A FAST MAPPING THEOREM FOR POLYGONS

CHRISTOPHER J. BISHOP

ABSTRACT. Given a simple n-gon P in the plane, normalized to contain the unit disk, we define a map from P to the unit circle so that (1) the map extends to be 8-quasiconformal on the interior, (2) it contracts arclength on the boundary and (3) the images of all n vertices can be computed in time $O(n)$. Thus we obtain a fast QC-approximation to the Riemann map with constant independent of the domain.

Date: July 2, 2008.

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 04-05578.

1. Introduction

The Schwarz-Christoffel formula gives the general form of a conformal map from the unit disk, $\mathbb{D}$, to a polygonal simply connected region, Ω, as

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

where $\alpha\pi = \{\alpha_1\pi, \ldots, \alpha_n\pi\}$ are the interior angles at the vertices $\mathbf{v} = \{v_1, \ldots, v_n\}$, and $\mathbf{z} = \{z_1, \ldots, z_n\} = f^{-1}(\mathbf{v})$ are the conformal preimages of the vertices (also know as the Schwartz-Christoffel parameters). This formula was published by Christoffel in 1867 and Schwarz in 1969, but even today computing the correct parameters $\mathbf{z}$ for a given polygon can be rather challenging (see [26] for some history and a survey of current methods). We will show, however, that there is a simple formula that gives approximations to the parameters with estimates that are independent of the particular polygon being considered.

We measure the distance between n-tuples, $\mathbf{z}, \mathbf{w} \in \mathbb{D}$, using the metric

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

Note that two n-tuples are distance zero apart if they are Möbius images of each other; this is natural, since we only expect to know the prevertices up to a Möbius transformation.

Theorem 1. *If Ω is simply connected plane domain bounded by a closed polygon P with n vertices, then we can compute n points $\mathbf{w} = \{w_1, \ldots, w_n\} \subset \mathbb{T}$ so that*

1. *$d_{QC}(\mathbf{w}, \mathbf{z}) < \log 8$ where $\mathbf{z}$ denotes the conformal prevertices of P,*
2. *All n points in $\mathbf{w}$ can be computed in $O(n)$ steps.*
3. *If $\mathbb{D} \subset \Omega$, then $|w_k - w_{k+1}| \leq |v_k - v_{k+1}|$ for $k = 1, \ldots n$ (mod n).*

This result lies in the intersection of geometric function theory, computational geometry and the theory of hyperbolic 3-manifolds. As we shall explain below, the points $\mathbf{w} = \iota(\mathbf{v})$ will be the images of the vertices of P under a map $\iota : \partial\Omega \to \partial\mathbb{D}$. The definition of this map and the fact that property (1) holds is based on a result of Dennis Sullivan concerning the convex hull of a hyperbolic 3-manifold. The fact $\mathbf{w}$ can be computed in linear time is based on re-interpreting the hyperbolic geometry in terms of a planar object known as the medial axis, which can be computed in linear

time by a result of Chin, Snoeyink and Wang. The contracting property follows from simply properties of Möbius transformations. Note that conformal maps do not have this property in general, so condition (3) actually prevents our particular approximation from being too close to a conformal map in some cases.

The medial axis of Ω consists of the centers of all disks in Ω whose boundary hits $\partial\Omega$ in two or more points. See Figure 1. It is a one dimensional object that is used in computer science to encode the shape of 2-dimensional objects and was introduced by Blum in 1967 [8] (the same set, with a different name, also appears in a 1945 paper of Erdös [19]). There is a large literature describing its mathematical properties and numerous applications including [12], [20], [21], [23], [27], [28]. It is a theorem of Chin, Snoeyink and Wang [10], [11] that the medial axis of polygon with n vertices can be computed in $O(n)$ steps. This depends on a deep and difficult result of Chazelle [9] that a polygon can be triangulated in linear time. However, special cases, such as convex polygons are easier (see [1]) and $O(n \log n)$ methods have been implemented for general polygons.

