Finite-Time Structural Descent in the Accelerated Collatz Map: Two-Step Vacancy Balances, Finite Compression, and Window Scores

Y.Ueoka
Independent Researcher

Tsumugi (ChatGPT)
AI Co-researcher

29 November 2026 — Preprint, version 1.0

Abstract

We organize previously established results in the ABN/CPE research programme on the accelerated Collatz map. Exact run dynamics give two-step vacancy-count identities and strict descent on unbounded families with short occupied runs and sufficiently long vacant runs. Finite-time compression of long vacant runs, and simultaneous compression of vacant and occupied runs, reduce certain infinite families to finite representative problems. A reproducible computer-assisted result gives strict vacancy descent within 54 steps when the normalized number of occupied cells is at most ten and the vacancy count is positive. These are structural descent results, not a proof of universal descent or of the Collatz conjecture. A finite-window score with a normalized occupied-count correction also admits exact increment-preserving compression. For window length six, exhaustive finite computations certify descent for every initial state with one, two or three vacant runs, respectively within 32, 46 and 103 steps. Positive-side upper connections also transport existing vacancy-descent certificates through arbitrarily many layers, including Mersenne upper blocks. These results are specific to the stated scores and domains; integration of different descent certificates is a separate problem.

Introduction and relation to earlier work

Ordinary stopping-time theory studies the first decrease of the integer itself, with finite parity vectors and residue classes as a natural way to organize finite segments of trajectories [2, 3]. Binary-run rewriting for the odd-to-odd map is also established [4, 5]; its relation to the present notation is explained in Paper A [1]. Our purpose here is different from inventing another encoding of the same map. The C/K coordinates expose the counts and lengths of occupied and vacant runs, and allow changes in these quantities to be calculated directly. In particular, a length extension can leave the entire finite-time increment of a structural score invariant. This permits a finite partition into cells with unbounded length coordinates, rather than a finite interval of starting integers. These are distinct forms of descent, and no global convergence or historical priority claim is inferred from the comparison.

1 Notation and inherited update rule

For positive odd n, set $F(n) = (3n+1)/2^{\nu_2(3n+1)}$ and $E(n) = n \,\mathrm{XOR}\, 2n$. In alternating binary notation (ABN), this mask records occupied cells of the alternating signed-power representation, not the ordinary binary digits of n. Its low-to-high maximal-run word (CPE) is

$$C(r_0)K(d_0)C(r_1)\cdots K(d_{m-1})C(r_m),$$

where every length is positive and $\sum r_i$ is even. Conversely any such word determines a unique positive odd integer by inverse adjacent XOR. Put $H_C = \sum r_i$, $Z = \sum d_i$, $B = H_C + Z$,

$p_i = r_i \bmod 2$, $\epsilon_i = p_i$ for $i > 0$, $\epsilon_0 = 1 - p_0$. Empty sums are zero. By Paper A (Theorem 4.1), $m = 0$ is the base set $F(n) = 1$; the division exponent is r_0.

The inherited full-word rule (Paper A, Theorem 7.1 and Section 7.1; Phase 6 P6-551/554) for $m > 0$ is

$$V_i = \begin{cases} 1^{1+\epsilon_i} 0^{d_i - 1 - \epsilon_i} 1\, 0^{r_{i+1}-1}, & d_i \geq 2, \\ 1^{2-\epsilon_i} 0^{r_{i+1}-1+\epsilon_i}, & d_i = 1. \end{cases} \tag{1}$$

Here $1^a, 0^a$ are repeated symbols and $a = 0$ means an empty word. Concatenate $V_0 \cdots V_{m-1} 1^{1+p_m}$ and regroup maximal runs. This exact rule is equivalent to the earlier ordinary-binary run rewrite discussed in Paper A; the present objective is to derive explicit descent theorems from it, not to claim a new integer map. In particular,

$$B(Fn) = B(n) - r_0 + 1 + p_m, \tag{2}$$

$$\Delta_1 Z = \sum_{i<m} [r_{i+1} - 2 + \epsilon_i (2\mathbf{1}_{d_i=1} - 1)]. \tag{3}$$

These equations apply for $m > 0$; the base case has output word 11.

Self-contained derivation of the update rule. Let x_j be the ordinary binary digits of n, extended by zero outside the finite support. The identity $n = \sum_{j \geq 0} (x_{j-1} - x_j) 2^j$ shows that the occupied signs alternate from -1 at the lowest cell. A chain beginning at k with sign σ contributes $\sigma 2^k (1 - (-2)^r)/3$. Multiplying by three and adding one cancels the initial constant from the lowest chain, leaving $(-2)^{r_0}$ and terms at strictly higher positions. Thus $\nu_2(3n+1) = r_0$; when $m = 0$ the even length gives $F(n) = 1$. For $m > 0$, after division the signed coefficients are at endpoints $a_0 = 0$, $b_i = a_i + d_i$, $a_{i+1} = b_i + r_{i+1}$. The coefficients at a_i, b_i are respectively $s_i = (-1)^{r_0 + \cdots + r_i}$ and $-s_i$ for $i < m$; the final coefficient at a_m is $+1$. Convert these signed coefficients to ordinary binary digits by scanning upward: at each cell add the incoming borrow to the coefficient, take the residue modulo two as the digit, and divide the difference by two for the next borrow. The borrow is always 0 or -1. Immediately below a_i, the digit and the borrow are v and $-v$ with $v \in \{0, 1\}$ (initially $v = 0$). Write $w = \mathbf{1}_{s_i = -1}$. At a_i the new digit is $1 - v$ and the outgoing borrow is $-w$. For $i > 0$, $v = \mathbf{1}_{s_{i-1} = +1}$, so $\mathbf{1}_{v=w} = \epsilon_i$ for all i, including the initial interval. The support mask records changes of consecutive ordinary digits. If $d_i \geq 2$, the first cell of the interval is occupied, the next is occupied exactly when $v = w$, and the digit is then constant up to b_i, where it changes from w to $1 - w$. It remains constant until a_{i+1}. If $d_i = 1$, the digit at b_i is $1 - w$, so the support there is $1 - \epsilon_i$. The outgoing borrow after b_i is $-(1 - w)$ in either case. These two cases yield exactly (1). At a_m the coefficient is positive and the preceding digit is 0 exactly when r_m is odd. The last support block is therefore 1^{1+p_m}. Counting the lengths and zeros of these blocks gives (2) and (3). This derivation also follows the independent R22 reviewer's argument, and is included here to remove dependence on Paper A for this identity.

