A Trichotomy for Rectangles Inscribed in Jordan Loops

Richard Evan Schwartz *

November 26, 2018

Abstract

Let γ be an arbitrary Jordan loop and let $G(\gamma)$ denote the space of rectangles R which are inscribed in γ in such a way that the cyclic order of the vertices of R is the same whether it is induced by R or by γ. We prove that $G(\gamma)$ contains a connected set S satisfying (at least) one of three properties.

1. S consists of rectangles of uniformly large area, including a square, and every point of γ is the vertex of a rectangle in S.
2. S consists of rectangles having all possible aspect ratios, and all but at most 4 points of γ are vertices of rectangles in S.
3. S contains rectangles of every sufficiently small diameter, and all but at most 2 points of γ are vertices of rectangles in S.

In particular, every Jordan loop has the property that at most 4 points on it are not vertices of inscribed rectangles.

1 Introduction

A *Jordan loop* is the image of the circle under a continuous injective map into the plane. O. Toeplitz conjectured in 1911 that every Jordan loop contains 4 points which are the vertices of a square. This is sometimes called the *Square Peg Problem.* There is a long literature on the problem. For historical details

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

and a long bibliography, we refer the reader to the excellent survey article [**M**] by B. Matschke, written in 2014.

An affirmative answer to the Square Peg Problem is known under various restrictions on the nature of the Jordan loop. For instance, the result is known for polygonal loops and smooth loops. A number of authors have widened the class of Jordan loops for which the result is true. See, for instance, the recent paper of T. Tao [**Ta**]. A recent paper of C. Hugelmeyer [**H**] shows that a smooth Jordan loop always has an inscribed rectangle of aspect ratio $\sqrt{3}$. The proof, though short, is far from elementary: It requires a result from Heegard Floer homology! The recent paper [**AA**] proves that any cyclic quadrilateral can (up to similarity) be inscribed in any convex smooth curve.

We insist that our Jordan loops are always oriented counter-clockwise around the regions they bound. Say that a rectangle R *graces* a Jordan loop γ if the vertices of R lie in γ and if the counter-clockwise cyclic ordering on the vertices induced by R coincides with the cyclic ordering induced by γ.

Let $G(\gamma)$ denote the space of counter-clockwise labeled gracing rectangles. The *aspect ratio* of such a rectangle is the length of the second side divided by the length of the first side. $G(\gamma)$ is naturally a subset of $\boldsymbol{R}^8$ and as such inherits a metric space structure.

Theorem 1.1 (Trichotomy) *Let γ be an arbitrary Jordan loop. Then $G(\gamma)$ contains a connected set S satisfying one of three properties.*

1. *S consists of rectangles of uniformly large area, including a square, and every point of γ is the vertex of a rectangle in S.*

2. *S consists of rectangles having all possible aspect ratios, and all but at most 4 points of γ are vertices of rectangles in S.*

3. *S contains rectangles of every sufficiently small diameter, and all but at most 2 points of γ are vertices of rectangles in S.*

Remarks:

(i) I would describe the three cases respectively as *elliptic*, *hyperbolic*, and *parabolic*, because the geometry of the situation seems to vaguely resemble the action of these kinds of linear transformations on $\boldsymbol{R}^2$. Note that more than one case could occur (e.g. for a circle).

(ii) Ruling out the parabolic case would resolve the Square Peg problem.

(iii) The elliptic case occurs for any curve with 4-fold rotational symmetry but I conjecture that the hyperbolic case also occurs at the same time.
(iv) The following result is an immediate consequence of our proof. For any non-atomic measure μ of total mass 1 on γ there is a rectangle R inscribed in γ such that the total μ-measure of each pair opposite sides of γ cut of by R is 1/2. See §5.7. This is tantalizingly close to the square-peg conjecture, but doesn't imply it.

Here is an immediate corollary.

Corollary 1.2 *Let γ be any Jordan loop. Then all but at most 4 points of γ are vertices of rectangles inscribed in γ.*

This result is sharp: Consider a non-circular ellipse.

Remark: A paper recently appeared on the arXiv, [**ACFSST**], which shows by completely different methods that every Jordan Loop contains a dense set of points which are vertices of inscribed rectangles. This result also has the result mentioned in Remark (iv) above, at least when γ is rectifiable and μ is arc-length measure normalized to have total length 1.

Our proof of the Trichotomy Theorem derives from taking the limit of what happens for a generic polygon. Now let γ be a polygon. By an *arc component* of $G(\gamma)$ we mean a connected component of $G(\gamma)$ which is homeomorphic to an arc. By *proper*, we mean that as one moves towards an endpoint of an arc component in $G(\gamma)$, the aspect ratio tends either to 0 or to ∞.

Theorem 1.3 *For an open dense set of polygonal loops γ the space $G(\gamma)$ is a piecewise smooth 1-manifold whose arc components are proper. The ends of $G(\gamma)$ are in naturally in bijection with the positively oriented critical points of the distance function $d : (\gamma \times \gamma) - \Delta \to \boldsymbol{R}$. Here Δ is the diagonal in $\gamma \times \gamma$.*

Each critical point of the distance function defines a chord χ of the polygon γ. We say that χ is *positively oriented* if γ meets the line extending χ in opposite orientationsat the two endpoints of χ. (The arrows in Figure 2 from §3 indicate what we mean.) Being positively oriented is a necessary and sufficient condition for gracing rectangles to accumulate on χ.

We say that an arc component A of $G(\gamma)$ is a *sweepout* if the aspect ratio of the rectangles in A tends to 0 as the rectangles move towards one endpoint of A endpoint of A and ∞ at the other. Figure 1 shows an example.

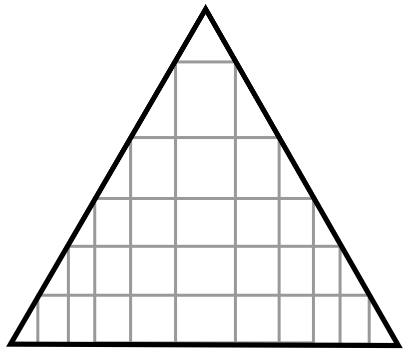

Figure 1: A sweepout on an equilateral triangle.

The operation of cyclically relabeling gives a $\boldsymbol{Z}/4$ action on the space $G(\gamma)$ which has no fixed points. We call a compact component of $G(\gamma)$ *elliptic* if it is stabilized by the $\boldsymbol{Z}/4$ action. The $\boldsymbol{Z}/4$ action freely permutes the arc components. Hence the labeled arc components all come in clusters of 4. When we forget the labelling, we divide by 4. Define the following quantities:

- $\Omega(\gamma)$ is the number of unlabelled gracing squares.
- $\Omega'(\gamma)$ is the number of unlabelled sweepouts.
- $\Omega^*(\gamma)$ is the number of elliptic components.

We will show that

$$\Omega(\gamma) + \Omega'(\gamma) + \Omega^*(\gamma) \equiv 0 \bmod 2 \tag{1}$$

for an open dense set of polygons.

The sweepouts and elliptic components are global in nature. For instance every point of γ except at most 4 is the vertex of a rectangle in a sweepout, and every vertex of γ is the vertex of a rectangle in an elliptic component. The fact that $G(\gamma)$ has one of these global components is what powers the Trichotomy Theorem. By showing that $\Omega(\gamma)$ is odd we establish the folliowing theorem.

Theorem 1.4 *For an open dense set of polygons γ the sum $\Omega'(\gamma) + \Omega^*(\gamma)$ is odd. Thus $G(\gamma)$ either has a sweepout or an elliptic component (or both).*

In §2 we will study the problem of taking four lines L_0, L_1, L_2, L_3, not necessarily all distinct, and finding points a_0, a_1, a_2, a_3 with $a_i \in L_i$ for all indices $i = 0, 1, 2, 3$, such that a_0, a_1, a_2, a_3 are the counterclockwise ordered vertices of a rectangle. We will identify a subset X_1 in the space X_0 of all such 4 tuples on which this problem has a nice solution. Lemma 2.3 states the result. Theorem 2.5 says that the complement $X_0 - X_1$ has codimension 2, and this is the key to our homotopy proof of Theorem 1.4.

In §3 we prove most of Theorem 1.3. The strategy is to define what we mean by a generic polygon precisely enough so that the manifold structure in Theorem 1.3 becomes an easy application of Lemma 2.3.

In §4 we finish the proof of Theorem 1.3 by showing that $\Omega(\gamma)$ is odd.

In §5 we prove the Trichotomy Theorem by taking a suitable limit of the either a sweepout or a elliptic component.

One thing I would like to mention is that I discovered all the results in this paper by computer experimentation. I wrote a Java program which computes the space $G(\gamma)$ in an efficient way for polygonal loops γ having up to about 20 sides.

I would like to thank Peter Doyle, Cole Hugelmeyer, and Sergei Tabachnikov for helpful and interesting conversations about the ideas in this proof. Some of this work was done while I was at the Isaac Newton Institute in Spring 2017, supported by the INI and also by a Simons Sabbatical Fellowship. I was also supported by the National Science Foundation while completing this work. I would like to thank all these institutions for their generous support.

2 Rectangles and Lines

2.1 The Generic Case

Given $V = (x, y) \in \boldsymbol{R}^2$, we define $V' = (-y, x)$. We get V' by rotating V counterclockwise by $\pi/2$ radians.

We consider 4 lines $L_0, L_1, L_2, L_3 \in \boldsymbol{R}^2$. These lines need not be distinct. We say that a rectangle R *graces* our lines if R has vertices a_0, a_1, a_2, a_3 with $a_j \in L_j$ for all j, and these points go in counter-clockwise order around R. Now we find formulas for the rectangles which grace our lines.

We rotate so that none of these lines is vertical. This lets us find constants m_j, b_j so that

$$a_j(x) = (x, m_j x + b_j) \in L_j \tag{2}$$

for all $x \in \boldsymbol{R}$.

