

The Flapping Birds in the Pentagon Zoo

Richard Evan Schwartz *

March 8, 2024

Abstract

We study the $(k + 1, k)$ diagonal map for $k = 2, 3, 4, \dots$. We call this map Δ_k . The map Δ_1 is the pentagram map and Δ_k is a generalization. Δ_k does not preserve convexity, but we prove that Δ_k preserves a subset B_k of certain star-shaped polygons which we call *k-birds*. The action of Δ_k on B_k seems similar to the action of Δ_1 on the space of convex polygons. We show that some classic geometric results about Δ_1 generalize to this setting.

1 Introduction

1.1 Context

When you visit the pentagram zoo you should certainly make the pentagram map itself your first stop. This old and venerated animal has been around since the place opened up and it is very friendly towards children. When defined on convex pentagons, this map has a very long history. See e.g. [15]. In modern times [19], the pentagram is defined and studied much more generally. The easiest case to explain is the action on convex n -gons. One starts with a convex n -gon P , for $n \geq 5$, and then forms a new convex n -gon P' by intersecting the consecutive diagonals, as shown Figure 1.1 below.

The magic starts when you iterate the map. One of the first things I proved in [19] about the pentagram map is the successive iterates shrink to a point. Many years later, M. Glick [3] proved that this limit point is an

*Supported by N.S.F. Grant DMS-21082802

algebraic function of the vertices, and indeed found a formula for it. See also [9] and [1].

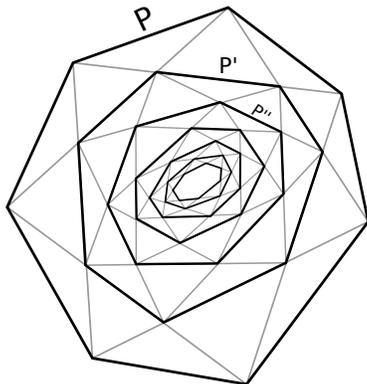


Figure 1.1: The pentagram map iterated on a convex 7-gon P .

Forgetting about convexity, the pentagram map is generically defined on polygons in the projective plane over any field except for $\mathbf{Z}/2$. In all cases, the pentagram map commutes with projective transformations and thereby defines a birational map on the space of n -gons modulo projective transformations. The action on this moduli space has a beautiful structure. As shown in [17] [18], and independently in [23], the pentagram map is a discrete completely integrable system when the ground field is the reals. ([23] also treats the complex case.) Recently, M. Weinreich [24] generalized the integrability result, to a large extent, to fields of positive characteristic.

The pentagram map has many generalizations. See for example [2], [14], [16], [10], [11], [6]. The paper [2] has the first general complete integrability result. The authors prove the complete integrability of the $(k, 1)$ diagonal maps, i.e. the maps obtained by intersecting successive k -diagonals. Figure 1.3 below shows the $(3, 1)$ diagonal map. (Technically, [2] concentrates on what happens when these maps act on so-called *corrugated polygons* in higher dimensional Euclidean spaces.) The paper [6] proves an integrability result for a very wide class of generalizations, including the ones we study below. (Technically, for the maps we consider here, the result in [6] does not establish the algebraic independence of invariants needed for complete integrability.) The pentagram map and its many generalizations are related to a number of topics: alternating sign matrices [20], dimers [5], cluster algebras [4], the KdV hierarchy [12], [13], spin networks [2], Poisson Lie groups [8], Lax pairs [23], [10], [11], [6], [8], and so forth. The zoo has many cages and sometimes you have to get up on a tall ladder to see inside them.

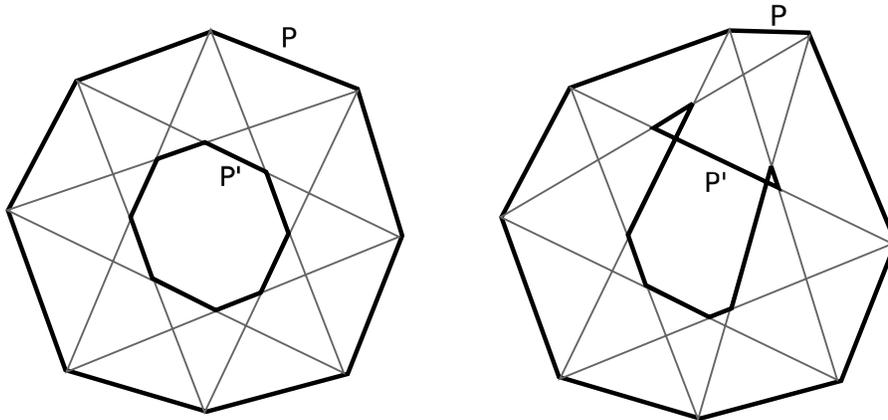


Figure 1.2: The $(3, 1)$ -diagonal map acting on 8-gons.

The algebraic side of the pentagram zoo is extremely well developed, but the *geometric side* is hardly developed at all. In spite of all the algebraic results, we don't really know, geometrically speaking, much about what the pentagram map and its relatives really do to polygons.

Geometrically speaking, there seems to be a dichotomy between convexity and non-convexity. The generic pentagram orbit of a projective equivalence class of a convex polygon lies on a smooth torus, and you can make very nice animations. What you will see, if you tune the power of the map and pick suitable representatives of the projective classes, is a convex polygon sloshing around as if it were moving through water waves. If you try the pentagram map on a non-convex polygon, you see a crazy erratic picture no matter how you try to normalize the images. The situation is even worse for the other maps in the pentagram zoo, because these generally do not preserve convexity. Figure 1.2 shows how the $(3, 1)$ -diagonal map does not necessarily preserve convexity, for instance. See [21], [22] for more details.

If you want to look at pentagram map generalizations, you have to abandon convexity. However, in this paper, I will show that sometimes there are geometrically appealing replacements. The context for these replacements is the $(k + 1, k)$ -diagonal map, which I call Δ_k , acting on what I call k -birds. Δ_k starts with the polygon P and intersects the $(k + 1)$ -diagonals which differ by k clicks. (We will give a more formal definition in the next section.) Δ_k is well (but not perfectly) understood algebraically [6]. Geometrically it is not well understood at all.

1.2 The Maps and the Birds

The Maps: Given a polygon P , let P_a denote the (a) th vertex of P . Let P_{ab} be the line through P_a and P_b . The vertices of $\Delta_k(P)$ are

$$P_{j,j+k+1} \cap P_{j+1,j-k}. \quad (1)$$

Here the indices are taken mod n . Figure 1.3 shows this for $(k, n) = (2, 7)$.

The Birds: We call a polygon P *planar* if some projective transformation maps P to a bounded subset of \mathbf{R}^2 , when it is considered as the affine patch of the projective plane. Given an n -gon P , we let $P_{a,b}$ denote the line containing the vertices P_a and P_b . We call P *k-nice* if $n > 3k$, and P is planar, and the 4 lines

$$P_{i,i-k-1}, \quad P_{i,i-k}, \quad P_{i,i+k}, \quad P_{i,i+k+1} \quad (2)$$

are distinct for all i .

We call P a *k-bird* if P is the endpoint of a path of k -nice n -gons that starts with the regular n -gon. We let $B_{k,n}$ be the subspace of n -gons which are k -birds. Note that $B_{k,n}$ contains the set of convex n -gons, and the containment is strict when $k > 1$. As Figure 1.3 illustrates, a k -bird need not be convex for $k \geq 2$. We will show that k -birds are always star-shaped, and in particular embedded. The homotopy part of our definition looks a bit strange, but it is necessary. For instance, a $2n$ -gon that wraps twice around a convex n -gon is 1-nice but not a 1-bird.

Figure 1.3 shows Δ_2 acting on what we call *2-birds*.

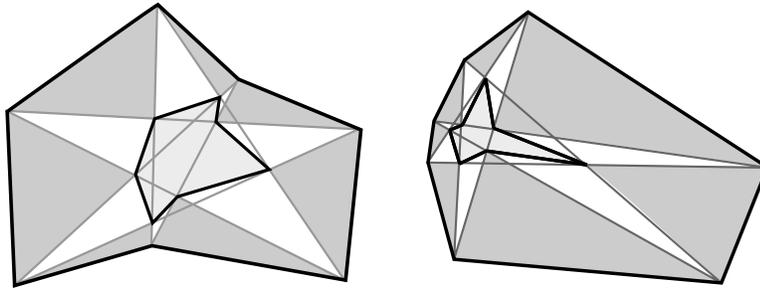


Figure 1.3: Δ_2 acting on 2-birds.

For us, a *polygon* is both the vertices and the edges. Δ_k only acts on the vertices but, given the homotopical way we have defined birds, we automatically have a way to define the edges of our birds. We just extend continuously from the regular n -gon.

1.3 The Main Result

Here is the main result.

Theorem 1.1 *Let $k \geq 2$ and $n > 3k$ and $P \in B_{k,n}$. Then*

1. P is strictly star-shaped.
2. $\Delta_k(P)$ lies in the interior of the region bounded by P .
3. $\Delta_k(B_{k,n}) = B_{k,n}$.

The statement that $n > 3k$ is present just for emphasis. $B_{n,k}$ is by definition empty when $n \leq 3k$. The restriction $n > 3k$ is necessary. Figure 1.4 illustrates what would be a counter-example to Theorem 1.1 for the pair $(k, n) = (3, 9)$. The issue is that a certain triple of 4-diagonals has a common intersection point. This does not happen for $n > 3k$. See Lemma 3.6.

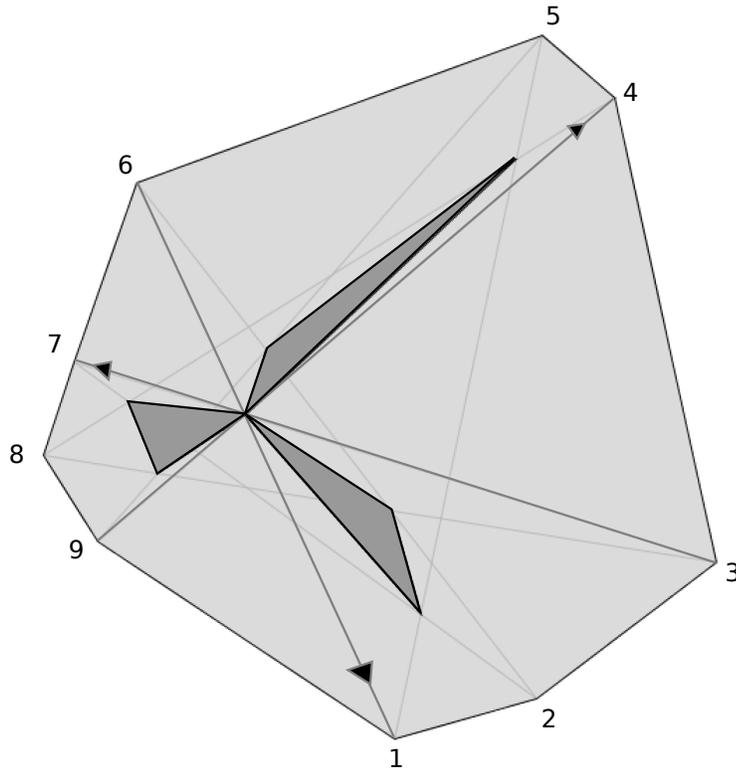


Figure 1.4: Δ_3 acting on a certain convex 9-gon.

1.4 The Energy

We will deduce Statements 1 and 2 of Theorem 1.1 in a geometric way. The key to proving Statement 3 is a natural quantity associated to a k -bird. We let $\sigma_{a,b}$ be the slope of the line $P_{a,b}$ and we define the *cross ratio*

$$\chi(a, b, c, d) = \frac{(a-b)(c-d)}{(a-c)(b-d)}. \quad (3)$$

We define

$$\chi_k(P) = \prod_{i=1}^n \chi(i, k, P), \quad \chi(i, k, P) = \chi(\sigma_{i,i-k}, \sigma_{i,i-k-1}, \sigma_{i,i+k+1}, \sigma_{i,i+k}) \quad (4)$$

Here we are taking the cross ratio the slopes the lines involved in our definition of k -nice. When $k = 1$ this is the familiar invariant $\chi_1 = OE$ for the pentagram map Δ_1 . See [19], [20], [17], [18]. When $n = 3k + 1$, a suitable star-relabeling of our polygons converts Δ_k to Δ_1 and χ_k to $1/\chi_1$. So, in this case $\chi_k \circ \Delta_k = \chi_k$. Figure 1.4 illustrates this for $(k, n) = (3, 10)$. Note that the polygons suggested by the dots in Figure 1.4 are not convex. Were we to add in the edges we would get a highly non-convex pattern.

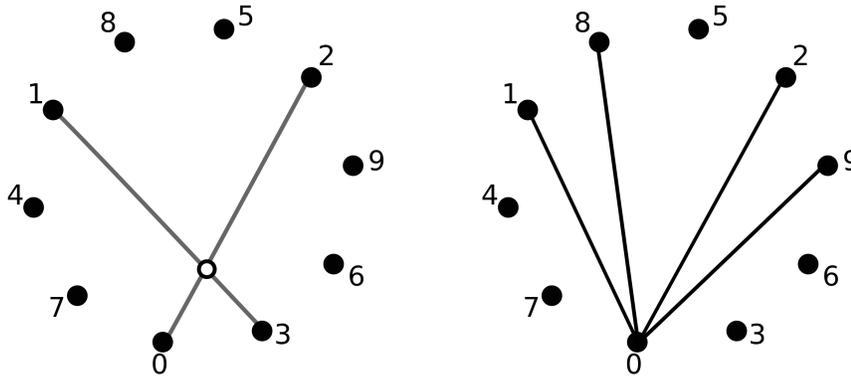


Figure 1.5: A star-relabeling converts Δ_1 to Δ_3 and $1/\chi_1$ to χ_3 .

In general, χ_k is not as clearly related to χ_1 . Nonetheless, we will prove

Theorem 1.2 $\chi_k \circ \Delta_k = \chi_k$.

Theorem 1.2 is meant to hold for all n -gons, as long as all quantities are defined. There is no need to restrict to birds.

1.5 The Collapse Point

When it is understood that $P \in B_{k,n}$ it is convenient to write

$$P^\ell = \Delta_k^\ell(P) \tag{5}$$

We also let \widehat{P} denote the closed planar region bounded by P . Figure 1.6 below shows $\widehat{P} = \widehat{P}^0, \widehat{P}^1, \widehat{P}^2, \widehat{P}^3, \widehat{P}^4$ for some $P \in B_{4,13}$.

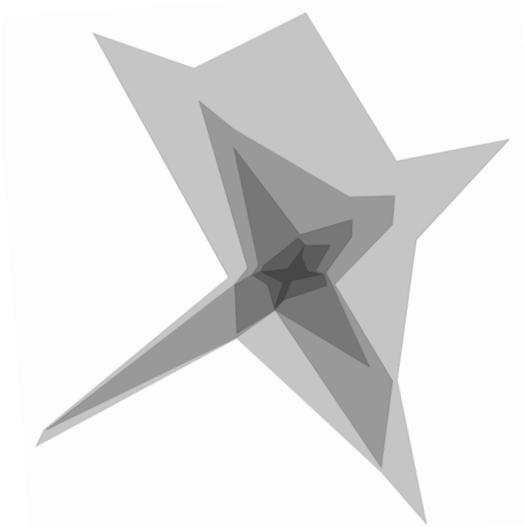


Figure 1.6: Δ_4 and its iterates acting on a member of $B_{4,13}$.

Define

$$\widehat{P}_\infty = \bigcap_{\ell \in \mathbf{Z}} \widehat{P}^\ell, \quad \widehat{P}_{-\infty} = \bigcup_{\ell \in \mathbf{Z}} \widehat{P}^\ell. \tag{6}$$

Theorem 1.3 *If $P \in B_{k,n}$ then \widehat{P}_∞ is a point and $\widehat{P}_{-\infty}$ is an affine plane.*

Our argument will show that $P \in B_{k,n}$ is strictly star-shaped with respect to all points in \widehat{P}^n . In particular, all polygons in the orbit are strictly star-shaped with respect to the collapse point \widehat{P}_∞ . See Corollary 7.3.

As we remarked above, $B_{n,k}$ contains the set of convex n -gons. Thus, if we fix some convex n -gon, we get one collapse point for each $k \in [1, \dots, n/3]$. (The case $k = 1$ gives the pentagram map collapse point.) I satisfied myself that these collapse points are generally distinct from each other.

One might wonder if some version of Glick's formula works for the \widehat{P}_∞ in general. I discovered experimentally that this is indeed the case for $n = 3k + 1$ and $n = 3k + 2$. See §9.2.

1.6 The Triangulations

In §7.1 we associate to each k -bird P a triangulation $\tau_P \subset \mathbf{P}$, the projective plane. Here τ_P is an embedded degree 6 triangulation of $P_{-\infty} - P_{\infty}$. The edges are made from the segments in the δ -diagonals of P and its iterates for $\delta = 1, k, k + 1$.

