

Every One-Black-Cell Langton Ant Reaches the Period-104 Highway: A Lean-Checked Proof

Hao Ke
Independent researcher
i@kehao.me

13 October 2026

Abstract

We prove that every classical Langton-ant state with exactly one black cell, at any displacement and ant pose, eventually follows the complete standard period-104 highway trace up to translation and quarter-turn rotation. The proof couples the perturbed orbit to the finite blank-orbit transient and reduces the unbounded part of the scattering problem to 22 affine frontier channels. Moving a defect farther down a channel adds translated XOR wakes, allowing finite scattering anchors to extend by induction to every depth. Twenty-one channels admit geometric separation from the blank orbit’s historical wake. The remaining channel creates a reverse highway that returns to a fixed historical cell; an affine collision-time law and a second archive induction then establish entry into the forward highway. The infinite reductions, geometric guards and certificate-soundness theorems are formalized in Lean 4, with three finite reports evaluated using `native_decide`. The accompanying artifact supplies the proof sources, witness tables and reproduction commands.

Contents

1	Introduction	3
2	Dynamics, symmetries, and the main theorem	4
2.1	States and updates	4
2.2	Equivariance and canonical states	5
2.3	The certified blank boundary	5
3	The finite prefix and first-contact coupling	6
3.1	Prefix contacts	6
3.2	Untouched defects	6
4	Entry geometry and the 22 frontier channels	6
5	Pristine single-defect scattering	8
5.1	XOR archives	8
5.2	All-depth pristine classification	9
6	Restoring the physical history: 21 ordinary channels	11
7	Exceptional backscattering and the affine hit law	12
7.1	The exact reverse block	12
7.2	First historical collision	13
7.3	The post-hit tail	13
8	Universal assembly	14
9	Lean formalization and finite computation	15
9.1	Proof cone	15
9.2	Finite leaves	16
9.3	Reflection and trust boundary	16
9.4	Reproduction	16
10	Related work	17
10.1	Classical ant dynamics	17
10.2	Complexity and generalized highways	17
10.3	A geometric method for unbounded perturbations	18
11	Conclusion	18
A	Certificate inventory	19
B	References	19
C	Version history and acknowledgements	20

1 Introduction

Langton’s ant is among the simplest deterministic systems in which a short local rule produces a long irregular transient followed by rigid macroscopic motion. A cell of the square lattice is white or black. On each update the ant reads its current cell, turns right on white and left on black, flips that cell, and advances one lattice edge. From the all-white plane, the orbit is observed to settle into a diagonal highway: every 104 updates the same local behaviour recurs, translated by two cells in each coordinate (Langton 1986). The central question is which finite perturbations provably enter this permanent regime.

The general finite-support Highway Conjecture remains unresolved. Classical unboundedness results rule out confinement to a finite region (Bunimovich and Troubetzkoy 1992), yet unbounded motion need not have the period, drift, or phase coherence of the standard highway. Complexity results sharpen the obstruction: finite-support prediction already contains P-hard instances, while universality and undecidability constructions use an infinite but finitely described circuit background (Gajardo, Moreira, and Goles 2002). These input contracts define complementary questions about spatial escape, computation and eventual highway convergence.

This paper isolates the smallest translation-invariant nonblank family: states with exactly one black cell, at an arbitrary displacement from an arbitrarily posed ant. The family is infinite in two independent senses. There is no bound on the defect coordinate, and a late defect can interact with an arbitrarily long highway history. A simulation cutoff therefore cannot close the theorem. Recent experiments reported one-black runs in a bounded radius-three domain and substantial transients (Gajardo, Lutfalla, and Rao 2024); our task is to remove the spatial cutoff by proving an all-depth reduction.

Theorem 1.1 (Universal one-black theorem). *For every Langton-ant state s on $\mathbb{Z}^2$, if the black set of s is a singleton, then there exists a time N such that its entire observation sequence from $\text{run}(N, s)$ agrees with a translation and quarter-turn rotation of the blank orbit from update 9,977 onward. In particular, this is the fixed standard period-104 read trace with a quarter-turn rotation of drift $(-2, -2)$.*

The exact observation-level meaning of “permanent period-104” is given in Definition 2.1. The proof is a scattering argument around a certified blank-orbit boundary. After translation and rotation normalization, let q be the single black cell and E the blank highway entry state. If the blank transient reads q , a finite witness settles the case. Otherwise a first-difference coupling gives exactly $\text{blacken}(q, E)$. The future P104 read set decomposes all remaining q into cells that are inert, finitely many active-support cells, and 22 affine channels $q = h + dv$. This turns an unbounded plane into 22 one-parameter scattering problems.

