On Spaces of Inscribed Triangles

Richard Evan Schwartz *

August 20, 2019

Abstract

Meyerson's Theorem says that all but at most 2 points of any Jordan loop J are vertices of equilateral triangles inscribed in J. We first prove an enhanced version of Meyerson's Theorem which has topological information in it. We then show that for each such J there is an uncountable set G(J) of triangle shapes such that the same result holds for any shape in G(J).

1 Introduction

A Jordan loop is the image of a circle under a continuous injective map into the plane. A polygon is *inscribed* in a Jordan loop if the vertices of the polygon lie in the loop. There has been a lot of interest over the years in the problem of inscribing polygons, especially triangles and quadrilaterals, in Jordan loops. See, for instance, [AA], [ACFSST], [H], [Mak1], [Mak2], [Ma2], [M], [N], [NW], [S], [Shn], [Ta]. Much of this interest stems from the famous Toeplitz Square Peg Problem, which asks if every Jordan loop has an inscribed square. See [Ma1] or [P] for a detailed discussion.

The purpose of this paper is to prove some results about inscribing triangles in Jordan loops. Half our motivation comes from M. Meyerson's theorem $([\mathbf{M}], 1980)$:

Theorem 1.1 (Meyerson) Let J be any Jordan loop. Then all but at most 2 points of J are vertices of inscribed equilateral triangles.

^{*} Supported by N.S.F. Research Grant DMS-1204471

This result is sharp in the sense that there are examples (e.g. an obtuse isosceles triangle) which have two points that are not vertices of any inscribed equilateral triangle.

One might wonder about other triangle shapes. Similarity classes of triangles are parametrized by the equilateral triangle \mathcal{T} consisting of positive triples $\lambda = (\theta_1, \theta_2, \theta_3)$ which sum to π . The center $\lambda_{\Delta} = (\pi/3, \pi/3, \pi/3)$ corresponds to the equilateral triangle shape. We will call a triangle with angles λ a λ -triangle. Meyerson's Theorem is not known for any $\lambda \neq \lambda_{\Delta}$. One step in this direction is due to M. Neilson ([**N**, 1991): For any Jordan loop J and any $\lambda \in \mathcal{T}$, there is a dense set of points of J which are vertices of inscribed λ -triangles.

The other half of our motivation comes from a variant of Meyerson's Theorem we proved in $[\mathbf{S}]$.

Theorem 1.2 For any Jordan loop J there is a connected subset β of gracefully inscribed rectangles such that all but at most 4 points of J are vertices of rectangles in β .

Here is the meaning of the terminology. We say that a polygon P is gracefully inscribed in J if the cyclic order on the vertices of P is the same whether it is computed with respect to the counter-clockwise cyclic order on P or the counter-clockwise cyclic order on J. Theorem 1.2 is sharp in the same way that Meyerson's Theorem is sharp, and it has additional topological information.

Our first result in this paper is a kind of marriage between Meyerson's Theorem and Theorem 1.2.

Theorem 1.3 Let J be any Jordan loop. There is a connected set β of gracefully inscribed equilateral triangles such that all but at most 2 points of J are vertices of triangles in β .

Our result of course implies Meyerson's Theorem, but the proof is very different from Meyerson's own short and beautiful argument. We get Theorem 1.3 by approximating arbitrary Jordan loops with sequences of polygons and then taking the limit of a structural result about triangles inscribed in polygons. Our structural result works for any shape parameter, and we are able to squeeze some additional information out of it. Here is a preliminary definition needed for our main result. **Definition:** Let $G(J) \subset \mathcal{T}$ denote the set of parameters λ with the following property. There is a connected set β_{λ} of λ -triangles gracefully inscribed in J such that all but at most 2 points of J are vertices of triangles in β_{λ} .

The operation of cyclically permuting the labels on a triangle corresponds to order 3 isometric rotation of the parameters space \mathcal{T} . We let ρ be this order 3 rotation. Here is our main result.

Theorem 1.4 Let J be an arbitrary Jordan loop. If Λ is a ρ -invariant, compact, connected subset of \mathcal{T} , then $\Lambda \cap G(J) \neq \emptyset$.

We point out that Theorem 1.4 does *not* imply that there is some $\lambda \neq \lambda_{\Delta}$ which lies in G(J) for all J. At first it seemed that we could use the structural result we describe below to show that $G(J) = \mathcal{T}$ but Theorem 1.4 is the best we can do without new ideas.

We mention two corollaries of Theorem 1.4. Applying Theorem 1.4 to the triangle in \mathcal{T} whose vertices are the 3 cyclic permutations of $(\theta, \theta, \pi - 2\theta)$, we get the following corollary.

Corollary 1.5 Let J be an arbitrary Jordan loop. For any angle $\theta \in (0, \pi/2)$ the set G(J) contains a parameter which specifies a triangle having some angle equal to θ . We can take θ to be the smallest angle when $\theta \leq \pi/3$ and the largest angle when $\theta \geq \pi/3$.

We also have the following immediate corollary.

Corollary 1.6 With respect to any Jordan loop, the conclusion of Meyerson's Theorem holds for an uncountable collection of triangle shapes.

Now we explain our structural result about polygons. Let J be a polygon. We call J obtuse if all the angles at the vertices of J exceed $\pi/2$ degrees. We say that J is angle-adapted to λ if the minimum angle between any two edges of J exceeds the maximum angle of a λ -triangle. (These concepts are related: If J is angle-adapted to a parameter which specifies an obtuse triangle, then J is automatically obtuse.) Let Ω_3 denote the set of ordered triples $(p_1, p_2, p_3) \in J^3$ consisting of distinct points. The space Ω_3 is a disjoint union of two open solid tori. Let $I(J, \lambda) \subset \Omega_3$ be the set of triangles of shape λ that are inscribed in J. **Theorem 1.7** For each obtuse polygon J there is a finite union L_J of lines with the following property. If $\lambda \in \mathcal{T} - L_J$ and J is angle-adapted to λ then $I(J, \lambda)$ is a finite union of disjoint polygons and exactly one component of $I(J, \lambda)$ is homologically nontrivial.

Remarks:

(1) The set L_J depends linearly with the angles of J. For fixed λ we have $\lambda \in \mathcal{T} - C_J$ when J is generically chosen. Thus, Theorem 1.7 says something about every triangle shape.

(2) The crucial part of Theorem 1.7 is the single essential component; the rest is pretty easy. I discovered Theorem 1.7 experimentally.

(3) The condition that J is angle-adapted to λ is necessary. When λ specifies an acute triangle, the obtuseness condition is not necessary but makes the proof easier.

(4) In [S, Theorem 5.2] we proved that the triangles corresponding to any essential component of $I(J, \gamma)$ are gracefully inscribed. We use this fact to get the graceful nature of the triangles in Theorem 1.5.

(5) Theorem 1.7 is subtle and here we explain why the obvious approach does not work. The obvious approach would be to start with some J_0 where the result is clear – e.g., a convex polygon – and then analyze the local changes to the space $I(J_t, \lambda)$, where J_t is a family of polygons that makes a generic interpolation between J_0 and some polygon J_1 of interest. A fairly easy argument shows that the homology class of represented by $I(J_t, \lambda)$ in $H_1(\Omega_3)$ does not change with t.



Figure 1.1: A locally allowed but globally bad topological transition.

The difficulty is that a priori one inessential component could self-intersect at a critical point and turn into two oppositely oriented essential components. This would not change the picture homologically but it would change the number of essential components. Figure 1.1, which takes place in a cylinder made from identifying the vertical boundaries of a strip, shows what we mean. It seems hard to rule out this transition with just local information. Here is an outline of the rest of the paper.

- Taking a Limit: In §2 we deduce Theorems 1.3 and 1.5 from Theorem 1.7 and a related technical result, Lemma 2.8. We treat the Theorem 1.3 separately, again, for the sake of exposition. The fact that there is just one essential component in Theorem 1.7 is crucial to the proofs. Following §2, the rest of the paper concerns the proof of Theorem 1.7.
- Fiber Product Lemma: In §3 we prove an elementary result about the fiber product of two maps from the circle to itself. We call this the Fiber Product Lemma. It seems to me that this result must be known, but I haven't found it in the literature. The single essential component statement in Theorem 1.7 ultimately comes from the single essential component statement in the Fiber Product Lemma.
- Component Lemma: In §4 we discuss what we mean by a *folding* map. This is an almost-everywhere non-singular piecewise linear map from the torus into the plane whose folding set i.e. the edges over which the map reverses orientation is a finite union of pairwise disjoint embedded polygonal loops. We bound on the number of essential components in the pre-image of a Jordan loop such a map can have in terms of the number of essential components of its folding set. Our main result is what we call the Component Lemma.
- Folding Lemma: In §5 we relate the problem of inscribing triangles to the idea of a folding map, then get a bound on the number of essential components of the associated folding set by invoking the Fiber Product Lemma from §3. Our main result is the Folding Lemma.
- Putting it Together: In §6 we combine the Component Lemma from §4 and the Folding Lemma from §5 to prove Theorem 1.7. At the end of §6 we prove Lemma 2.8, the technical result left over from §2.