Given $s, t, \lambda \in \boldsymbol{R}$ we define the rectangle $R = R(s, t, \lambda)$ with vertices $a_0(s), a_1(t)$, $a_2(s, t)$ and $a_3(s, t)$, where

$$a_2(s,t) = a_1(t) + \lambda V', \quad a_3(s,t) = a_0(s) + \lambda V', \qquad V = a_1(t) - a_0(s). \tag{3}$$

The side connecting $a_0(s)$ to $a_1(t)$ is the first side of R. Hence, according to our definition of aspect ratio given in the introduction, R has aspect ratio λ.

The conditions that $a_j(s, t) \in L_j$ for $j = 2, 3$ lead to two equations in two unknowns:

$$\rho_j + \sigma_j s + \tau_j t, \qquad j = 2, 3. \tag{4}$$

Each coefficient is linear in r. When these equations do not describe parallel lines in (s, t)-space, their unique solution is given by

$$s = \frac{S_0 + S_1\lambda + S_2\lambda^2}{U_0 + U_1\lambda + U_2\lambda^2}, \qquad t = \frac{T_0 + T_1\lambda + T_2\lambda^2}{U_0 + U_1\lambda + U_2\lambda^2}, \tag{5}$$

where

$$S_0 = -(b_0 - b_3)(m_1 - m_2), \quad T_0 = -(b_1 - b_2)(m_0 - m_3), \quad U_0 = (m_1 - m_2)(m_0 - m_3)$$

$$S_1 = -b_0 + b_1 - b_2 + b_3 - b_0 m_1 m_2 + b_3 m_1 m_2 + b_0 m_1 m_3 - b_2 m_1 m_3 - b_0 m_2 m_3 + b_1 m_2 m_3.$$

$$T_1 = -b_0 + b_1 - b_2 + b_3 - b_1 m_0 m_2 + b_3 m_0 m_2 + b_1 m_0 m_3 - b_2 m_0 m_3 - b_0 m_2 m_3 + b_1 m_2 m_3$$

$$U_1 = m_0 - m_1 + m_2 + m_0 m_1 m_2 - m_3 - m_0 m_1 m_3 + m_0 m_2 m_3 - m_1 m_2 m_3$$

$$S_2 = -(b_0 - b_1)(m_2 - m_3) \quad T_2 = -(b_0 - b_1)(m_2 - m_3) \quad U_2 = (m_0 - m_1)(m_2 - m_3)$$

We found this equation using Mathematica [**W**]. Note that the denominators only depend on the slopes.

2.2 Definedness

In this generic situation, Equation 5 describes all the gracing rectangles. In the next section we will discuss the special case where it does not.

Lemma 2.1 *The functions $s(\lambda)$ and $t(\lambda)$ are well defined provided that the lines are not all parallel.*

Proof: For well-definedness, all we need is that some U_j does not vanish. If both U_0 and U_2 vanish then three of the slopes are the same. When we set three of the m_j variables to m and the other one to M we find that $U_1 = \pm(1+m^2)(M-m)$. The sign depends on which three we have chosen. So, U_1 is nonzero as long as $M \neq m$. ♠

We note certain *coincidences* that might happen for L.

1. Two lines of L are parallel or equal.
2. Two lines of L are perpendicular.
3. Three lines of L have a common intersection.
4. All 4 lines are distinct, and the line through $L_a \cap L_b$ and $L_c \cap L_d$ is perpendicular to the line L_a. These indices are meant to be distinct.

We let $\chi(L)$ denote the number of such coincidences which occur. For instance, if three lines of L are parallel then $\chi(L) \geq 3$.

Lemma 2.2 *$s(\lambda)$ and $t(\lambda)$ are nontrivial rational functions if $\chi(L) \leq 1$.*

Proof: We first work with $s(\lambda)$. The condition that $s(\lambda)$ is constant is a geometric one. It means that there is a single point on L_0 that is the vertex of infinitely many gracing rectangles. We will suppose that this happens and show that $\chi(L) \geq 2$. Given that the condition is geometric, we can translate the picture so that $(0,0)$ is the offending point. This gives $b_0 = 0$ and $S_0 = S_1 = S_2 = 0$. The condition $S_0 = 0$ gives $b_3 = 0$ or $m_1 = m_2$. The condition $S_2 = 0$ gives $b_1 = 0$ or $m_3 = m_2$. If $m_1 = m_2 = m_3$ then $\chi(L) \geq 2$. Now we consider the other 3 cases.

- Suppose $b_1 = 0$ and $b_3 = 0$. Then $S_1 = b_2(m_1 m_3 + 1)$. Note that $b_2 \neq 0$ because not all the lines contain the origin. This gives $m_1 m_3 = -1$, making L_1 and L_3 perpendicular. In this case L_0, L_1, L_3 contain the origin and L_1, L_3 are perpendicular. Hence $\chi(L) \geq 2$.

- Suppose that $b_3 = 0$ and $m_2 = m_3 = m$. This gives $\chi(L) \geq 1$. If $\chi(L) = 1$ then $b_1 \neq 0$ and no lines are perpendicular to L_2 or L_3. In particular, $L_0 \cap L_3$ and $L_1 \cap L_2$ are unique and distinct points. Thus, we can rotate so that $m = 0$. The condition $S_1 = 0$ now implies that $b_1 = b_2$. The intersection $L_1 \cap L_2$ lies on the y-axis. Since $b_0 = 0$ and $b_3 = 0$, the intersection $L_0 \cap L_3$ also lies on the y-axis. But then the y-axis, which contains both $L_0 \cap L_3$ and $L_1 \cap L_2$, is perpendicular to L_2. This gives a coincidence of Type 4. Hence $\chi(L) \geq 2$.

- Suppose that $b_1 = 0$ and $m_2 = m_1$. We get the same situation as in the previous case with the indices 1 and 3 swapped.

This completes the proof for $s(\lambda)$. The proof for $t(\lambda)$ is the same but with the indices suitably permuted. Alternatively, we can deduce this case from symmetry. Let ρ denote reflection in the y-axis. Let L^* denote the set of lines $\rho(L_1), \rho(L_0), \rho(L_2), \rho(L_3)$. Then $s(\lambda)$ is constant with respect to L^* if and only if $t(\lambda)$ is constant with respect to L. Thus, the result for $t(\lambda)$ follows from the result for $s(\lambda)$ and symmetry. ♠

The following result is the key to most of our analysis.

Lemma 2.3 *Let L_0, L_1, L_2, L_3 be 4 lines with $\chi(L) \leq 1$. Then there is a finite subset $\Sigma \subset \boldsymbol{R}$ with the following property. A rectangle R_λ of aspect ratio λ graces the lines if and only if $\lambda \notin \Sigma$. In this case, R_λ is unique.*

Proof: Since $s(\lambda)$ and $t(\lambda)$ are nontrivial rational functions, each choice of λ gives a unique rectangle of aspect ratio provided that these functions do not blow up at λ. The blow-up can only happen at finitely many points. ♠

2.3 The Coherent Case

The analysis above is incomplete when there is a value of λ such that the two equations in Equation 4 describe the same line. In this case, there is an extra line of solutions not captured by Equation 5, and we call our lines λ-*coherent*. We call λ the *coherent value*.

In the λ-coherent case, we have an infinite set of rectangles, all of aspect ratio λ, which grace the 4 lines. In particular, we can find 2 rectangles R_1 and R_2 of aspect ratio λ such that the line L_j contains the jth vertex of R_1 and the jth vertex of R_2, and all these vertices are unequal. Conversely, if we start with 2 such rectangles, and define the lines through their corresponding vertices, we produce a λ-coherent example. In short, all λ-coherent examples arise this way. It is worth recording this as a neat little theorem in similarity geometry.

Theorem 2.4 *Let R_1 and R_2 be two labeled rectangles of the same aspect ratio λ, having pairwise distinct vertices. Let L_1, L_2, L_3, L_4 denote the lines such that L_j contains the jth vertex of R_1 and R_2. Then there is a continuous 1-parameter family of rectangles, each of which graces L_1, L_2, L_3, L_4 and has aspect ratio λ.*

The λ-coherent case is very pretty, but we want to ignore it. Generically, the lines L_1, L_2, L_3, L_4 are not λ-coherent for any value of λ. Also, for each particular value of λ, the set of λ-coherent lines has codimension 2. We only need to know this fact when $\lambda = 1$, so in this case we will justify the codimension 2 claim.

We rotate and translate the picture so that L_0 is the x-axis. That is, $m_0 = b_0 = 0$. It suffices to show that, amongst the set of 4-tuples with $m_0 = b_0 = 0$, the set of 1-coherent 4-tuples has codimension 2. If our lines are 1-coherent, then we must have $\rho_2\sigma_3 = \rho_3\sigma_2$ and $\rho_2\tau_3 = \rho_3\tau_2$. This leads to the equation

$$m_1 = \frac{m_2 - m_3 + m_2 m_3}{1 + m_2 + m_2 m_3}, \qquad b_1 = \frac{b_2 - b_3 + b_2 m_3}{1 + m_2 + m_2 m_3}.$$

These are 2 independent algebraic equations. Hence, the solution set is contained in a finite union of manifolds of codimension 2.

2.4 Codimension Two Conditions

Here we prove a result that will be very useful in proving our parity result about the number of sweepouts. Let X_0 be the set of all 4-tuples of lines. This is naturally a connected smooth manifold. Let $X_1 \subset X_0$ denote the set of configurations L such that $\chi(L) < 1$ and L is not 1-coherent.

Theorem 2.5 *$X_0 - X_1$ is contained in a finite union of smooth submanifolds of codimension 2. In particular, X_1 is open dense in X_0 and path connected.*

We will prove this result through two smaller lemmas. Since the set of 1-coherent configurations is contained in a finite union of codimension 2 manifolds, it suffices to analyze the configurations in $X_0 - X_1$ which are not 1-coherent. In other words, we just have to show that the set of configurations L having $\chi(L) > 1$ is contained in a finite union of codimension 2 submanifolds.