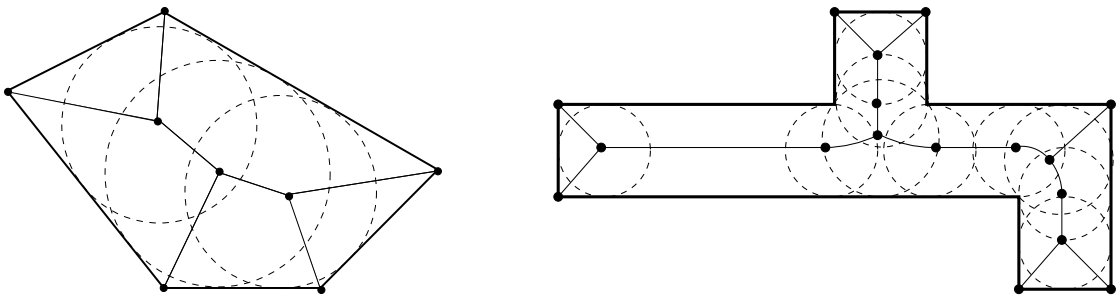

FIGURE 1. Examples of medial axes of polygons

How does the medial axis of $\partial\Omega$ give a map from $\partial\Omega$ to the boundary of a disk? If Ω is a finite union of disks, then its medial axis is a finite tree; vertices correspond to disk that hit the boundary in three or more points. We can rewrite Ω as the union of disks $\{D_k\}$ corresponding to vertices of this tree. Choose one of these, D_0, as the root. Then each non-root disk has a parent disk (the one closer to the root) and if we remove the parent from the disk, we are left with a crescent. Thus Ω may be written as the union of the root disk and a finite union of crescents.

Every crescent has two natural foliations by circular arcs: one by arcs passing through the two vertices (we call this the medial axis foliation) and the other by arcs

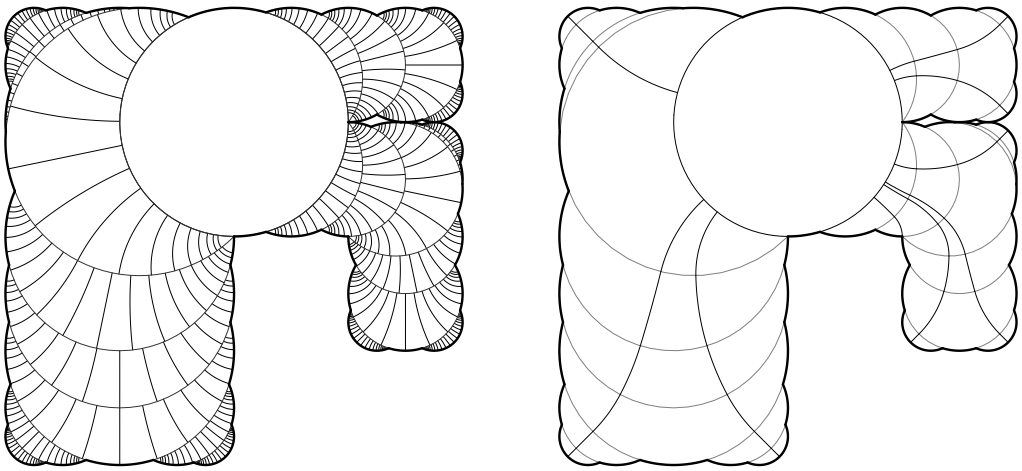

FIGURE 2. A finite union of disks written a union of a root disk D_0 and several crescents. Each crescent is foliated by circular arcs orthogonal to the boundary and following the foliation gives a map from $\partial\Omega$ to ∂D_0.

perpendicular to the first (the medial axis flow). Following leaves of the medial axis flow gives an identification between the two boundary arcs of the crescent (this is the same as applying an elliptic Möbius transformation which fixes each vertex and rotates one boundary arc to the other). Since Ω is written as a union of a root disk and a finite union of crescents, we can compose the maps on each crescent and get a map from $\partial\Omega$ to ∂D_0. We call this the medial axis flow from $\partial\Omega$ to ∂D_0. See Figure 2.

If Ω is bounded by an n-gon then the medial axis is still a finite tree and has $O(n)$ vertices and edges. The edges correspond to 1 parameter families of disks which meet $\partial\Omega$ at exactly two points and come in three types:

(1) a line segment that is equidistant from two vertices
(2) a line segment that is equidistant from two edges
(3) a parabolic arc that is equidistant from a edge and a vertex.

For each edge of the medial axis, consider the union of medial axis disks centered on that edge, minus the disk corresponding to the endpoint closer to the root. We call these subdomains "generalized crescents" (in case (1) it is an actual crescent). Together with the root disk, the generalized crescents decompose Ω into a finite

number of pieces and on each piece there is a foliation by boundary arcs of medial axis disks and a corresponding orthogonal flow; the medial axis flow for that piece. Figure 3 shows this foliation and flow for different types of generalized crescents, Figure 4 shows the decomposition, foliation and flow for a polygon and Figure 5 shows just the foliation and flow for two more examples.