I would like to thank Arseniy Akopyan, Ramin Naimi, Igor Rivin, and Sergei Tabachnikov for discussions and questions about topics related to this paper.

2 Taking a Limit

2.1 Limits and Convergence

In this chapter we deduce Theorem 1.5 from Theorem 1.7.

Let X be some metric space and let \mathcal{X} denote the set of compact subsets of X. The *Hausdorff distance* between two elements A, B of \mathcal{X} is the infimal $\epsilon > 0$ such that each point of A is within ϵ of B and vice versa. This notion of distance makes \mathcal{X} into a compact metric space.

Lemma 2.1 Suppose that a sequence $\{A_n\}$ of compact connected sets in X converges to a set A of X. Then A is connected.

Proof: If A is disconnected, there are disjoint open sets $U, V \subset C$ such that $A \subset U \cup V$, and $A \cap U$ and $A \cap V$ are both not empty. Since A is compact, there is some $\epsilon > 0$ such that every point in A is at least ϵ from every point in C - U - V. Since A_n is connected, this is only possible if C - U - V contains a point $x_n \in A_n$. But then x_n is at least ϵ from A, independent of the choice of n. Contradiction. See [**S**, Lemma 2.1] for a less terse proof.

Let J be a Jordan loop. We think of J as the image $\alpha(\mathbf{R}/\mathbf{Z})$ where $\alpha : \mathbf{R}/\mathbf{Z} \to \mathbf{C}$ is continuous and injective.

Lemma 2.2 The Jordan loop J can be approximated by a sequence of polygons that are obtuse and angle-adapted to shape λ . More precisely, there is a sequence $\{\alpha_n\}$ of maps from \mathbf{R}/\mathbf{Z} into \mathbf{C} which converges uniformly to α .

Proof: It is well known that any Jordan loop can be approximated by a sequence of polygons in the parametrized sense. See $[\mathbf{T}]$. To arrange the needed angle condition, we can cut off the corners of the polygons in our sequence, repeatedly, and increasingly near the vertices, until all the interior angles are large enough. At the same time we modify the maps to reflect this change. We then make small generic perturbations.

Let $J_n = \alpha_n(\mathbf{R}/\mathbf{Z})$. The sequence $\{J_n\}$ is a sequence of polygons converging to J in a parametrized sense.

2.2 The Equilateral Case

In this section we prove Theorem 1.3. Let $\lambda = \lambda_{\Delta}$, the equilateral triangle shape.

We take the polygonal approximation discussed above. Referring to Theorem 1.7, we perturb so that $\lambda \in \mathcal{T} - L_{J_n}$ for all n. Let $I_n = I(J_n, \lambda)$ denote the space of equilateral triangles inscribed in J_n . Let β_n be the unique essential component of I_n . Let ϵ_n denote the side length of the smallest triangle associated to β_n . If the sequence $\{\epsilon_n\}$ is uniformly bounded away from 0, then the Hausdorff limit $\beta = \lim \beta_n$ is a connected subset of $I(J, \lambda)$ such that every point of J is a vertex of some triangle in β . For later reference, we call this the *easy case*.

For the remaining case, we can pass to a subsequence so that $\epsilon_n \to 0$. Let Δ_j , for j = 1, 2, 3, be the vertices of an labeled equilateral triangle Δ . Let Δ_{12} denote the side of Δ which connects Δ_1 to Δ_2 . One of the two arcs of $J_n - (\Delta_1 \cup \Delta_2)$ does not contain Δ_3 . Let $\mu(\Delta)$ denote the measure of this arc, according to the parametrization α_n . What we mean is that there is some interval $I \subset \mathbf{R}/\mathbf{Z}$ such that $\alpha_n(I)$ connects Δ_1 to Δ_2 and avoids Δ_3 , and $\mu(\Delta)$ is the length of I.

Lemma 2.3 The range of μ on β_n converges to (0,1) as $n \to \infty$.

Proof: Here is the crucial part of the proof. The act of cyclically permuting the labels of a triangle acts on the space I_n . In particular, this action maps essential components to essential components. Since there is exactly one essential component, we see that β_n is invariant under cyclic relabeling.

The rest of the proof is continuity. Let τ_n be the smallest equilateral triangle associated to β_n . Let t_n be the 3-element subset of \mathbf{R}/\mathbf{Z} such that $\alpha_n(t_n) = \tau_n$. Given that the side length of τ_n tends to 0 and that the parametrizations converge, the diameter of the smallest interval of \mathbf{R}/\mathbf{Z} containing t_n tends to 0. But then $\mu(\tau_n)$ is either close to 0 or close to 1. In the former case, $\mu(\tau'_n)$ is close to 1 for a suitable relabeling τ'_n of τ_n . In the latter case, $\mu(\tau'_n)$ is close to 0 for a suitable cyclic relabelling τ'_n of τ_n . In either case, we find two triangles τ_n and τ'_n such that μ is close to 0 on one of them and close to 1 on the other. Since $\mu(\beta_n)$ is connected, this set achieves all values between these two extremes.

Let $I_n = [1/n, 1-1/n]$. We can pass to a subsequence so that $I_n \subset \mu(\beta_n)$. The next result is meant to hold for each fixed value of n. **Lemma 2.4** There are closed arcs $\beta_n(3) \subset ... \subset \beta_n(n) \subset \mathbf{R}/\mathbf{Z}$ such that μ maps $\beta_n(k)$ to I_k for k = 3, ..., n.

Proof: We start with a point $\tau \in \beta_n$ such that $\mu(\tau) = 1/2$ and then we let $\beta_n(k)$ be the smallest arc of β_n containing τ such that μ maps $\beta_n(k)$ to I_k . These arcs automatically have the desired containment properties.

Lemma 2.5 For fixed k there is some uniform $\epsilon_k > 0$ such that each point of $\beta_n(k)$ corresponds to an equilateral triangle of diameter greater than ϵ_k .

Proof: Let $\{\tau_n\}$ be a supposed equence of counter-examples. Let $\{t_n\}$ be the corresponding 3-element set of \mathbf{R}/\mathbf{Z} . The distance between the first 2 points of t_n tends to 0, forcing the side length of τ_n to 0. But then one of the sides of τ_n must subtend nearly the whole of J_n . This forces $\mu(\tau_n)$ either to 0 or to 1.

Now we are ready to take a limit.

Lemma 2.6 $I(J, \lambda)$ contains a connected set β such that $\mu(\beta) = (0, 1)$.

Proof: Passing to a subsequence we can assume that, for each k, the sequence $\{\beta_n(k)\}$ is a Cauchy sequence in the Hausdorff metric. Let $\beta(k)$ denote the limit. The set $\beta(k)$ is connected, by Lemma 2.1. Moreover, each point of $\beta(k)$ corresponds to an equilateral triangle inscribed in J. Finally, $\mu(\beta(k)) = I_k$ by continuity. We have $\beta(3) \subset \beta(4) \subset \beta(5)$ The nested union of connected sets is connected, so the union $\beta = \bigcup \beta(k)$ has the desired properties.

Suppose that there are three points of J which are not vertices of equilateral triangles associated to β . These three points divide J into three intervals. Since β is connected, the *j*th vertex of any triangle associated to β lies in the same interval, independent of the triangle. If each interval contains a vertex, we get a positive lower bound to the value of μ on β . If some interval is empty, we get a finite upper bound. Either way, we have a contradiction. Note that we are taking limits of gracefully inscribed triangles, so the triangles associated to β are graceful.

2.3 The General Case

In this section we will prove Theorem 1.4, which says that for any Jordan loop J and any ρ -invariant compact connected subset $\Lambda \subset \mathcal{T}$, the intersection $G(J) \cap \Lambda$ is nonempty. We begin with a lemma from point set topology.

Lemma 2.7 Let Λ be a planar set which has no isolated points. Let $\{L_t\}$ be a non-constant continuous family of lines. Some line L_s in the family has the property that $L_s \cap \Lambda$ is nowhere dense in Λ .