In the definition of χ, there are finitely many coincidences. We enumerate these coincidences in some order. For instance, Coincidence 0 could be the event that the lines L_0 and L_1 are equal or coincide. Let $M_i \subset X_0$ denote the set of those configurations which have coincidence i. There are 4 types of coincidences, and they are listed in §2.2. We define the *Type* of M_i to be the type of the coincidence associated to it.

Let $\zeta \in M$ be a point. Let L be the configuration corresponding to ζ. We define a *rotating deformation* as follows: We choose a line of L and rotate that line about some point in the plane. Such a deformation defines a nontrivial tangent vector in the tangent space $T_\zeta(X_0)$.

Lemma 2.6 *M_i is a smooth codimension 1 submanifold.*

Proof: We set $M = M_i$. Let $\zeta \in M$. We rotate so that none of the lines in the configuration L corresponding to ζ is vertical. We first show that M is the zero set of a smooth function in a neighborhood of ζ. There are 4 cases, corresponding to the types of coincidences.

1. If L_a and L_b are equal or parallel, then we set $F(\zeta) = m_a - m_b$.
2. If L_a and L_b are perpendicular, then we set $F(\zeta) = m_a m_b + 1$.
3. If L_a, L_b, L_c have a common point, we set $F(\zeta) = \det(V_a, V_b, V_c)$, where $V_a = (m_a, -1, b_a)$ is the vector dual to L_a, etc.
4. We have non-parallel lines L_a and L_b, and nonparallel lines L_c and L_d such that the line Λ through $L_a \cap L_b$ and $L_c \cap L_d$ is perpendicular to L_a. The slope μ of Λ is a rational function of the slopes m_a, m_b, m_c, m_d. Thus, we set $F(\zeta) = m_a \mu + 1$.

In all cases, a suitable rotating deformation destroys the coincidence. Therefore the differential dF is nonsingular at ζ. Our result now follows from the Implicit Function Theorem. ♠

Lemma 2.7 *If $i \neq j$ then $M_i \cap M_j$ is a smooth manifold of codimemsion* 2.

Proof: Let L be the configuration corresponding to ζ. The basic idea in the proof is to produce two rotating deformations V_i and V_j such that each deformation preserves one of the coincidences and destroys the other one. If M_i and M_j are defined by smooth functions in a neighborhood of ζ then the existence of V_i and V_j shows that the differentials dF_i and dF_j are linearly independent at ζ. This proves the result. What follows is a case-by-case analysis. Let (a, b) denote the pair such that M_i has Type a and M_j has Type b. We take $a \leq b$.

1. Suppose $b \leq 2$ or $(a, b) = (3, 3)$. In this case, there is a line in Coincidence i that is not involved in Coincidence j and *vice versa*. Any rotating deformations with respect to these lines do the job.

2. Suppose $a \leq 2$ and $b = 3$. The only way the argument in Case 1 breaks down is that both lines L_α and L_β involved in Coincidence i belong to the triple of lines $L_\alpha, L_\beta, L_\gamma$ involved in Coincidence j. In this case, we let V_i be the deformation that rotates L_α about a point of $L_\alpha \cap L_\beta \cap L_\gamma$ and let V_j be any rotation deformation of L_γ.

3. Suppose $a \leq 3$ and $b = 4$. In Coincidence j, the line L_α is perpendicular to the line through $p_{\alpha\beta} = L_\alpha \cap L_\beta$ and $p_{\gamma\delta} = L_\gamma \cap L_\delta$. We call L_α the *special line* and $p_{\alpha\beta}$ and $p_{\gamma\delta}$ the *special points*. Let $\alpha' \neq \alpha$ be an index of some line involved in Coincidence i. Let β' be the index of some line not involved in Coincidence i. We let V_i be the deformation that rotates $L_{\alpha'}$ about the special point that contins it and we let V_j denote the rotation of $L_{\beta'}$ about a generic point.

4. Suppose $(a, b) = (4, 4)$. Let α' be so that $L_{\alpha'}$ is not the special line with respect to either coincidence. Let V_i be the deformation obtained by rotating $L_{\alpha'}$ about the special point w.r.t. Coincidence j that it contains. Likewise define V_j. These deformations have the desired properties unless the two relevant special points coincide. So, we just

have to worry about the case when the special points w.r.t. Coincidence i coincide with the special points w.r.t. Coincidence j. In this case, the special lines are distinct. For $k \in \{i,j\}$ we let V_k be the rotation of the special line w.r.t. Coincidence k about the special point it contains.

This completes the proof. ♠

2.5 Compatible Perturbations

Remark: This section is rather technical. It will help us in §3.3 when we need to fix the local properties of squares which grace a polygon and share a vertex with that polygon. Later we will call these *critical rectangles*. We recommend that the reader skim the section on the first pass.

Let $L = (L_0, L_1, L_2, L_3)$ be as above. Let Σ be as in Lemma 2.3. Given $\lambda_0 \in \boldsymbol{R} - \Sigma$, there is a unique rectangle $R(\lambda_0)$ of aspect ratio λ_0 that graces L. As λ varies in a neighborhood of λ_0, the vertex $a_0(\lambda)$ varies as well. The point $a_0(\lambda)$ moves monotonically in a neighborhood of λ_0 provided that $s'(\lambda_0) \neq 0$. Here $s' = ds/d\lambda$. Since s' is a rational function of λ, there is a finite set $\Sigma_0 = \Sigma_0(L)$ such that $s'(\lambda_0) = 0$ if and only if $\lambda_0 \in \Sigma_0$.

Let L^* be another such configuration. We assume that $\chi(L) \leq 1$. This means that $\chi(L^*) \leq 1$ provided that L^* and L are sufficiently close. We also specify a point $a_0 \in L_0$ and $a_0^* \in L_0^*$ We call the pairs (L, a_0) and (L^*, a_0^*) *compatible* if the following is true.

- $a_0^* = a_0$.
- If $a_0 \in L_j$ and $j \neq 2$, then $L_j^* = L_j$.
- $L_i^* = L_j^*$ if and only if $L_i = L_j$.

L^* is a perturbation of L in which we do not split apart any equal lines and we do not touch L_0 or any line adjacent to L_0 that contains a_0.

We assume that there is some rectangle R gracing L which has a_0 as its first vertex. This means that $\lambda_0 \in \boldsymbol{R} - \Sigma(L)$, where λ_0 is the aspect ratio of R.

Lemma 2.8 *Suppose that $\lambda_0 \in \Sigma_0(L)$. Then there is a configuration L^*, arbitrarily close to L, with the following properties.*

- (L, a_0) *and* (L^*, a_0^*) *are compatible.*

- a_0^* *is the first vertex of a rectangle* R^* *that graces* L^**.*

- $\lambda_0^* \not\in \Sigma_0(L^*)$*, where* λ_0^* *is the aspect ratio of* R^**.*

Proof: Suppose first that the only lines in the set $\{L_3, L_1\}$ which contain L actually coincide with L_0. We choose $\lambda' \in \boldsymbol{R} - \Sigma - \Sigma_0(L)$ that is very close to λ and let R' be the corresponding rectangle that graces L. Letting τ be the translation along L_0 that moves the vertex a_0' of R' back to a_0, we define $L^* = \tau(L)$ and $R^* = \tau(R)$. These objects have the desired properties.

It cannot happen that L_3, L_0, L_1 all contain a_0 because then L_3 and L_1 extend two adjacent sides of R and hence are perpendicular. This yields $\chi(L) \geq 2$. So, a_0 can be in at most 2 of these lines. The only case left to consider is when a_0 lies in L_0 and exactly one consecutive line, and these lines are distinct. We consider the case when $a_0 = L_0 \cap L_1$. The case when $a = L_3 \cap L_0$ has a similar treatment and indeed follows from symmetry.

If $L_2 = L_3$ then this common line contains two points of R, and L_1 contains the other two points. But then L_1, L_2, L_3 are all either parallel or equal. This yields $\chi(L) \geq 2$. Hence $L_2 \neq L_3$. If $L_2 = L_1$ then the points a_0, a_1, a_2 all lie on L_1, and this is impossible for a three vertices of a rectangle. Hence $L_2 \neq L_1$. If $L_2 = L_0$ then $\chi(L) \geq 2$ because $L_2 = L_3$ and L_0, L_1, L_2 all contain a_0. Hence $L_2 \neq L_0$. Similar arguments show that $L_3 \neq L_0$ and $L_3 \neq L_1$. In short, all 4 lines are distinct and an arbitrary perturbation which just moves L_2 and L_3 is compatible with L provided that the perturbation is small enough.

Given the geometric nature of the situation, we can freely rotate and translate for the purpose of making the calculation easier. We rotate and translate so that $a_0 = (0, 0)$ and L_0 and L_1 have nonzero slopes. This gives $b_0 = b_1 = 0$, and both m_0 and m_1 are nonzero. Given the algebraic nature of the situation, we just have to find *some* variation of just L_2 and L_3 which has the property that $s'(\lambda)$ is either nonzero or infinite for all $\lambda \in \boldsymbol{R}$. To do this, we set $(m_2, m_2, m_3, b_3) = (0, 0, 0, 1)$. With these choices, we compute

$$s'(\lambda) = \frac{m_1^2}{(m_0 m_1 + m_0 \lambda - m_1 \lambda)^2}$$

This expression is either nonzero or infinite for any $\lambda \in \boldsymbol{R}$. ♠

3 The Moduli Space of Rectangles

3.1 The Basic Spaces

We fix an integer $N \geq 3$. Let $S_0 = S_0(N) \subset (\boldsymbol{R}^2)^N$ denote the open set of all embedded labeled N-gons. We often suppress the dependence on N in our notation. The space S_0 inherits a metric space from $(\boldsymbol{R}^2)^N$ and also a smooth manifold structure. Our first main goal of this chapter is to define an open dense subset $S_3 \subset S_0$ which consists of sufficiently generic polygons to make the proof of Theorem 1.3 a fairly easy consequence of Lemma 2.3.

