A FAST APPROXIMATION TO THE RIEMANN MAP

CHRISTOPHER J. BISHOP

Abstract. Given a planar region Ω bounded by a simple n-gon P we construct in $O(n)$ steps approximate locations for the conformal preimages of the vertices. These approximations are uniformly close to the actual prevertices in the sense that there is a K-quasiconformal self-map of the disk which maps our approximate vertices to the actual prevertices and K is independent of n and the geometry of P. The mapping that we use arises from the theory of hyperbolic 3-manifolds and has been previously studied by various authors including Thurston, Sullivan and Epstein and Marden. The fact we can construct it in $O(n)$ steps follows from the result of Chin, Snoeyink and Wang that the medial axis of a simple polygon can be computed in $O(n)$ steps. We also show that our map is closely related to the CRDT algorithm of Driscoll and Vavasis and we use it to resolve a conjecture of theirs.

Date: October 25, 2003.

1991 *Mathematics Subject Classification.* Primary: 30C35, Secondary: 30C85, 30C62 .

Key words and phrases. numerical conformal mappings, Schwarz-Christoffel formula, hyperbolic 3-manifolds, Sullivan's theorem, convex hulls, quasiconformal mappings, quasisymmetric mappings, medial axis, CRDT algorithm.

The author is partially supported by NSF Grant DMS 0103626.

1. Introduction

How fast can one approximate the Riemann mapping of a planar simply connected domain Ω to the unit disk $\mathbb{D}$? In this note we will combine results from hyperbolic 3-dimensional geometry, geometric function theory and computational geometry to show there is a map $\iota : \partial\Omega \to \mathbb{T} = \partial\mathbb{D}$ which is fast to compute and gives a rough but uniform approximation to the Riemann mapping.

Briefly, given a simply connected planar domain Ω, one defines the "dome" S_Ω over Ω as the upper envelope of all open semispheres in $\mathbb{R}^3_+$ whose base is in Ω. This is a surface in $\mathbb{R}^3_+$ which meets $\mathbb{R}^2$ along $\partial\Omega$. Thurston [100] observed that there is an isometry ι from S_Ω with its hyperbolic path metric to the hyperbolic unit disk. The restriction $\iota : \partial\Omega \to \partial\mathbb{D}$ is the map we will use as a fast approximation to the Riemann mapping.

In order to quantify what we mean by "fast" and "uniformly close" we will consider the case when Ω is bounded by a simple polygon P with n vertices $\mathbf{v} = \{v_k\}_1^n$. If $f : \mathbb{D} \to \Omega$ is conformal then we let $\mathbf{z} = f^{-1}(\mathbf{v})$ denote the "true" conformal preimages of the vertices and let $\mathbf{w} = \iota(\mathbf{v})$ be our approximate preimages. We will consider our approximation fast if we can compute the n points $\mathbf{w}$ with work $O(n)$. We will consider them uniformly close to the true preimages if $d_{QC}(\mathbf{w}, \mathbf{z}) \leq K$ for some $K < \infty$ independent of n and P, where

$$d_{QC}(\mathbf{w}, \mathbf{z}) = \inf\{\log K : \exists\ K\text{-quasiconformal } h : \mathbb{D} \to \mathbb{D} \text{ such that } h(\mathbf{z}) = \mathbf{w}.\}$$

Theorem 1. *There is a $C < \infty$ so that if Ω is bounded by a simply polygon P with n vertices we can find points $\mathbf{w} = \{w_1, \dots, w_n\} = \iota(\mathbf{v}) \subset \mathbb{T}$ so that*

1. *All n points in $\mathbf{w}$ can be computed in at most Cn steps.*
2. *$d_{QC}(\mathbf{w}, \mathbf{z}) < \log 8$ where $\mathbf{z}$ denotes the true conformal prevertices.*

A "step" is an infinite precision arithmetic operation or an evaluation of $\exp(x)$. The Schwarz-Christoffel formula gives a formula for the Riemann map of the disk onto a polygonal region Ω: if the interior angles of P are $\alpha\pi = \{\alpha_1\pi, \dots, \alpha_n\pi\}$, i.e.,

$$f(z) = A + C \int^z \prod_{k=1}^n (1 - \frac{w}{z_k})^{\alpha_k - 1} dw.$$

See e.g., [40], [78], [102]. The problem in applying the formula is that the points $\mathbf{z} = \{z_1, \dots z_n\}$ are unknown to us until we know f, so the formula seems circular.

However, there are various iterative methods for finding the points $\mathbf{z}$ starting from an initial guess (often taken to be n uniformly distributed points on $\mathbb{T}$), e.g., see [40], [67]. Theorem 1 says that an initial guess can be computed in $O(n)$ steps that is guaranteed to be in a fixed radius ball around the correct answer, independent of the geometry. I believe this is a new result, but I do not know if it offers any practical improvement over known methods. For example, in [7] Banjai and Trefethen give a $O(n)$ method using the fast multipole method (e.g., [27], [79], [54]) for finding the prevertices that is practical for tens of thousands of vertices (the bound, however is an average case analysis, not a uniform estimate for all polygons). For surveys of different numerical conformal mapping techniques see [51], [57], [63], [81], [101], [104].

The constant in part (2) of the theorem can be improved in various cases. For example, a result of Epstein, Marden and Markovic implies we can replace $\log 8$ by $\log 2$ if the region is convex. See Section 4. In addition to the conclusions of Theorem 1, the points $\mathbf{w} = \iota(\mathbf{v})$ have the following properties:

1. (Locally Lipschitz on boundary) If Ω contains a disk of radius 1, then we can choose $\mathbf{w}$ so that $|w_k - w_{k-1}| \leq |v_k - v_{k-1}|$ (i.e., the map decreases length along the boundary).
2. (Local dependence on boundary) If $\gamma \subset \partial\Omega$ is an arc whose endpoints lie on $\partial D \cap \partial\Omega$ for some open disk $D \subset \Omega$, then the images of vertices on γ can be chosen to depend only γ and not on the rest of $\partial\Omega$. For example, we can take ι to be the same on each of the bold arcs in Figure 1.

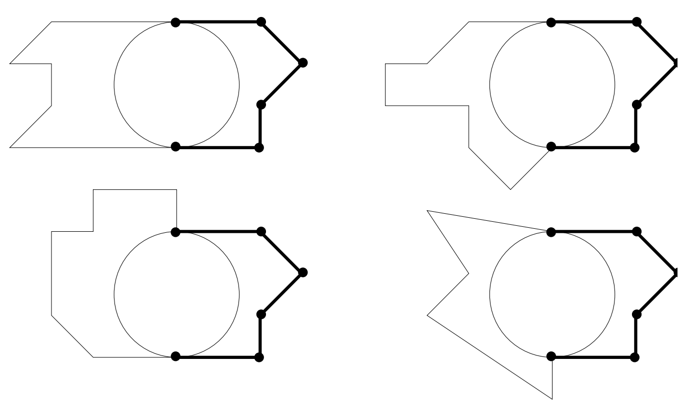

FIGURE 1. Vertices on the bold arc all have the same approximate preimages

3. (Preserves cross ratios of certain quadruples) If $\mathbf{v} = \{v_{j_1}, v_{j_2}, v_{j_3}, v_{j_4}\} \subset \partial D \cap \partial\Omega$ for some open disk $D \subset \Omega$ then the cross ratio of $\iota(\mathbf{v})$ equals the cross ratio of $\mathbf{v}$.
4. (Bounded modulus ratios) For any four vertices $v_{j_1}, v_{j_2}, v_{j_3}, v_{j_4}$ of P the corresponding true and approximate prevertices on $\mathbb{T}$ satisfy
$$\frac{1}{8} \leq \frac{\mathrm{Mod}_{\mathbb{D}}(z_{j_1}, z_{j_2}, z_{j_3}, z_{j_4})}{\mathrm{Mod}_{\mathbb{D}}(w_{j_1}, w_{j_2}, w_{j_3}, w_{j_4})} \leq 8,$$
where Mod denotes conformal modulus on the disk.

The medial axis of a domain is the set of points in the interior which are equidistant from two or more boundary points. It is a compact way of describing the shape of the region and there is now a large literature concerning its applications (see Section 6). The medial axis of Ω can be identified with the dual $\mathbb{R}$-tree of the bending lamination of the dome (see Section 6) and the dome of the polygon is easily written in terms of the medial axis, as is the boundary map ι. It is a theorem of Chin, Snoeyink, and Wang [29] that the medial axis of a simple n-gon can be computed in time $O(n)$. Using this we will show ι can be computed on all the vertices in time $O(n)$ and then use the theorem of Sullivan-Epstein-Marden (see Section 4) that ι has a uniform quasiconformal extension to a map from Ω to $\mathbb{D}$. For experts in computational and hyperbolic geometry, this sketch may already be a sufficient proof of Theorem 1. The rest of the paper is for the readers who would like to see more background and details. Since the result combines difficult theorems from two distinct areas, it seems worthwhile to provide a careful discussion of the definitions and results which we are using.

It is also worth noting that ι has a simple geometric description in terms of the medial axis that does not require passing to the dome in $\mathbb{R}^3_+$. We say an open disk in Ω is a medial axis disk if ∂D hits $\partial\Omega$ in at least two points. Fix some medial axis disk D_0. We can foliate $\Omega \setminus D_0$ by circular arcs which lie on the boundaries of medial axis disks in such a way that following the orthogonal trajectories to this "medial axis foliation" gives the map ι from $\partial\Omega$ to ∂D_0.

More precisely, suppose D is a medial axis disk other than D_0. Then $\Omega \setminus D$ has at least two components and exactly one of these intersects D_0. Take the boundary arcs of D corresponding to the other components (the ones that don't hit D_0) and put them in the foliation. Doing this for every medial axis disk foliates $\Omega \setminus D_0$ and

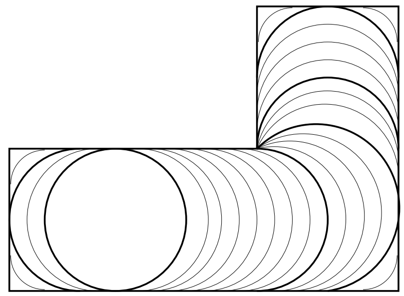

FIGURE 2. The medial axis foliation

the field of orthogonal directions to this foliation is Lipschitz with respect to the hyperbolic metric on Ω. Thus orthogonal trajectories exist and one can show they connect a point $x \in \partial\Omega$ with $\iota(x) \in \partial D_0$ (we will call this the medial axis flow).

The flow is particularly simple to understand if Ω is a finite union of disks. Then we can always write Ω as a union of a disk D_0 and a finite union of crescents. Inside each crescent the medial axis foliation consists of circular arcs passing through the vertices of the crescent. The orthogonal foliation consists of circular arcs which are orthogonal to both boundary arcs of the crescent. Following leaves of these foliations defines a map $\iota : \partial\Omega \to \partial D$ which is equivalent to the one we previously discussed. See Figure 3.

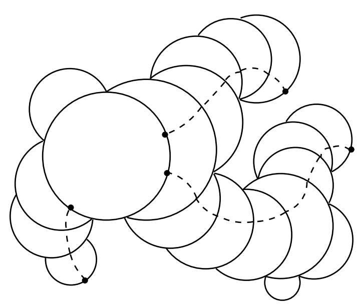

FIGURE 3. We can visualize ι by following orthogonal arcs in crescents

It is surprising (at least to the author) that such a simple geometric construction yields a map that is uniformly close to the Riemann map $\partial\Omega \to \partial D$, with an estimate independent of the geometry of Ω. One of the goals of geometric function theory is to understand conformal mappings and conformal invariants (such as conformal modulus) in geometric terms. Our observations about ι and its connection to both the dome of Ω and the medial axis indicates that these objects might be very interesting to investigate further. There is already a large literature on these topics in topology and computational geometry, but I believe the connections with complex function theory still remain to be investigated.

Instead of using the map $\iota : \partial\Omega \to \mathbb{T}$ to directly approximate the Riemann mapping, we can consider other variations. For example, if Ω is bounded by a polygon with vertices $\mathbf{v} = \{v_1, \dots, v_n\}$ we could replace Ω by an auxiliary domain Ω' which also has these points on its boundary and use the map $\iota : \partial\Omega' \to \mathbb{T}$. One such variation is the CRDT algorithm ("cross ratios of Delaunay triangulations") of Driscoll and Vavasis [41].

Driscoll and Vavasis present a method of avoiding the crowding problem in numerical conformal mappings (i.e., the prevertices can be much closer together than the vertices of P). Part of their method is to use cross ratios to choose an initial guess $\mathbf{w}$ for the conformal prevertices $\mathbf{z} = f^{-1}(\mathbf{v})$ (see Section 8). We will prove,

Theorem 2. *There is a $K < \infty$ (independent of P) so that the initial guess $\mathbf{w}$ of the CRDT algorithm satisfies $d_{QC}(\mathbf{w}, \mathbf{z}) \le K$.*

This does not quite give Theorem 1 because CRDT is not $O(n)$; the algorithm involves adding new vertices of angle π to the polygon P and the number of additions depends on the geometry of P (if P has long, narrow channels the number of new vertices is large). However, it does show the initial guess in CRDT is uniformly close to the correct prevertices in the sense discussed before. If we allow the new vertices to become dense in the sides of P, the CRDT prevertices $\mathbf{w}$ will converge to the points $\iota(\mathbf{v})$ given by Theorem 1. Thus we can think of the initial guess of CRDT as an approximation to the map ι.

It follows immediately from Theorem 2 and the fact that K-quasiconformal maps can change conformal modulus by at most a factor of K that

$$\frac{1}{K} \le \frac{\mathrm{Mod}_{\mathbb{D}}(\mathbf{z}')}{\mathrm{Mod}_{\mathbb{D}}(\mathbf{w}')} \le K,$$

for any 4 vertices $\mathbf{z}' = \{z_{j_1}, z_{j_2}, z_{j_3}, z_{j_4}\} \subset \mathbf{z}$, $\mathbf{w}' = \{w_{j_1}, w_{j_2}, w_{j_3}, w_{j_4}\} \subset \mathbf{w}$. Equivalently,

$$-\log K \le \log \mathrm{Mod}_{\mathbb{D}}(\mathbf{z}') - \log \mathrm{Mod}_{\mathbb{D}}(\mathbf{w}') \le \log K. \tag{1}$$

In [41] Driscoll and Vavasis asked if this held with the conformal modulus replaced by the cross ratio of the four points, in the case when the four points correspond to the vertices of two adjacent triangles in the Delaunay triangulation of P. We will see in Section 11 that the cross ratio version is false, so (1) is the appropriate modification of their conjecture.

The set of n-tuples of distinct points on $\mathbb{T}$ (modulo Möbius transformations) can be parameterized by using the conformal modulus of $n-3$ generalized quadrilaterals with vertices chosen from among the points (see Section 12). With these coordinates the space of n-tuples is identified with $\mathbb{R}^{n-3}$. Consider the map $\Psi : \mathbb{R}^{n-3} \to \mathbb{R}^{n-3}$ given by mapping an n-tuple to a polygon using Schwarz-Christoffel (with some fixed set of angles), followed by the ι map of the vertices to an n-tuple on the circle. Our results imply that Ψ satisfies

$$\|\Psi(z) - z\| \leq \log 8. \tag{2}$$

This may offer a partial answer to a question from [41], where Driscoll and Vavasis ask why an iteration similar to $\mathbf{w}_{n+1} = \mathbf{w}_n - \Psi(w_n)$ always seems to converge linearly to the correct conformal prevertices $\mathbf{z}$. This would be correct if the derivative of Ψ was close to the identity near $\mathbf{z}$. Equation (2) says that Ψ looks like the identity on large scales, but does this imply a similar statement on small scales?

The proof of Theorem 2 comes in two steps. To be precise, we state each step as a theorem.

Theorem 3. *If P is a planar polygon with vertices $\mathbf{v} = \{v_1, \dots, v_n\}$ then there exists a Riemann surface R_P containing P such that $\mathbf{v} \subset \partial P \cap \partial R_P$ and a K_1-quasiconformal mapping of $\psi : R_P \to \mathbb{D}$ so that $\iota(\mathbf{v}) = \mathbf{w}$, the CRDT approximation. K_1 is independent of n and the polygon P.*

Theorem 4. *If P is a planar polygon and we form a new polygon P' by adding vertices of angle π according to the CRDT algorithm, then there is a K_2-quasiconformal map $\phi : P' \to R_{P'}$ which fixes each vertex of P' ($R_{P'}$ is the surface from Theorem 3). K_2 is independent of n and the polygon P.*

Thus the CRDT guess $\mathbf{w} = \psi \circ \phi(\mathbf{v}) = \psi \circ \phi \circ f(\mathbf{z})$ is a K_1K_2-quasiconformal image of the true conformal prevertices and so Theorem 2 holds. Theorem 3 is a minor variation of a known result for planar domains; Theorem 4 is quite simple, based on the construction of Driscoll and Vavasis.

The paper breaks into two main parts: Sections 2 to 7 give basic background and the proof of Theorem 1. Sections 8 to 10 describe the CRDT algorithm and prove Theorems 3 and 4. A little more precisely:

Section 2: Definitions and basic properties of conformal modulus and cross ratios.

Section 3: Definitions and basic properties of quasiconformal mappings.

Section 4: Definition of the ι map for general simply connected domains and its connection to hyperbolic 3-dimensional geometry.

Section 5: More about ι in the special case when the domain is a finite union of disks.

Section 6: The medial axis of a polygon.

Section 7: How to build ι using the medial axis.

Section 8: We review the CRDT algorithm and explain why it is an ι map.

Section 9: We prove Theorem 3.

Section 10: We prove Theorem 4.

Section 11: We show the Driscoll-Vavasis conjecture is not correct for cross ratios.

Section 12: We discuss some variations, questions and conjectures.

2. Conformal modulus and cross ratios

Suppose Γ is a family of locally rectifiable paths in a planar domain Ω and ρ is a non-negative Borel function on Ω. We say ρ is admissible for Γ if

$$\ell(\Gamma) = \inf_{\gamma\in\Gamma} \int_\gamma \rho ds \geq 1,$$

and define the modulus of Γ as

$$\mathrm{Mod}(\Gamma) = \inf_{\Omega} \int_M \rho^2 dxdy,$$

where the infimum is over all admissible ρ for Γ. This is a well known conformal invariant whose basic properties are discussed in many sources such as Ahlfors' book [1]. A generalized quadrilateral Q is a Jordan domain in the plane with four specified boundary points x_1, x_2, x_3, x_4 (in counterclockwise order). We define the modulus of Q, $M_Q(x_1, x_2, x_3, x_4)$ (or just M_Q or $M(Q)$ if the points are clear from context), as the modulus of the path family in Q which connects the arc (x_1, x_2) to the arc (x_3, x_4). This is also the unique positive real number M such that Q can be conformally mapped to a $1 \times M$ rectangle with the arcs (x_1, x_2), (x_3, x_4) mapping to the opposite sides of length 1. In this paper, we will be particularly concerned with the case when $Q = \mathbb{D}$ and we are given four points in counterclockwise order on the unit circle.