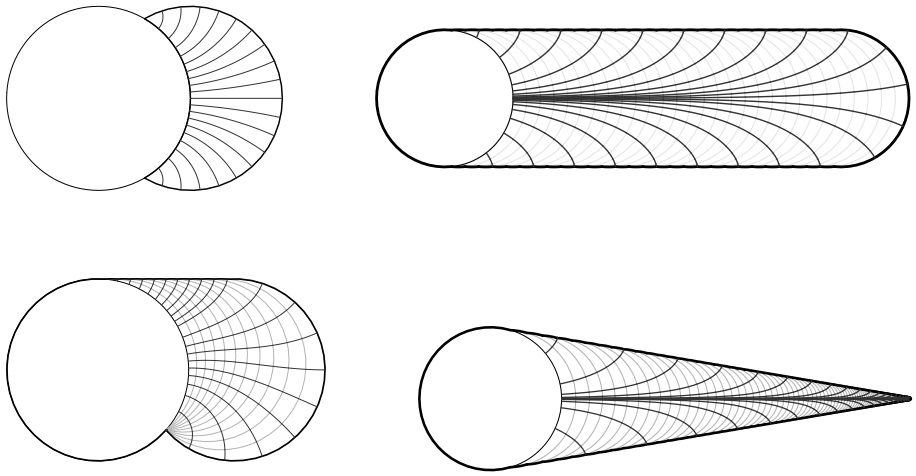

FIGURE 3. The medial axis foliation and flow corresponding to different types of medial axis edges. The white disk in each picture corresponds to the endpoint closer to the root.

To visualize where the vertices of the polygon are mapped we simply follow the medial axis flow from $\partial\Omega$ to ∂D_0. However, to actually compute these images, we start by defining a map from each medial axis vertex disk to its parent.

Suppose e is an edge of the medial axis, D, D^* are the disks corresponding to its endpoints and D^* is the parent of D (i.e., is closer to the root). Let Ω_e be the corresponding generalized crescent. The medial axis flow in Ω_e defines a map from $\partial D \cap \partial\Omega_e$ to $\partial D^* \cap \partial\Omega_e$. For type 1 edges this is just an elliptic Möbius transformation (with fixed points $\partial D \cap \partial D^*$), and we shall see that it is a Möbius transformation $\tau_D : D \to D^*$ in all cases. Moreover, we shall give explicit formulas τ_D in terms of D, D^* and the type of edge. The medial axis flow map from D to the root, $\iota_D : D \to D_0$ is the composition of the maps τ_{D_k}, where $\{D_k\}$ is a path from D to

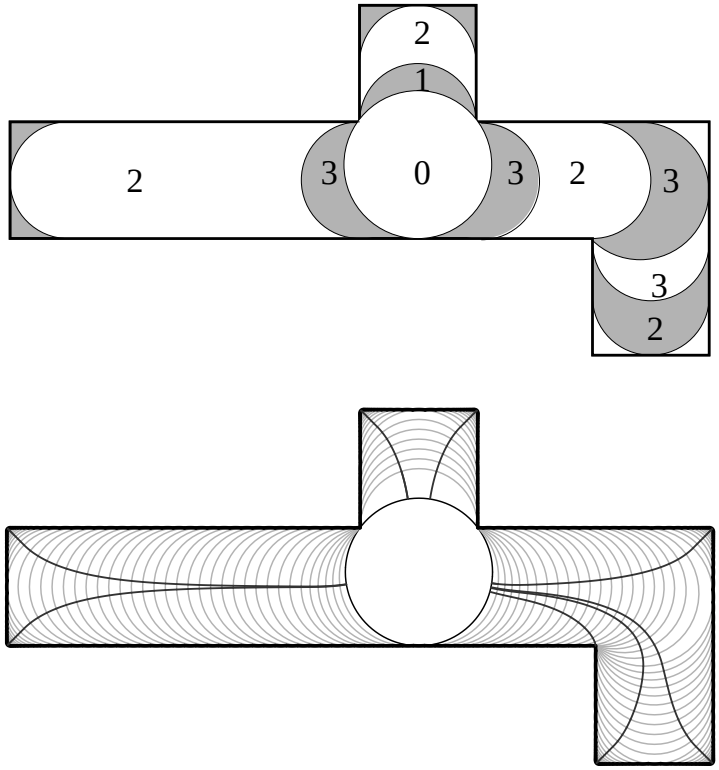

FIGURE 4. This shows the medial axis flow for the polygon illustrated on the right of Figure 1. The top picture shows the a root disk and the generalized crescents, labeled by their type. The bottom figure shows the medial axis flow. The landing points of the vertices are the parameter guesses.

