msp

Geometry & Topology 19 (2015) 953–1014

The centered dual and the maximal injectivity radius of hyperbolic surfaces

JASON DEBLOIS

We give sharp upper bounds on the maximal injectivity radius of finite-area hyperbolic surfaces and use them, for each $g \geq 2$, to identify a constant $r_{g-1,2}$ such that the set of closed genus-g hyperbolic surfaces with maximal injectivity radius at least r is compact if and only if $r > r_{g-1,2}$. The main tool is a version of the centered dual complex that we introduced earlier, a coarsening of the Delaunay complex. In particular, we bound the area of a compact centered dual two-cell below given lower bounds on its side lengths.

52C15, 57M50

This paper analyzes the centered dual complex of a locally finite subset $\mathcal{S}$ of $\mathbb{H}^2$, first introduced in our prior preprint [6], and applies it to describe the maximal injectivity radius of hyperbolic surfaces. The centered dual complex is a cell decomposition with vertex set $\mathcal{S}$ and totally geodesic edges. Its underlying space contains that of the geometric dual to the Voronoi tessellation. We regard it as a tool for understanding the geometry of packings.

The rough idea behind the construction is that geometric dual 2–cells that are not centered (see Definition 0.2) are hard to analyze individually but naturally group into larger cells that can be treated as units. Our first main theorem bears the fruit of this approach, turning a lower bound on edge lengths into a good lower bound on area for centered dual 2–cells.

Theorem 3.31 *Let C be a compact two-cell of the centered dual complex of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$, such that for some fixed $d > 0$ each edge of ∂C has length at least d. If C is a triangle then its area is at least that of an equilateral hyperbolic triangle with side lengths d. If ∂C has $k > 3$ edges, then*

$$\mathrm{Area}(C) \geq (k-2)A_m(d).$$

Here $A_m(d)$ is the maximum of areas of triangles with two sides of length d, that of a semicyclic *triangle, whose third side is a diameter of its circumcircle.*

Published: 10 April 2015 DOI: 10.2140/gt.2015.19.953

The bounds of Theorem 3.31 do not hold for arbitrary Delaunay or geometric dual cells, even triangles. The theorem further gives explicit form to the assertion that $\mathcal{S}$ has low density in a centered dual two-cell of high combinatorial complexity. We prove an analog of Theorem 3.31 for centered dual 2–cells of finite complexity that are not compact in Theorem 4.16.

Our next main theorem, which uses Theorems 3.31 and 4.16, illustrates the sort of application we have in mind for the centered dual complex. Below let $\mathit{injrad}_x\, F$ denote the *injectivity radius* of a hyperbolic surface F at $x \in F$, half the length of the shortest non-constant geodesic arc in F that begins and ends at x.

Theorem 5.11 *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 $2r$:*

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

A complete, oriented, finite-area hyperbolic surface F 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

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

Moreover, the collection of such surfaces with injectivity radius $r_{g,n}$ at some point is a non-empty finite subset of the moduli space $\mathfrak{M}_{g,n}$ of complete, oriented, finite-area hyperbolic surfaces of genus g with n cusps.

The closed (ie, $n = 0$) case of Theorem 5.11 was proved by Christophe Bavard [1]. It follows from Böröczky's theorem [3], which bounds the local density of constant-radius packings of $\mathbb{H}^2$, since a disk embedded in a hyperbolic surface has as its preimage a packing of the universal cover $\mathbb{H}^2$ with constant local density. We reproduce this argument in Lemma 5.10.

The general case does not follow in the same way, since the preimage of a maximal-radius embedded disk on a noncompact hyperbolic surface is not a maximal-density packing of $\mathbb{H}^2$.

By basic calculus, α and β are decreasing functions of r with $\alpha(r) < \beta(r)$ for each $r > 0$. Thus if $g' \leq g$ and $n' \leq n$ then $r_{g',n'} \leq r_{g,n}$. It also happens that $2\beta(r) < 3\alpha(r)$ for each $r > 0$ (see Corollary 5.15), whence $r_{g-1,n+2} < r_{g,n}$ for any $g > 0$ and $n \geq 0$. Therefore

$$r_{0,2g} < r_{1,2g-2} < \cdots < r_{g-1,2} < r_{g,0} \tag{0.0.1}$$

for any $g \geq 2$. This relates the upper bounds of Theorem 5.11 on maximal injectivity radius of surfaces with a fixed even Euler characteristic. It implies compactness results for certain subsets of moduli space. Below we use the topology of geometric convergence on $\mathfrak{M}_g \doteq \mathfrak{M}_{g,0}$ (see Benedetti and Petronio [2, Section E.1]). This is the usual, algebraic topology on $\mathfrak{M}_g$ (compare eg [8, Section 10.3]).

Corollary 0.1 *For $g \geq 2$, the collection of surfaces of maximal injectivity radius at least r,*

$$\mathfrak{C}_{\geq r,g} \doteq \{F \text{ orientable, closed and hyperbolic} \mid \mathit{injrad}_x\, F \geq r \text{ for some } x \in F\},$$

is a compact subset of $\mathfrak{M}_g$ if and only if $r > r_{g-1,2}$.

Corollary 0.1 contrasts with Mumford's compactness criterion [10], which asserts compactness for any $\epsilon > 0$ of the subset of $\mathfrak{M}_g$ consisting of surfaces with *minimal* injectivity radius at least ϵ. However it is a standard consequence of the Margulis lemma that $\mathfrak{C}_{\geq \epsilon_2,g} = \mathfrak{M}_g$ (and hence is noncompact), where ϵ_2 is the 2–dimensional Margulis constant. On the other hand, by Theorem 5.11 $\mathfrak{C}_{\geq r_g,0}$ is finite and hence compact.

We will sketch a proof of Corollary 0.1 below that uses Theorem 5.11 and standard results on geometric convergence (eg from [2, Chapter E]). Details can be easily filled in.

Proof of Corollary 0.1 It is a key fact that if (F, x) is a pointed geometric limit of $\{(F_n, x_n)\}$, then $\mathit{injrad}_x\, F = \lim_{n\to\infty} \mathit{injrad}_{x_n}\, F_n$. This implies that $\mathfrak{C}_{\geq r,g}$ is closed in $\mathfrak{M}_g$. For $r > r_{g-1,2}$ we will show that it is also bounded; ie contained in one of the Mumford sets above.

Let $\{F_n\}$ be a sequence of closed, oriented, genus-g hyperbolic surfaces with (minimal) injectivity radius approaching 0, and for each n fix $x_n \in F_n$ at which injectivity radius attains a maximum. A subsequence of $\{(F_n, x_n)\}$ has a geometric limit (F, x), where F is a non-compact hyperbolic surface with $\mathrm{Area}(F) \leq \mathrm{Area}(F_n)$, hence $\chi(F) \geq 2-2g$, and $x \in F$. Then $\mathit{injrad}_x\, F \leq r_{g-1,2}$ by (0.0.1). Thus by the key fact the F_n are not all in $\mathfrak{C}_{\geq r,g}$ for any $r > r_{g-1,2}$.

Thus $\mathfrak{C}_{\geq r,g}$ is closed and bounded in $\mathfrak{M}_g$, hence compact, for $r > r_{g-1,2}$. Example 5.16 describes a sequence in $\mathfrak{C}_{\geq r_{g-1,2},g}$ with minimal injectivity radius approaching 0, showing that it is not compact. □

It is straightforward to extend Corollary 0.1 to moduli spaces of non-compact surfaces, or the bounds of Theorem 5.11 to multiple-disk, equal-radius packings on surfaces. In future work we will apply the centered dual machine to more subtle packing problems on surfaces.

We now give a brief overview of the paper. Section 1 recalls basic properties of the Voronoi tessellation of a locally finite subset $\mathcal{S}$ of $\mathbb{H}^n$ and its geometric dual complex, before pointing out some special features of the two-dimensional setting. Lemma 1.5 includes the key fact that every geometric dual 2–cell is *cyclic*: inscribed in a metric circle. Hence it is determined up to isometry by its collection of side lengths (Schlenker [12]).

The centered dual complex of $\mathcal{S}$ is defined in Section 2. This runs parallel to [6, Section 3], but the definitions are modified to accommodate non-compact Voronoi edges. The fact that motivates our definition is that among cyclic polygons in $\mathbb{H}^2$, increasing the length of an edge increases area if and only if that edge is not the longest of a non-centered polygon; see DeBlois [7]. Here is the definition of a centered polygon (cf Definition 1.3).

Definition 0.2 A polygon P inscribed in a circle S is *centered* if the center of S is in *int* P.

The centered/non-centered dichotomy has been previously considered in the literature, eg in Vanderzee, Hirani, Guoy and Ramos [13] (there centered goes by well-centered). Centered dual two-cells collect non-centered two-cells of the geometric dual in a natural way. Two fundamental observations here are Lemma 2.5, relating non-centeredness of geometric dual cells to non-centeredness of Voronoi edges (see Definition 2.1), and Lemma 2.7, describing the structure of the set of these edges.

Centered dual 2–cells are not determined by their edge lengths, but the set of possible centered dual two-cells with a given combinatorics and edge length collection is parametrized by a compact *admissible space*. This is defined in Section 3.2, which parallels Section 5 of [6]. The area of centered dual 2–cells determines a function on the admissible space. Theorem 3.31 is proved in Section 3.3 by bounding this function below.

Section 4 has the same structure as Section 3. It describes admissible spaces for non-compact centered dual 2–cells and finishes with a proof of Theorem 4.16. We finally consider hyperbolic surfaces in Section 5, proving Theorem 5.11 there and describing some examples.

Acknowledgements This work was prompted by an offhand observation of Marc Culler. Thanks to Hugo Parlier for pointing out [1] to me, and thanks to the anonymous referee for numerous helpful comments.

Partially supported by NSF grant number DMS-1240329.

1 The Voronoi tessellation and its geometric dual

In this section we will record some facts about the Voronoi tessellation of a locally finite subset $\mathcal{S}$ of hyperbolic space and its geometric dual, using DeBlois [5] as a general reference. We will also establish notation and collect some facts that hold only in the 2–dimensional setting.

The *Voronoi tessellation* has n–cells in bijection with $\mathcal{S}$. The assertions below are from [5, Lemma 5.2]. For $\boldsymbol{s} \in \mathcal{S}$, the corresponding Voronoi n–cell is the convex polyhedron

$$V_{\boldsymbol{s}} = \{\boldsymbol{x} \in \mathcal{H}^n \mid d_H(\boldsymbol{s}, \boldsymbol{x}) \le d_H(\boldsymbol{s}', \boldsymbol{x}) \text{ for all } \boldsymbol{s}' \in \mathcal{S}\}.$$

Here d_H is the hyperbolic distance. The collection of Voronoi n–cells is locally finite, and cells of lower dimension are by definition of the form $\bigcap_{i=0}^{n} V_{\boldsymbol{s}_i}$ for subsets $\{\boldsymbol{s}_0, \dots, \boldsymbol{s}_n\}$ of $\mathcal{S}$.

The result below, from [5, Corollary 5.5], identifies the geometric dual to a Voronoi cell.

Proposition 1.1 *Let $\mathcal{S} \subset \mathbb{H}^n$ be locally finite. For a k–cell V of the Voronoi tessellation, if $\mathcal{S}_0 \subset \mathcal{S}$ is maximal such that $V = \bigcap_{\boldsymbol{s} \in \mathcal{S}_0} V_{\boldsymbol{s}}$ then the closed convex hull C_V of $\mathcal{S}_0$ in $\mathbb{H}^n$ is the* geometric dual *to V, an $(n-k)$–dimensional, compact convex polyhedron in $\mathbb{H}^n$.*

For a locally finite set $\mathcal{S}$, say the *geometric dual complex* of $\mathcal{S}$ is the collection of geometric duals to Voronoi cells. The result below shows it is a *polyhedral complex* in the sense of De Loera, Rambau and Santos [4, Definition 2.1.5], and characterizes it by an empty circumspheres condition.

Theorem 1.2 [5, Theorem 5.9] *Suppose $\mathcal{S} \subset \mathbb{H}^n$ is locally finite. For any metric sphere S that intersects $\mathcal{S}$ and bounds a ball B with $B \cap \mathcal{S} = S \cap \mathcal{S}$, the closed convex hull of $S \cap \mathcal{S}$ in $\mathbb{H}^n$ is a geometric dual cell. Every geometric dual cell is of this form. Moreover, if C is the geometric dual to a Voronoi cell then so is every face of C, and any geometric dual cell $C' \neq C$ that intersects C does so in a face of each.*

We now specialize to dimension 2 and make some definitions.

Definition 1.3 A polygon $C \subset \mathbb{H}^2$ is *cyclic* if its vertex set is contained in a metric circle S, its *circumcircle*. The *center* $v \in \mathbb{H}^2$ and *radius* $J > 0$ of a cyclic n–gon C are respectively the center and radius of S (so $S = \{\boldsymbol{x} \mid d_H(v, \boldsymbol{x}) = J\}$). C is *centered* if $v \in \mathit{int}\, C$.

The vertex set of a cyclic polygon C is *cyclically ordered* $\mathcal{S}_0 = \{\mathbf{s}_0, \ldots, \mathbf{s}_{n-1}\}$ if with the boundary orientation from C, an edge points from $\mathbf{s}_i$ to $\mathbf{s}_{i+1}$ for each i (taking $i+1$ modulo n). With its vertices cyclically ordered as above, the *side length collection* of C is $(d_0, \ldots, d_{n-1})$, where $d_0 = d(\mathbf{s}_0, \mathbf{s}_{n-1})$ and $d_i = d(\mathbf{s}_{i-1}, \mathbf{s}_i)$ for each $i > 0$.

The collection of Voronoi edges containing a Voronoi vertex v is *cyclically enumerated* $e_0, \ldots, e_{n-1}$ if the vertex set of C_v can be cyclically ordered $\{\mathbf{s}_0, \ldots, \mathbf{s}_{n-1}\}$ so that $e_i = V_{\mathbf{s}_{i-1}} \cap V_{\mathbf{s}_i}$ for each i (taking $i-1$ modulo n).

Remark 1.4 We will often refer to [7] for results on cyclic polygons. Definition 1.1 there defines one as a cyclically ordered finite subset of a hyperbolic circle, but by Lemma 2.1 there such a cyclic polygon is the vertex set of one defined as above and vice-versa.

Lemma 1.5 *For a vertex v of the Voronoi tessellation of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$, the geometric dual C_v to v is a cyclic polygon with vertex set $\mathcal{S}_0 \subset \mathcal{S}$ such that for $\mathbf{s} \in \mathcal{S}$, $v \in V_{\mathbf{s}}$ if and only if $\mathbf{s} \in \mathcal{S}_0$. C_v has center v and radius $J_v \doteq d(v, \mathbf{s})$ for any $\mathbf{s} \in \mathcal{S}_0$ and:*

- *If $\mathcal{S}_0 = \{\mathbf{s}_0, \ldots, \mathbf{s}_{n-1}\}$ is cyclically ordered then the Voronoi 2–cell $V_{\mathbf{s}_i}$ shares an edge e_i with $V_{\mathbf{s}_{i+1}}$ for each i (taking $i+1$ modulo n).*
- *For each $i \in \{0, \ldots, n-1\}$, the geometric dual γ_i to e_i as above joins $\mathbf{s}_i$ to $\mathbf{s}_{i+1}$.*

For $v \neq w$, $\mathit{int}\, C_v \cap \mathit{int}\, C_w = \varnothing$, and C_v shares an edge with C_w if and only if v and w are opposite endpoints of a Voronoi edge.

That the geometric dual to a Voronoi vertex is cyclic follows from Theorem 1.2. Proposition 1.1 implies $\mathit{int}\, C_v \cap \mathit{int}\, C_w$ for $v \neq w$. Together with the definitions here, it also implies the fact below, which is useful to record separately:

Fact 1.6 Say the *radius* of a Voronoi vertex v is the radius J_v of its geometric dual C_v. For every $\mathbf{s} \in \mathcal{S}$, $d_H(v, \mathbf{s}) \geq J_v$ and equality holds if and only if $\mathbf{s}$ is a vertex of C_v.

In two dimensions the vertex set of any polygon admits a cyclic order. The remaining assertions of Lemma 1.5 follow from [5, Lemma 5.8]. The facts below are straightforward:

Facts Suppose $\mathcal{S} \subset \mathbb{H}^2$ is locally finite.

- Each Voronoi edge is the intersection of exactly two Voronoi 2–cells $V_{\mathbf{s}}$ and $V_{\mathbf{t}}$, for $\mathbf{s}, \mathbf{t} \in \mathcal{S}$, and its geometric dual is the arc $\gamma_{\mathbf{st}}$ joining $\mathbf{s}$ to $\mathbf{t}$.
- Each Voronoi vertex v is the intersection of at least three Voronoi 2–cells.

The geometric dual complex of a locally finite set $\mathcal{S}$ is a subcomplex of what we call the *Delaunay tessellation* in [5], whose underlying space contains the convex hull of $\mathcal{S}$. In important special cases (eg if $\mathcal{S}$ is finite or lattice-invariant; see respectively [5, Proposition 3.5 or Theorem 6.23]), the Delaunay tessellation is a locally finite polyhedral complex. It is important to note that the geometric dual may be a proper subcomplex even in good conditions; see below, which reproduces [5, Example 5.11].

Example 1.7 Figure 1 illustrates the Voronoi and Delaunay tessellations determined by three points in $\mathbb{H}^2$, using the upper half-plane model. In each case the Delaunay triangle spanned by x, y and z is shaded, with its edges dashed. The edges of the Voronoi tessellation are in bold. The Euclidean circumcircle for $\boldsymbol{x}$, $\boldsymbol{y}$ and $\boldsymbol{z}$ is also included in each case.

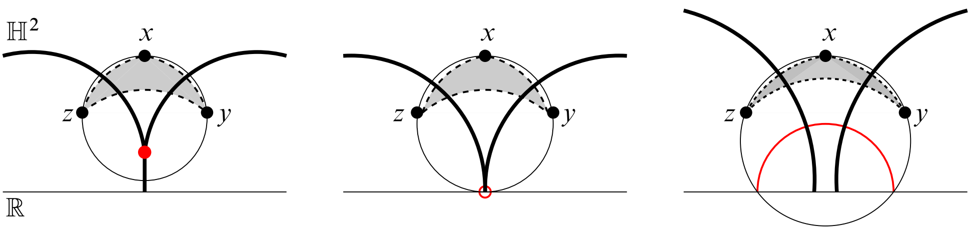

Figure 1: Delaunay and Voronoi tessellations of three-point sets in $\mathbb{H}^2$

In the left case the Delaunay tessellation and the geometric dual complex coincide. In particular, the Delaunay triangle is the geometric dual to the Voronoi vertex: the red dot. In the middle and on the right, the Voronoi tessellation has no vertex and the Delaunay triangle has no geometric dual; instead, the geometric dual to the Voronoi tessellation has cells $\boldsymbol{x}$, $\boldsymbol{y}$, $\boldsymbol{z}$, and the two edges containing $\boldsymbol{x}$.

This trichotomy reflects that the Euclidean circumcircle for $\boldsymbol{x}$, $\boldsymbol{y}$ and $\boldsymbol{z}$ is a metric hyperbolic circle in the left case, centered at the red dot, and intersects $\mathbb{H}^2$ in a horocycle and geodesic equidistant, respectively, in the middle and right cases. In particular, the triangle spanned by $\boldsymbol{x}$, $\boldsymbol{y}$ and $\boldsymbol{z}$ is cyclic only in the left-hand case.

Let us make some precise definitions connected with the upper half-plane model for $\mathbb{H}^2$.

Definition 1.8 The *upper half-plane model* for $\mathbb{H}^2$ is $\{\boldsymbol{z} \in \mathbb{C} \mid \Im \boldsymbol{z} > 0\}$, equipped with the inner product $\langle \boldsymbol{v}, \boldsymbol{w} \rangle = \frac{\boldsymbol{v} \cdot \boldsymbol{w}}{\Im z}$ for $\boldsymbol{v}, \boldsymbol{w} \in T_z\mathbb{H}^2$.

The *sphere at infinity* of $\mathbb{H}^2$ is $S_\infty = \mathbb{R} \cup \{\infty\}$. For $r \in \mathbb{R}$, a *horocycle* S with *ideal point* r is the non-empty intersection with $\mathbb{H}^2$ of a Euclidean circle in $\mathbb{C}$ tangent to $\mathbb{R}$ at r. The *horoball* B bounded by S is the intersection with $\mathbb{H}^2$ of the Euclidean ball that S bounds. A horocycle centered at ∞ is a horizontal line in $\mathbb{H}^2$, and the horoball that it bounds is the half-plane contained in $\mathbb{H}^2$.

Geodesics of the upper half-plane model are the intersections with $\mathbb{H}^2$ of Euclidean circles and straight lines that meet $\mathbb{R}$ perpendicularly. Every geodesic ray thus has a well-defined ideal endpoint in S_∞ (if it points up in a straight line, its ideal endpoint is ∞).

The isometry group of $\mathbb{H}^2$ is $\mathrm{PGL}_2(\mathbb{R})$, acting by Möbius transformations. It takes geodesics to geodesics and horocycles to horocycles and extends to a triply transitive action on S_∞.

We conclude this section with two technical lemmas on infinite-length Voronoi edges.

Lemma 1.9 *For a locally finite set $\mathcal{S} \subset \mathbb{H}^2$, if a Voronoi edge $e = V_{\boldsymbol{s}} \cap V_{\boldsymbol{t}}$ with $\boldsymbol{s}, \boldsymbol{t} \in \mathcal{S}$ has an ideal endpoint $v_\infty \in S_\infty$ then there is a unique horocycle S through $\boldsymbol{s}$ and $\boldsymbol{t}$ with ideal point v_∞, and the horoball B that it bounds satisfies $B \cap \mathcal{S} = S \cap \mathcal{S}$.*

Proof We work in the upper half-plane model. After moving $\mathcal{S}$ by an isometry, e is a subinterval $[iy_0, \infty)$ of $i\mathbb{R}^+$ and $v_\infty = \infty$. Each horocycle with ideal point ∞, being a horizontal line, is preserved by reflection ρ through $i\mathbb{R}^+$. Since $i\mathbb{R}^+$ perpendicularly bisects the geometric dual γ to e, ρ preserves γ and exchanges its endpoints $\boldsymbol{s}$ and $\boldsymbol{t}$. They thus lie on the same horocycle through ∞. Moving $\mathcal{S}$ again, by an isometry preserving $i\mathbb{R}^+$, we may assume this is $S_\infty = \mathbb{R} + i$; so $\boldsymbol{s} = -x_0 + i$, $\boldsymbol{t} = x_0 + i$ for some $x_0 > 0$.

For each $u \geq y_0$, the hyperbolic circle S_u centered at $\boldsymbol{u} = iu$ containing $\boldsymbol{x}$ and $\boldsymbol{y}$ has no points of $\mathcal{S}$ in the interior of the disk that it bounds, since $\boldsymbol{u} \in V_{\boldsymbol{s}} \cap V_{\boldsymbol{t}}$. Direct computation reveals that this hyperbolic circle is identical to the Euclidean circle of radius $u \sinh r_u$ centered at $(0, u\cosh r_u)$, where $r_u = d(\boldsymbol{u}, \boldsymbol{t})$ satisfies $\cosh r_u = (x_0^2 + u^2 + 1)/(2u)$. (Recall that circles of the upper half-plane model are Euclidean circles contained in $\mathbb{H}^2$.)

The convex complementary component to S_∞ is $\{x + iy \mid y > 1\}$. For $z = x + iy$ in this complementary component, we claim there exists $u_1 \geq y_0$ such that S_u encloses z for all $u > u_1$. This is obvious if $|x| \leq x_0$, taking $u_1 = y$, say, so assume that $|x| > x_0$. For a point $x + iy_u$ on S_u, the Euclidean distance formula gives

$$u^2 \sinh^2 r_u = x^2 + (u\cosh r_u - y_u)^2. \tag{1.9.1}$$

Solving for $y_u < u\cosh r_u$ and substituting for $\cosh u$, a little manipulation gives

$$\begin{aligned} y_u &= u\cosh r_u - \sqrt{u^2 \sinh^2 r_u - x^2} \\ &= (x_0^2 + u^2 + 1)/2 - \sqrt{(x_0^2 + u^2 - 1)^2/4 - (x^2 - x_0^2)}. \end{aligned}$$

Fixing any $b \in \mathbb{R}$ and taking $a \to \infty$,

$$a - \sqrt{a^2 - b^2} = \frac{b^2}{a + \sqrt{a^2 - b^2}} \to 0.$$

Thus taking $a = (x_0^2 + u^2 - 1)/2$ and $b = \sqrt{x^2 - x_0^2}$, we find that for any fixed $\epsilon > 0$, if u is large enough the square rooted quantity in the equation for y_u is at least $(x_0^2 + u^2 - 1)/2 - \epsilon$, hence $y_u < 1 + \epsilon$.

Such a solution y_u is bounded below by 1, so it is clear that $y_u \to 1$ as $u \to \infty$. A simpler argument shows that the solution $y_u > u \cosh r_u$ to (1.9.1) increases without bound as $u \to \infty$, and the claim follows. But the claim implies the result since for any $\boldsymbol{u} \in e$, no point of $\mathcal{S}$ has distance less than $d(\boldsymbol{u}, \boldsymbol{t})$ from $\boldsymbol{u}$. $\square$

Definition 1.10 If $\boldsymbol{s}$ and $\boldsymbol{t} \in \mathbb{H}^2$ lie on a horocycle S with ideal point v, the *horocyclic ideal triangle* with vertices $\boldsymbol{s}$, $\boldsymbol{t}$ and v is the convex hull in $\mathbb{H}^2$ of the geodesic rays from $\boldsymbol{s}$ and $\boldsymbol{t}$ with ideal endpoint v.

