Outer Billiards and The Pinwheel Map

Richard Evan Schwartz *

May 12, 2011

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 $\boldsymbol{R}^2 - P$, where P is a bounded convex planar set as in Figure 1.1 below. Given $x_1 \in \boldsymbol{R}^2 - P$, 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 $x_1 \to x_2$ is called *the outer billiards map*. The map is almost everywhere defined and invertible.

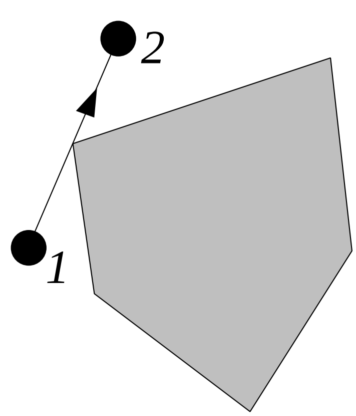

Figure 1.1: Outer billiards relative to P.

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

There has been a fair amount of work on outer billiards in the past 30 years. See, for instance, [**D**], [**DF**], [**G**], [**GS**], [**Ko**], [**M2**], [**S1**], [**S2**], [**T2**], [**VS**]. The book [**T1**] has a nice survey, though this survey misses some of the recent work.

Let ψ denote the second iterate of the outer billiards map on a convex polygon P. The purpose of this paper is to establish an equivalence between ψ and an auxilliary map ψ^* which we call the *pinwheel map*. For convenience, we work with convex polygons that have no parallel sides, and henceforth we make this assumption. The general case introduces tedious complications which we prefer to avoid. The pinwheel map is defined in terms of a certain family of strips associated to P. We first define these strips and explain their connection to ψ.

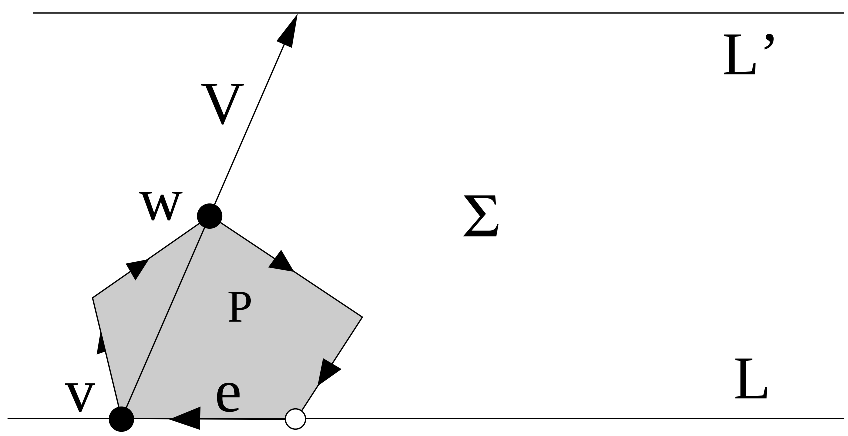

Figure 1.2: The strip associated to e.

We orient the edges of P so someone walking along an edge in the direction of the orientation would see P on the right. Given an edge e of P, we let L be the line extending e and we let L' be the line parallel to L so that the vertex w of P that lies farthest from L is equidistant from L and L'. We associate to e the pair (Σ, V), where Σ is the strip bounded by L and L', and $V = 2(w - v)$. See Figure 1.2. We call (Σ, V) a *pinwheel pair* and we call Σ a *pinwheel strip*.

We label the edges $e_1, ..., e_n$ in such a way that the lines $l_1, ..., l_n$ occur in counterclockwise cyclic order. Here l_k is the line through the origin parallel to e_k. We correspondingly label the pinwheel pairs $(\Sigma_1, V_1), ...(\Sigma_n, V_n)$. Note that these labellings do not typically correspond to the cyclic ordering of the edges around the boundary of P. See Figure 1.4 below for an example.

We work outside some large compact subset $K \subset \boldsymbol{R}^2$. Suppose we start with a point $p_1 \in \Sigma_1$. Then $\psi^k(p_1) = p_1 + kV_2$ for $k = 1, 2, 3...$ This general rule continues until we reach an an exponent k_1 such that $p_2 = \psi^{k_1}(p_1) \in \Sigma_2$. Then we have $\psi^k(p_2) = p_2 + kV_3$ for $k = 1, 2, 3$, until we reach an exponent k_2 such that $p_3 = \psi^{k_2}(p_2) \in \Sigma_3$. And so on. See Figure 1.3. We eventually reach a point $p_{n+1} \in \Sigma_1$, and the map $p_1 \to p_{n+1}$ is the first return map.

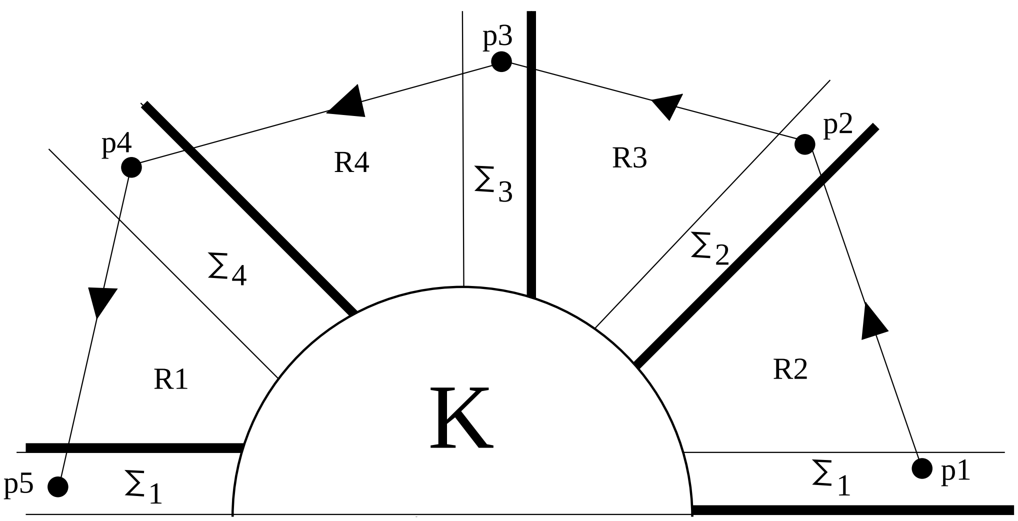

Figure 1.3: Far from the origin.

The connection between the first return map and ψ, for orbits *far away* from the polygon, appears in almost every paper on polygonal outer billiards. The pinwheel map ψ^* is a global map, which, far from the origin, carries the same information as the first return map we have already discussed.

Our main result, the Pinwheel Theorem, establishes a precise correspondence between the orbits of the ψ^* and the orbits of ψ. The correspondence is such that there is a canonical bijection between the unbounded orbits of ψ and the unbounded orbits of ψ^*. If we are interested in such questions as the existence of unbounded orbits, the Pinwheel Theorem allows us to essentially replace the outer billiards map with the pinwheel map.

The reason one might want to replace ψ by ψ^* is that ψ^* has a deep underlying structure which both simplifies and enhances the study of polygonal outer billiards. As evidence for this, we cite our work in [**S2**]. In [**S2**] we proved a version of the Pinwheel Theorem when P is a kite – i.e., a convex quadrilateral having a diagonal that is a line of symmetry. In that case, the pinwheel map has a kind of compactification as a higher dimensional polytope exchange map. See [**S2**, Master Picture Theorem]. Some version of this higher dimensional compactification of the Pinwheel Map works in general, though we have not yet worked it out. The Pinwheel Theorem, which allows one to relate outer billiards to this higher dimensional compact dynamical system, seems to be a good first step towards a general theory of polygonal outer billiards.

1.2 The Main Result

Now we will define the pinwheel map precisely. The n pinwheel strips $\Sigma_1, ..., \Sigma_n$ are ordered as above (according to their slopes). Given the pair (Σ, V), we define a map μ on $\boldsymbol{R}^2 - \partial\Sigma$ as follows.

- If $x \in \Sigma - \partial\Sigma$ then $\mu(x) = x$.
- If $x \not\in \Sigma$ 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 $\boldsymbol{R}^2_n = \boldsymbol{R}^2 \times \{1, ..., n\}$, with indices taken mod n. We define the *pinwheel map* $\psi^* : \boldsymbol{R}^2_n \to \boldsymbol{R}^2_n$ by the following conditions.

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

In other words, we try to move p by the $(j + 1)$st 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 $\pi : \boldsymbol{R}^2_n \to \boldsymbol{R}^2$ be the projection map. We introduce a map

$$\iota : \boldsymbol{R}^2 \to \boldsymbol{R}^2_n, \tag{1}$$

which we call the *pinwheel section*. This map has the property that $\pi \circ \iota$ is the identity.

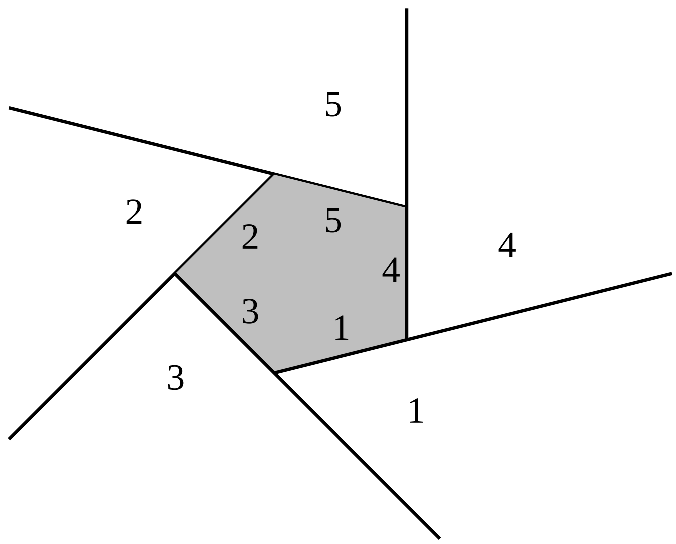

Figure 1.4: The Pinwheel Section in the case $n = 5$.