D_0. This map can easily be computed in linear time by induction and the formula

$$\iota_D = \iota_{D^*} \circ \tau_D,$$

if the disks are arranged in a list starting with the root (for which ι_{D_0} is the identity map) and so that every disk comes later that its parent on the list. Since each map is a Möbius transformation, we only have to record a 2×2 matrix for each disk and the composition above corresponds to matrix multiplication.

For each concave vertex v of the polygon (i.e., those with interior angle $\leq \pi$), we will choose a medial axis disk D with $v \in \partial D$ and define $\iota(v) = \iota_D(v)$. If v is a convex vertex then it is the endpoint of a medial axis edge whose other endpoint

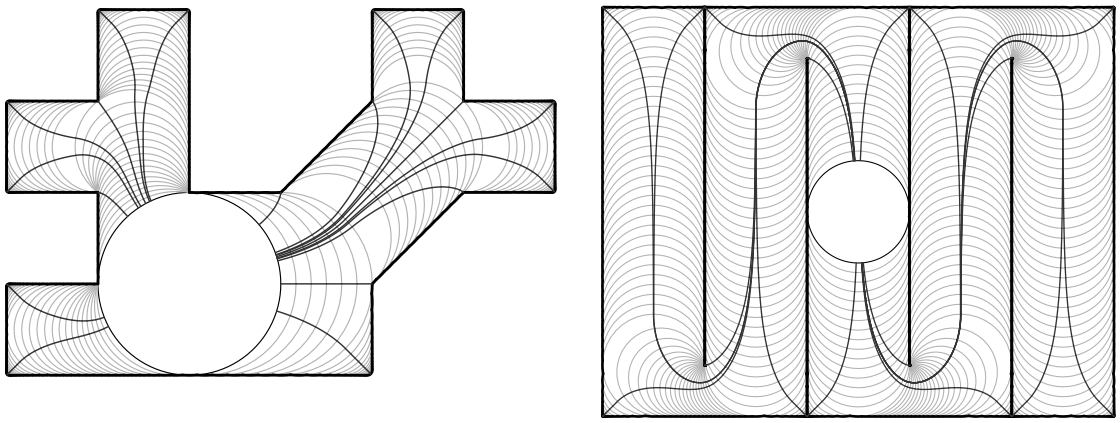

FIGURE 5. More examples of the medial axis foliation and flow.

corresponds to a disk D. We will map v to the closest point v^* of ∂D and then define $\iota(v) = \iota_D(v^*)$. This defines ι at every vertex and gives the parameter guesses $\mathbf{w}$ used in Theorem 1.

I was originally motivated to look at this problem by considering the CRDT algorithm (cross ratios and Delaunay triangulations) of Driscoll and Vavasis for computing conformal maps. CRDT is an iterative scheme for computing the Schwarz-Christoffel parameters which works well in practice, but is not proven to always converge. In [5], I show that the initial guess of CRDT also satisfies condition (1) of Theorem 1, although the method is not $O(n)$, as here. In [6], I extend the results of this paper to give a more complicated algorithm that will compute the Schwarz-Christoffel parameters to within ϵ in the d_{QC} metric in time $O(n|\log\epsilon\log\log\epsilon|)$.

The rest of the paper is structured as follows:

Section 2: We give the definition of the ι map and describe its connection to hyperbolic 3-dimensional geometry. We deduce part (1) of Theorem 1

Section 3: We prove the length decreasing property of ι.

Section 4: We give a formula for ι on each type of generalized crescent.

Section 5: We give explicit bounds for the d_{QC} distance of our guess to the true parameters in some exmaples and make some concluding remarks.

2. The Sullivan-Epstein-Marden theorem

The hyperbolic metric on $\mathbb{D}$ is given by $d\rho_{\mathbb{D}} = 2|dz|/(1-|z|^2)$. The hyperbolic metric on the upper half-space $\mathbb{R}^3_+$ is $d\rho_{\mathbb{R}^3_+} = |dz|/\mathrm{dist}(z, \mathbb{R}^2)$. In both cases, geodesics are circles orthogonal to the boundary, and orientation preserving isometries are exactly the Möbius transformations.

