FINITE-SUPPORT PERIODIC HIGHWAYS OF LANGTON'S ANT: NECESSARY CONDITIONS, TRANSVERSE EXCLUSIONS, AND EXACT SEARCH

ATHARVA JILLHEWAR

ABSTRACT. We study the classical two-colour Langton's ant on $\mathbb{Z}^2$, started from a colouring with finite black support, under the additional hypothesis that its turn trace is eventually periodic with nonzero drift. We give a necessary-and-sufficient, decidable criterion for a finite word to have such a finite-support periodic realisation, together with an explicit seed. We prove that every resulting highway has strictly positive net black growth divisible by four, and derive a signed mod-four residue identity for its growing wake and the bound $g \geq 2\max(|a|, |b|)$ for drift (a, b). We also give an exact medial (Tait)-graph formulation and a two-endpoint collision-chain parity identity. For diagonal drift we prove a transverse rigidity theorem: the two extremal level lines of the trace are horizontal, carry only the turn R, and every cell of them reached by the periodic tail is entered exactly once and never revisited within that tail, so the highway is bounded on both sides by guard rails of permanent wake; the transverse width is therefore even. From this we deduce that no such highway has transverse width two. We then give a computer-assisted all-period exclusion of width four: exact crossing sequences at successive untouched five-cell columns belong to a twelve-edge directed graph, and every edge strictly lowers an explicit rank. Two independently written local enumerators reproduce the complete edge table, while Lean checks its rank consequence. Thus every diagonal periodic highway has transverse width at least six. A residue-theorem-free audited enumeration applies the realisability criterion to every positive-growth, nonzero-drift leaf and excludes all periods at most 48. Selected finite algebraic kernels are checked in Lean 4; the geometric extraction of those kernels from an ant trace is proved conventionally here but is not yet formalised end to end. These results constrain periodic highways but prove neither that every finite-support orbit eventually becomes periodic nor that the standard period-104 highway is the unique positive-growth periodic trace.

1. INTRODUCTION

Langton's ant is the turmite on the square lattice $\mathbb{Z}^2$ defined as follows. Each cell is coloured white or black; a pointer, the *ant*, occupies a cell and faces one of the four cardinal directions. At each step the ant reads the colour of its current cell: on white it turns 90° right, on black it turns 90° left; it then flips the colour of the current cell and moves one unit forward in its new heading. In the convention used here, direct simulation from the all-white plane enters, after 9977 steps, a diagonal translating structure—the *highway*—whose turn trace has least period 104 and whose ant position translates by a vector $(\pm 2, \pm 2)$ each period. The rule was introduced by Langton [7]; the exact phase used here is reproduced in Appendix A.

Throughout we impose the *finite-support* condition: only finitely many cells are initially black. The empirical universality of the highway is the subject of a well-known open problem.

Conjecture 1.1 (Highway conjecture)**.** For every finite initial black set and every initial pose, there is a time after which the ant's future turn trace and path agree with the standard period-104 highway up to lattice translation, quarter-turn rotation, and cyclic phase shift. Finite cells outside the future path may remain as stationary debris; equality of the entire colouring is not asserted.

Date: 21 July 2026.

2020 *Mathematics Subject Classification.* Primary 37B15; Secondary 37B10, 05C10, 68Q80, 68R15.

Key words and phrases. Langton's ant, turmite, cellular automata, highway conjecture, symbolic dynamics, reversible dynamics, medial graph, Tait graph, interlacement.

1.1. **Prior art.** The rigorous facts about the classical ant are few. Every trajectory is *unbounded* (Bunimovich–Troubetzkoy [1]; the underlying checkerboard/parity structure is already in Gale–Propp–Sutherland–Troubetzkoy [6]). The ant is computationally universal, and associated prediction problems are undecidable, on suitable infinite backgrounds [4]; we draw no inference from this about Conjecture 1.1, which concerns finite configurations, beyond the evident point that no finite cutoff exhausts an infinite range. Gajardo [5] gives an algebraic characterisation of the forbidden factors of the ant's trace subshift. Recent work on *generalised* ants exhibits rules with infinitely many highways [3] and rules with non-highway emergent behaviour [8]. Quantitative confinement bounds for the classical ant are studied in [2]. The present paper concerns a different, conditional question: what restrictions follow once a classical finite-support orbit is assumed to have a periodic translating tail?

1.2. **Results.** A *repeating trace* is an infinite turn word w^ω realised from a finite-support colouring (Section 3), where the finite word w resets the heading; its *drift* $d = (a, b)$ is the per-period displacement and its *growth* g the per-period change in the number of black cells. The main results are:

(1) **Even-winding theorem** (Theorem 6.1): no finite-support clean translator with nonzero drift exists. Hence every finite-support periodic highway has growth $g \in 4\mathbb{Z}_{>0}$ (Corollary 6.5). Bunimovich–Troubetzkoy excludes only *bounded* repeatable orbits; a zero-growth glider is unbounded and is not covered by their theorem. (The divisibility by 4 alone is the elementary parity Lemma 2.5.)

(2) **Signed mod-four wake-residue identity.** Theorem 7.1; Corollary 7.3 gives the lower bound

$$g \geq 2\max(|a|, |b|).$$

(3) **Tait-graph conjugacy** (Theorem 4.1) and the cycle-rank surgery identity $E + \Delta C = 2\beta$ (Theorem 4.2), built on the checkerboard/HV partition of [6].

(4) **Collision-chain parity** (Theorem 8.1), a two-endpoint $\mathbb{F}_2$ boundary identity for colour defects.

(5) **Transverse rigidity and width-two/width-four exclusions** (Theorems 9.1, 9.3, and 9.5). For diagonal drift the two extremal level lines of a periodic highway are horizontal, every turn on them is R, and each cell reached on them is visited exactly once by the periodic tail and is left black; the transverse width is even. No such highway has transverse widths two or four; hence its width is at least six. Unlike (6), these statements are not bounded period searches. The width-four theorem is computer-assisted: two independent exact local enumerators certify a finite twelve-edge blank-column graph, and Lean certifies that every listed edge lowers rank. The local enumeration's completeness is not yet formalised in Lean. The proofs are independent of (2).

(6) **Small-period exclusion** (Theorem 10.1): a rank-audited exhaustive search excludes every finite-support periodic highway of nonzero drift and period at most 48. The search assumes no consequence of the residue theorem and applies the criterion to every positive-growth nonzero-drift leaf, so the exclusion does not depend on Section 7; faster variants that do assume it agree. The certifying search is deliberately separated from the faster conditional cross-checks described in Appendix C.

Section 3 also gives a decidable periodic-realisability criterion with an explicit finite seed, strengthening the forbidden-word program of [5] from local factor-admissibility to global periodic finite-support realisability. The elementary heading/black-count invariant (Lemma 2.5) is folklore, implicit in [1, 6].

1.3. **Scope.** Every theorem below is a *necessary condition* on a finite-support periodic highway, or an exact finite computation. None asserts that any orbit converges, and therefore none resolves Conjecture 1.1. What we prove is that any finite-support periodic highway must build a positive wake divisible by four and must satisfy the residue identities and the density bound $g \geq 2\|d\|_\infty$, and that no finite-support periodic highway has period at most 48 (equivalently, any such highway has period at least 50). Two logically separate gaps remain: an arbitrary finite-support orbit has not been shown to enter a translating periodic regime, and positive-growth finite-support periodic traces have not been classified. Section 11 states both problems explicitly.

2. Setup and elementary facts

A *state* is a triple $s = (B, p, q)$ with $B \subset \mathbb{Z}^2$ finite (the black set), $p \in \mathbb{Z}^2$ (the ant position) and $q \in \mathbb{Z}/4\mathbb{Z}$ the heading, numbered $0 = N$, $1 = E$, $2 = S$, $3 = W$, with unit vectors

$$u_0 = (0, 1), \quad u_1 = (1, 0), \quad u_2 = (0, -1), \quad u_3 = (-1, 0).$$

The *turn* at s is R (right) with $\varepsilon = +1$ if $p \notin B$, and L (left) with $\varepsilon = -1$ if $p \in B$. One step of the dynamics is

$$s = (B, p, q) \;\longmapsto\; (B \triangle \{p\},\ p + u_{q+\varepsilon},\ q + \varepsilon), \tag{1}$$

where $\triangle$ denotes symmetric difference and indices are read modulo 4. The *trace* of an orbit is the infinite word of turns encountered. Since (1) is invertible (step backward one cell, recolour, undo the turn), the dynamics is reversible and orbits never merge. Write $b_t = |B_t|$ and let $\chi(x, y) = (-1)^{x+y}$ be the checkerboard sign.

Definition 2.1 (Periodic terminology)**.** Let $w = w_0 \cdots w_{N-1}$ be a nonempty turn word and follow its induced lattice path from a specified initial heading. The word is *heading-resetting* if the heading after N turns equals the initial heading. In that case its *drift* is $p_N - p_0$, and its *growth* is $\#\{i : w_i = \mathrm{R}\} - \#\{i : w_i = \mathrm{L}\}$. An orbit has an *eventually periodic turn trace* if, from some time t_0, its turns are w^ω. If w is heading-resetting with nonzero drift, we call that tail a *periodic highway*; its *least trace period* is the smallest possible N. Highways are compared up to lattice translation, quarter-turn rotation, and cyclic phase shift of w.

Proposition 2.2 (Checkerboard/HV partition [6])**.** *The quantity* $\big(q_t - (p_{t,x} + p_{t,y})\big) \bmod 2$ *is constant along every orbit. Consequently each cell is entered by the ant always horizontally or always vertically. Calling a cell* vertical *(resp.* horizontal*) accordingly, at a vertical cell the two vertical incident edges are the only admissible arrival edges and the two horizontal ones the only departure edges, and vice versa; every admissible directed lattice edge thus carries a fixed orientation.*

Proof. Each step changes the heading by an odd quarter-turn, so $q_t \bmod 2$ alternates; the displacement $u_{q_{t+1}}$ has odd ℓ^1-length, so $(p_{t,x} + p_{t,y}) \bmod 2$ alternates as well. Their difference is therefore invariant. Fixing a cell z: at every visit the checkerboard parity of z is the same, so $q_t \bmod 2$ is the same, i.e. the arrival axis is fixed. The edge orientations follow because an arrival edge of one endpoint is a departure edge of the other. □

Proposition 2.3 (Alternating visits)**.** *At any fixed cell the encountered turns alternate strictly between* R *and* L*. All but finitely many cells begin with* R*.*

Proof. Each visit flips the cell's colour and the turn is a function of the colour, so the turn alternates. A cell white at its first visit begins with R; only the finitely many initially black cells can begin with L. □

Theorem 2.4 (Unboundedness [1])**.** *Every finite-support orbit visits infinitely many cells.*

Proof sketch. Were the visited set finite, determinism and reversibility would make the orbit periodic in the full state. Choose a visited cell z maximising $x + y$; its north and east neighbours are unvisited. If z is horizontal it can only be entered from the west. During a full state period its colour returns, so its alternating visits include both a white/R and a black/L visit; these two turns leave respectively south and north, forcing a visit to the unvisited north neighbour. If z is vertical it can only be entered from the south, and the two colours force departures east and west, including the unvisited east neighbour. Either case is a contradiction. □

Lemma 2.5 (Heading/black-count invariant; folklore)**.** *For every t,*

$$q_t - b_t \equiv q_0 - b_0 \pmod 4, \qquad p_{t+1} - p_t = u_{\,q_0 - b_0 + b_{t+1}}. \tag{2}$$

A heading-resetting period changes b by a multiple of 4.

Proof. A white/R step increments both q and b by 1; a black/L step decrements both. Hence $q_t - b_t$ is constant. Since the move follows the turn, $q_{t+1} \equiv q_0 - b_0 + b_{t+1}$, giving the displacement formula. If $q_N = q_0$ then $b_N - b_0 \equiv 0 \pmod 4$. □

3. A constructive periodic-realisability criterion

We first record the exact test that turns "is w^ω a highway?" into a finite computation. Fix an initial heading q_0 and a proposed turn word $w = w_0 \cdots w_{N-1} \in \{\mathrm{R}, \mathrm{L}\}^N$. Compute the induced path $p_0, \ldots, p_N$ from w (Proposition 2.2 makes each step well defined). Suppose $q_N = q_0$ and $d = p_N - p_0 \neq 0$. Define an equivalence on phases by

$$i \sim j \iff p_i - p_j \in \mathbb{Z}d,$$

and in each class I write $p_i = b_I + a_i d$ for a fixed representative b_I and integers a_i. For $n \geq \min_I a_i$, let $S_{I,n}$ be the subword of all w_i $(i \in I)$ with $a_i \leq n$, listed in increasing order of the key $(n - a_i,\, i)$.