As shown in Figure 1.4 for the case $n = 5$, in general we extend the edges of the polygon and then label the regions $1, ..., n$. The region k corresponds to the edge e_k. As above, these edges are ordered according to their slopes. Our result takes its name from the pinwheel-like appearance of Figure 1.4.

If p lies in the interior of the region labelled a, we define $\iota(p) = (p, a-1)$. If p is not in the interior of any of the labelled regions, we define the second coordinate of $\iota(p)$ any way we like. For such points, the relevant maps are undefined.

Theorem 1.1 (Pinwheel) *Relative to the pinwheel section ι, we have*

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

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

Far from the origin, we have $k(p) = 1$ unless p lies in a pinwheel strip. In this case $k(p) = 2$ because one extra iterate of ψ^* is required to shift the index. Figure 1.5 shows the example of a regular pentagon, with p lying just above one of the vertices. In this case $k(p) = 3$ and the pinwheel map successively adds the vector $2A$, $2B$, and $2C$. At the same time, $\psi(p) = p + 2D$. In this case, we have the "cancellation" $A + B + C = D$. Similar cancellations always occur when $k(p) > 2$.

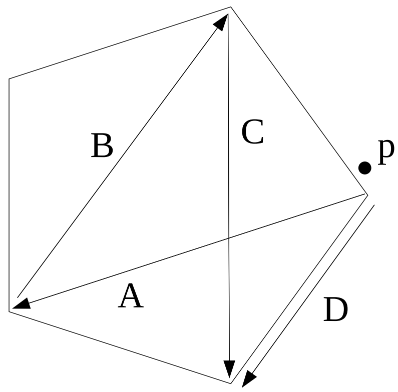

Figure 1.5: $A + B + C = D$.

The bound $k(p) \leq 3n$ is not sharp, but one can easily generalize the example in Figure 1.5 to show that $k(p) = n-2$ is attained, in some examples, for all n. We think that the bound $k(p) \leq n-2$ is sharp.

1.3 Corollaries

Just from knowledge of the relation between ψ and ψ^* far from the origin, one can conclude nothing about how the unbounded orbits of one system compare to the unbounded orbits of the other.

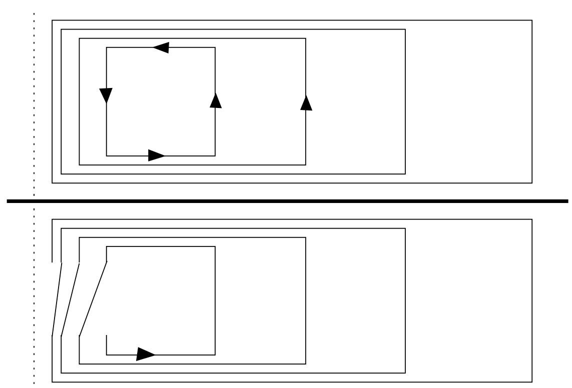

Figure 1.6: Systems which agree outside a compact set.

Figure 1.6 suggests an example of two dynamical systems which agree outside a compact set, in which one has unbounded orbits and the other does not. The top shows an infinite union of rectangles. The n-th rectangle intersects $\boldsymbol{R}$ at $-1 + 1/n$ and n. The bottom shows a spiral obtained by modifying the union of rectangles within a compact set. Both systems are defined as (say) the time one maps of the unit speed flows along the curves. We leave it to the interested reader to flesh out this example.

In light of the work in [**S2**], which exhibits unbounded outer billiards orbits that return infinitely often to a compact region of the plane, one might worry that the pinwheel map and the outer billiards map are somehow related as are the maps suggested in Figure 1.6. The Pinwheel Theorem rules out this possibility.

We say that an unbounded orbit of ψ^* is *natural* if it lies in $\iota(\boldsymbol{R}^2)$ sufficiently far from the origin.

Corollary 1.2 *Relative to any convex polygon having no parallel sides, there is a canonical bijection between the forward (respectively, backward or two-sided) unbounded ψ-orbits and the natural forward (respectively backward or two-sided) unbounded ψ^*-orbits. The bijection sending the orbit O to the orbit O^* is such that $O = \pi(O^*)$ outside a compact subset.*

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

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

Topologically, $\widehat{X}$ is the disjoint union of n strips. $\widehat{X}$ agrees with $\iota(\bigcup \Sigma_j)$ outside a compact set. We let $\widehat{\psi} : \widehat{X} \to \widehat{X}$ be the first return map of ψ^* to $\widehat{X}$.

Corollary 1.3 *There is a canonical bijection between the forward (respectively backward or two-sided) unbounded orbits of ψ and the forward (respectively backward or two-sided) unbounded orbits of $\widehat{\psi}$.*

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 one of the strips, say Σ_1, and consider the map $\widehat{\Psi} : \Sigma_1 \to \Sigma_1$. We call $\widehat{\Psi}$ the *pinwheel return map*.

At the same time, we can consider the first return map of ψ to Σ_1, as we did informally above. We call this map Ψ.

Corollary 1.4 *There is a canonical bijection between the forward (respectively backward or two-sided) unbounded orbits Ψ and the forward (respectively backward or two-sided) unbounded orbits of $\widehat{\Psi}$.*

1.4 Outline of the Paper

In §2 we prove a number of combinatorial and geometric results about constructions related to convex polygons. In §3 we prove all the results, modulo two lemmas, Lemma 3.7 and Lemma 3.8. In §4 we prove Lemma 3.8 and §5 we prove Lemma 3.7. 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.

1.5 Pinwheel Applet

My java applet, *Pinwheel*, to be found at
http://www.math.brown.edu/∼res/Pinwheel/Main.html
illustrates the Pinwheel theorem in great detail. The applet has its own summary of the result, though doesn't go so far as to give a proof. The reader might want to first play with the applet to get a general sense of what is going on visually.

1.6 Acknowledgements

I would like to thank John Smillie and Sergei Tabachnikov for helpful conversations about outer billiards and the pinwheel map. In particular, the formulation we give of the Pinwheel Theorem emerged in conversations with Smillie. Finally, I'd like to thank the anonymous referee for many helpful expository suggestions.

I would also like to thank the Clay Mathematics Institute for their support during my sabbatical two years ago, when I worked out the material in this paper.

2 Preliminaries

2.1 Basic Definitions

Throughout the chapter, P is a convex polygon with no parallel sides. Let ψ' be the outer billiards map and let $\psi = (\psi')^2$.

The Forward Partition: We have already mentioned that ψ is a piecewise translation. For almost every point $p \in \boldsymbol{R}^2 - P$, there is a pair of vertices (v_p, w_p) of P such that v_p is the midpoint of the segment connecting p and $\psi'(p)$ and w_p is the midpoint of the segment connecting $\psi'(p)$ and $\psi(p)$. In this case, we have

$$\psi(p) = p + V_p; \qquad V_p = 2(w_p - v_p). \tag{3}$$

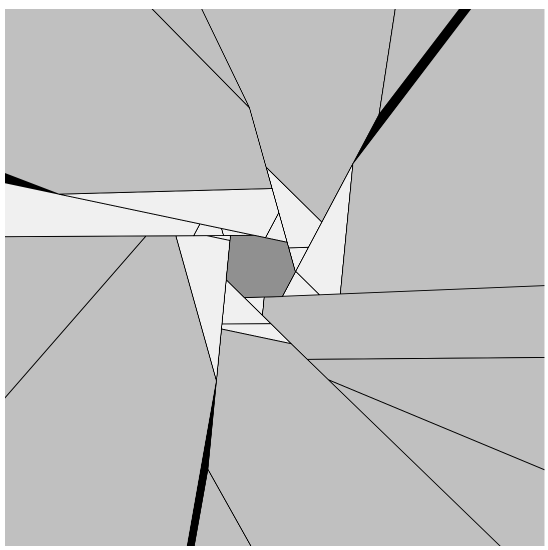

Figure 2.1: The forward partition for an octagon.

We have a partition of $\boldsymbol{R}^2 - P$ into maximal open convex polygonal sets on which the map $p \to (v_p, w_p)$ is constant. We call this partition the *forward partition* associated to P. Figure 2.1 shows an example. The white regions in Figure 2.1 are compact and the grey regions are noncompact All the grey regions, and one of the white ones, have been cut off by the bounding box. We say that the pair of vertices (v, w) *labels* a tile T in the forward partition if $v = v_p$ and $w = w_p$ for all $p \in T$.

Spokes: We have already associated n vectors to our n-gon P, namely the pinwheel vectors $V_1, ..., V_n$ defined in connection with Figure 1.2. We let S_k be the line segment that joins the head and tail of V_k. We call S_k a *spoke*. The spoke S_k is essentially an undirected version of V_k. Figure 2.2 shows an example, with some of the edges highlighted.