Given a generalized quadrilateral Q with four boundary points x_1, x_2, x_3, x_4, the quadrilateral Q' with vertices x_2, x_3, x_4, x_1 is called the reciprocal of Q and it is easy to see that $\mathrm{Mod}(Q') = 1/\mathrm{Mod}(Q)$.

Given four distinct points a, b, c, d in the plane we define their cross ratio as

$$\mathrm{cr}(a,b,c,d) = \frac{(d-a)(b-c)}{(c-d)(a-b)}.$$

Note that $\mathrm{cr}(a,b,c,z)$ is the unique Möbius transformation which sends a to 0, b to 1 and c to ∞. This makes it clear that cross ratios are invariant under Möbius transformations; that $\mathrm{cr}(a,b,c,d)$ is real valued iff the four points lie on a circle; and is negative iff in addition the points are labeled in counterclockwise order on the circle.

If the four points lie on $\mathbb{T}$, then since cr and $M_{\mathbb{D}}$ are both invariant under Möbius transformations of the disk to itself, each must be a function of the other in this case. The usual way to represent this function (e.g., as in Ahlfors' book [1]) is to map the disk to the upper half plane, $\mathbb{H}$, sending the points a, b, c to $0, 1, \infty$ respectively and d to $-P = \mathrm{cr}(a,b,c,d) \in (-\infty, 0)$. Then $M_{\mathbb{D}}(a,b,c,d)$ is the same as the modulus of the path family in $\mathbb{H}$ connecting $(-\infty, P)$ to $(0,1)$. By symmetry, this is twice the modulus of the path family of closed curves in the plane which separate $[P, 0]$ from $[1, \infty]$. We will denote this modulus by $M(P)$. The transformation $z \to (z-1)/(z+P)$ sends $0, 1, \infty, -P$ to $-\frac{1}{P}, 0, 1, \infty$, so by Möbius invariance of modulus and the fact that conjugate quadrilaterals have reciprocal moduli, we see that

$$M_{\mathbb{H}}(\frac{1}{P}) = \frac{1}{M_{\mathbb{H}}(P)}, \tag{3}$$

and hence $M_{\mathbb{H}}(1) = 1$ and $M(1) = 1/2$.

In [1] Ahlfors gives a formula (page 46, equation (16)) relating P and M

$$P + 1 = \exp(2\pi M) \frac{1}{16} \prod_{n=1}^{\infty} (\frac{1 + \exp((1-2n)2\pi M)}{1 + \exp((-2n)2\pi M)})^8.$$

For $M > 0$ the infinite product converges and for M large (say $M \geq 1$) we have

$$\prod_{n=1}^{\infty} (\frac{1 + \exp((1-2n)2\pi M)}{1 + \exp((-2n)2\pi M)})^8 = 1 + 8e^{-2\pi M} + O(e^{-4\pi M}).$$

Thus for $P \geq 1$, (equivalently $M \geq 1$), we have

$$\log(P+1) = 2\pi M - \log 16 + 8e^{-2\pi M} + O(e^{-4\pi M}),$$

which implies

$$P \simeq \exp(2\pi M).$$

For $0 < P \leq 1$, (equivalently $0 < M \leq 1$), we can use (3) to deduce

$$\log(\frac{1}{P}+1) = \frac{\pi}{2M} - \log 16 + 8e^{-\pi/(2M)} + O(e^{-\pi/M}),$$

which implies

$$P \simeq \exp(-\frac{\pi}{2M}).$$

In other words,

$$M \simeq \frac{1}{2\pi}\log P, \qquad P \gg 1,$$

$$M \simeq \frac{\pi}{2|\log P|}, \qquad P \ll 1,$$

Thus for $\mathbf{x} = \{x_1, x_2, x_3, x_4\} \subset \mathbb{T}$, since $\mathrm{Mod}_{\mathbb{D}} = 2M$,

$$M_{\mathbb{D}}(\mathbf{x}) \simeq \frac{1}{\pi}\log|\mathrm{cr}(\mathbf{x})|, \qquad |\mathrm{cr}(\mathbf{x})| \gg 1,$$

$$M_{\mathbb{D}}(\mathbf{x}) \simeq \frac{\pi}{|\log|\mathrm{cr}(\mathbf{x})||}, \qquad |\mathrm{cr}(\mathbf{x})| \ll 1,$$

Another elegant connection between modulus and cross ratios is given in [6] (Bagby shows that conformal modulus for a ring domain is given by minimizing an integral involving logarithms of cross ratios).

3. Quasiconformal mappings

Quasiconformal mappings are a generalization of conformal mappings which play an important role in modern analysis and a central role in the current paper. There are (at least) three equivalent definitions of a K-quasiconformal mapping between planar domains. Suppose $f : \Omega \to \Omega'$ is a homeomorphism.

Geometric definition: for any generalized quadrilateral $Q \subset \Omega$, $\mathrm{Mod}(Q)/K \leq \mathrm{Mod}(f(Q)) \leq K\mathrm{Mod}(Q)$.

Analytic definition: f is absolutely continuous on almost every vertical and horizontal line and the partial derivatives of f satisfy $|f_{\bar{z}}| \leq k|f_z|$ where $k = (K-1)/(K+1)$.

Metric definition: For every $x \in \Omega$

$$\limsup_{r\to 0} \frac{\max_{y:|x-y|=r}|f(x)-f(y)|}{\min_{y:|x-y|=r}|f(x)-f(y)|} \leq K.$$

For a proof of the equivalence of the first two, see [1] and for a discussion of the third and a generalization to metric spaces see [56] and its references. A mapping is conformal iff it is 1-quasiconformal and the composition of a K_1-quasiconformal map with a K_2-quasiconformal map is (K_1K_2)-quasiconformal. Thus the distance used in Theorem 1 satisfies the triangle inequality.

Although quasiconformal maps must have derivatives almost everywhere, they can be non-differentiable on smaller sets. For example, the famous von Koch snowflake (Figure 4) is a non-differentiable curve which is the image of the unit circle under a quasiconformal map of the plane (such an image is called a quasicircle; they have a simple geometric characterization, e.g., see [1]). Even though they don't have

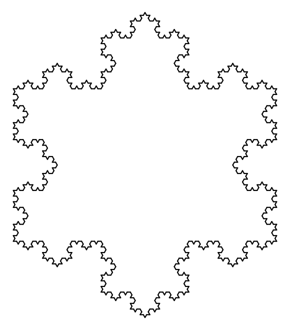

FIGURE 4. The von Koch snowflake is a quasicircle.

to be differentiable everywhere, by Mori's theorem every K-quasiconformal map is Hölder continuous of order $1/K$. Moreover, any quasiconformal map of $\mathbb{D}$ to itself extends continuously to the boundary. The extension is also a homeomorphism and its restriction to the boundary is a quasisymmetric homeomorphism, i.e., there is an $M < \infty$ (depending only on K) so that $1/M \leq |f(I)|/|f(J)| \leq M$, whenever $I, J \subset \mathbb{T}$ are adjacent intervals of equal length. Conversely, any quasisymmetric homeomorphism of $\mathbb{T}$ can be extended to a K-quasiconformal selfmap of the disk, where K depends only on M.

Quasiconformal maps are a generalization of biLipschitz maps, i.e., maps which satisfy

$$\frac{1}{K} \leq \frac{|f(x) - f(y)|}{|x - y|} \leq K.$$

For K-quasiconformal selfmaps of the disk, there is almost a converse. Although a quasiconformal map $f : \mathbb{D} \to \mathbb{D}$ need not be biLipschitz, it is a quasi-isometry of the disk with its hyperbolic metric ρ (see Section 4), i.e., there are constants A, B such that

$$\frac{1}{A}\rho(x, y) - B \leq \rho(f(x), f(y)) \leq A\rho(x, y) + B.$$

This says f is biLipschitz for the hyperbolic metric at large scales. A quasi-isometry is also called a rough isometry in some sources. In [46] Epstein, Marden and Markovic show that any K-quasiconformal selfmap of the disk is a quasi-isometry respect to the hyperbolic metric with $A = K$ and $B = K \log 2$ if $1 \leq K \leq 2$ and $B = 2.37(K - 1)$ if $K > 2$.

A quasi-isometry of the disk to itself need not be quasiconformal (indeed, it need not even be continuous), but there is a close connection in terms of boundary values. If $h : \mathbb{D} \to \mathbb{D}$ is a quasi-isometry then a result of Vaisala [103] implies that h has a continuous extension to the boundary and that this extension is quasisymmetric, with a constant that depends only on rough isometry constants of h. Hence there is a K-quasiconformal self-map H of the disk with these same boundary values, and with K depending only on rough isometry constants of h. Thus to prove a circle homeomorphism has a quasiconformal extension to the disk, it suffices to prove it has a quasi-isometric extension. This is often easier in practice and we will make use of this fact later.

One particular kind of quasiconformal map which we will use later in the paper is a "crescent map". A crescent is a simply connected plane domain that is bounded by two circular arcs which meet at two distinct points. A crescent is determined up to Möbius transformations by the interior angle at the vertices. We will use the observation that if C_1 and C_2 are crescents with angles $\alpha_1 \geq \alpha_2$ respectively, then there is a (α_1/α_2)-quasiconformal map between then which sends the vertices to vertices and is equals a Möbius transformation one each boundary arc. One can prove this easily by conjugating each crescent by a Möbius transformation so the vertices become 0 and ∞ and the crescents map to wedges of the form $\{r\exp(i\theta) : 0 < \theta < \alpha\}$. Then a map of the form $r\exp(i\theta) \to r\exp(\alpha_2\theta/\alpha_1)$ has the desired properties.

4. The Sullivan-Epstein-Marden theorem

In this and the next section we will summarize some basic facts connecting hyperbolic geometry, quasiconformal mappings and convex hulls. I am not aware of a single source for all the results we will discuss, but almost everything we need is covered in recent papers of Epstein, Marden and Markovic [43], [44], [45], [46], [47] and myself [14], [15], [16], [17]. Also related are the papers of Bridgeman [23], [24], and Bridgeman and Canary [25]. The books [9], [42] [90], [100], each deal with hyperbolic geometry and 3-manifolds, and explain how this material fits into that theory, although we will not use that aspect of the theory.

The hyperbolic metric on $\mathbb{D}$ is given by $d\rho_{\mathbb{D}} = 2|dz|/(1-|z|^2)$. Geodesics for this metric are circles orthogonal to the boundary and the orientation preserving isometries are exactly the Möbius transformations which preserve the disk, which all have the form $z \to e^{i\theta}(z-a)/(1-\bar{a}z)$, for some $\theta \in \mathbb{R}$ and $a \in \mathbb{D}$. The hyperbolic metric ρ_Ω on a simply connected domain Ω (or Riemann surface) is defined by transferring the metric on the disk to Ω by the Riemann map. We will sometimes write ρ instead of ρ_Ω when the domain is clear from context.

The hyperbolic metric on the upper half space $\mathbb{R}^3_+$ is $d\rho_{\mathbb{R}^3_+} = |dz|/\text{dist}(z, \mathbb{R}^2)$. As before, geodesics are circles orthogonal to the boundary, and orientation preserving isometries are exactly the Möbius transformations (every such transformation acting on $\mathbb{R}^2$ extends to an orientation preserving isometry of $\mathbb{R}^3_+$, and conversely).

Given a closed set $E \subset \mathbb{R}^2$ we let $C(E) \subset \mathbb{H}^3$ be the hyperbolic convex hull of E. This is the smallest convex set in $\mathbb{R}^3_+$ which contains all the infinite hyperbolic geodesics (i.e., circles orthogonal to $\mathbb{R}^2$) with both endpoints in E. The complement of $C(E)$ is the union of all hyperbolic half-spaces in $\mathbb{R}^3_+$ which do not intersect the set of geodesics; i.e., it is the union of all hemispheres centered on $\mathbb{R}^2$ whose bases on $\mathbb{R}^2$ miss E. We are most interested in the case when Ω is a simply connected plane domain and $E = \Omega^c$ is its complement. Let $S = S_\Omega$ be the boundary component of $C(\Omega^c)$ which separates Ω from $C(\Omega^c)$. This is called the "dome" of Ω. The name is appropriate because S is boundary of the union of all hemispheres centered on $\mathbb{R}^2$ whose bases are disks contained in Ω. Let ρ_S denote the intrinsic path metric on S

(using hyperbolic arclength) and let ρ denote the usual hyperbolic metric on the unit disk $\mathbb{D}$, the upper half space $\mathbb{H}^3$ or Ω. The most important facts about the dome of Ω are the following two theorems.

Theorem 5 (Thurston). *There is an isometry ι from (S, ρ_S) to $(\mathbb{D}, \rho)$.*

Theorem 6 (Sullivan-Epstein-Marden). *There is a $K_0 < \infty$ so that for any hyperbolic simply connected domain $\Omega \subset \mathbb{R}^2$ (other than the complement of a circular arc), there is a K_0-biLipschitz map σ from (Ω, ρ) to (S, ρ_S) which extends continuously to the identity on $\partial\Omega$. Consequently, there is a universal $K < \infty$ so that σ is K-quasiconformal. Moreover, the map σ is conformally natural in the sense that if Ω is invariant under a group of Möbius transformations then σ commutes with the group action.*

The first result appears in Thurston's notes [100] (with more detailed proofs in the papers of Epstein-Marden [43] and Rourke [92]). The second was apparently known to Thurston and appeared in Sullivan's paper [99] in the case when Ω is invariant under a convex, co-compact group of Möbius transformations. Epstein and Marden [43] proved the more general statement quoted above. Alternate proofs are given in [14], [17], [44]. The complement of a circular arc is only exceptional because in that case S needs to be interpreted as a two-sided "folded" surface. This is carefully explained in [45].

The Sullivan-Epstein-Marden convex hull theorem is part of Thurston's hyperbolization theorem for 3-manifolds that fiber over the circle. If G a discrete group of Möbius transformations acting on $S^2 = \mathbb{R}^2 \cup \{\infty\}$ then the group action extends to a discrete group of isometries on $\mathbb{R}^3_+$ with its hyperbolic metric. Thus $\mathbb{R}^3_+/G$ is a hyperbolic manifold (assuming G has no elliptic elements, i.e., no elements with fixed points in $\mathbb{R}^3_+$). If we take any point $z \in \mathbb{R}^3_+$ the orbit of z under G is a discrete sequence that only accumulates on S^2. The accumulation set is denoted Λ and is called the limit set of the group G (and is independent of the base point z, except for a few trivial cases). Moreover, Λ is mapped to itself by every element of G. The complement $\Omega = S^2 \setminus \Lambda$ is called the ordinary set. In the case when Λ is connected, Ω is a union of simply connected components and Ω/G is a union of Riemann surfaces and are interpreted as the "boundary at infinity" for the manifold M. Note that the convex hull $C(\Lambda) \subset \mathbb{R}^3_+$ must also be G-invariant, so the quotient $C(M) = C(\Lambda)/G$ is a well defined subset of $M = \mathbb{R}^3_+/G$. It is called the convex core of M and contains all the essential topology of M (e.g., every homotopy class has a representative loop in $C(M)$). The boundary components of $C(M)$ in M corresponds to the quotients of the domes for each complementary component of the limit set Λ. The Sullivan-Epstein-Marden theorem says that for a hyperbolic 3-manifold, the boundary of the convex core is homeomorphic (indeed, biLipschitz equivalent to) the boundary component at ∞ of M.

Epstein and Marden's proof of Theorem 6 in [43] gives biLipschitz constant $K_0 \approx 88.2$ and quasiconformal constant $K \approx 82.6$ and they conjectured $K_0 = K = 2$ is correct. In [17] it is proven that one can take $K = 7.82$ and $K_0 = 13.3$, but the construction there is not group invariant. In [46], Epstein, Marden and Markovic showed that $K > 2$ for the group invariant case using a result of McMullen's on the Teichmüller space of a punctured torus. See [44] for a different proof of this. More recently Epstein and Markovic [47] have shown that $K > 2.1$ in the non-invariant case, by showing the K is at least this large when Ω is the complement of a certain logarithmic spiral. Approximation theorems (also given in [47]) imply that there are polygons for which the constant is also at least 2.1. Most recently Epstein, Marden and Markovic [45] have lowered the group invariant constant to $K = 13.88$ and given a new, shorter proof of Theorem 6. It would be of great interest to determine even sharper bounds for the constant K in the convex hull theorem (we will discuss some possible applications in Section 12).

One way to prove Theorem 6 is to consider the nearest point retraction from Ω to S_Ω. Since S_Ω is the boundary of a convex set in hyperbolic space, each point outside this set has a nearest point on S_Ω and this defines a Lipschitz retraction r onto S_Ω. This map can be extended continuously to a mapping $r : \Omega \to S_\Omega$. A more geometric description of the same map is to take a point $z \in \Omega$ and consider the family of balls in $\mathbb{R}^3_+$ which are tangent to $\mathbb{R}^2$ at z. Expand the ball until it first makes contact with S_Ω; by convexity one can show there is a unique point of contact and this defines $r(z)$. It is also not hard to see that

$$\mathrm{dist}(r(z), \mathbb{R}^2) \simeq \mathrm{dist}(z, \partial),$$

with constants that are independent of z and Ω (see [14] or Lemma 18 of this paper). Moreover, $r : \Omega \to S_\Omega$ extends to the identity mapping on $\partial\Omega$.