The decisive structure is a $21 + 1$ channel classification. In 21 ordinary channels, translated XOR wakes become inert archives and all sufficiently deep collisions reduce inductively to a single pristine anchor; a second geometric coupling transfers this result through the finite historical wake of E . The exceptional channel creates an exact reverse highway with drift $(2, 2)$, travels back through the accumulated history, and hits the same historical cell for every depth. Adding one unit of depth inserts one forward and one reverse P104 block, so the hit time changes affinely by 208 updates. A final archive induction transfers one checked post-hit transient to every depth.

The proof is formalized in Lean 4 (Moura and Ullrich 2021). Analytic proof terms handle symmetry, coupling, affine geometry, recurrence, archive induction, and universal assembly. Closed Boolean reports handle three finite leaves and are reflected into typed certificates. They are evaluated using




Figure 1: Proof architecture. Finite replay closes early contact; all unbounded late-contact cases are discharged by channel induction.

`native_decide`; according to the Lean reference, that mechanism uses a dedicated axiom and relies on the native compiler and runtime (Lean developers 2026). Section 9 records the proof components and their dependencies.

The contributions are:

1. a proof of Theorem 1.1 for every exactly-one-black state, without coordinate or time bounds;
2. a geometric reduction of all late defects to 22 unique affine frontier channels and an all-depth $21 + 1$ scattering classification;
3. an explicit analysis of the exceptional reverse-highway collision, including its affine hit law; and
4. a reproducible Lean artifact whose finite/native trust boundary is explicit and whose proof checking performs no generation-time search.

Sections 2–4 define the dynamics and reduce the plane to frontier channels. Sections 5–7 prove pristine scattering, ordinary historical transfer, and the exceptional family. Section 8 assembles the theorem. Section 9 describes the formal artifact and checked computation. Sections 10 and 11 position and summarize the result.

2 Dynamics, symmetries, and the main theorem

2.1 States and updates

Write a lattice point as $p = (x, y) \in \mathbb{Z}^2$. A state is

$$s = (b, p, d), \quad b : \mathbb{Z}^2 \rightarrow \{0, 1\}, \quad d \in \{N, E, S, W\}.$$

Here $b(x) = 1$ denotes black. Let R and L be clockwise and counterclockwise quarter turns, and let e_d be the unit vector in heading d . One update reads $c = b(p)$ and sets

$$d' = \begin{cases} R(d), & c = 0, \\ L(d), & c = 1, \end{cases} \quad b' = b \oplus \mathbf{1}_{\{p\}}, \quad p' = p + e_{d'}.$$

Thus turning and flipping occur before movement. This convention fixes every reported time and coordinate. Write $\text{step}(s)$ for this map and $\text{run}(n, s)$ for its n -fold iterate, with time zero equal to the supplied state.

Define the local observation $\text{obs}(s) = (p(s), d(s), b_s(p(s)))$. For $u \in \mathbb{Z}^2$, translation of an observation changes only its position component.

Definition 2.1 (Permanent P104). Let $v_0 = (-2, -2)$ and let D be its four quarter-turn rotations. A state s has *permanent P104* if some $v \in D$ satisfies, for every $k \geq 0$ and $0 \leq r < 104$,

$$\text{obs}(\text{run}(104k + r, s)) = kv + \text{obs}(\text{run}(r, s)).$$

In addition its first 104 read colours must equal the fixed word

$$w_r = b_{\text{run}(9977+r, \text{white})}(p(\text{run}(9977 + r, \text{white}))), \quad 0 \leq r < 104.$$

A state *reaches P104* if $\text{run}(N, s)$ has permanent P104 for some $N \geq 0$.

This is exactly the Lean predicate `PermanentP104`; `ReachesP104` existentially quantifies N . It fixes every later read, turn, and translated ant position while allowing inert historical cells outside the future read trace.

Every accepting terminal certificate couples to the fixed blank boundary, so the fixed-word clause propagates through the scattering proof. The module `StandardTrace` further proves full future observation equality with a placement of the actual blank state at update 9,977: periodicity extends the fixed first word to all times, equality of read colours and aligned initial poses determines the same positions and headings inductively.

2.2 Equivariance and canonical states

Translations and quarter-turn rotations act simultaneously on the grid, position, and heading. Direct calculation gives

$$\text{step}(T_u s) = T_u \text{step}(s), \quad \text{step}(\rho s) = \rho \text{step}(s),$$

and hence the same identities for every iterate. Both actions preserve `ReachesP104`; a rotation merely rotates the standard drift.

