

Finite-Memory Prediction and Remote Feedback in Langton's Ant

Hao Ke

16 October 2026 · Version 0.1.0

Abstract

We study the failure of finite-recency prediction in the classical Langton ant. For each memory length D , the next old-black read is decidable on every finite input, with polynomial bit complexity for fixed D . Finite arithmetic segments retain the complete black field as an XOR expression across long predicted flights, making the event interface exactly iterable. We prove quadratic finite-event time budgets and an endpoint-tube theorem. These imply decidability of spatially filtered events and force every non-eventually-translating orbit to consume newly created, arbitrarily old black memory outside every fixed finite region. A birth-filtered total query returns the first such lifetime created since its source checkpoint, together with exact birth, consumption and complete endpoint. Uniform deadlines yield $\Omega(\sqrt{H}/\log H)$ late-born lifetimes at simultaneously growing radius and age on every non-eventually-translating orbit. An explicit twenty-black family has unbounded feedback lifetime with exactly 31 intervening S_{128} events, showing why spatial scale cannot be removed from a uniform lifetime bound.

1 Introduction

A Langton ant can read a cell long after it last wrote there. Predicting the next turn from a finite suffix therefore raises two questions: when can an old black cell first invalidate the prediction, and what complete state remains after the resulting long excursion? We answer both questions by combining a finite suffix predictor with exact arithmetic on periodic rays and finite XOR expressions for the entire black field.

The resulting event interface also exposes the geometry of memory feedback. Every interval ending in a prediction failure lies near the segment joining its endpoints. A fixed finite region can consequently host only a bounded consecutive run of failures. Spatially filtered event queries are decidable, and every non-eventually-translating orbit must repeatedly consume long-lived black memory outside every fixed region. Tracking last writers strengthens this to an exact first-lifetime query: initial inventory can be distinguished from newly generated memory even when the same coordinate is reused. Uniform event budgets then provide a simultaneous growing-scale lower bound. A twenty-black construction shows why a spatial parameter is essential even when initial mass and the number of intervening events are fixed.

Reversibility, unboundedness and highway geometry are classical features of the ant [1,2]. The periodic read test uses full-drift chronology [3], and the finite-region argu-

ment uses the square escape theorem [4]. The contributions here are the quantitative event interface, exact continuation with full history, geometric event budgets, exact birth-filtered feedback and its growing-scale consequence, and the construction separating those budgets from bounds based only on mass and event count.

2 Model and results

A complete state $X = (B, p, h)$ has finite $B \subset \mathbb{Z}^2$. White turns right, black turns left; flip then move. Let F be the invertible update. For $D \geq 1$, write $Q = 2^D$ and define the intrinsic event

$$S_D(X) \iff p \in B \quad \text{and} \quad \text{pos}(F^{-j}X) \neq p \quad (1 \leq j \leq D).$$

Its inverse-history definition applies at the input time as well as later. It records an old black read. The factor-language event J_{10} and the physical gate segmentation are different predicates; none is substituted for S_D in the following bounds.

Prediction and continuation. The next S_D , if present, occurs before $Q(R + 2Q + 10)$, where R is the initial head-relative radius. Existence is decidable in polynomial bit time for fixed D on sparse explicit input. Finite XOR segments allow the exact full endpoint to be returned without expanding every skipped write. The first n events and their full encoded endpoint are computable in $f(D) \text{poly}(m + n + b)$, where m is the initial mass, b the coordinate bit length, and n is measured numerically.

Geometry. With $W = \text{diam}_\infty(B \cup \{p\})$, the time just after event n , whenever it exists, is at most

$$T_D(n, W) = Qn(W + 2Q + 10) + 5Q^2n(n - 1).$$

Between a starting head p and its next event cell e , the entire finite trajectory stays within $3Q + 1$ of $[p, e]$. After joining successive event locations the corresponding tube radius is $3Q + 2$.

Remote feedback. For every fixed finite region K , existence of an S_D outside K is decidable. If a trajectory is not eventually a nonzero translator, then for every D, K, T there are infinitely many black lifetimes born after T , lasting more than D , and consumed outside K . The properties hold simultaneously for the same lifetimes.

Newly generated feedback. The first exterior old lifetime born since the input checkpoint is totally computable, including its exact birth, consumption and complete post-state. On a non-eventually-translating orbit, Section 7 gives $\Omega(\sqrt{H}/\log H)$ lifetimes with birth after $H^{1/4}$ up to integer rounding, radius of order $\log H$, and age of order $\log \log H$.

Spatial-scale obstruction. For every even $n \geq 20$, a complete 20-black initial state writes a cell at time 60 and next reads it at $624n - 1516$, with exactly 31 intervening S_{128} events. Its first permanent P104 time is $624n - 742$. Thus no lifetime bound depending only on D , intervening event count and initial mass can hold uniformly.

3 Predicting the next old-black read

3.1 Definitions and an intrinsic finite-local target

Write $X=(B,p,h)$, headings N,E,S,W . Let F be the literal update. Define

$$R(X)=\max(\{|q-p|_{\infty}:q \text{ in } B\} \cup \{0\}).$$

For integer $D \geq 1$, define the stale-black predicate

$$S_D(X) \text{ iff } p \text{ in } B \text{ and } \text{position}(F^{(-k)} X) \neq p \text{ for every } 1 \leq k \leq D. \quad (\text{SM1})$$

The inverse is literal and unique: $q=p\text{-dir}(h)$, flip q , and undo a right turn when q is black in the supplied post-state, otherwise undo a left turn. Only cells at distance at most D from p are read during D inverse steps. Consequently (SM1) is a translation-invariant FINITE LOCAL PREDICATE on the present colouring and heading. It does not use an infinite historical oracle.

Using the intrinsic backward orbit makes (SM1) defined even at input time zero. At forward times $t \geq D$, it says exactly: the current read is black and the same cell has not been read in the previous D updates. A very old WHITE reread is not an S_D event. A never-yet-read initial black cell can be an S_D event too. This convention is distinguished from an age defined only after time zero.

Let $\tau_D(X)=\min\{t \geq 0: S_D(F^t X)\}$, when it exists. This target comprises infinitely many full states; the result below is not just specified-full-state reachability.

3.2 Prediction identity: no event means one unique next turn

For a D -letter R/L word w , execute its geometric turns from a fixed pose. Compare the final position to the preceding D pre-update positions. If the last matching position has index j , set $f_D(w)=\text{opposite}(w_j)$; if there is no match, set $f_D(w)=R$. This depends only on w : translation and quarter turns preserve position equality.

Recover the actual previous D turns of X using $F^{(-1)}$, and call them $w(X)$. The following equality is exact for EVERY finite state:

$$\text{actual_turn}(X) \neq f_D(w(X)) \text{ iff } S_D(X). \quad (\text{SM2})$$

Moreover the only possible mismatch is predicted R versus actual L .

Proof: a recent previous read fixes the current colour to its flipped value. With no recent previous read, f_D predicts R . This is wrong precisely when the current cell is black. Old white cells and virgin white cells both give R . The backward suffix is part of the actual orbit, so inconsistent fake histories cannot enter this argument.

While S_D is absent, suffixes evolve autonomously as

$$w \rightarrow \text{suffix}_{(D-1)}(w) f_D(w).$$

This is a deterministic map on $Q=2^D$ states. Its generated infinite turn word has a transient shorter than Q and period at most Q . This alone does NOT prove that the physical ant follows the word forever: distant black cells may cause a first mismatch. The next section bounds WHEN any such mismatch can occur.

3.3 First mismatch against an ultimately periodic motion plan

Consider a complete finite state Y of radius r and a proposed periodic turn word u of length P , with matching start/end heading and displacement $d \neq 0$. Let q_i , $0 \leq i < P$, be its relative pre-update positions, $q_0 = 0$. Its proposed future reads are

$$q_i + k*d, \quad k \geq 0.$$

Earlier periodic reads of the same cell correspond exactly to

$$q_j - q_i = l*d, \quad l \geq 0, \text{ and } (l > 0 \text{ or } j < i), \text{ with } k \geq l. \quad (\text{SM3})$$

Each such l is at most P : the finite footprint's pairwise infinity-distance is at most P , and $\|d\|_\infty \geq 1$. Earlier-read parity can change only at these finitely many thresholds l .

The initial colour contribution is $1[q_i + k*d \text{ in } B_Y]$. If this is nonzero then

$$k * \|d\|_\infty \leq r + \|q_i\|_\infty \leq r + P,$$

so $k \leq r + P$. It is a one-epoch contribution, turning on at k and off at $k+1$. Thus the colour required by literal dynamics at each phase i is

$$1[q_i + k*d \text{ in } B_Y] \text{ XOR parity}(\text{number of eligible pairs in (SM3)}). \quad (\text{SM4})$$

All changes occur by epoch $r + P + 1$. Testing every breakpoint, including the last infinite constant interval, is necessary and sufficient for infinite execution. If there is a mismatch, its FIRST time is strictly below

$$P * (r + P + 2). \quad (\text{SM5})$$

This statement concerns the first mismatch, not a hypothetical trajectory after an earlier mismatch. Induction in time ensures that the proposed path agrees with the real path up to that first point.