Theorem 3.1 (Periodic-realisability criterion)**.** *The infinite word w^ω is the trace of Langton's ant from some finite-support colouring if and only if, for every class I:*

(i) *every $S_{I,n}$ with $\min_I a_i \leq n \leq \max_I a_i$ alternates strictly between* R *and* L*; and*
(ii) *the stabilised word $S_{I,\max_I a_i}$ begins with* R*.*

When these hold, an explicit finite black seed producing w^ω is

$$B = \{\, b_I + nd \;:\; \min_I a_i \leq n < \max_I a_i \text{ and } S_{I,n} \text{ begins with L} \,\}, \tag{3}$$

the union over all classes I; the bound $n < \max_I a_i$ together with $n \geq \min_I a_i$ makes B finite and $S_{I,n}$ well defined.

Proof. The occurrence of phase $i \in I$ in period $k \geq 0$ sits at $p_i + kd = b_I + (a_i + k)d$, which equals $b_I + nd$ exactly when $k = n - a_i \geq 0$; chronological order among these occurrences is the lexicographic order of $(n - a_i, i)$. Thus $S_{I,n}$ is precisely the chronological turn sequence at the physical cell $b_I + nd$ under w^ω.

Necessity. Each visit flips the cell, so its turn sequence alternates, proving (i). For $n \geq \max_I a_i$ all phases of I are present and the ordering no longer depends on n: the sequence has stabilised, and the infinitely many cells $b_I + nd$ $(n \geq \max_I a_i)$ share this first turn. By Proposition 2.3 only finitely many cells begin with L, so the stable first turn is R, proving (ii).

Sufficiency. Colour exactly the set (3) black. At each boundary cell $b_I + nd$ with $n < \max_I a_i$ the chosen colour reproduces the first symbol of $S_{I,n}$, and strict alternation supplies every subsequent symbol because the cell flips at each visit. At every stabilised cell $(n \geq \max_I a_i)$ condition (ii) makes the first encountered turn R, matching its white initial colour. Induction on chronological time shows the ant follows w^ω forever; since $q_N = q_0$ and the position advances by $d \neq 0$ each period, this is a highway. □

Remark 3.2. There is an equivalent single-sequence test: within a class, sort all entries by decreasing a_i, breaking ties by increasing i; every $S_{I,n}$ is a suffix of the result obtained by deleting whole high-level blocks. It therefore suffices to check that this one stabilised word alternates and begins with R, at the cost of one sort per class. To index the classes of $\mathbb{Z}^2/\mathbb{Z}d$ exactly, write $g = \gcd(a, b)$ and $d = g\,(a', b')$ with $\gcd(a', b') = 1$, and choose integers r, s with $ra' + sb' = 1$ (Bézout). Then

$$(\, b'x - a'y,\; (rx + sy) \bmod g \,)$$

is a *complete* invariant of the coset of (x, y): the matrix $\left(\begin{smallmatrix} b' & -a' \\ r & s \end{smallmatrix}\right)$ has determinant $b's + a'r = 1$, so it is unimodular, and in its coordinates translation by $d = g(a', b')$ becomes translation by $(0, g)$, leaving the first coordinate and the second coordinate modulo g fixed. (The naive residue $a'x + b'y \bmod g$ is *not* complete: for $d = (2, 2)$ the cells $(0, 0)$ and $(1, 1)$ share both $b'x - a'y = 0$ and $a'x + b'y \equiv 0 \pmod 2$ yet lie in different cosets, since $(1, 1) \notin \mathbb{Z}(2, 2)$; the Bézout coordinate $rx + sy$ separates them.) This is the checker used in all computations of Section 10; it strengthens the forbidden-word characterisation of the ant's trace subshift [5] to a decidable test for global periodic finite-support realisability with a constructive seed. A worked instance (the standard highway) is given in Appendix A.

4. The Tait-graph conjugacy

We recast the dynamics on a medial (Tait) graph, exposing topology invisible in the raw colouring. Normalise the HV partition of Proposition 2.2 so that even cells receive vertical arrivals. In the *frozen* routing (the permutation of directed states obtained by ignoring colour flips) the admissible directed arcs form clockwise directed four-cycles around one checkerboard class of unit plaquettes. Label a plaquette by its lower-left corner and take these plaquettes as the vertices of a rotated square lattice. Each ant cell $z = (x, y)$ lies at the crossing of exactly two white four-cycles and corresponds to the *Tait edge*

$$e_z = \begin{cases} \{(x, y-1), (x-1, y)\}, & x + y \text{ even}, \\ \{(x-1, y-1), (x, y)\}, & x + y \text{ odd}. \end{cases} \tag{4}$$

Let $H_B = \{e_z : z \in B\}$ be the finite edge subgraph induced by the black cells.

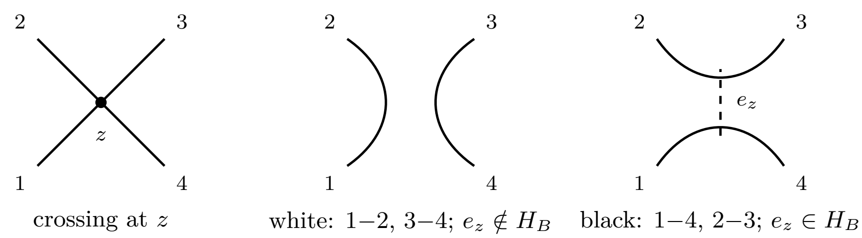

FIGURE 1. The two smoothings of the frozen crossing at an ant cell z. The four incident frozen arcs are labelled $1, 2, 3, 4$. Colour selects which of the two ways they reconnect: white joins $1{-}2$ and $3{-}4$, black joins $1{-}4$ and $2{-}3$. The black smoothing is exactly the occupied Tait edge e_z (dashed).

Theorem 4.1 (Boundary-loop conjugacy)**.** *White and black at z are the two smoothings of the directed crossing at the Tait edge e_z* (Figure 1)*. Consequently:*

- (a) *the frozen routing loops are exactly the boundary components of a regular neighbourhood of H_B, together with the untouched four-cycles;*
- (b) *the ant advances along its current frozen boundary loop and then toggles the edge e_z; and*
- (c) *each step either merges two frozen loops into one or splits one into two.*

Proof. Frozen (colour-ignoring) routing is the permutation ρ of directed edges of the ant lattice sending each directed edge to the next one the ant would follow under a fixed local turn rule; since a white cell turns right and a black cell turns left, at each ant cell z the local routing is one of the two ways of connecting the four directed edge-ends incident to z that avoids a U-turn. These are precisely the two *smoothings* of the 4-valent vertex z of the medial graph, and the Tait edge e_z of (4) records which diagonal pair of plaquette-corners the crossing joins; white is one smoothing and black the other, which is (a) at the level of a single crossing.

Now freeze a colouring B. Away from H_B every plaquette crossing is white, and the white smoothings assemble into the boundary circuits of a regular neighbourhood $\mathcal{N}$ of the edge subgraph H_B: each occupied edge e_z is thickened to a band, each Tait vertex to a disk, and the frozen routing runs along $\partial\mathcal{N}$, with the untouched white plaquettes contributing the trivial four-cycles. (This is the standard medial/Tait correspondence between a plane graph and the boundary of a ribbon neighbourhood.) This gives (a) globally: the frozen loops are the components of $\partial\mathcal{N}$ plus the untouched four-cycles. Statement (b) is immediate: the live ant sits on the current boundary loop and, on stepping off z, toggles the colour of z, i.e. replaces the smoothing at that single crossing—adding or deleting the band e_z in $\mathcal{N}$.

For (c), toggling e_z is the addition or deletion of one band at a 4-valent vertex, which is a single Dehn-twist-free band surgery on $\partial\mathcal{N}$. Locally it reconnects the two boundary arcs passing the vertex: if the two arcs belong to the same boundary component the surgery splits it into two, and if they belong to different components it merges them into one. There are no other cases, giving (c). □

Let $E = |B|$ be the number of occupied Tait edges, V the number of incident Tait vertices, k the number of nonempty components, and $\beta = E - V + k$ the cycle rank.

Theorem 4.2 (Cycle-rank surgery identity)**.** *A regular neighbourhood of H_B has $k+\beta = E-V+2k$ boundary components. Writing ΔC for the change in loop count relative to the V trivial plaquette loops,*

$$\Delta C = E - 2V + 2k, \qquad E + \Delta C = 2\beta. \tag{5}$$

Adding a bridge merges two loops and fixes β; adding a cycle edge splits a loop and raises β by one; removing a bridge splits a loop and fixes β; removing a non-bridge merges two loops and lowers β by one.

Proof. A regular neighbourhood of a plane graph with V vertices, E edges, and k components is a compact planar surface with k genus-zero pieces and Euler characteristic $V - E$. If it has b boundary circles, then $V - E = 2k - b$, so $b = E - V + 2k = k + \beta$, where $\beta = E - V + k$ is the first Betti number. Formula (5) is the difference from the V trivial plaquette loops. The four surgery cases are the standard effect on the boundary of adding or deleting a bridge (Betti number unchanged; boundary count changes by ∓ 1) versus a cycle edge (Betti number changes by ± 1; opposite merge/split), read off from Theorem 4.1(c). □

Proposition 4.3 (Advanced records are bridges)**.** *Once a new bounding-box record lies strictly beyond the bounding box of the finite initial support, its first visit is white, and a new east, north, west, or south record causes the ant to turn south, east, north, or west, respectively. In the Tait graph the corresponding edge has a previously untouched endpoint, hence is a bridge. Therefore sufficiently advanced record creation never raises β; any unbounded production of cycle rank occurs strictly inside already-explored territory.*

Proposition 4.4 (No monotone affine invariant)**.** *No nonconstant affine combination of (E, V, k, β) is monotone along the blank orbit. In particular, each of E, V, k, and β separately fails to be monotone. Thus no affine combination of these four graph statistics can serve as a monotone or Lyapunov potential along even the blank orbit.*

Proof. Since $\beta = E - V + k$ identically, $aE + bV + ck + d\beta + e = (a+d)E + (b-d)V + (c+d)k + e$, so the affine combinations of (E, V, k, β) are exactly those of (E, V, k), and such a combination is nonconstant precisely when $\lambda := (a+d,\ b-d,\ c+d) \in \mathbb{R}^3$ is nonzero.

One step flips one cell, hence inserts or deletes exactly one Tait edge $e_z = \{u, v\}$. Insertion admits exactly four cases, according to how many endpoints of e_z already meet an edge of H_B and, when both do, whether they lie in a common component; deletion reverses each. This is a complete case list, not a sample:

insertion of $e_z = \{u, v\}$	$(\Delta E, \Delta V, \Delta k)$	$\Delta\beta$
u, v both new	$(1, 2, 1)$	0
exactly one of u, v new	$(1, 1, 0)$	0
u, v both old, distinct components	$(1, 0, -1)$	0
u, v both old, same component	$(1, 0, 0)$	1

Hence every increment lies in

$$S_0 = \pm\{(1, 2, 1), (1, 1, 0), (1, 0, -1), (1, 0, 0)\}, \tag{6}$$

and, incidentally, β rises exactly when the inserted edge closes a cycle and is unchanged otherwise. All eight members of S_0 are realised on the blank orbit—every state of which is compact—within the first 12,000 steps, so the realised increment set is exactly S_0.