We call the vertices of $\gamma \in S_0$ by the names $v_1, ..., v_N$. We define the jth edge of γ to be the one which has v_j as its lagging vertex. Thus, going around counter-clockwise, we encounter $v_1, e_1, v_2, e_2, ...$

We consider cyclically ordered quadriples in $\{1, ..., N\}^4$ of the following form:

$$(a, b, c, d),\ (a, a, b, c),\ (a, b, b, c),\ (a, b, c, c),\ (a, b, c, a).$$

Different-lettered indices are meant to correspond to different edges of γ. We call such quadruples *relevant*. The cyclic ordering means, for instance, that $(N, 1, 2, 3)$ would be relevant. Given a relevant quadruple I, we let L_I be the quadruple of lines extending the corresponding edges of γ.

Say that a *chord* of γ is a line segment joining two distinct vertices of γ. Let S_1 denote the set of embedded N-gons γ with the following properties.

1. No two distinct chords of γ are parallel or perpendicular.
2. Any relevant list L of lines has $\chi(L) \leq 1$. Here χ is as in §2.2.

By construction, S_1 is open dense in S_0.

Lemma 3.1 *Any rectangle that graces γ also graces a relevant quadruple of lines of γ.*

Proof: Condition 1 guarantees that any rectangle R that graces $\gamma \in S_1$ has vertices on at least 3 distinct sides of γ. We order the vertices of R in increasing order and we let i_j be the smallest possible index of an edge that contains the jth vertex of R. By construction the quadruple (i_0, i_1, i_2, i_3) is of the kind considered above, the corresponding quadruple of lines is graced by R. ♠

Note that there might be other indices associated to our rectangle R than the one we defined in the proof of the preceding lemma. This can happen of some of the vertices of R are also vertices of $\gamma \in S_1$. We will study this situation carefully. We call R *critical* if at least one of the vertices of R is a vertex of γ.

Lemma 3.2 *Suppose $\{\gamma_n\}$ is a sequence in S_1 which converges to $\gamma \in S_1$. Suppose that R_n is a critical rectangle associated to γ_n. Then there is a uniform lower bound to the side length of a side of R_n.*

Proof: Since $\gamma_n \to \gamma$ there is a uniform lower bound to the distance between distinct vertices of γ_n. Hence there is a uniform positive lower bound to the diameter of R_n. So, we just have to worry about a pair of opposite sides of R_n shrinking to points. Let v_n be a vertex of R_n that is also a vertex of γ_n. Passing to a subsequence, we can assume that $v_n \to v$, a vertex of γ.

Once n is sufficiently large, the short side of R_n incident to v_n must lie in a single edge e_n of γ_n. Passing to a subsequence, we can assume that $e_n \to e$, an edge of γ.

Let w_n be the vertex of R_n opposite to v_n. Passing to a subsequence, we can assume that $w_n \to w \in \gamma$. The short side of R_n incident to w_n has its endpoints in two different edges of γ_n. Hence w is a vertex of γ. By constuction, the line $\overline{wv}$ is perpendicular to e. This violates Condition 1 above. ♠

3.2 Eliminating Doubly Critical Rectangles

Suppose that R is a rectangle that graces γ. We call R *doubly critical* if R is a square or if at least two of the vertices of R are vertices of γ. Let $S_2 \subset S_1$ denote the set of polygons which have no doubly critical rectangles.

Lemma 3.3 *S_2 is open in S_1.*

Proof: Suppose that $\gamma \in S_1$ is the limit of a sequence of polygons $\{\gamma_n\}$ in $S_1 - S_2$ having double critical rectngles. By Lemma 3.2 we can assume that $R_n \to R$, a rectangle that graces γ. By continuity, R_n is doubly critical. Hence $\gamma \in S_1 - S_2$. Hence $S_1 - S_2$ is closed in S_1. Hence S_2 is open in S_1. ♠

Lemma 3.4 *Any $\gamma \in S_1$ has only finitely many doubly critical rectangles.*

Proof: We consider the case when a critical rectangle has more than one vertex in common with γ. For any two distinct vertices $v, w \in \gamma$ we consider the two lines V and W through v and w respectively which are perpendicular to the line $\overline{vw}$. Condition 1 above quarantees that $\gamma \cap V$ and $\gamma \cap W$ are finite sets. If a doubly critical rectangle involves v and w as adjacent vertices, then there is a point of $\gamma \cap V$ that is the same distance from v as some point of $\gamma \cap W$ is from some point of w. There are only finitely many intersecton points like this. A similar argument, involving the circle with diameter $\overline{vw}$ through v and w, works when v and w are opposite vertices.

Now we consider the square case. Every rectangle that graces γ is associated to some multi-index. There are only finitely many such multi-indices and Lemma 2.3, applied to the case $\lambda = 1$, tells us that there are only finitely many squares gracing the set of lines corresponding to each one. Hence there are only finitely many gracing squares and in particular only finitely many critical squares. ♠

Lemma 3.5 *S_2 is dense in S_1.*

Proof: Now let $\gamma \in S_1$ be arbitrary. We will make a small perturbation of γ so as to eliminate any given doubly critical rectangle without creating any new ones. Doing this finitely many times, we have a perturbation, as small as we like, which eliminates all of them.

Consider the rectangle case first. We use the notation from the previous lemma. We simply perturb the edges of γ so that no point of $\gamma \cap V$ has the same distance from the vertex v as a point of $\gamma \cap W$ has from the vertex w. Likewise, we perturb so that no edge of γ intersects the circle with diameter $\overline{vw}$ through v and w at two antipodal points. This eliminates a doubly critical rectangle that involves both v and w.

We first eliminate all the critical rectangles which involve more than one vertex of γ. Now we turn to any critical squares that might remain. Let v be the vertex of a critical square and let w be the opposite vertex. From the order in which we have done things, w lies in a unique edge e of γ. We perturb γ by moving e parallel to itself slightly. This destroys the critical square. ♠

3.3 Eliminating Bad Indices

Remark: This is the section where we use the technical material from §2.5.

Let $\gamma \in S_2$. By construction, γ has the property that its critical rectangles are non-square and only involve a single vertex. Let R be a labeled rectangle that graces γ and let I be a multi-index as above. We call R and I *associates* if the jth vertex of R lies on the edge e_{I_j} of γ. In this case, and only in this case, we mention the pair (I, R). Let a_0, a_1, a_2, a_3 be the vertices of R. Let λ_0 be the aspect ratio of R. We call (I, R) *bad* if $\lambda_0 \in \Sigma_0(L)$, the set from the last chapter. We call I *bad* if there is some critical rectangle R such that (I, R) is bad.

Let $S_3(I)$ denote the subset of polygons in S_2 with respect to which the index I is not bad.

Lemma 3.6 *$S_3(I)$ is open in S_2.*

Proof: Suppose that $\gamma_n \in S_2 - S_3(I)$ converges to $\gamma_\infty \in S_2$. Define

- R_n, the bad critical rectangle associated to I.
- L_n, the set of lines extending sides of γ_n associated to I.
- λ_n, the aspect ratio of R_n.
- s'_n, the derivative $ds/d\lambda$ computed with respect to L_n.

By Lemma 3.2 we can pass to a subsequence so that $R_n \to R_\infty$, a critical rectangle gracing γ_∞ that is associated to I. We can pass to a further subsequence so that everything else in sight converges in the same way. By continuity, $s'_\infty(\lambda_\infty) = 0$. Hence (I, R) is bad. Hence $\gamma_\infty \in S_2 - S_3(I)$. ♠

Lemma 3.7 *$S_3(I)$ is dense in S_2*

Proof: Take an arbitrary $\gamma \in S_2$ for which I is a bad index. By Lemma 2.3 there are only finitely many critical rectangles associated to I. These are obtained by setting $s(\lambda)$ to be equal to the first coordinate of one of the two endpoints of the relevant edge of γ and solving. We will explain how to perturb γ by an arbitrarily small amount so as to eliminate one of the bad pairs (I, R). Doing the same thing finitely many times, we eliminate all the bad pairs of the form (I, R).

Suppose (I, R) is a bad pair. Let a_0 be the vertex of R which is also a vertex of γ. Let $L = L_I$ be the corresponding set of lines. Let L^* and R^* be the perturbed line configuration and perturbed rectangle guaranteed by Lemma 2.8. Note that a_0 is still a vertex of R^* and the remaining vertices of R^* are interior to edges of γ provided that the perturbation is small enough. Also, the aspect ratio of R^* does not belong to the set $\Sigma_0(L^*)$.

Taking a small enough perturbation, the new lines we get are not parallel to any of the lines of γ that we have left alone. A polygon without any parallel consecutive sides is defined by the lines extending its edges. We let γ^* be the polygon whose associated lines agree with those of γ for any index not in I and agree with those of L^* for those indices in I.

By construction, R^* is critical for γ^* and not bad. So, we have eliminated the bad pair (I, R) using an arbitrarily small perturbation and we have not created any new bad pairs. ♠

Let $S_3 = \bigcap S(I)$, the intersection being taken over all relevant multi-indices. By the results above, S_3 is open dense in S_2.

3.4 The Space of Gracing Rectangles

We will establish Theorem 1.3 for polygons in S_3. Let $G(\gamma)$ denote the set of (counter-clockwise) labeled rectangles which grace $\gamma \in S^3$. We identify $G(\gamma)$ as a subset of $\boldsymbol{R}^2 \times \boldsymbol{R}^2 \times \boldsymbol{R}$ as follows. Each rectangle R corresponds to the data (a_0, a_1, λ). Here a_0 and a_1 are the first two vertices of R and λ is the aspect ratio.

Theorem 3.8 *For any polygon $\gamma \in S_3$, the space $G(\gamma)$ is a non-empty piecewise smooth 1-manifold. Each arc component of $G(\gamma)$ is proper.*

We prove this result through a series of smaller lemmas.

