

Four-Cell Ports and Composable Highway Scattering for Langton's Ant

Hao Ke
Independent researcher
i@kehao.me

15 October 2026
Version 0.2.0

Abstract

Four black cells form a sharp one-sided interface for the classical Langton ant: they launch a permanent period-104 highway in a half-plane, whereas at most three force first exit within the sharp bound of 918 updates. The interface also completes any prescribed finite observed history using four remote cells. We extend this construction to composable highway scattering. A library of thirty protected rules closes under every compatible finite sequence with gaps of at least one hundred periods. A second interface permits arbitrarily long return excursions that read black memory written by the same ant and then restore the incoming highway orientation. Any finite number of these return packets can be installed in one finite input. Nested guard regions and exact XOR identities preserve the complete field through every encounter. Analytic composition proofs are accompanied by reconstructed finite kernels, affine inequalities and all-future tail checks.

1 A sharp outgoing interface

A local pattern can determine the ant's eventual motion while leaving an arbitrarily long initial observation unchanged. We construct such a pattern using four black cells in a half-plane. The supporting line separates its entire future from any configuration on the other side, making the pattern an outgoing interface for arbitrary finite histories. A matching lower bound shows that four cells are necessary for this one-sided interface.

We next compose this interface with protected scattering rules. The central closure condition is equality of the entire field in the next protected region. Increasing guards establish this condition for arbitrarily many encounters. A further return packet deliberately revisits old wake, then recovers a closed outgoing interface using a second exact comparison region. These two constructions give complementary finite-train theorems.

A colouring is a set $B \subseteq \mathbb{Z}^2$ of black cells; p and h denote the ant's position and cardinal heading. On an update, the ant turns left on black and right on white, toggles the current cell, and moves one unit in the new heading. Let (B_t, p_t, h_t) be the state before update t , and define its observation as

$$O_t = (p_t, h_t, \mathbf{1}_{B_t}(p_t)).$$

The half-plane result permits arbitrary exterior colouring. Finite support is required for the exterior-completion theorem.

Put $\mathcal{H}_+ = \{(x, y) : x \geq 0\}$. A *one-sided port* is a finite black seed in $\mathcal{H}_+$ whose trajectory, from $(0, 0)$ facing east, remains in $\mathcal{H}_+$ forever. The *standard P104 trace* is the fixed blank-orbit 104-turn word at update 9977, repeated forever, with a lattice translation, quarter-turn rotation and cyclic phase allowed. Positions, headings and read colours are part of this trace; black cells away from its future path remain in the state.

Theorem 1 (Sharp half-plane escape). *Start at $(0, 0)$ facing east. For every initial colouring with at most three black sites in $\mathcal{H}_+$, regardless of the colours in $x < 0$, the first-exit time*

$$\tau = \min\{t \geq 0 : (p_t)_x < 0\}$$

satisfies $\tau \leq 918$. Equality holds for $B_* = \{(0, -1), (2, 1), (3, 3)\}$ with all other cells white.

Theorem 2 (A four-cell outgoing port). *For $G = \{(0, -1), (1, 1), (1, 3), (2, 2)\}$ and initial pose $(0, 0), E$, every p_t lies in $\mathcal{H}_+$. At update 182 the position is $(5, 1)$, the heading is E , and there are 36 black cells. The entire future from this state is the standard P104 trace, with drift $(2, 2)$ per 104 updates. Consequently four is the minimum black-cell cardinality of a one-sided port.*




Figure 1: The sharp three-cell witness B_* and the four-cell port G . Cell coordinates refer to square centres. The ant starts at $(0, 0)$ facing east; the blue line separates columns -1 and 0 .

The port gives a completion theorem for any finite initial state. The preserved object is the observed trajectory: the four extra black cells are present in the modified colouring from time zero.

Theorem 3 (Four-cell exterior completion). *Let $X = (B, p, h)$ have finite black support. Given a finite protected set K , an integer $T \geq 0$, and any axis-aligned square Q containing B , K and the ℓ^∞ ball of radius T about p , one can compute a four-element set $A \subset \mathbb{Z}^2 \setminus Q$ such that the evolution of $X^+ = (B \cup A, p, h)$ satisfies*

$$O_t(X^+) = O_t(X) \quad (0 \leq t < T)$$

and eventually has the standard P104 trace. If Q has n cells per side, that trace starts by update $E(n) + 182$, where $E(n) = (13 \cdot 3^{n-1} - 2n^2 - 4n - 5)/2$.

The proof has three parts. First-contact branching reduces the unrestricted three-cell problem to a finite family. Exact parity then certifies the four-cell infinite trajectory. Finally, the port is placed beyond the first exit from Q , so its future avoids every earlier black cell.

2 First-contact coverage and the sharp bound

For a finite seed $S \subset \mathcal{H}_+$, let $t(S)$ be its first exit, when finite. Let $V(S)$ contain each cell whose first read before $t(S)$ is white. Define

$$\mathcal{F}_0 = \{\emptyset\}, \quad \mathcal{F}_{j+1} = \{S \cup \{q\} : S \in \mathcal{F}_j, q \in V(S)\}.$$

This construction uses the cells reached by exact trajectories, so the input coordinates require no bounding box.

Lemma 4 (First-contact coverage). *If every member of $\mathcal{F}_0, \dots, \mathcal{F}_k$ exits in finite time, every colouring with at most k black sites in $\mathcal{H}_+$ exits by $M = \max\{t(S) : S \in \mathcal{F}_j, 0 \leq j \leq k\}$.*

Proof. Suppose an actual trajectory stays in $\mathcal{H}_+$ through time M . Start with $S = \emptyset$, a subset of its initially black sites in $\mathcal{H}_+$. If the actual trajectory agrees with the trajectory from S until its exit, it exits by M , contrary to the supposition. Otherwise take the first differing read. All earlier reads and writes agree; hence their colour difference at this cell is its initial colour difference. The cell belongs to the actual initial black set but not to S . It was never read earlier, since an earlier read would already have revealed that difference. Thus it belongs to $V(S)$, and adjoining it gives a member of the next family, still a subset of the actual black set. Each repetition adds a distinct black cell. At most k repetitions exhaust that set, after which a further differing read is impossible. Cells in $x < 0$ are not read before the first exit. $\square$

Exact enumeration gives the complete families:

j	$ \mathcal{F}_j $	$\max t(S)$
0	1	2
1	2	38
2	20	106
3	261	918

The 284 records specify each seed, exit time and exit position. A verifier reconstructs all first-visit sets, checks equality with the next family, and replays every exit. A separate implementation branches on the initial colour at each first visit, with a three-black budget, and obtains the same maximum. The two enumerations have different counts because the online tree identifies redundant or unread black choices.