If $r : \Omega \to S_\Omega$ was quasiconformal with a uniform constant then Theorem 6 would be proven, but this is not true (it is true in some cases, e.g., if Ω is Euclidean convex then r is 2-quasiconformal, [44]). In general, r need not even be a homeomorphism; it is possible for an arc of points in Ω to be collapsed to a single point on S_Ω. This happens if Ω is the union of two overlapping (but distinct) disks. The surface S is the upper envelope of the two hemispheres with these disks as bases and each point on the intersection of the two round domes has a whole arc of points in the plane that retracts to it. See Figure 5 for a side view when the disks are $D(0,1), D(1,1)$. However, r is always a quasi-isometry from Ω with its hyperbolic metric to S_Ω with

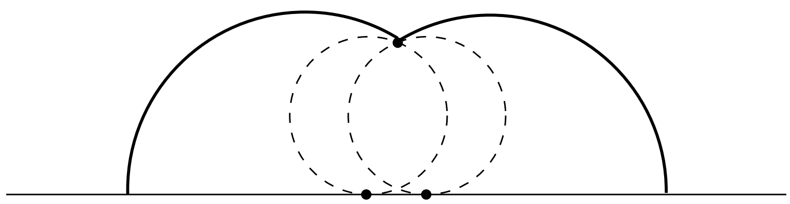

FIGURE 5. Two points can retract to same point on the dome

its hyperbolic path metric. This is Lemma 7.14 of [47] and is also proven in [14]. The paper [46] shows that $r : \Omega \to S_\Omega$ is 2-Lipschitz with respect to the hyperbolic metrics. [47] also contains basic facts about rough isometries and other useful results.

We can now deduce Theorem 6 using known results about quasi-isometries and quasiconformal maps. Suppose $f : \mathbb{D} \to \Omega$ is conformal. Since f and ι are isometries, $h = \iota \circ r \circ f$ is a rough isometry from the disk to the disk with the same constants as r. By Väisälä's theorem described in Section 3, this implies there is a K-quasiconformal self-map H of the disk with the same boundary values, and with K depending only on quasi-isometry constants of r. Thus $\sigma = \iota^{-1} \circ H \circ f^{-1}$ is the desired quasiconformal map from Ω to S_Ω.

This proof does not give a good estimate for K, but is simple to understand. The proofs of Theorem 6 in [43], [17] which give explicit constants involve taking averages of the retraction map $r : \Omega \to S_\Omega$ in order to make it a homeomorphism. The proof in [45] uses Teichmüller theory.

5. Finitely bent domains

When Ω is a finite union of disks, the dome $S = S_\Omega$ and the map $\iota : S \to \mathbb{D}$ have a particularly simple form. In this case the surface S is a finite union of geodesic faces. Two of these can meet along infinite geodesic at some angle (called the bending angle for the geodesic). Such a surface is called finitely bent and we will also refer to Ω as a finitely bent domain.

The map ι is an isometry from each face in S_Ω to a hyperbolic polygon in $\mathbb{D}$, each edge of which is an infinite geodesic. Theorems 8.13 and 8.14 of [44] show that we can approximate a general simply connected domain by finitely bent domains so that the corresponding maps σ and ι converge to the correct limits. Thus it often suffices to only consider the finitely bent case.

A crescent is a simply connected planar domain bounded by two circular arcs which meet at two distinct points (called the vertices). The angle θ of a crescent is the interior angle at these vertices. To any crescent and a choice of an edge, we associate the elliptic Möbius transformation that has the two vertices a, b as fix points and maps the given edge to the other edge. This is given by

$$\tau_{a,b,\theta}(z) = \frac{(be^{i\theta} - a)z + ab(1 - e^{i\theta})}{(e^{i\theta} - 1)z + (b - ae^{i\theta})}. \tag{4}$$

The formula can be easily derived by mapping a and b to 0 and ∞ by the map $w = (z - a)/(z - b)$, then multiplying by $e^{i\theta}$ and then applying the inverse map $z = (bw - a)/(w - 1)$. Each crescent is foliated by circular arcs orthogonal to both boundary arcs. See Figure 6. The elliptic map above can be interpreted geometrically by following the orthogonal foliation from one edge to the other.

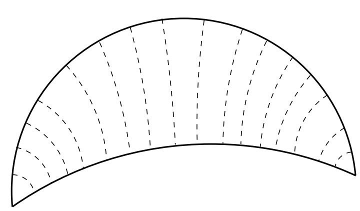

Figure 6. The orthogonal foliation of a crescent

If Ω is a disk, then S_Ω is a hyperbolic halfplane and a simple calculation shows that the retraction map $r : \Omega \to S_\Omega$ is an isometry. More generally, if Ω is a finite union of disks, then S_Ω is a finite union of geodesic faces $\{F_k\}$, each of which lies on a hyperbolic plane (i.e., a Euclidean hemisphere orthogonal to $\mathbb{R}^2$), which meets $\mathbb{R}^2$ along a Euclidean circle C_k whose interior D_k lies inside Ω. (These circles are the circumcircles for a Delaunay triangulation of the polygon formed by joining the vertices of Ω; see Section 8).

Two faces, F_k and F_j, can meet only along an infinite geodesic γ_{kj}, and the hyperbolic planes containing them meet at some angle $\pi + \alpha$. This is the same as the angle at which D_k and D_j meet. The number α is called the bending angle (or sometimes the bending measure) of the geodesic γ_{kj} (if it is zero, then both faces lie in the same plane; if $\alpha < 0$ the surface would not bound a convex region). By triangulating, we may (and will) assume every face is either an ideal hyperbolic triangle bounded by three infinite geodesics or a hyperbolic half-plane, bounded by a single infinite geodesic. See Figure 7.

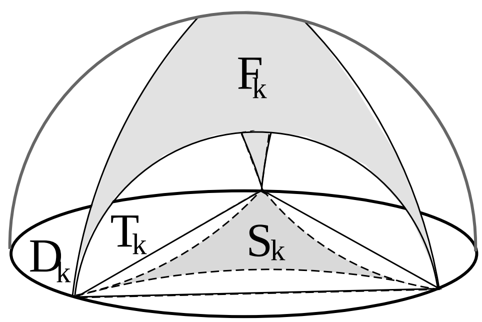

FIGURE 7. Geodesic face projects to triangles.

The retraction map $r : \Omega \to S_\Omega$, maps a hyperbolic polygon (either a hyperbolic triangle or half-plane) $S_k \subset D_k$ isometrically to F_k. If the faces F_k and F_j are adjacent in the dome along a geodesic γ_{kj}, then the corresponding polygons S_k and S_j are separated by a crescent C_{kj} which is mapped to γ_{kj} by r. Since the two circles bounding C_{kj} are orthogonal to C_k and C_j respectively, we see that the crescent C_{kj} has angle α at each of it vertices. See Figure 8. Each leaf of the orthogonal foliation of C_{kj} is collapsed to a single point by r and this is exactly what prevents r from being a homeomorphism. Note that all the crescents are disjoint since they project to disjoint sets under the retraction map to the dome.

We can now see how to map S_Ω to the disk isometrically in the finitely bent case. Since we will use almost exactly the same construction later we will describe it in general.

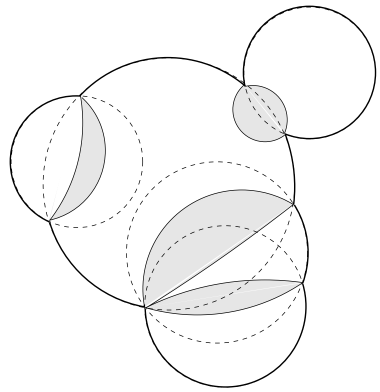

FIGURE 8. Crescents in a finitely bent domain.

Suppose we have a finite collection $\mathcal{D}$ of disks in the plane and an adjacency relation between them that makes the collection into the vertices of a tree. Suppose the disk D_0 has been designated the root of the tree. Then any other disk D has a unique "parent" D^* which is adjacent to D but closer to the root. Assume that for every (non-root) disk we are given a map $\tau_D : D \to D^*$. The we can define a map $\sigma_D : D \to D_0$ as follows. If $D = D_0$, the map is the identity. Otherwise, there is a unique shortest path of disks $D_0, \dots, D_k = D$ between D_0 and D. Note that each disk is preceded by its parent. Thus $\sigma = \tau_{D_1} \circ \cdots \circ \tau_{D_k}$ is a mapping from D to D_0 as desired. Later we will refer to this as a "tree-of-disks" map.

Lemma 7. *With notation as above, assume that every map τ_D is Möbius. Then given n points $\mathbf{v} = \{v_1, \dots, v_n\}$, with $v_k \in \partial D_k$ for $k \in \{1, \dots, n\}$, we can compute the n image points $\sigma(\mathbf{v})$ in at most $O(n)$ steps.*

Proof. If $D \in \mathcal{D}$ has positive radius, choose three distinct reference points z_1^D, z_2^D, z_3^D on ∂D; otherwise let this collection be empty. Every other point z on ∂D is uniquely determined by $\mathrm{cr}(z_1^D, z_2^D, z_3^D, z)$. Label each point v in $\mathbf{v}$ with the minimal k so that v is on the boundary of a kth generation disk D. For $k = 0$ we do nothing to the vertex. For $k > 0$, compute $\tau_D(v) \in \partial D^*$ and record the cross ratio $\mathrm{cr}(z_1^{D^*}, z_2^{D^*}, z_3^{D^*}, \tau_D(v))$. Also compute and record the images of the three reference points for D, i.e., $\mathrm{cr}(z_1^{D^*}, z_2^{D^*}, z_3^{D^*}, \tau_D(z_k^D))$ for $k = 1, 2, 3$.

If a vertex is on ∂D_0 then it maps to itself. If D is a first generation disk, then we just compute $\tau_D(v) \in \partial D_0$ and compute the τ_D images of the three reference points for D. For each child of D' of D, we can now compute $\sigma_{D'}$ for any associated vertices using the previously recorded cross ratios with respect to the reference points for D and we can also compute the images on ∂D_0 for the references points for D'. In general, if D is a disk and we have already computed where the reference points for its parent are mapped on ∂D_0, we can use the recorded cross ratio information to compute where the associated vertices and reference points for D map to. This allows us to map every point of $\mathbf{v}$ to ∂D_0 in $O(n)$ steps. □

We will apply this when Ω is a finitely bent domain and $\{D_k\}$ is the collection of bases of the geodesic hyperbolic half-planes which contain the faces of its dome S_Ω. Note that $\Omega = \cup_k D_k$, but that sometimes Ω can also be written as a union of disks in a different way. For example in Figure 9 we see a domain which is a union of three circles (dashed) but the dome also includes the hemisphere above a fourth circle (dotted).

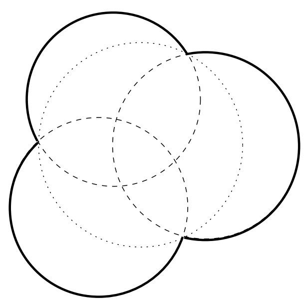

FIGURE 9. Ω is a union of 3 disks but its dome has 4 geodesic faces

Two disks will be considered adjacent if their corresponding faces on the dome are adjacent. This implies that two adjacent disks must intersect in two distinct points and our map $\tau_D : D \to D^*$ will be the unique elliptic Möbius transformation τ_{kj} which fixes both points of intersection and maps D to D^* (this is conjugate to a rotation by α, the bending angle of the corresponding geodesic). Note that if C_{kj} is the crescent separating $D = D_k$ and $D^* = D_j$ then the map τ_{kj} maps one edge of the crescent to the other edge. Applying the tree-of-disks construction above gives us a map $\iota = \iota_{\Omega,D_0} : \partial\Omega \to \partial D_0$.

We would like to extend this to the interior of Ω. The procedure above does not quite work since an interior point can be in more than one disk and the definitions for different disks need not agree (as they do on the boundary). However, we can define an extension to the interior as follows. Suppose D_k is the parent of D_j. First define a map $\phi_k : D_k \cup D_j \to D_j$ as follows. The domain $D_j \cup D_k$ is cut into two crescents by removing C_{kj}. Take the identity map on one of these (the one with an edge in ∂D_j), the map τ_{kj} on the other and interpolate between these by the map collapsing the leafs of the orthogonal foliation of C_{kj} to points. Restricted to D_j this gives a map of D_j into D_k. When we apply the tree-of-disks construction with these maps we get a map $\phi : \Omega \to D_0$ that agrees with ι on the boundary, is continuous on the closure of Ω and is Möbius on each component of $\Omega \setminus \cup C_{kj}$.

Thus $\phi \circ r^{-1}$ is a well defined continuous map of S_Ω to D_0 which is isometric on each face F_k. It is easy to see that such a map is an isometry from S_Ω to D_0 as follows. Take any two points on S_Ω and connect them by a geodesic γ on S_Ω. The image of γ is curve of equal length (but perhaps not a geodesic), so the images of z and w are no father apart than z and w are. The same argument works for the inverse map, so distances are preserved.

If we restrict the map to $\partial\Omega = \partial S_\Omega$, we have defined ι as a composition of elliptic Möbius transformations on each circular arc in $\partial\Omega$. Note that the crescents that we use are always of the form $W = D_2 \setminus D_1$ and that we are mapping the edge $\partial W \cap \partial D_2$ to the edge $\partial W \cap \partial D_1$. Thus we are in the case of the following lemma.

Lemma 8. *Suppose Ω is a crescent which lies on one side of the line L passing through the two vertices. Let γ_1, γ_2 be the circular arcs in $\partial\Omega$ with γ_1 between γ_2 and L. If τ is an elliptic Möbius transformation fixing the two vertices and mapping γ_2 to γ_1 then $|\tau'(z)| \leq 1$ on γ_2.*

Proof. To see this suppose $\tau(z) = (az+b)/(cz+d)$ where $ad - bc = 1$ (which we can always assume by normalizing). Then a simple calculation shows $|\tau'(z)| < 1$ iff $|1/c| < |z + d/c|$. Note that $-d/c = \tau^{-1}(\infty)$. By normalizing by a Euclidean similarity, we may assume the vertices are 1 and -1 and the crescent lies in the upper half-plane. See Figure 10. Then $-d/c$ is on the negative imaginary axis and $|\tau'(z)| < 1$ outside a circle C centered at $-c/d$ passing through -1 and 1. Let γ be the arc of this circle between 1 and -1 which lies in the upper half-plane. We claim that the upper edge of our crescent lies above γ. Since the circle C contains $-c/d = \tau^{-1}(\infty)$, its τ image contains ∞ and hence $\tau(\gamma)$ must lie below the real axis. But the $\tau(\gamma_2) = \gamma_1$ is above the real axis, so γ_2 is above γ and hence contained in the region where $|\tau'| < 1$. □

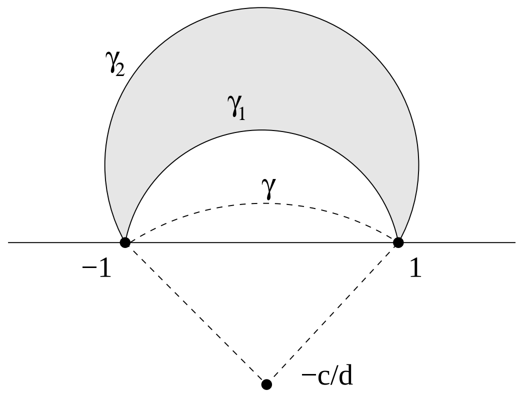

FIGURE 10. The setup in Lemma 8

This implies that the map $\iota : \partial\Omega \to \partial D_0$ can be chosen to have derivative < 1 on $\partial\Omega$ (except possibly at the vertices). It is proven in [16] that the quasiconformal extension of ι to the interior of Ω given by the Sullivan-Epstein-Marden theorem can be chosen with derivative bounded by M independent of Ω (but I don't know if $M = 1$ is possible).

We can now record several useful facts, by passing from finitely bent domains to general simply connected domains.

Lemma 9. *If Ω is simply connected and contains $\mathbb{D}$, then we can choose ι so that $\iota : \partial\Omega \to \mathbb{T}$ does not increase the length of any boundary arc.*

Proof. We can choose a sequence of approximating finitely bent domains and pass to the limit. □

Lemma 10. *If Ω is simply connected and contains a disk D and there is an arc $\gamma \subset \partial D \cap \partial\Omega$, then $\iota : \partial\Omega \to \mathbb{T}$ is a Möbius transformation when restricted to γ.*

Proof. We can approximate Ω using finite unions of disks which always include D. Then γ will be on the boundary of each approximating domain and so the approximating map will be Möbius on γ by our remarks above. Since a (non-degenerate) limit of Möbius transformations is Möbius, we are done. □

Given a domain Ω and an open disk $D \subset \Omega$ whose boundary hits $\partial\Omega$ in at least two points, we can define $\iota : \partial\Omega \to \partial D$ as above which is the identity on $\partial\Omega \cap \partial D$. We will denote this map by $\iota_{\Omega,D}$.

Lemma 11. *Suppose Ω is simply connected and contains a disk D so that $\partial D \cap \partial \Omega$ contains the endpoints of an arc $\gamma \subset \partial\Omega$ which is separated from the root disk D_0 in Ω by D. Let Ω' be the union of D and the component of $\Omega \setminus D$ intersecting D_0. Then on γ, $\iota_{\Omega,D_0} = \iota_{\Omega,D} \circ \iota_{\Omega',D_0}$.*

Proof. As above, one approximates by finitely bent domains which include D as one of the disks. For these the result is obvious and one passes to a limit as before. □

The point of stating this lemma is that when we go to compute ι for a polygonal domain Ω we will define a nested chain of domains $D_0 \subset \Omega_1 \subset \cdots \subset \Omega_k = \Omega$ where each inclusion is of the form in Lemma 11. Using Lemmas 10 and 11 we can then compute ι as a composition of Möbius transformations, each of which can be computed using the local geometry.

6. The medial axis of a polygon

We will start by describing the bending lamination of the dome of Ω and see how this leads to the medial axis.

When S_Ω, the dome of Ω, is finitely bent then every point on the dome is either in a geodesic face of S_Ω or it is on bending geodesic. The same fact is true for arbitrary simply connected Ω, except that there can be infinitely many bending lines (indeed, the whole surface could be a union of such lines, with no geodesic faces). The following is well known in the area, but we repeat it for the readers convenience.

