ON EXTREMAL PUNCTURED SPHERES

by

Marc J. Beauchamp

B.S. in Mathematics, University of Michigan, 2009

Submitted to the Graduate Faculty of

the Kenneth P. Dietrich School of Arts and Sciences in partial

fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2017

UNIVERSITY OF PITTSBURGH

KENNETH P. DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Marc J. Beauchamp

It was defended on

May 4, 2017

and approved by

J. DeBlois, Ph. D., Assistant Professor, University of Pittsburgh

P. Gartside, Ph. D., Professor, University of Pittsburgh

T. Hales, Ph. D., Professor, University of Pittsburgh

B. McReynolds, Ph. D., Associate Professor, Purdue University

Dissertation Director: J. DeBlois, Ph. D., Assistant Professor, University of Pittsburgh

ON EXTREMAL PUNCTURED SPHERES

Marc J. Beauchamp, PhD

University of Pittsburgh, 2017

We present a classification of extremal n-punctured spheres. We show that there are exactly three such surfaces which feature multiple extremal disks: the unique extremal 3- and 4-punctured spheres and a particular 6-punctured sphere as well. We prove that each of these surfaces admit precisely two extremal disks and in all cases the disks are exchanged by a self-isometry of the surface. We demonstrate that for all other n, each extremal n-punctured sphere has a unique extremal disk. We derive formulas to count the exact number of extremal punctured spheres and determine the asymptotic growth rate of this total. Finally, we establish an upper bound on the number of once-punctured extremal surfaces by determining the precise number of extremal disk - surface pairs in this case.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

PREFACE

First and foremost, I would like to thank my advisor Dr. DeBlois for his patient guidance, constant support and nurturing mentorship. No kind words I can write here could come close to doing him justice, but I remain eternally grateful and utterly humbled.

I would like to specifically thank Dr. Hales, whose out-reach email drew me towards Pitt in the first place and whose algebra courses were the best classes I have ever taken; Dr. Gartside, whose seminar on preparing for research proved immensely invaluable in unforseen ways; Dr. Pan, whose excellent courses on preparing for preliminary examinations were integral to my analytic development; Dr. Wheeler, for his friendship, advice and the outrageously enjoyable calculus-finals grading stints in the MAC; and Dr. Lewicka, for refusing to let me audit her analysis class without still doing the homework and whose unwavering devotion to rigorously celebrating the beauty of mathematics has been truly inspirational.

Special thanks to Derek Orr, whose observations on our recursive formulas greatly aided our extremal punctured sphere count. I'd also like to thank all the Pitt graduate students, current and former, who have helped me along this journey with their friendship, support and great mathematical discussions, including but not limited to: Corinne Brucato, Jeremy Harris, Alex Yarosh, Behnam Esmayli, Jake Mirra, Dan Grady, Victor DeCaria, Shuai Zhao, Kim Romanelli and Cezar Lupu.

Lastly, I would like to thank my immediate family: my wife Kate, all of my siblings and my mother, without whose love none of this would have been remotely possible.

I dedicate this dissertation to the man who first encouraged me to pursue a degree in pure mathematics - the Doc.

1.0 INTRODUCTION

A **hyperbolic surface** is an oriented 2-manifold equipped with a complete, finite-area Riemannian metric with constant sectional curvature -1. Hyperbolic surfaces of genus g and with n cusps are parameterized by a moduli space of real dimension $6g - 6 + 2n$ [9].

The **injectivity radius** at a point p on a hyperbolic surface S is the largest radius such that the exponential map [17]

$$\exp_p : T_pS \to S$$

is a diffeomorphism on the open ball of this radius in the tangent space to S at p. Among hyperbolic surfaces within a given topological type (i.e., of the same genus g and number of cusps n), those featuring at least one point at which the largest injectivity radius is attained (among injectivity radii of points on surfaces within that topological type) are known as **extremal surfaces**. We may view extremal surfaces as those which contain an embedded open disk of maximum possible radius (among all surfaces within their topological type) and any such disk is known as an **extremal disk**.

Bavard gave a combinatorial characterization of closed extremal surfaces [2], which in particular implies that there are only finitely many. These were thereafter counted in all genera by work of Girondo & Gonzales-Diez [13] [14] [15], Nakamura [20] [21], and Bacher & Vdovina [1] [26]. Moreover, for each such surface, its collection of extremal disks was classified.

Using a characterization for the non-compact setting, nearly analogous to Bavard's but supplied by DeBlois [8], we present the beginning of a classification for the non-compact case of extremal surfaces via extensive examination of surfaces of genus $g = 0$ with n cusps (the so-called n-punctured spheres). These n-punctured spheres admit complete, finite-area

hyperbolic metrics when $n \geq 3$ and for this special case of non-compact extremal surfaces we provide complete answers to the questions:

1. How many extremal surfaces exist?

and (since a natural approach for addressing this question is to count combinatorial objects determined by extremal surface-disk pairs),

2. Do extremal surfaces contain unique extremal disks?

by proving in this thesis that there are precisely three extremal punctured-sphere surfaces for which extremal disks are not unique:

Theorem. *The extremal* 3*-punctured and* 4*-punctured spheres each have exactly two extremal disks. In addition, there exists an extremal* 6*-punctured sphere which has exactly two extremal disks. Moreover, each of these admits a self-isometry which exchanges its extremal disks.*

This last assertion is analogous to results of Girondo & González-Diez [14] and Nakamura [20] which show there exist extremal surfaces of genus (respectively) 2 and 3 which have multiple extremal disks and that these disks are exchanged by an isometry of the surface.

Analogously to Girondo & González-Diez's demonstration that closed surfaces with genus $g > 3$ have unique disks [15], we then prove that all other extremal punctured spheres have unique extremal disks.

Theorem. *Extremal* n*-punctured spheres have unique extremal disks for* $n = 5$ *and all* $n > 6$*.*

Similarly, we mirror Bacher & Vdovina's closed extremal surface count [1] by procuring formulas to count the exact number of extremal n-punctured spheres for arbitrary n, and providing the explicit totals through $n = 15$.

Theorem. *In terms of Catalan numbers* $C(n) = \frac{1}{n+1}\binom{2n}{n} = \frac{(2n)!}{n!(n+1)!}$*, the exact number of extremal* n*-punctured spheres is given by*

$$C(n-2)/n + C(n/2-1)/2 + (2/3)C(n/3-1)$$

where Catalan terms are omitted if their arguments are not integers.

We show that the asymptotic growth rate of this total is given by $\frac{2^{2n-4}}{\sqrt{\pi n^5}}$ and provide the explicit totals for small n. In contrast, the number of closed extremal surfaces grows factorially with their genus. To emphasize this significantly slower comparative growth rate of the number of extremal n-punctured spheres versus the number of closed extremal surfaces of genus g (within the same Euler characteristic $\chi = 2-2g-n$), we present the totals alongside Bacher & Vdovina's count in a table on the following page.

Finally, we start an extension into the generic non-compact case by modifying Bacher & Vdovina's closed extremal surface count to obtain an upper bound on the number of once-punctured extremal surfaces.

χ	Extremal n-Punctured Spheres	Closed Extremal Surfaces of Genus g
-2	1	9
-4	4	1726
-6	19	1349005
-8	150	2169056374
-10	1424	5849686966988
-12	14924	23808202021448662
-14	167367	136415042681045401661
-16	1965058	1047212810636411989605202

Table 1: Comparison of extremal punctured sphere and closed surface totals

2.0 BACKGROUND

One especially useful perspective to take on a complete hyperbolic surfaces is to view it as the quotient of the hyperbolic plane $\mathbb{H}^2$ by a discrete group of hyperbolic isometries Γ. We treat as equivalent all points of $\mathbb{H}^2$ in the same orbit under the action as homeomorphisms of Γ on it (see Theorem 1.1.1 of Katok [16]), i.e., for any $x, y \in \mathbb{H}^2$, we let $x \sim y$ if and only if $y = T(x)$ for some $T \in \Gamma$. We describe the action itself in explicit detail below, but to make this all appropriately adapted to our situation, we first elaborate on the hyperbolic setting.

2.1 HYPERBOLIC GEOMETRY

Throughout our endeavors, we use two different models of hyperbolic geometry: the upper half-plane model and the Poincaré disk model. For the upper half-plane model, the hyperbolic plane is constructed by the equipping the upper half-plane of $\mathbb{C}$ (with boundary $\mathbb{R} \cup \{\infty\}$) with what becomes the hyperbolic metric [25]:

$$\mathbb{H}^2 = \{x + iy = z \in \mathbb{C} \mid y = \mathrm{Im}(z) > 0\} \quad \text{imbued with the metric} \quad ds = \frac{\sqrt{dx^2 + dy^2}}{y}.$$

We define the length of a piecewise differentiable path $\eta = \eta(t) = \langle x(t), y(t) \rangle : [0, 1] \to \mathbb{H}^2$ here by the integral

$$\ell(\eta) = \int_0^1 \frac{\sqrt{(\frac{dx}{dt})^2 + (\frac{dy}{dt})^2}}{y(t)}\, dt$$

Defining the distance between two points to be the infimum of the lenghts of paths joining them, it can be proved that geodesics in the upper half-plane model are given by (segments

of) circles orthogonal to $\mathbb{R}$ and vertical lines (see Figure 1) [16].

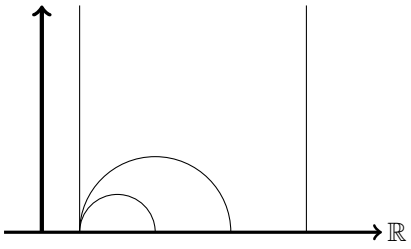

Figure 1: Geodesics in the upper half-plane model

Similarly, for a subset $D \subset \mathbb{H}^2$ we define its hyperbolic area $\mu(D)$ by (if the integral exists)

$$\mu(D) = \int_D \frac{dxdx}{y^2}.$$

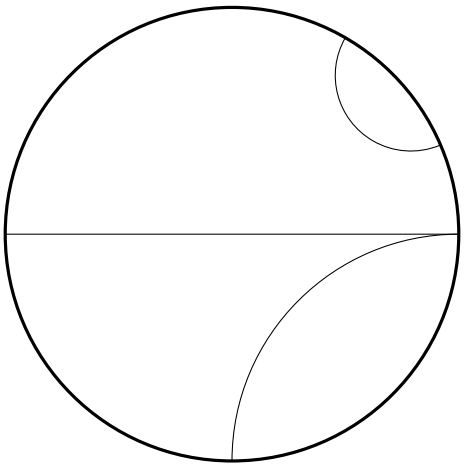

Figure 2: Geodesics in the Poincaré disk model

The second model for hyperbolic geometry that we consider, the aforementioned Poincaré disk model, identifies the hyperbolic plane with the interior of the unit disk in $\mathbb{R}^2$ (with boundary S^1_∞, the "circle at infinity,") and its own metric, specifically we take [6]

$$\mathbb{H}^2_{\mathfrak{p}} = \{(x, y) \in \mathbb{R}^2 \mid \sqrt{x^2 + y^2} < 1\} \quad \text{imbued with the metric} \quad ds = \frac{2\sqrt{dx^2 + dy^2}}{1 - x^2 - y^2}.$$

Here geodesics are given by Euclidean circles orthogonal to the boundary and straight lines through its center (see Figure 2), and as with the half-plane model, there is a unique geodesic between any points along which the distance between them is measured [6].

While we perform most computational affairs in terms of the upper half-plane model, we introduce the disk model here since we will also reference it in schematic diagrams and elsewhere when convenient. Except for where we explicitly mention otherwise however, we adopt the upper half-plane model.

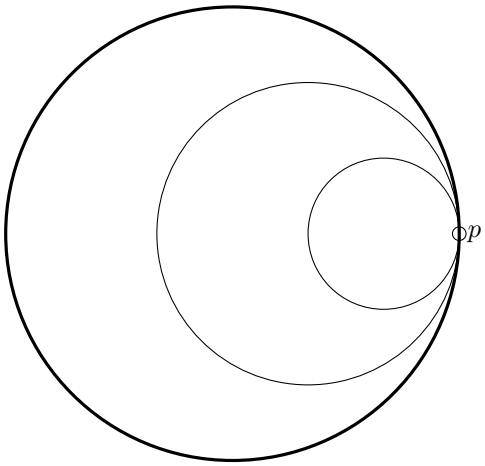

Figure 3: Horocycles with ideal point p, shown in the Poincaré disk model

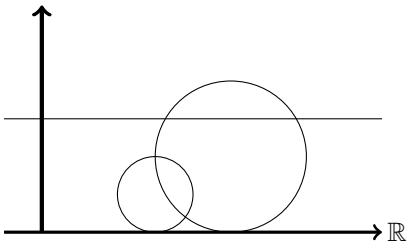

Figure 4: Horocycles shown in the upper half-plane model

To finish fleshing out the hyperbolic plane setting before we begin the discussion of its isometries, we introduce the notion of a horocycle. For any point $p \in \mathbb{H}^2$ and any geodesic ray $z(t)$ $0 \leq t < \infty$ starting from p, let $B_t(p)$ be the hyperbolic circle centered at $z(t)$ and

passing through the point p. Then $\lim_{t\to\infty} B_t(p)$ exists and is a Euclidean circle passing through p (which depends only on p itself and the initial direction of the ray), forming what we call a **horocycle** [16]. In the disk model, horocycles are euclidean circles tangent to the boundary S^1_∞, see Figure 3, and in the upper half-plane they are seen as euclidean circlces tangent to $\mathbb{R}$ and horizontal lines (for those with ideal point at ∞), see Figure 4. Rounding out the discussion on horocycles, we define a **horocyclic ideal triangle** to be the convex hull in $\mathbb{H}^2$ of three points, two of which lie on a fixed horocycle and one at the ideal point of that horocycle [7].

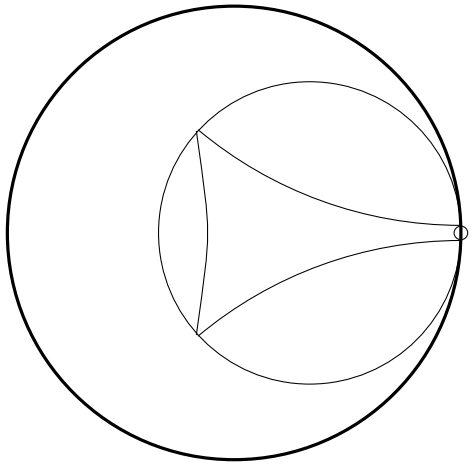

Figure 5: A horocyclic ideal triangle in the Poincaré disk model

2.2 FUCHSIAN GROUPS

The collection of all orientation-preserving isometries of $\mathbb{H}^2$ is isomorphic to the matrix group $\mathrm{PSL}(2,\mathbb{R})$, where $\mathrm{PSL}(2,\mathbb{R}) = \mathrm{SL}(2,\mathbb{R})/\{\pm \mathrm{I}\}$, and $\mathrm{SL}(2,\mathbb{R})$, the special linear group, is the group of 2×2 real matrices with determinant one (the quotient by $\{\pm \mathrm{I}\}$ renders any $\mathrm{SL}(2,\mathbb{R})$ matrix A equivalent to its additive inverse $-A$). The action as homeomorphisms of $\mathrm{PSL}(2,\mathbb{R})$ matrices on the upper half-plane can be explicitly seen as [16]:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot z = \frac{az+b}{cz+d}$$

and one can easily check that for an invertible matrix $A \in \mathrm{GL}(2,\mathbb{R})$ and any non-zero scalar $k \in \mathbb{R}$, the matrices A and kA induce the same transformation in this regard. In particular, the trace of a $\mathrm{PSL}(2,\mathbb{R})$ matrix is only well-defined up to the absolute value of the traces of its $\mathrm{SL}(2,\mathbb{R})$ representatives (and sometimes the also well-defined square of the trace is used). Furthermore, through a fixed-point determination, i.e., solving for conditions where

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot z = z$$

we distinguish between families of elements $T \in \mathrm{PSL}(2,\mathbb{R})$ (hence orientation-preserving isometries of $\mathbb{H}^2$), categorizing them as:

- **elliptic** if $|\mathrm{tr}(T)| < 2$. These transformations have a single fixed point in $\mathbb{H}^2$ and can be visualized as rotations in the Poincaré disk model.
- **parabolic** if $|\mathrm{tr}(T)| = 2$. These transformations have a single fixed point on the boundary and preserve every horocycle through it. For example, the parabolic isometry represented by the matrix

  $$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

  acts on the upper half-plane sending $z \mapsto z+1$, fixing only ∞, and translating all points horizontally, parallel to the real axis (thereby preserving all horocycles through ∞, i.e., horizontal lines). A fundamental region for the action is illustrated in Figure 6.

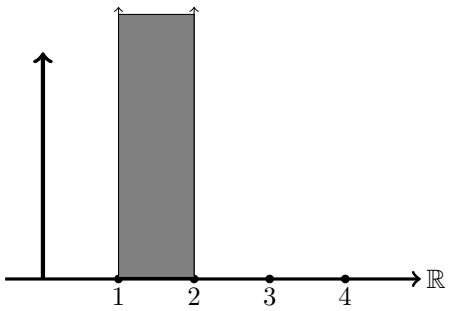

Figure 6: A fundamental domain for the parabolic transformation $z \mapsto z + 1$

- **hyperbolic** if $|\mathrm{tr}(T)| > 2$. These transformations have two fixed points on the boundary. For example, the hyperbolic isometry represented by the matrix

$$\begin{pmatrix} 2 & 0 \\ 0 & \frac{1}{2} \end{pmatrix}$$

  acts on the upper half-plane sending $z \mapsto 4z$, fixing 0 and ∞, and can be seen as a dialation.

Discrete subgroups $\Gamma \leq \mathrm{PSL}(2, \mathbb{R})$ are called **Fuchsian groups**. Alternatively and equivalently, we may define a Fuchsian group as any subgroup of $\mathrm{PSL}(2, \mathbb{R})$ in possession of a properly discontinuous action on the hyperbolic plane, i.e., one such that the Γ orbit of each point is locally finite (see Theorem 2.2.6 of Katok [16]). When we realize a given surface S as the quotient of $\mathbb{H}^2$ by the action of a Fuchsian group Γ, i.e., when $S = \mathbb{H}^2/\Gamma$, we say that S is **uniformized** by Γ and topologically speaking we have that $\pi_1(S) \cong \Gamma$.

As previewed in our description of parabolic isometries, we often observe such a Fuchsian model Γ for S in terms of a **fundamental region** (also called a **fundamental domain**) for the action, defined as a closed region $F \subset \mathbb{H}^2$ with nonempty interior satisfying [16]:

(i) $\bigcup_{T \in \Gamma} T(F) = \mathbb{H}^2$

(ii) $\text{int}(F) \cap T(\text{int}(F)) = \emptyset$ for all $T \in \Gamma - \{\text{Id}\}$

or in other words, a closed subset whose translates tile the hyperbolic plane but such that the translates of its interior are disjoint. We sometimes refer to the family $\{T(F) \mid T \in \Gamma\}$ as the **tessellation** of $\mathbb{H}^2$. A particularly nice type of fundamental region is a Dirichlet region: denoting hyperbolic distance by ρ, the **Dirichlet region** for Γ centered at p is defined to be the set

$$D_p(\Gamma) = \{z \in \mathbb{H}^2 \mid \rho(z,p) \leq \rho(z, T(p)) \text{ for all } T \in \Gamma\}$$

and for any point p not fixed by an element of $\Gamma - \{\text{Id}\}$, the Dirichlet region for Γ centered at p is a fundamental region for the action of Γ on $\mathbb{H}^2$.