Proof: We will assume that this is false and derive a contradiction. We note first of all that a basis for the topology on Λ is given by the intersections of Λ with open disks in the plane. If our result is false, then for each t there is some open disk D_t in the plane such that $D_t \cap \Lambda$ is contained in, and dense in, $D_t \cap L_t$. We can adjust D_t so that the center of D_t lies in Λ . Since Λ has no isolated points, there are points of $D_t \cap \Lambda$ arbitrarily close to the center of D_t but unequal to it.

Suppose that s and t are two different parameters such that L_s and L_t are distinct lines. Then the center of D_s cannot lie in D_t . Otherwise, there is a point ζ of $\Lambda \cap D_t$ (either the center of D_s or a nearby point of Λ) which does not lie in L_t . But then there is a positive distance between all points of $L_t \cap D_t$ and ζ . This contradicts the fact that $L_t \cap \Lambda$ is contained in, and dense in, $L_t \cap \Lambda$. The same argument works with the roles of s and t reversed. Hence, the distance between the centers of D_s and D_t is at least max $(|D_s|, |D_s|)$, where $|\cdot|$ is the radius function. But then $\frac{1}{2}D_s$ and $\frac{1}{2}D_t$ are disjoint.

For each parameter t we have constructed an open disk, namely $\frac{1}{2}D_t$, having the property that two of these disks are disjoint if they correspond to parameters indexing different lines. Since there are uncountably many distinct lines in our family, we have just put an uncountable number of disjoint open disks in the plane. This is obviously impossible.

Now we state a technical result whose proof we defer until §6.4. Let J and L_J be as in Theorem 1.7 and let A be a circular arc in \mathcal{T} with endpoints in $\mathcal{T} - L_J$. Define

$$I(J,A) = \bigcup_{\lambda \in A} I(J,\lambda).$$
(1)

Lemma 2.8 Suppose that J is obtuse and angle-adapted to all parameters in A. Then the essential components β_0 and β_1 corresponding to the endpoints of A lie in the same path component B of I(J, A), unless perhaps A contains a certain finite union P_J of exceptional points.

Remark: Excluding the points in P_J is not really necessary but it makes the proof of Lemma 2.8 easier.

Proof of Theorem 1.4: We ignore the case when Λ is a single point because Theorem 1.3 covers this case. Otherwise, Λ is infinite and connected, and hence has no isolated points. This means that Λ satisfies the conclusion of Lemma 2.7.

Let $\{J_n\}$ be a sequence of polygons approximating J. We can arrange that J_n is obtuse and angle-adapted to every point in some tubular neighborhood $\langle \Lambda \rangle$ of Λ . As discussed at the end of §6.3 we can vary J_n continuously so that each line of L_{J_n} either instantly disappears or moves in a generic family. Hence, by Lemma 2.7, we can perturb by as small an amount as we like so that $\Lambda \cap L_{J_n}$ is nowhere dense in Λ . Let Λ^* be the subset of Λ consisting of points not in any L_{J_n} . Note that Λ^* is dense in Λ . Also, L_{J_n} is ρ -equivariant, so Λ^* is also ρ -invariant.

We remind the reader that $\alpha_n : \mathbf{R}/\mathbf{Z} \to J_n$ is the parametrization of J_n . Given a triangle τ inscribed in J_n , and k = 1, 2, 3, we let $\mu_k(\tau)$ denote the arc length of the interval $\gamma \subset \mathbf{R}/\mathbf{Z}$ such that $\alpha_n(\gamma)$ is the arc of J_n subtended by the kth side of τ . Given $\lambda \in \Lambda^*$, define $\overline{\mu}_k(\lambda) \in [0, 1]$, where

$$\overline{\mu}_k(\lambda) = \limsup_{n \to \infty} (\sup_{\tau \in \beta_n} \mu_k(\tau)).$$
(2)

Here β_n is the essential component of $I(J_n, \lambda)$. Let $S_k \subset \Lambda^*$ denote the set of parameters $\lambda \in \Lambda^*$ such that $\overline{\mu}_k(\lambda) = 1$. Geometrically this means that for all large *n* there are very small triangles of shape λ inscribed in J_n where the *k*th side subtends practically all of J_n .

Suppose first that Λ^* contains a point $\lambda \in \mathcal{T} - S_1 - S_2 - S_3$. Then there is a uniform positive lower bound to the diameter of any λ -triangle inscribed in J_n , independent of n. The same argument as in the easy case of Theorem 1.3 gives $\lambda \in G(J)$.

Suppose now that $\Lambda^* \subset S_1 \cup S_2 \cup S_3$. Note that $p \in S_k$ iff $\rho(p) \in S_{k+1}$. Here indices are taken mod 3. Since Λ^* is ρ -invariant and dense in the compact connected set Λ , we can find sequences $\{\lambda_{1,n}\} \subset S_1$ and $\{\lambda_{2,n}\} \subset S_2$ such that $\lambda_{1,n}$ and $\lambda_{2,n}$ both converge to the same point $\lambda_{\infty} \in \Lambda$. Passing to a subsequence, we can assume that there is some triangle $\tau_{k,n}$ in the essential component $\beta_{\lambda_{k,n}}$ of $I(J_n, \lambda_{k,n})$ such that $\mu_k(\tau_{k,n}) > 1 - 1/n$. Here k = 1, 2. Setting $\mu = \mu_2$ we have

$$\mu(\tau_{1,n}) < 1/n, \qquad \mu(\tau_{2,n}) > 1 - 1/n.$$
 (3)

We can connect $\lambda_{1,n}$ to $\lambda_{2,n}$ by a circular arc A_n of length less than (say) $2|\lambda_{1,n} - \lambda_{2,n}|$ which avoids the set P_{J_n} from Lemma 2.8 and stays in the tubular neighborhood $\langle \Lambda \rangle$ of Λ .

By Lemma 2.8, there is a path component $B_n \subset I(J, A_n)$ which contains $\beta_{\lambda_{k,n}}$, and hence $\tau_{k,n}$, for k = 1, 2. So, we can connect $\tau_{1,n}$ to $\tau_{2,n}$ by a path $\beta_n \subset B_n$. By construction $I_n = [1/n, 1 - 1/n] \subset \mu(\beta_n)$. We now have the same situation that Lemma 2.3 establishes in the equilateral case, except that the shapes in β_n vary in a small neighborhood of the two parameters $\lambda_{1,n}$ and $\lambda_{2,n}$. This change makes no difference to our limiting arguments, because these parameters converge to λ_{∞} . Hence $\lambda_{\infty} \in G(J)$.

3 The Fiber Product Lemma

We call the map $f : \mathbf{R}/\mathbf{Z} \to \mathbf{R}/\mathbf{Z}$ a nice map if f has degree 1, is piecewise linear, and is not constant on any interval. We call the points where f locally reverses direction the fold points. We call t = f(s) a fold value if s is a fold point for f. Given two nice maps $f_1, f_2 : \mathbf{R}/\mathbf{Z} \to \mathbf{R}/\mathbf{Z}$ we can form the fiber product

$$H(f_1, f_2) = \{ (s_1, s_2) \in \mathbf{T} | f_1(s_1) = f_2(s_2) \}.$$
 (4)

We call two nice maps f_1 and f_2 unrelated if they have no common fold values. For the sake of completeness, we prove the following result in this chapter.

Lemma 3.1 (Fiber Product) Suppose that f_1 and f_2 are unrelated nice maps. Then $H(f_2, f_2)$ is a polygonal 1-manifold which has exactly one connected component that is homologically nontrivial in \mathbf{T} . When suitably oriented, the one nontrivial component represents (1, 1) in homology $H_1(\mathbf{T})$.

We will break the proof into 4 steps. Let $H = H(f_1, f_2)$ be the fiber product of f_1 and f_2 .

Lemma 3.2 *H* is a polygonal 1-manifold.

Proof: Given two partitions $\{I_i\}$ and $\{J_j\}$ of \mathbf{R}/\mathbf{Z} into intervals, we can take the product and get a partition of \mathbf{T} into rectangles $\{R_{ij}\}$ with $R_{ij} = I_i \times J_j$. Since f_1 and f_2 are unrelated, we can choose these partitions so that the restriction of each function to each interval is linear and injective, and no vertex of an R_{ij} belongs to H. The locations of the fold points force us to choose certain breaks in the partitions, but otherwise we choose the breaks generically.