Lemma 12. *Suppose S_Ω is the dome of a simply connected, hyperbolic plane domain Ω. Then for every $x \in S_\Omega$ there is a open hyperbolic half-space H disjoint from S_Ω so that $x \in \partial H \cap S_\Omega$. For any such half-space, $\partial H \cap S_\Omega$ contains an infinite geodesic.*

Proof. Let $E = C(\Omega^c)$ be the hyperbolic convex hull of Ω^c. Recall that $S_\Omega = \partial E$. By definition, E is the intersection of all closed half-spaces which contain it, and from this it is easy to see that any boundary point on E is on the boundary of some closed half-space which contains E. Thus x is also on the boundary of the complementary open half-space H (which must be disjoint from E). The base of H on $\mathbb{R}^2$ is a half-plane or a disk and by conjugating by a Möbius transformation, if necessary, we assume it is the unit disk $D = \mathbb{D}$ and that H contains the point $z = (0,0,1)\S_D$. Clearly ∂D hits $\partial\Omega$ in at least one point, for otherwise its closure would be contained in another open disk in Ω. Then the dome of this larger disk would separate ∂H from $z \in E$, which is impossible. In fact, ∂D must hit $\partial\Omega$ in at least two points. For suppose it only hit at one point, say $(1,0) \in \mathbb{R}^2$. Then we could choose a second disk D' which intersects D at the two points $(1/\sqrt{2}, \pm 1/\sqrt{2})$ and so that $D' \subset \Omega$. Then z lies beneath the dome of D'. Again, this is impossible since $z \in S_\Omega$. Thus ∂D hits $\partial\Omega$ in at least two points and the geodesic in $\mathbb{R}^3_+$ between these points lies on the $\partial H \cap S_D$, as desired. □

The boundary of the half-planes given by the lemma are called "support planes" for the dome. Every point of S_Ω is contained in a set of the form $\partial H \cap S_\Omega$ which is either (1) a convex subset of a hyperbolic plane bounded by infinite geodesics or (2) an infinite geodesic. Sets of the first type are called "flats" on S_Ω and the geodesics are called the "bending geodesics" of the surface. The union of these is called the "bending lamination" of the dome. In general, a lamination of a surface is a closed set which is a disjoint union of geodesics.

The bending lamination is also a measured lamination. For details, see [43], [100], but roughly speaking a transverse measure on a lamination is a function that assigns a non-negative number to each arc which is transverse to the leafs of the lamination and obeys the obvious addition law. See [13], [21], [77], [93], [94], (however, most of these deal with measured laminations when there is group action present, which is not our case). The transverse measure for bending lamination gives the amount of bending encountered by the arc as it crosses the bending geodesics (see [43] for the precise definition). It is easy to see what is going on in the finitely bent case and the general case is obtained by limits. Using the isometry ι, the bending lamination on S_Ω can be transfered to a measured lamination on the unit disk. The set of leafs, together with the transverse measure determine the original domain Ω (up to Möbius equivalence). It is an open (and probably difficult problem) to decide which measured laminations on the disk correspond to planar domains in this way, but there are some necessary and some sufficient conditions in terms of a certain norm on the measure. See [45] and its references.

An $\mathbb{R}$-tree is a metric space in which any two points are connected by a unique arc and this arc is isometric to a segment in $\mathbb{R}$. Any finite tree in graph theory is a

$\mathbb{R}$-tree, but so is $\mathbb{R}^2$ if we define $d(x, y) = |x - y|$ if x, y lie on the same ray from the origin and $|x| + |y|$ otherwise.

The support planes of the dome of a simply connected domain have the structure of an $\mathbb{R}$-tree. The support planes associated to a single bending geodesic form a 1-parameter family which is parameterized by the angle between them. More generally, the distance between support planes can be defined using the minimal bending measure of an arc connecting the associated bending geodesics. This is a special case of a general construction: every measured lamination has a dual object which is an $\mathbb{R}$-tree. See [13], [77].

If Ω is a simply connected planar domain, we say that a point $z \in \Omega$ is in the medial axis of Ω, $\text{MA}(\Omega)$, if z is the center of a disk $D \subset \Omega$ such that ∂D hits $\partial\Omega$ in two or more points. For $z \in \Omega$ let $d(z)$ be the distance of z to $\partial\Omega$. The medial axis, together with the distance function for points in the medial axis is called the medial axis transformation of Ω and denoted $\text{MAT}(\Omega)$.

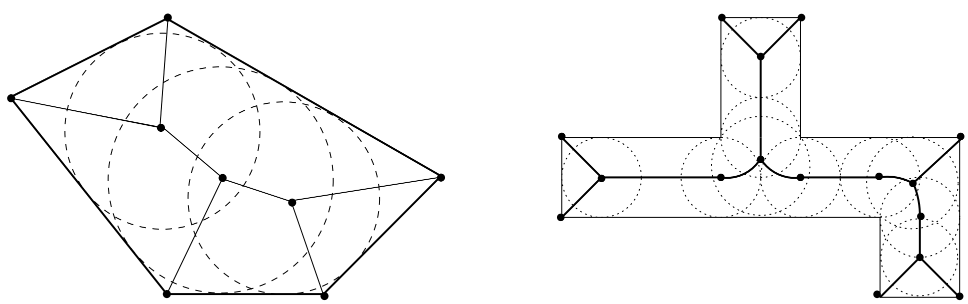

FIGURE 11. Examples of medial axes of polygons

Note that points of the medial axis are in 1-to-1 correspondence with support planes for the dome of Ω, and hence we can define a distance on the medial axis which makes it into an $\mathbb{R}$-tree. We will not use this observation here, but perhaps it will be useful for other problems (e.g., explaining how the medial axis changes with perturbations of the domain). So far as I know, the connection to bending laminations is new, but there have been previous attempts to better understand the medial axis using hyperbolic geometry [33], [34] (not hyperbolic geometry as used in this paper, but a space-time version where one direction contributes negative distances).

The first thing we should check is that the domain Ω can be recovered from the medial axis.

Lemma 13. *With notation as above,* $\Omega = \cup_{z \in \mathrm{MA}(\Omega)} D(z, d(z))$.

This follows from:

Corollary 14. *If* Ω *is simply connected then every* $z \in \Omega$ *is contained in a disk* $D \subset \Omega$ *so that* ∂D *hits* $\partial\Omega$ *in at least two points.*

Proof. Conjugating by a Möbius transformation, if necessary, we may assume $\infty \in \Omega^c$. Given $z \in \Omega$ let $r(z) \in S_\Omega$ be the nearest point retraction onto the dome. As we discussed before, $r(z)$ is the point of contact of a horoball B in $\mathbb{R}^3_+$ which is tangent to $\mathbb{R}^2$ at z and with interior disjoint from E. Let H be the unique open hyperbolic half-space which contains B and is tangent to it $r(z)$. Then the base of H is a disk which contains z. Moreover, H is disjoint from E, for if it contained a point w of E then by convexity E would contain the geodesic segment from w to $r(z)$, which would have to intersect B, a contradiction. Thus H is a halfspace which is disjoint from E, but contains a point of $S_\Omega = \partial E$ on its boundary. By Lemma 12, the base of H is a disk in Ω whose boundary hits $\partial\Omega$ in at least two points. □

The medial axis is a one dimensional object that has been used in computer science to encode the shape of 2-dimensional objects. For example, given the medial axis of a domain other problems can be quickly solved, e.g., finding the largest inscribed circle, computing the closest boundary point, and determining the region of all points within ϵ of the boundary. For domains whose boundary is a finite union of analytic arcs the medial axis is a finite tree, but this can fail even for C^∞ boundaries [31]. The medial axis can also be highly unstable under perturbations, e.g., is a point if Ω is a circle but has n arms of unit length if Ω is a regular n-gon inscribed in the unit circle.

The medial axis was introduced by Blum to describe biological shapes [18], [19], [20]. A few papers consider its mathematical properties (e.g., [31], [32], [33], [34], [35], [97], [107]), but most deal with algorithms for computing it and with applications to areas like pattern recognition, robotic motion, classification of biological shapes, control of cutting tools, sphere packing and mesh generation. A sample of such papers includes: [28], [36], [48], [53], [55], [58], [59], [60], [65], [69], [70], [71], [73], [74], [82], [86], [87], [91], [95], [96], [106], [108], [109].

The medial axis of a polygon is closely related to other important concepts in computational geometry. For example, given a finite collection of disjoint closed sets (called sites), the corresponding Voronoi diagram divides the plane according to which site a point is closest to. The medial axis of P is a Voronoi diagram for the interior of P where the sites are he complementary arcs in P of the convex vertices and distance is measured within P. Thus the medial axis can be obtained from algorithms for computing generalized Voronoi diagrams. Voronoi diagrams were defined by Voronoi in [105], but go back at least to Dirichlet [39] (indeed, in the theory of Kleinian groups the Voronoi cells of an orbit are called Dirichlet fundamental domains). For more about Voronoi diagrams see e.g., [4], [5], [10], [49], [50], [80], [88].

Before we use the medial axis to construct ι we will review its structure a little more carefully. Divide the the vertices of P into three collections $\mathcal{V}_0, \mathcal{V}_1, \mathcal{V}_2$ (also known as convex, flat and concave), depending on whether the interior angle is $< \pi$, $= \pi$ or $> \pi$. Some sources also refer to the convex vertices as "sharp" and the concave vertices as "dull", e.g., [34].

For each vertex $v \in \mathcal{V}_1$ (angle $= \pi$) take the largest disk in Ω which is tangent to the boundary at v. Call these type 1 disks. Since Ω is a polygon there is a largest such disk and its boundary hits $\partial\Omega$ in at least one point of $\partial\Omega$ other than v.

If $v \in \mathcal{V}_2$ (angle $> \pi$) we add two disks. For each edge e adjacent to v we take the largest disk in Ω which is tangent to the line containing e at v. Call these type 1 disks. As before, such a disk exists and its boundary hits the boundary of Ω somewhere other than v.

A type 0 disk is a point; i.e., the disk of radius 0 centered at a vertex in $\mathcal{V}_0$. The three collections of disks will be denoted $\mathcal{D}_0, \mathcal{D}_1, \mathcal{D}_2$. Note that the latter two classes may overlap. However, we assume that duplicates from within each class have been removed (there will be no overlaps or duplications if the vertices of P are in general position). We let $\mathcal{D}_3$ be the collection of open disks in Ω such that $\partial D \cap \partial\Omega$ has three or more points. Some of these may be in $\mathcal{D}_1 \cup \mathcal{D}_2$ but there could be others are not. Let $\mathcal{D} = \mathcal{D}_0 \cup \mathcal{D}_1 \cup \mathcal{D}_2 \cup \mathcal{D}_3$.

Since we are dealing with a polygonal region it is easy to see $\mathcal{D}_3$ is a finite collection, but we wish to have an explicit bound. Let n_k, $k = 0, 1, 2$ be the number or type k vertices in P. Clearly $n_0 + n_1 + n_2 = n$. Note $\#(\mathcal{D}_1) \leq n_1$ and $\#(\mathcal{D}_2) \leq 2n_2$. The centers of our disks are the vertices of the medial axis (which is a tree) and each disk D corresponds to a vertex of degree $\#(\partial D \cap P)$ in the medial axis.

Lemma 15. *Let* $n_3 = \sum_{D \in \mathcal{D}_3} (\deg(D) - 2)$. *Then* $n_3 = n_0 - 2$.

Proof. Given any tree T with n_0 vertices of degree 1, if we define n_3 as above then we can find another tree T' that has n_0 degree 1 vertices and n_3 degree 3 vertices and no vertices of any other degree (a degree 2 vertices and its two adjacent edges are identified to a single edge; vertices of degree $d > 3$ are "expanded" to a path of length $d - 3$ with old edges attached to the new vertices in order to make all the degrees 3. See Figure 12). In any tree with only degree 1 and degree 3 vertices we have $n_0 + 3n_3 = 2(n_0 + n_3 - 1)$, (since the sum of the degree is twice the number of edges). Thus $n_3 = n_0 - 2$, as desired. □

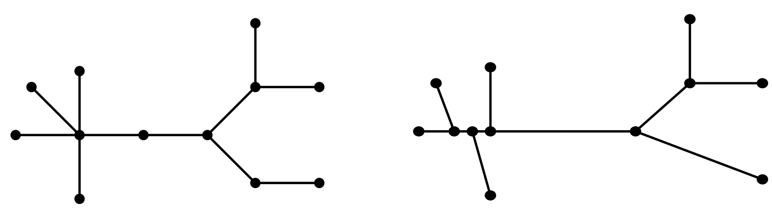

FIGURE 12. Converting a tree to all degree 1 or 3

Thus the medial axis of a simple n-gon can have at most $n_0 + n_1 + 2n_2 + (n_0 - 2) \leq 2n - 2$ vertices.

Suppose Ω is bounded by a simple polygon P and D is a disk in Ω. Then D is maximal iff ∂D hits $\partial\Omega$ in two or more points (this is not true in general, e.g., Ω is the interior of a parabola, but is true for polygons). If $D \subset \Omega$ is maximal but not in $\mathcal{D}$ then it is easy to check that D is part of a non-degenerate one-parameter family of maximal disks whose centers form an arc of one of the following three types:

1. a line segment that is equidistant from two vertices
2. a line segment that is equidistant from two lines
3. a parabolic arc that is equidistant from a line and a vertex.

These cases are illustrated in Figure 13. In that figure we have split case two into the subcases of parallel and non-parallel lines since the formulas for ι will look slightly different in these cases. Thus the medial axis is a finite tree whose vertices are the center if disks in $\mathcal{D}$ and whose edges are as above.

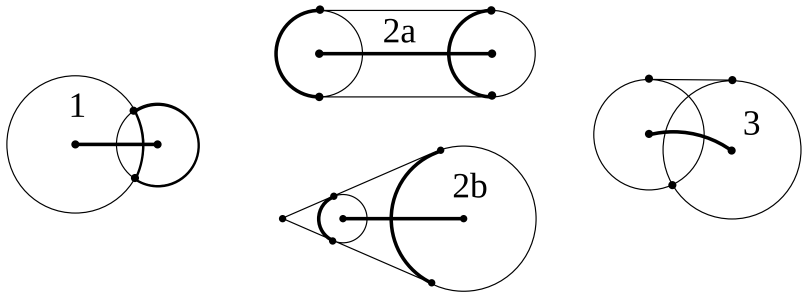

FIGURE 13. The different types of medial axis edges

It is a theorem of Chin, Snoeyink and Wang [29], [30] that the medial axis of polygon with n vertices can be computed in $O(n)$ steps. More precisely, we will assume that the following information can be computed in time $O(n)$:

1. a list of vertices for the medial axis and the distance of each from the boundary of the polygon. We assume the list starts with a given root vertex and every vertex occurs later in the list than its parent,
2. for each vertex of the medial axis, a list of adjacent vertices and the type of edge connecting them (including the the sites (vertices or edges) the edge is equidistant from,
3. for each vertex of P, a vertex of the medial axis it is closest to.

7. CONSTRUCTING ι FROM THE MEDIAL AXIS

In this section we will construct ι for a polygon using the medial axis. The construction another "tree of disks" construction as in Section 5. The disks in our tree will be the maximal disks corresponding to vertices of the medial axis, i.e., the collection $\mathcal{D}$ described in Section 6. Recall that this collection includes the convex angles (interior angle $< \pi$) and that the corresponding disks have zero radius. The map of such a disk to its parent will be the closest point of the parent. For non-degenerate disks, the map to the parent will be described below. Then Lemma 7 and the theorem of Chin, Snoeyink and Wang imply $\iota(\mathbf{v})$ for a simple n-gon can be computed in $O(n)$ steps. Combined with the estimate $K \leq 7.82$ for the Sullivan-Epstein-Marden constant given in [17], this will prove Theorem 1.

All that remains is to describe how to map a disk $D_1 \in \mathcal{D}$ to its parent $D_2 \in \mathcal{D}$. The form of the transformation depends on the type of edge connecting the two vertices:

1. A line segment that is equidistant from two vertices,
2. A line segment that is equidistant from two lines,
3. A parabolic arc that is equidistant from a line and a vertex.

In what follows we split case 2 into the subcases of parallel and non-parallel lines since the formulas will look slightly different in these subcases.

Case 1 (equidistant from two points): This is the simplest of the cases. If the two disks are D_1 and D_2 with D_2 being the parent, then the desired map is just the unique elliptic Möbius transformation which fixes the two points, a, b of $\partial D_1 \cap \partial D_2$ and maps $\partial D_1 \setminus D_2$ onto $\partial D_2 \cap \overline{D_1}$. If the two boundary circles meet at angle α then this map is given by the formula $\tau(z) = \sigma^{-1}(e^{i\alpha}\sigma)$, where $\sigma(z) = (z-b)/(z-a)$ sends a to ∞ and b to zero. Simplifying gives

$$\tau(z) = \frac{(ae^{i\alpha} - b)z - ab(1 - e^{i\alpha})}{(e^{i\alpha} - 1)z + (a - e^{i\alpha}b)}.$$

Case 2a (equidistant from two parallel lines): Consider Figure 14. We normalize so that $D_1 = D(0,1)$ and $D_2 = D(A,1)$. Note that on the dome of Ω, the points above the centers of the disks D_1, D_2 are exactly distance A apart (the straight line between them is a geodesic on the dome with this length). Since ι is an isometry from the dome to the disk, it will map these points to points which are exactly hyperbolic distance A apart in the disk.

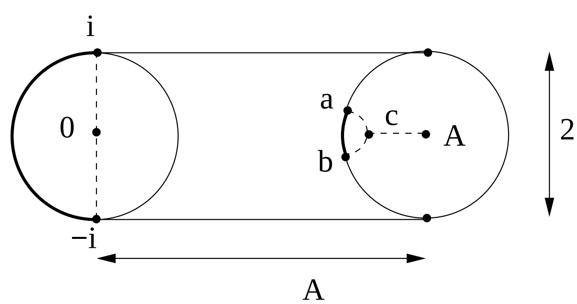

FIGURE 14. Case 2a: between parallel lines

With the points labeled as shown, i and $-i$ will map to a and b respectively and 0, the center of the disk D_1, will map to a point c which is hyperbolic distance A from the center of D_2. Thus the desired map is $\tau(z) = \sigma(z) + A$ where σ is the unique Möbius transformation which fixes 1 and -1 and maps 0 to the point $-x$ on the negative real axis at hyperbolic distance A from 0. Since the hyperbolic metric on the disk is given by

$$d\rho = \frac{2|dz|}{1-|z|^2},$$

the point w is the solution of

$$A = \int_0^w \frac{2}{1-r^2} dr = \log \frac{1+w}{1-w},$$

which is

$$w = \frac{1-e^{-A}}{1+e^{-A}}.$$

Thus

$$\sigma(z) = \frac{z-w}{1-wz},$$

and this determines τ.

