The maximum number of systoles for genus two Riemann surfaces with abelian differentials

Chris Judge[†] and Hugo Parlier[*]

Abstract. This article explores the length and number of systoles associated to holomorphic 1-forms on surfaces. In particular, we show that up to homotopy, there are at most 10 systolic loops on such a genus two surface and that the bound is realized by a unique translation surface up to homothety. We also provide sharp upper bounds on the the number of homotopy classes of systoles for a holomorphic 1-form with a single zero in terms of the genus.

1. Introduction

The *systolic length* of a length space (X, d) is the infimum of the lengths of non-contractible loops in X. If a non-contractible loop γ achieves this infimum, then we will call γ a *systole*. The systolic length and systoles have received a great deal of attention beginning with work of Loewner who is credited [Pu] with proving that among unit area Riemannian surfaces of genus one, the unit area hexagonal torus has the largest systolic length, $2/\sqrt{3}$, and is the unique such surface that achieves this value.

The hexagonal torus has another extremal property: Among all Riemannian surfaces of genus one, it has the maximum number of distinct homotopy classes of systoles, three. With respect to this property, the hexagonal torus is not the unique extremal among all genus one Riemannian surfaces, but it is the unique extremal among quotients of $\mathbb{C}$ by lattices Λ equipped with the metric $|dz|^2$.

The form dz on $\mathbb{C}/\Lambda$ is an example of a holomorphic 1-form on a Riemann surface. More generally, given a holomorphic 1-form ω on a Riemann surface X, one integrates $|\omega|$ over arcs to obtain a length metric d_ω on X. On the complement of the zero set of ω the metric is locally Euclidean, and each zero of order n is a conical singularity with angle $2\pi \cdot (n+1)$.

The length space (X, d_ω) determined by (X, ω) is the basic object of study in the burgeoning field of Teichmüller dynamics. See, for example, the recent surveys of [Forni-Matheus] and

[†] Research partially supported by a Simons collaboration grant.
[*] Research partially supported by Swiss National Science Foundation grant number PP00P2_153024.

2010 Mathematics Subject Classification: Primary: 32G15. Secondary: 30F10, 53C22.
Key words and phrases: systoles, translation surfaces, abelian differentials.

[Wright].

In this paper we prove the following.

Theorem 1.1. *Let ω be a holomorphic 1-form on a closed Riemann surface X of genus two. The number of distinct homotopy classes of systolic minimizers on (X, d_ω) is at most 10. Moreover, up to homothety, there a unique metric space of the form (X, d_ω) for which there exist exactly 10 distinct homotopy classes of systoles.*

In other words, among the unit area surfaces (X, d_ω) of genus two, there exists a unique surface $(X_{10}, d_{\omega_{10}})$ that has the maximum number of systolic homotopy classes. The surface obtained by multiplying this metric by $4\sqrt{3}$ is described in Figure 1. The surface $(X_{10}, d_{\omega_{10}})$ has two conical singularities each of angle 4π corresponding to the vertices of the polygon pictured in Figure 1. In other words, the 1-form ω_{10} has simple zeros corresponding to these vertices. Four of the ten systolic homotopy classes consists of geodesics that lie in an embedded Euclidean cylinders. Each of the other six systolic homotopy classes has a unique geodesic representative that necessarily passes through one of the two zeros of ω_{10}. It is interesting to note that some of the latter systoles intersect twice. Both intersections necessarily occur at zeros of ω_{10}. Indeed, if two curves intersect twice and one of the intersection points is a smooth point of the Riemannian metric, then a standard perturbation argument produces a curve of shorter length.

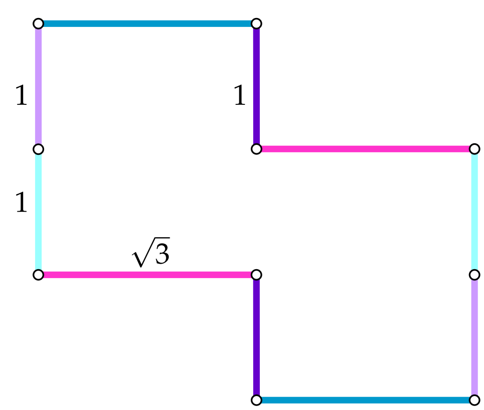

Figure 1: A pair (X, ω) that has ten systoles: By identifying parallel sides of the same color, we obtain a Riemann surface X. The one form dz in the plane defines a holomorphic 1-form on X.

Perhaps surprisingly, $(X_{10}, d_{\omega_{10}})$ does not maximize the systolic length among all unit area, genus two surfaces of the form (X, d_ω). To discuss this, it will be convenient to introduce the *systolic ratio*: the square of the systolic length divided by the area of the surface. A surface maximizes the systolic length among unit area surfaces if and only if it maximizes systolic ratio among all surfaces.

A genus two surface (X, d_ω) that has ten systoles has systolic ratio equal to $1/\sqrt{3} = .57735\ldots$. On the other hand, the surface described in Figure 2 has systolic ratio equal to

$$\frac{2 \cdot \left(\sqrt{13}-3\right)^2}{\sqrt{3} \cdot (1 - \frac{3}{4}(\sqrt{13}-3)^2)} = .58404\ldots \tag{1}$$

We believe that this surface has maximal systolic ratio.

Conjecture 1.2. *The supremum of the systolic ratio over surfaces (X, d_ω) of genus two equals the constant in (1). Moreover, up to homothety, the surface described in Figure 2 is the unique surface that achieves this systolic ratio.*

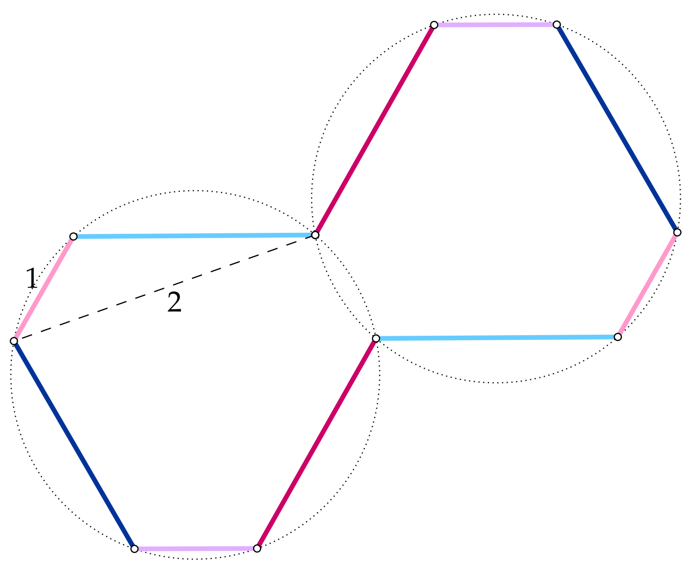

Figure 2: A surface (X, d_ω) whose systolic ratio equals the constant in (1). The surface is obtained from gluing parallel sides of two isometric cyclic hexagons in $\mathbb{C}$. Each hexagon has a rotational symmetry of order 3. The 1-form ω corresponds to dz in the plane.

By the Riemann-Roch theorem, the total number of zeros, including multiplicities, of a holomorphic 1-form on a Riemann surface of genus g equals $2g-2$. In particular, a 1-form ω on a genus two Riemann surface X consists of either two simple zeros or one double zero. Thus, we have a partition of the moduli space of pairs (X, ω) into the stratum, $\mathcal{H}(1,1)$, of those for which d_ω has two conical singularities of angle 4π and the complementary stratum, $\mathcal{H}(2)$, those for which d_ω has a single conical singularity of angle 6π.

In order to prove Theorem 1.1, we study both strata separately, and it turns out that the stratum $H(2)$ is considerably easier to analyse. Indeed, for $\mathcal{H}(2)$ we are able to prove sharp bounds on both the systolic ratio and on the number of systolic homotopy classes. We will say that a surface (X, d_ω) is *tiled by an equilateral triangle* T if there exists a triangulation of X such that each triangle is isometric to T and each vertex is a zero of ω.

Theorem 1.3. *If $(X, \omega) \in H(2)$, then (X, d_ω) has at most 7 homotopy classes of systoles, and the systolic ratio of (X, d_ω) is at most $2/(3\sqrt{3}) = .3849\ldots$ Furthermore, either inequality is an equality if and only if (X, d_ω) is tiled by an equilateral triangle.*

The unique surface that attains both optimal bounds is illustrated in Figure 3.

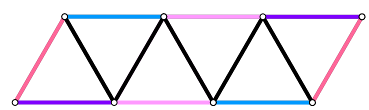

Figure 3: The genus two surface with a single cone point and tiled by equilateral triangles

To obtain the maximum systolic ratio, we adapt an argument of Fejes Tóth that he used to prove that the hexagonal lattice Λ_0 gives the optimal packing [Fejes Tóth]. This method applies to $(X, \omega) \in \mathcal{H}(2g-2)$, the space of holomorphic 1-forms on a genus g surface that have a single zero. The more general result about the systolic ratio of surfaces in these strata is the following.

Theorem 1.4. *The supremum of the systolic ratio of (X, d_ω) as (X, ω) varies over $\mathcal{H}(2g-2)$ equals $\frac{2\sqrt{3}}{9g-9}$. This supremum is achieved if and only if the surface is obtained by gluing equilateral triangles.*

Note that a more general (but non-optimal) bound for the systole was identified by Smillie and Weiss [Smillie-Weiss].

We are also able to identify optimal bounds for the number of homotopy classes of systoles of surfaces in $\mathcal{H}(2g-2)$. We also show that the optimal bounds cannot be attained by hyperelliptic surfaces in these strata. A condensed version of our results is the following (Proposition 3.1 and Theorem 3.3):

Theorem 1.5. *If ω be a holomorphic 1-form on X that has exactly one zero, then (X, d_ω) has at most $6g-3$ homotopy classes of systoles. If in addition ω is hyperelliptic, then (X, d_ω) has at most $6g-5$ homotopy classes of systoles. Both bounds are sharp.*

Note that as all genus two surfaces are hyperelliptic, this implies the bound on the number of homotopy classes in Theorem 1.3 above.

Although these questions have not been studied much in the context of translation surfaces, they have been studied in the context of hyperbolic and Riemannian surfaces. As hinted at above, smooth surfaces have systoles that intersect at most once, and from this one can deduce that there are at most 12 homotopy classes of systole in genus two (see for instance [Malestein-Rivin-Theran]). This bound is sharp. Indeed, among hyperbolic surfaces of genus two, there is a unique surface, called the Bolza surface, with exactly 12 systoles. It can be obtained by gluing opposite edges of a regular hyperbolic octagon with all angles $\frac{\pi}{4}$. This same surface is also optimal (again among hyperbolic surfaces) for systolic ratio, a result of Jenni [Jenni]. Either optimal quantities are unknown in higher genus although there are bounds. Interestingly, Katz and Sabourau [Katz-Sabourau] showed that among CAT(0) genus two surfaces, the optimal surface is an explicit flat surface with cone point singularities, conformally equivalent to the Bolza surface. This singular surface cannot be optimal among all Riemannian surfaces however, as by a result of Sabourau, the optimal surface in genus two necessarily has a region with positive curvature [Sabourau]. The optimal systolic ratio among all Riemannian surfaces is still not known.

Acknowledgements.

We thank Marston Condor for the examples in Remark 3.2. We thank Carlos Matheus Santos for kindly pointing out some mistakes and missing references in the first version of this paper. We also acknowledge support from U.S. National Science Foundation grants DMS 1107452, 1107263, 1107367 RNMS: Geometric structures And Representation varieties (the GEAR Network).

2. Facts concerning the geometry of (X, d_ω)

We collect here some relevant facts about the geometry of the surface (X, d_ω) sometimes called a 'translation surface'. Much of this material can be found in, for example, [Masur-Smillie], [Gutkin-Judge], and [Broughton-Judge].

2.1. Integrating the 1-form

By integrating the 1-form ω along a piecewise differentiable path $\alpha : [a, b] \to X$, we obtain a path in $\overline{\alpha} : [a, b] \to \mathbb{C}$ defined by

$$\overline{\alpha}(t) \;=\; \int_{\alpha|_{[a,t]}} \omega. \tag{2}$$

Since ω is closed, if two paths α, β in X are homotopic rel endpoints, then $\overline{\alpha}$ and $\overline{\beta}$ are homotopic rel endpoints. Thus, if $U \subset X$ is simply connected neighborhood of a point x, then

$$\mu_{x,U}(y) \;:=\; \int_{\alpha_y} \omega \tag{3}$$

is independent of the path α_y joining x to y. Note that $\mu_{x,U}$ is a holomorphic map from U into $\mathbb{C}$. If x is not a zero of ω, then it follows from the inverse function theorem that there exists a neighborhood U so that $\mu_{x,U}$ is a biholomorphism onto its image.

2.2. The metric

The norm, $|\omega|$, of ω defines an arc length element on X. We will let $\ell_\omega(\alpha)$ denote the length of a path on X, and we will let d_ω denote the metric obtained by taking the infimum of lengths of paths joining two points.

If x is not a zero of ω and U is a simply connected neighborhood of x, then $\mu_{x,U}$ is a local isometry from U into $\mathbb{C}$ equipped with its usual Euclidean metric $|dz|^2$. If, in addition, U is star convex at x, then $\mu_{x,U}$ is an isometry onto its image.

