A Bit Analyzer for the Collatz Carry Equation: Saturation Laws and Window Recycling Across the Sturmian Tower

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

June 2026

Abstract

We introduce a *bit analyzer* for the accelerated odd-only Collatz carry equation: the diagnostic $\ell(D) = v_2(Q(D) + \Xi_\alpha)$, which reads the first-disagreement partial-sum depth between the periodization D^∞ and the characteristic Sturmian word c_α directly off the 2-adic bits of the carry quotient. Deployed across grouped tower families, the analyzer detects exact *saturation laws*, enables *window recycling* (one centered residue nonsingularizing an infinite two-parameter family on its bounded-size region), and reads the Stern–Brocot recursion through its apparent exceptions, which are exact Christoffel standard-factorization identities. Certified families: $X_5^a X_6^b$ with $\ell = 83$ for $a \geq 2$ (the fourth note [1]); $X_6^a X_7^b$ with $\ell = 64$ for *all* $a, b \geq 1$; $X_7^a X_8^b$ with $\ell = 568$ for $a \geq 2$; and the semiconvergent family $X_5^a X_{(29,46)}^b$ with $\ell = 83$ for $a \geq 3$ — each by exact integer-power certification of two or three floor lemmas. The recycled residues are generic-sized ($|r_5| \approx 2^{81.79}$, $|r_6| \approx 2^{62.37}$, $|r_7| \approx 2^{563.396}$), and three 100×100 exact grids (30,000 members) contain no singular member. The analyzer pins 2-adic candidates only; odd-prime divisibility remains the independent attainment layer. No claim about the Collatz conjecture is made, and no n-uniform theorem across the whole tower is claimed.

1 Introduction

The first four notes of this sequence built, in order: a finite-block reduction theory with its near-critical renormalizing boundary [1]; the 2-adic Sturmian carry constant Ξ_α and its exact convergence law on standard blocks [1]; the standard-block nonsingularity ledger with its window/odd-prime architecture [1]; and the first grouped family, $X_5^a X_6^b$, closed on its bounded-size region by a single 83-bit residue. The previous four notes are collected in [1]. The fourth note's method was family-specific. This note turns it into an instrument. The central thesis: *the bit analyzer reads the Sturmian tower's recursion through* 2*-adic agreement depth* — saturation plateaus are floor-lemma horizons, and depth exceptions are not noise but the Stern–Brocot tree made visible, as exact word identities. The instrument is the object of this note; four families are its demonstration, three of them brought to certified closure on their bounded-size regions [2].

2 Definitions

A valuation word $D = (d_0, \ldots, d_{L-1})$, $d_i \geq 1$, has partial sums $s_j = \sum_{i<j} d_i$ (so letters carry indices $0, \ldots, L-1$ and s_j sums letters $0, \ldots, j{-}1$), carry $C(D) = \sum_{j<L} 3^{L-1-j} 2^{s_j}$, wall $\delta(D) = 3^L - 2^S$ (always odd), and quotient $Q(D) = -C(D)/\delta(D)$; D is *singular* if $Q(D) \in \mathbb{Z}$. The shell-(q,p) lower Christoffel block is $X_{(q,p)}$ with letters $\lfloor (i{+}1)p/q \rfloor - \lfloor ip/q \rfloor$; X_n abbreviates the convergent shells of $\alpha = \log_2 3$, and $s_j(c_\alpha) = \lfloor j\alpha \rfloor$. With $\Xi(w) = \sum_{j \geq 0} 3^{-(j+1)} 2^{s_j(w)} \in \mathbb{Z}_2$ and $\Xi_\alpha = \Xi(c_\alpha)$, the periodic-carry identity $Q(D) = -\Xi(D^\infty)$ holds [1].

Definition 2.1 (Bit analyzer)**.** The *bit depth* of D is $\ell(D) = v_2(Q(D) + \Xi_\alpha)$, and the *centered residue* at that depth is $r_D = \operatorname{cent}\big(-\Xi_\alpha \bmod 2^{\ell(D)}\big)$. Given an Archimedean bound $B_D \geq |Q(D)|$, the window is *active* if $2B_D + 1 < 2^{\ell(D)}$, and then *missed* if $|r_D| > B_D$ (nonsingularity certified) or *hit* if $|r_D| \leq B_D$ (fall through to the odd layer); it is *vacuous* if $2B_D + 1 \geq 2^{\ell(D)}$.