Case 2b (equidistant from two non-parallel lines): The situation is shown in Figure 15. We normalize so that $D_1 = D(x, x\sin\alpha)$ and $D_2 = D(y, y\sin\alpha)$. The points a, b are mapped to c, d respectively. The points on the dome above x and y are joined by a geodesic which projects to the straight line between x and y. Above a point $t \in \mathbb{R}_+$ the dome has height $t\sin\alpha$ so the length of this geodesic on the dome is

$$A = \int_x^y \frac{dt}{t\sin\alpha} = \frac{1}{\sin\alpha} \log \frac{y}{x}.$$

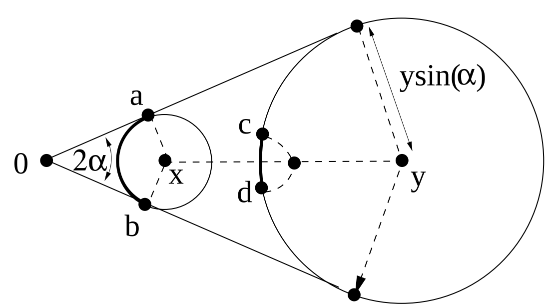

FIGURE 15. Case 2b: between non-parallel lines

Thus if $w = (1 - e^{-A})/(1 + e^{-A})$ and $\sigma(z) = \frac{z-w}{1-wz}$, the desired map $\tau : D_1 \to D_2$ is

$$\tau(z) = y(\sin\alpha)\sigma(\frac{1}{x\sin\alpha}(z-x)) + y.$$

Note that since $\sin\alpha = |a-x|/|x|$, there is no need to actually evaluate a trigonometric function, i.e., we could have written

$$\tau(z) = \frac{y|a-x|}{x}\sigma(\frac{z-x}{|a-x|}) + y.$$

Case 3 (equidistant from a point and a line): This case is pictured in Figure 16. Suppose a is the vertex in question and L is the line. Let a^* be the reflection of a across L. Then D_1 can be mapped to D_2 by an elliptic Möbius transformation which fixes a and a^* and sends L to itself. This elliptic element rotates around a by some angle θ. We will think of this as a composition of n separate rotations, each by angle θ/n and consider what happens. Applying each of these rotations in turn produces a sequence of disks $\{D_k\}$ intermediate between D_1 and D_2. Let $\tau_k : D_k \to D_{k+1}$ be the elliptic map that fixes the two points of intersection a, b_k on the boundary and $\sigma_k : D_k \to D_{k+1}$ be the elliptic rotation fixing a and a^*. Then $\tau_n \circ \cdots \circ \tau_1$ converges to the desired $\tau : D_1 \to D_2$ as $n \to \infty$. The map ι is invariant under compositions with

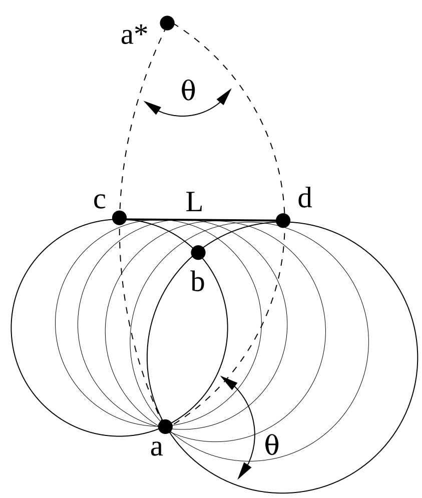

FIGURE 16. Case 3: equidistant from a point and a line.

linear fractional transformations, so we can conjugate by one to a situation where the elliptic rotation is easier to understand.

The map $\sigma(z) = -i(z-c)(a^*-a)/(z-a)(a^*-c)$ satisfies $\sigma(a) = \infty$, $\sigma(a^*) = -i$, $\sigma(c) = 0$ and maps D_1 to the upper half-plane. The point b_k maps to some point $-x$ on the negative real axis. Moreover, the line segment from $-i$ to $-x$ makes angle $\theta/(2n)$ with the vertical. Thus $-x = -\sin(\theta/(2n))$. Hence σ conjugates τ_k to a rotation by θ/n around $-x$ and σ_n conjugates to a rotation by θ/n around $-i$. Thus $\sigma_k^{-1} \circ \tau_k$ is an isometry of the plane which maps the real line into itself and maps the point $-x$ to x (the first maps fixes it and the second maps sends it to x by our comment about angles above). Thus the composition must be the map $z \to z + 2x$. For n large this is approximately $z \to z + \theta/n$. Thus if we iterate this map n times we get (in the limit) $z \to z + \theta$.

Thus $\sigma^{-1} \circ \tau$ is conjugate to the parabolic map $z \to z + \theta$. The actual map from D_1 to D_2 is

$$z \to \tau_{a,a^*,\theta}(\sigma^{-1}(\sigma(z) + \theta)),$$

where $\tau_{a,a^*,\theta}$ is the elliptic rotation of angle θ around with fixed points a, a^* as given by formula (4).

We have now completed the proof of Theorem 1.

8. The CRDT algorithm is an ι map

We now begin the second part of this paper, showing the connection between ι and the CRDT algorithm of Driscoll and Vavasis. We begin with a brief review of CRDT.

Suppose Ω is a simply connected planar domain which is bounded by a simply polygon P with vertices $\mathbf{v} = \{v_1, \dots, v_n\}$. Then Ω can be triangulated, i.e., written as a union of $n-2$ triangles with disjoint interiors and whose vertices are among the vertices of P. A triangulation is called Delaunay if whenever we form a quadrilateral Q by taking the union of two adjacent triangles T_1, T_2 from our triangulation which share a common edge e, the sum of the interior angles of Q at the two endpoints of e sum to $\geq \pi$. See Figure 17. More properly, this should be called a constrained Delaunay triangulation. A Delaunay triangulation of the vertices would actually be a triangulation of the plane which did not have to include the edges of P. We will only consider triangulations of the interior of P, and we force the edges of P to be included in the triangulation (which is why it is a constrained triangulation). However, we will continue to call these Delaunay triangulations since the other type are not considered here.

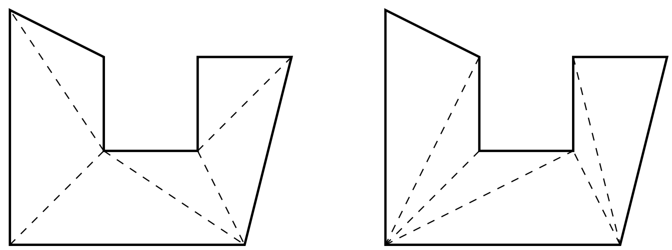

FIGURE 17. A Delaunay (on left) and non-Delaunay triangulation

It is well known that a Delaunay triangulation of a simple polygon always exists and is unique unless there are two adjacent triangles whose four vertices lie on a circle (in which case we can flip the common edge to join the other 2 vertices and get another Delaunay triangulation). Delaunay triangulations have many nice properties and have been intensively studied (e.g., they are dual to Voronoi diagrams; they maximize the minimum angle in the triangulation [68]; they minimize the largest circumcircle [38], [89]). The basic facts can be found in various sources such as [4], [10], [49], [50].

The CRDT algorithm is a method of finding the conformal prevertices of a given polygon P. The algorithm starts by adding new vertices of angle π to the edges of P to get a new polygon P' which has a Delaunay triangulation with certain nice properties. We will discuss how this is done in Section 10, and here we simply assume we are given P'. The second part of the algorithm uses a Delaunay triangulation of P' to produce guesses for the conformal prevertices of P' (and hence for P whose vertices are included among those of P'). It is this step we will reinterpret in terms of ι maps below. The third step of the CRDT algorithm is to solve for the actual prevertices using an iterative procedure that starts with the guess from the second step. We will not consider this part in detail except for a few general comments in Section 12.

How is the second step done? Suppose P is a simple n-gon and that $\mathcal{T} = \{T_1, \dots T_{n-2}\}$ is a Delaunay triangulation of P. Let D_k be the circumcircle associated to each triangle T_k, $k = 1, \dots, n-2$. The CRDT algorithm chooses points $\mathbf{w} = \{w_1, \dots w_n\} \subset \mathbb{T}$ corresponding to the vertices $\mathbf{v} = \{v_1, \dots, v_n\}$ as follows.

Choose some root triangle from the triangulation and map its vertices to any three points in $\mathbb{T}$ with the correct orientation. In general, suppose Q is the quadrilateral formed by two adjacent triangles T_1 and T_2 (which have vertices $v_{j_1}, v_{j_2}, v_{j_3}$ and $v_{j_1}, v_{j_3}, v_{j_4}$ in counterclockwise order respectively). Also suppose that we have already defined $w_{j_1}, w_{j_2}, w_{j_3}$. Then w_{j_4} is uniquely determined by the condition

$$\mathrm{cr}(w_{j_1}, w_{j_2}, w_{j_3}, w_{j_4}) = -|\mathrm{cr}(v_{j_1}, v_{j_2}, v_{j_3}, v_{j_4})|. \tag{5}$$

It is easy to see by induction that this uniquely determines the points $\mathbf{w}$ up to a Möbius transformation of the circle.

Next we will define a map $\iota : \mathbf{v} \to \mathbb{T}$ which gives the same points $\mathbf{w}$ up to a Möbius transformation. Given two adjacent triangles, the corresponding disks intersect at two points and we take the elliptic Möbius transformation as in (4) between them. This defines a "tree of disks" such as we considered in Sections 5 and 7 and so we can define a corresponding ι map by composing these maps along paths to the root. We claim $\iota(\mathbf{v})$ is the same as the initial guess of the CRDT algorithm. Suppose Q is the quadrilateral formed by T_k and its parent T_j (which have vertices v_1, v_3, v_4 and v_1, v_2, v_3 in counterclockwise order respectively). After conjugating by a Möbius transformation which sends $v_1 \to 0$, $v_2 \to 1$ and $v_3 \to \infty$, the elliptic map τ_k is conjugated to the Euclidean rotation around 0 which sends the image to v_4 onto the negative real axis. Thus

$$\mathrm{cr}(v_1, v_2, v_3, \tau_k(v_4)) = -|\mathrm{cr}(v_1, v_2, v_3, v_4)|.$$

The Möbius maps $\tau_1, \dots, \tau_{k-1}$ are then applied to all four points $v_1, v_2, v_3, \tau_k(v_4)$ and hence do not change the cross ratio. Thus

$$\mathrm{cr}(\iota(v_1), \iota(v_2), \iota(v_3), \iota(v_4)) = -|\mathrm{cr}(v_1, v_2, v_3, v_4)|.$$

Therefore $\mathbf{w} = \iota(\mathbf{v})$ is exactly the same as the CRDT guess for the prevertices.

If the union of the circumcircles $\{D_k\}$ is a planar simply connected domain Ω then Ω is finitely bent and the map we have just described is the ι map for Ω. In this case Theorem 2 and the modified Driscoll-Vavasis conjecture follow immediately. However, even if the union is not a simply connected planar domain (i.e., the union of disks surrounds one of the vertices of P as in the right side of Figure 18), we can interpret it as a simply connected Riemann surface. We can define this surface by induction, taking $R_0 = D_0$, and in general letting R_k be the surface obtained by attaching the crescent C_k to R_{k-1} along the edge $\gamma_k^* \subset \partial R_{k-1}$.

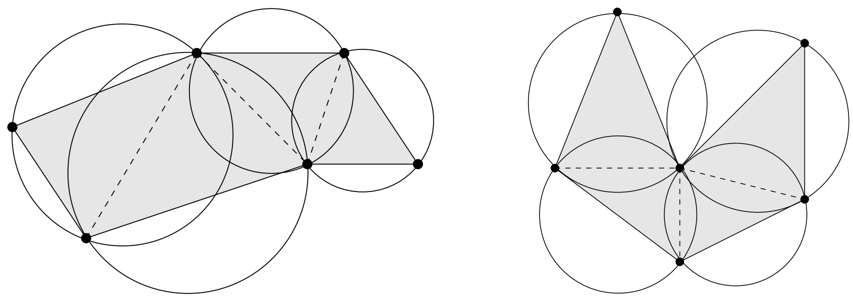

FIGURE 18. The union of circumcircles can be either planar or non-planar.

Thus the CRDT guess is an ι map, except that sometimes it corresponds to the boundary of a Riemann surface rather than a planar domain. We will prove in the next section that ι still has a uniform quasiconformal extension to the interior of this surface (but perhaps with a larger constant than in the planar case).

9. Proof of Theorem 3

In this section we will show that given a planar polygon P, the map defined in Section 8 has a K-quasiconformal extension to a map $\iota : R_P \to \mathbb{D}$, where K is bounded independent of n and P. Our proof is modeled on the proof using quasi-isometries in the planar case and will not give an explicit constant. Presumably the proofs giving explicit estimates in the planar case could be modified to give estimates in this case as well.

By the remarks on quasi-isometries and quasiconformal maps made in Section 4 it is enough to show the vertex map extends to a quasi-isometry from R_P with its hyperbolic metric to the disk with its hyperbolic metric. The first step is to show that the hyperbolic metric on R satisfies an estimate similar to that for planar domains. If Ω is a simply connected planar domain, then (e.g., [52], [85])

$$\frac{|dz|}{2\text{dist}(z, \partial\Omega)} \leq d\rho_\Omega \leq \frac{2|dz|}{\text{dist}(z, \partial\Omega)}.$$

Lemma 16. *There is an $M < \infty$ so that the following holds. If $R = R_P$ is the surface obtained from a Delaunay triangulation of a simple, planar polygon P, then*

$$\frac{|dz|}{M\text{dist}(z, \partial R)} \leq d\rho_R \leq \frac{2|dz|}{\text{dist}(z, \partial R)}.$$

Proof. The proof of the right hand inequality is exactly the same as the planar case: since R contains a disk D of radius $r = \text{dist}(z, \partial R)$ around z, the monotonicity property of hyperbolic distance [85] (also known as the Schwarz inequality) implies

$$d\rho_R(z) \leq d\rho_D(z) = \frac{2}{r}|dz|.$$

To prove the lower bound, it suffices to consider the disk of radius $r/2$ around z and show its hyperbolic diameter is bounded uniformly away from zero. This is equivalent to showing that the modulus of the path family connecting this disk to ∂R is bounded away from zero. First we need the following observation (which is the only place we use the special form of R).

Lemma 17. *Suppose R is obtained from the Delaunay triangulation of a simple polygon as above. If γ is a simple curve on R that projects into a circle $\{z : |z - x| = r\}$ then the projection is at most 3 to 1 (and hence γ has length $\leq 6\pi r$).*

Proof. R consists of P and a finite number crescents, one attached to each edge of P. If γ is a circle in R then its projection to the plane is 1-to-1. If γ is an arc in R which projects into a circle in the plane, then it leaves every crescent it enters, with at most two exceptions (the crescents containing it endpoints). If we remove the part of γ in these two exceptional crescents, the rest projects 1-to-1 to the plane because the projection of $P \subset R$ to $P \subset \mathbb{R}^2$ is 1-to-1. Thus the projection is at most 3-to-1. □

We can now return to estimating the modulus of the path family Γ connecting $D = D(z, \frac{1}{2}r) \subset R$ to ∂R. Let $v \in \partial R$ minimize $\text{dist}(z, \partial R)$ and let v also denote the projection of v to the plane. For $t \in [\frac{1}{2}r, \frac{3}{2}r]$, take a component γ_t of the lift of the circle $\{w : |w - v| = t\}$ which hits D. By Lemma 17 this arc has length $\leq 6\pi t$ and it must connect D to ∂R. Thus if ρ is an admissible metric for the path family Γ,

$\int_\gamma \rho|dz| \geq 1$ for all $\gamma \in \Gamma$ and hence

$$\begin{aligned}
\iint_R \rho^2 dxdy &\geq \int_{r/2}^{3r/2} \int_{\gamma_t} \rho^2 t dt d\theta \\
&\geq [\int_{r/2}^{3r/2} \int_{\gamma_t} t d\theta dt]^{-1} \int_{r/2}^{3r/2} \int_{\gamma_t} \rho^2 t d\theta dt \\
&\geq [6\pi \frac{3}{2} r]^{-2} \int_{r/2}^{3r/2} 1 t dt \\
&\geq [6\pi \frac{3}{2}]^{-2} r^{-2} \frac{1}{2} [\frac{9}{4} r^2 - \frac{1}{4} r^2] \\
&= (6\pi \frac{3}{2})^{-2}.
\end{aligned}$$

Thus the modulus of Γ is bounded away from zero uniformly, as desired. This completes the proof of Lemma 16. □

Suppose $R = R_P$ is the Riemann surface obtained from the simple polygon P and let $\mathcal{T}$ be a Delaunay triangulation of P. For each triangle $T_k \in \mathcal{T}$ let D_k be the corresponding disk determined by the three vertices, and let $S_k \subset T_k$ be the hyperbolic ideal triangle in D_k with the same vertices as T_k. For each edge e_j of P, suppose e_j is an edge of the triangle T_k and let U_k be the hyperbolic half-space in D_k bounded by the hyperbolic geodesic with same endpoints as e and disjoint from S_k. Let V_k be the dome of D_k (i.e., a hyperbolic plane in $\mathbb{R}^3_+$) and let F_k be the nearest point restriction of either a triangle S_k onto V_k or a half-plane U_k. Attaching the faces F_k according to the tree structure of the triangulation gives a surface S_R in $\mathbb{R}^3_+$ which we can think of as the dome of R_P, although it may have self-intersections, and no longer is the boundary of a convex region in $\mathbb{R}^3_+$. Nevertheless, the same proof as in the planar case shows that S_R is isomorphic to the hyperbolic disk. Moreover,

Lemma 18. *With notation as above,*

$$\text{dist}(z, \partial R_p) \leq \text{dist}(r(z), \mathbb{R}^2) \leq 2(\sqrt{2}-1)\text{dist}(z, \partial R_p).$$

Proof. First consider the case when $R_P = D$ is a disk in the plane. Without loss of generality we may assume it is the unit disk, z is on the segment between 0 and 1 and the line from 0 to $r(z) \in \mathbb{R}^3_+$ makes an angle $\theta \in (0, \pi/2)$ with the positive real axis. Then $\mathrm{dist}(r(z), \mathbb{R}^2) = \sin\theta$ and $\mathrm{dist}(z, \partial D) = 1 + \tan\theta - \sec\theta$ Thus

$$\frac{\mathrm{dist}(z, \partial D)}{\mathrm{dist}(r(z), \mathbb{R}^2)} = \frac{\cos\theta + \sin\theta - 1}{\sin\theta\cos\theta}.$$

A simple calculus exercise shows this is bounded above by 1 and below by $2^{-1/2} - 2^{-1}$.