If x is a zero of ω of order k, then there exists a neighborhood V of x and a chart $\nu : V \to \mathbb{C}$ such that $\omega = (k+1)\cdot \nu^*(z^k dz)) = \nu^*(d(z^{k+1}))$ and $\nu(x) = 0$. If V is sufficently small, the map ν is an isometry from (V, d_ω) to $(\nu(V), d_{d(z^{k+1})})$. In turn, the map $z \mapsto z^{k+1}$ is a local isometry from $(\nu(V) - \{0\}, d_{d(z^{k+1})})$ to a neighborhood of the origin with the Euclidean metric $|dz|^2$. Since the branched covering $z \mapsto z^{k+1}$ has degree $k+1$, the arc length of the boundary of an ϵ-neighborhood of x is $2\pi(k+1)\cdot\epsilon$. Therefore, we refer to x as a *cone point* of angle $2\pi(k+1)$. Thus, the set of zeros of ω, denoted Z_ω, will be regarded as the set of cone points of (X, d_ω).

2.3. Universal cover, developing map and holonomy

Let $p : \widetilde{X} \to X$ be the universal covering map, and let $\widetilde{\omega} = p^*(\omega)$. If we let $d_{\widetilde{\omega}}$ be the associated metric on $\widetilde{X}$, then p is a local isometry from $(\widetilde{X}, d_{\widetilde{\omega}})$ onto (X, d_ω). Since $\widetilde{X}$ is simply connected, we may fix $\widetilde{x}_0 \in \widetilde{X}$ and integrate $\widetilde{\omega}$ as in (3) to obtain a map $\mathrm{dev} : \widetilde{X} \to \mathbb{C}$ called the *developing map*. The restriction of dev to $\widetilde{X} - Z_{\widetilde{\omega}}$ is a local biholomorphism and a local isometry. Each zero of $\widetilde{\omega}$ is a branch point whose degree equals the order of the zero. If C is the closure of a convex subset of $\widetilde{X} - Z_{\widetilde{\omega}}$, then the restriction of dev to C is injective.

Let $x_0 = p(\widetilde{x}_0)$, and consider loops α in x based at x_0. The assignment $\alpha \to \overline{\alpha}$ defines a homomorphism from $\pi_1(X, x_0)$ to the additive group $\mathbb{C}$. Moreover, for each $[\alpha] \in \pi_1(X, x_0)$ and $\widetilde{x} \in \widetilde{X}$ we have

$$\mathrm{dev}([\alpha]\cdot\widetilde{x}) \;=\; \mathrm{dev}(\widetilde{x}) \;+\; \mathrm{hol}([\alpha]) \tag{4}$$

where $\alpha \cdot \widetilde{x}$ denotes action by covering transformations and $\mathrm{hol}([\alpha])$ is the holonomy of $[\alpha]$.

2.4. Geodesics

If a geodesic γ on (X, d_ω) passes through a zero of ω, then γ will be called *indirect* and otherwise *direct*. If γ is a direct simple geodesic loop, then, since Z_ω is finite, for sufficiently small $\epsilon > 0$, the ϵ-tubular neighborhood, N, of γ is disjoint from Z_ω. Each lift $\widetilde{N} \subset \widetilde{X}$ of N is convex and hence the restriction of the developing map to $\widetilde{N}$ is an isometry onto $\mathrm{dev}(\widetilde{N})$. Since $\widetilde{N}$ is stabilized by the cyclic subgroup $\langle\gamma\rangle$ of the deck group generated by γ, it follows from (4) that $\mathrm{dev}(\widetilde{N})$ is the convex hull of two parallel lines, and, moreover, the map dev determines an isometry from N to $\mathrm{dev}(\widetilde{N})/\langle\mathrm{hol}(\gamma)\rangle$. In particular, N is isometric to a Euclidean cylinder $[0, w] \times \mathbb{R}/\ell\mathbb{Z}$ where $\ell = |\mathrm{hol}(\gamma)|$ and w is the distance between the parallel lines. If $Z_\omega \neq 0$, then the union of all Euclidean cylinders embedded in $X - Z_\omega$ that contain γ is a cylinder called the *maximal cylinder* associated to γ. Each component of the frontier of a maximal cylinder consists of finitely many indirect geodesics.

Proposition 2.1. *If ω has at least one zero, then each homotopy class of loops is represented by a geodesic loop that passes through a zero of ω.*

Proof. Since X is compact, a homotopy class of simple loops has a geodesic representative γ. If γ does not pass through a zero, then γ lies in a maximal cylinder. The boundary of the maximal cylinder contains a geodesic representative that passes through a zero. $\square$

Proposition 2.2. *If two simple geodesic loops are homotopic, then they lie in the closure of the same maximal cylinder.*

Proof. Because the angle at each cone point $\widetilde{z} \in Z_{\widetilde{\omega}}$ is greater than 2π, the length space $(\widetilde{X}, d_{\widetilde{\omega}})$ is CAT(0). If two geodesic loops γ and γ' are homotopic, then they have have lifts that are asymptotic in $(\widetilde{X}, d_{\widetilde{\omega}})$. By the flat strip theorem [Bridson-Haefliger], the convex hull of the two lifts is isometric to a strip $[0, w] \times \mathbb{R}$. Thus, since each cone point has angle larger than 2π, the interior I of the convex hull contains no cone points. The developing map restricted to I is an isometry onto a strip in $\mathbb{C}$, and, moreover, it induces an isometry from $I/\langle g\rangle$ to the cylinder $\mathrm{dev}(I)/\langle\mathrm{hol}(g)\rangle$ where g is the deck transformation associated to the common homotopy class of γ and γ'. Since the lifts are boundary components of I, the loops γ and γ' lie in the boundary of the cylinder $\mathrm{dev}(I)/\langle\mathrm{hol}(g)\rangle$. $\square$

2.5. The Delaunay cell decomposition

The Delaunay decomposition is well-known in the context of complete constant curvature geometries. Thurston observed that the construction also applies to constant curvature metrics with conical singularities [Thurston].

We will first describe the Delaunay decomposition of the universal cover $\widetilde{X}$. Given $\widetilde{x} \in \widetilde{X} - Z_{\widetilde{\omega}}$, let $D_{\widetilde{x}}$ be the largest open disk centered at $\widetilde{x}$ that does not intersect $Z_{\widetilde{\omega}}$. Since $D_{\widetilde{x}}$ is convex, the restriction of dev to the closure $\overline{D}_{\widetilde{x}}$ is an isometry onto a closed Euclidean disk in $\mathbb{C}$. Since $Z_{\widetilde{\omega}}$ is discrete, the intersection $Z_{\widetilde{\omega}} \cap \overline{D}_{\widetilde{x}}$ is finite. Let $\mathcal{V}$ be the set of $\widetilde{x} \in \widetilde{X} - Z_{\widetilde{\omega}}$ such that $Z_{\widetilde{\omega}} \cap \overline{D}_{\widetilde{x}}$ contains at least three points. Because three points determine a circle, the set $\mathcal{V}$ is discrete.

For each $\widetilde{x} \in \mathcal{V}$, let $P_{\widetilde{x}}$ denote the convex hull of $Z_{\widetilde{\omega}} \cap \overline{D}_{\widetilde{x}}$. It is isometric to a convex polygon in the plane. Again, because three points determine a circle, if $\widetilde{x}, \widetilde{y} \in \mathcal{V}$ and $\widetilde{x} \neq \widetilde{y}$, then the set $Z_{\widetilde{\omega}} \cap \overline{D}_{\widetilde{x}} \cap \overline{D}_{\widetilde{y}}$ consists of at most two points, and hence $P_{\widetilde{x}} \cap P_{\widetilde{y}}$ is either empty, a point, or a geodesic arc lying in both the boundary of $P_{\widetilde{x}}$ and the boundary of $P_{\widetilde{y}}$. The interior of $P_{\widetilde{x}}$ is called a *Delaunay 2-cell* and the boundary edges are called *Delaunay edges*. The vertex set of this decomposition of $\widetilde{X}$ is the set of zeros of $\widetilde{\omega}$.

The deck group of the universal covering map p permutes the cells of the Delaunay decomposition, and so we obtain a decomposition of X. Note the restriction of p to each 2-cell P is an isometry onto its image. Indeed, if not then there exists a covering transformation γ, a lift $\widetilde{P}$ of P, and $\widetilde{x} \in \widetilde{P}$ such that $\gamma \cdot \widetilde{x} \in \widetilde{P}$. Since $\widetilde{P}$ is convex, it follows that for some vertex of $\widetilde{z} \in \widetilde{P}$, we would have $\gamma \cdot \widetilde{z} \in \widetilde{P}$. But γ maps $Z_{\widetilde{\omega}}$ to itself.

Our interest in the Delaunay decomposition stems from the following.

Proposition 2.3. *If α is a shortest non-null homotopic arc with both endpoints in Z_ω, then α is a Delaunay edge.*

Proof. Since the universal covering map p preserves the length of arcs, it suffices to prove that the analogous statement holds for the universal cover $\widetilde{X}$. Because α is a shortest arc, if m is the midpoint of α, then the largest disc D centered at m has diameter equal to $\ell(\alpha)$ and $\overline{D} \cap Z_{\widetilde{\omega}}$ consists of exactly two points, the endpoints z and z' of α. The circle $\text{dev}(\partial D)$ belongs to the pencil of circles containing $\text{dev}(z)$ and $\text{dev}(z')$. Since X is compact, by varying over this pencil, we find a disk D' so that $\overline{D}' \cap Z_{\widetilde{\omega}}$ contains z, z', and at least one other point. The center c of D' belongs to $\mathcal{V}$ and α is a boundary edge of the polygon P_c. $\square$

Proposition 2.4. *Let ω be a holomorphic 1-form on a closed surface of genus g. If ω has v zeros, then the Delaunay decomposition of X has at most $6g - 6 + 3 \cdot v$ edges and the number of 2-cells is $4g - 4 + 2 \cdot v$. Equality holds if and only if each 2-cell is a triangle.*

Proof. By dividing the Delaunay 2-cells (convex polygons) into triangles, we obtain a triangulation with v vertices. By Euler's formula and the fact that there are 3 oriented edges for each triangle, we find that each triangulation has $6g - 6 + 3v$ edges and $4g - 4 + 2 \cdot v$ triangles. $\square$

3. Systoles of 1-forms in $\mathcal{H}(2g-2)$

In this section, we consider holomorphic 1-forms with a single zero. In the first part of the section we give the optimal bound on the number of homotopy classes of systoles of such surfaces as well as the optimal bound for the hyperelliptic surfaces with one zero. In the second part, we provide the optimal estimate on the systolic ratio of such surfaces.

3.1. Bounds on the number of systoles

Proposition 3.1. *If ω be a holomorphic 1-form on X that has exactly one zero, then (X, d_ω) has at most $6g - 3$ homotopy classes of systoles.*

Proof. By Proposition 2.1, each homotopy class of systoles contains a representative that passes through the zero. Proposition 2.3 implies that each such systole is a Delaunay edge. By Proposition 2.4, there are at most $6g - 3$ Delaunay edges and hence at most $6g - 3$ homotopy classes of systoles. $\square$

The bound in Proposition 3.1 is sharp if the genus g of X is at least 3. For example, if $g = 3, 4, 5$, then consider the surfaces described in Figures 4, 5, and 6. Moreover, given a holomorphic 1-form ω_g on a surface X_g of genus g with one zero that achieves the bound $6g - 3$, one can construct a holomorphic 1-form ω_{g+3} with one zero on a surface X_{g+3} of genus $g + 3$ that achieves the bound $6(g + 3) - 3$. Indeed, remove a Delaunay edge from (X_g, d_{ω_g}) to obtain a surface X'_g with one boundary component that consists of two segments F and F', and then glue the surface described in Figure 7.

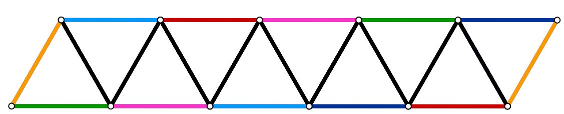

Figure 4: Glue the edges of the polygon according to the colors to obtain the Delaunay triangulation associated a holomorphic 1-form on a surface of genus three. Each edge is a systole, the 1-form ω has exactly one zero, and no two Delaunay edges are homotopic.

Remark 3.2. The problem of constructing surfaces that saturate the bound in Proposition 3.1 is equivalent to the problem of constructing two fixed-point free elements σ, τ in the symmetric group $S_{g-1} = \mathrm{Sym}(\{1, \ldots, 2g - 1\})$ such that $\sigma \cdot \tau$ has no fixed points and the commutator $[\sigma, \tau]$ is a $(2g - 1)$-cycle. Indeed, let $P_1, \ldots P_g$ be $2g - 1$ disjoint copies of the

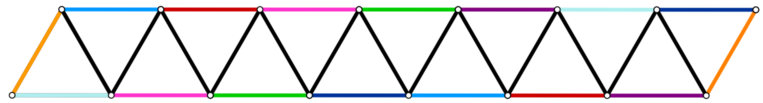

Figure 5: Glue the edges of the polygon according to the colors to obtain the Delaunay triangulation associated a holomorphic 1-form on a surface of genus four. Each edge is a systole, the 1-form ω has exactly one zero, and no two Delaunay edges are homotopic.

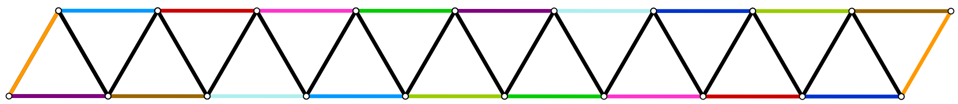

Figure 6: Glue the edges of the polygon according to the colors to obtain the Delaunay triangulation associated a holomorphic 1-form on a surface of genus five. Each edge is a systole, the 1-form ω has exactly one zero, and no two Delaunay edges are homotopic.

convex hull of $\{0, 1, e^{\pi i/3}, 1 + e^{\pi i/3}\}$. Given $\sigma, \tau \in S_{2g-1}$, glue the left side of P_i to the right side of $P_{\sigma(i)}$ and the top side of P_i to the bottom side of $P_{\tau(i)}$ to obtain a surface with a holomorphic 1-form ω. If $[\sigma, \tau]$ is an n-cycle, then it follows that ω has one zero, and if σ, τ, and $\sigma \cdot \tau$ have no fixed points, then it follows that (X, d_ω) has no cylinder with girth equal to the systole. Thus, by Proposition 2.2, no two systolic edges are homotopic.