Now S_0 is centrally symmetric by construction, since deleting an edge negates the increment of inserting it, and S_0 contains the linearly independent $(1, 1, 0), (1, 0, -1), (1, 0, 0)$, so $\operatorname{span} S_0 = \mathbb{R}^3$. Let $\lambda \neq 0$. Spanning gives $v \in S_0$ with $\lambda \cdot v \neq 0$; replacing v by $-v \in S_0$ if necessary, $\lambda \cdot v > 0$, while $\lambda \cdot (-v) < 0$ and $-v$ is realised as well. Hence λ strictly increases at one step of the orbit and strictly decreases at another, so it is monotone in neither direction. Equivalently, 0 lies in the interior of $\operatorname{conv} S_0$. The realisability witnesses are reproduced by `verify_nomonotone.py` (Appendix C). □

5. Additive charges and clean translators

Definition 5.1. A *clean finite translator with nonzero drift* is a finite state for which, for some $N > 0$ and $d \in \mathbb{Z}^2 \setminus \{0\}$,

$$B_N = B_0 + d, \qquad p_N = p_0 + d, \qquad q_N = q_0.$$

The word *clean* means the entire black set translates: no stationary residue and no newly emitted wake. This is not the ordinary highway, which grows a wake.

5.1. **Laurent charges and axis alignment.** Work in the Laurent polynomial ring $\mathbb{F}_2[X^{\pm 1}, Y^{\pm 1}]$, an integral domain. Let $V_t = 1$ if q_t is vertical and $H_t = 1 - V_t$, and set

$$A_t = \sum_{(x,y) \in B_t} X^x + V_t X^{x_t}, \qquad D_t = \sum_{(x,y) \in B_t} Y^y + H_t Y^{y_t}.$$

Proposition 5.2. *A_t and D_t are conserved. Hence a clean translator with $d = (a, b)$ satisfies $A = X^a A$ and $D = Y^b D$; since $A(1) = |B| + V$ and $D(1) = |B| + H$ cannot both vanish in $\mathbb{F}_2$ (because $V + H = 1$), the drift coordinates a and b are not both nonzero. Every clean translator is axis-aligned.*

Proof. Conservation of A_t: on a step from a vertical cell, the toggle at column x_t cancels the disappearing vertical correction X^{x_t}; on a step from a horizontal cell, the post-turn motion is vertical, so $x_{t+1} = x_t$ and the toggle cancels the newly appearing correction. Translation covariance then gives $A = X^a A$. In the integral domain $\mathbb{F}_2[X^{\pm 1}, Y^{\pm 1}]$, $a \neq 0$ forces $A = 0$ and $b \neq 0$ forces $D = 0$; but $A(1) + D(1) = 2|B| + 1 = 1 \neq 0$. □

Proposition 5.3 (Additive-charge cocycle)**.** *Let A be an abelian group and $w : \mathbb{Z}^2 \to A$. If there exist heading corrections $F_q : \mathbb{Z}^2 \to A$ (indexed by $q \in \mathbb{Z}/4\mathbb{Z}$) such that $\sum_{z \in B} w(z) + F_q(p)$ is conserved under*

$$F_{q+1}(z + u_{q+1}) - F_q(z) = -w(z) \quad \text{(R)}, \qquad F_{q-1}(z + u_{q-1}) - F_q(z) = +w(z) \quad \text{(L)}, \tag{7}$$

then w satisfies both the four-corner condition

$$w(x{+}1, y{+}1) + w(x{+}1, y) + w(x, y{+}1) + w(x, y) = 0 \quad \text{in } A \tag{8}$$

and $4w = 0$. Conversely, if $A = \mathbb{Z}/4\mathbb{Z}$ (so $4w = 0$ holds automatically) and w satisfies (8), then such F_q exist. Over $\mathbb{F}_2$ the solutions of (8) are exactly $w(x, y) = \alpha(x) + \beta(y)$.

Proof. Necessity. Eliminating the four F_q from (7) around a unit plaquette gives the four-corner condition (8). For $4w = 0$, combine the four relations at a single cell z: the R-at-S and L-at-N relations both express $F_W(z + u_W)$, forcing $F_S(z) = F_N(z) + 2w(z)$, while the L-at-S and R-at-N relations force $F_S(z) = F_N(z) - 2w(z)$; hence $4w(z) = 0$. Thus $w(x, y) = (-1)^{x+y}$ over $A = \mathbb{Z}$ satisfies (8) but admits no such F_q, since $4w \neq 0$.

Sufficiency for $A = \mathbb{Z}/4\mathbb{Z}$. Composing pairs of relations in (7) expresses F_N by its diagonal increments

$$\begin{aligned} F_N(x{+}1, y{+}1) - F_N(x, y) &= -w(x, y) + w(x{+}1, y), \\ F_N(x{+}1, y{-}1) - F_N(x, y) &= w(x, y{-}1) - w(x{+}1, y{-}1). \end{aligned} \tag{9}$$

which define F_N on each of the two $x{+}y$-parity sublattices by integration from a basepoint. Modulo four, the circulation of (9) around the elementary diamond $(x, y) \to (x{+}1, y{+}1) \to (x{+}2, y) \to (x{+}1, y{-}1) \to (x, y)$ equals the negative of the sum of the four-corner expressions (8) based at $(x, y{-}1)$ and $(x{+}1, y{-}1)$ (the two differ by $4\,[w(x{+}1, y) + w(x{+}1, y{-}1)] \equiv 0$); hence it vanishes by (8), so the integration is path-independent and F_N is well defined. Now set $F_S = F_N - 2w$ and define F_E, F_W from F_N by (7). Substituting these definitions back into the four R/L relations verifies each case in turn; the only possible ambiguity is between the two ways of reaching F_S—via the R-at-S/L-at-N pair, which gives $F_N + 2w$, and via the L-at-S/R-at-N pair, which gives $F_N - 2w$—and their discrepancy is exactly $4w$, which vanishes. Hence all four F_q are consistent. Over $\mathbb{F}_2$ the solution space of (8) is the separable weights $\alpha(x) + \beta(y)$. □

5.2. **The strand decomposition.** Let $g = \#\mathrm{R} - \#\mathrm{L}$ over one period. Lemma 2.5 gives $g \equiv 0 \pmod 4$, and $g \geq 0$ because g equals the number of *odd* stabilised translation classes of Theorem 3.1 (a class is odd iff its stabilised word has odd length, contributing one net black cell per period). We record the endpoint structure used in Section 7.

Proposition 5.4 (Canonical endpoint decomposition)**.** *The one-period signed toggle pattern $D : \mathbb{Z}^2 \to \mathbb{Z}$ (with $D(z) = \#\{\mathrm{R}$ at $z\} - \#\{\mathrm{L}$ at $z\}$ over one period) decomposes as*

$$D = S + (T_d - 1)A, \qquad S = \sum_{i=1}^{g} \delta_{c_i}, \tag{10}$$

where $(T_d A)(z) = A(z - d)$, A is a finitely supported moving widget*, and $c_1, \ldots, c_g \in \mathbb{Z}^2$ are the bases of the g growing strands, one per odd stabilised class. For the standard word $g = 12$ and $|B_n| = 11 + 12n$ at period boundaries.*

Proof. D is finitely supported (only finitely many cells are visited in one period). Partition the support into translation classes under $z \mapsto z + d$; there are finitely many nonempty classes. Fix a class with base point c and write $f(n) = D(c + nd) \in \mathbb{Z}$, a finitely supported function $\mathbb{Z} \to \mathbb{Z}$. Put $s = \sum_n f(n)$ and $h = f - s\,\delta_0$, so $\sum_n h(n) = 0$. With the shift $(Ta)(n) = a(n - 1)$, set

$$a(n) = \sum_{m>n} h(m).$$

Then $a(n - 1) - a(n) = h(n)$, i.e. $(T - 1)a = h$, so

$$f = s\,\delta_0 + (T - 1)a,$$

and a is finitely supported: $a(n) = 0$ for n beyond the support of f on the right, and $a(n) = \sum_m h(m) = 0$ for n far to the left (this is exactly where $\sum h = 0$ is needed). Summing over all classes gives (10) with A the (finite) sum of the class antiderivatives and $S = \sum_{\text{classes}} s_c\,\delta_c$. It remains to identify $\sum_n f$ for a class. Over one period the cells of a class are visited in level order, and the alternating-visit rule (Proposition 2.3) makes the per-class toggle total equal to $+1$ when the class's stabilised word has odd length (an "odd class") and 0 when it has even length; thus $\sum_n f \in \{0, 1\}$ and S is a sum of g unit masses, one per odd class, where $g = \#\{\text{odd classes}\} = \#\mathrm{R} - \#\mathrm{L}$ (the last equality because both count the net black cells created per period). The standard-word figures follow from $|B_0| = 11$ and $|B_{n+1}| - |B_n| = g = 12$. ☐

Corollary 5.5. *If $g = 0$ then a finite stationary set can be deleted to leave a clean finite translator with drift d.*

Proof. Here $S = 0$, so $D = (T_d - 1)A$. Let $b_n : \mathbb{Z}^2 \to \{0, 1\}$ be the indicator of B_n at the n-th period boundary. Since each cell is toggled a net $D(z) \in \{-1, 0, +1\}$ per period (Proposition 2.3) and the period-k toggle pattern is $T_{kd}D$ by translation covariance,

$$b_n = b_0 + \sum_{k=0}^{n-1} T_{kd}D = b_0 + \sum_{k=0}^{n-1} T_{kd}(T_d - 1)A = b_0 + (T_{nd} - 1)A = R + T_{nd}A, \qquad R := b_0 - A.$$

Both R and A are finitely supported, so for n large the supports of R and $T_{nd}A$ are disjoint; since b_n is $\{0, 1\}$-valued, R and A are then $\{0, 1\}$-valued and $B_n = R \sqcup (A + nd)$. Because the trace is periodic with nonzero drift, each cell is visited finitely often, so at a sufficiently late period boundary the ant never again visits the finite set R. Deleting R there changes no future turn and leaves a state whose black set is $A + nd$, translating exactly by d each period: a clean finite translator. ☐

Sections 6 and 7 prove, by two independent routes, that no such clean translator exists; hence $g \neq 0$.

6. The even-winding theorem

This section proves that no clean translator exists, via a cylinder homology argument. The one step that was only sketched in earlier accounts—the local parity identity relating the turn count at a cell to its horizontal edge parity—is proved in full as Lemma 6.4.

Theorem 6.1 (Even-winding theorem)**.** *No finite-support clean translator with nonzero drift exists.*

For the remainder of the section fix a hypothetical clean translator. Its drift is axis-aligned by Proposition 5.2. A clean translator has $g = 0$, so its period $N = \#\mathrm{R} + \#\mathrm{L} = 2\#\mathrm{R}$ is even; since each step changes $x + y$ modulo 2, the total displacement satisfies $a + b \equiv N \equiv 0 \pmod 2$, so the nonzero drift coordinate is even. After a quarter-turn rotation write $d = (L, 0)$ with $L > 0$ even—all that the argument below requires. Writing B_j, p_j for the black set and ant position at phase j of one period, reversibility and translation covariance give, for all $n \in \mathbb{Z}$,

$$B_{nN+j} = B_j + nd, \qquad p_{nN+j} = p_j + nd.$$

Quotient the bi-infinite trace by translation through d to obtain a closed walk on the cylinder $C_L = (\mathbb{Z}/L\mathbb{Z}) \times \mathbb{Z}$; because L is even the HV type of Proposition 2.2 descends to C_L. For an undirected cylinder edge, its *multiplicity* is the number of traversals by the projected one-period walk.