Lemma 3.9 *$G(\gamma)$ is a smooth 1-manifold in the neighborhood of any non-critical rectangle of $G(\gamma)$.*

Proof: Let R be such a rectangle. There is a unique multi-index I associated to I. Let $L = L_I$. Let λ_0 be the aspect ratio of I. Every rectangle in $G(\gamma)$

sufficiently close to R is associated uniquely to the same multi-index. Lemma 2.3 applied to L now tells us that the set of rectangles of $G(\gamma)$ sufficiently close to R is parametrized precisely by the curve

$$\lambda \to (s(\lambda), s(\lambda), \lambda), \qquad \lambda \in (\lambda_0 - \epsilon, \lambda_0 + \epsilon).$$

This gives a smooth and non-singular parametrization of $G(\gamma)$ in a neighborhood of R. The parametrization is non-singular because derivative of the last coordinate is always 1. ♠

Lemma 3.10 *$G(\gamma)$ is a piecewise smooth manifold in the neighborhood of any critical rectangle of $G(\gamma)$. The non-smooth point occurs precisely at this critical rectangle.*

Proof: Let R be a critical rectangle. We label R so that a_0 is the vertex of R that is also a vertex of γ. Let i_0 and j_0 be the indices of the two edges incident to a_0. There is a unique multi-index I associated to R which has i_0 as the first index and there is a unique multi-index associated to R having j_0 as the first index. The point is that the remaining vertices of R are comtained in the interiors of unique edges. Let $L_{I,0}$ be the line of L incident to a_0. Likewise define $L_{J,0}$.

We first consider the picture with respect to L_I. Let λ_0 be the aspect ratio of R. As we vary λ linearly away from λ_0, there is a family of rectangles $R_I(\lambda)$ such that the corresponding vertex $a_0(I, \lambda)$ varies monotonically away from a_0 in a non-singular way. Given the monotone motion, precisely one of two things is true for λ sufficiently close to λ_0:

- Negative case: $R_{I,\lambda} \in G(\gamma)$ if and only if $\lambda < \lambda_0$.
- Positive case: $R_{I,\lambda} \in G(\gamma)$ if and only if $\lambda > \lambda_0$.

In either case the subset of $G(\gamma)$ near R and associated to I is a smoothly parametrized half-open arc whose endpoint is R. The same goes for J. Finally, any member of $G(\gamma)$ sufficiently close to R is associated to one of I or J. Putting all this together, we see that in a neighborhood of R the space $G(\gamma)$ is the union of two smooth and half-open arcs which meet at their common closed endpoint, namely R. ♠

Remark: We want to be a bit more explicit about how our argument works with relabeling. The space $G(\gamma)$ really has 4 critical rectangles with the same unlabeled image as R. Each cyclic relabeling gives a smooth diffeomorphism from a neighborhood of R in the ambient space $\boldsymbol{R}^5$ to a relabeled version of R. Thus, our analysis above, for the specific choice of labelings, implies that $G(\gamma)$ looks the same up to ambient isomorphism around the 3 relabeled copies of R.

The last two results show that $G(\gamma)$ is a piecewise smooth 1-manifold provided that it is nonempty. We will deal with the non-emptiness at the very end. First we analyze the (hypothetical) components of $G(\gamma)$ that are arcs. As in the introduction, we call these *arc components*.

Lemma 3.11 *Every arc component of* $G(\gamma)$ *is proper.*

Proof: Let A be an arc component. Suppose that A is not proper. This would mean that there is a sequence of points $\{\zeta_n\} \in A$ that exits every compact subset of A but not does have aspect ratio tending to 0 or ∞. The only possibility is that the diameter of ζ_n tends to 0. Otherwise, we could pass to a subsequence and take a well-defined limiting rectangle which would belong to $G(\gamma)$. Note however that there is a lower bound to the diameter of any rectangle that graces γ because otherwise we could have all points of some rectangle on a union of two consecutive edges of γ. This is impossible. So, our sequence cannot shrink to a point. Hence A is proper. ♠

Let $d : \gamma \times \gamma \to \boldsymbol{R}_+$ denote the distance function. That is,

$$d(a, b) = \|a - b\|. \tag{6}$$

Say that a *stick* is a quadruple of points (a, a, b, b) or (a, b, b, a) where $a \neq b$ and (a, b) is a positively oriented critical point for d. Again, the positive orientaton means that γ intersects the line $\overline{ab}$ in opposite orientations at a and b. (See Figure 2 below.) The point (a, b) could either be a local max, a local min, or a saddle.

To make the next lemma precise we augment our space $\boldsymbol{R}^2 \times \boldsymbol{R}^2 \times \boldsymbol{R}$ so that the last coordinate can takes values in the compact arc $[0, \infty]$. That is, we compactify the last coordinate to include ∞.

Lemma 3.12 *$G(\gamma)$ contains rectangles arbitrarily close to any stick. More precisely, every stick is the end of a proper arc of $G(\gamma)$.*

Proof: This is a case-by-case analysis. There are 6 cases, as shown in Figure 2. In three of these cases, when one of the points of (a, b) lies in the interior of an edge of γ, the rectangles are easy to construct: They all have sides parallel to the stick.

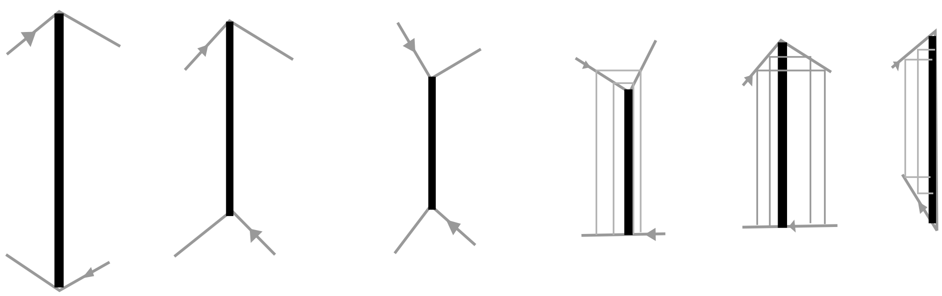

Figure 2: Rectangles near sticks.

The remaining cases are tricky. We will consider the first case shown in Figure 2. The remaining two cases have similar treatments. The crucial feature in all these examples is that the lines perpendicular to the stick, drawn at the endpoints of the stick, do not separate the edges of the polygon. Figure 3 shows this in more detail for the first example.

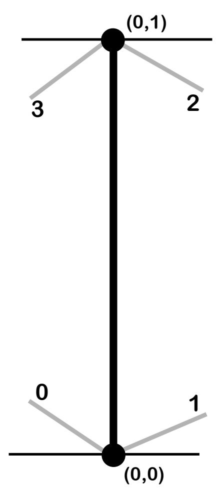

Figure 3: Rectangles near sticks.

We apply a similarity so that the endpoints of the stick are $(0,0)$ and $(1,0)$. Let L_0, L_1, L_2, L_3 be the lines shown. In terms of the data defining these lines we have $b_0 = b_1 = 0$ and $b_2 = b_3 = 1$. We also have $m_0, m_2 < 0$ and $m_1, m_3 > 0$. Referring to the equations in the previous chapter, we have

$$s(\lambda) = \frac{S_0 + S_1\lambda}{U_0 + U_1\lambda + U_2\lambda^2}, \qquad U_2 = (m_0 - m_1)(m_2 - m_3) \neq 0.$$

The coefficient S_2 vanishes. Hence, as $\lambda \to \infty$, the first coordinate of $a_0(\lambda)$ converges to 0. Hence $a_0(\lambda) \to (0,0)$ as $\lambda \to \infty$. The same kind of calculation shows that $a_1(\lambda) \to (0,0)$ as $\lambda \to \infty$. Cycling the indices by 2, we get the similar result that $a_2(\lambda) \to (0,1)$ and $a_3(\lambda) \to (0,1)$. This proves that the rectangle of aspect ratio λ converges to the stick as $\lambda \to \infty$. ♠

Lemma 3.13 *Every end of every proper arc of $G(\gamma)$ is a stick.*

Proof: Let A be such a proper arc. As we exit one end of A, the aspect ratio of the associated rectangles tends to either 0 or ∞. Hence, there must be some line segment which is a limit for these rectangles. Call this line segment σ. The perpendiculars to σ at the endpoints are tangent to γ in the following generalized sense. If the endpoint lies in the interior of an edge, then the perpendicular simply equals this edge. If the endpoint is a vertex, then the perpendicular does not separate the two edges incident to the vertex. This result follows essentially from the mean value theorem applied to the part of the polygon between the vertices bounding the short ends of the rectangles. The fact that the nearby rectangles are gracing forces σ to be positively oriented, and hence a stick.

There cannot be a second limiting stick because once the aspect is sufficiently large the rectangles are trapped in a vicinity of σ. They would have to open back up in order to get away from σ. ♠

Corollary 3.14 *For any $\gamma \in S_3$, the number of unlabeled proper arc components of $G(\gamma)$ equals half the number of unlabeled sticks.*

Remark: From the thousands of examples I computed, it seems that every arc component joins a stick which is a saddle to a stick which is a local maximum or a local minimum. I have no idea why.

4 The Number of Gracing Squares

4.1 A Parity Result

Here we establish Equation 1 for $\gamma \in S_3(N)$.

Theorem 4.1 $\Omega(\gamma) + \Omega'(\gamma) + \Omega^*(\gamma)$ *is even.*

Proof: We write

$$\Omega = \Omega_{\text{arc}} + \Omega_{\text{loop}}, \tag{7}$$

Where Ω_{arc} is the number of unlabeled squares on the union of the arc components and Ω_{loop} is the number of unlabeled squares on the union of loop components.

We first prove that Ω_{arc} and Ω' have the same parity. Here Ω' is the number of unlabeled sweepouts. It suffices to show that each arc component has an even number of squares if and only if it is a sweepout.

We say that the aspect ratio of a point of $G(\gamma)$ is the aspect ratio of the rectangle that this point represents. In particular, ζ has aspect ratio 1 iff ζ represents a square. Let $\zeta \in G(\gamma)$ such a point. Our analysis above shows that there is a small arc $\alpha \subset G(\gamma)$ containing ζ such α is parametrized monotonically by the aspect ratios of the corresponding points. When the aspect ratio of $\zeta' \in \alpha$ is slightly less than that of ζ, the point ζ' lies on one side of ζ. When the aspect ratio of ζ' is slightly greater than that of ζ the point ζ' is on the other side of ζ.