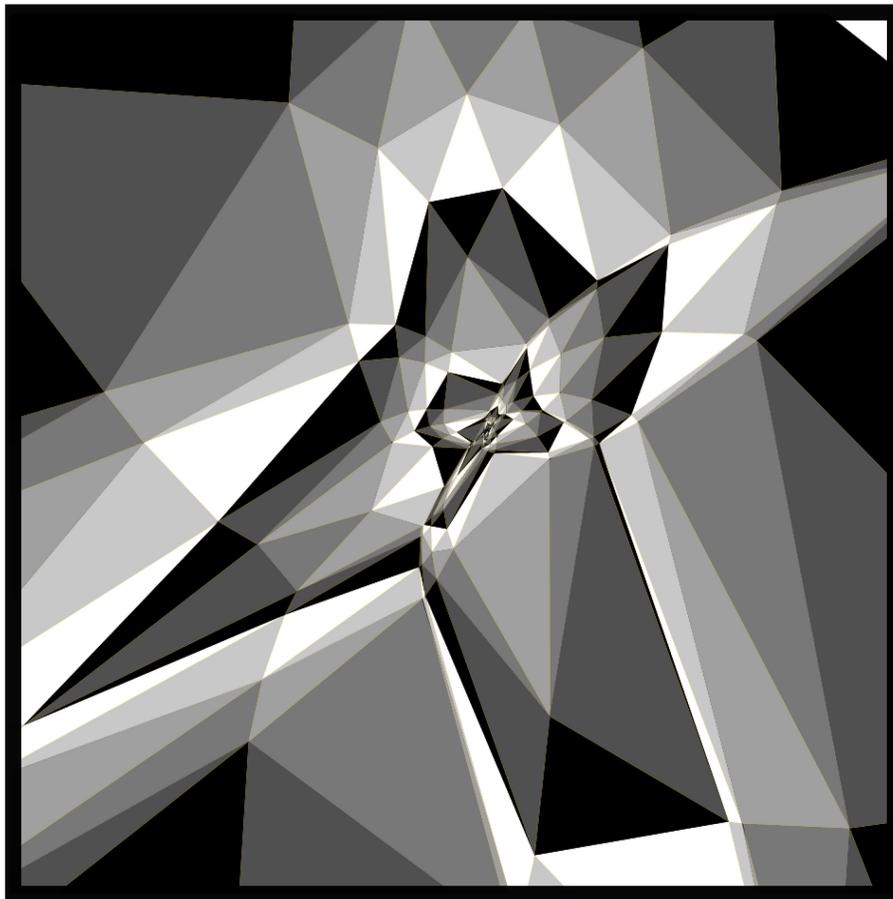


Figure 1.7: The triangulation associated to a member of $B_{5,16}$.

Figure 1.7 shows this tiling associated to a member of $B_{5,16}$. In this figure, the interface between the big black triangles and the big white triangles is some $\Delta_5^\ell(P)$ for some smallish value of ℓ . (I zoomed into the picture a bit to remove the boundary of the initial P .) The picture is normalized so that the line $P_{-\infty}$ is the line at infinity. When I make these kinds of pictures (and animations), I normalize so that the ellipse of inertia of P is the unit disk.

1.7 Paper Organization

This paper is organized as follows.

- In §2 we prove Theorem 1.2.
- In §3 we prove Statement 1 of Theorem 1.1.
- In §4 we prove Statement 2 of Theorem 1.1.
- In §5 we prove a technical result called the Degeneration Lemma, which will help with Statement 3 of Theorem 1.1.
- In §6 we prove Statement 3 of Theorem 1.1.
- In §7 we introduce the triangulations discussed above. They help with the proof of Theorem 1.3.
- In §8 we prove Theorem 1.3.
- In §9, an appendix, we sketch an alternate proof of Theorem 1.2 which Anton Izosimov kindly explained. We also discuss Glick's collapse formula and star relabelings of polygons.

1.8 Visit the Flapping Bird Exhibit

Our results inject some more geometry into the pentagram zoo. Our results even have geometric implications for the pentagram map itself. See §9.3. There are different ways to visit the flapping bird exhibit in the zoo. You could read the proofs here, or you might just want to look at some images:

<http://www.math.brown.edu/~reschwar/BirdGallery>

You can also download and play with the software I wrote:

<http://www.math.brown.edu/~reschwar/Java/Bird.TAR>

The software has detailed instructions. You can view this paper as a justification for why the nice images actually exist.

1.9 Acknowledgements

I would like to thank Misha Gekhtman, Max Glick, Anton Izosimov, Boris Khesin, Valentin Ovsienko, and Serge Tabachnikov for many discussions about the pentagram zoo. I would like to thank Anton, in particular, for extensive discussions about the material in §9.

2 The Energy

The purpose of this chapter is to prove Theorem 1.2. The proof, which is similar to what I do in [19], is more of a verification than a conceptual explanation. My computer program allows the reader to understand the technical details of the proof better. The reader might want to just skim this chapter on the first reading. In §9 I will sketch an alternate proof, which I learned from Anton Izosimv. Izosimov's proof also uses the first two sections of this chapter.

2.1 Projective Geometry

Let \mathbf{P} denote the real projective plane. This is the space of 1-dimensional subspaces of \mathbf{R}^3 . The projective plane \mathbf{P} contains \mathbf{R}^2 as the *affine patch*. Here \mathbf{R}^2 corresponds to vectors of the form $(x, y, 1)$, which in turn define elements of \mathbf{P} .

Let \mathbf{P}^* denote the dual projective plane, namely the space of lines in \mathbf{P} . The elements in \mathbf{P}^* are naturally equivalent to 2-dimensional subspaces of \mathbf{R}^3 . The line in \mathbf{P} such a subspace Π defines is equal to the union of all 1-dimensional subspaces of Π .

Any invertible linear transformation of \mathbf{R}^3 induces a *projective transformation* of \mathbf{P} , and also of \mathbf{P}^* . These form the *projective group* $PSL_3(\mathbf{R})$. Such maps preserve collinear points and coincident lines.

A *duality* from \mathbf{P} to \mathbf{P}^* is an analytic diffeomorphism $\mathbf{P} \rightarrow \mathbf{P}^*$ which maps collinear points to coincidence lines. The classic example is the map which sends each linear subspace of \mathbf{R}^3 to its orthogonal complement.

A *PolyPoint* is a cyclically ordered list of points of \mathbf{P} . When there are n such points, we call this an *n-Point*. A *PolyLine* is a cyclically ordered list of lines in \mathbf{P} , which is the same as a cyclically ordered list of points in \mathbf{P}^* . A projective duality maps PolyLines to PolyPoints, and *vice versa*.

Each *n-Point* determines 2^n polygons in \mathbf{P} because, for each pair of consecutive points, we may choose one of two line segments to join them. As we mentioned in the introduction, we have a canonical choice for k -birds. Theorem 1.2 only involves PolyPoints, and our proof uses PolyPoints and PolyLines.

Given a *n-Point* P , we let P_j be its j th point. We make a similar definition for *n-Lines*. We always take indices mod n .

2.2 Factoring the Map

Like the pentagram map, the map Δ_k is the product of 2 involutions. This factorization will be useful here and in later chapters.

Given a PolyPoint P , consisting of points P_1, \dots, P_n , we define $Q = D_m(P)$ to be the PolyLine whose successive lines are $P_{0,m}, P_{1,m+1}$, etc. Here $P_{0,m}$ denotes the line through P_0 and P_m , etc. We labeled the vertices so that

$$Q_{-m-i} = P_{i,m}. \quad (7)$$

This is a convenient choice. We define the action of D_m on PolyLines in the same way, switching the roles of points and lines. For PolyLines, $P_{0,m}$ is the intersection of the line P_0 with the line P_m . The map D_m is an involution which swaps PolyPoints with PolyLines. We have the compositions

$$\Delta_k = D_k \circ D_{k+1}, \quad \Delta_k^{-1} = D_{k+1} \circ D_k. \quad (8)$$

The *energy* χ_k makes sense for n -Lines as well as for n -Points. The quantities $\chi_k \circ D_k(P)$ and $\chi_k \circ D_{k+1}(P)$ can be computed directly from the PolyPoint P . Figure 2.1 shows schematically the 4-tuples associated to $\chi(0, k, Q)$ for $Q = P$ and $D_k(P)$ and $D_{k+1}(P)$. In each case, $\chi_k(Q)$ is a product of n cross ratios like these. If we want to compute the factor of $\chi_k(D_k(P))$ associated to index i we subtract (rather than add) i from the indices shown in the middle figure. A similar rule goes for $D_{k+1}(P)$.

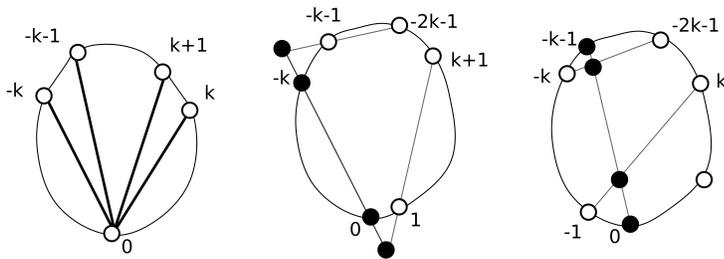


Figure 2.1: Computing the k -energy.

Theorem 1.2 follows from the next two results.

Theorem 2.1 $\chi_k \circ D_k = \chi_k$.

Theorem 2.2 $\chi_k \circ D_{k+1} = \chi_k$.

These results have almost identical proofs. We consider Theorem 2.1 in detail and then explain the small changes needed for Theorem 2.2.

2.3 Proof of the First Result

We study the ratio

$$R(P) = \frac{\chi_k \circ D_k(P)}{\chi_k(P)}. \quad (9)$$

We want to show that $R(P)$ equals 1 wherever it is defined. We certainly have $R(P) = 1$ when P is the regular n -Point.

Given a PolyPoint P we choose a pair of vertices a, b with $|a - b| = k$. We define $P(t)$ to be the PolyPoint obtained by replacing P_a with

$$(1 - t)P_a + tP_b. \quad (10)$$

Figure 2.2 shows what we are talking about, in case $k = 3$. We have rotated the picture so that P_a and P_b both lie on the X -axis.

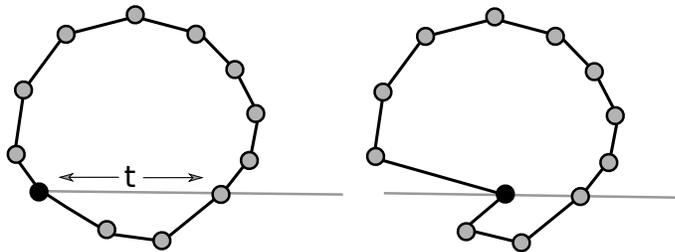


Figure 2.2: Connecting one PolyPoint to another by sliding a point.

The two functions

$$f(t) = \chi_k(P(t)), \quad g(t) = \chi_k \circ D_k(P(t)) \quad (11)$$

are each rational functions of t . Our notation does not reflect that f and g depend on P, a, b .

A *linear fractional transformation* is a map of the form

$$t \rightarrow \frac{\alpha t + \beta}{\gamma t + \delta}, \quad \alpha, \beta, \gamma, \delta \in \mathbf{R}, \quad \alpha\delta - \beta\gamma \neq 0.$$

Lemma 2.3 (Factor I) *If $n \geq 4k+2$ and P is a generically chosen n -Point, then $f(t)$ and $g(t)$ are each products of 4 linear fractional transformations. The zeros of f and g occur at the same points and the poles of f and g occur at the same points. Hence f/g is constant.*

The only reason we choose $n \geq 4k + 2$ in the Factor Lemma is so that the various diagonals involved in the proof do not have common endpoints. The Factor Lemma I works the same way for all k and for all choices of (large) n . We write $P \leftrightarrow Q$ if we can choose indices a, b and some $t \in \mathbf{R}$ such that $Q = P(t)$. The Factor Lemma implies that when P, Q are generic and $P \leftrightarrow Q$ we have $R(P) = R(Q)$. The result for non-generic choices of P follows from continuity. Any n -Point Q can be included in a finite chain

$$P_0 \leftrightarrow P_1 \leftrightarrow \cdots \leftrightarrow P_{2n} = Q,$$

where P_0 is the regular n -Point. Hence $R(Q) = R(P_0) = 1$. This shows that Theorem 2.1 holds for (k, n) where $k \geq 2$ and $n \geq 4k + 2$. (The case $k = 1$ is a main result of [19], and by now has many proofs.)

Lemma 2.4 *If Theorem 2.1 is true for all large values of n , then it is true for all values of n .*

Proof: If we are interested in the result for small values of n , we can replace a given PolyPoint P with its m -fold cyclic cover mP . We have $\chi_k(mP) = \chi_k(P)^m$ and $\chi_k(D_k(mP)) = \chi_k(D_k(p))^m$. Thus, the result for large n implies the result for small n . ♠

In view of Equation 4 we have

$$f(t) = f_1(t) \dots f_n(t), \quad f_j(t) = \chi(j, k, P(t)). \quad (12)$$

Thus $f(t)$ is the product of n “local” cross ratios. We call an index j *asleep* if none of the lines involved in the cross ratio $f_j(t)$ depend on t . In other words, the lines do not vary at all with t . Otherwise we call j *awake*.

As we vary t , only the diagonals $P_{0,h}$ change for $h = -k, -k - 1, k + 1, k$. From this fact, it is not surprising that there are precisely 4 awake indices. These indices are

$$j_0 = 0, \quad j_1 = k + 1, \quad j_2 = -k - 1, \quad j_3 = -k. \quad (13)$$

The index k is not awake because the diagonal $P_{0,k}(t)$ does not move with t .

We define a *chord* of $P(t)$ to be a line defined by a pair of vertices of $P(t)$. The point $P_0(t)$ moves at linear speed, and the 4 lines involved in the calculation of $f_{c_j}(t)$ are distinct unless $P_0(t)$ lies in one of the chords of $P(t)$.

Thus $f_{c_j}(t)$ only has zeros and poles at the corresponding values of t . It turns out that only the following chords are involved.

$$\begin{array}{cccccccc} -k & -k & -k & -k & -k-1 & -k-1 & k+1 & k+1 \\ -k-1 & k+1 & 1 & -2k-1 & -1 & -2k-1 & 1 & 2k+1 \end{array} \quad (14)$$

We call these c_0, \dots, c_7 . For instance, c_0 is the line through P_{-k} and P_{-k-1} . Let t_j denote the value of t such that $P(t_j) \in c_j$.

The PolyPoint $Q(t) = D_k(P(t))$ has the same structure as $P(t)$. Up to projective transformations $Q(t)$ is also obtained from the regular PolyPoint by moving a single vertex along one of the k -diagonals. The pattern of zeros and poles is not precisely the same because the chords of $Q(t)$ do not correspond to the chords of $P(t)$ in a completely straightforward way. The k -diagonals of $Q(t)$ correspond to the vertices of $P(t)$ and *vice versa*. The $(k+1)$ diagonals of $Q(t)$ correspond to the vertices of $\Delta_k^{-1}(P(t))$. This is what gives us our quadruples of points in the middle picture in Figure 2.1.

We now list the pattern of zeros and poles. We explain our notation by way of example. The quadruple $(f, 2, 4, 5)$ indicates that f_{c_2} has a simple zero at f_4 and a simple pole at t_5 .

$$(f, 0, 0, 1), \quad (f, 1, 6, 7), \quad (f, 2, 4, 5), \quad (f, 3, 2, 3). \quad (15)$$

$$(g, 0, 6, 5), \quad (g, 1, 0, 3), \quad (g, 2, 2, 1), \quad (g, 3, 4, 7). \quad (16)$$

Since these functions have holomorphic extensions to \mathbf{C} with no other zeros and poles, these functions are linear fractional transformations. This pattern establishes the Factor Lemma I.

Checking that the pattern is correct is just a matter of inspection. We give two example checks.

- To see why f_{c_2} has a simple zero at t_4 we consider the quintuple

$$(-k-1, -2k-1, -2k-2, 0, -1).$$

At t_4 the two diagonals $P_{-k-1,0}$ and $P_{-k-1,-1}$ coincide. In terms of the cross ratios of the slopes we are computing $\chi(a, b, c, d)$ with $a = b$. The point $P_0(t)$ is moving with linear speed and so the zero is simple.

- To see why g_{c_2} has a simple pole at t_1 we consider the 4 points

$$P_{2k+2,k+2} \cap P_{1,k+1}, \quad P_{k+1}, \quad P_1, \quad P_{1,k+1} \cap P_{-k,0}. \quad (17)$$

These are all contained in the k -diagonal $P_{1,k+1}$, which corresponds to the vertex $(-k-1)$ of $D_k(P)$. At $t = t_1$ the three points $P_0(t)$ and P_{-k} and P_{k+1} are collinear. This makes the 2nd and 4th listed point coincided. In terms of our cross ratio $\chi(a, b, c, d)$ we have $b = d$. This gives us a pole. The pole is simple because the points come together at linear speed.

The other explanations are similar. The reader can see graphical illustrations of these zeros and poles using our program.

2.4 Proof of the Second Result

The proof of Theorem 2.2 is essentially identical to the proof of Theorem 2.1. Here are the changes. The Factor Lemma II has precisely the same statement as the Factor Lemma I, except that

- When defining $P(t)$ we use points P_a and P_b with $|a - b| = k + 1$.
- We are comparing $P(t)$ with $D_{k+1}(P(t))$.