Lemma 6.2 (Visit count)**.** *Every vertex v of C_L is visited $n_v = 2k_v$ times in one quotient period, with exactly k_v turns* R *and k_v turns* L*.*

Proof. Let z be a physical representative of v. Each phase j with $p_j \in v + \mathbb{Z}d$ has a unique level m_j with $p_j = z + m_j d$, and phase j in period $-m_j$ visits z; conversely every visit to z over the bi-infinite orbit arises once this way. Hence the visits to v in one quotient period biject with the visits to z over the translated bi-infinite orbit. At phase $nN + j$ the black pattern is the finite translate $B_j + nd$, so z is white for all sufficiently early and all sufficiently late times. Its turn sequence therefore alternates, begins with R, ends with L, and contains equally many of each. ☐

At a vertex v let $H_{v,-}, H_{v,+}$ be the multiplicities of the west and east horizontal edges and $V_{v,-}, V_{v,+}$ those of the south and north vertical edges. By Proposition 2.2 one pair carries all arrivals and the other all departures, and each visit is one arrival plus one departure, so

$$H_{v,-} + H_{v,+} = V_{v,-} + V_{v,+} = 2k_v. \tag{11}$$

Lemma 6.3 (Edge parities)**.** *The horizontal-edge multiplicity has a constant parity $h_y \in \mathbb{F}_2$ along each row y. Every vertical-edge multiplicity is even; write $V_e = 2G_e$ with $G_e \in \mathbb{Z}_{\geq 0}$.*

Proof. Reducing the horizontal half of (11) modulo 2 gives $H_{v,-} \equiv H_{v,+}$, so consecutive horizontal edges in a row share parity; call it h_y. The vertical half gives $V_{v,-} \equiv V_{v,+}$, so vertical parity is constant up each column. The projected walk has finite vertical support, so a vertical edge far enough below or above it has multiplicity 0; constancy then forces every vertical multiplicity even. ☐

Lemma 6.4 (Local parity)**.** *At every vertex $v = (x, y)$ of C_L, $k_v \equiv h_y \pmod 2$.*

Proof. Suppose v is vertical (the horizontal case is the 90° rotation). By Proposition 2.2 the ant arrives at v along a vertical edge—from the south moving north, or from the north moving south—and departs along a horizontal edge. Record each visit by (arrival direction, turn). Using $u_{q+\varepsilon}$ from (1): an arrival from the south (R) departs east, from the south (L) departs west, from the north (R) departs west, from the north (L) departs east. Let $a_\mathrm{R}, a_\mathrm{L}$ count arrivals from the south with the two turns and $b_\mathrm{R}, b_\mathrm{L}$ those from the north. Because the east edge is only ever traversed *out* of v (its other endpoint is horizontal, hence cannot depart along it), its multiplicity equals the number of east departures:

$$H_{v,+} = a_\mathrm{R} + b_\mathrm{L}.$$

The north edge is only ever traversed *into* v, with multiplicity $V_{v,+} = b_\mathrm{R} + b_\mathrm{L}$. The total turn count is $k_v = a_\mathrm{R} + b_\mathrm{R}$. Hence

$$H_{v,+} - k_v = (a_\mathrm{R} + b_\mathrm{L}) - (a_\mathrm{R} + b_\mathrm{R}) = b_\mathrm{L} - b_\mathrm{R} \equiv b_\mathrm{L} + b_\mathrm{R} = V_{v,+} \pmod 2.$$

By Lemma 6.3, $H_{v,+}$ has parity h_y and $V_{v,+}$ is even; therefore $h_y \equiv k_v \pmod 2$. ☐

Proof of Theorem 6.1. Halving the vertical half of (11) and using $V_e = 2G_e$ gives the exact identity $k_{x,y} = G_{x,y-1/2} + G_{x,y+1/2}$, where $G_{x,y\pm 1/2}$ denote the halved multiplicities of the vertical edges below and above v. Fix a column x and sum Lemma 6.4 over all rows:

$$\sum_y h_y \;\equiv\; \sum_y k_{x,y} \;=\; \sum_y (G_{x,y-1/2} + G_{x,y+1/2}) \;\equiv\; 0 \pmod 2, \tag{12}$$

because each interior half-edge G appears exactly twice and the two exterior terms vanish by finite vertical support.

Now cut C_L along a vertical seam between two adjacent columns. The algebraic winding number w of the projected closed walk is the signed number of seam crossings; there is exactly one horizontal seam edge in each row y, of multiplicity $\equiv h_y \pmod 2$, and signs vanish modulo 2, so

$$w \equiv \sum_y h_y \equiv 0 \pmod 2$$

by (12). On the other hand the lift of the closed walk runs from p to $p + (L, 0)$, so it winds exactly once around C_L: $w = 1$. This contradiction proves the theorem for drift $(L, 0)$; quarter-turn rotations (including a half-turn) cover the remaining axis-aligned drifts. ☐

Corollary 6.5 (Growth is a positive multiple of four)**.** *Every finite-support periodic highway has net black growth*

$$g \in 4\mathbb{Z}_{>0}. \tag{13}$$

In particular every finite-support periodic highway builds a nonempty wake.

Proof. Zero drift confines the ant to the finitely many positions of one period, excluded by Theorem 2.4. Negative growth cannot recur forever since $b_t \geq 0$. A heading-resetting period has $g \equiv 0 \pmod 4$ (Lemma 2.5). If $g = 0$ then by Corollary 5.5 the state reduces, after deleting a finite stationary set never revisited, to a clean finite translator with nonzero drift, contradicting Theorem 6.1. Hence $g > 0$, and $g \in 4\mathbb{Z}_{>0}$. ☐

The standard highway has $g = 12$, consistent with (13). The telescope (12) is the finite algebraic core machine-checked in Lean (Appendix B).

7. The signed mod-four wake-residue theorem

We now give a second, independent proof that no zero-growth periodic highway exists, which additionally pins the arithmetic relation between growth and drift. Recall the strand bases $c_1, \dots, c_g$ of Proposition 5.4.

Theorem 7.1 (Signed mod-four wake-residue identity)**.** *Let a finite-support repeating trace have restored heading, nonzero drift $d = (a, b)$, growth g, and decomposition (10). Then:*

(i) *if $a \neq 0$, then for every residue $r \in \mathbb{Z}/|a|\mathbb{Z}$,*
$$\sum_{i:\, c_{i,x} \equiv r \pmod{|a|}} \chi(c_i) \;\equiv\; 2 \pmod 4; \tag{14}$$
(ii) *if $b \neq 0$, the analogous sum over each residue of $c_{i,y}$ modulo $|b|$ is $\equiv 2 \pmod 4$;*

(iii) *if $a = 0$, the signed sum $\sum_{i:\, c_{i,x} = x_0} \chi(c_i)$ over each exact column x_0 is $\equiv 0 \pmod 4$; symmetrically for exact rows when $b = 0$.*

The proof rests on a monodromy computation for the potentials of Proposition 5.3.

Lemma 7.2 (Diagonal increment and vertical periodicity)**.** *Fix $a \neq 0$ and any $\alpha : \mathbb{Z}/|a|\mathbb{Z} \to \mathbb{Z}/4\mathbb{Z}$, and set $w(x, y) = \chi(x, y)\, \alpha(x)$, a weight valued in $\mathbb{Z}/4\mathbb{Z}$. Then w satisfies the four-corner condition (8) in $A = \mathbb{Z}/4\mathbb{Z}$, so potentials F_q as in (7) exist, and*
$$F_N(x{+}1, y{+}1) - F_N(x, y) = -\chi(x, y)\big(\alpha(x) + \alpha(x{+}1)\big), \tag{15}$$
while F_N is 2-periodic in y.

Proof. Since $\chi(x{+}1, y{+}1) = \chi(x, y)$ and $\chi(x{+}1, y) = \chi(x, y{+}1) = -\chi(x, y)$, the four-corner sum (8) telescopes to $\chi(x, y)[\alpha(x{+}1) - \alpha(x{+}1) - \alpha(x) + \alpha(x)] = 0$; so F_q exist by Proposition 5.3. Compose the R step at (x, y) heading N (giving $F_E(x{+}1, y) - F_N(x, y) = -w(x, y)$) with the L step at $(x{+}1, y)$ heading E (giving $F_N(x{+}1, y{+}1) - F_E(x{+}1, y) = +w(x{+}1, y)$). Adding,
$$F_N(x{+}1, y{+}1) - F_N(x, y) = -w(x, y) + w(x{+}1, y) = -\chi(x, y)\alpha(x) - \chi(x, y)\alpha(x{+}1),$$
which is (15). The relation $F_S = F_N - 2w$ (from composing the two vertical steps) gives the same increment along $(1, -1)$; combining the $(1, 1)$ and $(1, -1)$ increments over a closed square shows $F_N(x, y{+}2) = F_N(x, y)$. □

Proof of Theorem 7.1. Fix $a \neq 0$, α, and w as in Lemma 7.2. Restored heading gives $g \equiv 0 \pmod 4$, so the period is even and $a + b$ is even; hence translation by d preserves χ, and $w(z + d) = w(z)$. Let $\varepsilon = \operatorname{sgn}(a)$, and set $\sigma = \operatorname{sgn}(b)$ if $b \neq 0$ and $\sigma = 1$ if $b = 0$. Move from z to $z + d$ by $|a|$ diagonal steps of vector (ε, σ), followed by the vertical displacement $(0,\, b - \sigma|a|)$, which is even because $a + b$ is even. The uniform diagonal increment of F_N (read off from (15) for $\varepsilon = +1$ and from its reverse for $\varepsilon = -1$) is
$$F_N(x + \varepsilon, y + \sigma) - F_N(x, y) = -\varepsilon\, \chi(x, y)\big(\alpha(x) + \alpha(x + \varepsilon)\big), \tag{16}$$
and χ is constant along the path because $\varepsilon + \sigma$ is even; the vertical part contributes nothing by the 2-periodicity of F_N in y. As x runs once through $\mathbb{Z}/|a|\mathbb{Z}$ each $\alpha(r)$ appears twice, so the increments telescope to $-2\varepsilon\, \chi(z) \sum_r \alpha(r)$, giving the monodromy
$$F_N(z + d) - F_N(z) \;\equiv\; 2 \sum_{r \in \mathbb{Z}/|a|\mathbb{Z}} \alpha(r) \pmod 4 \tag{17}$$
(since $-2\varepsilon\chi(z) \equiv 2 \pmod 4$ for either sign of a). The three other headings satisfy the same relation: $F_S = F_N - 2w$ and F_E, F_W differ from F_N by single applications of (7), all of which are themselves d-periodic because w is; so (17) holds with F_q in place of F_N for every q, and for either sign of a.

Charge conservation of $\sum_{z\in B} w(z) + F_q(p)$ over one period gives

$$\sum_z D(z)\,w(z) + F_q(p+d) - F_q(p) = 0, \tag{18}$$

where D is the signed toggle pattern of Proposition 5.4. Because w is d-periodic and $(T_dA)(z) = A(z-d)$, the widget terms cancel: $\sum_z (T_dA)(z)w(z) = \sum_z A(z)w(z)$. Substituting $D = S+(T_d-1)A$ and the monodromy (17) into (18),

$$\sum_{i=1}^{g} \chi(c_i)\,\alpha(c_{i,x}) + 2 \sum_{r\in\mathbb{Z}/|a|\mathbb{Z}} \alpha(r) \;\equiv\; 0 \pmod 4. \tag{19}$$

Taking $\alpha = \mathbf{1}_{\{r\}}$ turns (19) into $\sum_{i:\, c_{i,x}\equiv r} \chi(c_i) + 2 \equiv 0$, i.e. (14), using $-2 \equiv 2$. Rotating by 90° gives (ii). For $a = 0$, take a finitely supported exact-column $\alpha : \mathbb{Z} \to \mathbb{Z}/4\mathbb{Z}$; then b is even, w is d-periodic, the vertical 2-periodicity makes the monodromy vanish, and exact-column indicators give (iii). □

Corollary 7.3 (Strand-density bound)**.** *Every finite-support repeating trace with drift $d = (a, b)$ has*

$$g \;\geq\; 2\max(|a|, |b|). \tag{20}$$

In particular $g = 0$ is impossible—an independent proof of Theorem 6.1. If a residue fiber contains exactly two strand bases, they share checkerboard parity; for $d = (\pm2, \pm2)$ and $g = 4$ the 2×2 residue-count matrix is $\left[\begin{smallmatrix}2&0\\0&2\end{smallmatrix}\right]$ or $\left[\begin{smallmatrix}0&2\\2&0\end{smallmatrix}\right]$.