For $t \geq 0$ write $\Delta_t Z = Z(F^t n) - Z(n)$. A strict decrease is always compared with the same initial evaluation. It does not imply that the endpoint belongs to the initial family.

Lemma 1.1 (Lowest-chain normalization, P6-351–353)**.** *Replacing r_0 by $\bar{r}_0 = 1$ for odd r_0, or 2 for even r_0, defines a map ρ preserving Z and satisfying $F\rho = F$, $\rho^2 = \rho$. The normalized occupied count is $H = H_C - r_0 + \bar{r}_0$.*

Proof. For $R_k(x) = 4^k x + (4^k - 1)/3$ one has $3R_k(x) + 1 = 4^k(3x + 1)$ and $E(R_k(x)) = (4^k - 1) + 4^k E(x)$, the latter by iterating the disjoint-bit identity $E(4x + 1) = 3 + 4E(x)$. Thus exactly the lowest C-run is extended by $2k$, and the next odd integer is unchanged. The shortened word still has an even number of occupied cells and defines a unique positive odd x, so the assertion follows. This is the earlier Phase 6 normalization, not a new construction. □

2 Two-step balances

Lemma 2.1 (Elimination of the intermediate occupied count, P6-583)**.** *If $m(n), m(Fn) > 0$, denote the data of Fn by primes. Then*

$$\Delta_2 Z = \sum_{i=1}^{m} r_i + 1 + p_m - r'_0 - \sum_{j<m'} \{2 + \epsilon'_j(1 - 2\mathbf{1}_{d'_j=1})\}. \tag{4}$$

If $m(n) > 0$ and $m(Fn) = 0$, the first step already gives strict Z descent.

Proof. Equation (2) gives $Z' - Z + H'_C = H_C - r_0 + 1 + p_m$. Add (3) for the next step and cancel H'_C. The base-set case follows from $Z' = 0$. □

Theorem 2.2 (All K-runs of length at least three, P6-576)**.** *Suppose $m \geq 1$ and $d_i \geq 3$ for every i. Define*

$$P = \sum_{i=1}^{m} (\lceil r_i/2 \rceil - 2)_+, \quad A = \#\{1 \leq i < m : r_i = 1, d_i = 3\},$$
$$L = \mathbf{1}_{\epsilon_0=1,\ d_0 \geq 4}, \qquad I = \mathbf{1}_{r_m=1}.$$

Then

$$\Delta_2 Z = -2m + 1 + I + 2(P + A - L). \tag{5}$$

In particular $\Delta_2 Z < 0$ if and only if $P + A - L \leq m - 1 - I$.

Proof. By (1) each old K-run survives with length $v_i = d_i - 1 - \epsilon_i \geq 1$, and a new run of length $q - 1$ follows it when $q = r_{i+1} > 1$. The output C-run just below the surviving run has length $1 + \epsilon_0$ at $i = 0$ and $1 + \epsilon_i + \mathbf{1}_{r_i=1}$ at $i > 0$. The next-step flags at these runs are consequently $e_0 = \epsilon_0$ and $e_i = 1 - \epsilon_i + \mathbf{1}_{r_i=1}$ for $i > 0$. When the new run exists, its lower and upper C-runs have lengths 1 and $1 + p_{i+1}$; otherwise the upper C-run of the surviving run has length 3. All these statements follow by concatenating (1), including the top word.

For $q \geq 2$, the first-step contribution in (3) is $q - 2 - \epsilon_i$. The next-step contributions of the surviving and new runs are respectively $-1 + e_i(2\mathbf{1}_{v_i=1} - 1)$ and $p_{i+1} - 2 + 2\mathbf{1}_{q=2}$. For $q = 1$, the first-step contribution is $-1 - \epsilon_i$, and only the surviving run contributes next, with value $1 + e_i(2\mathbf{1}_{v_i=1} - 1)$. Thus the sum per initial interval is

$$g(q) - \epsilon_i + e_i(2\mathbf{1}_{v_i=1} - 1), \qquad g(1) = 0, \quad g(q) = -1 + 2(\lceil q/2 \rceil - 2)_+ \ (q \geq 2).$$

For $i = 0$ the last two terms sum to $-2L$. For $i > 0$ they sum to -1 if $r_i \neq 1$, and to $-1 + (2\mathbf{1}_{d_i=3} - 1)$ otherwise. Summing g gives $-m + \#\{1 \leq i \leq m : r_i = 1\} + 2P$; the internal counts of $r_i = 1$ cancel. This proves (5); the strict inequality is an integer comparison. □

Corollary 2.3 (Unbounded short-C family, P6-577)**.** *If $m \geq 1$, every $d_i \geq 3$ and $1 \leq r_i \leq 4$ for $i \geq 1$, then $\min(Z(Fn), Z(F^2 n)) < Z(n)$. Neither the number nor the lengths of K-runs, nor the lowest C-length, have an upper bound.*

Proof. Here $P = 0$, $A \leq m - 1$, $L \geq 0$, so (5) is nonpositive. Equality requires $I = 1, A = m - 1, L = 0$, hence $r_1 = \cdots = r_m = 1$. In that case (3) gives $\Delta_1 Z = \sum_i (-1 - \epsilon_i) < 0$. □

Theorem 2.4 (K-lengths at least two and upper C-lengths at least two, P6-591)**.** *Suppose $m \geq 1$, $d_i \geq 2$ for all i, and $r_i \geq 2$ for $i \geq 1$. Let*

$$P = \sum_{i=1}^{m} (r_i - 4 + p_i)_+, \qquad A = \#\{1 \leq i < m : d_i \geq 3\}.$$

Then $\Delta_2 Z = P - 2A + c$, where

$$c = \begin{cases} -2\mathbf{1}_{r_1=2}, & p_0 = 0,\ d_0 = 2, \\ -3, & p_0 = 0,\ d_0 \geq 4, \\ -1, & \text{otherwise.} \end{cases}$$

Proof. Put $a_i = d_i - 1 - \epsilon_i$, $b_i = r_{i+1} - 1 > 0$. Each V_i ends with 0^{b_i}, so adjacent intervals do not merge their occupied runs. If $a_i > 0$, $V_i = 1^{1+\epsilon_i}0^{a_i}10^{b_i}$; if $a_i = 0$ then $V_i = 1110^{b_i}$. In (4), denote the subtracted amount for each output K-run g by $S(g) = 2 + \epsilon'(1 - 2\mathbf{1}_{g=1})$, with $S(0) = 0$ for an absent run.