Conversely, suppose that a holomorphic 1-form surface saturates the bound, then the necessarily equilateral Delaunay triangles can be paired to form parallelograms as above that are glued according to permutations σ and τ. One verifies that σ and τ satisfy the desired properties.

The surface constructed in Figure 4 corresponds to the pair $\sigma = (12345)$, $\tau = (15243)$, the surface constructed in Figure 5 corresponds to the pair $\sigma = (1234567)$, $\tau = (1364527)$, and surface in Figure 4 corresponds to $\sigma = (123456789)$, $\tau = (146379285)$. We thank Marston Condor for finding these examples for us.

If the genus of the surface is two, then one can show that the maximum number of homotopy classes of systoles is $7 = 6g - 5$. More generally, the following is true.

Theorem 3.3. *Let ω be a holomorphic 1-form on a surface with a hyperelliptic involution τ. If ω has exactly one zero, then (X, d_ω) has at most $6g - 5$ homotopy classes of systoles. Moreover, (X, d_ω) has exactly $6g - 5$ homotopy classes of systoles if and only if each Delaunay edge is a systole and there exist exactly four Delaunay 2-cells each of which have two edges that are preserved by the hyperelliptic involution.*

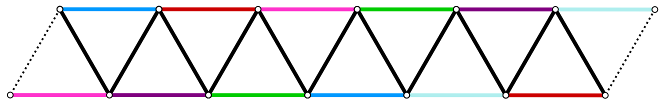

Figure 7: Glue the edges of the polygon according to the labels to obtain the Delaunay triangulation associated a holomorphic 1-form on a surface of genus two having two boundary components.

For each $g \geq 2$, the bound given in Theorem 3.3 is achieved by, for example, the surface described in Figure 8.

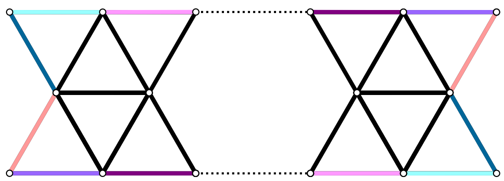

Figure 8: Glue the colored edges of the polygon according to the labels to obtain the Delaunay triangulation associated to a holomorphic 1-form on a surface of genus g. The surface is hyperelliptic, the 1-form ω has exactly one zero, and there are exactly $6g-5$ homotopy classes of systoles.

Proof. Each homotopy class of systole is represented by at least one systolic Delaunay edge. Since ω has exactly one zero z_0, the number of Delaunay edges is at most $6g-3$. Thus, we wish to show that if there are $6g-3$ or $6g-4$ systolic Delaunay edges, then there exist at least two homotopic pairs of systolic edges and that if there are $6g-5$ systolic edges, then there is at least one pair of homotopic edges.

<u>$6g-3$ systolic edges:</u> Suppose that there are exactly $6g-3$ systolic Delaunay edges. Then each Delaunay 2-cell is an equilateral triangle and there are $4g-2$ such cells. Since τ is an isometry, it preserves the Delaunay partition. In particular, since z_0 is the unique 0-cell, we have $\tau(z_0) = z_0$, and since an equilateral triangle has no (orientation preserving) involutive isometry, the involution τ has no fixed points on the interior of each 2-cell. Thus, the remaining $2g+1$ fixed points of τ lie on 1-cells. In particular, τ fixes exactly $2g+1$ Delaunay edges.

Suppose that T is a 2-cell with two fixed edges. Then $T \cup \tau(T)$ is a cylinder whose boundary components are the 'third' edges of T and $\tau(T)$, and, in particular, these 'third' edges are not fixed by τ. Thus, a 2-cell has either zero, one, or two fixed edges. Note that the number of 2-cells that have two fixed edges is even.

We claim that there exist at least four 2-cells that each have two fixed edges. Indeed, if, on the contrary, there are at most two such 2-cells, then there are at least $4g - 4$ remaining 2-cells that each have at most one fixed edge. There are at most $2g - 2$ of edges associated to these 2-cells, and at most 2 edges associated to the 2-cells that have two fixed edges. But, there are $2g + 1 > (2g - 2) + 2$ fixed edges, and we have a contradiction.

The four 2-cells form two cylinders each bounded by two systolic edges. Thus, there are at most $6g - 5$ homotopy classes of systoles. If there are exactly $6g - 5$ homotopy classes cylinders, then there are two maximal cylinders each bounded by two systolic edges. The integral of ω over the middle curve of each cylinder is nonzero, and hence the middle curve is not null-homologous. The induced action of a hyperelliptic involution on $H_1(X)$ is the antipodal map, and so τ preserves each cylinder and has exactly two fixed points on the interior of each cylinder. It follows that the there are four 2-cells each having two fixed edges.

$6g - 4$ systolic edges: Suppose that there are exactly $6g - 4$ systolic Delaunay edges. It follows that exactly $4g - 4$ Delaunay 2-cells are equilateral triangles. The complement, K, of the union of these equilateral triangles is (the interior of) a rhombus.

Since τ is an isometry, τ preserves K. In particular, the center c of the rhombus is a fixed point of τ, and so exactly $2g$ systolic edges are fixed by τ. Thus, since K is a rhombus, none of the edges of K are not fixed by τ. Indeed, if a boundary edge of K were fixed by τ, then the segment in K joining the midpoint of e to c would be 'rotated' by τ to a segment joining c to the midpoint of the edge e' opposite to e. Hence the midpoint of e would equal the midpoint of e', a contradiction.

As in in the case when there are $6g - 3$ systolic edges, each equilateral triangle has at most two fixed edges and the number of equilateral triangles that have two fixed edges is even. We claim that there are at least four equilateral triangles that each have two fixed edges. Indeed, if not, then there would be at least $4g - 6$ equilateral triangles with at most one fixed edge. The fixed edges of these triangles do not lie in the boundary of K, and so there are $2g - 3$ fixed edges associated to these triangles. There are at most two fixed edges associated to the triangles that each have two fixed edges. But there are $2g > 2g - 3 + 2$ fixed systolic edges, and we have a contradiction.

Each pair of equilateral triangles with two fixed edges determines a cylinder bounded by systolic edges. Hence, in this case, there are at most $6g - 6$ homotopy classes of systoles.

$6g - 5$ systolic edges: Suppose that there are exactly $6g - 5$ systolic edges. Then there are $4g - 6$ Delaunay 2-cells that are equilateral triangles. The complement, K, of the union of these equilateral triangles consists of either an equilateral hexagon or two disjoint rhombi.

Suppose that K is an equilateral hexagon. Then since τ preserves the Delaunay partition, we have $\tau(K) = K$. Hence K contains exactly one fixed point c and K is convex. Thus, arguing as above, we find that if a boundary edge of K is fixed by τ, then the edge equals an opposite edge. Since $X/\langle\tau\rangle$ is connected and the genus of X is at least two, all six edges can not be indentified, and hence there are at most 3 fixed points in $\overline{K}$.

We claim that at least one pair of equilateral triangles each have exactly two fixed edges. If not, then each of the $4g-6$ equilateral triangles contains at most one fixed edge. Thus, there are at most $2g-3$ such edges, and hence $2g(2g-3)+3+1 = 2g+1$ fixed points in total. But the total number of fixed points is $2g+2$. Thus, we have a pair of equilateral triangles that share a pair of fixed edges. The union is a cylinder bounded by two systolic edges, and so we have at most $6g-6$ homotopy classes of systoles in this case.

Finally suppose that K is the disjoint union of two rhombi R_+ and R_-. Since τ preserves the Delaunay partition, either $\tau(R_\pm) = R_\pm$ or $\tau(R_\pm) = R_\mp$.

If $\tau(R_\pm) = R_\pm$, then each rhombus contains a fixed point. If an edge of $R_\pm$ is fixed, then $\overline{R_\pm}$ is a cylinder bounded by systolic edges and so there are at most $6g-6$ homotopy classes of systoles. If neither rhombus has boundary edges fixed by τ, then $\overline{K}$ contains exactly two fixed points. If there is not a pair of equilateral triangles that share fixed boundary edges, then each of the $4g-6$ equilateral triangles would have at most one fixed edge, and so there would be at most $2g-3+2+1 = 2g$ fixed points, a contradiction. Hence we have a systolic cylinder and at most $6g-6$ homotopy classes of systoles.

If $\tau(R_\pm) = R_\mp$, then they do not contain fixed points. If a $R_\pm$ has a fixed boundary edge, then it shares this edge with $R_\mp$, and so there are at most three fixed points in $\overline{K}$. Arguing as in the case of the hexagon, we find that there are at most $6g-6$ homotopy classes of systoles. □

Since each genus two surface is hyperelliptic, we have the following corollary.

Corollary 3.4. *Let X be a surface of genus two. If ω is a holomorphic 1-form on X that has exactly one zero, then the number of homotopy classes of systoles of (X, d_ω) is at most 7.*

3.2. Lengths of systoles

Although our main concern is the number of systoles, we observe in this section that it is quite straightforward to find a sharp upper bound on the length of systoles of translation surfaces provided they have a single cone point singularity. One of the ingredients is the

Delaunay triangulation described in §2.5. The other ingredient is a result due to Fejes Tóth which we state in the form of the following lemma.

Lemma 3.5. *Let T be a Euclidean triangle embedded in the plane and let r be the maximal positive real number so that the open balls of radius r around the three vertices are disjoint. Then*

$$r^2 \leq \frac{\text{Area}(T)}{\sqrt{3}}$$

with equality if and only if T is equilateral.

This can be stated differently in terms of ratios of areas. Consider the area A_r of a triangle found at distance r from the vertices of T and so that the interior of the three wedges don't overlap. Then the ratio A_r/T never exceeds that of the equilateral triangle with r equal to half the length of a side.

With this in hand, the following is immediate.

Theorem 3.6. *Let $X \in H(2g-2)$. Then*

$$\frac{\text{sys}^2(X)}{\text{area}(X)} \leq \frac{2\sqrt{3}}{9g-9}$$

with equality if and only if X is obtained by gluing equilateral triangles.

Proof. Given $X \in H(2g-2)$, we consider a Delaunay triangulation of X formed of triangles $T_1, \ldots, T_{6g-6}$. All systolic paths passing through the cone point singularity belong to the triangulation. We now consider the r ball around the cone point singularity. For small enough r it is embedded and we set r_0 to be the first value for which the closed r_0 ball is no longer embedded. Clearly $r_0 = \frac{\text{sys}(X)}{2}$.

Maximizing systolic ratio is the same as minimizing its inverse, so we have

$$\frac{\text{area}(X)}{\text{sys}^2(X)} = \frac{\text{area}(X)}{4r_0^2} = \frac{1}{4}\sum_{k=1}^{6g-6} \frac{\text{area}(T_k)}{r_0^2}$$

which by the previous lemma satisfies

$$\frac{1}{4}\sum_{k=1}^{6g-6} \frac{\text{area}(T_k)}{r_0^2} \geq \frac{1}{4}(6g-6)\sqrt{3}$$

and thus

$$\frac{\text{sys}^2(X)}{\text{area}(X)} \leq \frac{2\sqrt{3}}{9g-9}$$

as claimed.

If one of the triangles has a smaller ratio, then the above inequality is strict and thus equality only occurs if the Delaunay triangulation consisted of equilateral triangles. $\square$

We note that there is a unique surface (up to homothety) in $H(2)$ tiled by equilateral triangles (illustrated previously in Figure 3). This property of having a unique surface be the maximum for both the number of systoles and the maximal ratio is something that is no longer true in $H(1,1)$ as we'll show in what follows.

4. Geodesics on a surface in $\mathcal{H}(1,1)$

In this section, X will denote a $\mathcal{H}(1,1)$ surface of genus two equipped with a translation structure with two cone points c_+ and c_- each of angle 4π. The tangent bundle of a translation surface is parallelizable. In particular, each oriented segment has a direction. The hyperelliptic involution $\tau : X \to X$ is an isometry that reverses the direction of each oriented segment. The isometry τ has exactly six fixed points, the *Weierstrass points*.

Lemma 4.1. *The hyperellipic involution τ interchanges cone points: $\tau(c_\pm) = c_\mp$*

Proof. Since τ is an isometry the set $\{c_+, c_-\}$ is permuted. If $\tau(c_+) = c_+$, then in a neighborhood of c_+, the isometry τ acts as a rotation of π radians. But the cone angle is 4π, and hence it is impossible for τ^2 to be the identity. □

By Lemma 4.1, the quotient $X/\langle\tau\rangle$ is a sphere with one cone point c^* with angle 4π and six cone points $\{c_1, \ldots, c_6\}$ each of angle π. Let $p : X \to X/\langle\tau\rangle$ denote the degree 2 covering map branched at $\{c_1, \ldots, c_6\}$. If γ is a simple geodesic loop, then either γ passes through two Weierstrass points in which case p maps γ onto a geodesic arc joining two distinct π cone points, or $p \circ \gamma$ is a simple geodesic loop that misses the π cone points.

A *flat torus* is a closed translation surface (necessarily of genus one). A *slit torus* is a flat torus with finitely many disjoint simple geodesic arcs removed. Each removed arc is called a *slit*. The completion of a slit torus (with respect to the natural length space structure) is obtained by adding exactly two geodesic segments for each removed disk. The interior angle between each pair of segments is 2π. This property characterizes slit tori.