This changes the definition of the functions f and g . With these changes made, the Factor Lemma I is replaced by the Factor Lemma II, which has an identical statement. This time the chords involved are as follows.

$$\begin{array}{cccccccc} -k-1 & -k-1 & -k-1 & -k-1 & -k & -k & k & k \\ -k & k & -1 & -2k-1 & 1 & -2k-1 & -1 & 2k+1 \end{array} \quad (18)$$

This time the 4 awake indices are:

$$j_0 = 0, \quad j_1 = k, \quad j_2 = -k - 1, \quad j_3 = -k. \quad (19)$$

Here is the pattern of zeros and poles.

$$(f, 0, 1, 0), \quad (f, 1, 7, 6), \quad (f, 2, 3, 2), \quad (f, 3, 5, 4). \quad (20)$$

$$(g, 0, 5, 6), \quad (g, 1, 3, 0), \quad (g, 2, 7, 4), \quad (g, 3, 1, 2). \quad (21)$$

The pictures in these cases look almost identical to the previous case. The reader can see these pictures by operating my computer program. Again, the zeros of f and g are located at the same places, and likewise the poles of f and g are located at the same places. Hence f/g is constant. This completes the proof the Factor Lemma II, which implies Theorem 2.2.

3 The Soul of the Bird

3.1 Goal of the Chapter

Given a polygon $P \subset \mathbf{R}^2$, let \widehat{P} be the closure of the bounded components of $\mathbf{R}^2 - P$ and let P^I be the interior of \widehat{P} .

Suppose now that $P(t)$ for $t \in [0, 1]$ is a path in $B_{n,k}$ starting at the regular n -gon $P(0)$. We can adjust by a continuous family of projective transformations so that $P(t)$ is a bounded polygon in \mathbf{R}^2 for all $t \in [0, 1]$. We orient $P(0)$ counter-clockwise around $P^I(0)$. We extend this orientation choice continuously to $P(t)$. We let $P_{ab}(t)$ denote the diagonal through vertices $P_a(t)$ and $P_b(t)$. We orient $P_{a,b}(t)$ so that it points from $P_a(t)$ to $P_b(t)$. We take indices mod n .

When P is embedded, we say that P is *strictly star shaped* with respect to $x \in P^I$ if each ray emanating from x intersects P exactly once.

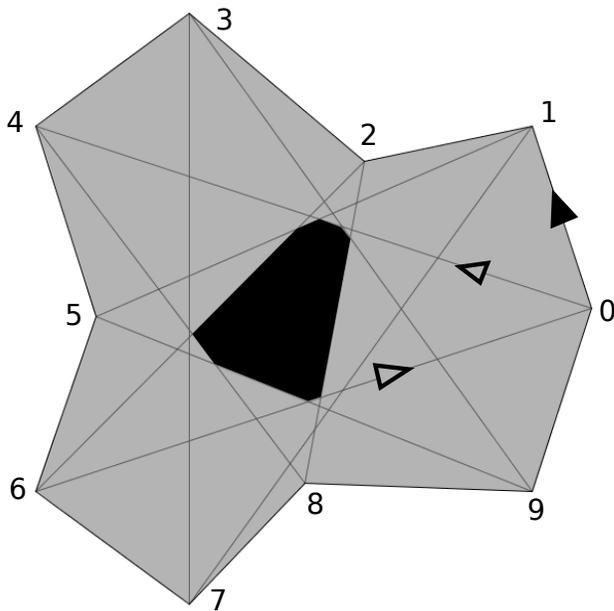


Figure 3.1: The soul of a 3-bird

Each such $(k + 1)$ -diagonal defines an oriented line that contains it, and also the (closed) *distinguished half plane* which lies to the left of the oriented line. These n half-planes vary continuously with t . The *soul* of $P(t)$, which we denote $S(t)$, is the intersection of the distinguished half-planes. Figure 3.1 shows the an example.

The goal of this chapter is to prove the following result.

Theorem 3.1 *Let P be a bird and let S be its soul. Then:*

1. S is has non-empty interior.
2. $S \subset P^I$.
3. P is strictly star-shaped with respect to any point in S .

Theorem 3.1 immediately implies Statement 1 of Theorem 1.1.

We are going to give a homotopical proof of Theorem 3.1. We say that a value $t \in [0, 1]$ is a *good parameter* if Theorem 3.1 holds for $P(t)$. All three conclusions of Theorem 3.1 are open conditions. Finally, 0 is a good parameter. For all these reasons, it suffices to prove that the set of good parameters is closed. By truncating our path at the first supposed failure, we reduce to the case when Theorem 3.1 holds for all $t \in [0, 1)$.

3.2 The Proof

For ease of notation we set $X = X(1)$ for any object X associated to $P(1)$.

Lemma 3.2 *If P is any k -bird, then P_0 and P_{2k+1} lie to the left of $P_{k,k+1}$. The same goes if all indices are cyclically shifted by the same amount.*

Proof: Consider the triangle with vertices $P_0(t)$ and $P_k(t)$ and $P_{k+1}(t)$. The k -niceness condition implies that this triangle is non-degenerate for all $t \in [0, 1]$. Since $P_0(t)$ lies to to the left of $P_{k,k+1}(t)$, the non-degeneracy implies the same result for $t = 1$. The same argument works for the triple $(2k + 1, k, k + 1)$. ♠

Lemma 3.3 *S is non-empty and contained in P^I .*

Proof: By continuity, S is nonempty and contained in $P \cup P^I$. By the k -niceness property and continuity, P_1 lies strictly to the right of $P_{0,k+1}$. Hence the entire half-open edge $[P_0, P_1)$ lies strictly to the right of $P_{0,k+1}$. Hence $[P_0, P_1)$ is disjoint from S . By cyclic relabeling, the same goes for all the other half-open edges. Hence $S \cap P = \emptyset$. Hence $S \subset P^I$. ♠

Lemma 3.4 P is strictly star-shaped with respect to any point of $S(1)$.

Proof: Since $P(t)$ is strictly star-shaped with respect to all points of $S(t)$ for $t < 1$, this lemma can only fail if there is an edge of $P(1)$ whose extending line contains a point $x \in S$. We can cyclically relabel so that the edge of $\overline{P_0P_1}$.

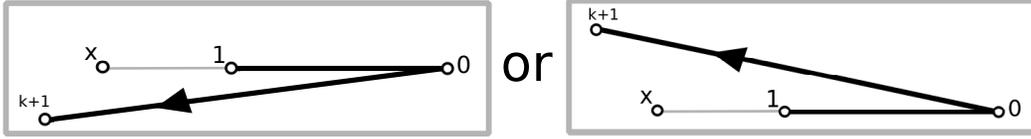


Figure 3.2: The diagonal $P_{0,k+1}$ does not separate 1 from x .

Since $x \notin P$, either P_1 lies between P_0 and x or P_0 lies in between x and P_1 . In the first case, both P_1 and x lie on the same side of the diagonal $P_{0,k+1}$. This is a contradiction: P_1 is supposed to lie on the right and x is supposed to lie on the left. In the second case we get the same kind of contradiction with respect to the diagonal $P_{-k,1}$. ♠

We say that P has *opposing $(k+1)$ -diagonals* if it has a pair of $(k+1)$ -diagonals which lie in the same line and point in opposite directions. In this case, the two left half-spaces are on opposite sides of the common line.

Lemma 3.5 P does not have opposing $(k+1)$ -diagonals.

Proof: We suppose that P has opposing diagonals and we derive a contradiction. In this case S , which is the intersection of all the associated left half-planes, must be a subset of the line L containing these diagonals. But then P intersects L in at least 4 points, none of which lie in S . But then P cannot be strictly star-shaped with respect to any point of S . This is a contradiction. ♠

We call three $(k+1)$ -diagonals of $P(t)$ *interlaced* if the intersection of their left half-spaces is a triangle. See Figure 3.3.

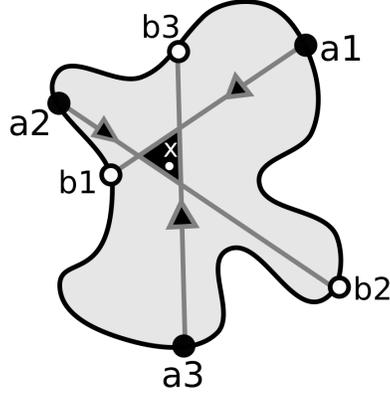


Figure 3.3: Interlaced diagonals on $P(t)$.

Given interlaced $(k + 1)$ -diagonals, and a point x in the intersection, the circle of rays emanating from x encounters the endpoints of the diagonals in an alternating pattern: $a_1, b_3, a_2, b_1, a_3, b_2$, where a_1, a_2, a_3 are the tail points and b_1, b_2, b_3 are the head points. Here a_1 names the vertex $P_{a_1}(t)$, etc.

Lemma 3.6 $P(t)$ cannot have interlaced diagonals for $t < 1$.

Proof: Choose $x \in S(t)$. The n -gon $P(t)$ is strictly star-shaped with respect to x . Hence, the vertices of P are encountered in order (mod n) by a family of rays that emanate from x and rotates around full-circle. Given the order these vertices are encountered, we have $a_{j+1} = a_j + \eta_j$, where $\eta_j \leq k$. Here we are taking the subscripts mod 3 and the vertex values mod n . This tells us that $n = \eta_1 + \eta_2 + \eta_3 \leq 3k$. This contradicts the fact that $n > 3k$. ♠

It only remains to show that S has non-empty interior. A special case of Helly's Theorem says the following: If we have a finite number of convex subsets of \mathbf{R}^2 then they all intersect provided that every 3 of them intersect. Applying Helly's Theorem to the set of interiors of our distinguished half-planes, we conclude that we can find 3 of these open half-planes whose triple intersection is empty. On the other hand, the triple intersection of the *closed* half-planes contains x . Since P has no opposing diagonals, this is only possible if the 3 associated diagonals are interlaced for t sufficiently close to 1. This contradicts Lemma 3.6. Hence S has non-empty interior.

4 The Feathers of the Bird

4.1 Goal of the Chapter

Recall that P^I is the interior of the region bounded by P . We call the union of shaded triangles in Figure 4.1 the *feathers* of the bird. the black region in the center is the soul.

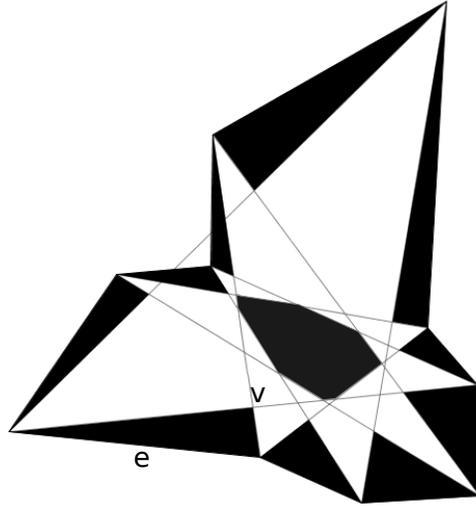


Figure 4.1 The feathers of a 3-bird.

Each feather F of a k -bird P is the convex hull of its *base*, an edge e of P , and its *tip*, a vertex of $\Delta_k(P)$.

The goal of this chapter is to prove the following result, which says that the simple topological picture shown in Figure 4.1 always holds.

Theorem 4.1 *The following is true.*

1. Let F be an feather with base e . Then $F - \{e\} \subset P^I$.
2. Distinct feathers can only intersect at a vertex of P .
3. The line segment connecting two consecutive feather tips lies in P^I .

When we apply Δ_k to P we are just specifying the points of $\Delta_k(P)$. We define the *polygon* $\Delta_k(P)$ so that the edges are the bounded segments connecting the consecutive tips of the feathers of P . With this definition, Statement 2 of Theorem 1.1 follows immediately from Theorem 4.1.

4.2 The Proof

There is one crucial idea in the proof of Theorem 4.1: The soul of P lies in the sector F^* opposite any of its feathers F . See Figure 4.2.

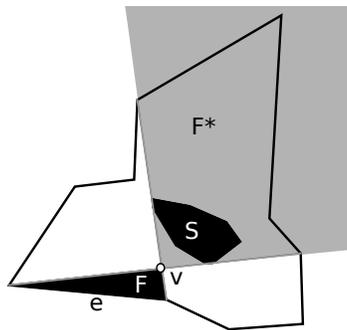


Figure 4.2 The soul lies in the sectors opposite the feathers.

We will give a homotopical proof of Theorem 4.1. By truncating our path of birds, we can assume that Theorem 4.1 holds for all $t \in [0, 1)$. We then want to rule out the various ways that Theorem 4.1 can fail for $t = 1$. As in the previous chapter we set $P = P(1)$, etc. Figure 4.3 shows the 2 ways that Statement 1 could fail:

1. The tip v of the feather F could coincide with some $p \in P$.
2. Some $p \in P$ could lie in the interior point of $\partial F - e$.

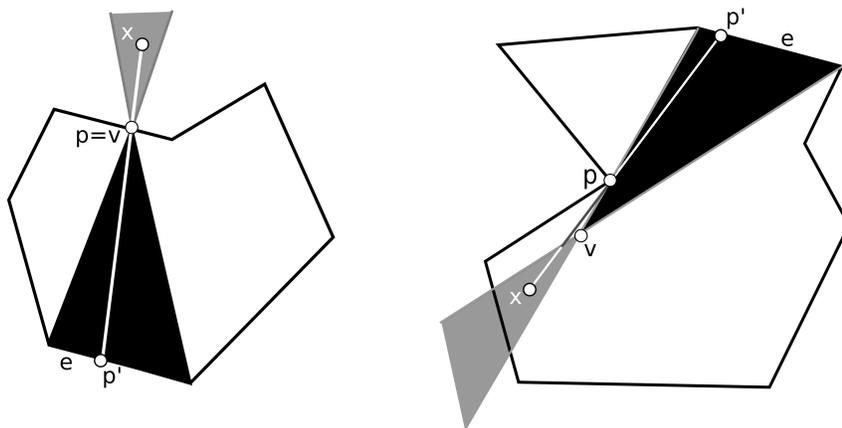


Figure 4.3: Case 1 (left) and Case 2 (right).

For all $x \in F^*$, the ray $\vec{x\hat{p}}$ intersects P both at p and at a point $p' \in e$. This contradicts the fact that for any $x \in S \subset F^*$, the polygon P is strictly star-shaped with respect to x . This establishes Statement 1 of Theorem 4.1.

Let F_1 and F_2 be two feathers of P , having bases e_1 and e_2 . For Statement 2, it suffices to prove that $F_1 - e_1$ and $F_2 - e_2$ are disjoint.

The same homotopical argument as for Statement 1 reduces us to the case when F_1 and F_2 have disjoint interiors but $\partial F_1 - e_1$ and $\partial F_2 - e_2$ share a common point x . If ∂F_1 and ∂F_2 share an entire line segment then, thanks to the fact that all the feathers are oriented the same way, we would have two $(k+1)$ diagonals of P lying in the same line and having opposite orientation. Lemma 3.5 rules this out.

In particular x must be the tip of at least one feather. Figure 4.4 shows the case when $x = v_1$, the tip of F_1 , but $x \neq v_2$. The case when $x = v_1 = v_2$ has a similar treatment.

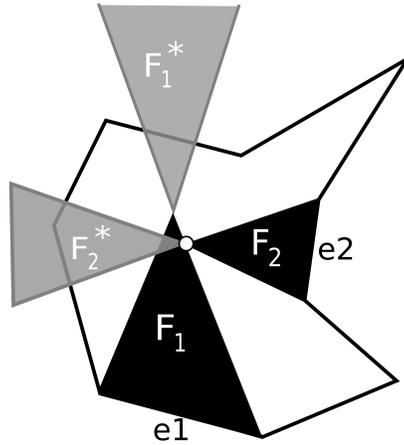


Figure 4.4: Opposiing sectors are disjoint

In this case, the two sectors F_1^* and F_2^* are either disjoint or intersect in a single point. This contradicts the fact that $S \subset F_1^* \subset F_2^*$ has non-empty interior. This contradiction establishes Statement 2 of Theorem 4.1.

Recall that $\hat{P} = P \cup P^I$. Let F_1 and F_2 be consecutive feathers with bases e_1 and e_2 respectively. Let f be the edge connecting the tips of F_1 and F_2 . Our homotopy idea reduces us to the case when $f \subset P$ and $f \cap P \neq \emptyset$. Figure 4.5 shows the situation.

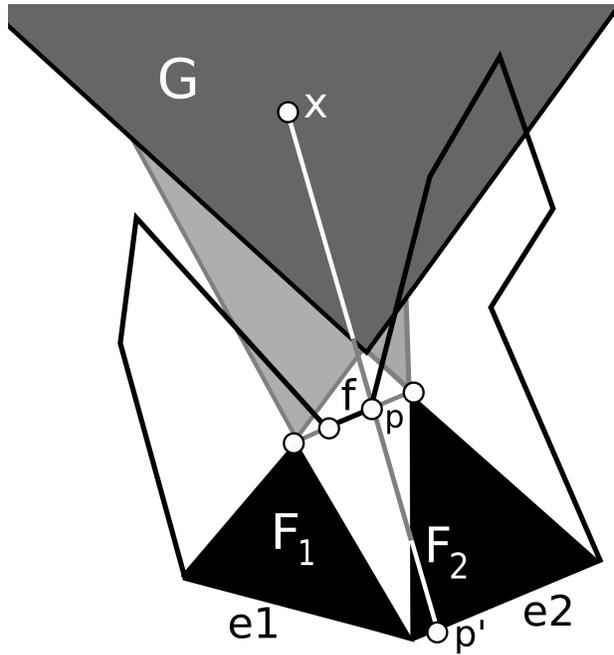


Figure 4.5: The problem a common boundary point

Note that $f \cap P$ must be strictly contained in the interior of f because (by Statement 1 of Theorem 4.1) the endpoints of f lie in P^I . But then, $f \cap P$ is disjoint from $F_1^* \cap F_2^*$, which is in turn contained in the shaded region G . For any $x \in G$ and each vertex p of f , the ray the ray $\vec{x\hat{p}}$ also intersects P at a point $p' \in e_1 \cup e_2$. This gives the same contradiction as above when we take $x \in S \subset F_1^* \cap F_2^* \subset G$. This completes the proof of Statement 3 of Theorem 4.1.

