

Outer Billiards and The Pinwheel Map

Richard Evan Schwartz *

July 24, 2009

1 Introduction

1.1 Background

B. H. Neumann [N] introduced outer billiards in the late 1950s and J. Moser [M1] popularized the system in the 1970s as a toy model for celestial mechanics. Outer billiards is a discrete self-map of $\mathbf{R}^2 - P$, where P is a bounded convex planar set as in Figure 1.1 below. Given $x_1 \in \mathbf{R}^2 - P$, one defines x_2 so that the segment $\overline{x_1x_2}$ is tangent to P at its midpoint and P lies to the right of the ray $\overrightarrow{x_1x_2}$. The map $x_1 \rightarrow x_2$ is called *the outer billiards map*. The map is almost everywhere defined and invertible.




Figure 1.1: Outer billiards relative to P .

The purpose of this paper is to establish an equivalence between polygonal outer billiards and an auxiliary map which we call the *pinwheel map*. We

* Supported by N.S.F. Research Grant DMS-0072607

worked out this equivalence for a restricted version of outer billiards on kites¹ in [S2, Pinwheel Lemma]. This equivalence allows one to answer questions about polygonal outer billiards just by understanding the pinwheel map, a simpler system that seems more amenable to analysis.

For instance, the pinwheel map has a geometrically appealing acceleration – i.e. speed up of the time parameter. This acceleration, in turn, has a higher dimensional compactification as a polytope exchange map. In the one case we have already worked out, that of the same restricted system on kites, the compactification is a polyhedron exchange map that involves a union of two flat 3-tori dissected into smaller convex polyhedra. The exchange depends on the kite, which (when properly normalized) depends on a single parameter. When all the exchanges are assembled into a kind of fibered product, the result is an exchange defined by 4-dimensional convex lattice polytopes. See [S2, Master Picture Theorem]. This is the simplest possible context for these constructions, and already it is quite rich.

For general convex n -gons, the pinwheel map and its compactification promise to be both extremely rich, and also useful for resolving the question about unbounded orbits in the polygonal case. This question, posed for general convex shapes, is known as the *Moser-Neumann problem*. For reference, here are some highlights of the progress made on the Moser-Neumann problem over the past 30 years. (See [T1], [T2], [DT1] and [S2] for additional references.)

- J. Moser [M2] sketches a proof that outer billiards on P has all bounded orbits provided that ∂K is at least C^6 smooth and positively curved. R. Douady gives a complete proof in his thesis [D].
- In Vivaldi-Shaidenko [VS], Kolodziej [Ko], and Gutkin-Simanyi [GS], it is proved (each with different methods) that outer billiards on a *quasirational polygon* has all orbits bounded. This class includes the regular polygons, and polygons with all rational vertices.
- D. Genin [G] shows that all orbits are bounded relative to trapezoids.
- In [S1] we prove that outer billiards on the Penrose kite (the convex tile in the Penrose tiling) has unbounded orbits, thereby answering the Moser-Neumann question in the affirmative.

¹A kite is a convex quadrilateral having a diagonal that is a line of symmetry.

- D. Dolgopyat and B. Fayad [DF] show that outer billiards on a half-disk has some unbounded orbits. Their proof also works for regions obtained from a disk by nearly cutting it in half with a straight line. This is a second affirmative answer to the Moser-Neumann question.
- In [S2] we proved that outer billiards has unbounded orbits when defined relative to any irrational kite. A kite is irrational if one of its diagonals divides P into two triangles whose areas are irrational multiples of each other. [S2, Pinwheel Lemma] and [S2, Master Picture Theorem] are two foundational results in this work.

1.2 The Main Result

We will prove a version of the Pinwheel Lemma for convex polygons that have no parallel sides. We call such polygons *nice*. There is probably a version of the result that works for any convex polygon, but allowing parallel sides introduces annoying complications.




Figure 1.2: The strip associated to e .

Let P be a nice n -gon. We orient the edges of P so that they go clockwise. To each edge e of P , one associates a pair (Σ, V) , where Σ is an infinite strip in the plane and V is a vector that points from one edge of Σ to the other. Let v be the head vertex of e . Since P is a nice polygon, there is a unique vertex w of P that is as far as possible from the line L containing e . Let $V = 2(w - v)$. Let L' be the line parallel to L such that w is equidistant from L and L' . We let Σ be the strip bounded by L and L' . This is a construction that comes up often in polygonal outer billiards. We call (Σ, V) a *pinwheel pair* and we call Σ a *pinwheel strip*.

The n pinwheel strips $\Sigma_1, \dots, \Sigma_n$ are cyclically ordered, according to their slopes. Our convention is that the strips rotate counterclockwise as we move forward through the indices.

Given the pair (Σ, V) , we define a map μ on $\mathbf{R}^2 - \partial\Sigma$ as follows.

- If $x \in \Sigma - \partial\Sigma$ then $\mu(x) = x$.
- If $x \notin \Sigma$ then $\mu(x) = x \pm V$, whichever point is closer to Σ .

The map μ moves points “one step” closer to lying in Σ , if they don’t already lie in Σ . Note that μ is not defined on the boundary $\partial\Sigma$.

Let $\mathbf{R}_n^2 = \mathbf{R}^2 \times \{1, \dots, n\}$, with indices taken mod n . We define the *pinwheel map* $\psi^* : \mathbf{R}_n^2 \rightarrow \mathbf{R}_n^2$ by the following conditions.

- $\psi^*(p, k) = (\mu_{k+1}(p), k + 1)$ if $\mu_{k+1}(p) = p$.
- $\psi^*(p, k) = (\mu_{k+1}(p), k)$ if $\mu_{k+1}(p) \neq p$.

In other words, we try to move p by the k th strip map. If the point doesn’t move, we increment the index and give the next strip map a chance to move the point.

Let ψ denote the second iterate of the outer billiards map. We define ψ to be the identity inside the polygon P . Let $\pi : \mathbf{R}_n^2 \rightarrow \mathbf{R}^2$ be the projection map. A *section* is a map $\iota : \mathbf{R}^2 \rightarrow \mathbf{R}_n^2$ such that $\pi \circ \iota$ is the identity.

Theorem 1.1 (Pinwheel) *There is a section $\iota : \mathbf{R}^2 \rightarrow \mathbf{R}_n^2$ such that*

$$\psi(p) = \pi \circ (\psi^*)^k \circ \iota; \quad k = k(p) \in \{1, \dots, 3n\}.$$

This relation holds on all points for which ψ is well defined.

Far from the origin, we have $k(p) = 1$ unless $\psi(p)$ lies in a pinwheel strip. In this case $k(p) = 2$ because one extra iterate of ψ^* is required to shift the index. The Pinwheel Theorem is well known, and very easy to prove, for points far from the polygon. (See Lemma 2.1.) For instance, some version of the “far away result” comes up in [K] and [GS].

The new information given by the Pinwheel Theorem is that correspondence extends in some way to the whole plane. When $k(p) > 2$ it means that there is a funny cancellation that happens in order to make the two systems line up. The fact that $k(p) \leq 3n$ puts a bound on the complexity of this cancellation. The bound we get on $k(p)$ probably isn’t sharp but it has the right order of dependence on n .

1.3 Corollaries

Just from knowledge of the relation far from the origin, one can conclude nothing about how the unbounded orbits of one system compare to the unbounded orbits of the other. It is easy to construct two maps of the plane that agree outside a compact set, such that one of the maps has unbounded orbits and the other one doesn't.

The Pinwheel Theorem adds the information that rules out such wierd pathologies. We say that an unbounded orbit of ψ^* is *natural* if it lies in $\iota(\mathbf{R}^2)$ sufficiently far from the origin. In the next result, the word *orbit* means both the forwards and backwards orbit. This makes sense because both ψ and ψ^* are invertible. We state our result for the *forward direction* of the orbit, though a similar statement holds for the backwards direction, and for both directions at the same time.

Corollary 1.2 *Relative to any nice polygon, there is a canonical bijection between the forward unbounded ψ -orbits and the natural forward unbounded ψ^* orbits. The bijection sending the orbit O to the orbit O^* is such that $O = \pi(O^*)$ outside a compact subset.*

A big reason why one might want to study ψ^* in place of ψ is that ψ^* has an appealing acceleration Define

$$\widehat{X} = \bigcup_{j=1}^n (\Sigma_j \times \{j\}) \subset \mathbf{R}_n^2. \quad (1)$$

Topologically, \widehat{X} is the disjoint union of n strips. \widehat{X} agrees with $\iota(\mathbf{R}^2)$ outside a compact set. We let $\widehat{\psi} : \widehat{X} \rightarrow \widehat{X}$ be the first return map of ψ^* to X . Geometrically, we start with a point in Σ_1 and iterate μ_2 until we land in Σ_2 , then iterate μ_3 until we land in Σ_3 , etc.

Corollary 1.3 *There is a canonical bijection between the forward unbounded orbits of ψ and the forward unbounded orbits of $\widehat{\psi}$.*

One can accelerate somewhat further. The map $\widehat{\Psi} = (\widehat{\psi})^n$ preserves each individual strip in \widehat{X} . The action on each strip is conjugate to the action on any of the other ones outside of a compact set. Thus, we can pick on of the strips, say Σ_1 , and consider the map $\widehat{\Psi} : \Sigma_1 \rightarrow \Sigma_1$. We call $\widehat{\Psi}$ the *pinwheel return map*.

At the same time, we can consider the first return map of ψ to Σ_1 . We call this map Ψ . Again, it is well known that $\Psi = \widehat{\Psi}$ outside a compact set, but this is not enough information to produce a correspondence between the unbounded orbits of the two systems. Our next corollary fills in this information.

Corollary 1.4 *There is a canonical bijection between the forward unbounded orbits Ψ and the forward unbounded orbits of $\widehat{\Psi}$.*

This last corollary is pretty close to a direct generalization of the Pinwheel Lemma we proved for kites in [S2]. In [S2] we actually proved that $\Psi = \widehat{\Psi}$ if Σ_1 was properly chosen from amongst the 4 possible strips. However, this stronger result does not seem to be true in general.