Any Dirichlet region $D_p(\Gamma)$ is a convex polygon whose sides are segments of geodesics which are moreover paired together by elements of Γ [19]. More precisely, for each side s of $D_p(\Gamma)$ there exists a unique side s' and an isometry $T_s \in \Gamma \leq \text{PSL}(2, \mathbb{R})$ such that $T_s(s) = s'$ and also $T_{s'} = T_s^{-1}$. We naturally refer to these T_s as side-pairing transformations.

Side-pairing transformations on a convex region also dictate what are referred to as cycle relations. We call two points $x, y \in D_p(\Gamma)$ congruent if they belong to the same Γ orbit (note that by the definition of a fundamental region, congruent points are necessarily on the boundary). If two vertices x, y on the boundary of a fixed Dirichlet region $D_p(\Gamma)$ are identified by the action of Γ (i.e., $y = T(x)$ for some $T \in \Gamma$) and x is fixed by an elliptic element S, then y is fixed by the elliptic element TST^{-1}; so any congruent vertices are fixed by conjugate elliptic elements [16]. This congruence dictates an equivalence relation among vertices of $D_p(\Gamma)$ and the corresponding equivalence classes are known as cycles. This is quite natural, as we can view such an equivalence class as a sequence of vertices

$$E = v_0 \to v_1 \to ... \to v_n \to v_0$$

where at each arrow we are applying a side-pairing transformation to a side containing the previous vertex to arrive at a side containing the next (note that any vertex is contained in two distinct sides). This is to say that for each $i \in \{0, 1, \ldots, n\}$ there exists elliptic T_i such that $T_i(v_i) = v_{i+1}$. The sequence of arrows indicated above hence give us a corresponding sequence of elliptic transformations T_i which satisfy [18]

$$(T_n \circ T_{n-1} \circ \cdots \circ T_1)(v_0) = v_0.$$

Furthermore, if a point has a nontrivial stabilizer in Γ, this stabilizer is a maximal finite cyclic subgroup (see Theorem 2.2.6 and Lemma 2.3.1 of Katok [16]). Hence, there is a one-to-one correspondence between the elliptic cycles of Dirichlet region and the conjugacy classes of non-trivial maximal finite cyclic subgroups. We get precisely one cycle relation for each cycle of vertices, consisting of the presentation of cyclic subgroup generated by an elliptic transformation fixing a vertex in it. More specifically, in terms of our ongoing example elliptic cycle, the corresponding cycle relation is given by

$$(T_n \circ T_{n-1} \circ \cdots \circ T_1)^m = 1.$$

Moreover, from Katok [16] we know:

Theorem. *Let $\{T_i\}$ be the subset of Γ consisting of those elements which pair the sides of some fixed Dirichlet region. Then $\{T_i\}$ is a set of generators for Γ.*

Hence understanding the side pairing isometries on any such domain is a means by which to understand Γ itself and the hyperbolic surface to which it gives rise. Towards the end of investigating hyperbolic surfaces through fundamental regions of the Fuchsian groups uniformizing them, we have the classical result of Poincaré, his Polygon Theorem [3] [19].

The essence of Poincaré's theorem is that a sufficiently nice, convex polygon region equipped with side-pairing isometries is necessarily the fundamental domain for a Fuchsian group and moreover one with a presentation determined entirely by the cycle relations from those very side-pairing transformations, i.e., that it prescribes a hyperbolic surface and in a very exact manner.

Theorem (Poincaré Polygon Theorem)**.** *Let F be a convex polygon region equipped with side-pairing transformations. Suppose that for each elliptic cycle $v_0 \to v_1 \to ... \to v_n$ there exists an integer m such that the corresponding sum of internal angle measures at those vertices satisfies $\angle v_0 + \cdots + \angle v_n = \frac{2\pi}{m}$. Then F is a fundamental domain for a Fuchsian group which is generated by those side-pairing transformations and which has a presentation given by its cycle relations.*

2.3 EXTREMAL SURFACES

The investigations into extremal surfaces were first thrust forward by Christophe Bavard [2], who elegantly characterized compact extremal surfaces with his beautiful result:

Theorem 2.1 (Bavard)**.** *Let R be the radius of a metric disk isometrically embedded into a compact Riemann surface S of genus g. Then*

$$R \leq R_g := \cosh^{-1}\left(1/\left(2\sinh\frac{\pi}{12g-6}\right)\right)$$

which is the radius of the inscribed circle to the regular $(12g-6)$-gon. When the latter is attained, S has a regular $(12g-6)$-gon as a Dirichlet domain and $D(R_g)$ is its inscribed disk.

Bavard's theorem gives an explicit formula for the maximal injectivity radius of a closed surface of genus g and a recipe for realizing an extremal surface enjoying it. This characterization spurred contributions from several other authors, ultimately culminating in a complete classification of the closed case. This complete classification includes the work of Girondo & González-Diez [15], who procure many equivalent descriptions of extremal surfaces themselves and answer the question about whether closed extremal surfaces have unique extremal disks by proving that

Theorem 2.2 (Girondo, González-Diez)**.** *A surface of genus $g > 3$ has at most one extremal disk.*

The closed-case classification also includes work of Bacher & Vdovina [1], who arrive at the precise number of closed extremal surfaces. While the combinatorial object they use to perform the count (the so-called oriented wicks forms) actually enumerate extremal disk - extremal surface pairs, this actually coincides with the number of extremal surfaces themselves: once $g > 3$ and extremal disks are unique per the theorem of Girondo & González-Diez and even for when $g = 3$ due to a result of Nakamura [20], which shows that the group of self-isometries of a genus-3 extremal surface acts transitively on its extremal disks (so that the a priori potential overcount due to enumerating disk-surface pairs when extremal disks are not unique is still in fact an exact count).

Despite the progress on the closed case and creative ingenuity used in resolving it, the techniques employed by the previous authors there do not port into the non-compact setting: Bavard's theorem does not hold and Girondo's equivalent descriptions are not retained. In order to then handle surfaces with cusps (i.e., with $n > 0$), DeBlois [8] has supplied a pseudo-analogue to Bavard's characterization for that non-compact setting by showing:

Theorem 2.3 (DeBlois)**.** *For $r > 0$, let $\alpha(r)$ be the angle of an equilateral hyperbolic triangle with sides of length $2r$, and let $\beta(r)$ be the angle at either endpoint of the finite side of a horocyclic ideal triangle with one side of length $2r$:*

$$\alpha(r) = 2\sin^{-1}\left(\frac{1}{2\cosh r}\right) \qquad\qquad \beta(r) = \sin^{-1}\left(\frac{1}{\cosh r}\right)$$

A complete, oriented, finite-area hyperbolic surface S with genus $g \geq 0$ and $n \geq 0$ cusps has injectivity radius at most $r_{g,n}$ at any point where $r_{g,n} > 0$ satisfies:

$$(4g + n - 2)3\alpha(r_{g,n}) + 2n\beta(r_{g,n}) = 2\pi$$

This bound is attained at $x \in S$ if and only if S has a decomposition into a non-overlapping union of $4g+n-2$ equilateral triangles and n horocyclic ideal triangles with its sole vertex at x. This decomposition is canonically obtained from the Delaunay tessellation of S determined by x by subdividing each horocyclic two cell with a single ray from its vertex that exits its cusp.

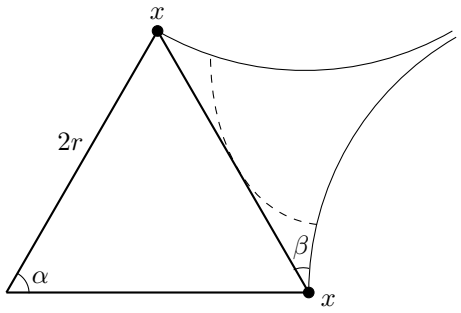

Figure 7: Equilateral triangle with horocyclic ideal triangle, as referenced in Theorem 2.3

3.0 PRELIMINARY RESULTS

Our first order of business is to demonstrate that DeBlois' characterization, Theorem 2.3, of non-compact extremal surfaces provides a Fuchsian model through which to investigate extremal punctured spheres. We prove that the decomposition of an extremal n-punctured sphere into equilateral and horocyclic ideal triangles dictates a fundamental region for the action of the Fuchsian group which uniformizes it.

To carefully forge this interpretation we begin by deriving a correspondence between such decompositions and trees with $(n-2)$ vertices of valence at most 3. This correspondence, later shown to be bijective, moreover proves to be a key ingredient in resolving the question pertaining to uniqueness of extremal disks.

Remark: For the remainder of our endeavors, any subsequent mention of "decomposition" refers to a decomposition of a hyperbolic extremal n-punctured sphere into equilateral and horocyclic ideal triangles according to Theorem 2.3. For ease of exposition, we may henceforth be fairly liberal with dropping the necessary but lengthy descriptors.

3.1 DUAL TREES AND FUNDAMENTAL DECOMPOSITIONS

For any n-punctured sphere decomposition, we can encode the layout of the equilaterals (and incidentally the layout of horocyclic ideal triangles as well) through the dual graph, which we formally define as follows:

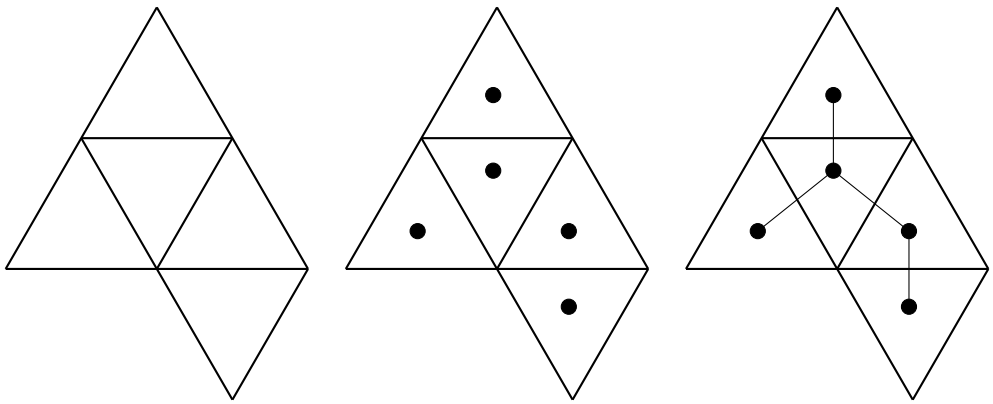

Figure 8: Constructuring the dual tree on a sample configuration of equilateral triangles

Definition 3.1 (Dual Graph)**.** *For any hyperbolic punctured sphere decomposed into equilateral and horocyclic ideal triangles, the dual graph is the topological dual to the subgraph of the* 1*-skeleton consisting of edges with an equilateral triangle on either side.*

In other words, the (unique) dual graph to any such decomposition has a unique vertex at the center of every equilateral triangle and a unique edge joining the vertices in any pair of adjacent equilateral triangles.

We demonstrate the naturality of this canonical identification procedure in the sequence of diagrams shown in Figure 8 and we capture the component of its essence relevant for our purposes with the following lemma, thereby coining the section's eponymous dual tree.

Lemma 3.2 (Dual Tree)**.** *The dual graph for any extremal* n*-punctured sphere decomposition is a tree on* $(n-2)$ *vertices of valence at most 3. We call this the* ***dual tree****.*

Proof. Note that connectedness is enjoyed by construction, following as an immediate consequence of the connectedness of the configuration of equilateral triangles from any decomposition. Now, since each edge of the dual graph crosses a 1-skeleton edge with a compact triangle on either side, we proceed by counting those edges. We have that three times the number of equilateral triangles, minus the number of those bordering a horocyclic ideal triangle, gives exactly twice the desired number e of edges:

$$2e = 3(n-2) - n = 2n - 6.$$

Hence, the dual graph has precisely $(n-3)$ edges. This is in contrast to its $(n-2)$ vertices by construction (exactly one for each equilateral triangle) so that it therefore has Euler characteristic 1 is hence a tree. As each equilateral triangle shares an edge with at most three others, that each vertex is of valence at most 3 holds naturally. □

In the parlance of Fuchsian groups, the decomposition of a surface into triangles from Theorem 2.3 can be seen as a fundamental region for its uniformizing group's action. Specifically:

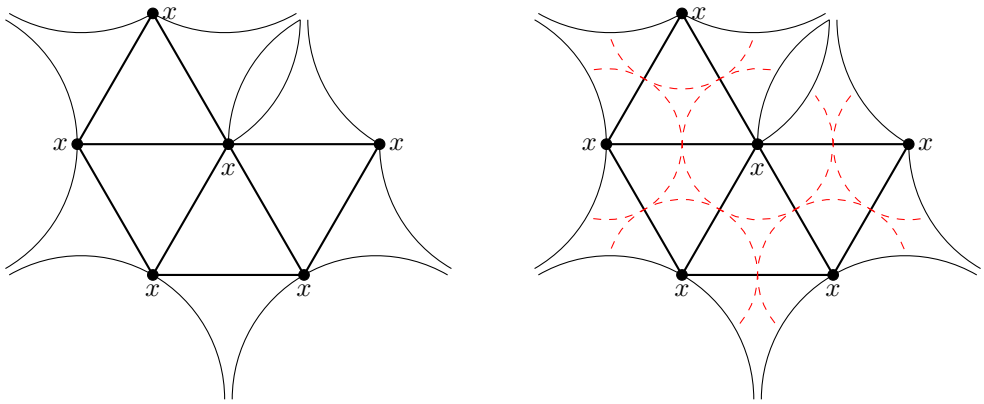

Figure 9: Fundamental region associated to decomposition of an extremal 6-punctured sphere, with extremal disk centered at x shown on the right

Lemma 3.3. *For any extremal n-punctured sphere S with a locally isometric universal cover $\pi : \mathbb{H}^2 \to S$ and extremal disk D, there exists a fundamental domain F in $\mathbb{H}^2$ that is the union of (n-2) equilateral triangles and n horocyclic ideal triangles, such that ∂F is a union of edges of horocyclic ideal triangles, S is obtained by pairing the edges of each horocyclic ideal triangle, and the center of D is the unique quotient vertex.*

Proof. With disk center x, the decomposition of DeBlois' Theorem 2.3 is obtained by projecting the Delaunay tessellation [10] of $\mathbb{H}^2$ determined by $\pi^{-1}(x)$. Let $T \subset S$ be the dual tree constructed in Lemma 3.2. Since T is a tree, each component of $\pi^{-1}(T)$ is a homeomorphic lift. Fix a component $\tilde{T}$ of $\pi^{-1}(T)$ and let F be the union of the equilateral triangles containing $\tilde{T}$ and all horocyclic ideal triangles with which they share an edge. Then F is a fundamental domain, as every triangle in the decomposition of S has a unique lift in F.

We get an n-punctured sphere back from such a decomposition into triangles (as in Figure 9) by identifying adjacent vertices and their corresponding non-compact edge on every compact edge shared by any horocyclic ideal triangle: the Fuchsian group which uniformizes an n-punctured sphere is generated by the n parabolic isometries which perform these identifications. In this way, all copies of x collapse into a single quotient vertex. □

Remark: With the interpretation of the decomposition of a hyperbolic extremal punctured sphere into equilateral and horocyclic ideal triangles as a fundamental region understood, we may hence refer to the two somewhat interchangeably throughout our exposition; when we make mention of a decomposition into triangles here, we are implicitly treating it as such a fundamental region. Reiterating this point: the schematic decomposition depictions herein will be thought of as fundamental regions of the Fuchsian groups which uniformize the surfaces to which they are associated (see, for example, Figure 9).

3.2 DISSECTED N-GONS

One of the principle motivations of the dual tree lens is to derive a bijective correspondence between extremal n-punctured sphere decompositions and some combinatorial objects, in order to achieve a count of the former by enumerating the latter. To enable this, we use the notion of a **dissected n-gon**, whose definition from Brown is a polygon triangulation of type $[0, n-3]$ [5]: a simplicial complex featuring n 0-cells in its boundary, but none in its interior, and with a homeomorphism of its polygon to a disk in the plane. Equivalently, a dissected n-gon can be thought of as a partition of a regular n-gon into $(n-2)$ triangles by non-crossing diagonals (or edges between distinct exterior vertices). With these notions in mind, we state:

Corollary 3.4. *For every pair consisting of an extremal n-punctured sphere and one of its extremal disks, the fundamental domain from Lemma 3.3 determines a dissected n-gon. Isometric surface-disk pairs determine isomorphic dissected n-gons.*

Proof. By Theorem 2.3 we know each extremal surface-disk pair yields a decomposition into equilateral and horocyclic ideal triangles, which determines its dual tree (from Lemma 3.2). This then gives rise to an associated fundamental domain (as shown in Lemma 3.3) which, by its very construction, naturally enjoys the structure of a dissected n-gon: by merely ignoring all edges of every horocyclic ideal triangle, we are left with a simplicial complex (where each triangle is a 2-simplex and each edge a 1-simplex). As its interior deformation retracts to a tree, it is hence simply connected and therefore homeomorphic to a disk by the Riemann Mapping Theorem.

Now, any orientation-preserving isometry of extremal surface-disk pairs preserves the corresponding decompositions into equilateral and horocyclic ideal triangles, so any lift to an isometry of their universal cover takes fundamental domains to fundamental domains. The simplicial complexes (i.e., dissected n-gons) thereby determined by the surfaces are thus indistinguishable (up to orientation-preserving isomorphism). □

We show the identification with an example decomposition in Figure 10.

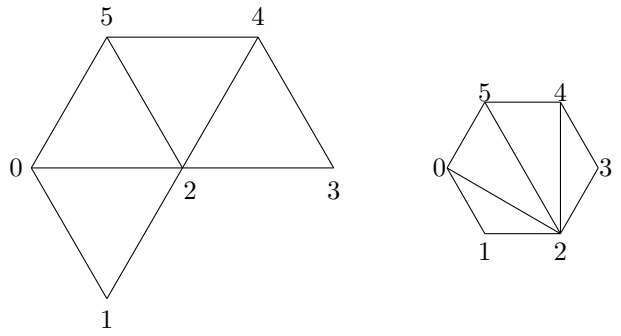

Figure 10: Fundamental region associated to decomposition of an extremal 6-punctured sphere on the left, interpretted as a dissected n-gon on the right

Lemma 3.5. *Every dissected n-gon determines an extremal n-punctured sphere surface-disk pair. Isomorphic dissected n-gons determine isometric surface-disk pairs.*

Proof. Take such a dissected n-gon and associate disjoint equilateral triangles $T_1, T_2, \ldots, T_{n-2}$ in $\mathbb{H}^2$ with side lengths $2r_{0,n}$, one for each distinct triangle formed by its non-crossing diagonals. Take separately enumerated disjoint horocyclic ideal triangles $H_1, H_2, \ldots, H_n$ in $\mathbb{H}^2$ each with compact side length $2r_{0,n}$ and observe that the given dissected n-gon structure dictates an attachment protocol for such a collection of equilateral and horocyclic ideal triangles: triangles T_i and T_j share a unique edge if and only if their associated 2-simplices do; similarly, each edge contained in a single 2-simplex must house a unique horocyclic ideal triangle H_k. Hence, we form a triangulated complex F as a quotient space of $\left(\bigsqcup T_i\right) \cup \left(\bigsqcup H_j\right)$ by identifying these edges of the geometric triangles (according to the manner described above) with an isometry that respects the corresponding pairings of the vertices of their triangles, and then isometrically identifying the two non-compact edges of each H_j.