Let $H_{ij} = H \cap R_{ij}$. By construction H_{ij} is either the empty set or a line segment which connects the interior point of some edge of R_{ij} to the interior point of some other edge of R_{ij} . Consider the picture around an endpoint pof H_{ij} . Let R' be the rectangle adjacent to R_{ij} across the edge containing p. Since $H \cap R'$ is not the emptyset, $H \cap R'$ has the structure just mentioned. In particular, H_{ij} meets a unique line segment of H at p. This shows that H is a polygonal 1-manifold. \blacklozenge **Lemma 3.3** *H* has an orientation with the following properties:

- Whenever the generic vertical geodesic $x = x_0$ intersects H at a point (x_0, y) , the relevant segment points to the right if and only if $f'_2(y) > 0$.
- Whenever the generic horizontal geodesic $y = y_0$ intersects H at a point (x, y_0) , the relevant segment points to the top if and only if $f'_1(x) > 0$.

Proof: Here is the construction. If $f_1(I_j)$ and $f_2(J_j)$ are not disjoint, then they overlap in one of 4 possible ways. At the same time, there are 4 possible orientations for these segments, depending on the signs of the derivatives f'_1 and f'_2 . All in all, there are 16 different possibilities. For each of these possibilities, we choose an orientation for the corresponding segment of H, according to the scheme shown in Figure 3.1.



Figure 3.1: The orientation on the fiber product

A case-by-case check shows that this scheme defines a consistent orientation on H. Figure 3.2 shows how Cases 1,2 fit together and how Cases 1,3 fit together.



Figure 3.2: Adjacent pairs of segments and their tiles.

Now we check how the geodesic $x = x_0$. intersects one of our tiles. In all cases, the arrow in Figure 3.1 points to the right if and only if the lower of the two segments (corresponding to $f_2(J_j)$ points to the right. Similarly, we check how the geodesic $y = y_0$ intersects our tiles. In all cases, the arrow in Figure 3.1 points up if and only if the upper of the two segments (corresponding to $f_1(I_i)$) points to the right.

From now on we equip H with the orientation given above, and we call it the *natural orientation*. Since H is oriented, it makes sense to ask which homology class H represents in $H_1(\mathbf{T})$.

Lemma 3.4 *H* represents the element (1,1) in $H_1(\mathbf{T})$.

Proof: If suffices to show that the geodesics $x = x_0$ and $y = y_0$ each intersect H once, counting the orientations. Consider $x = x_0$. Each intersection point with this geodesic corresponds to a parameter value y where $f_2(y) = f(x_0)$. The orientation points to the right if and only if $f'_2(y) > 0$. But the number of times $f'_2(y) > 0$ is one more than the number of times $f'_2(y) > 0$ because f_2 has degree 1. In other words, $f_2(\mathbf{R}/\mathbf{Z})$ crosses a point righwards one more time than it crosses leftwards. This proves our claim for the geodesic $x = x_0$. A similar argument works for the geodesic $y = y_0$.

Now we know that H represents (1, 1) in $H_1(\mathbf{T})$. Two distinct and nontrivial homology classes in \mathbf{T} intersect unless they represent the same homology classes or their sum is 0 in homology. Since H is an embedded 1-manifold, all the homologically nontrivial components of H represent either (1, 1) or (-1, -1). Moreover, the number of (1, 1) representatives is one more than the number of (-1, -1) representatives. The last step finishes the proof. **Lemma 3.5** An arbitrary non-trivial component h of H represents (1, 1) in homology.

Proof: Again, we are equipping h with its natural orientation. We can find a piecewise linear map $a : \mathbf{R}/\mathbf{Z} \to \mathbf{T}$ such that $a = (a_1, a_2)$ parametrizes h, and each a_j is a degree 1 map. Define $b = f_j \circ a_1$. This map is independent of j and has degree 1.

The parametrization a gives h a second orientation of h which we call the *forced orientation*. The component h represents the element (1, 1) in $H_1(\mathbf{T})$ with respect to the forced orientation. So, to finish the proof, we need to show that the forced orientation and the natural orientation coincide.

Given $t \in \mathbf{R}/\mathbf{Z}$ we can compare the signs of $f'_2(a_2(t))$ and $a'_1(t)$. The former quantity determines the direction that h points across the vertical line through a(t). The latter quantity determines the direction that h points across the vertical line through a(t). The two orientations agree iff the two quantities have the same sign. By the Chain Rule, $f'_2(a_2(t))$ is positive if and only if $a'_2(t)$ and b'(t) have the same sign. Therefore the two orientations agree if there is any point t such that

$$a_1'(t)a_2'(t)b'(t) > 0 (5)$$

Note that $a_j(s) = a_j(t)$ implies that b(s) = b(t). This is because $b = f_j \circ a_j$.

We will suppose that Equation 5 fails for all t and we will derive a contradiction. We can find lifts $A_1, A_2, B : \mathbf{R} \to \mathbf{R}$ of a_1, a_2, b respectively. Each function F is such that F(x+1) = F(x)+1. The lifted functions also satisfy the same property as above: If $A_j(s) = A_j(t)$ then B(s) = B(t). Moreover $A'_j = a'_j$ and B' = b'. So, $A'_1(t)A'_2(t)B'(t) < 0$ whenever all these derivatives are defined. In particular, these derivatives cannot all be positive.

Say that a point $t \in \mathbf{R}$ is a *peak* if the function B(t) - t has a global maximum at t. A peak exists because the function B(t) - t is periodic. Let t_0 be a peak. By construction, $B(s) < B(t_0)$ for all $s < t_0$. For $\epsilon > 0$ sufficiently small, we have $B'(t_0 - \epsilon) \ge 1 > 0$. We pick ϵ so small that no derivative changes sign on $[t_0 - \epsilon, t_0]$. Since not all derivatives are positive, have $A'_j(t_0 - \epsilon) < 0$ for some j. By the Fundamental Theorem of Calculus, $A_j(t_0 - \epsilon) > A_j(t_0)$. Since $A_j(t_0 - 1) < A_j(t_0)$ there is some $s \in (t_0 - 1, t_0)$ such that $A_j(s) = A_j(t_0)$. But then $B(s) = B(t_0)$. This is a contradiction.

4 The Component Lemma

4.1 Main Result

Let T be the square torus. Suppose that $\Delta : T \to C$ is a piecewise linear map that is almost-everywhere non-singular. So, T has a triangulation and the restriction of Δ to the interior of each triangle is an affine isomorphism onto its image. We define the folding set B to be the set of edges e in the triangulation of T such that Δ is orientation reversing on one side of e and orientation preserving on the other. Informally, the map Δ locally folds the domain over the folding set. We call Δ a *folding map* if B is a finite union of pairwise disjoint embedded polygonal loops.

Example: Let Υ be a planar annulus bounded by 2 polygonal curves. Let Υ^* denote the piecewise linear torus obtained by doubling Υ – i.e., gluing 2 copies along their common boundary. There is a natural map $\pi : \Upsilon^* \to \Upsilon$ which just forgets the name of the copy. A folding map is given by $\Delta = \pi \circ f$, where $f: \mathbf{T} \to \Upsilon^*$ is some piecewise linear homeomorphism.

Let Δ be a folding map, as above. We say that a primitive homology class $\Theta \in H_1(\mathbf{T})$ is a *characteristic* of Δ if, when suitably oriented, each essential component of B represents Θ . Note that Θ' is a characteristic Δ if and only if $\Theta' = \pm \Theta$. In our applications, $\Theta = \pm (1, 1)$ in the standard homology basis.

Suppose that Δ has a nonzero characteristic. Of the two nonzero characteristics, we make a choice of one of them, and call it Θ . We say that a polygon $K \subset \mathbf{C}$ is *adapted* to Δ if the following is true.

- $A = \Delta^{-1}(K)$ is a finite union of pairwise disjoint embedded polygonal loops, all transverse to B.
- When suitably oriented, each essential component of A represents Θ.
 We call this the *positive orientation*.
- The restriction of Δ to each essential component of A, given the positive orientation, is a degree 1 map onto K.

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

Lemma 4.1 (Component) The number of essential A-components is at most the number of essential B-components.

4.2 A Special Case

Here we prove the Component Lemma under the assumption that the essential components of A are disjoint from B. In this case, the restriction of $\Delta : a \to K$ is a homeomorphism for each essential component a of A. The reason is that Δ has degree 1 and makes no folds in a neighborhood of a.

We will assume that the conclusion of the Component Lemma is false and derive a contradiction. The essential components of A divide T into a number of annuli, one of which is shown at left in Figure 4.1. The B curves are drawn in black and the A curves are drawn in various shades of grey. Since the annuli boundaries are disjoint from B, each B component must be contained in some annulus. Since there are more annuli than B components, one of the annuli does not contain an essential B component. The annulus X shown in Figure 4.1 has this feature. We have colored T white or grey according as the map Δ is orientation preserving or reversing.