1.4 Outline of the Paper

In §2 we reduce the Pinwheel Theorem to two auxilliary results, Theorem 2.3 and Theorem 2.4. Roughly, Theorem 2.4 describes an infinite family of “local substitution rules” that allows one to convert between between the pinwheel map and the polygonal outer billiards map. Theorem 2.3 guarantees that all these local rules patch together globally.

In §3 and §4 we establish some geometric and combinatorial facts that we will need for the proofs of Theorems 2.3 and 2.4. In §5 we prove Theorem 2.3, and in §6-7 we prove Theorem 2.4. Ultimately, the argument boils down to robust general properties like convexity and induction, but only after we find the combinatorial structure of what is going on.

I would like to thank John Smillie and Sergei Tabachnikov for helpful conversations about outer billiards and the pinwheel map. This work is part of an ongoing conversation John Smillie about the pinwheel map. In particular, the formulation we give of the Pinwheel Theorem emerged in conversations with Smillie.

I would also like to thank the Clay Mathematics Institute for their support during my sabbatical this year.

2 Proof in Broad Strokes

2.1 The Forward Partition

Let P be a nice n -gon. Let ψ be the square outer billiards map. We have already mentioned that ψ is a piecewise translation. For almost every point $p \in \mathbf{R}^2 - P$, there is a pair of vertices (v_p, w_p) of P such that

$$\psi(p) = p + V_p; \quad V_p = 2(w_p - v_p). \quad (2)$$

See Figure 2.1.




Figure 2.1: A piecewise translation

The dependence of V_p on p is locally constant, and the regions where the map $p \rightarrow V_p$ is constant are convex and polygonal. Since there are only finitely many pairs of vertices of P , we have a partition of $\mathbf{R}^2 - P$ into convex polygonal sets. We call this partition the *forward partition* associated to P . Here we summarize some of the results we establish in the next chapter.

- P has $2n$ unbounded regions, which we provisionally call $R_i(\pm)$ for $i = 1, \dots, n$. The region $R_i(\pm)$ is labelled by the pair of vertices (v, w) such that $2(w - v) = \pm V_i$. Compare Figure 1.2.
- Sufficiently far from the origin, the regions $R_i(+)$ and $R_-(i)$ share a ray with Σ_i . A sufficiently large circle centered at the origin, oriented counterclockwise, encounters $R_1, \Sigma_1, R_2, \Sigma_2$, etc. The ordering of the pluses and minuses depends on the geometry of P . See Figure 2.3 below.

Figure 2.2 shows an example of the forward partition for a nice octagon. The central shaded figure is the octagon. The white regions surrounding the octagon are the compact tiles of the partition. The shaded regions on the outside, which have been cut off by the bounding box of the figure, are unbounded tiles. One of the compact tiles has also been cut off by the bounding box.

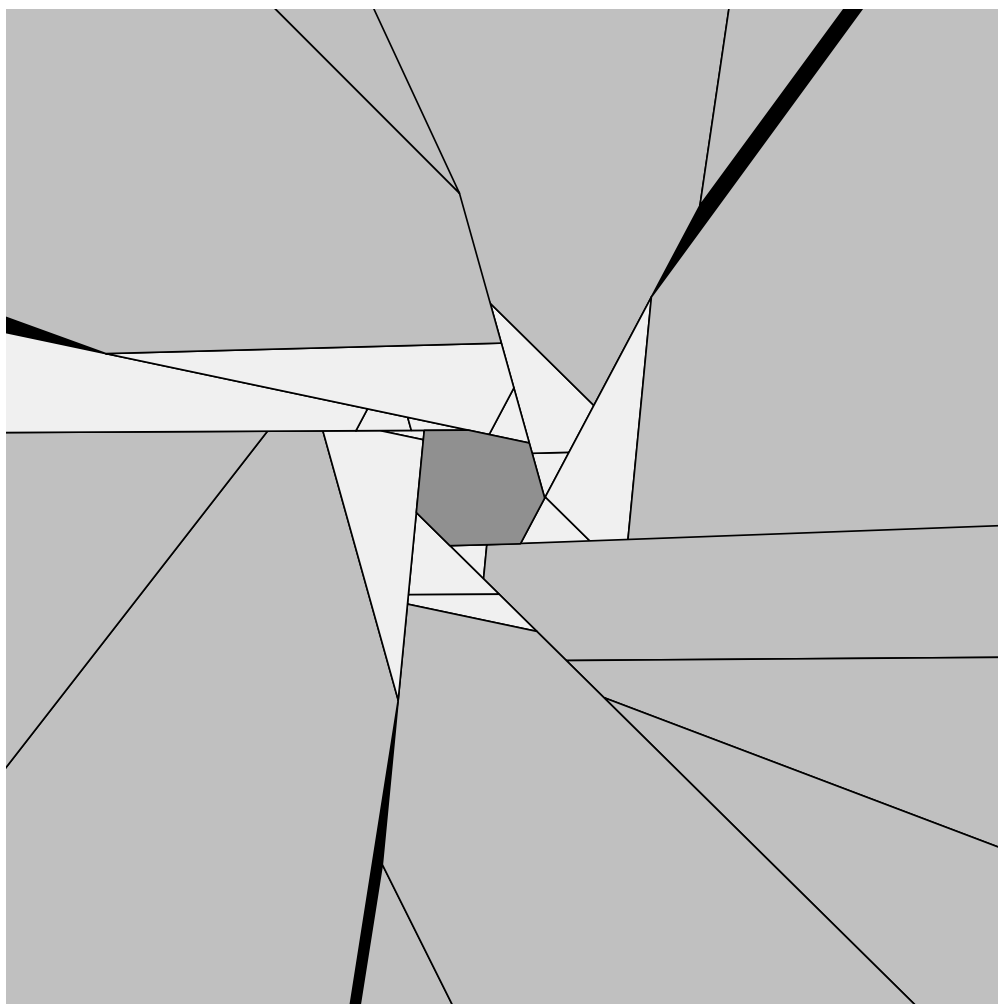


Figure 2.2: The forward partition for an octagon.

2.2 Far From the Origin

Let μ_1, \dots, μ_n be the strip maps defined in §1.2. Let \widehat{X} be as in Equation 1. Given a compact set K , let

$$\widehat{X}_K = \widehat{X} - \pi^{-1}(K); \quad X_k = (\Sigma_1 \cup \dots \cup \Sigma_n) - K. \quad (3)$$

If K is sufficiently large, the map $\pi : \widehat{X}_K \rightarrow X_K$ is a bijection (and a local isometry.) The point is that the pinwheel strips are all disjoint far from the origin, thanks to the fact that P is a nice polygon.

Lemma 2.1 $\psi = \pi \circ \psi^* \circ \pi^{-1}$ on X_K provided that K is sufficiently large.

Proof: Figure 2.3 shows the situation for $n = 4$. The central disk Δ covers up all the bounded tiles. The dark lines are parts of the pinwheel strip boundaries. We just show the top half of the picture.




Figure 2.3: Outside a compact set

Proof: We will take $k = 1$ and set $p_1 = p$. Provided that $p_1 \in \Sigma_1$ starts sufficiently far away from Δ , the map ψ simply adds V_1 to p_1 until the resulting orbit reaches a point $p_2 \in \Sigma_2$. The map M^* does exactly the same thing, and after the same number steps, we arrive at the point $(p_2, 2)$. And so on. ♠

2.3 The Pinwheel Identities

We have already associated n vectors to our nice n -gon P , namely V_1, \dots, V_n . Each V_k defines a line segment S_k that joins the two ends of S_k . Referring to Figure 1.2, the spoke S_k joins the vertices v and w . There are n spokes associated to P and we write these as S_1, \dots, S_n . This cyclic ordering is compatible with the ordering that comes from the slopes of the spokes. Figure 2.4 shows an example, with some of the edges highlighted.




Figure 2.4: The spokes of the polygon, and the path $7 \rightarrow 3$.

We say that an oriented, connected polygonal path γ is *admissible* if the following holds.

- γ consists of an odd number of spokes of P .
- The ordering on the spokes of γ is compatible with the cyclic order.
- Let γ' be the polygonal path in $\mathbf{R} \cup \infty$ obtained by connecting the slopes of the spokes of γ . Then γ' is a proper subset of $\mathbf{R} \cup \infty$.

The last condition means that γ does not wrap all the way around P . We use the notation $a \rightarrow b$ to name admissible paths. The first spoke is a and the last one is b . We use the notation $a \rightarrow b$ to name admissible paths. We often take the indices mod n and use $b + n$ in place of b in case $b < a$. Thus, $7 \rightarrow 10$ is another name for the the path in Figure 2.5.

We prove the following result in §4.2

Lemma 2.2 *There is a bijection between tiles in the forward partition and admissible paths. The tile corresponding to $a \rightarrow b$ is labelled by the vertex pair (v, w) , where v is the first vertex of $a \rightarrow b$ and w is the last one.*

We let $T(a \rightarrow b)$ denote the tile in the forward partition corresponding to the path $a \rightarrow b$. Here are our two main technical results.

Theorem 2.3 *Suppose $p \in T(a \rightarrow b)$ and $q = \psi(p) \in T(c \rightarrow d)$. Then $(M^*)^k(q, b - 1) = (q, c - 1)$ for some $k \in \{0, \dots, n\}$.*

Theorem 2.4 *Let $a \rightarrow b$ be an admissible path, labelled so that $b \geq a$. Then $(M^*)^k(p, a - 1) = (\psi(p), b - 1)$ for some $k \in \{1, \dots, 2n\}$.*

In §5 and §6-7 we prove Theorems 2.3 and 2.4 respectively. Combining these two results, we have

Corollary 2.5 *Let $T(a \rightarrow b)$ and $T(c \rightarrow d)$ be two tiles such that there is a point p such that $p \in T(a \rightarrow b)$ and $\psi(p) \in T(c \rightarrow d)$. Then there is some integer $k \in \{0, \dots, 3n\}$ such that $(M^*)^k(p, a - 1) = (\psi(p), c - 1)$.*

Now we deduce the Pinwheel Theorem from Corollary 2.5. The map ψ is defined precisely on the interiors of the tiles in the forward partition. For p in the interior of the tile $T(a \rightarrow b)$, we define

$$\iota(p) = (p, a - 1). \tag{4}$$

As usual, we take the indices mod n . The conclusion of Corollary 2.5 is just a restatement of the equation in the Pinwheel Theorem.