To then ensure that the resulting identification space object F enjoys the structure of a hyperbolic surface, we now describe a family of chart maps to $\mathbb{H}^2$, with reference to [7]. For each point x in the interior of any T_i or $H_j \subset F$, the local inverse of the quotient map is a homeomorphism from a neighborhood of x in the interior of its triangle to an open subset of $\mathbb{H}^2$. Furthermore, if $x \in F$ is obtained by identifying x_0 in an edge e_0 of triangle T_i with

some $x_1 \in e_1 \subset T_j$ then there exists a unique non-elliptic isometry f_{ij} of $\mathbb{H}^2$ taking e_0 to e_1 (with vertices paired appropriately) such that $f_{ij}(T_i) \cap T_j = e_1$. So take $\epsilon > 0$ which is less than the minimal distance (respectively) from x_0, x_1 to those edges of T_i, T_j that do not contain them, and let U be the projection to F of the disjoint union of $B(x_0, \epsilon) \cap T_i$ with $B(x_1, \epsilon) \cap T_j$. Then $U \subset F$ is an open neighborhood of x with a homeomorphism to $B(x_0, \epsilon)$ and $B(x_1, \epsilon)$ and the overlap maps with the interiors of T_i and T_j are respectively given by f_{ij}^{-1} and the identity. Now, by the definition of a dissected n-gon (specifically the 0 internal vertices denoted by the first entry of the type $[0, n-3]$ triangulation), the quotient space has a single quotient vertex. All vertices (including $3n-6$ from the T_i and $2n$ from the H_j) become identified (since each pair of subsequent vertices are identified by the edge pairing for their associated horocyclic ideal triangle) and the total angle around this quotient vertex is 2π by Theorem 2.3 (since the T_i and H_j were chosen with compact side lengths $2r_{0,n}$). So a chart for this single quotient vertex can be assembled from vertex neighborhoods in the T_i and H_j by composing the f_{ij} in a prescribed order.

To justify completeness, note that the identification of the non-compact edges of each H_j is performed by a parabolic isometry which fixes its ideal point. Hence, any cross-section of H_j by a horocycle with the same ideal point has its edges identified by that same parabolic, which implies that F is complete by Ratcliffe (Theorems 11.1.4 and 11.1.6) [22].

Then F is a complete, finite-area hyperbolic surface, which has a decomposition into $(n-2)$ equilateral and n horocyclic ideal non-overlapping triangles, so that the sufficiency condition of Theorem 2.3 ensures it is indeed an extremal punctured sphere. Any orientation-preserving isomorphism of dissected n-gons dictates an orientation-preserving isometry of the corresponding extremal surfaces triangle-by-triangle. □

We have thus forged our desired bijection between extremal n-punctured sphere surface-disk pairs and dissected n-gons, as the two previous results combine to prove that:

Theorem 3.6. *The collection of pairs, each consisting of an extremal n-punctured sphere and one of its extremal disks (up to orientation-preserving isometry of the surface), is in one-to-one correspondence with the collection of dissected n-gons (up to orientation-preserving isomorphism of the simplicial complex).*

3.3 COUNT OF EXTREMAL PAIRS

Having established the bijection between extremal n-punctured sphere surface-disk pairs and dissected n-gons (in Theorem 3.6), we can hence obtain the desired count of n-punctured sphere extremal surface-disk pairs as a special case of Brown's triangulation formula [5].

Theorem 3.7 (Brown)**.** *The number, $G_{k,m}$, of triangulations of type $[k, m]$ up to orientation-preserving isomorphism is given by:*

$$G_{k,m} = \frac{1}{m+3} \sum_{s|m+3} \phi(s) E_{s;k,m}$$

where ϕ denotes the Euler totient function and all $E_{s;k,m}$ vanish except the following:

$$E_{1;k,m} = \frac{2(2m+3)!(4k+2m+1)!}{m!(m+2)!k!(3k+2m+3)!},$$

$$E_{2;2s+j,2p-1} = \frac{2(2p)!(4s+2p+2j-1)!}{p!(p-1)!s!(3s+2p+2j)!},$$

$$E_{3;3s,3p} = \frac{(2p+1)!(4s+2p)!}{p!p!s!(3s+2p+1)!}.$$

The above statement, condensed for our purposes, can be extracted from [5] as formula (6.3) with $E_{s;k,m}$ defined by equations (4.7), (8.10), and (8.11) respectively. Since our extremal surface-disk pairs are in bijection with triangulations of type $[0, n-3]$:

Theorem 3.8. *The number of pairs consisting of an extremal n-punctured sphere and one of its extremal disks (up to orientation-preserving isometry of the surface) is given by*

$$C(n-2)/n + C(n/2-1)/2 + (2/3)C(n/3-1)$$

where $C(n) = \frac{1}{n+1}\binom{2n}{n} = \frac{(2n)!}{n!(n+1)!}$ (the Catalan numbers) and terms are omitted if their arguments are not integers.

Proof. Applying Brown, we see that the number of triangulations of type $[0, n-3]$ is

$$G_{0,n-3} = \frac{1}{n}\sum_{s|n}\phi(s)E_{s;0,n-3}$$

$$= \begin{cases} \frac{1}{n}\Big(E_{1;0,n-3} + E_{2;0,n-3} + 2E_{3;0,n-3}\Big) & \text{if } \ 2 \mid n, 3 \mid n \\ \frac{1}{n}\Big(E_{1;0,n-3} + E_{2;0,n-3}\Big) & \text{if } \ 2 \mid n, 3 \nmid n \\ \frac{1}{n}\Big(E_{1;0,n-3} + 2E_{3;0,n-3}\Big) & \text{if } \ 2 \nmid n, 3 \mid n \\ \frac{1}{n}E_{1;0,n-3} & \text{if } \ 2 \nmid n, 3 \nmid n \end{cases}.$$

i.e., the term $\frac{E_{2;0,n-3}}{n}$ is omitted from the equation if $\frac{n}{2}$ is not an integer; similarly, the term $\frac{2E_{3;0,n-3}}{n}$ is omitted if $\frac{n}{3}$ is not an integer. It therefore suffices to "match the terms" between Brown's figures and our desired Catalan expressions, or more precisely, to show that:

(1) $\frac{E_{1;0,n-3}}{n} = \frac{C(n-2)}{n}$
(2) $\frac{E_{2;0,n-3}}{n} = \frac{C(n/2-1)}{2}$
(3) $\frac{2E_{3;0,n-3}}{n} = \frac{2C(n/3-1)}{3}$.

We proceed sequentially, beginning with a demonstration of claim (1):

$$\begin{aligned} \frac{E_{1;0,n-3}}{n} &= \frac{2(2n-5)!}{n(n-3)!(n-1)!} \\ &= \frac{2(n-2)(2n-5)!}{n(n-2)(n-3)!(n-1)!} \\ &= \frac{(2n-4)(2n-5)!}{(n-2)(n-3)!n(n-1)!} \\ &= \frac{(2n-4)!}{(n-2)!n!} \\ &= \frac{C(n-2)}{n}. \end{aligned}$$

Next, we establish claim (2):

$$\begin{aligned}
\frac{E_{2;0,n-3}}{n} &= \frac{2(n-3)!}{n(n/2-1)!(n/2-2)!} \\
&= \frac{2(n-2)(n-3)!}{n(n-2)(n/2-1)!(n/2-2)!} \\
&= \frac{2(n-2)(n-3)!}{2(n/2)(n-2)(n/2-1)!(n/2-2)!} \\
&= \frac{(n-2)!}{(n/2)!(n-2)(n/2-2)!} \\
&= \frac{(n-2)!}{(n/2)!2(n/2-1)(n/2-2)!} \\
&= \frac{(n-2)!}{2(n/2)!(n/2-1)!} \\
&= \frac{C(n/2-1)}{2}.
\end{aligned}$$

Finally, we prove claim (3):

$$\begin{aligned}
\frac{E_{3;0,n-3}}{n} &= \frac{2\left(\frac{2}{3}(n-3)\right)!}{n\left(\frac{1}{3}(n-3)\right)!\left(\frac{1}{3}(n-3)\right)!} \\
&= \frac{2\left(\frac{2}{3}(n-3)\right)!}{3(n/3)\left(\frac{1}{3}(n-3)\right)!\left(\frac{1}{3}(n-3)\right)!} \\
&= \frac{2\left(\frac{2}{3}(n-3)\right)!}{3(n/3)!\left(\frac{1}{3}(n-3)\right)!} \\
&= \frac{2C(n/3-1)}{3}.
\end{aligned}$$

So $G_{0,n-3} = C(n-2)/n + C(n/2-1)/2 + (2/3)C(n/3-1)$ (with Catalan terms omitted when their arguments are not integers) and the proof is complete. □

Remark: Presently this formula comprises a count of pairs consisting of an extremal n-punctured sphere and one of its extremal disks (up to orientation-preserving isometry of the surface), rather than (a priori) a firm count of extremal n-punctured spheres themselves. We will however subsequently prove results (namely, the theorems pertaining to uniqueness of extremal disks) which enable us to immediately promote this upper bound to an exact count of extremal punctured spheres.

3.4 CANONICAL HOROBALL CUSP NEIGHBORHOODS

As we have mentioned, the surface triangulation scheme from the characterization of non-compact extremal surfaces is our principle piece of machinery for resolving the question of extremal disk uniqueness for punctured spheres. To get our hands on these extremal surfaces, we apply the triangulation and lift it to $\mathbb{H}^2$ to get a fundamental domain in which we perform the relevant computations directly (as formalized in the preceding section).

To explain how these computational affairs answer our extremal-surface questions, we now describe our methods and this begins by mentioning phenomena actually underlying the characterization theorem all together. Since they do however prove to be core tenants of our general strategy for attacking the disk-uniqueness problem, we pay them due expository respect.

Definition 3.9 (Horoball Cusp Neighborhood)**.** *A horoball cusp neighborhood in a complete hyperbolic surface $S = \mathbb{H}^2/\Gamma$ is a quotient B/P of a horoball $B \subset \mathbb{H}^2$ by the action of a parabolic subgroup $P \leq \Gamma$ fixing the ideal point of B such that the projection $B \to S$ factors through an embedding of B/P.*

This definition allows us to state that first strategic core tenant, which is that:

Proposition 3.10. *On any punctured extremal surface, there is a canonical packing by equal-area horoball cusp neighborhoods. The canonical horoball neighborhoods associated to the decomposition of any extremal surface-disk pair do not depend on the extremal disk.*

Proof. The canonical horoball neighborhood of the cusp enclosed by a horocyclic ideal triangle H is B/P where B is the horoball centered at the ideal vertex of H with boundary tangent to the compact edge of H and P is generated by the parabolic isometry whose fixed point is the ideal vertex of H and which pairs its sides.

If F is a fundamental region, then the quotient of $B \cap F$ obtained by identifying the two sides of intersection with H injects into the surface (by the definition of a fundamental region). The parabolic subgroup P which stabalizes the ideal vertex of H is generated by the side-pairing transformation which identifies the sides of H so that, since $B \cap F$ contains

the entire region in B between these two sides, it is also necessarily a fundamental domain for P.

As the area of the region in B bounded by two vertical geodesics is given by the length of its arc along ∂B, this arc length is in turn determined by the finite side length of H, which in turn depends only on n. Moreover, for a fixed horoball B and a parabolic subgroup P fixing the ideal point of B, if B/P has area A, then for any horoball B' properly contained in B, the area of B'/P is less than A; similarly, if B' properly contains B, then the area of B'/P is greater than A, so that the area of any canonical horoball cusp neighborhood does not depend on the extremal disk determining the decomposition of an extremal surface into equilateral and horocyclic ideal triangles. □

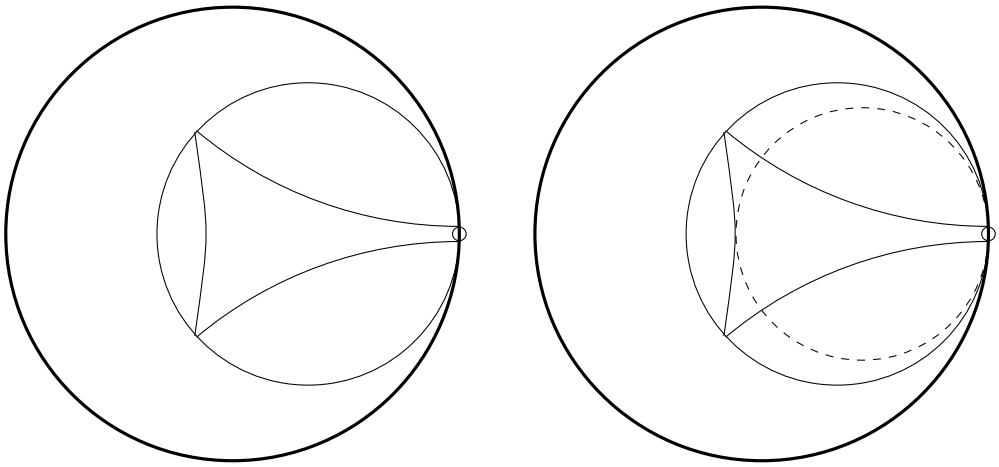

Figure 11: Horocyclic ideal triangle with dashed horoball cusp neighborhood

Note that the canonical horoball collection and any maximal-radius disk's size is preserved by every isometry of the surface. Moreover, courtesy of Bowditch & Epstein we know [4]:

Fact 3.11 (Bowditch, Epstein)**.** *Given two disjoint horocycles in the hyperbolic plane, the locus of points equidistant from the two is geodesic.*

3.5 COMPUTATION LEMMAS

The following lemma is a key ingredient for rigorously establishing our n-punctured sphere uniqueness condition, acting there as the computational link between our dual tree framework and the symmetry reasoning that ultimately guarantees unique extremal disks once $n > 6$. While that component of its significance requires some additional context coming further down the road to fully appreciate, our discourse immediately following it (Lemma 3.14) likewise employs it.

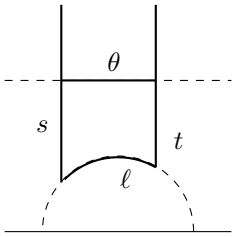

Figure 12: A triangle with one ideal vertex, and a horocycle centered there.

Lemma 3.12 (Bridge Lemma)**.** *For a hyperbolic triangle T with one ideal vertex v and a horocycle C centered at v that intersects the geodesics containing the infinite-length edges of T in an arc of length θ, such that the other vertices of T are at signed distance s and t from C (as depicted in Figure 12), the length ℓ of the compact side of T satisfies*

$$\cosh \ell = \frac{e^{s+t}}{2}\theta^2 + \cosh(s-t)$$

Moreover, if α is the interior angle at the vertex at distance s, then

$$\cosh \ell - \cos \alpha \sinh \ell = e^{t-s}$$

Proof. Applying an isometry, we may take $v = \infty$ and $C = \mathbb{R} + i$. Then we may view the finite vertices of T as $x_0 + iy_0$ and $x_0 + \theta + iy_1$ for some x_0, y_0 and $y_1 \in \mathbb{R}$ which per the $\mathbb{H}^2$ metric must satisfy

$$s = \log 1 - \log y_0 = \log \frac{1}{y_0} \implies e^s = \frac{1}{y_0}$$
$$t = \log 1 - \log y_1 = \log \frac{1}{y_1} \implies e^t = \frac{1}{y_1}.$$

Substituting these into Theorem 1.2.6(ii) of Katok [16] then gives

$$\cosh \ell = 1 + \frac{\theta^2 + (y_0 - y_1)^2}{2y_0y_1} = 1 + \frac{e^{s+t}}{2}\left(\theta^2 + e^{-2s} + e^{-2t} - 2e^{-(s+t)}\right) = \frac{e^{s+t}}{2}\theta^2 + \cosh(s-t)$$

which demonstrates the first assertion. For the second assertion, let β be the interior angle at the vertex at distance t. Then the geodesic arc from β to α is parameterized by

$$x \mapsto (x_c + r\cos x, r\sin x) \qquad \text{for } x \in [\beta, \pi - \alpha]$$

where $(x_c, 0)$ is the center and r the radius of the Euclidean circle containing that geodesic arc. Now, integrating the hyperbolic arc length element yields:

$$\ell = \int_\beta^{\pi-\alpha} \frac{r}{r\sin x}\, dx = \log|\csc\alpha + \cot\alpha| - \log|\csc\beta - \cot\beta|.$$

Then combining and rewriting this gives $e^\ell = \frac{\sin\beta}{\sin\alpha} \cdot \frac{1+\cos\alpha}{1-\cos\beta}$ so that we hence obtain

$$\begin{aligned}
\cosh\ell - \cos\alpha\sinh\ell &= \frac{1}{2}\left(\frac{\sin\beta}{\sin\alpha}\cdot\frac{1+\cos\alpha}{1-\cos\beta}(1-\cos\alpha) + \frac{\sin\alpha}{\sin\beta}\cdot\frac{1-\cos\beta}{1+\cos\alpha}(1+\cos\alpha)\right) \\
&= \frac{1}{2}\left(\frac{\sin\alpha\sin\beta}{1-\cos\beta} + \frac{\sin\alpha}{\sin\beta}(1-\cos\beta)\right) \\
&= \frac{1}{2}\frac{\sin\alpha}{\sin\beta}\left(\frac{\sin^2\beta}{1-\cos\beta} + 1 - \cos\beta\right) \\
&= \frac{1}{2}\frac{\sin\alpha}{\sin\beta}(1+\cos\beta+1-\cos\beta) \\
&= \frac{\sin\alpha}{\sin\beta}
\end{aligned}$$

and since $\sin\alpha = y_0/r = e^{-s}/r$, $\sin\beta = y_1/r = e^{-t}/r$, the result follows. □

Moving towards our next lemma, we mention a result of Bowditch & Epstein [4], which enables us to perform that computation.

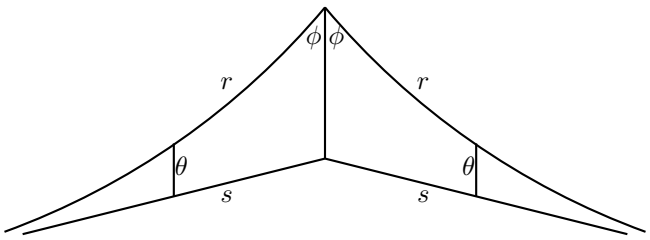

Figure 13: Depiction of the quantities referenced by Theorem 3.13

Theorem 3.13 (Bowditch, Epstein)**.** *With quantities as depicted in Figure 13, we have*

$$\cot\phi = (\theta^2 e^r + e^{r-2s} - e^{-r})/2\theta$$

Since it will be convenient to adopt one perspective or another (context depending) in our forthcoming calculations, we stress here that Figure 15 is a different visualization of the (nearly) same situation depicted in Figure 14, now displayed with more geometric accuracy, in the upper half-plane model of $\mathbb{H}^2$ with its ideal point at ∞ (but with the left half of the original equilateral triangle now not shown).

With this and the result of Bowditch & Epstein in mind, we now state a few computational lemmas, applicable to n-punctured spheres for arbitrary n, but actually used in the subsequent chapter for formally resolving the question of extremal disk uniqueness in the $n = 3, 4, 5$ and 6 extremal n-punctured sphere cases.

