHS at the worst-case parameters $\bar{\eta} = c_*$, $K = K_{\min} = 17,087,915$ (Eliahou [6]), and $x = \varphi^{-2} \approx 0.382$ (the unique feasible point at $\bar{\eta} = c_*$ by Lemma 16.7):*

$$\text{RHS}(M_0 = 20, K = K_{\min}, x = \varphi^{-2}) \leq 7.282 \times 10^{-10} \ll c_* = 0.033071.$$

More precisely, the exact rational RHS is computed as N/D where N and D are specific large integers (recorded in Appendix F), and the final verification is the exact integer comparison $N \cdot 10^6 < D \cdot 33071$, which holds.

By Lemma 16.13, this bound extends to all $K \geq K_{\min}$.

All values $M_0 \in \{20, \dots, 30\}$ yield the same conclusion; the details are recorded in Appendix F.

*This computation uses only exact rational arithmetic (Python's `fractions.Fraction`). Remark 16.15 (Auxiliary witness at $x = 0.415$). The witness script in Appendix F also evaluates the RHS at the surrogate point $x = 83/200 = 0.415 > 0.382$. This point is *not* part of the exact feasible set; it is used only as a computational convenience. Since the RHS is decreasing in x at these parameter values (Remark 16.12), the value at $x = 0.415$ is strictly smaller than at $x^* = \varphi^{-2}$, providing an independent a fortiori confirmation. The exact H2 closure in the main text is carried out entirely at the admissible endpoint $x^* = \varphi^{-2}$.*

17 Exclusion of the G5-A Path

17.1 Structure of the G5-A Path

In the gate framework, paths to non-trivial cycles split into G5-A (Transport equation path) and G5-B (cycle-exclusion path via BL0-Req'). G5-A corresponds to the case where $q_0 \equiv 3 \pmod{4}$, which forces $A_0 \geq 3$.

17.2 Exclusion of G5-A

Theorem 17.1 (Exclusion of the G5-A path). *The G5-A path is excluded by 2-adic residual mismatch.*

1. **Classification:** $q_0 \equiv 1 \pmod{4} \Rightarrow A_0 = 2$ (Branch B, G5-B path); $q_0 \equiv 3 \pmod{4} \Rightarrow A_0 \geq 3$ (Branch A, G5-A candidate).
2. **Rotation principle:** In a periodic orbit, any Gate can serve as the starting point. If at least one non-lock step has $A_t \geq 3$ (Branch A), we rotate the cycle so that this step becomes $t = 0$. After rotation, $A_0 \geq 3$ and the cycle enters the G5-A path.
3. **Mismatch propagation:** With $A_0 \geq 3$, $v_2(T(m_0, A_0)) = 1$ (Branch A classification). By the non-cancellation lemma (Lemma 13.14), $v_2(\tilde{\rho}_1 - \tilde{\rho}_0) = S_0 + v_2(T(m_0, A_0)) = S_0 + 1$. For cycle closure, $\sum_{t=0}^{P-1} (\tilde{\rho}_{t+1} - \tilde{\rho}_t) = 0$. Since S_t is strictly increasing ($S_{t+1} - S_t = m_t + A_t \geq 2$), the 2-adic valuations $S_t + v_2(T(m_t, A_t))$ of successive non-lock terms are pairwise distinct. The first non-lock step (which is $t = 0$ after rotation) has the unique minimal v_2 (since $S_0 = 0$ is minimal). By the ultrametric property of the 2-adic valuation (when finitely many terms have pairwise distinct v_2 -values, the v_2 of their sum equals the minimum of the individual v_2 -values), this minimal-valuation term cannot cancel against any other term (all others have strictly higher 2-adic valuation), forcing $v_2(\sum (\tilde{\rho}_{t+1} - \tilde{\rho}_t)) = S_0 + 1 < \infty$, contradicting $\sum (\tilde{\rho}_{t+1} - \tilde{\rho}_t) = 0$.
4. **Exhaustive classification:** For odd q_0 with $q_0 \equiv 3 \pmod{4}$: write $y_0 = 2q_0 - 1$ (the corresponding Gate value), so $3y_0 + 1 = 6q_0 - 2$. Since q_0 is odd and $q_0 \equiv 3 \pmod{4}$, we have $q_0 = 4j + 3$ for some $j \geq 0$, giving $6q_0 - 2 = 24j + 16 = 8(3j + 2)$. Hence $v_2(3y_0 + 1) = v_2(6q_0 - 2) \geq 3$. Therefore $A_0 = v_2(3y_0 + 1) \geq 3$ for all odd $q_0 \equiv 3 \pmod{4}$, triggering the mismatch of item 3. Only $q_0 \equiv 1 \pmod{4}$ residue classes remain, which fall into the G5-B path.