2.4 Dynamical Corollaries

2.4.1 Proof of Corollary 1.2

Let $p \in \mathbf{R}^2$ be a point with an unbounded ψ -orbit O . Let $p^* = \iota(p)$. Let O^* be the ψ^* orbit of p^* . By the Pinwheel Theorem, there is an infinite sequence t_1, t_2, \dots such that

$$\iota \circ \psi^k(p) = (\psi^*)^{t_k}(p^*).$$

This clearly shows that O is unbounded if and only if O^* is unbounded. The analysis in Lemma 2.1 shows that $\pi(O^* - K^*) = O - K$ if K is a sufficiently large compact set. Here $K^* = \pi^{-1}(K)$.

Since π is a bijection outside of K^* , there is only one orbit O^* such that $\pi(O^* - K) = O - K$. Hence, the assignment $O \rightarrow O^*$, which first seems to depend on the choice of p , is well defined independent of the choice. If O_1 and O_2 are different unbounded orbits, they differ outside of K . Hence O_1^* and O_2^* differ as well. This shows that the assignment $O \rightarrow O^*$ is injective.

Finally, let O^* be some unbounded natural orbit. We just choose some $p \in O^* - K^*$ and let $p = \pi(p^*)$. Then the argument above shows that O^* is the image of O under our correspondence. Hence, our correspondence is a bijection.

2.4.2 Proof of Corollary 1.3

In light of Corollary 1.2, we just have to construct a bijection between the set of forward unbounded natural orbits of ψ^* and the set of forward unbounded orbits of $\hat{\psi}$. But $\iota(\mathbf{R}^2)$ and \hat{X} agree outside a compact set. So, suppose that O^* is a forward unbounded natural orbit. The set $\hat{O} = O^* \cap \hat{X}$ is a forward unbounded orbit of \hat{X} . The nature this construction makes it clear that the correspondence $O^* \rightarrow \hat{O}$ is a bijection.

2.4.3 Proof of Corollary 1.4

For each orbit O of ψ , the intersection $O \cap \Sigma_1$ is the corresponding orbit of Ψ . In light of the analysis in Lemma 2.1 this gives a bijection between the set of unbounded orbits of ψ and the set of unbounded orbits of Ψ . Similarly, there is a canonical bijection between the set of unbounded orbits of $\hat{\psi}$ and the set of unbounded orbits of $\hat{\Psi}$. Finally, Corollary 1.3 gives us a canonical bijection between the unbounded orbits of ψ and the set of unbounded orbits of $\hat{\psi}$. Composing all these bijections gives us the desired result.

3 Spokes

3.1 Basic Definitions

Throughout the chapter, P is a nice n -gon. We say that a *minimal strip* is a strip Σ that contains the interior of P in its interior and is not a proper subset of any other such strip. (The pinwheel strips are twice as fat as certain minimal strips.) We call an ordered pair (v, w) of vertices of P a *maximal pair* if there is a minimal strip Σ such that v and w lie in distinct components of $\partial\Sigma$. The spokes we defined in §2.3 are precisely the line segments joining vertices of maximal pairs. As in the previous chapter, we let S_k denote the spoke corresponding to the vector V_k .

Lemma 3.1 *Let S_1 and S_2 be two consecutive spokes. Then S_1 and S_2 share a common vertex. Moreover, there is a minimal strip that contains all three vertices of $\partial S_1 \cup \partial S_2$.*

Proof: For each $\theta \in \mathbf{R} \cup \infty$, there is a unique minimal strip Σ_θ having slope θ . For all but n value of θ , the strip Σ_θ contains 2 vertices of P on its boundary. For such value, there is a unique spoke S_θ selected by Σ_θ . The assignment $\theta \rightarrow S_\theta$ is locally constant and changes only at one of the n special values. At these special values, Σ_θ contains 3 vertices of P in its boundary, and these are precisely the endpoints of two consecutive spokes. Our lemma is a restatement of this geometric picture. ♠

Recall that e_1, \dots, e_n are the edges of P , ordered according to slope. The correspondence $e_j \rightarrow S_j$ is what we call the *clockwise correspondence*. Figure 2.4 shows an example.

Lemma 3.2 *The clockwise correspondence is a bijection between edges and spokes.*

Proof: By the Pidgeonhole Principle, it suffices to prove that the correspondence $e \rightarrow P_e$ is an injection. So, for the sake of contradiction, suppose that there are two edges e and f such that $P_e = P_f = P$. Let a, b, c, d be the edges of e and f , ordered as in Figure 3.1. The vertex b is farther from \overline{cd} than is a and the vertex c is farther from \overline{ab} than is d . This is a geometrically impossible situation. ♠




Figure 3.1: An impossible situation

3.2 Smooth Approximation

We say that an *oval* is a smooth and convex simple closed curve of strictly positive curvature. The condition of strict positive curvature guarantees that the tangent lines to the curve vary strictly monotonically in small neighborhoods. This monotonicity is all we really need.

Minimal strips and spokes are defined for ovals just as we defined them for nice polygons. The spokes are easier to understand for ovals than for polygons. Here are two properties of spokes associated to ovals.

1. Any oval has a unique spoke of a given slope.
2. Any two distinct spokes cross in their interior.

These properties follow from the monotonicity of the tangent lines and continuity. We omit the easy proofs.

We are really only interested in the spokes of a nice polygon, but one efficient way to understand certain things about them is through smooth approximations. We can find a sequence $\{P_n\}$ of ovals that converges to P in the Hausdorff topology. More precisely, we mean that P_n is contained in the $(1/n)$ -tubular neighborhood of P , and *vice versa*. We omit the straightforward proof of this fact.

Lemma 3.3 *For any spoke S of P there is a sequence $\{S_n\}$ that converges to S . Here S_n is a spoke of P_n .*

Proof: There is a spoke S_n of P_n that has the same slope as S . Evidently S_n converges to S . ♠

3.3 Admissible Paths and Vertex Pairs

Here we give a quick illustration of the utility of smoothly approximating nice polygons by ovals. (We will give somewhat more substantial applications in later sections of this chapter.)

Lemma 3.4 *Any two spokes of a nice polygon P intersect.*

Proof: The corresponding result for ovals is fairly obvious. We mentioned it above. The corresponding result for nice polygons follows by continuity. ♠

We say that a pair (v, x) of vertices of P is *admissible* if there is some vertex w such that (v, w) is a maximal pair and The clockwise path joining v to w contains x . Figure 3.2 shows a picture.




Figure 3.2: Admissible pairs of vertices.

Lemma 3.5 *Suppose that (v, w) and (w, v) are both admissible pairs. Then (v, w) is a maximal pair.*

Proof: If (v, w) is an admissible pair that is not maximal, then there are maximal pairs (v, w') and (w, v') that relate to (v, w) and (w, v) as described in the definition of admissible pairs. By construction, these two maximal pairs do not intersect. This contradicts. Lemma 3.4. ♠

The next lemma refers to admissible paths, which we defined in §2.3.

Lemma 3.6 *A pair of vertices (v, w) is admissible if and only if v and w respectively are the starting and endpoint points of an admissible path.*

Proof: Each maximal pair of vertices is clearly admissible. These correspond to single spokes. Conversely, each admissible path of length 1 is an oriented spoke and hence corresponds to a maximal pair of vertices. If we start with an admissible path $a \rightarrow b$ and minimally lengthen it to the new admissible path $a \rightarrow b'$, the new endpoint is a vertex adjacent to, and counterclockwise from, the old endpoint. At the same time, the admissible vertices are obtained from the maximal vertices by moving the endpoints counterclockwise. Our result follows from these facts and from induction. ♠

3.4 Interpolation Properties of Spokes

Suppose that S_a and S_b are two spokes of P . The indices correspond to the ordering on the spokes that comes from their slopes. We take indices mod n and arrange that $a < b$. Let $S_a(1)$ and $S_a(2)$ be the two vertices of S_a . Let $S_b(1)$ and $S_b(2)$ be the two vertices of S_b . We choose these labels so the vertices

$$S_a(1); \quad S_b(1); \quad S_a(2); \quad S_b(2)$$

are counterclockwise cyclically ordered, as in Figure 3.3.




Figure 3.3: Two spokes and two arcs.

We distinguish 2 arcs of P . The first one, $V(a, b, 1)$, connects $S_a(1)$ to $S_b(1)$ while avoiding $S_a(2)$ and $S_b(2)$. We define $V(a, b, 2)$ the same way, with the roles of (1) and (2) interchanged. These two paths are highlighted in grey in Figure 3.3. The way we have assigned spokes to edges in the previous section is not symmetric with respect to order reversal. This subtle asymmetry makes itself felt in the “lopsided” condition $a \leq j < b$ in the next result.

Lemma 3.7 *Suppose that $a \leq j < b$. Then $e_j \subset V(a, b, 1) \cup V(a, b, 2)$. Here e_j is the edge associated to the spoke S_j .*

Proof: The same kind of limiting argument as in Lemma 3.4 shows that the endpoints of S_j lie in $V(a, b, 1) \cup V(a, b, 2)$ for each relevant j . This lemma fails only if an endpoint of S_j is also an endpoint of $V(a, b, 1) \cup V(a, b, 2)$ and the edge e_j sticks out the side. Figure 3.4 shows one of the two possibilities.




Figure 3.4: The edge sticking out the side.