Lemma 1.11 *For a locally finite set $\mathcal{S} \subset \mathbb{H}^2$, if a Voronoi edge $e = V_{\boldsymbol{s}_0} \cap V_{\boldsymbol{t}_0}$ has an ideal endpoint v_∞, let $\Delta(e, v_\infty)$ be the horocyclic ideal triangle with vertices at $\boldsymbol{s}_0$, $\boldsymbol{t}_0$ and v_∞. If e has an endpoint $v_0 \in \mathbb{H}^2$ then $C_{v_0} \cap \Delta(e, v_\infty) = \gamma$, where γ is the geometric dual to e. For any other Voronoi vertex v, $C_v \cap \Delta(e, v_\infty) \subset \partial\gamma$.*

Proof Working in the upper half-plane model and moving $\mathcal{S}$ by an isometry, we will take $v_\infty = \infty$, $\boldsymbol{s}_0 = -x_0 + i$ and $\boldsymbol{t}_0 = x_0 + i$ for some $x_0 > 0$. The horocycle through $\boldsymbol{s}_0$ and $\boldsymbol{t}_0$ with ideal point v_∞ is $S_\infty = \mathbb{R} + i$, so by Lemma 1.9 every $z = x + iy \in \mathcal{S}$ has $y \leq 1$.

The geodesic through $\boldsymbol{s}_0$ and $\boldsymbol{t}_0$ is the intersection with $\mathbb{H}^2$ of the Euclidean circle through them centered at the origin. It intersects the horoball B bounded by S_∞ in γ, and separates all other vertices of C_{v_0} from $\Delta(e, v_\infty)$. This is because they lie outside B on the circumcircle of C_{v_0}, a circle in $\mathbb{H}^2$ containing $\boldsymbol{s}_0$ and $\boldsymbol{t}_0$, hence with Euclidean center on the positive imaginary axis. It follows that $C_{v_0} \cap \Delta(e, v_\infty) = \gamma$.

Theorem 1.2 implies that a geometric dual 2–cell C_v intersects the interior of γ only if γ is a face of C_v, hence only if v is an endpoint of e. Therefore for any v not in e, if C_v intersects $\Delta(e, v_\infty)$ outside $\partial\gamma$ then C_v intersects $\Delta(e, v_\infty) - \gamma$. It follows that an edge λ of C_v also intersects $\Delta(e, v_\infty) - \gamma$. Since λ does not cross γ it lies in a Euclidean circle centered in $\mathbb{R}$ with the property that at least one of $\boldsymbol{s}_0$ and $\boldsymbol{t}_0$ lies in the interior of the disk it bounds. We claim that the circumcircle of C_v has the same property, contradicting the empty circumcircles condition of Theorem 1.2.

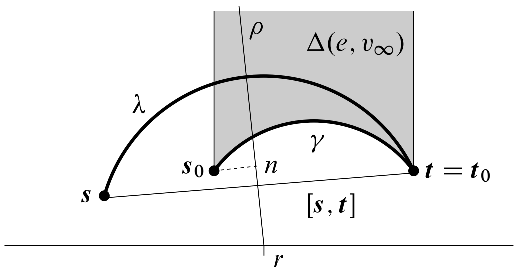

Figure 2: Some objects from the proof of Lemma 1.11

Let s and t be the endpoints of λ, and assume s_0 is in the interior of the disk bounded by the Euclidean circle containing λ. The line segment $[s,t]$ is a chord of this circle that separates its center $r \in \mathbb{R}$ from s_0, since by Lemma 1.9 each of s and t has imaginary part at most one. Since the circumcircle of C_v lies in $\mathbb{H}^2$ and contains s and t its Euclidean center lies on the ray ρ from r that perpendicularly bisects $[s,t]$. (See Figure 2.)

Let n be the nearest point on ρ to s_0, and let $d_0 = \mathrm{dist}(s_0, n)$ and $\ell_0 = \mathrm{dist}(n, r)$. Let $d = \mathrm{dist}(s, \rho \cap [s,t])$ and $\ell = \mathrm{dist}(\rho \cap [s,t], r)$. (All distances measured in the Euclidean metric.) Since s_0 is inside the disk centered at r and containing s, we have $d_0^2 + \ell_0^2 < d^2 + \ell^2$. Also, since $[s,t]$ separates s_0 from r we have $\ell_0 > \ell$, whence $d > d_0$.

If the Euclidean center c of the circumcircle of C_v is on ρ between its intersection with $[s,t]$ and r then the Euclidean distance squared to s (respectively, s_0) is $(\ell-\epsilon)^2 + d^2$ (resp. $(\ell_0-\epsilon)^2 + d_0^2$) for some $\epsilon > 0$. But since $\ell_0 > \ell$, $\ell_0^2 - (\ell_0-\epsilon)^2 > \ell^2 - (\ell-\epsilon)^2$ so the Euclidean distance from s to c is still larger than the distance from s_0 to c. If c is between $\rho \cap [s,t]$ and n, then its distance to s_0 is at most

$$d_0^2 + (\ell_0 - \ell)^2 < d_0^2 + \ell_0^2 - \ell^2 < d^2$$

which is less than its distance to s. If c is past the nearest point to s_0 then it is clearly closer to s_0 than to s. This proves the claim and hence the lemma. □

2 The centered dual to the Voronoi tessellation

The ultimate goal of this section is to show how geometric dual 2–cells that are not centered (in the sense of Definition 1.3) can be grouped to form 2–cells of the coarser centered dual decomposition; see Definition 2.26. We will later describe some advantages of the centered dual. A key tool in defining it is the notion of a (non-)centered Voronoi edge.

Definition 2.1 For a locally finite set $\mathcal{S} \subset \mathbb{H}^2$, we will say an edge e of the Voronoi tessellation of $\mathcal{S}$ is *centered* if e intersects its geometric dual edge $\gamma_{\boldsymbol{st}}$ at a point in $\mathit{int}\, e$. If e is not centered, we orient it pointing away from $\gamma_{\boldsymbol{st}}$.

We will refer to the one-skeleton of the Voronoi tessellation as the *Voronoi graph*, and to the union of its non-centered edges as the *non-centered Voronoi subgraph*.

Section 2.1 describes the structure of the non-centered Voronoi subgraph. In Section 2.2 we use it to organize the centered dual decomposition and prove its basic properties.

2.1 Non-centeredness in the Voronoi graph

The key results of this section are Lemma 2.5, which gives a dictionary between non-centered geometric dual 2–cells and non-centered Voronoi edges, and Lemma 2.7, which asserts that each component of the non-centered Voronoi subgraph is a tree, with a canonical root vertex if finite.

Fact 2.2 For locally finite $\mathcal{S} \subset \mathbb{H}^2$ and $\boldsymbol{s} \in \mathcal{S}$, an edge e of the Voronoi 2–cell $V_{\boldsymbol{s}}$ is non-centered with initial vertex v if and only if the angle α at v, measured in $V_{\boldsymbol{s}}$ between e and the geodesic segment joining v to $\boldsymbol{s}$, is at least $\pi/2$.

This is because there is a right triangle with vertices at $\boldsymbol{s}$ and v and edges contained in $\gamma_{\boldsymbol{st}}$ and $\gamma_{\boldsymbol{st}}^{\perp}$, where $\gamma_{\boldsymbol{st}}$ is the geometric dual to e. This triangle has angle equal to either α or $\pi-\alpha$ at v, depending on the case above; see Figure 3.

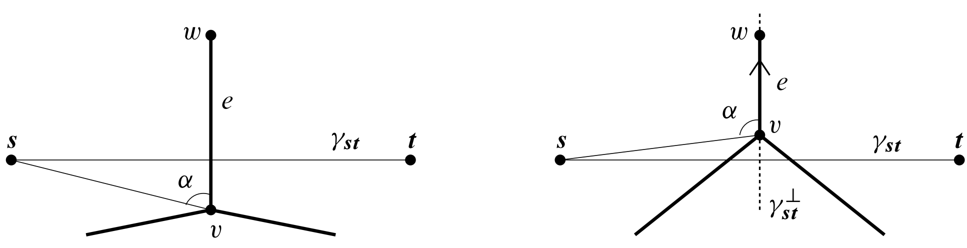

Figure 3: Centered and non-centered edges

If e has another endpoint w then since $\boldsymbol{s} \in C_v \cap C_w$, with α as in Fact 2.2, the hyperbolic law of cosines implies that the respective radii J_v and J_w of v and w (see Fact 1.6) satisfy

(2.2.1) $$\cosh J_w = \cosh \ell(e) \cosh J_v - \sinh \ell(e) \sinh J_v \cos\alpha.$$

Because $\cos\alpha \leq 0$ if $\alpha \geq \pi/2$, we have the following.

Lemma 2.3 *Suppose v is the initial and w the terminal vertex of a non-centered edge, oriented as prescribed in Definition 2.1, of the Voronoi tessellation of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$. Then $J_v < J_w$.*

Remark 2.4 While every non-centered edge of the Voronoi tessellation has an initial vertex, note that not every such edge has a terminal vertex in $\mathbb{H}^2$, as in the left-hand case of Figure 1. There all Voronoi edges are non-compact, and $V_y \cap V_z$ is non-centered.

Below we relate centeredness of edges of V to that of geometric dual 2–cells.

Lemma 2.5 *Let v be a vertex of the Voronoi tessellation of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$. Its geometric dual C_v is non-centered if and only if v is the initial vertex of a non-centered edge e of V. If this is so then the geometric dual γ to e is the unique longest edge of C_v.*

Proof Suppose first that v is the initial vertex of a non-centered edge $e = V_{\boldsymbol{s}} \cap V_{\boldsymbol{t}}$ with geometric dual γ joining $\boldsymbol{s}$ and $\boldsymbol{t}$, and let $\mathcal{H}'$ be the half-space containing e and bounded by the geodesic containing γ. If e is non-compact with ideal vertex v_∞ then the triangle $\Delta(e, v_\infty)$ of Lemma 1.11 intersects C_v in γ. But $\Delta(e, v_\infty)$ contains e and hence v in this case, since $e \subset \mathcal{H}'$, so C_v is not centered (recall Definition 1.3).

If e is compact we claim that the distance from the other endpoint w of e to any $z \in S \cap \mathcal{H}'$ is less than J_w, where S is the circle centered at v through $\boldsymbol{s}$ and $\boldsymbol{t}$. Thus applying the empty circumcircles condition Fact 1.6 to w ensures that no point of $\mathcal{S}$, in particular no vertex of C_v, lies on $S \cap \mathcal{H}' - \{\boldsymbol{s}, \boldsymbol{t}\}$. This implies that C_v is contained in the half-space $\mathcal{H}$ opposite $\mathcal{H}'$, and hence is non-centered.

The claim's proof is an exercise in hyperbolic trigonometry. If the angle at v between e and $\boldsymbol{s}$ or $\boldsymbol{t}$ is $\alpha \geq \pi/2$ then $S \cap \mathcal{H}'$ consists of $z \in S$ such that the angle α' at v between e and z is less than α (compare Figure 3). The formula (2.2.1) determines J_w, and applying the hyperbolic law of cosines to such $z \in S$ yields

$$\cosh d(z, w) = \cosh \ell(e) \cosh J_v - \sinh \ell(e) \sinh J_v \cos \alpha'.$$

Since $\alpha' < \alpha$, $\cos \alpha' > \cos \alpha$, and it follows that $d(z, w) < J_w$, proving the claim.

If C_v is not centered, then by [7, Proposition 2.2] its unique longest side is characterized by the fact that the geodesic containing it has C_v and v in opposite half-spaces. Thus assuming e is not centered its geometric dual γ is the longest side of C_v.

We now assume that C_v is not centered, take $\mathcal{H}$ and $\mathcal{H}'$ as above, and let e be the geometric dual to the longest side γ of C_v. The perpendicular bisector $\gamma^\perp$ of γ, which contains e, is divided by v into rays ρ_+ and ρ_-, with ρ_+ being the points of $\gamma^\perp$ further from $\mathcal{H}$ than v. We claim that the interior of e is contained in ρ_+, hence e is non-centered with initial vertex v.

If γ joins vertices $\boldsymbol{s}$ and $\boldsymbol{t}$ of C_v then every point in the interior of $e = V_{\boldsymbol{s}} \cap V_{\boldsymbol{t}}$ is closer to $\boldsymbol{s}$ and $\boldsymbol{t}$ than to any other point of $\mathcal{S}$, in particular, to the other vertices of C_v. All vertices of C_v lie in $S \cap \mathcal{H}$, where S is the circle centered at v through $\boldsymbol{s}$ and $\boldsymbol{t}$. Applying the hyperbolic law of cosines in an analogous way to the previous case shows that every point of ρ_+ is closer to $\boldsymbol{s}$ and $\boldsymbol{t}$ than to other points of $S \cap \mathcal{H}$, and this is reversed for points of ρ_-. The claim follows. □

If v is the initial vertex of a non-centered Voronoi edge e, the fact that the geometric dual to e is the *unique* longest edge of C_v immediately implies the following.

Corollary 2.6 *For a locally finite set $\mathcal{S} \subset \mathbb{H}^2$, no vertex of the Voronoi tessellation of $\mathcal{S}$ is the initial vertex of more than one non-centered edge.*

Below, given a graph G we will say that $\gamma = e_0 \cup e_1 \cup \cdots \cup e_{n-1}$ is an *edge path* if e_i is an edge of G for each i and $e_i \cap e_{i-1} \neq \varnothing$ for $i > 0$. An edge path γ as above is *reduced* if $e_i \neq e_{i-1}$ for each $i > 0$, and γ is *closed* if $e_0 \cap e_{n-1} \neq \varnothing$.

Lemma 2.7 *Each component T of the non-centered Voronoi subgraph determined by locally finite $\mathcal{S} \subset \mathbb{H}^2$ is a tree. Each compact reduced edge path γ of T has a unique vertex v_γ such that $J_{v_\gamma} > J_v$ for all vertices $v \neq v_\gamma$ of γ, and every edge of γ points toward v_γ.*

Proof Suppose that such a component T admits closed, reduced edge paths, and let $\gamma = e_0 \cup e_1 \cup \cdots \cup e_{n-1}$ be shortest among them. Orienting the e_i as in Definition 2.1, we may assume (after re-numbering if necessary) that e_0 points toward $e_0 \cap e_{n-1}$. We claim that then e_i points to $e_i \cap e_{i-1}$ for each $i > 0$ as well. Otherwise, for the minimal $i > 0$ such that e_i points toward e_{i+1} it would follow that the vertex $e_i \cap e_{i-1}$ was the initial vertex of both e_i and e_{i-1}, contradicting Corollary 2.6.

Let $v_0 = e_0 \cap e_{n-1} \in V^{(0)}$, and for $i > 1$ take $v_i = e_i \cap e_{i-1}$. Applying Lemma 2.3 to e_i for each i, we find that $J_{v_i} > J_{v_{i+1}}$. By induction this gives $J_{v_0} > J_{v_{n-1}}$; but since e_{n-1} points to v_{n-1} Lemma 2.3 implies that $J_{v_{n-1}}$ must exceed J_{v_0}, a contradiction. Thus T contains no closed, reduced edge paths, so it is a tree.

Let $\gamma = e_0 \cup \cdots \cup e_{n-1}$ be a reduced edge path, and let v_γ be a vertex with J_{v_γ} maximal. Assume for now that v_γ is on the boundary of γ, say the endpoint of e_0 not in e_1. Lemma 2.3 implies that e_0 points toward v_γ; thus if $i > 0$ were minimal such that e_i did not point toward v_γ then $v_i = e_i \cap e_{i-1}$ would the initial endpoint of e_i and e_{i-1}, contradicting Corollary 2.6. It follows that each edge of γ points toward v_γ, and by repeated application of Lemma 2.3, that $J_{v_T} > J_v$ for all vertices $v \neq v_\gamma$. The case that v_γ is in the interior of γ follows by applying the argument above to the compact subpaths obtained by splitting γ along v_γ. □

Definition 2.8 If a component T of the non-centered Voronoi subgraph determined by locally finite $\mathcal{S} \subset \mathbb{H}^2$ has a vertex v_T with maximal radius, we call it the *root vertex* of T.

If v_T is a root vertex of T, Lemma 2.7 immediately implies that $J_{v_T} > J_v$ for all $v \in T^{(0)} - \{v_T\}$. In particular, v_T is unique.

Proposition 2.9 *A component T of the non-centered Voronoi subgraph determined by locally finite $\mathcal{S} \subset \mathbb{H}^2$ has at most one non-compact edge.*

(1) *If one exists then its initial vertex is the root vertex v_T of T, and C_{v_T} is non-centered.*

(2) *If all edges are compact and there is a root vertex v_T, then C_{v_T} is centered.*

For every non-root vertex v of T, the geometric dual C_v is non-centered.

Remark 2.10 The left case of Figure 1 is an example of the phenomenon (1) above.

Proof A vertex v of T is contained in at least one non-centered Voronoi edge. If v is the initial point of a non-centered edge e, then by Lemma 2.5, C_v is non-centered.

If v is the initial vertex of a non-compact edge e of T then by Corollary 2.6, v is the terminal vertex of every other edge of T that contains it. In particular, for any $w \in T^{(0)} - \{v\}$, each edge of the unique reduced edge path γ in T joining v to w is compact, so the edge of γ that contains v points towards it. By Lemma 2.7 every other edge of γ points toward v as well, and $J_v > J_w$. Since w was arbitrary, it follows that $v = v_T$ is the root vertex of T. The uniqueness of the root vertex now implies that e is the unique non-compact edge of T.

If every edge of T is compact and v_T is a root vertex, then by Lemma 2.3 v_T is the terminal point of every edge of T that contains it. Hence Lemma 2.5 implies that C_{v_T} is centered. □

2.2 Introducing the centered dual

Recall from Proposition 1.1 that the geometric dual complex is dual to the Voronoi tessellation. The basic idea of this section is to coarsen the Voronoi tessellation by thinking of components of the non-centered Voronoi subgraph as large vertices, and make the centered dual complex dual to the result. In particular:

Definition 2.11 For a component T of the non-centered Voronoi subgraph of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$, we define the *centered dual* 2*–cell* C_T *dual to* T as follows:

(1) If T has a non-compact edge e_0 with ideal endpoint v_∞ (recall Definition 1.8), take
$$C_T = \Delta(e_0, v_\infty) \cup \Bigg(\bigcup_{v \in T^{(0)}} C_v \Bigg),$$
where $\Delta(e_0, v_\infty)$ is the horocyclic ideal triangle defined in Lemma 1.11.

(2) Otherwise, let $C_T = \bigcup_{v \in T^{(0)}} C_v$.

Define the *boundary* ∂C_T of C_T, in case (2) above, as the union of geometric duals γ to Voronoi edges e that are not in T but have an endpoint there, or, in case (1), the union of such γ with the infinite edges of $\Delta(e_0, v_\infty)$. Let the *interior* $\mathit{int}\, C_T$ of C_T be $C_T - \partial C_T$.

See Figure 4 for an example. Though the definition above applies to each component T of the non-centered Voronoi subgraph, we can only guarantee that it produces a true cell (a copy of $\mathbb{D}^2$ embedded on its interior) in the case that $T^{(0)}$ is finite. Indeed, Lemmas 2.13 and 2.18 and Proposition 2.23 as well as Corollary 2.24 below only hold in this case. This is the relevant case for the main results of this paper.

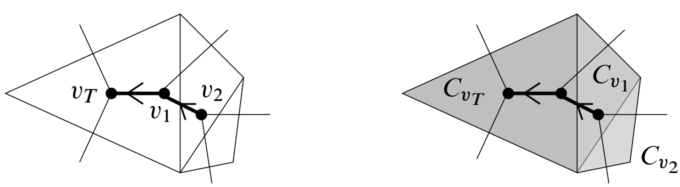

Figure 4: A two-edged component T of the non-centered Voronoi subgraph (in bold), and the geometric duals to its vertices (shaded). $C_T = C_{v_T} \cup C_{v_1} \cup C_{v_2}$.

Lemma 2.12 *Let T be a component of the non-centered Voronoi subgraph of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$. Then ∂C_T contains each $\mathbf{s} \in \mathcal{S} \cap C_T$, and every geometric dual edge $\gamma \subset C_T$ whose dual Voronoi edge is centered.*

Proof For $s \in \mathcal{S} \cap C_T$ we claim that each component I of $T \cap V_s$ has a minimal-radius vertex. For a fixed vertex v of I, it follows from Lemma 2.3 that the initial vertex w of an edge pointing toward v in I has $J_w < J_v$, so w is contained in the ball about s of radius J_v since $d(s, w) = J_w$. By local finiteness of the Voronoi tessellation there are only finitely many such vertices. The claim follows.

For a minimal-radius vertex v of such a component I, Lemma 2.3 and Corollary 2.6 imply that a centered edge of V_s contains v. Its geometric dual lies in ∂C_T and contains s, so $s \in \partial C_T$. Any geometric dual edge γ contained in C_T is by definition an edge of C_v for some $v \in T^{(0)}$, so the geometric dual e to γ has v as a vertex. If e is centered then it does not lie in T, so $\gamma \subset \partial C_T$ by definition. □

Lemma 2.13 *For a component T of the non-centered Voronoi subgraph of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$, the interior of its geometric dual C_T is connected, open in $\mathbb{H}^2$ and dense in C_T. If $T^{(0)}$ is finite then C_T is closed, and its topological frontier is contained in ∂C_T.*

Remark 2.14 In fact, the proof below will reveal that an edge γ of ∂C_T is entirely contained in the topological frontier of C_T unless its geometric dual has both endpoints in T.

Proof For any vertex v of T the geometric dual C_v is a convex polyhedron and therefore closed in $\mathbb{H}^2$, with dense interior that is the complement of the union of its edges. This also holds for $\Delta(e_0, v_\infty)$, if applicable. Since ∂C_T is defined in Definition 2.11 as a union of edges, the interior $C_T - \partial C_T$ of C_T is therefore dense in C_T.

It is also connected: For points x and y in the interior of C_T there is a path ρ in T joining v and w, where $x \in C_v$ and $y \in C_w$ respectively. For any edge e of ρ, the geometric duals to the endpoints of e intersect in the geometric dual γ to e by Lemma 1.5. Each point in the interior of γ is in the interior of C_T, so one easily produces a path from x to y in the interior of C_T that is contained in the union of geometric duals to vertices of ρ.

For any $x \in \mathit{int}\, C_v \subset C_T$, x is in the interior of C_T and has an open neighborhood in $\mathbb{H}^2$ with this property. If x is in the interior of the geometric dual to an edge e of T then x has an open neighborhood in $\mathbb{H}^2$ that is contained in $\mathit{int}\, C_v \cup \mathit{int}\, C_w$ and hence the interior of C_T, where v and w are the endpoints of e. By Lemma 2.12, no point of $\mathcal{S}$ is in the interior of C_T. Therefore C_T is the union of points already described, hence open in $\mathbb{H}^2$.

If $T^{(0)}$ is finite then C_T is closed in $\mathbb{H}^2$, being a finite union of polygons. Any convergent sequence in C_T has an infinite subsequence in C_v for some fixed $v \in T^{(0)}$, so if it converges outside the interior of C_T the accumulation point lies in an edge of $C_v \cap \partial C_T$. □

To establish finer properties of C_T we will re-decompose it in a couple of different ways. We first use the collection of triangles defined below.

Definition 2.15 For an edge e of the Voronoi tessellation of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$, and a vertex v of e, let $\Delta(e, v)$ be the triangle in $\mathbb{H}^2$ with a vertex at v, and the geometric dual γ to e as an edge; ie $\Delta(e, v)$ is the convex hull in $\mathbb{H}^2$ of v and γ.

If e is non-compact and $v_\infty \in S_\infty$ is an ideal endpoint, let $\Delta(e, v_\infty)$ be the horocyclic ideal triangle with vertices at v_∞ and the endpoints of the geometric dual to e. (Recall Lemma 1.9 and Definition 1.10; this case agrees with the definition in Lemma 1.11.)

The endpoints of the geometric dual to e are points $\boldsymbol{s}, \boldsymbol{t} \in \mathcal{S}$ such that $e = V_{\boldsymbol{s}} \cap V_{\boldsymbol{t}}$. Thus $\Delta(e, v)$ is isosceles: its edges joining v to $\boldsymbol{s}$ and $\boldsymbol{t}$ each have length J_v. If v and w are opposite endpoints of e, then $\Delta(e, v)$ and $\Delta(e, w)$ share the edge γ. Whether their intersection is larger than this depends on whether e is centered; see Figure 5. In particular, we have the following lemma.

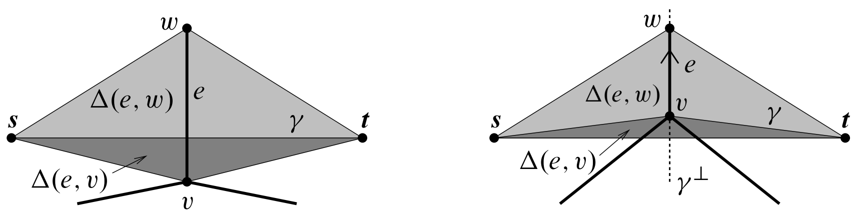

Figure 5: Triangles $\Delta(e, v)$ and $\Delta(e, w)$ when e is centered (on the left) and not centered

Lemma 2.16 *If e is a non-centered edge of the Voronoi tessellation of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$, with initial vertex v and terminal vertex w, then $\Delta(e, v) \subset \Delta(e, w)$, and $\Delta(e, v) \cap \partial\Delta(e, w)$ is the geometric dual γ to e. The same holds if e is non-compact and $w = v_\infty$ is its ideal endpoint.*

Proof Since e is non-centered it is contained on one side of the geodesic in $\mathbb{H}^2$ containing its geometric dual γ. Since v is the nearest point on e to γ, for any point w of $e - \{v\}$ the triangle T_w determined by w and the geometric dual γ to e has v

in its interior. Hence by convexity $\Delta(e,v)\subset T_w$, and $\Delta(e,v)\cap\partial T_w$ is their common edge γ.