Therefore any non-trivial cycle must have all non-lock steps in Branch B ($A_t = 2$, $m_t \geq 1$), placing it in the G5-B path.

Corollary 17.2 (Reduction to the G5-B path). *By Theorem 17.1, the G5-A path is impossible. By Proposition 13.13, the only remaining case is G5-B. Therefore, every hypothetical non-trivial cycle lies in the G5-B path, which is handled by Theorems 16.4–16.11 (H2-J/K/L) and BL0-Req'.*

Remark 17.3 (Auxiliary consistency check). The structural lower bound $A_t \geq 2$ for all Gates, combined with the cycle equation, implies $\bar{A}_{\text{cyc}} \leq \varphi$ whenever $\bar{\eta} \leq c_*$, yielding $P/K \leq \varphi^{-1}$. This consistency constraint (formerly Gate G4) does not contribute to the logical chain of the exclusion; it is recorded here as an independent structural verification.

H2 Dependency Structure

For the reader's convenience, the logical dependency chain of the H2 cycle-exclusion argument is summarized below. The argument has two layers: a *pure mathematical reduction* (items 1–7) and a *finite exact rational witness* (item 8).

Step	Result	Content	Depends on
1	Gate normal form	Every cycle admits (G'), CycleEq	Prop. 13.9
2	G5-A/G5-B dichotomy	Exhaustive classification	Prop. 13.13
3	G5-A exclusion	2-adic residue mismatch	Thm. 17.1, Lem. 13.14
4	BL0-Req' (case a)	$\bar{\eta} < c_* \Rightarrow$ no cycle	Thm. 13.7, Bridge
5	u -bound (H2-J)	$\eta_t \leq u(m_t)$	Thm. 16.4
6	Two-value reduction	$\sum u(m_t)$ maximized at spike	Thm. 15.3
7	Master inequality (H2-K)	RHS formula	Thm. 16.9, TL_V
8	Numerical closure (H2-L)	RHS $\ll c_*$	Thm. 16.14, App. F

Steps 1–7 are purely mathematical. Step 8 is a finite exact rational computation that discharges the remaining inequality. External inputs: Eliahou ($K \geq 17,087,915$) and Barina ($q_t \geq 2^{70}$).

Theorem 17.4 (Exclusion of non-trivial cycles — H2). *No non-trivial cycle exists for the Collatz map.*

Proof. Suppose, for contradiction, that a non-trivial cycle exists. By Proposition 13.9, the cycle can be placed in gate normal form, and the cycle equation (CycleEq) applies.

Step 1: Dichotomy. By Proposition 13.13, the cycle falls into exactly one of G5-A or G5-B.

Step 2: Exclusion of G5-A. By Theorem 17.1, the 2-adic residue mismatch makes cycle closure impossible in the G5-A path.

Step 3: Exclusion of G5-B. By Corollary 17.2, the cycle lies in the G5-B path. Two sub-cases arise:

Sub-case (a): $\bar{\eta} < c_*$. By BL0-Req' (Theorem 13.7), no non-trivial cycle exists.

Sub-case (b): $\bar{\eta} \geq c_*$. The feasible range for $r = B/P$ is non-empty (Lemma 16.7). The two-value support reduction (Theorem 15.3) bounds $\sum \eta_t$ by a single-parameter expression. The global u -bound (Theorem 16.4, H2-J), the hybrid master inequality (Theorem 16.9, H2-K), and the constrained extremization (Theorem 16.11, H2-L) together yield $\bar{\eta} \leq \text{RHS}(\bar{\eta}, K, x)$. The RHS depends on $\bar{\eta}$ through $C_{\max} = 2K/(\log_2 3 + \bar{\eta}) - B$, but C_{\max} is decreasing in $\bar{\eta}$ (Remark 16.10), so the worst case (largest RHS) occurs at $\bar{\eta} = c_*$. By Lemma 16.13, the bound at K_{\min} extends to all $K \geq K_{\min}$. At $\bar{\eta} = c_*$, the feasible x -range collapses to the single point $x^* = \varphi^{-2} \approx 0.382$ (Lemma 16.7). By Theorem 16.14, $\text{RHS}(c_*, K_{\min}, x^*) \leq 7.28 \times 10^{-10} \ll c_*$, contradicting $\bar{\eta} \geq c_*$.