Theorem 1 claims that we can quickly produce guesses for the Schwarz-Christoffel parameters that are only a bounded distortion of the true parameters. What is a good way to measure the amount of this distortion? There is no canonical way to choose the parameters; given one n-tuple $\mathbf{z} \in \mathbb{T}$, any image of $\mathbf{z}$ under a Möbius transformation of the disk (i.e., a linear fractional transformation) is an equally good set of parameters for the same polygon, since these transformations are exactly the conformal selfmaps of the disk. Thus we measure the distance between two n-tuples by seeking the "most conformal" map of the disk to itself whose boundary extension sends $\mathbf{z}$ to $\mathbf{w}$. By our previous remark, this is the same as seeking the "most isometric" map with respect to the hyperbolic metric.

One possibility is to consider biLipschitz (BL) maps with respect to the hyperbolic metric, i.e., maps $f : \mathbb{D} \to \mathbb{D}$ that satisfy

$$\frac{1}{A}\rho(x,y) \le \rho(f(x), f(y)) \le A\rho(x,y).$$

A strictly larger class of maps are the quasi-isometries (QI) which are "biLipschitz at large scales", i.e., satisfy

$$\frac{1}{B}\rho(x,y) - C \le \rho(f(x), f(y)) \le B\rho(x,y) + C,$$

for some constants B, C. A homeomorphism $f : \mathbb{D} \to \mathbb{D}$ is quasiconformal (QC) if for every x

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

Clearly, this class contains the biLipschitz maps and (less clearly) is contained in the quasi-isometries. See [16].

A homeomorphism of the unit circle to itself is called quasisymmetric (QS) if it maps adjacent arcs of equal length to arcs of comparable length, i.e.,

$$\frac{1}{M} \le \frac{|f(x) - f(y)|}{|f(y) - f(z)|} \le M,$$

whenever $x, y, z \in \mathbb{T}$ and $|x - y| = |y - z|$. Each of the three classes (BL, QC, QI) have well defined boundary extensions which are QS and every QS homeomorphism of the circle can be extended to a BL mapping of the disk (which is automatically also QC and QI). Therefore, when we want to measure how close an n-tuple $\mathbf{w} \in \mathbb{T}$ is to being a conformal image of another n-tuple $\mathbf{z} \in \mathbb{T}$ we can try to minimize the constant among maps sending $\mathbf{z}$ to $\mathbf{w}$ in any of these classes. Theorem 1 is stated in terms of the quasiconformal constant K, since the the best explicit estimates are available in this case.

If f is C^1, then we can define the partial derivatives $f_z = \frac{1}{2}(f_x - if_y)$ and $f_{\bar{z}} = \frac{1}{2i})f_x + if_y)$ and the Beltrami coeficient $\mu_f = f_{\bar{z}}/f_z$. Then f is K-quasiconformal if and only if $|\mu_f|$ is bounded above by $(K-1)/(K+1)$. See [2]. We will use this in the last section in order to compute the QC constant for some piecewise affine maps between polygons.

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 hyperbolically 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 base disks 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. S is the boundary of the union of all hemispheres centered on $\mathbb{R}^2$ whose bases are disks contained in Ω. In fact, we do not have to consider all subdisks of Ω. It is not hard to show that every point on the dome of Ω is on a unique hemisphere whose base is a medial axis disk of Ω (e.g., see Lemma 3 of [6]). Thus each point on the dome corresponds to some point of the medial axis. See Figure 6.

Let ρ_S denote the intrinsic path metric on S (using hyperbolic arclength). The most important facts about the dome of Ω are the following two theorems.

Theorem 2 (Thurston)**.** *There is an isometry ι from (S, ρ_S) to $(\mathbb{D}, \rho)$.*

Theorem 3 (Sullivan-Epstein-Marden)**.** *Suppose Ω is a hyperbolic simply connected domain (other than the complement of a ray) and suppose $f : \mathbb{D} \to \Omega$ is conformal. Let S denote the dome of Ω. There is a map $\sigma : \Omega \to S$ which extends continuously*