As j increases to b , the spoke S_j turns counterclockwise, but its endpoints cannot leave $V(a, b, 1) \cup V(a, b, 2)$. One of the endpoints is already stuck at the right endpoint of $V(a, b, 1)$ and so this endpoint cannot move at all. The only possibility is that the spokes S_{j+1}, \dots, S_b all share the endpoint $S_b(1)$. In particular, we may assume that $j = b - 1$. In this case, the two unequal endpoints of S_j and S_b are the endpoints of an edge, say f . Since S_j and S_b are consecutive, it follows from Lemma 3.1 that there is a minimal strip whose boundary contains $\partial S_j \cup \partial S_b$. One of the sides of this strip must contain f . But then, the clockwise correspondence assigns S_j to f . On the other hand, the clockwise correspondence assigns S_j to e_j and $e_j \neq f$. This contradicts Lemma 3.2, which says that the clockwise correspondence is a bijection. ♠

Let $p_1, p_2, p_3 \in \mathbf{R}^2 - P$ be a portion of an outer billiards orbit. We say that the triple (p_2, S_a, S_b) is *harmonious* if the ray $\overrightarrow{p_1 p_2}$ is tangent to P at an endpoint of S_a and the ray $\overrightarrow{p_2 p_3}$ is tangent to P at an endpoint of S_b . Figure 3.5 shows the situation.




Figure 3.5: A harmonious triple.

Figure 3.5 also shows a distinguished arc $A(p)$ having endpoints $S_a(1)$ and $S_b(2)$. We say that this $A(p)$ is the arc *subtended* p . The arc $A(p)$ consists of those points $q \in P$ such that the line segment joining q to p only intersects P at q . Say that two closed arcs of P are *almost disjoint* if they share an endpoint.

Lemma 3.8 *Suppose (p, S_a, S_b) are a harmonious triple. Then $A(p)$ is almost disjoint from $V(a, b, 1)$ and $V(a, b, 2)$.*

Proof: All the same definitions make sense for ovals, and the result there is obvious. The result now follows for nice polygons by continuity. ♠

3.5 The Pinwheel Orientation

We define a canonical orientation on the spokes of a nice n -gon P . Each spoke S_j corresponds to a unique pinwheel strip Σ_j in the sense that the vector V_j associated to Σ_j points from one endpoint of S_j to the other. We orient S_j by saying that $\text{oriented}(S_j) = V_j$,

We say that a spoke S_j is *special* if the three spokes S_{j-1}, S_j, S_{j+1} share a common vertex, and otherwise *ordinary*. For instance, in Figure 2.2 the spokes S_1 and S_3 are special and the rest are ordinary.

Lemma 3.9 *Let γ be an admissible path of length at least 3. The orientation on the spokes of γ induced by the orientation on γ coincides with the pinwheel orientation on all edges but the last one. On the last spoke, the two orientations agree if and only if the spoke is ordinary.*

Proof: Consider a subpath of γ that contains two spokes, say, S_1 and S_b . Here $b \in \{2, 3, 4, \dots\}$. Figure 3.6 shows the possible pictures, depending on b .




Figure 3.6: The possibilities for $b = 2, 3, 4$.

In all cases, it follows from Lemma 3.1 that there exists a minimal strip Σ' such that $\partial\Sigma'$ contains $\partial S_1 \cup \partial S_2$. In particular, $\partial\Sigma'$ contains the edge e shown in Figure 3.6. But then the strip Σ that is twice as wide as Σ' is a pinwheel strip containing e in its boundary. Comparing Figures 1.2 and 3.6, we see that the orientation given to S_1 by γ coincides with the pinwheel orientation on S_1 .

It remains to describe what happens for the last spoke of γ . The argument we just gave never uses the third property satisfied by an admissible path. In case the last spoke of γ is ordinary, we can prolong γ by two more spokes in such a way that the longer path γ' satisfies the first two properties for an admissible path. The previous argument then shows that the pinwheel

orientation and the γ -induced orientations coincide the third-to-last spoke of γ' , which is the last spoke of γ .




Figure 3.7: Ending on a special spoke.

Suppose that the last spoke of γ is special. Just for the sake of argument, let's suppose that the last spoke of γ is labelled S_2 . Figure 3.7 shows the picture of the last two spokes of γ when S_2 is special. In this case $e \in \Sigma_2$. Comparing Figures 1.2 and 3.7 we see this time that the orientation on S_2 given by γ is $-V_j$, as claimed. ♠

We also record the following result.

Lemma 3.10 *The edges e_0 and e_1 are adjacent if and only if S_1 is a special spoke.*

Proof: If S_1 is a special spoke, then the spokes S_0, S_1, S_2 all share a common point. The other three points are the vertices of the arc $e_0 \cup e_1$. Hence e_0 and e_1 are adjacent. Conversely, if e_0 and e_1 are adjacent, then it follows from Lemma 3.1 that e_0 is the edge connecting two points of $\partial S_0 \cup \partial S_1$ and e_1 is the edge connecting two points of $S_1 \cup S_2$. But then S_0, S_1, S_2 share a common vertex. Hence S_1 is special. ♠

4 The Forward Partition

4.1 The Combinatorics of the Partition

We say that a *forward wall* is a line segment or ray that lies in the closure of 2 tiles of the forward partition. We say that a forward wall W is *primary* if the outer billiards map is not defined on interior points of W . We say that W is *secondary* if W is not primary. For a point in the interior of a secondary wall, the outer billiards map is defined but the next iterate is not defined. Put another way, the secondary walls are obtained by pulling back the union of the primary walls under the outer billiards map. Figure 4.1 shows pattern for a pentagon. The primary walls divide $\mathbf{R}^2 - P$ into n distinct *primary cones*.




Figure 4.1: Primary walls

Each secondary wall is either a line segment or a ray.

Lemma 4.1 *An endpoint of a secondary wall lies on a primary wall and cannot be the vertex of the cone that contains it.*

Proof: Let p_0 be the endpoint of a secondary wall w_0 . Let p_1 be the image of p_0 under the outer billiards map. If p_0 is in the interior of a cone then p_1 must lie on an interior point of a primary wall w_1 . Pulling back a neighborhood of p_1 in w_1 , we see that w_0 contains a 2-sided neighborhood of p_0 . Hence p_0 is not an endpoint of w_0 . The second statement follows from the fact that the square-outer billiards map is clearly defined in a neighborhood of the vertex of a cone. ♠

Lemma 4.2 *No two secondary walls intersect.*

Proof: Suppose that $p_0 \in w_1 \cap w_2$ is a point of intersection between two secondary walls. We can approximate p_0 by a sequence $\{p_0(k, n)\}$ of points in w_k for $k = 1, 2$. We can arrange that both sequences lie in the interior of the same cone. Hence, the outer billiards map is defined on both sequences. Let $p_1(k, n)$ denote the image of $p_0(k, n)$ under the outer billiards map. Since the outer billiards map is an isometry on the interior of each cone, the two sequences $\{p_1(k, n)\}$ converge to the same point q_1 . By construction, q_1 lies on two different primary walls. But the primary walls only intersect at the vertices of P . Hence q_1 is a vertex of P . Given the definition of the outer billiards map, p_0 must be the same vertex of P . This is ruled out by the previous lemma. ♠

According to the lemmas above, the secondary walls cut across the cones like crosshatching, as shown schematically in Figure 4.2. Every bounded tile is either a triangle or a quadrilateral. The triangular bounded tiles contain a neighborhood of a cone apex. The quadrilateral bounded tiles have one pair of opposite sides that are primary walls and another pair of opposite sides that are secondary walls. As we justify in the proof of Lemma 4.3 below, each secondary wall is parallel to an edge of P , and the pattern is as indicated by the labels in Figure 4.2. Compare Figure 2.2.




Figure 4.2: Secondary walls.

4.2 Proof of Lemma 2.2

Just for this section we say that the pair (v, w) of vertices is a *label* if it labels some nonempty tile in the forward partition.

Lemma 4.3 *Every maximal pair of vertices is a label.*

Proof: Let (v, w) be a maximal pair, and let Σ be the corresponding pin-wheel strip. We rotate so that Σ is horizontal, as in Figures 4.3 and 4.4 below. If $p \in \mathbf{R}^2$ lies just below L and far to the right, then

$$\psi(p) = p + V. \tag{5}$$




Figure 4.3: (w, v) is a label.

If p lies just above L' and far to the left, the map ψ has the form

$$\psi(p) = p - V. \tag{6}$$




Figure 4.4: (v, w) is a label.

It is a consequence of Equations 5 and 6 that there is an unbounded region labelled by (w, v) and an unbounded region labelled by (v, w) . ♠

Now we show that the set of labels coincides with the set of admissible pairs. Let's first show that such any admissible pair (v, x) is a label. First

choose a point $p_0 \in \mathbf{R}^2 - P$ in the unbounded tile labelled by (v, w) . Consider the outer billiards orbit p_0, p_1, p_2 , with $p_2 = \psi(p_0)$ as usual. The situation is shown in Figure 4.4.




Figure 4.5: (v, x) is a label.

Let $p_1(t)$ denote a point that varies along the line segment connecting p_1 to v . Let $p_0(t)$ and $p_2(t)$ be the backwards and forwards images of $p_1(t)$ under the outer billiards map. As $p_1(t)$ varies from p_1 to v , the moving tangent line from $p_1(t)$ to P encounters every vertex along the oriented path from v to w . Hence, there is some value of t for which $p_0(t)$ lies in a region labelled by (v, x) . Hence (v, x) is a label. This argument, incidentally, justifies the comments about the secondary walls mentioned at the end of the last section.

To show that all the labels are admissible, we observe that we can reverse the process we just described, starting with a point p_0 in the region labelled by (v, x) and then moving outward along the ray pointing from v to p_0 until we arrive at a point that lies in a region labelled by the maximal pair (v, w) . This shows that (v, x) is admissible.

Combining the result here with Lemma 3.6, we establish Lemma 2.2.

4.3 Some Geometric Estimates

The bounded tiles in a single cone C all have a common boundary with one of the strips Σ , but are disjoint from that strip. If we rotate so that Σ is horizontal, as in Figure 4.6, then the shaded region T beneath Σ is the union of the bounded tiles in C .




Figure 4.6: Bounding the tiles.