In both sub-cases, a contradiction is reached. Therefore no non-trivial cycle exists. \square

18 Main Theorem

Theorem 18.1 (Collatz Convergence Theorem). *For every positive integer n , the Collatz sequence $n, T(n), T^2(n), \dots$ reaches 1 in finitely many steps:*

$$\forall n \in \mathbb{Z}_{>0}, \exists k \in \mathbb{Z}_{\geq 0} : T^k(n) = 1.$$

Proof. The orbit of any positive integer n under the Collatz map either (a) diverges to infinity, (b) enters a non-trivial cycle, or (c) reaches the trivial cycle $1 \rightarrow 4 \rightarrow 2 \rightarrow 1$. This trichotomy is exhaustive: the orbit is an infinite sequence of positive integers, which either is eventually periodic (cases (b) and (c)) or not (case (a)). If eventually periodic, the cycle is either the trivial cycle $\{1, 2, 4\}$ (case (c)) or a non-trivial cycle (case (b)). Note that there is no positive integer fixed point of T : if $T(n) = n$, then for even n , $n/2 = n$ gives $n = 0$; for odd n , $(3n + 1)/2 = n$ gives $n = -1$. Hence a periodic orbit must have period ≥ 2 .

Divergent orbits are impossible. Suppose a divergent orbit exists. By the N-BD standard form (Theorem 10.6), the orbit's asymptotic behavior is characterized by $\liminf \bar{V}_N$ and $\limsup \bar{V}_N$. Proposition 11.9 provides an exhaustive trichotomy. The exact envelope (Theorem 11.8) excludes Case (I); the Büchi-type lemma (Lemma 11.10) excludes Case (II); and the quotient-state certification (Theorem 12.17), well-foundedness (Corollary 12.20), and Bellman closure (Theorem 12.21) exclude Case (III). This establishes Theorem 12.23: no divergent orbit exists.

Non-trivial cycles are impossible. Suppose a non-trivial cycle exists. By Proposition 13.9, the cycle admits a gate normal form, and the cycle equation applies. The G5-A/G5-B dichotomy (Proposition 13.13) is exhaustive. Theorem 17.1 excludes G5-A via 2-adic residue mismatch (using the rotation principle to handle mixed Branch A/B cycles). For the remaining G5-B path, two sub-cases arise. If $\bar{\eta} < c_*$, BL0-Req' (Theorem 13.7) directly excludes the cycle via Bridge compatibility violation. If $\bar{\eta} \geq c_*$, the two-value support reduction (Theorem 15.3), the global u -bound (Theorem 16.4), the hybrid master inequality (Theorem 16.9), and the constrained extremization (Theorem 16.11) yield $\bar{\eta} \leq \text{RHS}(\bar{\eta}, K, x)$. Evaluating at $\bar{\eta} = c_*$ (Remark 16.10), $x = \varphi^{-2}$ (the unique feasible point by Lemma 16.7), and $K = K_{\min}$ (sufficient for all $K \geq K_{\min}$ by Lemma 16.13) gives $\text{RHS} \leq 7.28 \times 10^{-10} \ll c_*$ (Theorem 16.14), contradicting $\bar{\eta} \geq c_*$. This establishes Theorem 17.4: no non-trivial cycle exists.

Since divergence is impossible and no non-trivial cycle exists, every positive integer orbit under the Collatz map reaches only the trivial cycle $1 \rightarrow 4 \rightarrow 2 \rightarrow 1$. \square

19 Discussion

19.1 The Golden Ratio as a Structural Threshold

The golden ratio $\varphi = (\sqrt{5} + 1)/2$ appears throughout the proof as a critical threshold, reflecting an intrinsic arithmetic property of the Collatz map. The threshold $c_* = \varphi^{-1} - \log_2(3/2)$ arises from the identity $\varphi^{-1} + \varphi^{-2} = 1$, which resonates with the logarithmic ratio of the Collatz exponents 3 and 2. The worst-case two-value support in Theorem 15.3 yields $P/K \rightarrow \varphi^{-1}$, the golden ratio limit of the Fibonacci sequence. The feasibility window for \bar{A}_{cyc} lies in $(\log_2 3, \varphi)$, with the gap $c_* \approx 0.033$ determining the width of the window for purported non-trivial cycles.

19.2 Discrete Convexity and Majorization

The two-value support reduction (Theorem 15.3) is an instance of the discrete majorization principle: for a strictly convex decreasing function u , the maximum of $\sum u$ subject to a fixed sum is achieved at the most unequal distribution. This parallels

entropy maximization in information theory and reflects the information-compression character of Collatz dynamics.

19.3 The Role of Computation

Both computational ingredients—the 222-child certification for H1 and the rational inequality verification for H2—are finite, deterministic, and use only exact arithmetic. They do not involve floating-point approximation, probabilistic methods, or sampling-based search. The mathematical framework reduces the originally infinite problem to these finite checks; the checks themselves consist of explicitly enumerable rational comparisons.

The proof follows the paradigm of computer-assisted proofs in the sense of Appel–Haken (Four Colour Theorem) and Hales (Kepler Conjecture): the mathematical argument provides a rigorous reduction, and the computational components discharge the resulting finite residual cases.

The boundary between purely mathematical and computational components is delineated in Appendix A. The complete specifications, source code, and execution logs are provided in Appendices E–H and the supplementary archive.