Figure 4.1: One of the annuli.

Given that X contains no essential component of B, one can connect the two boundary components of X by a polygonal path β that avoids the B curves and therefore remains entirely in a monochrome (white or grey) region. We have indicated β with a dotted path that remains in the white region. Let $\gamma = \Delta(\beta)$. The curve γ starts out on K, moves away, crosses K an even number of times, then returns to K. These crossings come from the places where β crosses inessential components of $\Delta^{-1}(K)$. The right side of Figure 4.1 shows the situation. If β stays in the white (orientation preserving) region, then $\Delta(\beta)$ must start by moving into the bounded region of C - K, as shown in Figure 4.2. In other words, we can say that $\Delta(\beta)$ starts out by moving *inside*. Counting the crossings, we see that γ must reach the second endpoint on K from the inside. That is, all points of γ sufficiently near the second endpoint must also be in the bounded component of C - K. But this is a contradiction: The map Δ would have to be orientation reversing at the second endpoint of β . A similar argument with *outside* replacing *inside*.

4.3 The General Case

Now we consider the general case. The proof goes by induction on the number of intersections between the essential components of A and the components of B. The special case above is the base case for the induction. For the inductive step, the idea is to find an "innermost intersection" and modify Δ near the relevant segment of A so as to eliminate 2 intersection points and retain the number of essential A and B components and the general nature of the map. We will do the construction in small steps.

Notation and Terminology: Let a be an essential component of A. The restriction $\Delta : a \to K$ is a degree 1 map, but in general it need not be monotone. The map $\Delta|_a$ is monotone on the intervals bounded by the finitely many points of $a \cap B$. Because a lies in the same homology class as each essential component of B, there are an even number a_1, \ldots, a_n of such intervals. There are numbers w_1, \ldots, w_n such that the restriction $\Delta|a_i$ is monotone, and winds $4w_i$ units around K, which we normalize below to have length 4. The number w_i is positive iff $\Delta(a_i)$ winds counter-clockwise around K. The only constraint we have is that $w_1 + \ldots + w_n = 1$ and the signs alternate We call these numbers the weights. There are n/2 positive weights and n/2 negative weights.

Cleaning up the Map. We pre-compose Δ with a piecewise linear homeomorphism so that the essential components of A are geodesics, and we postcompose by a piecewise linear homeomorphism so that K is the unit square centered at the origin. The transversality condition guarantees that we can make small modifications to Δ near a so that it has a product structure in the sense that

$$(\Delta \circ t_{\epsilon})|_{a} = (d_{1+\epsilon} \circ \Delta)|_{a}.$$
(6)

Here t_{ϵ} is translation by ϵ normal to a, and $d_{1+\epsilon}$ is dilation by $1 + \epsilon$ about the origin. The cleanest way to do this is to slice T open along a and graft in the desired product map.

A Surgery Operation: We now explain a modification of Δ where we change Δ inside σ . Let σ_+ and σ_- be the two components of $\sigma - a$. Let Δ_{\pm} be the restriction of Δ to $\partial \sigma_{\pm}$. We first define a new piecewise linear and degree 1 map $f : a \to K$. We then take a new map Δ_f in σ_{\pm} so that it maps $t_{\epsilon}(a)$ to $d_{1+\epsilon}(K)$ and interpolates between f and Δ_{\pm} . That is, we use σ to implement homotopies between f and Δ_{\pm} . Finally, we set $\Delta_f = \Delta$ outside σ . By construction, Δ_f is a fold map and $A_f = A$. Here $A_f = \Delta_f^{-1}(K)$. The folding set B_f agrees with B outside σ but inside σ the two sets can differ. We call f the surgery core. After we do a surgery we modify the map again, as above, to retain the product structure near a.

Changing the Weights: If necessary, we first choose the surgery core f so that it has the following properties.

- $f|_a$ and $\Delta|_a$ have the same branch points.
- The weights $w'_1, ..., w'_n$ for f satisfy $sign(w'_j) = sign(w_j)$ for all j.
- If $w'_j > 0$ then $|\frac{2}{n} w'_j| < \frac{1}{100n}$.
- If $w'_j < 0$ then $|w'_j| < \frac{1}{100n}$.
- When $w'_j < 0$, the image $f(a_j)$ is at least $\frac{1}{2n}$ from any vertex of K.

The choice of $\frac{1}{100n}$ is an arbitrary but convenient cutoff. We do the homotopies so to that the resulting weights linearly interpolate between the new ones and the old ones and the branch points do not move. This surgery does not change A or B.

Innermost Digon: Since all the essential loops are in the same homology class, we can find a digon D whose interior is disjoint from $A \cup B$ and whose boundary has one arc in an essential component a of A and the other edge in B, as shown in Figure 4.2.



Figure 4.2: An innermost digon and the region of modification.

Let a_1 be the interval of a which is the boundary component of D. Once we do this, we can find a rectangular neighborhood R of a_1 , as shown on the right side of Figure 4.2, so that

- $\Delta(a_1) \subset \Delta(R \cap a)$, and $\Delta(R \cap a)$ lies in a single side of K.
- The top and bottom of R lie in the top and bottom of the kind of product-behavior strip we have been considering.

Note that $\Delta \cap B$ is just a union of 2 segments that cut across R.

Eliminating the Fold: We choose a surgery core f, so that $f = \Delta$ on a - R and f is monotone in $a \cap R$. We do the surgery using homotopies that are constant outside R. These homotopies just undo the little fold in the most straightforward way. The new map has the same number of essential A components and essential B components but overall 2 fewer intersections between essential A components and essential B components. Our result now follows from induction on the number of these intersection points.



5 The Folding Lemma

5.1 Main Result

Recall that a polygon is *obtuse* if every angle between consecutive edges is obtuse. Suppose, for j = 1, 2, that $\alpha_j : \mathbf{R}/\mathbf{Z} \to J_j \subset \mathbf{C}$ is a piecewise linear parametrization of an embedded obtuse polygonal loop J_j . Recall that $\mathbf{T} = (\mathbf{R}/\mathbf{Z})^2$. We introduce the *difference map* $\Delta : \mathbf{T} \to \mathbf{C}$. The map is

$$\Delta(s_1, s_2) = \alpha_1(s_1) + \alpha_2(s_2).$$
(7)

Lemma 5.1 (Folding) Suppose that both J_1 and J_2 are obtuse and no side of J_1 is parallel to a side of J_2 . Then Δ is a folding map and the fold set Bhas precisely 2 essential components. When suitably oriented, these essential components each represent (1, 1) in $H_1(\mathbf{T})$.

Remark: Since we can replace J_2 by the reflected curve $-J_2$, the truth of the Folding Lemma for the map Δ implies the truth of the Folding Lemma for the map Δ' where $\Delta'(s_1, s_2) = s_1 - s_2$. We are really interested in Δ' but we make the arguments for Δ because the picture is easier to see.

5.2 The Folding Map

Lemma 5.2 The map Δ is a folding map.

Proof: For j = 1, 2 let e_j be an edge of J_j . Since e_1 and e_2 are not parallel. The restriction $\Delta|_{e_1 \times e_2}$ agrees with an invertible affine map onto a parallelogram. So, T has a tiling by rectangles such that the restriction of Δ to the interior of each rectangle is locally affine and invertible. Hence Δ is a piecewise linear map and the fold set B is a union of horizontal and vertical edges of the rectangle tiling. It remains to show that B is a finite union of embedded polygonal loops. We just need to show that each vertex of the rectangle tiling is incident to either 0 or 2 edges in the fold set.

The set $\Delta(J_1^* \times J_2^*)$ is a union of 4 parallelograms, with the property that every two consecutive ones share an edge. For this reason, there are an even number edges in *B* meeting at a point. We just have to rule out the case of 4 such edges. If 4 edges of *B* come together at a point then all 4 edges of $J_1^* \cup J_2^*$ lie in the same half-plane, and the edges of J_1^* do not interlace with the edges of J_2^* . This is impossible with the obtuse angle condition.

5.3 Another View of the Folding Set

It remains to analyze the essential components of the fold set B. If we make piecewise linear changes to α_1 and α_2 we do not change the truth of the result. So, we take α_1 and α_2 to be constant speed parametrizations. We now think about the folding set in a different way. Let J be either J_1 or J_2 . We orient J counterclockwise. Say that a line ℓ is *tangent* to J if ℓ intersects the solid region bounded by J in either a single point or in an edge of J. So, ℓ is either the extension of an edge of J or else a line through a vertex of J. In the former case, we say that ℓ is tangent to J at each of the points on the relevant edge. In either case, the counterclockwise orientation on J induces an orientation on ℓ .