Consider the aspect ratio as a function of an arc component A. As we trace along A the aspect ratio passes through 1 an odd number of times if and only if the limiting aspect ratios at either end of A are different – i.e. if and only if A is a sweepout. Hence A is a sweepout if and only if it contains an odd number of points representing squares.

Now we prove that Ω_{loop} has the same parity as Ω^*, the number of cycling components. There are three cases to consider for the loop components. The same argument as above shows that that each loop component α has an even number of labeled squares on it. There are three cases to consider for the loop components. Let ϕ denote the operation of cyclically relabeling by one click, so to speak.

Case 1: Suppose that the orbit of α, namely $\{\phi^k(\alpha)\}$, consists of 4 loops, then the total number of labeled squares is a multiple of 8. Hence α contributes an even number fo $\Omega_{\rm loop}$.

Case 2: Suppose that the orbit consists of 2 loops. Then ϕ^2 fixes α and preserves the aspect ratios. The same argument as in the arc case shows that the quotient loop α/ϕ^2 has an even number of equivalence classes of labeled squares. But then α has a multiple of 4 labeled squares. Hence $\alpha \cup \phi(\alpha)$ contains a multiple of 8 labeled squares. Again, α contributes an even number to $\Omega_{\rm loop}$.

Case 3: Suppose that ϕ preserves α. Then α is a cycling component. α must contain a square because one can connect a rectangle R of aspect ratio less than 1 to the rectangle $\phi(R)$ of aspect ratio greater than 1. So, let Q be a square and let Q' be a rectangle which lies just a bit before Q in the cyclic order of Q'. When we go around α from Q to $\phi(Q)$ we reach $\phi(Q')$ just before reaching $\phi(Q)$. If the aspect ratio of Q' is less than 1 then the aspect ratio of $\phi(Q')$ is greater than 1. This means that we must have encountered an even number of labeled squares on the path between Q and $\phi(Q')$. Hence α contributes an odd number to $\Omega_{\rm loop}$.

These three cases establish the fact that $\Omega_{\rm loop}$ and Ω^* have the same parity. Hence Ω and $\Omega' + \Omega^*$ have the same parity. ♠

4.2 A Clean Case

The rest of the chapter is devoted to proving that $\Omega(\gamma)$ is odd when $\gamma \in S_3(N)$. We first consider a clean special case of the result. We fix some $N \geq 8$.

We call a multi-index $I = (i_1, i_2, i_3, i_3)$ *well separated* if consecutive indices are unequal and not cyclically adjacent. Let $S_3' \subset S_0$ denote the set of those polygons γ with the following conditions.

1. $\chi(L_I) \leq 1$ whenever I is well-separated.
2. A critical square of $G(\gamma)$ has just one vertex in common with γ.
3. $G(\gamma)$ has at most one critical square.

4. No bad pair of γ involves a critical square. See §3.3.

5. Any multi-index associated to a square in $G(\gamma)$ is well-separated.

We call a polygon *ordinary* if it has no critical squares.

Lemma 4.2 *Suppose that $\gamma_{-1}, \gamma_1 \in S_3'$ are ordinary polygons that are connected by a continuous path in S_3'. Then $\Omega(\gamma_{-1})$ and $\Omega(\gamma_1)$ have the same parity.*

Proof: Let $u \to \gamma_u$ for $u \in [-1, 1]$ be our path. We first prove the result when γ_u is ordinary for all $u \in [-1, 1]$.

Let Q_u be the square that graces γ_u. There is a unique multi-index I that is associated to U, and it is well separated. Let L_u be the configuration of lines associated to I. We rotate so that no lines of L_u are vertical. By hypothesis, $\chi(L_u) \leq 1$. Let $s_u(\lambda)$ and $t_u(\lambda)$ be the two functions from Lemma 2.3 that are associated to L_u. Given the existence of Q_u, we know that $s_u(1)$ and $t_u(t)$ specify the first coordinates of the two vertices $Q_{u,0}$ and $Q_{v,0}$ and these vertices lie on the interiors of the edges of γ_u.

Note that the functions s_u and t_u are continuous as functions of u and also the edges of γ_u are continuous as functions of u. That means that for v sufficiently near u, there is a gracing square Q_v near Q_u associated to I as well. Indeed, the map $v \to Q_v$ is continuous as a function of v. Hence there is a bijection between the squares gracing γ_u and the squares gracing γ_v provided that $|v - u|$ is small enough. Hence the squares of $G(\gamma_u)$ vary continuously with u and their number does not change. For later reference we call this the *continuity argument*.

Since $S_3'(N)$ is an open subset of $(\boldsymbol{R}^2)^N$, we can replace our path by one that is analytic. In this case, there are only finitely many non-ordinary points along the path. In light of the special case above, it suffices to prove the result for the restriction of our path to a tiny neighborhood of a non-ordinary point.

Composing with a continuous family of isometries, and cyclically relabeling, we can arrange hat $v_{1,0}$ is the critical vertex and, for all $u \in [-1, 1]$, we have

- $v_{1,u} = (0, 0)$.
- $e_{0,u}$ lies to the left of the y-axis.
- $e_{1,u}$ lies to the right of the y-axis.

There are exactly 2 associated multi-indices I_0 and I_1 such that $v_{1,u}$ lies on the first line of the associated configuration. Both of these are well-separated and the first line of $L_{I_j,u}$ extends $e_{j,u}$. We set $L_{j,u} = L_{I_j,u}$. Let $s_{j,u}$ denote the version of the function s with respect to the configuration $L_{j,u}$.

Shrinking the range of the path if necessary, using properties of S_3' together with compactness and continuity, and finally scaling, we can find some $\epsilon > 0$ so that

$$|s'_{j,u}(\lambda)| > 1, \quad \forall \lambda \in (1-\epsilon, 1+\epsilon), \quad s_{j,u}(1) \in (-\epsilon, +\epsilon) \qquad \forall\ (j,u). \qquad (8)$$

There are 4 cases, depending on the signs of $s'_{j,u}$ for $j = 0, 1$.

Suppose $u \neq 0$. By construction, there is some $\lambda_j \in (1-\epsilon, 1+\epsilon) - \{1\}$ such that $s_{j,u}(\lambda_j) = 0$. The key point is that $\lambda_1 = \lambda_2 = \lambda$, because these numbers both describe the aspect ratio of the same critical rectangle.

Suppose first that $\lambda < 1$. The functions $s_{j,u}$ are monotone throughout $(1-\epsilon, 1+\epsilon)$. Hence $s_{j,u}(1) > 0$ if and only if $s'_{j,i} > 0$. If $s_{0,u}(1) < 0$ then this point is the vertex of the unique square that graces the configuration $L_{0,u}$. If we pick ϵ small enough, then this square also graces γ_u. If $s_{0,u}(1) > 0$ then there is no such square. Similar statements hold for the index $j = 1$ with the signs reversed. We conclude from this that the total number of squares gracing γ_u that are associated to I_0 or I_1 has the opposite parity as the number of indices j for which $s'_{j,u}$ is greater than 0. The same final count works when $\lambda_u > 1$.

In short, the number of squares gracing γ_u that are associated to I_0 or I_1 is independent of u at least if we count mod 2. Also, the analysis for the remaining non-critical squares is as in the first special case we considered. So, the total count is the same mod 2 for all parameters. ♠

4.3 Connecting Polygons

We take $N \geq 8$. We will work our way up to finding paths in S_3'. Let S_1' denote the subset of S_0 consisting of those polygons γ such which satify Condition 1 from the previous section.

Lemma 4.3 *Let $\gamma \subset S_0(N)$ be a continuous path connecting two points in $S_1'(N)$. Then there is an arbitrarily small smooth perturbation γ' of γ which has the same endpoints and lies entirely in $S_1'(N)$.*

Proof: We can perturb γ so that it is smooth. Let X_0 and X_1 be the line configurations from Theorem 2.5. We say that a well-separated index I is *good* for γ if $L_I \in X_1$ whenever L_I is a configuration of lines associated to the index I with respect to a configuration along γ. Otherwise we say that I is *bad* for γ.

If I is good for γ then I is also good for any sufficiently small perturbation of γ. This follows from compactness and from the fact that X_1 is open in the set the fact that X_1 is open dense in X_0. So, we just have to show that there are arbitrarily small perturbations of γ with respect to which I is good. If we know this, then we can take an arbitrarily small perturbation of γ which decreases the number of bad indices until there are none.

The path γ determines a smooth path $\zeta \subset X_0$. We simply take the configurations L_I with respect to polygons along γ. Since $X_0 - X_1$ has codimension 2 in X_0, we can find a smooth path ζ' arbitrarily close to ζ which remains in X_1.

Say that a *pseudo-polygon* is a cyclically ordered collection of lines. A pseudo-polygon determines a polygon unless there is a pair of successive lines which are either parallel or agree. In this case, we call the pseudo-polygon *degenerate*. On the other hand, a polygon always determines a pseudo-polygon which may or may not be degenerate.

We define a path of γ' of pseudo-polygons as follows: For an index $i \not\in I$ we vary the lines according to what γ does. When $i \in I$ we vary the lines according to what ζ' does. There are two cases to consider. If all the polygons along γ give rise to non-degenerate pseudo-polygons, then by continuity and compactness we can arrange the same thing for γ'. In this case, γ' determines a polygon path with all the desired properties.

Suppose that the pseudo-polygon path defined by γ has some degenerate members. We can perturb γ so that there are only finitely many degenerate members and so that each degenerate member only involves a single pair of consecutive lines. By considering finitely many pieces of γ separately, we reduce to the case where γ just has a single degenerate member. We cyclically relabel so that e_0 and e_1 are the offending edges and v_1 is the vertex between them.

If I contains neither 0 nor 1 then we simply use the knowledge of the location of v_1 to reconstruct the polygon from the pseudo-polygon. In other words, we have no reason to forget the location of v_1 because the edges around it are doing the same thing with respect to γ' as they are with respect to γ.

