

SWITCH CORRIDORS AND BOUNDARY-RIGID LENSES IN EVEN DUMBBELL THRACKLES: A PROOF STRATEGY FOR CONWAY'S CONJECTURE

LUCA BLANCHI

ABSTRACT. We prove Conway's thrackle conjecture by reducing it to the even dumbbell case and then showing that every irreducible thrackle drawing of an even dumbbell is a T_3 -thrackle.

The proof proceeds through the following chain:

$$\begin{aligned} \text{minimal counterexample} &\Rightarrow \text{even dumbbell} \\ &\Rightarrow \text{irreducible even dumbbell} \\ &\Rightarrow T_3 \\ &\Rightarrow |E| \leq |V|, \end{aligned}$$

while an even dumbbell satisfies

$$|E| = |V| + 1.$$

The new part is the implication

$$\text{irreducible even dumbbell thrackle} \Rightarrow T_3.$$

This is proved by contradiction. If the drawing is not T_3 , then a finite domain-incidence argument gives two ordinary vertices p, q on the same cycle with disjoint incident-domain pairs:

$$I(p) \cap I(q) = \emptyset.$$

The two paths between p and q on that cycle generate switch corridors. A switch corridor is either reducible or regular. A regular corridor produces a local ladder or a closed alternating train. Complementary ladders close into a train. Closed alternating trains are impossible by a braid/full-twist argument and a support-shift reduction. Hence no non- T_3 irreducible even dumbbell thrackle exists.

The technical core is the switch-corridor regularization theorem. Its key local ingredient is boundary rigidity for reduced quadrilateral lenses: if a local lens contains no reducible configuration, then its domain lanes are pairwise disjoint properly embedded arcs in a disk, hence preserve order; therefore the lens induces the identity boundary relation and may be removed from a minimal relational corridor.

1. EXTERNAL INPUTS, CONVENTIONS, AND STANDING DEFINITIONS

1.1. **General conventions.** All graphs in this paper are finite and simple unless explicitly stated otherwise. A drawing of a graph in the plane is

Date: June 2026.

assumed to be in general position: vertices are represented by distinct points, edges are represented by Jordan arcs joining the corresponding endpoints, no edge passes through a vertex to which it is not incident, no three edges meet at a common interior point, and all intersections between interiors of edges are proper crossings.

A thrackle drawing of a graph $G = (V, E)$ is a drawing in the plane such that every pair of distinct edges meets exactly once. More explicitly:

- (1) if two edges are adjacent, their unique meeting point is their common endpoint;
- (2) if two edges are non-adjacent, they meet in exactly one proper crossing.

A graph is called thrackleable if it admits a thrackle drawing.

Conway's thrackle conjecture asserts that if $G = (V, E)$ is thrackleable, then

$$|E| \leq |V|.$$

The goal of the paper is to prove this inequality.

The proof uses three standard external inputs from the literature. They are stated below in the precise forms needed later.

1.2. External Input I: reduction to even dumbbells. A dumbbell graph consists of two vertex-disjoint cycles connected by a path, with the endpoints of the path attached to one vertex of each cycle. If the connecting path has length zero, the two cycles share a single vertex. We denote this special case by

$$DB(c', c'', 0).$$

Thus $(DB(c', c'', 0))$ is the graph obtained by identifying one vertex of a cycle of length c' with one vertex of a cycle of length c'' .

The form of the known reduction used in this paper is the following.

Theorem 1.1 — Even dumbbell reduction. To prove Conway's thrackle conjecture, it is enough to prove that no graph

$$DB(c', c'', 0)$$

with c' and c'' both even is thrackleable.

Equivalently, if Conway's conjecture is false, then there exists a thrackle drawing of a graph consisting of two even cycles sharing exactly one vertex.

In this paper we write such a graph as

$$F_{2a, 2b} = C_{2a} \vee C_{2b}.$$

The two cycles are denoted by

$$A = C_{2a}, \quad B = C_{2b},$$

and their common vertex is denoted by

$$v.$$

The notation $C_{2a} \vee C_{2b}$ means the one-point union of the two cycles: one vertex of C_{2a} is identified with one vertex of C_{2b} , and no other vertices or edges are identified.

The graph $F_{2a,2b}$ has

$$|E(F_{2a,2b})| = 2a + 2b$$

and

$$|V(F_{2a,2b})| = (2a) + (2b) - 1.$$

Therefore

$$|E(F_{2a,2b})| = |V(F_{2a,2b})| + 1.$$

Consequently, every thrackle drawing of such a graph would violate Conway's inequality unless it is shown to be impossible.

1.3. External Input II: the T_3 -thrackle theorem. Let T be a plane drawing of a graph. The domains of T are the connected components of the complement of the drawing in the plane.

A drawing is called a T_d -drawing if all original vertices of the graph lie on the boundaries of at most d connected domains of the complement.

Thus a drawing is T_3 if there exist three connected domains

$$D_1, D_2, D_3$$

such that every original vertex of the graph belongs to

$$\partial D_1 \cup \partial D_2 \cup \partial D_3.$$

The external result needed here is the following.

Theorem 1.2 — T_3 -thrackles satisfy Conway. If a graph $G = (V, E)$ admits a T_3 -thrackle drawing, then

$$|E| \leq |V|.$$

The proof below will therefore focus on showing that every irreducible thrackle drawing of an even dumbbell is T_3 . Since an even dumbbell has $|E| = |V| + 1$, such a drawing cannot exist.

1.4. External Input III: edge removal and irreducibility. We also use the standard edge-removal principle for thrackles.

A 3-path is a path

$$P = (x_0, x_1, x_2, x_3)$$

of length (3) in the graph, usually with the middle vertices x_1, x_2 of degree (2) in the relevant subconfiguration. Associated with such a path in a thrackle drawing is an edge-removal triangle, denoted

$$\Delta(P).$$

The precise geometric construction of $\Delta(P)$ is standard: it is the triangular region bounded by the three drawn edge-arcs of the path and the direct replacement arc from x_0 to x_3 , chosen in the locally prescribed way.

The form needed here is the following.

Theorem 1.3 — Edge-removal criterion. If

$$\Delta(P) \cap V_{\text{orig}} = \emptyset,$$

then the 3-path P can be removed and replaced by a single edge joining x_0 to x_3 , while preserving the thrackle property.

A thrackle drawing is called irreducible if no such edge-removal operation is possible.

In a minimal counterexample obtained through the even-dumbbell reduction, we may assume irreducibility: if an edge-removal were possible, it would produce a strictly smaller thrackle drawing in the same reduction scheme, contradicting minimality.

Thus, throughout the internal proof, we work with an irreducible thrackle drawing of

$$F_{2a,2b} = C_{2a} \vee C_{2b}.$$

1.5. Summary of the external framework. Combining the three inputs, the overall route is as follows.

To prove Conway's conjecture, it is enough to exclude thrackle drawings of even dumbbells

$$F_{2a,2b} = C_{2a} \vee C_{2b}.$$

Such a graph satisfies

$$|E| = |V| + 1.$$

If every irreducible thrackle drawing of $F_{2a,2b}$ were T_3 , then Theorem 1.2 would give

$$|E| \leq |V|,$$

contradicting

$$|E| = |V| + 1.$$

Therefore the key internal theorem is:

Every irreducible thrackle drawing of an even dumbbell is T_3 .

The rest of the paper is devoted to proving this internal theorem.

2. REDUCTION OF CONWAY'S CONJECTURE TO THE INTERNAL THEOREM

2.1. Statement of Conway's conjecture.

Theorem 2.1 — Conway's thrackle conjecture. Let $G = (V, E)$ be a finite graph admitting a plane thrackle drawing. Then

$$|E| \leq |V|.$$

We prove Theorem 2.1 by reducing it to the internal theorem stated in Section 3.

2.2. The internal theorem used in the reduction. The internal theorem is the following.

Theorem 2.2 — Internal even-dumbbell theorem. Let

$$F_{2a,2b} = C_{2a} \vee C_{2b}$$

be an even dumbbell. Every irreducible thrackle drawing of $F_{2a,2b}$ belongs to T_3 .

The proof of Theorem 2.2 is the main work of the paper. It is developed from Section 3 onward.

We now show that Theorem 2.2 implies Conway's conjecture.

2.3. Proof that the internal theorem implies Conway. Assume, for contradiction, that Conway's thrackle conjecture is false.

Then there exists a finite graph $G = (V, E)$ admitting a thrackle drawing with

$$|E| > |V|.$$

By the even-dumbbell reduction, there exists a thrackle drawing of an even dumbbell

$$F_{2a,2b} = C_{2a} \vee C_{2b}.$$

Choose such a drawing minimal with respect to the edge-removal reduction. By the edge-removal criterion, this drawing is irreducible.

By Theorem 2.2, this irreducible drawing belongs to T_3 .

Therefore, by the T_3 -thrackle theorem,

$$|E(F_{2a,2b})| \leq |V(F_{2a,2b})|.$$

But direct counting gives

$$|E(F_{2a,2b})| = 2a + 2b,$$

whereas

$$|V(F_{2a,2b})| = 2a + 2b - 1.$$

Thus

$$|E(F_{2a,2b})| = |V(F_{2a,2b})| + 1.$$

This contradicts

$$|E(F_{2a,2b})| \leq |V(F_{2a,2b})|.$$

Therefore no counterexample to Conway's conjecture exists. Hence every planar thrackle drawing satisfies

$$|E| \leq |V|.$$

This proves Theorem 2.1, assuming Theorem 2.2.

3. THE INTERNAL EVEN-DUMBBELL THEOREM

3.1. Statement. Let

$$F_{2a,2b} = C_{2a} \vee C_{2b}$$

be an even dumbbell. Let $A = C_{2a}$ and $B = C_{2b}$, and denote their common vertex by v .

Let T be a thrackle drawing of $F_{2a,2b}$. We assume T is irreducible with respect to the edge-removal operation.

The internal theorem is:

Theorem 3.1 — Irreducible even dumbbell thrackles are T_3 . If T is an irreducible thrackle drawing of

$$F_{2a,2b},$$

then

$$T \in T_3.$$

Equivalently, there exist three connected domains

$$D_1, D_2, D_3$$

of the complement of the drawing such that

$$V_{\text{orig}}(F_{2a,2b}) \subseteq \partial D_1 \cup \partial D_2 \cup \partial D_3.$$

3.2. General strategy. The proof is by contradiction.

Assume that T is an irreducible thrackle drawing of $F_{2a,2b}$ and that

$$T \notin T_3.$$

We planarize the drawing. Let

$$P(T)$$

be the planarized graph obtained by replacing every proper crossing by a vertex of degree (4).

The complement of $(P(T))$ has finitely many connected components, called domains.

For every original vertex x , define

$$I(x) = D : x \in \partial D,$$

the set of domains incident to x .

The condition $T \in T_3$ is equivalent to the existence of three domains meeting all the sets $(I(x))$. Thus $T \notin T_3$ gives a finite hitting-set obstruction.

The first step is to convert this obstruction into a pair of ordinary vertices on the same cycle whose incident-domain pairs are disjoint.

The second step is to use those two vertices to build switch corridors.

The third step is to show that every minimal switch corridor is either reducible or regular.

The fourth step is to show that regular corridors produce ladders or closed alternating trains.

The fifth step is to show that complementary ladders close into a train and that closed trains are impossible.

This contradiction proves that the assumption

$$T \notin T_3$$

was false.

3.3. Step I: from $T \notin T_3$ to disjoint domain pairs. The first main lemma is the domain-pair separation lemma.

Lemma 3.2 — Domain-pair separation. If

$$T \notin T_3,$$

then there exist two ordinary original vertices

$$p, q \neq v$$

lying on the same cycle A or B such that

$$I(p) \cap I(q) = \emptyset.$$

The proof uses the ordinary domain graph. For every ordinary vertex $x \neq v$, the vertex has degree (2), hence is incident to exactly two domains:

$$I(x) = D_x^-, D_x^+.$$

We build a graph G_D whose vertices are domains and whose edges are the pairs

$$D_x^- D_x^+.$$

Because the planarized graph is Eulerian, its face-dual is bipartite. Hence G_D is bipartite. Covering all ordinary vertices by domains is equivalent to finding a vertex cover of G_D . König's theorem then translates failure of 3-coverability into the existence of a matching.

If the ordinary vertices are not 3-coverable, G_D has a matching of size at least (4). Among the four corresponding vertices, two lie on the same cycle and have disjoint incidence pairs.

If the ordinary vertices are 3-coverable but the whole drawing is not T_3 , then the common vertex v is responsible for the obstruction. In that case, deleting any domain incident to v leaves a bipartite graph with no vertex cover of size (2), hence with a matching of size (3). Among those three vertices, two lie on the same cycle and have disjoint incidence pairs.

Thus the lemma follows.

This lemma converts a global failure of T_3 into a local switching problem along a single cycle.

3.4. Step II: switch corridors. Let p, q be the two vertices supplied by Lemma 3.2. Suppose, without loss of generality, that both lie on A .

There are two paths in A from p to q . Consider one such path γ .

Along every edgelet of the planarized path γ , there are two lateral domains. Near p , the lateral pair is $(I(p))$. Near q , the lateral pair is $(I(q))$.

Because

$$I(p) \cap I(q) = \emptyset,$$

the pair of lateral domains must change essentially along γ .

A switch corridor is a disk-like region in the planarization that records this change of lateral domain data.

The precise object used in the proof is a relational switch corridor. It consists of finitely many normal quadrilateral chambers

$$Q_1, \dots, Q_N$$

and a total boundary relation

$$R_C = R_{Q_N} \circ \dots \circ R_{Q_1}.$$

The purpose of the corridor is to encode the finite topological data of how domain lanes pass from the initial lateral pair to the final lateral pair.

3.5. Step III: regularization of switch corridors. The central technical theorem is the switch-corridor regularization theorem.

Theorem 3.3 — Switch-corridor regularization. Every minimal switch corridor is either reducible or regular.

Here “reducible” means that the corridor contains one of the following standard forbidden configurations:

- (1) an edge-removal;
- (2) a double crossing;
- (3) a repeated support;
- (4) a pure ear;
- (5) a relevant inner domain or smaller T_3 -critical core;
- (6) a smaller non-identity relational subcorridor;
- (7) a gate identification;
- (8) a terminal lane;
- (9) a returning lane;
- (10) a branching or merging lane;
- (11) a closed alternating train.

A corridor is “regular” if it decomposes into alternating cells

$$H_{2m}, \quad m \geq 3,$$

glued along gates with unique continuation of lanes.

The proof of Theorem 3.3 has two main ingredients.

First, every geometric switch corridor admits a relational decomposition into normal quadrilaterals. This is obtained by choosing a generic PL height function and slicing the corridor into finitely many transverse strips.

Second, reduced quadrilaterals are boundary-rigid. Namely, if a normal quadrilateral contains no reducible feature, then its domain lanes are pairwise disjoint properly embedded arcs in a disk. Such arcs preserve the order of boundary states. Therefore the boundary relation is the identity.

This boundary-rigidity principle eliminates monochromatic lenses. If a monochromatic lens is reduced, then it is identity and may be removed from a minimal relational corridor. If it is not reduced, it already contains a reducible configuration.

Thus any non-reducible minimal corridor must be regular.

3.6. Step IV: regular corridors produce ladders or trains. Once a corridor is regular, its structure is controlled by its gate graph.

The gate graph has:

- vertices equal to gates;
- edges equal to lane continuations.

In a regular corridor, every internal gate has unique continuation. Thus every internal gate has degree (2).

Therefore each connected component of the gate graph is either:

- (1) a path; or

(2) a cycle.

A cyclic component is a closed alternating train.

A path component connecting the two endpoint sections is a local ladder.

Thus:

Lemma 3.4 — Gate graph dichotomy. A regular switch corridor contains either a local ladder or a closed alternating train.

3.7. Step V: complementary ladders close into trains. The two paths from p to q in the cycle A are complementary.

If one of their corridors already contains a closed alternating train, we are in the train case.

If both corridors produce local ladders, then the two ladders glue at p and q . Indeed, p and q are ordinary vertices of degree (2). Locally, each has exactly two incident domains. Hence the terminal gates of the two ladders are paired uniquely by those two domains.