5 The Degeneration of Birds

5.1 Statement of the Result

Let $k \geq 2$ and $n > 3k$. Let $B_{k,n}$ denote the space of n -gons which are k -birds. Let χ_k denote the k -energy. In this chapter we will prove a technical result which will help us prove, in the next chapter, that $\Delta_k(B_{k,n}) = B_{k,n}$. The reader should probably just use the statement as a black box on the first reading. Our argument is a pretty tedious case-by-case analysis.

We say that a *degenerating path* is a path $Q(t)$ of n -gons such that

1. $Q(t) \in B_{k,n}$ for all $t \in [0, 1)$ but $Q(1) \notin B_{k,n}$.
2. $\chi_k(Q(t)) > \epsilon_0 > 0$ for all $t \in [0, 1]$.
3. $Q_j(1) \neq Q_{j+k}(1)$ for all $j = 1, \dots, n$.
4. $Q_j(1) \neq Q_{j+k+1}(1)$ for all $j = 1, \dots, n$.

Lemma 5.1 (Degeneration) *If $Q(\cdot)$ is a degenerating path, then all but at most one vertex of $Q(1)$ lies in a line segment.*

Let us first explain that this kind of degeneration can actually occur. Consider the case where $Q(t)$ is projectively equivalent to the regular n -gon for all $t \in [0, 1)$. That is $Q(t) = T_t(P)$ where P is the regular n -gon and T_t is some projective transformation that depends on t . We can choose T_t to be quite drastic, so that the points $Q_{-1}(t), Q_0(t), Q_1(t)$ are the vertices of a triangle and the remaining vertices converge to the line segment joining $Q_{\pm 1}(t)$. The limit $Q(1)$ will have the shape of an equilateral triangle, with all the vertices distinct and all but one vertex contained in the same line.

Remarks: (1) For our application of the Degeneration Lemma, all the vertices of $Q(1)$ are distinct. However, for a later application to Theorem 1.3 we prove the result under the weaker hypotheses we have stated.

(2) The Degeneration Lemma also works, with the same proof, in case we have a sequence $\{Q^\ell\} \subset B_{k,n}$, rather than a path, which converges to some $Q^\infty \notin B_{k,n}$ and has the uniform lower bound on χ_k . We will invoke the sequence case of the result when we prove Theorem 1.3.

5.2 Distinguished Diagonals

We orient $Q(t)$ so that it goes counter-clockwise around the region it bounds. We orient the diagonal Q_{ab} so that it points from Q_a to Q_b . For $t < 1$ the vertices $Q_1(t)$ and $Q_k(t)$ lie to the right of the diagonal $Q_{0,k+1}(t)$, in the sense that a person walking along this diagonal according to its orientation would see that points in the right. This has the same proof as Lemma 3.2. The same rule holds for all cyclic relabelings of these points. The rule holds when $t < 1$. Taking a limit, we get a weak version of the rule: Each of $Q_1(1)$ and $Q_k(1)$ either lies to the right of the diagonal $Q_{0,k+1}(1)$ or on it. The same goes for cyclic relabelings. We call this the *Right Hand Rule*.

Say that a *distinguished diagonal* of $Q(t)$ is either a k -diagonal or a $(k+1)$ -diagonal. There are $2n$ of these, and they come in a natural cyclic order:

$$Q_{0,k}(t) \quad Q_{0,k+1}(t), \quad Q_{1,k+1}(t), \quad Q_{1,k+2}(t), \dots \quad (22)$$

The pattern alternates between k and $(k+1)$ -diagonals. We say that a *diagonal chain* is a consecutive list of these.

We say that one oriented segment L_2 lies *ahead* of another one L_1 if we can rotate L_1 by $\theta \in (0, \pi)$ radians counter-clockwise to arrive at a segment parallel to L_2 . In this case we write $L_1 \prec L_2$. We have

$$Q_{0,k+1}(t) \prec Q_{1,k+1}(t) \prec Q_{1,k+2}(t) \prec Q_{2,k+2}(t). \quad (23)$$

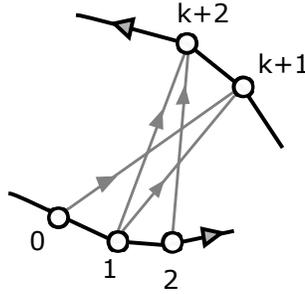


Figure 5.1: The turning rule

This certainly holds when $t = 0$. By continuity and the Right Hand Rule, this holds for all $t < 1$. Taking a limit, we see that the k -diagonals of $Q(1)$ weakly turn counter-clockwise in the sense that either $L_1 \prec L_2$ for consecutive diagonals or else L_1 and L_2 lie in the same line and point in the same direction. Moreover, the total turning is 2π . If we start with one distinguished diagonal and move through the cycle then the turning angle increases by jumps in $[0, \pi]$ until it reaches 2π . We call these observations *the Turning Rule*.

5.3 Subdivision into Cases

We set $X = X(1)$ for each object X associated to $Q(1)$. The situation is that $Q(t)$ is k -nice for all $t < 1$ but Q is not. Figure 5.2 shows the distinguished diagonals incident to Q_0 . We always take indices mod n . Thus $-k - 1 = n - k - 1 \pmod n$.

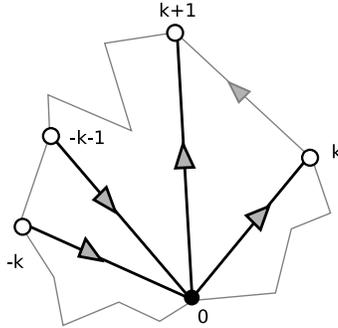


Figure 5.2: The 4 distinguished diagonals incident to $Q_0(t)$.

We say that Q has *collapsed diagonals* at Q_k if the 4 distinguished diagonals incident to Q_k do not all lie on distinct lines.

Lemma 5.2 *If Q has collapsed diagonals at Q_0 then $Q_{-k-1,0}$ and $Q_{0,k+1}$ point in opposite directions or $Q_{-k,0}$ and $Q_{0,k}$ point in the same direction.*

Proof: Associated to each diagonal incident to Q_0 is the ray which starts at Q_0 and goes in the direction of the other endpoint of the diagonal. (Warning: The ray may have the opposite orientation than the diagonal it corresponds to.) If the angle between any of the rays tends to π as $t \rightarrow 1$ then the angle between the outer two rays tends to π . In this case $Q_{-k,0}$ and $Q_{0,k}$ point in the same directions. If the angle between non-adjacent rays tends to 0 then $Q_{-k-1,0}$ and $Q_{0,k+1}$ are squeezed together and point in opposite directions.

Suppose that the angle between adjacent rays tends to 0. If the two adjacent rays are the middle ones, we have the case just considered. Otherwise, either the angle between the two left rays tends to 0 or the angle between the two right rays tends to 0. In either case, the uniform lower bound on the cross ratio forces a third diagonal to point either in the same or the opposite direction as these adjacent diagonals when $t = 1$. Any situation like this leads to a case we have already considered. ♠

5.4 The Case of Aligned Diagonals

We say that Q has *aligned diagonals* if there are 2 parallel distinguished diagonals of Q such that both diagonal chains which start and stop with these have length at least $2k$. This is one of the cases of Lemma 5.2.

Lemma 5.3 *If Q has aligned diagonals, then $2k + 1$ consecutive points of Q are collinear.*

Proof: The total turning of the diagonals is 2π , so one of the two chains defined by our diagonals turns 2π and the other turns 0. Hence, we can find $2k$ consecutive parallel distinguished diagonals. We will suppose that our chain starts with $Q_{-k,0}$ and ends with $Q_{0,k}$. The proof is essentially the same if the chain starts with a $(k + 1)$ -diagonal rather than a k -diagonal.

Given that $Q_{-k,0}$ and $Q_{-k,1}$ are parallel, Q_{-k}, Q_0, Q_1 are collinear. In other words, $(-k, 0, 1)$ is a triple of indices for collinear points. Likewise $(-k, -k + 1, 0)$ is such an index. Continuing this way, we get collinear triples

$$(-k, 0, 1), \quad (-k, -k + 1, 1), \quad (-k + 1, 1, 2), \quad \dots, \quad (-1, 0, k).$$

This implies that the points $Q_{-k}, \dots, Q_0, \dots, Q_k$ are all collinear. ♠

Now we forget about the aligned diagonals and we just use the property that Q has a consecutive run of $2k + 1$ collinear points. Let L be the line containing these points. If all points of Q lie in L , we are done. Otherwise there is some smallest index $i > k$ such that $Q_i \notin L$ but the preceding $2k + 1$ points are in L . Cyclically relabelling, we can assume that $i = k + 1$. Once we make this relabeling, we lose control over where our aligned diagonals are. Now we regain some control.

Lemma 5.4 *The length $2k$ -diagonal chain $Q_{-k,0} \rightarrow \dots \rightarrow Q_{0,k}$ consists entirely of parallel diagonals. There is no turning at all.*

Proof: The diagonals $Q_{-k,0}$ and $Q_{0,k}$ are either parallel or anti-parallel. If they are anti-parallel, then the angle between the corresponding rays incident $Q_0(t)$ tends to 0 as $t \rightarrow 1$. But these are the outer two rays. This forces the angle between all 4 rays incident to $Q_0(t)$ to tend to 0. The whole picture just folds up like a fan. But one of these rays corresponds to $Q_{0,k+1}(t)$. This picture forces $Q_{k+1} \in L$. But this is not the case.

Now we know that $Q_{-k,0}$ and $Q_{0,k}$ are parallel. All the diagonals in our chain are either parallel or anti-parallel to the first and last ones in the chain. If we ever get an anti-parallel pair, then the diagonals in the chain must turn 2π around. But then none of the other distinguished diagonals outside our chain turns at all. That is, $Q_{0,k}, Q_{1,k+1}, \dots, Q_{n-k,n}$ are all parallel. In this situation the argument proving Lemma 5.3 shows that $Q \subset L$, a contradiction. ♠

We rotate the picture so that L coincides with the X -axis and so that $Q_{0,k}$ points in the positive direction. Since we are already using the words *left* and *right* for another purpose, we say that $p \in L$ is *forward of* $q \in L$ if p has larger X -coordinate. Likewise we say that q is *backwards of* p in this situation. We say that $Q_{0,k}$ points *forwards*. We have established that $Q_{-k,0}, \dots, Q_{0,k}$ all point forwards.

Lemma 5.5 $Q_{k+2} \in L$ and both $Q_{1,k+2}$ and $Q_{2,k+2}$ point backwards.

Proof: Let us first justify the fact that Q_{k+1} lies above L . This follows from Right Hand Rule applied to $Q_{0,k+1}$ and Q_k and the fact that $Q_{0,k}$ points forwards. Since Q_{-k-1}, Q_{-k}, Q_1 are collinear, Q has collapsed diagonals at Q_1 . But Q cannot have aligned diagonals because $Q_{1,k+1}$ is not parallel to $Q_{-k,1}$. Hence Q has folded diagonals at 1. Since $Q_{-k,1}$ points forwards $Q_{1,k+2}$ points backwards.

We have $Q_2 \in L$ because $2 \leq k$. Suppose $Q_{2,k+2}$ points forwards. We consider the 3 distinguished diagonals

$$Q_{0,k}, \quad Q_{1,k+2}, \quad Q_{2,k+2}.$$

These diagonals respectively point forwards, backwards, forwards. But then, in going from $Q_{0,k}$ to $Q_{2,k+2}$, the diagonals have already turned 2π . Since the total turn is 2π , the diagonals $Q_{2,k+2}, Q_{3,k+3}, \dots, Q_{n,n+k}$ are all parallel. But then $Q_2, \dots, Q_n \in L$. This contradicts the fact that $Q_{k+1} \notin L$. ♠

Lemma 5.6 For at least of the two indices $j \in \{2k+2, 2k+3\}$ we have $Q_j \in L$ and $Q_{k+2,j}$ points forwards.

Proof: Since Q_1, Q_2, Q_{k+2} are collinear, Q has collapsed diagonals at Q_{k+2} . So, by Lemma 5.2, we either have folded diagonals at Q_{k+2} or aligned diagonals at Q_{k+2} . The aligned case gives $Q_{2k+2} \in L$ and the folded case gives $Q_{2k+3} \in L$. We consider the cases in turn.

Consider the aligned case. Suppose $Q_{k+2,2k+2}$ points backwards, as shown in Figure 5.3.

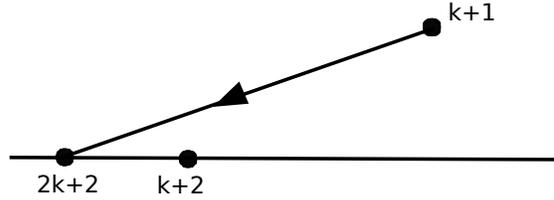


Figure 5.3: Violation of the Right Hand Rule

This violates the Right Hand Rule for Q_{k+2} and $Q_{k+1,2k+2}$ because Q_{k+1} lies above L .

Consider the folded case. Since $Q_{k+2,2k+3}$ and $Q_{1,k+2}$ point in opposite directions, and $Q_{1,k+2}$ points backwards (by the previous lemma), $Q_{k+2,2k+3}$ points forwards. ♠

Let $j \in \{2k + 2, 2k + 3\}$ be the index from Lemma 5.6. Consider the 3 diagonals

$$Q_{0,k}, \quad Q_{1,k+1}, \quad Q_{k+2,j}.$$

These diagonals are all parallel to L and respectively point in the forwards, backwards, forwards direction. This means that the diagonals in the chain $Q_{0,k} \rightarrow \dots \rightarrow Q_{k+2,j}$ have already turned 2π radians. But this means that the diagonals

$$Q_{k+2,2k+3}, \quad Q_{k+3,2k+3}, \quad Q_{k+3,2k+4}, \quad \dots \quad Q_{0,k} = Q_{n,n+k}$$

are all parallel and point forwards along L . Hence $Q_{k+2}, Q_{k+3}, \dots, Q_n \in L$. Hence all points but Q_{k+1} lie in L .

5.5 Separating the Soul from the Polygon

Recall that $Q = Q(1)$ and $Q_0 = Q_0(1)$, etc. It remains to analyze the case of folded diagonals, but before doing that we discuss some properties of the limiting soul. We define S to be the set of all accumulation points of sequences $\{p(t_n)\}$ where $p(t_n) \in S(t_n)$ and $t_n \rightarrow 1$. Since $S(t)$ is non-empty and closed for all $t < 1$, we see by compactness that S is also a non-empty closed subset of the closed region bounded by Q . Also, S lies to the left of all the half-planes defined by the oriented $(k + 1)$ diagonals.

Lemma 5.7 *The Degeneration Lemma is true for Q if $S \cap Q \neq \emptyset$.*

We prove this result through several smaller results.

Lemma 5.8 *The Degeneration Lemma is true for Q if $S \cap Q$ contains a point in the interior of an edge of Q .*

Proof: Suppose this happens. We relabel so that S contains an interior point of the edge $\overline{Q_0 Q_1}$. We rotate and scale so that $Q_0 = (0, 0)$ to $Q_1 = (1, 0)$. So $(x, 0) \in S$ for some $x \in (0, 1)$. By the Right Hand Rule, the vertex Q_{k+1} lies either on or above the X -axis. The only way for $(x, 0)$ to lie on or to the left of $Q_{0,k+1}$ is if $Q_{k+1} = (x_0, 0)$ for some $x_0 > 0$. Similar considerations show that $Q_{-k} = (x_1, 0)$ for some $x_1 < 1$. Figure 5.4 shows the relevant points of $Q(t)$ for t very near 1.

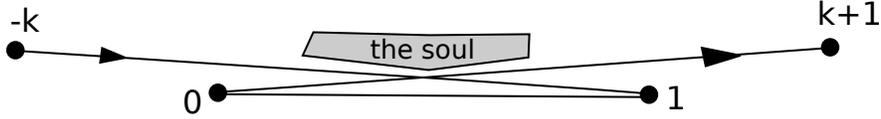


Figure 5.4: The relevant points and lines.

The diagonals $Q_{-k,1}$ and $Q_{0,k+1}$ are parallel. Given the indices involved, they are aligned. The proof in §5.4 finishes the proof. ♠

Lemma 5.9 *The Degeneration Lemma is true for Q if $S \cap Q$ contains a vertex of Q .*

Proof: We again relabel so that $Q_0 \in S$. We first give the proof when Q_{-1}, Q_0, Q_1 are all distinct. Figure 5.5 below shows the situation. This picture is meant to depict $Q(t)$ for t near 1. Figure 5.5 shows two cases, depending on whether the interior angle of Q at Q_0 is acute or obtuse. The interior angle of Q could also be 0 or 2π . We explain the various degenerate cases at the end. The shaded region in Figure 5.5 contains the soul, though this fact is not relevant to our argument.