The layer separation of [1] is unchanged: the analyzer performs 2-adic candidate pinning only; odd-prime divisibility of the odd wall decides actual integrality, and neither layer subsumes the other.

3 Fundamental Lemma

Lemma 3.1. *Suppose D^∞ and c_α first differ at letter index J, so their partial sums agree through s_J, and let d be the smaller of the two differing letters. Then*

$$\ell(D) = s_J + d.$$

In particular, on the alphabet $\{1,2\}$ relevant throughout this note, $\ell(D) = s_J + 1$.

Proof. By the periodic-carry identity, $Q(D) + \Xi_\alpha = \Xi(c_\alpha) - \Xi(D^\infty) = \sum_{j \geq J+1} 3^{-(j+1)}(2^{s_j(c_\alpha)} - 2^{s_j(D^\infty)})$. The $j = J+1$ term has $v_2 = \min(s_{J+1}) = s_J + d$; every later term has both partial sums at least $s_J + d + 1$ (letters are ≥ 1), hence strictly larger valuation; and powers of 3 are units in $\mathbb{Z}_2$. The ultrametric inequality is therefore an equality. $\square$

The analyzer thus converts a combinatorial quantity (first disagreement of two infinite words) into one 2-adic valuation of one rational number, computable either way — and the agreement of the two computations is itself a consistency check (Section 12).

4 Standard factorization: the degeneracy identities

Theorem 4.1 (Borel–Laubie; see Reutenauer)**.** *For lower Christoffel words $u = X_{(q,p)}$ and $v = X_{(q',p')}$ with $pq' - p'q = \pm 1$ and $p/q < p'/q'$, the concatenation uv is the lower Christoffel word of the mediant shell:*

$$X_{(q,p)}\, X_{(q',p')} = X_{(q+q',\,p+p')}.$$

This is Proposition 1 of Borel–Laubie [3] in the multiplicative form of Reutenauer [4, Cor. 2.4.2]; see also the monograph [5]. Convention alignment: our shell notation (q,p) records word length and valuation sum; the lower-slope-first orientation of the cited statement matches our lower-Christoffel floor convention, as pinned independently by the following exact word identities, each verified letter-for-letter:

$$X_5 X_6 = X_7, \qquad X_7 X_8 = X_{(359,569)}, \qquad X_7 X_8^2 = X_9,$$

$$X_5X_s = X_6, \qquad X_5^2X_s = X_7, \qquad X_3X_4 = X_{(7,11)},$$

with $X_s = X_{(29,46)}$ (wrong-order controls fail, as they must). The moral, which the analyzer discovered empirically before the theorem named it: *apparent bit-depth exceptions are not noise; they are the Stern–Brocot tree made visible.* A family member whose depth jumps off the saturation plateau has, in every observed case, collapsed by Theorem 4.1 into a standard or semiconvergent block of a deeper shell, where it obeys that shell's own law.

5 Family A: $X_5^aX_6^b$

The fourth note proved [1]: $\ell = 83 = p_5 + p_6 - 1$ for all $a \geq 2$, $b \geq 1$, with the family-wide window certificate at $|r_5| \approx 2^{81.79}$ closing the bounded-size region. The $a = 1$ branch: $(1, 1)$ is the identity $X_5X_6 = X_7$ (Tier 1) and obeys the level-7 standard law; the sub-law for $b \geq 2$ (first mismatch $J = 93$, depth 148) is measured at $b \leq 6$ with the proof template clear (Tier 3); its strip residue is computed in Section 10.