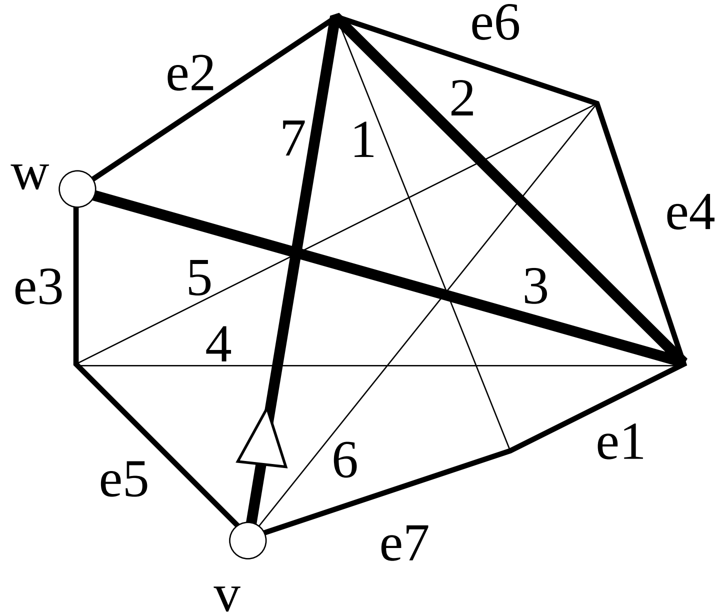

Figure 2.2: The spokes of the polygon, and the path $7 \to 3$.

Admissible Paths: Recall that the spokes are cyclically ordered according to their slopes. 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 $\boldsymbol{R} \cup \infty$ obtained by connecting the slopes of the spokes of γ. Then γ' is a proper subset of $\boldsymbol{R} \cup \infty$.

The last condition means that γ does not wrap all the way around P. We use the notation $a \to b$ to name admissible paths. The first spoke is a and the last one is b. We often take the indices mod n and use $b+n$ in place of b in case $b < a$. Thus, $7 \to 10$ is another name for the the path in Figure 2.2.

Maximal Pairs: We say that a pair of vertices (v, w) of P is *maximal* if and only they are the endpoints of a spoke. Here is an equivalent characterization: (v, w) is maximal if and only if there is an infinite strip T such that $P \subset T$ and $\partial T \cap P = \{v, w\}$. This equivalence uses the fact that P has no parallel sides.

Admissible Pairs: A pair (v, w) of vertices of P is *admissible* if there is a maximal pair (v, x) such that v, w, x occur in clockwise order on ∂P. We allow the possibility that $w = x$, so that maximal pairs count as admissible. We mention, without giving the easy proof, that if (v, w) and (w, v) are both admissible than (v, w) is maximal.

The Clockwise Correspondence: As we illustrated in Figure 1.2, we associate the pinwheel pair (Σ_j, V_j) to the edge e_j of P. We let S_j be the spoke corresponding to V_j. We call the correspondence $e_j \to S_j$ the *clockwise correspondence*. This correspondence is a bijection. (See Lemma 2.2 for a proof.) The segments S_j, S_{j+1}, and e_j form the sides of a triangle.

Pinwheel Orientation: Recall that S_j is just an undirected version of V_j. When we want to assign an orientation to S_j, we say that oriented$(S_j) = V_j$. We call this the *pinwheel orientation* on S_j.

Special and Ordinary Spokes: We say that a spoke S_2 of P is *special* if there are spokes S_1 and S_3 such that S_1, S_2, S_3 share a common vertex. We call the remaining spokes *ordinary*. In Figure 2.2, the spokes S_1 and S_3 are special and the rest are ordinary.

The Edge Set: Let $e_1, ..., e_n$ be the edges of P and let $S_1, ..., S_n$ be the corresponding spokes. We take indices mod n, so that the objects e_{a+kn} and S_{a+kn} refer to the same objects e_a and S_a for all $k \in \boldsymbol{Z}$. Suppose that $a < b < a + n$ are two integers. We define the *edge set*

$$V(a, b) = E \cup X; \qquad X = \bigcup_{i=a}^{b-1} e_i; \qquad X = S_a \cap S_b \cap \partial P. \tag{4}$$

X is empty unless S_a and S_b share a common vertex. Figure 2.3 shows $V(a, b)$ in two cases. $V(a, b)$ is highlighted in grey.

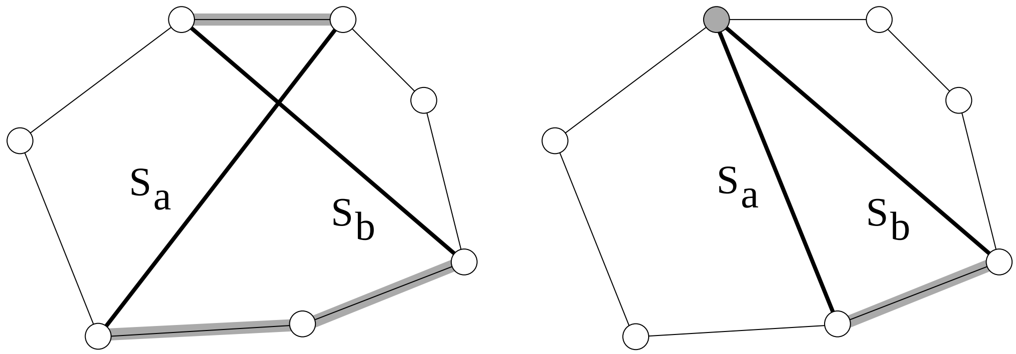

Figure 2.3: The set $V(a,b)$.

2.2 The Structure of Spokes

Here we prove a few combinatorial results about spokes. Figure 2.2 gives a nice illustration of these results.

Lemma 2.1 *Every two spokes intersect.*

Proof: Suppose that S and S' are disjoint spokes. Let T and T' be infinite strips corresponding to S and S' respectively. Both points of S' lie on the same side of $\partial P - \partial T$. This fact, and the convexity of P, forces ∂T and $\partial T'$ to be parallel. But then $T = T'$, and $\partial P \cap \partial T$ consists of two edges. This is a contradiction. ♠

Lemma 2.2 *When the spokes are ordered according to their slopes, S_j and S_{j+1} are consecutive. In particular, the clockwise correspondence between edges and spokes is a bijection.*

Proof: We identify the set of slopes of lines with $\boldsymbol{RP}^1$, the projective line. The edges S_j, S_{j+1}, and e_j form a triangle τ and S_{j+1} is obtained from S_j by rotating S_j counterclockwise about the vertex u_j of τ opposite e_j. Let $I_j \subset \boldsymbol{RP}^1$ be the set of slopes of lines joining points of e_j to u_j. Any other spoke S_k intersects both S_j and S_{j+1}, and from this property we see that the intervals I_j and I_k have disjoint interiors when $j \neq k$. Furthermore, the union $\bigcup_j I_j$ is exactly $\boldsymbol{RP}^1$. Our lemma follows from this structure. ♠

Lemma 2.3 *The edges e_j and e_{j+1} are adjacent if and only if S_{j+1} is a special spoke.*

Proof: Without loss of generality, we set $j = 1$. If S_2 is a special spoke, then the spokes S_1, S_2, S_3 all share a common point. The other three points are the vertices of the arc $e_1 \cup e_2$. Hence e_1 and e_2 are adjacent. Conversely, suppose e_1 and e_2 are adjacent. Then the each of the 3 vertices of $e_1 \cup e_2$ is a vertex of some S_j for $j = 1, 2, 3$. This forces the vertex $S_1 \cap S_2$ to equal the vertex $S_2 \cap S_3$. Hence S_2 is special. ♠

Lemma 2.4 *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: We make two observations: First, for the first edge of an admissible path, the pinwheel orientation coincides with the path orientation. See Figure 2.2. Second, suppose that $S_1, ..., S_k$ are consecutive spokes all sharing a common vertex v. When these spokes are given the pinwheel ordering, the common vertex v is the head vertex of $S_2, ..., S_k$ and the tail vertex of S_1. This lemma follows from the two observations and induction. ♠

2.3 The Structure of the Edge Set

For this whole section, we choose two integers a and b with $a < b < a + n$. Here n is the number of sides of P.

Lemma 2.5 *$V(a, b)$ has two connected components, both having their endpoints in $S_a \cup S_b$.*

Proof: The proof goes by induction on $b - a$. When $b - a = 1$, we have $V(a, b) = e_a \cup (S_a \cap S_{a+1})$, the union of an edge and a vertex. In this case, the result is clear from the basic fact that V_a, V_{a+1} and e_a make a triangle. When $b - a > 1$, we have $V(a, b) = V(a, b - 1) \cup e_{b-1}$. We simply adjoin the edge e_{b-1} to one component of $V(a, b)$ or the other, and use the fact that S_{b-a}, S_b, and e_{b-1} make a triangle. ♠

Lemma 2.6 *Let e be an edge of $V(a,b)$. Let v be the vertex of P that is farthest from the line extending e. Then v is a point of the component of $V(a,b)$ that does not contain e.*

Proof: Let S be the spoke corresponding to e. Then S intersects both S_a and S_b. One endpoint of S is a vertex of e. So, by Lemma 2.5, the other endpoint, namely v, must lie in the other component of $V(a,b)$. ♠

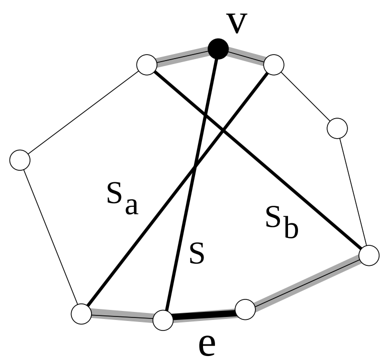

Figure 2.4: The set $V(a,b)$ and the spoke S.

Lemma 2.7 *Let p_1, p_2, p_3 be three consecutive points of an outer billiards orbit. Let w_1 be the midpoint of $\overline{p_1p_2}$ and let w_2 be the midpoint of $\overline{p_2p_3}$. Suppose w_1 is an endpoint of S_a and w_2 is an endpoint of S_b. Then $V(a,b)$ does not separate w_1 from w_2.*

Proof: We will treat the case when S_a and S_b intersect at an interior vertex, so that both components of $V(a,b)$ are finite unions of edges. See Figure 2.5 below. The other case has a similar treatment.