Lemma 5.3 The set B is the set of points (p_1, p_2) such that there are parallel lines ℓ_1 and ℓ_2 respectively tangent to J_1 and J_2 at $\alpha_1(p_1)$ and $\alpha_2(p_2)$.

Proof: Let B' be the set described above in terms of tangent lines. Both B and B' are defined by the property that $\det d\Delta$, where defined, takes both signs at points arbitrarily close to a point if and only if it belongs to the set in question. Hence B = B'.

Keeping the notation from Lemma 5.3, we call a point $(p_1, p_2) \in B$ positive if ℓ_1 and ℓ_2 point in the same direction and negative if ℓ_1 and ℓ_2 point in opposite directions. This partitions B into two sets B_+ and B_- , and each of the sets is a finite union of components of B. If we were to reverse the orientation of one of the curves, and keep the same definitions, we would interchange B_+ and B_- . We will show that B_+ has a single essential component which, when suitably oriented, represents (1, 1) in $H_1(\mathbf{T})$. The same argument works for B_- .

The Approximation Idea: We would like to apply our fiber product result. Morally speaking, Lemma 5.3 reveals the set B_+ to be the fiber product of the unit tangent maps " f_1 " and " f_2 " associated to α_1 and α_2 . Here the map " f_j " applied to a point s is "the unit tangent vector" at $\alpha_j(s)$. This notion does not quite make sense when we work with polygons, and that is why we put the things in quotes. To make the idea work, we need to replace our polygons with Jordan loops that have a well-defined (and easy to understand) unit tangent map.

5.4 Piecewise Circular Approximation

We say that a *piecewise circular loop* is a Jordan loop that is made from a finite union of circular arcs. We insist that the tangent lines match when two different circular arcs meet. Thus, a piecewise circular loop is a C^1 curve. Any constant speed parameterization of a PCL is C^1 .

Lemma 5.4 J_1 and J_2 can each be approximated by a sequence of PCLs.

Proof: Let J be one of J_1 or J_2 . We replace each vertex v of J by a small circular arc V of curvature n which has endpoints on the edges of J incident to v and makes an angle 1/n with these edges at the endpoints. Now let e be an edge of J. There are two consecutive circular arcs V_1 and V_2 which have their endpoints on J. We connect these endpoints either by a single circular arc of very small curvature or a C^1 union of 2 circular arcs having very small curvature. (The choice depends on whether the tangent vectors to V_1 and V_2 lie on the same side of e or on opposite sides.) In either case, these interpolating arcs are replacements for the edges of J. We adjust so that the resulting union J_n is the desired approximation.

For j = 1, 2, let $\{J_{j,n}\}$ be the sequence constructed in the previous lemma. We fix constant speed parametrizations $\alpha_{j,n} : \mathbf{R}/\mathbf{Z} \to J_{j,n}$ which converge uniformly to α_j . The unit tangent map $f_{j,n} : \mathbf{R}/\mathbf{Z} \to S^1$ is defined, continuous, and piecewise linear once we identify the range with \mathbf{R}/\mathbf{Z} . The maps $f_{j,n}$ are nice maps, by construction. By perturbing if necessary, we can arrange that $f_{1,n}$ and $f_{2,n}$ are unrelated. Define the fiber product

$$B_{+,n} = H(f_{1,n}, f_{2,n}).$$
(8)

Lemma 5.5 On a subsequence, the set $B_{+,n}$ converges in the Hausdorff topology to B_+ .

Proof: Since $\{B_{+,n}\}$ is a sequence of compact subsets of the compact metric space T, we can pass to a convergent subsequence. The limit $B' \subset T$ is again a compact subset of T. We want to show that $B' = B_+$.

A point $(p_1, p_2) \in B_+$ has the property that some tangent line to J_1 at $\alpha_1(p_1)$ is parallel to, and points in the same direction as, some tangent line to J_2 at $\alpha_2(p_2)$. But any point in B' also has this property by continuity. Hence $B' \subset B_+$.

Consider a point $(p_1, p_2) \in B_+$. We will consider the case when $\alpha_1(p_1)$ is a vertex of J_1 . The other case, when $\alpha_2(p_2)$ is a vertex of J_2 , has the same treatment. Let ℓ_1 and ℓ_2 be the parallel tangent lines at $\alpha_1(p_1)$ and $\alpha_2(p_2)$ respectively. The points $\alpha_{1,n}(p_1)$ and $\alpha_{2,n}(p_2)$ respectively are very close to $\alpha_1(p_1)$ and $\alpha_2(p_2)$ when n is large. By construction, the corresponding tangent lines $\ell_{1,n}$ and $\ell_{2,n}$ are also close to ℓ_1 and ℓ_2 respectively. These lines may not be pointing precisely in the same direction, so it may not be true that $(p_1, p_2) \in B_{+,n}$. Note, however, that ℓ_1 is tangent to neither of the edges incident to $\alpha_1(p_1)$. Given the nature of our construction, we can move p_1 very slightly to a new point p'_1 such that the corresponding tangent line $\ell'_{n,1}$ at $\alpha_{n,1}(p_1)$ does point in the same direction as $\ell_{2,n}$. In other words, the point $(p'_1, p_2) \in B_{+,n}$ is very close to (p_1, p_2) . Hence $B_+ \subset B'$.

We pass to a subsequence ¹ so that $B_{+,n} \to B_+$.

Lemma 5.6 (Short Arcs) If $\{x_n\}$ and $\{y_n\}$ are two sequences of points in $B_{+,n}$ that converge to the same point of B_+ , then some arc in $B_{+,n}$ of length L_n connects x_n to y_n , and $L_n \to 0$.

Proof: Let $(p_1, p_2) \in B_+$ be the limit point. There are 3 cases, depending on whether $\alpha_1(p_1)$ is a vertex of J_1 or $\alpha_2(p_2)$ is a vertex of J_2 , or both. We will treat the case when $\alpha_2(p_1)$ is a vertex of J_2 and $\alpha_1(p_1)$ is not a vertex of J_1 . The other cases have similar treatments.

Let us write $x_n = (p_{1,n}(0), p_{2,n}(0))$ and $y_n = (p_{1,n}(1), p_{2,n}(1))$. Let $p_{1,n}(t)$ be the path which interpolates linearly between $p_{1,n}(0)$ and $p_{1,n}(1)$. The path $t \to p_{1,n}(t)$ is very short, and the corresponding tangent line $\ell_{1,n}(t)$ also changes very little. At the same time, when we move $p_{2,n}(0)$ slightly the tangent corresponding line $\ell_{2,n}$ moves quite a bit, given the fact that this line is tangent to a circular arc of curvature n. Moreover, the dependence is monotonic. Hence, there is a unique point $p_{1,n}(t)$ very near both $p_{1,n}(0)$ and $p_{1,n}(1)$, such that $\ell_{n,1}(t)$ points in the same direction as $\ell_{2,n}(t)$. Given the uniqueness, the path $t \to p_{1,n}(t)$ is a continuous path which interpolates between $p_{1,n}(0)$ and $p_{1,n}(1)$. We have constructed a short path in $B_{n,+}$ connecting our two points x_n and y_n .

¹In the previous result we did not really need to pass to a subsequence to get convergence, but we don't want to take the trouble to show that the original sequence converges.

5.5 Applying the Fiber Product Lemma

Now we finish the proof of the Folding Lemma.

Lemma 5.7 B_+ has exactly one essential component, and when this component is suitably oriented it represents (1,1) in $H^1(\mathbf{T})$.



Figure 5.1: Some bad Hausdorff convergence

Proof: The Short Arc Lemma rules out the kinds of Hausdorff convergence depicted in Figure 5.1. The pictures are meant to take place inside a cylinder – the vertical sides in the picture are identified. On the left, a sequence of inessential loops converges to a proper arc of some loop. On the right, a sequence of inessential loops converges to an essential loop. Both kinds of convergence involve a kind of folding which brings together points that are not joined by short arcs.

By the Fiber Product Lemma, $B_{+,n}$ has one essential component β_n . The sequence $\{\beta_n\}$ converges on a subsequence to a subset of some component β of B_+ . The component β must be essential, because otherwise for large n the essential loop β_n would fail to intersect some representative of (say) (1,-1) in $H_1(\mathbf{T})$ that avoids β . Since β_n represents (1,1) in $H_1(\mathbf{T})$ for all n, the component β , suitably oriented, must also represent (1,1).