6 Family B: $X_6^aX_7^b$ — upper-leading saturation

Lemma 6.1. *$\lfloor j\alpha\rfloor = \lfloor 65j/41\rfloor$ for $1 \leq j \leq 40$, with first failure at $j = 41$: $\lfloor 41\alpha\rfloor = 64$ while $s_{41}(X_6) = 65$. Each instance is the exact integer statement $2^k \leq 3^j < 2^{k+1}$ and was certified by exact power comparison.*

Theorem 6.2 (Upper-leading saturation)**.** *For all $a \geq 1$, $b \geq 1$: $\ell(X_6^aX_7^b) = p_6 - 1 = 64$.*

Proof. By Lemma 6.1, the first X_6 copy agrees with c_α through letter index 39 and differs at letter index 40 — its own last letter, where X_6 steps to 65 while c_α steps to 64. The mismatch is *internal to the first block*, so no exponent can affect it: every member of the family, for every $a, b \geq 1$, first differs from c_α at letter index 40 with depths $(65, 64)$, and Lemma 3.1 gives $\ell = \min(65, 64) = 64$. Confirmed on the full grid $a \leq 4$, $b \leq 3$ and at $(1, 20)$, $(10, 2)$. There is no recursion branch: the upper-leading orientation is exceptionless. ☐

The recycled residue is $|r_6| \approx 2^{62.37}$; the window-recycling theorem for this family is Theorem 9.3 with the upper-leading size bound of Section 10.

7 Family C: $X_7^aX_8^b$ — lower-leading saturation

Lemma 7.1. *$X_8 = X_7^5X_6$ exactly (an instance of Theorem 4.1 iterated along $a_8 = 5$); equivalently $\lfloor 485j/306\rfloor = \lfloor 84j/53\rfloor$ for $1 \leq j \leq 305$ with first failure at $j = 306$ (485 vs 484): the block X_8 follows the X_7-periodic floor pattern through its first 305 letters and departs only at its last letter. Verified exactly.*

Lemma 7.2. *$\lfloor j\alpha\rfloor = \lfloor 84j/53\rfloor$ for $1 \leq j \leq 358$, with first failure at $j = 359$: $\lfloor 359\alpha\rfloor = 569$ while $\lfloor 84 \cdot 359/53\rfloor = 568$. All 359 instances certified by exact integer power comparison.*

Theorem 7.3 (Lower-leading saturation)**.** *For all $a \geq 2$, $b \geq 1$: $\ell(X_7^aX_8^b) = p_7 + p_8 - 1 = 568$.*

Proof. By the periodic floor extension and Lemma 7.1, the grouped word holds the $\lfloor 84j/53\rfloor$-floors through $j \leq 53a + 305 \geq 411 > 359$ when $a \geq 2$; by Lemma 7.2, c_α departs those floors first, at $j = 359$, with depths $s_{359}(u) = 568$ and $s_{359}(c_\alpha) = 569$. Lemma 3.1 gives $\ell = 568$. Confirmed at $(2, 1), (2, 2), (3, 1), (4, 2)$. ☐

The $a = 1$ branch is the recursion: $X_7 X_8 = X_{(359,569)}$ (the semiconvergent word; measured depth $1053 = p_9 - 1$) and $X_7 X_8^2 = X_9$ (the next standard block, $a_9 = 2$), both Tier 1 word identities; for $b \geq 3$ the measured depth is $1538 = p_9 + p_8 - 1$ (Tier 3, proof template as above one level up). The recycled residue is $|r_7| \approx 2^{563.40}$.

8 Family D: the semiconvergent family $X_5^a X_s^b$

Lemma 8.1. *$\lfloor 46j/29 \rfloor = \lfloor 19j/12 \rfloor$ for $1 \leq j \leq 28$, with first failure at $j = 29$ (46 vs 45); equivalently $X_s = X_5^2 X_4$ exactly. Certified this session by direct integer evaluation.*

Theorem 8.2. *For all $a \geq 3$, $b \geq 1$: $\ell(X_5^a X_s^b) = 83$.*

Proof. By Lemma 8.1 the grouped word holds the $\lfloor 19j/12 \rfloor$-floors through $j \leq 12a + 28 \geq 64 > 53$ when $a \geq 3$; the c_α comparison is Lemma C of [1] (departure at $j = 53$, certified there), and Lemma 3.1 gives $\ell = \min(83, 84) = 83$. Confirmed on the grid $a \leq 6$, $b \leq 3$. □

The boundary members are again the recursion: $X_5 X_s = X_6$ ($a = 1$: the family member *is* X_6, upper law $\ell = 64$, observed) and $X_5^2 X_s = X_7$ ($a = 2$, $b = 1$: level-7 law, observed), both Tier 1 identities; the $(2, b \geq 2)$ sub-law at measured depth 129 is Tier 3.