Examining our proof of Lemma 2.5, we note that we can describe $V(a,b)$ as follows. As we rotate S_a counterclockwise, through a suitable family of chords of P, until it reaches S_b, the set $V(a,b)$ is the curve swept out by the endpoints of the rotating chord. We make this rotating family precise as follows. Since S_a, S_{a+1}, e_a make a triangle, we can interpolate from S_a to S_{a+1} by considering the chords that join $S_a \cap S_{a+1}$ to some point of e_a. Then we do the same thing for S_{a+1} and S_{a+2}. And so on.

At the same time, (w_1, w_2) is an admissible pair, so the clockwise arc A joining w_1 to w_2 avoids the interior of $V(a,b)$. Indeed, A is the arc swept out by one of the endpoints of S_a as we rotate it clockwise until it reaches S_b. ♠

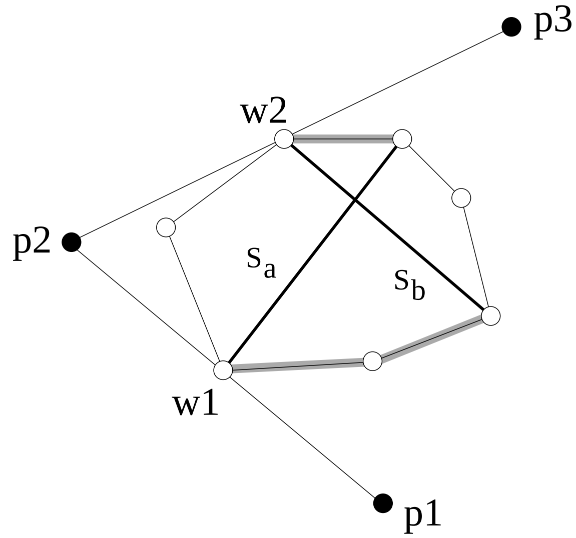

Figure 2.5: The set $V(a,b)$ and part of an outer billiards orbit.

2.4 Admissible Paths and Tiles

Here is the main result in this section.

Lemma 2.8 *There is a bijection between tiles in the forward partition and admissible paths. The tile corresponding to $a \to b$ is labelled by the vertex pair (v,w), where v is the first vertex of $a \to b$ and w is the last one.*

We will prove this result through a series of three smaller results.

Lemma 2.9 *A pair of vertices (v,w) is admissible if and only if v and w respectively are the starting and ending 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 $a \to b$ and minimally lengthen it to the new admissible path $a \to b'$, 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 claim follows from these facts and from induction. ♠

Lemma 2.10 *There is a bijection between unbounded tiles in the forward partition and the maximal pairs of vertices.*

Proof: Suppose that p_1, p_2, p_3 are three consecutive points in an outer billiards orbit. Let w_1 be the midpoint of $\overline{p_1p_2}$ and let w_2 be the midpoint of $\overline{p_2p_3}$. We claim that there is some compact set K such that (w_1, w_2) is not maximal only if $p_1, p_2, p_3 \in K$. If this was false, then we could take a limit of counter examples and produce an infinite strip T such that $P \subset T$ and $\partial T \cap \partial S = \{w_1, w_2\}$. Hence, once p_1 is sufficiently far from the origin, the tile containing p is labelled by a maximal pair of vertices.

Conversely, let e be an edge of P and let (Σ, V) be the associated pinwheel pair. Let v and w be the tail and head vertices of V. We scale so that the lines $y = 0$ and $y = 1$ bound Σ. If R is sufficiently large, then the tiles containing $(R, -1)$ and $(-R, 2)$ are respectively labelled by (v, w) and (w, v). These points each lie just outside the strip Σ. ♠

Lemma 2.11 *There is a bijection between bounded tiles in the forward partition and admissible pairs which are not maximal.*

Proof: We use the notation from Lemma 2.10. We write $p_j(0) = p_j$ and we let $p_1(t)$ be the point that is t units away from $p_1(0)$ along the ray $\overrightarrow{p_2(0)p_1(0)}$. Let $w_1(t)$ and $w_2(t)$ be the vertices that depend on the points $p_j(t)$.

As t increases, and $p_1(t)$ moves away from the origin, the line $\overline{p_1(t)p_2(t)}$ does nothing and the line $\overline{p_2(t)p_3(t)}$ rotates clockwise. The vertex $w_1(t)$ is independent of t but the vertex $w_2(t)$ moves clockwise away from $w_1(t)$ in discrete jumps. A jump occurs every time $p_2(t)$ lies on one of the rays shown in Figure 1.4. When t is sufficiently large $(w_1(t), w_2(t))$ is maximal. Hence $(w_1(0), w_2(0))$ is admissible.

Since $p_1(0)$ lies in a bounded tile and $p_1(t)$ lies in an unbounded tile for t sufficiently large, the vertex $w_2(t)$ must have made at least 1 jump along the way. Hence the initial pair $(w_1(0), w_2(0))$ is not maximal. This shows that each bounded tile corresponds to an admissible but not maximal pair.

For the converse, our construction shows that each admissible pair arises as the pair associated to some tile in the forward partition. We first choose p_1 so that w_1 and w_2 are adjacent vertices. We then pull p_1 away from P, as described above, and observe that we encounter every admissible pair whose initial vertex is w_1. ♠

2.5 A Result about Strips

Let P be a convex polygon with no parallel sides. Let Σ the pinwheel strip associated to an edge e of P. We rotate so that e is horizontal, and lies on the bottom edge of Σ, as in Figure 2.6. Let T be the triangle formed by lines extending the two edges of P incident to v and the top boundary component of Σ. Let R be the ray extending the left edge of T.

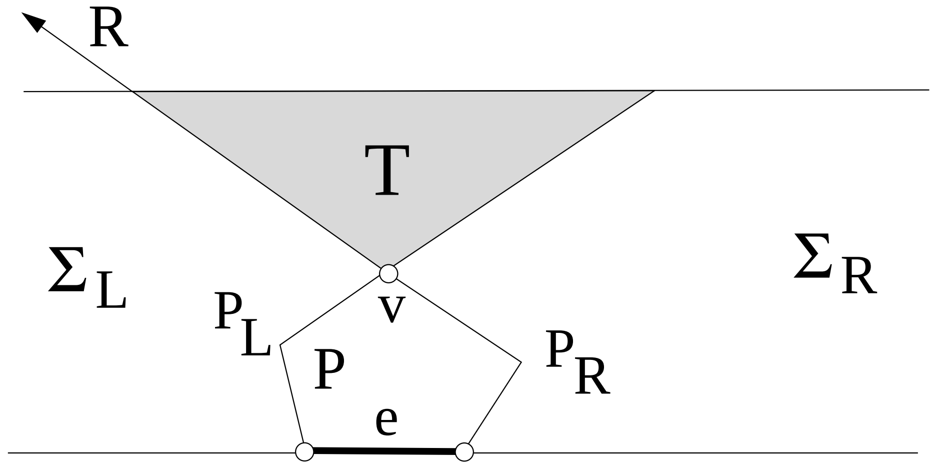

Figure 2.6: Modified Strip

Let P_L and P_R repectively be the closures of the left and right halves of $P - e - \{v\}$. Let Σ_L and Σ_R denote the left and right halves of $\Sigma - P - T$.

Lemma 2.12 *Let p_1, p_2, p_3 be a portion of the outer billiards orbit. Let w_j be the vertex of P that bisects $\overline{p_j p_{j+1}}$. Suppose that $w_1, w_2 \in P_L$ and the clockwise arc from w_2 to w_3 does not contain e. Then $p_2 \in \Sigma_L$. On the other hand, suppose that $w_1, w_2 \in P_R$ and the counterclockwise arc from w_1 to w_0 does not contain e. Then $p_2 \in \Sigma_R$.*

Proof: We will prove the former statement. The latter statement has the same kind of proof. The conditions $w_1, w_2 \in P_L$ force p_2 to lie in the region Σ_L' bounded by P_L, by R, and by the negative x-axis. Σ_L is exactly the portion of Σ_L' that lies below the line $y = 1$. If $p_2 \in \Sigma_L' - \Sigma_L$ then $w_2 = v$ and p_3 lies below the line $y = 0$. But then $w_3 \in P_L$ and the clockwise arc from w_2 to w_3 contains e. This contradiction shows that $p_2 \in \Sigma_L$. ♠

2.6 Structure of the Forward Partition

The rays in Figure 1.4 are the sets where the outer billiards map is not defined. We call these the *primary rays*. The primary rays separate $\boldsymbol{R}^2 - P$ into n distinct cones. Figure 2.7 shows two of these rays and one of the cones. We label the cones $C_1, ..., C_n$, so that C_a contains all the tiles of the form $T(a \to b)$. The apex of C_a is a vertex of the spoke S_a, as the labelling in Figure 1.4 indicates.

Figure 2.7 illustrates what the forward partition looks like within a single cone. The thin segments $A, B, C, ...$ separating the tiles within the cone are mapped, by the outer billiards map, to the primary rays extending the edges we have also labelled $A, B, C....$ All the tile except one are quadrilaterals, and the tile containing the apex of the cone is a triangle.