For $i \geq 1$, the occupied run below b_i has odd length, so $S(b_i) = 3 - 2\mathbf{1}_{r_{i+1}=2}$. Directly separating $a_i = 0, 1, \geq 2$ gives $S(a_i) = 1 - p_i + 2\mathbf{1}_{d_i \geq 3}$. Writing $v = r_{i+1}$, $P_i = (v - 4 + p_{i+1})_+$, one has $v - S(b_i) = P_i + 1 - p_{i+1}$. The net internal contribution is therefore $P_i + p_i - p_{i+1} - 2\mathbf{1}_{d_i \geq 3}$, whose parity terms telescope to $p_1 - p_m$.

For $a_0 = 0$, the lowest output C-run is of length 3 and $S(b_0) = 2$. Adding the boundary term $1 + p_m - r'_0$ and the telescoped $p_1 - p_m$ to $v - S(b_0)$ gives $v - 4 + p_1 = P_0 - 2\mathbf{1}_{v=2}$. For $a_0 > 0$, $r'_0 = 1 + \epsilon_0$, $S(b_0) = 3 - 2\mathbf{1}_{v=2}$, and $S(a_0) = 2 + \epsilon_0(1 - 2\mathbf{1}_{a_0=1})$. The corresponding sum is $P_0 + 1 - \epsilon_0 - S(a_0)$. Substituting $\epsilon_0 = 1 - p_0$ gives the stated cases. There is always a nonempty b_0, so Lemma 2.1 applies. □

Corollary 2.5 (Short upper C-runs with K-length two allowed; part of P6-592)**.** *Under Theorem 2.4, assume $r_i \in \{2, 3, 4\}$ for $i \geq 1$. Then $\Delta_2 Z < 0$ except when $p_0 = 0$, $r_1 \in \{3, 4\}$ and all $d_i = 2$. In the exceptional case $\Delta_2 Z = 0$ and $F(n) < n$.*

Proof. Here $P = 0$; the nonpositive terms $-2A$ and c both vanish precisely under the exceptional conditions. The division exponent r_0 is then even and at least two; $n > 1$ since $m > 0$, so $F(n) \leq (3n + 1)/4 < n$. □

The exceptional integer descent is not a Z-descent certificate. The additional endpoint bound and transfer to χ_G in P6-592 will be treated separately from this corollary.

3 Finite compression

Theorem 3.1 (K-extension invariance, P6-565)**.** *Fix $T \geq 1$. Extend any collection of K-runs by nonnegative lengths a_i, subject to $a_i > 0 \Rightarrow d_i \geq 2T + 1$. Denote the extended state by $\widetilde{n}$ and $a = \sum a_i$. For $0 \leq t \leq T$, the ordered C-lengths are identical in both orbits, and only the descendants of marked K-runs differ in length, by their respective a_i. Their unextended lengths are at least $2(T - t) + 1$. In particular the differences between the extended and unextended states in (Z, H_C, B) are $(a, 0, a)$ at every such time.*

Proof. Existence of the extended state follows from the unchanged even occupied count. At $t < T$, each marked run has length at least three in both states. In (1) only the factor $0^{d_i - 1 - \epsilon_i}$ changes length; its unextended length is at least $2(T - t) - 1 \geq 1$. It is bounded by occupied cells, so cannot disappear or merge with another K-run. All other factors and the top word agree; maximal C-runs and the correspondence of K-runs are preserved. Induction on t proves the assertions. □

Thus clipping each d_i to $\min(d_i, 2T+1)$ preserves all $\Delta_t Z$, $\Delta_t H_C$ and $\Delta_t B$ for $t \leq T$, including the first strict-descent time if it lies in that window. Constant-coefficient affine combinations inherit this increment invariance; arbitrary state-dependent weights do not automatically inherit it.

Theorem 3.2 (Simultaneous C/K extension, P6-568)**.** *In addition to the K-extensions of Theorem 3.1, extend upper C-runs ($i \geq 1$) by even $b_i \geq 0$ with $b_i > 0 \Rightarrow r_i \geq 2T$. Set $b = \sum b_i$. Then the differences in $(Z, H, H + Z)$, extended minus unextended, are $(a, b, a + b)$ at time zero and $(a + b, 0, a + b)$ at all $1 \leq t \leq T$.*

Proof. The even C-extensions do not alter flags, the lowest division exponent or the top word. By (1) a marked upper C-run extends its new K-run by b_i; its original length is at least $2T - 1$.

The surviving marked K-runs also have length at least $2T-1$. They remain distinct positive-length runs bounded by occupied cells, and the first-step C-length lists agree. For $T>1$, apply Theorem 3.1 to these descendants for the remaining $T-1$ steps. Normalizing the lowest C-run does not alter the differences since its length is unchanged initially and agrees thereafter. □

Consequently, for a fixed m, T, clip K-lengths at $2T+1$ and upper C-lengths at $2T$ (even) or $2T+1$ (odd), preserving parity. For $Z+\lambda H$ the increment difference at $t \geq 1$ equals $(1-\lambda)b$. In particular $\lambda = 1$ gives exact invariance and $\lambda \geq 1$ transports strict descent from the compressed representative to all extensions. This is not an assertion that all representatives descend.

4 A computer-assisted infinite-family theorem

Theorem 4.1 (Bounded normalized occupied count, P6-566)**.** *Every positive odd n with $Z(n) > 0$ and $H(n) \leq 10$ satisfies $Z(F^t n) < Z(n)$ for some $1 \leq t \leq 54$. This result is computer-assisted; there is no upper bound on the initial integer or K-lengths.*

Proof. By Lemma 1.1, use only $r_0 \in \{1, 2\}$. The possible even occupied counts are $H = 2, 4, 6, 8, 10$. For each, enumerate every ordered composition with $r_0 \leq 2$ and at least two parts. For each composition, at horizon $T = 1$ take every K-coordinate in $\{1, 2, 3\}$: coordinate 3 represents $[3, \infty)$, others are singleton sets.