If u does not restore heading, repeat it two or four times to restore heading. The resulting displacement is zero, since the quarter-rotated displacement vectors sum to zero. If heading is restored but $d=0$ already, use u itself. In either case let v be the closed-pose block, of length at most $4P$. If two consecutive copies of v were both physically valid, all flips would occur twice and the complete state would recur. Such an orbit would remain in finitely many cells forever, impossible for the classical ant. Hence a first mismatch occurs before $2|v| \leq 8P$ updates.

For clarity, the elementary no-bounded-orbit argument used here is as follows. If the set I of infinitely read cells were nonempty and finite, choose a cell maximizing $x+y$. After the north and east neighbours have ceased being read, the checkerboard/heading-axis invariant fixes the only possible entering and leaving directions at that cell. This requires the same read colour on every sufficiently late visit, contradicting colour alternation. No periodicity, finite-support HC or new computational evidence is assumed in this argument.

3.4 Uniform waiting theorem

Theorem 1. For every finite state X and $D \geq 1$, with $Q = 2^D$, put

$$H_D(X)=Q*(R(X)+2Q+10) . \quad (SM6)$$

Exactly one of the following occurs:

- $\tau_D(X)$ exists and $0 \leq \tau_D(X) < H_D(X)$;
- S_D never occurs at any nonnegative time; the ant has a nonzero translating periodic trace with period $\leq Q$, beginning at a time $\leq Q$.

Thus observing no S_D event at times $0, \dots, H_D(X)-1$ proves that none ever occurs.

Proof: construct the autonomous suffix trajectory. If a physical mismatch occurs before the first repeated suffix, (SM2) gives the event within Q updates. Otherwise at a repetition time $j \leq Q$, the physical trajectory has followed the plan and a periodic suffix cycle of length $P \leq Q$ begins from its current phase. The current complete radius is at most $R(X)+j$: initial black cells move relative to the ant by at most j ; any newly black cell is also within j steps of it.

For nonzero translational drift, (SM5) bounds a first future mismatch by

$$j+P*(R(X)+j+P+2) \leq Q*(R(X)+2Q+3) .$$

For a closed geometric plan, the other bound is $j+8P \leq 9Q$. Both are strictly covered by (SM6). If no mismatch exists, the planned periodic trace is the exact infinite future. A closed or rotating periodic plan was excluded above, so its drift is nonzero. By (SM2), no mismatch also means no S_D event. This proves both the waiting bound and the dichotomy.

No width, cone, retreat, mass, visit-count, future-active set, or hypothesized computation-time bound is supplied as an extra premise. The ONLY spatial size in (SM6) is the computable radius of the actual finite input.

3.4.1 What is the corresponding decision problem?

The family

$$L_D = \{X : \text{there exists } t \geq 0 \text{ with } S_D(F^t X)\} \quad (SM7)$$

is recursive, uniformly in D . This is a genuine all-input decision theorem for a family of local predicates, not a HIGHWAY algorithm. Even an input-dependent computable $D(x)$ cannot make a reduction of an undecidable source language to (SM7). An arbitrary HALT pattern is NOT automatically of this form.

A simple decision algorithm simulates the explicit finite bound (SM6). The implementation uses the stronger symbolic version: only the first suffix cycle is produced, then all-future ray/flip arithmetic finds the first event, or verifies a permanent event-free trace. No simulation to a late arithmetic collision is necessary.

3.4.2 Bit complexity of the implemented first-event decision

With an explicit sparse list of m black coordinates of bit length at most b , the symbolic first-event algorithm uses time $2^{O(D)}$ times a polynomial in the input length. In particular, for every fixed D , language (SM7) is decidable in polynomial time in the sparse BINARY-coordinate input size.

Indeed inverse-history recovery uses D steps; the suffix search uses at most $Q=2^D$ states; its explicit black field has at most $m+Q+D$ cells. The candidate word has length at most Q , and the ray verifier checks $O(Q*m+Q^2)$ phase/cell and phase/phase relations. All coordinate, epoch and offset arithmetic has bit length polynomial in $b+D$. Closed-pose control replay costs at most $8Q$ updates. Ordinary polynomial-time set/dictionary implementations suffice. The parameter factor grows exponentially in D .

For fixed D the upper bound is linear in R . This order cannot in general be replaced by a radius-independent constant: take a certified event-free P104 state with $D=128$ and insert one remote black cell at a virgin frontier site k highway periods ahead. Until its first read the original periodic trajectory is unchanged. The first event is at $104k+O(1)$, while $R=2k+O(1)$. The phase-ray calculation gives the constants exactly. The full input coordinates supply the growing scale.

4 Exact continuation with complete history

4.1 Finite XOR segments, not infinite periodic background

For integer a, v and $n \geq 1$ put

$$\text{Seg}(a, v, n) = \{a+j*v: 0 \leq j < n\}.$$

If $n > 1$ require $v \neq 0$. An encoded colouring consists of a finite list of segments and means their SYMMETRIC DIFFERENCE. A singleton is a segment of length one. Thus an explicit finite black set embeds exactly as singleton terms. Every encoded state still has finite support, including after a huge skipped prefix. Overlaps cancel modulo two; they are never silently interpreted as union.

Membership of a point is the parity of the segments containing it. For a nontrivial segment, containment is one integer-collinearity test followed by a finite-index interval test. A literal update and its inverse therefore remain exactly computable without expanding any segment. Identical singleton terms may be cancelled syntactically; no cell is removed because it is presumed irrelevant. The list is not canonical: distinct expressions may encode the same finite set.

4.2 Exact first mismatch against a periodic motion plan

Let a supplied word u have length P , restore heading, and have nonzero drift d . Its pre-update phase positions are q_i . The proposed future reads are

$$q_{i+k*d}, \quad 0 \leq i < P, \quad k \geq 0.$$

As in the earlier finite-field verifier, an earlier periodic read at this same position exists for each pair (j, r) satisfying

$$q_j - q_i = r*d, \quad r \geq 0, \quad \text{and } (r > 0 \text{ or } j < i), \quad \text{with } k \geq r.$$

Each pair toggles read parity at the one threshold r . These are finite in number.

Now intersect a phase ray with ONE old finite segment:

$$q_{i+k*d}=a+j*v, \quad k \geq 0, \quad 0 \leq j < n. \quad (\text{XOR1})$$

If d and v are not parallel, determinants give at most one integer pair (k,j) , which is tested against the bounds. If they are parallel but $a-q_i$ is not, there is no intersection. Otherwise reduce a nonzero coordinate of d to

$$A*k - B*j = C, \quad A > 0.$$

Writing $g = \gcd(A, |B|)$, there is no solution unless g divides C . When it does, solutions for k form one residue class modulo $|B|/g$. The finite bounds on j supply an interval $[lo, hi]$, intersected with $k \geq 0$. Thus (XOR1) is either empty, a singleton, or a finite arithmetic progression of epochs. The singleton case also handles length-one terms independently of the unused stride.

For a phase i , its initial-black contribution is the XOR of membership in these finitely many epoch progressions. Partition $k \geq 0$ at all progression starts, ends+1, and earlier-read thresholds. Within each partition interval, the active progressions are fixed; their combined parity is periodic in k , with period dividing the least common multiple of their moduli. The earlier-read parity is constant on that interval. Test a whole residue period, or the shorter finite interval. The final unbounded interval has no active finite segment and needs one parity test.

This is an exhaustive finite algorithm, not sampling a few epochs. It decides whether u^ω is the exact future; otherwise it finds the first mismatch. Until the first mismatch the proposed and literal paths agree by induction, so its position and time are real. The running time can depend on a large lcm. No polynomial-time claim is made for arbitrary segment-encoded input strides.

4.3 Exact algebra of the skipped writes

Suppose the first n proposed updates have been verified (including $n=0$). For each phase i put

$$n_i = \max(0, \text{floor}((n-1-i)/P)+1).$$

Its updated cells are exactly $\text{Seg}(q_i, d, n_i)$, omitting phases with $n_i=0$. Therefore the complete endpoint is

$$B_n = B_0 \text{ XOR } \bigoplus_{(0 \leq i < P, n_i > 0)} \text{Seg}(q_i, d, n_i). \quad (\text{XOR2})$$

If $n = kP + i$, the endpoint position is $q_i + k*d$ and its heading is the recorded heading before phase i . Equation (XOR2) accounts for every flip, including self-overlap, overlap between phases, and overlap with every old segment. It adds at most P terms regardless of n . The helper applying (XOR2) does NOT assert that an arbitrary geometric word is physically executable; its caller must first establish validity through the chosen prefix.