For B_* , the ant is at $(0, 7)$ facing south before update 917, reading white. The resulting right turn moves it to $(-1, 7)$ at time 918. The finite certificates and Lemma 4 prove Theorem 1. A check that reaches its safety cap fails; it never treats an unfinished trajectory as evidence of escape or eternal residence.

3 Certifying the outgoing infinite trace

Exact replay of G for 182 updates gives the endpoint in Theorem 2, with every prefix position in $\mathcal{H}_+$. Let B' be its complete 36-cell black set. Fix the standard word w , obtained by replaying the blank ant from $(0, 0), N$ to time 9977 and reading its next 104 turns. The same word, started from $(5, 1), E$, restores its heading and has drift $d = (2, 2)$.

Write q_i for the *absolute* pre-update position in phase $i \in \{0, \dots, 103\}$ of this proposed block. Its infinite motion plan reads $q_i + kd$ at time $104k + i$, and the x -coordinates of the q_i range from 4 through 9. It therefore suffices to verify the read colour at every such position and time. An earlier visit at phase j in epoch ℓ reaches the same cell precisely when $q_j - q_i = (k - \ell)d$. The integer $s = k - \ell$ below records this epoch difference; each earlier visit contributes one flip.

Lemma 5 (Finite parity test for an infinite trace). *Let a finite word w restore its heading with nonzero drift d , and let q_i be its phase positions. From a finite initial black set B' , the proposed repeated trace is an actual ant trajectory if and only if, for every phase i and integer $k \geq 0$,*

$$\mathbf{1}_{B'}(q_i + kd) \oplus \left| \{(j, s) : q_j - q_i = sd, 0 \leq s \leq k, s > 0 \text{ or } (s = 0 \text{ and } j < i)\} \right| \bmod 2 = \mathbf{1}_{\{w_i=L\}}. \quad (1)$$

It is sufficient to check $k = 0$, every nonnegative overlap threshold s , and every nonnegative initial-black hit epoch and its successor.

Proof. A previous read at $q_j + \ell d$ affects the proposed read at $q_i + kd$ exactly when $q_j - q_i = (k - \ell)d$. Setting $s = k - \ell$, its time is earlier exactly when $s > 0$, or $s = 0$ and $j < i$; nonnegative ℓ gives $s \leq k$. Each preceding read toggles the cell once, giving the left side of (1). Necessity follows from the rule, and sufficiency follows by induction on chronological time.

For fixed i, j , nonzero d gives at most one integer s , whose parity contribution starts at that threshold. Each initial black cell lies on the phase ray at at most one epoch k ; its contribution is present only there, so both k and $k + 1$ are breakpoints. Between consecutive breakpoints every term is constant. The last breakpoint also represents the final unbounded interval. These finitely many tests therefore cover every $k \geq 0$. $\square$

For the port endpoint, the test uses 170 phase/epoch evaluations and accepts w^ω . Applied independently to the actual blank endpoint, it certifies that w is the standard infinite word. The prefix guard, the phase-position range and the positive x -drift now prove the port stays in $\mathcal{H}_+$ forever. Together with Theorem 1, this proves four-cell optimality.

An additional check uses chronological visits to each physical cell, as in the periodic-realisability criterion of Jillhewar [3]. Partition phases modulo the full group $\mathbb{Z}d$, and write $q_i = a + \ell_i d$ in each class. The cell $a + nd$ is visited at times $104(n - \ell_i) + i$ for all $\ell_i \leq n$. At these times the prescribed colours must alternate, beginning with its initial colour. The visit list changes only when n passes a phase level ℓ_i ; initial black cells add singleton changes. Checking the resulting finite intervals, including the last, is another exact all-future verification. The independent audit checks 40 cell intervals in 22 translation classes, both for the port and the blank reference.

The smallest black set required on the port's future read set consists of 13 cells:

$$K_0 = \{(4, 3), (4, 4), (5, 1), (5, 5), (5, 6), (5, 7), \\ (6, 1), (6, 2), (6, 3), (6, 6), (7, 1), (7, 4), (7, 5)\}.$$

It is contained in B' . The remaining 23 black cells miss every phase ray $q_i + \mathbb{N}d$, and remain as stationary physical history.

4 Completing an arbitrary finite history

Proof of Theorem 3. Simulate the original state to its first exit from Q , at time τ , exterior position $e = p_\tau \notin Q$ and heading $g = h_\tau$. Finite-domain escape ensures termination; the integer-clock bound [5] gives $\tau \leq E(n)$. Since Q contains the radius- T ball and each step has unit length, $\tau > T$.

The exiting move crosses one side of Q , so g is its outward normal. Let R be the quarter-turn rotation mapping east to g , and set

$$A = e + RG.$$

The outward closed half-plane through e is disjoint from Q . Every cell of A lies in it, so A has four distinct, initially white cells outside Q . None is read before time τ .

Compare the original and augmented runs. Induction gives equal positions, headings and read colours before the exit, together with the exact complete-field relation

$$B_t^+ = B_t \triangle A = B_t \cup A \quad (0 \leq t \leq \tau).$$

All of B_τ is inside Q , because the original support and every earlier write lie there. At the exit the augmented state thus consists of the placed port, with pose (e, g) , and physical debris B_τ strictly behind its supporting line.