At horizon T represent each surviving cell by its tuple with coordinates at most $c = 2T+1$, where c denotes $[c, \infty)$. Evaluate the representative by exact integer iteration and direct XOR mask, comparing every Z with its own initial value. A successful representative certifies its entire cell by Theorem 3.1. In every failed cell split each tail $[c, \infty)$ into $\{c\}, \{c+1\}, [c+2, \infty)$ and take the product; singleton coordinates stay fixed. The resulting cells are disjoint and cover exactly the previously uncertified set. By induction this is an exhaustive cover, not random sampling. The original exact-integer code R103_VERIFY.py was recovered from the Phase 6 archive and rerun without modification for this draft. Its validation JSON matches the original byte for byte. Results:

H	Checks	Certified cells	Final horizon	Unresolved
2	3	3	1	0
4	77	64	5	0
6	1554	1114	40	0
8	21518	18152	54	0
10	335288	296262	51	0
Total	358440	315595	54	0

Every cell is certified by horizon 54, and the normalization preserves the original initial Z and the subsequent orbit. This proves the theorem conditional on the exact finite enumeration, which is reproducible from the archived code. No optimality of the bound is asserted. □

5 Finite-window structural scores

The finite-window construction from Phase 6 R171 uses a sum of future vacancy counts, together with an initial occupied-count correction. It is not a probabilistic assumption about trajectories. Write $n_j = F^j(n)$ and define, for every integer $L \geq 1$,

$$W_L(n) = \sum_{j=0}^{L-1} Z(n_j), \qquad \Xi_L(n) = W_L(n) + H(n) - 2. \tag{6}$$

The normalization H is as in Lemma 1.1; in particular $H(1) = 2$. Define $\widehat{B} = H + Z$.

Lemma 5.1 (Finite-window sums, P6-779)**.** *The quantity W_L is a nonnegative integer, vanishes exactly on $Z = 0$, and satisfies $W_L \circ \rho = W_L$ and*

$$W_L(Fn) - W_L(n) = Z(n_L) - Z(n). \tag{7}$$

For $L \geq 2$, $\widehat{B}(n) \leq 3W_L(n) + 3$. Thus its sublevel sets on normalized states are finite. Along a trajectory that is not eventually periodic, $W_L(n_j) \to \infty$ for $L \geq 2$.

Proof. The first statements follow by normalization, the base-set characterization, and telescoping. For the width bound we may first replace n by $\rho(n)$, since W_L and $\widehat{B}$ are invariant and $F\rho = F$. For $m > 0$, every factor in (1) has at most three occupied cells and the top has at most two. Hence $H_C(Fn) \leq 3m + 2$. The width equation and $\bar{r}_0, 1 + p_m \in \{1, 2\}$ give

$$\widehat{B}(n) \leq B(Fn) + 1 \leq 3m + Z(Fn) + 3 \leq 3Z(n) + Z(Fn) + 3 \leq 3W_L(n) + 3.$$

For $m = 0$ we have $\widehat{B} = 2$. On the normalized states a width bound is a bound on the integer itself. If a trajectory not eventually periodic had a bounded subsequence of W_L, the corresponding normalized states would lie in a finite set. Their one-step images would then also lie in a finite set because $F\rho = F$; a repeated image implies eventual periodicity, a contradiction. □

Theorem 5.2 (Finite-window compression, P6-779–780)**.** *For every $L \geq 1$, Ξ_L is a nonnegative integer, vanishes exactly on $Z = 0$, is invariant under ρ, and has finite sublevel sets on normalized states. Fix $T \geq 1$ and put $U = T + L - 1$. Apply the simultaneous extensions of Theorem 3.2 with horizon U, of total K-extension a and upper C-extension b. For $1 \leq t \leq T$,*

$$\begin{aligned} W_L(\widetilde{n}) - W_L(n) &= La + (L-1)b, \\ W_L(F^t\widetilde{n}) - W_L(F^t n) &= L(a+b), \\ \Xi_L(F^t\widetilde{n}) - \Xi_L(F^t n) &= L(a+b), \end{aligned}$$

where the last equality holds at $t = 0$ as well. In particular, every Ξ_L increment through T is preserved by clipping the K-lengths at $2U + 1$ and upper C-lengths at $2U$ or $2U + 1$ according to parity.

Proof. Since $H \geq 2$ and $H = 2$ on the base set, the vanishing and nonnegativity assertions follow from Lemma 5.1. Also $\widehat{B} \leq \Xi_L + 2$, and all terms are normalization-invariant. Theorem 3.2 gives a difference a in Z at time zero and $a + b$ at each positive time through U; the difference in H is b at time zero and zero thereafter. Sum the L terms of the window at $t = 0$ and at $1 \leq t \leq T$ and add the initial H correction. The largest time needed is $T + L - 1 = U$. The stated clipping produces an admissible unextended word with the same occupied parity. Its increments therefore equal the original increments. □

The correction is consequential: for the uncorrected score the increment difference is b, rather than zero. In particular, for any fixed L, T the two-chain word with $r_0 = 2$, $r_1 = 2U$, $d_0 = 1$ permits arbitrarily large even b; choosing b larger than the magnitude of every negative increment of the unextended word rules out strict W_L descent through T. This is a finite-horizon statement, not perpetual non-descent for a fixed initial state.

Lemma 5.3 (One-step width balance for the finite window, P6-803)**.** *Write $r_{0,j}$ for the lowest C-length of n_j and $g_j = 1 + (r_{m,j} \bmod 2)$ if $m(n_j) > 0$, and $g_j = 2$ otherwise. Then*

$$\Xi_L(Fn) - \Xi_L(n) = H(n_1) - H_C(n_L) + \sum_{j=0}^{L-1} g_j - \bar{r}_{0,0} - \sum_{j=1}^{L-1} r_{0,j}. \tag{8}$$

In particular, $H_C(n_L) + \bar{r}_{0,0} + \sum_{j=1}^{L-1} r_{0,j} > H(n_1) + 2L$ suffices for strict one-step descent.

Proof. Equation (2) extends to the base case with the stated g_j, since then $B(n_j) = r_{0,j}$ and $B(n_{j+1}) = 2$. Summing it gives

$$B(n_L) - \widehat{B}(n_0) = \sum_{j=0}^{L-1} g_j - \bar{r}_{0,0} - \sum_{j=1}^{L-1} r_{0,j}.$$

Use (7) and $Z(n_L) = B(n_L) - H_C(n_L)$ in the difference of (6). The sufficient condition follows from $g_j \le 2$. □

When $m(n) > 0$, the sufficient condition also holds if its right-hand side is replaced by $3m(n) + 2 + 2L$, using the bound in Lemma 5.1. Neither sufficient condition is asserted to hold universally.