Lemma 4.2. *Let Y be a topological torus with finitely many disjoint closed discs removed. If Y is equipped with a translation structure such that each boundary*[1] *component consists of at most two geodesic segments, then Y is isometric to a slit torus.*

Remark 4.3. Figure 9 shows that Lemma 4.2 is false if one replaces the assumption of translation structure with the assumption of flat structure.

[1]By *boundary* we mean the set of points added by taking the metric completion of the length structure associated to the translation structure.

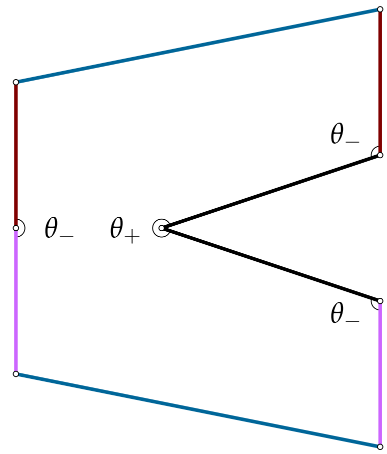

Figure 9: Identify the edges of the same color (except for black) via elements of $\mathrm{Isom}(\mathbb{R}^2)$ to obtain a torus with a disc removed equipped with a flat structure such that the boundary consists of exactly two geodesics. The angles between the geodesics are not both π though they sum to 4π.

Proof of Lemma 4.2. Let Z be a connected component of ∂Y. Let A be the intersection of the maximal geodesic segments in Z. By assumtion A is either empty, contains one point, or contains two points. Let $\alpha : [0,1] \to Z$ be a parameterization of Z such that if A is nonempty, then $\alpha(0) = \alpha(1) \in A$. Let $\overline{\alpha}$ be the develoment of α into the plane $\mathbb{C}$ as discussed in §2.

Since $[\alpha] \in \pi_1(Y)$ is a commutator and $\mathbb{C}$ is abelian, the holonomy of $[\alpha]$ equals 0. Hence by (4), we have $\overline{\alpha}(1) - \overline{\alpha}(0) = \mathrm{dev}([\alpha] \cdot \widetilde{\alpha}(0)) = 0$, and therefore $\overline{\alpha}(1) = \overline{\alpha}(0)$.

If A is empty or consists of one point, then $\overline{\alpha}$ is a line segment, but this is impossible nontrivial line segments in $\mathbb{C}$ have distinct endpoints. If A consists of two points, then the curve $\overline{\alpha}$ consists of two line segments. Since $\overline{\alpha}(1) = \overline{\alpha}(0)$, the line segments coincide. Removing this segment and its translates by $\mathrm{hol}(\pi_1(Y))$ and quotienting it by $\mathrm{hol}(\pi_1(Y))$ gives a surface isometric to Y. $\square$

As a corollary, we have the following sharpening of Theorem 1.7 in [McMullen].

Corollary 4.4. *If α is a separating simple closed geodesic on X, then $X - \alpha$ is the disjoint union of two slit tori. Moreover, each slit torus contains exactly three Weierstrass points, and the hyperelliptic involution τ preserves α.*

Proof. Since α is separating and X is closed of genus two, the complement of α consists of two one-holed tori Y_+ and Y_-. Since α is geodesic, the boundaries of Y_+ and Y_- are piecewise geodesic. Since α is simple and there are only two cone points, the number of geodesic pieces of $Y_{\pm}$ is at most two. Lemma 4.2 implies that each component is a slit torus.

The restriction of τ to a slit torus component determines an elliptic involution $\underline{\tau}$ of the torus. The endpoints of each slit correspond to the cone points c_+ and c_-, and so the are preserved by the induced elliptic involution. Since τ preserves the cone points, the map $\underline{\tau}$ preserves the slit, and hence α is preserved by τ. In particular, the midpoint of the slit is fixed by $\underline{\tau}$ and the three other fixed points of $\underline{\tau}$ are fixed points of τ. $\square$

A *cylinder* of girth ℓ and width w is an isometrically embedded copy of $(\mathbb{R}/\ell\mathbb{Z}) \times [-w/2, w/2]$. Each cylinder is foliated by geodesics indexed by $t \in [-w/2, w/2]$. We will refer to the geodesic that corresponds to $t = 0$ as the *middle geodesic*. By Corollary 4.4, if a simple closed geodesic lies in a cylinder, then it is nonseparating.

A cylinder C is said to be *maximal* if it is not properly contained in another cylinder. If a closed translation surface has a cone point, then each geodesic that does not pass through a cone point lies in a unique maximal cylinder.

Because the hyperelliptic involution τ reverses the orientation of isotopy classes of simple curves, the map τ restricts to an orientation reversing isometry of each maximal cylinder C, and thus it restricts to an orientation reversing isometry of the middle geodesic $\gamma \subset C$. In particular, it contains two Weierstrass points.

Proposition 4.5. *If γ is a nonseparating simple closed geodesic, then γ is homotopic to a unique geodesic γ' such that the restriction of τ to γ' is an isometric involution of γ'.*

Proof. If γ does not contain a cone point, then γ belongs to a maximal cylinder. If γ belongs to a maximal cylinder C, then it is homotopic to the middle geodesic $\gamma' \subset C$.

If γ does not belong to a cylinder, then γ is the unique geodesic in its homotopy class. Since τ reverses the orientation of the homotopy classes of simple loops, it acts like an orientation reversing isometry on γ. $\square$

Proposition 4.5 reduces the counting of homotopy classes of nonseparating systoles to a count of nonseparating systoles that pass through exactly two Weierstrass points. In the next two sections we analyse such geodesics.

5. Direct Weierstrass arcs

If γ is a simple closed geodesic on X that passes through two Weierstrass points, then the projection $p(\gamma)$ is an arc on $X/\langle\tau\rangle$ that joins one angle π cone point to another angle π

cone point. We will call each such an arc a *Weierstrass arc*. Note that the p inverse image of a Weierstrass arc is a geodesic and so we obtain a one-to-one correspondence between homotopy classes of nonseparating simple geodesic loops on X and Weierstrass arcs on $X/\langle\tau\rangle$. A Weierstrass arc that is the image of a systole will be called a *systolic Weierstrass arc*.

The Weierstrass arcs come in two flavors. We will say that a Weierstrass arc is *indirect* if it passes through the angle 4π cone point, and otherwise we will call it *direct*.

Lemma 5.1. *There is at most one direct systolic Weierstrass arc joining two angle π cone points.*

Proof. Suppose to the contrary that there exist two distinct direct systolic Weirestrass arcs that both join the angle π cone point c to the angle π cone point $c' \neq c$. These arcs lift to closed systoles γ_+ and γ_- that interesect tranversally at two Weierstrass points corresponding to c and c'. In particular, the Weierstrass points divide each geodesic into two arcs. By concatenating a shorter[2] arc of γ_+ with a shorter arc of γ_- we construct a piecewise geodesic closed curve α that has length at most the systole. Since the angle between the arcs is strictly between 0 and π, we can perturb α to obtain a shorter curve whose length is strictly less than the systole. This contradicts the assumption that γ_+ and γ_- are both systoles. $\square$

Proposition 5.2. *If c is a cone point on $X/\langle\tau\rangle$ with angle π, then at most two direct systolic Weierstrass arcs have an endpoint at c. Thus, there are at most six direct systolic Weierstrass arcs.*

Proof. Suppose to the contrary that there exist three direct systolic Weierstrass arcs each having c as an endpoint. Let $\theta_1 \leq \theta_2 \leq \theta_3$ denote the angles between the arcs at c. Since c is an angle π cone point, we have $\theta_1 + \theta_2 + \theta_3 = \pi$. Label the arcs α_i, $i \in \mathbb{Z}/3\mathbb{Z}$, so that the angle between α_{i-1} and α_i equals θ_i. By Lemma 5.1, the other endpoints of the α_i are all distinct. Label the other endpoint of α_i with c_i. Let c_4 and c_5 denote the two remaining angle π cone points.

The lift, $\widetilde{\alpha}_i$, of each α_i to X is a non-separating direct simple closed geodesic on X. The involution preserves $G := \widetilde{\alpha}_1 \cup \widetilde{\alpha}_2 \cup \widetilde{\alpha}_3$ and hence the complement $A := X - G$. We have $\chi(A) = \chi(X) - \chi(G) = 2 - 2 = 0$, and since A contains the fixed points c_4 and c_5, it follows that A is connected and, moreover, is homeomorphic to an annulus.

Let γ be a shortest geodesic in X that represents the free homotopy class corresponding to a generator of $\pi_1(A) \subset \pi_1(X)$. Because $\theta_i < \pi$ and each $\widetilde{\alpha}_i$ is a geodesic, the geometric

[2] If the arcs have the same length, then choose either arc.

intersection number of γ and each $\widetilde{\alpha}_i$ is zero. In particular, γ can not coincide with some $\widetilde{\alpha}_i$ as the intersection number $i(\widetilde{\alpha}_i, \widetilde{\alpha}_j) = 2$ for $i \neq j$. Therefore, $\widetilde{\alpha}_i$ and γ are disjoint for each $i \in \mathbb{Z}/3\mathbb{Z}$, and γ lies in A.

In the remainder of the proof, we will consider separately the two cases: (1) the closed geodesic γ is direct and (2) γ passes through an angle 4π cone point.

$\underline{\gamma \text{ is direct}}$: If γ is direct, then it belongs to a maximal cylinder C. Without loss of generality, γ is the middle geodesic of this cylinder. Since γ is nonseparating, τ preserves C and γ, and in particular, the fixed points c_4 and c_5 lie on γ. To obtain the desired contradiction in this case, it suffices to show that the length of γ is less than $\mathrm{sys}(X)$.

Each component of ∂C consists of a direct geodesic segment $\beta_\pm$ joining an angle 4π cone point $c^*_\pm$ to itself. The geometric intersection number of $\beta_\pm$ and each $\widetilde{\alpha}_i$ equals zero, and hence $\beta_\pm$ does not intersect any of the $\widetilde{\alpha}_i$. Hence the complement $A - C$ consists of two topological annuli K_+ and K_- with $\beta_\pm \subset \partial K_\pm$. Because τ preserves each maximal cylinder as well as A, we have $\tau(K_\pm) = \tau(K_\mp)$. Thus, we will now limit our attention to only one of the two annuli, $K := K_+$. One boundary component of K is the direct geodesic segment $\beta := \beta_+$ joining an angle 4π cone point, $c^* := c^*_+$, to itself. The other boundary component, β', of K consists of three geodesic segments $\overline{\alpha}_1$, $\overline{\alpha}_2$, and $\overline{\alpha}_3$ corresponding respectively to $\widetilde{\alpha}_1$, $\widetilde{\alpha}_2$, and $\widetilde{\alpha}_3$. Moreover, the interior angle between $\overline{\alpha}_{i-1}$ and $\overline{\alpha}_i$ is equal to θ_i. See the left hand side of Figure 10.

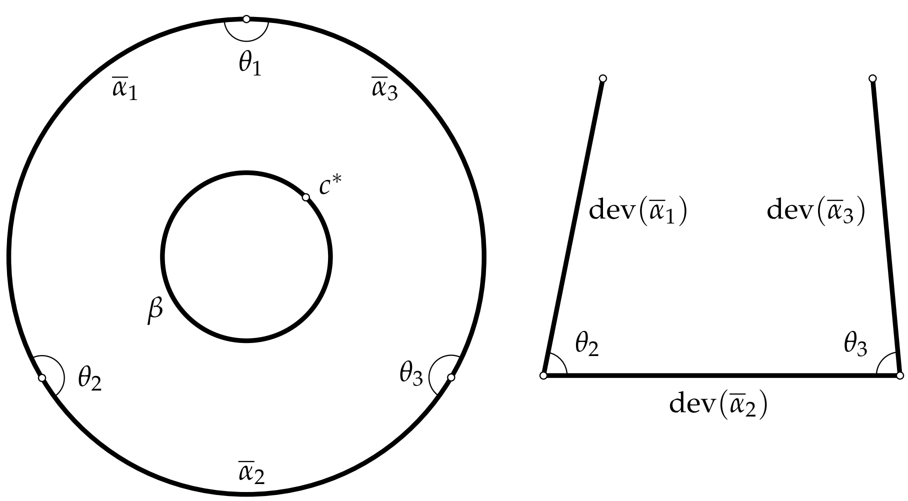

Figure 10: On the left is the topological annulus K case when the closed geodesic γ is direct. The right side shows the development of $\beta' = \overline{\alpha}_1 \cup \overline{\alpha}_2 \cup \overline{\alpha}_3$.

Since β and γ are parallel geodesics in the same cylinder C, it suffices to show that the length of β is less than $\mathrm{sys}(X)$. Since β is a direct geodesic segment, the length of β equals the length of the holonomy vector associated to β. Since β and β' are homotopic, their holonomy vectors have the same length. Thus, it suffices to show that the length of the holonomy vector associated to β' is less than $\mathrm{sys}(X)$.

Since, by assumption, each $\widetilde{\alpha}_i$ is a systole, the length of β' is $b := 3 \cdot \mathrm{sys}(X)$. Let $\beta' : [0, b] \to \partial_\pm$ be a parameterization of β' so that $\beta'(0) = \overline{\alpha}_3 \cap \overline{\alpha}_1 = \beta'(1)$. The development, $\overline{\beta}'$, consists of three line segments each of length $\mathrm{sys}(X)$ joined end to end with consecutive angles θ_2 and θ_3. See the right hand side of Figure 10.