The placed port's entire future remains in the outward half-plane by Theorem 2. Induction on updates couples its trajectory to the augmented run: every read avoids B_τ , so both runs read the same colour and make the same turn, write and move. The inside debris remains present and unchanged. The standard trace begins by time $\tau + 182$, hence by $E(n) + 182$. The earlier observation equality includes the first T updates, and adding A preserves every initial colour in K . $\square$

The construction has the quantifier form $\forall X \forall K \forall T \exists A$ with $|A| = 4$: its conclusion concerns the augmented input X^+ . The matching lower bound characterizes one-sided ports from a boundary pose. It does not assert minimality among all possible exterior modifications of an individual input.

5 Guarded composition of highway scattering

5.1 Literal convention and finite certificate data

A state is the complete finite black set B , position p , and heading h in $N, E, S, W = 0, 1, 2, 3$. White turns right, black turns left, then flip and move. Time t is before update t . All markers are installed at time zero. There is no intervention between encounters and no physical deletion of old cells.

Rotate the repository's thirteen-cell canonical template by 180 degrees:

$K=\{(0,6),(1,5),(0,5),(2,4),(0,4),(2,3),(-1,3),$
 $(1,2),(-1,2),(1,1),(2,0),(1,0),(0,0)\}.$

The start pose is (0,0), EAST. Its standard word w , length 104, and displacement $d=(2,2)$ are reproduced and infinitely checked in check.py. Let P_i, H_i be its pre-update phase positions and headings, $0 \leq i < 104$, with $P_0=(0,0)$. Let R_0 be identity and $R_1(x,y)=(y,-x)$, a clockwise rotation, NOT reflection. The two incoming orientations have displacement $d_k=R_k d$, $k=0,1$.

A phase j is fresh when its cell has no earlier visit in the two-sided planned standard geometry. For $P_j - P_i = a*d$, an earlier visit exists exactly when $a < 0$, or $a=0$ and $i < j$.

The finite test yields 22 fresh phases:

13,14,15,27,28,29,32,38,45,46,47,
 55,56,57,60,61,62,70,71,72,75,89.

This is geometry of a certified periodic plan; it does not assume an infinite physical seed or replace the intrinsic history of an arbitrary input.

5.1.1 The 30 checked rules

Each table entry in check.py specifies incoming k,j , guard $g \leq 0$, recovery length t INCLUDING the first black marker update, and the exact output pose and standard phase. The graph is:

Incoming	Outgoing	Allowed collision phases
NE	NE	13,14,27,28,60,61,70,75,89
NE	SE	29,32,38,55,56,62,71
SE	SE	13,14,27,28,60,61,70,75,89
SE	NE	15,45,46,47,57

There are 16 NE-input rules and 14 SE-input rules. Recovery lengths range from 312 to 32656 updates. These are supplied sufficient recovery boundaries, NOT necessarily first permanent-P104 times. The finite incoming kernels contain 24 to 252 black cells; some recovered kernels contain 1797 black cells. The rank below does not require that this black mass decrease.

5.2 Exact mature-wake restriction and guarded recovery

For a rule (k,j,g) , put $A_i=R_k(P_i-P_j)$, $d_k=(2,+/-2)$, and $u_i=0$ if $i < j$, $u_i=-1$ otherwise. Define the finite set by XOR:

$$M(k,j,g) = \text{XOR over } i=0..103, \text{ceil}((g-A_i.x)/2) \leq a \leq u_i \\ \text{of } \{A_i + a*d_k\}.$$

It is the restriction to $x \geq g$ of all writes before the collision in the mature incoming standard plan. The positive x drift makes the displayed ranges finite. Freshness implies 0 is not in M . The certificate kernel is

$$X_r = (M \cup \{0\}, (0,0), (H_j+k) \bmod 4).$$

For each of the 30 rules the checker proves the following finite obligations:

1. The literal t-update execution ends at the declared complete kernel field and pose; a separately written integer-vector simulator agrees on every turn and on the complete endpoint.
2. Every recovery read has $x \geq g$.
3. Both exact infinite-tail verifiers accept the declared standard phase from that endpoint. Its drift has positive x component 2, and every phase read in the first tail period has $x \geq g$.

The two infinite tests are not finite repetition heuristics. One evaluates, for each phase ray, all onset thresholds for previous writes and all exceptions from the finite initial black field. Parity is constant between those finitely many breakpoints, including the last infinite interval. The other orders all visits to a fixed translation-class cell: writing its coordinate as $m \cdot d$, a phase at longitudinal index z_i visits at time $104 \cdot m + (i - 104 \cdot z_i)$, when $m \geq z_i$. Those visits must alternate colours beginning with the actual initial colour. It checks every change of the active list and every finite-field exception. The two tests share the integer geometry primitives; they are two independently formulated all-future checks, not two independently maintained software stacks.

Guard lemma. Add any finite black field G contained in $x < g$ to X_r . The entire future read/turn sequence remains that of X_r , and G is never touched.

Proof. Until a first $x < g$ read, first-difference induction gives the certified recovery and then its certified tail. Both stay in $x \geq g$, so no first such read exists. This proves invisibility instead of assuming it. G remains physically present. The same reasoning holds until a later, explicitly planned marker interrupts the tail. Permanent protection across such interruptions requires the nesting argument in the next section. $\square$

5.3 Finite arithmetic certificate for every compatible pair

Suppose rule r has output displacement d_k , phase a and endpoint v_r relative to its marker. To place a later phase- j marker after n full periods plus the phase adjustment, its position relative to the previous marker is

$$q' = v_r + (n + 1_{[j < a]}) \cdot d_k + R_k(P_j - P_a),$$

$$\text{flight updates} = 104 \cdot n + ((j - a) \bmod 104). \quad (1)$$

Use $n \geq 100$. This deliberately loose threshold is not asserted optimal. For the next rule s with guard g_s , its absolute half-plane starts at $q' \cdot x + g_s$. The certificate checks that this boundary is strictly to the right of:

- every previous recovery read and every black cell of its complete kernel endpoint;
- the previous guard;
- the entire synthetic standard past before that endpoint.