At the first mismatch of the recency- D predictor, the planned turn is R and the actual turn is L . Toggle that point in (XOR2), turn left and move once. This is the exact full input for the next macro, with no separate history oracle: the needed preceding D positions are recovered by the literal inverse.

4.4 Total single-event macro and finite iteration

For $D \geq 1$ put $Q = 2^D$. The previous suffix identity says that no S_D event means one deterministic next turn from the latest D -turn word. There are Q words. Execute at most Q matching turns until an event or a repeated suffix occurs. In the latter case the prediction has a period $P \leq Q$.

If its geometric period is nontranslating or rotates the heading, two copies of a closed-pose block of length at most $4P$ expose a mismatch: otherwise the complete state would recur and the ant would be permanently confined. Literal no-bounded-orbit excludes that. Replaying at most $8P$ steps handles this case.

For nonzero translation apply Subsection 4.2. A valid infinite word gives a certified permanent event-free tail. Otherwise (XOR2) and one real black update give the next full post-event state. Every input is finite and every arithmetic search here is finite, so this single macro is total without caller resource caps.

If the input expression has M terms, an event output has at most

$$M + 9Q + 1 \text{ terms.} \quad (\text{XOR3})$$

The loose uniform value covers Q preliminary writes, at most $8Q$ closed-plan writes, or at most Q new periodic segment terms, and the final event. Inverse history recovery is temporary and does not alter the returned physical field. Nontrivial new segment strides are period drifts; both coordinate magnitudes are at most Q . The same statement applies to a terminating tail's entry state.

After n successive event outputs from an explicit m -black input, the expression has at most $m + n(9Q + 1)$ terms. Every individual call terminates; the iteration over ALL events is not asserted to terminate. `event_prefix` computes the first n events, or stops earlier on a certified event-free tail. Reaching its finite requested event budget is neither HALT nor NO-HIGHWAY. Optional resource caps return INCONCLUSIVE and do not change the mathematical total algorithm.

4.5 Parameterized complexity of a FINITE event prefix

For an explicit sparse input with m black cells, coordinate bit length b , fixed D , and an event count n measured numerically (or in unary), the first n events and their complete compressed endpoint are computable in

$$f(D) * \text{poly}(m + n + b). \quad (\text{XOR4})$$

This is not a polynomial bound in $\log(n)$, an ordinary-step cutoff, or a claim that large D is practical. f may be doubly exponential in D .

Here is the bound rather than an assumption of compactness. All initial terms are singletons, and every nontrivial generated stride has components of magnitude at most Q by (XOR3). A parallel intersection modulus is $|B|/g \leq Q$. Hence all residue periods divide $\text{Lambda}_Q = \text{lcm}(1, \dots, Q)$, a function only of D . The numbers of old terms, phase pairs and interval endpoints are polynomial in $m + n$ for fixed Q . One whole residue enumeration therefore costs a factor at most Lambda_Q , not a factor equal to the huge skipped physical time.

For bit lengths use the already proved event deadline, or rederive it by $R(F^t X) \leq R(X) + t$. For this bit-length argument the elementary preliminary bound suffices. With $C = R(X) + 2Q + 10$, the n th event, if present, occurs before

$$C * ((Q+1)^{n-1}).$$

All segment endpoints come from the initial support or actual reads before this time. Their binary coordinate and length encodings thus use $O(b+n(D+1))$ bits. No cancellation-dependent radius computation or hypothesized mass bound is required. Membership, determinants, gcd and residue arithmetic consequently have polynomial bit cost; combining these bounds proves (XOR4).

General arbitrary-stride symbolic inputs retain totality but are outside (XOR4)'s explicit-input/generated-stride premise. The lcm cost must not be hidden.

5 Quadratic event budgets

Put $C(X) = \text{conv}(B \cup \{p\})$. A finite macro ends just after its first S_D update. We prove the convex-halo enclosure

$$C(X_{\text{next}}) \subset C(X) + [-5Q, 5Q]^2,$$

including every intermediate head position. Iterating gives $C(X_n) \subset C(X_0) + [-5Qn, 5Qn]^2$ and the stated $T_D(n, W)$. In the proof below, (A) denotes this one-macro enclosure, and (B1)–(B3) denote the iterated enclosure, $R(X_n) \leq W + 10Qn$, and the quadratic time bound, respectively. The first-event waiting estimate is the predictor bound proved above, denoted (GEO1) in this section.

5.1 A geometric lemma for ANY translating periodic motion plan

Start at $Y=(B,p,h)$. Let a nonempty word u of length P restore heading and have nonzero drift d . Write its pre-update phase positions as $q_0, \dots, q_{(P-1)}$, $q_0=p$. The prescribed infinite positions are q_i+k*d for $k \geq 0$. These are at first just a MOTION PLAN: executability is not assumed.

Until its first read mismatch with literal LR dynamics, the plan is the actual orbit, by induction on updates. All earlier plan writes then legitimately contribute to its current read colour.

5.1.1 3.1 Only finitely many old-field and self-read changes matter

A previous phase j visits the position q_i+k*d precisely when

$$q_j - q_i = s*d, \quad s \geq 0, \quad \text{and } (s > 0 \text{ or } j < i), \quad k \geq s.$$

For each phase pair this toggles parity at one threshold s . Let K_{self} be the maximum of all these s and zero. Since the two phase points are connected by at most P unit moves within the block,

$$K_{\text{self}} * ||d||_{\text{infinity}} \leq P. \quad (\text{GEO3})$$

Let K_B be the greatest epoch $k \geq 0$ for which some $q_i + k*d$ is in the ACTUAL finite set B . If there is no such pair, put $K_B = -1$. This maximum is finite. If $K_B \geq 0$, retain one corresponding black point

$$z = q_j + K_B * d \text{ in } B. \quad (\text{GE04})$$

Put

$$K = \max(0, K_self, K_B + 1). \quad (\text{GE05})$$

At every epoch $k \geq K$ the old field contributes zero at every phase and every self-read threshold has stabilized. The prescribed read-colour vector is thus constant in k there. If a mismatch exists at all, a FIRST mismatch must occur within epochs $0, \dots, K$, hence at an offset strictly below $(K+1)*P$.

This remains true if early old-black cells temporarily mask a self-inconsistent motion plan. Such masking ends by $K_B + 1$ and has not been silently discarded. The argument works for either direction of colour mismatch; its later suffix instantiation has the special R-to-L mismatch property.

5.1.2 3.2 The same cutoff bounds SPACE by a constant halo, not by travel time

We prove that every prescribed read in epochs $0, \dots, K$ is in

$$\text{conv}(B \text{ union } \{p\}) + [-3P, 3P]^2. \quad (\text{GE06})$$

If $K = K_self$, (GEO3) gives $\|k*d\| \leq P$ for all $k \leq K$, while $\|q_i - p\| \leq P$. These points are within $2P$ of p , already sufficient for (GEO6). If $K = 0$ the statement is immediate.

Otherwise take $K = K_B + 1$ and the old black point z from (GEO4). For $0 \leq k \leq K_B$ and $K_B > 0$, put $\lambda = k/K_B$. Then

$$\begin{aligned} q_i + k*d &= (1 - \lambda)*p + \lambda*z \\ &\quad + (q_i - p) - \lambda*(q_j - p). \end{aligned}$$

The first two terms lie on the segment $[p, z]$; the remainder has l_∞ norm at most $2P$. For $k = K_B + 1$, add d to the $k = K_B$ expression; $\|d\| \leq P$, giving $3P$. When $K_B = 0$, the two epochs $k = 0, 1$ are within $2P$ of p directly. Thus (GEO6) holds in every case, including ties between maxima in (GEO5).

In particular, if the plan has a first physical mismatch, ALL positions before and AT that read satisfy (GEO6). Executing the actual mismatched move instead of the planned one adds at most one unit of halo. The old field and every new black cell at the resulting complete endpoint also lie inside (GEO6) enlarged by one, since each new black cell is a previously updated location.

There is no claim that an arbitrarily long VALID event-free translating tail stays in (GEO6). The finite mismatch is an essential hypothesis.

5.2 Proof of Theorem A, including rotating/closed plans

Run the order-D predictor from X until an immediate S_D event or a repeated suffix. Before the repetition $j \leq Q$, all written cells and the head are within Q of the original

head, so

$$C(X_j) \subset C(X) + [-Q, Q]^2. \quad (\text{GE07})$$

An immediate event and its outgoing move have the claimed bound trivially. At a repeated suffix the plan cycle has length $P \leq Q$.