Now given z in a general R_p, z lies in some disk D such that $r(z)$ lies on the dome of D. Thus the inequality of the previous paragraph holds and we only have to compare $\mathrm{dist}(z, \partial D)$ to $\mathrm{dist}(z, \partial R_P)$. Since $D \subset R_p$ we clearly have $\mathrm{dist}(z, \partial D) \leq \mathrm{dist}(z, \partial R_P)$. If $r(z)$ lies in a half-plane on the dome of R_p then we clearly have equality. Otherwise $r(z)$ lies in a geodesic triangle or on a bending line. In either case, z is separated from its radial projection onto ∂D be a geodesic in D, the endpoints of which are vertices of P. Sliding z along its radius towards the boundary until it hits the geodesic decreases the distance to the circle faster than the distance to the endpoints of the geodesic. Thus without loss of generality we may assume z is on a geodesic. Simple geometry shows that the ratio of the distance to the circle to the distance to the endpoints is at least $1/\sqrt{2}$. This means

$$\mathrm{dist}(z, \partial D) \geq \frac{1}{\sqrt{2}}\mathrm{dist}(z, \partial R_P),$$

which proves the lemma. □

If we map the S_k's and U_j's into $\mathbb{D}$ using $\tau_1 \circ \cdots \circ \tau_k$ (these are the maps in the definition of ι) and collapse the leaves of the crescents to points, then we get a continuous map ϕ of R into $\mathbb{D}$ which is 2-Lipschitz and such that $\phi \circ r^{-1} : S_R \to \mathbb{D}$ is well defined and an isometry. To complete the proof, we need only show $r : R \to S_R$ is a quasi-isometry. We already know r is Lipschitz, so we only need to prove the lower bound, i.e.,

$$\rho_{S_R}(r(z), r(w)) \geq A\rho_R(z, w) - B,$$

or equivalently,

$$\rho_R(z, w) \leq (A)^{-1}\rho_{S_R}(r(z), r(w)) + B.$$

We claim it suffices to show that there is a constant $b < \infty$, so that given $z, w \in R$ such that $\rho_{S_R}(r(z), r(w)) \leq 1$ implies $\rho_R(z, w) < b$. If this holds, then given two points z, w with $\rho_{S_R}(r(z), r(w)) = d > 1$ apart, we can connect them on S_r by a geodesic γ and choose $\leq 2 + d$ points on γ which are distance ≤ 1 apart (first equal to $r(z)$ and last equal to $r(w)$). Choosing one r-preimage for each intermediate point and using the triangle inequality gives

$$\rho_(z, w) \leq (2 + d)b = bd + 2b,$$

which is the rough isometry estimate with $A_1 = b$ and $B_1 = 2b$.

So suppose we are given $z, w \in R$ such that $\rho_{S_R}(r(z), r(w)) \leq 1$. The hyperbolic path distance between two points on S_R is longer (or equal) than the hyperbolic distance in $\mathbb{R}^3_+$ between them, so $r(w)$ lies in a radius 1 hyperbolic ball around $r(w)$ (in terms of $\mathbb{R}^3_+$). This implies

$$e^{-1} \leq \frac{h(z)}{h(w)} \leq e,$$

where $h(z) = \text{dist}(r(z), \mathbb{R}^2)$ (h is for "height"). It also implies that $|r(z) - r(w)| \leq eh(z)$ This, in turn, implies $|z - w| \leq Ch(z)$ for some constant C that doesn't depend on z (this is because $|z - r(z)| \leq Ch(z)$ is clear from the definition of r; similarly for w).

Consider the grid $\mathcal{G}$ of squares $\{Q_j\}$ of sidelength $Cr(z)$ and take the subcollection $\mathcal{C}$ such that $\text{dist}(z, Q_j) \leq Mh(z)$. Note that $\mathcal{C}$ has at most $M\pi^2/C$ elements (compare areas). Let γ be the geodesic on S_R connecting $r(z)$ and $r(w)$ and let E be a connected component of $r^{-1}(\gamma)$ containing z and w. Then E is a compact, connected set on R. Since a geodesic segment is convex, E can intersect at most two boundary crescents of R. Thus E can be covered by a collection of squares on R each of which projects to some element of $\mathcal{C}$ and which satisfies $\text{dist}(Q, \partial R) \geq \text{diam}(Q)$. By the upper bound in Lemma 16 the latter condition implies each square has uniformly bounded diameter for the hyperbolic metric on R. Moreover, each element of $\mathcal{C}$ is used at most three times (once by a square hitting P and twice by squares hitting the boundary crescents). Thus E is covered by a uniformly bounded number of sets, each of which has uniformly bounded hyperbolic diameter. Thus any two points of E (in particular, z, w) are a uniformly bounded distance apart.

This completes the proof that the map ι defined using a Delaunay triangulation of a polygon has a K-quasiconformal extension to a map $R_P \to \mathbb{D}$ with a K that is bounded independent of P.

10. Proof of Theorem 4

It would be nice if the ι map of the vertices of P defined in Section 8 had a K-quasiconformal extension to a map $P \to \mathbb{D}$, but this is not true with K independent of P. Instead, we will see how to add new vertices of angle π to P to obtain a polygon P' for which it is true. This is basically the first step of the CRDT algorithm as described in [41].

To see that quasiconformal extension need not hold uniformly for all P, simply consider a $1 \times n$ rectangle. The four corners lie on the boundary of a disk D and the Delaunay triangulation consists of adding a diagonal. The modulus of the path family in the rectangle connecting the two sides of length 1 is $1/n$. The modulus of the generalized quadrilateral with domain D and the four corners as vertices is approximately the logarithm of their cross ratio, i.e., is approximately $1/\log n$. However, if ι had a K-quasiconformal extension to P with K independent of P, these two moduli would be comparable within a factor of K of each other, so we deduce there is no such extension for large n.

The problem with the previous example is that the quadrilateral formed by taking adjacent triangles in the Delaunay triangulation was too long and skinny. However, if all the triangles in the triangulation are "round" (e.g., no small angles) then there is no problem. The rest of this section is devoted to making this more precise.

Lemma 19. *Suppose P has a rooted Delaunay triangulation such that each triangle either (1) is isosceles with the base adjacent to its parent and base angle $\leq \pi/4$, (2) is isosceles with any angle and is a leaf of the triangulation tree with the base adjacent to its parent or (3) has all angles in $[\theta, \pi - \theta]$ for some $\theta > 0$. Then there is a K-quasiconformal map $f : P \to R_P$ which fixes each of the vertices of P, and K depends only on θ.*

Proof. First consider the case when T is isosceles. We will let the two equal base angles be denoted α and the opposite angle it then $\pi - 2\alpha$. The circumcircle of the triangle defines two crescents attached to T along the non-base sides of T. We claim that the interior angles of these crescents are also α. This is simple plane geometry as illustrated in Figure 19. By definition $\angle cad = \alpha$ and $\angle adc = \frac{\pi}{2} - \alpha$. Since $\triangle abd$ is isosceles with base ad we have $\angle bad = \frac{\pi}{2} - \alpha$ and $\angle abd = 2\alpha$. Since triangles eab

and eca are similar, this means $\angle eac = \angle eba = 2\alpha$ and hence $\angle ead = 2\alpha - \alpha = \alpha$, as claimed.

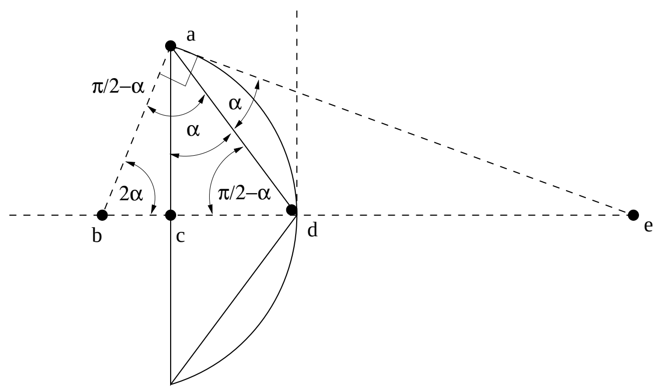

FIGURE 19. The isosceles case: small α

If $\alpha \leq \frac{\pi}{4}$, then the reflection of each crescent attached to T (across the edge shared with T) is a crescent inside T. Moreover the two crescents we get are disjoint. The union of a crescent C and its reflection is a crescent C' of twice the angle and so we can map C to C' be a 2-quasiconformal map which extends to the identity on the rest of T.

If T is isosceles but $\alpha > \frac{\pi}{4}$, the picture is illustrated in Figure 20. By assumption the triangle is attached to its parent along the base, but not adjacent to any other triangle. Now we map the big crescent with angle 2α (this is the angle by the previous paragraph) and base ab to the smaller crescent with angle α and the same base by a 2 quasiconformal map. Then map the small crescent to the triangle abc by a K-quasiconformal map which is the identity on ab and send d to c. This can be done with a uniform K, independent of α. We could either build such a map explicitly in this case, or we can make use of the fact that we know ι has uniform quasiconformal extension and note that the inverse of the ι map corresponding to $T \cup D \to D$ (where D is the disk determined by the points a, b, d) has the desired properties.

Finally consider a triangle T with all its angles bounded below by $\theta > 0$. Consider a crescent C of angle α attached to T along edge e. Then the angle bisectors of T split T into three subtriangles, each of which contains a crescent with angle $\geq \theta/2$ and shares an edge with T. See Figure 21. Let C' be the crescent of angle $\theta/2$ adjacent to C. We can map $C \cup C' \to C'$ by a quasiconformal map with constant

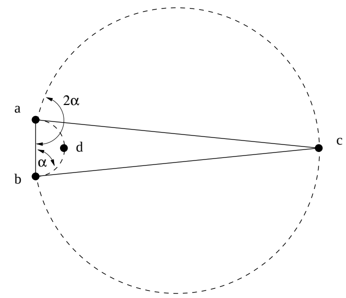

FIGURE 20. The isosceles case: large α

$K = (\alpha + \theta/2)/(\theta/2)$ which extends to the identity of the rest of T. Since $\alpha \le \pi$, we have $K \le 1 + 2\pi/\theta$, as desired.

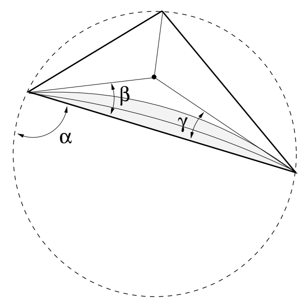

FIGURE 21. Mapping R_P to R quasiconformally

□

Lemma 20 (Driscoll-Vavasis, [41]). *Given any polygon P we can add vertices of angle π to form a new polygon which satisfies the hypothesis of Lemma 19.*

We will not repeat the proof from [41] here, but we will sketch their construction, which has two steps. In the first step, every vertex with interior angle $\le \pi/4$ is "chopped off" by taking the largest isosceles triangle T in P formed with v and subsets of its two adjacent edges and adding two new vertices to P at the midpoints of the two sides of T. A new polygon P' is formed by replacing v by these two new vertices and the edge between them. This edge is protected: no new vertices may be added to it later. The second step is iterative. For each edge e in the current polygon compute its length L and the minimal distance D in P from e to any vertex which is not one of its endpoints (distance is the path distance within the polygon). If $D < L/(3\sqrt{2})$, the edge e is split into three equal edges and the process continues.

See Figure 22. The filled dots indicate the original vertices and the open dots the vertices added by the algorithm.

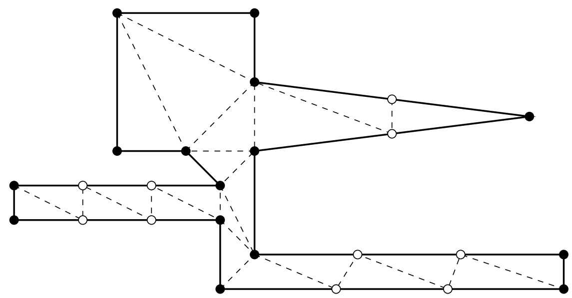

FIGURE 22. Adding extra points to a polygon in the CRDT algorithm to remove small angles from the triangulation

Driscoll and Vavasis give a lower bound r_0 on the shortest edge that can be produced: r_0 is the minimum over all unprotected edges of the path distance in P from that edge to any non-adjacent edge. Since every step reduces the length of some edge by a third, this proves the process terminates in a finite number of steps (and gives an estimate for the number of vertices that have been added). Thus the resulting polygon has a triangulation in which every triangle T is either (1) isosceles with angle $\leq \pi/4$ or (2) $D \geq L/(3\sqrt{2})$ for any side of T. The later condition easily implies the angles of T are bounded away from 0, so Lemma 20 holds.

As Driscoll and Vavasis note in their paper [41], the number of new points needed can be arbitrarily large, depending on the geometry of P (if P has a long, narrow corridor, then many new vertices have to be added).

Corollary 21. *Suppose P is a planar polygon satisfying the conditions of Lemma 20 (e.g., P is a Driscoll-Vavasis refinement of another polygon). Suppose P has vertices $\mathbf{v}$, $f : \mathbb{D} \to P$ is conformal and $\mathbf{z} = f^{-1}(\mathbf{v})$, are the conformal prevertices. Let $\mathbf{w} = \iota(\mathbf{v})$ be the approximate prevertices computed by the CRDT algorithm. Then there is a $K < \infty$ (independent of n and P) so that $d_{QC}(\mathbf{z}, \mathbf{w}) \leq \log K$. In particular, for any j_1, j_2, j_3, j_4,*

$$\frac{1}{K} \leq \frac{\mathrm{Mod}_{\mathbb{D}}(z_{j_1}, z_{j_2}, z_{j_3}, z_{j_4})}{\mathrm{Mod}_{\mathbb{D}}(w_{j_1}, w_{j_2}, w_{j_3}, w_{j_4})} \leq K.$$

This result proves, in a weaker form, the conjecture stated on page 1792 of [41]. On the other hand, it holds for all 4-tuples, not just the ones coming from adjacent Delaunay triangles.

11. A counterexample

Here we shall show that Corollary 21 does not hold with modulus replaced by cross ratios. Consider the domain Ω bounded by a polygon P in Figure 23. It is the square $[-1,1]\times[-1,1])$ with thickened slit removed, $[0,1]\times[-\epsilon,\epsilon]$. The dotted lines in the figure are the boundaries of a Delaunay triangulation of P and the shaded quadrilateral Q is formed by the union of two adjacent triangles. The absolute value of the cross ratio of the four vertices is

$$\mathrm{cr}_1 = \frac{|D-A||B-C|}{|C-D||A-B|} = \frac{\sqrt{1+(1-\epsilon)^2}\cdot\sqrt{1+(1-\epsilon)^2}}{2\epsilon\cdot 2} = \frac{1}{2\epsilon} - \frac{1}{2} + \frac{\epsilon}{4}.$$

On the other hand, the conformal map f of Ω to the disk which sends $-1/2$ to 0 and has positive derivative at $-1/2$ will map the edge $[A,B]$ to an interval on $\mathbb{T}$ centered at -1 and with length $\simeq 1$, but will map the edge $[C,D]$ to an arc centered at 1 of length $\simeq \sqrt{|C-D|}$. The cross ratios of these points will therefore be

$$\mathrm{cr}_2 = \frac{|f(D)-f(A)||f(B)-f(C)|}{|f(C)-(D)||f(A)-f(B)|} \leq \frac{M}{\sqrt{\epsilon}},$$

for some constant M independent of ϵ, if ϵ is small enough. Hence

$$\log \mathrm{cr}_1 - \log \mathrm{cr}_2 > \frac{1}{2}\log\frac{1}{\epsilon} - O(1),$$

which is as large as we wish if ϵ is small enough. This polygon does not satisfy the

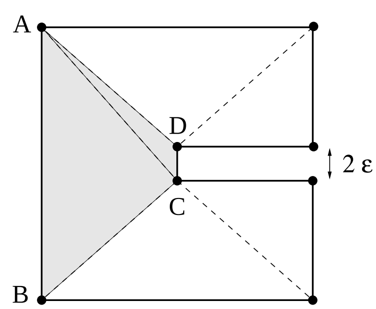

Figure 23. Cross ratios of $f(Q)$ and $\iota(Q)$ scale at different powers of ϵ

edge separation condition of the CRDT algorithm, but applying that method to this polygon will only add vertices to the horizontal edges adjacent to C and D and the same quadrilateral will be present. Thus the conjecture fails even if we add the extra vertices.

12. Variations, questions and conjectures

Modifying ι: There are several ways that the ι map might be improved to give a better approximation to the Riemann map, but it is not clear which of these (if any) would consistently give a better result, so it would be interesting to perform some numerical experiments to test these changes. For example, consider the infinite horizontal strip $\Omega = (-\infty, \infty) \times (-1, 1)$. Given points s, t on the x-axis we easily see $\rho_\Omega(s,t) = \frac{\pi}{2}|t-s|$ but the corresponding points on the dome satisfy $\rho_{S_\Omega}(r(s), r(t)) = |t-s|$. Thus ι will underestimate hyperbolic distance in long narrow channels. Equivalently, ι makes the channel look wider than it really is. Would it be a good idea to replace a long channel with a narrower one before computing ι? We could try to account for this changing our formula for τ in the second case of Section 7 or by actually constructing an auxiliary domain in which the channel has been "pinched".

Discretizing the medial axis: One systematic way to approximate Ω by smaller domains is to choose a finite subset of the medial axis and take the union of the corresponding disks. This gives a finitely bent domain which is inside P. If we include the vertices of the medial axis then all the concave vertices of P will be vertices of the approximating domain and so we can compute ι for them (the convex vertices require a separate but easy computation). Is there a systematic way of choosing a finite approximation to the medial is which gives good results?

Solving a Beltrami equation: Here is one way to improve ι. The map $\iota : \partial\Omega \to \partial\mathbb{D}$ has a quasiconformal extension to the interior and we can can even compute what the dilatation of this mapping is [17]. If we transfer this to the disk and solve the corresponding Beltrami equation then we can get a quasiconformal self-map h of the disk which sends our approximate prevertices $\mathbf{w} = \iota(\mathbf{v})$ to the true prevertices $\mathbf{z}$. Thus an approximation to h should map $\mathbf{w}$ closer to $\mathbf{z}$.

If the L^∞ norm of the dilatation is small, then one can approximately solve the Beltrami equation by an integral (see [1], [15]). More precisely, there is a $0 < k < 1$ and a $C < \infty$ so that the following holds. Suppose that f is a quasiconformal mapping of the plane to itself which preserves $\mathbb{H}$, fixing $0, 1$ and ∞ and the Beltrami coefficient of f is μ with $\|\mu\|_\infty \leq k$. Then

$$|f(w) - [w - \frac{1}{\pi}\int_{\mathbb{R}^2} \mu(z)R(z,w)dxdy]| \leq C\|\mu\|_\infty^2,$$

for all $|w| \leq 1$, where

$$R(z,w) = \frac{1}{z-w} - \frac{w}{z-1} + \frac{w-1}{z} = \frac{w(w-1)}{z(z-1)(z-w)}.$$

The dilatation μ that occurs for ι can be thought of as a sum of dilatations $\mu = \sum \mu_k$ with disjoint supports where each μ_k corresponds to shrinking a single crescent down to a geodesic. For each component the Beltrami equation can be solved explicitly because we can find an explicit conformal map from a union of overlapping disks $D_1 \cup D_2$ to one of them (map the two vertices to $0, \infty$ and use a power function). Thus the integral above is a sum of explicitly computable functions. Subtracting this function from ι should give an even better approximation to the conformal map than ι does (with rigorous bounds if ι is already sufficiently close to conformal).