The last maximum is finite and, relative to v_r , equals

$$\max_i [R_k(P_i) \cdot x - R_k(P_a) \cdot x + (0 \text{ if } i < a \text{ else } -2)]. \quad (2)$$

Earlier epochs only decrease x by further multiples of 2. Formula (2), rather than a simulation horizon, proves that all omitted pre-output planned writes are outside the next half-plane.

There are exactly 448 compatible ordered rule pairs. At $n=100$ all pass, with minimum strict-boundary slack **122** and minimum marker x displacement **184**. For larger n , the left-hand boundary increases by $2(n-100)$ and the earlier objects do not change. Thus the same inequalities hold for EVERY $n \geq 100$. The 16 initial NE choices also pass; the minimum initial boundary slack is 160.

Future markers cannot interrupt an earlier recovery because they lie strictly beyond its full read footprint. Nor can they interrupt an earlier free flight: the x overshoot of a periodic flight before its target is bounded by twice the maximum of the two finite canonical x spans plus 2, which is less than 184. The checker tests this bound against every pair. Consequently chronological order of the planned encounters is proved; it is not an extra all-future assumption hidden in the definition of an admissible train.

5.3.1 Why the output really re-enters the same input family

At the next collision half-plane, all previous endpoint black cells and all previous retired debris are outside. Inside, only the newly generated standard writes and the remaining preinstalled markers occur. By (2), adding back the synthetic pre-output past changes nothing there. Therefore those newly generated writes restrict to EXACTLY $M_{(k,j,g_s)}$. The current marker adds its single black cell; more distant markers remain untouched. This is equality of the complete inside field, not equality on a fixed-radius sample or merely on reference rays. It proves closure of the interface needed by the next rule.

All later guards strictly increase. Between guards, the certified tail stays to the right of the previous one; each later recovery stays to the right of its newer guard. Hence every retired western region is permanently protected under the whole composed run, even though later deflections occur.

5.4 All-length, multi-encounter convergence theorem

Choose any finite directed word $r_1, \dots, r_m$ in the displayed two-state graph, starting NE, and ANY integers $n_1, \dots, n_m \geq 100$. Starting with $p_0 = (0,0)$, phase 0 and NE orientation, apply (1) recursively to obtain q_i , then set

$$p_i = q_i + v_{-}(r_i),$$

and use rule r_i 's certified outgoing phase and orientation for the next step. This is a finite, effective prescription using integer arithmetic and the fixed rule table. Define the literal initial state

$$B_0 = K \cup \{q_1, \dots, q_m\} \cup G, \\ p_0 = (0,0), \quad h_0 = E,$$

where G is any finite set in $x < -1$. The markers, K and G are disjoint, so $|B_0| = 13 + m + |G|$. No marker is installed later during the run.

Theorem (guarded scattering trains). Every such state encounters the markers in the prescribed order, executes all prescribed recoveries with its complete history retained, and reaches permanent standard P104 by time

$$T = \sum_{i=1..m} [104 * n_i + ((j_i - a_{-}(i-1)) \bmod 104) + t_{-}(r_i)]. \quad (3)$$

For $m=0$ take $T=0$. The final direction is the graph's final NE or SE orientation.

Proof. The canonical tail has minimum read $x=-1$, so G cannot affect its first flight. The initial-boundary inequalities remove K and every synthetic pre-start write from the first collision half-plane. Freshness and the flight overshoot bound give the first encounter and its complete restricted kernel. Apply the guard lemma up to the next encounter. Subsection 5.3 proves both that this next encounter is not premature and that it again has exactly the required restricted kernel. Induct on the finite rule word. Nested guards keep every retired part invisible throughout all remaining recoveries. After the last marker there is no further perturbation ahead and the final infinite-tail certificate applies with the entire western history still present. Summing the actual flight and recovery lengths gives (3). $\square$

This is not just checking more finite seeds: m , all gap parameters, and the finite western history are unbounded. The proof is the closed-interface induction, not extrapolation from the replay tests.

5.4.1 A concrete repeated-turn instance

Repeat the pair (NE phase 29, SE phase 15) twelve times, with every $n_i=100$. The initial field has **37 black cells**, including the thirteen-cell base. The first six markers are

(204,201), (422,-20), (653,194),
(871,-27), (1102,187), (1320,-34).

The 24 prescribed encounters complete at time **451214**, with head at (5387,-81), NE output and standard phase 94. The retained complete field then has **43669** black cells. Both infinite-tail tests accept it. This illustrates why a decreasing black-cell count is unnecessary, but is not a statement about all 37-black-cell inputs or the first time this state becomes P_{now} .

5.5 Transfer of the strict clock into an actual renewal rank

The following rank is defined on the certified scattering train, whose checkpoints are its finite recovery boundaries.

During a recovery choose its finite read rectangle Q , including its endpoint. Exterior colours remain unchanged. Let $s=x+y-c$, with integer c centering the rectangle's diagonal range, and $M=\max_{(p \text{ in } Q)} |s(p)|$. With the invariant sigma from the clock proof, use the scaled integer clock

$$C_Q(X) = -5 * \sigma * \sum_{(q \text{ in } B \text{ intersect } Q)} \chi(q) \\ * (4^{(M+s(q))} + 4^{(M-s(q))}) + G(p, h),$$

where $G=2*(4^{(M+s)}-4^{(M+1-s)})$ for $h=N,E$ and $G=2*(4^{(M-s)}-4^{(M+1+s)})$ for $h=S,W$. It is $5*4^M$ times the $r=4$ clock, with the unchanged outside contribution subtracted. For every internal literal update,

$$C_Q(FX) - C_Q(X) \geq 2*4^M. \quad (4)$$

A computable upper bound U_Q for all internal colourings and poses is

$$5 * \sum_{(q \text{ in } Q)} (4^{(M+s(q))} + 4^{(M-s(q))}) + \max_{(p, h \text{ in } Q)} |G(p, h)|.$$

Thus

$$\text{rho_Q}(X) = \text{ceil}((U_Q - C_Q(X)) / (2 \cdot 4^M)) \quad (5)$$

is a nonnegative integer which strictly decreases on every internal recovery update. Each stored finite recovery checks (4), (5), and the complete endpoint clock again. The universal identity is imported from the existing clock proof; these finite checks validate the particular representation and convention.