If w above is the other endpoint of e then $T_w=\Delta(e,w)$ and the conclusion of the lemma holds. If e is non-compact with ideal endpoint w_∞ then $\Delta(e,w_\infty)\supset\bigcup_{w\in e}T_w$, and the conclusion again holds. $\square$

Lemma 2.17 *For a vertex v of the Voronoi tessellation of a locally finite set $\mathcal{S}\subset\mathbb{H}^2$:*

(1) *If C_v is centered then $C_v=\bigcup\{\Delta(e,v)\mid v\in e\}$.*

(2) *Otherwise, $C_v\cap\Delta(e_v,v)=\gamma_v$ and $C_v\cup\Delta(e_v,v)=\bigcup\{\Delta(e',v)\mid e'\neq e_v, v\in e'\}$, where e_v is the non-centered edge of V with initial vertex v and γ_v is its geometric dual.*

The decompositions $\bigcup\{\Delta(e,v)\mid v\in e\}$ of case (1) and $\bigcup\{\Delta(e',v)\mid e'\neq e_v, v\in e'\}$ of case (2) are non-overlapping.

Proof Upon cyclically enumerating the Voronoi edges containing v as $e_0,\dots,e_{n-1}$, each $\Delta(e_i,v)$ is identical to the triangle T_i defined in [7, Proposition 2.2]. By Lemma 2.5, C_v is non-centered if and only if v is the initial vertex of a non-centered edge e_v, and in this case if $e_v=e_{i_0}$ its geometric dual γ_{i_0} is the unique longest edge of C_v. This result is thus a direct application of [7, Proposition 2.2]. $\square$

Lemma 2.18 *Let C_T be a centered dual 2–cell, dual to a component T of the non-centered Voronoi subgraph determined by locally finite $\mathcal{S}\subset\mathbb{H}^2$. Assume $T^{(0)}$ is finite.*

(1) *If T has a noncompact edge e_0 with ideal endpoint v_∞, then*

$$C_T=\Delta(e_0,v_\infty)\cup\Bigg(\bigcup_{v\in T^{(0)},e\ni v}\Delta(e,v)\Bigg).$$

(2) *Otherwise, $C_T=\bigcup_{v\in T^{(0)},e\ni v}\Delta(e,v)$.*

Proof Lemma 2.17 and the definition of C_T together directly imply that C_T is contained in the union above. For the other inclusion we will we note the key fact that since $T^{(0)}$ is finite there is a root vertex v_T (recall Definition 2.8). We first suppose that T is compact and claim, for any $v\in T^{(0)}$ and edge e containing v, that $\Delta(e,v)\subset\bigcup_{w\in\gamma^{(0)}}C_w$, where γ is the unique reduced edge path joining v to v_T.

The proof is by induction on the number of edges in γ. The base case $v=v_T$ follows directly from Lemma 2.17 (since C_{v_T} is centered; see Proposition 2.9), so we assume

γ has $n \geq 1$ edges. By Lemma 2.17, $\Delta(e,v) \subset C_v \cup \Delta(e_v,v)$, where e_v is the non-centered edge with initial vertex v. Lemma 2.7 implies that the edge of γ containing v points toward v_T, so v is its initial vertex; hence by Corollary 2.6 this edge is e_v. The claim follows upon applying the inductive hypothesis to the terminal vertex w of e_v, since $\Delta(e_v,v) \subset \Delta(e_v,w)$ and w is connected to v_T by the reduced edge path consisting of all edges of γ but e_v.

In the case that T has a non-compact edge e_0, we change the claim to assert that $\Delta(e,v) \subset \Delta(e_0,v_\infty) \cup \left(\bigcup_{w \in \gamma^{(0)}} C_w\right)$. The proof is unchanged, except that in the base case Lemma 2.17 gives $\Delta(e,v_T) \subset C_{v_T} \cup \Delta(e_0,v_T)$, and we appeal to Lemma 2.16 to show that this is contained in $C_{v_T} \cup \Delta(e_0,v_\infty)$. □

That the $\Delta(e,v)$ overlap is a problem that we deal with by decomposing again.

Definition 2.19 For an edge e, with (possibly infinite) endpoints v and w, of the Voronoi tessellation of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$, define

- $Q(e) = \Delta(e,v) \cup \Delta(e,w)$ if e is centered;
- $Q(e) = \overline{\Delta(e,w) - \Delta(e,v)}$ if e is non-centered and v is its initial vertex.

Here $\Delta(e,v)$ and $\Delta(e,w)$ are as in Definition 2.15. See Figure 6.

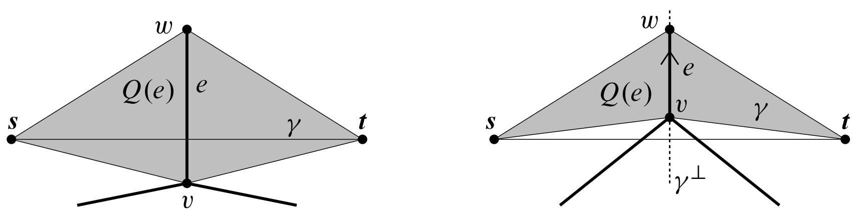

Figure 6: Quadrilaterals $Q(e)$ (shaded) when e is centered and not centered.

Lemma 2.20 *For distinct edges e and f of the Voronoi tessellation of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$, with geometric duals γ_e and γ_f,*

$$Q(e) \cap Q(f) = \begin{cases} [\boldsymbol{x}, v] & \text{if } e \cap f = v \text{ and } \gamma_e \cap \gamma_f = \boldsymbol{x}, \\ v = e \cap f & \text{if } e \cap f = v \text{ and } \gamma_e \cap \gamma_f = \varnothing, \\ \boldsymbol{x} & \text{if } e \cap f = \varnothing \text{ and } \gamma_e \cap \gamma_f = \boldsymbol{x}, \\ \varnothing & \text{otherwise.} \end{cases}$$

Above, $[\boldsymbol{x}, v]$ is the geodesic arc joining $\boldsymbol{x}$ to v. In particular, $Q(e)$ does not overlap $Q(f)$.

Proof For $\boldsymbol{x}$ and $\boldsymbol{y} \in \mathcal{S}$ such that $e = V_{\boldsymbol{x}} \cap V_{\boldsymbol{y}}$, inspection of Figure 5 reveals that $Q(e)$ is the union of the arcs joining $\boldsymbol{x}$ to points of e, together with those joining $\boldsymbol{y}$ to points of e. In particular, $Q(e) \subset V_{\boldsymbol{x}} \cup V_{\boldsymbol{y}}$, and it intersects the boundary of this union only at the endpoints of e. For $\boldsymbol{x}'$ and $\boldsymbol{y}'$ such that $f = V_{\boldsymbol{x}'} \cap V_{\boldsymbol{y}'}$, it is clear that $\{\boldsymbol{x}, \boldsymbol{y}\} \neq \{\boldsymbol{x}', \boldsymbol{y}'\}$, and if these sets are disjoint then $Q(e)$ can intersect $Q(f)$ only at a shared endpoint of e and f. Therefore suppose $\boldsymbol{x}' = \boldsymbol{x}$ (and hence $\boldsymbol{y}' \neq \boldsymbol{y}$). It is now easy to see from the description above that $Q(e)$ intersects $Q(f)$ only at $\boldsymbol{x}$, if e and f are not adjacent edges of $V_{\boldsymbol{x}}$, or along the arc joining $\boldsymbol{x}$ to v if $e \cap f$ is a vertex v. □

Definition 2.21 For distinct vertices v and w of a component T of the non-centered Voronoi subgraph determined by locally finite $\mathcal{S} \subset \mathbb{H}^2$, say $w < v$ if J_v is maximal among radii of vertices of the unique edge arc of T joining v to w. (Recall Lemma 2.7.)

Lemma 2.22 *Let T be a component of the non-centered Voronoi subgraph determined by locally finite $\mathcal{S} \subset \mathbb{H}^2$, with edge set $\mathcal{E}$. For $v \in T^{(0)}$, if C_v is non-centered, then*

$$\begin{aligned}(2.22.1)\quad & C_v \cup \Delta(e_v, v) \\ & \subset \bigcup_{w<v} \Big(Q(e_w) \cup \bigcup \{\Delta(e,w) \mid w \in e, e \notin \mathcal{E}\}\Big) \cup \bigcup \{\Delta(e,v) \mid v \in e, e \notin \mathcal{E}\}.\end{aligned}$$

Here e_v is the edge of T with initial vertex v. The analog holds if C_v is centered, replacing $C_v \cup \Delta(e,v)$ with C_v on the left side above.

Proof Below for $k \in \mathbb{N}$, let $v - k$ refer to the set of $w < v$ joined to v by an edge arc of T with length k. In particular, $v - 1$ is the set of initial endpoints of edges $e \in \mathcal{E} - \{e_v\}$ such that $v \in e$. Applying this to the decomposition $C_v \cup \Delta(e_v, v) = \bigcup \{\Delta(e,v) \mid v \in e, e \neq e_v\}$ from Lemma 2.17, then noting for $w \in v - 1$ that $\Delta(e_w, v) = Q(e_w) \cup \Delta(e_w, w)$ by Definition 2.19, yields

$$\begin{aligned}(2.22.2)\quad C_v \cup \Delta(e_v, v) &= \Big(\bigcup_{w \in v-1} \Delta(e_w, v)\Big) \cup \bigcup \{\Delta(e,v) \mid v \in e, e \notin \mathcal{E}\} \\ &= \bigcup_{w \in v-1} \big(Q(e_w) \cup \Delta(e_w, w)\big) \cup \bigcup \{\Delta(e,v) \mid v \in e, e \notin \mathcal{E}\}\end{aligned}$$

We can apply the same strategy to $C_w \cup \Delta(e_w)$ for each $w \in v - 1$, so an analog of (2.22.2) holds for $\Delta(e_w, w)$ with equality replaced by containment, v replaced everywhere by w, and $w \in v - 1$ by $u \in w - 1$. Iterating and applying an inductive argument gives, for any $k \in \mathbb{N}$, that

$$\begin{aligned}(2.22.3)\quad C_v \cup \Delta(e_v, v) \subset \bigcup_{v-k \leq w < v} & \Big(Q(e_w) \cup \bigcup \{\Delta(e,w) \mid w \in e, e \notin \mathcal{E}\}\Big) \\ & \cup \bigcup \{\Delta(e,v) \mid v \in e, e \notin \mathcal{E}\} \cup \bigcup_{w \in v-k} \Delta(e_w, w).\end{aligned}$$

Here we say $v-k \le w < v$ if $w \in v-j$ for some $j \le k$.

We claim that $C_v \cup \Delta(e_v, v)$ intersects $Q(e_w) \cup \Delta(e_w, w) \cup C_w$ for only finitely many $w < v$. $C_v \cup \Delta(e_v, v)$ is contained in the ball $B(v, J_v)$ of radius J_v around v, and if w' is the terminal endpoint of e_w then $Q(e_w)$ is in $B(w', J_{w'})$ while $\Delta(e_w, w) \cup C_w \subset B(w, J_w)$. For any $w < v$, since $J_w < J'_w \le J_v$, it follows that if $Q(e_w) \cup \Delta(e_w, w) \cup C_w$ intersects $C_v \cup \Delta(e_v, v)$ then w is in $B(v, 2J_v)$. The claim thus follows from local finiteness of Voronoi vertices.

The claim implies that there is some k_0 such that for any $k \ge k_0$ and $w \in v-k$,

$$(C_v \cup \Delta(e_v, v)) \cap (Q(e_w) \cup \Delta(e_w, w) \cup C_w) = \varnothing.$$

Taking k to be this k_0 in (2.22.3), we note that the intersection of $C_v \cup \Delta(e_v, v)$ with the union on the second line is empty, so the inclusion there holds with this union omitted. This immediately implies (2.22.1). The case that $v = v_T$ and C_v is centered is analogous. $\square$

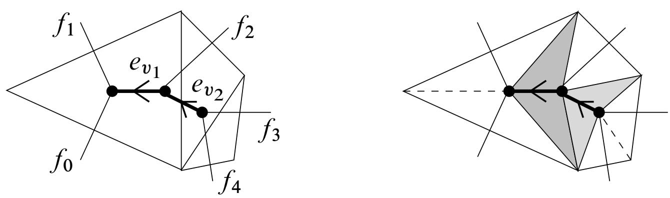

Figure 7: For T from Figure 4, with $Q(e_{v_1})$ and $Q(e_{v_2})$ shaded at right: $C_T = Q(e_{v_1}) \cup Q(e_{v_2}) \cup \bigcup_{i=0}^{4} \Delta(f_i, v_{j_i})$, where $j_i = T$ for $i = 0, 1$, $j_2 = 1$ and $j_3 = j_4 = 2$.

Proposition 2.23 *Let C_T be a centered dual 2–cell, dual to a component T of the non-centered Voronoi subgraph determined by locally finite $\mathcal{S} \subset \mathbb{H}^2$ with $T^{(0)}$ finite. Then*

$$C_T = \left(\bigcup_{e \in \mathcal{E}} Q(e) \right) \cup \bigcup_{v \in T^{(0)}} \{ \Delta(e, v) \mid v \in e, e \notin \mathcal{E} \},$$

where $\mathcal{E}$ is the edge set of T. This union is non-overlapping.

A simple case of this decomposition is illustrated in Figure 7.

Proof Lemma 2.22 implies that the right-hand side contains the left (compare with Definition 2.11). For an edge e of T with terminal vertex v, $Q(e) \subset \Delta(e, v)$, so Lemma 2.18 implies the other containment. Lemma 2.20 implies that the union is non-overlapping, upon recalling that if e is centered then $\Delta(e, v) \subset Q(e)$ for either vertex v of e. $\square$

Corollary 2.24 *Let C_T be a centered dual 2–cell, dual to a component T of the non-centered Voronoi subgraph determined by locally finite $\mathcal{S} \subset \mathbb{H}^2$ with $T^{(0)}$ finite. Then $T \subset int\, C_T$, and there is a homeomorphism $U \to int\, C_T$, where U is the interior of the unit disk $\mathbb{D}^2$ that extends to an onto map $\mathbb{D}^2 \to C_T$; or $\mathbb{D}^2 \to C_T \cup \{v_\infty\}$ if T has a non-compact edge with ideal vertex v_∞.*

Proof For a compact Voronoi edge $e = V_{\boldsymbol{x}} \cap V_{\boldsymbol{y}}$, $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{S}$, every point of $Q(e) - \{\boldsymbol{x}, \boldsymbol{y}\}$ is on a unique geodesic arc joining one of $\boldsymbol{x}$ or $\boldsymbol{y}$ to a point of e. There is a deformation retract taking this entire arc to its intersection with e; parametrizing each such by arclength determines a continuous deformation retract of $Q - \{\boldsymbol{x}, \boldsymbol{y}\}$ to e. If e has an endpoint v_∞ on S_∞ then we must also exclude the arcs joining $\boldsymbol{x}$ and $\boldsymbol{y}$ to v_∞.

For a finite vertex v of a Voronoi edge e, the deformation retract defined above takes (the complement of $\{\boldsymbol{x}, \boldsymbol{y}\}$ in) each edge of $Q(e)$ that contains v to v. Hence if $Q(e)$ intersects $Q(f)$ these homotopies agree on their overlap (recall Lemma 2.20).

Note also that if e is centered then its geometric dual γ is a flow line of the deformation retract $Q(e) - \{\boldsymbol{x}, \boldsymbol{y}\} \to e$. This thus restricts to a deformation retract $Q(e) - \gamma \to e - (e \cap \gamma)$. (Here again if e is non-compact we also exclude the arcs joining $\boldsymbol{x}$ and $\boldsymbol{y}$ to any infinite vertices.)

For a component T of the non-centered Voronoi subgraph with edge set $\mathcal{E}$, these two observations and Proposition 2.23 imply that the deformation retracts described above combine to determine a well-defined homotopy on a set that includes the interior of C_T. Recall that all points of $\mathcal{S} \cap C_T$ are in ∂C_T by Lemma 2.12, and all geometric duals γ to centered edges intersecting T are contained in ∂C_T by definition, as are edges with infinite endpoints. The image of this deformation retract is

$$T \cup \bigcup \{[v, e \cap \gamma) \mid v \in e \cap T^{(0)}, e \notin \mathcal{E}, \gamma \text{ geometrically dual to } e\}.$$

Here $[v, e \cap \gamma)$ refers to the sub arc of e joining v to $e \cap \gamma$, but not including the latter point. Each such arc deformation retracts to v, so the set above deformation retracts to T.

It follows that C_T is simply connected, since T is a tree. The Riemann mapping theorem thus asserts the existence of a homeomorphism $f\colon int\, \mathbb{D}^2 \to int\, C_T$. (Recall from Lemma 2.13 that $int\, C_T$ is a connected, open subset of $\mathbb{H}^2$, which we may take in $\mathbb{C}$ using the Poincaré disk model.)

If $T^{(0)}$ is finite then by Lemma 2.13 C_T is closed in $\mathbb{H}^2$ and the closure of its interior. Either C_T is compact, therefore also closed in $\mathbb{C}$, or it is compactified by the addition of the ideal point v_∞ of the non-compact edge of T (recall Proposition 2.9). It is not hard to show that *int* C_T is finitely connected along its boundary in the sense of [11, Section IX.4.4], so the results there on conformal mapping imply that f extends to a map from $\mathbb{D}^2$ to C_T or $C_T \cup \{v_\infty\}$. □

Corollary 2.25 *Let C_T be a centered dual 2–cell, dual to a component T of the non-centered Voronoi subgraph determined by locally finite $\mathcal{S} \subset \mathbb{H}^2$. Then:*

- *If $T' \neq T$ is a component of the non-centered Voronoi subgraph then $C_T \cap C_{T'} \subset \partial C_T$.*
- *For a Voronoi vertex v outside the non-centered Voronoi subgraph, $C_T \cap C_v = \partial C_T \cap \partial C_v$.*

Proof Let T be a component of the non-centered Voronoi subgraph and v a Voronoi vertex outside T. Then for any $w \in T^{(0)}$, $C_v \cap C_w$ is either a vertex of each or the geometric dual to an edge e with endpoints v and w. In the former case $C_v \cap C_w \subset \partial C_T$ by Lemma 2.12. In the latter e is centered since $v \notin T$, so again $C_v \cap C_w \subset \partial C_T$ (recall Definition 2.11).

The paragraph above implies the lemma's second assertion if T has no non-compact edges. If there is a non-compact edge e_0 we appeal to Lemma 1.11. The first assertion follows as well, upon noting that for distinct non-compact edges e_0 and f_0 with respective ideal endpoints v_∞ and w_∞, $\Delta(e_0, v_\infty)$ and $\Delta(f_0, w_\infty)$ intersect along their boundaries. □

Definition 2.26 The *centered dual complex* of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$ has vertex set

$$\mathcal{S} \cup \{v_\infty \mid v_\infty \text{ is the ideal endpoint of a non-centered Voronoi edge}\},$$

one-skeleton consisting of all

- geometric dual edges whose dual Voronoi edges are centered
- rays $[\boldsymbol{x}, v_\infty] \doteq [\boldsymbol{x}, v_\infty) \cup \{v_\infty\}$, where v_∞ is the ideal endpoint of a non-centered Voronoi edge, $\boldsymbol{x}$ is an endpoint of its geometric dual and $[\boldsymbol{x}, v_\infty)$ is the geodesic ray from $\boldsymbol{x}$ with ideal endpoint v_∞

and two-skeleton consisting of all

- geometric dual two-cells C_v, where v is a Voronoi vertex outside the non-centered Voronoi subgraph (in particular such a C_v is centered; see Lemma 2.5)
- cells C_T or $C_T \cup \{v_\infty\}$ of Definition 2.11, where T is a component of the non-centered Voronoi subgraph and v_∞ is an ideal endpoint of its noncompact edge (if applicable).

The results of this section imply that if each component T of the non-centered Voronoi subgraph has finite vertex set then the centered dual is indeed a cell decomposition of a subspace of $\mathbb{H}^2 \cup S_\infty$, in the sense that each cell above is the image of a disk by a map that restricts on the interior to a homeomorphism. By construction, its underlying topological space contains every geometric dual two-cell.

3 Admissible spaces and area bounds: The compact case

The ultimate goal of this section is to prove Theorem 3.31, which bounds the area of a compact centered dual 2–cell below given a uniform lower bound on its edge lengths. There is no corresponding result for cyclic polygons (at least no good one) because of a non-monotonicity property of the area of those that are non-centered. See Section 3.1 below, where we will collect useful results from [7] on cyclic polygons.

The price we pay for passing from the geometric dual to the centered dual complex is that a two-cell is no longer determined by its collection of boundary edge lengths. In Section 3.2 we will define an admissible space that parametrizes all possibilities for a centered dual two-cell with a given combinatorics and edge length collection, and prove some of its basic properties. Finally in Section 3.3 we will prove the theorem, by bounding values of the area functional on admissible spaces.

3.1 The geometry of cyclic polygons

Up to isometry there is a unique cyclic polygon with a given set of edge lengths (see eg [12, Theorem C]). Given this it is natural to parametrize the set of cyclic n–gons by a subset of $(\mathbb{R}^+)^n$ representing their side length collections. The result below describes this space and some of its geometrically important subspaces.

Proposition 3.1 *For $n \geq 3$, a cyclic n–gon is* marked *by fixing a vertex. The collection*

$$\mathcal{AC}_n = \left\{ (d_0, \dots, d_{n-1}) \in (\mathbb{R}^+)^n \;\middle|\; \sinh(d_i/2) < \sum_{j \neq i} \sinh(d_j/2) \text{ for each } i \in \{0, \dots, n-1\} \right\}$$

parametrizes marked cyclic n–gons by their side length collections. Below let $\theta(d, J) = 2\sin^{-1}(\sinh(d/2)/\sinh J) \in (0, \pi]$ for $d > 0$ and $J \geq d/2$. The collection

$$\mathcal{C}_n = \left\{ (d_0, \dots, d_{n-1}) \in (\mathbb{R}^+)^n \;\middle|\; \sum_{i=0}^{n-1} \theta(d_i, d_{i_0}/2) > 2\pi, \text{ where } d_{i_0} = \max\{d_i\}_{i=0}^{n-1} \right\} \subset \mathcal{AC}_n$$

parametrizes marked, centered n–gons, where a cyclic n–gon P is centered *if the center v of its circumcircle lies in its interior. The collection*

$$\mathcal{BC}_n = \left\{ (d_0, \dots, d_{n-1}) \in (\mathbb{R}^+)^n \,\middle|\, \sum_{i=0}^{n-1} \theta(d_i, d_{i_0}/2) = 2\pi, \text{ where } d_{i_0} = \max\{d_i\}_{i=0}^{n-1} \right\} \subset \mathcal{AC}_n$$

parametrizes marked, semicyclic n–gons, where an n–gon P is semicyclic *if its circumcircle radius is $d_{i_0}/2$; or equivalently, if v is in its longest edge. The collection*

$$\mathcal{HC}_n = \left\{ (d_0, \dots, d_{n-1}) \in (\mathbb{R}^+)^n \,\middle|\, \sinh(d_i/2) = \sum_{j \neq i} \sinh(d_j/2) \text{ for some } i \in \{0, \dots, n-1\} \right\}$$

parametrizes marked horocyclic *n–gons, those with vertices on a horocycle (recall Definition 1.8).*

The description of $\mathcal{AC}_n$ is in Corollary 1.6 of [7], of $\mathcal{C}_n$ and $\mathcal{BC}_n$ in Proposition 1.7 there and of $\mathcal{HC}_n$ in Corollary 3.5. The geometric characterizations of centeredness and semicyclicity used above are from Proposition 2.2 there.

Remark 3.2 Note the following easy consequences of Proposition 3.1: $\mathcal{C}_n$ and $\mathcal{AC}_n$ are open in $(\mathbb{R}^+)^n$, and $\mathcal{C}_n \cup \mathcal{BC}_n$ and $\mathcal{AC}_n \cup \mathcal{HC}_n$ are closed there.

We now record some differential formulas that we proved in [7], treating the area and radius of cyclic polygons as functions on $\mathcal{AC}_n$ (with the smooth structure inherited from $\mathbb{R}^n$).

Proposition 3.3 [7, Proposition 1.13] *For $n \geq 3$, the function $J\colon \mathcal{AC}_n \to \mathbb{R}^+$ that records circumcircle radius is smooth and symmetric. For $\boldsymbol{d} = (d_0, \dots, d_{n-1}) \in \mathcal{AC}_n$,*

$$\begin{cases} 0 < \frac{\partial J}{\partial d_i}(\boldsymbol{d}) < \frac{1}{2} & \text{if } \boldsymbol{d} \in \mathcal{C}_n \text{ for any } i, \\ \frac{\partial J}{\partial d_{i_0}}(\boldsymbol{d}) > \frac{1}{2} & \text{if } \boldsymbol{d} \in \mathcal{AC}_n - (\mathcal{C}_n \cup \mathcal{BC}_n) \text{ and } d_{i_0} = \max\{d_i\}_{i=0}^{n-1}, \\ \frac{\partial J}{\partial d_j}(\boldsymbol{d}) < 0 & \text{if } \boldsymbol{d} \in \mathcal{AC}_n - (\mathcal{C}_n \cup \mathcal{BC}_n) \text{ and } d_j \neq \max\{d_i\}_{i=0}^{n-1}. \end{cases}$$

Furthermore, if $d_i > d_j$ then

$$\left| \frac{\partial J}{\partial d_i}(\boldsymbol{d}) \right| > \left| \frac{\partial J}{\partial d_j}(\boldsymbol{d}) \right|.$$

By continuity,

$$\frac{\partial J}{\partial d_{i_0}}(\boldsymbol{d}) = \frac{1}{2} \quad \text{and} \quad \frac{\partial J}{\partial d_j}(\boldsymbol{d}) = 0$$

if $\boldsymbol{d} \in \mathcal{BC}_n$, for i_0 and j as above.

By [7, Proposition 3.6], values of J approach infinity on any sequence in $\mathcal{AC}_n$ approaching $\mathcal{HC}_n$.

The next result, on area of cyclic n–gons, is something like a Schläfli formula but in terms of side lengths.