If the heading is restored and drift is nonzero, apply Subsection 5.1 from the full state X_j . A finite next event is exactly its first mismatch. By (GEO6), (GEO7) and the real final update, the entire macro stays within halo

$$Q + 3P + 1 \leq 4Q + 1 \leq 5Q. \quad (\text{GE08})$$

If the cycle rotates heading, repeat it at most four times to get a closed-pose plan; if its displacement is already zero use one copy. This closed plan has length $L \leq 4P$. Each point of a closed unit-step walk is reachable from its origin in either time direction, so its distance from the origin is at most $L/2 \leq 2P$. If TWO copies of that plan were fully executable, they would flip each cell an even total number of times and restore the pose, repeating the complete state. Consequently a mismatch occurs before $2L \leq 8P$. All positions up to its read are in the same closed footprint, and the true mismatched move adds one unit. The resulting source-relative halo is at most $Q + 2P + 1 \leq 5Q$.

Every unchanged black cell was already in $C(X)$, and every newly black cell was on the bounded read footprint. This proves both the path and complete-support parts of (A). Neither the geometric convex-envelope theorem nor deleting any archive is needed for this new bound.

5.3 Proof of Theorem B and the improved coordinate bit bound

Minkowski sums of l_∞ balls add their radii. Applying (A) successively proves (B1), and also covers the complete trajectory between checkpoints. The l_∞ diameter of $C(X_0) + [-5Qn, 5Qn]^2$ is $W + 10Qn$. Both the current head and each current black cell lie in that set, proving (B2).

The n th macro starts after $n-1$ events. Because event offset $<$ (GEO1), and the post-event duration is offset+1, its duration is at most

$$Q * (W + 10Q * (n-1) + 2Q + 10).$$

Summing from 1 through n proves (B3). These are integer non-strict bounds for post-event time; the event itself occurred one update earlier.

The old exponential event deadline remains sound; this is a stronger conversion, not a retraction of its proof. No claim of a best constant or sharp lower bound is made. Smaller D can produce more events, and the constants are exponential in D ; the comparison is for each FIXED D as n grows.

5.3.1 Complete XOR histories still have a growing number of terms

In the prior symbolic representation, a skipped periodic phase adds a finite arithmetic segment of actual read positions, and ordinary updates add singletons. Starting from an explicit m -black input, every generated term endpoint up to n events is an actual read location covered by (B1); original singleton terms also lie in that hull. This is

true BEFORE syntactic cancellations, since a term is produced from a genuine finite path, not from invented archive filling. Each nontrivial segment has nonzero integer stride, so its length is at most $W+10Qn+1$ (distance between its two endpoints is length-1 times a norm ≥ 1). If initial absolute coordinate bit length is b , endpoints and lengths thus have bit length

$$O(b+D+\log(n+1)), \quad (\text{GEO9})$$

improving the earlier $O(b+n(D+1))$ estimate from the exponential deadline. The NUMBER of terms can still grow as $m+n(9Q+1)$. Thus (GEO9) is NOT a constant-size memory theorem, not a finite-state quotient, and not logarithmic-in- n runtime. Temporary inverse-history reconstruction adds at most D additional local moves; it changes neither the returned state nor this asymptotic bit bound.

6 Endpoint tubes and spatially filtered feedback

Throughout this section $Q = 2^D$, and T_D is the quadratic budget above, denoted (REMOTE2). For the square argument one may use the current bound [4]

$$E(\ell) = \frac{13 \cdot 3^{\ell-1} - 2\ell^2 - 4\ell - 5}{2}.$$

The accompanying remote-feedback implementation retains the conservative $8\ell 3^\ell + 1$ bound, which suffices for the same termination argument; it does not evaluate huge powers during event filtering. Neither choice changes the tube lemma or the qualitative feedback conclusion.

6.1 Two-endpoint tube for a finite event macro

Start at X with head p . Suppose its first S_D event is at offset τ , at cell e . The macro includes that read and its actual outgoing move, so its duration is $L=\tau+1$. Let $[p,e]$ denote the real line segment between its endpoints.

Theorem 1 (endpoint tube and distance-sensitive time). Every head position from time 0 through time $\tau+1$ is in

$$[p,e] + [-(3Q+1), 3Q+1]^2. \quad (\text{REMOTE3})$$

Every changed cell is in the same tube, and

$$L \leq Q \cdot ||e-p||_{\infty} + 5Q^2. \quad (\text{REMOTE4})$$

The tube does not depend on the radius or cardinality of B . Unchanged black cells outside it remain physically present; (REMOTE3) does NOT enclose the full input black field. A forever event-free trajectory is expressly outside the premise.

6.1.1 Proof

Until the first event, the actual turns agree with the autonomous suffix predictor. Its word has a prefix of length $j \leq Q$ and a cycle u of length $P \leq Q$. If the event occurs in

that prefix, its whole path and actual outgoing move are within $Q+1$ of p , sufficient for (REMOTE3)-(REMOTE4).

Otherwise let a be the head after the j -step prefix, and v_i the relative pre-update phase positions of u . Then $\|a-p\| \leq j$ and $\|v_i\| \leq P$.

First suppose u restores heading and has nonzero displacement d . Write the event as phase i , epoch k , so

$$e = a + v_i + k*d.$$

When $k \geq 1$, any earlier phase position in epoch $\ell \leq k$ has

$$a + v_s + \ell*d = (1 - \lambda)*a + \lambda*e + v_s - \lambda*v_i, \quad \lambda = \ell/k \text{ in } [0, 1]. \quad (\text{REMOTE5})$$

The error has l_∞ norm at most $2P$. Moving a to p costs at most j ; the prefix is within j of p as well. For $k=0$ all cycle-prefix positions are within $j+P$ of p directly. Thus all positions through the event read lie in $[p, e] + [-3Q, 3Q]^2$. The TRUE mismatched outgoing move costs at most one more.

Also, since d is a nonzero integer vector,

$$k \leq k*\|d\| \leq \|e-p\| + j + P.$$

Hence

$$\begin{aligned} \tau_{u+1} &= j + k*P + i + 1 \\ &\leq Q*\|e-p\| + 2Q^2 + 2Q \\ &\leq Q*\|e-p\| + 5Q^2. \end{aligned} \quad (\text{REMOTE6})$$

If the cycle rotates heading, repeat it twice or four times to restore heading. The resulting geometric block is closed, since its successively rotated drift vectors sum to zero. If heading is already restored with $d=0$, use one copy. The closed block has length at most $4P$; its footprint is within $2P$ of a , because a point on a closed unit-step walk is reached by either side of the loop. Thus (REMOTE3) again holds. Two physically valid copies of a closed block would repeat the complete state (every cell is toggled an even number of times), impossible by classical unboundedness or the strict clock. Consequently $\tau_{u+1} \leq j + 8P \leq 9Q$, which is at most $5Q^2$ because $Q \geq 2$. This proves (REMOTE4) also in the closed cases. Every changed cell was a read position, proving the final assertion. $\square$

This argument needs only a finite macro to its first event, not a presumption that any candidate predictor cycle is physically realizable forever.

6.2 The event skeleton carries the large-scale geometry

List the actual S_D event times $t_1 < t_2 < \dots$ and cells e_i . Put

$$\begin{aligned} H_D &= 3Q + 2, \\ J_m &= \|e_1 - p_0\| + \sum_{i=1}^{m-1} \|e_{i+1} - e_i\|, \\ s_m &= t_{m+1}. \end{aligned}$$

Corollary 2. Up to the complete post- m th-event state, the entire head path is contained in the H_D -neighbourhood of the polygonal chain

$p_0 \rightarrow e_1 \rightarrow \dots \rightarrow e_m,$ (REMOTE7)

and

$J_{m+1} \leq s_m \leq Q \cdot J_m + 6Q^2 \cdot m.$ (REMOTE8)

In particular, all changed cells are in that tube; the complete state still contains every unaltered initial black cell outside it.

Proof: each later macro starts one unit from e_i . Replacing that start by e_i in (REMOTE3) adds at most one to the tube radius. Sum (REMOTE4), bounding each later source-to-event distance by $\|e_{i+1} - e_i\| + 1$. The lower bound follows from unit-speed motion between the actual event cells. $\square$

For every real linear functional a , the maximum of $a.p_s$ over this prefix lies between $\max(a.p_0, a.e_1, \dots, a.e_m)$ and that value plus $H_D \cdot \|a\|_1$. Thus, on an infinite event chain, any direction in which the head is unbounded above is also a direction in which these event locations are unbounded above. This is a statement about directional ranges, not limiting speeds or angular convergence.

This is stronger than using the convex hull of the WHOLE black field: far, untouched black debris does not enlarge the event-to-event tube. It is not an assertion that the event skeleton itself is bounded or eventually straight.