Proof. A sum of m signs ±1 congruent to $2 \pmod 4$ is nonzero, with m positive and even. Hence each of the $|a|$ residue classes modulo $|a|$ contains at least two strand bases, giving $g \geq 2|a|$; symmetrically $g \geq 2|b|$. If $g = 0$ some class is empty while $d \neq 0$, contradicting (14). A two-element fiber has χ-sum ±2, forcing equal signs. When $d = (\pm2, \pm2)$ and $g = 4$, every row and column of the 2×2 residue matrix therefore has sum 2. Equal checkerboard signs in either fixed-x row force its two bases into the same y-parity column, ruling out the matrix with all entries 1; the only possibilities are the two matrices displayed in the statement. □

Remark 7.4. The bound (20) is a genuine lower bound but is not claimed to be attained by a chronological highway: the standard highway has $g = 12$ against $2\max(2, 2) = 4$. Determining which residue profiles can be realised by a single ant tour is part of the chronology problem left open in Section 11.

Unlike the even-winding argument, Theorem 7.1 uses neither axis alignment nor the clean-packet reduction; it applies to every repeating trace and additionally fixes the drift–growth arithmetic. Independent computation supports the theorem in the ranges where it is testable (Appendix C): the underlying toggle-charge identity (18) was confirmed on all 172,092 heading-reset nonzero-drift words through length 18. The strand identity (14) and the bound (20) were confirmed on the standard highway, its powers, quarter-turn rotations, and cyclic phase shifts; the direct strand-level test is vacuous through length 18 because no realisable nonzero-drift trace occurs in that range.

8. Collision chains

We encode a finite cell set as edges of the bipartite row–column incidence graph: the cell (x, y) is the edge C_xR_y between a column vertex C_x and a row vertex R_y. To a directed state $s = (p, q)$ assign the token $\eta(s) = C_{p_x}$ if q is vertical and $\eta(s) = R_{p_y}$ if q is horizontal.

Theorem 8.1 (Two-endpoint collision-chain parity)**.** *Over $\mathbb{F}_2$, the boundary (set of odd-degree vertices) of the row–column incidence graph of the symmetric difference $B_s\triangle B_t$ equals $\eta(s)+\eta(t)$. Intermediate tokens telescope across any collision checkpoints; consequently at most one defect component of $B_s\triangle B_t$ is non-Eulerian, and a same-pose return ($\eta(s) = \eta(t)$) has only Eulerian defect components.*

Proof. One step toggles the single edge $C_{p_x}R_{p_y}$, changing the degree parity of exactly C_{p_x} and R_{p_y}. If the pre-step heading is vertical then $\eta(s) = C_{p_x}$ and the post-step heading is horizontal with unchanged row, so $\eta(s') = R_{p_y}$; for a horizontal pre-step heading the roles are reversed. Hence $\eta(s)+\eta(s') = C_{p_x}+R_{p_y}$, exactly the one-step boundary contribution. Summing over a run telescopes to $\eta(s) + \eta(t)$. By the handshaking lemma, every connected component has an even number of odd-degree vertices. If $\eta(s) \neq \eta(t)$, the entire graph has exactly those two odd vertices, so they lie in the same component; that component is the unique non-Eulerian one. If $\eta(s) = \eta(t)$, the boundary is empty and every component is Eulerian. □

This yields a four-cell lower bound on completed lifetimes. Take s immediately after an R visit to $z = (x, y)$ and t immediately before its next (L) visit. Then z is black at both times, so its edge C_xR_y is absent from $G := B_s\triangle B_t$. If arrivals at z are vertical, the R step leaves z horizontally along row y, so $\eta(s) = R_y$, while $\eta(t) = C_x$; if arrivals are horizontal the roles are reversed. Thus in either case the tokens are the two ends of the missing edge: $\eta(s) + \eta(t) = C_x + R_y$. By Theorem 8.1, $\partial G = C_x + R_y$, so C_x and R_y are joined in G by a path avoiding the edge C_xR_y; in the simple bipartite incidence graph such a path has at least three edges. Adjoining C_xR_y closes a cycle of length at least four, so the completed R-to-L lifetime involves at least four distinct cells.

Circuit-partition and interlacement formalisms may provide a useful compression of the chronological constraints, but we make no touch-graph rank claim here. The two-endpoint boundary identity above is the only result of this section used later.

9. Transverse rigidity and period-unbounded exclusions

Every exclusion obtained so far is bounded: Theorem 10.1 below rules out periods up to a fixed cutoff, and the residue identity constrains arithmetic rather than geometry. This section gives three statements that hold at *every* period. Extremal rigidity and the width-two theorem use the ordering content of Theorem 3.1—condition (ii), that the stabilised word begins with R, and condition (i), that it alternates. The computer-assisted width-four theorem then combines extremal rigidity with exact local ant dynamics. None uses Section 7.

Throughout this section the drift is *diagonal*: $d = (a, b)$ with $|a| = |b| = m \geq 1$. A quarter-turn rotation of the plane sends $d = (a, b)$ to $(-b, a)$, hence sends ab to $-ab$, and commutes with the ant's rule; since Definition 2.1 identifies highways up to quarter-turn rotation, we may and do *normalise so that* $ab < 0$, and then set

$$t(x, y) = x + y, \tag{21}$$

so that $d = me$ with $e = \pm(1, -1)$. The standard highway is already in this frame, with $d = (2, -2)$, $m = 2$ and $e = (1, -1)$.

With this normalisation $t(d) = 0$, so t is constant on translation classes and the range of t along the trace is the same in every period; and *every unit step changes t by exactly* ± 1, the four increments being $+1$ for east and north and -1 for west and south. Diagonality is what supplies the latter: for an axis drift the corresponding transverse functional is x or y, which two of the four moves leave unchanged, and the arguments below fail.

The normalisation is not cosmetic. In the rotated frame $ab > 0$ one must instead take $t = x - y$, whose increments are $+1$ for east and *south* and -1 for west and *north*; the roles of the two axes are exchanged, and every horizontal/vertical assertion below is exchanged with it. Only the conclusions that do not name an axis —that every extremal turn is R, that extremal cells are visited once, and that W is even—are stated the same way in both frames. Remark 9.2 records the unnormalised form.

Put

$$T = \max_i t(p_i), \qquad T' = \min_i t(p_i), \qquad W = T - T',$$

and call W the *transverse width*. Call a lattice cell *horizontal* or *vertical* according to whether the arrivals at it are east/west or north/south; by Proposition 2.2 this is determined by its checkerboard parity. Since all cells on a level line $t = c$ share the parity of c, each level line is entirely horizontal or entirely vertical.

Theorem 9.1 (Extremal-line rigidity)**.** *Let w^ω be a finite-support periodic highway with diagonal drift, normalised as above so that $t = x + y$. Then*

(i) *both extremal level lines $t = T$ and $t = T'$ are horizontal;*
(ii) *every turn on those two lines is* R*; the arrival heading is east on the line $t = T$ and west on the line $t = T'$;*
(iii) *every cell of those two lines that the periodic trace visits at all is visited* exactly once*; equivalently, every translation class meeting an extremal line is a singleton, so each such cell is entered once, is left black, and is never returned to;*
(iv) *W is even.*

Proof. Let I be a translation class meeting the line $t = T$, and consider its cells $z_n = b_I + nd$; since $t(d) = 0$ every z_n lies on that same line. Take n large enough that z_n has been reached by every phase of I, so that the chronological turn sequence at z_n is the stabilised word of I; in particular every visit to z_n has a predecessor. The ant can reach $t = T$ only from $t = T - 1$, so each arrival at z_n increases t; and it can never occupy $t > T$, so no departure from z_n increases t.

We use the explicit turn images: R is clockwise, so it sends north, east, south, west to east, south, west, north; L is anticlockwise and reverses this. With $t = x + y$ the increments are $+1$ for east and north and -1 for west and south.

Consider $t = T$. The arrival at z_n increases t, so it is east or north, and the departure must decrease t. If the line were vertical the arrival would be north, whose clockwise image is east—which increases t and is therefore forbidden—while its anticlockwise image is west, which is allowed; every turn at z_n would then be L, so the stabilised word of I would begin with L, contradicting Theorem 3.1(ii). Hence the line is horizontal and the arrival is east. Now the anticlockwise image of east is north, which increases t and is forbidden, while its clockwise image is south, which is allowed; so every turn at z_n is R.

At $t = T'$ the arrival decreases t, so it is west or south, and the departure must increase t. If the line were vertical the arrival would be south, whose clockwise image is west decreases t and is forbidden, while its anticlockwise image east is allowed; every turn would be L, the same contradiction. Hence that line is horizontal too, the arrival is west, its clockwise image north is allowed and its anticlockwise image south is forbidden, so again every turn is R. This proves (i) and (ii).

In both cases the stabilised word of I is a block of R's, and by Theorem 3.1(i) it alternates, so that block has length one: I is a singleton. A singleton class is visited once at each of its cells and is odd, so every cell of I that the trace reaches is entered exactly once and left black. As I was an arbitrary class meeting an extremal line, this proves (iii). Both extremal lines being horizontal, they carry the same checkerboard parity, so $T \equiv T' \pmod 2$, which is (iv). □

Theorem 9.1 says that a periodic highway is bounded transversally by two guard rails of permanent wake, touched once and never returned to. The mechanism deserves a remark, since it is what the incidence arguments of the preceding sections cannot reach: the vertical case is eliminated purely by the *first-turn* condition, and the single visit purely by *alternation*. Both are assertions about the order in which a cell is visited, and the countermodels recorded in Section 11 satisfy every residue equation while violating the conclusion, so neither is a consequence of the arithmetic developed there.

Remark 9.2 (the unnormalised statement)**.** Without the normalisation $ab < 0$ the conclusion must be read in the frame determined by the drift, and it is easy to state it wrongly. In the rotated frame

$ab > 0$ one takes $t = x - y$; the headings increasing t are then east and *south*, the argument of Theorem 9.1 eliminates the *horizontal* case, and both extremal lines are *vertical*, with arrival south on $t = T$ and north on $t = T'$. What is common to the two frames, and is what the sequel uses, is that the two extremal lines have the *same* type, that every extremal turn is R, that extremal cells are entered at most once, and that W is even. Concretely, the standard highway read from a north-facing start has $d = (2, -2)$ and horizontal extremal lines with arrivals east and west; the same highway read from an east-facing start—the same object, rotated—has $d = (-2, -2)$ and *vertical* extremal lines with arrivals south and north. In both cases every extremal turn is R.

The rigidity is strong enough to collapse a narrow strip to a one-dimensional system, which can then be refuted at all periods simultaneously.

Theorem 9.3 (No highway of transverse width two)**.** *No finite-support periodic highway with diagonal drift has transverse width $W = 2$. Combined with Theorem 9.1(iv), every such highway has even width $W \geq 4$.*

Proof. Suppose $W = 2$ and label the three occupied level lines $0, 1, 2$, meaning $t = T$, $T - 1$, $T - 2$. By Theorem 9.1 lines 0 and 2 are horizontal, each of their cells is entered at most once and turns R; line 1 is vertical. Because each step changes t by ± 1, the ant alternates between line 1 and lines $\{0, 2\}$, so consecutive visits to line 1 are separated by exactly one excursion to line 0 or to line 2.

Theorem 9.1 makes those excursions rigid. Orient the longitudinal axis so that $e = (1, -1)$; recall east and north increase t. From a cell z on line 1 the upward excursion enters line 0 heading east and leaves it heading south, so it carries the ant from z to $z + (1, 0) + (0, -1) = z + e$; the downward excursion enters line 2 heading west and leaves heading north, carrying z to $z - e$. Index the cells of line 1 by $\ell \in \mathbb{Z}$ along e. Every excursion changes ℓ by exactly ± 1.

Now bound how often the ant can cross between adjacent sites. An upward excursion from site ℓ passes through the line-0 cell determined by ℓ; call it U_ℓ, and let D_ℓ be the line-2 cell used by a downward excursion from ℓ. Distinct sites give distinct cells, and by Theorem 9.1(iii) each of U_ℓ and D_ℓ is entered at most once by the whole periodic trace. Since the gap between sites ℓ and $\ell + 1$ is crossed upward only by an upward excursion from ℓ, and downward only by a downward excursion from $\ell + 1$,

$$\text{each gap is crossed at most once upward and at most once downward.} \tag{22}$$

Choose a cyclic phase at a visit to line 1, and let s_0 be that initial site. Because the drift is $+m$ along ℓ, the ant's site tends to $+\infty$; so for every $\ell \geq s_0$ the walk starts weakly left of the gap $(\ell, \ell + 1)$ and is eventually permanently to its right, whence its net number of upward minus downward crossings of that gap is exactly 1. By (22) the only way to achieve a net of 1 with each count at most 1 is one upward crossing and *no* downward crossing. Hence for every $\ell > s_0$ the ant enters site ℓ exactly once, necessarily from $\ell - 1$, and leaves it towards $\ell + 1$, never returning.