9 Window recycling

Definition 9.1. $\Delta_{a,b} = aR_n + bR_{n+1}$ for the leading/second blocks; $\Delta^+ = \max(0, \Delta)$, $\Delta^- = \max(0, -\Delta)$.

Proposition 9.2 (Orientation-dependent size bounds)**.** *With L the word length and within-block Christoffel excursion at most 2 per side: lower-leading ($R_n < 0$): $\mathrm{osc} \leq a|R_n| + \Delta^+ + 4$; upper-leading ($R_n > 0$): $\mathrm{osc} \leq aR_n + \Delta^- + 4$; and in both cases*

$$|Q_{a,b}| \leq B_{a,b} = \frac{L}{3} \cdot \frac{2^{\mathrm{osc}+1}}{|1 - 2^{\Delta_{a,b}}|}.$$

Proof. The block-level prefix path runs from 0 along $\{iR_n\}$ to aR_n, then along $\{aR_n + iR_{n+1}\}$ to $\Delta_{a,b}$; in the lower-leading case the trough is at least $aR_n - 2$ and the peak at most $\Delta^+ + 2$; in the upper-leading case the peak is at most $aR_n + 2$ and the trough at least $-\Delta^- - 2$. Validated numerically on- and off-strip in both orientations with no violations. □

Theorem 9.3 (Window recycling)**.** *Let a family satisfy a saturation law at depth ℓ on a parameter region, let $r = \mathrm{cent}(-\Xi_\alpha \bmod 2^\ell)$ with $|r| < 2^{\ell-1}$, and let $B_{a,b}$ be the bound of Proposition 9.2. Every member of the region with $B_{a,b} < |r|$ is nonsingular.* Certified instances: *family B on $\log_2 L + aR_6 + \Delta^- + \log_2 |1 - 2^\Delta|^{-1} < 62.3$ (exact $\log_2 |r_6| = 62.375$, so the region constant is strictly below), and family C on the same form < 563.3 (exact $\log_2 |r_7| = 563.396$, strictly above the region constant) — the latter region covering $a \lesssim 2 \times 10^5$ and $|\Delta| \gtrsim 2^{-540}$: deeper level, deeper window, larger certificate.*

Proof. Singularity gives $Q_{a,b} \in \mathbb{Z}$ with $Q_{a,b} \equiv -\Xi_\alpha \pmod{2^\ell}$ by saturation; both $Q_{a,b}$ and r lie in the open centered interval $(-2^{\ell-1}, 2^{\ell-1})$, whose integers are pairwise distinct modulo 2^ℓ, so $Q_{a,b} = r$; but $|Q_{a,b}| \leq B_{a,b} < |r|$. □

residue	modulus	$\lvert r\rvert$	margin to ceiling (bits)
r_5	2^{83}	$2^{81.79}$	1.21
r_6	2^{64}	$2^{62.375}$	0.63
r_7	2^{568}	$2^{563.396}$	3.60
r_8	2^{484}	$2^{481.14}$	1.86
r_9	2^{25780}	$2^{25776.37}$	2.63
$r(148)$	2^{148}	$2^{145.64}$	1.36
$r(129)$	2^{129}	$2^{122.34}$	5.66
$r(1538)$	2^{1538}	$2^{1532.56}$	4.44

All eight computed residues are generic-sized. The *residue-genericity conjecture* — $|r_n| \geq 2^{\ell_n - c}$ for a moderate absolute c (all data fit $c = 6$) — is stated as Tier 4 only; a single small residue would shrink one family's certificate region without affecting the framework.