Failure of this pairing would produce a terminal lane, a branching, or a gate identification, each of which is reducible.

Therefore:

Lemma 3.5 — Complementary ladder closure. If both complementary p -to- q corridors produce ladders, then the union of those ladders contains a closed alternating train.

3.8. Step VI: closed alternating trains are impossible. The final structural theorem is the train-impossibility theorem.

Theorem 3.6 — No closed alternating trains. A closed alternating train produced by a regular switch corridor cannot occur in a thrackle drawing.

The proof has two parts.

First, if the train contains a uniform block of arity m consisting of m consecutive cells on the same m original supports, then the braid represented by that block is

$$(\sigma_1\sigma_2\cdots\sigma_{m-1})^m = \Delta^2.$$

The full twist Δ^2 forces some pair of original supports to cross twice, contradicting the thrackle condition.

Second, if a support-shift occurs before such a uniform full-twist block appears, follow a lane around the closed train and record the cyclic itinerary of original supports. A first repeated support gives a closed walk. An innermost simple cycle extracted from that walk gives a pure ear, a relevant inner domain, a smaller train, a repeated support, or a double crossing. If no support repeats before the final return, the lane itself is a simple closed curve and the disk it bounds gives the same alternatives.

Each alternative contradicts minimality or the thrackle condition.

Thus no closed alternating train exists.

3.9. Completion of the internal theorem. We now assemble the contradiction.

Assume

$$T \notin T_3.$$

By Lemma 3.2, choose ordinary vertices p, q on the same cycle such that

$$I(p) \cap I(q) = \emptyset.$$

Consider the two switch corridors corresponding to the two p -to- q paths on that cycle.

By Theorem 3.3, each corridor is reducible or regular.

If a corridor is reducible, then it contains edge-removal, double crossing, pure ear, relevant inner domain, smaller corridor, or closed train. These contradict irreducibility, the thrackle condition, minimality, or Theorem 3.6.

Thus both corridors are regular.

By Lemma 3.4, each regular corridor contains either a ladder or a closed train.

If a closed train occurs, Theorem 3.6 gives a contradiction.

If both corridors yield ladders, Lemma 3.5 glues them into a closed train, again contradicting Theorem 3.6.

Therefore the assumption

$$T \notin T_3$$

is impossible.

Hence

$$T \in T_3.$$

This proves Theorem 3.1.

3.10. Logical dependencies. The proof of Theorem 3.1 depends on the following internal results:

(1) Domain-pair separation:

$$T \notin T_3 \Rightarrow \exists p, q \text{ on the same cycle with } I(p) \cap I(q) = \emptyset.$$

- (1) Geometric realization of switch corridors as relational corridors.
- (2) Boundary rigidity of reduced quadrilaterals.
- (3) Switch-corridor regularization.
- (4) Gate graph dichotomy:

regular corridor \Rightarrow ladder or train.

- (1) Complementary ladder closure.
- (2) Impossibility of closed alternating trains.

These are proved in the subsequent sections.

4. PLANARIZATION, DOMAINS, AND THE T_3 CONDITION

Throughout this section let

$$T$$

be a thrackle drawing of the even dumbbell

$$F_{2a,2b} = C_{2a} \vee C_{2b}.$$

We denote the two cycles by

$$A = C_{2a}, \quad B = C_{2b},$$

and their common vertex by

$$v.$$

All other original vertices are called ordinary vertices.

The purpose of this section is to pass from the drawing T to a finite planar object, its planarization, and to express the T_3 -condition as a finite hitting problem for domains.

4.1. **Planarization.** Let

$$P(T)$$

be the planarization of the drawing T . Thus every proper crossing point between two non-adjacent original edges is replaced by a new vertex of degree (4). The original vertices remain vertices of $(P(T))$.

The vertices of $(P(T))$ are therefore partitioned as

$$V(P(T)) = V_{\text{orig}} \sqcup V_{\times},$$

where:

- V_{orig} is the set of original vertices of $F_{2a,2b}$;
- V_{\times} is the set of crossing vertices introduced by planarization.

The edges of $(P(T))$ are called edgelets. Each edgelet is a subarc of an original edge of $F_{2a,2b}$, cut at crossing points and at original endpoints.

Every edgelet inherits a color:

$$A \quad \text{or} \quad B,$$

according to whether the original edge containing it belongs to the cycle A or the cycle B .

The underlying subset of the plane occupied by $(P(T))$ is the same as the underlying subset occupied by the original drawing T . The planarization only changes the combinatorial structure at crossing points. Therefore the connected components of the complement are the same for T and for $(P(T))$.

4.2. Domains. The domains of T , or equivalently of $(P(T))$, are the connected components of

$$\mathbb{R}^2 \setminus P(T).$$

The set of domains is denoted by

$$\mathcal{F} = \mathcal{F}(P(T)).$$

Since $(P(T))$ is a finite planar graph, \mathcal{F} is finite.

For an original vertex

$$x \in V_{\text{orig}},$$

define the domain-incidence set

$$I(x) = D \in \mathcal{F} : x \in \partial D.$$

Thus $(I(x))$ is the set of domains whose boundary contains x .

The common vertex v may be incident to up to four local sectors, and consequently to up to four domains globally. Ordinary vertices have degree (2), and their local structure is simpler.

4.3. Ordinary vertices and their two incident domains. Let

$$x \in V_{\text{orig}} \setminus v$$

be an ordinary vertex.

In the graph $F_{2a,2b}$, the vertex x has degree (2). Therefore, in the planarized graph $(P(T))$, the vertex x also has degree (2). Its two incident edgelets are the two edge germs of the corresponding cycle through x .

Locally near x , the embedded graph looks like two arcs meeting at a common endpoint. Therefore the punctured neighborhood of x is divided into two local sectors.

We shall use the following elementary fact.

Lemma 4.1 — Ordinary vertices have two incident domains. For every ordinary vertex

$$x \neq v,$$

the set $(I(x))$ consists of exactly two domains.

Thus we may write

$$I(x) = D_x^-, D_x^+.$$

\{\}\text{Proof.}

The vertex x has degree (2) in $(P(T))$. Hence a sufficiently small closed disk B_x centered at x , meeting $(P(T))$ only in the two incident edge germs, has

$$B_x \setminus P(T)$$

equal to two connected local sectors.

It remains only to note that these two local sectors belong to two distinct global domains.

Suppose, for contradiction, that the two local sectors belonged to the same global domain. Then one could join them by an arc in

$$\mathbb{R}^2 \setminus P(T)$$

outside B_x . Together with a small arc around x inside B_x , this would show that the two sides of one of the incident edgelets at x lie in the same face. In a plane graph this can happen only if that edgelet is a bridge of the embedded graph.

But $(P(T))$ is connected and Eulerian. Indeed:

- every crossing vertex has degree (4);
- every ordinary original vertex has degree (2);
- the common vertex v has degree (4).

A connected Eulerian graph has no bridges. Therefore no edgelet of $(P(T))$ is a bridge.

Thus the two local sectors at x lie in two distinct domains.

Hence

$$|I(x)| = 2.$$

□

4.4. The T_3 -condition as a hitting problem. The drawing T is called T_3 if there exist three domains

$$D_1, D_2, D_3 \in \mathcal{F}$$

such that every original vertex lies on the boundary of at least one of them:

$$V_{\text{orig}} \subseteq \partial D_1 \cup \partial D_2 \cup \partial D_3.$$

Equivalently, if

$$S = D_1, D_2, D_3,$$

then

$$S \cap I(x) \neq \emptyset$$

for every

$$x \in V_{\text{orig}}.$$

Thus $T \in T_3$ if and only if the finite family

$$I(x) : x \in V_{\text{orig}}$$

has a transversal of size at most (3) inside \mathcal{F} .

Equivalently,

$$T \notin T_3$$

means:

no set of at most three domains intersects all the sets $I(x)$.

This is the finite domain-cover obstruction that will be exploited in the next section.

4.5. Eulerian planar graphs and bipartite duals. We shall need a standard fact about plane graphs.

Let G be a connected plane graph. Its face-dual G^* has one vertex for each face of G , and one dual edge crossing each edge of G . If G has multiple edges, the dual is understood as a multigraph. This causes no difficulty for bipartiteness or matching arguments.

Lemma 4.2 — Eulerian plane graphs have bipartite face-duals. Let G be a connected plane graph. If every vertex of G has even degree, then the faces of G admit a proper two-coloring. Equivalently, the face-dual G^* is bipartite.

$\{\}\text{textbf{Proof.}}$

Because every vertex of G has even degree, G is Eulerian. An Eulerian plane graph has the property that every closed walk in the dual has even length.

One way to see this is as follows. Choose a face F_0 and color it black. For any other face F , choose a path in the dual graph from F_0 to F , and color F according to the parity of the length of that path. This is well-defined if and only if every cycle in the dual has even length.

A cycle in the dual corresponds to a cut in the primal graph. Since every vertex of the primal graph has even degree, every cut has even cardinality. Hence every dual cycle has even length.

Therefore the parity coloring is well-defined and gives a bipartition of the dual graph.

Thus G^* is bipartite.

□

Applying this to $G = P(T)$, we conclude that the face-dual of $(P(T))$ is bipartite.

This bipartiteness is the key combinatorial input in the domain graph argument.

5. THE DOMAIN GRAPH AND THE DOMAIN-PAIR SEPARATION LEMMA

The goal of this section is to prove that if $T \notin T_3$, then there exist two ordinary vertices on the same cycle whose domain-incidence pairs are disjoint.

This is the first major reduction in the proof.

5.1. The ordinary domain graph. For every ordinary vertex

$$x \in V_{\text{orig}} \setminus v,$$

Lemma 4.1 gives

$$I(x) = D_x^-, D_x^+.$$

Define the ordinary domain graph

$$G_D$$

as follows.

The vertex set of G_D is the set of domains:

$$V(G_D) = \mathcal{F}.$$

For every ordinary vertex

$$x \neq v,$$

add an edge

$$e_x$$

joining the two domains in $I(x)$:

$$e_x = D_x^- D_x^+.$$

Thus edges of G_D are indexed by ordinary original vertices.

The graph G_D is allowed to be a multigraph: two different ordinary vertices may determine the same unordered pair of domains. This is harmless. A matching and a vertex cover are defined in the usual way for multigraphs.

By Lemma 4.1, no edge of G_D is a loop.

5.2. G_D is bipartite.

Lemma 5.1. The ordinary domain graph G_D is bipartite.

\{\}\textbf{Proof.}

The graph G_D is a subgraph of the face-dual graph of $(P(T))$.

Indeed, an edge $e_x = D_x^- D_x^+$ of G_D records the fact that the two domains D_x^- and D_x^+ meet at the ordinary vertex x . In the local rotation at x , they lie on the two sides of the two edgelets incident to x . Since x is a degree-two vertex and the incident edgelets are not bridges, the two domains are adjacent through the local structure of the embedded graph. Equivalently,

they are separated by the local passage of the planarized drawing through x , and they inherit opposite colors in the face two-coloring.

By Lemma 4.2, the face-dual of $(P(T))$ is bipartite. Therefore every subgraph of it, including G_D , is bipartite.

□

A more direct formulation is: color all domains black and white using the face two-coloring of $(P(T))$. At every ordinary vertex, the two incident domains lie on opposite local sides of the embedded graph and hence have opposite colors. Therefore every edge of G_D joins opposite colors.

5.3. Covering ordinary vertices as a vertex cover problem. A set of domains

$$S \subseteq \mathcal{F}$$

covers all ordinary vertices if and only if for every ordinary vertex $x \neq v$,

$$S \cap I(x) \neq \emptyset.$$

Because

$$I(x) = D_x^-, D_x^+,$$

this is equivalent to saying that S contains at least one endpoint of the edge

$$e_x = D_x^- D_x^+$$

of G_D .

Therefore:

Lemma 5.2. A set

$$S \subseteq \mathcal{F}$$

covers all ordinary vertices if and only if S is a vertex cover of G_D .

\{\}\textbf{Proof.}

Suppose S covers all ordinary vertices. Let

$$e_x = D_x^- D_x^+$$

be any edge of G_D . Since S covers x ,

$$S \cap D_x^-, D_x^+ \neq \emptyset.$$

Thus S contains at least one endpoint of e_x . Therefore S is a vertex cover.

Conversely, suppose S is a vertex cover of G_D . Let $x \neq v$ be an ordinary vertex. The corresponding edge of G_D is

$$e_x = D_x^- D_x^+.$$

Since S is a vertex cover, S contains at least one endpoint of e_x . Hence

$$S \cap I(x) \neq \emptyset.$$

Thus S covers all ordinary vertices.

□

5.4. König's theorem. We use König's theorem in the following standard form.

Theorem 5.3 — König. If H is a finite bipartite graph, then

$$\tau(H) = \nu(H),$$

where $\tau(H)$ is the size of a minimum vertex cover of H , and $\nu(H)$ is the size of a maximum matching of H .

This theorem applies to G_D , since G_D is bipartite.

5.5. The domain-pair separation lemma. We now prove the main result of this section.

Lemma 5.4 — Domain-pair separation. Assume

$$T \notin T_3.$$

Then there exist two ordinary vertices

$$p, q \in V_{\text{orig}} \setminus v$$

lying on the same cycle, either both on A or both on B , such that

$$I(p) \cap I(q) = \emptyset.$$

\{\}\textbf{Proof.}

Let G_D be the ordinary domain graph.

There are two cases.

\{\}\textbf{Case 1: the ordinary vertices are not coverable by three domains}

Suppose first that no set of three domains covers all ordinary vertices.

By Lemma 5.2, this means that G_D has no vertex cover of size (3). Therefore

$$\tau(G_D) > 3.$$

Since G_D is bipartite, König's theorem gives

$$\nu(G_D) = \tau(G_D) > 3.$$

Thus

$$\nu(G_D) \geq 4.$$

So G_D contains a matching of size (4).
Let the four matched edges be

$$e_{x_1}, e_{x_2}, e_{x_3}, e_{x_4}.$$

Each edge e_{x_i} corresponds to an ordinary vertex x_i . Since the four edges form a matching, they are pairwise vertex-disjoint in G_D . Therefore the corresponding domain-incidence pairs are pairwise disjoint:

$$I(x_i) \cap I(x_j) = \emptyset \quad (i \neq j).$$

The four ordinary vertices x_1, x_2, x_3, x_4 lie on the two cycles A and B . By the pigeonhole principle, at least two of them lie on the same cycle. Call those two vertices

$$p, q.$$

Then

$$I(p) \cap I(q) = \emptyset.$$

This proves the lemma in Case 1.

\{\}\textbf{Case 2: the ordinary vertices are coverable by three domains}

Now suppose the ordinary vertices are coverable by three domains.

Since

$$T \notin T_3,$$

no set of three domains covers all original vertices, including the common vertex v . Thus the obstruction must involve v .

We first show that the ordinary vertices are not coverable by two domains.

Suppose, for contradiction, that there exist two domains

$$D_1, D_2$$

covering all ordinary vertices. Since v is an original vertex, it is incident to at least one domain. Choose

$$D_3 \in I(v).$$

Then the three domains

$$D_1, D_2, D_3$$

cover all ordinary vertices and also cover v . Hence they cover all original vertices, contradicting

$$T \notin T_3.$$

Therefore the ordinary vertices have minimum cover number exactly (3). In terms of G_D ,

$$\tau(G_D) = 3.$$

Next we claim that no minimum vertex cover of G_D contains a domain incident to v .

Indeed, suppose S is a vertex cover of G_D with

$$|S| = 3$$

and

$$S \cap I(v) \neq \emptyset.$$

Then S covers all ordinary vertices, by Lemma 5.2. Since S contains a domain incident to v , it also covers v . Therefore S covers all original vertices. This contradicts

$$T \notin T_3.$$

Thus every minimum vertex cover of G_D is disjoint from $(I(v))$. Fix any domain

$$d \in I(v).$$

Consider the graph

$$G_D - d,$$

obtained from G_D by deleting the vertex d and all edges incident to it.