Lemma 2.2 (Pose normalization). *Every state with exactly one black cell is a translation and a quarter-turn rotation of a state σ_q whose ant is at the origin, points north, and whose unique black cell is $q \in \mathbb{Z}^2$.*

Consequently Theorem 1.1 reduces to proving `ReachesP104`(σ_q) for every $q \in \mathbb{Z}^2$.

2.3 The certified blank boundary

Let ω be the all-white canonical state. Exact replay supplies the following data.

Proposition 2.3 (Blank highway boundary). *At time $T = 9,977$, the blank orbit is in a state E whose ant is at $(-15, 10)$ heading west. The next 104 updates read a support S of 40 distinct cells, restore the heading, and translate the local observation by $v = (-2, -2)$. Thirteen cells form the active black pattern. The state E has 715 black cells in total: 13 active cells and a 702-cell historical wake. Moreover, the 104-step block iterates permanently.*

To prove permanent iteration, a finite XOR comparison between one block and its translate produces a wake W ; affine ray checks prove that all later translated read supports avoid all archived copies of W . Induction then proves Definition 2.1 for every future cycle.

3 The finite prefix and first-contact coupling

Let $P = \{p(\text{run}(t, \omega)) : 0 \leq t < T\}$ be the set of cells read before the blank highway boundary. Deduplication of the exact trace gives $|P| = 1,376$.

3.1 Prefix contacts

If $q \in P$, the orbit of σ_q first differs from the blank orbit within the finite prefix. For each possible cell, the checked table stores a terminal witness time and exact replay verifies that the resulting state reaches a certified P104 boundary.

Proposition 3.1 (Prefix certificate). *For every $q \in P$, σ_q reaches P104.*

Completeness follows from equality between the stored deduplicated list and the read set of the first T updates.

3.2 Untouched defects

For a point q , say that an orbit avoids q for n steps when its current position differs from q at every time $0, \dots, n-1$. Two states are the same except at q when their poses agree, all other colours agree, and their colours at q are opposite.

Lemma 3.2 (First-difference coupling). *If states a and b are the same except at q , and the orbit of a avoids q for n steps, then $\text{run}(n, a)$ and $\text{run}(n, b)$ remain the same except at q .*

The proof is induction on n . Away from q the two ants read the same colour, make the same turn, flip the same cell, and move together. Since the initial white state and σ_q differ only at q , this gives an exact state identity.

Corollary 3.3 (Untouched entry). *If $q \notin P$, then*

$$\text{run}(T, \sigma_q) = \text{blacken}(q, E).$$

Figure 2 displays the first-contact split established by the read-set equality and coupling statement above.

4 Entry geometry and the 22 frontier channels


The future clean highway reads exactly $S, S+v, S+2v, \dots$. A defect untouched by the prefix can matter only if it lies in their union. Define

$$H = \{h \in S : h+v \notin S\}, \quad \text{obstacle}(h, d) = h + dv \quad (d \geq 1).$$


The sign is important: depth increases in the forward drift direction. Checked diagonal, parity, and order conditions on the 40 points of S give an exhaustive decomposition.

Proposition 4.1 (Entry partition). *For every $q \notin P$, an ordered case split gives one of the following:*

1. q is already black in E , so blackening changes nothing;
2. q is one of the 27 white cells in the 40-cell active support, handled by finite replay;
3. $q \notin \bigcup_{k \geq 0} (S + kv)$, so the blank highway never reads it; or
4. there are unique $h \in H$ and $d \geq 1$ with $q = h + dv$.



(a) Untouched placements.



(b) Prefix-hit placements.

Figure 2: The finite-prefix dichotomy, exported from Stages 01 and 02 of the interactive app. In (a), gray is the blank-orbit read footprint and blue is its complement in the displayed window; the diagonal arrow indicates that the gray P104 corridor continues indefinitely. In (b), the 1,376 blue cells are exactly the deduplicated prefix set P , while gray shows the subsequent P104 corridor. The red outline marks the canonical time-zero ant position.

Furthermore, $|H| = 22$.

The raw predicates need not be pairwise disjoint before earlier branches are removed; the formal theorem uses the stated order. Uniqueness of (h, d) in the infinite remainder replaces two unrestricted coordinates by a 22-element index and one positive natural number. For $h \in H$ and $d \geq 1$, define

$$\Sigma_E(h, d) = \text{blacken}(h + dv, E).$$

It remains to prove that every $\Sigma_E(h, d)$ reaches P104.