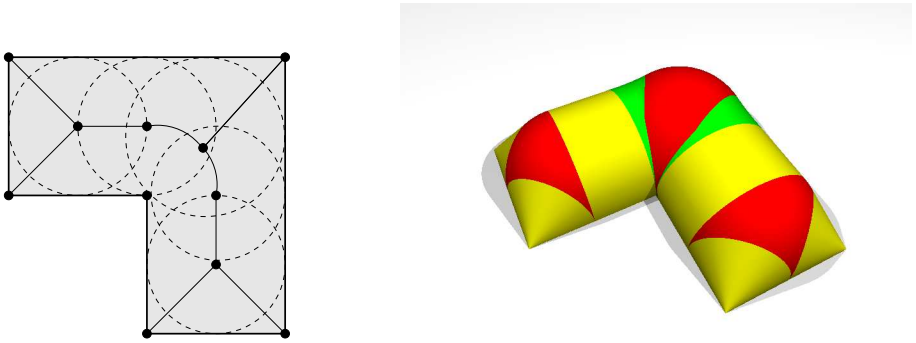

FIGURE 6. On the left is a polygon and its medial axis. The dashed circles correspond to vertices of the medial axis. On the right is the dome of the polygon, with sections shaded according to the part of the medial axis they correspond to.

to the identity on $\partial\Omega$ *and so that* $\varphi = \iota \circ \sigma \circ f : \mathbb{D} \to \Omega \to S \to \mathbb{D}$ *is a* K_0 *biLipschitz with respect to the hyperbolic metric and* K_0 *is independent of* Ω*. Consequently, there is a universal* $K < \infty$ *so that* φ *is* K*-quasiconformal.*

The map ι restricted to $S \cap \mathbb{R}^2 = \partial\Omega$ is the map used in Theorem 1, i.e., we take $\mathbf{w} = \iota(\mathbf{v})$ as out guess for the conformal prevertices. The true prevertices are given by $\mathbf{z} = f^{-1}(\mathbf{v})$ and hence $\varphi(\mathbf{z}) = \mathbf{w}$. Thus Theorem 3 implies part (1) of Theorem 1 with constant K.

Theorem 2 appears in Thurston's notes [25], with more detailed proofs in the papers of Epstein-Marden [13] (updated and reprinted in [14]) and Rourke [22]. The second was apparently known to Thurston and appeared in Sullivan's paper [24] in the case when Ω is invariant under a convex co-compact group of Möbius transformations. Epstein and Marden [13] proved the more general statement quoted above. Alternate proofs are given in [3], [7], [15]. The complement of a ray is only exceptional because in that case S needs to be interpreted as a two-sided "folded" surface. This is carefully explained in [17]. Epstein and Marden's proof of Theorem 3 in [13] 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 [7] it is proven that one can take $K = 7.82$. This gives the "8" in Theorem 1. More recently Epstein and Markovic [18] have shown that $K > 2.1$, by showing the K is at least this large when Ω is the complement of a certain logarithmic spiral. This bound is for the particular map given by Theorem

2. Possibly a different definition would lead to a map with a better quasiconformal estimate. However, it is impossible to take $K < 2$ and still have (3) of Theorem 1, as a slit disk shows. It is shown in [4] that $K = 2$ and (3) together imply the well known Brennan conjecture.

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. If two faces meet, they meet along infinite geodesic and so the faces form the vertices of a tree. Each face lies on the dome of a disk which is a vertex of the medial axis, and the tree structure of the faces agrees with the tree structure of the medial axis.

It is easy to visualize the ι map for a finite union of disks. Start with the case of two overlapping disks, as shown in Figure 7. The dome is the union of two geodesic faces, each sitting on a hemisphere, and joined along an infinite hyperbolic geodesic. There is a 1-parameter family of elliptic Möbius transformations which leave this geodesic fixed and rotate around it. Using these maps, we can leave one face fixed and rotate the other until it is "flush" with the first one, i.e, we map the dome to a hemisphere by a map which is the identity on one face and a hyperbolic isometry on the other. This map is therefore an isometry of the hyperbolic path metric on the dome to the path metric on the hemisphere (which equals the usual hyperbolic metric and is isometric to the hyperbolic disk). Note that points on the boundary are moved by an elliptic transformation of the plane.

If Ω is a finite union of disks, then its dome is finitely bent. Choose a root face and starting at the leaves of the tree of faces, rotate each face of the dome until it becomes flush with its parent. This gives an isometry of the dome of Ω to the dome of a disk, which is isometric to the hyperbolic disk, thus giving ι. The elliptic transformation which rotates each face also acts on the plane and agrees with the medial axis flow we described earlier. Thus for finite unions of disks, the medial axis flow gives the same boundary map as Thurston's ι map of the dome. This flow was illustrated in Figure 2. The corresponding dome is shown in Figure 8.