The whole certified computation has a lexicographic rank in $N \times \{0,1\} \times N$:

(number of still-unread input markers, mode, local fuel),

with recovery mode 1 and flight mode 0. During flight, local fuel is the exact number of updates to the next marked read, computed by (1). Reading that marker decreases the first coordinate, permitting a reset to recovery mode. During recovery, use rho_Q; returning to flight decreases mode, permitting the next flight fuel to be reset. After the last recovery the certificate program halts. This is a rank on the state with its finite checked train description, not a claimed all-input physical-state rank. It stops at certification; the actual ant then continues forever along its highway.

The logical role of closure is crucial: generated historical black cells do not create further unaccounted marker obligations because their possible interactions have been absorbed in the verified recovery and then guarded. Without that proof, counting original markers would repeat the failed general seed-size induction. The clock alone cannot supply that guard.

6 Composable excursions with genuine old-wake read-back

6.1 Convention, standard word, and finite mature restrictions

Use the unchanged HC contract: complete finite black set B, integer head p, headings N,E,S,W=0,1,2,3. Read white/right or black/left, toggle, then move. State t is before update t. Rotations are clockwise quarter-turns, never an unannounced reflection of the LR rule.

Use exactly the standard word w and thirteen-cell NE seed K from verification/guarded_scattering/check.py. Let P_i, H_i be its 104 pre-update positions and headings, with P_0=0, H_0=E and d=(2,2). Its minimum future value of x+y is zero. Write w_a for its cyclic word beginning at phase a, and R_k for clockwise rotation through k quarter-turns. The return checker pins the earlier source hash; the paper reproduction replays both the 30-rule library and the return module.

Put

$$s(x,y)=x+y, \quad r(x,y)=x-y, \quad ell(x,y)=y-x, \quad G=-160.$$

For a phase j, orientation k, and a linear form f with $f(R_k d) > 0$, define

$$\begin{aligned} M(j,k,f,G) = & \text{XOR}_{(i=0..103)} \text{XOR}_{(L_i \leq a \leq u_i)} \\ & \{ R_k(P_i - P_j) + a R_k d \}, \\ L_i = & \text{ceil}((G - f(R_k(P_i - P_j))) / f(R_k d)), \\ u_i = & 0 \text{ if } i < j, \text{ and } -1 \text{ otherwise.} \quad (2.1) \end{aligned}$$

This is finite. It is the exact parity of all planned standard writes before a phase-j collision, restricted to $f \geq G$. It is a description of a finite restriction, not an infinite physical input. A finite approach realizes it once the old physical endpoint and the

synthetic past omitted before that approach lie outside the comparison half-plane. Those obligations are proved below, not assumed from observed repetition.

Our first collision interface has head 0, heading $H_{57}=S$ and field

$$\begin{aligned} M1 & \text{ union } \{0, q2(n), q3\} \text{ union } H, \\ M1 & = M(57, 0, s, G), \quad H \text{ subset } \{s < G\}, \quad (2.2) \end{aligned}$$

where H is any finite complete historical black field. The two-marker version omits $q3$. Future packets add further markers above the current packet's entire read extent; their noninterference is proved in Subsection 6.7.

6.2 Finite A and C blocks, and the first long flight

6.2.1 3.1 First turn A

The exact 832-update orbit from $(M1 \text{ union } \{0\}, 0, H_{57})$ ends at

$$v_A = (-12, 12), \quad h_A = S,$$

with 580 black cells and a certified standard NW tail w_{58} . Every A read has $-10 \leq s \leq 9$ and every tail read has $s \geq G$. Both infinite checkers accept this complete reference endpoint. Its initial field has 484 cells.

Let $Q_{B,i}$ be the phase positions of w_{58} from this endpoint. Then

$$q2(n) = Q_{B, 14+n*(-2, 2)},$$

and the first reference read there is exactly $104n+14$ updates later. The finite ray calculation enumerates every i and integer z with $q2(0) - Q_{B,i} = z*(-2, 2)$. All $n+z$ are nonnegative for $n \geq 100$, and $\min(104z+i) = 14$. The expected letter is R. This checks first contact, not just that the proposed marker eventually lies on the reference route.

The marker $q2(n)$ is strictly left of all A reads, for every $n \geq 100$. The fixed $q3$ is absent from A and from every future NW phase ray. The old field H is outside $s \geq G$; the A prefix and its NW reference remain inside. First-difference coupling proves the actual A block and the subsequent $104n+14$ -update flight B, ending at the black marker $q2$. Its absolute packet time is $104n+846$.

6.2.2 3.2 Distant reversal C

Let $M2 = M(72, 3, \text{ell}, G)$. The exact 51272-update orbit from

$$(M2 \text{ union } \{0\}, 0, (H_{72}+3) \bmod 4)$$

ends at $v_C = (-24, -44)$, heading N, with 2569 black cells and a certified SE reference tail w_{18} . Its initial field has 485 black cells. All its finite reads have $\text{ell} \geq G$ and, after translation by $q2(n)$, $s \geq G$. Its bounds are $-75 \leq s \leq 63$ and $-76 \leq r \leq 54$ in LOCAL coordinates.

The true field before this C block agrees with this kernel on

$$\Omega_C(n) = \{s \geq G \text{ and } \text{ell}(v - q2(n)) \geq G\}. \quad (3.1)$$

Indeed every black cell of the A endpoint and $q3$ lies strictly below the ell boundary. So does every point of the synthetic standard past preceding the A endpoint. Their maxima are finite; subtracting $\text{ell}(q2(n)) = 4n+26$ makes each required strict inequality an integer affine inequality in n . They hold at $n=100$ with nonnegative slope thereafter. Hence inside (3.1), the actual B writes equal exactly the mature M2 writes

from (2.1). The marker adds {0}. The arbitrary H contributes nothing because of the s boundary.