It remains to read off the turn at such a site. Arriving from $\ell - 1$ means completing an upward excursion, which leaves line 0 heading south; so the ant arrives at site ℓ heading south. Departing towards $\ell + 1$ means beginning an upward excursion, which leaves site ℓ heading east. The anticlockwise image of south is east and the clockwise image is west, so the unique turn at site ℓ is L.

Finally, let I be the translation class of such a site with $\ell > s_0$. Its cells are the sites $\ell + nm$, all but finitely many of which exceed s_0 and are therefore entered exactly once. So I has a single phase, its stabilised word is the single letter L, and Theorem 3.1(ii) is violated. Since there are infinitely many such sites, the contradiction is unavoidable. ∎

The next case admits a different reduction. We give the finite computation as a separate lemma so that the computer-assisted boundary of the theorem is explicit. Write a finite sequence such as $(2, 0, 2)$ compactly as 202.

Lemma 9.4 (blank-column classification). *Consider an initially white column consisting of the five cells on transverse lines* $0, 1, 2, 3, 4$*. At each of its left and right boundaries record, in chronological order, the transverse line of the destination cell immediately after every horizontal crossing. Suppose each sequence starts and ends with a rightward crossing, its directions alternate, every rightward label belongs to* $\{2, 4\}$*, every leftward label belongs to* $\{0, 2\}$*, and on each boundary the rightward label* 4 *and leftward label* 0 *occur at most once. The compatible ordered pairs of left and right sequences are exactly the following twelve:*

left sequence	*possible right sequences*
202	4
4220222	2, 202, 20222, 2022224
2042222	2, 202, 20222, 2022224
202242222	422, 42202, 4220222.

The directed graph formed by these pairs is acyclic; its longest directed path has three edges.

Exact finite verification. A column state consists of its five colour bits, the transverse entry line and heading, and four Boolean flags recording whether the exceptional top-rightward or bottom-leftward event has already occurred on either boundary. Starting from the all-white mask and a first rightward entry on line 2 or 4, literal application of the ant rule reaches 22 such entry states, has no reachable state cycle and stops after depth at most eight. Exhausting the permitted return entry (2 or the one available extreme) gives precisely the twelve pairs displayed above.

This enumeration is implemented twice. The primary implementation uses a five-bit integer mask and numeric headings; an independently written verifier uses a set of black rows and compass-letter headings and imports no code from the first. Both return the same twelve pairs. As a convention check, the second implementation extracts two far-ahead crossing sequences from the documented standard-highway seed and exactly replays the intervening initially white column. Finally, assigning rank 3 to 202242222, rank 2 to 4220222 and 2042222, rank 1 to 202, and rank 0 to every other displayed vertex makes every edge strictly decrease rank. The rank calculation and the impossibility of a four-edge path are also checked in Lean 4. The accompanying certificate records the edge table, source hashes and the boundary between the enumerated and formalised parts. □

Theorem 9.5 (No highway of transverse width four). *No finite-support periodic highway with diagonal drift has transverse width* $W = 4$*. Consequently every such highway has even transverse width* $W \geq 6$*.*

Proof. Suppose $W = 4$. Rotate by $180°$ if necessary so that the normalised drift is $(m, -m)$ with $m > 0$, shift the transverse coordinate so that $t = x + y \in \{0, 1, 2, 3, 4\}$, and use x as longitudinal coordinate. A north or south step remains in one column $x = k$, whereas an east or west step crosses a boundary $x = k + \frac{1}{2}$.

At the beginning of the periodic tail, only finitely many cells have been visited, and the initial black support is finite. Hence all sufficiently large longitudinal columns are untouched and all five of their cells are white. For such a far-ahead boundary let Γ_k be its complete future crossing sequence, recording the destination transverse line. The longitudinal position in successive periods is nm plus one of finitely many phase offsets. Thus the ant begins to the left of the boundary, crosses it only finitely many times, and is eventually permanently to its right. Horizontal crossing directions therefore alternate right, left, ..., right, so Γ_k is nonempty and has odd length.

On a rightward crossing the destination label is 2 or 4; on a leftward crossing it is 0 or 2. Direction is therefore implicit in the parity of the event index. Moreover, two rightward labels 4 at the same cut would enter the same top-extremal cell twice, and two leftward labels 0 would enter the same bottom-extremal cell twice. Both are forbidden by Theorem 9.1(iii).

The untouched column $x = k$ has left sequence Γ_{k-1} and right sequence Γ_k, so Lemma 9.4 makes $\Gamma_{k-1} \to \Gamma_k$ an edge of its directed graph. Infinitely many consecutive untouched columns would give a directed path of arbitrary length, while the lemma permits no path of four edges. This contradiction excludes $W = 4$. Together with Theorem 9.1(iv) and Theorem 9.3, it gives $W \geq 6$. Appendix D records an independent arithmetic route that excludes width four for bounded drift without the blank-column table (Lemma 9.4), and pins the admissible odd-level masks of any width-four highway. □

Remark 9.6 (why the same proof stops at width six)**.** The standard highway has $W = 10$. The blank-column recursion is acyclic at width four, but an extra interior horizontal line creates a reachable two-state local oscillator at widths six and eight. In one state the black rows are $\{1, 2\}$ and the ant enters row 2 eastward; two turns change the mask to $\{1, 3\}$ and exit eastward. A westward return on row 2 restores the first state and exits westward. The two reused cells alternate their turns correctly, and neither extreme is reused. This is not a highway; it proves only that the acyclic one-column argument does not extend unchanged. A wider-width proof must couple columns or exploit the order of whole translation classes. In particular, a finite alphabet on an unbounded strip does not by itself make the highway question finite-state.

The hypotheses of this section are exactly Proposition 2.2 and Theorem 3.1. In particular Theorems 9.1, 9.3, and 9.5 are independent of Section 7 and of the computations of Section 10.

10. Small-period exclusion

All searches are exact: integer coordinates, exact finite sets, and no probabilistic equality tests. Three search variants are used. The *residue-free variant* uses no consequence of Theorem 7.1 and applies the exact criterion of Theorem 3.1 to every positive-growth nonzero-drift leaf; its zero-hit outcome is what certifies Theorem 10.1 below, and that certification is unconditional on the residue theorem. The *strand-pruned* and *residue-pruned* variants additionally prune with Corollary 7.3 and with the full identity of Theorem 7.1 respectively; they are faster, and their agreement with the residue-free variant is reported as a cross-check. All three accept the standard word. We stress that the three variants share one enumeration framework, one normalisation, and one representation, so their mutual agreement is a consistency check and not implementation independence; a separate verifier, written from the statement of Theorem 3.1 with a different class construction, is described in Appendix C. The methodology, per-shard records, and cross-audit are described there as well.

Theorem 10.1 (No periodic highway of period at most 48)**.** *No finite-support periodic highway of nonzero drift has period at most* 48*. Consequently any nonstandard finite-support periodic highway has even period at least* 50 *(the standard highway has period* 104*; periods* $50, 52, \ldots, 102$ *are not addressed here).*

Verification. Odd periods restore no heading (an odd number of quarter turns is never 0 mod 4), and Corollary 6.5 excludes zero growth at every period, so it suffices to search even periods for positive-growth traces. The evidence is:

(a) a complete first-R, cyclic-phase-normalised enumeration through period 32 (46,185,421 nodes at depths 6–32 over 32 prefix shards, plus a 59-node unprefixed baseline through depth 5; zero highways). Every positive-growth heading-resetting word contains R, so a cyclic shift places it in this enumeration; shifting only changes the chosen phase on the same periodic orbit and preserves finite-support realisability. No growth or endpoint pruning is applied, so the search is residue-free by construction; independently, a brute-force Python verifier (Appendix C) reproduces the exclusion over all $\sum_{L=1}^{18} 2^L = 524{,}286$ nonempty words of length at most 18;

(b) rank-audited positive-growth searches at periods $34, \ldots, 48$ with the residue-free variant, which uses no consequence of Theorem 7.1 and applies the exact criterion of Theorem 3.1 to *every* positive-growth nonzero-drift leaf. At the largest, period 48, it visited 41,710,394,384 nodes and reached 7,446,719,550 such leaves over all 2^{16} prefix ranks with no node cap, evaluated the exact criterion at every one of them, and returned zero certificates. Every prefix rank is covered exactly once, every shard reports completion, and the aggregate counters equal the sums of the per-rank counters (Table 1).

The zero-hit outcome of the residue-free variant is the certificate of nonexistence, and it does not assume Theorem 7.1. The strand-pruned and residue-pruned engines of Appendix C, which do assume it, return zero certificates on the strictly smaller trees they explore; that agreement is a cross-check only. It is, however, evidence *for* Corollary 7.3: the residue-free variant reaches the leaves P16 discards—at period 48, 6,994,756,680 of them—and finds no realisable word among them, so within this range the pruning discarded nothing. Table 1 summarises. Together these exclude every finite-support periodic highway of nonzero drift and period at most 48. □

Period	Nodes	Leaves	Criterion evaluations	Highways
≤ 32	46,185,480	—	—	0
34	40,418,033	6,989,418	6,989,418	0
36	108,995,287	18,928,333	18,928,333	0
38	293,763,689	51,386,767	51,386,767	0
40	791,801,484	139,000,900	139,000,900	0
42	2,133,302,151	376,781,423	376,781,423	0
44	5,748,138,964	1,018,293,035	1,018,293,035	0
46	15,483,269,352	2,757,152,898	2,757,152,898	0
48	41,710,394,384	7,446,719,550	7,446,719,550	0

TABLE 1. Exhaustive exclusion of finite-support periodic highways of nonzero drift. All rows from period 34 up are the output of the *residue-free* variant, which assumes no consequence of Theorem 7.1; the "criterion evaluations" column equals the leaf column, i.e. the exact criterion of Theorem 3.1 is applied to every positive-growth nonzero-drift leaf, with no leaf discarded unexamined. Each period-34–48 entry is complete over its parameter range with no node cap, and every prefix rank is covered exactly once. The period-≤ 32 row enumerates all first-R cyclic-phase representatives, plus the unprefixed depths below six. Every positive-growth heading-resetting word is represented, and no growth or endpoint pruning is applied, so this row is residue-free by construction. Its prefix shards were configured with a 10^8-node safety cap, but every shard completed below that cap (the largest used 2,623,020 nodes). Node counts are not comparable across different prefix lengths (Appendix C).

A finite-period exclusion does not by itself establish eventual convergence for arbitrary finite initial configurations: Theorem 10.1 bounds the minimal period of any nonstandard highway from below, but Conjecture 1.1 quantifies over all finite configurations and all periods, and no finite computation of the present kind exhausts that range.

11. The remaining gap

The results above constrain the arithmetic and topology of any finite-support periodic highway, but do not force convergence. Two distinct obstructions delimit the gap.

Problem 11.1 (Bounded retreat / bounded active core)**.** Show that for every finite-support initial configuration there are a finite window $W \subset \mathbb{Z}^2$, a time t_0, and translations v_t such that, for every $t \geq t_0$, the ant and every cell or Tait-boundary component capable of influencing the future orbit lie in $v_t + W$; all components outside that translated window are permanently inert and never reabsorbed. This is a bounded-active-core formulation, stronger than merely bounding the ant's distance behind a record frontier.

Under this formulation the active window has finitely many states modulo translation, determinism forces eventual periodicity, and Theorem 3.1 decides whether the resulting cycle is finite-support realisable and constructs a seed when it is. We do not claim these are the only possible routes to a solution, but no proof along them is known, and the following show why several easy attempts fail: unboundedness (Theorem 2.4) permits an orbit that keeps enlarging a two-dimensional region; no nonconstant affine combination of (E, V, k, β) is monotone (Proposition 4.4); and a local pumping argument is invalid because equal local cores may connect through arbitrarily distant explored territory, turning a later merge into a split (Theorem 4.1).