We claim that $G_D - d$ has no vertex cover of size (2).

Suppose, for contradiction, that

$$S \subseteq V(G_D) \setminus d$$

is a vertex cover of $G_D - d$ with

$$|S| = 2.$$

Then

$$S \cup d$$

is a vertex cover of G_D . Indeed:

- every edge of G_D incident to d is covered by d ;
- every edge not incident to d lies in $G_D - d$ and is covered by S .

Thus $S \cup d$ covers all ordinary vertices.

Moreover, since $d \in I(v)$, the set $S \cup d$ also covers v .

Therefore $S \cup d$ is a set of three domains covering all original vertices.

This contradicts

$$T \notin T_3.$$

Hence $G_D - d$ has no vertex cover of size (2). Since $G_D - d$ is bipartite, König's theorem gives

$$\nu(G_D - d) \geq 3.$$

Thus $G_D - d$ contains a matching of size at least (3).

Let three matched edges be

$$e_{x_1}, e_{x_2}, e_{x_3}.$$

They correspond to three ordinary vertices

$$x_1, x_2, x_3.$$

Because the edges form a matching, the domain-incidence pairs are pairwise disjoint:

$$I(x_i) \cap I(x_j) = \emptyset \quad (i \neq j).$$

The three vertices x_1, x_2, x_3 lie on the two cycles A and B . Therefore two of them lie on the same cycle. Call them

$$p, q.$$

Then

$$I(p) \cap I(q) = \emptyset.$$

This proves the lemma in Case 2.

Since the two cases exhaust all possibilities, the lemma is proved.

□

5.6. Interpretation. Lemma 5.4 is the bridge from the global T_3 -obstruction to a local switching problem.

The failure of T_3 gives two ordinary vertices p, q on the same cycle whose domain-incidence pairs are disjoint:

$$I(p) \cap I(q) = \emptyset.$$

Since p and q lie on the same cycle, there are two paths in that cycle joining them. Along either path, the lateral domains near p must be transformed into the lateral domains near q . The disjointness condition says that this transformation is essential: neither initial domain survives as a final domain.

The rest of the proof studies the finite planar mechanism responsible for this transformation. That mechanism is the switch corridor.

6. SWITCH TRACES AND SWITCH CORRIDORS

In this section we introduce the geometric objects associated with the pair of vertices supplied by Lemma 5.4.

The purpose is not yet to prove regularity. Rather, this section defines the raw switching data and the class of corridors that will be regularized in later sections.

6.1. Setup. Let

$$p, q$$

be ordinary vertices on the same cycle, obtained from Lemma 5.4, such that

$$I(p) \cap I(q) = \emptyset.$$

Without loss of generality suppose

$$p, q \in A.$$

There are two paths in the cycle A from p to q . Let

$$\gamma$$

be one of them.

The drawn image of γ need not be embedded: non-adjacent edges of the same cycle may cross each other in a thrackle drawing. Therefore one must distinguish:

- (1) the parameter path γ in the abstract graph;
- (2) its drawn image in the plane;
- (3) the corresponding walk in the planarized graph $(P(T))$.

This distinction is essential.

The planarized image of γ is a finite walk in $(P(T))$, possibly visiting some crossing vertices more than once.

6.2. Edgelets along γ . Traverse γ from p to q . The planarization subdivides the drawn image of γ into finitely many oriented edgelets:

$$e_1, e_2, \dots, e_m.$$

Each e_i is an open edgelet of $(P(T))$, and all of them lie on original edges of the cycle A .

For each edgelet e_i , there are exactly two domains adjacent to it, one on each side. Let

$$P_i = L_i, R_i$$

be the unordered pair of domains adjacent to e_i .

The order of L_i, R_i is not important at this stage. Later, inside a corridor, oriented transverse sections will give ordered boundary states. For now the lateral data are unordered pairs.

Near the initial vertex p , the lateral pair is

$$P_1 = I(p).$$

Near the terminal vertex q , the lateral pair is

$$P_m = I(q).$$

Since

$$I(p) \cap I(q) = \emptyset,$$

we have

$$P_1 \cap P_m = \emptyset.$$

Thus the sequence

$$P_1, P_2, \dots, P_m$$

cannot be constant. At least one lateral domain changes, and in fact both domains of P_1 must be replaced before reaching P_m .

6.3. Switch events. A switch event along γ is an index

$$i \in 1, \dots, m - 1$$

such that

$$P_i \neq P_{i+1}.$$

Switch events occur at vertices of the planarized graph encountered by the traversal of γ . Such a vertex may be:

- (1) a crossing vertex where the traversed original edge crosses another original edge;
- (2) an original vertex of the cycle A ;
- (3) a repeated crossing vertex of the drawn path, visited at different parameter times.

A switch event is essential if it contributes to changing the initial pair ($I(p)$) into the final pair ($I(q)$). More precisely, consider the reduced sequence obtained from

$$P_1, \dots, P_m$$

by deleting consecutive repetitions. A switch event is essential if it corresponds to a change in this reduced sequence.

The reduced lateral sequence is denoted

$$Q_0, Q_1, \dots, Q_r,$$

where

$$Q_0 = I(p), \quad Q_r = I(q),$$

and

$$Q_j \neq Q_{j+1}$$

for all

$$0 \leq j < r.$$

Since

$$I(p) \cap I(q) = \emptyset,$$

we have

$$r \geq 1.$$

The reduced sequence records the finite switching problem along γ .

6.4. Raw switch trace. The raw switch trace associated with (p, q, γ) is the finite combinatorial object consisting of:

- (1) the parameterized planarized walk determined by γ ;
- (2) the sequence of lateral domain pairs

$$Q_0, \dots, Q_r;$$

- (1) the planarized vertices at which the transitions

$$Q_j \rightarrow Q_{j+1}$$

occur;

- (1) small transverse endpoint sections near p and q , meeting the two lateral domains of $(I(p))$ and $(I(q))$, respectively.

The raw switch trace is not assumed to be embedded and is not assumed to be a disk.

This is important. A path in an abstract cycle may cross itself in the drawing. Consequently, a regular neighborhood of its drawn image may fail to be a disk. The later corridor formalism is designed precisely to extract disk-like relational pieces from this raw trace.

6.5. Geometric corridors. A geometric switch corridor associated with (p, q, γ) is a disk-like PL region

$$C$$

together with two transverse boundary sections

$$\Sigma_-, \quad \Sigma_+$$

such that:

- (1) Σ_- lies near the initial side of a switch subtrace;
- (2) Σ_+ lies near the terminal side of that subtrace;
- (3) the intersection $P(T) \cap C$ is finite and transverse to the sections;
- (4) the components of

$$C \setminus P(T)$$

define a nontrivial boundary relation between states on Σ_- and states on Σ_+ ;

- (1) this boundary relation is part of the transformation of lateral domain data from $(I(p))$ to $(I(q))$.

A geometric switch corridor should be thought of as a disk-like witness to part of the switching process. It need not contain the entire drawn image of γ . Rather, it is extracted from the raw switch trace and records a finite portion of the domain-lane transport.

Later sections prove that minimal such corridors may be treated as relational corridors and then regularized.

6.6. Boundary states of a corridor. Let

$$(C, \Sigma_-, \Sigma_+)$$

be a geometric switch corridor.

The states on Σ_- are the connected components of

$$\Sigma_- \setminus P(T).$$

They form a finite set

$$\mathcal{S}_-.$$

Similarly, the states on Σ_+ form a finite set

$$\mathcal{S}_+.$$

A component of

$$C \setminus P(T)$$

whose closure meets both sections determines a connection between a state on Σ_- and a state on Σ_+ .

Thus C induces a boundary relation

$$R_C \subseteq \mathcal{S}_- \times \mathcal{S}_+.$$

We call C nontrivial if R_C is not the identity relation after the natural ordering of the two sections is fixed.

The relevant corridors in the proof are nontrivial: they record genuine changes in lateral domain data.

6.7. Minimal corridors. Among all nontrivial geometric switch corridors extracted from the raw trace of (p, q, γ) , we choose corridors minimal with respect to the following lexicographic complexity:

$$\kappa(C)(f(C), c(C), e(C), b(C), s(C)),$$

where:

- $f(C)$ is the number of relative domains of $C \setminus P(T)$;
- $c(C)$ is the number of crossing vertices of $(P(T))$ lying in the interior of C ;
- $e(C)$ is the number of edgelets of $P(T) \cap C$;
- $b(C)$ is the combinatorial length of $\partial C \cap P(T)$;
- $s(C)$ is the total number of boundary states on $\Sigma_- \cup \Sigma_+$.

Since $(P(T))$ is finite, there are only finitely many relevant subcomplexes. Hence minimal corridors exist.

A minimal switch corridor is a nontrivial geometric switch corridor minimal with respect to this complexity.

6.8. Relational corridors. A geometric switch corridor is later encoded as a finite sequence of normal quadrilaterals. This motivates the following definition.

A normal quadrilateral is a disk

$$Q$$

whose boundary is decomposed as

$$\partial Q = \Sigma_- \cup U \cup \Sigma_+ \cup V,$$

where:

- Σ_- and Σ_+ are transverse sections;
- U, V are longitudinal boundary arcs;
- $(P(T))$ meets the transverse sections transversely;
- no vertex of $(P(T))$ lies in the interior of either transverse section.

The quadrilateral Q induces a relation

$$R_Q \subseteq \mathcal{S}_- \times \mathcal{S}_+$$

exactly as above: two states are related if they belong to the closure of the same component of

$$Q \setminus P(T).$$

A relational corridor is a finite sequence of normal quadrilaterals

$$\mathcal{C} = (Q_1, \dots, Q_N)$$

with compatible consecutive sections. Its total relation is

$$R_{\mathcal{C}} R_{Q_N} \circ \dots \circ R_{Q_1}.$$

Later, the Geometric Realization Lemma proves that minimal geometric switch corridors can be faithfully represented in this way.

6.9. Reducibility and regularity preview. The later proof distinguishes between reducible corridors and regular corridors.

A corridor is reducible if it contains one of a finite list of configurations that contradict irreducibility, the thrackle condition, minimality, or lead to a smaller corridor. Examples include:

- edge-removal;
- double crossing;
- repeated support;
- pure ear;
- relevant inner domain;
- gate identification;
- terminal lane;
- returning lane;
- branching or merging lane;
- closed alternating train.

A corridor is regular if it decomposes into alternating cells

$$H_{2m}, \quad m \geq 3,$$

glued along gates with unique continuation of lanes.

The next major theorem will state:

every minimal switch corridor is reducible or regular.

The present section only establishes the objects to which that theorem applies.

6.10. Summary of the switch-corridor setup. Starting from

$$T \notin T_3,$$

Lemma 5.4 gives ordinary vertices p, q on the same cycle with

$$I(p) \cap I(q) = \emptyset.$$

Choosing one path γ between p and q gives a finite lateral sequence

$$Q_0, \dots, Q_r$$

with

$$Q_0 = I(p), \quad Q_r = I(q),$$

and

$$Q_0 \cap Q_r = \emptyset.$$

Thus the lateral domains must switch essentially along γ .

The raw switch trace records this finite switching data. From the raw trace one extracts minimal geometric switch corridors, which are disk-like witnesses to nontrivial boundary relations. These corridors are then encoded as relational corridors.

The rest of the proof shows that such corridors cannot survive in an irreducible non- T_3 dumbbell thrackle: they either reduce, or regularize into ladders and trains, and trains are impossible.

7. RELATIONAL CORRIDORS

The preceding section introduced geometric switch corridors as disk-like witnesses to nontrivial transport of domain-lane data. In this section we give the formal relational language used to analyze such corridors.

The central idea is simple: a corridor is cut into finitely many elementary quadrilateral chambers. Each chamber induces a finite relation between boundary states. The whole corridor is then described by the composition of those finite relations.

This section is purely topological and combinatorial. We work in the PL category.

7.1. Ambient notation. Let

$$\Gamma$$

be a finite graph embedded PL in the plane. In the application,

$$\Gamma = P(T)$$

is the planarization of a thrackle drawing.

Let $Q \subset \mathbb{R}^2$ be a closed PL disk. We always assume that $\Gamma \cap Q$ is a finite embedded graph in Q , possibly with vertices on ∂Q , and that all intersections with designated transverse boundary intervals are transverse.

Whenever $S \subset Q$ is a transverse interval, the states on S are the connected components of

$$S \setminus \Gamma.$$

Equivalently, states are the open subintervals of S lying between consecutive intersection points with Γ , together with the two end intervals if the endpoints of S do not lie in Γ .

If

$$S \cap \Gamma = z_1, \dots, z_k$$

with the z_i ordered along S , then $S \setminus \Gamma$ has $(k+1)$ states.

7.2. Normal quadrilaterals. A normal quadrilateral is a tuple

$$Q = (|Q|, \Sigma_-, U, \Sigma_+, V)$$

where $|Q| \subset \mathbb{R}^2$ is a closed PL disk and the boundary is decomposed cyclically as

$$\partial|Q| \Sigma_- \cup U \cup \Sigma_+ \cup V.$$

The four pieces satisfy:

- (1) Σ_- and Σ_+ are closed PL intervals;
- (2) U and V are closed PL arcs, called the longitudinal sides;
- (3) the interiors of the four boundary pieces are pairwise disjoint;
- (4) $\Sigma_- \cap \Sigma_+$ is empty;
- (5) Σ_- and Σ_+ meet Γ transversely;
- (6) no vertex of Γ lies in the interior of $\Sigma_- \cup \Sigma_+$;
- (7) $\Gamma \cap |Q|$ is finite.

We usually suppress the vertical bars and write simply Q for the underlying disk.

The transverse intervals Σ_- and Σ_+ are called the input section and output section. The labels ‘\{\}texttt{input}’ and ‘\{\}texttt{output}’ do not carry orientation data; they only indicate the direction in which relations will later be composed.

The components of

$$\Sigma_- \setminus \Gamma$$

form the finite set of input states

$$\mathcal{S}_-(Q).$$

Similarly, the components of

$$\Sigma_+ \setminus \Gamma$$

form the finite set of output states

$$\mathcal{S}_+(Q).$$

7.3. Lanes in a normal quadrilateral. Let

$$\Gamma_Q = \Gamma \cap Q.$$

A relative domain of Q is a connected component of

$$Q \setminus \Gamma_Q.$$

A relative domain Λ is called a lane if

$$\bar{\Lambda} \cap (\Sigma_- \cup \Sigma_+) \neq \emptyset.$$

A lane is:

- passing if its closure meets both Σ_- and Σ_+ ;
- terminal if its closure meets exactly one of the two sections;
- returning if its closure meets the same section in two or more distinct states.

Thus a passing lane may still be returning if it meets, for example, Σ_- in two distinct states and Σ_+ in at least one state.

The distinction matters because a single connected relative domain could, a priori, touch a transverse section several times. Such a configuration is precisely one of the degeneracies later excluded in reduced quadrilaterals.

7.4. Boundary relation of a quadrilateral. The boundary relation of Q is the relation

$$R_Q \subseteq \mathcal{S}_-(Q) \times \mathcal{S}_+(Q)$$

defined as follows.

For

$$s \in \mathcal{S}_-(Q), \quad t \in \mathcal{S}_+(Q),$$

we declare

$$(s, t) \in R_Q$$

if and only if there exists a relative domain Λ of $Q \setminus \Gamma_Q$ such that

$$s \subset \bar{\Lambda} \cap \Sigma_-,$$

and

$$t \subset \bar{\Lambda} \cap \Sigma_+.$$

Equivalently, s and t are related if a path exists inside

$$Q \setminus \Gamma_Q$$

joining a point of s to a point of t .

This second characterization is often useful.

Lemma 7.1 — Path characterization of the boundary relation. For states $s \in \mathcal{S}_-(Q)$ and $t \in \mathcal{S}_+(Q)$,

$$(s, t) \in R_Q$$

if and only if there is a PL arc

$$\eta \subset Q \setminus \Gamma_Q$$

joining an interior point of s to an interior point of t .

\{\}\textbf{Proof.}

If such an arc η exists, then η is contained in a single connected component of $Q \setminus \Gamma_Q$. Therefore s and t lie in the closure of the same relative domain, and hence $(s, t) \in R_Q$.