Let β' be any other component of B_+ . We want to show that β' is inessential. Since there is a uniform upper bound on the number of components of $B_{+,n}$ (in terms of the number of sides of the polygons) and since all of β' is contained in the Hausdorff limit of $B_{+,n}$, there is some sequence $\{\beta'_n\}$ of components converges to an uncountable closed subset β'' of β' . By the Fiber Product Lemma, β'_n is inessential because $\beta'_n \neq \beta_n$. Either $\beta'' = \beta'$ or β'' is a proper arc. In the latter case, we have the first kind of convergence depicted in Figure 5.1. This is impossible, so $\beta'' = \beta'$. If β' is essential, we have the convergence shown on the right side of Figure 5.1. Hence β' is inessential.

6 Putting it Together

The main goal of this chapter is to prove Theorem 1.7. After we finish this, we will prove Lemma 2.8.

6.1 Manifold Structure

Say that a λ -triangle is a triangle of shape λ . Each λ triangle (p_1, p_2, p_3) determines 3 similarities S_1, S_2, S_3 , where

$$S_j(p_j) = p_j, \qquad S_j(p_{j+1}) = p_{j+2}.$$
 (9)

Here the indices are taken mod 3. The linear part of these similarities does not depend on the representative triangle.

Let $E = (E_1, E_2, E_3)$ be a triple of segments, not necessarily distinct, and let $L = (L_1, L_2, L_3)$ be the triple of lines extending them. We say that Lis *unrelated* to λ if $S_j(L_{j+1}) \cap L_{j+2}$ is a single point for at least one index j. This condition only depends on λ and not on the chosen representative triangle. If we fix L, then L is unrelated to all but at most one $\lambda \in \mathbf{C} - \mathbf{R}$.

We say that a λ -triangle is *inscribed* in E (respectively L) if the kth vertex of the triangle lies in L_k (respectively E_k) for k = 1, 2, 3. We allow the degenerate possibility that the the triangle is just 3 identical points. Let $I(L, \lambda)$ denote the space of triangles of shape λ inscribed in L. Likewise define I(L, E). We say that a *brick* is the product of 3 segments.

Lemma 6.1 If L is unrelated to λ then $I(L, \lambda)$ is a straight line. Hence $I(E, \lambda)$ is the intersection of a straight line and the brick $E_1 \times E_2 \times E_3$.

Proof: Given a λ -triangle with vertices (p_1, p_2, p_3) , let S_1 be the similarity such that $S_1(p_1) = p_1$ and $S_1(p_2) = p_3$. Note that S_1 only depends on the point p_1 and the shape λ . We cyclically relabel so $S_1(L_2) \cap L_3$ is a single point. If we place $p_1 \in L_1$ then there are unique points $p_2 \in L_2$ and $p_2 \in L_3$ such that p_1, p_2, p_3 is inscribed in L, namely $p_3 = S_1(L_2) \cap L_3$ and $p_2 = S_1^{-1}(p_3)$. The point p_2, p_3 vary linearly with p_1 . This constructs a line of inscribed triangles, and every other inscribed triangle is among the ones constructed. **Lemma 6.2** Let J be some obtuse polygon. There is a finite union $L_{J,1}$ of lines such that if $\lambda \in \mathcal{T} - L_{J,1}$ and J is angle-adapted to λ , then $I(J, \lambda)$ is a finite union of pairwise disjoint embedded polygons in \mathbb{C}^3 .

Proof: Say that the *lines* of J are the lines extending the edges of J. There is a finite list of *exceptional* parameters $\lambda_1, ..., \lambda_m$ such that some triple of lines of J is not unrelated to λ_j . We arbitrarily pick lines $\alpha_1, ..., \alpha_m$ with $\lambda_j \in \alpha_j$.

We call a λ -triangle singularly inscribed in J if at least 2 of the vertices of the triangle are also vertices of J. For each triple (v_1, v_2, e) , where v_1, v_2 are vertices of J and e is an edge of J, there is a union α of 3 lines with the following property. If τ is a triangle of shape λ having vertices v_1, v_2 and the third vertex in e, then $\lambda \in \alpha$. Taking the union over all triples, we define

$$L_{J,1} = \alpha_1 \cup \ldots \cup \alpha_n \cup \bigcup \alpha(v_1, v_2, e).$$

Now we take $\lambda \in \mathcal{T} - L_{J,1}$. We say that the *bricks* of J are the products $E_1 \times E_2 \times E_3$ where E_1, E_2, E_3 are edges of J. We call the brick good if $E_1 \cap E_2 \cap E_3 = \emptyset$ and otherwise bad. In the bad case, two of the edges coincide and the third one is adjacent to these. The product $J \times J \times J \subset \mathbb{C}^3$ has a decomposition into bricks, each one the product of edges of J. Since J is angle adapted to λ , the set $I(J, \lambda)$ is disjoint from all the bad bricks.

Let *B* be a good brick that intersects $I(J, \lambda)$ and let $b = B \cap I(J, \lambda)$. By Lemma 6.1, the set *b* is a line segment, the intersection of a straight line with *B*. Since *J* has no singularly inscribed triangles of shape λ , the endpoints of *b* must lie in the interiors of the 2-dimensional faces of *B*. Let b_0 be one of the endpoints of *b*. There is a unique second brick *B'* such that $b_0 \in B'$. The brick *B'* is also good because it intersects $I(J, \lambda)$. But then $b' = B' \cap I(J, \lambda)$ is also a line segment that has b_0 as an endpoint. This analysis shows that $I(J, \lambda)$ is an embedded graph with straight line edges and all vertex degrees equal to 2.

Remark: If (J, λ) satisfy all the conditions above except that there is exactly 1 singularly inscribed triangle sharing exactly 2 vertices with J, then the corresponding point of $I(J, \lambda)$ lies in a 1 dimensional edge of our brick tiling, and either 0, or 2, or 4 edges of $I(J, \lambda)$ meet at this point. The number of edges must be even because the sum of degrees of any graph is even, and $I(J, \lambda)$ has degree 2 at every other vertex. When there are 4 edges, this is the simplest kind of non-manifold point in $I(J, \lambda)$.

6.2 The Parity Argument

Let (J, λ) be as in Lemma 6.2. We know from Lemma 6.2 that $I(J, \lambda)$ is a finite union of polygons. We just have to establish the fact that $I(J, \lambda)$ has exactly 1 essential component. In this section we show that $I(J, \lambda)$ has an odd number of essential components.

Lemma 6.3 $I(J, \lambda)$ has an odd number of essential components.

Proof: Let Ω_3^{\pm} be the two components of Ω_3 . Let Ω_2 be the set of ordered and unequal pairs $(p_1, p_2) \in J \times J$. Let

$$\pi: \Omega_3 \to \Omega_2$$

be the map which forgets the third point. The map π induces an isomorphism from $H_1(\Omega_3^{\pm}) \to H_1(\Omega_2)$ and moreover π is injective on $I = I(J, \lambda)$. Here The reason here is that the first two points of a triangle of shape λ determine the whole triangle.

We compactify Ω_2 by adding two boundary components. The points near one component have the form (p_1, p_2) where p_2 immediately follows p_1 in the cyclic order of J, and the points near the other component have the reverse property. Let $\overline{\Omega}_2$ denote this compactification. Note that an embedded polygon in Ω_2 is homologically nontrivial if and only if it separates the two boundary components of $\overline{\Omega}_2$.

We partition Ω_2 into two sets as follows: Given $(p_1, p_2) \in \Omega_2$ we choose $p_3 \in \mathbb{C}$ so that (p_1, p_2, p_3) is a λ -triangle. We color (p_1, p_2) red (respectively blue) if p_3 lies in the unbounded (respectively bounded) component of $\mathbb{C} - J$. With this scheme, Ω_2 is partitioned into the red points, the blue points, and the points of $\pi(I)$.

Given that J is adapted to λ , all points sufficiently near one boundary component of $\overline{\Omega}_2$ have one color and all points sufficiently near the other boundary component have the other color. A generic polygonal path connecting one component to the other must therefore intersect $\pi(I)$. But such a path intersects each inessential component an even number of times and each essential component an odd number of times. Hence there are an odd number of essential components.

6.3 One Essential Component

Again let (J, λ) be as in Lemma 6.2. Let κ be the number of essential components of $I(J, \lambda)$. We know that κ is odd.

We define

$$J_1 = J, \qquad J_2 = \frac{\lambda}{\lambda - 1}J, \qquad K = \frac{1}{1 - \lambda}J. \tag{10}$$

We let Δ be the map from Lemma 5.1, the Folding Lemma. This time we change the sign so that $\Delta(s_1, s_2) = s_1 - s_2$. As we remarked in the proof of the Folding Lemma, this sign change makes no difference for the truth of the result. We call Δ the *difference map*.