5 Pristine single-defect scattering

To extend a finite scattering calculation to arbitrary depth, we first separate the repeating highway mechanism from the fixed blank-orbit history. The physical state E contains both. Its pristine counterpart B retains the 13-cell active pattern and the same ant pose, with the remaining plane white. In this setting, an extra clean highway block contributes a translated XOR wake, which can be tracked uniformly as the defect moves down a channel. This gives the archive induction below. The following two sections restore the physical history, using geometric separation for the ordinary channels and an explicit return-collision analysis for the exceptional channel. Put

$$\Sigma_B(h, d) = \text{blacken}(h + dv, B).$$

For each head h , the finite certificate supplies a positive stable depth P_h and a witness time at depth P_h . Depths $1 \leq d < P_h$ are direct replays. The stable case is extended by a generic archive theorem.

5.1 XOR archives

The coupling principle is local: two states with the same pose generate the same observations as long as their common path avoids the cells on which their grids differ. We express agreement outside a set by the *same-outside* relation and prove avoidance separately using the read sets. Thus a growing archive of differences can be carried through an induction while the ant continues along the same path.

For Boolean grids, finite changes compose by exclusive-or. Let W be the finite difference between one clean 104-step block and its translated target. If the obstacle is moved one additional period down a channel, the ant first executes one additional clean block and deposits one additional translated copy of W . After n clean blocks the difference has the form

$$\mathcal{A}_n(W) = W \oplus (W + v) \oplus \cdots \oplus (W + (n - 1)v). \quad (1)$$

The XOR expression retains cancellation between overlapping wakes. The following lemma combines this algebra with avoidance of the accumulated archive during both scattering and the final highway motion.

Lemma 5.1 (Archive induction). *Fix a clean state, drift v , obstacle anchor, period, anchored scattering trace, wake W , and terminal P104 corridor. Suppose that:*

1. *the clean block is exact up to W ; 2. all translated clean and anchored read sets avoid the relevant copies of W ; and 3. the accumulated archive avoids the translated infinite terminal corridor.*

Then the depth- $P+n$ scattering state reaches the translated anchored terminal state for every $n \geq 0$, up to the inert archive $\mathcal{A}_n(W)$, and hence reaches P104.

The Lean proof packages the three assumptions as `Induction.Data`. It first proves an exact normal form for the finite trace and then uses untouched-region coupling for the infinite corridor. The separation premises are discharged by exact affine-ray predicates for all translated copies.

5.2 All-depth pristine classification

For every head, the certificate checks the nonempty direct row, the stable replay, its deduplicated read set, terminal orientation, and the ray guards required by Lemma 5.1. There are 152 pristine replay witnesses in total, including one stable witness per channel.

Proposition 5.2 (Pristine scattering). *For all $h \in H$ and $d \geq 1$, the pristine state $\Sigma_B(h, d)$ reaches P104. For 21 heads the stable output has the forward standard drift. For the remaining head h_{72} , the stable output has reverse drift $-v = (2, 2)$.*

Figure 3 displays one common-depth replay for every frontier head. These panels are exported directly from Stage 03 of the interactive app: the app runs the exact ant rule to the stored terminal witness and then continues the resulting P104 orbit for 20 cycles. The checked 22-head list and archive induction extend this classification to all depths.



Figure 3: The complete 21 + 1 pristine scattering spectrum, grouped by outgoing orientation. Beige cells are the complete disturbed footprint, gray cells are the undisturbed P104 reference, green cells show 20 terminal P104 cycles, the black cell with a white outline is the selected perturbation q , and the triangle is the ant. Each app canvas is independently fitted to its complete footprint; lattice spacing remains visible in every panel. The boxed phase-72 panel is the unique reverse response and the exceptional channel in the physical-history proof.

The subscript records the phase label used by the artifact, not a count. Its coordinate relative to the pristine base pose is

$$h_{72} = p(B) + (-2, -8), \quad P_{h_{72}} = 11.$$

At this stage the reverse highway is already a valid P104 terminal state. The difficulty appears only when the removed history is restored: the reverse ant eventually travels toward that history rather than away from it.

6 Restoring the physical history: 21 ordinary channels

For an ordinary channel, moving the defect sufficiently far down its ray separates the translated pristine scattering trace and its final highway corridor from the fixed historical wake. This gives a cutoff depth: shallower cases are checked directly, while every deeper case transfers from the pristine induction by untouched-region coupling. We now specify the cutoff and the geometric guards that make this transfer uniform in depth.