Conversely, suppose $(s, t) \in R_Q$. Then s and t lie in the closure of the same relative domain Λ . Since Λ is a component of the complement of a finite PL graph in a disk, it is an open polygonal region. In particular, it is path-connected. Therefore there exists a PL arc in Λ connecting a point of s to a point of t .

□

7.5. Ordered states and identity relations. The longitudinal sides U and V determine compatible orders on the two state sets.

On Σ_- , order the states from the endpoint adjacent to U toward the endpoint adjacent to V .

On Σ_+ , order the states in the same way: from the endpoint adjacent to U toward the endpoint adjacent to V .

Thus both sections are ordered from the U -side to the V -side.

If

$$|\mathcal{S}_-(Q)| = |\mathcal{S}_+(Q)|,$$

there is a canonical order-preserving bijection

$$\iota_Q : \mathcal{S}_-(Q) \longrightarrow \mathcal{S}_+(Q).$$

We call the graph of this bijection the identity relation on Q , and denote it by

$$\text{id}_Q.$$

Thus

$$R_Q = \text{id}_Q$$

means that Q connects the k -th state on Σ_- to the k -th state on Σ_+ , for every k .

This convention is crucial. The two sections are not literally the same set; “identity” always means the canonical order-preserving identification induced by the two longitudinal sides of the quadrilateral.

7.6. Composition of relations. Suppose Q_1 and Q_2 are normal quadrilaterals glued along a common transverse section

$$\Sigma = \Sigma_+(Q_1) = \Sigma_-(Q_2),$$

and suppose the induced states on Σ agree. Then

$$R_{Q_1} \subseteq \mathcal{S}_-(Q_1) \times \mathcal{S}(\Sigma),$$

and

$$R_{Q_2} \subseteq \mathcal{S}(\Sigma) \times \mathcal{S}_+(Q_2).$$

The composed relation is

$$R_{Q_2} \circ R_{Q_1} \subseteq \mathcal{S}_-(Q_1) \times \mathcal{S}_+(Q_2),$$

where

$$(s, u) \in R_{Q_2} \circ R_{Q_1}$$

if and only if there exists a state

$$t \in \mathcal{S}(\Sigma)$$

such that

$$(s, t) \in R_{Q_1}$$

and

$$(t, u) \in R_{Q_2}.$$

This is the ordinary composition of relations.

7.7. Relational corridors. A relational corridor is a finite sequence of normal quadrilaterals

$$\mathcal{C} = (Q_1, Q_2, \dots, Q_N)$$

such that for every $i = 1, \dots, N-1$, the output section of Q_i is identified with the input section of Q_{i+1} , and their state sets agree.

The initial state set of the corridor is

$$\mathcal{S}_-(\mathcal{C}) = \mathcal{S}_-(Q_1).$$

The terminal state set is

$$\mathcal{S}_+(\mathcal{C}) = \mathcal{S}_+(Q_N).$$

The total relation of \mathcal{C} is

$$R_{\mathcal{C}} R_{Q_N} \circ R_{Q_{N-1}} \circ \cdots \circ R_{Q_1}.$$

Thus

$$R_{\mathcal{C}} \subseteq \mathcal{S}_-(\mathcal{C}) \times \mathcal{S}_+(\mathcal{C}).$$

If $N = 1$, then

$$R_{\mathcal{C}} = R_{Q_1}.$$

7.8. Subcorridors. Let

$$\mathcal{C} = (Q_1, \dots, Q_N)$$

be a relational corridor.

For

$$1 \leq i \leq j \leq N,$$

the consecutive block

$$\mathcal{C}[i, j] = (Q_i, Q_{i+1}, \dots, Q_j)$$

is called a subcorridor.

Its total relation is

$$R_{\mathcal{C}[i, j]} R_{Q_j} \circ \cdots \circ R_{Q_i}.$$

It is a proper subcorridor if

$$(i, j) \neq (1, N).$$

We shall repeatedly use the following elementary fact.

Lemma 7.2 — Identity subcorridors are removable. Let

$$\mathcal{C} = (Q_1, \dots, Q_N)$$

be a relational corridor, and suppose

$$\mathcal{C}[i, j]$$

is a proper subcorridor whose total relation is the identity relation between its input and output states:

$$R_{\mathcal{C}[i, j]} = \text{id}.$$

Then deleting the block

$$Q_i, \dots, Q_j$$

does not change the total relation of \mathcal{C} , after identifying the input and output state sets of the deleted block by the canonical order-preserving bijection.

\{\}\textbf{Proof.}

Write

$$A = R_{Q_{i-1}} \circ \cdots \circ R_{Q_1}$$

for the relation before the block, with the convention that A is the identity on the initial states if $i = 1$.

Write

$$B = R_{Q_N} \circ \cdots \circ R_{Q_{j+1}}$$

for the relation after the block, with the convention that B is the identity on the terminal states if $j = N$.

Then

$$R_C B \circ R_{C[i,j]} \circ A.$$

If

$$R_{C[i,j]} = \text{id},$$

then

$$R_C B \circ A,$$

after identifying the state sets at the two ends of the deleted block by the order-preserving identity.

Thus the total relation is unchanged.

□

This deletion changes only the witness corridor, not the underlying drawing. The corridor is an auxiliary object used to encode connectivity of relative domains.

7.9. Minimal relational corridors. A relational corridor \mathcal{C} is called non-trivial if

$$R_C$$

is not the identity relation between its initial and terminal state sets.

Fix a nontrivial relation

$$R_*$$

between two ordered state sets. A relational corridor \mathcal{C} is said to realize R_* if

$$R_C = R_*.$$

Among all corridors realizing R_* , we order by the lexicographic complexity

$$\kappa(\mathcal{C})(N, F, C, E, B, S),$$

where:

- N is the number of quadrilaterals in \mathcal{C} ;
- F is the total number of relative domains in the quadrilaterals;
- C is the total number of crossing vertices contained in the quadrilaterals;
- E is the total number of edgelets of Γ contained in the quadrilaterals;
- B is the total combinatorial length of the longitudinal boundaries;
- S is the total number of states on all internal transverse sections.

A corridor realizing R_* is minimal if it has minimal κ among all corridors realizing R_* .

Because all underlying graphs and subdivisions are finite, this lexicographic minimum exists whenever at least one realizing corridor exists.

7.10. Minimal corridors contain no identity subcorridors.

Lemma 7.3. Let \mathcal{C} be a minimal relational corridor realizing a fixed nontrivial relation R_* . Then no proper subcorridor of \mathcal{C} has identity total relation.

$\backslash\}\text{textbf{Proof.}}$

Suppose that a proper subcorridor

$$\mathcal{C}[i, j]$$

has identity total relation.

By Lemma 7.2, deleting this block does not change the total relation of \mathcal{C} . Therefore the shorter corridor obtained after deleting the block still realizes R_* .

This contradicts the minimality of \mathcal{C} , since the number N of quadrilaterals strictly decreases.

Therefore no proper subcorridor has identity relation.

□

This lemma is the algebraic reason why reduced identity lenses disappear in minimal corridors.

8. GEOMETRIC REALIZATION OF SWITCH CORRIDORS

The preceding section defined relational corridors abstractly. We now prove that every geometric switch corridor admits such a relational representation.

This section is the bridge between the geometry of the planarized drawing and the finite relational formalism.

8.1. Geometric two-terminal corridors. Let

$$\Gamma = P(T)$$

be the planarized graph.

A geometric two-terminal corridor is a triple

$$(C, \Sigma_-, \Sigma_+)$$

such that:

- (1) $C \subset \mathbb{R}^2$ is a closed PL disk;
- (2) $\Sigma_-, \Sigma_+ \subset \partial C$ are disjoint closed PL intervals;
- (3) Σ_- and Σ_+ meet Γ transversely;
- (4) no vertex of Γ lies in the interior of $\Sigma_- \cup \Sigma_+$;
- (5) the remaining two components of

$$\partial C \setminus (\text{int } \Sigma_- \cup \text{int } \Sigma_+)$$

are called the longitudinal sides.

The states on Σ_{\pm} are the components of

$$\Sigma_{\pm} \setminus \Gamma.$$

The corridor induces a boundary relation

$$R_C \subseteq \mathcal{S}_- \times \mathcal{S}_+$$

exactly as in Section 7: two states are related if they lie in the closure of the same component of

$$C \setminus (\Gamma \cap C).$$

8.2. Generic PL height functions. The main tool is slicing the corridor by a generic height function.

Lemma 8.1 — Existence of a generic height. Let

$$(C, \Sigma_-, \Sigma_+)$$

be a geometric two-terminal corridor. There exists a PL homeomorphism

$$\phi : C \longrightarrow [0, 1] \times [0, 1]$$

such that:

$$\phi(\Sigma_-) = 0 \times [0, 1],$$

$$\phi(\Sigma_+) = 1 \times [0, 1],$$

and, for

$$h = \text{pr}_1 \circ \phi,$$

the following hold:

- (1) no edgelet of $\Gamma \cap C$ is contained in a level set of h ;
- (2) each vertex of $\Gamma \cap \text{int } C$ has a distinct h -value;
- (3) no vertex of Γ lies on the transverse levels used for slicing;
- (4) all intersections of Γ with regular levels of h are transverse.

\{\}\textbf{Proof.}

Choose any PL homeomorphism

$$\phi_0 : C \rightarrow [0, 1] \times [0, 1]$$

sending Σ_- and Σ_+ to the left and right vertical sides.

Because $\Gamma \cap C$ is finite, there are finitely many edgelets and vertices to control. A sufficiently small PL perturbation of ϕ_0 , fixed on $\Sigma_- \cup \Sigma_+$, makes the first coordinate function nonconstant on every edgelet and gives distinct first-coordinate values to all interior vertices.

The conditions are open and dense in the finite-dimensional space of PL embeddings compatible with the prescribed boundary data. Since only finitely many constraints are involved, such a perturbation exists.

Set

$$h = \text{pr}_1 \circ \phi.$$

Then h has the required properties.

□

8.3. Slicing into normal quadrilaterals. Let h be a generic height function as above.

Let

$$0 = t_0 < t_1 < \cdots < t_N = 1$$

be a sequence satisfying:

- (1) each t_i is a regular value of $h|_{\Gamma \cap C}$;
- (2) every interior vertex of $\Gamma \cap C$ has height lying in exactly one interval

$$(t_{i-1}, t_i);$$

- (1) each interval (t_{i-1}, t_i) contains at most one vertex-height.

Define

$$Q_i = h^{-1}([t_{i-1}, t_i]).$$

Each Q_i is a closed PL disk. Its input and output sections are

$$h^{-1}(t_{i-1}), \quad h^{-1}(t_i),$$

and its longitudinal sides are the corresponding subarcs of the longitudinal sides of C .

Because the chosen levels are regular, Q_i is a normal quadrilateral.

Thus we obtain a finite sequence

$$\mathcal{C}_h = (Q_1, \dots, Q_N).$$

This is the relational slicing of the geometric corridor with respect to h .

8.4. Equality between geometric and relational boundary relations.

Proposition 8.2. Let

$$(C, \Sigma_-, \Sigma_+)$$

be a geometric two-terminal corridor, and let

$$\mathcal{C}_h = (Q_1, \dots, Q_N)$$

be a relational slicing by a generic height function. Then

$$R_C = R_{Q_N} \circ \dots \circ R_{Q_1}.$$

\{\}\textbf{Proof.}

Let

$$\mathcal{S}_i$$

denote the set of states on the level

$$h^{-1}(t_i).$$

Thus

$$\mathcal{S}_0 = \mathcal{S}_-, \quad \mathcal{S}_N = \mathcal{S}_+.$$

We prove both inclusions.

First suppose

$$(s_0, s_N) \in R_C.$$

Then s_0 and s_N lie in the closure of the same component of

$$C \setminus \Gamma.$$

Equivalently, by Lemma 7.1, there is a PL arc

$$\eta \subset C \setminus \Gamma$$

joining an interior point of s_0 to an interior point of s_N .

After a small perturbation of η , fixing its endpoints and keeping it in $C \setminus \Gamma$, we may assume that η is transverse to all intermediate levels

$$h^{-1}(t_i).$$

For each i , the intersection of η with $h^{-1}(t_i)$ may consist of several points. Choose one point along η at the first time η crosses from the portion $h \leq t_i$ to the portion $h \geq t_i$. This point lies in a unique state

$$s_i \in \mathcal{S}_i.$$

The subarc of η between the chosen point on $h^{-1}(t_{i-1})$ and the chosen point on $h^{-1}(t_i)$ lies in $Q_i \setminus \Gamma$. Therefore

$$(s_{i-1}, s_i) \in R_{Q_i}.$$

Thus

$$(s_0, s_N) \in R_{Q_N} \circ \cdots \circ R_{Q_1}.$$

Conversely, suppose

$$(s_0, s_N) \in R_{Q_N} \circ \cdots \circ R_{Q_1}.$$

Then there exist intermediate states

$$s_i \in \mathcal{S}_i$$

for

$$i = 1, \dots, N - 1$$

such that

$$(s_{i-1}, s_i) \in R_{Q_i}$$

for every i .

By the path characterization of boundary relations, for every i there exists an arc

$$\eta_i \subset Q_i \setminus \Gamma$$

joining a point of s_{i-1} to a point of s_i .

The arcs η_i may meet different points of the same intermediate state s_i . Since s_i is a connected interval contained in

$$h^{-1}(t_i) \setminus \Gamma,$$

we may connect those endpoints inside s_i .

Concatenating the arcs η_i and these short connecting intervals inside intermediate states gives a PL arc in

$$C \setminus \Gamma$$

joining s_0 to s_N .

Hence

$$(s_0, s_N) \in R_C.$$

Therefore

$$R_C = R_{Q_N} \circ \cdots \circ R_{Q_1}.$$

□

8.5. Subdisks and subcorridors. The next point is crucial for the induction used later. Any internal subdisk that arises from a local degeneracy is contained in a proper consecutive block of the relational slicing.

Lemma 8.3 — *Internal subdisks are contained in proper subcorridors.* Let

$$D \subset \text{int } C$$

be a compact PL subdisk. Then, after possibly refining the slicing by adding two regular levels, there exist indices

$$1 \leq i \leq j \leq N$$

such that

$$D \subset Q_i \cup Q_{i+1} \cup \cdots \cup Q_j,$$

and

$$(Q_i, \dots, Q_j)$$

is a proper subcorridor of the slicing.

\{\}\text{Proof.}

Because $D \subset \text{int } C$ is compact and h is continuous,

$$h(D)$$

is a compact interval

$$[a, b] \subset (0, 1).$$

Choose regular values

$$\alpha, \beta$$

such that

$$0 < \alpha < a \leq b < \beta < 1.$$

Refine the slicing by inserting α and β among the slicing levels. Since α and β are regular values, the refined slices are again normal quadrilaterals.

Let

$$C[\alpha, \beta] = h^{-1}([\alpha, \beta]).$$

This is a consecutive block of the refined slicing, and it contains D .

Since

$$0 < \alpha < \beta < 1,$$

this block is proper: it does not contain the initial section Σ_- or the terminal section Σ_+ of the whole corridor.

Thus D is contained in a proper subcorridor.

□

8.6. Identity subcorridors and deletion.

Proposition 8.4. Let

$$\mathcal{C}_h = (Q_1, \dots, Q_N)$$

be a relational slicing of a geometric corridor. Let

$$\mathcal{C}_h[i, j] = (Q_i, \dots, Q_j)$$

be a proper subcorridor.

If

$$R_{\mathcal{C}_h[i, j]} = \text{id},$$

then deleting this subcorridor does not change the total boundary relation of the geometric corridor.

\{\}\textbf{Proof.}

This is exactly Lemma 7.2 applied to the relational slicing. The total relation is

$$R_{\mathcal{C}_h} R_{Q_N} \circ \dots \circ R_{Q_{j+1}} \circ R_{\mathcal{C}_h[i, j]} \circ R_{Q_{i-1}} \circ \dots \circ R_{Q_1}.$$

If

$$R_{\mathcal{C}_h[i, j]} = \text{id},$$

then the middle factor may be removed without changing the composite relation.