Proposition 3.4 [7, Proposition 2.3] *For $n \geq 3$, the function $D_0\colon \mathcal{AC}_n \to \mathbb{R}^+$ that records hyperbolic area is smooth and symmetric. For $\boldsymbol{d} = (d_0, \dots, d_{n-1}) \in \mathcal{AC}_n$,*

$$\frac{\partial D_0}{\partial d_i}(\boldsymbol{d}) = \begin{cases} -\sqrt{\dfrac{1}{\cosh^2(d_i/2)} - \dfrac{1}{\cosh^2 J(\boldsymbol{d})}} & \text{if } \boldsymbol{d} \in \mathcal{AC}_n - \mathcal{C}_n \text{ and } d_i = \max\{d_j\}_{j=0}^{n-1}, \\ \sqrt{\dfrac{1}{\cosh^2(d_i/2)} - \dfrac{1}{\cosh^2 J(\boldsymbol{d})}} & \text{otherwise.} \end{cases}$$

Corollary 3.5 [7, Corollary 2.4] *For $n \geq 3$ and $(d_0, \dots, d_{n-1})$, $(d_0', \dots, d_{n-1}') \in \mathcal{C}_n \cup \mathcal{BC}_n$, if after a permutation $d_i \leq d_i'$ for all i and $d_i < d_i'$ for some i, then $D_0(d_0, \dots, d_{n-1}) < D_0(d_0', \dots, d_{n-1}')$.*

Remark 3.6 Since the radius and area functions are symmetric, we will not worry much in practice about the particular cyclic order on edge or vertex sets of geometric dual polygons.

$\mathcal{BC}_n$ and $\mathcal{HC}_n$ are smoothly parametrized, disjoint, codimension-one submanifolds of $(\mathbb{R}^+)^n$. The result below combines Proposition 1.11 and Corollary 3.5 of [7].

Proposition 3.7 *For each $n \geq 3$, there are smooth, positive-valued functions b_0 and h_0 on $(\mathbb{R}^+)^{n-1}$ such that $\mathcal{BC}_n$ and $\mathcal{HC}_n$ are the respective orbits of graph(b_0) and graph(h_0) under the $\mathbb{Z}_n$–action on $(\mathbb{R}^+)^n$ by cyclic permutation of entries, where*

$$\operatorname{graph}(b_0) = \left\{(b_0(\boldsymbol{d}), \boldsymbol{d}) \mid \boldsymbol{d} \in (\mathbb{R}^+)^{n-1}\right\}, \quad \operatorname{graph}(h_0) = \left\{(h_0(\boldsymbol{d}), \boldsymbol{d}) \mid \boldsymbol{d} \in (\mathbb{R}^+)^{n-1}\right\}.$$

The functions b_0 and h_0 have the following additional properties:

(1) *For any $(d_1, \dots, d_{n-1}) \in (\mathbb{R}^+)^{n-1}$,*

$$\max\{d_i\}_{i=1}^{n-1} < b_0(d_1, \dots, d_{n-1}) < h_0(d_1, \dots, d_{n-1}).$$

(2) [7, Corollary 4.10] *If $\boldsymbol{d} = (d_0, \dots, d_{n-1}) \in (\mathbb{R}^+)^n$ has largest entry d_0, then*

$$\boldsymbol{d} \in \mathcal{C}_n \Leftrightarrow d_0 < b_0(d_1, \dots, d_{n-1}) \quad \text{and} \quad \boldsymbol{d} \in \mathcal{AC}_n \Leftrightarrow d_0 < h_0(d_1, \dots, d_{n-1}).$$

(3) *If* $0 < d_i \le d_i'$ *for each* $i \in \{1, \dots, n-1\}$ *then*

$$b_0(d_1, \dots, d_{n-1}) \le b_0(d_1', \dots, d_{n-1}'),$$

and the same holds for h_0.

It is not convenient to attempt explicit formulas for the functions D_0 and J, but it is useful to know explicit values in a few cases.

Lemma 3.8 *For* $n \ge 3$ *and* $d > 0$, $\boldsymbol{d} \doteq (d, \dots, d) \in (\mathbb{R}^+)^n$ *is in* $\mathcal{C}_n$, *and*

$$D_0(\boldsymbol{d}) = (n-2)\pi - 2n \sin^{-1}\left(\frac{\cos(\pi/n)}{\cosh(d/2)}\right), \quad \sinh J(\boldsymbol{d}) = \frac{\sinh(d/2)}{\sin(\pi/n)}.$$

For $n \ge 3$ *and* $(B_0, d, \dots, d) \in \mathcal{BC}^n$,

$$D_0(B_0, d, \dots, d) = (n-2)\pi - (2n-2)\sin^{-1}\left(\frac{\cos(\pi/(2n-2))}{\cosh(d/2)}\right).$$

Proof It follows directly from the definitions that $\boldsymbol{d} \in \mathcal{C}_n$. A cyclic n–gon with all sides of length d is divided into n isometric isosceles triangles by arcs joining its vertices to its center v. Each of the resulting triangles thus has angle $2\pi/n$ at v, with each edge containing v of length $J(\boldsymbol{d})$, and opposite edge of length d. Applying [7, Lemma 1.3] and rearranging gives $\sinh J = \sinh(d/2)/\sin(\pi/n)$.

Applying the hyperbolic law of sines now yields the following formula for the angle α between the sides of length $J(\boldsymbol{d})$ and d:

$$\sin\alpha = \frac{\sinh(d/2)}{\sinh d} \cdot \frac{\sin(2\pi/n)}{\sin(\pi/n)} = \frac{\cos(\pi/n)}{\cosh(d/2)}.$$

The latter equation again follows from half-angle formulas. $D_0(\boldsymbol{d})$ is n times the area of one of these triangles, the angle defect $\pi - 2\pi/n - 2\alpha$. This gives the first formula above.

The circumcircle center of a semicyclic n–gon P with side-length collection

$$(B_0, d, \dots, d)$$

is at the midpoint of its longest side, the union of P with its reflection $\overline{P}$ across the longest side is a cyclic $(2n-2)$–gon with the same circumcircle and all sides of length d. Thus

$$\mathrm{Area}(P \cup \overline{P}) = \mathrm{Area}(P) + \mathrm{Area}(\overline{P}) = 2\,\mathrm{Area}(P).$$

The second area formula therefore follows from the first. □

3.2 Admissible spaces

By the results of Section 3.1, a centered dual 2–cell C_T is determined by the edge lengths of its constituent geometric dual polygons, together with their combinatorial arrangement. The latter data are captured by the corresponding component T of the non-centered Voronoi subgraph. Recall from Definition 2.11 that the boundary of C_T is the union of geometric duals to edges in the frontier of T. Our goal here is to understand the geometry of C_T using only its combinatorial structure and edge length data.

It is not hard to see that this is insufficient to determine C_T, but in this section we will describe properties of an *admissible space* that, given this data collection, parametrizes all possibilities for such a cell. We focus on the case that C_T is compact, so T is as well; in particular, all its edges are compact and $T^{(0)}$ is finite.

Blanket hypothesis In this subsection we take V to be a graph, perhaps with some non-compact edges, such that each vertex v has valence n_v satisfying $3 \le n_v < \infty$. $T \subset V$ is a compact, rooted subtree with root vertex v_T, edge set $\mathcal{E}$, and frontier $\mathcal{F}$ in V. The sole exception to this rule is Lemma 3.14, where explicit hypotheses are given.

Here the *frontier* of T in V is the set of pairs (e, v) such that e is an edge of V but not of T, and v is a vertex in $e \cap T$. We may refer to an edge of the frontier of T, without reference to its vertices, but note that such an edge may contribute up to two elements to $\mathcal{F}$.

Definition 3.9 Partially order $T^{(0)}$ by setting $v < v_T$ for each $v \in T^{(0)} - \{v_T\}$, and $w < v$ if the edge arc in T joining $w \in T^{(0)} - \{v_T, v\}$ to v_T runs through v. Let $v - 1$ be the set of $w < v$ joined to it by an edge, and say v is *minimal* if $v - 1 = \varnothing$. For $v \in T^{(0)} - \{v_T\}$, let e_v be the initial edge of the arc in T joining v to v_T, and say $e \to v$ for each edge $e \neq e_v$ of V containing v.

Definition 3.10 Let $(\mathbb{R}^+)^{\mathcal{F}}$ be the set of tuples of positive real numbers indexed by the elements of $\mathcal{F}$, and define $(\mathbb{R}^+)^{\mathcal{E}}$ analogously. For any elements $\boldsymbol{d}_{\mathcal{E}} = (d_e \mid e \in \mathcal{E}) \in (\mathbb{R}^+)^{\mathcal{E}}$ and $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$, let $\boldsymbol{d} = (\boldsymbol{d}_{\mathcal{E}}, \boldsymbol{d}_{\mathcal{F}})$ and $P_v(\boldsymbol{d}) = (d_{e_0}, \dots, d_{e_{n-1}})$ for $v \in T^{(0)}$, where the edges of V containing v are cyclically ordered as $e_0, \dots, e_{n-1}$. We say the *admissible set* $Ad(\boldsymbol{d}_{\mathcal{F}})$ determined by $\boldsymbol{d}_{\mathcal{F}}$ is the collection of $\boldsymbol{d} \in (\mathbb{R}^+)^{\mathcal{E}} \times \{\boldsymbol{d}_{\mathcal{F}}\}$ such that:

(1) For $v \in T^{(0)} - \{v_T\}$, $P_v(\boldsymbol{d}) \in \mathcal{AC}_{n_v} - \mathcal{C}_{n_v}$ has largest entry d_{e_v}.

(2) $P_{v_T}(\boldsymbol{d}) \in \mathcal{C}_{n_T}$, where we refer by n_T to the valence n_{v_T} of v_T in V.

(3) $J(P_v(\boldsymbol{d})) > J(P_w(\boldsymbol{d}))$ for each $w \in v-1$, where $J(P_v(\boldsymbol{d}))$ and $J(P_w(\boldsymbol{d}))$ are the respective radii of $P_v(\boldsymbol{d})$ and $P_w(\boldsymbol{d})$.

Remark 3.11 $Ad(\boldsymbol{d}_{\mathcal{F}})$ is empty for certain $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$. For instance if T has one edge and vertices of valence 3 in V then for any $d > 0$ and $\boldsymbol{d}_{\mathcal{F}} = (d,d,d,d)$, $Ad(\boldsymbol{d}_{\mathcal{F}}) = \varnothing$.

Remark 3.12 If $T = \{v_T\}$ then $Ad(\boldsymbol{d}_{\mathcal{F}})$ is either empty or $\{\boldsymbol{d}_{\mathcal{F}}\}$ for any $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$; the latter if and only if $P_{v_T}(\boldsymbol{d}_{\mathcal{F}}) \in \mathcal{C}_{n_T}$. (Note that the valence n_T of v_T in V is $|\mathcal{F}|$.)

Definition 3.13 Fix $\boldsymbol{d}_{\mathcal{F}} = (d_e \mid e \in \mathcal{F}) \in (\mathbb{R}^+)^{\mathcal{F}}$ such that $Ad(\boldsymbol{d}_{\mathcal{F}}) \neq \varnothing$. For each $\boldsymbol{d} \in Ad(\boldsymbol{d}_{\mathcal{F}})$ and $R \geq 0$, define

$$D_T(\boldsymbol{d}) = \sum_{v \in T^{(0)}} D_0(P_v(\boldsymbol{d})),$$

where $P_v(\boldsymbol{d})$ is as in Definition 3.10 and $D_0(P)$ is as in Proposition 3.4.

Lemma 3.14 *Let C_T be a compact centered dual two-cell, dual to a component T of the non-centered Voronoi subgraph determined by locally finite $\mathcal{S} \subset \mathbb{H}^2$. Let $\mathcal{E}$ be the edge set of T and $\mathcal{F}$ its frontier in the Voronoi graph V, and for each edge e of V that intersects T let d_e be the length of the geometric dual to e. Then $\boldsymbol{d} = (d_e \mid e \in \mathcal{E}) \in Ad(\boldsymbol{d}_{\mathcal{F}})$, where $\boldsymbol{d}_{\mathcal{F}} = (d_e \mid (e,v) \in \mathcal{F}$ for some $v \in T^{(0)})$, and C_T has area $D_T(\boldsymbol{d})$.*

Proof Since C_T is compact so is T; in particular, $T^{(0)}$ is finite. It follows that T has a root vertex v_T (recall Definition 2.8). By Proposition 2.9(2), the geometric dual C_{v_T} to v_T is centered, and C_v is non-centered for each $v \in T^{(0)} - \{v_T\}$. It further follows from Lemma 2.7 that for each $v \in T^{(0)} - \{v_T\}$, e_v as defined in Definition 3.9 is the edge of T with initial vertex v.

If $e_0, \dots, e_{n-1}$ is the cyclically ordered collection of edges of V containing $v \in T^{(0)}$, then C_v is represented by $(d_{e_0}, \dots, d_{e_{n-1}}) \in \mathcal{AC}_n$ (recall Proposition 3.1). Criterion (2) from Definition 3.10 follows, as does (1) upon observing that for each $v \in T^{(0)} - \{v_T\}$, C_v has longest side length d_{e_v} by Lemma 2.5.

For $v \in T^{(0)}$ and $w \in v-1$, since w is the initial vertex of e_w and v is its terminal vertex Lemma 2.3 yields $J_v > J_w$. Definition 3.10(3) follows, upon noting that $J_v = J(P_v)$ and $J_w = J(P_w)$, where the left-hand quantities are described in Lemma 1.5 and the others in Proposition 3.3.

That C_T has area $D_T(\boldsymbol{d})$ is a direct consequence of Definitions 2.11 and 3.13, since the union $C_T = \bigcup_{v\in T^{(0)}} C_v$ is non-overlapping and $D_0(P_v(\boldsymbol{d}))$ is the area of C_v for each $v \in T^{(0)}$. □

It is not hard to see that $Ad(\boldsymbol{d}_{\mathcal{F}})$ is generally not closed in $(\mathbb{R}^+)^{\mathcal{E}} \times \{\boldsymbol{d}_{\mathcal{F}}\}$. We will find it convenient to enlarge it slightly, since our main goal here is to compute minima of D_T.

Definition 3.15 For $\boldsymbol{d}_{\mathcal{F}} = (d_e \mid e \in \mathcal{F}) \in (\mathbb{R}^+)^{\mathcal{F}}$ let $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ consist of those $\boldsymbol{d} = (\boldsymbol{d}_{\mathcal{E}}, \boldsymbol{d}_{\mathcal{F}})$ for $\boldsymbol{d}_{\mathcal{E}} \in (\mathbb{R}^+)^{\mathcal{E}}$ such that:

(1) For $v \in T^{(0)} - \{v_T\}$, $P_v(\boldsymbol{d}) \in \mathcal{AC}_{n_v} - \mathcal{C}_{n_v}$ has largest entry d_{e_v}.

(2) $P_{v_T}(\boldsymbol{d}) \in \mathcal{C}_{n_T} \cup \mathcal{BC}_{n_T}$, where we refer by n_T to the valence n_{v_T} of v_T in V.

(3) $J(P_v(\boldsymbol{d})) \geq J(P_w(\boldsymbol{d}))$ for each $w \in v-1$, where $J(P_v(\boldsymbol{d}))$ and $J(P_w(\boldsymbol{d}))$ are the respective radii of $P_v(\boldsymbol{d})$ and $P_w(\boldsymbol{d})$.

It is immediate from its definition that $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ contains $Ad(\boldsymbol{d}_{\mathcal{F}})$. We will show in Lemma 3.21 that it is compact and in particular closed, so it contains the closure of $Ad(\boldsymbol{d}_{\mathcal{F}})$. However:

Remark 3.16 If T has one edge and vertices of valence 3 in V then for any $d > 0$ and $\boldsymbol{d}_{\mathcal{F}} = (d,d,d,d)$, $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}}) = \{(B, \boldsymbol{d}_{\mathcal{F}})\}$ where $B = b_0(d,d)$.

With Remark 3.11 this shows that the inclusion $\overline{Ad(\boldsymbol{d}_{\mathcal{F}})} \subset \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ is proper in some cases.

Remark 3.17 If $T = \{v_T\}$ then $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ is either empty or $\{\boldsymbol{d}_{\mathcal{F}}\}$ for any $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$; the latter if and only if $P_{v_T}(\boldsymbol{d}_{\mathcal{F}}) \in \mathcal{C}_{n_T} \cup \mathcal{BC}_{n_T}$. (Here $n_T = |\mathcal{F}|$ is the valence of v_T in V.)

Remark 3.18 Definition 3.15(1) implies that for any $v \in T^{(0)} - \{v_T\}$, $d_{e_v} > d_e$ for each $e \to v$ (cf Proposition 3.7(2)). It follows that $d_{e_v} > d_e$ for each $e \to w$ such that $w < v$. In particular, for some fixed $d > 0$ if $d_e \geq d$ for all $e \in \mathcal{F}$ then $d_e > d$ for all $e \in \mathcal{E}$.

The lemma below expands on Remark 3.18.

Lemma 3.19 *Collections $\{b_e\colon (\mathbb{R}^+)^{\mathcal{F}} \to \mathbb{R}^+\}_{e\in\mathcal{E}}$ and $\{h_e\colon (\mathbb{R}^+)^{\mathcal{F}} \to \mathbb{R}^+\}_{e\in\mathcal{E}}$ are determined by the following properties: For $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$ and $\boldsymbol{d}_{\mathcal{E}} \in (\mathbb{R}^+)^{\mathcal{E}}$ with $\boldsymbol{d} = (\boldsymbol{d}_{\mathcal{E}}, \boldsymbol{d}_{\mathcal{F}})$,*

- $d_e = b_e(\boldsymbol{d}_{\mathcal{F}})$ *for each* $e \in \mathcal{E}$ *if and only if for each* $v \in T^{(0)} - \{v_T\}$, $P_v(\boldsymbol{d})$ *is in* $\mathcal{BC}_{n_v}$ *and has largest entry* d_{e_v}.
- $d_e = h_e(\boldsymbol{d}_{\mathcal{F}})$ *for each* $e \in \mathcal{E}$ *if and only if for each* $v \in T^{(0)} - \{v_T\}$, $P_v(\boldsymbol{d})$ *is in* $\mathcal{HC}_{n_v}$ *and has largest entry* d_{e_v}.

For $e \in \mathcal{E}$, *the functions* b_e *and* h_e *have the following properties:*

(1) *For* $\boldsymbol{d}_{\mathcal{F}}$ *and* $v \in T^{(0)} - \{v_T\}$,
$$b_{e_v}(\boldsymbol{d}_{\mathcal{F}}) > \max\{b_e(\boldsymbol{d}_{\mathcal{F}}) \mid e \to v \in \mathcal{E}\} \cup \{d_e \mid e \to v \in \mathcal{F}\}.$$

(2) *If* $\boldsymbol{d} = (\boldsymbol{d}_{\mathcal{E}}, \boldsymbol{d}_{\mathcal{F}}) \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ *then for each* $e \in \mathcal{E}$, $b_e(\boldsymbol{d}_{\mathcal{F}}) \le d_e < h_e(\boldsymbol{d}_{\mathcal{F}})$.

(3) *If* $d'_e \ge d_e$ *for each* $e \in \mathcal{F}$ *then* $b_e(\boldsymbol{d}'_{\mathcal{F}}) \ge b_e(\boldsymbol{d}_{\mathcal{F}})$ *for each* $e \in \mathcal{E}$, *where* $\boldsymbol{d}'_{\mathcal{F}} = (d'_e)_{e \in \mathcal{F}}$.

Proof We construct by induction, the key point being that for $v \in T^{(0)} - \{v_T\}$, $b_{e_v}(\boldsymbol{d}_{\mathcal{F}})$ is determined by $\boldsymbol{d}_{\mathcal{F}}$ and $\{b_{e_w}(\boldsymbol{d}_{\mathcal{F}}) \mid w < v\}$, and similarly for $h_{e_v}(\boldsymbol{d}_{\mathcal{F}})$. Fix $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$.

Suppose first that $v \in T^{(0)}$ is minimal, so each $e \to v$ is in $\mathcal{F}$. Cyclically enumerate the edges of V containing v as $e_0, \dots, e_{n-1}$ so that $e_0 = e_v$, and for each $i > 0$ let $d_i = d_{e_i}$. Let $b_{e_v}(\boldsymbol{d}_{\mathcal{F}}) = b_0(d_1, \dots, d_{n-1})$ and $h_{e_v}(\boldsymbol{d}_{\mathcal{F}}) = h_0(d_1, \dots, d_{n-1})$, for b_0 and h_0 taking $(\mathbb{R}^+)^{n-1}$ to $\mathbb{R}^+$ as in Proposition 3.7. That result implies that $b_{e_v}(\boldsymbol{d}_{\mathcal{F}})$ is the unique real number with the property that $(b_{e_v}(\boldsymbol{d}_{\mathcal{F}}), d_1, \dots, d_{n-1})$ is in $\mathcal{BC}_n$ and has its largest entry first; and it implies the analog for $h_{e_v}(\boldsymbol{d}_{\mathcal{F}})$ and $\mathcal{HC}_n$.

Note also that if $\boldsymbol{d} \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$, then Definition 3.15(1) implies that
$$P_v(\boldsymbol{d}) = (d_{e_v}, d_1, \dots, d_{n-1}) \in \mathcal{AC}_n - \mathcal{C}_n$$
has largest entry d_{e_v}, so $b_{e_v}(\boldsymbol{d}_{\mathcal{F}}) \le d_{e_v} < h_{e_v}(\boldsymbol{d}_{\mathcal{F}})$ by Proposition 3.7(2). This implies property (2) above for b_{e_v}. Property (1) and property (3) above also follow from Proposition 3.7, respectively using assertions (1) and (3) there.

Now fix $v \in T^{(0)} - \{v_T\}$ non-minimal, and assume that $b_{e_w}(\boldsymbol{d}_{\mathcal{F}})$ and $h_{e_w}(\boldsymbol{d}_{\mathcal{F}})$ are defined, for each $w < v$, uniquely such that for $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$ and $\boldsymbol{d}_{\mathcal{E}} \in (\mathbb{R}^+)^{\mathcal{E}}$, with $\boldsymbol{d} = (\boldsymbol{d}_{\mathcal{E}}, \boldsymbol{d}_{\mathcal{F}})$:

- $P_w(\boldsymbol{d}) \in \mathcal{BC}_{n_w}$ with largest entry d_{e_w} for all $w < v$ if and only if $d_{e_w} = b_{e_w}(\boldsymbol{d}_{\mathcal{F}})$ for each $w < v$.
- $P_w(\boldsymbol{d}) \in \mathcal{HC}_{n_w}$ with largest entry d_{e_w} for all $w < v$ if and only if $d_{e_w} = h_{e_w}(\boldsymbol{d}_{\mathcal{F}})$ for each $w < v$.

- Property (2) holds for each b_{e_w} and h_{e_w}, and (1) and (3) hold for each b_{e_w}, $w < v$.

Cyclically order the edges containing v as $e_0, \dots, e_{n-1}$ so that $e_0 = e_v$, and for $i > 0$ take

$$d_i = \begin{cases} d_{e_i} & e_i \in \mathcal{F}, \\ b_{e_i}(\boldsymbol{d}_{\mathcal{F}}) & e_i \in \mathcal{E}, \end{cases} \qquad d'_i = \begin{cases} d_{e_i} & e_i \in \mathcal{F}, \\ h_{e_i}(\boldsymbol{d}_{\mathcal{F}}) & e_i \in \mathcal{E}. \end{cases}$$

Proposition 3.7 again implies that $b_{e_v}(\boldsymbol{d}_{\mathcal{F}}) \doteq b_0(d_1, \dots, d_{n-1})$ is unique among $b > \max\{d_i\}$ such that $(b_{e_v}(\boldsymbol{d}_{\mathcal{F}}), d_1, \dots, d_{n-1}) \in \mathcal{BC}_n$, and $h_{e_v}(\boldsymbol{d}_{\mathcal{F}}) \doteq h_0(d'_1, \dots, d'_{n-1})$ is unique among $h > \max\{d'_i\}$ such that $(h_{e_v}(\boldsymbol{d}_{\mathcal{F}}), d'_1, \dots, d'_{n-1}) \in \mathcal{HC}_n$.

Now let $\boldsymbol{d} \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$. Since property (2) holds by hypothesis for each $e_i \in \mathcal{E}$, $d_{e_i} \geq d_i$ for such i (and otherwise $d_{e_i} = d_i$ by construction). Thus Proposition 3.7(3) implies that $b_{e_v}(\boldsymbol{d}_{\mathcal{F}}) \leq b_0(d_{e_1}, \dots, d_{e_{n-1}})$, and Definition 3.15(1) and Proposition 3.7(2) imply that $b_0(d_{e_1}, \dots, d_{e_{n-1}}) < d_{e_v}$. Analogously, $h_{e_v}(\boldsymbol{d}_{\mathcal{F}}) > h_0(d_{e_1}, \dots, d_{e_{n-1}}) > d_{e_v}$. To summarize,

$$b_{e_v}(\boldsymbol{d}_{\mathcal{F}}) \leq b_0(d_{e_1}, \dots, d_{e_{n-1}}) \leq d_{e_v} < h_0(d_{e_1}, \dots, d_{e_{n-1}}) < h_{e_v}(\boldsymbol{d}_{\mathcal{F}}).$$

This proves property (2) for e_v. Properties (1) and (3) again follow from the corresponding assertions of Proposition 3.7, along with the inductive hypothesis.

The lemma now follows by induction. (Recall in particular that there is a unique e_v for each $v \in T^{(0)} - \{v_T\}$, and that $\mathcal{E}$ is the set of all such e_v.) □

Remark 3.20 For any given tree T with frontier $\mathcal{F}$, the proof of Lemma 3.19 is easily adapted (using formulas from [7]) to produce a recursive algorithm that takes $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$ and computes the values $b_e(\boldsymbol{d}_{\mathcal{F}})$ or $h_e(\boldsymbol{d}_{\mathcal{F}})$ from the outside in.