6.3 A finite-region burst bound, independent of the outside field

Let K be a nonempty FIXED finite set of lattice cells, with its l_∞ diameter, and choose an axis-aligned square enclosing $K + [-H_D, H_D]^2$. A sufficient side is

$\text{ell}(K, D) = w + 2H_D + 1,$
 $b(K, D) = E(\text{ell}(K, D)).$ (REMOTE9)

Theorem 3 (bounded internal bursts). A consecutive run in the EVENT LIST whose event cells all belong to K has length at most $b(K, D)$. The bound is independent of the rest of the complete black field and of the starting time.

Proof: start immediately before the first event of that run. Corollary 2 keeps the whole actual path, through every completed internal event in the run, inside the square in (REMOTE9). Its first exit is no later than $E(\text{ell})$. Thus the elapsed time through any such consecutive completed run is strictly less than $E(\text{ell})$, and the number of distinct event reads is at most that elapsed time. The slightly weaker integer bound b is sufficient. Exterior colours can be arbitrary because they stay unchanged before first exit. $\square$

The word ‘consecutive’ refers to events, not ordinary updates. The ant may execute many non-event updates between them. Neither infinite local oscillation nor arbitrarily much exterior storage invalidates this uniform bound.

Corollary 4 (joint, not merely separate, escape). For each fixed D , an orbit has infinitely many S_D events if and only if their positions form an unbounded set. If the events are infinite, infinitely many lie outside EVERY fixed finite K .

An infinite tail confined to K would contradict Theorem 3 after its last external event. A finite event list has a finite position set, giving the reverse direction. This is not

a claim that e_i tends to infinity: returns to K can occur infinitely often, interleaved with external events.

6.4 A genuinely total spatially filtered event algorithm

Theorem 5. Uniformly in a complete finite X , integer $D \geq 1$ and explicitly supplied finite K , the following predicate is decidable:

exists $t \geq 0$: $S_D(F^t X)$ and $\text{pos}(F^t X)$ not in K . (REMOTE10)

The algorithm returns the FIRST such event, when one exists. A negative answer also certifies eventual nonzero translating periodicity, but need not certify standard P104.

Algorithm: call the already-total next-event procedure. If it certifies no future S_D , return NO_EXTERNAL_EVENT and its periodic certificate. If its event cell is outside K , return that event and accumulated time. Otherwise execute the finite prefix through that event with the complete field retained and repeat.

Every individual query and finite replay terminates. Were the outer loop infinite, there would be infinitely many S_D events all in K , contradicting Corollary 4. This supplies the previously missing loop-termination proof; it is not a chosen resource cap. More quantitatively, use $K \cup \{p_0\}$ in (REMOTE9): every completed inside-only prefix has its whole path inside that square, so fewer than its $E(\text{ell})$ literal updates, hence at most $E(\text{ell})$ internal events, can precede the final query. The final external event itself may be arbitrarily distant in time. For K empty, the original next-event procedure suffices directly.

The implementation leaves $E(\text{ell})$ symbolic, instead of attempting to construct an integer such as $3^{(6 \cdot 2^{128})}$. It uses literal replay of inside-event intervals; no efficient full symbolic continuation is claimed for huge inside intervals. Optional resource limits return INCONCLUSIVE, NEVER a false negative answer.

6.4.1 An explicit first-external-event deadline

Let $R_0 = \max\{q \text{ in } B_0 \mid |q - p_0|\}$ (zero for empty B), and use $\text{ell}_0 = \text{ell}(K \cup \{p_0\}, D)$. If an external event exists, its FIRST time is strictly less than

$(Q+1) \cdot E(\text{ell}_0) + Q \cdot (R_0 + 2Q + 10)$. (REMOTE10a)

Indeed, the head first exits the guard square at some $\tau \leq E(\text{ell}_0)$. Either an external event has already occurred, or all preceding event cells are in K . In the latter case the NEXT S_D after τ , if present, cannot be in K : otherwise (REMOTE7) would enclose the whole prefix through it inside that square, contradicting the exit at τ . Apply the existing first-event waiting bound at X_τ , using $R(X_\tau) \leq R_0 + \tau$. If there is no next event, there is no external event at all. This is also a direct two-stage decision proof, independent of the loop argument.

For fixed D and fixed K in the normalized $p_0=0$ frame, the filtered predicate is polynomial-time in the sparse binary-coordinate input encoding, using the existing fixed- D polynomial-time next-event procedure. There are only constantly many internal events, and the TOTAL literal inside-prefix replay is bounded by $E(\text{ell}_0)$, a

constant in this fixed-parameter assertion. These constants can be enormous; no practical efficiency or uniform polynomial bound in D, K is claimed.

Why this is more than the old first-event test: the first many S_D events can occur inside K and must be ignored. The earlier fixed- n event theorem alone does not say when that filtering loop stops. The endpoint tube plus finite-domain escape now proves that it always does.

6.5 Quantitative exterior feedback on every infinite event chain

Let $N_D(t)$ count all S_D reads at offsets $0 \leq s < t$, and $M_D(K)(t)$ count those outside fixed K . Internal events occur in at most $M+1$ runs. Theorem 3 gives

$$\begin{aligned} N_D(t) &\leq b+(b+1)*M_D(K)(t), \\ M_D(K)(t) &\geq \text{floor}(N_D(t)/(b+1)). \end{aligned} \quad (\text{REMOTE11})$$

For an infinite S_D chain, define the computable integer

$$n_D(t, W) = \max\{n \geq 0 : T_D(n, W) \leq t\}, \quad (\text{REMOTE12})$$

with the retained deadline (REMOTE2). It follows that

$$M_D(K)(t) \geq \text{floor}(n_D(t, W)/(b+1)). \quad (\text{REMOTE13})$$

For fixed X, D, K , this is an $\Omega(\sqrt{t})$ lower bound as t tends to infinity, with a very large fixed denominator $b+1$. It is not an asserted optimal density. No independent-bit count or computational power conclusion follows from it. For empty K , use $M=N$ and $b=0$. No finite observation can establish the premise that the entire chain is infinite.

6.5.1 A stronger exterior-only quadratic deadline

The factor $b+1$ in (REMOTE13) can be avoided in the LEADING quadratic coefficient. Augment the state hull by the observation region:

$$\begin{aligned} C_K(X) &= \text{conv}(B \cup \{p\} \cup K), \\ W_K &= \text{diam_infinity}(C_K(X_0)). \end{aligned}$$

A finite macro ending just after the FIRST EXTERNAL S_D event satisfies

$$\begin{aligned} C_K(X_{\text{new}}) &\subset C_K(X) + [-9Q, 9Q]^2, & (\text{REMOTE13a}) \\ \text{duration} &\leq 2Q * \text{diam_infinity}(C_K(X)) + 20Q^2 + b. & (\text{REMOTE13b}) \end{aligned}$$

To prove (REMOTE13a), if there are internal events, their entire prefix is in $\text{conv}(\{p\} \cup K) + [-H_D, H_D]^2$ by (REMOTE7). The full field at the last internal post-event state is therefore in $C_K(X) + [-H_D, H_D]^2$, WITHOUT dropping the initial field. Apply the retained $5Q$ next-event convex-halo theorem to the final ordinary event macro. Since $H_D + 5Q = 8Q + 2 \leq 9Q$, (REMOTE13a) follows. With no internal event, the original $5Q$ bound is already sufficient.

For (REMOTE13b), (REMOTE4) bounds the first internal event duration by

$$Q \text{diam}(C_K(X)) + 5Q^2.$$

From its pre-read state through the last internal post-event state there are fewer than b further updates, by the burst proof. At that boundary the complete head-relative

radius is at most $W + 2H_D = W + 6Q + 4$. The first-event estimate makes the last macro at most $Q(W + 8Q + 14)$ updates. Adding these conservative bounds gives

$$2QW + 13Q^2 + 14Q + b \leq 2QW + 20Q^2 + b,$$

since $Q \geq 2$. With no internal event, the original $Q(W + 2Q + 10)$ estimate is smaller. Overcounting the first internal read only weakens the bound. For empty K , use $b = 0$.

After n EXTERNAL events, (REMOTE13a) gives augmented width at most $W_K + 18Qn$. Consequently the post- n th-external-event time, IF PRESENT, satisfies

$$\begin{aligned} U_{-}(D, K)(n, W_K) \\ = n*(2QW_K + 20Q^2 + b) + 18Q^2 * n * (n - 1). \end{aligned} \quad (\text{REMOTE13c})$$

On an infinite chain,

$$M_{-}(D, K)(t) \geq \max\{n \geq 0: U_{-}(D, K)(n, W_K) \leq t\}. \quad (\text{REMOTE13d})$$

For every fixed K this lower bound has leading asymptotic coefficient $\sqrt{t}/(\sqrt{18}Q)$, INDEPENDENT of K . The linear coefficient and the onset of that asymptotic regime can grow enormously with K through b . There is no assertion of an estimate uniform in regions K that grow with time. The new spatial expansion is in EXTERNAL event count; the weaker bound obtained by multiplying the old all-event deadline by the internal burst count is not used.