Since Proposition 8.2 identifies

$$R_{\mathcal{C}_h}$$

with the geometric relation R_C , the deletion leaves the geometric boundary relation unchanged.

Again, this operation deletes a portion of the witness corridor, not a portion of the original drawing.

□

8.7. Non-identity subcorridors.

Proposition 8.5. Let

$$D \subset \text{int } C$$

be a subdisk arising from a local degeneracy. If every proper subcorridor containing D has non-identity boundary relation, then D is contained in a proper nontrivial relational corridor of smaller complexity.

\{\}\text{Proof.}

By Lemma 8.3, after refinement D is contained in a proper subcorridor

$$\mathcal{C}_h[i, j].$$

By hypothesis,

$$R_{\mathcal{C}_h[i, j]} \neq \text{id}.$$

Thus $\mathcal{C}_h[i, j]$ is a nontrivial relational corridor.

It is proper because it is contained strictly between two interior regular levels

$$\alpha, \beta$$

with

$$0 < \alpha < \beta < 1.$$

Its complexity is strictly smaller than the complexity of the whole corridor: it contains fewer slices after refinement, and in particular it does not contain the portions of the corridor between Σ_- and $h^{-1}(\alpha)$, nor the portions between $h^{-1}(\beta)$ and Σ_+ .

Therefore it is a proper smaller nontrivial relational corridor.

□

This is the precise form needed in the regularization induction. A subdisk need not itself be a new corridor between original vertices. It is enough that it is contained in a smaller nontrivial relational corridor.

8.8. Geometric Realization Lemma. We now collect the preceding results.

Theorem 8.6 — Geometric Realization Lemma. Every geometric two-terminal switch corridor admits a relational slicing

$$\mathcal{C}_h = (Q_1, \dots, Q_N)$$

such that:

- (1) each Q_i is a normal quadrilateral;
- (2) the geometric boundary relation equals the composed relational boundary relation:

$$R_C = R_{Q_N} \circ \cdots \circ R_{Q_1};$$

- (1) every compact internal subdisk is contained, after refinement if necessary, in a proper relational subcorridor;
- (2) an identity subcorridor may be removed from the witness without changing the total boundary relation;
- (3) a non-identity subcorridor is a proper smaller relational corridor.

\{\}\textbf{Proof.}

Parts (1) and (2) are Proposition 8.2.

Part (3) is Lemma 8.3.

Part (4) is Proposition 8.4.

Part (5) is Proposition 8.5.

□

This theorem is the bridge from geometric switch corridors to the relational formalism.

9. REDUCED QUADRILATERALS AND BOUNDARY RIGIDITY

We now prove the central local rigidity statement.

Informally, a quadrilateral that contains no reducible local feature must transmit its boundary states without permutation. The proof is purely planar: its lanes are disjoint arcs in a disk, and disjoint arcs preserve order.

9.1. Lanes and degeneracies. Let Q be a normal quadrilateral.

Recall that a lane is a connected component of

$$Q \setminus \Gamma_Q$$

whose closure meets at least one of the two transverse sections.

We now define the lane degeneracies that will be excluded.

A lane Λ is terminal if its closure meets exactly one transverse section.

A lane Λ is returning if its closure meets the same transverse section in more than one state.

A collection of lanes has a branching or merging degeneracy if, at some point of the planarized graph inside Q , the local continuation of relative-domain components is not one-to-one. Equivalently, in any sufficiently small neighborhood of that point, one incoming lane-germ is adjacent to two outgoing lane-germs, or two incoming lane-germs are adjacent to one outgoing lane-germ.

A gate identification occurs when two distinct boundary intervals through which lanes are intended to continue represent the same topological interval of the planarized graph. In the quadrilateral setting, this means that two combinatorially distinct lane-continuation intervals are supported by the same interval of Γ_Q .

These definitions are local. Their significance is that each of them gives rise, in the full corridor, either to a reduction or to a smaller subcorridor. In the present section we simply exclude them in the definition of reducedness.

9.2. Reduced normal quadrilaterals. A normal quadrilateral Q is called reduced if none of the following configurations occurs inside Q :

- (1) an edge-removal triangle empty of original vertices;
- (2) a double crossing of two original supports;
- (3) an essential repeated support;
- (4) a pure ear;
- (5) a relevant inner domain or smaller T_3 -critical core;
- (6) a proper non-identity relational subcorridor;
- (7) a gate identification;
- (8) a terminal lane;
- (9) a returning lane;
- (10) branching or merging of lanes;
- (11) a closed alternating train.

For the boundary-rigidity theorem below, only items 7–10 are strictly needed. The other items are included because reduced quadrilaterals will later arise inside the global regularization argument.

9.3. Every boundary state belongs to a passing lane.

Lemma 9.1. Let Q be a reduced normal quadrilateral. Then every state on $\Sigma_- \cup \Sigma_+$ belongs to a passing lane.

\{\}\textbf{Proof.}

Let

s

be a state on Σ_- . Let Λ be the relative domain of $Q \setminus \Gamma_Q$ whose closure contains s .

If Λ does not meet Σ_+ , then Λ is a terminal lane. This is excluded by reducedness.

Therefore Λ meets Σ_+ , and hence Λ is passing.

The same argument applies to a state on Σ_+ .

□

9.4. A passing lane meets each section in exactly one state.

Lemma 9.2. Let Q be reduced. A passing lane Λ meets Σ_- in exactly one state and Σ_+ in exactly one state.

\{\}\textbf{Proof.}

Since Λ is passing, it meets both sections in at least one state.

Suppose Λ meets Σ_- in two distinct states. Then Λ is a returning lane. This is excluded by reducedness.

Therefore Λ meets Σ_- in at most one state. Since it meets Σ_- at least once, it meets it exactly once.

The same argument applies to Σ_+ .

□

9.5. The boundary relation is a bijection.

Lemma 9.3. Let Q be reduced. Then

$$R_Q$$

is the graph of a bijection

$$\pi_Q : \mathcal{S}_-(Q) \longrightarrow \mathcal{S}_+(Q).$$

\{\}\textbf{Proof.}

Let

$$s \in \mathcal{S}_-(Q).$$

By Lemma 9.1, s belongs to a passing lane Λ . By Lemma 9.2, Λ meets Σ_+ in exactly one state, call it t .

Thus s is related to exactly one state $t \in \mathcal{S}_+(Q)$.

Conversely, let

$$t \in \mathcal{S}_+(Q).$$

Again by Lemma 9.1, t belongs to a passing lane, and by Lemma 9.2 that lane meets Σ_- in exactly one state.

Therefore every output state has exactly one input preimage.

It follows that R_Q is the graph of a bijection

$$\pi_Q : \mathcal{S}_-(Q) \rightarrow \mathcal{S}_+(Q).$$

□

9.6. Core arcs inside lanes.

Lemma 9.4. Let Q be reduced. For every passing lane Λ , there exists a properly embedded PL arc

$$\lambda_\Lambda \subset \Lambda$$

joining the unique state of Λ on Σ_- to the unique state of Λ on Σ_+ .

Moreover, if Λ_1 and Λ_2 are distinct passing lanes, the arcs λ_{Λ_1} and λ_{Λ_2} may be chosen disjoint.

\{\}\textbf{Proof.}

A lane Λ is a connected component of

$$Q \setminus \Gamma_Q.$$

Since Q is a PL disk and Γ_Q is a finite PL graph, every component of $Q \setminus \Gamma_Q$ is a polygonal surface region. In particular, it is path-connected and locally path-connected.

By Lemma 9.2, $\bar{\Lambda}$ meets Σ_- in exactly one state s_- and Σ_+ in exactly one state s_+ .

Choose points

$$x_- \in s_-, \quad x_+ \in s_+.$$

Since Λ is path-connected, there is a PL path in Λ from a point arbitrarily close to x_- to a point arbitrarily close to x_+ . By adjoining tiny subarcs inside s_- and s_+ , we obtain a PL arc from x_- to x_+ whose interior lies in Λ .

If the arc self-intersects, remove loops in the usual way to obtain a simple PL arc. Because Λ is open and locally planar, this loop-removal can be done inside Λ .

This gives a properly embedded arc

$$\lambda_\Lambda.$$

If $\Lambda_1 \neq \Lambda_2$, then Λ_1 and Λ_2 are distinct connected components of $Q \setminus \Gamma_Q$. They are disjoint. Therefore the chosen arcs lie in disjoint sets and are disjoint.

□

9.7. Noncrossing arcs preserve order.

Lemma 9.5. Let D be a closed disk with two disjoint boundary intervals

$$\Sigma_-, \quad \Sigma_+.$$

Let

$$\lambda_1, \dots, \lambda_n$$

be pairwise disjoint properly embedded arcs in D , each joining a point of Σ_- to a point of Σ_+ .

Order the endpoints on both Σ_- and Σ_+ from the same longitudinal side of D to the other.

Then the induced bijection between endpoints is order-preserving.

{\textbf{Proof.}}

Suppose the induced bijection is not order-preserving. Then there exist two arcs

$$\lambda_i, \lambda_j$$

such that their endpoints occur in one order on Σ_- and in the opposite order on Σ_+ .

Equivalently, the four endpoints are alternating on the boundary of the disk.

Consider λ_i . Together with one of the two boundary arcs joining its endpoints, it forms a simple closed curve. By the Jordan curve theorem, this curve separates the disk into two components.

Because the endpoints of λ_j are alternating with those of λ_i , the two endpoints of λ_j lie in different components of the complement of that simple closed curve.

Therefore λ_j must intersect λ_i .

This contradicts the assumption that the arcs are pairwise disjoint.

Thus no inversion is possible, and the induced bijection is order-preserving.

□

9.8. Boundary-rigidity theorem.

Theorem 9.6 — Boundary rigidity of reduced quadrilaterals. If Q is a reduced normal quadrilateral, then

$$R_Q = \text{id}_Q .$$

\{\}\textbf{Proof.}

By Lemma 9.3, the boundary relation R_Q is the graph of a bijection

$$\pi_Q : \mathcal{S}_-(Q) \rightarrow \mathcal{S}_+(Q).$$

By Lemma 9.4, for every passing lane we may choose a properly embedded core arc, and the chosen arcs are pairwise disjoint.

These arcs realize the bijection π_Q .

By Lemma 9.5, the bijection induced by pairwise disjoint arcs in a disk, with both boundary sections ordered from the same longitudinal side to the other, is order-preserving.

Therefore π_Q is order-preserving.

Since $\mathcal{S}_-(Q)$ and $\mathcal{S}_+(Q)$ are finite linearly ordered sets of the same cardinality, the only order-preserving bijection between them is the canonical identity bijection.

Hence

$$R_Q = \text{id}_Q .$$

□

9.9. Consequence for reduced lenses. The main application of Theorem 9.6 will be to monochromatic lenses in later sections.

If such a lens is reduced, then it induces identity boundary relation. In a minimal relational corridor, an identity subcorridor is removable. Therefore a reduced monochromatic lens cannot survive in a minimal corridor.

This is the local rigidity mechanism that replaces the invalid principle “locally T_3 -coverable implies globally neutral.”

The correct principle is instead:

reduced disk \Rightarrow disjoint lanes \Rightarrow order preservation \Rightarrow identity relation.

This is purely topological and does not rely on local T_3 -coverability.

10. MONOCHROMATIC LENSES

This section treats the only local configuration in which a compensation chamber may fail to alternate between the two cycles A and B . The key result is that a monochromatic lens cannot survive inside a minimal non-reducible corridor. If it is reduced, it has identity boundary relation by the boundary-rigidity theorem of Section 9; if it is not reduced, then it already contains one of the allowed reductions.

10.1. Colored support arcs. Let Q be a disk-like chamber inside a switch corridor. Its boundary is a finite concatenation of two types of arcs:

- (1) support arcs, which are subarcs of original edge supports of the drawing;
- (2) gate arcs, which are transverse intervals through which lanes enter or leave the chamber.

Every support arc inherits a color

$$A \text{ or } B$$

according to whether the original edge containing it belongs to the cycle A or the cycle B .

A support arc is always taken maximal with respect to the property of lying on a single original support and containing no gate endpoint in its interior.

Thus, when we traverse ∂Q , the support arcs appear in a cyclic sequence

$$c_1, c_2, \dots, c_k,$$

where each

$$c_i \in A, B.$$

This cyclic sequence is called the colored boundary word of Q , and is denoted

$$\omega(Q).$$

Gate arcs are not recorded in $\omega(Q)$; they separate consecutive support arcs.

10.2. Alternating and non-alternating chambers. A chamber Q is called alternating if its colored boundary word alternates between A and B :

$$A, B, A, B, \dots$$

or

$$B, A, B, A, \dots$$

Thus no two consecutive entries of $\omega(Q)$ are equal.

A chamber is non-alternating if its colored boundary word contains two consecutive equal entries. Equivalently, there exists a pair of consecutive support arcs on ∂Q with the same color:

$$A, A$$

or

$$B, B.$$

A monochromatic run is a maximal consecutive block in $\omega(Q)$ of the form

$$A, A, \dots, A$$

or

$$B, B, \dots, B.$$

A monochromatic run is nontrivial if it has length at least (2).

10.3. Monochromatic lenses. Let Q be a chamber whose boundary contains a nontrivial monochromatic run. Let

$$\rho$$

be such a maximal run.

The run ρ is represented geometrically by a connected boundary arc

$$U_\rho \subset \partial Q$$

which is a concatenation of support arcs of the same color, separated by gate endpoints or crossing vertices.

Let the endpoints of U_ρ be

$$x, y.$$

The complementary boundary arc of ∂Q from y back to x is denoted

$$V_\rho.$$

Since Q is a disk, the two arcs U_ρ and V_ρ form the full boundary of a disk. We call this disk the monochromatic lens associated with ρ and denote it by

$$L_\rho.$$

In the simplest case $L_\rho = Q$. More generally, if a chamber has been refined by transverse cuts, L_ρ is treated as a normal quadrilateral whose input and output sections are small transverse intervals near x and y .

The two longitudinal sides of L_ρ are:

- (1) the monochromatic side U_ρ ;
- (2) the complementary side V_ρ .

The input and output sections are chosen in small neighborhoods of x and y , transverse to the local lanes crossing the lens.

Once these sections are chosen, L_ρ has boundary states and induces a boundary relation exactly as in Section 7.

10.4. Lens refinement. The following observation makes precise why a monochromatic lens may be treated as a factor in a relational corridor.

Lemma 10.1 — Lens refinement. Let L_ρ be a monochromatic lens contained in a geometric switch corridor C . After refining the relational slicing of C , the lens L_ρ is contained in a consecutive proper subcorridor. If L_ρ carries a nontrivial boundary relation, then this subcorridor is nontrivial. If L_ρ carries the identity boundary relation, then this subcorridor is relationally removable.

\{\}\textbf{Proof.}

The lens L_ρ is a compact subdisk of the interior of the switch corridor, except possibly along portions of chamber boundary. Choose small transverse sections near the two endpoints x, y of the monochromatic run so that the lens lies between them.

By the Geometric Realization Lemma, after adding the corresponding transverse levels to the slicing, the region containing L_ρ is contained in a consecutive block

$$Q_i \cup Q_{i+1} \cup \cdots \cup Q_j.$$

This block is a proper subcorridor unless the entire corridor is itself exactly the lens. In that exceptional case the same argument applies to the corridor itself.

If the relation of the block is identity, then by the identity-subcorridor deletion lemma it may be removed from the witness corridor without changing the total relation.

If the relation is non-identity, then the block is a nontrivial relational subcorridor. Since it is strictly smaller than the original corridor, it is available for induction on corridor complexity.

□

10.5. Reduced monochromatic lenses. A monochromatic lens L is called reduced if, as a normal quadrilateral, it is reduced in the sense of Section 9. That is, it contains none of the following:

- (1) edge-removal;
- (2) double crossing;
- (3) repeated support;
- (4) pure ear;
- (5) relevant inner domain or smaller T_3 -critical core;
- (6) proper non-identity relational subcorridor;
- (7) gate identification;
- (8) terminal lane;
- (9) returning lane;
- (10) branching or merging lane;
- (11) closed alternating train.