Lemma 3.21 *For any $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$, $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ is compact.*

Proof This is vacuous if $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ is empty, so fix $\boldsymbol{d}_{\mathcal{F}}$ such that $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}}) \neq \varnothing$. It is enough to show that $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ is closed in $\mathbb{R}^{\mathcal{E}} \times \{\boldsymbol{d}_{\mathcal{F}}\}$, since Lemma 3.19(2) implies it is bounded. Note also that Lemma 3.19(1) implies for fixed $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$ that if $d = \min\{d_e \mid e \in \mathcal{F}\}$ then $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}}) \subset [d, \infty)^{\mathcal{E}} \times \{\boldsymbol{d}_{\mathcal{F}}\}$.

It is clear from the definition of $P_v(\boldsymbol{d}) \in \mathcal{AC}_{n_v}$ that it varies continuously with $\boldsymbol{d}$ (to this point, recall that $\mathcal{AC}_n$ takes the subspace topology from $\mathbb{R}^n$). Since $\mathcal{C}_{n_T} \cup \mathcal{BC}_{n_T}$ is closed in $(\mathbb{R}^+)^{n_T}$ (see [7, Proposition 1.11]), and no sequence in $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ approaches the frontier of $(\mathbb{R}^+)^{\mathcal{E}} \times \{\boldsymbol{d}_{\mathcal{F}}\}$ in $\mathbb{R}^{\mathcal{E}} \times \{\boldsymbol{d}_{\mathcal{F}}\}$ (see above), condition (2) is preserved under any limit of points in $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$. By Proposition 3.3, $J(P_v(\boldsymbol{d}))$ varies continuously with $\boldsymbol{d}$ on $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ for each $v \in T^{(0)}$, so (3) is also preserved by such a limit.

Since $\mathcal{AC}_n$ is open in $(\mathbb{R}^+)^n$ it is *a priori* possible that (1) is not preserved; ie that for some sequence $\{\boldsymbol{d}_i\} \subset \overline{Ad}(\boldsymbol{d}_\mathcal{F})$ limiting to $\boldsymbol{d} \in (\mathbb{R}^+)^\mathcal{E} \times \{\boldsymbol{d}_\mathcal{F}\}$ there exists $v \in T^{(0)} - \{v_T\}$ such that $P_v(\boldsymbol{d}) \in \mathcal{HC}_{n_v}$, where v has valence n_v in V. For such $\{\boldsymbol{d}_i\} \to \boldsymbol{d}$, let v be a closest vertex to v_T such that $P_v(\boldsymbol{d}) \in \mathcal{HC}_{n_v}$. In particular $P_w(\boldsymbol{d}) \in \mathcal{AC}_{n_w}$ for the endpoint w of e_v (note that $P_{v_T}(\boldsymbol{d}) \in \mathcal{AC}_{n_T}$ by preservation of (2)). Proposition 3.3 implies on the one hand that $J(P_w(\boldsymbol{d}_i)) \to J(P_w(\boldsymbol{d}))$, since $P_w(\boldsymbol{d}_i) \to P_w(\boldsymbol{d})$, and on the other that $J(P_v(\boldsymbol{d}_i)) \to \infty$, since $P_v(\boldsymbol{d}_i) \to P_v(\boldsymbol{d}) \in \mathcal{HC}_n$. But then for some $\boldsymbol{d}_i$ the inequality of Definition 3.15(3) fails, a contradiction. Therefore (1) is preserved under taking limits, and $\overline{Ad}(\boldsymbol{d}_\mathcal{F})$ is closed. □

Lemma 3.22 *Fix $\boldsymbol{d}_\mathcal{F} = (d_e \mid e \in \mathcal{F}) \in (\mathbb{R}^+)^\mathcal{F}$ such that $\overline{Ad}(\boldsymbol{d}_\mathcal{F}) \neq \varnothing$. Then $D_T(\boldsymbol{d})$ is continuous on $\overline{Ad}(\boldsymbol{d}_\mathcal{F})$ and attains a minimum there.*

Proof Since $P \mapsto D_0(P)$ is continuous on $\mathcal{AC}_n$ (by Proposition 3.4) and $P_v(\boldsymbol{d}) \in \mathcal{AC}_n$ for each $\boldsymbol{d} \in \overline{Ad}(\boldsymbol{d}_\mathcal{F})$, $D_T(\boldsymbol{d})$ is continuous on $\overline{Ad}(\boldsymbol{d}_\mathcal{F})$. Since this is compact by Lemma 3.21, $D_T(\boldsymbol{d})$ attains a minimum on it. □

Finally, we observe that D_T attains a minimum only at one of a short list of special locations.

Proposition 3.23 *For $\boldsymbol{d}_\mathcal{F} \in (\mathbb{R}^+)^\mathcal{F}$ with $\overline{Ad}(\boldsymbol{d}_\mathcal{F}) \neq \varnothing$, at a minimum point $\boldsymbol{d} = (\boldsymbol{d}_\mathcal{E}, \boldsymbol{d}_\mathcal{F})$ for $D_T(\boldsymbol{d})$ one of the following holds:*

(1) *$P_v(\boldsymbol{d}) \in \mathcal{BC}_{n_v}$ for each $v \in T^{(0)} - \{v_T\}$, where v has valence n_v in V.*

(2) *$P_{v_T}(\boldsymbol{d}) \in \mathcal{BC}_{n_T}$, where v_T has valence n_T in V.*

(3) *$J(P_v(\boldsymbol{d})) = J(P_w(\boldsymbol{d}))$ for some $v \in T^{(0)}$ and $w \in v - 1$.*

Proof Suppose that none of the criteria above hold at $\boldsymbol{d}$, and fix $v \in T^{(0)} - \{v_T\}$ such that $P_v(\boldsymbol{d}) \notin \mathcal{BC}_{n_v}$, where v has valence n_v in V. We will show that for the edge e_v of T with initial point v, reducing d_{e_v} while keeping the remaining entries of $\boldsymbol{d}$ constant produces new points of $\overline{Ad}(\boldsymbol{d}_\mathcal{F})$ at which D_T takes smaller values.

We first observe that $D_T(\boldsymbol{d})$ is reduced by reducing d_{e_v}. Changing only the length of e_v affects only $P_v(\boldsymbol{d})$ and $P_{v'}(\boldsymbol{d})$, where v' is its terminal vertex. $P_v(\boldsymbol{d}) \in \mathcal{AC}_{n_v} - \mathcal{C}_{n_v}$ has largest side length d_{e_v}, but either $P_{v'}(\boldsymbol{d})$ has $d_{e_{v'}}$ as its largest or, if $v' = v_T$,

$P_{v'}(\boldsymbol{d}) \in \mathcal{C}_{n_T} \cup \mathcal{BC}_{n_T}$. Thus Proposition 3.4 implies that

$$\text{(3.23.1)} \quad \frac{\partial}{\partial d_{e_v}} D_T = \frac{\partial}{\partial d_{e_v}} [D_0(P_{v'}(\boldsymbol{d})) + D_0(P_v(\boldsymbol{d}))]$$
$$= \sqrt{\frac{1}{\cosh^2(d_{e_v}/2)} - \frac{1}{\cosh^2 J(P_{v'}(\boldsymbol{d}))}}$$
$$- \sqrt{\frac{1}{\cosh^2(d_{e_v}/2)} - \frac{1}{\cosh^2 J(P_v(\boldsymbol{d}))}}.$$

Since condition (3) above does not hold by hypothesis, but condition (3) of Definition 3.15 does, $J(P_{v'}(\boldsymbol{d})) > J(P_v(\boldsymbol{d}))$. Therefore the quantity above is positive. Since this is also $\frac{\partial}{\partial d_{e_v}} D_T(\boldsymbol{d})$, reducing d_{e_v} reduces the value of D_T near $\boldsymbol{d}$.

Our hypothesis and Definition 3.15(1) imply that $P_v(\boldsymbol{d})$ is in the open subset $\mathcal{AC}_{n_v} - (\mathcal{C}_{n_v} \cup \mathcal{BC}_{n_v})$ of $\mathbb{R}^{n_v}$. Thus small deformations of d_{e_v} keep it there. It is possible that $v' = v_T$; if so then because (2) above does not hold but the corresponding criterion from Definition 3.15 does, $P_{v'}(\boldsymbol{d})$ is in the open set $\mathcal{C}_n$. It follows again in this case that small deformations of d_{e_v} keep it here.

If $v' \neq v_T$ then it is possible that $P_{v'}(\boldsymbol{d}) \in \mathcal{BC}_{n'}$, where v' has valence n' in V. However in this case, direct appeal to Proposition 3.1 shows that reducing d_{e_v} keeps $P_{v'}$ in $\mathcal{AC}_{n'} - \mathcal{C}_{n'}$. Recall in particular that d_{e_v} is not the largest side length of $P_{v'}(\boldsymbol{d})$ by Definition 3.15(1); one easily shows that $\theta(d, D/2)$ increases with d for any fixed $D > d$.

Criterion (3) from Definition 3.15 holds for any small deformation of $\boldsymbol{d}$. This is because $J(P_v(\boldsymbol{d})) > J(P_w(\boldsymbol{d}))$ for all $v \in T^{(0)}$ and $w \in v - 1$, as we pointed out above, and $J(P_v(\boldsymbol{d}))$ varies continuously with $\boldsymbol{d}$. Thus by Definition 3.15, any small deformation of $\boldsymbol{d}$ that reduces d_{e_v} and leaves every other entry constant lies in $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$. □

3.3 A lower bound on area

Here we will prove Theorem 3.31, by induction on the number of vertices of the component T of the non-centered Voronoi subgraph contained in a centered dual 2–cell C_T. For the purposes of this argument we will give each Voronoi vertex v that is not contained in the non-centered Voronoi subgraph honorary status as a component of it. Thus $T = \{v\}$ is a tree with no edges, and the case $C_T = C_v$ is the base case of the induction. Note that C_v is centered for such v, by Lemma 2.5.

Proposition 3.24 *For $d > 0$ and $(d, \dots, d) \in (\mathbb{R}^+)^n$, where $n \geq 4$, $D_0(d, \dots, d) \geq (n-2) D_0(B_0, d, d)$, where $B_0 = b_0(d, d)$ for $b_0\colon \mathbb{R}^2 \to \mathbb{R}$ as in Proposition 3.7.*

Proof For $d^n \doteq (d,\dots,d) \in (\mathbb{R}^+)^n$, Lemma 3.8 implies that

$$D_0(d,\dots,d) = (n-2)\pi - 2n\sin^{-1}\left(\frac{\cos(\pi/n)}{\cosh(d/2)}\right)$$

and also that

$$(n-2)D_0(B_0,d,d) = (n-2)\left[\pi - 4\sin^{-1}\left(\frac{1/\sqrt{2}}{\cosh(d/2)}\right)\right].$$

Fixing $d>0$, for $n \geq 4$ we define $f_d(n) = D_0(d^n) - (n-2)D_0(B_0,d,d)$, so

$$f_d(n) = 2\left[2(n-2)\sin^{-1}\left(\frac{1/\sqrt{2}}{\cosh(d/2)}\right) - n\sin^{-1}\left(\frac{\cos(\pi/n)}{\cosh(d/2)}\right)\right].$$

Note that $f_d(4) = 0$ for each d. This reflects the fact that a cyclic quadrilateral with all sides of length d is the union of two triangles in $\mathcal{BC}_3$, each with two sides of length d. Now allowing n to take arbitrary values in $[4,\infty)$, we record the first and second derivatives of f_d:

$$f_d'(n) = 2\left[2\sin^{-1}\left(\frac{1/\sqrt{2}}{\cosh(d/2)}\right) - \sin^{-1}\left(\frac{\cos(\pi/n)}{\cosh(d/2)}\right) - \frac{\pi}{n}\frac{\sin(\pi/n)}{\sqrt{\cosh^2(d/2)-\cos^2(\pi/n)}}\right].$$

$$f_d''(n) = 2\frac{\pi^2}{n^3}\frac{\cos(\pi/n)\sinh^2(d/2)}{(\cosh^2(d/2)-\cos^2(\pi/n))^{3/2}}.$$

From this we find in particular that for any fixed d, f is concave up. For fixed d we have

$$f_d'(4) = 2\left[\sin^{-1}\left(\frac{1/\sqrt{2}}{\cosh(d/2)}\right) - \frac{\pi/4}{\sqrt{2\cosh^2(d/2)-1}}\right].$$

We claim that the quantity above is positive for each $d>0$. To this end, we compute

$$\frac{\partial}{\partial d}(f_d'(4)) = \frac{\sinh(d/2)}{\sqrt{2\cosh^2(d/2)-1}}\left[\frac{(\pi/2)\cosh(d/2)}{2\cosh^2(d/2)-1} - \frac{1}{\cosh(d/2)}\right].$$

The quantity in brackets above is positive at $d=0$, and one easily finds the unique $d_0>0$ at which it vanishes. Thus $\frac{\partial}{\partial d}f_d'(4)$ is positive on $(0,d_0)$ and negative on (d_0,∞). It is not hard to see that $f_0'(4) = 0 = \lim_{d\to\infty} f_d'(4)$, so $f_d'(4)$ is positive on $(0,\infty)$.

For any fixed $d>0$, we showed above that $f_d'(4)>0$ and that $f_d'(n)$ increases in n on $(4,\infty)$, so in particular $f_d'(n)>0$ for all n. Therefore $f_d(n)>0$ for every $n\in(4,\infty)$. The result follows. □

We will address the case when T has more than one vertex using Proposition 3.23. Of the three conditions there, (1) and (2) may each be addressed directly in different ways. We will use the lemma below to reduce complexity in case (3) and thereby apply induction.

Lemma 3.25 *For $\boldsymbol{c}_0 = (c_0, \dots, c_{m-1}) \in \mathcal{AC}_m - \mathcal{C}_m$ and $\boldsymbol{d}_0 = (d_0, \dots, d_{n-1}) \in \mathcal{AC}_n$, suppose:*

- $J(\boldsymbol{c}_0) = J(\boldsymbol{d}_0)$.
- $c_0 = d_0$ *is maximal among the* c_i.
- *Either* $\boldsymbol{d}_0 \in \mathcal{C}_n \cup \mathcal{BC}_n$, *or* $\boldsymbol{d}_0 \in \mathcal{AC}_n - \mathcal{C}_n$ *and* d_0 *is not maximal among the* d_i.

Then $\boldsymbol{d} \doteq (c_1, \dots, c_{m-1}, d_1, \dots, d_{n-1})$ is in $\mathcal{AC}_{m+n-2}$, and in $\mathcal{C}_{m+n-2} \cup \mathcal{BC}_{m+n-2}$ if and only if $\boldsymbol{d}_0 \in \mathcal{C}_n \cup \mathcal{BC}_n$. Also, $D_0(\boldsymbol{c}_0) + D_0(\boldsymbol{d}_0) = D_0(\boldsymbol{d})$.

Proof Cyclic polygons P_0 and Q_0 with side length collections $\boldsymbol{c}_0$ and $\boldsymbol{d}_0$, respectively, can be moved by an isometry so that they share a circumcircle C and a side γ_0 with length $c_0 = d_0$. It can be further arranged that P_0 and Q_0 lie in opposite half-spaces bounded by the geodesic through γ_0, so that $P_0 \cap Q_0 = \gamma_0$, with Q_0 in the half-space containing the center v of C. In fact this must hold, by [7, Proposition 2.2], unless $\boldsymbol{c}_0 \in \mathcal{BC}_m$. In this case v is the midpoint of γ_0, so also $\boldsymbol{d}_0 \in \mathcal{BC}_n$, and if P_0 and Q_0 are on the same side of γ_0 then one of them can be rotated about v by an angle of π in order to correct this.

Upon arranging P_0 and Q_0 as above, since $P_0 \cap Q_0 = \gamma_0$ the area of $P_0 \cup Q_0$ is the sum of their areas. Moreover, if the vertex sets of P_0 and Q_0 are cyclically ordered $\{x_0, \dots, x_{m-1}\}$ and $\{y_0, \dots, y_{n-1}\}$, respectively, (recall Definition 1.3) so that $x_0 = y_{n-1}$ and $y_0 = x_{m-1}$, then $\{x_1, \dots, x_{m-1}, y_1, \dots, y_{n-1}\}$ is cyclically ordered on C (recall Remark 1.4). Therefore by [7, Lemma 2.1] its convex hull P is a cyclic polygon.

It is clear that P contains P_0 and Q_0, and that P is contained in the disk B bounded by C. But every point of $B - (P_0 \cup Q_0)$ is separated from the vertices of P by the geodesic through a side of P_0 or Q_0, so $P = P_0 \cup Q_0$. Its side length collection is thus $\boldsymbol{d}$ as described above, and $D_0(\boldsymbol{d}) = D_0(\boldsymbol{c}_0) + D_0(\boldsymbol{d}_0)$ since D_0 measures area. Finally, P contains v if and only if Q_0 does, so $\boldsymbol{d} \in \mathcal{C}_{m+n-2} \cup \mathcal{BC}_{m+n-2}$ if and only if $\boldsymbol{d}_0 \in \mathcal{C}_n \cup \mathcal{BC}_n$ by [7, Proposition 2.2] again. □

Blanket hypothesis Until the proof of Theorem 3.31, each definition and result below uses the following hypothesis: V is a finite graph with vertices of valence at least 3, T is a rooted subtree of V with root vertex v_T, $\mathcal{E}$ is the edge set of T and $\mathcal{F}$ is its frontier in V.

Definition 3.26 For an edge e of T, let $p_e\colon V \to V_e$ be the quotient map that identifies e to a point, and let $T_e = p_e(T)$.

Remark 3.27 It is easy to see that T_e is a tree, and that p_e maps $\mathcal{E} - \{e\}$ and $\mathcal{F}$ bijectively to the edge set $\mathcal{E}_e$ and frontier $\mathcal{F}_e$ of T_e, respectively. In particular, if the endpoints v and w have valences n_v and n_w in V, respectively, then $p_e(v) = p_e(w)$ has valence $n_v + n_w - 2$.

Lemma 3.28 *For $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$, suppose that some $\boldsymbol{d} = (\boldsymbol{d}_{\mathcal{E}}, \boldsymbol{d}_{\mathcal{F}}) \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ satisfies condition (3) of Proposition 3.23. Then $D_T(\boldsymbol{d}) = D_{T_f}(\boldsymbol{d}_f)$ for T_f as in Definition 3.26, where:*

- *$f \in \mathcal{E}$ has initial vertex v and terminal vertex w such that $J(P_v(\boldsymbol{d})) = J(P_w(\boldsymbol{d}))$.*
- *$\boldsymbol{d}_f = (\boldsymbol{d}_{\mathcal{E}_f}, \boldsymbol{d}_{\mathcal{F}_f}) \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}_f})$ for*

$$\boldsymbol{d}_{\mathcal{E}_f} = (d_{p_f(e)} \mid e \in \mathcal{E} - \{f\}) \quad \text{and} \quad \boldsymbol{d}_{\mathcal{F}_f} = (d_{p_f(e)} \mid e \in \mathcal{F}),$$

where $d_{p_f(e)} = d_e$ for each e in $\mathcal{E} - \{f\}$ or occurring in $\mathcal{F}$.

This follows directly from Lemma 3.25. The result below will allow us to address condition (2) of Proposition 3.23, by varying $\boldsymbol{d}_{\mathcal{F}}$ and tracking the changes in $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$.

Lemma 3.29 *The set SAd_T, consisting of $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$ such that $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}}) \neq \varnothing$, is closed in $(\mathbb{R}^+)^{\mathcal{F}}$, and the function*

$$\boldsymbol{d}_{\mathcal{F}} \mapsto \min\{D_T(\boldsymbol{d}) \mid \boldsymbol{d} \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})\}$$

is lower-semicontinuous on SAd_T.

Proof Suppose $\boldsymbol{d}_{\mathcal{F}}^{(i)}$ is a sequence in SAd_T converging to $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$, and for each i let

$$\boldsymbol{d}^{(i)} \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}}^{(i)})$$

be a point at which $D_T(\boldsymbol{d})$ attains its minimum over all $\boldsymbol{d} \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}}^{(i)})$. We claim that there exist $0 < d < D$ such that $\boldsymbol{d}^{(i)} \subset [d, D]^{\mathcal{E}} \times [d, D]^{\mathcal{F}}$ for each i.

Since $\{\boldsymbol{d}_{\mathcal{F}}^{(i)}\}$ converges, $D_0 = \sup\{d_e^{(i)} \mid e \in \mathcal{F}, i \in \mathbb{N}\}$ is finite, and since no side of a polygon has length greater than the sum of the lengths of the other sides, one finds easily that $d_e^{(i)} \le D \doteq q^p D_0$ for each $e \in \mathcal{E}$ and $i \in \mathbb{N}$, where $p = |\mathcal{E}|$ and $q+1$ is the maximal valence in V of a vertex of T. Furthermore, $d = \inf\{d_e^{(i)} \mid e \in \mathcal{F}, i \in \mathbb{N}\} > 0$, and Lemma 3.19(1) implies that $d_e^{(i)} \ge d$ for each $e \in \mathcal{E}$ and $i \in \mathbb{N}$. The claim follows.

The claim implies that a subsequence of $\{\boldsymbol{d}^{(i)}\}$ converges to $\boldsymbol{d} = (\boldsymbol{d}_{\mathcal{E}}, \boldsymbol{d}_{\mathcal{F}})$ for some $\boldsymbol{d}_{\mathcal{E}} \in (\mathbb{R}^+)^{\mathcal{E}}$ and $\boldsymbol{d}_{\mathcal{F}}$ as fixed at the beginning. We next claim that $\boldsymbol{d} \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$. The proof of this claim is essentially identical to the proof of Lemma 3.21, replacing $\boldsymbol{d}_i$ there with $\boldsymbol{d}^{(i)}$.

The second claim immediately implies that SAd_T is closed. Furthermore, for $\boldsymbol{d} = (\boldsymbol{d}_{\mathcal{E}}, \boldsymbol{d}_{\mathcal{F}})$, $D_T(\boldsymbol{d})$ is an upper bound for the value at $\boldsymbol{d}_{\mathcal{F}}$ of the minimum function in question here. Proposition 3.4 implies $D_T(\boldsymbol{d}) = \lim_{i\to\infty} D_T(\boldsymbol{d}_i)$, and lower semicontinuity follows. □

Proposition 3.30 *For fixed $d > 0$ and $\boldsymbol{d}_{\mathcal{F}} \in SAd_T$ with all entries at least d,*

$$\min\{D_T(\boldsymbol{d}) \mid \boldsymbol{d} \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})\} \geq (|\mathcal{F}| - 2)\, D_0(B_0, d, d),$$

where $B_0 = b_0(d, d)$ for b_0 as described in Proposition 3.7.

Proof We prove this by induction on the number of vertices of T. The base case $T = \{v_T\}$ follows directly from Proposition 3.24 and Corollary 3.5 (cf Remark 3.17), so let us suppose that T has $n > 1$ vertices and the result holds for all trees with fewer than n vertices. Since T is a tree, $|\mathcal{E}| = n - 1$ (by Euler characteristic), so in particular T has at least one edge.

Fix an arbitrary $D > d$ and consider the intersection of SAd_T (as in Lemma 3.29) with $[d, D]^{\mathcal{F}}$. Lemma 3.29 implies that this set is compact, so the function

$$\boldsymbol{d}_{\mathcal{F}} \mapsto \min\{D_T(\boldsymbol{d}) \mid \boldsymbol{d} \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})\}$$

attains a minimum on it by lower-semicontinuity. Fix $\boldsymbol{d}_{\mathcal{F}}$ at which the minimum occurs, and let $\boldsymbol{d}$ be a minimum point for D_T on $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$. This satisfies at least one of the three conditions described in Proposition 3.23. We claim that because of our choice of $\boldsymbol{d}_{\mathcal{F}}$, $\boldsymbol{d}$ in fact satisfies at least one of conditions (1) or (3).

Assume by way of contradiction that $\boldsymbol{d}$ satisfies only (2), and cyclically enumerate the edges of V containing v_T as $e_0, \dots, e_{n_T-1}$ so that d_{e_0} is maximal. By the hypothesis and Remark 3.18, $d_{e_i} \geq d$ for each $i > 0$. Since $P_{v_T}(\boldsymbol{d}) \in \mathcal{BC}_{n_T}$ we have

$$d_{e_0} = b_0(d_{e_1}, \dots, d_{e_{n_T-1}})$$

for b_0 as in Proposition 3.7. That result implies in particular that $d_{e_0} > \max\{d_{e_i}\}_{i=1}^{n_T} \geq d$, so d_{e_0} can be reduced slightly while preserving the inequality $d_{e_0} > d$.

We note that it is not the case that $e_0 \in \mathcal{E}$: since $P_{v_T}(\boldsymbol{d}) \in \mathcal{BC}_{n_T}$ by (2), $J(P_{v_T}(\boldsymbol{d})) = d_{e_0}/2$ (recall Proposition 3.1). But for any $v \in T^{(0)}$, $J(P_v(\boldsymbol{d})) \geq \max\{d_e/2 \mid e \ni v\}$ (see [7, Proposition 1.5]); in particular, if e_0 were in $\mathcal{E}$ it would follow that $J(P_{v_0}(\boldsymbol{d})) \geq$

$d_{e_0}/2$ for the other endpoint $v_0 \in T^{(0)}$ of e_0. But Definition 3.15(3) implies that $J(P_{v_0}(\boldsymbol{d})) \leq J(P_{v_T}(\boldsymbol{d}))$, so in fact equality would hold, violating our assumption that Proposition 3.23(3) does not.

Thus e_0 is in $\mathcal{F}$; or, more precisely, $(e_0, v_T) \in \mathcal{F}$. Moreover, the other endpoint v_0 of e_0 is not in T: if it were, Remark 3.18 would imply that $d_{e_{v_0}} > d_{e_0}$, and applying Definition 3.15(1) inductively along the path joining v_0 to v_T would yield $i > 0$ such that $d_{e_i} > d_{e_0}$, a contradiction. Therefore changing d_{e_0} while fixing the other entries of $\boldsymbol{d}$ changes only $P_{v_T}(\boldsymbol{d})$.