Since $2\pi/3 \leq \theta_2 + \theta_3 < \pi$ and the three sides of $\overline{\beta}'$ have the same length, an elementary fact from Euclidean geometry applies to give that the distance between $\mathrm{dev}(\beta'(0))$ and $\mathrm{dev}(\beta'(1))$ is less than $\mathrm{sys}(X)$. Thus the holonomy vector of β' has length less than $\mathrm{sys}(X)$ as desired.

<u>γ is indirect:</u> In the remainder of the proof we consider the case in which $\pi_1(A)$ is not generated by a direct simple closed geodesic. In this case, the shortest geodesic γ that generates $\pi_1(A)$ is unique in its homotopy class. In particular, since τ induces a nontrivial automorphism of $\pi_1(A) \cong \mathbb{Z}$, the isometry τ preserves γ and reverses its orientation. It follows that γ is a union of two geodesic segments each joining the two 4π angle cone points, and each segment contains as its midpoint one of the remaining two Weierstrass points. Let σ_+ denote the segment containing c_4, and let σ_- denote the segment containing c_5.

The complement of γ consists of two topological annuli K_+ and K_- that are isometric via τ. We limit our attention to one of the annuli, K. One boundary component of K consists of the geodesic segments $\overline{\alpha}_1$, $\overline{\alpha}_2$, and $\overline{\alpha}_3$ with the interior angle between $\overline{\alpha}_{i-1}$ and $\overline{\alpha}_i$ equal to θ_i. The other boundary component consists of σ_+ and σ_-. See Figure 11.

Let c^*_+ and c^*_- denote the angle 4π cone points. Let $\theta_\pm$ denote the interior angle between σ_+ and σ_- at $c^*_\pm$. Because τ interchanges the two components of $A - \gamma$, we have $\theta_+ + \theta_- = 4\pi$. Since γ is not direct, there is no direct geodesic segment joining c_4 and c_5 inside K. Indeed, if there were such a segment δ, then $\delta \cup \tau(\delta)$ would be a direct simple closed geodesic that generates $\pi_1(A)$ contradicting our assumption. It follows that $\theta_\pm \geq \pi$.

We claim that $\theta_1 < \pi/3$. Indeed if not, then since $\theta_1 + \theta_2 + \theta_3 = \pi$ and $\theta_1 \leq \theta_2 \leq \theta_3$, we would have $\theta_i = \pi/3$ for each i and in particular, the holonomy of $\beta = \widetilde{\alpha}_1 \cup \widetilde{\alpha}_2 \cup \widetilde{\alpha}_3$ would be zero. Thus, since $\sigma_+ \cup \sigma_-$ is homotopic to β, the holonomy of $\sigma_+ \cup \sigma_-$ would be trivial. Since $\sigma_\pm$ is a geodesic segment, the angle at $c^*_\pm$ would equal 2π and the lengths of $\sigma_=$ and σ_- would be equal. It would follow that the developing map would map $\overline{K}$ onto the an

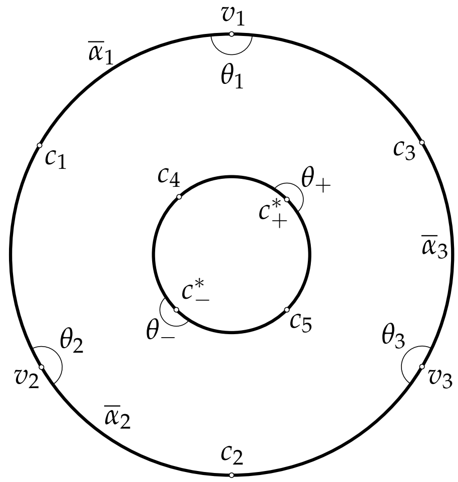

Figure 11: The topological annulus K.

equilateral triangle T having sidelengths $\mathrm{sys}(X)$. Moreover, $\mathrm{dev}(\sigma_+) = \mathrm{dev}(\sigma_-)$ would be a segment σ in the interior of T and the restriction of dev to $\overline{K} - (\sigma_+ \cup \sigma_-)$ would be injective. By elementary Euclidean geometry, the distance from each interior point of T to the set of midpoints of the sides of T is less than $\mathrm{sys}(X)/2$. In particular, it would follow that there would be a direct geodesic segment in $\overline{K}$ joining the set $\{c_4, c_5\}$ and $\{c_1, c_2, c_3\}$ having length less than $\mathrm{sys}(X)/2$. This would contradict the definition of $\mathrm{sys}(X)$.

Thus, in the remainder of the proof, we may assume that $\theta_1 < \pi/3$. Our next goal is the show that this implies that there exists a direct geodesic joining v_1 to one of the two 4π cone points, $c_\pm^*$.

Let V be the set of points $x \in K$ such that there exists a direct geodesic segment in K joining v_1 to x. By lifting to $\widetilde{X}$ and applying the developing map, the set V is mapped injectively onto a subset of the Euclidean sector S of angle θ_1. In particular, v_1 is mapped to the vertex $\overline{v}_1$ of S. The bounding rays of S contain the respective images, $\overline{c}_1$ and $\overline{c}_3$, of the points c_1 and c_3.

Let T be the convex hull of $\{\overline{v}_1, \overline{c}_1, \overline{c}_3\}$ The set T is an isoceles triangle with $|\overline{v}_1\overline{c}_1| = \mathrm{sys}(X)/2 = |\overline{v}_1\overline{c}_3|$, and the angle $\angle \overline{c}_1\overline{v}_1\overline{c}_3$ is less than $\pi/3$. In particular, the side of T that joins $\overline{c}_1$ and $\overline{c}_3$ has length less than $\mathrm{sys}(X)/2$, and the distance from $\overline{v}_1$ to any other point of T is at most $\mathrm{sys}(X)/2$.

Let $x^* \in \overline{S - V}$ be a point such that $\mathrm{dist}(x^*, v_1)$ equals the distance between $\overline{v}_1$ and the $\overline{S - V}$. We claim that x^* is the image of an angle 4π cone point, and hence that there exists a direct geodesic joining v_1 and this angle 4π cone point. See Figure 12.

To verify the claim, we first note that x^* lies in the interior of T. Indeed if it did not, then

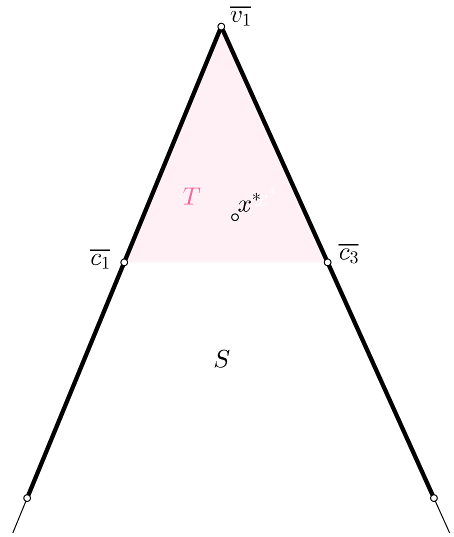

Figure 12: The point x^* in the triangle T.

since the developing map is injective on V, the side of T that joins $\bar{c}_1$ to $\bar{c}_3$ would be the image of a direct geodesic segment joining c_1 and c_3 having length less than $\text{sys}(X)/2$. This would contradict the definition of $\text{sys}(X)$.

Because $\theta_1 \le \theta_2 \le \theta_3$, the distance between v_1 and $\tilde{\alpha}_2$ is at least $\text{sys}(X)/2$, and hence the point x^* can not belong to $\text{dev}(\tilde{\alpha}_i)$. Thus, x^* is the image of a point in σ_+ or σ_-. Thus, it suffices to show that x^* is not the image of an interior point of $\sigma_\pm$.

Suppose to the contrary that x^* were the image of an interior point. Then $\text{dev}(\sigma_\pm)$ would be perpendicular to the segment joining v_1 and x^*, and hence parallel to the side of T that opposes v_1. The segment $\text{dev}(\sigma_\pm)$ does not intersect either $\text{dev}(\tilde{\alpha}_1)$ or $\text{dev}(\tilde{\alpha}_3)$, and hence the midpoint of $\text{dev}(\sigma_\pm)$ would lie in T. The segment joining the midpoint and $\bar{v}_1$ corresponds to a direct geodesic segment joining v_1 to either c_4 or c_5. Since this segment has length less than $\text{sys}(X)/2$, we would obtain a contradiction.

Thus, x^* is the image of either c_-^* or c_+^*. By relabeling if necessary, we may assume that $\text{dev}(c_+^*) = x^*$. Let δ denote the direct geodesic joining v_1 and c_-^*.

Let P denote the metric completion of $X - \delta$. The space P is a topological disk bounded by seven geodesic segments. In particular, we have two vertices, $p_\pm$, corresponding to c_+^*, one vertex, q, corresponding to c_-^*, two vertices, $v_\pm$, corresponding to v_1, as well as vertices v_2 and v_3.

We claim that the developing map $\text{dev} : P \to \mathbb{C}$ is an injection. To see this, let $x, x' \in P$, and

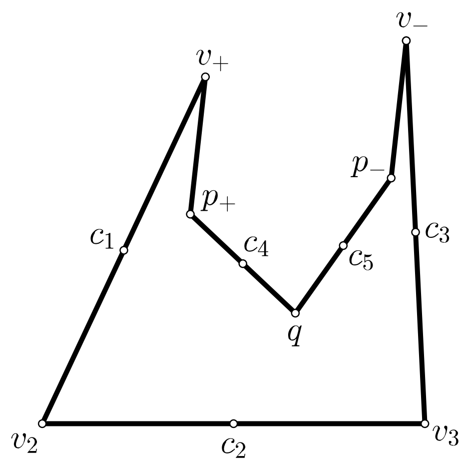

Figure 13: The polygon P.

let η be the unique geodesic in P joining x and x'. To prove the claim it suffices to show that the endpoints of $\text{dev} \circ \eta$ are distinct. If η is a direct geodesic segment, then $\text{dev} \circ \eta$ is a single Euclidean line segment and so $\text{dev}(x) \neq \text{dev}(x')$. If η is not direct, then η is a concatenation of a finite number direct geodesic segments, $\gamma_1, \ldots, \gamma_n$, such that $\gamma_i \cap \gamma_{i+1}$ is a vertex v_i and the angle ψ_i between γ_i and γ_{i+1} satisfies $\pi \leq \psi \leq \theta_v$ where θ_v is the angle between boundary segments at v. Since the angles at $v_\pm$, v_2, and v_3 are less than π, the geodesic η can only pass through the vertices p_+, p_-, or q. Since η is minimal, each of these vertices can appear at most once. Since the angles $\theta_{p_\pm}$ and θ_q are strictly less than 2π, we have $\pi \leq \psi_i < 2\pi$. Since the angle at c^* is 4π, the sum $\theta_{p_+} + \theta_{p_-} + \theta_q = 4\pi$. It follows that if η passes through two vertices, then $\psi_1 + \psi_2 < 3\pi$ and if η passes through three vertices, then $\psi_1 + \psi_2 + \psi_3 < 4\pi$. An elementary Euclidean geometry argument shows that $\text{dev} \circ \gamma$ has distinct endpoints. Thus the claim holds.

In what follows, we will identify the polygon P with its image in $\mathbb{C}$. See Figure 13.

Our next goal is to show that the minimal geodesic joining c_3 to c_5 is direct. Towards proving this, we first show that that the shortest geodesic γ_1 joining c_1 to p_+ is direct. To do this, we refer to the triangle T described in Figure 12. The point p_+ corresponds to $x^* = c_+^*$, and so if the shortest geodesic joining c_4 and p_+ were not direct, then the shortest geodesic in X joining c_1 to would also pass through c_-^*. Hence c_-^* would also belong to the triangle T described above, and so either the image of σ_+ or the image of σ_- would lie in T. But then the midpoint c_4 of σ_+ or the midpoint c_5 of σ_- would belong to T. Hence $|v_1 c_4|$ or $|v_1 c_5|$ would be less than $\text{sys}(X)/2$, a contradiction. (A similar argument shows that the shortest geodesic from c_3 to p_+ is direct.)

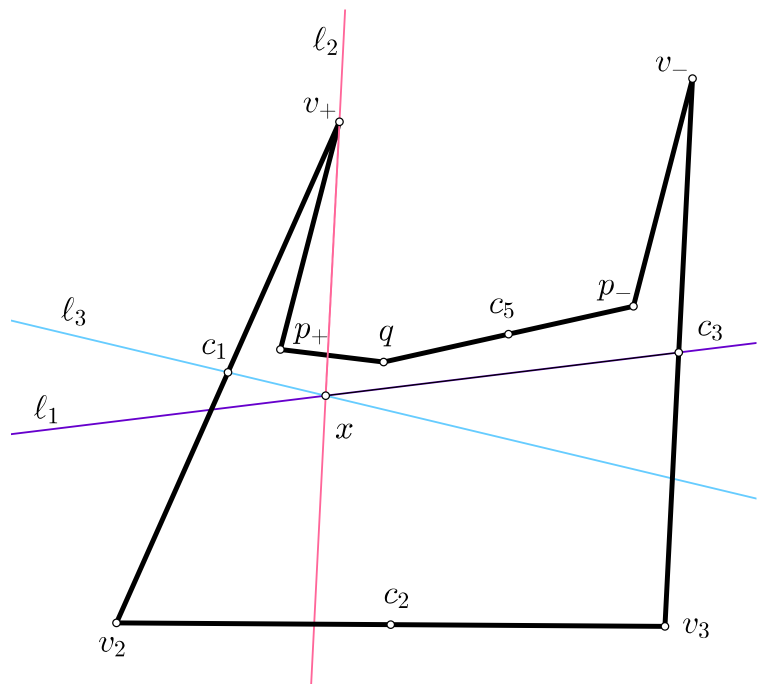

Figure 14: The segment that joins c_3 to c_5 belongs to P.