6 Computer-assisted descent for one, two and three K-runs

Theorem 6.1 (Fixed-window descent, P6-781, P6-782, P6-802)**.** *For every positive odd integer n with $m(n) = 1, 2, 3$, respectively, there exists an integer $1 \le t \le T_m$ such that*

$$\Xi_6(F^t n) < \Xi_6(n), \qquad (T_1, T_2, T_3) = (32, 46, 103).$$

There is no restriction on the initial run lengths. This is a computer-assisted theorem, whose finite checks are specified below.

Proof. Normalize the lowest chain. For each choice of the m upper C-lengths the lowest normalized length is 1 if their sum is odd, and 2 otherwise. At horizon T put $c = 2(T + 5) + 1$. Theorem 5.2 shows that it is enough to consider every coordinate among $1, \dots, c$. For a K-coordinate, c denotes the entire tail $[c, \infty) \cap \mathbb{N}$. For a C-coordinate, $c-1, c$ denote the two tails of their respective parities, and smaller values are singletons. Thus the cells are disjoint and cover the entire normalized domain. Their representatives have the same increments as all of their members through T.

For $m = 1$ use $T = 32$, $c = 75$, and test all $75^2 = 5625$ representatives. For $m = 2, 3$ start instead with $T = 1$, $c = 13$, and test 13^{2m} representatives. For a representative form the occupied mask by concatenating the runs, reconstruct the ordinary binary digits by low-to-high prefix XOR (or evaluate the alternating signed-power sum), and iterate the exact integer map. Compute (6) through T using states up to $F^{T+5}n$. If a strict decrease occurs, certify that entire cell. Otherwise refine the cell for the next horizon: each K-tail with lower endpoint c splits into the singletons $c, c+1$ and the tail starting at $c+2$; each C-tail with lower endpoint $v = c-1$ or c splits into the singleton v and the same-parity tail starting at $v+2$. Take the Cartesian product of these disjoint refinements. Singleton coordinates are unchanged. Normalized lowest-chain parity is unchanged by C-tail refinements. Thus the uncertified cells continue to partition exactly the part of the domain not yet certified. There is no implicit cutoff of infinite tails.

The finite computation yields the following counts. "Certified" counts disjoint cells, not integers or a density; "checks" counts every evaluation of a representative, including those subsequently refined.

m	Checks	Certified	Final horizon	Unresolved
1	5,625	5,625	32	0
2	30,709	28,674	46	0
3	4,890,896	4,832,236	103	0

For the three-run-count case, the initial pass has $13^6 = 4,826,809$ representatives, of which $4,811,375$ pass and $15,434$ remain. The initial mask width is at most $2 + 6 \cdot 13 = 80$; $F(n) \le 2n$

for every positive odd n, so 128-bit unsigned arithmetic suffices for the initial six steps, including masks and the multiplication before division. All subsequent steps in the computation use arbitrary-precision integers. A fresh C++ initial pass and an explicit-endpoint Python refinement (including independent signed-power reconstruction and direct window evaluation) reproduce the original counts at every horizon and the entire first-descent histograms. The one- and two-band checks use the same Python evaluation throughout. These computations certify every cell at the stated horizons; Theorem 5.2 transfers the result to all members and Lemma 1.1 to all original lowest-chain lengths. □

Corollary 6.2 (Early two-band coverage, P6-783)**.** *For $m = 2$, if $r_1 \geq 9$ or $r_2 \geq 7$ or $d_0 \geq 12$ or $d_1 \geq 8$, then strict Ξ_6 descent occurs by time eight.*

Proof. The computation in Theorem 6.1 leaves at horizon eight exactly 67 cells, all singletons, with coordinate-wise maxima $(r_1, r_2, d_0, d_1) = (8, 6, 11, 7)$. These facts were rechecked with explicit cell bounds; thus every state outside that box has already been certified. The normalized lowest C-length does not constrain the original one. □

For reproducibility, the five two-band representatives still uncertified at horizon 32 are displayed below. The first column is the normalized integer, not a restriction on the original initial value.

n	r_1	r_2	d_0	d_1	First descent
55	2	1	2	1	34
27	2	1	1	1	35
97	1	1	3	1	36
6121	3	3	1	5	38
1055	1	2	4	4	46

For $m = 3$, at horizon 19 only 317 singleton cells remain; the last is the normalized integer 1563647, with first descent at 103. For $m = 1$, the first-descent histogram is $1 : 5551$, $2 : 71$, $4 : 1$, $31 : 1$, $32 : 1$. The original R172 initial implementation had a delayed-evaluation error in refinement; the corrected algorithm and explicit endpoint implementation, not the erroneous result, support the theorem. In this draft the new checks were written without importing the original research modules. The full stage logs are retained in the archive. Appendix A supplies a single standalone verifier of the two main exhaustive covers. The early-coverage assertion is also checked by the coordinate-maxima test at horizon eight.

7 Transport of descent certificates through upper connections

Finite compression enlarges a certificate by varying lengths while retaining a fixed number of initial runs. A distinct mechanism, developed in Phase 6 R131–134, allows the number of occupied cells and runs to increase without bound. We present its positive-side version. In this section $L_2(x)$ denotes the ordinary binary length of $x > 0$ (not the window length L), and all occupied counts in the balance are the unnormalized H_C.

Lemma 7.1 (Positive-side separation and first-contact balance; P6-655, 658, 662)**.** *Let b, q be positive odd integers, $k = B(b) + d$ with $d \geq 1$, and $n = 2^k q + b$. Write $b_t = F^t(b)$, $S_t = \sum_{j<t} \nu_2(3b_j + 1)$, $X_t = 2^{S_t} b_t$, and $u_t = 3^t q$. Define $d_t = k - L_2(X_t) - 1$ and $\tau = \min\{t \geq 1 : d_t \leq 0\}$. Then $d_0 = d$, $1 \leq d_t - d_{t+1} \leq 2$, and $d_\tau \in \{0, -1\}$. For $0 \leq t \leq \tau$,*

$$\begin{aligned} F^t(n) &= 2^{k - S_t} u_t + b_t, \\ \Delta_t Z(n) &= \Delta_t Z(b) + L_2(u_t) - L_2(q) - L_2(X_t) + L_2(b) \\ &\quad + H_C(q) - H_C(u_t) + \mathbf{21}_{d_t = -1}. \end{aligned}$$

In particular $d \geq 2T - 1$ guarantees $T \leq \tau$, and $d \geq 2T$ also guarantees $d_T \geq 0$.