Reducing d_{e_0} while fixing all other entries of $\boldsymbol{d}$ takes $P_{v_T}(\boldsymbol{d})$ into $\mathcal{C}_{n_T}$, by Proposition 3.7, while reducing $D_0(P_{v_T}(\boldsymbol{d}))$, by Proposition 3.4. (Note that a such a deformation $\boldsymbol{d}(t)$ would have $\frac{d}{dt}D_0(P_{v_T}(\boldsymbol{d}(t)) = 0$ at $t = 0$ but negative thereafter.) Since all other $P_v(\boldsymbol{d})$ are unaffected by such a deformation, and since $J(P_{v_T}(\boldsymbol{d}))$ varies continuously with $\boldsymbol{d}$, by Definition 3.15 such a family $\boldsymbol{d}(t)$ would produce new $\boldsymbol{d}_{\mathcal{F}}(t) \in SAd_T \cap [d, D]^{\mathcal{F}}$ with $\boldsymbol{d}(t) \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}}(t))$ and $D_T(\boldsymbol{d}(t)) < D_T(\boldsymbol{d})$. This contradicts our minimality hypothesis, and it follows that one of (1) or (3) must hold.

Suppose first that $\boldsymbol{d}$ satisfies (1), so $P_v(\boldsymbol{d}) \in \mathcal{BC}_{n_v}$ for each $v \in T^{(0)} - \{v_T\}$, where n_v is the valence of v in T. Fix such v and cyclically enumerate the edges containing v as $e_0, \dots, e_{n_v-1}$ so that $e_0 = e_v$. Remark 3.18 implies that $d_{e_i} \geq d$ for all $i > 0$, so $d_{e_0} = b_0(d_{e_1}, \dots, d_{e_{n_v-1}})$ is at least $b_0(d, \dots, d)$ (recall Proposition 3.7(3)). If $n_v = 3$ then we conclude from Corollary 3.5 that $D_0(P_v(\boldsymbol{d})) > D_0(B_0, d, d)$. If $n_v \geq 4$ then we only need the bound $d_{e_v} > d$ (and Corollary 3.5) to conclude that $D_0(P_v(\boldsymbol{d})) \geq (n_v - 2) D_0(B_0, d, d)$ using Proposition 3.24.

For $v = v_T$ we argue as above to show that $D_0(P_{v_T}(\boldsymbol{d})) \geq (n_T - 2) D_0(B_0, d, d)$, where n_T is the valence of v_T in V. If $n_T \geq 4$ then this follows from Proposition 3.24. If $n_T = 3$, since T has at least one edge (by hypothesis), at least one $e \to v$ is of the form e_v for some $v \in T^{(0)}$ so $d_{e_v} \geq b_0(d, \dots, d) \geq b_0(d, d, 0, \dots, 0) = B_0$. The latter inequality here follows from the fact that $b_0(d, \dots, d) > b_0(d, d, x, \dots, x)$ for $x < d$, by Proposition 3.7(3), upon taking a limit as $x \to 0$ (see Lemma 5.2 of [7]). Thus in this case Corollary 3.5 implies that $D_0(P_{v_T}(\boldsymbol{d})) \geq D_0(B_0, d, d)$.

For $\boldsymbol{d}$ satisfying (1) the above implies that

$$D_T(\boldsymbol{d}) \geq \sum_{v \in T^{(0)}} (n_v - 2) D_0(B_0, d, d) = \left[\left(\sum_{v \in T^{(0)}} n_v \right) - 2n \right] D_0(B_0, d, d),$$

since T has n vertices. Since T is a tree, its Euler characteristic is one so $|\mathcal{E}| = n - 1$. We also have $\sum_{v \in T^{(0)}} n_v = 2|\mathcal{E}| + |\mathcal{F}|$, recalling here that each edge of V that is not in $\mathcal{E}$ but has both endpoints in T contributes two distinct elements to $\mathcal{F}$ (see above

Definition 3.9). Thus the quantity in brackets above is $|\mathcal{F}|-2$, and the result follows in case (1).

It remains only to consider the case that $\boldsymbol{d}$ satisfies condition (3); ie that $J(P_v(\boldsymbol{d})) = J(P_w(\boldsymbol{d}))$ for some $v \in T^{(0)}$ and $w \in v-1$. This case follows directly from Lemma 3.28 and the induction hypothesis. The conclusion thus holds for each $\boldsymbol{d}_{\mathcal{F}} \in SAd_T \cap [d, D]^{\mathcal{F}}$ and hence, since $D > d$ is arbitrary, for each $\boldsymbol{d}_{\mathcal{F}} \in SAd_T$ with all entries at least d. Since the conclusion is vacuous for $\boldsymbol{d}_{\mathcal{F}} \notin SAd_T$, the result follows. □

Theorem 3.31 *Let C be a compact two-cell of the centered dual complex of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$, such that for some fixed $d > 0$ each edge of ∂C has length at least d. If C is a triangle then its area is at least that of an equilateral hyperbolic triangle with side lengths d. If ∂C has $k > 3$ edges, then*

$$\text{Area}(C) \geq (k-2)A_m(d).$$

Here $A_m(d)$ is the maximum of areas of triangles with two sides of length d, that of a semicyclic *triangle, whose third side is a diameter of its circumcircle.*

Proof If C is a triangle then it is centered (recall Definition 2.26), so the result follows directly from Corollary 3.5. Therefore assume below that ∂C has $k > 3$ edges.

Proposition 3.4 implies that $A_m(d)$, as defined above, equals $D_0(b_0(d,d),d,d)$ for b_0 as defined in Proposition 3.7. If C is a centered geometric dual cell, the conclusion thus follows by combining Proposition 3.24 with Corollary 3.5. We may therefore assume that $C = C_T$ is dual to a component T of the non-centered Voronoi subgraph (recall Definition 2.11). In this case C_T has area $D_T(\boldsymbol{d})$ by Lemma 3.14, where the entries of $\boldsymbol{d}$ are lengths of geometric duals to edges of T or its frontier in the Voronoi graph. Since each such edge has length at least d by hypothesis, the result follows directly from Proposition 3.30. □

4 Admissible spaces and area bounds with mild noncompactness

The goal of this section is to produce and prove a result analogous to Theorem 3.31 for centered dual 2–cells that are not compact, but for which the associated Voronoi subtree still has finite vertex set. The development follows a parallel track: we introduce an admissible space parametrizing all possible cells with a given edge length collection in Section 4.2, and minimize the area functional on it in Section 4.3.

Section 4.1 collects some useful results on horocyclic and horocyclic ideal polygons.

4.1 Horocyclic ideal polygons

Recall that *horocycles* of $\mathbb{H}^2$ are defined in Definition 1.8; in particular, a horocycle has a single *ideal point* on the sphere at infinity S_∞ of $\mathbb{H}^2$.

Definition 4.1 A *horocyclic polygon* is the convex hull in $\mathbb{H}^2$ of a locally finite subset of a horocycle. An *infinite* horocyclic polygon C is the convex hull of an infinite, locally finite subset of a horocycle. A *horocyclic ideal polygon* is the convex hull P of the union of geodesic rays joining a finite subset of a horocycle to its ideal point, the *ideal vertex* of P.

Note this agrees with Definition 1.10 in the special case of horocyclic ideal triangles. In particular, the triangles $\Delta(e_0, v_\infty)$ of Lemma 1.11 and Definition 2.11 fit this description.

Remark 4.2 As in the cyclic case (compare Remark 1.4), horocyclic and horocyclic ideal polygons are defined differently in our main reference [7, Definition 3.3] than above, but Proposition 3.8 there implies the definitions are equivalent.

If a horocyclic ideal polygon P has ideal vertex v then $\overline{P} = P \cup \{v\}$, where P is taken in $\mathbb{C}$ via the upper half-plane model (recall Definition 1.8), and the closure $\overline{P}$ of P is taken in the one-point compactification $\mathbb{C} \cup \{\infty\}$ of $\mathbb{C}$. Cyclically ordering the vertices of $\overline{P}$ as $\{x_0, \dots, x_{n-1}\}$ along the lines of Definition 1.3, we take the *side length collection* of P to be $(d_0, \dots, d_{n-1})$, where $d_i = d(x_{i-1}, x_i)$ unless x_i or x_{i-1} is v. In this case we define $d_i = \infty$.

The set of marked, oriented horocyclic n–gons is parametrized up to orientation-preserving isometry by their side length collections, determining a subset $\mathcal{HC}_n$ of $(\mathbb{R}^+)^n$ (recall Proposition 3.1). For horocyclic ideal n–gons, [7, Corollary 3.5] similarly gives the following.

Proposition 4.3 *For $n \geq 3$, the set of marked horocyclic ideal n–gons is parametrized by*

$$\mathcal{HI}_n = \{(d_0, \dots, d_{n-1}) \in (0, \infty]^n \mid d_{i_0} = d_{i_0+1} = \infty \text{ for a unique } i_0,\ 0 \leq i_0 < n\}.$$

It is the orbit of $\{(\infty, \infty)\} \times \mathbb{R}^{n-2}$ under cyclic permutation of entries.

The areas of horocyclic and horocyclic ideal polygons have nice explicit expressions recorded below from [7, Proposition 3.7].

Proposition 4.4 *For $n \geq 3$, the formulas below define a symmetric, continuous extension of D_0 to $\mathcal{AC}_n \cup \mathcal{HC}_n \cup \mathcal{HI}_n$. For $(d_0, \dots, d_{n-1}) \in \mathcal{HC}_n$ with maximal entry d_{i_0}, define*

$$D_0(d_0, \dots, d_{n-1}) = (n-2)\pi + 2\left[\sin^{-1}\left(\frac{1}{\cosh(d_{i_0}/2)}\right) - \sum_{i \neq i_0} \sin^{-1}\left(\frac{1}{\cosh(d_i/2)}\right)\right].$$

For $(d_0, \dots, d_{n-1}) \in \mathcal{HI}_n$ with $d_{i_0} = d_{i_0+1} = \infty$, take

$$D_0(d_0, \dots, d_{n-1}) = (n-2)\pi - 2 \sum_{i \neq i_0, i_0+1} \sin^{-1}\left(\frac{1}{\cosh(d_i/2)}\right).$$

Given $(d_0, \dots, d_{n-1})$ and $(d_0', \dots, d_{n-1}') \in \mathcal{HC}_n \cup \mathcal{HI}_n$, if up to a fixed permutation $d_i \leq d_i'$ for each i, and $d_i < d_i'$ for some i, then $D_0(d_0, \dots, d_{n-1}) < D_0(d_0', \dots, d_{n-1}')$.

By [7, Proposition 3.8], the respective formulas above give the area of the horocyclic or horocyclic ideal n–gon with side length collection $(d_0, \dots, d_{n-1})$.

Recall that for $d > 0$, the maximal-area triangle with two sides of length d has a third with length $b_0(d, d)$, where b_0 is as defined in Proposition 3.7. This is still less than the area of a horocyclic ideal triangle with finite side length d.

Corollary 4.5 *For any $d > 0$, $D_0(\infty, d, \infty) > D_0(b_0(d, d), d, d)$, for b_0 as in Proposition 3.7.*

Proof For any $x > d$, Corollary 3.5 implies that

$$D_0(b_0(d, x), d, x) > D_0(b_0(d, d), d, d).$$

Note also that $b_0(d, x) > x$. Taking a limit as $x \to \infty$, the result follows from continuity of the extension to $\mathcal{HI}_3$, by Proposition 4.4. □

The lemma below is the analog, in the context of horocyclic polygons, to Lemma 3.25.

Lemma 4.6 *Suppose $(c_0, \dots, c_{m-1}) \in \mathcal{HC}_m$ and $(d_0, \dots, d_{n-1}) \in \mathcal{HC}_n$ have largest entries c_0 and d_0, respectively, such that $c_0 = d_i$ for some $i > 0$. Then*

$$\boldsymbol{d} = (d_0, d_1, \dots, d_{i-1}, c_1, \dots, c_{m-1}, d_{i+1}, \dots, d_{n-1}) \in \mathcal{HC}_{m+n-2},$$

and $D_0(c_0, \dots, c_{m-1}) + D_0(d_0, \dots, d_{n-1}) = D_0(\boldsymbol{d})$. Analogously, if $(c_0, \dots, c_m) \in \mathcal{HC}_m$ has c_0 maximal, and $(d_1, \dots, d_{n-1}) \in (\mathbb{R}^+)^{n-1}$ has $d_i = c_0$ for some i, then

$$\begin{aligned} &D_0(c_0, \dots, c_{m-1}) + D_0(\infty, d_1, \dots, d_{n-1}, \infty) \\ &\qquad = D_0(\infty, d_1, \dots, d_{i-1}, c_1, \dots, c_{m-1}, d_{i+1}, \dots, d_{n-1}, \infty). \end{aligned}$$

The proof of Lemma 4.6 follows the lines of Lemma 3.25, with the references to Lemma 2.1 and Proposition 2.2 of [7] replaced by Proposition 3.8 there. Here is the geometric picture: if the largest side length of a horocyclic polygon P equals a side length of a (say) horocyclic ideal polygon Q, then upon moving P by an isometry so that $P \cap Q$ is the longest side of P, $P \cup Q$ is itself a horocyclic ideal polygon with area the sum $D_0(P) + D_0(Q)$.

4.2 Admissible spaces: The case of a non-compact edge

This is the analog of Section 3.2 for centered dual two-cells C_T such that the dual tree T has a non-compact edge e_0 but $T^{(0)}$ finite. We recall Proposition 2.9, which motivates this section's blanket hypothesis.

Blanket hypothesis Except where explicitly noted, here V is a finite graph with vertices of valence at least 3 and (possibly) some non-compact edges, and $T \subset V$ is a rooted subtree with a single non-compact edge e_0 and root vertex $v_T \in e_0$. Let $\mathcal{E}$ be the edge set of T and $\mathcal{F}$ its frontier in F. We let n_v denote the valence in V of a vertex v of T.

The major definitions and results of this section closely parallel those of Section 3.2, though almost all will require some revision. We will compare and contrast as appropriate. Below is the analog of Definition 3.9, differing from the original only in that we define e_v for $v = v_T$.

Definition 4.7 Partially order $T^{(0)}$ by setting $v < v_T$ for each $v \in T^{(0)} - \{v_T\}$, and $w < v$ if the edge arc in T joining $w \in T^{(0)} - \{v_T, v\}$ to v_T runs through v. Let $v-1$ be the set of $w < v$ joined to it by an edge, and say v is *minimal* if $v - 1 = \varnothing$. Let $e_{v_T} = e_0$, and for $v \in T^{(0)} - \{v_T\}$, let e_v be the initial edge of the arc in T joining v to v_T. For each $v \in T^{(0)}$, say $e \to v$ for each edge $e \neq e_v$ of V containing v.

Again the definition below differs from its predecessor Definition 3.10 only in treating v_T like other vertices of T.

Definition 4.8 Let $(\mathbb{R}^+)^{\mathcal{F}}$ be the set of tuples of positive real numbers indexed by the elements of $\mathcal{F}$, and define $(\mathbb{R}^+)^{\mathcal{E}}$ analogously. For any elements $\boldsymbol{d}_{\mathcal{E}} = (d_e \mid e \in \mathcal{E}) \in (\mathbb{R}^+)^{\mathcal{E}}$ and $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$, let $\boldsymbol{d} = (\boldsymbol{d}_{\mathcal{E}}, \boldsymbol{d}_{\mathcal{F}})$ and $P_v(\boldsymbol{d}) = (d_{e_0}, \dots, d_{e_{n-1}})$ for $v \in T^{(0)}$, where the edges of V containing v are cyclically ordered as $e_0, \dots, e_{n-1}$. We say the *admissible set* $Ad(\boldsymbol{d}_{\mathcal{F}})$ determined by $\boldsymbol{d}_{\mathcal{F}}$ is the collection of $\boldsymbol{d} \in (\mathbb{R}^+)^{\mathcal{E}} \times \{\boldsymbol{d}_{\mathcal{F}}\}$ such that:

1. For each $v \in T^{(0)}$, $P_v(\boldsymbol{d}) \in \mathcal{AC}_{n_v} - \mathcal{C}_{n_v}$ has largest entry d_{e_v}.
2. $J(P_v(\boldsymbol{d})) > J(P_w(\boldsymbol{d}))$ for each $w \in v - 1$, where $J(P_v(\boldsymbol{d}))$ and $J(P_w(\boldsymbol{d}))$ are the respective radii of $P_v(\boldsymbol{d})$ and $P_w(\boldsymbol{d})$.

Definition 4.9 Fix $\boldsymbol{d}_{\mathcal{F}} = (d_e \mid e \in \mathcal{F}) \in (\mathbb{R}^+)^{\mathcal{F}}$ such that $Ad(\boldsymbol{d}_{\mathcal{F}}) \neq \varnothing$. For $\boldsymbol{d} \in Ad(\boldsymbol{d}_{\mathcal{F}})$ define

$$D_T(\boldsymbol{d}) = \pi - 2\sin^{-1}\left(\frac{1}{\cosh(d_{e_0}/2)}\right) + \sum_{v \in T^{(0)}} D_0(P_v(\boldsymbol{d})),$$

where $P_v(\boldsymbol{d})$ is as in Definition 4.8 and $D_0(P)$ is as in Proposition 3.4.

Lemma 4.10 *Let C_T be a centered dual two-cell, dual to a component T of the non-centered Voronoi subgraph determined by locally finite $\mathcal{S} \subset \mathbb{H}^2$ with a non-compact edge e_0 and $T^{(0)}$ finite. Let $\mathcal{E}$ be the edge set of T and $\mathcal{F}$ its frontier in the Voronoi graph V, and for each edge e of V that intersects T let d_e be the length of the geometric dual to e. Then $\boldsymbol{d} = (d_e \mid e \in \mathcal{E}) \in Ad(\boldsymbol{d}_{\mathcal{F}})$, where $\boldsymbol{d}_{\mathcal{F}} = (d_e \mid (e, v) \in \mathcal{F}$ for some $v \in T^{(0)})$, and C_T has area $D_T(\boldsymbol{d})$.*

Proof The proof is analogous to that of Lemma 3.14, with a couple of differences. Again the main point is that for each $v \in T^{(0)}$, C_v is represented in $\mathcal{AC}_{n_v}$ by $P_v(\boldsymbol{d})$. In contrast with that case, here C_{v_T} is non-centered, by Proposition 2.9, and its longest side is the geometric dual γ_0 to e_0, by Lemma 2.5.

By Definition 2.11, $C_T = \Delta(e_0, v_\infty) \cup \bigcup_{v \in T^{(0)}} C_v$ in this case, where v_∞ is the ideal endpoint of e_0. Lemma 1.11 implies that the union above is non-overlapping, so the area of C_T is the sum of the areas of the C_v with that of $\Delta(e_0, v_\infty)$. But $\Delta(e_0, v_\infty)$ is a horocyclic ideal triangle with vertices on the unique horocycle through the endpoints of γ_0 with ideal point v_∞ (see Lemma 1.11), so its area is $\pi - 2\sin^{-1}(1/\cosh(d_{e_0}/2))$ by Proposition 4.4. The lemma follows. □

As with Definition 3.15, we will compactify $Ad(\boldsymbol{d}_{\mathcal{F}})$ here by expanding it somewhat.

Definition 4.11 For $\boldsymbol{d}_{\mathcal{F}} = (d_e \mid e \in \mathcal{F}) \in (\mathbb{R}^+)^{\mathcal{F}}$ let $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ consist of those $\boldsymbol{d} = (\boldsymbol{d}_{\mathcal{E}}, \boldsymbol{d}_{\mathcal{F}})$, for $\boldsymbol{d}_{\mathcal{E}} \in (\mathbb{R}^+)^{\mathcal{E}}$, such that:

(1) For each $v \in T^{(0)}$, with valence n_v in V, $P_v(\boldsymbol{d}) \in (\mathcal{AC}_{n_v} \cup \mathcal{HC}_{n_v}) - \mathcal{C}_{n_v}$ has largest entry d_{e_v}.

(2) $J(P_v(\boldsymbol{d})) \geq J(P_w(\boldsymbol{d}))$ for each $w \in v - 1$, where $J(P_v(\boldsymbol{d}))$ and $J(P_w(\boldsymbol{d}))$ are the respective radii of $P_v(\boldsymbol{d})$ and $P_w(\boldsymbol{d})$, and $J(P) \doteq \infty$ if $P \in \mathcal{HC}_n$ (see the final assertion of Proposition 3.3 below).

Note that $\overline{Ad}(\mathcal{F})$ above is somewhat larger than its analog from Definition 3.15, since it includes points of $\mathcal{HC}_n$. We nonetheless require only subtle changes to the analog of Lemma 3.19. In particular, the properties below apply to all vertices of T, including v_T, and in (1) below an inequality ceases to be strict.

Lemma 4.12 *Collections $\{b_e\colon (\mathbb{R}^+)^{\mathcal{F}} \to \mathbb{R}^+\}_{e\in\mathcal{E}}$ and $\{h_e\colon (\mathbb{R}^+)^{\mathcal{F}} \to \mathbb{R}^+\}_{e\in\mathcal{E}}$ are determined by the following properties. For $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$ and $\boldsymbol{d}_{\mathcal{E}} \in (\mathbb{R}^+)^{\mathcal{E}}$, with $\boldsymbol{d} = (\boldsymbol{d}_{\mathcal{E}}, \boldsymbol{d}_{\mathcal{F}})$:*

- *$P_v(\boldsymbol{d}) \in \mathcal{BC}_{n_v}$, with largest entry d_{e_v}, for all $v \in T^{(0)}$ if and only if $d_e = b_e(\boldsymbol{d}_{\mathcal{F}})$ for each $e \in \mathcal{E}$.*
- *$P_v(\boldsymbol{d}) \in \mathcal{HC}_{n_v}$, with largest entry d_{e_v}, for all $v \in T^{(0)}$ if and only if $d_e = h_e(\boldsymbol{d}_{\mathcal{F}})$ for each $e \in \mathcal{E}$.*

The collections $\{b_e\}$ and $\{h_e\}$ have the following properties:

1. *If $\boldsymbol{d} = (\boldsymbol{d}_{\mathcal{E}}, \boldsymbol{d}_{\mathcal{F}}) \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ then for each $e \in \mathcal{E}$, $b_e(\boldsymbol{d}_{\mathcal{F}}) \le d_e \le h_e(\boldsymbol{d}_{\mathcal{F}})$.*
2. *For $\boldsymbol{d}_{\mathcal{F}}$ and $v \in T^{(0)}$, $b_{e_v}(\boldsymbol{d}_{\mathcal{F}}) > \max\{b_e(\boldsymbol{d}_{\mathcal{F}}) \mid e \to v \in \mathcal{E}\} \cup \{d_e \mid e \to v \in \mathcal{F}\}$.*
3. *If $d'_e \ge d_e$ for each $e \in \mathcal{F}$ then $b_e(\boldsymbol{d}'_{\mathcal{F}}) \ge b_e(\boldsymbol{d}_{\mathcal{F}})$ for each $e \in \mathcal{E}$, where $\boldsymbol{d}'_{\mathcal{F}} = (d'_e)_{e\in\mathcal{F}}$.*

The proof directly follows that of Lemma 3.19, and there is no need to rehash it. We merely point out that the reason d_e as at most (instead of less than) $h_e(\boldsymbol{d}_{\mathcal{F}})$ in (1) here is that now $P_v(\boldsymbol{d})$ may be in $\mathcal{HC}_{n_v}$ for $\boldsymbol{d} \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$.

Lemma 4.13 *For any $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$, $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ is compact.*

The proof of this result is easier than the proof of its antecedent Lemma 3.21 because $(\mathcal{AC}_n \cup \mathcal{HC}_n) - \mathcal{C}_n$ is closed in $(\mathbb{R}^+)^n$, unlike $\mathcal{AC}_n - \mathcal{C}_n$; recall Remark 3.2. For this reason the criterion (1) of Definition 4.11 is preserved in limits of points in $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$. That criterion (2) is preserved in limits, and that $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ is bounded in $(\mathbb{R}^+)^{\mathcal{F}}$ away from its frontier in $\mathbb{R}^{\mathcal{F}}$, follow as before.

We also observe the analog of Lemma 3.22.

Lemma 4.14 *For $\boldsymbol{d}_{\mathcal{F}} = (d_e \mid e \in \mathcal{F}) \subset (\mathbb{R}^+)^{\mathcal{F}}$ such that $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}}) \ne \varnothing$, $D_T(\boldsymbol{d})$ is continuous on $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ and attains a minimum there.*

The only thing worth adding to the proof here is that $\sin^{-1}(1/\cosh(d_{e_0}/2))$ clearly varies continuously with $\boldsymbol{d}$ (compare Definitions 3.13 and 4.9). Finally, a version of Proposition 3.23 for the current context:

Proposition 4.15 *For $\boldsymbol{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$ with $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}}) \ne \varnothing$, at a minimum point $\boldsymbol{d} = (\boldsymbol{d}_{\mathcal{E}}, \boldsymbol{d}_{\mathcal{F}})$ for $D_T(\boldsymbol{d})$ on of the following holds:*

1. *$P_v(\boldsymbol{d}) \in \mathcal{BC}_{n_v}$ for each $v \in T^{(0)}$.*
2. *$P_{v_T}(\boldsymbol{d}) \in \mathcal{HC}_{n_{v_T}}$.*
3. *$J(P_v(\boldsymbol{d})) = J(P_w(\boldsymbol{d}))$ for some $v \in T^{(0)}$ and $w \in v - 1$.*

Proof The proof follows the strategy of Proposition 3.23: we suppose none of the above criteria holds at $\boldsymbol{d} \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$, fix $v \in T^{(0)}$ such that $P_v(\boldsymbol{d}) \notin \mathcal{BC}_{n_v}$, and show that reducing d_{e_v} while keeping all other entries of $\boldsymbol{d}$ constant produces a deformation through $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ that lowers the value of D_T.