Lemma 4.4 *No point of T is farther from $\partial\Sigma$ than half the width of Σ .*

Proof: The boundary wall w of T opposite the cone apex v_0 is the largest bounded secondary wall in C . The secondary wall w is parallel to a line L which extends one of the sides of P . The fact that w is bounded gives control on the slope of L , and the basic geometric fact is that the point p_1 necessarily lies in the lower half of Σ . On the other hand, p_0 and p_1 are equidistant from the lower boundary of Σ . Here p_1 is the image of p_0 under the outer billiards map. But this shows that the distance from p_0 to the lower boundary of Σ is less than half the width of Σ . Since p_0 is the extreme point of T , the same statement can be made for any point of T . ♠

Let e_1, \dots, e_n be the edges of P . Let $\partial_1 \Sigma_j$ be the component of $\partial \Sigma_j$ that contains e_j . Let $\partial_2 \Sigma_j$ be the other component.

Lemma 4.5 *If e_0 is adjacent to e_1 then $T(1 \rightarrow b)$ and P lie on the same side of $\partial_1 \Sigma_0$. Otherwise, $T(1 \rightarrow b)$ and P lie on the same side of $\partial_2 \Sigma_0$.*

Proof: If e_0 and e_1 are adjacent, then $e_0 = f_0$ in Figure 4.7. The lines extending e_0 and e_1 contain the boundary of the primary cone that contains $T(a \rightarrow b)$. This is the shaded region in Figure 4.7. In this case, the result is obvious. In the second case, $e_0 = g_0$, and we want to rule out the possibility that $T(a \rightarrow b)$ intersects the black cone. But one checks easily that points lie in the black cone lie in tiles labelled by vertex pairs (x, y) where y lies in the interior of the clockwise arc connecting the two endpoints of the spoke S_a . This is a contradiction. ♠




Figure 4.7: The excluded region.

Recall that the vectors V_1, \dots, V_n are associated to the pinwheel strips $\Sigma_1, \dots, \Sigma_n$. In our proof of Theorem 2.4 we will be interested in sets such as

$$T(a \rightarrow b; a) = T(a \rightarrow b) + 2V_a. \quad (7)$$

Corollary 4.6 *If e_0 is adjacent to e_1 then the line $\partial_2\Sigma_0$ separates P from $T(1 \rightarrow b; 1)$. Otherwise, the line $\partial_1\Sigma_0$ separates P from $T(1 \rightarrow b; 1)$.*

Proof: Suppose that e_0 and e_1 are adjacent. Since e_0 and e_1 are adjacent, the vector V_1 joins a point on $\partial_1\Sigma_0$ to a point on the centerline of Σ_0 . Hence

$$\partial_1\Sigma_0 + 2V_1 = \partial_2\Sigma_0.$$

The first case of this lemma now follows from Lemma 4.5. The point is that adding V_1 ejects $T(1 \rightarrow b)$ outside of Σ_0 , and onto the correct side.

Suppose that e_0 and e_1 are not adjacent. This time V_1 joins a point on the centerline of Σ_0 to a point on $\partial_2\Sigma_0$. Hence $\partial_2\Sigma_0 + 2V_1 = \partial_1\Sigma_0$. The rest of the proof is the same in this case. ♠

Recall from §3.9 that a spoke S_a is *special* if the spokes S_{a-1}, S_a, S_{a+1} share a common vertex. Otherwise, S_a is ordinary. Combining Lemma 4.6 with Lemma 3.10 and relabelling, we have the following result.

Corollary 4.7 *If S_a is a special spoke, then the line $\partial_2\Sigma_{a-1}$ separates P from $T(a \rightarrow b; a)$. Otherwise, the line $\partial_1\Sigma_{a-1}$ separates P from $T(a \rightarrow b; a)$.*

The last result we prove is not needed anywhere in this paper. We include it because we think gives a nice piece of extra information about the geometry of the partition.

Lemma 4.8 *A partition tile T is unbounded if and only if $\psi(T) \cap T \neq \emptyset$.*

Proof: If T is unbounded, then T contains an infinite cone. Being a uniformly bounded piecewise translation, ψ cannot map an infinite cone off itself. It remains to show that $\psi(T) \cap T \neq \emptyset$ implies that T is unbounded. Let (v, w) be the pair of vertices labelling T . It suffices to prove that (v, w) is a maximal pair. Let p_0, p_1, p_2, \dots be the forwards outer billiards orbit of a point $p_0 \in T$, Suppose $p_2 = \psi(p_2) \in T$.




Figure 4.8: A special quadrilateral

By construction, $P \subset Q$, where Q is the quadrilateral with vertices v , w , p_2 , and $\overline{p_0p_1} \cap \overline{p_3p_4}$. Since v and w are opposite vertices of Q , there is an infinite strip that contains v and w on its boundary and contains the rest of Q in its interior. This strip picks out (v, w) as a maximal pair. ♠

5 Proof of Theorem 2.3

5.1 Reformulation of the Result

We can define the backwards partition for P just as we defined the forwards partition. We just use the inverse map ψ^{-1} . The map ψ sets up a bijection between the tiles in the forward partition and the tiles in the backward partition. A tile in the forward partition labelled by a pair of vertices (v, w) corresponds to a tile in the backward partition labelled by a pair of vertices (w, v) .

We change our notation so that $T_+(a \rightarrow b) = T(a \rightarrow b)$, and $T_-(a \rightarrow b)$ denotes a tile in the backwards partition corresponding to the “backwards admissible path” The backwards admissible paths have the same definition as the forwards admissible paths, except that the spokes are traced out in reverse cyclic order. In short, if we reverse a forwards admissible path, we get a backwards admissible path. From this structure, we have

$$\psi(T_+(a \rightarrow b)) = T_-(b \rightarrow a). \quad (8)$$

We define

$$T_-(b) = \bigcup_a T_-(b \rightarrow a); \quad T_+(c) = \bigcup_d T_+(c \rightarrow d). \quad (9)$$

If $c = b$ there is nothing to prove. If necessary, we add n to c so that $b < c$. In this chapter, we will prove the following result.

$$T_-(b) \cap T_+(c) \subset \Sigma_j; \quad j = b, \dots, (c - 1). \quad (10)$$

Theorem 2.3 is a quick corollary.

Proof of Theorem 2.3: Our point q in Theorem 2.3 lies $T_-(b) \cap T_+(c)$. Hence

$$q \in \Sigma_j; \quad j = b, \dots, (c - 1).$$

Hence

$$\psi^*(q, j) = (q, j + 1); \quad j = (b - 1), \dots, (c - 2).$$

Hence $(\psi^*)^k(q, b - 1) = (q, c - 1)$ for $k = b - a < n$. ♠

5.2 A Result about Strips

Let P be a nice polygon. Let Σ the pinwheel strip associated to an edge e of P . We rotate and scale so that Σ is bounded by the lines $y = 0$ and $y = 1$, as in Figure 5.2. Let v be the vertex of P farthest from the bottom edge of P . Let T be the triangle formed by lines extending the two edges of P incident to v and the top boundary component of Σ .

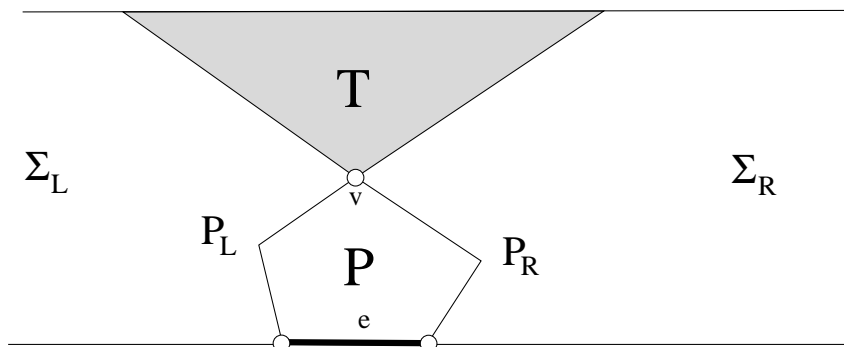


Figure 5.2: Modified Strip

Let P_L and P_R respectively be the closures of the left and right halves of $\Sigma - P - T$. Let Σ_L and Σ_R denote the left and right halves of $\Sigma - P - T$.

Let $\{p_i\}$ be an outer billiards orbit, with $p_{i+2} = \psi(p_i)$. Let $m_i \in P$ be the midpoint of the segment joining p_i to p_{i+1} .

Lemma 5.1 *Suppose that $q_1, q_2 \in P_L$ and the clockwise arc from q_2 to q_3 does not contain e . Then $p_2 \in \Sigma_L$.*

Proof: We rotate and scale so that Σ is bounded by the horizontal lines $y = 0$ and $y = 1$ and $v = (0, 1/2)$. Let R denote the ray that starts at v and extends the left edge of T . The conditions $q_1, q_2 \in P_L$ force p_2 to lie in the infinite sector Σ'_L bounded by P_L , by R , and by the negative x -axis. Σ_L is exactly the portion of Σ that lies below the line $y = 1$. If $p_2 \in \Sigma'_L - \Sigma_L$ then $q_2 = v$ and p_3 lies below the line $y = 0$. But then the clockwise arc from q_2 to q_3 contains e . This contradiction shows that $p_2 \in \Sigma_L$. ♠

Lemma 5.2 *Suppose that $q_1, q_2 \in P_R$ and the counterclockwise arc from q_1 to q_0 does not contain e . Then $p_2 \in \Sigma_R$.*

Proof: This follows from the previous result and reflection symmetry. ♠

5.3 The End of the Proof

We label the vertices of P so that $b < c$. Let

$$p \in T_-(b) \cap T_+(c); \quad j \in \{a, \dots, (b-1)\}. \quad (11)$$

Our goal is to show that $p \in \Sigma_j$. The sets of interest so far, such as P_L and P_R , all depend on an edge e_j , and we sometimes write $P_L(e_j)$ and $P_R(e_j)$ to denote this dependence. We will establish the stronger result that

$$p \in \Sigma_L(e_j) \cup \Sigma_R(e_j). \quad (12)$$

We set $p_2 = p$ and we let $\{p_k\}$ be the outer billiards orbit of p . We let q_k be the midpoint of the line segment joining p_k to p_{k+1} , as in the previous section. The next lemma refers to a definition made in §3.4.