Write

$$H_E = \{x : E(x) = 1, x \notin S\}.$$

The exact entry certificate gives $|H_E| = 702$. For an ordinary head h , let A_h be one plus the length of its checked physical direct row. The certificate verifies $P_h \leq A_h$. Thus

$$1 \leq d < A_h$$

is a finite direct band, while $d = A_h + n$ is the inductive region. The offset $L_h = A_h - P_h$ aligns physical depth $A_h + n$ with pristine depth $P_h + L_h + n$.

Each channel's cutoff is certified by separation guards. The finite report checks these guards for every historical cell and the ordinary stable snapshot; ray lemmas extend them to all deeper translations.

Lemma 6.1 (Historical transfer). *Let h be ordinary and $n \geq 0$. Up to the historical region H_E , the orbit from $\Sigma_E(h, A_h + n)$ agrees with the corresponding pristine orbit at depth $P_h + L_h + n$. The common path avoids H_E through the scattering trace and through its permanent terminal corridor.*

The key point is temporal as well as spatial. A finite check shows separation from the anchored trace; the affine guard shows that translating the obstacle farther cannot bring a historical cell into any new trace copy or output ray. Untouched-region coupling then turns geometric nonintersection into exact equality of observations.

Figure 4 shows why a channel-specific cutoff is needed: shallow channel points lie among the finite historical cloud, whereas each stable lane eventually extends past it in the common drift direction.

Proposition 6.2 (Ordinary physical channels). *For each of the 21 ordinary heads h and every $d \geq 1$, $\Sigma_E(h, d)$ reaches P104.*




Figure 4: The ordinary physical-history map, exported from Stage 04 of the interactive app. Gray is the finite time-9,977 history together with the active entry support. Orange cells are channel depths discharged by direct physical replay; blue cells extend the 21 stable ordinary lanes through depth 30 to display their affine continuation. Red and cyan outlines mark the time-zero ant origin and the time-9,977 P104 entry, respectively. The all-depth extension follows from Lemma 6.1.

For $d < A_h$ this is the physical replay table. For $d \geq A_h$, write $d = A_h + n$, apply Lemma 6.1, and invoke pristine archive induction at parameter $L_h + n$. The generated physical rows contain 179 replay witnesses across all 22 channels; the exceptional row is used separately below.

7 Exceptional backscattering and the affine hit law

For the exceptional head, the pristine stable depth is 11 but the physical base depth is

$$P_{h_{72}} = 11, \quad A_{h_{72}} = 15, \quad A_{h_{72}} - P_{h_{72}} = 4.$$

Depths 1 through 14 are direct physical cases. At depth 15, four clean forward translations carry the physical state into the affine reverse family.

7.1 The exact reverse block

Let R be the reverse base state obtained after the checked pristine scattering time of 52,262 updates. Let $v_r = -v = (2, 2)$, and let W_r be the finite difference between one reverse block and the translated state. The certificate and its soundness theorem establish

$$\text{run}(104, R) = T_{v_r} R \oplus W_r. \quad (2)$$

Positive-copy ray guards show that W_r begins affecting the representation only at the next translated copy and is never read by any later reverse block. Equation 2 therefore iterates indefinitely in the pristine setting.

7.2 First historical collision

With the historical wake restored, the reverse path eventually meets a cell that the blank transient had changed. The checked first-hit report proves that the base reverse trajectory avoids all 702 historical cells for ten complete reverse cycles and phases 0 through 88 of the next cycle, then reaches

$$c = (20, -22)$$

at phase 89. Thus the elapsed reverse time is $10 \cdot 104 + 89 = 1,129$ updates. The report also verifies $c \in H_E$.

Increasing physical obstacle depth from $15 + n$ to $15 + (n + 1)$ inserts one forward block before reversal and one reverse block before the collision. The remaining phase and collision cell are unchanged. Consequently

$$t_{\text{hit}}(n) = t_{\text{hit}}(0) + 208n. \quad (3)$$

More strongly, there is a finite XOR layer L and a checked collision state C such that

$$\text{run}(t_{\text{hit}}(n), \Sigma_E(h_{72}, 15 + n)) = C \oplus \bigoplus_{j=0}^{n-1} (L + jv). \quad (4)$$

Equation 4, formalized as `Scattering.Classification.exceptionalAffine`, is the normal form that closes the infinite family. It says more than equal collision coordinates: the complete state at collision is fixed modulo an explicit translated archive.