That (2) does not hold implies for each $w \in T^{(0)}$ that $P_w(\boldsymbol{d}) \notin \mathcal{HC}_{n_w}$. This is because if $P_w(\boldsymbol{d}) \in \mathcal{HC}_{n_w}$ then $J(P_w(\boldsymbol{d})) = \infty$, so criterion (2) of Definition 4.11 implies that $J(P_{v'}(\boldsymbol{d})) = \infty$, and hence $P_{v'}(\boldsymbol{d}) \in \mathcal{HC}_{n_{v'}}$, for all $v' \in T^{(0)}$ with $w < v'$, in particular for $v' = v_T$.

For $v < v_T$ the argument of Proposition 3.23 thus shows that the deformation described in the first paragraph acts as claimed there. If $v = v_T$ then $e_v = e_0$ is the non-compact edge of T, so $\partial D_T / \partial d_{e_v}$ is not quite as described in (3.23.1). Instead we have

$$\frac{\partial D_T}{\partial d_{e_0}} = \frac{1}{\cosh(d_{e_0}/2)} - \sqrt{\frac{1}{\cosh^2(d_{e_0}/2)} - \frac{1}{\cosh^2 J(P_{v_T}(\boldsymbol{d}))}}. \tag{4.15.1}$$

The right-hand quantity above is $\partial D_0(P_{v_T}(\boldsymbol{d}))/\partial d_{e_0}$ (by Proposition 3.4); on the left is $[\pi - 2\sin^{-1}(1/\cosh(d_{e_0}/2))]'$ (by direct computation). □

4.3 Another area bound

Here we will prove an analog of Theorem 3.31 for centered dual 2–cells C_T that are dual to components T of the non-centered Voronoi subgraph with a non-compact edge but finite vertex set (recall Definition 2.11).

Theorem 4.16 *Let C_T be a centered dual 2–cell, dual to a component T of the non-centered Voronoi subgraph determined by locally finite $\mathcal{S} \subset \mathbb{H}^2$ with finite vertex set but a noncompact edge. For $d > 0$, if ∂C_T has k edges and each has length at least d then*

$$\mathrm{Area}(C_T) \geq D_0(\infty, b_0(d,d), \infty) + (k-3) D_0(b_0(d,d), d, d),$$

where D_0 measures area of cyclic, horocyclic and horocyclic ideal polygons (see Propositions 3.4 and 4.4), and $(b_0(d,d), d, d)$ is the side length collection of a semicyclic triangle with two sides of length d (see Propositions 3.1 and 3.7).

The proof strategy is similar to that of Theorem 3.31. In particular, we again induct on the number of vertices. Here however, in the one-vertex case T has a single non-compact edge. We address this directly below.

Lemma 4.17 *Let $T = \{e_0\}$ be a non-compact edge of a finite graph V, with vertex v of valence $n \geq 3$ in V. Cyclically enumerate the edges containing v as $e_0, e_1, \dots, e_{n-1}$. Then for any $\boldsymbol{d}_{\mathcal{F}} = (d_{e_1}, \dots, d_{e_{n-1}}) \in (\mathbb{R}^+)^{n-1}$, $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}}) = [b_{e_0}(\boldsymbol{d}_{\mathcal{F}}), h_{e_0}(\boldsymbol{d}_{\mathcal{F}})] \times \{\boldsymbol{d}_{\mathcal{F}}\}$ (where b_{e_0} and h_{e_0} are as in Lemma 4.12). D_T takes its minimum and maximum values at the left and right endpoints of this interval, respectively. These are given by*

$$D_T(b_{e_0}(\boldsymbol{d}_{\mathcal{F}}), \boldsymbol{d}_{\mathcal{F}}) = D_0(\infty, b_{e_0}(\boldsymbol{d}_{\mathcal{F}}), \infty) + D_0(b_{e_0}(\boldsymbol{d}_{\mathcal{F}}), \boldsymbol{d}_{\mathcal{F}}),$$
$$D_T(h_{e_0}(\boldsymbol{d}_{\mathcal{F}}), \boldsymbol{d}_{\mathcal{F}}) = D_0(\infty, \boldsymbol{d}_{\mathcal{F}}, \infty).$$

Before proving the lemma we record a useful corollary pertaining to horocyclic ideal polygons.

Corollary 4.18 *For $n \geq 2$ and $\boldsymbol{d} = (d_1, \dots, d_n) \in (\mathbb{R}^+)^n$, we have*

$$D_0(\infty, \boldsymbol{d}, \infty) > D_0(\infty, b_0(\boldsymbol{d}), \infty) + D_0(b_0(\boldsymbol{d}), \boldsymbol{d}).$$

Here b_0 is as in Proposition 3.7 and D_0 is from Proposition 4.4. For $d > 0$, if $d_i \geq d$ for each i then

$$D_0(\infty, \boldsymbol{d}, \infty) > D_0(\infty, b_0(d, d), \infty) + (n-1) D_0(b_0(d, d), d, d).$$

The first assertion follows directly from Lemma 4.17; the only thing to note is that by its definition in Lemma 3.19, in this case $b_{e_0} = b_0\colon (\mathbb{R}^+)^{n-2} \to \mathbb{R}^+$. The second follows from the first, applying monotonicity of b_0 (see Proposition 3.7(3)) and D_0 (by Corollary 3.5 and Proposition 4.4), and Proposition 3.24.

Proof of Lemma 4.17 Lemma 4.12(1) asserts that $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ is contained in the set above, and by the definitions of b_{e_0} and h_{e_0}, $P_v(\boldsymbol{d}) \in (\mathcal{AC}_n - \mathcal{C}_n) \cup \mathcal{HC}_n$ for any $\boldsymbol{d} = (d, \boldsymbol{d}_{\mathcal{F}})$ where $b_{e_0}(\boldsymbol{d}_{\mathcal{F}}) \leq d \leq h_{e_0}(\boldsymbol{d}_{\mathcal{F}})$ (cf Proposition 3.7). Thus Definition 4.11(1) holds for such $\boldsymbol{d}$, and since (2) holds vacuously in this case,

$$\overline{Ad}(\boldsymbol{d}_{\mathcal{F}}) = [b_{e_0}(\boldsymbol{d}_{\mathcal{F}}), h_{e_0}(\boldsymbol{d}_{\mathcal{F}})] \times \{\boldsymbol{d}_{\mathcal{F}}\}.$$

We appeal to Definition 4.9 and Proposition 3.4 to compute the derivative of $D_T(\boldsymbol{d})$ at $\boldsymbol{d} = (d, \boldsymbol{d}_{\mathcal{F}})$ in the open interval $(b_{e_0}(\boldsymbol{d}_{\mathcal{F}}), h_{e_0}(\boldsymbol{d}_{\mathcal{F}})) \times \{\boldsymbol{d}_{\mathcal{F}}\}$

$$\frac{\partial}{\partial d} D_T(d, \boldsymbol{d}_{\mathcal{F}}) = \frac{1}{\cosh(d/2)} - \sqrt{\frac{1}{\cosh^2(d/2)} - \frac{1}{\cosh^2 J(d, \boldsymbol{d}_{\mathcal{F}})}}.$$

This is clearly positive for all such d, so $D_T(d, \boldsymbol{d}_{\mathcal{F}})$ attains its minimum on $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ at $(b_{e_0}(\boldsymbol{d}_{\mathcal{F}}), \boldsymbol{d}_{\mathcal{F}})$ and its maximum at $(h_{e_0}(\boldsymbol{d}_{\mathcal{F}}), \boldsymbol{d}_{\mathcal{F}})$.

That $D_T(b_{e_0}(\boldsymbol{d}_{\mathcal{F}}),\boldsymbol{d}_{\mathcal{F}})$ is as described above is a direct application of Definition 4.9. That $D_T(h_{e_0}(\boldsymbol{d}_{\mathcal{F}},\boldsymbol{d}_{\mathcal{F}})) = D_0(\infty,\boldsymbol{d}_{\mathcal{F}},\infty)$ follows from the definition and the second assertion of Lemma 4.6, since $P_{v_T}(h_{e_0}(\boldsymbol{d}_{\mathcal{F}}),\boldsymbol{d}_{\mathcal{F}}) \in \mathcal{HC}_n$ by its definition in Lemma 4.12. □

We are now in position to prove Theorem 4.16.

Proof of Theorem 4.16 Recall (from Definition 2.11) that ∂C_T is the union of geometric duals to edges in the frontier $\mathcal{F}$ of the tree T dual to C_T, together with the two infinite edges of $\Delta(e_0,v_\infty)$. Here e_0 is the noncompact edge of T and v_∞ is its ideal endpoint. In particular, ∂C_T has $k = |\mathcal{F}|+2$ edges.

Let $\boldsymbol{d}_{\mathcal{F}}$ collect the lengths of the geometric duals to edges of $\mathcal{F}$. By hypothesis, $d_e \geq d > 0$ for each $e \in \mathcal{F}$. By Lemma 4.10, C_T has area equal to $D_T(\boldsymbol{d})$ for some $\boldsymbol{d} \in Ad(\boldsymbol{d}_{\mathcal{F}}) \subset \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$. We will thus prove the result by showing that for every tree T with one non-compact edge,

$$D_T(\boldsymbol{d}) \geq D_0(\infty, b_0(d,d),\infty) + (|\mathcal{F}|-1)D_0(b_0(d,d),d,d)$$

for every $\boldsymbol{d}_{\mathcal{F}} = (d_e \mid (e,v) \in \mathcal{F} \text{ for some } v \in T^{(0)})$ with all $d_e \geq d$ and $\boldsymbol{d} \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$.

If T has one vertex the result follows directly from Lemma 4.17. We will thus assume that T has $n>1$ vertices, and that for all trees with fewer than n vertices, $\min\{D_T(\boldsymbol{d}), \boldsymbol{d} \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})\}$ satisfies the conclusion if $d_e \geq d$ for all $e \in \mathcal{F}$.

A minimum point $\boldsymbol{d}$ for D_T on $\overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ satisfies one of the cases described in Proposition 4.15. Cases (1) and (3) follow from lines of argument analogous to those of Proposition 3.30. In Case (1), a direct computation yields the conclusion here as well. In Case (3) we collapse the edge f shared by v and w and appeal to induction. A new possibility here is that $J(P_v(\boldsymbol{d})) = J(P_w(\boldsymbol{d})) = \infty$; ie $P_v(\boldsymbol{d}) \in \mathcal{HC}_{n_v}$ and $P_w(\boldsymbol{d}) \in \mathcal{HC}_{n_w}$. Lemma 3.28 still holds in this case, though, replacing the appeal to Lemma 3.25 with one to Lemma 4.6.

Our treatment of Case (2) from the conclusion of Proposition 4.15 departs from the analogous case in the proof of Proposition 3.30. We suppose henceforth that $\boldsymbol{d} \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}})$ satisfies this case; ie that $P_{v_T}(\boldsymbol{d}) \in \mathcal{HC}_{n_{v_T}}$.

Let $\mathcal{E}$ be the edge set of T and e_0 its non-compact edge. Removing e_0 from T yields a collection $T_1,\dots,T_l$ of subtrees, one for each edge of T that contains v_T. For each i, T_i has a single non-compact edge e_i whose closure contains v_T. Let $\mathcal{E}_i$ be the edge set of T_i and $\mathcal{F}_i$ its frontier in V, and define $\boldsymbol{d}_{\mathcal{F}_i} = (d_e \mid (e,v) \in \mathcal{F}_i \text{ for some } v \in T_i^{(0)})$, $\boldsymbol{d}_{\mathcal{E}_i} = (d_e \mid e \in \mathcal{E}_i)$ and $\boldsymbol{d}_i = (\boldsymbol{d}_{\mathcal{E}_i},\boldsymbol{d}_{\mathcal{F}_i})$. It is evident that the property $\boldsymbol{d}_i \in \overline{Ad}(\boldsymbol{d}_{\mathcal{F}_i})$ is inherited from the corresponding property of $\boldsymbol{d}$.

Suppose v_T has valence n_0 in V, and let $e_{l+1},\dots,e_{n_0-1}$ be the collection of edges in $\mathcal{F}$ that contain v_T. Then $\mathcal{E}=\{e_0\}\cup\bigcup_{i=1}^{l}\mathcal{E}_i$, and $\mathcal{F}=\{e_{l+1},\dots,e_{n_0-1}\}\cup\bigcup_{i=1}^{l}\mathcal{F}_i$. We claim that

$$D_T(\boldsymbol{d})=\sum_{i=1}^{l}D_{T_i}(\boldsymbol{d}_i)+\sum_{i=l+1}^{n_0-1}D_0(\infty,d_{e_i},\infty).$$

The main point here is simply that because $P_{v_T}(\boldsymbol{d})\in\mathcal{HC}_{n_{v_T}}$, applying Lemma 4.6 gives

$$D_0(\infty,d_{e_0},\infty)+D_0(P_{v_T}(\boldsymbol{d}))=D_0(\infty,d_{e_1},\dots,d_{e_{n_0-1}},\infty)=\sum_{i=1}^{n_0-1}D_0(\infty,d_{e_i},\infty).$$

That the second and third quantities above are equal is evident on its face from the latter formula of Proposition 4.4. Since each vertex of T other than v_T is in exactly one T_i, the claim follows.

The inductive hypothesis applies to T_i for each i, so using the claim we find

$$\begin{aligned}D_T(\boldsymbol{d})\geq\sum_{i=1}^{l}\Big[D_0(\infty,b_0(d,d),\infty)+(|\mathcal{F}_i|-1)D_0(b_0(d,d),d,d)\Big]&\\+\sum_{i=l+1}^{n_0-1}D_0(\infty,d,\infty)&\\\geq D_0(\infty,b_0(d,d),\infty)\qquad\qquad&\\+D_0(b_0(d,d),d,d)\sum_{i=1}^{l}(|\mathcal{F}_i|-1)+(n_0-1-l)D_0(\infty,d,\infty).&\end{aligned}$$

The latter inequality above follows from Corollaries 4.5 and 3.5. Applying Corollary 4.5 again, and the fact that $|\mathcal{F}|=(\sum_{i=1}^{l}|\mathcal{F}_i|)+n_0-1-l$, gives the result. □

5 On hyperbolic surfaces

The main goal of this section is to prove Theorem 5.11. Then in Section 5.4 we will describe families of hyperbolic surfaces with maximal injectivity radius approaching its upper bound, showing this bound is sharp. First, in Section 5.1 below we recall some facts from [5] on the Delaunay tessellation and geometric dual complex of a finite subset of a hyperbolic surface, and use these to produce a description of the centered dual.

5.1 When covering a surface

Below we interpret [5, Theorem 6.23] for surfaces.

Theorem 5.1 *For a complete, oriented, finite-area hyperbolic surface F with locally isometric universal cover $\pi\colon \mathbb{H}^2 \to F$, and a finite set $\mathcal{S} \subset F$, the Delaunay tessellation of $\widetilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$ is a locally finite, $\pi_1 F$–invariant decomposition of $\mathbb{H}^2$ into convex polygons (the* cells*) such that each edge of each cell is a cell, and distinct cells that intersect do so in an edge of each. For each circle or horocycle of $\mathbb{H}^2$ that intersects $\mathcal{S}$ and bounds a disk or horoball B with $B \cap \mathcal{S} = S \cap \mathcal{S}$, the closed convex hull of $S \cap \mathcal{S}$ in $\mathbb{H}^2$ is a Delaunay cell. Each Delaunay cell has this form.*

For each parabolic fixed point $u \in S_\infty$ there is a unique Γ_u–invariant 2–cell C_u, where Γ_u is the stabilizer of u in $\pi_1 F$, whose unique circumcircle (in the sense above) is a horocycle with ideal point u. Each other cell is compact and has a metric circumcircle.

Fixing a locally isometric universal cover $\pi\colon \mathbb{H}^2 \to F$ determines an isomorphic embedding of $\pi_1 F$ to a lattice in $\mathrm{PSL}_2(\mathbb{R})$, so that π factors through an isometry $\mathbb{H}^2/\pi_1 F \to F$. An element of $\pi_1 F$ is *parabolic* if it fixes a unique $u \in S_\infty$ (recall Definition 1.8); such a point u is a *parabolic fixed point*.

Corollary 6.26 of [5] describes the image of the Delaunay tessellation in F itself:

Corollary 5.2 *For a complete, oriented hyperbolic surface F of finite area with locally isometric universal cover $\pi\colon \mathbb{H}^2 \to F$, and a finite set $\mathcal{S} \subset F$, there are finitely many $\pi_1 F$–orbits of Delaunay cells of $\widetilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$. The interior of each compact Delaunay cell embeds in F under π. For a cell C_u with parabolic stabilizer Γ_u, $\pi|_{int\, C_u}$ factors through an embedding of $int\, C_u/\,\Gamma_u$ to a set containing a cusp of F.*

A *cusp* of F is a non-compact component of the ϵ*–thin part* of F,

$$F_{(0,\epsilon]} = \{x \in F \mid injrad_x\, F \leq \epsilon\}$$

for some $\epsilon > 0$ that is less than the two-dimensional Margulis constant. See eg [2, Chapter D]. Each cusp is of the form $B/\,\Gamma_u$, for a horoball B whose ideal point is a parabolic fixed point u with cyclic stabilizer Γ_u in $\pi_1 F$.

Remark 6.24 of [5] identifies the geometric dual as a subcomplex of the Delaunay tessellation:

Remark 5.3 For a complete, oriented hyperbolic surface F of finite area with locally isometric universal cover $\pi\colon \mathbb{H}^2 \to F$, and a finite set $\mathcal{S} \subset F$, the geometric dual complex of $\widetilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$ consists precisely of the non parabolic-invariant Delaunay cells. The interior of each geometric dual cell is embedded in F by π.

We now build on these results to describe the centered dual complex. The first observations below use the notions of centeredness from Lemma 1.5 and Definition 2.1.

Lemma 5.4 *For a complete, oriented hyperbolic surface F of finite area with locally isometric universal cover $\pi\colon \mathbb{H}^2 \to F$, and a finite set $\mathcal{S} \subset F$ with $\widetilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$:*

- *If Voronoi vertices v and w of $\widetilde{\mathcal{S}}$ satisfy $v = g.w$ for some $g \in \pi_1 F$ then $J_v = J_w$ (recall Fact 1.6), and the geometric dual $C_v = g.C_w$ is centered if and only if C_w is.*
- *If Voronoi edges e and f determined by $\widetilde{\mathcal{S}}$ satisfy $\pi(e) = \pi(f)$, then e is centered if and only if f is centered.*

This follows from the fact that $\pi_1 F$ acts *isometrically* by covering transformations. If Voronoi edges v and w project to the same point of F, then the covering transformation taking v to w takes the sphere of radius J_v centered at v to the sphere of radius J_v centered at w, and C_v to C_w. And if a centered Voronoi edge e has the same projection as f, then the covering transformation taking e to f takes the intersection of e with its geometric dual to the intersection of f with its geometric dual.

Lemma 5.5 *For a complete, oriented hyperbolic surface F of finite area with locally isometric universal cover $\pi\colon \mathbb{H}^2 \to F$ and a finite set $\mathcal{S} \subset F$, any component T of the non-centered Voronoi subgraph of $\widetilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$ is a tree, with $T^{(0)}$ finite, that embeds in F under π.*

Proof Corollary 5.2 implies the set of Voronoi vertices determined by $\widetilde{\mathcal{S}}$ has finitely many $\pi_1 F$–orbits, since Voronoi vertices are geometric duals to compact 2–cells of the Delaunay tessellation (Remark 5.3). Therefore the set $\{J_v \mid v \in T^{(0)}\}$ has only finitely many distinct elements (recall Lemma 5.4). Let v_T satisfy $J_{v_T} \geq J_v$ for all $v \in T^{(0)}$. By Lemma 2.7, T is a tree and v_T satisfies $J_{v_T} > J_v$ for all $v \in T^{(0)} - \{v_T\}$ (cf Definition 2.8 and below).

Since covering transformations exchange components of the union of non-centered edges, if $\gamma.T \cap T \neq \varnothing$ for some $\gamma \in \pi_1 F - \{1\}$ then $\gamma.T = T$. Since $J_{\gamma.v} = J_v$ for each $v \in T^{(0)}$, the claim above would imply that $\gamma.v_T = v_T$ for such γ, contradicting freeness of the $\pi_1 F$–action. Therefore T does not intersect its $\pi_1 F$–translates and thus projects homeomorphically to F. Moreover, each $\pi_1 F$–orbit of Voronoi vertices contains at most one point of $T^{(0)}$, so $T^{(0)}$ is finite. □

Corollary 5.6 *For a complete, oriented hyperbolic surface F of finite area with locally isometric universal cover $\pi\colon \mathbb{H}^2 \to F$ and a finite set $\mathcal{S} \subset F$, the centered dual complex of $\widetilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$ is $\pi_1 F$–invariant, and π embeds the interior of each centered dual cell in F.*

Proof The invariance of the centered dual follows directly from Lemma 5.4 (recalling Definition 2.26). Since $g \in \pi_1 F$ has no fixed points as a covering transformation of $\mathbb{H}^2$, it is not hard to see that g does not preserve any vertex or edge, or any geometric dual two-cell (each of which is compact; recall Lemma 1.5). For a cell of the form C_T, where T is a component of the non-centered Voronoi subgraph, applying Definition 2.11 then Lemma 5.5 shows that $g(C_T) = C_{g(T)} \neq C_T$. □

5.2 The centered dual plus

For a non-compact hyperbolic surface F and a finite subset $\mathcal{S} \subset F$ with preimage $\widetilde{\mathcal{S}}$ in the universal cover $\mathbb{H}^2$ of F, there is a parabolic-invariant Delaunay cell of $\widetilde{\mathcal{S}}$ for each parabolic fixed point $u \in S_\infty$, by Theorem 5.1. The geometric dual complex does not contain C_u (see Remark 5.3), so since it is a subcomplex its underlying space does not intersect the interior of C_u.

Because the centered dual complex may contain horocyclic ideal triangles of the form $\Delta(e, v_\infty)$ in addition to geometric dual cells, its underlying space may overlap such C_u. However, if for instance all Voronoi edges are centered then the centered dual coincides with the geometric dual. Even if there are non-centered edges it is not clear how the underlying space of the centered dual intersects the interior of a given parabolic-invariant Delaunay cell.

In this brief section we will try to clarify the situation, ultimately introducing a complex that we call the centered dual plus, with the centered dual a subcomplex, in which parabolic-invariant Delaunay cells have been decomposed into unions of horocyclic ideal triangles.

Lemma 5.7 *For a horocycle S of $\mathbb{H}^2$, a locally finite set $\mathcal{S}_0 \subset S$ that is invariant under a parabolic isometry g fixing the ideal point v of S can be enumerated $\{s_i \mid i \in \mathbb{Z}\}$ so that for each i, the compact interval of S bounded by s_i and s_{i+1} contains no other points of $\mathcal{S}_0$, and $g(s_i) = s_{i+k}$ for each i and some fixed $k \in \mathbb{Z}$.*

For such an enumeration, the closed convex hull of $\mathcal{S}_0$ in $\mathbb{H}^2$ is the non-overlapping union $\bigcup_i T_i$, where T_i is horocyclic ideal triangle with vertices at s_i, s_{i+1} and v for each i.

Proof Applying an isometry of $\mathbb{H}^2$, one can arrange that $S = \mathbb{R} + i$, so $v = \infty$ and for some fixed $r \in \mathbb{R}$, $g(z) = z + r$ for all $z \in \mathbb{H}^2$. We may thus simply enumerate the points of $\mathcal{S}_0$ in order of increasing real part, choosing an arbitrary $s \in \mathcal{S}_0$ to be s_0. By g–invariance there are no points of $\mathcal{S}$ in the interval between $s_k = g(s_0)$ and $g(s_1)$, so since g preserves order of real parts it follows that $g(s_1) = s_{k+1}$. An induction argument gives $g(s_i) = s_{i+k}$ for all i.

For fixed i and any $n \in \mathbb{Z}$, let $r_n = \frac{1}{2}(\Re s_n - \Re s_i)$. The geodesic arc through s and s_i is contained in the Euclidean circle in $\mathbb{C}$ with center at $\Re s_i + r_n \in \mathbb{R}$ and radius $\sqrt{r_n^2 + 1}$. It follows from g–invariance that $\Re s_n \to \infty$ as $n \to \infty$. This implies that the geodesic arcs from s_i to the s_n intersect the vertical line through s_{i+1} at a sequence of points whose imaginary parts go to infinity. Hence by convexity the closed convex hull C of $\mathcal{S}_0$ contains the entire geodesic ray $[s_{i+1}, \infty)$.

Since the above holds for any i it is not hard to see that C contains $\bigcup_i T_i$, which further is clearly a non-overlapping union. For any $x \in \mathbb{H}^2$ outside this union there is some i such that $\Re s_i \le \Re x \le \Re s_{i+1}$, and the geodesic arc joining s_i to s_{i+1} separates x from T_i in the region $\{\Re s_i \le \Re z \le \Re s_{i+1}\}$. It follows that the geodesic γ_i of $\mathbb{H}^2$ containing s_i and s_{i+1} separates x from $\mathcal{S}_0$, so since x was arbitrary $C = \bigcup_i T_i$. $\square$

Lemma 5.8 *Let F be a complete, non-compact hyperbolic surface of finite area with universal cover $\pi\colon \mathbb{H}^2 \to F$, and for finite $\mathcal{S} \subset F$ let $\widetilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$. If C_T is a non-compact centered dual two-cell and $\Delta(e_0, v_\infty) \subset C_T$ (recall Definition 2.11) then there is a Delaunay two-cell C_{v_∞} invariant under a parabolic element of $\pi_1 F$ fixing v_∞, with $\Delta(e_0, v_\infty) \subset C_{v_\infty}$ such that the geometric dual γ to e_0 is an edge of C_{v_∞}.*

For a Delaunay two-cell C invariant under a parabolic subgroup Γ of $\pi_1 F$ fixing some $v_\infty \in S_\infty$, the intersection of int C with the centered dual complex, if non-empty, is a Γ–invariant union of the form $\bigcup(\Delta(e, v_\infty) - \gamma)$, where e (with geometric dual γ) ranges over the set of non-centered non-compact Voronoi edges with ideal endpoint v_∞.