The same analysis as in the previous lemma shows that Q_1, Q_0, Q_{-k} are collinear and Q_1 is an extreme point. In other words, Q_1 is not between Q_0 and Q_{-k} . Likewise Q_{-1}, Q_0, Q_k are collinear and Q_{-1} is an extreme point. We have drawn the cases when Q_0 is between $Q_{\pm 1}$ and $Q_{\mp k}$. This is not essential for the argument.

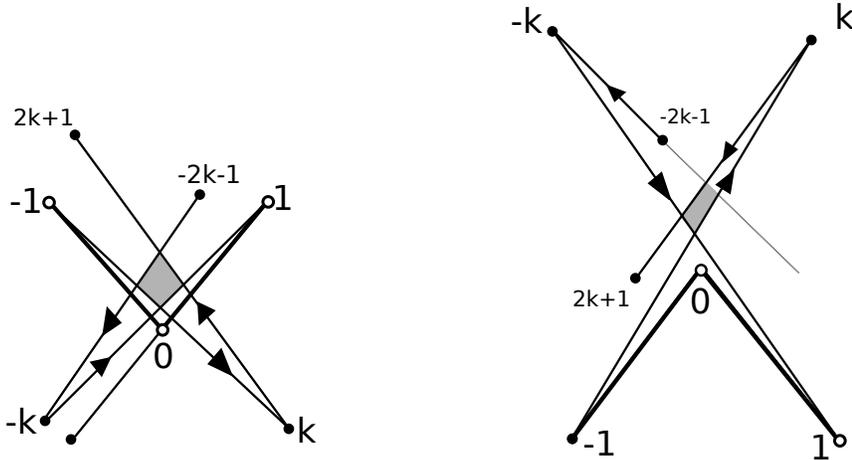


Figure 5.5: two pairs of folded diagonals

Since Q_{-1}, Q_0, Q_k are collinear, Q has collapsed diagonals at Q_k . If Q has aligned diagonals at Q_k then the proof in §5.4 finishes the job. So, we can assume Q has folded diagonals at Q_k . Likewise Q has folded diagonals at Q_{-k} . Now we look at the diagonal chain. We have

$$Q_{-1,k} \longrightarrow Q_{k,2k+1} \longrightarrow Q_{-2k-1,-k} \longrightarrow Q_{-k,1}. \quad (24)$$

If $n = 3k + 1$ then the inner two diagonals in Equation 24 coincide, and we have an impossible situation. So, we must have $n > 3k + 1$ in this case. But then these diagonals come in order in the diagonal chain. In going from successive diagonals on our list we turn respectively by π, θ, π degrees, where $\theta > 0$. This gives us $2\pi + \theta$ turning before we have completed the chain, violating the Turning Rule.

We now deal with the degenerate cases. When Q_{-1}, Q_0, Q_1 are distinct and lie in order on a single line, we have $\theta = 0$ above, and now we observe that the diagonals in the chain $Q_{-2k-1}, \dots, Q_{k,2k+1}$ are all parallel. This gives us more than $2k + 1$ consecutive parallel points on L , and the proof in §5.4 finishes the job.

if $Q_{-1} = Q_0$ then we directly see that the diagonals at Q_k are collapsed. This gives us our folded diagonals at Q_k . The same goes for Q_{-k} if $Q_0 = Q_1$. So, in all cases, we get the chain in Equation 24. Either we have $\theta > 0$ and an outright contradiction, or we have $\theta = 0$ and the proof finishes as above. ♠

5.6 Confining the Soul

We need one more lemma about the soul S of Q . This time we use the assumption about the folded diagonals. The diagonals $Q_{-k-1,0}$ and $Q_{0,k+1}$ point in opposite directions. We normalize, as above, so that both are contained in the X -axis and $Q_{0,k+1}$ points forwards. We treat the case when Q_{k+1} does not lie forwards of Q_{-k-1} . (The points could coincide.) We can always get to this case by dihedrally relabeling and then reflecting.

Lemma 5.10 *Suppose that Q does not have aligned diagonals. Then S is a subset of the closed line segment joining Q_0 to Q_{k+1} .*

Proof: Since S lies to the left of (or on) each $(k+1)$ diagonal, S is a subset of the line L common to the folded diagonals.

Now we consider the picture for $t < 1$. We rotate so that $Q_{0,k+1}(t)$ is horizontal and points forwards. We make a counter-clockwise turn of less than π to get from the ray $\overrightarrow{Q_0(t)Q_{k+1}(t)}$ to the ray $\overrightarrow{Q_0(t)Q_{-k-1}(t)}$. Hence $Q_{-k-1}(t)$ lies above L , as drawn in Figure 5.6.

Let $e(t)$ be the edge joining $Q_k(t)$ to $Q_{k+1}(t)$. The vertex $Q_k(t)$ lies below $Q_{0,k+1}(t)$ by the Right Hand Rule. Let $F(t)$ be the feather based at $e(t)$. The tip $v(t)$ of $F(t)$ lies on $Q_{0,k+1}(t)$, and $S(t)$ lies in the sector opposite $F(t)$ across $v(t)$.

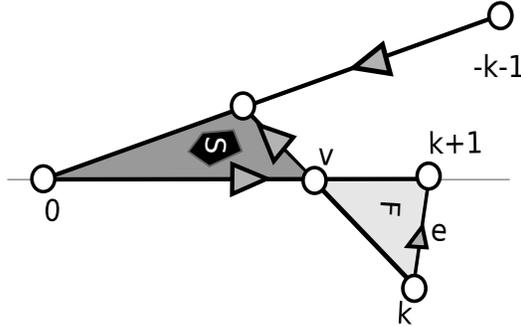


Figure 5.6: $S(t)$ is trapped in a triangle.

Hence $S(t)$ lies in the triangle $\Delta(t)$ bounded the lines

$$\overline{Q_k(t)v(t)}, \quad \overline{Q_0(t)Q_{k+1}(t)}, \quad \overline{Q_0(t)Q_{-k-1}(t)}.$$

This structure gives us the claim of the lemma. The only way $\Delta(t)$ can accumulate on points not on $0, k+1$ as $t \rightarrow 1$ is when $Q_{k,2k+1}$ is parallel to $Q_{0,k+1}$ in the limit. But then Q has aligned diagonals. ♠

5.7 Folded Edges

We say $Q(1)$ has a *fold* if two consecutive edges of $Q(1)$ are on the same line.

Lemma 5.11 (Folding) $Q(1)$ cannot have a fold.

Proof: Consider the lines extending the edges of $Q(t)$ incident to $Q_a(t)$. These lines bound 4 acute sectors. We let $C(t)$ be the sector which locally intersects $Q(t)$, as shown in Figure 5.7

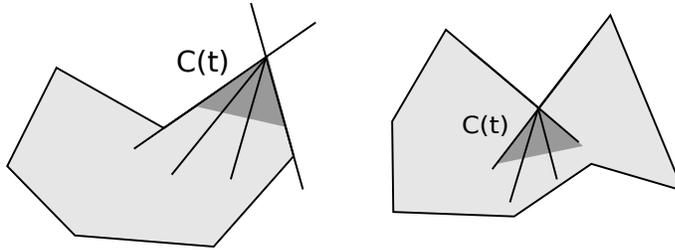


Figure 5.7: The cone

It follows from the k -niceness property and continuity that the $(k + 1)$ diagonals incident to $Q_a(t)$ lie in $C(t)$ for all $t < 1$. This is true even when the interior angle of $Q(t)$ at $Q_a(t)$ exceeds π . If this interior angle tends to either 0 or 2π then the angle of $C(t)$ tends to 0. This forces our two $(k + 1)$ diagonals to lie in the same line, violating the k -niceness of $Q(1)$. ♠

5.8 Good Folded Diagonals

Henceforth we assume that Q does not have aligned diagonals and also that we have $S \cap Q = \emptyset$.

We say that the folded diagonals $Q_{-k-1,0}$ and $Q_{0,k+1}$ are *good* if all the points $Q_{k+1}, Q_{k+2}, \dots, Q_{n-k-1}$ are collinear. We make a similar definition for other folded diagonals. Note that there are $n - 2k - 1$ of these points. If $n \geq 4k + 2$ then Q has at least $2k + 1$ consecutive collinear points, and the proof in Lemma 5.4 finishes this case. However, we might have $3k < n < 4k + 2$. In that case, we need the following lemma.

Lemma 5.12 *If $n > 3k + 2$ and some pair of folded diagonals of Q is good, then the Degeneration Lemma is true for Q .*

Proof: We normalize as in the previous section, so that $Q_{-k-1,0}$ and $Q_{0,k+1}$ are folded, and good.

Suppose $Q_1 \in L$. Since Q_0, Q_1, Q_{k+1} are collinear, Q has collapsed diagonals at Q_{k+1} . To avoid the proof in §5.4, we can assume that Q has folded diagonals at Q_{k+1} . Since $Q_{0,k+1}$ points forwards, $Q_{k+1,2k+2}$ points backwards. Hence $Q_{-k-1,0}$ and $Q_{k+1,2k+2}$ are parallel. Since $n > 3k + 2$ these two diagonals are aligned. This is a contradiction.

So, to finish the proof of this lemma, we show that $Q_1 \in L$. We first give the proof when all edges between Q_{k+1} and Q_{n-k-1} are nontrivial. We claim that Q_{k+2} is forwards of Q_{k+1} . Suppose not. Then there is some index $a \in \{k + 2, \dots, -k - 2\}$ such that Q_a is backwards of $Q_{a\pm 1}$. This gives us folded edges at Q_a .

By Lemma 5.11, we have folded diagonals at Q_a . But then $Q_{a,a+k+1}$ points forwards. Hence $Q_{0,k+1}$ and $Q_{a,a+k+1}$ are parallel. Given the indices, these diagonals are aligned. This is a contradiction.

Now we know that Q_{k+2} is forwards of Q_{k+1} . Suppose $Q_1 \notin L$. by the Right Hand Rule applied to the diagonal $Q_{0,k+1}$, the point Q_1 lies beneath L , as shown in Figure 5.8.

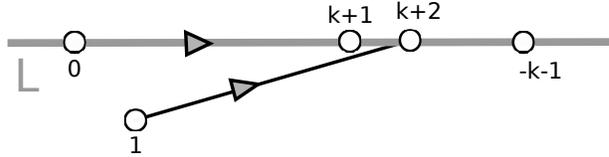


Figure 5.8: A contradiction involving Q_1 .

But then Q_{k+1} lies to the left of the diagonal $Q_{1,k+2}$. This violates the Right Hand Rule. Now we know that $Q_1 \in L$.

Now we give the proof that $Q_1 \in L$ when some edges are collapsed. That is, there is an index $a \in \{k + 2, \dots, n - k - 2\}$ such that $Q_a = Q_{a+1}$. Then the diagonals at Q_{a-k-1} are collapsed. In order to avoid the proof in §5.4 we must have folded diagonals at Q_{a-k-1} . Let L' be the line containing these folded diagonals. The proof in Lemma 5.10 shows that $S \subset L'$. But $S \subset L$. If $L' \neq L$ then $S = L \cap L' = \{Q_{a-k-1}\}$. But then S contains a vertex of S . This is a contradiction. Hence $L' = L$.

Since $S \cap Q = \emptyset$, the point Q_a lies forwards of Q_0 . If both Q_{a+k+1} and Q_{a-k-1} lie forwards of Q_a then Lemma 5.10 says that does not S lie backwards of Q_a . This is not the case. Hence Q_{a-k-1} lies backwards of Q_a .

Hence $Q_{0,k+1}$ and $Q_{a-k-1,a}$ are parallel. Consider the diagonal chain

$$Q_{0,k+1} \rightarrow \dots \rightarrow Q_{a-k-1,a}.$$

This chain either consists of parallel diagonals or else it twists 2π radians. The latter does not happen because the complementary diagonal chain contains a diagonal that is not parallel to $Q_{0,a}$, namely $Q_{-k-1,0}$. Hence $Q_{0,k+1}$ and $Q_{1,k+1}$ are parallel. Hence $Q_1 \in L$ in this case as well. ♠

Lemma 5.13 *If $n = 3k + 2$ and all folded diagonals of Q are good, then the Degeneration Lemma is true for Q .*

Proof: Suppose $n = 3k + 2$. We have all the same arguments as in the previous lemma. In particular, $Q_1 \in L$ and hence $Q_{0,k+1}$ and $Q_{k+1,2k+2}$ are folded. We need a different endgame because now the parallel diagonals $Q_{-k-1,0}$ and $Q_{k+1,2k+2}$ are not sufficiently well separated to be called aligned.

Again, $Q_{0,k+1}$ and $Q_{k+1,2k+2}$ are folded diagonals. Since these folded diagonals are good, the points $Q_{2k+2}, Q_{2k+3}, \dots, Q_0$ are collinear. We already know that $Q_{2k+2}, Q_0 \in L$. The collinearity gives $Q_{2k+2}, \dots, Q_0 \in L$. Since $n = 3k + 2$ we have $(2k + 2) = (-k) \pmod n$ we get $Q_{k+1}, \dots, Q_0 \in L$, which is a run of more than $2k + 1$ collinear points. The proof in §5.4 now shows that the Degeneration Lemma is true for Q . ♠

Lemma 5.14 *If $n = 3k + 1$ and all folded diagonals of Q are good, then the Degeneration Lemma is true for Q .*

Proof: We have the same set-up as in the previous result but this time all we can say is that the points

$$Q_{k+1}, \dots, Q_{2k}, Q_{2k+2}, \dots, Q_0 \in L.$$

Note that Q has collapsed diagonals at all these points except perhaps for Q_0 and Q_{k+1} . In order to avoid the proof in §5.4, we see that Q must have folded diagonals at all these points. Since all these folded diagonals are good, this suffices to show $Q_1, \dots, Q_k \in L$. ♠

5.9 Ungood Folded Diagonals

The only case left to consider is when Q has a pair of folded diagonals which are not good. Also, we can assume that Q has no aligned diagonals and $S \cap Q = \emptyset$. We normalize as above, so that Q_0, Q_{k+1}, Q_{-k-1} lie in forward order on L , which is the X -axis. Here we list some information we have.

- Not all of Q_{k+1}, \dots, Q_{-k-1} lie in L .
- Not all of $Q_{-k-1}, \dots, Q_0, \dots, Q_{k+1}$ lie in L . Otherwise we'd have $2k + 1$ consecutive collinear points, and the proof in §5.4 would finish the job.
- Q_{-k-1} and Q_{k+1} divide Q into 2 arcs, both of which start and end on L to the right of $x \in S$. This point x does not lie in Q .

We call an edge of Q *escaping* if $e \cap L$ is a single point. We call two different edges of Q , in the labeled sense, *twinned* if they are both escaping and if they intersect in an open interval. Even if two distinctly labeled edges of Q coincide, we consider them different as labeled edges.

Lemma 5.15 Q cannot have twinned escaping edges.

Proof: Consider $Q(t)$ for t near 1. This polygon is strictly star shaped with respect to a point $x(t)$ near x .

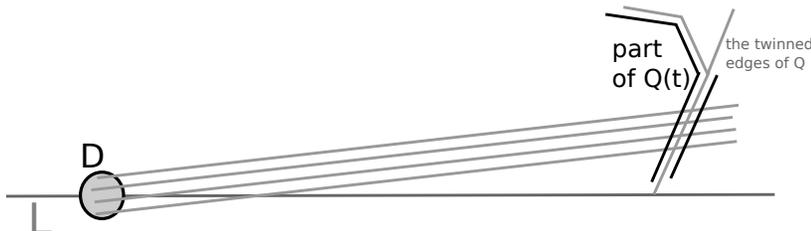


Figure 5.9: Rays intersecting the twinned segments

There is a disk D about x such that every $p \in D$ contains a ray which intersects the twinned edges in the middle third portion of their intersection. Figure 5.11 shows what we mean. Once t is sufficiently near 1, the soul $S(t)$ will intersect D , and for all points $p \in D$ there will be a ray which intersects $Q(t)$ twice. This contradicts the fact that $Q(t)$ is strictly star-shaped with respect to all points of $S(t)$. ♠

We say that an escape edge *rises above* L if it intersects the upper half plane in a segment.

Lemma 5.16 Q cannot have two escape edges which rise above L and intersect Q on the same side of the point x .

Proof: This situation is similar to the previous proof. In this case, there is a small disk D about x such that every point $p \in D$ has a ray which intersects both rising escape edges transversely, and in the middle third of each of the two subsegments of these escape edges that lie above L . Figure 5.10 shows this situation.

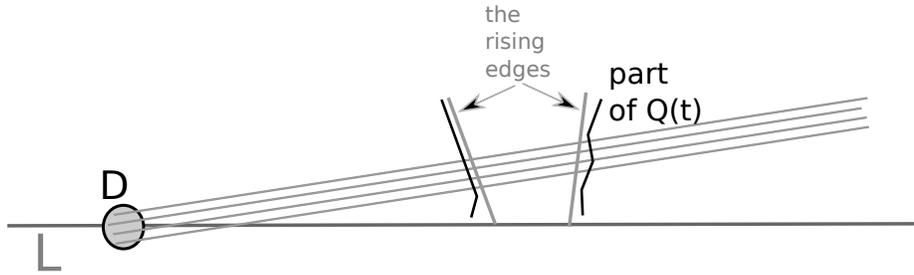


Figure 5.10: Rays intersecting the rising segments.

In this case, some part of $Q(t)$ closely shadows our two escape edges for t near 1. But then, once t is sufficiently near 1, each ray we have been talking about intersects $Q(t)$ at least twice, once by each escaping edge. This gives the same contradiction as in the previous lemma. ♠

We define falling escape segments the same way. The same statement as in Lemma 5.16 works for falling escape segments. Since $x \notin Q$ we conclude that Q can have at most 4 escaping segments total.