The boundary-rigidity theorem applies directly to such a lens.

Lemma 10.2 — *Reduced monochromatic lenses are identity.* If L is a reduced monochromatic lens, then

$$R_L = \text{id}.$$

\{\}\textbf{Proof.}

A monochromatic lens is a normal quadrilateral with two transverse sections and two longitudinal sides. Since L is reduced, Theorem 9.6 applies.

Therefore its boundary relation is the identity relation:

$$R_L = \text{id}.$$

□

10.6. Monochromatic lenses in minimal corridors. We now prove the key lemma of this section.

Lemma 10.3 — *Monochromatic lenses are reducible in minimal corridors.* Let \mathcal{C} be a minimal nontrivial relational corridor. Then every monochromatic lens contained in \mathcal{C} is reducible.

\{\}\textbf{Proof.}

Let

L

be a monochromatic lens contained in \mathcal{C} .

If L is not reduced, then by definition it contains one of the allowed reducible configurations. Therefore L is reducible.

Suppose instead that L is reduced. By Lemma 10.2,

$$R_L = \text{id}.$$

By Lemma 10.1, after refining the slicing if necessary, L is contained in a proper subcorridor with identity boundary relation. By the identity deletion lemma, this subcorridor can be removed without changing the total boundary relation of \mathcal{C} .

This contradicts the minimality of \mathcal{C} .

Therefore L cannot be reduced. Hence every monochromatic lens is reducible.

□

10.7. Consequence for non-alternating chambers.

Corollary 10.4. Let \mathcal{C} be a minimal nontrivial relational corridor containing no reducible configuration. Then every compensation chamber of \mathcal{C} is alternating.

\{\}\textbf{Proof.}

Suppose a compensation chamber Q is not alternating. Then its colored boundary word contains a nontrivial monochromatic run. This run determines a monochromatic lens L .

By Lemma 10.3, L is reducible. Hence \mathcal{C} contains a reducible configuration, contradicting the hypothesis.

Thus no compensation chamber is non-alternating.

Therefore every compensation chamber is alternating.

□

11. REGULARIZATION OF SWITCH CORRIDORS

In this section we prove the switch-corridor regularization theorem. The theorem says that a minimal switch corridor either contains a reducible configuration or else has a regular alternating-cell structure.

This is the structural heart of the proof.

11.1. Reducible configurations. A relational switch corridor is called reducible if it contains at least one of the following configurations:

- (1) edge-removal: an empty edge-removal triangle for a 3-path;
- (2) double crossing: two original supports meet more than once;
- (3) repeated support: an original support appears twice in an essential boundary or lane itinerary;
- (4) pure ear: a disk bounded by portions of the corridor and the planarized graph, containing no original vertices and no essential boundary relation;
- (5) relevant inner domain: an internal disk containing original vertices and producing either a smaller T_3 -critical core or a smaller relational corridor;
- (6) proper non-identity subcorridor: a proper subcorridor with non-identity boundary relation;

- (7) gate identification: two distinct gates represent the same topological boundary interval;
- (8) terminal lane: a lane meets exactly one transverse section and is not eliminable as a pure ear;
- (9) returning lane: a lane meets the same transverse section in more than one state;
- (10) branching or merging lane: the continuation of lanes is not locally one-to-one;
- (11) closed alternating train: a cyclic component of the gate graph;
- (12) proper subtrain: a smaller closed train contained in a train.

The reason for including closed trains among the reducible outcomes is that Section 14 later proves such trains impossible.

11.2. Compensation chambers. A compensation chamber is a minimal disk-like region in a corridor bounded by support arcs and gate arcs that realizes one elementary compensation of domain lanes.

Equivalently, it is a minimal chamber in which one domain lane that disappears from the lateral boundary is reconnected to the lane data downstream.

For the regularization theorem, the exact origin of the chamber is less important than the following formal properties:

- (1) its boundary is a finite cyclic concatenation of support arcs and gate arcs;
- (2) each support arc has color A or B ;
- (3) it induces a local boundary relation on the gate states incident to it;
- (4) it is minimal with respect to inclusion among chambers realizing that local compensation.

A chamber is embedded if its boundary is a simple closed curve in the planarized disk.

11.3. Closed-walk innermost lemma. We shall repeatedly use the following planar fact.

Lemma 11.1 — Innermost closed-walk reduction. Let W be a closed walk in the planarized graph inside a switch corridor. Then one of the following occurs:

- (1) W contains an embedded simple cycle Z ;
- (2) W contains a repeated support;
- (3) W contains a double crossing;
- (4) W contains a gate identification.

If Z is chosen innermost, the disk D_Z bounded by Z is one of:

- (1) a pure ear;
- (2) a relevant inner domain;
- (3) a proper identity subcorridor;
- (4) a proper non-identity subcorridor;

(5) a closed alternating train or proper subtrain.

\{\}\textbf{Proof.}

Every closed walk in a finite plane graph either contains a simple cycle or closes by repeating combinatorial data. In the present planarized setting, a nontrivial repetition of combinatorial data manifests as one of:

- repetition of an original support;
- repeated intersection of two original supports, i.e. double crossing;
- identification of gate intervals.

If a simple cycle Z exists, choose it so that the disk D_Z it bounds contains no smaller simple cycle produced by the same closed-walk construction.

If D_Z contains no original vertices and carries no nontrivial relation, it is a pure ear.

If D_Z contains original vertices, then either those vertices form a smaller T_3 -critical core, or the disk carries boundary relation data and is a proper subcorridor. If the subcorridor has identity relation, it is an identity subcorridor; if not, it is a non-identity subcorridor.

If the structure inside D_Z is a cyclic chain of alternating cells glued by gates, it is a closed alternating train or a proper subtrain.

This exhausts all possibilities.

□

11.4. Short alternating chambers. An alternating chamber may have colored boundary word of length (2), (4), or at least (6).

The following lemma eliminates the short cases.

Lemma 11.2 — Short alternating chambers are reducible. Let Q be an embedded alternating compensation chamber.

If its colored boundary word has length (2) or (4), then Q is reducible.

\{\}\textbf{Proof.}

If the colored boundary word has length (2), then the boundary has one A -arc and one B -arc. These two support arcs form a bigon. A bigon between two original supports either contains no original vertex, in which case it is a pure ear, or it forces the same two supports to meet twice, in which case it is a double crossing. Hence the chamber is reducible.

Now suppose the colored boundary word has length (4):

$$A, B, A, B$$

or

$$B, A, B, A.$$

If the chamber contains no original vertices and carries no nontrivial boundary relation, it is an identity strip or pure ear and is eliminable.

If it contains original vertices, then it contains a relevant inner domain or smaller T_3 -critical core.

If it carries nontrivial boundary relation, then it is a proper non-identity subcorridor.

If its boundary is not embedded, we apply Lemma 11.1 and obtain a reducible configuration.

Therefore in all cases a length-(4) alternating chamber is reducible.

□

11.5. Alternating cells. A non-reducible alternating chamber must therefore have colored boundary word of even length at least (6).

We denote such a chamber by

$$H_{2m}, \quad m \geq 3.$$

Its boundary word is

$$A, B, A, B, \dots, A, B$$

or the reverse.

More explicitly, after choosing a starting point, the boundary can be written as

$$a_0 b_0 a_1 b_1 \cdots a_{m-1} b_{m-1},$$

where each a_i is an A -support arc and each b_i is a B -support arc.

Such a chamber is called an alternating cell.

11.6. Gate regularity. A gate is a boundary interval through which a lane passes from one chamber to an adjacent chamber, or from a chamber to an endpoint section of the corridor.

A gate is internal if it is not contained in an endpoint section of the whole corridor.

Gate continuation is unique if every internal gate has exactly one outgoing lane-continuation on each side of the gate.

Lemma 11.3 — Gate continuation is unique in a non-reducible corridor. In a minimal non-reducible corridor, every internal gate has unique continuation.

\{\}\textbf{Proof.}

If an internal gate has no continuation on one side, then the lane terminating at the gate is a terminal lane. This is reducible.

If an internal gate has two continuations on one side, then a lane branches. This is reducible.

If two distinct internal gates share the same continuation interval, then the corridor contains a gate identification. This is reducible.

Therefore none of these possibilities occurs in a non-reducible corridor. Hence every internal gate has unique continuation.

□

11.7. Cell intersections.

Lemma 11.4 — Alternating cells meet only along gates. In a minimal non-reducible corridor, two distinct alternating cells are either disjoint or meet along a single gate. No three alternating cells share a gate.

\{\}\textbf{Proof.}

If two cells overlap in a two-dimensional region, the boundary of the overlap contains a closed walk. By Lemma 11.1, this yields a reducible configuration: pure ear, relevant inner domain, proper subcorridor, repeated support, double crossing, or gate identification.

If two cells meet along two distinct gates, then the portion of one cell boundary between the two gates and the corresponding portion of the other cell boundary form a closed walk. Applying Lemma 11.1 again gives a reducible configuration.

If three cells share the same gate, then the local continuation of lanes through that gate is not one-to-one: at least one lane branches or two lanes merge. This is reducible.

Therefore distinct cells meet only along a single gate, and no triple gate-sharing occurs.

□

11.8. Definition of regular corridor. A switch corridor is regular if it satisfies all of the following:

- (1) every compensation chamber is an alternating cell H_{2m} with $m \geq 3$;
- (2) distinct cells meet only along gates;
- (3) no three cells share a gate;
- (4) every internal gate has unique continuation;
- (5) the only terminal gates are on the two endpoint sections of the corridor.

11.9. Switch-corridor regularization theorem.

Theorem 11.5 — Regularization. Every minimal switch corridor is either reducible or regular.

\{\}\textbf{Proof.}

Let \mathcal{C} be a minimal switch corridor.

Assume \mathcal{C} is not reducible. We prove that it is regular.

First, no compensation chamber can be non-embedded. If a chamber boundary is a non-simple closed walk, Lemma 11.1 produces a reducible configuration or a proper subcorridor. A proper non-identity subcorridor is reducible by definition; an identity subcorridor contradicts minimality because it can be removed. Hence non-embedded chambers cannot occur.

Second, no compensation chamber can be non-alternating. If a chamber is non-alternating, it contains a monochromatic lens. By Corollary 10.4,

such a lens is reducible in a minimal nontrivial corridor. This contradicts the assumption that \mathcal{C} is not reducible.

Thus every compensation chamber is alternating.

Third, no alternating chamber can have boundary length (2) or (4), by Lemma 11.2. Hence every chamber is an alternating cell

$$H_{2m}, \quad m \geq 3.$$

Fourth, cell intersections are regular by Lemma 11.4.

Fifth, gate continuation is unique by Lemma 11.3.

Therefore all defining properties of a regular corridor are satisfied.

Hence \mathcal{C} is regular.

We have shown that every minimal switch corridor is reducible or regular.

□

12. REGULAR CORRIDORS, LADDERS, AND TRAINS

Once a corridor is regular, its structure becomes purely graph-theoretic. The key object is the gate graph.

12.1. The gate graph. Let \mathcal{C} be a regular corridor.

The gate graph

$$\mathfrak{G}(\mathcal{C})$$

is defined as follows.

The vertices of $\mathfrak{G}(\mathcal{C})$ are the gates of \mathcal{C} . There is an edge between two gates if a lane-continuation connects them through an alternating cell or through a strip of the corridor.

A gate lying on an endpoint section of the corridor is called an external gate. All other gates are internal gates.

Because \mathcal{C} is regular, every internal gate has unique continuation. Therefore every internal vertex of $\mathfrak{G}(\mathcal{C})$ has degree (2).

External gates may have degree (1).

12.2. Essential components. A connected component of the gate graph is called essential if it participates in the nontrivial boundary relation of the corridor. Equivalently, it carries at least one lane involved in transporting input states to output states.

Since the corridor is nontrivial, at least one essential component exists.

A component consisting entirely of internal gates and internal continuations is necessarily a cycle, because all its vertices have degree (2).

A component meeting the external gate set is a path, unless it also contains a cycle. If it contains a cycle, that cycle is already a closed train.

12.3. Closed alternating trains. A closed alternating train is a cyclic sequence of alternating cells

$$H_{2m_1}, H_{2m_2}, \dots, H_{2m_r}$$

glued successively along gates, with lane continuations forming a cyclic component of the gate graph.

Thus a cycle component of

$$\mathfrak{G}(\mathcal{C})$$

is precisely a closed alternating train.

12.4. Local ladders. A local ladder is a path component of the gate graph connecting an external gate on the input section of the corridor to an external gate on the output section.

Equivalently, a local ladder is a finite chain of alternating cells through which a lane travels from the input side of the corridor to the output side.

A path component connecting two external gates on the same endpoint section is not a local ladder. Such a component is reducible, as shown next.

12.5. Same-side path components are reducible.

Lemma 12.1. Let \mathcal{C} be a minimal non-reducible regular corridor. No essential path component of the gate graph connects two external gates on the same endpoint section.

\{\}\textbf{Proof.}

Suppose a path component P connects two external gates on the input section Σ_- .

Let the two endpoint gates be

$$g_1, g_2.$$

The path component P , together with the interval of Σ_- between g_1 and g_2 , forms a closed walk in the corridor.

Choose an innermost simple cycle Z extracted from that closed walk.

By Lemma 11.1, the disk bounded by Z is one of:

- (1) a pure ear;
- (2) a relevant inner domain;
- (3) a proper identity subcorridor;
- (4) a proper non-identity subcorridor;
- (5) a closed train or subtrain.

Each alternative is reducible or contradicts minimality.

Therefore no such path component exists.

The same argument applies to a path component connecting two external gates on the output section Σ_+ .

□

12.6. Gate graph dichotomy.

Proposition 12.2 — Ladder/train dichotomy. Let \mathcal{C} be a minimal non-reducible regular corridor. Then \mathcal{C} contains either:

- (1) a local ladder; or
- (2) a closed alternating train.

\{\}\textbf{Proof.}

Since \mathcal{C} is nontrivial, its gate graph has at least one essential component.

Let

$$K$$

be an essential connected component of

$$\mathfrak{G}(\mathcal{C}).$$

If K contains a cycle, then that cycle is a closed alternating train.

Suppose K contains no cycle. Since every internal vertex has degree (2), K must be a path whose endpoints are external gates.

By Lemma 12.1, the two endpoints cannot lie on the same endpoint section. Therefore one endpoint lies on Σ_- and the other lies on Σ_+ .

Hence K is a local ladder.

Thus every minimal non-reducible regular corridor contains either a local ladder or a closed alternating train.

□

12.7. Interpretation for switch corridors. Let the regular corridor arise from a path γ between vertices p, q with

$$I(p) \cap I(q) = \emptyset.$$

The nontriviality of the boundary relation ensures that the corridor has an essential gate component. Proposition 12.2 therefore applies.

Thus a regular switch corridor always produces one of two objects:

local ladder or closed alternating train.

The next sections analyze these two outcomes. Complementary ladders close into trains, and closed alternating trains are impossible in thrackle drawings.

13. COMPLEMENTARY LADDERS CLOSE INTO TRAINS

In Section 12 we showed that a minimal non-reducible regular switch corridor contains either a local ladder or a closed alternating train. In this section we prove that if the two complementary corridors between the same two ordinary vertices both produce local ladders, then those ladders glue into a closed alternating train.

This is the only point where we use the fact that the two corridors come from the two complementary paths in the same original cycle.

13.1. The setting. Let

$$p, q$$

be ordinary vertices on the same cycle, say A , satisfying

$$I(p) \cap I(q) = \emptyset.$$

Let

$$\gamma_0, \gamma_1$$

be the two A -paths from p to q . Their union is the whole cycle A :

$$\gamma_0 \cup \gamma_1 = A.$$

Let

$$\mathcal{C}_0, \mathcal{C}_1$$

be minimal non-reducible regular switch corridors associated with γ_0 and γ_1 , respectively.