Proof For C_T and $\Delta(e_0, v_\infty)$ as above, we recall from Lemma 1.9 that the endpoints of the geometric dual γ to e_0 are contained in a unique horocycle S with ideal point v_∞, and the horoball B bounded by S satisfies $B \cap \mathcal{S} = S \cap \mathcal{S}$. Theorem 5.1 therefore implies that the closed convex hull of $S \cap \mathcal{S}$ in $\mathbb{H}^2$ is a Delaunay two-cell C_{v_∞} invariant under a parabolic subgroup of $\pi_1 F$ fixing v_∞. The decomposition of Lemma 5.7 includes $\Delta(e_0, v_\infty)$.

For a Delaunay two-cell C invariant under a parabolic subgroup Γ of $\pi_1 F$ fixing some $v_\infty \in S_\infty$, Corollary 5.6 implies in particular that the intersection of C with the centered dual complex of $\widetilde{\mathcal{S}}$ is Γ–invariant. Since the centered dual is a union of cyclic Delaunay cells and triangles of the form $\Delta(e, v)$ as above (recall Definition 2.26), *int* C intersects only the ideal triangles. For any such $\Delta(e, v)$ the lemma's first assertion implies that $v = v_\infty$, $\Delta(e, v) \subset C$, and the geometric dual γ to e is an edge of C. $\square$

Proposition 5.9 *For a complete, oriented, non-compact hyperbolic surface F of finite area with locally isometric universal cover $\pi\colon \mathbb{H}^2 \to F$ and a finite set $\mathcal{S} \subset F$, there is a* centered dual complex plus *of $\widetilde{\mathcal{S}} \doteq \pi^{-1}(\mathcal{S})$ with underlying space*

$$\bar{\mathbb{H}}^2 \doteq \mathbb{H}^2 \cup \{v \in S_\infty \mid g(v) = v \text{ for some parabolic } g \in \pi_1 F\}.$$

Its vertex and edge sets include those of the centered dual and also, for each parabolic fixed point $v \in S_\infty$ that is not the endpoint of a non-centered Voronoi edge, the vertex v and an edge $[\boldsymbol{s}, v]$ for each vertex $\boldsymbol{s}$ of the corresponding Delaunay two-cell C_v (as in Theorem 5.1). The two-cells consist of

- *all centered dual two-cells; and*
- *for each Delaunay two-cell C that is invariant under a parabolic subgroup Γ of $\pi_1 F$ fixing $v \in S_\infty$, and each edge γ of C that does not intersect the interior of a centered dual two-cell, $T \cup \{v\}$, where T is the horocyclic ideal triangle spanned by v and γ.*

For each cell C of the centered dual plus, π is embedding on $\mathit{int}\, C$ and $\pi|_{C\cap\mathbb{H}^2}$ extends to a continuous map $C \to \bar{F}$, where $\bar{F}$ is the closed surface obtained by adding one point to each cusp of F. The images determine a cell decomposition of $\bar{F}$.

Proposition 5.9 follows directly from the prior results of this section, recalling that each geometric dual two-cell is contained in a centered dual two-cell by Definition 2.26.

5.3 Proof of the main theorem

For closed hyperbolic surfaces, the upper bound of Theorem 5.11 is assertion 1) of the main theorem of [1]. We reproduce this result below as Lemma 5.10, and prove it along the same lines using Böröczky's theorem. We first saw this kind of argument in the related result [9, Corollary 3.5].

Though it is not noted in [1], this strategy gives bounds for all surfaces. They are not sharp in the non-compact case. To obtain sharp bounds we use the centered dual machine.

Lemma 5.10 *For a complete, oriented hyperbolic surface F of finite area and any $x \in F$, $\mathit{injrad}_x\, F \le r_\chi$ for $r_\chi > 0$ defined by the equation $3\alpha(r_\chi) = \pi/(1-\chi(F))$, where $\alpha(r)$ is defined in Theorem 5.11 and $\chi(F)$ is the Euler characteristic of F.*

Proof Let $\pi\colon \mathbb{H}^2 \to F$ be the universal cover. For an open disk D of radius r embedded in F, $\pi^{-1}(D)$ is a packing of $\mathbb{H}^2$ by radius-r disks that is invariant under the $\pi_1 F$–action by covering transformations of $\mathbb{H}^2$.

For such a disk D let $s \in F$ be the center of D, and let $\mathcal{S} = \pi^{-1}(s) \subset \mathbb{H}^2$. $\mathcal{S}$ is locally finite at $\pi_1 F$–invariant, so its Voronoi tessellation is as well. The $\pi_1 F$–action is transitive on $\mathcal{S}$, so also is on Voronoi 2–cells. Thus any fixed two-cell is a fundamental domain for the $\pi_1 F$–action, hence by the Gauss–Bonnet theorem has area $-2\pi\chi(F)$.

Each component $\widetilde{D}$ of $\pi^{-1}(D)$ is contained in a Voronoi 2–cell $\widetilde{V}$, and the *local density* of $\pi^{-1}(D)$ at $\widetilde{D}$ is by definition the ratio $\mathrm{Area}(\widetilde{D})/\mathrm{Area}(\widetilde{V})$. Since $\mathrm{Area}(\widetilde{D}) = 2\pi(\cosh r - 1)$, Böröczky's theorem [3] asserts the bound

$$\frac{\mathrm{Area}(\widetilde{D})}{\mathrm{Area}(\widetilde{V})} = \frac{2\pi(\cosh r - 1)}{-2\pi\chi(F)} \le \frac{3\alpha(r)(\cosh r - 1)}{\pi - 3\alpha(r)}.$$

The quantity on the right-hand side above is interpreted as follows: it is the ratio of the area of intersection of an equilateral triangle T that has all side lengths $2r$ with the union of radius r disks centered at its vertices, divided by the area of T. Solving the inequality above yields $\pi \le 3\alpha(r)(1 - \chi(F))$, and the desired bound follows since α decreases with r. □

Theorem 5.11 *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 $2r$:*

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

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

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

Moreover, the collection of such surfaces with injectivity radius $r_{g,n}$ at some point is a non-empty finite subset of the moduli space $\mathfrak{M}_{g,n}$ of complete, oriented, finite-area hyperbolic surfaces of genus g with n cusps.

Proof If $n = 0$ (ie F is closed), the equation defining $r_{g,n}$ simplifies to

$$(2g - 1)3\alpha(r_{g,0}) = \pi.$$

This bound is supplied by Lemma 5.10, so we assume below that F has at least one cusp.

Let $\pi\colon \mathbb{H}^2 \to F$ be a locally isometric universal cover, fix $x \in F$ and let $\widetilde{\mathcal{S}} = \pi^{-1}\{x\}$. Let $\{C_1, \dots, C_k\}$ be a complete set of representatives in $\overline{\mathbb{H}}^2$ for $\pi_1 F$–orbits of two-cells of the centered dual plus. By Proposition 5.9 these project to the two-cells of a cell decomposition of $\overline{F}$, which is obtained from F by compactifying each cusp with a single point. Their interiors project homeomorphically to F, so by the Gauss–Bonnet theorem,

$$\text{Area}(C_1) + \cdots + \text{Area}(C_k) = -2\pi\chi(F).$$

(We should technically replace each C_i above by $C_i \cap \mathbb{H}^2$ above.) Since the centered dual plus has a vertex at each parabolic fixed point of $\pi_1 F$, its projection has a vertex at each point of $\bar{F} - F$, in addition to the vertex at x, for a total of $n+1$. Each edge of the projection either begins and ends at x or joins x to a point of $\bar{F} - F$. Edges in the former category have length at least $d \doteq 2\,injrad_x\,F$, and those in the latter intersect F in infinite-length arcs.

Each point of $\bar{F}-F$ is contained in at least one cell of the centered dual plus. Each horocyclic ideal triangle (including all those not in the centered dual; recall Proposition 5.9) has area at least $D_0(\infty, d, \infty)$ by Proposition 4.4. For a cell C_i with $n_i \geq 4$ edges, Theorem 4.16 asserts

$$\text{Area}(C_i) \geq D_0(\infty, b_0(d,d), \infty) + (n_i - 3) D_0(b_0(d,d), d, d).$$

Each non-triangular cell C_i of the centered dual plus that is entirely contained in F is compact and hence satisfies the bound of Theorem 3.31 (cf Proposition 3.24):

$$\text{Area}(C_i) \geq (n_i - 2) D_0(b_0(d,d), d, d).$$

A triangular such cell C_i satisfies $\text{Area}(C_i) \geq D_0(d,d,d)$ by Corollary 3.5.

Since F has genus g and n cusps, $\bar{F}$ is closed of genus g, and the projection of the centered dual plus is a cell decomposition with vertex set $\{x\} \cup (\bar{F} - F)$ of order $n+1$. It satisfies the Euler characteristic identity $v - e + f = \chi(\bar{F})$, where v, e and f are the numbers of vertices, edges and faces, respectively. Substituting $n+1$ for v and $2-2g$ for $\chi(\bar{F})$ yields

$$e - f = (n+1) - (2-2g) = 1 - \chi(F).$$

After renumbering if necessary, there exists $k_0 \leq k$ such that C_i has an ideal vertex if and only if $i \leq k_0$. Each such C_i has only one ideal vertex, so $k_0 \geq n$ since each of the n points of $\bar{F} - F$ is in the projection of such a cell. We will apply the area inequalities recorded above, together with

$$D_0(\infty, b_0(d,d), \infty) > D_0(\infty, d, \infty) > D_0(b_0(d,d), d, d) > D_0(d,d,d).$$

These follow respectively from Proposition 4.4, Corollary 4.5 and Corollary 3.5. Together with the above they imply that for $i \leq k_0$,

$$\text{Area}(C_i) \geq D_0(\infty, d, \infty) + (n_i - 3) D_0(d,d,d)$$

with equality holding if and only if C_i is a horocyclic ideal triangle with finite side of length d. For $k_0 < i \leq k$ we have

$$\mathrm{Area}(C_i) \geq (n_i - 2)\,D_0(d,d,d)$$

with equality again holding if and only if C_i is a triangle with all sides of length d. Applying these inequalities and the Gauss–Bonnet theorem yields

$$\begin{aligned} -2\pi\chi(F) &\geq k_0 \cdot D_0(\infty,d,\infty) + \left(\sum_{i=1}^{k}(n_i-2) - k_0\right)\cdot D_0(d,d,d) \\ &\geq n \cdot D_0(\infty,d,\infty) + \left(\sum_{i=1}^{k} n_i - 2k - n\right)\cdot D_0(d,d,d). \end{aligned}$$

Equality holds here if and only if $k_0 = n$, ie every ideal point of $\bar{F}$ is in a unique C_i, and every Delaunay edge has length d. The sum of n_i counts each edge of the centered dual plus twice, so $\sum_{i=1}^{k} n_i - 2k = 2e - 2f = 2 - 2\chi(F))$. Moreover, $D_0(\infty,d,\infty) = \pi - 2\beta(r)$ and $D_0(d,d,d) = \pi - 3\alpha(r)$, where $r = d/2$. It is not hard to check that α and β are strictly decreasing functions of r, so substituting above yields the desired inequality.

Examples 5.13 and 5.14 below describe some surfaces with injectivity radius $r_{g,0}$ and $r_{g,n}$ ($n > 0$), respectively. That there are only finitely many of these (for fixed g and n) follows from the fact that any such surface is triangulated by equilateral and horocyclic ideal triangles with all (finite) side lengths equal to d. Its isometry class is thus determined by the combinatorics of its triangulation, with only finitely many possibilities. □

5.4 Some examples

Below we describe a closed, oriented hyperbolic surface F of genus g with maximal injectivity radius r_g. The same examples were constructed in [1], but we give an alternate approach that extends easily to the non-compact case. We require a lemma.

Lemma 5.12 *Suppose C is a centered or semicyclic hyperbolic polygon with cyclically ordered vertex set $\{x_0,\dots,x_{n-1}\}$ and side length collection $(d_0,\dots,d_{n-1})$, and for $r \leq \min\{d_i/2\}$ let $B(x_i,r)$ be the closed metric disk of radius r centered at x_i. Then $B(x_i,r)\cap C$ is a full sector of $B(x_i,r)$ and $B(x_i,r)$ does not overlap $B(x_j,r)$ in C for $j \neq i$.*

Here a *sector* of a metric disk is its intersection with two half-planes whose boundaries contain its center. In particular, by the above $B(x_i,r)\cap\partial C$ is contained in the union of the edges containing x_i. This does not hold for all non-centered polygons.

Proof Suppose for the moment that C is centered, so it has the center v of its circumcircle in its interior. By Proposition 2.2 of [7], C is divided into isosceles triangles by arcs joining the x_i to v. Following that result, let T_i be the resulting triangle with vertices at x_{i-1}, x_i and v. The angle of C at x_i is $\alpha_i+\alpha_{i+1}$, where for each i, α_i is the angle of T_i at its vertices other than v. For each i, T_i is divided by a perpendicular ρ_i from v to the side opposite it into isometric right triangles, T_i^- containing x_{i-1} and T_i^+ containing x_i.

The result follows from the fact that T_i^+ (respectively, T_i^-) contains an entire sector of $B(x_i,r)$ (respectively, $B(x_{i-1},r)$) of angle measure α_i. This in turn follows from the fact that the sides of T_i^+ containing x_i have respective lengths $d_i/2\geq r$ and $J\geq\max\{d_i/2\}\geq r$ (where J is the circumcircle radius), and ρ_i opposite x_i intersects the first at a right angle. Thus the closest point to x_i on ρ_i is its intersection point with the first side above, $T_i^+\cap\partial C$.

The semicyclic case is similar but one can omit T_{i_0}, where d_{i_0} is maximal among the d_i, since this triangle is degenerate and contained in the union of the others. □

Example 5.13 Fix $g\geq 2$ and let $r_g=r_{g,0}$ from Theorem 5.11. Substituting into the defining equation for $r_{g,0}$ there and solving for $\alpha(r_g)$ yields

$$\alpha(r_g)=\frac{\pi}{3(2g-1)}.$$

Let $T_1,\dots,T_{4g-2}$ be a collection of equilateral triangles, each with all vertex angles equal to $\alpha(r_g)$, arranged in $\mathbb{H}^2$ sharing a vertex v so that for $1\leq i<j\leq 4g-2$, $T_i\cap T_j$ is an edge of each, if $j=i+1$, or else v. Then $P_g\doteq T_1\cup T_2\cup\cdots\cup T_{4g-2}$ is a $4g$–gon with all edge lengths equal and vertex angles that sum to

$$(4g-2)\cdot 3\cdot\frac{\pi}{3(2g-1)}=2\pi.$$

Label the edges of P as $a_1,b_1,c_1,d_1,a_2,b_2,\dots,d_{g-1},a_g,b_g,c_g,d_g$ in cyclic order. For each i let f_i be the orientation-preserving hyperbolic isometry such that $f_i(a_i)=\overline{c}_i$ and $c_i=f(P)\cap P$, and let g_i have $g_i(b_i)=\overline{d}_i$ and $d_i=g_i(P)\cap P$. Here the bar indicates that when a_i, b_i, c_i and d_i are given the boundary orientation from P, $f_i|_{a_i}$ and $g_i|_{b_i}$ reverse orientation.

One easily shows that the edge-pairing of P described above has a single quotient vertex, so since the vertex angles of P sum to 2π the Poincaré polygon theorem implies that the group $G = \langle f_1, g_1, \dots, f_g, g_g \rangle$ acts properly discontinuously on $\mathbb{H}^2$ with fundamental domain P.

For $r \le r_g$, if the open metric disk $D_r(v)$ of radius r centered at v does not embed in $F \doteq \mathbb{H}^2/G$ under the quotient map $\pi\colon \mathbb{H}^2 \to F$ then $D_r(v)$ intersects a translate $D_r(w)$ for some $w \in \pi^{-1}(\pi(v))$. For each vertex w of each T_i, T_i contains the entire sector of the open metric disk $D_r(w)$ that it determines, by Lemma 5.12. Moreover, disks of radius r centered at distinct vertices of T_i do not meet. It follows that $D_r(v) \cap D_r(w) = \varnothing$ for any distinct vertices v and w of P, and that P contains the full sector of any such $D_r(w)$ that it determines. We therefore find that $D_r(v)$ embeds in F, upon noting that $\pi^{-1}(\pi(v))$ is the set of vertices of G–translates of P. Thus F has injectivity radius r_g at P.

Example 5.14 Fix $g \ge 0$ and $n \ge 1$ (excluding $(g,n) = (0,1)$ or $(0,2)$), and take $r_{g,n}$ as in Theorem 5.11. Let $T_1, \dots, T_{4g+n-2}$ be equilateral triangles with side lengths $2r_{g,n}$, arranged as in Example 5.13 so that their union is a $(4g+n)$–gon P_0. Label the edges of P_0 in cyclic fashion as

$$e_1, \dots, e_n, a_1, b_1, c_1, d_1, a_2, \dots, d_{g-1}, a_g, b_g, c_g, d_g,$$

so that $v = e_1 \cap d_g$. Then append horocyclic ideal triangles $S_1, \dots, S_n$, each with finite side length $2r_{g,n}$, to P_0 so that $S_i \cap P_0 = e_i$ for each i. Let $P = P_0 \cup \bigcup S_i$.

For $1 \le i \le n$, let p_i be the parabolic isometry fixing the ideal point of S_i and taking one of its sides to the other, and for $1 \le i \le g$ let $f_i(a_i) = \overline{c}_i$ and $g_i(b_i) = \overline{d}_i$ as in Example 5.13. As in that example, each vertex of P is equivalent to v under the resulting edge-pairing. By definition of α and β, the vertex angles of P sum to

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

Therefore the Poincaré polygon theorem implies that

$$G = \langle p_1, \dots, p_n, f_1, g_1, f_2, \dots, f_g, g_g \rangle$$

acts properly discontinuously on $\mathbb{H}^2$ with fundamental domain P and quotient $F = \mathbb{H}^2/G$, a complete hyperbolic surface.

Inspecting the edge pairing one finds that F has n cusps. Its area is equal to that of P, $(4g+2n-2)\pi - 2\pi = 2\pi(2g-2+n)$, so by the Gauss–Bonnet theorem F has genus g. We claim that F has injectivity radius $r_{g,n}$ at the projection of v; the argument is completely analogous to that of Example 5.13.

Corollary 5.15 *For any $r > 0$, the function $x \mapsto D_0(2r, x, x)$ is continuous and increasing on $[2r, \infty]$. In particular,*

$$\pi - 3\alpha(r) = D_0(2r, 2r, 2r) < D_0(2r, \infty, \infty) = \pi - 2\beta(r),$$

whence $2\beta(r) < 3\alpha(r)$.

Proof For $2r \le x < \infty$, $(2r, x, x) \in \mathcal{C}_3$ since its maximal entry is not unique. The function above is therefore continuous and increasing on $[2r, \infty)$ by Corollary 3.5, and the result follows by taking a limit as $x \to \infty$, using Proposition 4.4. □

Example 5.16 For fixed $g \ge 2$, note that

$$(4g-2)3\alpha(r_{g-1,2}) \\ = [(4(g-1)+2-2)3\alpha(r_{g-1,2}) + 4\beta(r_{g-1,2})] + (6\alpha(r_{g-1,2}) - 4\beta(r_{g-1,2})).$$

The quantity in brackets above is 2π by definition of $r_{g-1,2}$ (see Theorem 5.11), so by Corollary 5.15 the entire sum is greater than 2π. Hence $r_{g-1,2} < r_{g,0}$, since $\alpha(r)$ decreases in r; moreover $(4g-2)3\alpha(r) > 2\pi$ for any $r \in (r_{g-1,2}, r_{g,0})$, and some rearrangement yields

$$(2g-2)2\pi > (4g-2)(\pi - 3\alpha(r)) = (4g-2)D_0(2r, 2r, 2r).$$

On the other hand, for such r we also have $(4g-4)3\alpha(r) + 4\beta(r) < 2\pi$, so an analogous rearrangement yields

$$(2g-2)2\pi < (4g-4)(\pi - 3\alpha(r)) + 2(\pi - 2\beta(r)) \\ = (4g-4)D_0(2r, 2r, 2r) + 2D_0(2r, \infty, \infty).$$

Thus by Corollary 5.15 and the intermediate value theorem there exists $x \in (2r, \infty)$ with

$$(2g-2)2\pi = (4g-4)D_0(2r, 2r, 2r) + 2D_0(2r, x, x). \tag{5.16.1}$$

We arrange triangles $T_1, T_2, \dots, T_{4g-2}$ as in Example 5.13; however in this case only the last $4g-4$ are equilateral, each with sides of length $2r$. We take T_1 and T_2 isosceles, each with one side of length $2r$ and others of length x, arranged so that b_1 and $T_2 \cap T_3$ have length $2r$, and a_1, $T_1 \cap T_2$ and c_1 have length x. Here the sides of $P = \bigcup T_i$ are cyclically labeled as in Example 5.13, starting at v, so that in particular a_1 and b_1 are sides of T_1, and $T_2 \cap \partial P = c_1 \sqcup \{v\}$.

As in Example 5.13, for each i there exists f_i taking a_i to $\bar{c}_i$, and g_i taking b_i to $\bar{d}_i$. The collection $\{f_i, g_i\}_{i=1}^{g}$ is an edge-pairing on P with a single quotient vertex. The sum of all vertex angles of P is the sum over i of the vertex angle sum of each T_i.

Since the area of T_i is the difference between π and its vertex angle sum, the choice of x and the equation above imply that the vertex angle sum of P is 2π. Therefore by the Poincaré polygon theorem $G = \langle f_i, g_i \mid 1 \le i \le g \rangle$ acts discontinuously on $\mathbb{H}^2$ with fundamental domain P.

The proof that $F = \mathbb{H}^2/G$ has injectivity radius r at the projection of v follows that of Example 5.13. The only additional note is that T_1 and T_2, being isosceles, are still centered, so the conclusions of Lemma 5.12 apply to them as well.

We claim that the minimal injectivity radius of F approaches zero as $r \to r_{g-1,2}^+$. This follows from two sub-claims: first, that the solution x to Equation (5.16.1) above goes to infinity as $r \to r_{g-1,2}$, and second, that the arc in T_1 joining points halfway up its sides with length x has length approaching 0. Toward the second, a hyperbolic trigonometric calculation shows that the length d of this arc satisfies

$$\cosh d = 1 + \frac{\cosh(2r) - 1}{2\cosh x + 2}.$$

The second sub-claim therefore follows from the first. The point p at the midpoint of a_1 has distance at most $2d$ from its image in c_1, so F has injectivity radius at most d at the projection of p, and the claim follows.

It remains to show the first sub-claim. Toward this end, let us recall from Theorem 5.11 that $r_{g-1,2}$ is defined by the equation

$$(2g-2)2\pi = (4g-4)D_0(2r_{g-1,2}, 2r_{g-1,2}, 2r_{g-1,2}) + 2D_0(2r_{g-1,2}, \infty, \infty).$$

Therefore by Corollary 5.15, for any fixed x_0 with $2r_{g-1,2} < x_0 < \infty$ we have

$$(2g-2)2\pi > (4g-4)D_0(2r_{g-1,2}, 2r_{g-1,2}, 2r_{g-1,2}) + 2D_0(2r_{g-1,2}, x_0, x_0).$$

For r near $r_{g-1,2}$ the function $r \mapsto (4g-4)D_0(2r, 2r, 2r) + D_0(2r, x_0, x_0)$ is continuous, so there exists $\epsilon > 0$ such that the inequality above holds with $r_{g-1,2}$ replaced with any $r \in (r_{g-1,2}, r_{g-1,2} + \epsilon)$. Since $D_0(2r, x, x)$ increases in x this implies for any such r that the solution to (5.16.1) lies outside $[2r, x_0]$. This proves the sub-claim.

References

[1] **C Bavard**, *Disques extrémaux et surfaces modulaires*, Ann. Fac. Sci. Toulouse Math. 5 (1996) 191–202 MR1413853

[2] **R Benedetti**, **C Petronio**, *Lectures on hyperbolic geometry*, Springer, Berlin (1992) MR1219310

[3] **K Böröczky**, *Packing of spheres in spaces of constant curvature*, Acta Math. Acad. Sci. Hungar. 32 (1978) 243–261 MR512399

[4] **J A De Loera**, **J Rambau**, **F Santos**, *Triangulations: Structures for algorithms and applications*, Algorithms and Computation in Mathematics 25, Springer, Berlin (2010) MR2743368

[5] **J DeBlois**, *The Delaunay tessellation in hyperbolic space* arXiv:1308.4899

[6] **J DeBlois**, *Tessellations of hyperbolic surfaces* arXiv:1103.4604

[7] **J DeBlois**, *The geometry of cyclic hyperbolic polygons* (2014) arXiv:1101.4971v4

[8] **B Farb**, **D Margalit**, *A primer on mapping class groups*, Princeton Mathematical Series 49, Princeton Univ. Press (2012) MR2850125

[9] **S Kojima**, **Y Miyamoto**, *The smallest hyperbolic 3–manifolds with totally geodesic boundary*, J. Differential Geom. 34 (1991) 175–192 MR1114459

[10] **D Mumford**, *A remark on Mahler's compactness theorem*, Proc. Amer. Math. Soc. 28 (1971) 289–294 MR0276410

[11] **B P Palka**, *An introduction to complex function theory*, Springer, New York (1991) MR1078017

[12] **J-M Schlenker**, *Small deformations of polygons and polyhedra*, Trans. Amer. Math. Soc. 359 (2007) 2155–2189 MR2276616

[13] **E Vanderzee**, **A N Hirani**, **D Guoy**, **E A Ramos**, *Well-centered triangulation*, SIAM J. Sci. Comput. 31 (2009/10) 4497–4523 MR2594991

Department of Mathematics, University of Pittsburgh
301 Thackeray Hall, Pittsburgh, PA 15260, USA

jdeblois@pitt.edu

http://www.pitt.edu/~jdeblois

Proposed: Danny Calegari
Seconded: Ian Agol, Dmitri Burago

Received: 5 September 2013
Revised: 19 March 2014