Since I is well-separated, it cannot contain both 0 and 1. We will consider the case when $1 \in I$. In this case we modify the path ζ', just by translating the configurations slightly, so that the line corresponding to the index 1 always contains the vertex v_1 of the corresponding polygon along γ. Call the new path ζ'' we can arrange that $\zeta'' \in X_1$ and that ζ'' is as close as we like to ζ. We now use γ'' in place of γ'. We reconstruct a smooth family of polygons from the pseudo-polygons along ζ'', and the vertex v_1, just as in the previous case. The new path of polygons has all the desired properties. ♠

Let S_2' denote the subset of S_1' consisting of those polygons γ such which satify Conditions 1-4 from the previous section.

Lemma 4.4 *Let $\gamma \subset S_1'(N)$ be a continuous path connecting two points in $S_2'(N)$. Then there is an arbitrarily small smooth perturbation γ' of γ which has the same endpoints and lies entirely in $S_2'(N)$.*

Proof: Our argument is similar to what we do in the proof of Lemma 3.5. The set of polygons which satisfy Conditions 1-k is an open subset of $(\boldsymbol{R}^2)^N$ for any choice of k. So, it suffices to show that we can arrange Conditions 2,3,4 one at a time by perturbing a path that lies in $S_1'(N)$ and has endpoints in $S_2'(N)$.

Consider Condition 2. The set of polygons in the ambient space $S_0(N)$ which have a gracing square with more than one vertex has codimension at least 2 because we can find an entire plane worth of deformations which destroy the condition. We can move the polygon off the one vertex or the other independently. Since the set of polygons not satisfying Condition 2 has codimension at least 2, we can perturb our path so that it avoids this set. We make this perturbation.

Consider Condition 3. This is a also a codimension 2 condition. If $G(\gamma)$ has two critical gracing squares, then the two critical vertices are distinct and we can independently perturb so as to eliminate the one critical square or the other. Now we perturb our path to avoid this codimension 2 set. We make this perturbation as well.

Consider Condition 4. A polygon γ which does not satisfy Condition 4 has a critical square Q. Let I be the multi-index corresponding to Q as in Lemma 3.7. Let L_0 be the first line of the line configuration associated to v. There are two independent variations which destroy the bad critical square.

- We can vary v along the line L_0. By Lemma 2.3 this produces new polygons which do not have nearby critical squares. The nearby critical rectangles have aspect ratios unequal to 1.

- We can take the variation from Lemma 3.7. This variation does not move v at all, and results in a polygons which do not have bad critical rectangles at v.

This shows that the set of polygons which fail to satisfy Condition 4 has codimension 2 in the ambient space $S_0(N)$. So, we can perturb our path to avoid this set. Once we make this last perturbation, we are done.

4.4 A Subdivision Trick

In this section we start with a polygon $\gamma \in S_3(N)$ and we produce a new ordinary polygon $\gamma^* \in S_3'(N+K_2)$ with 3 properties:

- No side of γ has length greater than ϵ.

- γ^* and γ are within ϵ of each other with respect to the Hausdorff metric.

- $G(\gamma)$ and $G(\gamma^*)$ have the same number of gracing squares.

The process is canonical once we fix the data (K_1, K_2, ρ). Here K_1 and K_2 are positive integers and $\rho \in (0,1)$.

We say that *ρ-position* on an edge $\overline{v_1v_2}$ is the point $\rho v_1 + (1-\rho)v_2$. Here is our procedure. At the kth step we place new vertex in ρ-position on an edge. Next, we move this vertex outward, normal to edge k, a distance 2^{-K_1-k}. Finally, we move to the next edge of whatever polygon we have created. We do this for $K-2$ steps. We call this the *bending process*.

We also mention the *placebo process*, where we place the new vertices but not make any motions. When we do the placebo process, we get the same underlying polygon as γ, except that many extra vertices are inserted. If we hold K_2 and ρ fixed and let $K_1 \to \infty$, then the bent polygon converges to the placebo polygon.

We pick some large value of K_2 so that each edge gets subdivided many times. Next, we pick ρ so that none of the subdivision points in the placebo polygon is the vertex of a gracing square. The same principle as our continuity argument in Lemma 4.2, combined with induction on K_2, shows that once K_1 is large enough the polygons γ^* and γ have the same number of gracing squares. Moreover, there is a canonical bijection between these gracing squares, and each square in $G(\gamma^*)$ converges to the corresponding square in $G(\gamma)$ as $K_1 \to \infty$.

Suppose now that Γ is a continuous path of the form $u \to \gamma_u$. Here $\gamma_u \in S_0(N)$ for all $u \in [0,1]$. Once we fix the seed (K_1, K_2, ρ), we can make our construction simultaneously for all γ_u, If our original path is a path Γ belongs to $S_0(N)$ then our new path Γ^* belongs to $S_0(N+K_2)$. The new path Γ^* is given by $u \to \gamma_u^*$. We don't have much control over how the inserted points interact with the critical squares during the path, but we can control what happens at the endpoints, as above, if those endpoints belong to $S_3(N)$. This lack of control in the middle won't bother us because we will make a further perturbation.

For each path Γ we let $\square\Gamma$ denote the infimal side length of a square in $G(\gamma_u)$ for any $u \in [0,1]$. By compactness, $\square\Gamma > 0$.

Lemma 4.5 *Let Γ be some path. Then there is some $\epsilon_0 > 0$ such that $\square\Gamma'$ for any sufficiently small perturbation Γ' of Γ.*

Proof: Let γ' be a polygon in a small perturbation Γ' of Γ. By compactness, there is a uniform positive lower bound to the distance between any two non-adjacent edges of γ'. Therefore, any sufficiently small square gracing γ' would have to have all its vertices on two consecutive sides. This is impossible. ♠

4.5 The End of the Proof

In this section we prove that $\Omega(\gamma)$ is odd for any $\gamma \in S_3(N)$. Suppose that γ_0 and γ_1 are two elements of $S_3(N)$. We first connect γ_0 to γ_1 by a path $\Gamma \subset S_0(N)$. By Lemma 4.5, there is some $\epsilon_0 > 0$ so that any sufficiently small perturbation Γ' of Γ satisfies $\square\Gamma' > \epsilon_0$. We will take all our perturbations to have this property without saying so explicitly.

We first choose data (K_1, K_2, ρ) and consider the subdivided path

$$\Gamma^* \subset S_0(N+K_2).$$

If we pick this data appropriately, then we arrange that every polygon edge in sight has length less that $\epsilon_0/100$. We can arrange, moreover, that the endpoints of Γ^* are ordinary points of $S_3'(N+K_2)$. We explained this at the beginning of the last section.

By the results in the preceding section we can perturb Γ^* to a new path Γ^{**} with the following virtues.

- $\Gamma^{**} \subset S_2'$.
- No edge of any polygon along Γ^{**} has length greater than $\epsilon_0/10$.
- $\Box\Gamma^{**} > \epsilon_0$.

Under these conditions, any multi-index associated to any gracing square of any $G(\gamma_u^{**})$ is well separated. This shows that $\Gamma^{**} \subset S_3'$. Lemma 4.2 now shows that $\Omega(\gamma_0)$ and $\Omega(\gamma_1)$ have the same parity.

For any N, we can give examples of $\Gamma \in S_3(N)$ such that $\Omega(\gamma)$ is odd. We simply take an equilateral triangle and suitably subdivide it as above. This completes the proof that $\Omega(\gamma)$ is always odd.

The fact that $\Omega(\gamma)$ is odd combines with Equation 1 to establish Theorem 1.4.

5 The Trihotomy Theorem

5.1 The Circular Invariant

Let S^1 be the unit circle. Let $\Sigma \subset (S^1)^4$ denote the subset of distinct quadruples, which go counterclockwise around S^1. Any member $\sigma \in \Sigma$ corresponds to points which cut S^1 into 4 arcs A_0, A_1, A_2, A_3. Let $|A_j|$ be the arc length of A_j. We define the *circular invariant* of the points to be

$$\Lambda(\sigma) = \frac{|A_0| + |A_2|}{|A_1| + |A_3|}. \tag{9}$$

When σ consists of the vertices of a rectangle $\Lambda(\sigma)$ is the aspect ratio of this rectangle. Otherwise $\Lambda(\sigma)$ is only vaguely like an aspect ratio.

For each integer $K \geq 1$ we let $\Sigma(K)$ denote those members σ such that $\Lambda(\sigma) \in [K^{-1}, K]$. Note that $\Sigma(K)$ is not compact. For instance, we can let two successive arcs shrink to points while keeping the other two large. This situation cannot happen, however, when these quadruples are tied to gracing rectangles by a homeomorphism.

Suppose that $\phi : S^1 \to \gamma$ is some homeomorphism from S^1 to a Jordan loop. In [**Tv**], Tverberg gives a way to approximate γ by a sequence $\{\gamma_n\}$ of parametrized embedded polygons so that the parametrizations $\phi_n : S^1 \to \gamma_n$ converge uniformly to ϕ.

Lemma 5.1 *Fix some positive integer K. Let $\{\sigma_n\}$ be a sequence of elements of Σ such that $\phi_n(\sigma_n)$ is a rectangle gracing γ_n and $\Lambda(\sigma_n) \in \Sigma[K]$ for all n. Then the set $\{\sigma_n\}$ is precompact.*

Proof: It suffices to show that there is a uniformly positive distance between consecutive points of σ_n. We let $\sigma_{n,k}$ denote the kth point of σ_n. Let R_n be the rectangle whose vertices are $\phi_n(\sigma_n)$. Suppose that $|\sigma_{n,0} - \sigma_{n,1}| \to 0$.