Suppose that neither corridor contains a closed alternating train. By Proposition 12.2, each corridor contains a local ladder. Let these ladders be

$$L_0 \subset \mathcal{C}_0, \quad L_1 \subset \mathcal{C}_1.$$

We shall prove that

$$L_0 \cup L_1$$

contains a closed alternating train.

13.2. Local structure near an ordinary vertex. We first record the local topology near an ordinary vertex.

Lemma 13.1 — Two-domain local pairing. Let $x \neq v$ be an ordinary original vertex. Then in a sufficiently small disk B_x around x , the planarized drawing consists of exactly two edge germs meeting at x . The complement

$$B_x \setminus P(T)$$

has exactly two local sectors. These two sectors belong to the two domains in

$$I(x) = D_x^-, D_x^+.$$

Consequently, any two corridor gates approaching x along the two complementary sides of the same cycle must pair according to the two domains D_x^- and D_x^+ , unless one of the following occurs:

- (1) a terminal lane;
- (2) a gate identification;
- (3) branching or merging of lanes;
- (4) a pure ear;
- (5) a relevant inner domain.

\{\}textbf{Proof.}

Since x is ordinary, its degree in the original dumbbell graph is (2). The planarization does not add additional edge germs at original vertices. Therefore, for a sufficiently small disk B_x , the intersection

$$P(T) \cap B_x$$

is exactly the union of two embedded arcs meeting only at x .

The complement of two arcs meeting at their common endpoint in a disk has exactly two local sectors. By Lemma 4.1 these sectors belong to two distinct global domains, namely the two domains in $(I(x))$.

Now consider two corridor structures approaching x from the two complementary sides of the cycle. A lane ending at x on one side must continue on the other side through the same local domain sector. There are only two such sectors, so the pairing is forced.

If the pairing fails, then one of the following must occur:

- a lane has no continuation, giving a terminal lane;
- two different gates represent the same local sector, giving a gate identification;
- one lane splits into two or two lanes merge into one, giving branching or merging;
- a closed disk is cut off near x with no original vertices inside, giving a pure ear;
- a closed disk is cut off containing original vertices, giving a relevant inner domain.

Each of these is one of the reducible configurations.

□

13.3. Endpoint gates of a local ladder. A local ladder is a path component of the gate graph of a regular corridor connecting an input external gate to an output external gate.

For a corridor associated with a path from p to q , its external gates occur near p and q . More precisely:

- the input terminal gates lie on a transverse section near p ;
- the output terminal gates lie on a transverse section near q .

Thus each local ladder has two terminal ends:

one near p , one near q .

Let

$$\partial_p L_i$$

denote the terminal gate of L_i near p , and

$$\partial_q L_i$$

the terminal gate of L_i near q .

At each of p, q , there are exactly two local domains. Hence the two ladder ends near that vertex must pair according to these two domains, unless a reducible configuration appears.

This is the mechanism that closes the two ladders into a cycle.

13.4. Complementary ladder closure.

Proposition 13.2 — Complementary ladders close. Let p, q be ordinary vertices on the same cycle. Let γ_0, γ_1 be the two complementary p -to- q paths in that cycle. Suppose the associated minimal non-reducible regular corridors contain local ladders

$$L_0, L_1.$$

Then

$$L_0 \cup L_1$$

contains a closed alternating train.

\{\}\textbf{Proof.}

The ladder L_0 has one terminal gate near p and one terminal gate near q . The same is true of L_1 .

Consider first the vertex p .

By Lemma 13.1, the local complement near p has exactly two domain sectors. The terminal gate of L_0 near p must continue through one of those sectors. The terminal gate of L_1 near p lies on the complementary side of the same cycle and must continue through the corresponding sector.

If the two terminal gates near p do not pair, then Lemma 13.1 gives one of the reducible configurations:

- terminal lane;
- gate identification;
- branching or merging;
- pure ear;
- relevant inner domain.

Since the corridors are assumed non-reducible, none can occur. Therefore the pairing at p is forced.

The same argument applies at q .

Thus the path component L_0 is glued to L_1 at both of its ends. After the two endpoint pairings are made, every gate in

$$L_0 \cup L_1$$

has degree (2): internal gates had degree (2) by regularity, and the formerly terminal gates now have their unique continuation through the local domain sectors at p and q .

Therefore

$$L_0 \cup L_1$$

contains a cyclic component of the gate graph.

By definition, a cyclic component of alternating cells glued along gates is a closed alternating train.

□

13.5. Consequence. For the pair p, q produced by the domain-pair separation lemma, the two complementary switch corridors have only two possible non-reducible outcomes:

- (1) one of them already contains a closed alternating train;
- (2) both contain local ladders, and the ladders close into a closed alternating train.

Thus:

Corollary 13.3. If both complementary switch corridors between p and q are non-reducible and regular, then a closed alternating train exists.

The rest of the proof is devoted to showing that such trains cannot occur in thrackle drawings.

14. CLOSED ALTERNATING TRAINS ARE IMPOSSIBLE

We now prove that the closed alternating trains produced by regular corridors cannot occur in a thrackle drawing.

The proof has two parts.

First, if the train proceeds for long enough on a fixed set of original supports, it creates a positive pure braid. Such a braid forces at least one pair of supports to cross twice.

Second, if support-shifts occur before such a fixed-support block appears, then following a lane around the train produces either a repeated support, a pure ear, a relevant inner domain, or a smaller train. All of these are reducible or contradict minimality.

14.1. Closed alternating trains. A closed alternating train is a cyclic sequence of alternating cells

$$H_{2m_1}, H_{2m_2}, \dots, H_{2m_r}$$

glued successively along gates so that the lane-continuations form a cyclic component of the gate graph.

The train is normal if it arises from a regular corridor. Thus:

- (1) every cell is an alternating cell H_{2m} with $m \geq 3$;

- (2) every internal gate has unique continuation;
- (3) there are no terminal lanes;
- (4) there is no branching or merging;
- (5) there are no gate identifications;
- (6) no two original supports cross twice inside the train;
- (7) no support is repeated in a way already classified as reducible.

A train is minimal if it has the smallest number of cells among all closed alternating trains occurring in the drawing after all reducible configurations have been excluded.

In the rest of this section we assume that

$$\mathcal{T}$$

is a minimal normal closed alternating train, and we derive a contradiction.

14.2. Orientation of cells. Each alternating cell has two possible orientations, depending on whether its boundary order induces a positive cyclic shift or a negative cyclic shift on the local ordering of lanes.

We call a cell positive if the induced local braid on its lanes is

$$\sigma_1\sigma_2\cdots\sigma_{m-1}$$

after a suitable ordering of the lanes.

It is negative if the inverse cyclic order is induced.

The sign convention itself is not important. What matters is that adjacent cells of opposite orientation form a local cancellation region.

14.3. Opposite orientations are reducible.

Lemma 14.1 — *No adjacent inverse orientations in a minimal train.* A minimal normal closed alternating train cannot contain two adjacent cells whose induced orientations are opposite.

\{\}\textbf{Proof.}

Suppose two adjacent cells H and H' have opposite orientations.

Consider the smallest subregion R consisting of the portion of $H \cup H'$ between the first lane-order inversion produced by H and the corresponding restoration produced by H' .

The boundary relation of R is identity: the first cell applies a local cyclic shift, and the second cell applies the inverse shift on the same local set of lanes.

If R is reduced, then by the boundary-rigidity theorem for reduced quadrilaterals it is identity and therefore removable from the witness train, contradicting the minimality of \mathcal{T} .

If R is not reduced, it contains one of the reducible configurations: pure ear, relevant inner domain, repeated support, gate identification, branching, terminal lane, or double crossing.

All are forbidden in a minimal normal train.

Therefore adjacent inverse orientations cannot occur.

□

14.4. Uniform orientation.

Corollary 14.2. Every minimal normal closed alternating train is uniformly oriented.

\{\}\textbf{Proof.}

If the cyclic sequence of cells contained both positive and negative cells, then at some place along the cyclic order a positive cell would be adjacent to a negative cell. This is forbidden by Lemma 14.1.

Therefore all cells have the same orientation.

After reversing the cyclic direction of the train if necessary, we may assume all cells are positive.

□

14.5. Fixed-support blocks and positive braids. A support is an original edge of the dumbbell, regarded as the union of all its planarized edgelets. A lane is said to be supported by a support s along a cell if the corresponding boundary arc of the cell lies on s .

A segment of a train is said to have fixed supports if the same ordered list of original supports carries its lanes through each cell of the segment.

Suppose a segment of the train consists of k consecutive positive cells of arity m on the same ordered set of m supports. Then the induced braid word is

$$(\sigma_1\sigma_2\cdots\sigma_{m-1})^k.$$

In particular, for $k = m$, we obtain

$$(\sigma_1\sigma_2\cdots\sigma_{m-1})^m = \Delta^2,$$

the full twist.

We shall use the following elementary braid fact.

Lemma 14.3 — Positive pure braid gives double crossing. Let β be a non-trivial positive braid on a fixed set of labelled strands. If the permutation induced by β is the identity, then at least one pair of labelled strands crosses at least twice.

\{\}\textbf{Proof.}

In a braid diagram, each crossing between a pair of labelled strands changes the relative order of that pair.

If the final permutation is the identity, then every pair of labelled strands has the same relative order at the end as at the beginning. Therefore each pair crosses an even number of times.

Since β is nontrivial and positive, at least one crossing occurs. The corresponding pair of labelled strands crosses a positive even number of times, hence at least twice.

□

14.6. Consequence for fixed-support trains.

Lemma 14.4 — Fixed-support train segments are impossible. A minimal normal closed alternating train cannot contain a nontrivial closed fixed-support segment.

In particular, it cannot contain m consecutive positive cells of arity m on the same m original supports.

\{\}\textbf{Proof.}

If a closed segment of the train uses a fixed set of labelled original supports and returns them to their initial order, then the corresponding braid is positive and pure.

It is nontrivial because the segment contains at least one alternating cell, hence at least one crossing in the braid.

By Lemma 14.3, some pair of original supports crosses at least twice.

But in a thrackle drawing, two original edges meet exactly once if non-adjacent, and adjacent edges meet only at their common endpoint. Thus two original supports cannot cross twice.

This contradiction proves the claim.

For m consecutive positive cells of arity m , the braid is explicitly

$$(\sigma_1 \sigma_2 \cdots \sigma_{m-1})^m = \Delta^2,$$

which is a positive pure braid. Hence the same contradiction applies.

□

14.7. Support-shifts. Since fixed-support closed behavior is impossible, any closed train must contain a support-shift.

A support-shift occurs when, following a lane through consecutive cells, the support carrying that lane changes from one original edge to another.

Because the train is closed, following a lane around the train yields a cyclic sequence of supports

$$s_0, s_1, \dots, s_{r-1}, s_0.$$

This is the support itinerary of the lane.

A support itinerary has a proper repetition if

$$s_i = s_j$$

for some

$$0 \leq i < j < r.$$

If no such repetition occurs, the itinerary is support-simple.

14.8. First repeated support.

Lemma 14.5 — *A proper repeated support is reducible.* If a support itinerary in a minimal normal closed train has a proper repeated support, then the train contains a reducible configuration or a smaller closed train.

\{\}\textbf{Proof.}

Let

$$s_i = s_j, \quad 0 \leq i < j < r,$$

be a proper repetition with (j-i) minimal.

Consider the segment of the lane between the two occurrences of the support $s_i = s_j$. Join its endpoints by the corresponding subarc of the original support s_i .

The union is a closed walk W in the planarized graph.

Apply the innermost closed-walk lemma.

If W contains a repeated support before the chosen repetition, this contradicts the minimality of (j-i).

If W contains a double crossing, we contradict the thrackle property.

If W contains a gate identification, terminal lane, or branching, the train is not normal.

Otherwise W contains an embedded simple cycle Z . Let D_Z be the disk bounded by Z , chosen innermost.

If D_Z contains no original vertices, then it is a pure ear or a removable identity region.

If D_Z contains original vertices, then it contains a relevant inner domain or a smaller T_3 -critical core, unless the part of the train inside D_Z is itself a closed train.

If the latter occurs, that train has fewer cells than \mathcal{T} , contradicting the minimality of \mathcal{T} .

Thus every possibility is reducible or contradicts minimality.

□

14.9. Support-simple itineraries.

Lemma 14.6 — *Support-simple closed lanes are reducible.* A support-simple closed lane in a minimal normal closed train produces a reducible configuration or a smaller closed train.

\{\}\textbf{Proof.}

Let the support itinerary be

$$s_0, s_1, \dots, s_{r-1}, s_0$$

with no proper repetition.

Consider the geometric trace of the lane in the plane.

If this trace self-intersects, then the self-intersection must occur in one of the following ways:

- (1) the same support is encountered twice, contradicting support-simplicity;
- (2) two distinct supports meet twice, giving a double crossing;
- (3) two gate intervals are identified;
- (4) a lane branches or merges.

All of these are forbidden in a normal train.

Therefore the trace of the lane is a simple closed curve.

Let

$$D$$

be the disk it bounds.

If

$$D \cap V_{\text{orig}} = \emptyset,$$

then D is a pure ear or a removable identity region.

If

$$D \cap V_{\text{orig}} \neq \emptyset,$$

then either D contains a relevant inner domain or a smaller T_3 -critical core, or the cells of the train inside D form a smaller closed alternating train.

In every case we obtain a reducible configuration or contradict the minimality of \mathcal{T} .

□

14.10. Impossibility of closed trains.

Theorem 14.7 — No closed alternating trains. A closed alternating train produced by a regular switch corridor cannot occur in a thrackle drawing.

\{\}\textbf{Proof.}

Suppose, for contradiction, that a closed alternating train occurs. Choose one with the minimum number of cells. It is normal because it arises from a regular corridor.

By Corollary 14.2, it is uniformly oriented. Assume all cells are positive.

If some nontrivial closed portion of the train has fixed supports, Lemma 14.4 gives a double crossing, impossible.

Therefore support-shift must occur.

Follow a lane with support-shift. Its cyclic support itinerary either has a proper repeated support or is support-simple.

If it has a proper repeated support, Lemma 14.5 gives a reducible configuration or a smaller closed train.

If it is support-simple, Lemma 14.6 gives a reducible configuration or a smaller closed train.

Both alternatives contradict the choice of a minimal normal closed train, or the thrackle condition itself.

Therefore no closed alternating train can occur.

□

15. PROOF OF THE INTERNAL EVEN-DUMBBELL THEOREM

We now assemble the internal proof.

Recall the statement.

Theorem 15.1 — Irreducible even dumbbell thrackles are T_3 . Let

$$T$$

be an irreducible thrackle drawing of

$$F_{2a,2b} = C_{2a} \vee C_{2b}.$$

Then

$$T \in T_3.$$

15.1. Proof by contradiction. Assume, for contradiction, that

$$T \notin T_3.$$

By Lemma 5.4, there exist ordinary vertices

$$p, q \neq v$$

lying on the same cycle such that

$$I(p) \cap I(q) = \emptyset.$$

Without loss of generality,

$$p, q \in A.$$

Let

$$\gamma_0, \gamma_1$$

be the two paths in A from p to q .

Each path gives a switch problem: the lateral pair near p is $(I(p))$, while the lateral pair near q is $(I(q))$, and these pairs are disjoint.

For each path γ_i , choose a minimal switch corridor

$$C_i.$$

15.2. Applying regularization. By the switch-corridor regularization theorem, each \mathcal{C}_i is either reducible or regular.

Suppose first that some \mathcal{C}_i is reducible.

Then one of the reducible configurations occurs:

- (1) edge-removal;
- (2) double crossing;
- (3) repeated support;
- (4) pure ear;
- (5) relevant inner domain or smaller T_3 -critical core;
- (6) proper non-identity subcorridor;
- (7) gate identification;
- (8) terminal lane;
- (9) returning lane;
- (10) branching or merging lane;
- (11) closed alternating train.

We analyze these alternatives.