Thus C is an actual literal block on the complete input. No all-future conclusion for the TRUE orbit is inferred from its isolated-kernel tail: the older field outside Omega_C will matter during the reverse flight.

6.3 Exact first return to old memory, not an assumed invisible archive

At C's end the head is

$$q2(n)+v_C=(-2n-39, 2n-33), \quad h=N, \quad (4.1)$$

at packet time $104n+52118$. Its isolated reference follows w_{18} southeast. Let $Q_{D,i}(0)$ be the 104 phase positions from $(-39,-33)$, heading N. For general n the phase positions are $Q_{D,i}(0)-n*(2,-2)$.

The true field and this reference differ by the following exact XOR:

$$A_endpoint \text{ XOR } B_writes \text{ XOR } (q2(n)+M2) \text{ XOR } \{q3\} \text{ XOR } H. \quad (4.2)$$

The finite C writes cancel because both trajectories execute them identically. The two copies of the marker $q2$ also cancel. This identity retains every older write; it does not identify the true state with a cleaned reference.

All reverse reference phase rays have s values strictly below both the NW B write stream and $q2(n)+M2$, while still having $s \geq G$. Therefore they cannot read those two terms or H. The remaining first-difference query is against the fixed finite A endpoint and $q3$.

For each phase i and black point b in $A_endpoint$ union $\{q3\}$, enumerate integers z satisfying

$$b - Q_{D,i}(0) = z*(2, -2).$$

There are 158 phase/cell candidates. Every $n+z$ is nonnegative for $n \geq 100$. The minimum of $104z+i$ is -170 , attained at phase index 38 and $b=q^*=(-38,-32)$. No candidate belongs to $q3$. The reference letter there is R and the actual cell is black. Hence the actual reverse flight has EXACT length

$$104n-170 = 104(n-2)+38, \quad (4.3)$$

and ends in the old black read at packet time $208n+51948$. Before it, all turns agree with the SE reference. At it, the real turn is L rather than R.

6.3.1 The black cell was created by the same actual trajectory

In the incoming NE plan relative to phase 57, solve

$$\begin{aligned} q^* &= P_i - P_{57} + a*d, \\ 104a+i-57 &< 0. \end{aligned}$$

The last such read is $i=89, a=-17$, at time -1736 , and $w_{89}=R$. The A block never reads q^* ; B is disjoint by its s range; C misses the return region below; and (4.3) proves that D does not read q^* earlier. Thus its black lifetime is exactly

$$(208n+51948) - (-1736) = 208n+53684. \quad (4.4)$$

In every train realization below, the incoming flight has more than 1736 updates, so this birth occurs AFTER the finite input time, even in the first packet. It is not an initially supplied marker, an intrinsic pre-input lifetime, or a cell that is simply presumed eventually relevant.

6.4 A constant full comparison field after an arbitrarily long excursion

Use the fixed region

$$\Omega = \{s \geq G \text{ and } r \geq G\}. \quad (5.1)$$

At the old read q^* , the true field restricted to Ω is independent of n . This is the closure-enabling fact. The arbitrary H is absent by its s guard. The distant q_2 marker and ALL C reads are outside Ω for every $n \geq 100$: their r values are $-4n-26$ plus a bounded local value, strictly below G . This does not say they were erased; they remain in the complete true endpoint.

For completeness, here is an explicit finite XOR formula for the field in Ω . Start with $A_{\text{endpoint}} \cup \{q_3\}$. For each $Q_{B,i}$ put

$$h_i = \text{floor}((r(Q_{B,i}) - G)/4).$$

Add by XOR all $Q_{B,i+j*(-2,2)}$ with $0 \leq j \leq h_i$ and $s(Q_{B,i}) \geq G$. All $h_i = 100$. The terminal 14 B reads at epoch n are outside Ω . Thus these are exactly the B writes inside Ω , not a truncation by simulation time.

For the reverse flight write its complete epochs as $v = j - n$, so $-n \leq v \leq -3$. For each $Q_{D,i(0)}$ put

$$l_i = \text{ceil}((G - r(Q_{D,i(0)}))/4).$$

Add by XOR $Q_{D,i(0)+v*(2,-2)}$ for $l_i \leq v \leq -3$ when $s(Q_{D,i(0)}) \geq G$. All $l_i \geq -n$ for $n \geq 100$. Finally add the first 38 phase positions at $v = -2$, if they lie in Ω . Restrict the resulting finite field to Ω .

The resulting field C_{old} has 1446 cells. Its exact hash and the 3618 B terms and 3893 D terms are recorded in the report. This is equality of the ENTIRE inside field, including the old A scar and both relevant write streams. It is not equality on the next few read positions or a bounded mask inferred from matching poses. Every lower/upper epoch omission is justified by the linear forms above.

6.4.1 5.1 Two-marker output

From $(C_{\text{old}} \text{ minus } \{q_3\}, q^*, S)$, an exact 9308-update block ends at $(-22, -60), S$, with a certified SE w_{14} tail. All these finite reads and the infinite tail stay in Ω . First-difference induction with (5.1) proves that the complete two-marker state is permanently standard by packet time $208n + 61256$.

This is already an all- n positive family with a real unbounded-age readback. Its whole-input realizations from K have FIFTEEN initial black cells when there is no extra debris. No uniform onset bound depending only on this cardinality is implied; coordinates still encode n .

6.4.2 5.2 Reset to the original NE interface

The marker q3 is not read in that 9308-update recovery. It is the phase-15 first-visit cell exactly 10401 updates along the SE reference, and is read white/R there by the unmarked reference. The third initial black cell changes that turn. The complete finite block from (C_old,q*,S) has 23661 updates and ends at

$$v_out=(203,-244), E, \text{ standard NE phase } 94. \quad (5.2)$$

All its reads and its infinite NE tail stay in Omega. Two infinite criteria accept its complete 3453-black comparison endpoint. The same coupling proves the actual THREE-marker packet permanently standard by $208n+75609$ when no later packet interrupts it.