Given a generalized crescent, Ω_e, corresponding to an edge e of the medial axis of a polygon, we can approximate it by a finite union of actual crescents by taking a union of medial axis disks with centers along the edge e. See Figure 9. For the

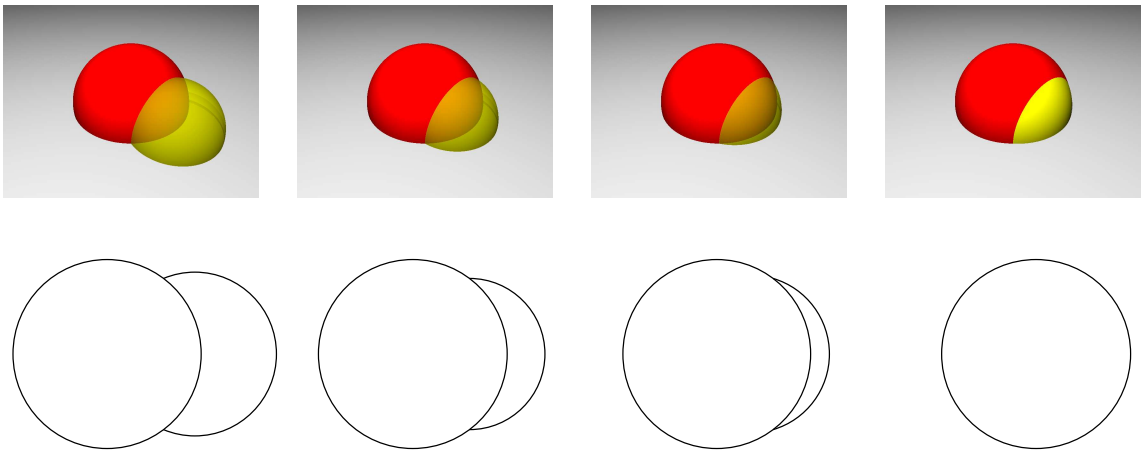

FIGURE 7. A dome consisting of two geodesic faces joined along an infinite geodesic. By bending the dome along the geodesic we get a one-parameter, isometric family of surfaces ending with a hemisphere, which is obviously isometric to the hyperbolic disk. Underneath we show the base domains in the plane. The crescent is being collapsed by an elliptic Möbius transformation.

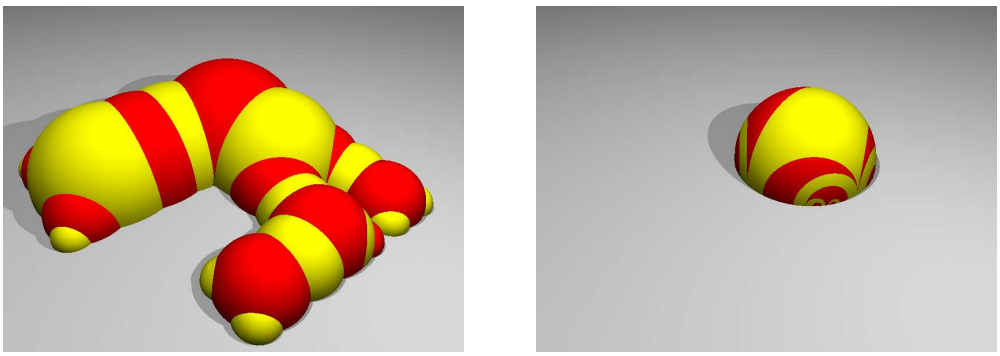

FIGURE 8. This is is the dome of a finite union of disks and its isometric image after all the faces have been rotated along bending geodesics to become flush with one another. The crescent decomposition and medial axis flow for this example were shown in Figure 2.

approximation, the flow map from the outer to inner boundary arcs is a composition of Möbius transformations and hence is Möbius. Since a (non-degenerate) limit of Möbius transformations is also Möbius, we see that the flow map for a generalized crescent is also a Möbius transformation between the medial axis disks corresponding to the endpoints of the medial axis edge e.