7.3 The post-hit tail

From C , exact replay for 7,994 updates reaches a forward P104 terminal state. The checked post-hit trace guard proves that every layer in Equation 4 avoids this finite replay. A second forward ray guard proves that every layer also avoids the infinite terminal corridor. Archive induction therefore transfers the depth-15 computation to all $15 + n$.

Proposition 7.1 (Exceptional physical channel). *For every $d \geq 1$, $\Sigma_E(h_{72}, d)$ reaches P104.*

The finite band proves $1 \leq d < 15$; the reverse-block, first-hit, post-hit, and archive lemmas prove $d = 15 + n$. Together with Proposition 6.2, this is the complete single-defect scattering spectrum.




Figure 5: The exceptional phase-72 history replay at the inductive base depth 15, exported directly from Stage 05 of the interactive app. Gray cells are the finite state present at the time-9,977 entry boundary; beige cells are the complete disturbed footprint; the coral diagonal is the reverse P104 segment returning toward the history; and the green diagonal is 20 cycles of the final forward P104 highway after the fixed historical collision and post-hit transient. The black square is the perturbation q and the triangle is the terminal ant. Cyan, red, and orange outlines mark the P104 entry, the original time-zero ant position, and the fixed historical hit, respectively.

8 Universal assembly

We can now prove the main theorem without any remaining search bound.

Proof of Theorem 1.1. Apply Lemma 2.2 and use translation and rotation equivariance to reduce to σ_q . If $q \in P$, apply the prefix certificate, Proposition 3.1. Otherwise Corollary 3.3 identifies the time-9,977 state with $\text{blacken}(q, E)$. Apply the entry partition, Proposition 4.1. The already-black branch is the blank highway; active-support cells are finite checked cases; a geometrically separated cell is never read by the permanent blank highway. In the remaining branch, write uniquely $q = h + dv$ with $h \in H$ and $d \geq 1$. If h is ordinary, apply Proposition 6.2; if $h = h_{72}$, apply Proposition 7.1. Every branch reaches P104, and equivariance restores the original pose. $\square$

Figure 6 combines those branches in one visible-domain map. It provides a spatial index to the preceding propositions and their exhaustive, inductive coverage.




Figure 6: Global one-black proof map, exported from Stage 06 of the interactive app. Every lattice cell in the visible window is assigned to one proof stage: untouched (gray-blue), prefix hit (blue), pristine P104 scattering (green), finite-history intersection (orange), or exceptional reverse highway (purple). A point belonging geometrically to both Stages 03 and 04 is assigned to Stage 04 exactly when its replay or terminal corridor meets the finite history. The colored rays are displayed finite segments of affine classes that the Lean proof extends to every depth.

9 Lean formalization and finite computation

9.1 Proof cone

In Lean, the same assembly is exposed at three levels. The theorem `Scattering.single_defect_scattering` constructs the $21 + 1$ classification. The entry theorem combines that classification with the geometric partition. Finally the exported statement is literally

```
theorem OneBlack.universal_one_black :
   $\forall s, \text{ExactlyOneBlack } s \rightarrow \text{ReachesP104 } s$ 
```

The Lean 4.30.0 development combines analytic proof terms with three finite computation leaves. Its proof cone includes the dynamics, certificate modules and explicit `StandardTrace` bridge; the release manifest records all source digests. Table 1 maps the mathematical stages to their principal modules.

Table 1: Proof-to-module map.

Mathematical obligation	Principal Lean modules
Dynamics, observations, symmetries	Core, Semantics
Blank P104 replay and induction	HighwayData, Highway
First-contact coupling and prefix	Coupling, Prefix, PrefixLeaf
Entry partition and affine channels	Geometry, Rays, Entry
Pristine archive induction	Induction, PristineChecks, Pristine
Physical ordinary channels	ActualChecks, ActualHistory, OrdinaryGeometry, Ordinary
Exceptional reverse family	the Phase* modules
Spectrum and universal theorem	Scattering, Universal, StandardTrace

9.2 Finite leaves

The complete build has three native computation leaves.

1. **Blank leaf.** It replays 9,977 blank updates, checks the 104-step boundary, and validates the finite separation data used by the permanent highway induction.
2. **Prefix leaf.** It checks all 1,376 first-contact cells against stored terminal witness times.
3. **Scattering leaf.** It checks the 27 active-support cases, direct pristine and physical channel rows, one stable pristine snapshot per head, ordinary historical guards, and the exceptional reverse, first-hit, post-hit, and layer guards.

The lane tables contain 152 pristine replay witnesses and 179 physical-entry replay witnesses. Their stored witness times total 1,380,384 and 2,761,211 updates of finite certificate work, respectively. Archive induction extends these witnesses to deeper channel members.