Lemma 5.3 *The triple (p_2, S_b, S_c) is a harmonious triple.*

Proof: The point p_2 lies in some tile $T_-(b, a)$. This means that b is the starting point of an admissible path whose first spoke is S_b . Since (b, a) labels the tile $T_-(b, a)$ in the backwards partition, the line connecting p_1 to p_2 is tangent to P at an endpoint of S_b . In short q_1 is an endpoint of S_a . Similarly, p_2 is in some tile $T_+(c, d)$. The same argument shows that q_2 is an endpoint of S_b . ♠

We will treat the case when the spokes S_b and S_c do not share a vertex, though this is mainly for the sake of drawing one picture rather than two. When S_b and S_c share a vertex, the proof is essentially identical.

Let $X = S_b \cap S_c$. In the case we are considering, X is a point interior to both S_b and S_c . We rotate so that S_b has positive slope and S_c has negative slope.

Let $\Sigma = \Sigma_j$ for some relevant index j . Let e_j be the edge of P contained in $\partial\Sigma_j$. We can further rotate P , if necessary, so that $e = e_j$ lies below the intersection $S_b \cap S_c$. Figure 5.3 shows the situation.

Referring to Lemma 3.7, we let $V(b, c, 1)$ denote the vertex set on the bottom of the picture and $V(b, c, 2)$ the vertex set on the top. We will treat the case when p_2 lies to the left of X . This case relies on Lemma 5.1. The case when p_2 lies to the right of X has practically the same proof, and relies on Lemma 5.2 instead.

Let H_b and H_c respectively denote the right halfplanes bounded by S_b and S_c . Figure 5.3 shows a circle in place of a nice polygon. This is our attempt to draw a somewhat schematic picture of what is going on.

1. Combining Lemma 5.3 and Lemma 3.8, we see that $q_1 = S_b(1)$ and $q_2 = S_c(2)$ both lie to the left of X .
2. By Lemma 3.7, we have $e_j \subset \text{closure}(H_b - H_c)$.
3. By Lemma 3.7 again, $v_j \subset H_c$. Here v_j is the vertex farthest from e_j .
4. By Steps 3 and 4, we have $A(p_2) \subset P_L(e_j)$. See §3.4 for a definition of $A(P_2)$. Hence $q_1, q_2 \in P_L(e_j)$.
5. The ordered pair (q_2, q_3) is an admissible pair. Hence, $q_3 \subset H_c$.
6. Combining Steps 2 and 5, we conclude that the clockwise arc connecting q_2 to q_3 does not contain e_j .
7. Combining Steps 4 and 6 and Lemma 5.1, we see that $p_2 \subset \Sigma_L(e_j)$, as desired. This completes the proof.




Figure 5.3: Case 1

6 Theorem 2.4 modulo a detail

6.1 Reformulation of the Result

The unbounded regions correspond to the admissible paths of length 1. When $a = b$, the result is a tautology. So, we assume that $a \neq b$. The admissible path in question has length at least 3. We are going to prove Theorem 2.4 by induction. As previously, we replace b by $b + n$ if necessary so that $a < b$.

The admissible path $a \rightarrow b$ does not necessarily involve all the spokes between a and b . If S_i is not involved in the path $a \rightarrow b$, let $W_i = 0$. Otherwise, let $W(a, b, i)$ be the vector that points from the first endpoint of S_i to the last one. By Lemma 3.9, we have $W_i = V_i$ for $i < b$. For $i = b$ we have $W_i = \pm V_i$, where the (+) option is taken if and only if S_b is an ordinary spoke. The vector W_i depends (mildly) on $a \rightarrow b$, but we suppress this from our notation.

Lemma 6.1 *For any $p \in T(a \rightarrow b)$ we have*

$$\psi(p) - p = \sum_{i=1}^b 2W_i.$$

Proof: Let (v, w) be the admissible pair of vertices associated to $T(a \rightarrow b)$. Recall that a and b respectively name the first and last spoke of the admissible path associated to $T(a \rightarrow b)$ whereas v and w respectively name the first endpoint of the path and the last endpoint of the path. The path $a \rightarrow b$ simply traces out the involved spokes. By definition

$$w - v = \sum_{i=a}^b W_i.$$

At the same time $\psi(p) - p = 2(w - v)$. Putting these two equations together gives the lemma. ♠

Lemma 6.1 establishes an identity between certain multiples of the vectors involved in the relevant strip maps. This is a good start. What connects the result in Lemma 6.1 to the pinwheel map is the claim that the multiples involved are precisely the ones that arise in the relevant strip maps. This amounts to showing that certain translates of the tile $T(a \rightarrow b)$ lie in the

right place with respect to the relevant strips. Here is the main construction in our proof. Define

$$T(a \rightarrow b; k) = T(a \rightarrow b) + \sum_{i=a}^k 2W_i; \quad k = a, \dots, b \quad (13)$$

The sets $T(a \rightarrow b; k)$ for $i = a, \dots, b$ are translates of $T(a \rightarrow b)$.

Recall that μ_b is the local strip map with index b . In this chapter we will prove

Lemma 6.2 $\mu_b(p) = p + 2W_b$ for all $p \in T(a \rightarrow b, b - 1)$.

In the next chapter, we will prove

Lemma 6.3 $T(a \rightarrow b; k) \subset \Sigma_k$ for all $k = a, \dots, (b - 1)$.

Remarks:

(i) Lemma 6.3 is the much more interesting of the two results. It turns out that Lemma 6.2 is just a disguised version of Corollary 4.7.

(ii) It would be easier to say simply that $T(a \rightarrow b; k) \subset \Sigma_k$ for all $k = a, \dots, b$, but this is not true. Lemma 6.2 makes the strongest statement we can make.

Proof of Theorem 2.4: We relabel so that $a = 1$. Let

$$T_0 = T(1 \rightarrow b); \quad T_k = T(1 \rightarrow b; k). \quad (14)$$

Let $p = p_0 \in T_0$ be an arbitrary point. Define

$$p_k = p_0 + \sum_{i=1}^k 2W_i; \quad k = 1, \dots, b. \quad (15)$$

This is just a pointwise version of Equation 13.

By Lemma 6.1, we have

$$\psi(p_0) = p_b. \quad (16)$$

Choose any $k = 0, 1, \dots, b - 2$ and consider the pair (p_k, k) . There are two cases to consider. Suppose first that the index k is involved in the path $1 \rightarrow b$. Then

$$p_k + 2W_k \in \Sigma_k, \quad (17)$$

by Lemma 6.3. Therefore

$$\mu_{k+1}(p_k) = p_k + 2W_k = p_{k+1}; \quad \mu_{k+1}(p_{k+1}) = p_{k+1}. \quad (18)$$

Hence

$$(\psi^*)^2(p_k, k) = \psi^* \circ \mu_{k+1}(p_k, k) = \psi^*(p_{k+1}, k) = (p_{k+1}, k+1). \quad (19)$$

Now suppose that the index k is not involved in the path $1 \rightarrow b$. Then $p_{k+1} = p_k$ and this common point lies in Σ_{k+1} . Hence

$$\psi^*(p_k, k) = (p_k, k+1) = (p_{k+1}, k+1). \quad (20)$$

In either case, we see that

$$(p_{k+1}, k+1) = (\psi^*)^e(p_k, k) \quad (21)$$

for some exponent $e = e_k$. Applying this argument for as long as we can, we see that

$$(p_{b-1}, b-1) = (\psi^*)^e(p_0, 0), \quad (22)$$

for some exponent e . Finally, by Lemma 6.2, we have

$$\psi^*(p_{b-1}, b-1) = (\mu_b^*(p_{b-1}), b-1) = (p_{b-1} + W_b, b-1) = (p_b, b-1). \quad (23)$$

Hence $(p_b, b-1)$ is in the forward ψ^* -orbit of $(p_0, 0)$. Combining this information with Equation 16, we see that there is some positive $k < 2n$ such that

$$(\psi(p), b-1) = (p_b, b-1) = (\psi^*)^k(p, 0). \quad (24)$$

This completes the proof. ♠

The rest of this chapter is devoted to the proof of Lemma 6.2. The material in the next section will also be used in §7.

6.2 The Conjugate Polygon

As in the proof of Theorem 2.3, we find it useful here to consider both the forward and backward partition of P . Since our labelling conventions have been geared towards the forward partition, we find it useful to set things up in a way that doesn't require us to deal directly with the backward partition.

Let R denote reflection in the x -axis. For any subset $A \subset \mathbf{R}^2$, let $\overline{A} = R(A)$. One could equally well think of R as complex conjugation.

A basic and easy fact is that R maps the backward partition of P to the forward partition of \overline{P} . For that reason, we will consider the forward partitions of P and \overline{P} at the same time. We rotate so that one edge of P is horizontal and lies in the x -axis. We call this edge e_0 . The associated pinwheel strip is Σ_0 and the associated spoke is S_0 . We also arrange that P lies in the upper half plane, so that e_0 is the bottom edge. The labellings of the remaining objects of P are then forced by our earlier conventions. To emphasize the dependence on P , we write $e_0(P)$, etc. We let $e_0(\overline{P})$ be the horizontal edge of \overline{P} .

The cyclic ordering forces two sets of equations.

$$\overline{e_k(P)} = e_{-k}(\overline{P}); \quad \overline{\Sigma_k(P)} = \Sigma_{-k}(\overline{P}); \quad (25)$$

$$\overline{S_k(P)} = S_{1-k}(\overline{P}); \quad \overline{V_k(P)} = V_{1-k}(\overline{P}). \quad (26)$$

This second set of equations is more subtle. Figure 6.3 illustrates why $\overline{S_0(P)} = S_1(\overline{P})$. There are several possible pictures, depending on the geometry of P , and we have picked one of the possibilities. The other possibilities have the same outcome. The rest of Equation 26 is then forced by the cyclic ordering.




Figure 6.3: The polygon and its conjugate

We can define the sets $T(a \rightarrow b; k)$ relative to \overline{P} just as well as for P . We tack a P or \overline{P} on the end of our notation to indicate which polygon we mean.

Lemma 6.4 $\overline{T(a \rightarrow b; k; P)} = T(1 - b, 1 - a; -k; \overline{P})$.