Lemma 3.14. *For a hyperbolic extremal n-punctured sphere decomposed into equilateral triangles (with sides of length $2r_{0,n}$) and horocyclic ideal triangles (each with compact edge length $2r_{0,n}$) let $d_{0,n}$ denote the distance between the sole vertex x and any of the horoball neighborhoods within the canonical collection along the geodesic edge of an adjacent horocyclic ideal triangle (as depicted in Figures 14, 15). Then $d_{0,n} = \log\cosh r_{0,n}$.*

Proof. Writing $d = d_{0,n}$, $r = r_{0,n}$ and applying the formula of Theorem 3.13 to the situation best depicted in Figure 15 we get that

$$0 = \cot\tfrac{\pi}{2} = \frac{1}{2\frac{\theta}{2}}((\tfrac{\theta}{2})^2 e^0 + e^{-2d} - e^0) = \frac{1}{\theta}\left(\frac{\theta^2}{4} - 1 + e^{-2d}\right)$$

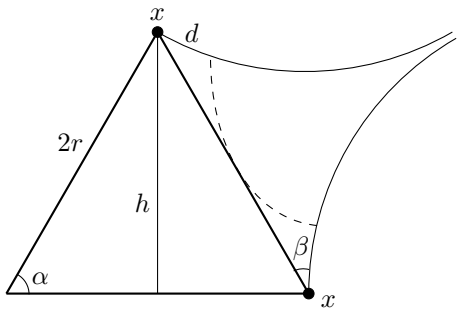

Figure 14: An equilateral hyperbolic triangle with horocyclic ideal triangle

from which we deduce that

$$\frac{\theta^2}{4} = 1 - e^{-2d}. \tag{3.1}$$

Now, by applying our bridge lemma 3.12 to Figure 15's quantities we get

$$\cosh(2r) = \frac{e^{2d}\theta^2}{2} + 1 \tag{3.2}$$

and naturally by the hyperbolic half-angle formulas we have

$$\sinh r = \sqrt{\tfrac{1}{2}(\cosh 2r\ - 1)}. \tag{3.3}$$

Substituting (3.2) into (3.3) we get

$$\sinh r = \sqrt{\frac{1}{2}\left(\frac{e^{2d}\theta^2}{2}\right)} = \frac{e^d\theta}{2}$$

and using (3.1) it follows that

$$\sinh^2 r = \frac{e^{2d}\theta^2}{4} = e^{2d}(1 - e^{-2d}) = e^{2d} - 1$$

so that classic identities yield

$$\cosh^2 r = \sinh^2 r + 1 = e^{2d}$$

and therefore that $e^d = \cosh r$, which proves the assertion. □

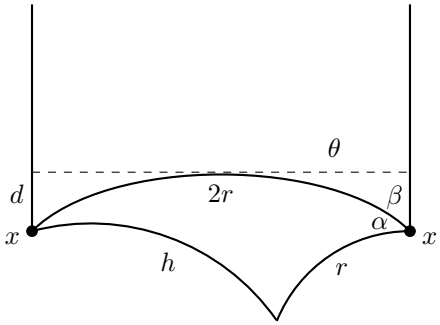

Figure 15: The previous figure taken in the upper half-plane with its ideal point at ∞

Then utilizing this derivation, we are then also afforded that

Lemma 3.15. *For a hyperbolic extremal n-punctured sphere decomposed into equilateral triangles (with sides of length $2r_{0,n}$) and horocyclic ideal triangles (each with compact edge length $2r_{0,n}$) let $h_{0,n}$ denote the shortest distance from the sole vertex x to an opposite equilateral triangle side. Then $d_{0,n} < h_{0,n}$.*

Proof. Writing $h = h_{0,n}$, $r = r_{0,n}$ and using the hyperbolic law of cosines on Figure 14 yields

$$\cosh 2r = \cosh r \cosh h \implies \cosh h = \frac{\cosh 2r}{\cosh r} = \frac{2\cosh^2 r - 1}{\cosh r}$$

Hence, we have that

$$\cosh h = 2\cosh r - \frac{1}{\cosh r}.$$

Now, since $e^d = \cosh r$ from Lemma 3.14 we have by the definition of hyperbolic cosine that

$$\cosh d = \frac{1}{2}\left(\cosh r + \frac{1}{\cosh r}\right)$$

This makes our desired comparison of h and d a simple endeavor. Using our formulas for $\cosh d$ and $\cosh h$ we therefore verify

$$\cosh d < \cosh h \iff \frac{1}{2}\left(\cosh r + \frac{1}{\cosh r}\right) < 2\cosh r - \frac{1}{\cosh r} \iff \frac{1}{\cosh r} < \cosh r$$

The final inequality above is equivalent to $1 < \cosh r$ which holds naturally for $r > 0$. Noting that hyperbolic cosine is increasing on $[0, \infty)$ alongside this explicit demonstration that $\cosh d < \cosh h$, we conclude that indeed $d < h$. □

Lemma 3.16. *For a hyperbolic extremal n-punctured sphere decomposed into equilateral triangles and horocyclic ideal triangles (of side lengths as in Lemma 3.14), the shortest distance from the sole vertex x to a canonical horoball neighborhood is attained exactly along edges of the fundamental domain F associated (as from Lemma 3.3).*

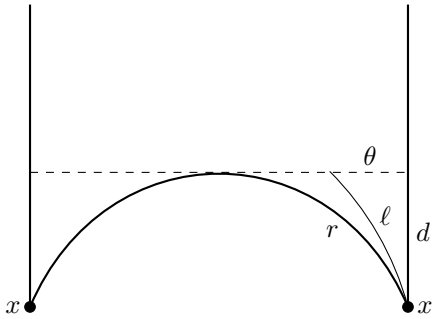

Figure 16: Demonstration that α does not minimize distance in Lemma 3.16, case 2

Proof. Let α be a geodesic arc from x to a horoball neighborhood B.

(Case 1) Suppose that α points into an equilateral triangle T. Since each horoball neighborhood is contained in a horocyclic ideal triangle, we see that α traverses a distance of at least $h_{0,n}$ before exiting T. As $h_{0,n} > d_{0,n}$ by Lemma 3.15, α hence fails to minimize distance.

(Case 2) Suppose that instead α points into a horocyclic ideal triangle H. Then as depicted in Figure 16, the distance along α to the nearest horoball neighborhood is ℓ which by the Bridge Lemma 3.12 satisfies

$$\cosh \ell = \frac{e^d}{2}\theta^2 + \cosh d$$

and since $\theta > 0$ we conclude $\ell > d$. □

Corollary 3.17. *For a hyperbolic extremal n-punctured sphere decomposed into equilateral triangles and horocyclic ideal triangles (of side lengths as in Lemma 3.14), each extremal disk center is at distance exactly $d_{0,n}$ from every member of the canonical horoball packing associated to the surface.*

Proof. By Lemma 3.16 the minimizing distance is attained along edges of the fundamental domain, with paths of length $d_{0,n}$. For a fixed extremal punctured sphere, the canonical horoball neighborhood collection associated to the extremal disk-surface pair does not depend on the disk (Proposition 3.10). □

Corollary 3.18. *For a hyperbolic extremal n-punctured sphere decomposed into equilateral triangles and horocyclic ideal triangles (of side lengths as in the previous Lemma 3.14) with associated fundamental region F (as in Lemma 3.3), the shortest distance from any point p on $\pi(\partial F) - \{x\}$ to the canonical horoball packing occurs only along the edge that contains it.*

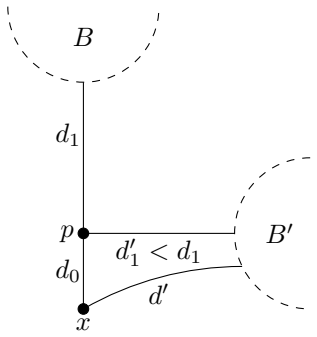

Figure 17: Depiction of Corollary 3.18

Proof. Otherwise, we contradict Lemma 3.16 by demonstrating a path between the known disk center x and a horoball neighborhood within the canonical collection of length less than $d_{0,n}$. To see this, take a point $p \in \pi(\partial F) - \{x\}$, at distance d_1 from horoball B along the edge of ∂F and suppose there is a horoball neighborhood B' at distance $d_1' \leq d_1$ from point p (portrayed in Figure 17). Then disk center x is at distance d' from B' where d' contradicts Lemma 3.16 by satisfying

$$d' < d_0 + d_1' \leq d_0 + d_1 = d$$

Note that the first inequality above holds by virtue of d' being a geodesic, hence locally length-minimizing, path from x to B' (whereas the path from x to B' through p, of length $d_0 + d_1'$, has a corner at p). □

Corollary 3.19. *For a hyperbolic extremal n-punctured sphere decomposed into equilateral triangles and horocyclic ideal triangles (of side lengths as in the previous Lemma 3.14) with associated fundamental region F (as in Lemma 3.3) and any point $p \in int(F)$, the distance from $\pi(p)$ to the canonical horoball packing occurs only along projections of arcs in F.*

Proof. Through repeating the argument from Corollary 3.18 at the boundary point where the path exits F we see that otherwise we contradict Lemma 3.16 by demonstrating a path between the known disk center x and a horoball neighborhood within the canonical collection of length less than $d_{0,n}$. □

The real upshot of these demonstrations is that we are warranted in performing all of the computational affairs pertaining to extremal disk uniqueness entirely within the scope of the fundamental region (as in Lemma 3.3) derived from an extremal punctured sphere's decomposition into equilateral and horocyclic ideal triangles. In particular, the locus of points equidistant from any pair of horoball neighborhoods is a geodesic segment contained within the interior of such a fundamental region and any distinct such geodesics intersect in a unique point there.

This justifies a straight-forward computational attack on the extremal disk uniqueness problem alongside our previously mentioned strategic tenant, that is, the mandatory condition described by Proposition 3.10 and Corollary 3.17: that any extremal disk center must be equidistant from every horoball neighborhood within the canonical collection associated.

4.0 NON-UNIQUENESS OF EXTREMAL DISKS

Now we demonstrate that some n-punctured spheres do admit multiple extremal disks, in a few particular cases featuring relatively low cusp counts. With careful computational analysis, we identify that extremal 3- and 4-punctured spheres always have precisely two extremal disks. Moreover, certain decompositions of extremal 6-punctured spheres (since now, unlike with the lower-n predecessors, there are multiple possible decompositions into equilateral and horocyclic ideal triangles) do in fact also allow for a second extremal disk.

As mentioned, the surface triangulation scheme from DeBlois' Theorem 2.3 is our principle piece of machinery for this computational approach for resolving the question of extremal disk uniqueness. Notice that

Fact 4.1. *The decomposition of DeBlois' Theorem 2.3 applied to an extremal n-punctured sphere yields $(n-2)$ equilateral triangles and n horocyclic ideal triangles.*

For each of the cases where $n = 3, 4, 5$ and 6, we will view the corresponding n-punctured sphere through the lens of this triangulation, namely, as and in terms of a corresponding fundamental region (Lemma 3.3) for a Fuchsian group generated by n parabolic isometries as we unveil its extremal-disk mysteries.

Along the way, performing the same type of analysis on 5-punctured spheres shows they must have unique disks (and the pattern we observe there generalizes to eventually comprise our wider-reaching uniqueness theorem), but we showcase this in the apropos chapter on uniqueness of extremal disks, and now present our specific findings for the special n-punctured sphere cases of $n = 3, 4$, and 6.

4.1 3-PUNCTURED SPHERES

Here we show that the unique extremal 3-punctured sphere has precisely two extremal disks. To this end, we first note that

Fact 4.2. *For each point on any 3-punctured sphere, there exists an isometry of the surface which does not fix it.*

Proof. This follows from the observation that the group of its self-isometries has no global fixed point. More specifically, we can uniformize the 3-punctured sphere through the (Fuchsian) principal level-2 congruence normal subgroup of the modular group $\mathrm{PSL}(2,\mathbb{Z})$ (i.e., the kernel of the map which reduces entries mod 2) and consider the $\mathrm{PSL}(2,\mathbb{Z})$ maps

$$\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}; \begin{pmatrix} 0 & 1 \\ -1 & 1 \end{pmatrix}$$

whose only (respective) fixed points are i and $\frac{1+i\sqrt{3}}{2}$ in the upper half-plane; one can show that in addition these maps share no fixed points on the uniformized surface itself. □

Alongside the utilization of any such self-isometry not fixing the extremal 3-punctured sphere's known extremal disk center (which in and of itself detects a second extremal disk), we have a computational means by which to locate an alternate extremal disk: within the fundamental domain (Lemma 3.3) induced by the extremal 3-punctured sphere's decomposition into equilateral and horocyclic ideal triangles, it is enough to identify another point equidistant from the corresponding horoball neighborhood collection (Proposition 3.10): as proven in the previous chapter, the center of any extremal disk must be located at a spot equidistant from the canonical collection of horoball neighborhoods.

We begin with an explicit computation of the maximal injectivity radius at a point on the 3-punctured sphere: the quantity $r = r_{0,3}$ from DeBlois' Theorem 2.3. Recall that by Corollary 3.17, this quantity dictates the distance $d = d_{0,3}$ between any extremal disk center and any horoball neighborhood through the relation $d = \log\cosh r$ supplied by Lemma 3.14.

Lemma 4.3. $\cosh r = \cosh r_{0,3} = \sqrt{\frac{7}{3}}$

Proof. Defined in terms of $\alpha = 2\sin^{-1}(\frac{1}{2\cosh r})$ and $\beta = \sin^{-1}(\frac{1}{\cosh r})$, the angle formula from Theorem 2.3 for r on a genus g surface with n cusps states:

$$(4g + n - 2)3\alpha + 2n\beta = 2\pi$$

so that in the case of the 3-punctured sphere we have that $\cosh r$ satisfies

$$(4(0) + 3 - 2)3\alpha + 2(3)\beta = 2\pi$$

$$3\alpha + 6\beta = 2\pi$$

$$\tfrac{1}{2}\alpha + \beta = \frac{\pi}{3}$$

$$\text{i.e.,} \quad \sin^{-1}\left(\tfrac{1}{2\cosh r}\right) + \sin^{-1}\left(\tfrac{1}{\cosh r}\right) = \frac{\pi}{3} \tag{4.1}$$

Now, from standard trigonometric identities we know that

$$\sin^{-1} x + \sin^{-1} y = \sin^{-1}\left(x\sqrt{1-y^2} + y\sqrt{1-x^2}\right).$$

When applied above while adopting the convention $c := \cosh r$, equation (4.1) becomes

$$\sin^{-1}\left(\frac{1}{2c}\sqrt{1-\frac{1}{c^2}} + \frac{1}{c}\sqrt{1-\frac{1}{4c^2}}\right) = \frac{\pi}{3}$$

and after applying sine it reduces to

$$\frac{\sqrt{1-\frac{1}{c^2}}}{2c} + \frac{\sqrt{1-\frac{1}{4c^2}}}{c} = \frac{\sqrt{3}}{2}$$

$$\sqrt{c^2-1} + \sqrt{4c^2-1} = \sqrt{3}c^2$$

$$4(c^2-1)(4c^2-1) = (3c^4 - 5c^2 + 2)^2$$

which, under the requirement that $r > 0$, yields that $c = \cosh r = \sqrt{\frac{7}{3}}$. [27] □

As demonstrated in Section 3.5, the collection of points equidistant from any pair of horoball neighborhoods is a geodesic arc entirely contained within the fundamental domain associated to the decomposition, and any pair of such geodesic arcs has a single point of intersection within that fundamental domain.

We hence have a natural candidate for an alternate extremal disk location: in the interior of the single equilateral triangle of the fundamental region (Lemma 3.3) induced by the associated decomposition, where any two of the equidistant loci running between distinct pairs of horoball neighborhoods intersect (and by symmetry of that equilateral triangle, all three equidistant loci necessarily intersect in but a single point). We depict this candidate in Figure 18 as point p_3, at distance labelled y from the canonical horoball neighborhoods, and we show that this location houses the only alternate extremal disk in the 3-punctured sphere setting by proving that $y = d$; as there are no other points equidistant from each neighborhood within the canonical horoball collection, there are no other possible disk-center candidates to check.

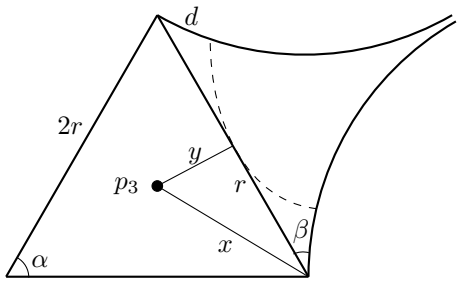

Figure 18: An equilateral triangle with horocyclic ideal triangle in 3-punctured sphere

To summarize: since the extremal 3-punctured sphere admits a self-isometry that does not fix the center of its known extremal disk, Corollary 3.17 guarantees p_3 is the center of the extremal 3-punctured sphere's second (and final) extremal disk upon verification that $y = d$.

Lemma 4.4 ($y = d$ for $n = 3$)**.** *With quantities as depicted in Figure 18, $y = d$.*

Proof. We verify that the following system, derived through applications of Theorem 1.5.2 of Katok [16] (specifically Theorem 1.5.2(i) for (4.3) and Theorem1.5.2(ii) for (4.2)) on the situation depicted in Figure 18 is satisfied when $y = d$. Recall that per Lemma 3.14, $d_{0,n}$ is determined by $r_{0,n}$ through the relation $d_{0,n} = \log\cosh r_{0,n}$ for all n.

$$\cosh r \cosh y = \cosh x \tag{4.2}$$

$$\frac{\sinh x}{\sin\frac{\pi}{2}} = \frac{\sinh r}{\sin\frac{\pi}{3}}. \tag{4.3}$$

Evaluating these before combining and rearranging them results in the equation

$$\cosh y = \frac{\cosh x}{\cosh r} = \frac{\sqrt{1+\frac{4}{3}\sinh^2 r}}{\cosh r} = \frac{\sqrt{1+\frac{4}{3}(\cosh^2 r - 1)}}{\cosh r}$$

which, using our previous computation $\cosh r = \sqrt{\frac{7}{3}}$, gives that $\cosh y = \frac{5}{\sqrt{21}}$. Therefore,

$$y = \cosh^{-1}\left(\frac{5}{\sqrt{21}}\right) = \log\left(\frac{5}{\sqrt{21}} + \sqrt{\frac{25}{21} - 1}\right) = \log\frac{7}{\sqrt{21}} = \log\sqrt{\frac{7}{3}} = d.$$

□

For 3-punctured spheres, the DeBlois tessellation yields but a single equilateral triangle (per Fact 4.1). Thus the existence of a second extremal disk that we explicitly witness through the above demonstration exhaustively covers the 3-punctured sphere case: there are no alternative decomposition configurations to consider.