Proof. The lifted recursion is $X_{t+1} = 3X_t + 2^{S_t}$, with $3X_t < X_{t+1} \le 4X_t$; hence binary length grows by one or two per step. Before contact, $k - S_j \ge B(b_j) + 1$. The valuation $s_j = \nu_2(3b_j + 1)$ is at most $B(b_j)$ (equivalently the lowest C-length), so the upper term in $3F^j(n) + 1$ is divisible by 2^{s_j+1}. Exact division by 2^{s_j} gives the asserted formula inductively, including the step entering contact. At every $t \le \tau$, $K = k - S_t \ge 1$ and $b_t < 2^K$, because $d_t = K - L_2(b_t) - 1 \ge -1$. The ordinary binary blocks are disjoint. Adjacent XOR gives

$$E(F^t(n)) = E(b_t)\,\mathrm{XOR}\,2^K E(u_t).$$

The only possible overlap is at cell K when $d_t = -1$; both occupied cells then cancel. The top bit remains unchanged since $E(u_t)$ has width at least two. Thus $Z(F^t(n)) = Z(b_t) + Z(u_t) + d_t + 2\mathbf{1}_{d_t=-1}$. Subtract the initial identity and use $B = L_2 + 1$ and $d_t - d = -L_2(X_t) + L_2(b)$. Finally, $d_j \ge d - 2j$ gives the stated horizon guarantees. □

The formula holds through the first contact only; it makes no assumption of renewed separation afterwards. The upper trajectory is pure multiplication by 3, not the accelerated Collatz trajectory of q.

Theorem 7.2 (Transport through a top C(2), P6-663)**.** *In Lemma 7.1, take $q = 1$, so the added upper word is $K(d)C(2)$. For $1 \le t \le \tau$,*

$$\Delta_t Z(n) \le \Delta_t Z(b) - 2\mathbf{1}_{t\ge 2} + 2\mathbf{1}_{d_t=-1}.$$

If $1 \le t \le \tau$, $\Delta_t Z(b) < 0$, and either $t \ge 2$ or $d \ge 2$, the same time certifies strict descent of n.

Proof. Since $X_t \ge 3^t b$, $L_2(X_t) - L_2(b) \ge L_2(3^t) - 1$: indeed $b \ge 2^{L_2(b)-1}$ and $3^t \ge 2^{L_2(3^t)-1}$. This cancels the upper length increase in the balance for $q = 1$. The occupied count of 1 is two, and an odd positive integer has occupied count two exactly when it is $2^a - 1$ for some $a \ge 1$ (one ordinary binary run of ones). If $3^t + 1$ is a power of two, then for even $t \ge 2$ it is 2 (mod 8) and larger than two, and for odd $t \ge 3$ it is 4 (mod 8) and larger than four. Thus $H_C(3^t) \ge 4$ for $t \ge 2$, while $H_C(3) = 2$. The bound follows. At $t = 1, d \ge 2$ there is no overlap; at $t \ge 2$ the loss of two occupied cells at contact is compensated by the upper occupied-count increase. □

Corollary 7.3 (Repeated C(2) connections, P6-664)**.** *Suppose $\Delta_T Z(b) < 0$ for a positive odd b and an integer $T \ge 1$. Let $n_0 = b$ and, for any finite number of layers, set*

$$n_j = 2^{B(n_{j-1})+d_j} + n_{j-1}, \qquad d_j \ge \max(2, 2T - 1).$$

Then $\Delta_T Z(n_j) < 0$ for every layer. In particular, if $Z(b) > 0$, $H(b) \le 10$ and each $d_j \ge 107$, strict Z descent is guaranteed within 54 steps, with no bound on the number of layers. If $Z(b) = 0$, $d_1 \ge 4$ and subsequent $d_j \ge 3$, every n_j with $j \ge 1$ has strict descent at time two.

Proof. The gap condition and Lemma 7.1 allow Theorem 7.2 to be applied at T at each layer. Induct on j. The 54-step assertion uses Theorem 4.1 for the initial block and the inequality $\max(2, 2T - 1) \le 107$ for $T \le 54$. In the base case $F(b) = 1$ and $\Delta_2 Z(b) = 0$. With $d_1 \ge 4$ there is no overlap at time two, so the bound in Theorem 7.2 is at most -2. Transfer this strict witness to all subsequent layers with $d_j \ge 3$. □

For example, $27 \mapsto 41$ decreases Z from two to one. Adding any finite sequence of upper $K(d)C(2)$ blocks with $d \ge 2$ preserves this one-step certificate. The 54-step corollary is computer-assisted because it uses Theorem 4.1; the transport itself is analytic.

Lemma 7.4 (Pure tripling of Mersenne blocks, P6-665)**.** *For $q = 2^a - 1$, $a \ge 1$, and $t \ge 1$, one has $H_C(3^t q) \ge 4$ except at $(a, t) = (1, 1), (3, 2)$, when the count is two. Initially $H_C(q) = 2$.*

Proof. It is enough to solve $3^t(2^a-1)=2^s-1$. The case $a=1$ was handled in Theorem 7.2. For $a\geq 2$ one has $s>a$ and $a\mid s$: division of s by a leaves a remainder $v<a$ and reduction modulo 2^a-1 gives $2^v-1=0$. Put $A=2^a$, $h=s/a\geq 2$. Then $3^t=1+A+\cdots+A^{h-1}$. If $A\equiv 1 \pmod 3$, the valuation of this sum is $\nu_3(h)$. To see this, write $h=3^v w$ with $3\nmid w$; the factor for the initial w terms is nonzero modulo three, and each subsequent cubing contributes exactly one power of three, since $1+C+C^2\equiv 3 \pmod 9$ for $C\equiv 1 \pmod 3$. A pure power of three with that valuation is at most h, whereas the sum is larger than h. If $A\equiv -1 \pmod 3$, the sum is divisible by three only for even $h=2v$. Factor it as $(1+A)\sum_{i=0}^{v-1}A^{2i}$. The second factor must itself be a power of three; the preceding argument for $A^2\equiv 1 \pmod 3$ forces $v=1$. Thus $2^a+1=3^t$ with odd $a\geq 3$. Modulo eight, t is even. The two positive powers of two $3^{t/2}-1$ and $3^{t/2}+1$ differ by two, hence are two and four; so $a=3, t=2$. Both exceptional cases are direct identities. All other occupied counts are even and at least four. □

Theorem 7.5 (Multilayer balance with general positive upper blocks, P6-666–667)**.** *For $T\geq 1$ define the integer contribution of a positive odd upper block q by*

$$\gamma_T(q)=L_2(3^Tq)-L_2(q)-L_2(3^T)+1+H_C(q)-H_C(3^Tq).$$

For $n_0=b$ and $n_j=2^{B(n_{j-1})+d_j}q_j+n_{j-1}$ with $d_j\geq 2T$, every finite J satisfies

$$\Delta_T Z(n_J)\leq \Delta_T Z(b)+\sum_{j=1}^{J}\gamma_T(q_j).$$

If the right-hand side is negative, the final state has strict Z descent at T, without requiring descent of each intermediate layer.