- Edge-removal contradicts irreducibility of T .
- Double crossing contradicts the definition of a thrackle drawing.
- Repeated support, gate identification, returning lane, terminal lane, and branching or merging produce, by the innermost closed-walk reduction, either a pure ear, a relevant inner domain, a smaller corridor, a double crossing, or a closed train.
- A pure ear is eliminable from the witness corridor and contradicts minimality.
- A relevant inner domain yields either a smaller T_3 -critical core or a smaller nontrivial relational corridor, contradicting minimality of the chosen obstruction.
- A proper non-identity subcorridor contradicts minimality of the chosen corridor.
- A closed alternating train contradicts Theorem 14.7.

Thus no reducible corridor can occur.

Therefore both

$$\mathcal{C}_0 \quad \text{and} \quad \mathcal{C}_1$$

are regular.

15.3. Regular corridors yield ladders or trains. By Proposition 12.2, each regular corridor contains either a local ladder or a closed alternating train.

If either \mathcal{C}_0 or \mathcal{C}_1 contains a closed alternating train, then Theorem 14.7 gives a contradiction.

Therefore neither contains a closed train.

Hence both corridors contain local ladders:

$$L_0 \subset \mathcal{C}_0, \quad L_1 \subset \mathcal{C}_1.$$

15.4. **Complementary ladders close.** The paths

$$\gamma_0, \gamma_1$$

are complementary paths in the same cycle A .
By Proposition 13.2, the two ladders

$$L_0, L_1$$

glue at p and q to form a closed alternating train.

Indeed, p and q are ordinary degree-two vertices. Their local domain structure forces unique pairing of the terminal gates of the ladders. Failure of pairing would produce a reducible configuration, already excluded.

Thus a closed alternating train exists.

But this contradicts Theorem 14.7.

15.5. **Conclusion.** The contradiction arose from the assumption

$$T \notin T_3.$$

Therefore

$$T \in T_3.$$

This proves Theorem 15.1.

□

15.6. **Consequence for even dumbbells.** Let

$$F_{2a,2b} = C_{2a} \vee C_{2b}.$$

If T were an irreducible thrackle drawing of $F_{2a,2b}$, then by Theorem 15.1,

$$T \in T_3.$$

By the T_3 -thrackle theorem,

$$|E(F_{2a,2b})| \leq |V(F_{2a,2b})|.$$

But

$$|E(F_{2a,2b})| = 2a + 2b,$$

and

$$|V(F_{2a,2b})| = 2a + 2b - 1.$$

Therefore

$$|E(F_{2a,2b})| = |V(F_{2a,2b})| + 1,$$

a contradiction.

Thus no irreducible thrackle drawing of an even dumbbell exists.

This completes the internal proof.

16. COMPLETION OF THE PROOF OF CONWAY'S THRACKLE CONJECTURE

In this section we complete the proof of Conway's thrackle conjecture from the internal even-dumbbell theorem proved in Section 15.

Recall the two external results used in the reduction:

- (1) it is enough to exclude thrackle drawings of even dumbbells $(DB(c', c'', 0))$;
- (2) every T_3 -thrackle satisfies Conway's inequality.

The internal theorem proved above states that every irreducible thrackle drawing of an even dumbbell is T_3 .

We now combine these ingredients.

16.1. Minimal counterexamples and irreducibility. Suppose, for contradiction, that Conway's thrackle conjecture is false.

Then there exists a finite graph $G = (V, E)$ admitting a thrackle drawing with

$$|E| > |V|.$$

By the even-dumbbell reduction, there exists a thrackle drawing of an even dumbbell

$$DB(c', c'', 0),$$

where both c' and c'' are even.

We write

$$c' = 2a, \quad c'' = 2b,$$

and identify

$$DB(c', c'', 0) = F_{2a, 2b} = C_{2a} \vee C_{2b}.$$

Choose such a thrackle drawing minimal with respect to the standard edge-removal operation.

By the edge-removal criterion, this minimal drawing is irreducible. Indeed, if an edge-removal were possible, then the drawing could be replaced by a strictly smaller thrackle drawing in the same reduction class, contradicting minimality.

Thus we may assume that we have an irreducible thrackle drawing

$$T$$

of

$$F_{2a, 2b}.$$

16.2. Applying the internal theorem. By Theorem 15.1, every irreducible thrackle drawing of an even dumbbell belongs to T_3 .

Therefore

$$T \in T_3.$$

By the T_3 -thrackle theorem, every T_3 -thrackle satisfies

$$|E| \leq |V|.$$

Hence

$$|E(F_{2a,2b})| \leq |V(F_{2a,2b})|.$$

But direct counting gives:

$$|E(F_{2a,2b})| = 2a + 2b,$$

and

$$|V(F_{2a,2b})| = 2a + 2b - 1.$$

Thus

$$|E(F_{2a,2b})| = |V(F_{2a,2b})| + 1.$$

This contradicts

$$|E(F_{2a,2b})| \leq |V(F_{2a,2b})|.$$

Therefore no such even dumbbell thrackle drawing exists.

16.3. Final contradiction. The even-dumbbell reduction says that if Conway's conjecture were false, then such an even dumbbell thrackle drawing would exist.

We have just shown that no irreducible even dumbbell thrackle drawing exists.

Since a minimal even dumbbell counterexample would be irreducible, no minimal counterexample exists.

Hence Conway's thrackle conjecture is true:

$$|E| \leq |V|$$

for every finite graph admitting a planar thrackle drawing.

□

17. COMPATIBILITY OF DEFINITIONS AND EXTERNAL INPUTS

The proof above uses two external structural inputs: the even-dumbbell reduction and the T_3 -thrackle theorem. This section records the compatibility checks needed to ensure that the present notation matches those inputs.

17.1. Compatibility with the even-dumbbell notation. The reduction used in the literature refers to graphs of the form

$$DB(c', c'', 0).$$

Here the last parameter (0) means that the connecting path between the two cycles has length zero. Therefore the two cycles share a single vertex.

Thus

$$DB(c', c'', 0)$$

is the one-point union of two cycles of lengths c' and c'' .

When c' and c'' are even, write

$$c' = 2a, \quad c'' = 2b.$$

Then

$$DB(c', c'', 0) = C_{2a} \vee C_{2b} = F_{2a, 2b}.$$

This is exactly the graph considered throughout the internal proof.

17.2. Counting vertices and edges. For

$$F_{2a, 2b} = C_{2a} \vee C_{2b},$$

the edge count is immediate:

$$|E(F_{2a, 2b})| = |E(C_{2a})| + |E(C_{2b})| = 2a + 2b.$$

The vertex count is:

$$|V(F_{2a, 2b})| = |V(C_{2a})| + |V(C_{2b})| - 1,$$

because one vertex is identified.

Therefore

$$|E(F_{2a, 2b})| = |V(F_{2a, 2b})| + 1.$$

This is the numerical contradiction obtained once the drawing is shown to be T_3 .

17.3. Compatibility of the T_3 convention. In the present paper, a drawing is called T_3 if all original vertices lie on the boundaries of at most three connected domains of the complement of the drawing.

Equivalently, after planarization, there exist three connected components

$$D_1, D_2, D_3$$

of

$$\mathbb{R}^2 \setminus P(T)$$

such that

$$V_{\text{orig}} \subseteq \partial D_1 \cup \partial D_2 \cup \partial D_3.$$

This agrees with the convention used in the theorem on pants or T_3 -thrackles: the relevant domains are connected complementary domains, and the vertices considered are the original graph vertices, not the crossing vertices introduced by planarization.

Crossing vertices are artifacts of planarization and are not required to lie on the boundaries of the three domains in the definition of T_3 .

Thus the internal theorem

$$\text{irreducible even dumbbell thrackle} \Rightarrow T_3$$

is directly compatible with the external result

$$T_3\text{-thrackle} \Rightarrow |E| \leq |V|.$$

17.4. Compatibility of irreducibility and minimality. The proof uses irreducibility with respect to the standard edge-removal operation.

The role of irreducibility is limited: it is used only to exclude the reducible configuration called edge-removal. Every time the internal proof produces an empty edge-removal triangle, the assumed irreducibility gives a contradiction.

A minimal even dumbbell counterexample may be assumed irreducible because the edge-removal operation preserves the thrackle property and reduces the size of the drawing. If a minimal counterexample admitted such a reduction, then a smaller counterexample of the same reduction type would exist.

Thus the irreducibility hypothesis in Theorem 15.1 is exactly the one needed for the final reduction.

17.5. Compatibility of planarization. The planarization (P(T)) used throughout the proof does not change the complementary domains of the drawing. It only converts crossings into degree-four vertices to make the drawing into an embedded planar graph.

Therefore:

- (1) the domains of T and (P(T)) are the same connected regions of the plane;
- (2) the boundary-incidence relation

$$x \in \partial D$$

for original vertices x and domains D is unchanged;

- (1) the T_3 -condition may be checked on the planarized graph.

This justifies using (P(T)) throughout the proof while invoking a theorem stated for the original drawing.

18. DEPENDENCY SUMMARY AND END-OF-PROOF CHECKLIST

This final section summarizes the logical dependencies of the proof and identifies where each major ingredient is used.

18.1. Main dependency chain. The proof has the following dependency structure.

First, the external reduction gives:

$$\text{Conway false} \Rightarrow \text{even dumbbell thrackle exists.}$$

Choosing one minimally gives:

$$\text{irreducible even dumbbell thrackle.}$$

The internal theorem gives:

$$\text{irreducible even dumbbell thrackle} \Rightarrow T_3.$$

The external T_3 -theorem gives:

$$T_3 \Rightarrow |E| \leq |V|.$$

But the even dumbbell satisfies:

$$|E| = |V| + 1.$$

Contradiction.

18.2. Internal proof dependency chain. The proof of the internal theorem has the following structure.

Step 1: domain-pair separation

$$T \notin T_3 \Rightarrow \exists p, q \text{ on the same cycle with } I(p) \cap I(q) = \emptyset.$$

This uses:

- planarization;
- domains;
- the ordinary domain graph G_D ;
- bipartiteness of the face-dual;
- König's theorem.

Step 2: switch corridors

The pair p, q gives two complementary switch problems along the two paths between p and q .

Each switch problem is represented by a minimal relational corridor.

This uses:

- geometric switch traces;
- PL slicing;
- relational corridors;
- the Geometric Realization Lemma.

Step 3: corridor regularization

Each minimal corridor is either reducible or regular.

This uses:

- reduced quadrilaterals;
- boundary rigidity;
- monochromatic lens elimination;
- innermost closed-walk reductions;
- short-cell elimination.

Step 4: ladder/train dichotomy

Each regular corridor contains either:

local ladder or closed alternating train.

This uses the gate graph.

Step 5: complementary ladders close

If both complementary corridors yield ladders, the ladders glue at p and q to form a closed alternating train.

This uses the fact that p and q are ordinary degree-two vertices.

Step 6: trains are impossible

Closed alternating trains are impossible by:

- uniform-orientation reduction;
- positive braid/full-twist obstruction;
- support-shift analysis;
- first repeated support and innermost disk arguments.

Therefore $T \notin T_3$ is impossible.

18.3. Reducible configurations and their contradictions. The proof uses the following reducible configurations.

Reducible configuration	Contradiction or reduction
Edge-removal	Contradicts irreducibility
Double crossing	Contradicts thrackle condition
Repeated support	Produces closed walk, hence ear/domain/smaller train/double crossing
Pure ear	Eliminable from the corridor witness
Relevant inner domain	Gives smaller T_3 -critical core or smaller corridor
Non-identity subcorridor	Contradicts minimality of the chosen corridor
Identity subcorridor	Removable without changing relation
Gate identification	Produces closed walk or branching
Terminal lane	Produces ear/domain or smaller corridor
Returning lane	Produces closed walk and hence a reduction

Branching/merging lane	Produces gate identification, closed walk, or domain reduction
Closed alternating train	Impossible by train theorem
Proper subtrain	Contradicts minimality of train

18.4. **Key local topological facts.** The local topology of the proof rests on four facts.

Fact 1: disjoint arcs preserve order

Pairwise disjoint properly embedded arcs in a disk connecting two ordered boundary intervals induce an order-preserving bijection.

This proves boundary rigidity for reduced quadrilaterals.

Fact 2: closed walks contain innermost disks

Every closed walk in a plane graph either yields a simple cycle or exhibits a repeated support, double crossing, or gate identification.

The innermost disk bounded by a simple cycle gives a pure ear, relevant inner domain, subcorridor, or train.

Fact 3: identity factors are removable

A relational subcorridor with identity boundary relation may be deleted from the witness corridor without changing the total relation.

Fact 4: positive pure braids force repeated crossings

A nontrivial positive pure braid forces some pair of strands to cross at least twice. Since original supports in a thrackle may meet only once, such a braid cannot occur.

18.5. **Conclusion.** The proof reduces the global conjecture to the exclusion of irreducible even dumbbell thrackles, proves that every such drawing is T_3 , and then invokes the known T_3 -thrackle theorem.

Thus every planar thrackle drawing satisfies

$$|E| \leq |V|.$$

This completes the proof of Conway's thrackle conjecture.

□

DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this work, ChatGPT, by OpenAI, was used to assist with mathematical drafting, formalization, review, and editing. This work is shared as a preliminary AI-assisted mathematical note. The mathematical content may have been only partially reviewed and may contain errors; it should not be treated as peer-reviewed or as a fully verified manuscript.

REFERENCES

- [Woo71] D. R. Woodall, Thrackles and deadlock, in D. J. A. Welsh, ed., *Combinatorial Mathematics and its Applications: Proceedings of the Conference held at the Mathematical Institute, Oxford, 1969*, Academic Press, London, 1971, pp. 335–347.
- [CK99] G. Cairns and D. M. King, The answer to Woodall’s musquash problem, *Discrete Mathematics* 207(1–3) (1999), 25–32.
- [CK01] G. Cairns and D. M. King, All odd musquashes are standard, *Discrete Mathematics* 226(1–3) (2001), 71–91.
- [LPS97] L. Lovász, J. Pach, and M. Szegedy, On Conway’s thrackle conjecture, *Discrete & Computational Geometry* 18(4) (1997), 369–376.
- [CN00] G. Cairns and Y. Nikolayevsky, Bounds for generalized thrackles, *Discrete & Computational Geometry* 23(2) (2000), 191–206.
- [CMN04] G. Cairns, M. McIntyre, and Y. Nikolayevsky, The thrackle conjecture for K_5 and $K_{3,3}$, in *Towards a Theory of Geometric Graphs, Contemporary Mathematics* 342, American Mathematical Society, Providence, RI, 2004, pp. 35–54.
- [CN09] G. Cairns and Y. Nikolayevsky, Generalized thrackle drawings of non-bipartite graphs, *Discrete & Computational Geometry* 41(1) (2009), 119–134.
- [PS11] J. Pach and E. Sterling, Conway’s conjecture for monotone thrackles, *The American Mathematical Monthly* 118(6) (2011), 544–548.
- [CN12] G. Cairns and Y. Nikolayevsky, Outerplanar thrackles, *Graphs and Combinatorics* 28(1) (2012), 85–96.
- [CKN15] G. Cairns, T. J. Koussas, and Y. Nikolayevsky, Great-circle spherical thrackles, *Discrete Mathematics* 338(12) (2015), 2507–2513.
- [FP11] R. Fulek and J. Pach, A computational approach to Conway’s thrackle conjecture, *Computational Geometry: Theory and Applications* 44(6–7) (2011), 345–355.
- [GX17] L. Goddyn and Y. Xu, On the bounds of Conway’s thrackles, *Discrete & Computational Geometry* 58(2) (2017), 410–416.
- [FP19] R. Fulek and J. Pach, Thrackles: an improved upper bound, *Discrete Applied Mathematics* 259 (2019), 226–231.
- [Xu21] Y. Xu, A new upper bound for Conway’s thrackles, *Applied Mathematics and Computation* 389 (2021), Article 125573.
- [MN18] G. Misereh and Y. Nikolayevsky, Annular and pants thrackles, *Discrete Mathematics & Theoretical Computer Science* 20(1) (2018), Article #16.
- [Mis18] G. Misereh, Thrackles containing a standard musquash, *Australasian Journal of Combinatorics* 70(2) (2018), 168–182.
- [HVKS25] C. Hernández-Vélez, J. Kynčl, and G. Salazar, Thrackles on nonplanar surfaces, arXiv:2506.11808, 2025.