6.6 Simultaneously remote, late-born, arbitrarily long black lifetimes

An orbit that is NOT eventually a nonzero translating periodic trace has infinitely many S_D for EVERY D , by the existing stale-memory theorem.

Theorem 6 (remote long feedback necessity). On every such orbit, for every $D \geq 1$, finite K , and time T , there are infinitely many cells/lifetimes described by times b, e with

$$\begin{aligned} T < b < e, \quad e - b > D, \quad x \text{ not in } K, \\ \text{update } b \text{ reads } x \text{ white and makes it black,} \\ \text{update } e \text{ is its NEXT read, black.} \end{aligned} \quad (\text{REMOTE14})$$

These lifetimes can be chosen at infinitely many DISTINCT, unbounded cells. This does not assert that the whole event sequence eventually avoids each cell.

Proof: restart at the complete state at time $T+1$. Corollary 4 supplies infinitely many later S_D events outside K . At most $|B_{-}(T+1)|$ can consume black lifetimes already present at the restart. Every other such read consumes a lifetime born at a real forward white update $b \geq T+1$, at the same exterior cell. The intrinsic S_D condition then forces $e - b > D$. There was no intermediate read because black lifetimes end at their next read. Removing the finitely many initial-lifetime events from an unbounded event-position set leaves an unbounded set, so the surviving lifetimes can be chosen at infinitely many distinct cells. Equations (REMOTE13) and (REMOTE13d) also give count bounds after subtracting the finite stock. $\square$

This is a genuine quantifier strengthening. Separately knowing ‘some blockers are arbitrarily distant’ and ‘some lifetimes are arbitrarily long’ does NOT imply

(REMOTE14); those could have been different families. The endpoint-tube argument is what forces long-lived feedback to recur outside every finite region.

It concerns literal S_D consumptions. It must not be relabelled as a theorem that the same long-lived cells are PG primary/secondary blockers: some S_D events can occur inside one PG macro, and no transfer theorem is proved here.

7 Exact detection of newly generated remote feedback

Spatial filtering can be refined to distinguish newly generated memory from black lifetimes already present in the input. This refinement returns the first qualifying lifetime, its provenance, and a complete state from which the query can be repeated. It also makes the dependence of remote feedback on the observation scale quantitative.

7.1 The first birth-filtered query

Fix a complete source state $X = (B, p, h)$, $D \geq 1$, and a finite set K . A qualifying lifetime consists of consecutive reads of a cell $q \notin K$ at times b, e , with

$$0 \leq b < e, \quad e - b > D,$$

where update b turns right and update e turns left. The question asks for the least consumption time e , its birth b , the cell q , and the complete state $F^{e+1}X$. An initially black coordinate can qualify after its initial lifetime has been consumed and the site has been written black again.

Put $M = |B|$, $Q = 2^D$, and $W = \text{diam}_\infty(B \cup \{p\} \cup K)$. To keep one explicit bound throughout this section, use the conservative square escape estimate

$$\ell = \text{diam}_\infty K + 2(3Q + 2) + 1, \quad b_K = 8\ell 3^\ell + 1$$

for nonempty K , and $b_K = 0$ for empty K . The stronger escape bound used earlier can decrease this budget. The exterior-event estimate is

$$U(n, W) = n(2QW + 20Q^2 + b_K) + 18Q^2n(n - 1).$$

Theorem (first newly generated exterior lifetime). The birth-filtered query is total on every finite input. It uses at most $(M + 1)(b_K + 1)$ ordinary next- S_D calls. A positive answer satisfies $e + 1 \leq U(M + 1, W)$. A negative answer certifies eventual nonzero translating periodicity.

Proof. Process the exact successive S_D events in chronological order. Ignore those in K . For an exterior black read, find the last read of the same cell at a nonnegative time strictly before the event. If one exists, it is a white read and hence the birth of the consumed lifetime; return it. If none exists, the event consumes one of the M source lifetimes. Each source lifetime can be skipped only once. In particular, consuming a source lifetime does not permanently exclude its coordinate.

There are at most M skipped exterior events. Each intervening run of internal events has length at most b_K , by the fixed-region burst bound. Either a qualifying exterior

event or a permanently event-free suffix is therefore reached within the stated number of calls. The successful event is no later than exterior event $M + 1$, which proves the deadline. In the negative case only finitely many S_D occur, so the finite-suffix theorem gives eventual translating periodicity. The ordered search returns the first qualifying consumption. $\square$

The negative certificate concerns translating periodicity; its translating word is not presumed standard. Optional implementation limits return an inconclusive status. The unlimited algorithm terminates by the proof above, without expanding the possibly enormous integer 3^ℓ .

7.2 Complete endpoints and exact birth provenance

Alongside the XOR field, store each actual read sequence as a finite arithmetic segment with its time progression:

$$(a, v, L, t_0, \delta_t, c), \quad \text{read } a + jv \text{ at } t_0 + j\delta_t, \text{quad} 0 \leq j < L.$$

A repeated heading-restoring word contributes one such segment per phase; literal updates contribute singleton records. These records partition all executed updates, including those skipped by arithmetic acceleration. Intersecting a queried cell with every segment and selecting the largest recorded time below e recovers its exact last writer. The same records retain every field toggle, giving the complete reusable post-state.

The arithmetic applies to nonprimitive strides as well. For parallel segments write $d = gu, v = hu, a - q = cu$, with u primitive. The intersection equation is $gk = c + hj$; its congruence and finite index bounds give a finite arithmetic progression of epochs k . Nonparallel segments use an exact determinant solve. Splitting epochs modulo the least common multiple of progression steps reduces the colour parity to finite interval sweeps, including the final infinite interval. Thus neither a remote initial marker nor a remote write is lost during acceleration.

For fixed D and fixed K in the translated source-pose frame, the query has polynomial bit complexity on explicitly listed sparse input. There are $O_{D,K}(M + 1)$ calls, each adding $O_D(1)$ segments. Generated strides and the required residue modulus are bounded in terms of D ; the existing finite-macro bounds control coordinate and time bit lengths. For arbitrary compressed initial fields, totality still holds, while the call bound uses actual black cardinality rather than the number of XOR terms.

7.3 Adaptive checkpoints

On a trajectory that is never eventually translating periodic, the query is positive from every actual checkpoint s . One may choose D and a finite K using the known history, then obtain a lifetime born at or after s , consumed outside K , with age greater than D . Restarting after its consumption produces temporally nonoverlapping lifetimes. Including previously selected cells in the next K makes their sites distinct. The region is fixed during each call and may change between calls. Each round has its own complete-source mass and deadline.

7.4 A simultaneous growing-scale lower bound

Normalize $p_0 = 0$, and put $M_0 = |B_0|$ and $W_0 = \text{diam}_\infty(B_0 \cup \{0\})$. For a large integer horizon H , define

$$\begin{aligned} L &= \lfloor \log_3 H \rfloor \geq 64, & D_H &= \lfloor \log_2 L \rfloor - 5, & Q_H &= 2^{D_H}, \\ R_H &= \lfloor L/16 \rfloor, & K_H &= [-R_H, R_H]^2 \cap \mathbb{Z}^2, & s_H &= \lfloor H^{1/4} \rfloor. \end{aligned}$$

Use $\ell_H = 2R_H + 6Q_H + 5$ and the preceding definition of b_K . Let n_H be the largest nonnegative integer satisfying

$$U_{D_H, K_H}(n_H, W_0 + 2R_H) \leq H.$$

Theorem (growing-scale feedback). On every non-eventually-translating finite-support orbit, the number $F(H)$ of black lifetimes satisfying simultaneously

$$s_H \leq b < e < H, \quad e - b > D_H, \quad \|q\|_\infty > R_H$$

is at least $\max(0, n_H - M_0 - 2s_H)$. Moreover,

$$\liminf_{H \rightarrow \infty} \frac{F(H) \lfloor \log_3 H \rfloor}{\sqrt{H}} \geq \frac{32}{\sqrt{18}}.$$

Proof. For each horizon the orbit has infinitely many S_{D_H} , so the uniform exterior deadline supplies at least n_H exterior consumptions before H . At most s_H precede the cutoff s_H . At that cutoff there are at most $M_0 + s_H$ black lifetimes, so at most that many later consumptions can have earlier births. Subtracting these supplies proves the finite bound. The counted objects are distinct lifetimes, and may occur at the same coordinate at different times.