This only works for maps onto the disk, but variations may work for other domains, as might arise in a continuation method. In that case it would be harder to compute the solutions by hand, but given the form of the integral, it seems reasonable that some variation of the Fast Multipole Method ([79], [54], [27]) could evaluate it quickly and that one could prove some uniform improvement in the d_{QC} distance.

Extending to non-simple polygons: Given n points on the unit circle and an angle $\alpha_k\pi$ assigned to each one, the Schwarz-Christoffel formula defines a locally univalent map of the disk. The map need not be globally 1-to-1, but it can always be interpreted as a 1-to-1 conformal map onto a Riemann surface whose boundary consists of n straight line segments. Such a surface can be triangulated (i.e., written as a union of $n-2$ Euclidean triangles) and one can even find a Delaunay triangulations (the angle condition is satisfied for the union of any two adjacent triangles). Using this one can define an ι of the vertices of the surface to an n-tuple on the circle (which is well defined modulo Möbius transformations).

Iterations of Ψ: Fix n angles and let S denote the Schwarz-Christoffel map from n-tuples on $\mathbb{T}$ to polygons determined by these angles, Iterate the map $\Psi = \iota \circ S$ on n-tuples in the circle. Do the iterates converge? Is there a fixed point? Similarly investigate the closely related map $S \circ \iota$ which maps polygons to polygons. Do they remain bounded under iteration? How do sides degenerate? Near an internal angle of size α the conformal map to the disk acts like $z^{\pi/\alpha}$ and ι behaves line z if $\alpha > \pi$ and $z^{1/\sin(\alpha)}$ if $\alpha \leq \pi$. Thus $\iota \circ S$ will look repelling nears points with angle $< \pi$ and attracting near angles $> \pi$.

Coordinates on n-tuples: Let X_n denote the space of n-tuples of distinct points on the unit circle (counter-clockwise order) and let $\tilde{X}_n$ be X_n with any two Möbius equivalent n-tuples identified. Let S_α be the mapping of X_n to polygonal surfaces given by the Schwarz-Christoffel formula with angles $\alpha = \{\alpha_1, \dots, \alpha_n\}$. If we change the positions of the prevertices by a Möbius transformation of the disk, the corresponding Schwarz-Christoffel mapping sends the disk conformally onto a new polygonal surface with the same angles as before. Moreover, there is clearly a conformal map f between the two surfaces which maps vertices to vertices and hence $\arg(f')$ is constant on the boundary. This implies f is linear and hence the ι map for the two surfaces agree up to a Möbius transformation. Thus Ψ is a well defined map $\tilde{X}_n \to \tilde{X}_n$. Is it 1-to-1? There is a related result of Snoeyink [98], but his result seems to require a fixed triangulation whereas our triangulations vary over the whole space. His result should imply that Ψ is at most finite to one since there are only finitely many possible triangulations to consider.

Conformal moduli coordinates: We can give $\tilde{X}_n$ coordinates which identify it with $\mathbb{R}^{n-3}$ by taking a triangulation of the disk with the given vertices, fixing a root triangle and using the $n-3$ numbers $\log \mathrm{Mod}_{\mathbb{D}}(z_{j_1}, z_{j_2}, z_{j_3}, z_{j_4})$ for each 4-tuple corresponding to a pair of adjacent triangles in our triangulation. Theorem 1 implies

$$\|\Psi(\mathbf{w}) - \mathbf{w}\| \leq \log 8,$$

on the subset $Y_n \subset \tilde{X}_n$ corresponding to planar polygons. On the other hand, any polygonal surface R we get is at worse a $n-2$ to 1 cover of the plane and this implies that ι has a $K(n)$-quasiconformal extension to R. Thus

$$\|\Psi(\mathbf{w}) - \mathbf{w}\| \leq \log K(n).$$

Since Ψ is continuous, this condition implies it is onto from $\mathbb{R}^{n-3}$ to itself. Is Ψ 1-to-1? Is it a diffeomorphism? Does $\|\Psi(\mathbf{w}) - \iota(\mathbf{v})\|^2$ have any local minimums besides the true conformal prevertices?

Rough monotonicity: Suppose we take coordinates for $\tilde{X}_n$ as above, fix one of the quadrilaterals Q, and vary its modulus. How does the modulus of the image quadrilateral $\Psi(Q)$ vary (assuming all the other coordinates are held fixed)? According to our remarks above, we have

$$-K \leq \text{Mod}(Q) - \text{Mod}(\Psi(Q)) \leq K,$$

so that the coordinates of Ψ roughly monotone on large scales (if $\text{Mod}(Q)$ increases by more than $2K$ then $\text{Mod}(\Psi(Q))$ must increase). Unfortunately, the relationship need not be monotone on small scales, we we shall see next.

Consider a polygon P bounding a region Ω and two adjacent Delaunay triangles v_1, v_2, v_3 and v_3, v_4, v_1 with the same circumdisk D. Let Q be the generalized quadrilateral with domain D and vertices v_1, v_2, v_3, v_4. Then $\text{Mod}_Q(v_1, v_2, v_3, v_4)$ is preserved by the ι map to the circle, but $\text{Mod}_\Omega(v_1, v_2, v_3, v_4)$ is preserved by the conformal map to the disk. Thus it suffices to show the second quantity is not a monotone function of the first. By first taking a Möbius transformation that takes v_1 to 0 and v_3 to ∞ and then a logarithm, we can map Ω to a domain Ω' where D is mapped to the strip $S = \{x + iy : |y| < 1\}$ and varying v_4 corresponds to translating the upper boundary of Ω' horizontally.

Let

$$\Omega = \{x + iy : |y| < 1 + \epsilon \text{ and } |y| < 1 \text{ if } x \in \mathbb{Z}\}.$$

Thus Ω is a strip with some vertical slits removed. See Figure 24. Let Ω_t be the

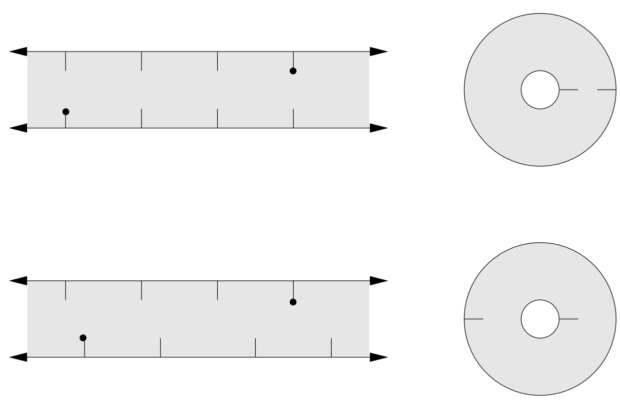

FIGURE 24. Coordinates of Ψ need not be monotone

domain obtained by translating the upper boundary component of Ω t units. Let $M(t) = \text{Mod}_{\Omega_t}(-\infty, -i, +\infty, i + t)$. We wish to show that this is not a monotone function of t. The domain Ω_t is clearly a covering space of a doubly connected domain

A_t which is a round annulus with two slits removed, one with argument 0 and the other with argument $2\pi t$ (see Figure 24). It is easy to check that the modulus of the path family separating the two boundary components of this domain is not constant it t: it is minimized when the slits are on the same radius and maximized when they are opposite one another. Let m_1 and m_2 denote these two extreme values of the modulus. If n is an integer $M(n) = m_1 n + O(1)$ and $M(n + \frac{1}{2}) = m_2 n + O(1)$. Thus for n large $M(n+1) - M(n + \frac{1}{2})$ and $M(n + \frac{1}{2}) - M(n)$ must have opposite signs and hence M is not monotone.

Solving the parameter problem: Theorem 1 gives a rough approximation of the conformal prevertices, but we really want to compute them to any given accuracy. In [41] Driscoll and Vavasis observe that the iteration

$$\mathbf{w}_{n+1} = \mathbf{w}_n - \Psi(\mathbf{w}_n),$$

seems to converge linearly to the correct prevertices starting from the initial guess of CRDT (actually they consider an iteration where the coordinates are logarithms of cross ratios rather conformal modulus, but we will ignore the difference). This would occur if Ψ had a derivative that approximates the identity. Is this true in some sense? Our estimates above show that Ψ looks like the identity on large scales, so perhaps it also looks like the identity on small scales at many points.

Another method of solving for the correct prevertices is to define a function on $\tilde{X}_n$ which attains a unique minimum at the correct prevertices and then use a general non-linear solver to find them (this approach is described in [40] and [61]). Following [41] let $\mathbf{v}_Q$ be the four vertices corresponding to two adjacent triangles in a Delaunay triangulation of P we could define

$$F(\mathbf{w}) = \|\Psi(\mathbf{w}) - \mathbf{w}_0\|^2 = \sum_Q |\mathrm{Mod}_\Omega(\mathbf{v}_Q) - \mathrm{Mod}_{\mathbb{D}}(\iota(\mathbf{v}_Q))|^2,$$

where $\mathbf{w}_0 = \iota(\mathbf{v})$ and the sum is over quadrilaterals associated to a Delaunay triangulation of P. It would be good if this function had some monotonicity. As we saw above, the individual terms do not have to be monotone on small scales, although they do have a rough monotonicity at large scales. In particular, the level sets of the function are compact. The approach of Davis [37] uses an iteration that is based on the heuristic that increasing the length of an edge increases the distance between the corresponding prevertices (i.e., harmonic measure is monotone in arclength). This is false in some cases, and as was pointed out in [61] this method need not converge even if we start near the correct answer. Using the ι map and these coordinates would seem to have a better chance, although I do not know a proof that non-linear solver will converge to the correct answer.

Parameterizing by the medial axis: Earlier we indicated how one might use non-linear solvers to find the prevertices by parameterizing our polygons by conformal moduli of quadrilaterals obtained from the ι. Since ι is derived from the medial axis, could we instead define a distance between polygons based directly on the medial axis and use this? This would be like the approach of Davis [37] of using an iteration based on the geometry of the polygon. Some authors have already considered distances between medial axes based on hyperbolic metrics [33], [34].

Continuation by angle scaling: Another possible advantage of the ι maps is that there is a natural way to write our domain Ω as part of a one parameter family which starts at the disk and ends at Ω. This is called "angle scaling" and is most easily described when Ω is finitely bent. This is explained in detail by Epstein, Marden and Markovic [44], so we will only sketch their results here. In the finitely bent case, Ω is a union of crescents (which project to bending geodesics on the dome) and the regions between crescents (which project to geodesic faces on the dome). There is a natural way to shrink the angle of each crescent by a factor of $t \in [0,1]$. When $t = 1$ we have the original domain. When $t = 0$ the crescents have all collapsed to circular arcs and Ω has been mapped to the disk (this is exactly the same as quasi-isometric extension of ι which we constructed in Section 5). Let Ω_t denote the result of scaling the angles by a factor of t, and let $\iota_t : \Omega_t \to \mathbb{D}$ be the corresponding map. Then $(\iota_t)^{-1} \circ \iota_s$ is a map of $\partial\Omega_s$ to $\partial\Omega_t$ and the important fact is that it has a quasiconformal extension to a map $\Omega_s \to \Omega_t$ with constant $C|s-t|$. This means that if the s and t are close, then the conformal prevertices for Ω_s and Ω_t are close in the d_{QC} metric.

Now suppose we could always solve for the true prevertices if we were within d_{QC} distance ϵ to start. To be precise, suppose we have a map F so that $d_{QC}(z,w) < \epsilon$ implies $d_{QC}(z, F(w)) < \lambda\epsilon$ for some $\lambda < 1$, independent of z and w. Now choose $t_0 = 0 < t_1 < \dots t_N = 1$ and let $\Omega_k = \Omega_{t_k}$. We choose these points close enough together that the true conformal prevertices $\mathbf{z}_k$ for Ω_k and Ω_{k+1} are within $\epsilon/2$ of each other for $k = 0, \dots, N-1$. Let w_0 be our initial guess for Ω. Then $\mathbf{w}_0 = \mathbf{z}_0$ (since Ω_0 is the disk, ι_0 is the identity). By assumption $d_{QC}(\mathbf{z}_0, \mathbf{z}_1) < \epsilon$, so we can find $\mathbf{z}_1$

starting from $\mathbf{w}_0 = \mathbf{z}_0$. In particular, after at most $\log 2/\log(1/\lambda)$ applications of F we have a point $\mathbf{w}_1$ with $d_{QC}(\mathbf{z}_1, \mathbf{w}_1) < \epsilon/2$. Hence this point is within ϵ of $\mathbf{z}_2$. Continuing this way we obtain after N steps a point $\mathbf{w}_{N-1}$ within ϵ of $\mathbf{z}_N = \mathbf{z}$, the true conformal prevertices of Ω, and iteration of F from this point converges to $\mathbf{z}$. Thus the ι map and angle scaling reduce the global convergence problem to proving local convergence (with a radius of convergence estimate that is uniform over the compact family Ω_t).

Angle scaling can be performed on any simply connected domain using the bending lamination, although the intermediate domains need not be planar (nor are they in the finitely bent case above). Moreover, angle scaling of polygons need not be polygonal.

This type of angle scaling shrinks all the crescents at the same time. The geometric description of ι given in the introduction (by following orthogonal paths to the medial axis foliation) is very similar, except that we collapse crescents farthest from the root first, and them move towards the root. This gives subsets of the original domain.

Optimal constants: What is the best constant in Theorem 1? If one allows the constant in the work estimate to grow, then probably one can get the d_{QC} distance as close to 0 as desired. If we only consider ι maps then this is related to finding the best K in the Sullivan-Epstein-Marden theorem. As noted in Section 4, the best current results are $2.1 \leq K \leq 7.82$ and I expect the correct value is closer to the smaller value.

It would be very interesting to try to estimate the best K numerically. I do not know a way to compute K directly, but given a polygon we could compute the conformal and ι images for various 4-tuples of vertices and comparing these would give a lower bound for K. It is known that given a homeomorphism h of the circle, if the modulus of any four points on the circle is changed by a factor of at most K, then h has a C-quasiconformal extension to the disk where C depends only on K (this is well known, e.g., [72]). One must take $C > K$ in general, [2], but $C = K$ is true if one considers more general path families than just quadrilaterals [64].

A sharp estimate for K would be of interest for several reasons. First of all, the best explicit estimates for the constant currently come from estimates of the amount of bending a surface in the hyperbolic space $\mathbb{R}^3_+$ can have without crossing itself (see [17], [23], [24], [43]) which is a topic of independent interest in low dimensional topology. The sharp K in Sullivan's theorem also has numerous potential applications in classical complex analysis and dynamics. Consider the well known Brennan's conjecture. Suppose Ω is a simply connected plane domain and $F = f^{-1} : \Omega \to \mathbb{D}$ is a conformal map. It is obvious that $\int_\Omega |F'|^2 dxdy = \text{area}(\mathbb{D}) = \pi$ so that $F' \in L^2(\Omega, dxdy)$, but it is not clear what other L^p spaces F' must belong to. Gehring and Hayman (unpublished) showed that $F' \in L^p$ for $p \in (\frac{4}{3}, 2]$ and showed the lower bound is sharp. Metzger [76] improved this to $p \in (\frac{4}{3}, 3)$. In 1978 James Brennan [22] improved this by showing one can take $p \in (\frac{4}{3}, p_0)$ for some $p_0 > 3$ and conjectured that $p_0 = 4$ is possible (this is sharp since the Koebe function mapping $\mathbb{D} \to \mathbb{C} \setminus [\frac{1}{4}, \infty)$ gives an $F' \notin L^4$). The best estimate (so far as I know) is currently due to Bertilsson [11], [12] who showed $p_0 \geq 3.422$. This is a slight improvement of the earlier result of Pommerenke [84, 83], that $p_0 \geq 3.399$. In addition to its intrinsic interest, the Brennan conjecture has interesting consequences and is currently under intense investigation, e.g., [8], [26], [62], [66], [75].

The connection between the Sullivan-Epstein-Marden theorem and Brennan's conjecture is explained in detail in [16]. Briefly, Astala [3] proved that if h is a K-quasiconformal map of the disk to itself, then $|h'|$ is in weak L^p where $p = 2K/(K-1)$. Thus, if the Sullivan-Epstein-Marden theorem holds K, then the Brennan conjecture is true for $p < 2K/(K-1)$. As noted above, Epstein and Markovic have produced a simply connected domain where $K > 2.1$, so a direct proof of Brennan's conjecture using $K = 2$ is not possible. However, in order to improve Bertilsson's result we would only need

$$K = \frac{p}{p-2} = \frac{3.422}{3.422 - 2} \approx 2.4064.$$

Moreover, Brennan's conjecture clearly holds for the Epstein-Markovic example, so we can still hope to prove Brennan's conjecture using some of these ideas. For example, can we estimate L^p norms of the quasiconformal mappings arising from ι directly, without passing though Astala's theorem?

References

[1] L. V. Ahlfors. *Lectures on quasiconformal mappings.* The Wadsworth & Brooks/Cole Mathematics Series. Wadsworth & Brooks/Cole Advanced Books & Software, Monterey, CA, 1987. With the assistance of Clifford J. Earle, Jr., Reprint of the 1966 original.

[2] J. M. Anderson and A. Hinkkanen. Quadrilaterals and extremal quasiconformal extensions. *Comment. Math. Helv.*, 70(3):455–474, 1995.

[3] K. Astala. Area distortion of quasiconformal mappings. *Acta Math.*, 173(1):37–60, 1994.

[4] F. Aurenhammer. Voronoi diagrams - a survey of a fundamental geometric data structure. *ACM Comp. Surveys*, 23:345–405, 1991.

[5] F. Aurenhammer and R. Klein. Voronoi diagrams. In *Handbook of computational geometry*, pages 201–290. North-Holland, Amsterdam, 2000.

[6] T. Bagby. The modulus of a plane condenser. *J. Math. Mech.*, 17:315–329, 1967.

[7] L. Banjai and L.N. Trefethen. A multipole method for Schwarz-Christoffel mapping of polygons with thousands of sides. *SIAM J. Sci. Comp.* to appear.

[8] K. Barański, A. Volberg, and A. Zdunik. Brennan's conjecture and the Mandelbrot set. *Internat. Math. Res. Notices*, (12):589–600, 1998.