There is a finite union $L_{J,2}$ of lines such that if $\lambda \in \mathcal{T} - L_{J,2}$ then no side of any one of these 3 polygons is parallel to a side of any of the others. This implies that Δ is a folding map. By the Component Lemma (Lemma 4.1) The fold set B of Δ has exactly 2 essential components.

Let $L_J = L_{J,1} \cup L_{J,2}$. Below we will prove that when $\lambda \in \mathcal{T} - L_J$ the polygon K is adapted to Δ , in the sense of the Component Lemma, and that A has $\kappa + 1$ essential components. But now the Component Lemma (Lemma 5.1) says that the number of essential components of A is at most the number of essential components of B. That is, $\kappa + 1 \leq 2$. Since κ is odd, we have $\kappa = 1$.

Lemma 6.4 The polygon K is adapted to Δ and $\Delta^{-1}(K)$ has $\kappa + 1$ essential components.

Proof: The no-parallel-sides condition guarantees that $A = \Delta^{-1}(K)$ is transverse to the fold set B.

Define

$$T_p(z) = \lambda(z-p) + p. \tag{11}$$

A triple of points (p_1, p_2, p_3) has shape λ if and only if the points are unequal and $T_{p_1}(p_2) = p_3$. Consider now a pair $(p_1, q_2) \in J_1 \times J_2$. We let

$$p_2 = \frac{\lambda - 1}{\lambda} q_2 \in J_1. \tag{12}$$

We compute

$$T_{p_1}(p_2) = (1 - \lambda)(p_1 - q_2).$$
(13)

Therefore, p_1 , p_2 , and $p_3 = T_{p_1}(p_2)$ are the vertices of a triangle in $I(J, \lambda)$ if and only if

$$q_2 \neq \frac{\lambda}{\lambda - 1} p_1, \qquad p_1 - q_2 \in K.$$
 (14)

Assuming that we have picked parametrizations α_1 and α_2 of J_1 and J_2 respectively, we let s_1, s_2 be such that $p_1 = \alpha_1(s_1)$ and $q_2 = \alpha_2(s_2)$. Equation 14 tells us that $\Delta(s_1, s_2) \in K$ if and only if one of two conditions holds:

- 1. $q_2 = \lambda/(\lambda 1)p_1$, or
- 2. p_1 and p_2 are the first two vertices of a triangle in $I(J, \lambda)$.

The pairs (s_1, s_2) satisfying Condition 1 correspond to a single essential component of $\Delta^{-1}(K)$. The remaining pairs, the ones which satisfy Condition 2, correspond to points in $I(J, \lambda)$. So, $\Delta^{-1}(K)$ has 1-essential component coming from Condition 1 and κ essential components coming from Condition 2. This shows that $\Delta^{-1}(K)$ has $\kappa + 1$ essential components.

The map

 $(s_1, s_2) \to (p_1, q_2) \to (p_1, p_2)$

gives an affine homeomorphism between all but one of the components of $\Delta^{-1}(K)$ and the components of $I(J,\lambda)$. The one remaining component is the one coming from Condition 1 above. It is an embedded polygon disjoint from all these. This polygon represents (1,1) in $H_1(\mathbf{T})$; it is essentially the "main diagonal" when J_2 is identified with J_1 . Therefore $\Delta^{-1}(K)$ is a finite union of pairwise disjoint embedded polygons. Since one of the essential components (namely the exceptional one just considered) represents (1,1) in $H_1(\mathbf{T})$ and since the other essential components are disjoint from this one, they also represent (1,1) in $H^1(\mathbf{T})$ when suitably oriented.

As we trace out a (1,1) component in the positive orientation, the corresponding point p_1 winds once around J_1 counter-clockwise. But then so do p_2 and p_3 . Since $p_3 = (1 - \lambda)(p_1 - p_2)$, we see that $p_1 - p_2$ winds counterclockwise around K. Hence Δ has degree 1 with respect to the preferred orientations.

This establishes that K is adapted to the folding map Δ .

Remark: Inspecting the conditions for the lines in L_J , we see that there exist polygons J' as close as like to J such that L_J and L_J have no lines in common. This means that we can vary J continuously so that each line in L_J either varies in a non-constant continuous family or instantly disappears.

6.4 Proof of Lemma 2.8

In his section we prove Lemma 2.8, the technical result we used in §2. Let J be some polygonal Jordan loop and let L_J be the finite union of lines defined in the previous section. Let (J, A) be as in Lemma 2.8.

Let $t \to \lambda_t$ be a constant speed parameterization of A. The two endpoints are λ_0 and λ_1 . Since a line can intersect a circle (or an unequal line) at most twice, there are finitely many values $0 < t_1, ..., t_k < 1$ such that (J, t_j) does not satisfy the conclusion of Theorem 1.7.

Let P_J be the union of the exceptional parameters from the proof of Lemma 6.2 and the points that lie in more than one line of L_J . The set P_J finite. If A avoids P_J , the space $I(J, t_j)$ has the structure discussed in the remark after the proof of Lemma 6.2 for each j = 1, ..., k.

Let $\Upsilon_0, ..., \Upsilon_k$ be components of $[0, 1] - (t_1 \cup ... \cup t_k)$. Let

$$B_j = \bigcup_{t \in \Upsilon_j} \beta_t.$$

Here β_t is the essential component of $I_t = I(J, \lambda_t)$.

Inspecting the proof of Lemma 6.2 we see that β_t varies continuously for $t \in \Upsilon_j$. Hence B_j is a topological (and indeed piecewise algebraic) cylinder. These sets are path connected. Given the structure of $I(J, t_j)$, and the result of Theorem 1.7, here are the only possible topological changes to I_t as t passes through a special value.

- An inessential component is born or dies.
- Two inessential components merge into one, or the reverse.
- The essential component merges with an inessential component to form the new essential component, or the reverse.

From this analysis of the topological transitions, we see that $\overline{B}_j \cap \overline{B}_{j+1}$ is either a polygonal loop in I_{t_j} or two polygonal loops of I_{t_j} joined together at the exceptional vertex. Hence

$$B = \bigcup_{j=0}^{k} \overline{B}_k$$

is a path connected subset of I(J, A) which contains both β_0 and β_1 . This completes the proof of Lemma 2.8.

7 References

[**AA**] A. Akopyan and S Avvakumov, Any cyclic quadrilateral can be inscribed in any closed convex smooth curve. arXiv: 1712.10205v1 (2017)

[ACFSST] J. Aslam, S. Chen, F. Frick, S. Saloff-Coste, L. Setiabrate, H. Thomas, *Splitting Loops and necklaces: Variants of the Square Peg Problem*, arXiv 1806.02484 (2018)

[H] C. Hugelmeyer, Every Smooth Jordan Curve has an inscribed rectangle with aspect ratio equal to $\sqrt{3}$. arXiv 1803:07417 (2018)

[Mak1] V. Makeev, On quadrangles inscribed in a closed curve, Math. Notes 57(1-2) (1995) pp. 91-93

[Mak2] V. Makeev, On quadrangles inscribed in a closed curve and vertices of the curve, J. Math. Sci. 131(1) (2005) pp 5395-5400

[Ma1] B. Matschke, A survey on the Square Peg Problem, Notices of the A.M.S. Vol 61.4, April 2014, pp 346-351.

[Ma2] B. Matschke, *Quadrilaterals inscribed in convex curves*, arXiv 1801:01945v2

[**M**] M. Meyerson *Equilateral Triangles and Continuous Curves*, Fundamenta Mathematicae, 110.1, 1980, pp 1-9.

[**N**] M. Neilson, *Triangles Inscribed in Simple Closed Curves*, Geometriae Dedicata, 1991

[**NW**] M. Neilson and S. E. Wright, *Rectangles inscribed in symmetric continua*, Geometriae Dedicata **56(3)** (1995) pp. 285-297

[**P**] I. Pak, *Lectures on Discrete and Polyhedral Geometry*, online book, www.math.ucla.edu/~pak/geompol8.pdf, (2018)

[S] R. E. Schwartz, A Trichotomy for Rectangles Inscribed in Jordan Loops, preprint 2018.

[Shn], L. G. Shnirelman, On certain geometric properties of closed curves (in Russian), Uspehi Matem. Nauk 10 (1944) pp 34-44;

[**Ta**], T. Tao, An integration approach to the Toeplitz square peg conjecture Fourm of Mathematics, Sigma, 5 (2017)