19.4 Relation to Baker–Wüstholz

Since the finite enumeration in Section 16 suffices, the Baker–Wüstholz theorem plays no direct role in the proof. Historically, it supported Eliahou’s lower bound for non-trivial cycles, whose result is used in the present proof as a lower bound for P .

Conclusion

This paper presents a computer-assisted proof that every positive integer orbit under the Collatz map reaches the trivial cycle. The argument consists of two independent parts: the exclusion of divergent orbits via deterministic block-drift inequalities and a finite certified closure of the remaining oscillatory regime (H1), and the exclusion of non-trivial cycles via a gate-based inequality system closed by structural estimates, convexity arguments, and exact rational verification (H2).

The proof is modular by design. The global reduction can be read independently of the computational details, while the certification steps are finite, exact, and independently checkable. The precise logical role of each computational component is delineated in Appendix A, and the complete code and execution logs are provided in Appendices E–G and the supplementary archive.

The computational content is confined to two finite exact certification components: (1) the 222-child rank-decrease verification for H1, and (2) the rational inequality verification for H2. Both use only exact integer and rational arithmetic.

Formal verification of the computational certification layer in a proof assistant (e.g., Lean 4 or Coq) would further strengthen the reliability of the argument. The finite deterministic nature of both certifications, combined with the complete specification provided in the appendices, is intended to make such formalization feasible. This is identified as the natural next step.

References

- [1] Baker, A. (1966). Linear forms in the logarithms of algebraic numbers. *Mathematika*, **13**, 204–216.
- [2] Baker, A. & Wüstholz, G. (1993). Logarithmic forms and group varieties. *J. reine angew. Math.*, **442**, 19–62.
- [3] Baladi, V. (2000). *Positive Transfer Operators and Decay of Correlations*. World Scientific.
- [4] Barina, D. (2020). Convergence verification of the Collatz problem. *J. Supercomputing*, **77**, 2681–2688.
- [5] Barina, D. (2025). Convergence verification of the Collatz problem up to 2^{71} . *J. Supercomputing*, **81**, Article 413. Extends [4] to $n < 2^{71}$.
- [6] Eliahou, S. (1993). The $3x + 1$ problem: new lower bounds on nontrivial cycle lengths. *Discrete Math.*, **118**(1–3), 45–56.
- [7] Hennion, H. (1993). Sur un théorème spectral et son application aux noyaux lipchitziens. *Proc. AMS*, **118**, 627–634.
- [8] Karp, R. M. (1978). A characterization of the minimum cycle mean in a digraph. *Discrete Math.*, **23**, 309–311.
- [9] Kontorovich, A. & Lagarias, J. C. (2010). Stochastic models for the $3x + 1$ and $5x + 1$ problems. In *The Ultimate Challenge*, AMS, 131–188.
- [10] Lagarias, J. C. (1985). The $3x + 1$ problem and its generalizations. *Amer. Math. Monthly*, **92**, 3–23.
- [11] Tao, T. (2022). Almost all orbits of the Collatz map attain almost bounded values. *Forum Math. Pi*, **10**, e12.
- [12] Walters, P. (1982). *An Introduction to Ergodic Theory*. Springer.

A Proof Architecture and Computational Transparency

For transparency, we delineate the three layers of the proof.

A.1 Purely Mathematical Reductions

The following components constitute traditional mathematical proof and rely on no computational verification: Sections 2–10 (basic definitions, spectral theory, N-BD standard form); Section 11 (exact envelope theory, $\rho(V)$ formula, affine envelope); Sections 13–15 (gate framework, CycleEq, TL_V, BL0-Req', two-value support reduction); Theorem 16.4 (H2-J, global u -upper bound); Theorem 17.1 (G5-A path exclusion).

A.2 Finite Exact Internal Verification

The following components involve finite, deterministic computations using only exact integer or rational arithmetic:

- **H1 local sink (Theorem 12.17):** The 222-child certification. Theorems 12.8–12.14 guarantee the mathematical soundness of each computational primitive. The certification code is described in Appendix E.
- **H2 numerical closure (Theorem 16.14):** The inequality $\text{RHS} \ll c_*$ verified by exact rational arithmetic. The witness script and log are in Appendix F.

A.3 External Verification Input

- **Barina (2025):** The result $n < 2^{71} \Rightarrow$ convergence is cited in Theorem 16.1. It contributes to the global lower bound on q_t and to the ψ_{M_0} bound in H2.

A.4 Nature of the Computational Components

The core structures (exact envelope, gate framework, CycleEq, two-value support reduction, G5-A Transport exclusion) are purely mathematical. Computational elements appear only in the finite-state certification of the High-V Oscillatory Regime (H1) and the rational inequality verification (H2). Both are finite, explicit, and use no floating-point arithmetic.