Because x^* belongs to the interior of T, we have $\angle \overline{v}_1 \overline{c}_1 x^* < \angle \overline{v}_1 \overline{c}_1 \overline{c}_3$. Since T is isoceles, we have $2\angle \overline{v}_1 \overline{c}_1 \overline{c}_3 + \theta = \pi$. Thus, it follows that

$$\angle\, v_+\, c_1\, p_+ \ < \ \frac{\pi - \theta}{2}. \tag{5}$$

(A similar argument shows that $\angle\, v_- c_3\, p_- < (\pi - \theta)/2$.)

We now prove that the minimal geodesic joining c_3 to c_5 is direct. Let ℓ_1 be the line parallel to $\overline{p_+ p_-}$ that passes through c_3, and let ℓ_2 be the line ℓ_2' parallel to $\overline{v_- v_3}$ and passing through v_+. Since $\theta_2 < \pi/2$, the points v_2 and v_3 lie in distinct components of $\mathbb{C} - \ell_2$. Because p_- lies in the component of $\mathbb{C} - \ell_2'$ that contains v_2 and $\overline{p_+ p_-}$ is a translate of $\overline{v_+ v_-}$, the point p_+ lies in the component H_2 of $\mathbb{C} - \ell_2$ that contains v_2. See Figure 14.

Let x be the point of intersection of ℓ_1 and ℓ_2, and let ℓ_3 be the line passing through c_1 and x. Since $|v_+ x| = |v_- c_3| = \mathrm{sys}(X)/2 = |v_+ c_1|$, the triangle $\triangle c_1 x v_+$ is isoceles. Moreover, $\angle c_1 v_+ x = \theta_1$, and so $\angle v_+ c_1 x = (\pi - \theta)/2$. Therefore, if follows from (5) that p_+ lies in the component H_3 of $\mathbb{C} - \ell_3$ that lies

Because $\theta_2 \leq \theta_3$ and $\theta_2 + \theta_3 < \pi$, the intersection $H_2 \cap H_3$ lies in the component H_1 of $\mathbb{C} - \ell_1$ that contains v_+. Thus, $p_+ \in H_1$ and since $\overline{p_+ p_-}$ is parallel to ℓ_1, we have that $p_- \in H_1$. Hence, the angle $\angle c_3 p_- p_+$ is less than π. Therefore, because the interior angle θ_-

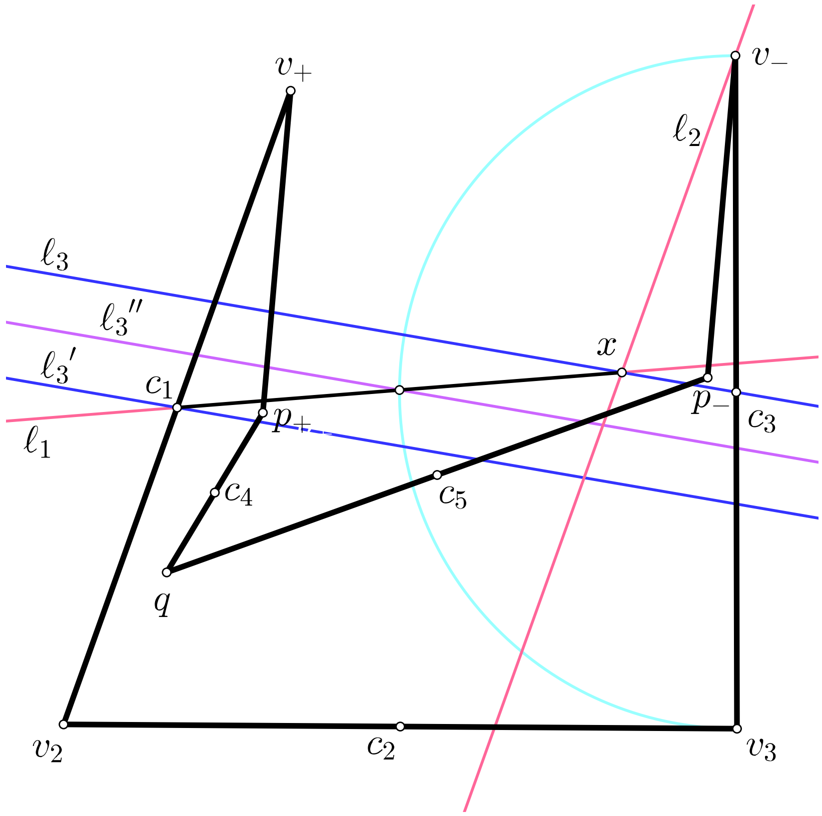

Figure 15: The segment that joins c_1 to c_4 belongs to P.

at q is greater than π, we find that the angle $\angle c_3 p_- q < \pi$. It follows that there is a direct segment from c_3 to c_5 as desired.

Our next goal is to show that the shortest geodesic that joins c_1 to c_4 is direct. Let ℓ_1 be the line passing through c_1 that is parallel to $\overline{v_+ v_-}$. Let H_1 be the the component $\mathbb{C} - \ell_1$ that contains v_+. If the point $p_+ \in H_1$, then since $\overline{p_+ p_-}$ is parallel to ell_1 and the angle $\theta_- < \pi$, it follows that the minimal geodesic joining c_1 to c_4 is direct.

Suppose that p_+ belongs to $\mathbb{C} - H_1$. Then then since $\overline{p_- p_+}$ is parallel to $\ell_1 = \partial(\mathbb{C} - H_1)$, the point p_- also belongs to $\mathbb{C} - H_1$. Moreover, since the angle θ_- at q is larger than π, we also have $c_5 \in \mathbb{C} - H_1$.

Let ℓ_2 be the line through v_- that is parallel to $\overline{v_+ p_+}$, and let x be the intersection point of ℓ_1 and ℓ_2. Let ℓ_3 be the line that passes through x and c_3. The triangle $\triangle x c_3 v_-$ is isoceles, and in particular, $\angle x c_3 v_-$ equals $(\pi - \theta)/2$. The argument analogous to that used to derive (5) gives the inequality $\angle p_- c_3 v_- < (\pi - \theta)/2$. Therefore, p_- lies in the component H_3 of $\mathbb{C} - \ell_3$ that contains v_-.

If we let ℓ_3' denote the line parallel to ℓ_3 that passes through c_1, then, since $\overline{p_+p_-}$ is a translate of $\overline{c_1x}$, the point p_+ lies in the component H_3' of $\mathbb{C} - \ell_3'$ that contains v_-. Thus, to prove that there is a direct segment from c_1 to c_4, it suffices to show that q lies in $\mathbb{C} - H_3$ for then $\angle c_1p_+c_4 < \pi$.

Let m be the midpoint of $\overline{c_1x}$, and let ℓ_3'' be the line parallel to ℓ_3 that passes through m. To show that $q \in \mathbb{C} - H_3'$, it suffices to show that c_5 lies in the closure of the component H_3'' of $\mathbb{C} - \ell_3''$ that contains v_2. Indeed, c_5 is the midpoint of $\overline{p_-q}$ and we know that p_- lies in H_3.

Since there is a direct segment joining c_5 to c_3, the point c_5 lies outside the ball B of radius $\text{sys}(X)/2$ with center at c_3. We also know that c_5 lies in Q, the convex hull of $\{v_+, v_2, v_3, v_-\}$ and that c_5 belongs to $\mathbb{C} - H_1$. An elementary geometric argument shows that $(Q - B) \cap (\mathbb{C} - H_1)$ lies in $\overline{H_3''}$. Thus, $c_5 \in \overline{H_3''}$ and there exists a direct segment joining c_1 to c_4 as desired.

Given that there are direct segments between c_1 and c_4 and between c_3 c_5, we derive a contradiction and thus complete the proof as follows.

Let ℓ_+ be the line that passes through $v_=$ and c_4, let ℓ_- denote the line that passes through v_- and c_5, and let x be the intersection of ℓ_+ and ℓ_-. See Figure 16. An elementary argument shows that the point c_4 is the midpoint of $\overline{v_+x}$. Since c_1 is the midpoint of v_+v_2, we have $|xv_2| = 2 \cdot |c_1c_4|$. Since the geodesic from c_1 to c_4 is direct, we have $|c_1c_4| \geq \text{sys}(X)/2$ and hence $|xv_2| \geq \text{sys}(X)/2$. Similarly, since the geodesic from c_3 to c_5 is direct, we find that $|xv_3| \geq \text{sys}(X)$.

In other words, if we let B^+ (resp. B_-) be the ball of radius $\text{sys}(X)$ about v_2 (resp. v_3), then x lies outside $B_+ \cup B_-$. Note that $\{v_+, v_2, v_3, v_-\}$ belongs to $B_+ \cup B_-$. It follows that P is contained in the convex hull of $B_+ \cup B_-$.

Let $y : \mathbb{C} \to \mathbb{R}$ denote the real linear 1-form such that $|y(z)|$ is the distance from z to the line ℓ_{23} that joins v_2 to v_3 and such that $y(v_1) > 0$. Because $\theta_2 \leq \theta_3 < \pi$, we also have $y(v_-) > 0$. It follows that $y(z) \geq 0$ for each $z \in P$.

Note that that $y(x) < y(q)$. Indeed, since $\angle c_1v_+p_+ < \theta_1$ and $\theta_1 + \theta_2 < \pi$, it follows that $y(v_+) > y(p_+)$. The segment $\overline{xq}$ is the reflection of $\overline{v_+p_+}$ about the point c_4, and hence $y(x) < y(q)$.

Let ℓ_{23} be the line passing through v_2 and v_3, and let x' be the intersection point of ℓ_{23} and the line passing through x and q. The point x' lies in the line segment $\overline{v_2v_3}$. Indeed, because $\theta_2 + \theta_3 < \pi$, line through v_+ and v_2 and the line through v_- and v_3 intersect at a unique point z, and moreover, the polygon P lies in the convex hull T' of $\{z, v_2, v_3\}$. Because $\overline{p_-v_-}$

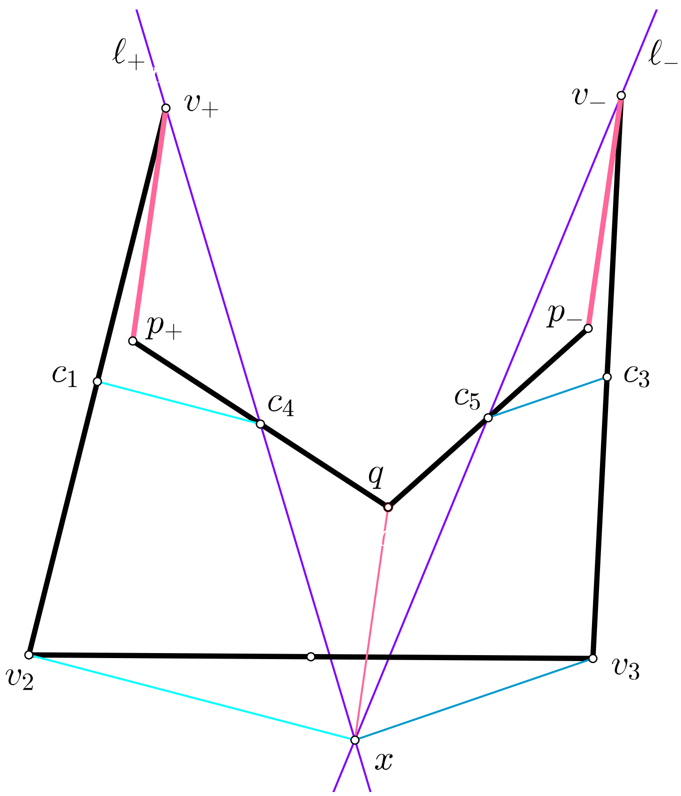

Figure 16: The distances $|xv_2|$ and $|xv_3|$ are at least $\mathrm{sys}(X)$.

and $\overline{p_+v_+}$ are parallel, p_+ and p_- lie in T', and v_+ lies in $\overline{zv_2}$ and v_- lies in $\overline{zv_3}$, any line parallel to p_+v_+ that intersects T' must intersect ℓ_{23} at a point in the segment $\overline{v_2v_3}$. In particular, the point x' lies in $\overline{v_2v_3}$.

We claim that $y(x) > 0$. Indeed, suppose not. Then x' would lie in the segment $\overline{xq}$. Thus, $|x'x| \leq |xq| = |v_\pm p_\pm| \leq \mathrm{sys}(X)/2$, and hence x would belong to the set, A, of points whose distance from $\overline{v_2v_3}$ is at most $\mathrm{sys}(X)/2$. Elementary geometry shows that $A \subset B \cup B_+$, but x lies in the complement of $B_- \cup B_+$, a contradiction.

Let Q be the convex hull of $\{v_+, v_2, v_3, v_-\}$. We have $P \subset Q$ and hence $q \in Q$. Since $0 < y(x) < y(q)$ and the line through x and q meets $\ell_{23} = \ker(y)$ at $x' \in \overline{v_2v_3}$, the point x also belongs to Q. The set Q is contained in the convex hull of $B_+ \cup B_-$. Therefore, x lies inside the convex hull of $B_+ \cup B_-$ and outside $B_+ \cup B_-$. Since $x' \in \overline{v_2v_3}$ it follows that $\pi/4 \leq \angle v_2x'x \leq 3\pi/4$, and, therefore, since $y(q) > y(x)$, we find that q is also outside $B_+ \cup B_-$.