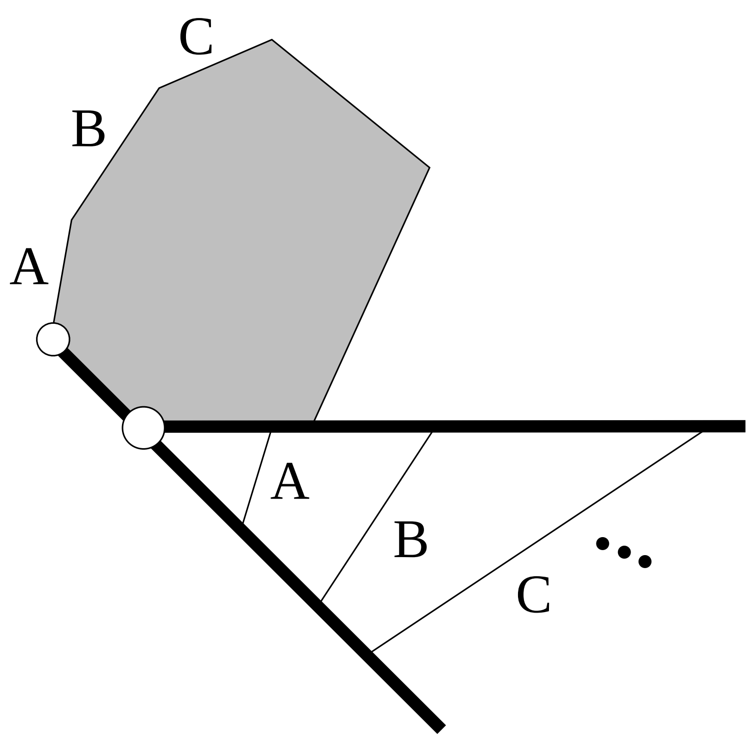

Figure 2.7: The forward partition within a cone.

The bounded tiles in the cone $C = C_a$ all have a common boundary with the pinwheel strip $\Sigma = \Sigma_a$, but are disjoint from Σ. Our next result gives an estimate on the geometry of these tiles.

Lemma 2.13 *Let τ be any bounded tile in C. No point of τ is farther from $\partial\Sigma$ than half the width of Σ.*

Proof: We rotate so that Σ is horizontal, as in Figure 2.8. The line through v_1 is the centerline of Σ. Let T be the union of all the bounded tiles in C. Let w be the edge of T opposite the apex v_0 of C. Let p_0 be the bottom vertex of T. Let p_1 be the image of p_0 under the outer billiards map. Note that w is parallel to a line L which extends one of the sides of P. The basic fact is that p_1 below the centerline of Σ. But p_0 and p_1 are equidistant from the lower boundary of Σ. Hence the distsance 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 holds for all points in T. ♠

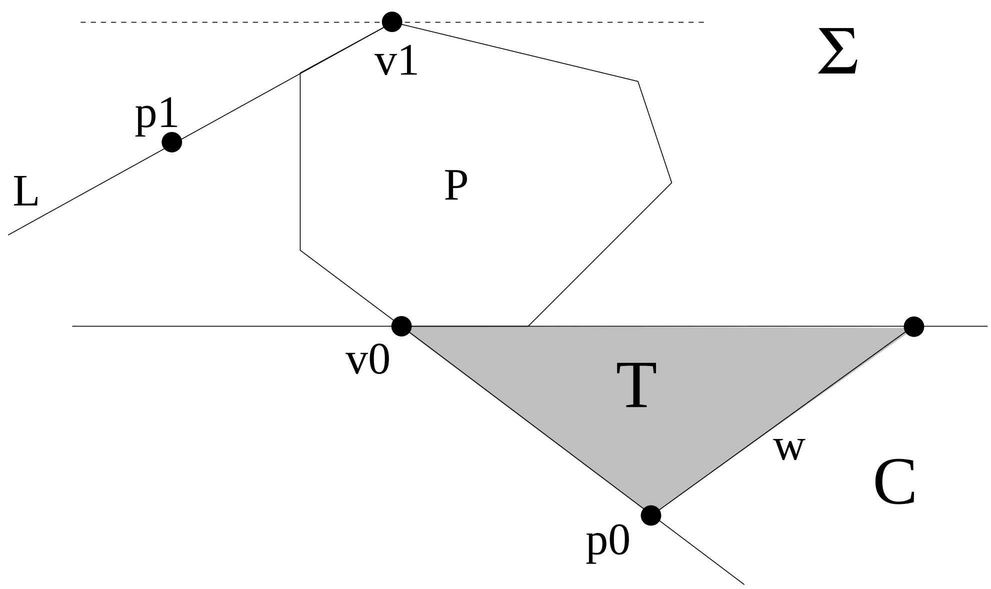

Figure 2.8: Bounding the tiles.

Recall that (Σ_j, V_j) is the pinwheel pair associated to the edge e_j 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 of $\partial\Sigma_j$. Let $T(1 \to b)$ be the tile in the forward partition associated to the admissible path $1 \to b$.

Lemma 2.14 *Suppose $1 < b < n+1$. If e_0 is adjacent to e_1 then $T(1 \to b)$ and P lie on the same side of $\partial_1\Sigma_0$. Otherwise, $T(1 \to b)$ and P lie on the same side of $\partial_2\Sigma_0$.*

Proof: Since $1 < b < n+1$, the tile $T(1 \to b)$ is bounded. If e_0 and e_1 are adjacent, then $e_0 = f_0$ in Figure 2.9. The shaded cone in Figure 2.9 contains $T(1 \to b)$. In this case, the whole shaded cone lies on one side of $\partial_1 \Sigma_0$, and the result is obvious. In the second case, $e_0 = g_0$, and we want to rule out the possibility that $T(1 \to b)$ intersects the black cone. But one checks easily that points in the black cone correspond to the pair (x, y). In this case, these are the endpoints of the spoke S_0. But such points lie in the tiles of the form $T(0 \to a)$. This is a contradiction. ♠

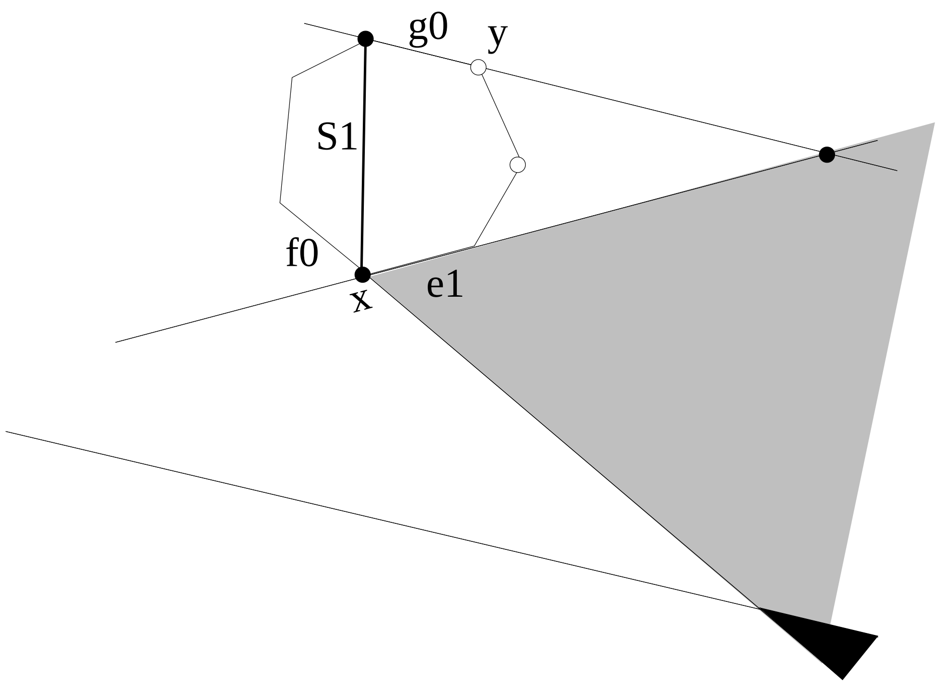

Figure 2.9: The excluded region.

3 Proof modulo Two Details

3.1 The Main Argument

Let ι be the pinwheel section. According to Lemma 2.8, there is a bijective correspondence between tiles in the forward partition and admissible paths. We let $T(a \to b)$ denote the tile in the forward partition corresponding to the path $a \to b$. Referring to the discussion in §2.6, we recall that the tile $T(a \to b)$ lies in the cone C_a. But then, by definition of the pinwheel section,

$$p \in T(a \to b); \qquad \Longrightarrow \qquad \iota(p) = (p, a-1). \tag{5}$$

Let ψ^* denote the pinwheel map, defined in §1.2. Below we will prove the following result.

Theorem 3.1 *Suppose $p \in T(a \to b)$ and $q = \psi(p) \in T(c \to d)$. Then $(\psi^*)^k(q, b-1) = (q, c-1)$ for some $k \in \{0, ..., n\}$.*

We note in Theorem 3.1 that the conclusion of the result only involves the point q, but the choice of b in the conclusion comes from the hypothesis that $q = \psi(p)$ and $p \in T(a \to b)$.

We will also prove the following result.

Theorem 3.2 *Let $a \to b$ be an admissible path, labelled so that $a \leq b$. If $p \in T(a \to b)$ then $(\psi^*)^k(p, a-1) = (\psi(p), b-1)$ for some $k \in \{1, ..., 2n\}$.*

Combining Theorems 3.1 and 3.2, we have the following result.

Corollary 3.3 *Let $T(a \to b)$ and $T(c \to d)$ be two tiles. Suppose that we have $p \in T(a \to b)$ and $\psi(p) \in T(c \to d)$. Then there is some integer $k \in \{0, ..., 3n\}$ such that $(\psi^*)^k(p, a-1) = (\psi(p), c-1)$.*

Proof of the Pinwheel Theorem: The map ψ is defined precisely on the interiors of the tiles in the forward partition. For p in the interior of the tile $T(a \to b)$, we have $\iota(p) = (p, a-1)$. As usual, we take the indices mod n. The conclusion of Corollary 3.3 is just a restatement of the equation in the Pinwheel Theorem. ♠

Now we prove the corollaries. We will consider the forward orbits. The other cases are similar.