10 Odd-prime layer and shell-reusable blockers

The analyzer pins candidates; odd-prime divisibility of the odd wall decides attainment, and the trilogy's level-4 vacuous-window case [1] remains the standing proof that neither layer subsumes the other. Shell reuse operates on the odd side too: on the shell $(665, 1054)$, the prime $37 \mid \delta_9$ blocks the standard word ($C \equiv 24$), the grouped witness $X_5^{11}X_6^{13}$ ($C \equiv 17$) [1], and — by the identity $X_7X_8^2 = X_9$ — the family-C boundary member, which *is* the standard word. Same shell and wall, different carries: shell-level arithmetic is reusable, word-level arithmetic is not.

11 Computation protocol

Log schema per member: family, parameters (a, b), depth ℓ, first-disagreement letter index J, centered residue (per family, once), window status, blocking prime if recorded. Every reported ℓ is computed two independent ways — modular evaluation of $Q + \Xi_\alpha$ (deep-level proxy) and direct floor-disagreement search — and the two must agree; a discrepancy is treated as a bug detector, and none occurred. Exact grid scans: three 100×100 grids (families B, C, D), 30,000 members, *zero singular members*. Every scanned member of families B and C lies far inside the certified regions of Theorem 9.3 ($|Q|$ of order 2^{10}–2^{16} against thresholds 2^{62} and 2^{563}), so for those families the scans are confirmation of theorems; for family D and the sub-law strips they are evidence.

12 Theorem targets and conjectures

General adjacent-pair saturation (Tier 4): upper-leading $\ell = p_n - 1$ exceptionless; lower-leading $\ell = p_n + p_{n+1} - 1$ for $a \geq 2$ with the $a = 1$ branch governed by Theorem 4.1; three pairs confirmed, each new pair requiring only its own finite power certification. *General degeneracy law*: Theorem 4.1 already supplies it; the open part is only its systematic interaction with depth laws at semiconvergent shells. *Window recycling*: Theorem 9.3, extended family-by-family as size bounds are certified. *Residue genericity*: Tier 4 as stated. *Deep-compensation caution*: lattice points with $|\Delta|$ below the region floor renormalize to deeper shells and are excluded, not resolved.

13 Limitations

No Collatz proof is claimed. The analyzer does not cover arbitrary positive-entropy words, whose agreement depths need not stabilize. No n-uniform theorem across the tower is claimed: each certified family rests on its own finite, exact floor-lemma certification. The odd-prime layer remains independent and undecided in general. Small residues or deep compensation points, should they occur, would shrink certificate regions.

14 Tier ledger

Tier 1	Fundamental Lemma; family A theorem [1]; family B theorem (Thm. 6.2); family C theorem (Thm. 7.3); family D theorem (Thm. 8.2); the six degeneracy identities; all eight residues; all reported exact computations.
Tier 2	window-recycling certificates for families B and C (Thm. 9.3), through the archived Ξ_α [1] and the size bounds.
Tier 3	sub-law depths (148, 129, 1538, and the $a = 1$ measured laws); family-D and sub-strip grid evidence beyond theorem coverage.
Tier 4	general n-uniform taxonomy; residue-genericity conjecture; deep-compensation renormalization.

References

[1] E. De Jesús, *Near-Critical and Grouped Structures in the Collatz Carry Equation: A Series of Notes on Finite-Block Reduction, 2-Adic Windows, and Tower Products*, Zenodo, 2026. `https://doi.org/10.5281/zenodo.20557441`

[2] J. C. Lagarias, *The* $3x + 1$ *problem and its generalizations*, Amer. Math. Monthly **92** (1985), 3–23.

[3] J.-P. Borel, F. Laubie, *Quelques mots sur la droite projective réelle*, J. Théor. Nombres Bordeaux **5** (1993), 23–51.

[4] C. Reutenauer, *From Christoffel Words to Markoff Numbers*, Oxford University Press (2019), Cor. 2.4.2.

[5] J. Berstel, A. Lauve, C. Reutenauer, F. Saliola, *Combinatorics on Words: Christoffel Words and Repetitions in Words*, CRM Monograph Series **27**, Amer. Math. Soc. (2009).