For the asymptotic estimate,

$$L/64 < Q_H \leq L/32, \quad \ell_H \leq 5L/16 + 5, \quad b_{K_H} = O(LH^{5/16}).$$

The linear coefficient of U is $O_X(L^2 + LH^{5/16})$, negligible at its positive quadratic root. Hence $n_H = (1 + o(1))\sqrt{H}/(\sqrt{18}Q_H)$. Finally, $M_0 + 2s_H = o(\sqrt{H}/L)$, giving the displayed lower limit. $\square$

Thus one uniform estimate forces late births, a growing excluded radius, and a growing memory age on the same collection of lifetimes. The conclusion is a necessary condition on non-eventually-translating motion. It supplies $\Omega_X(\sqrt{H}/\log H)$ lifetimes, while radius, age and birth cutoff scale respectively as $\log H$, $\log \log H$ and $H^{1/4}$.

The dedicated controls compare complete endpoints and exact birth times against literal dynamics on 120 inputs, include source-coordinate rebirth, and exercise arithmetic at epoch 10^{60} . On the blank orbit, repeated queries with $D = 128$, $K = \emptyset$ return 31 positive rounds before a periodic negative answer at checkpoint 9990. This illustrates how exact renewal detection fits an orbit that eventually reaches the standard highway.

8 A twenty-black construction and the role of distance

The following construction realizes the spatial-scale obstruction with a fixed number of input cells and a fixed number of intervening events. Its lifetime has birth $b = 60$, consumption $e = 624n - 1516$, and age $e - b = 624n - 1576$. The exact first standard time (CIRCUIT1) is $e + 774$. All twenty cells are specified at input time. We give the full affine construction and the guard argument for every even $n \geq 20$. The outgoing port is the four-cell theorem [5]; the present paper owns the fixed-mass lifetime and event-count conclusion.

8.1 Exact seed and the inherited reference

The literal rule is white/right, black/left, toggle then move, on Z^2 . A state at time t precedes update t . Headings are N,E,S,W=0,1,2,3. All unspecified input cells are white.

Use the already certified outgoing four-cell port

$$G = \{(0, -1), (1, 1), (1, 3), (2, 2)\}$$

at the origin facing E. Let ρ be clockwise quarter rotation. At a pose (p, h) , write $G[p, h] = p + \rho^{(h-1)}G$. For even $n \geq 20$ the complete input black set is

$$B_n = \text{union of } G[p_i, h_i], \quad i=0, \dots, 4,$$

with the following five poses:

i	p_i	h_i
0	(0,0)	E
1	($n-1, n$)	N
2	($-n, 3n-4$)	W
3	($-5n+5, -n$)	S
4	($-3n, -3n+8$)	E

Start the one physical ant at (0,0),E. Subsection 8.2 proves that these five four-cell sets are pairwise disjoint, so the input has EXACTLY twenty black sites.

8.1.1 Reference data, checked rather than inferred from a long run

From G the 182-step prefix A reaches position (5,1),E with 36 black sites. Its entire subsequent word is U^ω , where U is the fixed standard 104-word with drift (2,2). Define q_i and h_i to be the absolute reference pose at time $182+i$. The finite phase data used below are

$$\begin{aligned} q_{14} &= (8, 2), \quad h_{14} = E, \quad U_{14} = R; \\ q_{61} &= (7, 8), \quad h_{61} = N; \\ q_{71} &= (6, 9), \quad h_{71} = N. \end{aligned}$$

The reference reads (0,2) white at update 60, does not read it again in the 182-step prelude, and never reads it on the tail (all tail x-coordinates are at least four). It is

therefore a newly born, permanent black lifetime on the UNPERTURBED reference. The constructed circuit will read it again.

The four initial G sites are first read at offsets 1,22,32,58. The only later S_128 reads of this reference are at offsets 157,280,281, with respective previous reads 9,144,145. The no-later-events claim is checked for all future: for each L phase of U, there is a same-cell preceding phase at distance ≤ 128 already available by tail epoch one. The finite prelude/first epoch is checked literally. The stationary maximum black reread gap is 116, as in the retained standard-memory analysis.

The reference's complete 36-cell state passes both the all-epoch breakpoint parity verifier and the independent stable-class/clean-seed-and-ray verifier. The latter criterion retains its attribution to Jillhewar's 2026 supplied-word result, already cited in `periodic-candidate-certification.md`; neither it nor the four-cell port is claimed new here.

8.2 Universal autonomous execution of the five stages

Put $n=2k$, $k \geq 10$. The five claimed stage words, in their respective rotated poses, are

$A \ U^{(k-4)} \ U[:61],$
 $A \ U^{(2k-5)} \ U[:71],$
 $A \ U^{(4k-6)} \ U[:61],$
 $A \ U^{(2k-7)} \ U[:71],$
 $A \ U^{(3k-4)} \ U[:14].$
(CIRCUIT2)

The endpoint identities from the phase data give the start times

$t_0=0,$
 $t_1=52n-173,$
 $t_2=156n-440,$
 $t_3=364n-821,$
 $t_4=468n-1296,$
 $e=624n-1516.$
(CIRCUIT3)

Stages 0 through 3 end at the next table pose. Stage 4 ends at (0,2),E, BEFORE its next update. This final update is not included in (CIRCUIT2).

Endpoint algebra alone would not prove these words executable: previous wakes and future initial ports could cause an earlier divergence. The next certificate is the missing physical guard.

8.2.1 3.1 Pairwise separation of all stage READ sets

Let $V_i(k)$ be the full set of pre-update read positions in stage i of (CIRCUIT2), not its black set and not just its net flips. The following ten strict separations cover every pair. A row states $\max l(V_A) < \beta \leq \min l(V_B)$.

A	B	$l(x,y)$	$\beta(k)$
0	1	y	$2k$
2	0	x	0

A	B	l(x,y)	beta(k)
3	0	y	-2
4	0	x	0
2	1	x	-2k+1
3	1	y	2k
4	1	y	2k
3	2	y	-2k+1
2	4	x-y	-11
3	4	x	-6k

Here is the all-parameter checking argument. Each read set is exactly a union of 182 affine points, 104 affine arithmetic segments, and the stated 14/61/71 partial-epoch points. For a segment

$$p_i(k) + \rho^r q_j + a \cdot d_i, \quad 0 \leq a < R_i(k),$$

its extrema under l are at $a=0$ or $a=R_i(k)-1$, according to the sign of $l(d_i)$. They are affine integer functions of k . Partial-epoch positions use $a=R_i(k)$. Thus each table row reduces to FINITELY many inequalities $\alpha \cdot k + \gamma \geq 0$. Such an inequality holds for all integers $k \geq 10$ iff $\alpha \geq 0$ and $10 \cdot \alpha + \gamma \geq 0$. These two conditions are checked exactly, for every point and segment; no sampling or floating-point optimization is involved.

The regenerated certificate contains 6,832 such inequalities and has minimum integer slack zero. Every $R_i(k)$ is also checked to be positive throughout the domain. In particular the table establishes $V_i(k) \cap V_j(k) = \text{empty}$. Every initial port cell belongs to the corresponding reference's first 59 reads, hence to its V_i . The twenty input cells are consequently distinct, and a port belonging to another stage is outside the whole current read set.

8.2.2 3.2 Why the disjoint read sets suffice

At the start of stage i , previous stages have changed only cells in their V_j . Every other initial port is also outside V_i . On V_i the complete true field therefore agrees with a fresh rotated G , with white everywhere else on that read set. First-difference induction makes the entire stage word in (CIRCUIT2) literal. Induction over the five stages proves (CIRCUIT2)-(CIRCUIT3) for the ONE complete input B_n . All previous black residue is still present outside the current stage reads.

The apparent online placing of gates used during discovery is not the theorem's execution model. Formula B_n places all gates at time zero, and the read-set separation proves precisely why these remote initial additions do not alter an earlier stage.

8.2.3 3.3 The actual long-memory return and the count 31

The target $x=(0,2)$ is not initially black. Stage 0 writes it black at $b=60$. It is not read again in that stage, and the read-set separations exclude it from stages 1-4 up to time e . At time e the ant reaches x , E , and reads it black; its next turn is L , while

the fresh stage-4 reference would require R. This is an exact first returned lifetime, not a geometric near-return.

In (60,e), stage 0 contributes its three offsets 157,280,281. Each of stages 1–4 contributes all seven reference offsets 1,22,32,58,157,280,281. No other stage has read any of its cells. Its four initial black first reads occur at absolute times greater than 128, and so are genuine intrinsic S₁₂₈ events. All other gaps are those of the reference. Therefore the count is exactly

$$3+4*7=31. \quad (\text{CIRCUIT4})$$

The consumption at time e itself is NOT included. This convention matters.