Proof of Corollary 1.2: Let $p \in \boldsymbol{R}^2$ be a point with an unbounded forward ψ-orbit O. We define $p^* = \iota(p)$. Let O^* be the ψ^* orbit of p^*. By the Pinwheel Theorem, there is an infinite sequence $t_1, t_2, ...$ such that $\iota \circ \psi^k(p) = (\psi^*)^{t_k}(p^*)$. This clearly shows that O is forward unbounded if and only if O^* is forward unbounded. As we mentioned after stating the Pinwheel Theorem, we have $k(p) = 1$ when p is sufficiently far from the origin and not contained in a pinwheel strip. We have $k(p) = 2$ when p is sufficiently far from the origin and contained in a pinwheel strip. From these properties, it is easy to see that the following is true for all sufficiently large compact sets K.

$$\pi(O^* - K^*) = O - K; \qquad K^* = \pi^{-1}(K). \tag{6}$$

It follows from Equation 6 that there exists only one orbit O^* such that $\pi(O^* - K) = O - K$. Hence, the assignment $O \to O^*$, 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 $O \to O^*$ is injective.

Recall that an unbounded orbit of ψ^* is *natural* if it lies in $\iota(\boldsymbol{R}^2)$ sufficiently far from the origin. 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. ♠

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 $\widehat{\psi}$. But $\iota(\boldsymbol{R}^2)$ and $\widehat{X}$ agree outside a compact set. So, suppose that O^* is a forward unbounded natural orbit. The set $\widehat{O} = O^* \cap \widehat{X}$ is a forward unbounded orbit of $\widehat{X}$. The nature this construction makes it clear that the correspondence $O^* \to \widehat{O}$ is a bijection. ♠

Proof of Corollary 1.4: For each orbit O of ψ, the intersection $O \cap \Sigma_1$ is the corresponding orbit of Ψ. The assignment $O \to (O \cap \Sigma_1)$ 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 $\widehat{\psi}$ and the set of unbounded orbits of $\widehat{\Psi}$. Finally, Corollary 1.3 gives us a canonical bijection between the unbounded orbits of ψ and the set of unbounded orbits of $\widehat{\psi}$. Composing all these bijections gives us the desired result. ♠

3.2 Proof of Theorem 3.1

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 $T_+(a \to b) = T(a \to b)$, and $T_-(a \to b)$ 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. Tautologically, we have

$$\psi(T_+(a \to b)) = T_-(b \to a). \tag{7}$$

If $b = c$ in Theorem 3.1, then we can take $k = 0$ and there is nothing to prove. So, assume $b \neq c$. If necessary, we add n to c so that $b < c < b + n$.

Lemma 3.4 *Let p_1, p_2, p_3 be a portion of the outer billiards orbit. Let w_j be the vertex of P that bisects $\overline{p_j p_{j+1}}$. Suppose also that $p_2 \in T_-(a, b) \cap T_+(c, d)$ for some indices a, d. Then w_1 is an endpoint of the spoke S_b and w_2 is an endpoint of the spoke S_c.*

Proof: By definition the admissible path $c \to d$ has w_2 as its first endpoint and w_3 as its second endpoint. Hence w_2 is an endpoint of S_c. The same argument, applied to the inverse of the outer billiards map, shows that w_1 is an endpoint of S_b. ♠

Lemma 3.5 *We have $T_-(a, b) \cap T_+(c, d) \subset \Sigma_j$ for all $j = b, ..., c - 1$.*

Proof: By definition, $e = e_j \subset V(b, c)$, the set defined in §2.1. We rotate so that e is as in Lemma 2.12. Lemma 2.1 justifies our picture of S_b and S_c as crossing. Lemma 2.5 justifies our picture of $V(b, c)$ as the union of two arcs whose endpoints lie in $S_b \cup S_c$. We have drawn $V(b, c)$ thickly in Figure 4.2. Let v be the vertex of P farthest from the line extending e. Lemma 2.6 says that v is a subset of the component of $V(b, c)$ that does not contain e.

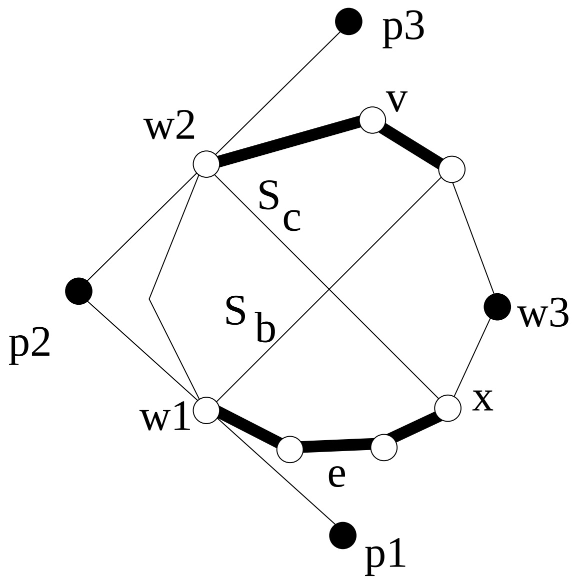

Figure 3.1:

By Lemma 3.4, the point w_1 is an endpoint of S_b and the point w_2 is an endpoint of S_c. We let P_L and P_R and Σ_L and Σ_R be as in Lemma 2.12. We mean for these sets to be defined relative to e and v. Note that one component of $\partial P - V(b,c)$ lies in P_L and the other component lies in P_R. Hence, by Lemma 2.7, either $w_1, w_2 \in P_L$ (as shown) or $w_1, w_2 \in P_R$. We will consider the former case. The latter case is similar.

By Lemma 2.8, the pair (w_2, w_3) is admissible. Hence w_2, w_3, x occur in clockwise order. Here w_2 and x are the two endpoints of S_c. In particular, the arc of ∂P that goes clockwise from w_2 to w_3 does not contain e. Lemma 2.12 now says that $p \in \Sigma_L \subset \Sigma_j$. ♠

Proof of Theorem 3.1: Our point q in Theorem 3.1 lies in $T_-(a,b) \cap T_+(c,d)$ for some choice of indices a and d. Hence

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

Hence

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

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

3.3 Theorem 3.2 modulo Two Details

We first prove Theorem 3.2 in the case that $a = b$. In this case $a \to a$ is an admissible path of length 1, and the corresponding tile $T(a \to a)$ in the forward partition is unbounded. Let $p \in T(a \to a)$. By construction

$$\psi(p) = p + 2V_a \tag{8}$$

We scale so that e_a lies in the x-axis and Σ_a lies above the x-axis. The region $T(a \to a)$ contains the point $(R, -1)$ once R is sufficiently large. In particular $T(a \to a)$ lies beneath Σ_a. But then $\psi^*(p, a-1) = (p+2V_a, a-1)$ for all $p \in T(a \to a)$. Hence $\psi^*(p, a-1) = (\psi(p), a-1)$, as desired.

Now we assume that $a < b < a+n$. For convenience, we label so that $a = 1$. The admissible path $1 \to b$ does not necessarily involve all the spokes between 1 and b. Let V_i be the ith pinwheel vector and let S_i be the ith spoke. If S_i is not involved in the path $1 \to b$, let $W_i = 0$. Otherwise, let W_i be the vector that points from the first endpoint of S_i to the last one. Here we are orienting S_i according to the admissible path $1 \to b$.

If W_k is nonzero, and $1 \leq k < b$ then, according to Lemma 2.4, we have $W_k = V_k$. We also have $W_b = \pm V_b$, with the sign depending on whether or not S_b is a special spoke.

Lemma 3.6 *For any $p \in T(1 \to b)$ we have*

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

Proof: Let (v, w) be the admissible pair of vertices associated to $T(1 \to b)$. Recall that 1 and b respectively name the first and last spoke of the admissible path associated to $T(1 \to b)$ whereas v and w respectively name the first endpoint of the path and the last endpoint of the path. The path $1 \to b$ simply traces out the involved spokes. By definition

$$w - v = \sum_{i=1}^{b} W_i.$$

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

Lemma 3.6 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 3.6 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 $T(a \to b)$ lie in the right place with respect to the relevant strips. Define

$$T(1 \to b; k) = T(1 \to b) + \sum_{i=1}^{k} 2W_i; \qquad k = 1, ..., b \tag{9}$$

Recall that μ_b is the strip map associated to the pinwheel pair (Σ_b, V_b).

Lemma 3.7 *$T(1 \to b; k) \subset \Sigma_k$ for all $k = a, ..., (b-1)$.*

Proof: See §4 ♠

Lemma 3.8 *$\mu_b(p) = p + 2W_b$ for all $p \in T(1 \to b, b-1)$.*

Proof: See §5 ♠

Proof of Theorem 3.2: Let $p = p_0 \in T(1 \to k)$ be an arbitrary point. Define

$$p_k = p_0 + \sum_{i=1}^{k} 2W_i; \qquad k = 1, ..., b. \tag{10}$$

We have $p_k \in T(1 \to b; k)$.