Each stable snapshot computes its terminal state, deduplicated read set, cached read list, terminal orientation, and acceptance result once. Multiple proof guards consume the same values. Verification directly replays the checked-in terminal times and checks their acceptance conditions.

9.3 Reflection and trust boundary

Each Boolean report is connected to a typed certificate by a soundness theorem of the form

`theorem certificate_of_report (verified : report = true) : Certificate`

and the typed certificates feed the analytic theorems. The closed equalities `report = true` are discharged with `native_decide`. This design makes the finite obligations inspectable and reproducible. The trusted computing base consists of Lean’s kernel and foundational axioms together with the native compiler, runtime, code-generation backend and any externally linked native code used by `native_decide` (Lean developers 2026). The final theorems depend on three native-evaluation axioms, one for each finite report.

9.4 Reproduction

The baseline source code, checked witness data, and Lean proof are available in the public repository github.com/kehao95/langtons-ant-core (Ke 2026b). The accompanying revision archive contains

the exact strengthened sources audited here. The repository pins `leanprover/lean4:v4.30.0`. From its root, the complete theorem is rebuilt with

```
python3 one_black/check.py
```

The command builds both closed theorem targets from checked-in sources and witness tables by direct replay. For a fresh build without cached proof objects, run

```
python3 preprint/verify.py
```

This copies the Lean sources into a new temporary directory and records build output, source SHA256 digests and the axiom dependencies of both final theorems. The strengthened conclusion is exposed as

```
theorem OneBlack.universal_one_black_standard_trace
  (s : State) (one : ExactlyOneBlack s) :
  ∃ n, HasStandardTrace (run n s)
```

Here `HasStandardTrace` quantifies a symmetry placement of the fixed actual blank state at update 9977 and equality of every future position, heading and read colour. Both final theorems have exactly the same three native computation leaves; the new bridge adds no native evaluation or proof assumption. The accompanying source package records exact file digests in its manifest and supplies the strengthened theorem and clean-build evidence for this revision of the preprint (Ke 2026a). The interactive website illustrates the proof geometry and channel classification.

10 Related work

10.1 Classical ant dynamics

Langton introduced the model in the context of artificial life and reported the emergence of organized motion from simple local rules (Langton 1986). Bunimovich and Troubetzkoy proved recurrence and unboundedness results for the associated Lorentz lattice-gas dynamics (Bunimovich and Troubetzkoy 1992). Unboundedness is a qualitatively weaker conclusion than Theorem 1.1: it excludes finite confinement but neither selects the period-104 orbit nor gives a time after which translated phase equality holds. Gale, Propp, Sutherland, and Troubetzkoy studied further ant rules and symmetry phenomena (Gale et al. 1995).

10.2 Complexity and generalized highways

Gajardo, Moreira, and Goles embedded Boolean-circuit computation in finite configurations and obtained P-hardness; their universality and undecidability construction uses an infinite, finitely described trapezoidal circuit array (Gajardo, Moreira, and Goles 2002). Their circuit construction can route an ant to a highway seed, giving an infinite family of highway-reaching configurations. The present theorem covers the exhaustive natural family defined solely by “exactly one black cell”, with arbitrary relative position.

For generalized reversible ants, recognition of repeatable configurations can be PSPACE-hard (Tsukiji and Hagiwara 2011). Recent work shows that generalized rules may possess unboundedly or infinitely many distinct highways (Gajardo, Lutfalla, and Rao 2024), and may also exhibit qualitatively different finite-configuration asymptotics (Lutfalla 2025). These results motivate separating the theorem’s classical LR rule and exact P104 predicate from broader claims about generalized ants.

10.3 A geometric method for unbounded perturbations

The proof architecture separates a finite baseline transient from a translated renewal boundary, classifies the first-contact frontier into affine channels, and promotes finite anchors with geometric archive guards. In the exceptional channel, a reverse renewal orbit and an affine collision law account for the return through the historical wake. Together these mechanisms extend finitely checked anchors to every channel depth and suggest a route to studying interactions among several defects.

11 Conclusion

Every classical Langton-ant state with exactly one black cell reaches the standard period-104 highway, in the precise observation-level sense of Definition 2.1. The proof has no coordinate cutoff. It reduces the 1,376 early-contact cells to finite replay and every later relevant cell to one of 22 affine channels. Twenty-one channels become insensitive to the blank-orbit history after channel-specific cutoffs. The remaining channel creates a reverse highway, returns to a fixed historical cell with the affine law $t_{\text{hit}}(n) = t_{\text{hit}}(0) + 208n$, and then enters forward P104 through an archived post-hit tail.