Theorem 4.5. *The unique extremal 3-punctured sphere has precisely two extremal disks and these are exchanged by a self-isometry of the surface.* □

4.2 4-PUNCTURED SPHERES

As our next exemption from extremal disk uniqueness, we now examine the particular case of the 4-punctured sphere. We adopt similar methodology to the 3-punctured case and include the details of that rigorous computational demonstration, but first note that totally naive arguments again guarantee the existence of multiple disks in this case. More specifically, from Farb & Margalit [11], alternatively Ruberman [23], we know that

Fact 4.6. *Any 4-punctured sphere admits multiple orientation-preserving self-isometries (known as hyperelliptic involutions), each of which fix a distinct pair of points.*

So indeed, the hyperelliptic involution which does not fix the known extremal disk center sends it to what must be a newly discovered extremal disk (as such a maximal-radius disk's size is necessarily preserved by any isometry of the surface). To numerically verify that second extremal disk's location (within the fundamental domain given by the surface's decomposition into equilateral and horocyclic ideal triangles) and determine there is no third, we compute the maximal injectivity radius of a point on a 4-punctured sphere, $r = r_{0,4}$, showing that it satisfies

Lemma 4.7. $\cosh r = \cosh r_{0,4} = \frac{5+\sqrt{17}}{4}$

Proof. Applying the formula from Theorem 2.3 with $g = 0$ and $n = 4$ shows $\cosh r$ satisfies

$$(4(0) + 4 - 2)3\alpha + 2(4)\beta = 2\pi \implies 6\alpha + 8\beta = 2\pi \implies 3\alpha + 4\beta = \pi$$

$$\text{i.e., } 3\sin^{-1}\left(\tfrac{1}{2\cosh r}\right) + 2\sin^{-1}\left(\tfrac{1}{\cosh r}\right) = \frac{\pi}{2} \tag{4.4}$$

Now, from trigonometric identities (e.g., double-angle formulas) we know that

$$3\sin^{-1} x = \sin^{-1}\left(3x - 4x^3\right)$$

$$2\sin^{-1} x = \sin^{-1}\left(2x\sqrt{1-x^2}\right).$$

Continuing our convention of letting $c = \cosh r$ while applying these transforms (4.4) into

$$\sin^{-1}\left(\frac{3}{2c} - \frac{1}{2c^3}\right) + \sin^{-1}\left(\frac{2}{c}\sqrt{1-\frac{1}{c^2}}\right) = \frac{\pi}{2}$$

but we also know (and in fact witnessed in our 3-punctured investigations) that

$$\sin^{-1} x + \sin^{-1} y = \sin^{-1}\left(x\sqrt{1-y^2} + y\sqrt{1-x^2}\right)$$

so that we can rewrite our equation yet again and apply sine on both sides to get

$$\frac{3c^2-1}{2c^3} \cdot \frac{\sqrt{c^4-4c^2+4}}{c^2} + \frac{\sqrt{4c^2-4}}{c^2} \cdot \frac{\sqrt{4c^6-9c^4+6c^2-1}}{2c^3} = 1$$

$$\text{or} \quad (3c^2-1)\sqrt{c^4-4c^2+4} + \sqrt{4c^2-4}\sqrt{4c^6-9c^4+6c^2-1} = 2c^5$$

whose only solution (under the stipulation that r be positive) is $c = \frac{5+\sqrt{17}}{4}$. [27] □

Now that we have identified d numerically (recall that by Lemma 3.14, d is determined by r through the relation $d = \log\cosh r$) we can use this as an explicit means by which to verify (or rule out) candidates for alternate extremal disk centers. Knowing that any disk center must lie at distance precisely d from each of the canonical horoball neighborhoods again establishes an obvious candidate to check: where the neighborhoods' equidistant loci intersect on the shared edge of the equilateral triangles in this case (see Figures 19 and 20, where the candidate point is labeled as p_4). We hence call that unknown distance (between p_4 and the horoball collection) y and describe it through of a system of equations utilizing Theorem 3.13, hoping it proves to be consistent with the d we determined above [4].

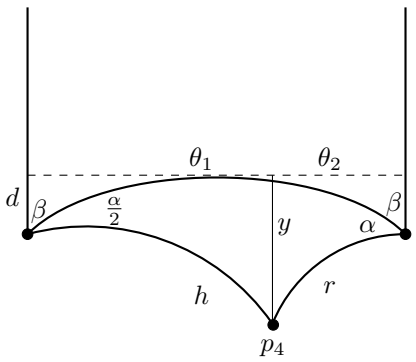

Figure 19: Horocycle view of a cusp in 4-punctured sphere case

Lemma 4.8 ($y = d$ for $n = 4$)**.** *With quantities as depicted in Figures 19, 20 $y = d$.*

Proof. By Theorem 3.13 we know that $y = d$ if and only if $d = \log\left(\frac{5+\sqrt{17}}{4}\right)$, as directly computed above, satisfies the system of equations

$$\cot\left(\frac{\alpha}{2} + \beta\right) = \frac{e^d\theta_1}{2} \tag{4.5}$$

$$\cot(\alpha + \beta) = \frac{e^d\theta_2}{2} \tag{4.6}$$

$$\theta_1 + \theta_2 = 2e^{-d}\sinh r \tag{4.7}$$

Using the fact that $\sinh r = \sqrt{c^2 - 1}$ and (4.7) in order to combine (4.5) with (4.6) we get

$$\cot(\alpha + \beta) = \frac{e^d}{2}\theta_2 = \frac{e^d}{2}\left(\frac{2\sqrt{c^2-1}}{c} - \theta_1\right) = \sqrt{c^2-1} - \frac{e^d\theta_1}{2} = \sqrt{c^2-1} - \cot\left(\frac{\alpha}{2} + \beta\right).$$

Now, using trigonometric identities and simplifying expressions we get

$$\cot\left(\frac{\alpha}{2} + \beta\right) = \frac{\sqrt{4c^2-1}\sqrt{c^2-1} - 1}{\sqrt{4c^2-1} + \sqrt{c^2-1}}$$

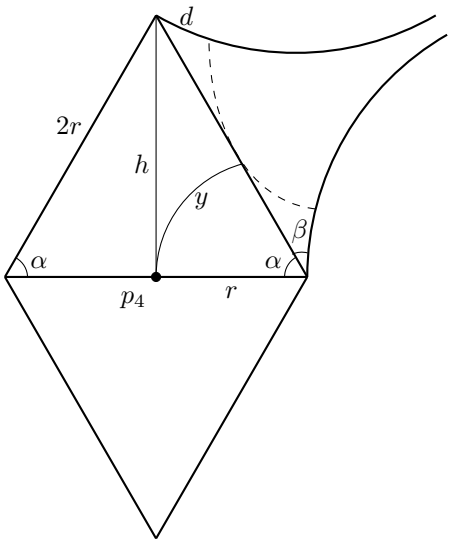

Figure 20: Alternate view of horocyclic ideal triangle for a 4-punctured sphere

$$\cot(\alpha+\beta)=\frac{(2c^2-1)\sqrt{c^2-1}-\sqrt{4c^2-1}}{2c^2-1+\sqrt{(4c^2-1)(c^2-1)}}$$

so that our system then reduces to verifying with our $c=\frac{5+\sqrt{17}}{4}$ that

$$\frac{(2c^2-1)\sqrt{c^2-1}-\sqrt{4c^2-1}}{2c^2-1+\sqrt{(4c^2-1)(c^2-1)}}=\frac{c^2}{\sqrt{4c^2-1}+\sqrt{c^2-1}}$$

which Mathematica verifies, establishing candidate point p_4 as a bona fide disk center. □

So our direct computation formally identifies the hidden disk whose existence our naive argument detected. As only a single decomposition into equilateral and horocyclic ideal triangles is possible here (like when $n=3$ and until $n>5$), we conclude that

Theorem 4.9. *The unique extremal* 4*-punctured sphere has precisely two extremal disks and these are exchanged by a self-isometry of the surface.* □

4.3 6-PUNCTURED SPHERES

As a final exception to extremal disk uniqueness, we zoom in on the case of 6-punctured spheres. Reminiscent of the 3- and 4-punctured cases, some 6-punctured spheres do have multiple extremal disks; however, unlike in those previous cases, there are multiple possible triangulations of 6-punctured spheres (see Figure 25). As it turns out, only one admits multiple disks (and perhaps unsurprisingly by this point, the one that presents the most symmetry). We focus our attention on that highly symmetric decomposition, dubbed 6_S and shown in Figure 22, beginning the search for multiple disks by explicitly solving for $r_{0,6}$.

Lemma 4.10. $\cosh r = \cosh r_{0,6} = 1 + 2^{1/3} + 2^{2/3}$

Proof. Once again applying the angle formula of Theorem 2.3, now with $n = 6$, gives that

$$12\alpha + 12\beta = 2\pi \implies \alpha + \beta = \frac{\pi}{6}$$

$$\text{i.e., } \; 2\sin^{-1}\left(\tfrac{1}{2\cosh r}\right) + \sin^{-1}\left(\tfrac{1}{\cosh r}\right) = \frac{\pi}{6} \tag{4.8}$$

and using identities as in the 4-punctured sphere (and continuing with the convention of taking $c = \cosh r$), we solve for c. For this endeavor, we reuse the facts

$$2\sin^{-1} x = \sin^{-1}\left(2x\sqrt{1-x^2}\right) \quad \text{and} \quad \sin^{-1} x + \sin^{-1} y = \sin^{-1}\left(x\sqrt{1-y^2} + y\sqrt{1-x^2}\right)$$

Then formula (4.8) hence becomes

$$\sin^{-1}\left(\frac{\sqrt{4c^2-1}\sqrt{c^2-1}}{2c^3} + \frac{2c^2-1}{2c^3}\right) = \frac{\pi}{6}$$

which after applying sine, nicely reduces to the equation

$$\sqrt{4c^2-1}\sqrt{c^2-1} + 2c^2 - 1 = c^3$$

whose only solution (when subjected the constraint that r be real and positive) is that $c = 1 + 2^{1/3} + 2^{2/3}$ [27]. □

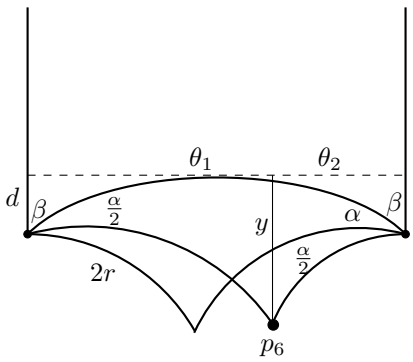

Figure 21: Horocyclic ideal triangle for extremal 6-punctured sphere under 6_S

Lemma 4.11 ($y = d$ for $n = 6$)**.** *As depicted in Figure 21, 22 $y = d$.*

Proof. With quantities as depicted in the upper half-plane Figure 21, we establish a candidate disk center at point p_6 through an explicit demonstration that $y = d$ (as the proposition's namesake suggests). Using Bowditch-Epstein again, this amounts to verifying that the system of equations below is satisfied by the $d = \log\left(1 + 2^{1/3} + 2^{2/3}\right)$ we just computed.

$$\cot\left(\frac{\alpha}{2} + \beta\right) = \frac{e^d \theta_1}{2} \tag{4.9}$$

$$\cot\left(\frac{3\alpha}{2} + \beta\right) = \frac{e^d \theta_2}{2} \tag{4.10}$$

$$\theta_1 + \theta_2 = 2e^{-d} \sinh r \tag{4.11}$$

First we address (4.10) by using trigonometric identities. Here we get:

$$\cot\left(\frac{3\alpha}{2} + \beta\right) = \frac{\cot(\frac{3\alpha}{2})\cot(\beta) - 1}{\cot(\frac{3\alpha}{2}) + \cot(\beta)} = \frac{\cot(3\sin^{-1}(\frac{1}{2c}))\cot(\sin^{-1}(\frac{1}{c})) - 1}{\cot(3\sin^{-1}(\frac{1}{2c})) + \cot(\sin^{-1}(\frac{1}{c}))}$$

$$= \frac{\frac{\sqrt{4c^2-1}(c^2-1)}{3c^2-1}\sqrt{c^2-1} - 1}{\frac{\sqrt{4c^2-1}(c^2-1)}{3c^2-1} + \sqrt{c^2-1}} = \frac{\sqrt{4c^2-1}(c^2-1)^{3/2} - 3c^2 + 1}{\sqrt{4c^2-1}(c^2-1) + \sqrt{c^2-1}(3c^2-1)}$$

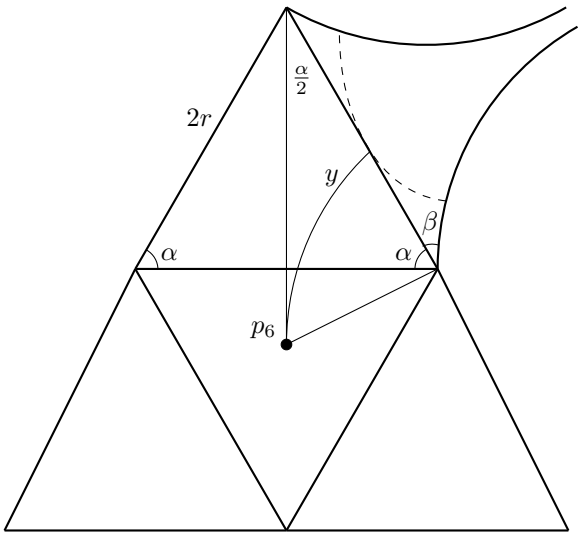

Figure 22: Highly symmetric decomposition 6_S of an extremal 6-punctured sphere

Fortunately we can inherit both the simplification of (4.9) and the means of combining with (4.10) through (4.11) from the 4-punctured case. Accordingly, $y = d$ here is equivalent to the below equation being satisfied with $c = 1 + 2^{1/3} + 2^{2/3}$

$$\frac{\sqrt{4c^2-1}(c^2-1)^{3/2} - 3c^2 + 1}{\sqrt{4c^2-1}(c^2-1) + \sqrt{c^2-1}(3c^2-1)} = \frac{c^2}{\sqrt{4c^2-1} + \sqrt{c^2-1}}$$

which holds [27], showing that candidate point p_6 is in fact also at distance d from each neighborhood in the canonical horoball collection, like the known disk center x. □

Unlike its punctured sphere predecessors featuring non-unique extremal disks, which entered the fray sporting known self-isometries (the hyperelliptic involutions of the 3- and 4-punctured cases), 6-punctured spheres do not a priori enjoy any such isometries. Hence, different from the previous instances of extremal disk non-uniqueness, we now argue for an isometry taking the known disk center x to our candidate p_6 in order to definitively conclude we have a second extremal disk here.

To this end, we take a Fuchsian model for a 6-punctured sphere whose fundamental domain respects this particular symmetric decomposition, and manually construct the desired isometry relative to it. Specifically, let $\Lambda = \langle \phi_0, \ldots, \phi_5 \rangle$ be the Fuchsian group with a

fundamental domain shown in Figure 23, generated by 6 parabolic isometries, each with a single fixed point x_i for ϕ_i on the boundary ($\mathbb{H}^2/\Lambda$ is homeomorphic to a 6-punctured sphere, $S^2 - \{x_0, \ldots, x_5\}$). Then we have:

Lemma 4.12 (Existence of Disk Isometry for Symmetric 6-punctured Sphere)**.** *There exist a self-isometry of* $\mathbb{H}^2/\Lambda$ *taking the known disk center to the candidate center point* p_6*.*

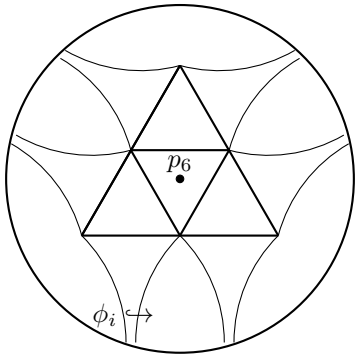

Figure 23: Schematic depiction of a fundamental domain for Λ shown in the disk model

Proof. Let $\Lambda = \langle \phi_0, \ldots, \phi_5 \rangle$ be the Fuchsian group with a fundamental domain shown in Figure 23, generated by the 6 parabolic isometries which identify the noncompact sides of each horocyclic ideal triangle (enumerated counter-clockwise). Take ρ to be the order-3, counter-clockwise rotation about the center of the fundamental domain for Λ, and consider a fundamental region for this order-3 action of ρ, depicted in Figure 24. Since we observe

$$\rho \phi_i \rho^{-1} = \phi_{i+2} \pmod 6$$

we have that ρ normalizes Λ. Hence, we consider Λ as an index-3 normal subgroup $\Lambda \trianglelefteq \Gamma$ where by the Poincaré Polygon Theorem Γ has presentation wholly dictated by its cycle relations. Namely, as x is fixed by the order-3 element $\phi_1 \rho^{-1} \phi_2$ and p_6 by ρ, we have that

$$\Gamma = \langle \rho, \phi_1, \phi_2 \mid (\phi_1 \rho^{-1} \phi_2)^3 = \rho^3 = 1 \rangle.$$

It is on this quotient orbifold $\mathbb{H}^2/\Gamma$ that we uncover our desired isometry.

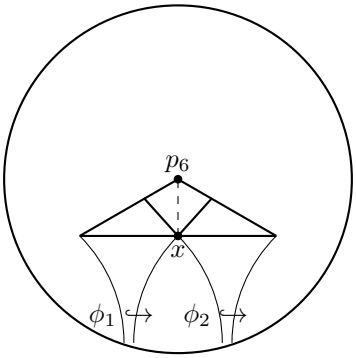

Figure 24: Fundamental domain for Γ shown in the Poincaré disk model

Let τ be the 180° degree rotation about the midpoint of the geodesic arc joining x and p_6. From our $y = d$, part 6 computation (Lemma 4.11) we see that the action of τ on $\mathbb{H}^2/\Gamma$ exchanges the cusps of quotient orbifold. Hence we have that

$$\tau\rho\tau^{-1} = \phi_1\rho\phi_2$$
$$\tau\phi_1\tau^{-1} = \phi_2$$
$$\tau\phi_2\tau^{-1} = \phi_1$$

Hence, τ normalizes Γ and thereby induces a self-isometry of the quotient orbifold $\mathbb{H}^2/\Gamma$, which sends x to p_6. Here we make use of the fact that an $\mathbb{H}^2$ isometry ι induces a self-isometry of a given surface S precisely if it normalizes the Fuchsian group uniformizing S. In other words, $\iota \in \mathrm{PSL}(2,\mathbb{R})$ induces an isometry on the quotient surface $\mathbb{H}^2/\Gamma$ if and only if $\iota T\iota^{-1} \in \Gamma$ for every $T \in \Gamma$. In this way, any points $x, y \in \mathbb{H}^2$ identified in the quotient surface (i.e., such that there exists $T \in \Gamma$ such that $y = Tx$) are still identified following the action by ι, as

$$\iota y = \iota T x = \iota T \iota^{-1}\iota x = T'\iota x \quad \text{for some} \quad T' \in \Gamma$$

So in order to show that the quotient orbifold isometry τ, which sends x to p_6, respects the 6-punctured sphere uniformized by Λ and therefore provides the isometry we seek, we need

to ensure that τ normalizes Λ. Now, as we have already observed (Figure 23)

$$\begin{aligned}\phi_3 &= \rho\phi_1\rho^{-1}\\ \phi_4 &= \rho\phi_2\rho^{-1}\\ \phi_5 &= \rho^2\phi_1\rho^{-2}\\ \phi_0 &= \rho^2\phi_2\rho^{-2}\end{aligned}$$

hence, τ does normalize ϕ_3:

$$\begin{aligned}\tau\phi_3\tau^{-1} &= \tau(\rho\phi_1\rho^{-1})\tau^{-1}\\ &= (\tau\rho\tau^{-1})(\tau\phi_1\tau^{-1})(\tau\rho^{-1}\tau^{-1})\\ &= (\phi_1\rho\phi_2)\phi_2(\phi_1\rho\phi_2)^{-1}\\ &= \phi_1(\rho\phi_2\rho^{-1})\phi_1^{-1}\\ &= \phi_1\phi_4\phi_1^{-1} \in \Lambda\end{aligned}$$

and similarly τ normalizes ϕ_4:

$$\begin{aligned}\tau\phi_4\tau^{-1} &= \tau(\rho\phi_2\rho^{-1})\tau^{-1}\\ &= (\tau\rho\tau^{-1})(\tau\phi_2\tau^{-1})(\tau\rho^{-1}\tau^{-1})\\ &= (\phi_1\rho\phi_2)\phi_1(\phi_1\rho\phi_2)^{-1}\\ &= \phi_1\rho(\phi_2\phi_1\phi_2^{-1})\rho^{-1}\phi_1^{-1} \quad \text{and } \rho \text{ normalizes } \Lambda\\ &= \phi_1\lambda_4\phi_1^{-1} \quad \text{for some } \lambda_4 \in \Lambda\\ &\in \Lambda\end{aligned}$$