Since x and q both lies inside the convex hull of $B_+ \cup B_-$ but outside $B_+ \cup B_-$, we have $y(q) - y(x) < (1 - \sqrt{3}/2) \cdot \mathrm{sys}(X)$. Since $\pi/4 \leq \angle v_2x'x \leq 3\pi/4$, we have $|xq| \leq \sqrt{2} \cdot |y(q) - y(x)|$ and hence

$$|v_\pm p_\pm| \;\leq\; \sqrt{2} \cdot \left(1 - \frac{\sqrt{3}}{2}\right) \cdot \mathrm{sys}(X) \;<\; \frac{\mathrm{sys}(X)}{4}. \tag{6}$$

Let ℓ_p be the line through p_+ and p_- and let ℓ_v be the line through v_+ and v_-. Let $\ell_\pm$ denote the line passing through $v_\pm$ and $p_\pm$. Because the interior angle θ_- at $q \in P$ is greater than π, the point q lies in the component of $\mathbb{C} - \ell_p$ that contains the segment $\overline{v_2v_3}$, and hence q lies in the component of $\mathbb{C} - \ell_v$ that contains $\overline{v_2v_3}$. Since q lies outside $B_+ \cup B_-$, it follows that q lies in the bounded component of $\mathbb{C} - (\ell_{23} \cup \ell_v \cup \ell_+ \cup \ell_-)$.

Let ℓ_q be the line through q that it parallel to ℓ_v. Let A be the parallelogram that is the bounded component of $\mathbb{C} - (\ell_q \cup \ell_v \cup \ell_+ \cup \ell_-)$. Let $b_\pm$ be the intersection of $\ell_\pm$ and q. Then A is the convex hull of $\{b_+, b_-, v_+, v_-\}$. Because q lies in the component of $\mathbb{C} - \ell_p$ that contains $\overline{x_2x_3}$, the point $p_\pm$ lies in $\overline{v_\pm b_\pm}$.

The line m through x and q is parallel to the sides corresponding to ℓ_+ and ℓ_-. Let x'' be the intersection of m with the side $\overline{v_+v_-}$ of A corresponding to ℓ_v. Since $v_\pm \in B_+ \cup B_-$, the point x'' lies in the convex hull of $B_+ \cup B_-$. By applying the argument that led to (6) to this situation, we find that $|x''q| < \mathrm{sys}(X)/4$.

We have $|bb'| = |vv'| < \mathrm{sys}(X)$ and hence either $|b_+q| < \mathrm{sys}(X)/2$ or $|b_-q| < \mathrm{sys}(X)/2$. Suppose that $|b_+q| < \mathrm{sys}(X)/4$. The midpoint, c_4, of $\overline{p_+q}$ lies in A. Let a_+ be the point of intersection of ℓ_+ and the line through c_4 that is parallel to ℓ_q. Then a_+ lies in the segment $\overline{p_+b_+}$.

By the triangle inequality, we have

$$|v_+p_+| + |p_+c_4| \;\leq\; |v_+a_+| \;+\; |a_+c_4| \;<\; \frac{\mathrm{sys}(X)}{4} + \frac{\mathrm{sys}(X)}{4} = \frac{\mathrm{sys}(X)}{2}$$

But v_+ and c_4 are both Weierstrass points, and hence we would have a curve of length less than $\mathrm{sys}(X)/2$. A similar contradiction is obtained in the case when $|b_-q| < \mathrm{sys}(X)/2$. □

The following is immediate.

Corollary 5.3. *There are at most six homotopy classes of nonseparating systoles.*

6. Indirect Weierstrass arcs

The angle 4π cone point c^* divides each systolic indirect Weierstrass arc into two subarcs. We will call each such subarc a *prong*. The prongs cut a radius ϵ circle about c^* into disjoint arcs. Two prongs are said to be *adjacent* if they are joined by one of these arcs, and the *angle* between two adjacent arcs is $2\pi \cdot \epsilon$ divided by the length of the arc that joins them.

If a systolic indirect Weierstrass arc is the union of two adjacent prongs then the angle between the two prongs must be at least π. Indeed, otherwise one can shorten the arc by perturbing it near c^*.

If the length of the shortest prong is ℓ, then the other prongs have length $\mathrm{sys}(X)/2 - \ell$. Thus, all but at most one of the prongs have the same length. If all of the prongs have the same length, then each pair of adjacent prongs determines a systolic Weierstrass arc. Since the angle between each adjacent pair is at least π and c^* has total angle 4π, there are at most four adjacent pairs and if there are exactly four pairs, then each angle equals π. In sum, we have

Proposition 6.1. *If all of the prongs have the same length, then the number of prongs is at most four and the angle between each pair of adjacent prongs is exactly π.*

We will show below that if one of the prongs is shorter than the the others then there are at most five prongs. To do this we will use the following lemma.

Lemma 6.2. *Two distinct prongs can not end at the same angle π cone point, c'.*

Proof. Suppose not. Then the concatenation, α, of the two prongs would be a closed curve that divides the sphere $X/\langle\tau\rangle$ into two discs. Since there are five other cone points, one of the discs, D, would contain at most two cone points. There are no Euclidean bigons and so D would have to contain at least one cone point.

If D were to contain two angle π cone points, then α would be homotopic to the concatenation of the two oriented minimal arcs joining the two cone points. The length of the unoriented minimal arc is at least $\mathrm{sys}(X)/2$, and hence, since the length of each prong is less than $\mathrm{sys}(X)/2$, we would have a contradiction.

If D were to contain one angle π cone point c, then α would be homotopic to the the concatenation of the two oriented minimal arcs joining c and c'. We would then arrive at a contradiction as in the case of two cone points. □

Since there are exactly six Weierstrass points, Lemma 6.2 implies that there are at most five prongs. In fact, we have the following.

Proposition 6.3. *There are at most five prongs.*

Proof. Suppose to the contrary that there are six prongs. Let e_1 denote the unique shortest prong, let ℓ be its length, and let c_1 denote its endpoint. Let $e_1, \ldots, e_6$ be a cyclic ordering of the remaining prongs, let $L = \mathrm{sys}(X)/2 - \ell$ denote their common length, and let $c_2, \ldots, c_6$ denote their respective endpoints.

Since $\ell(e_1 + e_2) = \mathrm{sys}(X)/2 = \ell(e_1 + e_6)$, the angles $\angle c_1 c^* c_2$ and $\angle c_1 c^* c_2$ are each at least π. (Otherwise, by perturbation near the 4π cone point we could construct a direct Weierstrass arc with length less than $\mathrm{sys}(X)/2$.) Each of the other four angles between adjacent prongs is greater than $\pi/3$. Indeed, otherwise, since $L < \mathrm{sys}(X)/2$, we would have a segment joining two angle π cone points having length less than $\mathrm{sys}(X)/2$ which contradicts the definition of systole. Since $\angle c_1 c^* c_2 + \angle c_1 c^* c_6 \geq 2\pi$ it follows that each of these four angles is less than π. Moreover, since the angle at c^* equals 4π, the sum $\angle c_1 c^* c_2 + \angle c_1 c^* c_6 < 8\pi/3 < 3\pi$ and individually $\angle c_1 c^* c_2 < 5\pi/3$ and $\angle c_1 c^* c_6 < 5\pi/3$.

By cutting along the prongs and taking the length space completion, we obtain a closed topological disc D whose boundary consists of a topological disc bounded by six geodesic segments. The midpoint of each segment corresponds to an end point of a prong. The developing map provides an immersion of D into the Euclidean plane. Since $\angle c_k c^* c_{k+1} + \angle c_1 c^* c_2 < 3\pi$ and $\angle c_i c^* c_{i+1} < \pi$ for $i = 2, \ldots, 5$, this immersion is an embedding. In other words, we may regard D as Euclidean hexagon.

Let v_i donote the vertex of D corresponding to c^* that lies between c_{i-1} and c_i. The length of the side $\overline{v_1 v_2}$ is 2ℓ, and the common length of the other sides is $2L$. From above, the interior angles at v_1 and v_6 are between π and $5\pi/3$, and the angles at the other four vertices lie between $\pi/3$ and π. Without loss of generality, $c_1 = (0,0)$, $v_1 = (\ell, 0)$, $v_6 = (-\ell, 0)$ and an H-neighborhood of c_1 lies in the upper half plane (see Figure **??**).

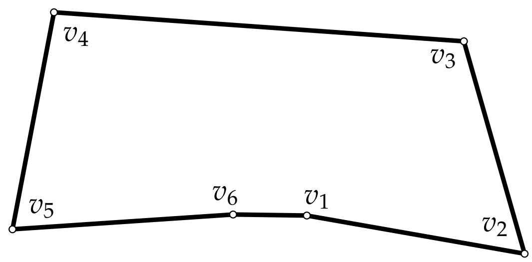

Figure 17: The points v_i, $i = 1, \ldots, 6$

Since the angle at v_2 (resp. v_5) is greater than $\pi/3$, and the edges $\overline{v_1v_2}$ and $\overline{v_2v_3}$ (resp.$\overline{v_4v_5}$ and $\overline{v_5v_6}$) have length $2L$, the vertex v_3 (resp. v_4) lies outside the ball of radius $2L$ centered at v_1 (resp. v_6). It follows that if both v_4 and v_3 both lie in the lower half plane then the shortest arc in H that joins v_4 to v_3 has distance at least $2L + 2\ell$. This contradicts the equality $|v_3v_4| = 2L$.

Since the angle at v_1 (resp. v_6) is at least π and the angle at v_2 (resp. v_5) is greater than $\pi/3$, if v_3 (resp. v_4) lies in the upper half plane, then v_3 (resp. v_4) lies in the half plane $V_+ = \{(x_1, x_2) \mid x_1 > \ell + L\}$ (resp. $V_- = \{(x_1, x_2) \mid x_1 < -\ell - L\}$). Since the distance between U_+ and U_- equals $2L + 2\ell$, if v_3 and v_4 both lie in U, then we contradict $|v_3v_4| = 2L$.

If v_3 lies in the upper half plane and that v_4 lies in the lower half plane but not in U_-, then v_4 lies in the half plane that is bounded by the line trough v_3 and v_6 and contains v_1. In particular, the shortest path in D between v_3 and v_4 passes through v_6. But the distance from v_6 to U_+ is equal to $2L + \ell$, and the distance form v_6 to v_4 is greater than $2L$. Thus, we contradict $|v_3v_4| = 2L$.

A symmetric argument rules out the remaining case in which the rôles of v_3 and v_4 are reversed. $\square$

Theorem 6.4. *There are at most six systolic indirect Weierstrass arcs. Equality occurs if and only if the angle 4π cone point bisects each arc.*

Proof. If the prongs are not all of the same length, then one prong has length less than $\mathrm{sys}(X)/4$ and hence the others have length greater than $\mathrm{sys}(X)/4$. Therefore, concatenations of none of the others constitute a systolic Weierstrass arcs. By Proposition 6.3, there are at most five prongs and hence at most five systolic indirect Weierstrass arcs.

If the prongs all have the same length—namely $\mathrm{sys}(X)/4$—then by Proposition 6.1 there are at $n \leq 4$ prongs. Each concatentation of a pair prongs constitutes a systolic Weierstrass arc, and so there are exactly $n \cdot (n-1)/2$ prongs and hence at most six. $\square$

7. A separating systole

In this section we wish to prove the following:

Theorem 7.1. *If X has a separating systole, then X has at most nine homotopy classes of closed curves with systolic representatives.*

We first observe the following.

Lemma 7.2. *X has at most one separating systole.*

Proof. The angle of intersection between two systoles must be at least π for otherwise one could construct a shorter curve in the same homotopy class. Thus any two systoles must intersection occurs at an angle 4π cone point. But there are only two angle 4π cone points. $\square$

Lemma 7.3. *If α is a separating systole, and γ is direct systolic Weierstrass arc, then $p(\alpha) \cap \gamma = \emptyset$.*

Proof. Suppose not. The lift, $\widetilde{\gamma}$, of γ to X is a systole that does not pass through an angle 4π cone point. Since α is separating, the curve $\widetilde{\gamma}$ intesersects α at least twice. Let p_- and p_+ be two of the intersection points. The points p_+ and p_- divides α (resp. $\widetilde{\gamma}$) into a pair of arcs. One of the arcs, α_- (resp. $\widetilde{\gamma}_-$), has length at most $\mathrm{sys}(X)/2$. By concatenating α_- and $\widetilde{\gamma}_-$, we obtain a non null homotopic closed curve β of length at most $\mathrm{sys}(X)$. Since each intersection point is not a cone point and the geodesics are distinct, the angle at each intersection point $\widetilde{\gamma}_-$ is less than π. Thus, a perturbation of β near an intersection point produces a curve homotopic to β that has shorter length, a contradiction. $\square$

Lemma 7.4. *If X has a separating systole α, then each prong has length equal to $\mathrm{sys}(X)/4$. Moreover, the angle between $p(\alpha)$ and each prong is at least π.*

Proof. If not, then by the discussion at the beginning of §6, there would exist a prong of length strictly less than $\mathrm{sys}(X)/4$. The preimage of a prong under p is an arc γ of length $\mathrm{sys}(X)/2$ that joins one angle 4π cone point c_-^* to the other angle 4π cone point c_+^*. By Corollary 4.4, the separating systole α passes through both c_-^* and c_+^*, and the complement

$\alpha \setminus \{c_-^*, c_+^*\}$ consists of two arcs α_+ and α_- each of length $\text{sys}(X)/2$. By concatenating $\alpha_\pm$ with γ we would obtain a non-null homotopic closed curve having length less than $\text{sys}(X)$, a contradiction.