Proof: We've already remarked that ψ sets up a bijection between the tiles in the forward partition of P to the tiles in the backward partition. Briefly using the notation from the previous chapter, we have

$$\psi(T_+(a \rightarrow b)) = T_-(b \rightarrow a). \quad (27)$$

The path $b \rightarrow a$ is the same as the path $a \rightarrow b$ but it is given the opposite orientation. We call this the *reversal property*. We will use it below.

The composition $R \circ \psi$ carries the tiles in the forward partition of P to the tiles in the forward partition of \overline{P} . Combining Equations 26 and 27, we see that

$$R \circ \psi(T(a \rightarrow b; P)) = T(1 - b \rightarrow 1 - a; \overline{P}). \quad (28)$$

Note that

$$\psi(T(a \rightarrow b; P)) = T(a \rightarrow b; b; P). \quad (29)$$

Combining Equations 29 and 28, we have

$$\overline{T(a \rightarrow b; b; P)} = T(1 - b, 1 - a; \overline{P}). \quad (30)$$

By the reversal property and Equation 26, we have

$$\overline{W_k(P)} = -W_{1-k}(\overline{P}). \quad (31)$$

To consider the case $k = b$ in detail, we have

$$\begin{aligned} \overline{T(a \rightarrow b; b - 1; P)} &= \overline{T(a \rightarrow b; b; P)} - \overline{W_b(P)} = \\ &T(1 - b \rightarrow 1 - a; \overline{P}) + W_{1-b}(\overline{P}) = T(1 - b, 1 - a; 1 - b; \overline{P}). \end{aligned}$$

in short

$$\overline{T(a \rightarrow b; b - 1; P)} = T(1 - b, 1 - a; 1 - b; \overline{P}). \quad (32)$$

Applying the same argument, inductively, to each of the vectors $W_{b-1}(P)$, $W_{b-2}(P)$, etc., we establish the lemma. ♠

6.3 Proof of Lemma 6.2

We are going to apply Corollary 4.7 to \overline{P} . There are two cases for us to consider, depending on whether or not the spoke S_b is ordinary. We'll first consider the case when S_b is ordinary.

It is convenient to set

$$\alpha = 1 - b; \quad \beta = 1 - a. \quad (33)$$

By Lemma 6.4,

$$\overline{T(a \rightarrow b, b - 1; P)} = T(\alpha \rightarrow \beta; \alpha; \overline{P}). \quad (34)$$

By definition and Lemma 3.9,

$$T(\alpha \rightarrow \beta; \alpha; \overline{P}) = T(\alpha \rightarrow \beta; \overline{P}) + V_\alpha(\overline{P}). \quad (35)$$

By Corollary 4.7, the line

$$\partial_1 \Sigma_{\alpha-1}(\overline{P})$$

separates $T(\alpha \rightarrow \beta; \overline{P}) + V_\alpha(\overline{P})$ from \overline{P} . Applying the reflection R and using Equation 26, we see that the line

$$\partial_1 \Sigma_b(P)$$

separates $T(a, b; b - 1; P)$ from P . But then

$$\mu_b(p) = p + V_b$$

for any $p \in T(1 \rightarrow b; b - 1; P)$. Since S_b is an ordinary spoke, $V_b = W_b$ by Lemma 3.9. Now we know that $\mu_b(p) = p + W_b$, as desired.

When S_b is special, the proof is the same except for some sign changes. This time, the line

$$\partial_2 \Sigma_b(P)$$

separates $T(a, b; b - 1; P)$ from P . But then

$$\mu_b(p) = p - V_b$$

for any $p \in T(1 \rightarrow b; b - 1; P)$. This time $-V_b = W_b$, and we get the same result as in the previous case.

This completes the proof of Lemma 6.2.

7 Proof of Lemma 6.3

7.1 Some Combinatorial Definitions

Abstract Admissible Paths: We say that an *abstract admissible path* is a finite tree τ with the following structure. First, there is a distinguished maximal path γ in τ having odd length at least 3. Every other edge of τ is incident to γ . We call the edges of $\tau - \gamma$ *special*. We call the edges of γ *ordinary*, except perhaps for the first and last edge of γ . The first and last edges of γ can be either special or ordinary. We draw γ as a zig-zag with lines of alternating negative and positive slope. We insist that every edge of τ intersects x -axis.

We orient γ from left to right. We draw the ordinary edges with thick lines and the special edges with thin lines. Figure 7.1 shows an example. A dotted line represents the x -axis. γ is the path 13678. Here, the first edge of γ is ordinary and the last edge is special. The *initial vertex* is the left endpoint of the γ . The *final edge* is the last edge of γ .




Figure 7.1: Abstract admissible paths.

Linear Order: There is a natural linear ordering on the edges of τ , induced from the order in which they intersect the x -axis, going from left to right. The numerical labels in Figure 7.1 indicate the ordering. We see that $\tau' \subset \tau$ is a *prefix* of τ if τ' and τ share the same initial set of edges and if τ' is an abstract admissible path in its own right.

Flags and Sites: To each edge e of τ , except the last one, we assign a vertex v_e . If e is an edge of γ , then v_e is the leading vertex of e . If e is not an edge of γ , then v_e is the vertex of γ incident to e . We say that a *flag* is a

pair (e, v_e) . For instance, the flags in Figure 7.1 are

$$(1, a); (2, a); (3, b); (4, b); (5, b); (6, c); (7, d).$$

We say that a *site* is a pair (τ, f) , where f is a flag of τ .

Natural Involution: There is a natural involution R on the set of abstract admissible paths: Simply rotate the path about the origin by 180 degrees and you get another one. We call this map R . The map R carries flags of τ to flags of $R(\tau)$ in a slightly nontrivial way. We first create *reverse flags* of τ by interchanging the notion of left and right, and then we apply R to these reverse flags to get ordinary flags of $R(\tau)$. In Figure 7.1, the reverse flags are

$$(8, d); (7, c); (6, b); (5, b); (4, b); (3, a); (2, a).$$

R maps the leftmost flag of τ to the image under rotation of the leftmost reverse flag. For instance $(1, a)$ corresponds to the rotation of $(2, a)$.

Reduction: We say that the site (τ', f) is a *direct reduction* of (τ, f) if τ' is a prefix of τ . The flag f is the same in both cases. We say that (τ', f') is an *indirect reduction* of (τ, f) if (τ', f') is a direct reduction of $R(\tau, f)$. We say that one site (τ_2, f_2) a *reduction* of another site (τ_1, f_1) if (τ_2, f_2) is either a direct or an indirect reduction of (τ_1, f_1) . In this case, we write $(\tau_1, f_1) \rightarrow (\tau_2, f_2)$.

Hereditary Properties:

Let \mathcal{C} be a collection of sites. We say that \mathcal{C} is *hereditary* if it has the following properties.

- \mathcal{C} is closed under the natural involution.
- \mathcal{C} is closed under reduction.

Say that a site (τ, f) is *prototypical* if f is the first flag of τ . Let Ω be a map from $\{0, 1\}$. We say that Ω is *hereditary* if Ω has the following properties.

- Ω evaluates to 1 on all prototypical sites in \mathcal{C} .
- $\Omega \circ R = \Omega$. Here R is the natural involution.
- If $\Omega(\tau_1, f_1) = 1$ and $(\tau_2, f_2) \rightarrow (\tau_1, f_1)$ then $\Omega(\tau_2, f_2) = 1$.

7.2 The Reduction Lemma

In this section we will prove the following result.

Lemma 7.1 (Reduction) *Suppose that \mathcal{C} is a hereditary collection of sites and Ω is a hereditary function on \mathcal{C} . Then $\Omega \equiv 1$ on \mathcal{C} .*

Proof: It suffices to prove that, through the two operations of R and reduction, every site can be transformed into a prototypical site. It henceforth goes without saying that all sites belong to \mathcal{C} .

Let (τ, f) be a site. Let γ be the maximal path of τ . Let $f = (e, v)$. Either v lies in the left half of γ or the right half. (There are an even number of vertices.) Applying R if necessary, we can assume that v lies in the left half of τ . If γ has length 5 we let τ' denote the subtree of τ obtained by deleting the last two vertices of γ and all incident edges. Then τ' is a prefix of τ and $(\tau, f) \rightarrow (\tau', f)$.




Figure 7.2: Abstract admissible paths.

We just have to worry about the case when γ has length 3. Let $(\tau_1, f_1) = (\tau, f)$ and let $(\tau_2, f_2) = R(\tau, f)$. Also, let e_k be the edge of f_k for $k = 1, 2$. If e_1 is not the first edge of τ_1 then (τ_2, f_2) has the following two properties.

1. e_2 is neither of the last two edges of τ_2 .
2. At least 3 edges of τ_2 are incident to the right vertex of γ .

Figure 7.2 shows a typical situation. The thick grey lines represent e_1 and e_2 . The upshot is that after applying R , we can assume that $f = (e, v)$, where e is neither of the last two edges of τ . We let τ' be the prefix obtained by cutting off these last two edges. The second property mentioned above guarantees that τ' is a prefix of τ . Again we have $(\tau, f) \rightarrow (\tau', f)$.

In summary, the process above only stops when we reach a prototypical site. ♠

7.3 The Pinwheel Collection

Let P be a nice polygon. Each admissible path associated to P gives rise to an abstract admissible path. This path encodes the way the spokes in the path (and the skipped spokes) meet at their endpoints. Figure 7.3 shows an example. The unusual labelling on the abstract admissible path is present only to underscore the correspondence.




Figure 7.3: An admissible path and its abstraction

We call the abstract admissible path produced in this way the *abstraction* of the admissible path. The maximal path in the abstract admissible path corresponds to the actual admissible path. We let \mathcal{C} be the class of all sites (τ, f) , where τ is the abstraction of an admissible path for some nice polygon.

Lemma 7.2 \mathcal{C} is a hereditary collection.

Proof: The fact that \mathcal{C} is closed under reduction comes from the fact that we have built in the basic properties of admissible paths into our definition of abstract admissible paths. If τ is the abstraction of $a \rightarrow b$, then any prefix τ' is the abstraction of some admissible path $a \rightarrow b'$, where $b' < b$.