Proof. The lower-block lifted length satisfies $L_2(X_T)-L_2(b)\geq L_2(3^T)-1$, as in Theorem 7.2. The gaps ensure no overlap at T, so Lemma 7.1 bounds each increment by the lower increment plus $\gamma_T(q_j)$. Iterate this inequality. The length-only part of γ_T is zero or one, by the binary-length bounds on a product of positive integers. □

Corollary 7.6 (Mersenne connections and compensation, P6-667)**.** *If every q_j is Mersenne, then $\gamma_T(q_j)\leq 0$ for all T, and $\gamma_T(q_j)\leq -1$ for $T\geq 3$. Thus a strict lower Z witness transfers through arbitrarily many such blocks with $d_j\geq 2T$. If $T\geq 3$, the upper bound improves by at least one per layer. In particular, $Z(b)>0$, $H(b)\leq 10$, and $d_j\geq 108$ guarantee descent within 54 steps. For $Z(b)=0$, $d_j\geq 6$ guarantees descent at time three for any $J\geq 1$.*

Proof. Use Lemma 7.4 and the length-only bound of one; in the two exceptional pairs the length-only term is zero by direct calculation. Apply Theorem 7.5, Theorem 4.1 for the bounded-H statement, and $F(b)=1$ for the base case. □

The general bound also permits compensation between unlike blocks: at $T=1$, $\gamma_1(5)=2$ and $\gamma_1(3)=-1$. Starting with $b=3$, whose one-step increment is -1, append a copies of $q=5$ and c copies of $q=3$ in any order with gaps at least two. If $c\geq 2a$, the total increment is at most $-1+2a-c\leq -1$. This is a sufficient analytic condition even though individual layers need not descend.

The gap hypotheses and block restriction matter (P6-668). For $b=3, q=5, d=2$, the lower increment is -1 at $T=1$ but $163\mapsto 245$ has zero increment. With $b=3, q=7, d=3, T=2$, the lower increment is again -1, whereas $451\mapsto 677\mapsto 127$ has zero increment. Thus the general Mersenne transport cannot replace $2T$ by $2T-1$ indiscriminately. These examples concern the specified witness time, not perpetual non-descent. The sharper top-C(2) bound remains valid. The general signed decomposition, negative-side dynamics, and multi-generation isolation are reserved for a separate treatment; the transfer to the common score χ_{132} is reserved for the gluing paper.

8 Scope, provenance and next material

These theorems establish strict descent of specified quantities on specified domains, not universal descent, invariance of the domains under the endpoint maps, or convergence of each member by induction on the same restricted domain. The base set $Z = 0$ and well-ordering of the nonnegative integers would permit a global descent induction only after a universal descent statement (or an appropriate common induction) is proved.

The Z and Ξ_6 certificates are distinct. Even though both scores have the same zero set, a strict decrease of one is not automatically a strict decrease of the other. The endpoint of Theorem 6.1 is not required to remain in $m \leq 3$. These domain-restricted theorems do not give universal descent. Existing χ and Ψ certificates, including P6-1063, are preserved for the companion work on gluing. No claim about their coverage under the present scores is made.

Research provenance and contributions. This paper continues the ABN/CPE programme introduced by Y.Ueoka and the companion Paper A [1]. It reconstructs Phase 6 results from R47, R99–100, R103–104, R107, R109, R112, R131–134, R171–172 and R181; Appendix B records the correspondence and documented roles. Y.Ueoka proposed the structural research directions, including local CPE analysis of coalescence, isolation and pure powers of three, the finite-window averaging direction, and the general gluing and variational-induction framework. Historical ChatGPT sessions derived and proved the specific results assembled here and performed the original finite computations, as recorded in those sources. Tsumugi (ChatGPT), an AI co-researcher, reconstructed the proofs, implemented and verified the finite covers, and prepared the present manuscript with Y.Ueoka. The separate R22 ChatGPT reviewer independently reproduced the finite-cover computations and provided the endpoint derivation in Section 1. The historical sessions are not retrospectively identified with the present pen name. These are research attributions, not a claim of historical priority or an independent human peer review. Y.Ueoka handles submission and human-side publication responsibilities; the research roles above are distinct from platform metadata.