We claim that the distance between $\sigma_{n,2}$ and $\sigma_{n,3}$ converges to 0 as well. Suppose not. Since $\phi_n \to \phi$ uniformly, the first side of R_n is converging to a point. However, the opposite side is not converging to a point. This is impossible for a rectangle. Hence $|\sigma_{n,2} - \sigma_{n-3}| \to 0$ as well. Given the bound on $\Lambda(\sigma_n)$ this can only happen all the associated arcs shrink to points. This is impossible because the sum of the arc lengths is 2π. ♠

5.2 Hausdorff Limits

Suppose that C is a compact metric space. Let X_C denote the set of closed subsets of X. We define the *Hausdorff distance* between closed subsets $A, B \subset C$ to be the infimal ϵ such that each of the two sets is contained in the ϵ-tubular neighborhood of the other one. This definition makes X_C into a compact metric space.

Lemma 5.2 *Suppose that a sequence $\{A_n\}$ of connected sets in C converges to a set A. Then A is connected.*

Proof: Suppose A is disconnected. Then there are 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 positive $\epsilon > 0$ such that

- All points of $A \cap U$ are at least ϵ from all points of $C - U$.
- All points of $A \cap V$ are at least ϵ from all points of $C - V$.

In particular, each point of A is at least ϵ from each points in $C - U - V$.

For all sufficiently large n, the arc A_n intersects both U and V and therefore contains a point $x_n \in C - U - V$. But then x_n is at least ϵ from A. This contradicts the fact that $A_n \to A$ in the Hausdorff metric. ♠

Note that Y need not be path connected. Consider a sequence of path approximations to the topologist's sine curve.

5.3 Limits of Sweepouts

let $\phi : S^1 \to \gamma$ be an arbitrary homeomorphism. Let $\phi_n : S^1 \to \gamma_n$ be as in the previous section. Slightly perturbing these polygons, we can assume that Theorems 1.3 and 1.4 hold for all n. In this section we will consider the case when γ_n has a sweepout for all n. In the next section we will consider the case when γ_n has an elliptic component for all n.

We call a subset $Y \subset \Sigma$ *extensive* if Y is connected and if $\Lambda(Y) = (0, \infty)$. That is, every possible circular invariant is achieved for points in Y.

Lemma 5.3 *There is an extensive subset $Y \subset \Sigma$ with the property that for each $\sigma \in Y$ the rectangle with vertices $\phi(\sigma)$ belongs to $G(\gamma)$.*

Proof: γ_n has a sweepout A_n. This sweepout defines a continuous path $Y_n \subset \Sigma$ such that $\Lambda(Y_n) = (0, \infty)$. Here Y_n is such that ϕ maps each member of Y_n to the vertices of a rectangle in the sweepout.

Even though $\Sigma(K)$ is not compact, Lemma 5.1 tells us that there is some compact subset $C(K) \subset \Sigma(K)$ such that $Y_n \cap \Sigma(K) \subset C(K)$ for all n. We can take $C(K)$ to be all the accumulation points of sequences in $Y_n \cap C(K)$.

For each n and for each K we can construct nested continuous paths

$$Y_n(n,1) \subset Y_n(n,1) \subset ... \subset Y_n(n,n). \tag{10}$$

such that the endpoints of $Y_n(n,K)$ have circular invariants $1/K$ and K and $Y_n(n,K) \subset C(K)$. Using the Cantor diagonal trick we can pass to a subsequence to that the Hausdorff limits

$$Y(K) = \lim_{n\to\infty} Y_n(n,K) \subset C(K) \tag{11}$$

exist simultaneously. By the lemma in the previous section, $Y(K)$ is connected for all K. By construction $Y(K) \subset Y(K+1)$ for all K. The nested union of connected sets is connected. Therefore $Y = \bigcup_K Y(K)$ is connected. By construction Y is extensive. Finally, $\phi(\sigma)$ is the set of vertices of a rectangle that graces γ, for every $\sigma \in Y$. ♠

5.4 Limits of Elliptic Components

Now we will suppose that the polygon γ_n is such that $G(\gamma_n)$ has an elliptic component ζ_n for all n.

Lemma 5.4 *Suppose there is a uniform positive lower bound to the side length of a rectangle in ζ_n. Then there is a connected subset $S \subset G(\gamma)$ consisting of rectangles having uniformly large side lengths, such that every point of γ is the vertex of some member of S and a member of S is square.*

Proof: Let $Y_n \subset \Sigma$ be the subset corresponding to ζ_n. By hypotheses there is a single compact subset $K \subset \Sigma$ such that $Y_n \subset K$ for all n. Passing to a subsequence, take the Hausdorff limit $Y = \lim Y_n$. The set Y is connected. We let $S = \phi(Y)$. By construction, $S \subset G(\gamma)$.

Every vertex of γ_n is the vertex of a rectangle corresponding to ζ_n. Let v be a point of γ. Let $v_n \in \gamma_n$ be a point that converges to v. Let R_n be one of the rectangles in ζ_n which has v_n as vertex. Given the diameter bound, the limit $R = \lim R_n$ exists and belongs to S.

Since ζ_n has a square for all n, we can take the limit of these squares to get a point in S corresponding to a square that graces S. ♠

The preceding result peels off the elliptic case of the Trichotomy Theorem. Let ϵ_n denote the infimal side length of a rectangle corresponding to ζ_n. From now on we suppose that $\epsilon_n \to 0$.

Lemma 5.5 *There is an extensive subset $Y \subset \Sigma$ with the property that for each $\sigma \in Y$ the rectangle with vertices $\phi(\sigma)$ belongs to $G(\gamma)$.*

Proof: Let $Y_n \subset \Sigma$ be the set which corresponds to ζ_n. Since there is no lower bound on the side length of rectangles in ζ_n, it must be the case that there are points in ζ_n whose circular invariant tends to 0. But then, given the invariance under relabeling, we see that Y_n also contains points whose circular invariant tends to ∞. But then the sets $\{Y_n\}$ behave exactly as they to in the sweepout case. The rest of the proof is the same as in Lemma 5.3. ♠

5.5 The Hyperbolic Case

We work with the set Y constructed in Lemma 5.3 or Lemma 5.5. Suppose there is a positive lower bound to the diameter of members of Y.

Lemma 5.6 *$G(\gamma)$ has a connected subset which contains points representing rectangles of every aspect ratio.*

Proof: The set $S = \phi(Y)$ is a connected subset of $G(\gamma)$. The set of aspect ratios achieved by members of S is connected because S is connected. We just have to show that S contains members having aspect ratio arbitrarily close to 0 and aspect ratio arbitrarily close to ∞. If $\sigma \in \Sigma$ is a set with large diameter and very small circular invariant, then two consecutive points of σ are very close together and the adjacent two consecutive points are not. But then the corresponding rectangle with vertices $\phi(\sigma)$ must have a very short side and also a long side. Keeping track of the labeling, we see that the aspect ratio of R is close to 0. Hence there are rectangles corresponding to points in S having aspect ratio arbitrarily close to 0. Likewise for ∞. ♠

Let $Z \subset S^1$ denote those points which are not members of Y. Each point of $\phi(S^1 - Z)$ is the vertex of a rectangle that graces γ.

Lemma 5.7 *Z contains at most 4 points.*

Suppose Z has cardinality at least 5. We choose any 5 points of Z and let J_0, J_1, J_2, J_3, J_4 be the complementary intervals. We label so that these intervals go around counter-clockwise. Because Y is connected, the kth vertex of any $\sigma \in Y$ is trapped in some fixed interval J_{j_k} for all $\sigma \in Y$. We label so that $i_0 = 0$. Given our counter-clockwise ordering, we have $i_0 \leq i_1 \leq i_2 \leq i_3$.

Given that Y members where the points a_0 and a_1 are arbtrarily close together, we must have $i_1 \leq i_0 + 1$. The same goes for the points a_1 and a_2. This forces $i_2 \leq i_1 + 1$. Likewise $i_3 \leq i_2 + 1$. But then the arc J_4 is empty, which prevents σ_3 and σ_0 from bounding the very small circular arc just mentioned. This is a contradiction. ♠

5.6 The Parabolic Case

Finally, we suppose that there is no uniform lower bound to the diameter of subsets of Y.

Lemma 5.8 *$G(\gamma)$ has a connected set S which contains gracing rectangles of arbitrarily small diameter. All but at most 2 points of γ are vertices of rectangles corresponding to points in S*

Proof: Let $S = \phi(Y)$, as usual. The first statement follows immediately from the fact that Y is connected and contains sets of arbitrarily small diameter.

For the second statement we use the same notation as in Lemma 5.7. This time we have arcs J_0, J_1, J_2. Let $\sigma \in Y$ be a set with very small diameter. Then one of the arcs of S^1 must be nearly all of S^1 and all the points are bunched together. Call the index of the special arc *special*. Let's say that the special index is 3. Then we must have $i_0, i_1, i_2, i_3 \in \{0, 1\}$. The arc J_2 is empty.

On the other hand, given that the circular invariant ranges all the way from 0 to ∞ there are other points of Y in which the arc between a_3 and a_0 is arbitrarily small. But J_2 separates these points. This is a contradiction.

We get similar contradictions if we suppose that a different index is special. ♠

We have considered all the cases. This completes the proof of the Trichotomy Theorem.

5.7 Non-Atomic Measures

Here we explain remark (iv) made just after the statement of our Main Theorem. Suppose that μ is a non-atomic probability measure on γ. Then we can parametrize γ continuously so that μ is the pushforward of linear measure on the circle. In this case, the quantities in Equation 9 are just the μ-measures of the respective arcs.

In the case of the elliptic component, for every rectangle $\sigma' \in S$ with $\lambda(\sigma') = \lambda_0$ there is another (relabeled) rectangle $\sigma'' \in S$ with $\lambda(\sigma'') = 1/\lambda_0$. Hence, by connectedness, there is some $\sigma \in \Sigma$ with $\lambda(\sigma) = 1$. In the remaining cases, we have $\lambda(S) = (0, \infty)$ and so we have $\lambda(\sigma) = 1$ for some $\sigma \in S$.

6 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)

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

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

[**Tv**], H. Tverberg, *A Proof of the Jordan Curve Theorem*, Bulletin of the London Math Society, 1980, pp 34-38.

[**W**] S. Wolfram, *The Mathematica Book*, 4th ed. Wolfram Media/Cambridge University Press, Champaign/Cambridge (1999)