But $Q = Q_+ \cup Q_-$, where Q_{\pm} is an arc of Q that starts at Q_{k+1} and ends at Q_{-k-1} . Since both these arcs start and end on L , and since both do not remain entirely on L , we see that each arc has at least 2 escape edges. and none of these are twinned. This means that both Q_+ and Q_- have exactly two escape edges.

Now for the moment of truth: Consider Q_+ . Since Q_+ just has 2 escape edges, they both have to be either rising or falling. Also, since Q_+ starts and ends on the same side of x , and cannot cross x (By Lemma 3.1) we see that Q_+ has 2 of the same kind of escape arc on the same side of x . This is a contradiction. The same argument would work for Q_- but we don't need to do it twice.

6 The Persistence of Birds

In this chapter we prove Statement 3 of Theorem 1.1, namely the fact that $\Delta_k(B_{n,k}) = B_{n,k}$. First we use the Degeneration Lemma to prove that $\Delta_k(B_{n,k}) \subset B_{n,k}$. Then we deduce the opposite containment from projective duality and the factoring we discussed in §2.2.

6.1 Containment

Suppose for the sake of contradiction that there is some $P \in B_{k,n}$ such that $\Delta(P) \notin B_{k,n}$. Recall that there is a continuous path $P(t)$ for $t \in [0, 1]$ such that $P(0)$ is the regular n -gon.

Define $Q(t) = \Delta_k(P(t))$. There is some $t_0 \in [0, 1]$ such that $Q(t_0) \notin B_{k,n}$. We can truncate our path so that $t_0 = 1$. In other words, $Q(t) \in B_{k,n}$ for $t \in [0, 1)$ but $Q(1) \notin B_{k,n}$.

Lemma 6.1 $Q(\cdot)$ is a degenerating path.

Proof: Note that $Q(\cdot)$ satisfies Property 1 for degenerating paths. The energy χ_k is well-defined and continuous on $B_{k,n}$. Hence, by compactness, $\chi_k(P(t)) > \epsilon_0$ for some $\epsilon_0 > 0$ and all $t \in [0, 1]$. Now for the crucial step: We have already proved that

$$\chi_k \circ \Delta_k = \chi_k. \tag{25}$$

Hence $\chi_k(Q(t)) > \epsilon_0$ for all $t \in [0, 1]$. That is, $Q(\cdot)$ satisfies Property 2 for degenerating paths. Finally, if any two vertices of $Q(1)$ coincide, we violate Theorem 4.1 for the k -bird $P(1)$. Hence $Q(\cdot)$ satisfies Properties 3 and 4 for degenerating paths. In short, $Q(\cdot)$ is a degenerating path. ♠

We conclude that all but at most 1 vertex of $Q(1)$ lies in a line L . Stating this in terms of $P(1)$, we can say that all but at most one of the feathers of $P(1)$ have their tips in a single line L . Call an edge of $P(1)$ *ordinary* if the feather associated to it has its tip in L . We call the remaining edge, if there is one, *special*. Thus, all but at most one edge of P is ordinary.

Let $S(t)$ be the soul of $P(t)$. We know that $S(1)$ has non-empty interior by Theorem 3.1. The contradiction in our argument is to be that the arrangement of tips of $P(1)$ is going to force $S(1)$ to either be empty or to have empty interior.

Lemma 6.2 $P(1)$ cannot have ordinary edges e_1 and e_2 that lie on opposite sides of L and are disjoint from L .

Proof: Suppose this happens. Let F_1 and F_2 be the two associated feathers. Then the opposite sector F_1^* lies above L , and the opposite sector F_2^* lies below L and the two tips are distinct. But then $S(1)$ is empty. ♠

Lemma 6.3 $P(1)$ cannot have more than 2 ordinary edges crossing L .

Proof: As we trace along L in one direction or the other we encounter the first crossing edge and then the last one. Let v and v' be the tips of the corresponding feathers. Then the line segment from v to v' crosses all the other edges crossing L . But this contradicts Statement 3 of Theorem 4.1. ♠

We know that $P(1)$ cannot have ordinary edges on both sides of L and disjoint from L . We know also that at most 2 ordinary edges can cross L by Lemma 6.3. Finally, an ordinary edge cannot lie in L because then the tip would not. Hence, all but at most 2 of the ordinary edges of $P(1)$ lie on one side of L . Call this the *abundant side* of L . Call the other side the *barren side*. From this structure, we see that $P(1)$ has at most 2 vertices on the barren side of L that are not contained in L . At the same time, each ordinary edge on the abundant side contributes two vertices to the barren side: We just follow the diagonals comprising the corresponding feather. These diagonals cross L from the abundant side into the barren side. Two different ordinary edges contribute at least 3 distinct vertices to the barren side. This is a contradiction.

We have ruled out all possible behavior for $P(1)$ assuming that $Q(1)$ is degenerate. Hence, $Q(1)$ is not degenerate. This means that $Q(1)$ is a bird. This completes the proof that

$$\Delta_k(P_{k,n}) \subset P_{k,n}. \tag{26}$$

6.2 Equality

Now we show that $\Delta_k(B_{k,n}) = B_{k,n}$. We follow the setup and notation from §2.2. A *polygon* is a PolyPoint together with additional data specifying an edge in \mathbf{P} joining each consecutive pair of points. Dually, we get a polygon

from a PolyLine by specifying, for each pair of consecutive lines L_j, L_{j+1} , an arc of the pencil of lines through the intersection point which connects L_j to L_{j+1} .

Equation 8 implies that

$$\Delta_k^{-1} = D_{k+1} \circ \Delta_k \circ D_{k+1}. \quad (27)$$

Since the dual projective plane \mathbf{P}^* is an isomorphic copy of \mathbf{P} , it makes sense to define $B_{k,n}^*$. This space is just the image of $B_{k,n}$ under any projective duality. Below we prove

Theorem 6.4 $D_{k+1}(B_{k,n}) \subset B_{k,n}^*$.

It follows from Theorem 6.4 and projective duality that $D_{k+1}(B_{k,n}^*) \subset B_{k,n}$. This combines with Equation 27 and Equation 26 to show $\Delta_k^{-1}(B_{k,n}) \subset B_{k,n}$.

Implicit in the statement that $D_{k+1}(B_{k,n}) \subset B_{k,n}^*$ is that statement that we have a way to enhance D_{k+1} so that it maps polygons to polygons. We will explain this below. Theorem 6.4 combines with what we know already to prove that $\Delta(B_{k,n}) = B_{k,n}$.

Remark: In view of Equation 8, an alternate way of proving Statement 3 of Theorem 1.1 would be to show that $D_k(B_{k,n}) \subset B_{k,n}^*$. If we knew this, we could bypass the painful Degeneration Lemma. I couldn't see how to prove directly that $D_k(B_{k,n}) \subset B_{k,n}$.

Lemma 6.5 D_{k+1} maps a member of $B_{k,n}$ to an n -gon which is k -nice.

Proof: Let $Q = D_{k+1}(P)$. A $(k+1)$ -diagonal of Q is just a vertex of P . A k diagonal of Q is a vertex of $\Delta_k(p)$. Thus, to check the k -nice property for Q we need to take n -collections of 4-tuples of points and check that they are distinct. In each case, the points are collinear because the lines of Q are coincident.

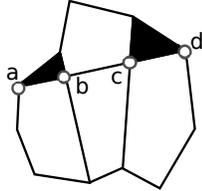


Figure 6.1 One of the n different 4-tuples we need to check.

Once we make this specification, there is really combinatorially only possibility for which collections we need to check. Figure 6.1 shows one such 4-tuple, a, b, c, d . The shaded triangles are the two feathers of P whose tips are b, c . But a, b, c, d are distinct vertices of $P \cup \Delta_k(P)$ and so they are distinct. That is all there is to it. ♠

Lemma 6.6 *If $P \in B_{k,n}$, then we can enhance $D_{k+1}(P)$ in such a way that $D_{k+1}(P)$ is a planar polygon in \mathbf{P}^* . The enhancement varies continuously.*

Proof: Specifying an enhancement of $D_{k+1}(P)$ is the same as specifying, for each consecutive pair L_1, L_2 of $(k+1)$ diagonals of P , an arc of the pencil through their intersection that connects L_1, L_2 . There are two possible arcs. One of them avoids the interior of the soul of P and the other one sweeps through the soul of P . We choose the arc that avoids the soul interior. Figure 6.2 shows that we mean for a concrete example.

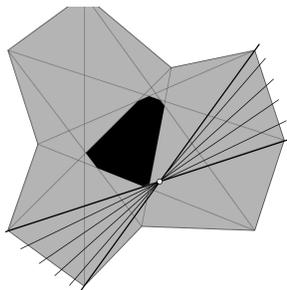


Figure 6.2: Enhancing a PolyLine to a polygon: Avoid the soul.

Since the soul of P has non-empty interior, there exists a point $x \in P$ which is disjoint from all these pencil-arcs. Applying duality, this exactly says that there is some line in \mathbf{P}^* which is disjoint from all the edges of our enhanced $D_{k+1}(P)$. Hence, this enhancement makes $D_{k+1}(P)$ planar. Our choice also varies continuously on $B_{n,k}$. ♠

To show that $Q = D_{k+1}(P)$ is a k -bird, we consider a continuous path $P(t)$ from the regular n -gon $P(0)$ to $P = P(1)$. We set $Q(t) = P(t)$. By construction, $Q(0)$ is a copy of the regular n -gon in \mathbf{P}^* , and $Q(t)$ is k -nice for all $t \in [0, 1]$, and $Q(t)$ is a planar polygon for all $t \in [0, 1]$. By definition $Q = Q(1)$ is a k -bird.

This completes the proof that $D_{k+1}(B_{k,n}) \subset B_{k,n}^*$, which in turn completes the proof that $\Delta_k(B_{k,n}) = B_{k,n}$. Our proof of Theorem 1.1 is done.

7 The Triangulation

7.1 Basic Definition

In this section we gather together the results we have proved so far and explain how we construct the triangulation τ_P associated to a bird $P \in B_{k,n}$.

Since $\Delta_k(B_{k,n}) \subset B_{k,n}$, we know that $\Delta_k(P)$ is also a k -bird. Combining this with Theorem 3.1 and Theorem 4.1 we can say that $\Delta_k(P)$ is one embedded n -gon contained in P^I , the interior of the region bounded by the embedded P . The region between P and $\Delta_k(P)$ is a topological annulus. Moreover, $\Delta_k(P)$ is obtained from P by connecting the tips of the feathers of P . The left side Figure 7.1 shows how this region is triangulated. The black triangles are the feathers of P and each of the white triangles is made from an edge of $\Delta_k(P)$ and two edges of adjacent feathers.

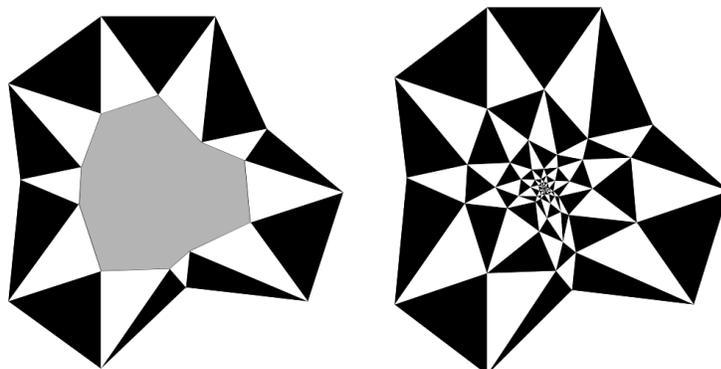


Figure 7.1: The triangulation of the annulus

Lemma 7.1 *For every member $P \in B_{k,n}$, the associated $2n$ triangles have pairwise disjoint interiors, and thus triangulate the annular region between P and $\Delta_k(P)$.*

Proof: As usual, we make a homotopical argument. If this result is false for some P , then we can look at path which starts at the regular n -gon (for which it is true) and stop at the first place where it fails. Theorem 4.1 tells us that nothing goes wrong with the feathers of P . The only thing that can go wrong is $\Delta_k(P)$ fails to be an embedded polygon. Since this does not happen, we see that in fact there is no counter-example at all. ♠

We can now iterate, and produce $2n$ triangles between $\Delta_k(P)$ and $\Delta_k^2(P)$, etc. The right side of Figure 7.1 shows the result of doing this many times. The fact that $\Delta_k(B_{k,n}) = B_{k,n}$ allows us to extend outward as well. When we iterate forever in both directions, we get an infinite triangulation of a (topological) cylinder that has degree 6 everywhere. This is what Figure 1.6 is showing. We call this bi-infinite triangulation τ_P .

7.2 Some Structural Results

Theorem 7.2 *Let $P \in B_{n,k}$. Let S be the soul of B . Then for $\ell \geq n$ we have $\Delta_k^\ell(P) \subset S$.*

Proof: We first note the existence of certain infinite polygonal arcs in τ_P . We start at a vertex of P and then move inward to a vertex of $\Delta_k(P)$ along one of the edges. We then continue through this vertex so that 3 triangles are on our left and 3 on our right. Figure 7.2 below shows the two paths like this that emanate from the same vertex of P .

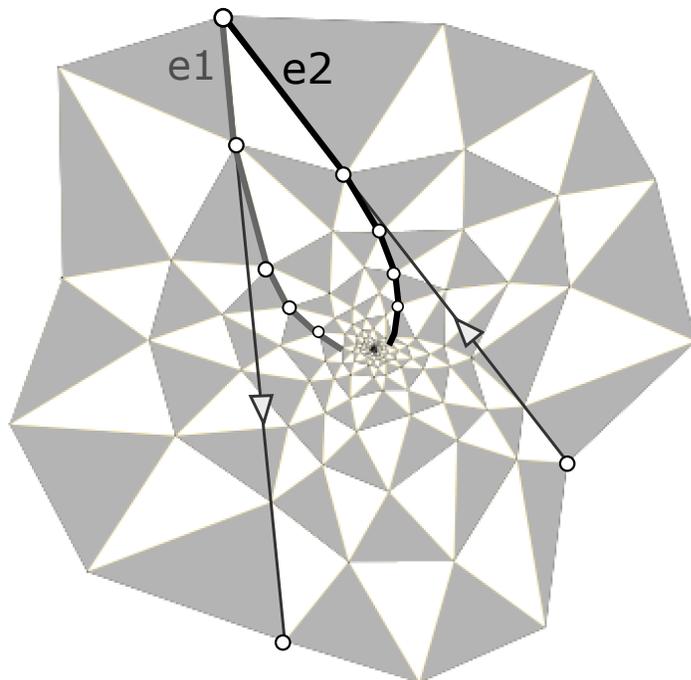


Figure 7.2: The spiral paths.

The usual homotopical argument establishes the fact that the spiral paths are locally convex. One can understand their combinatorics, and how they relate to the polygons in the orbit, just by looking at the case of the regular n -gon. We call the two spiral paths in Figure 7.2 *partners*. In the regular n -gon the partners intersect infinitely often. So this is true in general. Each spiral path has an initial segment joining the initial endpoint on P to the first intersection point with the partner. We define a *petal* to be the region bounded by the initial paths of the two partners.

It is convenient to write $P^\ell = \Delta_k^\ell(P)$. In the regular case, P^ℓ is contained in the petal for $\ell > n - 1$. Hence, the same goes in the general case. Because the initial segments are locally convex, the petal lies to the left of the lines extending the edges e_1 and e_2 when these edges are oriented according to the $(k + 1)$ -diagonals of P . But this argument works for every pair of partner spiral paths which start at a vertex of P . We conclude that for $\ell \geq n$, the polygon P^ℓ lies to the left of all the $(k + 1)$ -diagonals of P . But the soul of P is exactly the intersection of all these left half planes. ♠

Theorem 7.2 in turn gives us information about the nesting properties of birds within an orbit. Let S_ℓ denote the soul of P^ℓ . Let

$$S_\infty = \bigcap_{\ell \in \mathbf{Z}} S_\ell, \quad S_{-\infty} = \bigcup_{\ell \in \mathbf{Z}} S_\ell. \quad (28)$$

It follows from Theorem 7.2 that $\widehat{P}_\infty = S_\infty$ and $\widehat{P}_{-\infty} = S_{-\infty}$, because

$$S_{\ell+n} \subset P^{\ell+n} \subset S_\ell \subset P^\ell. \quad (29)$$

Hence these sets are all convex subsets of an affine plane.

Corollary 7.3 *Any $P \in B_{k,n}$ is strictly star-shaped with respect to all points in the convex hull of $\Delta_k^n(P)$.*

Proof: Since $P^{\ell+n} \subset S_\ell$, and P^ℓ is strictly star shaped with respect to all points of S^ℓ , we see that P^ℓ is strictly star shaped with respect to all points of $P^{\ell+n}$. Since S_ℓ is convex, we can say more strongly that P^ℓ is strictly star-shaped with respect to all points of the convex hull of $P^{\ell+n}$. Now we just set $\ell = 0$ and recall the meaning of our notation, we get the exact statement of the result. ♠

An immediate corollary is that P is strictly star-shaped with respect to \widehat{P}_∞ . (Theorem 1.3 says that this is a single point.)

8 Nesting Properties of Birds

In this chapter we prove Theorem 1.3. We just keep peeling away cases until the proof is done.