By Lemma 3.6, we have

$$\psi(p_0) = p_b. \tag{11}$$

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

$$p_{k+1} = p_k + 2W_k \in T(1 \to b; k+1) \subset \Sigma_{k+1}. \tag{12}$$

The second containment is Lemma 3.8. Therefore

$$\mu_{k+1}(p_k) = p_k + 2V_k = p_k + 2W_k = p_{k+1}; \qquad \mu_{k+1}(p_{k+1}) = p_{k+1}. \tag{13}$$

Hence

$$(\psi^*)^2(p_k,k) = \psi^* \circ \mu_{k+1}(p_k,k) = \psi^*(p_{k+1},k) = (p_{k+1},k+1). \tag{14}$$

Now suppose that the index k is not involved in the path $1 \to b$. Then $p_{k+1} = p_k$ and,

$$p_k = p_{k+1} \in T(1 \to b; k+1) \subset \Sigma_{k+1}. \tag{15}$$

Again, the second containment is Lemma 3.8. Hence, by definition,

$$\psi^*(p_k,k) = (p_k,k+1) = (p_{k+1},k+1). \tag{16}$$

In either case, we see that

$$(p_{k+1},k+1) = (\psi^*)^e(p_k,k); \qquad e \in \{1,2\}. \tag{17}$$

Applying this argument for as long as we can, we see that

$$(p_{b-1},b-1) = (\psi^*)^e(p_0,0); \qquad e \in \{1,...,2n-2\}. \tag{18}$$

Finally, by Lemma 3.7.

$$\psi^*(p_{b-1},b-1) = (\mu_b^*(p_{b-1}),b-1) = (p_{b-1}+W_b,b-1) = (p_b,b-1). \tag{19}$$

Hence $(p_b,b-1)$ is in the forward ψ^*-orbit of $(p_0,0)$. Combining this information with Equation 11, we see that there is some positive $k < 2n$ such that

$$(\psi(p),b-1) = (p_b,b-1) = (\psi^*)^k(p,0). \tag{20}$$

This completes the proof. ♠

4 Proof of Lemma 3.7

4.1 A Technical Lemma

Let $W_1, ..., W_b$ be the vectors associated to the admissible path $1 \to b$, as in the statement of Lemma 3.7. We have $W_1 = V_1$. As a special case of Equation 9, we consider

$$T(1 \to b, 1) = T(1 \to b) + 2W_1 = T(1 \to b) + 2V_1. \tag{21}$$

Lemma 4.1 *If S_1 is a special spoke, then the line $\partial_2\Sigma_0$ separates P from $T(1 \to b; 1)$. If S_1 is an ordinary spoke, then the line $\partial_1\Sigma_0$ separates P from $T(1 \to b; 1)$.*

Proof: If S_1 is a special spoke, then e_0 and e_1 are adjacent by Lemma 2.3. Since e_0 and e_1 are adjacent, the vector V_1 joins a 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 2.14. The point is that adding V_1 ejects $T(1 \to b)$ outside of Σ_0, and onto the correct side.

If S_1 is an ordinary spoke then e_0 and e_1 are not adjacent, Lemma 2.3. 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. ♠

4.2 The Conjugate Polygon

Let $R : \boldsymbol{R}^2 \to \boldsymbol{R}^2$ denote reflection in the x-axis. For any subset $A \subset \boldsymbol{R}^2$, let $\overline{A} = R(A)$. 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 the edge e_0 of P is horizontal and lies in the x-axis, below the rest of P. To emphasize the dependence on P, we write $e_0(P)$, etc. We let $e_0(\overline{P})$ be the horizontal edge of $\overline{P}$. Note that $e_0(\overline{P})$ lies above the rest of $\overline{P}$.

The cyclic ordering forces two sets of equations.

$$\overline{e_k(P)} = e_{-k}(\overline{P}); \qquad \overline{\Sigma_k(P)} = \Sigma_{-k}(\overline{P}). \tag{22}$$

$$\overline{S_k(P)} = S_{1-k}(\overline{P}); \qquad \overline{V_k(P)} = V_{1-k}(\overline{P}). \tag{23}$$

This second set of equations is more subtle. Figure 4.1 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 possiblities. The other possibilities have the same outcome. The rest of Equation 23 is then forced by the cyclic ordering.

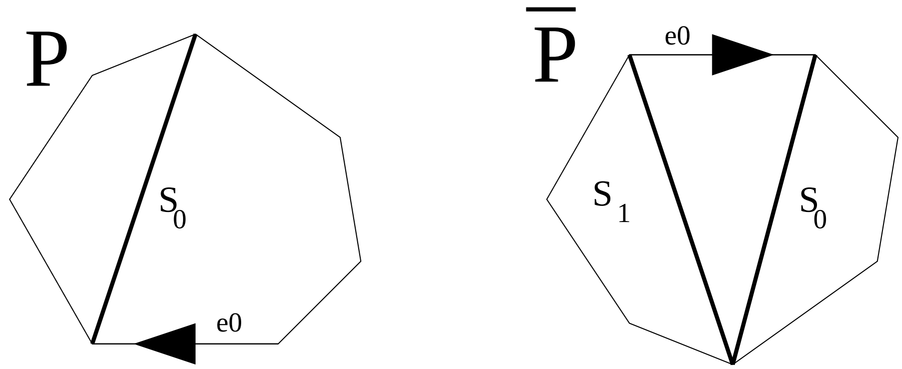

Figure 4.1: The polygon and its conjugate

We can define the sets $T(a \to b; k)$ relative to $\overline{P}$ just as well as for P. We add a P or $\overline{P}$ on the end of our notation to indicate which polygon we mean. We first record a more convenient version of Lemma 3.6.

Lemma 4.2

$$\psi(T(a \to b; P)) = T(a \to b; b; P). \tag{24}$$

Proof: Lemma 3.6 was stated for the case $a = 1$, but the same result holds for arbitrary $a < b < a + n$, just by cyclic relabelling. Hence for any $p \in T(a \to b)$ we have

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

Our lemma now follows from this equation and from the fact that

$$T(a \to b; b; P) = T(a \to b; P) + \sum_{i=a}^{b} W_i.$$

This completes the proof ♠

Before we prove the next result, we note that $T(a \to b, a-1; P)$ is just $T(a \to b; P)$. We also note that the sets $T(1-b \to 1-a, -k; \overline{P})$ make sense because $-k$ runs from $-b$ to $-a$. In particular we have the equation $T(1-b \to 1-a; -b; \overline{P}) = T(1-b \to 1-a; \overline{P})$. We mention this just to make sure that all the objects in the next result are really defined.

Lemma 4.3 $\overline{T(a \to b; k; P)} = T(1-b, 1-a; -k; \overline{P})$ *for all* $k = a-1, ..., b-1$.

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

$$\psi(T_+(a \to b)) = T_-(b \to a). \tag{25}$$

The path $b \to a$ is the same as the path $a \to b$ 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 23 and 25, we see that

$$R \circ \psi(T(a \to b; P)) = T(1-b \to 1-a; \overline{P}) = T(1-b \to 1-a; -b; \overline{P}). \tag{26}$$

Combining Equations 24 and 26, we have

$$\overline{T(a \to b; b; P)} = T(1-b, 1-a; -b; \overline{P}). \tag{27}$$

By the reversal property and Equation 23, we have

$$\overline{W_k(P)} = -W_{1-k}(\overline{P}). \tag{28}$$

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

$$\overline{T(a \to b; b-1; P)} = \overline{T(a \to b; b; P)} - \overline{W_b(P)} =$$

$$T(1-b \to 1-a; \overline{P}) + W_{1-b}(\overline{P}) = T(1-b, 1-a; 1-b; \overline{P}).$$

in short

$$\overline{T(a \to b; b-1; P)} = T(1-b, 1-a; 1-b; \overline{P}). \tag{29}$$

The above argument allows us to deduce the case $k = b-1$ from the case $k = b$. Repeating the same argument, we get the cases $k = b-2, b-3, ...$ ♠

4.3 The End of the Proof

We are going to apply Lemma 4.1 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

$$\alpha = 1 - b; \qquad \beta = 1 - a. \tag{30}$$

By Lemma 4.3,

$$\overline{T(a \to b, b-1; P)} = T(\alpha \to \beta; \alpha; \overline{P}). \tag{31}$$

By definition and Lemma 2.4,

$$T(\alpha \to \beta; \alpha; \overline{P}) = T(\alpha \to \beta; \overline{P}) + V_\alpha(\overline{P}). \tag{32}$$

By Corollary 4.1, the line

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

separates $T(\alpha \to \beta; \overline{P}) + V_\alpha(\overline{P})$ from $\overline{P}$. Applying the reflection R and using Equation 23, 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 $p \in T(1 \to b; b-1; P)$. Since S_b is an ordinary spoke, $V_b = W_b$ by Lemma 2.4. 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 $p \in T(1 \to b; b-1; P)$. This time $-V_b = W_b$, and we get the same result as in the previous case.

This completes the proof of Lemma 3.7.

5 Proof of Lemma 3.8

5.1 Some Combinatorial Definitions

In this section we introduce some special terminology for the proof of Lemma 3.8. the main point of these definitions is to abstract out some of the features of the spokes of a convex polygon without parallel sides.

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 5.1 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* if the last edge of γ.

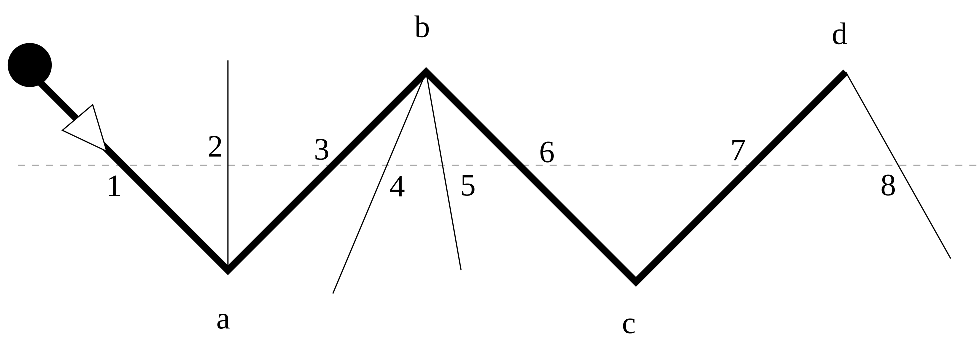

Figure 5.1: Abstract admissible paths.