If the angle between the prong and $p(\alpha)$ were less than π, then one could perturb the concatenation of $\alpha_\pm$ and γ to obtain a non-null homotopic closed curve whose length would be less than $\text{sys}(X)/2$, a contradiction. □

Proof of Theorem 7.1. Let α denote the separating systole to $X/\langle\tau\rangle$ which is unique by Lemma 7.2. By Lemma 7.4, each prong has length equal to $\text{sys}(X)/4$ and the angle between $p(\alpha)$ and each prong is at least π. Thus, since the total angle at c^* is 4π, there are at most two prongs. Hence there are at most two indirect systolic Weierstrass arcs.

By Proposition 5.2, there are at most six direct systolic Weierstrass arcs. Thus, by Proposition 4.5 and the discussion at the beginning of §5, there are at most eight homotopy classes of non-separating closed curves that have systolic representatives. Since α is the unique separating systole, the claim is proven. □

8. Crossing systoles

In this section we prove the following:

Theorem 8.1. *If $X/\langle\tau\rangle$ has four prongs of equal length, then X has at most ten homotopy classes of closed curves that are represented by systoles. Moreover, if X has ten homotopy classes of systoles, then X is homothetic to the surface described in Figure 1, and otherwise X has at most eight homotopy classes of systoles.*

Proof. By Lemma 7.4, the surface X has no separating systole. By Theorem 6.4, there are exactly six indirect systolic Weierstrass arcs. Thus, by Proposition 4.5 and the discussion at the beginning of §5, it suffices to show that there are at most four direct systolic Weierstrass arcs.

By Proposition 6.1, the angle between adjacent prongs equals π. Thus, by cutting along the four prongs we obtain a topological disc D bounded by a geodesic β with no corners. The geodesic β has length $8 \cdot (\text{sys}(X)/4) = 2 \cdot \text{sys}(X)$ and contains one point corresponding to each of the angle 4π cone points that are endpoints of the four prongs. Label those cone points in cyclic order c_1, c_2, c_3, and c_4. For each i, there is a unique point c_i^* on β lying between c_i and c_{i+1} that corresponds to c^*. The distances satisfy $\text{dist}(c_i, c_i^*) = \text{sys}(X)/4 = \text{dist}(c_i^*, c_{i+1})$. The interior angle at each c_i, c_i^* is π.

The two remaining angle π cone points, c_5 and c_6, lie in the interior of the disc D. Because β is a geodesic (without corners), the disk in geodesically convex, and there exists a direct Weierstrass arc γ joining c_5 and c_6. By cutting along γ we obtain a topological annulus A with geodesic boundary components β and β'. Since X is a translation surface, A is a Euclidean cylinder isometric to $[0, h] \times (\mathbb{R}/\ell \cdot \mathbb{Z})$ where $\ell = 2 \cdot \mathrm{sys}(X)$ is the common length of β and β'.

The length of γ equals $(1/2) \cdot \ell$, and hence γ is not systolic. The distance between c_5 (resp. c_6) and $\{c_1, c_2, c_3, c_4\}$ is at least $\mathrm{sys}(X)/2$. It follows that the height h of the cylinder A is at least $(\sqrt{3}/4) \cdot \mathrm{sys}(X)$. As a consequence, there does not exist a direct systolic Weierstrass arc joining two distinct points in $\{c_1, c_2, c_3, c_4\}$.

In sum, if δ is a direct systolic Weierstrass arc, then δ joins a point in $\{c_5, c_6\}$ to a point in $\{c_1, c_2, c_3, c_4\}$. Since A is a Euclidean annulus, there are at most two direct systolic Weierstrass arcs joining c_5 (resp. c_6) to $\{c_1, c_2, c_3, c_4\}$, and hence at most ten systolic Weierstrass arcs in total.

Moreover, since the points $\{c_1, c_2, c_3, c_4\}$ are evenly spaced around β and $\{c_1, c_2\}$ are evenly spaced about β', there are exactly four systolic arcs only if the respective shortest segments, σ_- and σ_6, joining c_5 and c_6 to β bisect arcs joining successive points in $\{c_1, c_2, c_3, c_4\}$, that is, only if σ_- and σ_6 have endpoints in $\{c_1^*, c_2^*, c_3^*, c_4^*\}$. In this case, $h = (\sqrt{3}/4) \cdot \mathrm{sys}(X)$. It follows that X is homothetic to the surface described in Figure 1.

Finally, if there is only one direct systolic Weierstrass arc joining c_5 (resp. c_6) to $\{c_1, c_2, c_3, c_4\}$, then there is only one direct systolic Weierstrass arc joining c_6 (resp. c_5). Hence, if X is not homothetic to the surface described in Figure 1, then X has at most eight homotopy classes of simple closed curves with systolic representatives. □

9. One short prong

In this section we prove the following:

Theorem 9.1. *If $X/\langle\tau\rangle$ has one short prong, then X has at most nine homotopy classes of closed curves that are represented by sytoles.*

Proof. By Lemma 7.4, the surface X has no separating systole. By Proposition 6.3, there are at most five prongs, and so by assumption there is one prong of length $\ell < \mathrm{sys}(X)/4$ and four prongs of length $L = \mathrm{sys}(X)/2 - \ell$. Thus, there are at most four indirect systolic Weierstrass arcs. Thus, it suffices to show that X has at most five direct systolic Weierstrass arcs.

By cutting $X/\langle\tau\rangle$ along the five prongs, we obtain a topological disc D with one angle π cone point in the interior. The boundary consists of five geodesic arcs each of whose endpoints—vertices—corresponds to the angle 4π cone point. The midpoint of each arc corresponds to an angle π cone point on $X/\langle\tau\rangle$. Choose an orientation of the boundary, and let c_1^* and c_2^* denote the endpoints of the oriented arc that corresponding to the short prong. Label the other vertices c_3^*, c_4^*, and c_5^* according to the orientation. Denote by c_i, the midpoint of the arc with endpoints c_i^* and c_{i+1}^*. There remains one angle π cone point, c_6, that belongs to the interior of D.

By Lemma 5.2, for each angle π cone point c_i, there are at most two direct systolic Weierstrass arcs ending at c_i. Thus, to prove the claim, it suffices to show that c_1 is the endpoint of at most one direct systolic Weierstrass arc. We will show that if c_1 is the endpoint of a direct systolic Weierstrass arc, then the other endpoint must be c_6.

Since systolic Weirstrass arcs can not intersect except at a cone point, a direct Weierstrass arc joing c_1 to another angle π cone point can not pass through the boundary of D. In particular, if α is a direct Weierstrass arc joining c_1 to either c_2, c_3, c_4, or c_5, then the complement of α consists of two disks, one that contains c_6 and one that does not.

Suppose that α is a direct geodesic segment that joins c_1 and c_2. Consider the component, D', of $D \setminus \alpha$, containing c_1^*. If D' does not contains c_6, then D' is a flat surface bounded by three geodesic segments. Since the angle at c_1^* is at least π, the Gauss-Bonet formula implies that the angles at c_1 and c_2 are both zero, and hence α is not direct.

If D' contains c_6, then by cutting D' along the geodesic segment joining c_6 and c_1^* we obtain a quarilateral Q with a side corresponding to α. The endpoints of α correspond to c_1 and c_2. Let x_- and x_+ denote the vertices of Q distinguished by $|x_-c_1| = \ell$ and $|x_+c_2| = L$. If α is systolic, then, by the triangle inequality, $|c_1x_+| \leq L + \mathrm{sys}(X)/2$ with equality if and only if c_1, c_2 and x_+ are colinear. The midpoint of $\overline{x_-x_+}$ is c_6, and thus by the triangle inequality

$$|c_1c_6| \;\leq\; \frac{|c_1x_-|}{2} \;+\; \frac{|c_1x_+|}{2} \;\leq\; \ell + L \;=\; \frac{\mathrm{sys}(X)}{2}$$

with equality c_1, c_6, and c_2 are colinear. Thus, either $|c_1c_6| < \mathrm{sys}(X)/2$ or $|c_2c_6| < \mathrm{sys}(X)/2$, a contradiction. Therefore, there is no direct systolic Weierstrass arc joining c_1 and c_2. Similarly, there is no direct systolic Weierstrass arc joining c_1 and c_5.

Suppose that α is a direct geodesic segment that joins c_1 to c_3. Let D' denote the component of $X \setminus \alpha$ that contains c_2. If D' does not contain c_6, then D' is a quadrilateral with vetices c_1, c_1^*, c_2^*, and c_3. Since $|c_2c_2^*| = L = |c_2^*c_3|$, the angle $\angle c_3c_2c_2^*$ is less than $\pi/2$, and thus $\angle c_1^*c_2c_3 > \pi/2$. Therefore $|c_1^*c_3| > |c_2c_3| \geq \mathrm{sys}(X)/2$. Because $|c_1^*c_2^*| = 2L$ and $|c_2^*c_3| = L$, the angle $\langle c_2c_1^*c_3$ is acute. Thus, since the interior angle at c_1^* is at least π, the angle $\angle c_1c_1^*c_3$

greater than π. In particular, $|c_1c_3| > |c_1^*c_3|$, and so, in sum, the length of α is greater than $\text{sys}(X)/2$.

If D' contains c_6, then the other component of $D \setminus \alpha$, is a pentagon with vertices c_1, c_3, c_3^*, c_4^*, and c_5^*. Using the triangle inequality, we have

$$L + |c_3c_5^*| \;\geq\; |c_3c_3^*| + |c_3^*c_5^*| \;\geq\; |c_3^*c_5^*| \;=\; 2|c_4c_5| \;\geq\; \text{sys}(X) \;=\; 2\ell + 2L,$$

and therefore $|c_3c_5^*| \geq 2\ell + L > \ell + L = \text{sys}(X)/2$. Since $|c_3^*c_5^*| \geq \text{sys}(X) > 2L = |c_3^*c_4^*| = |c_4^*c_3^*|$, the angle $\angle c_3^*c_5^*c_3$ is less than $\pi/3$. Because $|c_3^*c_5^*| > 2L = |c_3^*c_3|$, we have $\angle c_3^*c_5^*c_3 < \pi/6$. Thus, since the interior angle at c_5^* is at least π, the angle $\angle c_1c_5^*c_3$ is greater than $\pi/2$. Therefore, $|c_1c_3| > |c_3c_5^*|$. In sum, $|c_1c_3| > \text{sys}(X)/2$, and hence α is not systolic. Therefore, there is no direct systolic Weierstrass arc joining c_1 to c_3. A similar argument shows that there is no direct systolic Weierstrass arc joining c_1 to c_4. □

References

[Broughton-Judge] S. Allen Broughton and Chris Judge, *Ellipses in translation surfaces.* Geom. Dedicata **157** (2012), 111-151.

[Bridson-Haefliger] Martin Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften **319**. Springer-Verlag, Berlin, 1999.

[Fejes Tóth] László Fejes Tóth, *Regular figures.* A Pergamon Press Book. The Macmillan Co., New York (1964).

[Forni-Matheus] Giovanni Forni and Carlos Matheus, *Introduction to Teichmller theory and its applications to dynamics of interval exchange transformations, flows on surfaces and billiards*. J. Mod. Dyn. **8** (2014), no. 3-4, 271-436.

[Gutkin-Judge] Eugene Gutkin and Chris Judge, *Affine mappings of translation surfaces: geometry and arithmetic.* Duke Math. J. **103** (2000), no. 2, 191-213.

[Jenni] Jenni, Felix, *Über den ersten Eigenwert des Laplace-Operators auf ausgewählten Beispielen kompakter Riemannscher Flächen.* Comment. Math. Helv. **59** (1984), no. 2, 193-203.

[Katz-Sabourau] Katz, Mikhail G. and Sabourau, Stéphane, *An optimal systolic inequality for CAT(0) metrics in genus two.* Pacific J. Math. **227** (2006), no. 1, 95-107.

[Malestein-Rivin-Theran] Malestein, Justin, Rivin, Igor and Theran, Louis, *Topological designs.* Geom. Dedicata **168** (2014), 221-233.

[Masur-Smillie] Masur, Howard and Smillie, John, *Hausdorff dimension of sets of nonergodic measured foliations.* Ann. of Math. (2) **134** (1991), no. 3, 455-543.

[McMullen] McMullen, Curtis T., *Dynamics of $SL_2(R)$ over moduli space in genus two.* Annals of Mathematics, **165** (2007), no. 2, 397-456.

[Pu] Pu, P. M. *Some inequalities in certain nonorientable Riemannian manifolds.* Pacific J. Math. **2**, (1952). 55-71.

[Sabourau] Sabourau, Stéphane. *Systoles des surfaces plates singulières de genre deux.* Math. Z. **247** (2004), no. 4, 693-709.

[Smillie-Weiss] Smillie, John and Weiss, Barak. *Characterizations of lattice surfaces.* Invent. Math. **180** (2010), no. 3, 535-557.

[Thurston] Thurston, William P. *Shapes of polyhedra and triangulations of the sphere.* The Epstein birthday schrift, pp 511-549, Geom. Topol. Monogr. **1**, Geom. Topol. Publ., Coventry, 1998.

[Thurston Book] Thurston, William P., *Three-dimensional geometry and topology.* Vol. 1. Edited by Silvio Levy. Princeton Mathematical Series, 35. Princeton University Press, Princeton, NJ, 1997.

[Wright] A. Wright, *From rational billiards to dynamics on moduli spaces.* Bull. Amer. Math. Soc. (N.S.) **53** (2016), no. 1, 41-56.

Addresses:

Department of Mathematics, Indiana University, Bloomington, IN, USA

Mathematics Research Unit, University of Luxembourg, Esch-sur-Alzette, Luxembourg

Emails:

cjudge@indiana.edu

hugo.parlier@uni.lu