8.1 Discussion

This discussion is not part of our proof, but some people who are used to the pentagram map and its connection to projective geometry might appreciate it. Other readers can safely ignore this discussion.

Usually when one studies the pentagram map, say on convex polygons, one takes the quotient by the group of projective transformations. One can also take the quotient by the group of affine transformations, but this seems like a somewhat different thing. I want to relate the two concepts.

There is an asymptotic sense in which the affine quotient and projective quotient are quite closely related. Let Δ_1 be the pentagram map. We start with a convex polygon P and look at $P^\ell = \Delta_1^\ell(P)$ for very large $\ell > 0$. Suppose we choose an affine transformation T_ℓ so that $Q^\ell = T_\ell(P^\ell)$ remains bounded and in some sense uniformly fat. For instance, we could normalize so that a certain three vertices make an equilateral triangle. For large ℓ , the map T^ℓ has projective significance and not just affine significance.

Recall that as $\ell \rightarrow \infty$ the union $\bigcup \widehat{P}^{-\ell}$ converges to an affine plane $\widehat{P}_{-\infty}$. Here $\widehat{P}^{-\ell}$ is the closure of the region bounded by $P^{-\ell}$. The projective significance is that $T^\ell(\widehat{P}_\infty)$ is very nearly the affine patch of \mathbf{P} . Making such a normalization is usually something you would do with a projective transformation, but here we are doing it in an automatic way with an affine transformation. In the asymptotic forward direction, the best affine normalization really coincides with the best projective normalization.

To put this another way, the compactness of the full orbit $\{P^\ell, \ell \in \mathbf{Z}\}$ modulo projective transformations is very closely related to the compactness of the *forward* orbit modulo affine transformations. In the case of convex polygons, we can freely convert between one kind of compactness and the other. For the birds, I do not know how to do this and I find it easier to work with affine transformations.

For our proof, we will concentrate on Statement 1 of Theorem 1.3, and then bring back the projective geometry, in the form of projective duality, to deduce Statement 2 from Statement 1.

8.2 The Compact Case

Suppose that $S_1 \subset S_2$ are two compact convex sets. The ratio of diameters, $\text{diam}(S_1)/\text{diam}(S_2)$, is not affine invariant. We come up with a replacement notion. For each direction v in the plane, we let $\|S\|_v$ denote the maximum length of $L \cap S$ where L is a straight line parallel to v . We then define

$$\delta(S_1, S_2) = \sup_v \frac{\|S_1\|_v}{\|S_2\|_v} \in [0, 1]. \quad (30)$$

The quantity $\delta(S_1, S_2)$ is affine invariant, and (choosing a direction μ which realizes the diameter of S_1) we have

$$\frac{\text{diam}(S_1)}{\text{diam}(S_2)} \leq \frac{\|S_1\|_\mu}{\|S_2\|_\mu} \leq \delta(S_1, S_2). \quad (31)$$

Let $P \in B_{n,k}$ and suppose that P is normalized so as to be a subset of \mathbf{R}^2 . Let $P^\ell = \Delta_k^\ell(P)$. We define

$$\delta(P) = \delta(S(P^n), S(P)). \quad (32)$$

Here $S(P)$ is the soul of P . By Theorem 5.8, we have $S(P^n) \subset P^n \subset S(P)$, so our definition makes sense. Also, $\delta(P) < 1$ because $S(P^n)$ is contained in the interior of P^n .

We equip the space of compact convex subsets of \mathbf{R}^2 with the *Hausdorff metric*: The distance between two such subsets A_0, A_1 is the infimal ϵ such that A_j is contained in the ϵ -tubular neighborhood of A_{1-j} for $j = 0, 1$. This metric lets us talk about the convergence in an easy way. The function $\delta(\cdot, \cdot)$ is continuous with respect to this metric.

Suppose we have a sequence $\{Q^\ell\}$ of k -birds which converges to some other bird Q^∞ in the sense that the vertices and edges converge. Then $\Delta_k^n(Q^\ell)$ converges to $\Delta_k^n(Q^\infty)$ and the corresponding souls converge in the Hausdorff metric. This means that $\delta(Q^\ell) \rightarrow \delta(Q^\infty) = \delta_0 < 1$ for some δ_0 . We conclude that there is some $\delta_1 < 1$ so that $\delta(Q^\ell) < \delta_1 < 1$ for all $\ell > 0$.

If our forward orbit $\{P^\ell\}$ is compact modulo affine transformations, then we can find a sequence of affine transformations $\{T_\ell\}$ such that $Q^\ell = T_\ell(P^\ell)$ converges to another bird Q^∞ . The affine invariance of δ then tells us that $\delta(P^\ell) < \delta_1 < 1$. But then

$$\text{diam}(S(P^{\ell+n})) < \delta_1 \text{diam}(S(P^\ell)).$$

This shows that $\widehat{P}_\infty = \bigcap S(P^\ell)$ is a single point.

8.3 Normalizing by Affine Transformations

Henceforth we assume that the forward orbit $\{P^\ell\}$ of P under Δ_k is not compact modulo affine transformations. Our first step is to normalize as much as we can.

Lemma 8.1 *There is a sequence $\{T_\ell\}$ of affine transformations such that*

1. $T_\ell(P^\ell)$ has (the same) 3 vertices which make a fixed equilateral triangle.
2. T_ℓ expands distances on P^ℓ for all ℓ .
3. $T_\ell(P^\ell)$ is contained in a uniformly bounded subset of \mathbf{R}^2 .

Proof: To P^ℓ we associate the triangle τ_ℓ made from 3 vertices of P^ℓ and having maximal area. The diameter of τ_ℓ is uniformly small, so we can find a single equilateral triangle T and an expanding affine map $T_\ell : \tau_\ell \rightarrow T$. Let d be the side length of T . Every vertex of $T_\ell(P^\ell)$ is within d of all the sides of T , because otherwise we'd have a triangle of larger area. The sequence $\{T_\ell\}$ has the advertised properties. ♠

Let $Q^\ell = T_\ell(P^\ell)$. By compactness we can pass to a subsequence so that the limit polygon Q exists, in the sense that the vertices and the edges converge. Note that some vertices might collapse in the limit. This does not bother us. Let Q_0, Q_1 , etc. be the vertices of Q .

Each distinguished diagonal of Q^ℓ defines the unit vector which is parallel to it. Thus Q^ℓ defines a certain list of $2n$ unit vectors. We can pass to a subsequence so that all these unit vectors converge. Thus, to each distinguished diagonal of Q we still have a well-defined direction, even if the diagonal is trivial. We are keeping track of the 1-jet. Given a point and a unit direction, we have a well-defined oriented line which contains the point and is oriented along the direction, and a corresponding left half-plane. Thus we associate to each $(k+1)$ -diagonal, trivial or not, a left-half plane. We define \widehat{S} to be the intersection of all these half-planes.

We also define the limiting souls. We set $S^\ell = S(Q^\ell)$. We define the soul S of Q just as we defined it in §5.5: It is the set of accumulation points of sequences $\{p^\ell\}$ with $p^\ell \in S^\ell$. We will see that $S \subset \widehat{S}$.

8.4 The Limiting Soul

Lemma 8.2 $S \subset \widehat{S}$.

Proof: Fix $\epsilon > 0$. If this is not the case, then by compactness we can find a convergent sequence $\{p^\ell\}$, with $p^\ell \in S^\ell$, which does not converge to a point of \widehat{S} . But p^ℓ lies in every left half plane associated to Q^ℓ . But then, by continuity, the accumulation point p lies in every left half plane associated to Q . Hence $p \in \widehat{S}$. This is a contradiction. ♠

Corollary 8.3 *Suppose that \widehat{P}_∞ is not a single point. Then $\delta(S, H_Q) = 1$. Here H_Q is the convex hull of Q .*

Proof: Suppose not. Note that $H_{Q^\ell} \subset S^{\ell-n}$ by Theorem 7.2 and convexity. Then for ℓ large we have

$$\delta(Q^{\ell-n}) = \delta(S^\ell, S^{\ell-n}) \leq \delta(S^\ell, H_{Q^\ell}) < \delta(S, H_Q) + \epsilon,$$

and we can make ϵ as small as we like. This gives us a uniform $\delta < 1$ such that $\delta(Q^\ell) < \delta$ once ℓ is large enough. The argument in the compact case now shows that \widehat{P}_∞ is a single point. ♠

Corollary 8.3 gives us a powerful structural result. It says in particular that S and Q have the same diameter. Hence there is a chord $S^* \subset S$ which has the same diameter as Q . Since Q is a polygon, this means that Q must have vertices at either endpoint of S^* . We normalize so that S^* is the unit segment joining $(0, 0)$ to $(1, 0)$.

Lemma 8.4 *Let $Q' \subset Q$ be an arc of Q that joins $(0, 0)$ to $(1, 0)$.*

1. *The vertices of Q' must have non-decreasing x -coordinates.*
2. *If consecutive vertices of Q' have the same x -coordinate, they coincide.*
3. *Either $Q' \subset S^*$ or Q' intersects S^* only at $(0, 0)$ and $(1, 0)$.*

Proof: Suppose the Statement 1 is false. Then we can find a vertical line Λ which intersects S^* at a relative interior point and which intersects Q' transversely at 3 points. But then once ℓ is sufficiently large, Q^ℓ will intersect all vertical lines sufficiently close to Λ in at least 3 points and moreover some of these lines will contain points of S^ℓ . This contradicts the fact that Q^ℓ is strictly star-shaped with respect to all points of Q^ℓ .

For Statement 2, we observe that Q' does not contain any point of the form $(0, y)$ or $(1, y)$ for $y \neq 0$. Otherwise Q has larger diameter than 1. This is to say that once Q' leaves $(0, 0)$ it immediately moves forward in the X -direction. Likewise, once Q' (traced out the other way) leaves $(1, 0)$ it immediately moves backward in the X -direction. If Statement 2 is false, then we can find a non-horizontal line Λ' which intersects S^* in a relative interior point and which intersects Q' transversely at 3 points. The slope of Λ' depends on which of the two vertices of Q' lies above the other. Once we have Λ' we play the same game as for the first statement, and get the same kind of contradiction.

Suppose Statement 3 is false. We use the kind of argument we had in §5.9. By Statements 1 and 2 together, Q' must have an escape edge which touches S^* in a relative interior point. Moreover, this one escape edge is paired with another escape edge. Thus we can find a point $x \in S^*$ which strictly lies on the same side of both of these same-type escape edges. The argument in §5.9 now shows that Q^ℓ is not strictly star-shaped with respect to points of S^ℓ very near x . ♠

Corollary 8.5 *Up to adding repeated indices, Q is embedded.*

Proof: Lemma 8.4 implies that up to adding repeated vertices, Q is the union of two embedded arcs which connect $(0, 0)$ to $(1, 0)$, both of which are graphs of functions on $[0, 1]$. Call these functions f and g . Since $Q \not\subset S^*$, at least one of these functions is strictly nonzero. If it ever happens that $f(p) = g(p)$ for $p \in (0, 1)$, then the two arcs intersect at some point not on S^* . But then some vertical ray through a relative interior point of S^* intersects Q transversely at two points, giving the same contradiction. Hence our two arcs are disjoint. ♠

Corollary 8.6 *Suppose $0 \leq a < b < n$ and $Q_a = Q_b$. Then either we have $Q_a = Q_{a+1} = \dots = Q_b$ or else we have $Q_b = Q_{b+1} = \dots = Q_{a+n}$.*

8.5 The Triangular Limit Case

Here we prove a more general result that covers the triangular limit case. Suppose that there is a line L such that $Q_0 \notin L$ and $Q_j \in L$ unless $j \in \{-k+1, \dots, 0, \dots, k-1\}$. This is a run of $2k-1$ consecutive indices.

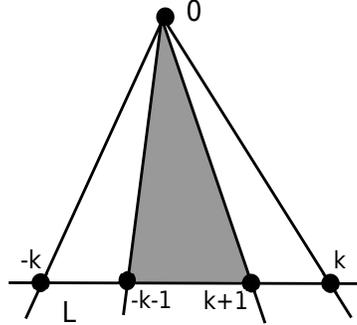


Figure 8.1: The triangular limit Q .

The cross ratio of the lines

$$Q_{0,k}, Q_{0,k+1}, Q_{n-k-1,0}, Q_{n-k,0}$$

is at least ϵ_0 . Also, these lines are cyclically ordered about 0 as indicated in Figure 8.1, thanks to the k -niceness property and continuity. Also, the two lines containing $Q_{0,k}$ and $Q_{-k,0}$ are not parallel because $Q_0 \notin L$. Hence S is contained in the shaded region in Figure 8.1, namely the triangle with vertices Q_0 and $Q_{\pm(k+1)}$. But this shaded region has diameter strictly smaller than the triangle τ with vertices Q_0 and $Q_{\pm k}$. Hence $\text{diam}(S) < \text{diam}(\tau) \leq \text{diam}(Q)$. This is a contradiction.

8.6 The Case of Folded Diagonals

In this section we suppose Q has a pair of folded diagonals. We relabel so that the folded diagonals are $Q_{-k-1,0}$ and $Q_{0,k+1}$.

The left half planes defined by these folded diagonals intersect in a line. Hence \widehat{S} is contained in a line L . Since $S \subset \widehat{S}$, we see that $S \subset L$. Indeed, we must have $S^* = S$, where S^* is as in §8.4. We will repeatedly use the fact that S realizes the diameter of Q , so that points of Q not in S do not lie in the line L containing S . There are 3 cases.

Case 1: Suppose that Q_{k+1} is not an endpoint of S . Then by Lemma 3.6 the arc $Q_0 \rightarrow \dots \rightarrow Q_{k+1}$ lies in S . Likewise the arc $Q_{k+1} \rightarrow \dots \rightarrow Q_{n-k-1}$ also lies in S . Hence $Q_j \in L$ unless $j \in \{-k, \dots, -1\}$. We can cyclically relabel so that this case is covered by the result in §8.5. Exactly the same argument works when Q_{-k-1} is not an endpoint of S .

Case 2: Suppose Q_{k+1} and Q_{-k-1} are the same endpoint of S and that $Q_0 \neq Q_{k+1}$. In this case, Corollary 8.6 says that $Q_{k+1} = \dots = Q_{-k-1} \in S$. If $Q_{\pm k} \in S$ then the case in §8.5 would cover us. So, one of these points does not lie in S . We consider the case when $Q_k \notin S$. The other case has the same treatment.

Suppose $n > 3k + 1$. Then $2k, 2k + 1 \in \{k + 1, \dots, n - k - 1\}$. Hence $Q_{k,2k}$ and $Q_{k,2k+1}$ are nontrivial and contained in some line $L' \neq L$. The notions of collapsed diagonals, folded diagonals, and aligned diagonals from §5 make sense for Q because the concepts just involve the directions of the diagonals. Likewise, Lemmas 5.2 and 5.3 hold for Q . By construction Q has collapsed diagonals at Q_k . If Q has folded diagonals at Q_k then $S \subset L'$, a contradiction. Hence Q has aligned diagonals at Q_k . Lemma 5.3 applies either to the short diagonal chain and gives us $Q_0, \dots, Q_{2k} \in L'$ or else it applies to the long diagonal chain and gives us all points of Q in L' . In either case $Q_0 \subset L'$. But then L' contains both Q_0 and $Q_{2k+1} = Q_{k+1} \neq Q_0$. So does L . Hence $L = L'$. This is a contradiction.

Suppose $n = 3k + 1$. In the next section we will show that when Q has a folded diagonal, Q always has a run of $k + 1$ repeating points (which is stronger than the k repeating points we could get from the argument above in this case.) We cyclically relabel so that $Q_{k+1} = \dots = Q_{2k+1}$. If $Q_k \in S$ then §8.5 covers the case. Otherwise $Q_k \notin S$ and we can run the argument as above.

Case 3: Suppose that Q_{k+1} and Q_{-k-1} are each endpoints of S . Either $Q_{k+1} = Q_{-k-1}$ or $Q_0 = Q_{k+1}$ or $Q_0 = Q_{-k-1}$. We cannot have Q_0 strictly between these points because then $Q_{-k-1,0}$ and $Q_{0,k+1}$ are not folded. If $Q_{k+1} = Q_{-k-1}$ then to avoid Case 2 we have $Q_0 = Q_{k+1} = Q_{-k-1}$. In all cases, Lemma 8.6 then gives us at least $k + 1$ consecutive points of Q which coincide. So, we have some index a such that $\nu := Q_a = \dots = Q_{a+k} \in S$. Using Lemma 3.6 we can assume that either $Q_{a-1} \notin L$ or $Q_{a+k+1} \notin L$. Again L is the line containing S . We consider the second case. The first case has the same treatment.

Let $L' \neq L$ denote the line containing ν and Q_{a+k+1} . Two of the distinguished diagonals connect Q_{a+k+1} to ν and hence point along L' . Hence Q has collapsed diagonals at Q_{a+k+1} . If Q has folded diagonals at Q_{a+k+1} then $S \subset L'$. This is contradiction. Hence $Q_{a+1,a+k+1}$ and $Q_{a+k+1,a+2k+1}$ are aligned and point along L' . The same argument as in Case 2 now tells gives us $2k$ parallel diagonals $Q_{a+1,a+k+1}, \dots, Q_{a+k+1,a+2k+1}$ which point along L' .