τ normalizes ϕ_5:

$$\begin{aligned}
\tau\phi_5\tau^{-1} &= \tau(\rho^2\phi_1\rho^{-2})\tau^{-1} \\
&= \tau\rho(\rho\phi_1\rho^{-1})\rho^{-1}\tau^{-1} \\
&= \tau\rho\phi_3\rho^{-1}\tau^{-1} \\
&= (\tau\rho\tau^{-1})(\tau\phi_3\tau^{-1})(\tau\rho^{-1}\tau^{-1}) \\
&= \phi_1\rho(\phi_2\phi_1\phi_4\phi_1^{-1}\phi_2^{-1})\rho^{-1}\phi_1^{-1} \quad \text{and } \rho \text{ normalizes } \Lambda \\
&= \phi_1\lambda_5\phi_1^{-1} \quad \text{for some } \lambda_5 \in \Lambda \\
&\in \Lambda
\end{aligned}$$

and finally τ normalizes ϕ_0:

$$\begin{aligned}
\tau\phi_0\tau^{-1} &= \tau(\rho^2\phi_2\rho^{-2})\tau^{-1} \\
&= \tau\rho(\rho\phi_2\rho^{-1})\rho^{-1}\tau^{-1} \\
&= \tau\rho\phi_4\rho^{-1}\tau^{-1} \\
&= (\tau\rho\tau^{-1})(\tau\phi_4\tau^{-1})(\tau\rho^{-1}\tau^{-1}) \\
&= (\phi_1\rho\phi_2)(\phi_1\lambda_4\phi_1^{-1})(\phi_2^{-1}\rho^{-1}\phi_1^{-1}) \\
&= \phi_1\rho(\phi_2\phi_1\lambda_4\phi_1^{-1}\phi_2^{-1})\rho^{-1}\phi_1^{-1} \quad \text{and } \rho \text{ normalizes } \Lambda \\
&= \phi_1\lambda\phi_1^{-1} \quad \text{for some } \lambda \in \Lambda \\
&\in \Lambda
\end{aligned}$$

so that τ normalizes each generator of Λ and thus Λ itself, thereby inducing a self-isometry of our 6-punctured sphere $\mathbb{H}^2/\Lambda$ which sends the known disk center to the once-candidate-now-extremal-disk-center point p_6. □

Theorem 4.13. *If a* 6*-punctured sphere decomposes according to* 6_S *then it has two extremal disks which are exchanged by a self-isometry of the surface.* □

Although this demonstration explicitly shows that some 6-punctured spheres do have multiple extremal disks, the configuration utilized in our 6-punctured endeavors thus far is not the only possibility (unlike the $n = 3, 4, 5$ cases where but a single configuration is possible). Indeed, there are three other possible decompositions (whose equilateral triangle configurations are displayed in Figure 25), none of which allow for multiple extremal disks.

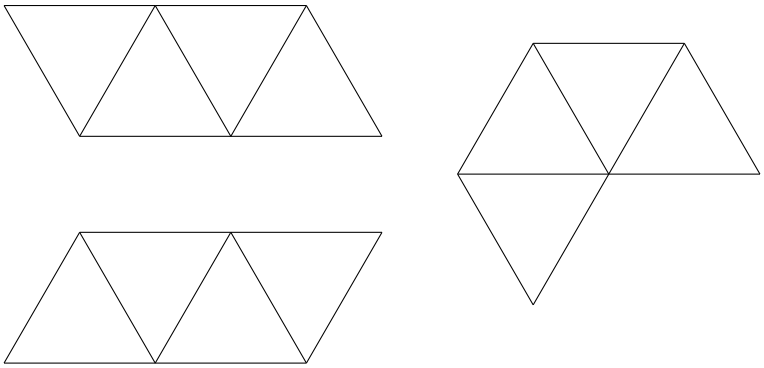

Figure 25: Alternate equilateral triangulations of 6-punctured spheres

Since, however, uniqueness of extremal disks for these configurations follows as an immediate consequence of more powerful results forthcoming in the following chapter (see Corollary 5.8), we suspend the proof for now and instead turn our efforts towards achieving those further-reaching results to which we have alluded.

5.0 UNIQUENESS OF EXTREMAL DISKS

Recall that, on the flipside of the very technique we used to discover additional disks when $n = 3, 4$, and 6, we can likewise rule out the possibility of alternate extremal disks by examining the equidistant locus of the canonical collection of horoball neighborhoods. More precisely, since any disk center must reside at a known distance (represented by d throughout our previous computations) from the canonical horoball neighborhood collection, it is enough to show that no other point enjoys precisely that proximity in order to definitively conclude we have a unique extremal disk. We now utilize this strategy in order to show that the unique extremal 5-punctured sphere has a unique extremal disk.

5.1 5-PUNCTURED SPHERES

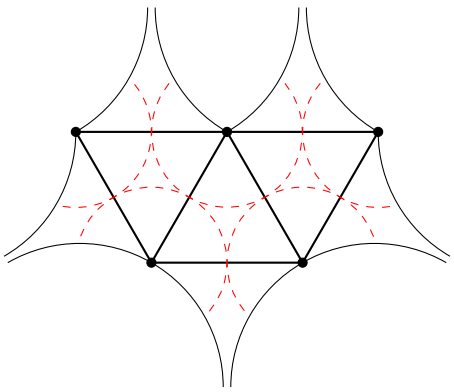

Figure 26: Triangulation of an extremal 5-punctured sphere with extremal disk shown

The crux of the idea is that the asymmetry of the only possible decomposition here renders an alternate disk center all together impossible. The reasoning applied in this demonstration (namely, using symmetry to rule out alternate disk centers) in conjunction with the bridge lemma 3.12 prove to be a successful strategy in ultimately proving that once $n > 6$, extremal n-punctured spheres must have unique extremal disks. Without further adieu, we state

Theorem 5.1. *The unique extremal* 5*-punctured sphere has a unique extremal disk.*

Proof. Take an extremal 5-punctured sphere, decompose it according to Theorem 2.3 and as usual take the known disk center to be at distance d from each horoball neighborhood in the canonical collection. Since any alternate disk center must also be of distance precisely d from the horoball neighborhoods, we look for such a point p.

We first notice that any candidate center point must reside within the interior of the central triangle, where the equidistant loci running between the pair of horocyclic ideal triangles on the right and the pair on the left intersect. But by the bridge lemma, such a potential disk center point p is closer to the lowermost neighborhood than it is to either the right- or leftmost. To see this, consider the construction featured in Figure 27, where t_1 denotes the distance between p and the leftmost horocycle neighborhood and t_2 its distance from the lowermost. The two corresponding 4-gons that arise (shown with red edges in Figure 27) then agree on their remaining three sides but each feature a pair of unequal angles. In terms of our oft-utilized $\alpha = 2\sin^{-1}(\frac{1}{2\cosh r})$ and $\beta = \sin^{-1}(\frac{1}{\cosh r})$ from Theorem 2.3, the unequal angle opposite side t_1 is given by $\phi_1 = \beta + (\alpha + \xi)$ whereas the unequal angle opposite t_2 is $\phi_2 = \beta + (\alpha - \xi)$, for some $\xi \in (0, \alpha)$.

Remark: Since the fundamental domain provided by the punctured sphere's decomposition into triangles here can obfuscate this application of the bridge lemma depicted in Figure 27, we also include a hopefully easier-to-digest and appreciate visualization, Figures 28 and 29, to aid in parsing the distortion.

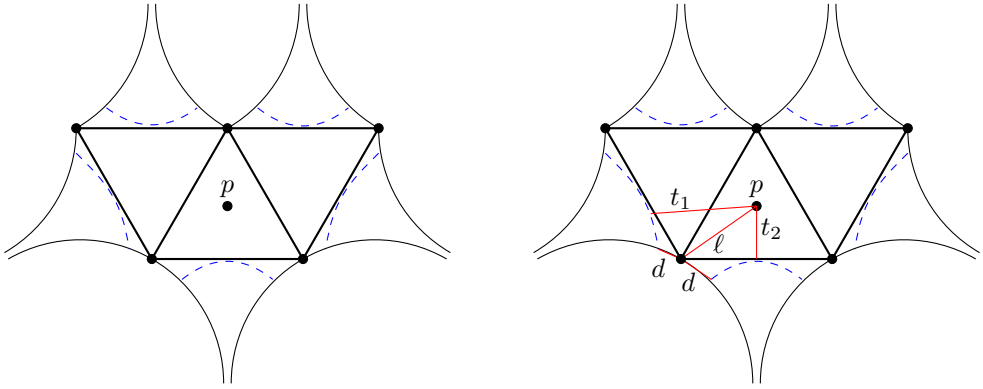

Figure 27: Identifying potential center for a second extremal disk

Therefore, the bridge lemma yields that

$$e^{t_1-d} = \cosh \ell - \cos(\phi_1)\sinh(\ell)$$

$$e^{t_2-d} = \cosh \ell - \cos(\phi_2)\sinh(\ell)$$

and since $\phi_2 < \phi_1 < \pi$, we have

$$\begin{aligned}
\cos(\phi_1) < \cos(\phi_2) &\implies -\cos(\phi_2) < -\cos(\phi_1) \\
&\implies -\cos(\phi_2)\sinh(\ell) < -\cos(\phi_1)\sinh(\ell) \\
&\implies \cosh(\ell) - \cos(\phi_2)\sinh(\ell) < \cosh(\ell) - \cos(\phi_1)\sinh(\ell)
\end{aligned}$$

hence, $e^{t_2} < e^{t_1}$ so that $t_2 < t_1$ and the candidate point is therefore not equidistant from each horoball neighborhood within the canonical collection (and hence cannot be the center of a second extremal disk). So there are in fact no other points at distance exactly d from the collection of horoball neighborhoods and therefore a second extremal disk is not possible. $\square$

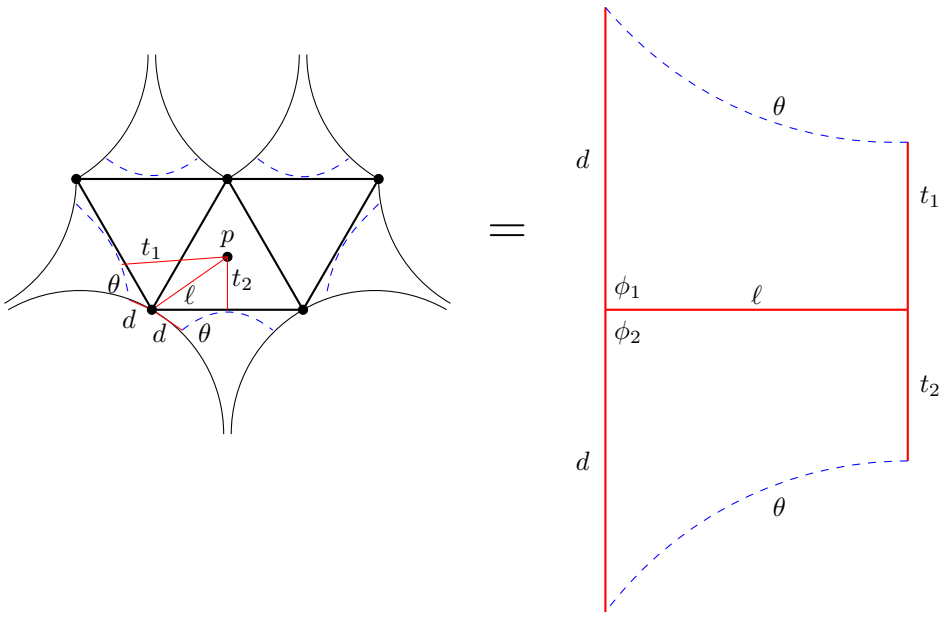

Figure 28: Applying the bridge lemma to extremal 5-punctured spheres

In some sense, uniqueness of extremal disks in this case is quite clear, since any extremal disk center must be an intersection point of every edge within the canonical horoball neighborhood collection's cut locus, of which the extremal 5-punctured sphere's decomposition obviously admits none. Moreover, in the next section we prove a more general result that incidentally covers this particular case. Nevertheless, because the observations here act as the conceptual seed for that forthcoming more-general solution (and that proof largely mirrors this one while generalizing it), we retain this section to act as a strategic blueprint of sorts.

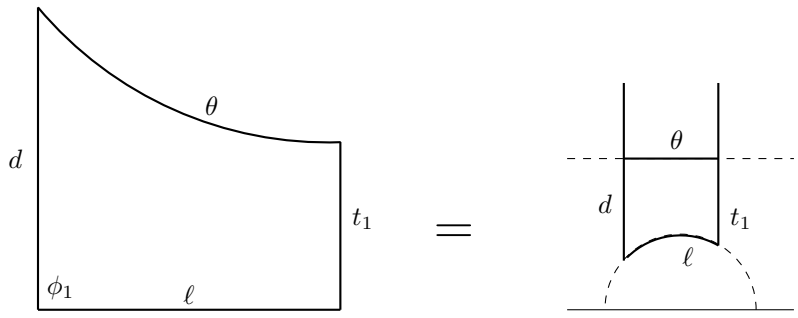

Figure 29: Applying the bridge lemma to extremal 5-punctured spheres

5.2 KILLER TREES

Recall that we can rule out the possibility of alternate extremal disk locations by examining the equidistant locus of the canonical collection of horoball neighborhoods. More precisely, since any disk center must reside at a known distance (represented by d throughout our previous computations) from the canonical horoball neighborhood collection, it is enough to show that no other point enjoys precisely that proximity in order to definitively conclude we have a unique disk. For this goal, we invoke symmetry arguments: reminiscent of the strategy used to show 5-punctured spheres must have unique extremal disks, we likewise through a given decomposition's asymmetry guarantee the same once $n > 6$. Here our dual tree framework pays us off in spades by allowing us to efficiently yet meaningfully translate simple, quantifiable conditions on equilateral triangle layout (as encoded and represented by dual trees) into symmetry arguments which eliminate the possibility of multiple extremal disks for n sufficiently large.

From observing cases of increasing non-compactness (i.e., a greater number of punctures) by hand it seems intuitively clear that the lack of symmetry exhibited by, say, any 5-punctured sphere is only exacerbated as n grows large - it quickly becomes increasingly more difficult to satisfy the evidently stringent requirement of being equidistant from all horoball cusp neighborhoods. To transfer this loose intuition into a rigorous theorem, we look for subtrees that can manifest sufficient decomposition asymmetry in and of themselves

(and whose presence thereby renders alternate disk centers impossible), ultimately resulting in the affectionately dubbed Killer Trees.

Definition 5.2 (Killer Tree 1)**.** *Killer Tree 1 consists of a valence-*1 *vertex attached to a valence-*2 *vertex.*

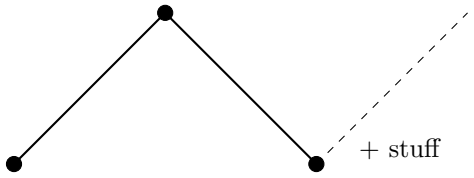

Figure 30: Killer Tree 1

Definition 5.3 (Killer Tree 2)**.** *Killer Tree 2 consists of exactly two valence-*1 *vertices attached to a single valence-*3 *vertex.*

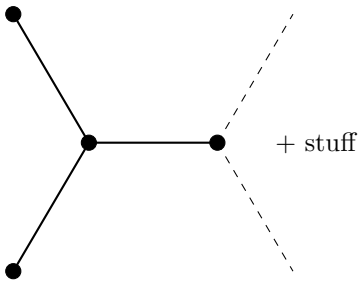

Figure 31: Killer Tree 2

Remark: Note that three valence-1 vertices attached to a single valence-3 vertex does not constitute a Killer Tree (consistent with what we witnessed throughout the explorations of Section 4.3, which discovered a 6-punctured sphere with multiple extremal disks).

Our first order of business is to show that these Killer Trees appear as subtrees in any decomposition of an extremal n-punctured sphere for $n > 6$.

Lemma 5.4. *For $n > 6$, the dual tree to any n-punctured sphere decomposition which does not contain Killer Tree 1 as a subtree must contain Killer Tree 2 as a subtree.*

Proof. First note that any tree with a finite number of vertices must have a vertex of valence 1. To see this, take a tree with finite number, say m, of vertices and choose a vertex arbitrarily. If this vertex is of valence 1 the claim is satisfied; if not, we can move to an adjacent new vertex. But before m repetitions, this procedure halts at a valence-1 vertex or we have contradicted having a tree in the first place (as there are only m distinct vertices).

Now, suppose our tree does not contain Killer Tree 1 as a sub-tree: that is, suppose that no valence-1 vertex is attached to a valence-2 vertex. Then every valence-1 vertex is attached to a valence-3 vertex. Denoting the number of valence-i vertices as v_i and the number of edges as e we have that

$$2e = v_1 + 2v_2 + 3v_3$$

and from the tree-defining condition $v - e = 1$ we deduce that

$$e = v_1 + v_2 + v_3 - 1$$

Combining these yields $v_1 = v_3 + 2$ so that the pigeon hole principle then guarantees at least one valence-3 vertex has exactly two valence-1 vertices attached to it (i.e., our tree contains Killer Tree 2 as a subtree). □

Next we show that the presence of either Killer Tree as a subtree within the dual tree associated to any decomposition of an extremal n-punctured sphere with $n > 6$ renders a second extremal disk all together impossible.

Lemma 5.5. *Any extremal n-punctured sphere whose decomposition's dual tree contains either Killer Tree 1 or Killer Tree 2 as a subtree has a unique extremal disk.*

Proof. (Killer Tree 1 Case) We essentially mimic the proof presented to tackle the $n = 5$ case. Since the center of any potential second extremal disk must be equidistant from each horoball neighborhood of the canonical collection, any such candidate must in particular lie on the intersection of the equidistant loci running between (1) the two neighborhoods on either side of the equilateral triangle which contains the valence-1 vertex on the end of our Killer Tree 1 subtree and (2) the single neighborhood from the equilateral triangle containing the valence-2 vertex of our Killer Tree 1 subtree and the adjacent neighborhood from the previously

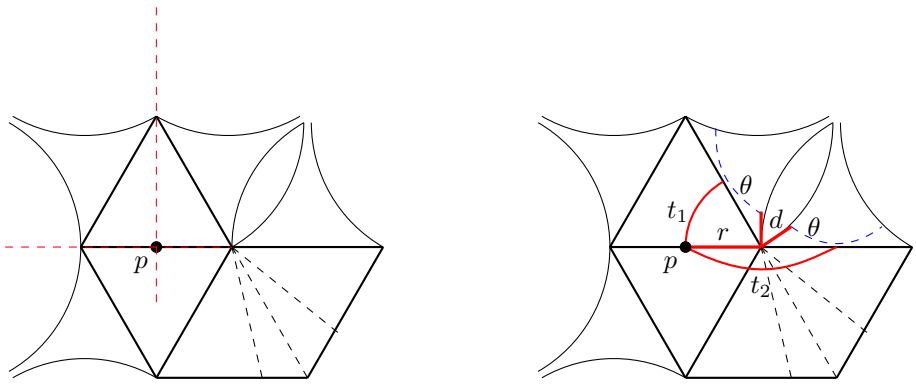

Figure 32: Identifying candidate point by observing equidistant loci (left) and then comparing its distances between horoball neighborhoods with the bridge lemma 3.12 (right)

mentioned valence-1 (from the Killer Tree 1 subtree) containing equilateral triangle (depicted in Figure 32).