8.3 The loop does not renew: a full-history P104 recovery certificate

To decide whether the return might start another enlarged circuit, retain the FULL endpoint B_e(n). It has all five old wakes. We do not replace the physical input by a cleaned configuration.

Use the fixed unbounded region

$$\Omega = \{ (x,y) : x \geq -16 \text{ and } -14 \leq x+y \leq 12 \}. \quad (\text{CIRCUIT5})$$

We prove that the restriction of B_e(n) to Ω is the SAME 90-cell set C for every $k \geq 10$. This is an equality of restrictions for this constructed family, not a confinement hypothesis on arbitrary HC inputs.

8.3.1 4.1 All-parameter equality of the protected field

The complete endpoint is the XOR of all twenty initial sites and every actual read position in (CIRCUIT2), counted with its visit parity. The stage-1 read set has $x+y > 12$. Stage 3 has $x+y < -14$. Stage 2 has $x-y < -44$, whereas Ω entails $x-y = 2x - (x+y) \geq -44$. Those three stages and their initial ports miss Ω .

Stage 0 has a fixed prefix and NE period streams. Every epoch $j \geq 6$ has $x+y > 12$; all $k \geq 10$ include the first six epochs. Its terminal partial epoch is outside Ω . Thus its restriction is constant.

For stage 4, let $r = 3k - 4$ and change epoch variable from j to $v = j - r$. Every tail point becomes

$$p_4(k) + q_{i+j}(2,2) = q_{i+(-8,0)} + v(2,2), \quad (\text{CIRCUIT6})$$

independent of k . Complete epochs have $-r \leq v \leq -1$; the partial epoch has $v = 0$ and $i < 14$. At $k = 10$, $r = 26$. Every additional epoch $v \leq -27$ lies below $x+y = -14$. The stage-4 initial port and 182-step prefix also lie below that line for all $k \geq 10$.

Consequently C is the following finite XOR, restricted to Ω :

$$\begin{aligned} &G; \\ &\text{all 182 stage-0 prefix reads;} \\ &q_{i+j}(2,2), \quad 0 \leq i < 104, \quad 0 \leq j < 6; \\ &q_{i+(-8,0)} + v(2,2), \quad 0 \leq i < 104, \quad -26 \leq v < 0; \\ &q_{i+(-8,0)}, \quad 0 \leq i < 14. \end{aligned} \quad (\text{CIRCUIT7})$$

The checker evaluates (CIRCUIT7), not an assumed invisible-history oracle. It obtains 90 cells. It additionally verifies all 1,247 affine and 208 constant inequalities used above. The genuine full endpoints for eight parameters are checked against an independent literal replay as a control, not as the universal argument.

8.3.2 4.2 Protected scattering and exact first standard time

From the reference state $(C, (0, 2), E)$, execute 774 literal updates. All positions remain in Omega. The complete reference endpoint has 162 black cells and pose $(5, -3), E$. Its next 104-word is the cyclic phase-103 shift of the inherited U, with drift $(2, -2)$. Both independent infinite checks certify its entire future. Every tail phase lies in Omega, and adding $(2, -2)$ preserves $x+y$ and increases x . Thus the whole infinite reference future stays in Omega.

The TRUE state at time e agrees with this reference on all of Omega. By first-difference coupling its whole future read trace is the same, for every n , irrespective of the physically retained exterior wakes. This proves standardness from $e+774$.

The reference turn at offset 773 is R; the preceding letter in the backwards 104-periodic extension of the certified tail is L. Therefore a standard future cannot start at or before $e+773$: any such 104-periodic word would agree with the known tail over a whole period and hence at the preceding offset as well. The mismatch contradicts that. Together with the positive certificate, this proves the exact FIRST time (CIRCUIT1), including the exclusion of starts before e . It does not rely on a finite observed repetition as an infinite certificate.

This closes this particular self-renewal attempt in the positive direction. At the return it does NOT recreate the original five-port input. The certified future is P104, not a second expanding circuit.

8.4 Precise consequence for the attempted global bound

Let $N_D(b, e)$ count S_D reads strictly between a black lifetime's birth and its consumption. On all complete finite-support classical inputs there is no finite-valued function $F(D, m, M)$, computable or otherwise, satisfying

$$e - b \leq F(D, N_D(b, e), |B_0|) \quad (\text{CIRCUIT8})$$

for every post-input-created consumed black lifetime. In this family the three arguments are fixed at $(128, 31, 20)$, while $e - b = 624n - 1576$ is unbounded. The same examples refute a bound in terms of an upper budget $m = 31$ and $M = 20$. Their excursion radius is unbounded as well.

By contrast, the zero-intervening-event estimate remains true. At a white birth b there is no S_D event. If $e - b > D$ and there is none in (b, e) , the death is the FIRST S_D from X_b and its source/event positions are equal. The endpoint-tube time estimate gives $e - b + 1 \leq 5 \cdot 4^D$. If $e - b \leq D$ the weaker bound $e - b \leq 5 \cdot 4^D$ is immediate. Thus

$$N_D(b, e) = 0 \Rightarrow e - b \leq 5 \cdot 4^D. \quad (\text{CIRCUIT9})$$

Equation (CIRCUIT9) is a direct corollary of the retained theorem, not a separate rediscovery of its predictor proof. The false step was extending it across an arbitrary finite number of intervening events without carrying spatial scale.

No bound depending on the complete input, its coordinate bit length or spatial diameter is refuted. Here $|B_0|=20$ but $\text{diam}(B_0)=\Theta(n)$. The retained quadratic event deadline explicitly includes input diameter and is fully consistent with these linear-in- n delays. Nor does this show that one fixed seed has unbounded lifetimes, or refute a uniform bound on event count for each input.

8.4.1 The same obstruction for the concurrent J10 event

The concurrent factor-language import defines L10 as the ten-letter cyclic factors of U , and J10 as a true ten-turn window outside L10. We use only this definition, not its separate rigidity theorem. There are 84 allowed factors. On every X_n the number of J10 events strictly in $(60, e)$ is exactly **264**.

To check this uniformly, replace each U^r in (CIRCUIT2) by one U . All r are positive. Every removed complete standard cycle contains only allowed factors; both length-nine boundary contexts are unchanged and one retained cycle is longer than their combined lengths. Therefore no forbidden-factor occurrence is lost or created at a splice. Direct counting on the fixed compressed word gives 264. The same calculation through the protected recovery and nine following standard turns gives 532 J10 events from absolute update 9 onwards, with the last at $e+782$. Later factors are standard. Eight uncompressed physical runs corroborate the interior count.

Thus a bound using only J10 count and initial black count cannot bound this newly created lifetime's age either. This does NOT refute a bound on the TOTAL number of J10 events as a function of the complete input; that is a different potential route to HC. The concurrent import and its evidence remain unchanged.

9 Reproduction and perspective

Run `python3 memory_feedback/reproduce.py` to replay the six scoped modules: first-event prediction, exact XOR continuation, event geometry, remote filtering, birth-filtered feedback, and the twenty-cell circuit. The checks compare full finite fields, reconstruct integer and rational inequalities, and verify permanent periodic tails at all parity breakpoints. Parameter families are covered by the affine and geometric proofs above; finite controls test their numerical premises.

The predictor supplies a total query, while the XOR representation supplies its reusable physical output. The tube and escape arguments make fixed-region filtering total and force nonperiodic behaviour to renew memory remotely. The construction shows that mass and event count alone cannot replace spatial information. Dependence on full coordinates or diameter remains available, as does the separate question of whether one fixed seed can renew feedback indefinitely.

10 References

1. Anahí Gajardo. *A symbolic projection of Langton's Ant (Extended Abstract)*. DMTCS Proceedings AB, 57–68 (2003). [doi:10.46298/dmtcs.2312](https://doi.org/10.46298/dmtcs.2312).
2. Anahí Gajardo, Victor Lutfalla and Michaël Rao. *Ants on the highway* (2024). [arXiv:2409.10124v1](https://arxiv.org/abs/2409.10124v1).
3. Atharva Jillhewar. *Finite-Support Periodic Highways of Langton's Ant: Necessary Conditions, Transverse Exclusions, and Exact Search* (2026). [doi:10.21203/rs.3.rs-10516269/v1](https://doi.org/10.21203/rs.3.rs-10516269/v1). Theorem 3.1 and Remark 3.2 supply full-drift visit chronology and seed reconstruction.
4. Hao Ke. *An Ant on the Clock: Exponential Escape Bounds for Langton's Ant*, version 0.4.0 (2026). [doi:10.5281/zenodo.23402379](https://doi.org/10.5281/zenodo.23402379).
5. Hao Ke. *Four Black Cells Force an Outgoing Highway: A Sharp Half-Plane Threshold for Langton's Ant*, version 0.1.0 (2026). [doi:10.5281/zenodo.23402414](https://doi.org/10.5281/zenodo.23402414).