The initial q3 has not been manufactured by this process: omitting it gives the different, permanent SE future just proved. This explicit distinction prevents treating the reset as an autonomous instruction generator.

6.5 Single-packet theorem and protection under later interruptions

Theorem 1. Every complete state (2.2), for $n \geq 100$ and arbitrary finite H subset $\{s < G\}$, has the three-marker behavior just described and reaches NE standard phase 94 by time $208n+75609$. The two-marker state has the SE output from Subsection 6.4.1. All supplied black history remains physically present.

Proof. Subsections 6.2 and 6.3 give the exact A, B, C, D execution and first old read. Subsection 6.4 gives the exact whole-field restriction and protected finite E continuation and tail. The comparison paths remain within the regions on which field agreement was established, so a first divergent read is impossible. This is a finite-stage induction plus an infinite-tail certificate, not extrapolation from the sampled n values. $\square$

For the three-marker packet, every read BEFORE (5.2) satisfies

$$-109 \leq s \leq 59. \quad (6.1)$$

The upper bound also covers all initial M1 cells and the three markers. The NW and SE flights have constant s on each phase ray; their arbitrarily long lengths do not enlarge (6.1). The finite A, C, E blocks supply the other extrema. At the C boundary, (4.1) gives $x = -2n-39$, so no n -independent backward x -displacement or bounded-diameter packet footprint can replace this argument.

A later marker may interrupt the final NE tail. The single-packet theorem alone would not guarantee old-history protection after that interruption. The next section supplies the necessary nested-guard induction for the FULL composition.

6.6 Finite arithmetic connector and marker ordering

A packet outputs pose v_out , orientation NE, phase 94. After $m \geq 100$ full periods and phase adjustment to 57, its next primary marker is displaced by

$$\begin{aligned} \Delta(m) &= v_out + (m+1) \cdot d + P_{57} - P_{94} \\ &= (2m+208, 2m-242). \end{aligned} \quad (7.1)$$

Thus $s(\Delta(m)) = 4m - 34 \geq 366$. The next diagonal guard starts at $s(\Delta(m)) + G \geq 206$, while all previous packet black cells and reads not in inherited history have $s \leq 59$.

The minimum guard slack is 147, strictly positive. Every synthetic pre-output standard write is also below the next boundary; this is checked for each phase at its latest possible pre-output epoch, with all earlier epochs decreasing s further.

It follows that, at the next primary, the full inside field consists exactly of its M1, its three still-unread markers, and still later preinstalled markers. The actual NE flight writes and the mature synthetic writes agree there because the omitted pre-flight past is entirely outside. Inherited earlier history is below an even older guard and hence below the new one.

The companion markers of this next packet cannot interrupt that NE approach: $q_2(n)$ and q_3 have r respectively $-4n-26$ and 437 relative to the primary, outside ALL two-sided incoming NE phase lines. This is a geometric disjointness test, not merely checking a finite flight. Nor can markers from a still later packet cause a premature contact. Their minimum relative s is -83 ; (7.1) places them strictly above both the current packet's entire read maximum and the maximum s of an incoming flight before its primary. The latter is at most primary $s+1$, as checked from the 104 phases with their latest pre-collision epochs.

These inequalities prove the chronological order of all planned contacts, not just compatibility of consecutive endpoint poses. For $m > 100$ their left sides increase by $4(m-100)$; no further trajectory simulation is required.

6.7 Arbitrarily many genuine returns from one complete finite input

Choose ANY integer $k \geq 0$ and ANY integer lists

$m_1, \dots, m_k \geq 100$, $n_1, \dots, n_k \geq 100$.

Start with pose $p_0=0$, phase $a_0=0$, orientation NE. For $i=1, \dots, k$ set

$Q_i = p_{i-1} + (m_{i+1} [57 < a_{i-1}]) * d + P_{57} - P_{a_{i-1}};$
 $p_i = Q_i + v_{out}$, $a_i = 94$. (8.1)

Install at INPUT time all three cells

Q_i , $Q_i + q_2(n_i)$, $Q_i + q_3$.

The complete input is K plus these cells plus any finite black G_0 subset $\{s < 0\}$; its ant starts at 0,E. All these sets are disjoint, and the input has exactly $13 + 3k + |G_0|$ black cells. Markers are never installed during execution. For $k=0$ the initial state is already on the certified NE future.

Theorem 2 (closed arbitrary-return trains). Every such finite input follows the prescribed packets with all past writes retained and reaches permanent standard P104 by

$T = \sum_{i=1..k} [104m_i + (57 - a_{i-1} \bmod 104) + 208n_i + 75609]$. (8.2)

The output after $k > 0$ packets is p_k , heading E, NE phase 94. Each packet contains the genuine post-input-created black lifetime (4.4).

Proof. The canonical initial ray has $s \geq 0$, so it cannot read G_0 . The first primary guard is beyond every K cell and the entire synthetic pre-start past; at $m_1=100$ its slack is at least 243. Its companion markers miss the incoming NE rail, and later markers are beyond the finite read envelope. Therefore the first collision has the correct complete restricted field (2.2).

Apply Theorem 1 up to its output, but keep all later markers in the physical state. Subsection 6.6 proves that they have not altered this prefix, that the next collision has exactly the required restricted field, and that its guard is strictly above all preceding history. Induct on i . The guards strictly increase in the SAME absolute linear form s . Between collisions the previous certified tail remains above its previous guard; within each new packet all reads stay above the new guard. Consequently no later reset reactivates a retired region. After the last packet there is no remaining forward perturbation and the certified NE tail persists. Summing the exact block lengths proves (8.2). $\square$

The assertion about arbitrary finite G_0 is a proved initial guard, not a deletion of an unknown future-active set. Complete-state train replays retain G_0 and every black point from every earlier excursion.