The analysis in §6.2 shows that \mathcal{C} is closed under the natural involution. The basic reason, as we have discussed already, is that ψ carries the forward partition to the backward partition. ψ maps the forward tile associated to the path $a \rightarrow b$ to the backward tile associated to the path $b \rightarrow a$. The abstraction of $b \rightarrow a$ is exactly the image of the abstraction of $a \rightarrow b$ under the natural involution. ♠

7.4 The Binary Function

Now we are going to define a function $\Omega : \mathcal{C} \rightarrow \{0, 1\}$. We will first consider the situation for a given nice polygon P , and then we will take into account all nice polygons at the same time.

Let (τ, f) be a site in \mathcal{C} . This means that there is a nice polygon P and an admissible path $a \rightarrow b$ such that τ is the abstraction of $a \rightarrow b$. Moreover, f is just one of the sites. The edges of τ are naturally in correspondence with the strips $\Sigma_a, \dots, \Sigma_b$. Moreover, there is a natural correspondence between the sets $T(a \rightarrow b; a), \dots, T(a \rightarrow b; b - 1)$ and the sites of τ . The correspondence is set up in such a way that each site (τ, f) corresponds to a pair

$$T(a \rightarrow b; k); \quad \Sigma_k. \quad (36)$$

All of this depends on P . We define $\Omega'(\tau, f; P) = 1$ if Lemma 6.3 is true for the pair in Equation 36. Finally, we define $\Omega(\tau, f) = 1$ if and only if $\Omega'(\tau, f; P) = 1$ for every instance in which the site (τ, f) arises.

Lemma 6.3 is equivalent to the statement that $\Omega \equiv 1$ on \mathcal{C} . Accordingly, we will prove Lemma 6.3 by showing that Ω is a hereditary function and then invoking the Hereditary Lemma.

7.5 Property 1

The prototypical sites correspond to the case $k = a$ in Equation 36. Cyclically relabelling, we take $a = 1$. The following lemma implies that $\Omega = 1$ on all prototypical sites.

Lemma 7.3 $T(1 \rightarrow b; 1) \subset \Sigma_1$.

Proof: We rotate so that Σ_1 is horizontal, and $T(1 \rightarrow b)$ lies beneath the lower boundary of Σ_1 . Let $L\Sigma_1$ and $U\Sigma_1$ denote the lower and upper boundaries of Σ_1 respectively. $T(1 \rightarrow b)$ is contained in the shaded region T shown in Figure 4.3. By Lemma 4.4, every point of $T(1 \rightarrow b)$ is closer to $L\Sigma_1$ than (half) the distance between $L\Sigma_1$ and $R\Sigma_1$. We have the following 2 properties.

- The vector W_1 points from $L\Sigma_1$ to $U\Sigma_1$.
- $L\Sigma_1$ contains the top edge of $T(1 \rightarrow b)$.

It follows from these two properties that $T(1 \rightarrow b) + W_1 \subset \Sigma_1$. ♠

7.6 Property 2

That $\Omega \circ R = \Omega$ is a consequence of the relations between P and \overline{P} worked out in §6.2. Let (τ_1, f_1) be the site corresponding to the pair in Equation 36. We label the edges of τ_1 as a, \dots, b . Let $(\tau_2, f_2) = R(\tau_1, f_1)$. We label the edges of τ_2 as $(1 - b), \dots, (1 - a)$.

We label the sites of τ_1 by symbols of the form $\langle k \rangle_1$. Here k names the label of the edge involved in the site. Likewise we label the sites of τ_2 with the label $\langle k \rangle_2$.

Lemma 7.4 *The natural involution carries $\langle k \rangle_1$ to $\langle -k \rangle_2$.*

Proof: Given that the natural involution reverses the order of the sites, it suffices to check the claim for a single site. We choose the first site, with $k = a$. The natural involution carries this site to the that involves the next-to-last edge of τ_2 . But this edge is labelled $-a$. ♠

According to this result, if the site (τ_1, f_1) corresponds to the pair

$$T(a \rightarrow b; k); \quad \Sigma_k.$$

then the site (τ_2, f_2) corresponds to the pair

$$T(1 - b \rightarrow 1 - a; -k); \quad \Sigma_{-k}.$$

But, by Lemma 6.4, reflection in the x -axis carries the one pair to the other. Hence, the desired containment holds in the one case if and only it holds in the other. In other words $\Omega(\tau_1, f_1) = \Omega(\tau_2, f_2)$. This establishes the second property.

7.7 Property 3

This property has the most interesting proof.

Recall from §4.1 that the primary walls of the forward partition divide $\mathbf{R}^2 - P$ into primary cones. The beginning vertex v of the admissible path $a \rightarrow b$ is the apex of the cone. If we consider all the tiles of the form $T(a \rightarrow b)$ with a fixed and b increasing, then Lemma 6.3 involves increasingly many containments. On the other hand, the tiles involved get smaller and smaller in the sense that they shrink down to v . One might expect that the vertex v itself satisfies all the identities we can form.

Let $a \rightarrow B$ denote the maximal admissible path that starts with S_a . Then $a \rightarrow B$ corresponds to the triangular tile $T(a \rightarrow B)$ that v as a vertex. Define

$$p_0 = v; \quad p_k = p_0 + \sum_{i=1}^k 2W_i. \quad (37)$$

Note the similarity between this equation and Equation 13. In the next section we will prove the following result. Assume this result for now.

Lemma 7.5 $p_k \subset \Sigma_k$ for $k = 1, \dots, B$.

When $k \leq b - 1$ we let $\Omega(a \rightarrow b; k)$ be statement that $T(a \rightarrow b; k) \subset \Sigma_k$.

Corollary 7.6 *Suppose $c > b$ is that that $a \rightarrow c$ is an admissible path. Then $\Omega(a \rightarrow b; k)$ implies $\Omega(a \rightarrow c, k)$.*

Proof: The admissible paths $a \rightarrow b$ and $a \rightarrow c$ agree except for possibly the last edge of $a \rightarrow b$. Hence the sets

$$T(a \rightarrow b; k); \quad T(a \rightarrow c; k); \quad p_k$$

are respectively translates, by the same vector X , of the sets

$$T(a \rightarrow b); \quad T(a \rightarrow c); \quad p_0.$$

Let $\widehat{T}(a \rightarrow b)$ denote the convex hull of $T(a \rightarrow b)$ and p_0 . Let $\widehat{T}(a \rightarrow b; k)$ denote the convex hull of $T(a \rightarrow b; k)$ and p_k . First of all,

$$T(a \rightarrow c) \subset \widehat{T}(a \rightarrow b), \quad (38)$$

by the analysis done in connection with Figure 4.2. Translating the whole picture by X , we have

$$T(a \rightarrow c; k) \subset \widehat{T}(a \rightarrow b; k) \subset \Sigma_k. \quad (39)$$

The second equality follows from the convexity of Σ_k . ♠

Corollary 7.6 is just a restatement of Property 3.

7.8 Proof of Lemma 7.5

Now we take care of the final piece of business.

We relabel so that $a = 1$. Let $v_0 = v = p_0$ and let $\{v_k\}$ be the successive vertices of our admissible path, where we only advance the point if the spoke is actually involved in the path. Put another way, the vertex v_k is incident to the spokes S_k and S_{k+1} . Figure 7.4 shows an example.




Figure 7.4: The vertices of the path

The strip Σ_k contains the three points of $\partial S_k \cup \partial S_{k+1}$ and $v_k = S_k \cap S_{k+1}$. Compare Lemma 3.1. From this structure, we see that v_k lies on the centerline of the strip Σ_k for all k .




Figure 7.5: The vertices of the path

The fixed point of H , namely v_0 , belongs to every strip. In particular, $v_0 \in \Sigma_k$. This situation forces $p_k = H(v_k) \in \Sigma_k$. ♠

References

- [D] R. Douady, *These de 3-eme cycle*, Université de Paris 7, 1982.
- [DF] D. Doljopyat and B. Fayad, *Unbounded orbits for semicircular outer billiards*, Annales Henri Poincaré, to appear.
- [DT1] F. Dogru and S. Tabachnikov, *Dual billiards*, Math. Intelligencer **26**(4):18–25 (2005).
- [DT2] F. Dogru and S. Tabachnikov, *Dual billiards in the hyperbolic plane*, Nonlinearity **15**:1051–1072 (2003).
- [G] D. Genin, *Regular and Chaotic Dynamics of Outer Billiards*, Pennsylvania State University Ph.D. thesis, State College (2005).
- [GS] E. Gutkin and N. Simanyi, *Dual polygonal billiard and necklace dynamics*, Comm. Math. Phys. **143**:431–450 (1991).
- [Ko] Kolodziej, *The antibilliard outside a polygon*, Bull. Pol. Acad. Sci. Math. **37**:163–168 (1994).
- [M1] J. Moser, *Is the solar system stable?*, Math. Intelligencer **1**:65–71 (1978).
- [M2] J. Moser, *Stable and random motions in dynamical systems, with special emphasis on celestial mechanics*, Ann. of Math. Stud. 77, Princeton University Press, Princeton, NJ (1973).
- [N] B. H. Neumann, *Sharing ham and eggs*, Summary of a Manchester Mathematics Colloquium, 25 Jan 1959, published in Iota, the Manchester University Mathematics Students' Journal.
- [S1] R. E. Schwartz, *Unbounded Orbits for Outer Billiards*, J. Mod. Dyn. **3**:371–424 (2007).
- [S2] R. E. Schwartz, *Outer Billiards on Kites*, Annals of Mathe-

ematics Studies **171**, Princeton University Press (2009)

[**T1**] S. Tabachnikov, *Geometry and billiards*, Student Mathematical Library 30, Amer. Math. Soc. (2005).

[**T2**] S. Tabachnikov, *Billiards*, Société Mathématique de France, “Panoramas et Synthèses” 1, 1995

[**VS**] F. Vivaldi and A. Shaidenko, *Global stability of a class of discontinuous dual billiards*, Comm. Math. Phys. **110**:625–640 (1987).