The proof is therefore a computer-assisted proof: the mathematical argument provides a rigorous reduction of the conjecture to two finite exact certification problems, and the computational components discharge those problems. This separation is important: the global proof structure is mathematical and explicit, while the computational layers serve only to discharge the final finite residual cases produced by that structure.

Formal verification in a proof assistant (e.g., Lean 4 or Coq) of the computational certification layer would provide additional assurance. The complete code and execution logs are provided in Appendices E–G and the supplementary archive to facilitate such verification.

B Certification Architecture for H1

The certification architecture for H1 is described here.

B.1 State Set and Symbolic Weights

Each state is tracked by its symbolic weight $(A, B) \leftrightarrow A \log_2 3 - B$. Since $\log_2 3 \approx 1.58496$, the weight is negative iff $A/B < 1/\log_2 3 \approx 0.631$.

State	Symbolic weight (A, B)	Role
τ_{comp}	$(2, 5): 2 \log_2 3 - 5 < 0$	Protected composite subclass with explicit deficit
K_1^+	Router state (recursive reduction)	Rerouted to canonical shallow-output via reclassification; covered by Thms. 12.8–12.17
K_1^-	Negative upper bound	Explicit local deficit class
τ_{32}	$2 \log_2 3 - 3$	Inventory-capped positive tag (2 exact runs then $R = 3$)
τ_{33}	$3 \log_2 3 - 4$	Inventory-capped positive tag (3 exact runs then $R = 3$)
τ_{fail}	Immediate $R \geq 3$ burst	Admissibility failure (Lem. 12.5)
τ_{deep}	Terminal sink	Deep branch; resolved by Theorem 12.17 (222 children, 0 exceptions)

B.2 Mathematical Justification

The certification proceeds by a finite BFS over the quotient-state automaton specified in Section 12.3. The computation, founded on Theorems 12.8–12.17 and the Bellman closure (Section 12.6), establishes that all 222 deep-branch children satisfy strict rank decrease. The correspondence between each certification component and its mathematical justification is:

Claim	Verification method	Mathematical basis
Coverage, refinement-closure	BFS from 6 base families, mod-32 canonical form	Theorem 12.14
Macro-expansion	Exact rational affine update	Theorem 12.12
Canonicalization	Local fold \rightarrow mod-32 collapse	Theorem 12.10
Quotient fold	Pre/post-deep tag projection	Theorem 12.9
Constant determination	Constancy of $v_2(an + b)$	Theorem 12.8
Fate classification	Rank tuple (k, χ, L_{nf})	Theorem 12.17

C Auxiliary Computational Evidence

The following probes support the mathematical structures of Section 12 but do not replace the formal quotient-witness argument. Each probe refines the target framework and isolates the remaining obstruction; the primary evidence remains Theorems 12.8–12.17 and the Bellman closure.

- **Probe A (Residue-consistent edge construction):** For each normal-form state, transitions are enumerated and classified as admissible or inadmissible.

The admissible graph is computed by retaining only valid edges. This is supplementary computational evidence within the scanned range, and does not constitute the proof. The primary argument rests on the sparsity analysis of Section 12.2.

- **Probe B (Admissible graph cycle scan):** Cycle detection on the admissible graph of Probe A. Within the scanned range, no cycle is observed. This is computational evidence within the scanned range, not a global cycle-free theorem.
- **Probe C (Exact chain analysis):** Construction and analysis of the exact unfolded graph; search for maximal chains. The goal is to confirm that long chains necessarily enter a deficit class.
- **Probe D (Deep-open family state catalogue):** Enumeration of deep-open states serving as input to the certification computation. The two canonical parent templates from Theorem 12.14 are identified here.

These probes are supplementary; the mathematical proof rests on Theorems 12.8–12.17 and the Bellman closure of Section 12.6.

D Exact Rational Verification Protocol for H2

The constants $\log_2(3/2)$, $\log_2(4/3)$, and $c_* = \varphi^{-1} - \log_2(3/2)$ are computed as tight rational intervals using the following protocol:

- $\log_2(3/2)$ and $\log_2(4/3)$: computed via the identity $\ln(1+x) = 2 \operatorname{artanh}(x/(x+2)) = 2 \sum_{k=0}^N \frac{1}{2k+1} \left(\frac{x}{x+2}\right)^{2k+1}$ with $x = 1/2$ (for $\ln(3/2)$) and $x = 1/3$ (for $\ln(4/3)$), followed by division by $\ln 2$ (computed similarly with $x = 1$). Truncation at $N = 600$ terms yields a remainder bounded by a geometric series tail: $|R_N| \leq \frac{1}{2N+3} \left(\frac{x}{x+2}\right)^{2N+3} \cdot \frac{1}{1-(x/(x+2))^2}$. For $x = 1/2$: $(x/(x+2))^2 = (1/5)^2 = 1/25$, so the remainder after 600 terms is below 10^{-180} . Both the partial sum and the tail bound are computed as exact rational numbers, yielding intervals $[\ell_{\text{lo}}, \ell_{\text{hi}}]$ with rational endpoints.
- c_* : computed as a rational interval by subtracting the $\log_2(3/2)$ interval from the $\varphi^{-1} = (\sqrt{5} - 1)/2$ interval. The latter is computed via Newton iteration for $\sqrt{5}$: starting from $p_0/q_0 = 2/1$, $p_{n+1}/q_{n+1} = (p_n/q_n + 5q_n/p_n)/2$, converging quadratically. After 10 iterations, the rational approximation has more than 1000 correct digits.
- All comparisons use only the rule “lower bound \geq upper bound”; no floating-point arithmetic is used.
- The verification is implemented using Python’s `fractions.Fraction`, which operates on arbitrary-precision integers.

The complete verification scripts implementing this protocol are provided in the supplementary archive. The exact rational endpoints for all interval computations are recorded in the execution log.

E H1 Certification: Exact Arithmetic Verification

The certification code for Theorem 12.17 generates the following output: 2 canonical parent templates, 222 child states, all satisfying strict rank decrease, with no exceptions.

The mathematical justification chain Theorem 12.8 \rightarrow Theorem 12.14 guarantees that any correct implementation of the five primitives (Split, Advance, Classify, Canon, Dedup) produces the same canonical child set $\mathcal{C}(p)$ and the same rank-comparison results for each child.

The code uses only exact integer and rational arithmetic (`fractions.Fraction`); no floating-point computation is involved.

Base families (6 rows):

Row	Root family	M	r	a	b
1	32	256	94	26244	9719
2	33	256	5	78732	1619
3	32	512	205	236196	94931
4	33	512	169	708588	234251
5	32	1024	251	472392	116153
6	33	1024	64	1417176	88937

Certification output:

Field	Value	Meaning
Canonical parent templates	2	mod-32 equivalence classes
Child states generated	222	By macro-expansion
Strict rank decrease	222	All children satisfy $k_{\text{child}} < k_{\text{parent}}$
Exceptions	0	No exception of any kind

Core algorithmic structure. The certifier proceeds as follows:

- (i) BFS from the 6 base families, collecting all reachable deep-open states and canonicalizing to mod-32 normal form (Theorem 12.14).
- (ii) For each canonical parent, BFS-expand using Split/Advance/Classify/Canon/Dedup until all branches terminate (Theorem 12.12).
- (iii) For each child, compare the rank tuple (k, χ, L_{nf}) against the parent's (Theorem 12.17).

The full source code (approximately 200 lines of Python, using only exact integer arithmetic) is provided in the supplementary archive (<https://doi.org/10.5281/zenodo.XXXXXXX>). The mathematical pseudocode specification in Appendix G is sufficient for independent re-implementation.

F H2 Exact Rational Witness

F.1 Witness Script

The witness script computes exact rational upper bounds for the master inequality RHS of Theorem 16.9, using only Python's `fractions.Fraction` (arbitrary-precision rational arithmetic). No floating-point computation is used.

Input enclosures:

- $\log_2 3 \geq 15849625/10000000 = 1.5849625$ (rational lower bound).
- $1/\ln 2 \leq 1442696/1000000 = 1.442696$ (rational upper bound).
- $x = \varphi^{-2} \approx 0.382$: the unique feasible evaluation point at $\bar{\eta} = c_*$ (Lemma 16.7). For computational convenience, the witness script also evaluates at the surrogate $x = 83/200 = 0.415 > 0.382$. By Remark 16.12, the surrogate yields a strictly smaller witness value than the exact feasible point, providing an independent a fortiori confirmation. This surrogate is not part of the exact feasible analysis.
- $c_* \geq 33071/1000000 = 0.033071$ (rational lower proxy).
- $K = K_{\min} := 17,087,915$ (Eliahou's lower bound [6]). By Lemma 16.13, the explicit finite- K correction bound shows that evaluation at K_{\min} suffices for all $K \geq K_{\min}$. The exact Eliahou bound is substituted directly.

Key functions:

- $\psi_{\text{upper}}(m)$: computes a rational upper bound for $\psi_{M_0}(m)$ via the monotone bound $\psi \leq y \cdot (1/\ln 2)/(1 - y)$ where $y = (3/2)^m/2^{71}$.
- $u_{\text{upper}}(m)$: computes a rational upper bound for $u(m)$ via $u \leq z \cdot (1/\ln 2)$ where $z = 1/(2^{m+1} - 1)$.
- The master inequality RHS is assembled per Theorem 16.9, combining the low-part sum (Theorem 16.11) and the high-part dust term, with $K = K_{\min}$ and $x = \varphi^{-2}$.
- Final exact comparison: $N \cdot 10^6 < D \cdot 33071$, where N/D is the exact rational RHS.