Let us now cyclically relabel so that $a = 0$. This means that our parallel diagonal chain is $Q_{0,k}, \dots, Q_{k,2k}$. This is all we are going to use in our proof for the rest of this case. Lemma 5.3 tells us that $Q_0, \dots, Q_{2k} \in L'$.

The two diagonals $Q_{0,k+1}$ and $Q_{1,k+1}$ are on our list and hence parallel. Hence Q has collapsed diagonals at Q_{k+1} . To avoid $S \subset L'$, as in Case 2, we must have $Q_{1,k+1}$ and $Q_{k+1,2k+1}$ parallel to each other and also parallel to the $2k$ diagonals we already have. this extends our parallel diagonal chain from length $2k$ to length $2k + 2$. We now repeat this argument indefinitely, showing that all distinguished diagonals point along L' . Hence $Q \subset L'$, a contradiction.

8.7 No Folded Diagonals

If Q has a trivial distinguished diagonal, then after relabeling, we can say that $Q_0 = Q_{k+i}$ for one of $i = 0, 1$. In all cases, Lemma 8.6 gives us a run of $k + 1$ repeated points. Cyclically relabeling again, if necessary, we can arrange that $Q_0 = \dots = Q_k$ but $Q_{k+1} \neq Q_0$. The only case we have not considered is when Q has no folded diagonals. Let L' be the line through Q_0 and Q_{k+1} . Then Q has collapsed diagonals at Q_{k+1} . Since Q has no folded diagonals, Q has aligned diagonals at Q_{k+1} . These diagonals point along L' . As in Case 3 above, we get a parallel diagonal chain of length $2k$. Since we have no folded diagonals, we can repeat the argument at the end of Case 3 indefinitely to show $Q \subset L'$, a contradiction.

8.8 Applying Duality

We have finished proving that \widehat{P}_∞ is a point. Now we prove that $\widehat{P}_{-\infty}$ is an affine plane. We take $\ell \geq 0$ and consider $P^{-\ell} = \Delta_k^{-\ell}(P)$.

Lemma 8.7 *There is a line $L \subset \mathbf{P}$ which is disjoint from $\widehat{P}^{-\ell}$ for all ℓ .*

Proof: Let Ω_ℓ be the set of lines in \mathbf{P} which are disjoint from the interior of $P^{-\ell}$. Since $P^{-\ell}$ is planar, this set is nonempty. The sets $\{\Omega_\ell\}$ are nested

and hence have a non-empty intersection. Let L be a line in the intersection. Since the interior of the region bounded by $P^{-\ell-1}$ contains $P^{-\ell}$ we see that L is actually disjoint from $P^{-\ell}$ as well. ♠

We normalize so that L is the line at infinity and P_∞ is the origin. Let

$$\Pi^\ell = \Delta_k^\ell(D_{k+1}(P)). \quad (33)$$

Then P^ℓ is planar polygon that is strictly star-shaped with respect to the origin. The map D_{k+1} conjugates Δ_k to Δ_k^{-1} and maps k -birds to k -birds. $\{\Pi^\ell\}$ shrinks to a point in the dual projective plane \mathbf{P}^* . Because the vertices of Π^ℓ shrink to a single point, all the $(k+1)$ -diagonals of $P^{-\ell}$ converge to a single line L' . This is enough to prove that $\widehat{P}_{-\infty}$ is either all of \mathbf{R}^2 or an infinite strip in \mathbf{R}^2 .

To rule out the strip case we note that the edges of Π^ℓ shrink to a point as well. We rotate so that the supposed strip is vertical. Then L' is a vertical line. Consider the soul $S^{-\ell}$ of $P^{-\ell}$. Let v_ℓ be the vertex of $S^{-\ell}$ with the largest y -coordinate. Let Υ_ℓ be the set of lines through v_ℓ which avoid the interior of $S^{-\ell}$. Compare Figure 6.6. This set corresponds to an edge of Π^ℓ . Hence Υ_ℓ converges to L' as well. But Υ_ℓ always contains a horizontal line. This horizontal line either converges on a subsequence to the line L at infinity or else to some horizontal line in \mathbf{R}^2 . In either case we get a limiting line that does not equal L' . This is a contradiction.

This completes the proof of Statement 2 of Theorem 1.3. Our proof of Theorem 1.3 is done.

9 Appendix

9.1 The Energy Invariance Revisited

In this section we sketch Anton Izosimov's proof that $\chi_k \circ \Delta_k = \chi_k$. This proof is more conceptual than the one in §2 but it is not self-contained. It requires the machinery from [6]. (The perspective comes from [8], but the needed result for Δ_k is in the follow-up paper [6].)

Let P be an n -gon. We let V_1, \dots, V_n be points in \mathbf{R}^3 representing the consecutive vertices of P . Thus the vertex P_j is the equivalence class of V_j . We can choose periodic sequences $\{a_i\}, \{b_i\}, \{c_i\}, \{d_i\}$ such that

$$a_i V_i + b_i V_{i+k} + c_i V_{i+k+1} + d_i V_{i+2k+1} = 0, \quad \forall i. \quad (34)$$

Recall from §2.2 that $\Delta_k = D_k \circ D_{k+1}$.

Lemma 9.1 *One of the cross ratio factors of $\chi_k \circ D_{k+1}$ is $(a_0 d_{-k}) / (c_0 b_{-k})$.*

Proof: One of the factors is the cross ratio of P_0, y, x, P_{k+1} , where

$$x = P_{0,k+1} \cap P_{k,2k+1}, \quad y = P_{-k,1} \cap P_{0,k+1}.$$

(Compare the right side of Figure 2.1, shifting all the indices there by $k+1$.)

The points x and y respectively are represented by vectors

$$X = a_0 V_0 + c_0 V_{k+1} = -b_0 V_k - d_0 V_{2k+1},$$

$$Y = -a_{-k} V_{-k} - c_{-k} V_1 = b_{-k} V_0 + d_{-k} V_{k+1}.$$

The point here is that the vector X lies in the span of $\{V_0, V_{k+1}\}$ and in the span of $\{V_k, V_{2k+1}\}$ and projectively this is exactly what is required. A similar remark applies to Y .

Setting $\Omega = V_0 \times V_{k+1}$, we compute the relevant cross ratio as

$$\frac{V_0 \times Y}{V_0 \times X} \cdot \frac{X \times V_{k+1}}{Y \times V_{k+1}} = \frac{d_{-k} \Omega}{c_0 \Omega} \times \frac{a_0 \Omega}{b_{-k} \Omega} = \frac{d_{-k} a_0}{b_{-k} c_0}, \quad (35)$$

which is just a rearrangement of the claimed term. ♠

The other cross ratio factors are obtained by shifting the indices in an obvious way. As an immediate corollary, we see that

$$\chi_k(D_{k+1}(P)) = \prod_{i=1}^n \frac{a_i d_i}{b_i c_i}. \quad (36)$$

Let us call this quantity $\mu_k(P)$.

Lemma 9.2 *If $\mu_k \circ \Delta_k = \mu_k$ then $\chi_k \circ \Delta_k = \chi_k$.*

Proof: If $\mu_k \circ \Delta_k = \mu_k$ then $\mu_k \circ \Delta_k^{-1} = \mu_k$. Equation 36 says that

$$\chi_k \circ D_{k+1} = \mu_k, \quad \mu_k \circ D_{k+1} = \chi_k. \quad (37)$$

The first equation implies the second because D_{k+1} is an involution. Since D_{k+1} conjugates Δ_k to Δ_k^{-1} we have

$$\chi_k \circ \Delta_k = \chi_k \circ D_{k+1} \circ \Delta_k^{-1} \circ D_{k+1} = \mu_k \circ \Delta_k^{-1} \circ D_{k+1} = \mu_k \circ D_{k+1} = \chi_k.$$

This completes the proof. ♠

Let $\tilde{P} = \Delta_k(P)$. Let $\{\tilde{a}_i\}$, etc., be the sequences associated to \tilde{P} . We want to show that

$$\prod_{i=1}^n \frac{a_i d_i}{b_i c_i} = \prod_{i=1}^n \frac{\tilde{a}_i \tilde{d}_i}{\tilde{b}_i \tilde{c}_i}. \quad (38)$$

This is just a restatement of the equation $\mu_k \circ \Delta_k = \mu_k$.

Now we use the formalism from [6] to establish Equation 38. We associate to our polygon P operator D on the space \mathcal{V} of bi-infinite sequences $\{V_i\}$ of vectors in \mathbf{R}^3 . The definition of D is given coordinate-wise as

$$D(V_i) = a_i V_i + b_i T^k(V_i) + c_i T^{k+1}(V_i) + d_i T^{2k+1}(V_i). \quad (39)$$

Here T is the shift operator, whose action is $T(V_i) = V_{i+1}$. If we take $\{V_i\}$ to be a periodic bi-infinite sequence of vectors corresponding to our polygon P , then D maps $\{V_i\}$ to the 0-sequence.

Next, we write $D = D_+ + D_-$ where coordinate-wise

$$D_+(V_i) = a_i V_i + c_i T^{k+1}(V_i), \quad D_-(V_i) = b_i T^k(V_i) + d_i T^{2k+1}(V_i). \quad (40)$$

The pair (D_+, D_-) is associated to the polygon P .

Let \tilde{D} and $(\tilde{D}_+, \tilde{D}_-)$ be the corresponding operators associated to \tilde{P} . One of the main results of [6] is that the various choices can be made so that

$$\tilde{D}_+ D_- = \tilde{D}_- D_+. \quad (41)$$

This is called *refactorization*. Equating the lowest (respectively highest) terms of the relation in Equation 41 gives us the identity $\tilde{a}_i b_i = \tilde{b}_i a_{i+k}$ (respectively $\tilde{c}_i d_{i+k+1} = \tilde{d}_i c_{i+2k+1}$.) These relations hold for all i and together imply Equation 38.

9.2 Extensions of Glick's Formula

Let me first review Glick's formula for Δ_1 , the pentagram map. Let P be a convex n -gon. Let (x^*, y^*) denote the accumulation point of the forward iterates of P under Δ_1 . Let $\hat{P}_\infty = (x^*, y^*, 1)$ be the collapse point. In somewhat different notation, Glick introduces the operator

$$T_P = nI_3 - G_P, \quad G_P(v) = \sum_{i=1}^n \frac{|P_{i-1}, v, P_{i+1}|}{|P_{i-1}, P_i, P_{i+1}|} P_i. \quad (42)$$

Here $|a, b, c|$ denotes the determinant of the matrix with rows a, b, c and I_3 is the 3×3 identity matrix. It turns out T_P is a Δ_1 -invariant operator, in the sense that $T_{\Delta_0(P)} = T_P$. Moreover P_∞ is an eigenvector of T_P . This is Glick's formula for \hat{P}_∞ . Actually, one can say more simply that G_P is a Δ_0 -invariant operator and that \hat{P}_∞ is an eigenvector of G_P . The more complicated expression $nI_3 - G_P$ is easier to work with geometrically.

Define $G_{P,a,b}$ by the formula

$$G_{P,a,b}(v) = \sum_{i=1}^n \frac{|P_{i-a}, v, P_{i+b}|}{|P_{i-a}, P_i, P_{i+b}|} P_i. \quad (43)$$

Let $\hat{P}_{\infty,k}$ be the limit point of the forward iterates of P under Δ_k . It seems that when $k \geq 1$ and $n = 3k + 1$ the operator $G_{P,k,k}$ is Δ_k invariant and has $\hat{P}_{\infty,k}$ for an eigenvector. It seems that when $k \geq 1$ and $n = 3k + 2$ the operator $G_{P,k+1,k+1}$ is Δ_k invariant and has $\hat{P}_{\infty,k}$ for an eigenvector. I was not able to find any similar formulas when $n > 3k + 2$.

Anton Izosimov kindly explained the following lemma.

Lemma 9.3 *These operators are Δ_k -invariant.*

Proof: These operators are Glick’s operator in disguise. When $n = 3k + 1$ we can relabel our n -gons in a way that converts Δ_k to the pentagram map. The corresponding space of birds $B_{n,k}$ corresponds to some strange set of “re-labeled k -birds”. This relabeling converts $G_{P,k,k}$ respectively to Glick’s original operator. This proves the invariance of $G_{P,k,k}$ under Δ_k when $n = 3k + 1$. A similar thing works for $n = 3k + 2$, but this time the relabeling converts Δ_k to the inverse of the pentagram map. ♠

This result does not explain why the collapse point \widehat{P}_∞ is an eigenvector, but there is plenty of food for thought here. Glick’s formula is an analytic expression for the collapse point, and perhaps what is going on here is some kind of analytic continuation. I wonder if this means that the collapse point exists for *all* starting points of the pentagram map. Even if the iterations go completely crazy under the map, perhaps they still collapse to the point predicted by Glick’s operator. The idea of a completely general collapse point has always seemed absurd to me, but maybe it is not. Nobody knows. Even though the algebra of the pentagram is quite well understood, the geometry is not.

9.3 Star Relabelings

Let us further take up the theme in the proof of Lemma 9.3. Given an n -gon P and some integer r relatively prime to n , we define a new n -gon P^{*r} by the formula

$$P_j^{*r} = P_{rj}. \tag{44}$$

Figure 1.5 shows the $P^{*(-3)}$ when P is the regular 10-gon.

As we have already mentioned, the action of Δ_1 on the $P^{*(-k)}$ is the same as the action of Δ_k on P when $n = 3k + 1$. So, when $n = 3k + 1$, the pentagram map has another nice invariant set (apart from the set of convex n -gons), namely

$$B_{k,n}^{*(-k)} = \{P^{*(-k)} \mid P \in B_{k,n}\}.$$

The action of the pentagram map on this set is geometrically nice. If we suitably star-relabel, we get star-shaped (and hence embedded) polygons. A similar thing works when $n = 3k + 2$.

References

- [1] Q. Aboud and A. Izosimov, *The Limit Point of the Pentagon Map and Infinitesimal Monodromy*, I.M.R.N., Vol 7, (2022)
- [2] M. Gekhtman, M. Shapiro, S. Tabachnikov, A. Vainshtein, *Higher pentagram maps, weighted directed networks, and cluster dynamics*, Electron. Res. Announc. Math. Sci. **19** (2012), pp. 1–17
- [3] M. Glick, *The Limit Point of the Pentagon Map*, International Mathematics Research Notices 9 (2020) pp. 2818–2831
- [4] M. Glick, *The pentagram map and Y-patterns*, Adv. Math. **227**, 2012, pp. 1019–1045.
- [5] A. B. Goncharov and R. Kenyon, *Dimers and Cluster Integrable Systems*, (2011) ArXiv 1107.5588
- [6] A. Izosimov and B. Khesin, *Long Diagonal Pentagon Maps*, Bulletin of the L.M.S., vol. 55, no. 3, (2023) pp. 1-15
- [7] A. Izosimov, *The pentagram map, Poncelet polygons, and commuting difference operators*, arXiv 1906.10749 (2019)
- [8] A. Izosimov, *Pentagram maps and refactorization in Poisson-Lie groups*, Advances in Mathematics, vol. 404 (2022)
- [9] A. Izosimov, *Intersecting the Sides of the Polygon*, Proc. A.M.S. **150** (2022) 639-649.
- [10] B. Khesin, F. Soloviev *Integrability of higher pentagram maps*, Mathem. Annalen. Vol. 357 no. 3 (2013) pp. 1005–1047
- [11] B. Khesin, F. Soloviev *The geometry of dented pentagram maps*, J. European Math. Soc. Vol 18 (2016) pp. 147 – 179
- [12] G. Mari Beffa, *On Generalizations of the Pentagon Map: Discretizations of AGD Flows*, Journal of Nonlinear Science, Vol 23, Issue 2 (2013) pp. 304–334
- [13] G. Mari Beffa, *On integrable generalizations of the pentagram map* Int. Math. Res. Notices (2015) (12) pp. 3669-3693

- [14] R. Felipe and G. Mari-Beffa, *The pentagram map on Grassmannians*, arXiv 1507.04765 (2015)
- [15] Th. Motzkin, *The pentagon in the projective plane, with a comment on Napier's rule*, Bull. Amer. Math. Soc. **52**, 1945, pp. 985–989.
- [16] N. Ovenhouse, *The Non-Commutative Integrability of the Grassman Pentagram Map*, arXiv 1810.11742 (2019)
- [17] V. Ovsienko, R. E. Schwartz, S. Tabachnikov, *The pentagram map: A discrete integrable system*, Comm. Math. Phys. **299**, 2010, pp. 409–446.
- [18] V. Ovsienko, R. E. Schwartz, S. Tabachnikov, *Liouville-Arnold integrability of the pentagram map on closed polygons*, Duke Math. J. Vol 162 No. 12 (2012) pp. 2149–2196
- [19] R. E. Schwartz, *The pentagram map*, Exper. Math. **1**, 1992, pp. 71–81.
- [20] R. E. Schwartz, *Discrete monodromy, pentagrams, and the method of condensation*, J. of Fixed Point Theory and Appl. **3**, 2008, pp. 379–409.
- [21] R. E. Schwartz, *A Textbook Case of Pentagram Rigidity*, arXiv 2108-07604 preprint (2021)[
- [22] R. E. Schwartz, *Pentagram Rigidity for Centrally Symmetric Octagons*, arXiv 2111:08358 preprint (2022)
- [23] F. Soloviev *Integrability of the Pentagram Map*, Duke Math J. Vol 162. No. 15, (2012) pp. 2815 – 2853
- [24] M. Weinreich, *The Algebraic Dynamics of the Pentagram Map*, (2021) arXiv: 2104.06211