Labeling that candidate point p, we then compare its distance from different horoball neighborhoods in the collection: one of the neighborhoods from our Killer Tree 1's valence-1 containing equilateral triangle and a horoball neighborhood off the valence-1 containing equilateral triangle closest. Labeling those distances as t_1 and t_2 respectively, we may view each of the corresponding edges as living on two distinct 4-gons which share two edge lengths (those shared lengths are exactly $r = r_{0,n}$ and $d = \log\cosh r$) and both of which have our candidate point as a vertex (shown in Figure 32 more geometrically and in Figure 33 more schematically). Then the one associated with t_1 has an angle $\phi_1 = \alpha + \beta$ but the t_2 one has an angle $\phi_2 = \alpha + k\alpha + \beta$ for some $k \geq 1$, with angles α and β as described in DeBlois [8] (see the schematic portrayel in Figure 33). Then applying our bridge lemma on each (with $\ell = r_{0,n} = r$ and $s = d_{0,n} = d$) yields:

$$e^{t_1 - d} = \cosh r - \cos(\phi_1)\sinh(r)$$

$$e^{t_2 - d} = \cosh r - \cos(\phi_2)\sinh(r).$$

Now, since $\phi_1 < \phi_2 < \pi$, we have (arguing similarly to as in 5.1)

$$\begin{aligned}\cos(\phi_2) < \cos(\phi_1) &\implies -\cos(\phi_1) < -\cos(\phi_2) \\ &\implies -\cos(\phi_1)\sinh(\ell) < -\cos(\phi_2)\sinh(\ell) \\ &\implies \cosh(\ell) - \cos(\phi_1)\sinh(\ell) < \cosh(\ell) - \cos(\phi_2)\sinh(\ell)\end{aligned}$$

Hence, $e^{t_1} < e^{t_2}$ so that $t_1 < t_2$ and the candidate point is therefore not equidistant from each horoball neighborhood within the canonical collection (and hence cannot be the center of a second extremal disk). □

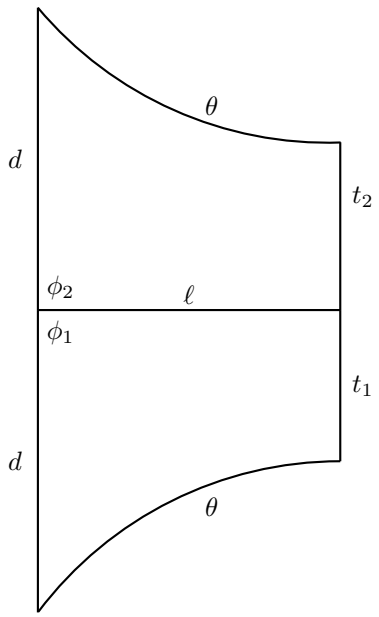

Figure 33: Schematic visualization of applying the Bridge Lemma on the Killer Trees

Proof. (Killer Tree 2 Case) We proceed similarly to the first case. By using the equidistant loci running between the three pairs of horoball neighborhoods around the Killer Tree 2 subtree, we identify the only potential candidate point where those three loci intersect (see left of Figure 34).

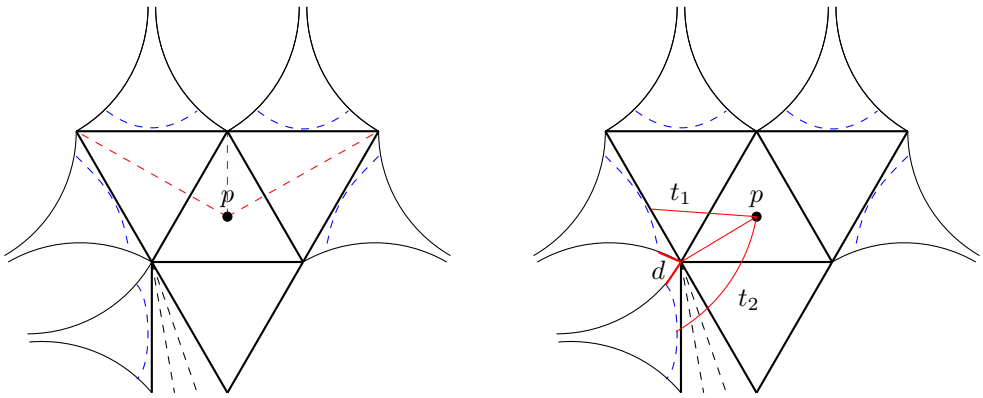

Figure 34: Visualization of Killer Tree 2 case

Labeling that candidate point p, we then compare its distance from different horoball neighborhoods in the collection: one of the valence-1 neighborhoods off our Killer Tree 2 subtree and a horoball neighborhood off the valence-1 containing equilateral triangle closest (see right of Figure 34). Labeling those distances as t_1 and t_2 respectively, we again view each of the corresponding edges as living on two distinct 4-gons which again share two edge lengths and both of which have our candidate point as a vertex (displayed in Figure 34). However, just as in Theorem 5.1, these 4-gons feature a pair of unequal angles, whose comparison showns $t_1 < t_2$. Specifically, the one associated with t_1 has an angle $\phi_1 = \alpha + \frac{\alpha}{2} + \beta$ but the t_2 one has an angle $\phi_2 = \alpha + \frac{\alpha}{2} + k\alpha + \beta$ for some $k \geq 1$, in terms of angles α and β as described in DeBlois 2.3 [8] (see the schematic portrayel in Figure 33). Then applying our bridge lemma 3.12 on each yields:

$$e^{t_1 - d} = \cosh \ell - \cos(\phi_1) \sinh(\ell)$$

$$e^{t_2 - d} = \cosh \ell - \cos(\phi_2) \sinh(\ell).$$

And once again since $\phi_1 < \phi_2 < \pi$, we have $e^{t_1} < e^{t_2}$ so that $t_1 < t_2$ and the candidate point is therefore not equidistant from each horoball neighborhood within the canonical collection (and hence cannot be the center of a second extremal disk). $\square$

Then the Killer Tree Lemmas combine for our capstone result towards answering the extremal disk uniqueness question for extremal n-punctured spheres.

Theorem 5.6. *Extremal n-punctured spheres have unique extremal disks for $n > 6$.*

Proof. The dual tree of the decomposition associated to any such extremal n-punctured sphere must contain either Killer Tree 1 or Killer Tree 2 by Lemma 5.4. Therefore any such surface has a unique extremal disk by Lemma 5.5. □

Corollary 5.7. *Extremal 5-punctured spheres have unique extremal disks.*

Proof. The extremal 5-punctured sphere decomposition has Killer Tree 1 as its dual tree and hence has a unique extremal disk by Lemma 5.5. □

Remark: Despite this result's redundancy with our previous section, we mention it here to reiterate that proof generalizes to the generic setting and emphasize that this more-general solution even covers the previously-solved $n = 5$ case.

Similar to the above, but duly stated in order to fulfill our promise made at the end of the 6-punctured escapades (Section 4.3):

Corollary 5.8. *If an extremal 6-punctured sphere does not decompose according to 6_S, then it has a unique extremal disk.*

Proof. The dual tree to any such decomposition contains Killer Tree 1. □

6.0 AN EXACT COUNT

Earlier in our endeavors, we found the number of extremal n-punctured sphere surface-disk pairs, up to orientation-preserving isometry of the surface. In the directly subsequent chapters however, we showed that every extremal punctured sphere featuring multiple extremal disks also admits a self-isometry which exchanges them. We proved that extremal disks are unique (via Theorem 5.1, Theorem 5.6, and Corollary 5.8) outside of three special cases, and that in each of those remaining three cases, the corresponding surface has exactly two extremal disks which are swapped by a self-isometry of the corresponding surface: the unique extremal 3-punctured sphere in Theorem 4.5, the unique extremal 4-punctured sphere in Theorem 4.9, and an extremal 6-punctured sphere in Theorem 4.13.

This then allows us to immediately promote our previous total from a mere upper bound to a firm and exact count of extremal punctured spheres. To reiterate: as our initial count of surface-disk pairs was only up to orientation-preserving isometry of the surface, the existence of such isometries in every single instance of an extremal punctured sphere surface with multiple extremal disks hence guarantees that what was once a potential overcount is in fact a precise count. We thus state a new and improved, slightly stronger form of Theorem 3.8:

Theorem 6.1. *The number extremal n-punctured spheres (up to orientation-preserving isometry) is given by*

$$C(n-2)/n + C(n/2-1)/2 + (2/3)C(n/3-1)$$

where $C(n) = \frac{1}{n+1}\binom{2n}{n} = \frac{(2n)!}{n!(n+1)!}$ (the Catalan numbers) and terms are omitted if their arguments are not integers.

In particular, the totals for $3 \leq n \leq 15$ are given on the following table:

n	Number of Extremal n-punctured Spheres
3	1
4	1
5	1
6	4
7	6
8	19
9	49
10	150
11	442
12	1424
13	4522
14	14924
15	49536

Table 2: Exact count of extremal punctured spheres

6.1 ASYMPTOTICS

Having achieved our exact count of extremal n-punctured spheres, a natural follow-up is to wonder about its asymptotics: at roughly what rate is this total growing? The Online Encyclopedia of Integer Sequences (OEIS) [24], which lists the sequence detailing the total number of extremal n-punctured spheres (up to orientation-preserving isometry) as the number of unlabeled plane trivalent trees (with $n-2$ internal vertices, all of degree 3, and n leaves) in addition to the number of dissected n-gons (which we have shown to be in bijection with extremal n-punctured spheres), suggests that this total grows at an exponential rate comparable to 4^n.

While no proof is cited for this assertion, we are able to reproduce that asymptotic claim with deeper scrutiny of the Catalan formulation of the sequence and an appeal to Stirling's approximation [12]. Recall that the number of extremal n-punctured spheres (via their bijection with dissected n-gons) is given by Brown's $G_{0,n-3}$ [5]. Then we have:

Theorem 6.2. $G_{0,n-3} \sim \frac{2^{2n-4}}{\sqrt{\pi n^5}}$ *i.e., the ratio of the two sides tends to* 1 *as* n *tends to* ∞.

Proof. Recall that in terms of Catalan numbers $C(n) = \frac{(2n)!}{n!(n+1)!}$ we have (from Theorem 3.8)

$$G_{0,n-3} = C(n-2)/n + C(n/2-1)/2 + (2/3)C(n/3-1)$$

(and terms are omitted if their inputs are not integers). Then by Stirling's approximation that $n! \sim \sqrt{2\pi n}\left(\frac{n}{e}\right)^n$, we consequently obtain:

$$\begin{aligned}
\frac{1}{n}C(n-2) = \frac{(2n-4)!}{n!(n-2)!} &\sim \frac{\sqrt{2\pi(2n-4)}(\frac{2n-4}{e})^{2n-4}}{\sqrt{2\pi n}(\frac{n}{e})^n\sqrt{2\pi(n-2)}(\frac{n-2}{e})^{n-2}} \\
&= \sqrt{\frac{2n-4}{2\pi n(n-2)}}e^2\frac{(2n-4)^{2n-4}}{n^n(n-2)^{n-2}} \\
&= \frac{1}{\sqrt{\pi n}}\frac{e^2}{n^2}\frac{(n-2)^{n-2}\cdot 2^{2n-4}}{n^{n-2}} \\
&= \frac{1}{\sqrt{\pi n}}\frac{e^2}{n^2}\left(\frac{n-2}{n}\right)^{n-2}2^{2n-4}.
\end{aligned}$$

Here we pause our asymptotic derivation for a crucial computational lemma:

Lemma 6.3. $\lim_{n\to\infty}\left(\frac{n-2}{n}\right)^{n-2} = e^{-2}$.

Proof. Let $y = \left(\frac{n-2}{n}\right)^{n-2}$ so that $\log y = (n-2)\log\left(\frac{n-2}{n}\right) = (n-2)\log\left(1-\frac{2}{n}\right)$.

Then $\lim_{n\to\infty}\log y = \lim_{n\to\infty}(n-2)\log\left(1-\frac{2}{n}\right) = \lim_{n\to\infty}\dfrac{\log\left(1-\frac{2}{n}\right)}{1/(n-2)}$

$\overset{\text{L'H}}{=} \lim_{n\to\infty}\dfrac{(2/n^2)/\left(1-\frac{2}{n}\right)}{-1/(n-2)^2} = \lim_{n\to\infty} -\dfrac{2}{1-2/n}\dfrac{(n-2)^2}{n^2} = \lim_{n\to\infty} -\dfrac{2}{1-2/n}\cdot\lim_{n\to\infty}\dfrac{(n-2)^2}{n^2} = -2.$

Hence, $\lim_{n\to\infty}\left(\frac{n-2}{n}\right)^{n-2} = \lim_{n\to\infty} y = \lim_{n\to\infty} e^{\log y} = e^{-2}$. □

Returning to our asymptotic derivation and applying Lemma 6.3 we see that

$$\frac{1}{n}C(n-2) \sim \frac{1}{\sqrt{\pi n}}\frac{1}{n^2}\cdot 2^{2n-4} = \frac{2^{2n-4}}{\sqrt{\pi n^5}}.$$

Now for the other terms, $C(n/2-1) = \frac{(2k-2)!}{k!(k-1)!}$ when $n = 2k$. Therefore,

$$\begin{aligned}\lim_{n\to\infty}\frac{n\cdot C(n/2-1)}{C(n-2)} &= \lim_{k\to\infty}\frac{(2k)!(2k-2)!(2k-2)!}{(4k-4)!k!(k-1)!}\\ &\le \lim_{k\to\infty}(k+1)(k+2)(k+3)(2/3)^{2k-3}\\ &= \lim_{k\to\infty}\frac{(k+1)(k+2)(k+3)}{(3/2)^{2k-3}} \overset{\text{L'H}}{=} 0.\end{aligned}$$

Similarly, $C(n/3-1) = \frac{(2p-2)!}{p!(p-1)!}$ when $n = 3p$ so that

$$\lim_{n\to\infty}\frac{n\cdot C(n/3-1)}{C(n-2)} = \lim_{p\to\infty}\frac{(3p)!(3p-2)!(2p-2)!}{(6p-4)!p!(p-1)!} = 0$$

and we obtain that $G_{0,n-3} \sim \frac{1}{n}C(n-2) \sim \frac{2^{2n-4}}{\sqrt{\pi n^5}}$. □

So the number of disk triangulations of type $[0, n-3]$, i.e., the number of dissected n-gons, i.e., the number of extremal n-punctured spheres grows at an exponential rate, comparable to 4^n.

7.0 COUNTING TREES

While we originally intended to achieve our count of extremal punctured spheres as a biproduct of a bijective correspondence between them and the dual trees of Lemma 3.2, we ultimately found it easier to rigorously arrive at said count via a bijection (established in Theorem 3.6) between extremal punctured spheres and n-gon dissections, whose count is explicitly detailed through formulas from Brown (as shown by Theorem 3.7) [5].

Nevertheless, as mentioned in Chapter 6, dissected n-gons are in one-to-one correspondence with trivalent planar trees with n leaves, which are in bijection with dual trees to extremal n-punctured spheres (modulo suitable notions of equivalence). Through the dual tree lens, we were able to derive some interesting recursive counting technology which we now introduce as an alternate attack on the extremal punctured sphere counting problem (and perhaps a curio in its own regard).

7.1 GEOMETRIC TREES

Dual trees to n-punctured sphere decompositions actually enjoy more structure than our preceding endeavors demonstrate: by virtue of being constructed as the dual to equilateral triangles, all its edges are necessarily the same length and all intersecting edges meet at an angle of $\frac{2\pi}{3}$.

Moreover, for any dual tree to an n-punctured sphere decomposition, its equal edge lengths are determined solely by n (recall that each equilateral triangle in any such decomposition has edges of length $2r_{0,n}$) and exactly through the system of equations we used in our 3-punctured investigations (specifically, Lemma 4.4). Precisely, calling that n-dependent

edge length ℓ_n, we know it must satisfy

$$\cosh\frac{\ell_n}{2} = \frac{\sqrt{1+\frac{4}{3}\sinh^2 r_{0,n}}}{\cosh r_{0,n}} = \frac{\sqrt{1+\frac{4}{3}(\cosh^2 r_{0,n}-1)}}{\cosh r_{0,n}}.$$

For instance, $\ell_4 \approx 0.9961$ and $\ell_6 \approx 1.0641$. (Perhaps somewhat unfortunately, the dual tree to a 3-punctured sphere's decomposition has no edges at all; but, per that forthcoming computation, any would have to be of length exactly $2\log\sqrt{\frac{7}{3}}$.) The rigidity of these trees motivates a new definition:

Definition 7.1 (Geometric Tree)**.** *We say a tree is geometric if its vertices are of valence at most 3, its edges are all the same length and any intersecting edges meet at angle* $\frac{2\pi}{3}$.

Continuing to denote the number of vertices with valence j by v_j, recall that the equation $v_1 = v_3 + 2$ is always satisfied in any geometric tree (as we witnessed in Lemma 5.4). This relation enables us to describe each v_j for an n-vertex tree in terms of just n and v_2. Taking v_2 to be i we get:

$$v_1 = \frac{n-i}{2}+1, \qquad v_2 = i, \qquad v_3 = \frac{n-i}{2}-1$$

and incidentally (but importantly) we note that $n \equiv v_2 \pmod 2$. We now introduce new notation towards establishing an upper bound for our tree count.