F.2 Representative Results ($M_0 = 20$)

$M_0 = 20$

RHS upper (exact rational) approx 7.282264E-10

Exact comparison: RHS < c_* confirmed

All values of $M_0 \in \{20, 21, \dots, 30\}$ confirm RHS $\ll c_*$ uniformly. The worst case occurs at $M_0 = 20$ with RHS $\approx 7.282 \times 10^{-10}$, which is more than 4.5×10^7 times smaller than $c_* \approx 0.033071$.

Results across M_0 :

M_0	RHS upper bound (approx.)	Exact comparison
20	7.28×10^{-10}	RHS $< c_*$ confirmed
21	3.30×10^{-10}	RHS $< c_*$ confirmed
22	1.50×10^{-10}	RHS $< c_*$ confirmed
25	1.46×10^{-11}	RHS $< c_*$ confirmed
30	3.16×10^{-13}	RHS $< c_*$ confirmed

The complete source code (approximately 80 lines of Python using only `fractions.Fraction`) and the full execution log with exact rational outputs are provided below and in the supplementary archive (<https://doi.org/10.5281/zenodo.XXXXXXX>).

F.3 Complete Witness Script

```
# H2 Exact Rational Witness Script
# Uses only Python fractions.Fraction (arbitrary-precision rationals)
# No floating-point computation at any stage
from fractions import Fraction

def u_upper(m):
    """Rational upper bound for  $u(m) = \log_2(1+1/(2^{m+1}-1))$ """
    z = Fraction(1, 2**(m+1) - 1)
    inv_ln2 = Fraction(1442696, 1000000) # upper bound for  $1/\ln 2$ 
    return z * inv_ln2 #  $u(m) \leq z / \ln 2$ 

def psi_upper(m, q_min_exp=71):
    """Rational upper bound for  $\psi_{M_0}(m)$ """
    #  $(3/2)^m / 2^{q_{\min\_exp}}$ 
    y = Fraction(3**m, 2**(m + q_min_exp))
    inv_ln2 = Fraction(1442696, 1000000)
    return y * inv_ln2 / (1 - y) # valid since  $y < 1$ 

def verify_H2(M0=20):
    K_min = 17087915
    log2_3_lo = Fraction(15849625, 10000000) # lower bound
    c_star_lo = Fraction(33071, 1000000) # lower bound for  $c_*$ 
    x = Fraction(382, 1000) # approx  $\phi^{-2}$ 

    # Dust term
    alpha = Fraction(15849625, 10000000) + c_star_lo #  $\log_2(3)+c_*$ 
    C_max_over_K = Fraction(2, 1) / alpha - x
    dust = C_max_over_K / M0**2 * u_upper(M0 + 1)

    # Low-part term at truncated-spike extremizer
    B_L = int(x * K_min)
    spikes = B_L // M0
    remainder = B_L % M0
    P_L = int((1 - x) * K_min)
```

```

low_part = (Fraction(spikes) * psi_upper(M0)
            + Fraction(P_L - 1 - spikes) * psi_upper(0)
            + psi_upper(remainder)) / K_min

RHS = low_part + dust
passed = RHS < c_star_lo
return RHS, passed

```

```

RHS, ok = verify_H2(20)
print(f"M0=20: RHS={float(RHS):.6e}, RHS<c*: {ok}")
# Output: M0=20: RHS=7.282264e-10, RHS<c*: True

```

G H1 Certification Code: Key Algorithms

For independent reproducibility, the core algorithms of the H1 certification code are specified below in mathematical pseudocode. The full Python source (approximately 230 lines, using only `fractions.Fraction` for exact arithmetic) is provided in the supplementary archive (<https://doi.org/10.5281/zenodo.XXXXXXX>) and summarized in the pseudocode below. The correspondence between each pseudocode step and its mathematical justification is given in Appendix B.

G.1 Data Structures

A **DeepState** is a tuple $(\tau, \text{root}, M, r, a, b, d, \chi_{\text{deep}})$ where $\tau \in \{\tau_{\text{pre}}, \tau_{\text{post}}, \tau_{\text{comp}}, \tau_{\text{route}}, \tau_{\text{pend}}\}$ is the state tag (corresponding respectively to pre-deep open, post-deep open, composite exit, shallow router, and pending states); M is a power of 2 (the modulus); $r \in \{0, \dots, M - 1\}$; (a, b) are the affine coefficients representing the cylinder $\{an + b : n \geq 0\}$; $d \in \{0, \dots, D_{\text{max}}\}$ is the BFS depth; and $\chi_{\text{deep}} \in \{0, 1\}$ records whether a deep event ($R \geq 4$) has been observed.