[9] C. Benedetti, R.and Petronio. *Lectures on hyperbolic geometry*. Universitext. Springer-Verlag, Berlin, 1992.

[10] M. Bern and D. Eppstein. Mesh generation and optimal triangulation. In *Computing in Euclidean geometry*, volume 1 of *Lecture Notes Ser. Comput.*, pages 23–90. World Sci. Publishing, River Edge, NJ, 1992.

[11] D. Bertilsson. Coefficient estimates for negative powers of the derivative of univalent functions. *Ark. Mat.*, 36(2):255–273, 1998.

[12] D. Bertilsson. *On Brennan's conjecture in conformal mapping*. PhD thesis, Royal Institute of Technology, Stockholm, Sweden, October 1999.

[13] M. Bestvina. $\mathbb{R}$-trees in topology, geometry, and group theory. In *Handbook of geometric topology*, pages 55–91. North-Holland, Amsterdam, 2002.

[14] C. J. Bishop. Divergence groups have the Bowen property. *Ann. of Math. (2)*, 154(1):205–217, 2001.

[15] C. J. Bishop. BiLipschitz approximations of quasiconformal maps. *Ann. Acad. Sci. Fenn. Math.*, 27(1):97–108, 2002.

[16] C. J. Bishop. Quasiconformal Lipschitz maps, Sullivan's convex hull theorem and Brennan's conjecture. *Ark. Mat.*, 40(1):1–26, 2002.

[17] C.J. Bishop. An explicit constant for Sullivan's convex hull theorem. In *Proceedings of the Ahlfors-Bers Colloqium, 2001*. Amer. Math. Soc., Providence, RI. to appear.

[18] H. Blum. A transfomation for extracting new descriptors of shape. In W.W. Dunn, editor, *Proc. Symp. Models for the perception of speech and visual form*, pages 362–380, Cambridge, 1967. MIT Press.

[19] H. Blum. Biological shape and visual science. *Journal of Theoretical Biology*, 38:205–287, 1973.

[20] H. Blum and R.N. Nagel. Shape descriptors using weighted symmetric axis functions. *Pattern Recognition*, 10(3):167–180, 1978.

[21] Francis Bonahon. Geodesic laminations on surfaces. In *Laminations and foliations in dynamics, geometry and topology (Stony Brook, NY, 1998)*, volume 269 of *Contemp. Math.*, pages 1–37. Amer. Math. Soc., Providence, RI, 2001.

[22] J.E. Brennan. The integrability of the derivative in conformal mapping. *J. London Math. Soc. (2)*, 18(2):261–272, 1978.

[23] M. Bridgeman. Average bending of convex pleated planes in hyperbolic three-space. *Invent. Math.*, 132(2):381–391, 1998.

[24] M. Bridgeman. Average curvature of convex curves in H^2. *Proc. Amer. Math. Soc.*, 126(1):221–224, 1998.

[25] M. Bridgeman and R. D. Canary. From the boundary of the convex core to the conformal boundary. *Geom. Dedicata*, 96:211–240, 2003.

[26] L. Carleson and N. G. Makarov. Some results connected with Brennan's conjecture. *Ark. Mat.*, 32(1):33–62, 1994.

[27] J. Carrier, L. Greengard, and V. Rokhlin. A fast adaptive multipole algorithm for particle simulations. *SIAM J. Sci. Statist. Comput.*, 9(4):669–686, 1988.

[28] C.-S. Chiang and C.M. Hoffmann. The medial axis transform for 2d regions. *ACM Transactions on graphics*, 1982.

[29] F. Chin, J. Snoeyink, and C. A. Wang. Finding the medial axis of a simple polygon in linear time. In *Algorithms and computations (Cairns, 1995)*, volume 1004 of *Lecture Notes in Comput. Sci.*, pages 382–391. Springer, Berlin, 1995.

[30] F. Chin, J. Snoeyink, and C. A. Wang. Finding the medial axis of a simple polygon in linear time. *Discrete Comput. Geom.*, 21(3):405–420, 1999.

[31] H. I. Choi, S. W. Choi, and H. P. Moon. Mathematical theory of medial axis transform. *Pacific J. Math.*, 181(1):57–88, 1997.

[32] H.I. Choi, C.Y. Han, and J.-H. Yoon. Medial axis tranform distance and its applications 2000, april 2000 kyung moon, seoul,korea. In *Geometric modeling and computer graphics*, pages 65–69.

[33] S. W. Choi and H.-P. Seidel. Hyperbolic Hausdorff distance for medial axis transformation. *Graphical Models*, 63:369–384, 2001.

[34] S. W. Choi and H.-P. Seidel. Linear one-sided stability of MAT for weakly injective domain. *J. Math. Imaging Vision*, 17(3):237–247, 2002.

[35] S.W. Choi and S.-W. Lee. Stability analysis of medial axis transform. In *Proc. 15th ICPR Brcelona, Spain*, volume 3, pages 139–142, 2000.

[36] T. Culver, J. Keyser, and D. Manocha. Accurate computation of the medial axis of a polyhedron. In *Proceedings of the fifth ACM symposium on Solid modeling and applications, June 8-11, 1999, Ann Arbor, MI USA*, pages 179–190, 1999.

[37] R.T. Davis. Numerical methods for coordinate generation based on Schwarz-Christoffel transformations. In *4th AIAA Comput. FLuid Dynamics Conf., Williamsburg VA*, pages 1–15, 1979.

[38] E. F. D'Azevedo and R. B. Simpson. On optimal interpolation triangle incidences. *SIAM J. Sci. Statist. Comput.*, 10(6):1063–1075, 1989.

[39] G.L. Dirichlet. Über die reducktion der positiven quadratischen formen mit drei unbestimmten ganzen zahlen. *J. Reine Angew. Math.*, 40:209–227, 1850.

[40] T. A. Driscoll and L. N. Trefethen. *Schwarz-Christoffel mapping*, volume 8 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, 2002.

[41] T. A. Driscoll and S. A. Vavasis. Numerical conformal mapping using cross-ratios and Delaunay triangulation. *SIAM J. Sci. Comput.*, 19(6):1783–1803 (electronic), 1998.

[42] D. B. A. Epstein, editor. *Analytical and geometric aspects of hyperbolic space*, volume 111 of *London Mathematical Society Lecture Note Series*, Cambridge, 1987. Cambridge University Press.

[43] D. B. A. Epstein and A. Marden. Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces. In *Analytical and geometric aspects of hyperbolic space (Coventry/Durham, 1984)*, volume 111 of *London Math. Soc. Lecture Note Ser.*, pages 113–253. Cambridge Univ. Press, Cambridge, 1987.

[44] D. B. A. Epstein, A. Marden, and V. Markovic. Complex angle scaling.

[45] D. B. A. Epstein, A. Marden, and V. Markovic. Complex earthquakes and deformations of the unit disk.

[46] D. B. A. Epstein, A. Marden, and V. Markovic. Quasiconformal homeomorphisms and the convex hull boundary.

[47] D. B. A. Epstein and V. Markovic. The logarithmic spiral: A counterexample to the K=2 conjecture.

[48] G. Evans, A. Middleditch, and N. Miles. Stable computation of the 2D medial axis transform. *Internat. J. Comput. Geom. Appl.*, 8(5-6):577–598, 1998.

[49] S. Fortune. Voronoï diagrams and Delaunay triangulations. In *Computing in Euclidean geometry*, volume 1 of *Lecture Notes Ser. Comput.*, pages 193–233. World Sci. Publishing, River Edge, NJ, 1992.

[50] S. Fortune. Voronoi diagrams and Delaunay triangulations. In *Handbook of discrete and computational geometry*, CRC Press Ser. Discrete Math. Appl., pages 377–388. CRC, Boca Raton, FL, 1997.

[51] D. Gaier. *Konstruktive Methoden der konformen Abbildung.* Springer Tracts in Natural Philosophy, Vol. 3. Springer-Verlag, Berlin, 1964.

[52] J.B. Garnett and D. Marshall. *Harmonic measure.* Cambridge University Press, Cambridge. to appear.

[53] C. Gaudeau, M. Boiron, and J. Thouvenot. Squelettisation et anamorphose dans l'étude de la dynamique des déformations des structures: application à l'analyse de la motricité gastrique. In *Recognition of shapes and artificial intelligence (Second AFCET-IRIA Cong., Toulouse, 1979), Vol. III (French)*, pages 57–63. IRIA, Rocquencourt, 1979.

[54] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. *J. Comput. Phys.*, 73(2):325–348, 1987.

[55] H. N. Gursoy and N. M. Patrikalakis. Automated interrogation and adaptive subdivision of shape using medial axis transform. *Advances in Engineering Software and Workstations*, 13(5/6):287–302, 1991.

[56] J. Heinonen and P. Koskela. Quasiconformal maps in metric spaces with controlled geometry. *Acta Math.*, 181(1):1–61, 1998.

[57] P. Henrici. *Applied and computational complex analysis. Vol. 3.* Pure and Applied Mathematics. John Wiley & Sons Inc., New York, 1986. Discrete Fourier analysis—Cauchy integrals—construction of conformal maps—univalent functions, A Wiley-Interscience Publication.

[58] C.M. Hoffmann. Computer vision, descriptive geometry and classical mechanics. In *Computer Graphics and Mathematics*, pages 229–244. Springer Verlag, Eurographics Series, 1992.

[59] C.M. Hoffmann. *Geometric Approaches to Mesh Generation*, volume 75 of *IMA Volumes in Mathematics and its Applications*, pages 31–52. Springer Verlag, 1995.

[60] C.M. Hoffmann and D. Dutta. On the skeleton of simple CSG objects. *Trans. ASME*, 115:87–94, 1993.

[61] L. Howell. *Computation of Conformal Maps by Modified Schwarz-Christoffel Transformations.* PhD thesis, 1990. citeseer.nj.nec.com/howell90computation.html.

[62] R. Hurri-Syrjänen and S. G. Staples. A quasiconformal analogue of Brennan's conjecture. *Complex Variables Theory Appl.*, 35(1):27–32, 1998.

[63] V. I. Ivanov and M. K. Trubetskov. *Handbook of conformal mapping with computer-aided visualization.* CRC Press, Boca Raton, FL, 1995. With 1 IBM-PC floppy disk (5.25 inch; HD).

[64] J. A. Jenkins. On quasiconformal mappings with given boundary values. *Indiana Univ. Math. J.*, 37(4):929–934, 1988.

[65] R. A. Jinkerson, S. L. Abrams, L. Bardis, C. Chryssostomidis, A. Clement, N. M. Patrikalakis, and F. E. Wolter. Inspection and feature extraction of marine propellers. *Journal of Ship Production,*, 9(2):88–106, 1993.

[66] P. Kraetzer. Experimental bounds for the universal integral means spectrum of conformal maps. *Complex Variables Theory Appl.*, 31(4):305–309, 1996.

[67] P. K. Kythe. *Computational conformal mapping.* Birkhäuser Boston Inc., Boston, MA, 1998.

[68] C.L. Lawson. *Software for C^1 surface interpolation*, pages ix+388. Academic Press [Harcourt Brace Jovanovich Publishers], New York, 1977. Publication of the Mathematics Research Center, No. 39.

[69] D.-T. Lee. The medial axis transform of a planar shape. *IEEE Transactions on pattern analysis and machine intelligence*, 4(4):363–369, 1982.

[70] D. T. Lee and R. L. Drysdale, III. Generalization of Voronoĭ diagrams in the plane. *SIAM J. Comput.*, 10(1):73–87, 1981.
[71] Y.-H. Lee and S.-J. Horng. The equivalence of the chessboard distance transform and the medial axis transform. *Int. J. Comput. Math.*, 65(3-4):165–177, 1997.
[72] O. Lehto and K. I. Virtanen. *Quasiconformal mappings in the plane*. Springer-Verlag, New York, second edition, 1973. Translated from the German by K. W. Lucas, Die Grundlehren der mathematischen Wissenschaften, Band 126.
[73] T. Maekawa and N. M. Patrikalakis. Computation of singularities and intersections of offsets of planar curves. *Computer Aided Geometric Design*, 10(5):407–429, 1993.
[74] T. Maekawa and N. M. Patrikalakis. Interrogation of differential geometry properties for design and manufacture. *The Visual Computer*, 10(4):216–237, 1994.
[75] N.G. Makarov. Fine structure of harmonic measure. *Algebra i Analiz*, 10(2):1–62, 1998. English translation in St. Petersburg J. Math, 1999.
[76] T. A. Metzger. On polynomial approximation in $A_q(D)$. *Proc. Amer. Math. Soc.*, 37:468–470, 1973.
[77] J. W. Morgan. Λ-trees and their applications. *Bull. Amer. Math. Soc. (N.S.)*, 26(1):87–112, 1992.
[78] Z. Nehari. *Conformal mapping*. Dover Publications Inc., New York, 1975. Reprinting of the 1952 edition.
[79] S. T. O'Donnell and V. Rokhlin. A fast algorithm for the numerical evaluation of conformal mappings. *SIAM J. Sci. Statist. Comput.*, 10(3):475–487, 1989.
[80] J. O'Rourke. *Computational geometry in C*. Cambridge University Press, Cambridge, second edition, 1998.
[81] N. Papamichael and E. B. Saff, editors. *Computational complex analysis*. North-Holland Publishing Co., Amsterdam, 1993. J. Comput. Appl. Math. **46** (1993), no. 1-2.
[82] N. M. Patrikalakis and T. Maekawa. *Shape Interrogation for Computer Aided Design and Manufacturing*. Springer Verlag, 2002.
[83] Ch. Pommerenke. On the integral means of the derivative of a univalent function. *J. London Math. Soc. (2)*, 32(2):254–258, 1985.
[84] Ch. Pommerenke. On the integral means of the derivative of a univalent function. II. *Bull. London Math. Soc.*, 17(6):565–570, 1985.
[85] Ch. Pommerenke. *Boundary behavior of conformal maps*. Grundlehren der Math. Wissenschaften, 299. Springer-Verlag, 1992.
[86] H. Pottmann and J. Wallner. *Computational line geometry*. Mathematics+Visualization. Springer-Verlag, Berlin, 2001.
[87] F. P. Preparata. The medial axis of a simple polygon. In *Mathematical foundations of computer science (Proc. Sixth Sympos., Tatranská Lomnica, 1977)*, pages 443–450. Lecture Notes in Comput. Sci., Vol. 53. Springer, Berlin, 1977.
[88] F. P. Preparata and M. I. Shamos. *Computational geometry*. Texts and Monographs in Computer Science. Springer-Verlag, New York, 1985. An introduction.
[89] V. T. Rajan. Optimality of the Delaunay triangulation in $\mathbf{R}^d$. *Discrete Comput. Geom.*, 12(2):189–202, 1994.
[90] J. G. Ratcliffe. *Foundations of hyperbolic manifolds*, volume 149 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994.
[91] G. X. Ritter. Topology of computer vision. In *Proceedings of the 1987 Topology Conference (Birmingham, AL, 1987)*, volume 12, pages 117–158, 1987.
[92] C. Rourke. Convex ruled surfaces. In *Analytical and geometric aspects of hyperbolic space (Coventry/Durham, 1984)*, pages 255–272. Cambridge Univ. Press, Cambridge, 1987.

[93] Peter B. Shalen. Dendrology of groups: an introduction. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 265–319. Springer, New York, 1987.

[94] Peter B. Shalen. Dendrology and its applications. In *Group theory from a geometrical viewpoint (Trieste, 1990)*, pages 543–616. World Sci. Publishing, River Edge, NJ, 1991.

[95] E. C. Sherbrooke, N. M. Patrikalakis, and E. Brisson. Computation of the medial axis transform of 3-d. *Symposium on Solid Modeling and Applications*, pages 187–200, 1995.

[96] E.C. Sherbrooke, N. M. Patrikalakis, and E. Brisson. An algorithm for the medial axis transform of 3d polyhedral solids. *IEEE Transactions on Visualization and Computer Graphics*, 2(1):44–61, 1996.

[97] E.C. Sherbrooke, N. M. Patrikalakis, and F.-E. Wolter. Differential and topological properties of medial axis transforms. *CVGIP: Graphical Model and Image Processing*, 58(6):574–592, 1996.

[98] J. Snoeyink. Cross-ratios and angles determine a polygon. *Discrete Comput. Geom.*, 22(4):619–631, 1999. 14th Annual ACM Symposium on Computational Geometry (Minneapolis, MN, 1998).

[99] D. Sullivan. Travaux de Thurston sur les groupes quasi-fuchsiens et les variétés hyperboliques de dimension 3 fibrées sur S^1. In *Bourbaki Seminar, Vol. 1979/80*, pages 196–214. Springer, Berlin, 1981.

[100] W.P. Thurston. *The geometry and topology of 3-manifolds.* The Geometry Center, University of Minnesota, 1979.

[101] L. N. Trefethen, editor. *Numerical conformal mapping.* North-Holland Publishing Co., Amsterdam, 1986. Reprint of J. Comput. Appl. Math. **14** (1986), no. 1-2.

[102] L. N. Trefethen and T. A. Driscoll. Schwarz-Christoffel mapping in the computer era. In *Proceedings of the International Congress of Mathematicians, Vol. III (Berlin, 1998)*, number Extra Vol. III, pages 533–542 (electronic), 1998.

[103] J. Väisälä. Free quasiconformality in Banach spaces. II. *Ann. Acad. Sci. Fenn. Ser. A I Math.*, 16(2):255–310, 1991.

[104] W. von Koppenfels and F. Stallmann. *Praxis der konformen Abbildung.* Die Grundlehren der mathematischen Wissenschaften, Bd. 100. Springer-Verlag, Berlin, 1959.

[105] G.M. Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. recherches sur les parallélloèdres primitifs. *J. Reine Angew. Math.*, 134:198–287, 1908.

[106] J. Wang. Medial axis and optimal locations for min-max sphere packing. *J. Comb. Optim.*, 4(4):487–503, 2000.

[107] E.-F. Wolter. Cut locus and the medial axis in global shape interrogation and representation. 1993. MIT, Dept. of Ocean Engineering, Design Laboratoy Memorandum 92-2.

[108] Q. J. Wu. Sphere packing using morphological analysis. In *Discrete mathematical problems with medical applications (New Brunswick, NJ, 1999)*, volume 55 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 45–54. Amer. Math. Soc., Providence, RI, 2000.

[109] C.-K. Yap. An $O(n \log n)$ algorithm for the Voronoĭ diagram of a set of simple curve segments. *Discrete Comput. Geom.*, 2(4):365–393, 1987.

C.J. Bishop, Mathematics Department, SUNY at Stony Brook, Stony Brook, NY 11794-3651

E-mail address: bishop@math.sunysb.edu