6.7.1 Small input and repeated-return examples

With $m_1=n_1=100$ and no G_0 , the three markers are $(204, 205)$, $(-11, 416)$, $(381, -55)$.

Together with K this is a SIXTEEN-black finite seed. It writes $(166, 173)$ at time 8721 and next reads it black at time 83205, an exact lifetime of 74484 updates. It has a certified NE output at time 106866, at $(407, -39)$, with 7796 black cells.

The retained six-packet example has $m_i=100$ and $n_i = (100, 101, 150, 103, 211, 100)$. Its THIRTY-ONE-black input executes six such returns and has a certified NE output at time 675566, at $(2447, -249)$, with 50671 black cells. These are examples of the all- k proof, not a claim about all sixteen- or thirty-one-black initial states.

6.8 A finite-stage rank for return packets

The packet phases are an exact segmentation of a restricted true orbit, NOT PG secondary gates or intrinsic S_D reads. Use a finite-stage lexicographic rank (packets not yet completed, remaining stage index, local fuel).

The stage order is incoming NE flight, A, B, C, D, E. Each flight has the exact remaining update count from the formulas above. In each finite A, C or E block, choose its finite read rectangle, INCLUDING the last endpoint, and use the scaled integer $r=4$ clock already proved in clock.md. The unchanged exterior contribution cancels. Its computable upper bound over all internal fields and poses yields the nonnegative integer fuel

$$\text{ceil}((U_Q - C_Q(X))/\text{epsilon}_Q),$$

which decreases at every internal literal update. This remains valid for the TRUE complete field: unread differences inside that rectangle merely contribute constant terms, and the universal clock inequality holds for every colouring. Transitions to the next stage decrease the middle coordinate, allowing fuel to reset. Completing a packet decreases the first coordinate and starts the next incoming flight; completing the last packet halts the certificate program. The actual ant continues forever along its highway.

Unbounded excursion n makes the two flight counters large but does not create new unbudgeted stages: the exact XOR return query and full return-region equality

proved that there is exactly the certified old collision and recovery. That is the closure fact on which the rank relies. The clock did not prove the guard, the return-field equality, or all-input coverage.

This rank carries the finite packet description. It is not a new all-state physical rank, a global PG/S_D event bound, or a proof that arbitrary finite seeds enter this interface.

7 Verification and relation to earlier work

The verification package separates the finite data from the coverage proof. The original integer checker reconstructs the 284-seed table, the complete 182-step endpoint, and two all-future certificates. A complex-coordinate implementation independently checks the online colour tree and the blank and port prefixes. A third integer-vector implementation reads the immutable certificates, reconstructs family coverage, and checks per-cell chronological intervals. Corrupted exit data and missing branches are rejected, as are changes to the 13 active black cells and a blocker 10^{100} periods ahead. An off-trace marker is correctly accepted as dormant history.

The exterior constructor is also tested on 96 finite states, covering all four exit sides and including 16 cases translated by coordinates of magnitude 10^{30} . The tests verify the exact complete-field difference before exit, retention of the inside debris, and the all-future certificate after the placed port. These cases test the implementation; the universal statement is the geometric coupling proof above.

Finite circuit constructions for the classical ant were given by Gajardo, Moreira and Goles [1]. Etse studies confinement in finite domains [2], while our clock paper supplies the square bound used for the explicit completion time [5]. The one-black theorem establishes eventual standard-highway behaviour for unchanged singleton inputs [4]. The present contribution is the sharp initial cardinality threshold for a one-sided port, the exact three-cell exit bound, and a four-cell completion for arbitrary finite histories. The general periodic-realisability criterion [3] is used as a verification method; the new threshold is established by first-contact coverage.

8 Conclusion

The sharp four-cell port connects a local resource threshold with effective completion of finite histories. Guarded scattering then turns the outgoing trace into a composable interface. Its proof requires agreement on a complete protected region and a strictly advancing guard; matching endpoint poses alone would not suffice. The return construction extends that interface to unbounded excursions and actual readback of old self-written memory, while retaining the full field at every stage.

Each finite train has its markers fixed at input time. The number of packets is supplied by that finite input, while every gap and excursion parameter may be arbitrarily large. The separate feedback manuscript studies a twenty-black family with fixed intermediate event count; its lifetime obstruction belongs to that paper rather than being a second main claim here.

Reproducibility. The source, finite certificates and checkers are in [public project directory](#). From the repository root, run:

```
python3 four_cell_port/reproduce.py
make -C four_cell_port/paper
```

The checks use Python 3.10+ and the standard library. The proofs combine analytic arguments and exact finite computation. Input hashes and the scope of each check are recorded with the verification evidence.

AI assistance. Generative AI assisted with developing and checking arguments, implementing verification programs, and revising the exposition.

References

- [1] A. Gajardo, A. Moreira and E. Goles, “Complexity of Langton’s ant,” *Discrete Applied Mathematics* 117 (2002), 41–50. [doi:10.1016/S0166-218X\(00\)00334-6](#).
- [2] K. R. Etse, “How Long Can The Escaping Ant Be Confined?,” *MFCS 2026*, LIPIcs 386, 49:1–49:17. [doi:10.4230/LIPIcs.MFCS.2026.49](#).
- [3] A. Jillhewar, “Finite-Support Periodic Highways of Langton’s Ant: Necessary Conditions, Transverse Exclusions, and Exact Search,” preprint (2026). [doi:10.5281/zenodo.21381637](#).
- [4] H. Ke, “Every One-Black-Cell Langton Ant Reaches the Period-104 Highway: A Lean-Checked Proof,” Zenodo, version 0.2.0 (2026). [doi:10.5281/zenodo.23341609](#).
- [5] H. Ke, “An Ant on the Clock: Exponential Escape Bounds for Langton’s Ant,” Zenodo, version 0.4.0 (2026). [doi:10.5281/zenodo.23402379](#).