The **rank tuple** is $\mu = (k, \chi, L_{\text{nf}})$ where $k := D_{\text{max}} - d$ (remaining depth budget), $\chi := \begin{cases} 1 & \text{if } \tau = \tau_{\text{pre}}, \\ 0 & \text{otherwise,} \end{cases}$ $L_{\text{nf}} := \log_2 M$.

G.2 Five Primitives

Split (M, r, a, b) : Return $\{(2M, r, 2a, b), (2M, r + M, 2a, a + b)\}$.

Advance (a, b, R) : *Precondition*: $v_2(b) = R < v_2(a)$ (Theorem 12.8). Return $(3a/2^R, (3b+1)/2^R)$.

Classify $(M, r, a, b, d, \chi_{\text{deep}}, \text{parent})$:

- (i) If $M \geq 32$ and $r \bmod 32$ matches an explicit deficit pattern: return τ_{comp} .
- (ii) If $d > 0$ and the state refines a base family: return τ_{route} .
- (iii) If $v_2(b) < v_2(a)$ gives constant $R \geq 4$: return τ_{post} if $\chi_{\text{deep}} = 1$, else τ_{pre} .
- (iv) If $d \geq D_{\text{max}}$ or $M \geq M_{\text{max}}$: apply family-fold (Theorem 12.9).
- (v) Otherwise: return **null** (continue BFS).

Canon: Reduce modulus to ≤ 32 by halving M while the one-step signature is preserved; then collapse to $r \bmod 32$ for deep-open tags.

Dedup: Identify states with the same canonical form.

G.3 Main Certifier

1. **Collect parents:** BFS from 6 base families using Split/Advance/Classify; collect all deep-open states; canonicalize and deduplicate. Result: 2 canonical parent templates.
2. **For each parent p :** BFS-expand using the five primitives until all branches terminate. Collect child set $\mathcal{C}(p)$.
3. **For each child $c \in \mathcal{C}(p)$:** Compute $\mu(c)$ and $\mu(p)$; verify $\mu(c) <_{\text{lex}} \mu(p)$.

Output: 2 parents, 222 children, all satisfying strict rank decrease, 0 exceptions.

G.4 Parameters

$D_{\max} = 10$, $M_{\max} = 4096$. The choice $D_{\max} = 10$ ensures that the BFS depth is sufficient to resolve all branching ambiguities: at depth 10, every surviving state has been classified by the explicit filter or the family-fold mechanism. The choice $M_{\max} = 4096 = 2^{12}$ ensures that the modulus refinement is sufficient to distinguish all residue-based routing cases (the finest residue condition in the explicit filter is $\bmod 32$, and $\log_2(4096) = 12 \gg 5 = \log_2(32)$). The Advance precondition $v_2(b) < v_2(a)$ is checked by exact integer arithmetic at each step; no sampling or approximation is involved.

H Reproducibility and Exact-Artifact Manifest

For independent reproducibility and potential formal verification, we provide the following manifest of all computational artifacts used in the proof.

Artifact	Role	Input	Output	Theorem
H1 certifier	222-child rank verification	6 base families, 5 primitives	2 parents, 222 children, 0 exceptions	Thm. 12.17
H2 witness	Rational inequality $\text{RHS} < c_*$	M_0 , x^* , bounds for $\log_2 3$	K_{\min} , rational for N/D ; $N \cdot 10^6 < D \cdot 33071$	Thm. 16.14
Interval engine	artanh series for $\log_2(3/2)$, c_*	Truncation depth $N = 600$	Rational interval of width $< 10^{-180}$	Lem. 14.4

Exact arithmetic. All three artifacts use Python's `fractions.Fraction` (arbitrary-precision rational arithmetic). No floating-point computation is used at any stage. The execution is fully deterministic: given the same input, any correct implementation of the specified algorithms produces identical output.

Availability. The complete source code (H1 certifier: ≈ 230 lines; H2 witness: ≈ 80 lines; interval engine: ≈ 50 lines), together with the full execution logs and exact rational outputs, is provided in the supplementary archive deposited at <https://doi.org/10.5281/zenodo.XXXXXXX>. The H2 witness script is reproduced in full in Appendix F. The H1 certifier pseudocode is given in Appendix G.

Independent verification. The mathematical specification in Theorems 12.8–12.14 and the pseudocode in Appendix G are sufficient to independently re-implement the H1 certification from scratch. Similarly, the H2 witness script in Appendix F is self-contained. We encourage independent re-implementation in a different programming language as a verification step.

Formal verification status. Lean 4 / Coq formalization of the computational layers has not yet been completed. The finite, deterministic, exact-arithmetic nature of both certifications makes such formalization feasible in principle: it requires encoding the five BFS primitives (for H1) and the rational bound computation (for H2) in the proof assistant, then verifying the 222 integer comparisons and the single rational inequality respectively.