Linear Order: There is a natural linear ordering on the edges of τ, induced from the order in which they intersect the x-axis, goind from left to right. The numerical labels in Figure 5.1 indicate the ordering. We see that $\tau' \subset \tau$ is a *prefix* of τ of 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 5.1 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 τ is an abstract admissible path and 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 5.1, 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 $(\tau_1, f_1) \to (\tau_2, f_2)$.

Hereditary Properties: Let $\mathcal{C}$ be a collection of sites. We say that $\mathcal{C}$ is *hereditary* if $\mathcal{C}$ is closed under the natural involution, and also under reduction. Say that a site (τ, f) is *initial* 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 initial sites in $\mathcal{C}$.
- $\Omega \circ R = \Omega$. Here R is the natural involution.
- If $\Omega(\tau_1, f_1) = 1$ and $(\tau_2, f_2) \to (\tau_1, f_1)$ then $\Omega(\tau_2, f_2) = 1$.

5.2 The Reduction Lemma

In this section we will prove the following result.

Lemma 5.1 (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 initial 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 left 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 $(\tau, f) \to (\tau', f)$.

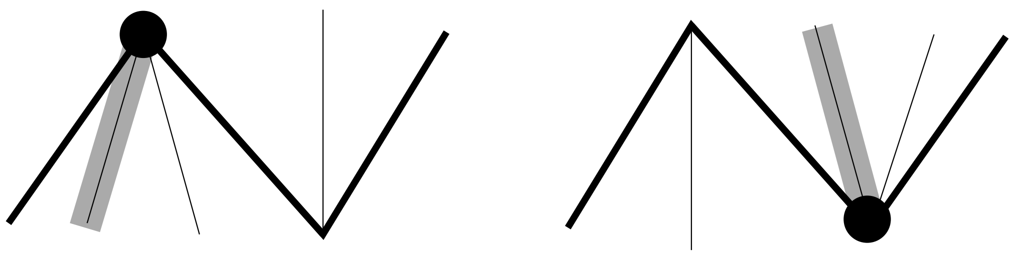

Figure 5.2: Abstract admissible paths.

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 5.2 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 $(\tau, f) \to (\tau', f)$.

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

5.3 The Pinwheel Collection

As usual, all the polygons we consider are convex and have no parallel sides. Let P be a 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 5.3 shows an example.

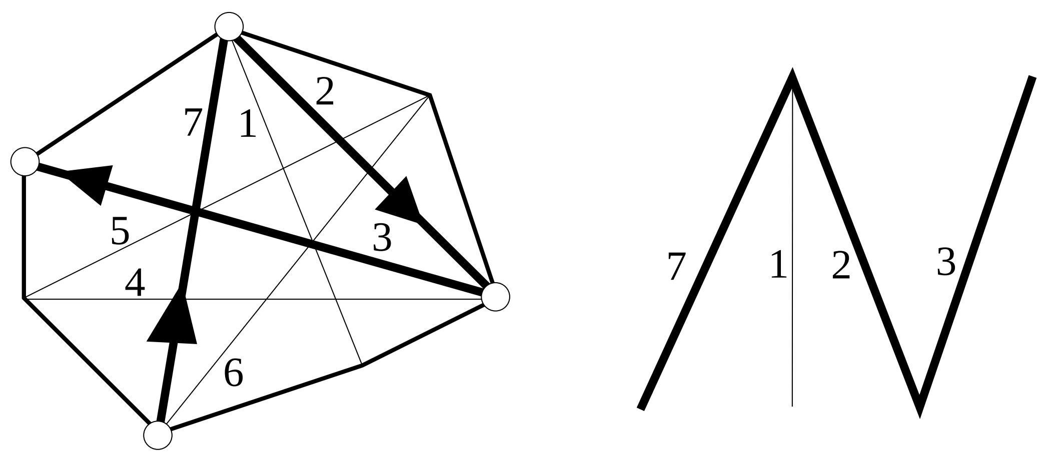

Figure 5.3: An admissible path and its abstraction

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 polygon.

Lemma 5.2 *$\mathcal{C}$ 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 $a \to b$, then any prefix τ' is the abstraction of some admissible path $a \to b'$, where $b' < b$.

The analysis in §4.2 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 $a \to b$ to the backward tile associated to the path $b \to a$. The abstraction of $b \to a$ is exactly the image of the abstraction of $a \to b$ under the natural involution. ♠

5.4 The Binary Function

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

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

$$T(a \to b; k); \qquad\qquad \Sigma_k. \tag{33}$$

All of this depends on P. We define $\Omega'(\tau, f; P) = 1$ if Lemma 3.8 is true for the pair in Equation 33 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 3.8 is equivalent to the statement that $\Omega \equiv 1$ on $\mathcal{C}$. Accordingly, we will prove Lemma 3.8 by showing that Ω is a hereditary function and then invoking the Hereditary Lemma.

5.5 Property 1

The initial sites correspond to the case $k = a$ in Equation 33. Cyclically relabelling, we take $a = 1$. The following lemma implies that $\Omega = 1$ on all initial sites.

Lemma 5.3 $T(1 \to b; 1) \subset \Sigma_1$.

Proof: We rotate so that Σ_1 is horizontal, and $T(1 \to b)$ 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. $T(1 \to b)$ is contained in the shaded region T shown in Figure 4.3. By Lemma 2.13, every point of $T(1 \to b)$ 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 $T(1 \to b)$.

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

5.6 Property 2

That $\Omega \circ R = \Omega$ is a consequence of the relations between P and $\overline{P}$ worked out in §4.2. Let (τ_1, f_1) be the site corresponding to the pair in Equation 33. We label the edges of τ_1 as $a, ..., b$. Let $(\tau_2, f_2) = R(\tau_1, f_1)$. We label the edges of τ_2 as $(1-b), ..., (1-a)$.

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 5.4 *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

$$T(a \to b; k); \qquad \Sigma_k.$$

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

$$T(1-b \to 1-a; -k); \qquad \Sigma_{-k}.$$

But, by Lemma 4.3, 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.

5.7 Property 3

Recall from §2.6 that the primary rays divide $\boldsymbol{R}^2 - P$ into cones. The beginning vertex v of the admissible path $a \to b$ is the apex of the cone. If we consider all the tiles of the form $T(a \to b)$ with a fixed and b increasing, then Lemma 3.8 involves increasingly many containments. On the other hand, the tiles involved get smaller and smaller in the sense that the shrink down to v. One might expect that the vertex v itself satisfies all the identities we can form.

Let $a \to b$ denote the maximal admissible path that starts with S_a. Then $a \to b$ corresponds to the triangular tile $T(a \to b)$ that has v as a vertex. Similar to Equation 9, we define

$$p_0 = v; \qquad p_k = p_0 + \sum_{i=1}^{k} 2W_i. \tag{34}$$

Lemma 5.5 *$p_k \subset \Sigma_k$ for $k = 1, ..., b$.*

Proof: We relabel so that $a = 1$. Let $v_0 = v = p_0$ and let $\{v_k\}$ be the successive vertices of our admissible path, labelled (redundantly) so that v_k is incident to the spokes S_k and S_{k+1}. Figure 5.4 shows an example.

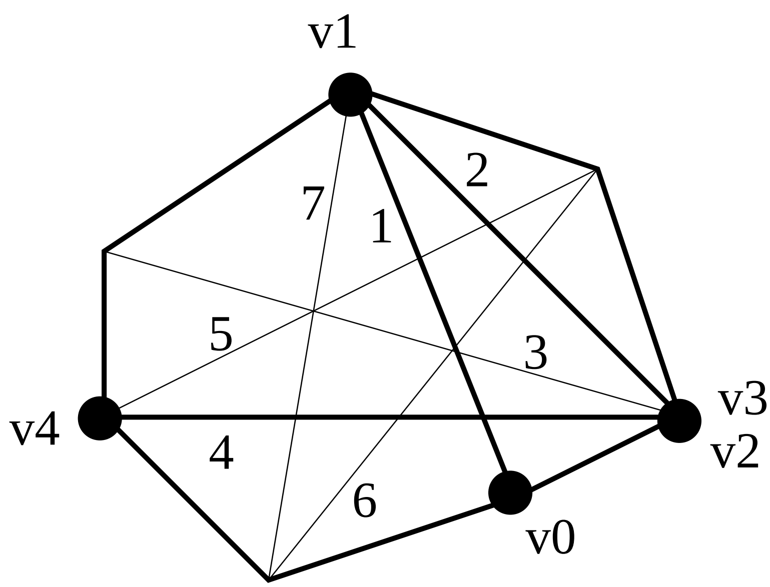

Figure 5.4: The vertices of the path

Since e_k, S_k, S_{k+1} form a triangle, v_k lies on the centerline of Σ_k for all k. Since all pinwheel strips contain P, we have $v_0 \in \Sigma_k$. But v_k is the midpoint of $\overline{v_0 p_k}$. Hence $p_k \in \Sigma_k$, as desired. ♠

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

Corollary 5.6 *Suppose $c > b$ is that that $a \to c$ is an admissible path. Then $\Omega(a \to b; k)$ implies $\Omega(a \to c, k)$.*

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

$$T(a \to b; k); \qquad T(a \to c; k); \qquad p_k$$

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

$$T(a \to b); \qquad T(a \to c); \qquad p_0.$$

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

$$T(a \to c) \subset \widehat{T}(a \to b), \tag{35}$$

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

$$T(a \to c; k) \subset \widehat{T}(a \to b; k) \subset \Sigma_k. \tag{36}$$

The second equality follows from the convexity of Σ_k. ♠

Corollary 5.6 is just a restatement of Property 3.

References

[**D**] R. Douady, *These de 3-eme cycle*, Université de Paris 7, 1982.

[**DF**] D. Dolyopyat 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 Mathematics 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 Syntheses" 1, 1995

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