Problem 11.2 (Periodic classification)**.** Show that every positive-growth periodic highway admitted by finite support agrees in path and turn trace with the standard period-104 highway up to lattice translation, quarter-turn rotation, and cyclic phase shift, as in Definition 2.1.

At fixed growth $g = 4$ the residue skeletons of Theorem 7.1 leave unbounded free parameters, and explicit heading-reset countermodels satisfy every incidence and residue condition yet fail the cross-period criterion. A resolution therefore needs chronological input valid for unbounded period—genuinely more than the incidence data used here. Section 9 supplies the first such input in this paper: Theorems 9.1 and 9.3 are proved from the ordering content of Theorem 3.1, and the computer-assisted Theorem 9.5 combines that rigidity with exact column dynamics. All three hold at every period and are not consequences of the residue equations. They pin the boundary of the strip and exclude widths two and four, but Problem 11.2 still requires control of all wider strips, non-diagonal periodic traces, and uniqueness within any surviving geometry. Remark 9.6 also shows why "fixed width is finite-state" is not presently a theorem: at width six the local column dynamics already has recurrent interior memory. Widths six and eight are concrete next targets, not claims of a completed decision procedure. A *disproof* of Conjecture 1.1 would instead exhibit an infinite certificate: a nonstandard periodic word passing Theorem 3.1, or a finite self-enlarging Tait gadget whose state recurs with a parameter increased. No such certificate has been found.

12. Concluding remarks

We have developed a structural and arithmetic theory of the classical finite-support periodic highway: a decidable realisability criterion, an exact medial-graph conjugacy with surgery calculus, two independent proofs that no zero-growth clean translator exists (so every highway grows a wake of size $\in 4\mathbb{Z}_{>0}$), the strand-density bound $g \geq 2\|d\|_\infty$, a two-endpoint collision-chain parity theorem, a transverse rigidity theorem with the resulting all-period exclusions of diagonal transverse widths two and four, and an exhaustive exclusion of all periodic highways of period at most 48. Selected algebraic kernels used in the even-winding and residue arguments are machine-checked in Lean 4.

Three of these results are of a different character from the rest. The arithmetic constraints are incidence statements: they count strands, edges and charges, and Section 11 records countermodels satisfying all of them. Theorems 9.1 and 9.3 instead use the order in which a cell is visited, while Theorem 9.5 turns that rigidity into an acyclic exact crossing-sequence graph. They apply at every period rather than up to a search bound. The width-six interior oscillator shows, however, that the one-column acyclicity mechanism ends sharply at width four. Widths six and eight are concrete next targets, but their unbounded longitudinal memory has not been reduced to a finite decision problem.

One identified route to a full solution is a bounded-retreat or bounded-active-core lemma followed by a chronology theorem for unbounded period. The arithmetic constraints proved here do not supply either step, and other routes are not excluded.

Acknowledgements. The author thanks the maintainers of the Lean 4 community for the toolchain used in Appendix B. AI systems were used for exploratory computation, drafting, and adversarial review. Their output was treated as untrusted: promoted claims were rechecked against the displayed

arguments, independent implementations, and released certificates. The author is responsible for all remaining errors.

Appendix A. Worked example: the standard highway

The blank orbit first enters its repeating trace at step 9977. Normalised so that the ant faces north at the origin and the trace begins with R at a global minimum of the black-count walk, the period-104 word (with R =R, L =L) is

```
RRRRLLRLLRRRRLLRRRRLLRLRRR
RLRLLLLRRRRLRRLRRRRLLLLRLR
RRRLRRRRLLLLRLRRRRLRLLRRLL
LLRRLLRRRRLLRRLRLLRLLRLRLL
```

(read as one word of 104 symbols, four rows of 26). It has $\#\mathrm{R} - \#\mathrm{L} = 12$ and normalised drift $(2, -2)$, in agreement with Corollary 6.5 ($g = 12 \in 4\mathbb{Z}_{>0}$) and Corollary 7.3 ($12 \geq 2 \max(2, 2) = 4$). Applying the seed construction (3) of Theorem 3.1 to this exact word (in this exact phase and orientation) produces the 11-cell black seed

$$(-2, -2),\ (-2, -1),\ (-1, 0),\ (0, 1),\ (1, -2),\ (1, 1),$$
$$(2, -2),\ (2, 0),\ (3, -1),\ (3, 0),\ (4, 0),$$

which, with a north-facing ant at the origin, produces the exact 104-step word above immediately, restoring the north heading and translating by $(2, -2)$; the leading corridor is white and the gateway pattern recurs translated by the drift each period. (Both the word and this seed were verified by direct simulation. A different but equally standard normalisation—the phase obtained by cyclically shifting this word by 100 symbols, whose trace begins with L and has drift $(-2, 2)$—is generated by a 13-cell seed; the two must not be interchanged, as they are different phases of the same highway.)

Appendix B. Lean 4 formalization

A dependency-light Lean 4 project (pinned to Lean 4.32.0, standard library only) machine-checks the exact transition kernel and selected algebraic kernels used above. Every audited theorem depends only on the standard axioms `propext`, `Classical.choice`, and `Quot.sound`; there are no `sorry`s, no custom axioms, and no `native_decide`. The checked statements include the finite-support step kernel and its toggle lemmas, the heading/black-count reset of Lemma 2.5, the four-corner admissibility and mod-four fiber arithmetic of Theorem 7.1 (including "one residue identity per class $\Rightarrow g \geq 2n$"), the checker-cycle monodromy $2\sum\alpha$ of Lemma 7.2, translated-widget cancellation, and the collision-chain endpoint parity of Theorem 8.1. It also checks the finite rank certificate for Lemma 9.4: every one of the twelve listed edges lowers rank, so no four-edge chain exists. Completeness of the local Python enumeration is explicitly outside the present Lean boundary.

The finite parity telescope (12)—the combinatorial heart of the even-winding theorem—is checked in the following self-contained form (`evenWindingParityCore`); `xorAll` is the $\mathbb{F}_2$ sum of a list, `incidences` takes successive XORs, and the two `false` endpoints are the vanishing exterior half-flows:

```
def bxor : Bool -> Bool -> Bool
  | false, b => b
  | true, false => true
  | true, true => false

def xorAll : List Bool -> Bool
  | [] => false
  | b :: bs => bxor b (xorAll bs)

def incidences : List Bool -> List Bool
  | [] => []
```

```
  | [_] => []
  | a :: b :: rest => bxor a b :: incidences (b :: rest)

-- XOR of all adjacent incidences of a list bracketed by two ‘false‘
-- endpoints is ‘false‘: the seam-crossing count is even.
theorem evenWindingParityCore (interior : List Bool) :
    xorAll (incidences (false :: interior ++ [false])) = false := by
  rw [xorAll_incidences, lastFrom_append_singleton]; rfl
```

The formalization isolates, as an explicit interface, the endpoint normal form and lifted potential geometry of a genuine periodic ant trace. In particular, `StabilizedP3Word.toOrbitEndpointData` constructs all algebraic endpoint data (first positive strand, balanced sum, and binary scan prefix) from an explicitly supplied stabilised P3 word. What remains outside Lean is the extraction of that word and of the local geometric charge certificates from an arbitrary periodic ant orbit. Thus Theorem 7.1 is certified conditional on a structured trace certificate rather than end to end; neither it nor Conjecture 1.1 is asserted as an unconditional Lean theorem. The project and build instructions accompany the paper [9].

Appendix C. Computational methodology and reproducibility

The periodic search space is the tree of legal prefixes. The following cycle-lemma normalisation justifies the prefix-growth pruning used by the search. Write the word weights as $u_0, \dots, u_{N-1} \in \{+1, -1\}$, with R assigned $+1$ and L assigned -1. Define $H_0 = 0$ and $H_j = \sum_{i=0}^{j-1} u_i$ for $1 \le j \le N$, and suppose $H_N = g > 0$. Choose $j \in \{0, \dots, N-1\}$ at which H_j is minimal. For a length-r prefix of the rotation beginning at j, where $1 \le r \le N$, the weight is $H_{j+r} - H_j \ge 0$ if $j + r \le N$, and is $g + H_{j+r-N} - H_j \ge g > 0$ if $j + r > N$. Thus every prefix of this rotation has nonnegative weight, and its first symbol is necessarily R. Restarting a realisable orbit at that phase preserves finite support, and a quarter-turn rotation restores the fixed initial heading. Hence rejecting a prefix with negative $\mathrm{R} - \mathrm{L}$ balance loses no positive-growth periodic trace. By Corollary 6.5 only positive-growth heading-reset words need be considered. To distribute a period-N search, the half-open interval of length-k prefix ranks $[0, 2^{k-1})$ is partitioned into disjoint worker intervals; each worker performs a depth-first enumeration of its subtree, applying the exact criterion of Theorem 3.1 at every leaf. A complete exclusion requires that every rank be covered exactly once, that every shard report completion without reaching any configured node cap, and that the aggregate node, leaf, and prune counters equal the sums of the per-rank counters.

What certifies the exclusion, and what merely cross-checks it. The residue-free variant is run for every even period 34–48 certification; the two conditional pruned variants provide selected cross-checks at periods 44, 46, and 48. It matters exactly which pruning rule each one assumes, so we spell this out.

The *residue-free variant* (`PositiveGrowthSearchIndep.java`) uses only rules that are free of Theorem 7.1: the cyclic normalisation above, same-cell alternation (Proposition 2.3), the growth congruence of Corollary 6.5 (a consequence of the even-winding theorem of Section 6, not of the residue theorem), and L_1 endpoint reachability to a nonzero drift. It applies the exact criterion of Theorem 3.1 to *every* positive-growth nonzero-drift leaf. Its zero-hit outcome is the certificate of Theorem 10.1, and that certificate is unconditional on Theorem 7.1. It is, however, conditional on Corollary 6.5; we make no claim that the search is unconditional on the even-winding theorem, which is what reduces the problem to positive growth in the first place. Every released certifying shard has `deficit_depths=[]`; the optional odd-ending deficit routine is inactive, and all of its recorded counters are zero.

The *strand-pruned engine* (`PositiveGrowthSearch.java`) additionally contracts endpoint reachability to the P16 square of Corollary 7.3 and rejects leaves violating $g \ge 2\max(|a|, |b|)$. The *residue-pruned engine* (`PositiveGrowthResidueSearch.java`) adds the full identity of Theorem 7.1. Both are therefore *conditional* on the residue theorem, and neither can certify Theorem 10.1 on its own. Because Corollary 7.3 prunes at the node level, the strand-pruned and residue-pruned engines do not merely discard leaves: they never reach them, which is why their reported leaf counts are smaller than the residue-free variant's rather than equal to them, and why their agreement on node counts is a statement about the P16-contracted tree only.

Period 46 illustrates the scale of the conditional pruning. The residue-free variant reaches 2,757,152,898 leaves and evaluates the criterion at each; the strand-pruned engine reaches 98,568,824, i.e. 3.6% of them, and the residue-pruned engine evaluates the criterion on 6,132,192, i.e. 0.22%. The 2,658,584,074 leaves the strand-pruned engine never reaches are exactly the ones a residue-free certificate must inspect. Every period-exclusion count reported here is output of the residue-free variant.

Finally, the three variants share one enumeration framework, one normalisation, one representation, and one exact-criterion routine, so a shared bug would affect all three identically; their agreement is a consistency check, not implementation independence. To test the criterion itself we use a separate verifier (`verify_criterion_indep.py`) written from the statement of Theorem 3.1 rather than from the Java source: it builds translation classes by explicit pairwise membership tests and union-find instead of the Java's arithmetic reduction and global sort, and it evaluates the *literal* criterion, checking every $S_{I,n}$ for $\min_I a_i \le n \le \max_I a_i$ rather than the single stabilised word that Remark 3.2 licenses—so it also exercises Remark 3.2, which the Java engines assume. On the 172,092 heading-reset nonzero-drift words of length at most 18 enumerated below, the literal criterion and the one-sort shortcut agree on every word, with zero disagreements; both accept the standard period-104 word, and the Java binary agrees with the Python verifier on every word sampled directly from it. For period at most 18 this is a genuine end-to-end cross-language replication: the Python verifier brute-forces all $\sum_{L=1}^{18} 2^L = 524{,}286$ nonempty words with no pruning whatsoever and independently reproduces the exclusion. Beyond that length brute force is infeasible and the check is confined to the criterion, which is the component a shared bug would most plausibly corrupt.