A Standalone finite-cover verifier

The following Python 3.10+ program uses exact integer arithmetic, no external files, and no external packages. The R22 independent reviewer ran both the extracted listing and a separately written verifier under Python 3.13.5, reproducing every reported count. The author-side listing and its R22 verification log are also embedded as PDF file attachments (where supported by the PDF viewer). The reported independent replay concerns R22; the present version retains its mathematical results and executable listing. The author-side complete replay predates the final editorial revision. Save the listing as a Python file and run it without the `-O` flag. It enumerates the cells in Theorems 4.1 and 6.1, checks every representative and the two-band early-coverage box, and raises an exception if any remain unresolved at the stated horizon or if the counts disagree. The score computation omits the constant -2, which cancels in every comparison. The finite computation alone is not the proof of the infinite-domain statements: the partition and increment-preservation theorems in the main text supply that implication.

```
# Python 3.10+, no external data or packages. Do not use -O.
from itertools import product

def encode(r, d):
    e = p = 0
    for i, v in enumerate(r):
        e |= ((1 << v) - 1) << p
        p += v + (d[i] if i < len(d) else 0)
    n, s = e, 1
```

```
    while s < p:
        n ^= n << s
        n &= (1 << p) - 1
        s *= 2
    assert n > 0 and n & 1 and n ^ (n << 1) == e
    return n

def step(n):
    x = 3*n + 1
    return x // (x & -x)

def info(n):
    e = n ^ (n << 1)
    h = e.bit_count()
    r = (e ^ (e + 1)).bit_length() - 1
    return e.bit_length() - h, h-r+(1 if r%2 else 2)

def witness(n, T, window):
    if not window:
        z = info(n)[0]
        for t in range(1, T+1):
            n = step(n)
            if info(n)[0] < z:
                return t
    else:
        a = [info(n)]
        for _ in range(5):
            n = step(n)
            a.append(info(n))
        z = sum(v[0] for v in a)+a[0][1]
        for t in range(1, T+1):
            n = step(n)
            a.append(info(n))
            if sum(v[0] for v in a[t:t+6])+a[t][1] < z:
                return t
    return 0

def compositions(h):
    if h == 0:
        yield ()
    else:
        for v in range(1, h+1):
            for tail in compositions(h-v):
                yield (v,) + tail

def split(r, d, cap, window):
    # Upper C tails have fixed parity. K tails are unrestricted.
    ra = [(v, v+2) if window and i and v >= cap-1
          else (v,) for i,v in enumerate(r)]
    da = [(v,v+1,v+2) if v == cap else (v,) for v in d]
    for rr in product(*ra):
        for dd in product(*da):
            yield rr, dd

def cover(active, start, stop, window):
    checks = certified = 0
    for T in range(start, stop+1):
        cap = 2*(T+5)+1 if window else 2*T+1
        bad = []
```

```
        for r,d in active:
            checks += 1
            if witness(encode(r,d), T, window):
                certified += 1
            else:
                bad.append((r,d))
        if window and T == 8 and bad and len(bad[0][0]) == 3:
            assert len(bad) == 67
            coords = [r[1:] + d for r,d in bad]
            assert tuple(map(max,zip(*coords))) == (8,6,11,7)
            # Maxima below both tail thresholds imply singletons.
        if not bad:
            return checks, certified, T
        active = [v for r,d in bad for v in split(r,d,cap,window)]
    raise AssertionError('Unresolved cells')

def window_cells(m, cap):
    for upper in product(range(1,cap+1), repeat=m):
        r = (1 if sum(upper)%2 else 2,) + upper
        for d in product(range(1,cap+1), repeat=m):
            yield r,d

if __name__ == '__main__':
    for h,expected in zip((2,4,6,8,10),
          ((3,3,1),(77,64,5),(1554,1114,40),
           (21518,18152,54),(335288,296262,51))):
        a = ((r,d) for r in compositions(h) if len(r)>1
             and r[0]<=2 for d in product(range(1,4),repeat=len(r)-1))
        result = cover(a,1,54,False)
        assert result == expected
        print('H',h,result,flush=True)
    for m,start,stop,expected in (
        (1,32,32,(5625,5625,32)),
        (2,1,46,(30709,28674,46)),
        (3,1,103,(4890896,4832236,103))):
        result = cover(window_cells(m,2*(start+5)+1),start,stop,True)
        assert result == expected
        print('m',m,result,flush=True)
```

B Source concordance and verification record

The following identifiers refer to the historical Phase 6 research archive, not to independent publications. The original proof and attribution records are preserved separately by the researchers; the statements and proofs needed for the results in this paper are supplied in the present PDF. A source identifier denotes the relevant result or restricted part of it, not a claim that the earlier result is reproduced in full. Equations (1)–(3) also draw on R99–100 (P6-551, 554), as in Paper A.

Present result	Phase 6 result (P6-)	Research revision (R)
1.1	351–353	47
2.1	583	109
2.2	576	107
2.3	577	107
2.4	591	112
2.5	592 (part)	112
3.1	565	103
3.2	568	104
4.1	566	103
5.1	779	171
5.2	779–780	171
5.3	803	181
6.1	781, 782, 802	171, 172, 181
6.2	783	172
7.1	655, 658, 662 (positive side)	131, 132, 133
7.2	663	133
7.3	664	133
7.4	665	134
7.5	666–667	134
7.6	667	134

Documented division of research. R47 records Y.Ueoka's proposal to study coalescence from partial CPE words and ChatGPT's explicit connection to the support mask and structural counts. R99–112 build on Y.Ueoka's structural-analysis, isolation and variational-induction directions; the complete run update, compression, two-step balances and finite-cover derivations and checks are recorded as ChatGPT's work. R131–134 record Y.Ueoka's pure-3^t and lower-balance integration direction and ChatGPT's separation/contact formulas, transport theorems and Mersenne classification. R171–181 record Y.Ueoka's finite-window averaging and band-count research directions, and ChatGPT's occupied-count correction, compression connection, window identities, exact finite-cover implementations and checks. This summary does not reassign earlier results or infer individual contributions of other collaborators from incomplete records.

Independent replay and version identification. The R22 reviewer used Python 3.13.5 without optimization flags, executed the original listing and an independently written verifier, and reproduced all eight rows of the finite-cover tables. The reviewer also checked the early two-band box. The total of the reported representative evaluations is 5,285,670; these are cell evaluations, not counts of integers. The independent review did not audit historical priority or attribution. This preprint retains the R23 mathematical statements, proofs and executable listing; no new full replay or review of version 1.0 is claimed. The SHA-256 of the embedded executable source is

`32a479e51f950d7e863a2d3684252f348b8148527ff947665c88e4a7900ea4cf`

References

[1] Y.Ueoka and Tsumugi (ChatGPT), *Alternating Binary Notation and Collatz Phase Expression: Canonical Representations and Exact Structural Dynamics*, aiXiv:261128.000003 (2026). https://aixiv.science/abs/aixiv.261128.000003.

[2] R. Terras, A stopping time problem on the positive integers, *Acta Arithmetica* 30 (1976), 241–252. https://doi.org/10.4064/aa-30-3-241-252.

[3] J. C. Lagarias, The $3x + 1$ problem and its generalizations, *American Mathematical Monthly* 92 (1985), 3–23. https://doi.org/10.1080/00029890.1985.11971528.

[4] J. Sinyor, The $3x + 1$ problem as a string rewriting system, *International Journal of Mathematics and Mathematical Sciences* (2010), Article 458563. https://doi.org/10.1155/2010/458563.

[5] R. Kaufman, A reduced forward Collatz algorithm: How binary strings change their length under $3x + 1$, arXiv:2301.07466 (2023). https://arxiv.org/abs/2301.07466.