$\mathcal{T}_n := \{\text{geometric trees with } n \text{ vertices of valence at most } 3\}$

$\mathcal{T}_n^i := \{\text{geometric trees with } n \text{ vertices of valence at most } 3, i \text{ of which are of valence } 2\}$

$k_n := |\mathcal{T}_n|$

$k_n^i := |\mathcal{T}_n^i|$

So if n is even we have that

$$\mathcal{T}_n = \mathcal{T}_n^0 \cup \mathcal{T}_n^2 \cup \cdots \cup \mathcal{T}_n^{n-2} = \bigcup_{j=0}^{\frac{n-2}{2}} \mathcal{T}_n^{2j}$$

and hence (since $l \neq p$ implies that $\mathcal{T}_n^l \cap \mathcal{T}_n^p = \emptyset$ and $k_n^n = 0$ naturally),

$$k_n = k_n^0 + k_n^2 + \cdots + k_n^{n-2} = \sum_{j=0}^{\frac{n-2}{2}} k_n^{2j}.$$

On the other hand, if n is odd we get that

$$\mathcal{T}_n = \mathcal{T}_n^1 \cup \mathcal{T}_n^3 \cup \cdots \cup \mathcal{T}_n^{n-2} = \bigcup_{j=0}^{\frac{n-3}{2}} \mathcal{T}_n^{2j+1}$$

and similarly

$$k_n = k_n^1 + k_n^3 + \cdots + k_n^{n-2} = \sum_{j=0}^{\frac{n-3}{2}} k_n^{2j+1}.$$

7.2 EDGE ATTACHMENT AND RECURSIVE COUNTING

One way we may view producing an n-vertex tree is by attaching an edge to an $(n-1)$-vertex tree at a permissible vertex (i.e., one not already of valence 3). To get an element of $\mathcal{T}_n^i$ in this recursive fashion however, we are restricted to one of either

Definition 7.2 (Edge Attachments)**.**

$$(1)\quad \mathcal{T}_{n-1}^{i-1} \xrightarrow{\text{add edge at valence-1 vertex}} \mathcal{T}_n^i$$

$$(2)\quad \mathcal{T}_{n-1}^{i+1} \xrightarrow{\text{add edge at valence-2 vertex}} \mathcal{T}_n^i$$

This construction does not however come equipped with a natural mechanism for detecting the possible tree-building redundancies that may occur within it. For instance, we can achieve an identical 5-vertex descendant tree from two distinct predecessors and through two nominally different attachment procedures (shown in Figure 35):

Moreover, it continues to over count in the sense of allowing for edge attachments that would actually result in a graph with a complete cycle (which thereby ceases to be a tree all together). Nevertheless, recognizing that elements of $\mathcal{T}_n^i$ must arise from attaching an edge to an element of either $\mathcal{T}_{n-1}^{i-1}$ or $\mathcal{T}_{n-1}^{i+1}$ does enable us to produce an upper bound on k_n^i inductively. Specifically, we can say that k_n^i is bounded above by the number of permissible edge attachments as described in (1) and (2) above. Noting also that any attachment of form (1) offers 2 choices for its corresponding descendant, we have

$$k_n^i \leq 2k_{n-1}^{i-1}\Big(\frac{n-i}{2}+1\Big) + k_{n-1}^{i+1}\Big(i+1\Big).$$

Then if n is even, our earlier formula yields

$$k_n = \sum_{j=0}^{\frac{n-2}{2}} k_n^{2j} \leq \sum_{j=0}^{\frac{n-2}{2}} 2k_{n-1}^{2j-1}\Big(\frac{n-2j}{2}+1\Big) + k_{n-1}^{2j+1}\Big(2j+1\Big)$$

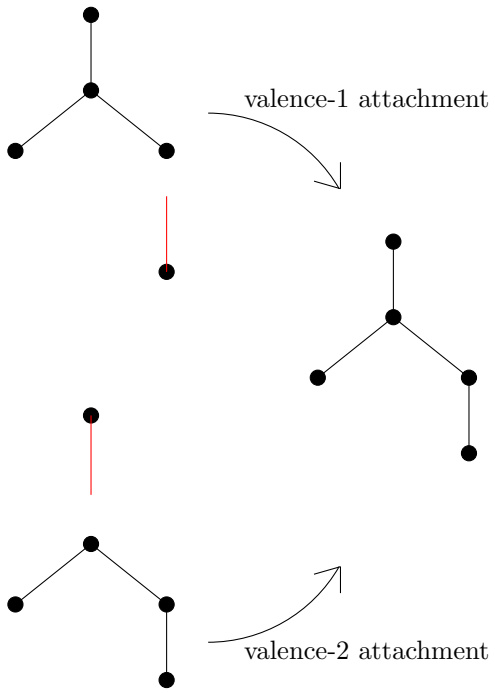

Figure 35: Two different edge-attachment processes resulting in the same descendant tree

and if n is odd, we instead get

$$k_n = \sum_{j=0}^{\frac{n-3}{2}} k_n^{2j+1} \leq \sum_{j=0}^{\frac{n-3}{2}} 2k_{n-1}^{2j}\Big(\frac{n-2j+1}{2}\Big) + k_{n-1}^{2j+2}\Big(2j+2\Big).$$

To the end of improving this upper bound to a precise count, we refine our notation. The key idea behind this refinement is the recognition that for our purposes, the ridgid rotational symmetry group of any such geometric tree is of order exactly 1, 2, or 3. So we formally define

$\mathfrak{T}_n^i(1) := \{\text{elements of } \mathfrak{T}_n^i \text{ whose symmetry group is of order 1}\}$

$\mathfrak{T}_n^i(2) := \{\text{elements of } \mathfrak{T}_n^i \text{ whose symmetry group is of order 2}\}$

$\mathfrak{T}_n^i(3) := \{\text{elements of } \mathfrak{T}_n^i \text{ whose symmetry group is of order 3}\}$

which enables us to, per the aforementioned observation, then write that

$$|\mathfrak{T}_n^i| = |\mathfrak{T}_n^i(1)| + |\mathfrak{T}_n^i(2)| + |\mathfrak{T}_n^i(3)|.$$

The real upshot of this further breakdown is that the symmetry itself dictates the manner in which the tree enjoying it can be realized. More precisely, by repurposing our recursive edge attachment procedures already mentioned, we have that:

Theorem 7.3 (Order-2 Symmetry Trees)**.** *Any geometric tree with n vertices whose rigid rotational symmetry group has order 2 is the union of a subtree with $\frac{n+2}{2}$ vertices and the image of that subtree under the order-2 rotation.*

Proof. For any geometric tree, its group of rigid rotational symmetries acts on its collection of n vertices. While an order-2 rotation has a unique fixed point in $\mathbb{H}^2$, the nontrivial rotation enjoyed by a tree whose rigid rotational symmetry group is of order 2 fixes none of its vertices (since the edge structure surrounding it would need to be preserved, which by the definition of a geometric tree requires a rotation doing so be of order 1 or 3) so all are exchanged in pairs and n is necessarily even. The group of rigid rotational symmetries also similarly acts

on its collection of $n-1$ edges; the order-2 rotation has a unique fixed edge and exchanges the remaining $n-2$ edges in pairs. Moreover, the path between distinct representatives of any given edge orbit necessarily includes the fixed edge, since we can demonstrate such a path by unioning the path between one representative and the fixed edge with its rotational image, and then conclude it is the unique such path by the definition of a tree.

So there exists a subtree with $\frac{n-2}{2}$ edges such that none of the images (under the order-2 rotation) of its edges are contained in it, namely, either remaining connected component when the single fixed edge is removed; such a maximal fixed-edge-free subtree necessarily consists of exactly one representative from each non-singleton edge orbit under the rotational group's action on them (and therefore must contain $\frac{n-2}{2}$ edges) or we contradict the above assertion by demonstrating a path between distinct representatives of the same edge orbit which does not include the edge fixed by the rotation. It is moreover connected to the unique edge fixed by the rotation by construction and comprises a geometric tree on $\frac{n}{2}$ edges (hence $\frac{n+2}{2}$ vertices) when considered with it. Attaching this augmented subtree with its image under the order-2 rotation (along the edge fixed by it) recovers the original geometric tree with n vertices. □

As a consequence, we then get that the number of geometric trees with an order-2 rigid rotational symmetry group is the number of constituent subtree attatchers (as described above) times the unique valence-1 points of attachment. We therefore have:

Corollary 7.4. $|\mathcal{T}_n^i(2)| = \left(\frac{n-i+6}{4}\right)\left(|\mathcal{T}^{\frac{i}{2}}_{\frac{n+2}{2}}(1)| + \frac{1}{2}|\mathcal{T}^{\frac{i}{2}}_{\frac{n+2}{2}}(2)| + \frac{1}{3}|\mathcal{T}^{\frac{i}{2}}_{\frac{n+2}{2}}(3)|\right)$

Similarly, for geometric trees with an order-3 symmetry group we have:

Theorem 7.5 (Order-3 Symmetry Trees)**.** *Any geometric tree with n vertices whose rigid rotational symmetry group has order 3 is the union of a subtree with $\frac{n+2}{3}$ vertices and the images of that subtree under both order-3 rotations.*

Proof. For any geometric tree with n vertices whose rigid rotational symmetry group is of order 3 we know (by noting that the group of rigid rotational symmetries of a geometric tree acts on its set of edges in general, and here with none fixed) that $n \equiv 1 \pmod 3$ so there is a unique trivalent vertex fixed by both its order-3 rotations. Moreover, any path

between distinct representatives of any given edge orbit must include the unique vertex fixed by order-3 rotation (similarly to the previous case).

Then sharing exactly one of any of that fixed vertex's three edges is a geometric subtree with $\frac{n+2}{3}$ vertices comprised of a representative of each edge orbit under the rotational group's action on them, namely, the connected component now containing the single fixed trivalent vertex when two of its edges are removed. As before, the union of this subtree with its order-3 rotation images forms the original tree by construction. □

Analogously to before, the number of geometric trees with an order-3 rigid rotational symmetry group is again given by the number of constituent subtree attatchers times the unique valence-1 points of attachment. Hence:

Corollary 7.6. $|\mathfrak{T}_n^i(3)| = \left(\frac{n-i+8}{6}\right)\left(|\mathfrak{T}^{\frac{i}{3}}_{\frac{n+2}{3}}(1)| + \frac{1}{2}|\mathfrak{T}^{\frac{i}{3}}_{\frac{n+2}{3}}(2)| + \frac{1}{3}|\mathfrak{T}^{\frac{i}{3}}_{\frac{n+2}{3}}(3)|\right)$

In conjunction with this symmetrical breakdown of Corollaries 7.4 and 7.6, we get our hands on the remaining $\mathfrak{T}_n^i(1)$ via Theorem 7.7 below, which allows to compute that missing quantity recursively, purely in terms of previously or otherwise known data, and thereby determine $|\mathfrak{T}_n^i|$ exactly.

Theorem 7.7 (Arrows In, Arrows Out Equality)**.**

$$2\left(\tfrac{n-i+2}{2}\right)\left(|\mathfrak{T}_{n-1}^{i-1}(1)| + \tfrac{1}{2}|\mathfrak{T}_{n-1}^{i-1}(2)| + \tfrac{1}{3}|\mathfrak{T}_{n-1}^{i-1}(1)|\right)$$
$$+(i+1)\left(|\mathfrak{T}_{n-1}^{i+1}(1)| + \tfrac{1}{2}|\mathfrak{T}_{n-1}^{i+1}(2)| + \tfrac{1}{3}|\mathfrak{T}_{n-1}^{i+1}(3)|\right)$$
$$= \left(\tfrac{n-i+2}{2}\right)\left(|\mathfrak{T}_n^i(1)| + \tfrac{1}{2}|\mathfrak{T}_n^i(2)| + \tfrac{1}{3}|\mathfrak{T}_n^i(3)|\right)$$

Proof. Any geometric tree on n vertices can be formed from several $(n-1)$-vertex predecessors as described in Definition 7.2, the total amount of which is simultaneously the number of ways to form a distinct $(n-1)$-vertex tree from an n-vertex tree through removing the single edge attached to a valence-1 vertex. □

8.0 THE GENERAL NON-COMPACT CASE

Here we begin an extension of the classification of extremal surfaces to the general non-compact setting by achieving a count of extremal disk-surface pairs in the case of once-punctured extremal surfaces. Assimilating our observations here and prior, we formulate a conjecture which would upgrade the upper bound on once-punctured extremal surfaces (given by the total of surface-disk pairs mentioned) to a firm count.

8.1 ONCE-PUNCTURED EXTREMAL SURFACES

As mentioned in our earlier descriptions of the complete classification of the closed case of extremal surfaces, Bacher & Vdovina produce an exact count of closed extremal hyperbolic surfaces [1]. To achieve their precise count, they enumerate what are referred to as oriented wicks forms. These are dual objects to 1-vertex triangulations and inequivalent wicks forms correspond to distinct closed extremal surface-disk pairs. Since any once-punctured surface may be viewed as a closed surface with a puncture added, equivalently an oriented wicks form with a bigon inserted at one edge, we can obtain an upper bound on the number of extremal surfaces of genus g with $n = 1$ cusp by merely modifying their count. The idea is to simply multiply the number of wicks forms they assemble by twice the number of edges appearing in them (each edge representing a spot where we can insert a bigon in one of two orientations, i.e., place a puncture on the surface), then dividing by the order of the automorphism group of that form. Note however that this a priori only comprises an upper bound rather than a sharp count because this details extremal surface-disk pairs rather than extremal surfaces themselves (see, however, Conjecture 8.3).

We now detail this upper bound in its full generality before providing the explicit totals for small g. Noting that the group of automorphisms of an oriented Wicks form is cyclic of order $1, 2, 3$, or 6, Bacher & Vdovina [1] first define (for fixed genus g):

$m_1^g = \frac{2}{1}(\frac{1}{12})^g \frac{(6g-5)!}{g!(3g-3)!}$
$m_2^g(r) = \frac{2}{2}(\frac{2^2}{12})^f \frac{(6f+2r-5)!}{r!f!(3f+r-3)!}$ with $f = \frac{2g+1-r}{4}$
$m_3^g(s,t) = \frac{2}{3}(\frac{3^2}{12})^f \frac{(6f+2s+2t-5)!}{s!t!f!(3f+s+t-3)!}$ with $f = \frac{g+1-s-t}{3}$
$m_6^g(3r;2s;2t) = \frac{2}{6}(\frac{6^2}{12})^f \frac{(6f+2r+2s+2t-5)!}{r!s!t!f!(3f+r+s+t-3)!}$ with $f = \frac{2g+5-3r-4s-4t}{12}$

and then correspondingly define:

$m_2^g = \sum_{r\in\mathbb{N},\ (2g+1-r)/4\in\mathbb{N}\cup\{0\}} m_2^g(r)$
$m_3^g = \sum_{s,t\in\mathbb{N},\ (g+1-s-t)/3\in\mathbb{N}\cup\{0\},\ s\equiv 2g+1 \pmod 3} m_3^g(s,t)$
$m_6^g = \sum_{r,s,t\in\mathbb{N},\ (2g+5-3r-4s-4t)/12\in\mathbb{N}\cup\{0\},\ 2s\equiv 2g+1 \pmod 3} m_6^g(3r;2s;2t)$

in order to prove that:

Theorem 8.1 (Bacher, Vdovina)**.** *There are*
$6m_6^g$ *forms with 6 automorphisms*
$3m_3^g - 3m_6^g$ *forms with 3 automorphisms*
$2m_2^g - 2m_6^g$ *forms with 2 automorphisms*
$m_1^g - m_2^g - m_3^g + m_6^g$ *forms without nontrivial automorphisms.*

Hence, we therefore obtain

Theorem 8.2. *For genus $g \geq 1$ with $n = 1$ cusp, the number of pairs consisting of an extremal surface and one of its extremal disks is given by $\frac{(2g-1)(6g-5)!}{12^{g-1}g!(3g-3)!}$.*

Proof. As each oriented wicks form of genus g possesses $6g-3$ edges, multiplying the number of distinct oriented wicks forms by twice the number of distinct edges appearing in them and dividing by the order of its automorphism group yields:

$$(12g-6)(m_6^g + m_3^g - m_6^g + m_2^g - m_6^g + m_1^g - m_2^g - m_3^g + m_6^g) = (12g-6)m_1^g.$$

□

g	Number of Extremal Surface-Disk Pairs with Genus g and $n = 1$ Cusp
1	1
2	105
3	50050
4	56581525
5	117123756750
6	386078943500250
7	1857039718236202500
8	12277353837189093778125
9	1068157066843978245571937 50
10	11831975829430747026200351 68750
11	16259070931137207808967206912537500
12	271431639969559736697533380065719781250
13	5410885346008569674243521188002406254687500

Table 3: Count of once-punctured extremal surface-disk pairs

8.2 AN EXTREMAL CONJECTURE

Girondo & González-Diez show [14] that in the genus $g = 2$ closed case there are many extremal surfaces which contain multiple extremal disks: several surfaces with 2 extremal disks and one extremal surface with 4 extremal disks. They demonstrate that, for each such surface with non-unique extremal disks, its group of self-isometries acts transitively on its collection of extremal disks. Furthermore, Nakamura [20] similarly observes in the $g = 3$ closed case that, for each surface that admits multiple extremal disks, there exists a self-isometry of the surface taking one to the other.

The consistency of these findings with our n-punctured sphere discoveries (which also demonstrate a self-isometry linking the extremal disks of every punctured sphere containing multiples) heuristically suggests the following conjecture.

Conjecture 8.3. *The group of self-isometries of an extremal surface acts transitively on the set of its extremal disks.*

As alluded to before, the veracity of this conjecture would in particular immediately promote our upper bound of once-punctured genus g extremal surfaces (given by $\frac{(2g-1)(6g-5)!}{12^{g-1}g!(3g-3)!}$) into an exact count.

BIBLIOGRAPHY

[1] Roland Bacher and Alina Vdovina. Counting 1-vertex triangulations of oriented surfaces. *Discrete Math.*, 246(1-3):13–27, 2002. Formal power series and algebraic combinatorics (Barcelona, 1999).

[2] Christophe Bavard. Disques extrémaux et surfaces modulaires. *Ann. Fac. Sci. Toulouse Math. (6)*, 5(2):191–202, 1996.

[3] Alan F. Beardon. *The geometry of discrete groups*, volume 91 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995. Corrected reprint of the 1983 original.

[4] B. H. Bowditch and D. B. A. Epstein. Natural triangulations associated to a surface. *Topology*, 27(1):91–117, 1988.

[5] William G. Brown. Enumeration of triangulations of the disk. *Proceedings of the London Mathematical Society*, s3-14(4):746–768, 1964.

[6] Andrew J. Casson and Steven A. Bleiler. *Automorphisms of surfaces after Nielsen and Thurston*, volume 9 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1988.

[7] Jason DeBlois. Bounds for several-disk packings of hyperbolic surfaces. Preprint. arXiv:1701.07770.

[8] Jason DeBlois. The centered dual and the maximal injectivity radius of hyperbolic surfaces. *Geom. Topol.*, 19(2):953–1014, 2015.

[9] Norman Do. Moduli spaces of hyperbolic surfaces and their weil-petersson volumes. arXiv:1103.4674.

[10] Herbert Edelsbrunner. *Geometry and topology for mesh generation*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 1 edition, 2001.

[11] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.

[12] William Feller. *An introduction to probability theory and its applications. Vol. I.* Third edition. John Wiley & Sons, Inc., New York-London-Sydney, 1968.

[13] Ernesto Girondo and Gabino González-Diez. On extremal discs inside compact hyperbolic surfaces. *C. R. Acad. Sci. Paris Sér. I Math.*, 329(1):57–60, 1999.

[14] Ernesto Girondo and Gabino González-Diez. Genus two extremal surfaces: extremal discs, isometries and Weierstrass points. *Israel J. Math.*, 132:221–238, 2002.

[15] Ernesto Girondo and Gabino González-Diez. On extremal Riemann surfaces and their uniformizing Fuchsian groups. *Glasg. Math. J.*, 44(1):149–157, 2002.

[16] Svetlana Katok. *Fuchsian groups.* Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 1992.

[17] Wolfgang Kühnel. *Differential geometry*, volume 77 of *Student Mathematical Library.* American Mathematical Society, Providence, RI, 2015. Curves—surfaces—manifolds, Third edition [of MR1882174], Translated from the 2013 German edition by Bruce Hunt, with corrections and additions by the author.

[18] Bernard Maskit. On Poincaré's theorem for fundamental polygons. *Advances in Math.*, 7:219–230, 1971.

[19] Bernard Maskit. *Kleinian groups*, volume 287 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences].* Springer-Verlag, Berlin, 1988.

[20] Gou Nakamura. Extremal disks and extremal surfaces of genus three. *Kodai Math. J.*, 28(1):111–130, 2005.

[21] Gou Nakamura. Corrections to "Extremal disks and extremal surfaces of genus three" [mr2122195]. *Kodai Math. J.*, 37(2):475–480, 2014.

[22] John G. Ratcliffe. *Foundations of Hyperbolic Manifolds.* Graduate Texts in Mathematics 149. Springer New York, 2nd ed edition, 1994.

[23] Daniel Ruberman. Mutation and volumes of knots in S^3. *Invent. Math.*, 90(1):189–215, 1987.

[24] N. J. A. Sloane, editor. *The On-Line Encyclopedia of Integer Sequences.* Published electronically at https://oeis.org.

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

[26] Alina Vdovina. On the number of optimal surfaces. In *The Zieschang Gedenkschrift*, volume 14 of *Geom. Topol. Monogr.*, pages 557–567. Geom. Topol. Publ., Coventry, 2008.

[27] Wolfram Research, Inc. Mathematica 10.4.