The analytic transverse results, Theorems 9.1 and 9.3, are proved rather than computed, while Theorem 9.5 uses the explicitly delimited finite classification in Lemma 9.4. The analytic inputs are also checked: `verify_extremal.py` confirms Theorem 9.1 on the standard word (transverse range $[-5, 5]$, width 10; each extremal line carrying two once-visited cells, all arrivals horizontal, all turns R), on its square, and on its R/L reflection, where the extremal turns are all L and the word is correctly rejected—the mirror image of the vertical case excluded in the proof. More usefully it checks the contrapositive exhaustively: of the 60,854 diagonal-drift heading-resetting words of length at most 18, exactly 18,946 have a repeated cell on an extremal line, and *none* of those is realisable, as Theorem 9.1 requires. `verify_width.py` checks the geometric inputs of Theorem 9.3—the two forced excursions and the arrival/departure headings from which the final turn is read—and then runs an adversarial search in which the colour of every not-yet-visited site is chosen by an opponent rather than assumed white. With at most k black seeds no opponent survives more than $k + 4$ steps before some gap must be crossed twice in one direction, which is the budget Theorem 9.1(iii) forbids. That the opponent is given this freedom matters: an earlier version of both the proof and the script assumed that a site's first visit turns R, which is *false*—condition (ii) of Theorem 3.1 constrains only stabilised words, and (3) blackens cells below their class's maximal level precisely so that some first turns are L. The argument given above uses no such assumption.

For Lemma 9.4, `width4_crossing_graph.py` supplies both the unbounded 22-state recursion and a pairwise cross-check over all 910 admissible signatures through length 25. The independently written `verify_width4_crossing_graph_indep.py` shares neither its state representation nor its turn implementation, reproduces the same twelve edges, and performs the standard-highway column replay described in the lemma. The machine- readable certificate `width4_crossing_graph_certificate.json` contains the edge list, statistics, source hashes and formalisation boundary. A separate set of combined exact searches, using the proved bound $P \leq 7m$, covers drift magnitudes $1 \leq m \leq 9$ with no hit. The fixed-drift macro engine supplies all periods for $m \leq 6$ and the upper ranges for $m = 7, 8, 9$; the general and unrestricted width-four searches supply their lower ranges. The $m = 10$ record is explicitly partial through $P = 66$. These searches are independent bounded cross-checks, not inputs to Theorem 9.5.

One further assumption of the engines is worth isolating, since it is what makes the search exhaustive. The engines never test the heading directly; they test $g \equiv 0 \pmod 4$. These agree by Lemma 2.5, but the lemma as usually stated gives only the forward implication, and it is the *converse* the search relies on: a word rejected for $g \not\equiv 0$ must genuinely fail to reset the heading, or a candidate would be silently skipped. For an arbitrary turn word, the heading increment modulo 4 is directly $\#\mathrm{R} - \#\mathrm{L} = g$, so the two conditions are equivalent even before a colour realisation is considered. We also confirmed the equivalence by enumeration over all 524,286 words of length at most 18: the heading-resetting words number 174,762, and so do the words with $g \equiv 0 \pmod 4$, with no exceptions in either direction.

The residue theorem was audited separately, with the following counts over all $\sum_{L=1}^{18} 2^L = 524{,}286$ nonempty turn words of length at most 18: 174,762 restore the heading; 172,092 of those have nonzero drift; and *none* of those 172,092 passes the realisability criterion of Theorem 3.1, consistent with the period-≤ 48 exclusion. The residue theorem (14)–(20) is a statement about *realisable* traces, so its direct test is vacuous in this length range; what the enumeration does verify on all 172,092 words is the underlying additive toggle-charge identity (18) (which holds for every heading-reset word, realisable or not). The strand-level identity (14) and the bound (20) are instead confirmed on the standard word, its powers, quarter-turn rotations, and cyclic phase shifts. (In particular a word such as RL, with drift $(1, 1)$ and growth 0, restores the heading and lies among the 172,092 but is *not* realisable and is excluded from the strand-level test, as it must be.)

Environment and archived records. The searches reported here were run on a single workstation, not on a cluster: one Intel Core i5-1334U (10 physical cores, 12 logical, 1.3 GHz base) with 16 GB of RAM under Windows 11 (build 26200), on Eclipse Temurin OpenJDK 25.0.1+8 (64-bit Server VM), with one JVM per shard and at most 12 concurrent shards; the Python verifier ran on CPython 3.13.14. The period-48 run of the residue-free variant is the largest single job: its shard records total 51,856.023 worker-seconds (14.40 aggregate worker-hours). With at most 12 one-JVM workers, the observed wall time was on the order of one to two hours; the exact per-shard times are archived.

The engines' source, an SHA-256 manifest of the released snapshot, the aggregation and cross-audit scripts, and the exact commands to regenerate every shard are publicly available in the accompanying repository [9]. For periods 34–48, the *complete* per-rank shard records—each rank's node, leaf, prune, and exact-evaluation counters, its wall time, and its hit list—are archived with the paper rather than left to be regenerated. For period at most 32, the archive instead contains the unprefixed depth-5 baseline and all 32 complete six-symbol prefix-shard logs; these certify prefix coverage but do not contain per-rank subrecords. Compressed, the full set is a few megabytes. The archive also carries the aggregation output, the period-34–48 rank-coverage verification (every rank covered exactly once), the check that every configured cap was left unreached, the raw records needed for per-rank comparison of the residue-free and pruned variants where both were run, the exact commands, and the environment strings above. The period-34–48 certifying shards have `node_cap=null`; the older period-≤ 32 shards record a 10^8 safety cap but all completed far below it. The exact search-record archive used for this revision has SHA-256 `7052f3070829bf804ae317a40379f6bb045b00775aff826b4e347880747bcc1a`.

Appendix D. Width-four structure and an independent partial exclusion

Theorem 9.5 excludes every diagonal width-four highway using the completeness of the twelve-edge blank-column table (Lemma 9.4), which is the one computer-assisted premise of Section 9. This appendix records a logically independent arithmetic route that reaches the same conclusion for bounded drift without that premise, and along the way pins the exact odd-level structure of any hypothetical width-four highway. It rests only on the residue identity, extremal-line rigidity, and the exact criterion.

Proposition D.1 (width-four masks and period bound)**.** *Let a finite-support periodic highway have transverse width four and normalised drift $d = (m, -m)$, $m > 0$. For each longitudinal residue, the set of levels carrying odd translation classes is one of*

$$\{0, 2\},\ \{1, 3\},\ \{0, 4\},\ \{2, 4\},\ \{0, 1, 2, 4\},\ \{0, 2, 3, 4\},$$

and for each diagonal-coordinate residue the corresponding set belongs to the same six-element list. If h of the m longitudinal residues use a four-element mask, then $g = 2m + 2h$ and $h \equiv m \pmod 2$. With P the macro period and $N = 2P$ the ant period,

$$P \le 7m, \qquad N \le 14m, \qquad P \equiv m \pmod 2.$$

Proof. Let $o_{j,r}$ indicate that the class on level j and longitudinal residue r is odd. The signed residue identity of Theorem 7.1 reads $o_{0,r} - o_{1,r} + o_{2,r} - o_{3,r} + o_{4,r} \equiv 2 \pmod 4$. Exhausting the five binary values gives exactly the six displayed masks; the other coordinate residue identity gives the diagonal version. Summing mask sizes gives $g = 2m + 2h$, and $g \in 4\mathbb{Z}$ (Corollary 6.5) fixes the parity of h.

For the period bound, start the cycle immediately after a bottom-extremal visit and observe the ant only on the odd levels. The forced bottom turn puts it in state LN, and one translated period returns to LN; width four ensures a top-line transition also occurs. Write the four states LN, LS, UN, US for lower/upper level and north/south arrival. The twelve possible two-turn macro transitions are

$$\begin{array}{r|lll} LN & (LS, +1, \mathrm{RR}) & (UN, +1, \mathrm{RL}) & (LN, -1, \mathrm{LR}) \\ LS & (LN, -1, \mathrm{RR}) & (LS, +1, \mathrm{LR}) & (UN, +1, \mathrm{LL}) \\ UN & (US, +1, \mathrm{RR}) & (UN, -1, \mathrm{LR}) & (LS, -1, \mathrm{LL}) \\ US & (UN, -1, \mathrm{RR}) & (LS, -1, \mathrm{RL}) & (US, +1, \mathrm{LR}). \end{array}$$

Let x count $LN \to LN$, y count $US \to US$, c count $LS \to LN$, and f count $UN \to US$. If a, b count $LN \to LS$, $LN \to UN$ and i, j count $US \to UN$, $US \to LS$, flow conservation gives $c = a + b$ and $f = i + j$. Extremal singleton rigidity (Theorem 9.1) gives $x + c \le m$ and $f + y \le m$. Let L_1, L_3 be the lower and upper odd-level L totals. The criterion of Theorem 3.1, that each stabilised class alternates and begins with R, makes the R-minus-L excess on each odd level nonnegative; the lower right-turn total is $a + b + c = 2c$ and the upper one is $f + i + j = 2f$, so $L_1 \le 2c$ and $L_3 \le 2f$. Flow conservation and the longitudinal displacement read from the table give

$$P = 2c + 2f + L_1 + L_3, \qquad m = L_1 - L_3 - 2x + 2y.$$

Hence $L_1 + L_3 \le 3m - 6x + 2y$, and substitution yields $P \le 7m - 8x \le 7m$. Every macro move changes the longitudinal coordinate by ± 1, so $P \equiv m \pmod 2$. □

Corollary D.2 (independent width-four exclusion for bounded drift)**.** *Fix $m \ge 1$. By Proposition D.1 a width-four highway of drift $(m, -m)$ has ant period $N \le 14m$, so only finitely many turn words can realise one, and Theorem 3.1 decides each. Performing this finite search excludes width-four highways of drift $(m, -m)$ for every $1 \le m \le 9$ (`results/width4_drift_exclusion_summary.json`). This conclusion uses only Proposition D.1 and the exact criterion, not the blank-column table of Lemma 9.4, so it corroborates Theorem 9.5 on that range by a logically separate route. The crossing-graph argument of Section 9 supersedes it by removing the drift bound.*

References

[1] L. A. Bunimovich and S. E. Troubetzkoy. Recurrence properties of Lorentz lattice gas cellular automata. *J. Stat. Phys.* **67** (1992), no. 1–2, 289–302.

[2] K. R. Etse. How long can the escaping ant be confined? arXiv:2606.26677 (2026).

[3] A. Gajardo, V. H. Lutfalla, and M. Rao. Ants on the highway. arXiv:2409.10124 (2024); *Natural Computing* **24** (2025), 497–509.

[4] A. Gajardo, A. Moreira, and E. Goles. Complexity of Langton's ant. *Discrete Appl. Math.* **117** (2002), no. 1–3, 41–50. doi:10.1016/S0166-218X(00)00334-6; arXiv:nlin/0306022.

[5] A. Gajardo. A symbolic projection of Langton's ant. *Discrete Math. Theor. Comput. Sci. Proc.* AB (DMCS'03) (2003), 57–68. doi:10.46298/dmtcs.2312.

[6] D. Gale, J. Propp, S. Sutherland, and S. Troubetzkoy. Further travels with my ant. *Math. Intelligencer* **17** (1995), no. 3, 48–56.

[7] C. G. Langton. Studying artificial life with cellular automata. *Physica D* **22** (1986), no. 1–3, 120–149. doi:10.1016/0167-2789(86)90237-X.

[8] V. H. Lutfalla. Sideways on the highways. arXiv:2505.05426 (2025).

[9] A. Jillhewar. Langton's ant: exact searches, a Lean formalization, and supplementary data. GitHub, `https://github.com/Atharva12456/langtons-ant-highway`; Zenodo concept DOI: 10.5281/zenodo.21381637. The search archive used here is pinned in Appendix C; file-level hashes are in `results/artifact_hashes.json`.

Independent researcher

Email address: `Atharvajil124@gmail.com`