The Lean development checks both the infinite reductions and the completeness of the case split. Its three finite leaves are evaluated natively, with the corresponding compiler/runtime dependencies recorded in the artifact. The result establishes a fully quantified convergence theorem for an unbounded perturbation family through scattering geometry and induction.

The next mathematical question is whether the frontier-channel method survives multiple defects. A second defect introduces interactions between independent contact times and between accumulated wakes, so a successful extension will likely require a compositional renewal invariant.

A Certificate inventory

Table 2 collects numerical constants used in the proof. Every value is asserted by a checked report or derived from a checked list; none is a parameter of the theorem statement.

Table 2: Checked certificate inventory.

Quantity	Value	Role
Blank boundary time	9,977	start of certified P104
Distinct prefix read cells	1,376	complete early-contact branch
P104 period and drift	104; $(-2, -2)$	standard highway convention
One-period support	40	entry geometry
Active white support cells	27	finite entry cases
Entry black cells	715	active pattern plus history
Historical cells	702	physical-history guards
Frontier heads	22	all unbounded future contacts
Ordinary / exceptional heads	21 / 1	scattering classification
Exceptional pristine / physical depth	11 / 15	reverse-family anchors
Exceptional reverse scattering time	52,262	exact pristine anchor replay
Reverse first-hit elapsed time	1,129	$10 \cdot 104 + 89$
Historical hit point	$(20, -22)$	depth-independent collision
Post-hit replay	7,994	terminal forward P104 witness
Pristine / physical replay witnesses	152 / 179	finite scattering leaves

B References

- Bunimovich, Leonid A., and Serge E. Troubetzkoy. 1992. “Recurrence Properties of Lorentz Lattice Gas Cellular Automata.” *Journal of Statistical Physics* 67 (1–2): 289–302. <https://doi.org/10.1007/BF01049035>.
- Gajardo, Anahí, Victor Lutfalla, and Michaël Rao. 2024. “Ants on the Highway.” <https://arxiv.org/abs/2409.10124>.
- Gajardo, Anahí, Andrés Moreira, and Eric Goles. 2002. “Complexity of Langton’s Ant.” *Discrete Applied Mathematics* 117 (1–3): 41–50. [https://doi.org/10.1016/S0166-218X\(00\)00334-6](https://doi.org/10.1016/S0166-218X(00)00334-6).
- Gale, David, James Propp, Scott Sutherland, and Serge Troubetzkoy. 1995. “Further Travels with My Ant.” *The Mathematical Intelligencer* 17 (3): 48–56. <https://arxiv.org/abs/math/9501233>.
- Ke, Hao. 2026a. “Every One-Black-Cell Langton Ant Reaches the Period-104 Highway: A Lean-Checked Proof.” Zenodo. <https://doi.org/10.5281/zenodo.22086968>.

- . 2026b. “Langton’s Ant: Proof and Research Core.” <https://github.com/kehao95/langtons-ant-core>.
- Langton, Christopher G. 1986. “Studying Artificial Life with Cellular Automata.” *Physica D: Nonlinear Phenomena* 22 (1–3): 120–49. [https://doi.org/10.1016/0167-2789\(86\)90237-X](https://doi.org/10.1016/0167-2789(86)90237-X).
- Lean developers. 2026. *Lean Language Reference: Axioms*. <https://lean-lang.org/doc/reference/latest/Axioms/>.
- Lutfalla, Victor. 2025. “Sideways on the Highways.” <https://arxiv.org/abs/2505.05426>.
- Moura, Leonardo de, and Sebastian Ullrich. 2021. “The Lean 4 Theorem Prover and Programming Language.” In *Automated Deduction – CADE 28*, edited by André Platzer and Geoff Sutcliffe, 12699:625–35. Lecture Notes in Computer Science. Springer. https://doi.org/10.1007/978-3-030-79876-5_37.
- Tsukiji, Tatsue, and Takeo Hagiwara. 2011. “Recognizing the Repeatable Configurations of Time-Reversible Generalized Langton’s Ant Is PSPACE-Hard.” *Algorithms* 4 (1): 1–15. <https://doi.org/10.3390/a4010001>.

C Version history and acknowledgements

An earlier version of this manuscript is publicly available as [doi:10.5281/zenodo.22086968](https://doi.org/10.5281/zenodo.22086968). This revision strengthens the exported semantic statement and documents a fresh proof build. Generative AI assisted in proof development, code, auditing and manuscript preparation.