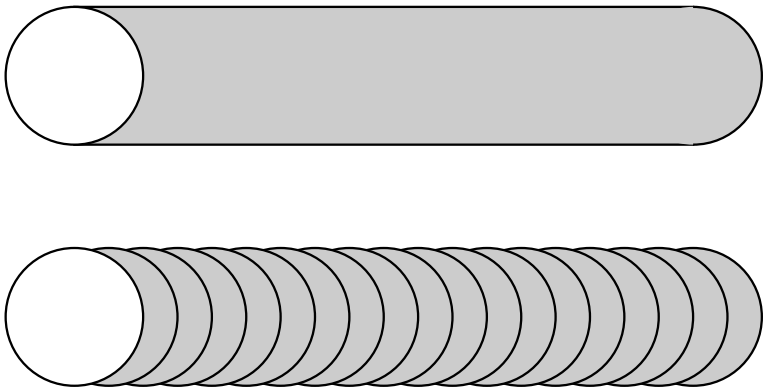

FIGURE 9. Each generalized crescent can be approximated by a union of actual crescents and the corresponding flow map is the limit of a composition of elliptic Möbius transformations, hence is Möbius.

3. ι DECREASES LENGTH

First suppose that Ω is a finite union of disks (the general case will follow from a limiting argument). We have defined ι as a composition of elliptic Möbius transformations on each circular arc in $\partial\Omega$. Each of these transformations is associated to a crescent as described earlier. 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 4. *Suppose Ω is a crescent which lies on one side of the line L passing through its 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 the elliptic Möbius transformation fixing the two vertices and mapping γ_2 to γ_1 then $|\tau'(z)| \le 1$ 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 $-d/c$ passing through -1 and 1 (since the derivative of an elliptic transformation has modulus one at the two fixed points). Let γ be the arc of this circle between 1 and -1 which lies in the upper half-plane. We claim that γ_2, the upper edge of our crescent, lies above γ.

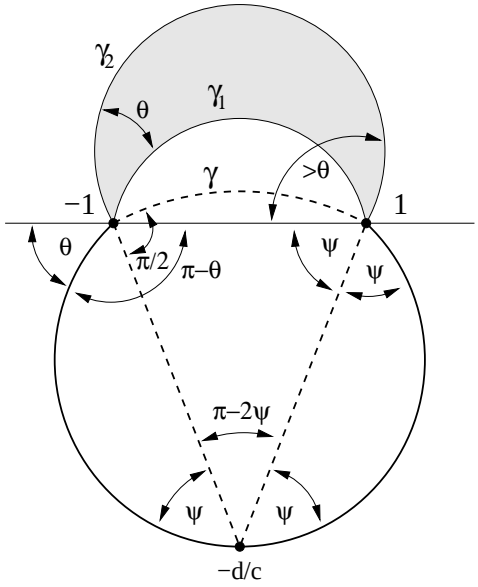

FIGURE 10. The setup in Lemma 4. We prove that γ_2 lies above γ by showing that γ makes angle $\theta/2$ with $[-1,1]$, but γ_2 makes angle $> \theta$ with the same segment.

Suppose the elliptic transformation is a rotation by θ around the points $-1, 1$. Since γ_2 and its image are both in the upper half-plane, $\theta < \pi$. Therefore $-d/c$ lies on a circle which makes angle $\pi - \theta$ with the segment $[-1,1]$. See Figure 10. Hence the isosceles triangle with base $[-1,1]$ and vertex $-d/c$ has two base angles of $\psi = (\pi - \theta)/2$ and the circle C makes angle $\pi/2 - \psi = \theta/2$ with $[-1,1]$. Since γ_1 lies in $\mathbb{H}$, γ_2 makes angle of at least θ with $[-1,1]$ and hence lies above C. □

This implies that for finite unions of disks, the map $\iota : \partial\Omega \to \partial D_0$ can be chosen to have derivative < 1 on $\partial\Omega$ (except possibly at the vertices). By taking approximations and passing to the limit we obtain:

Lemma 5. *If Ω is simply connected polygonal domain and D is a medial axis disk of Ω, the map $\iota : \partial\Omega \to \partial D$ does not increase the length of any boundary arc.*

4. Formulas for ι

As noted in the introduction, to compute ι for a polygon, we only have to compute the Möbius transformation for each generalized crescent. Let $\mathcal{D}$ denote the disks that correspond to vertices of the medial axis and assume we have chosen a root disk, D_0.

Below, we will describe how to map a disk $D_1 \in \mathcal{D}$ to its parent $D_2 \in \mathcal{D}$. The map from D_1 to the root disk can then be computed by composing with the corresponding map for the parent (which we may assume has already been computed). To map the vertices of $\partial\Omega$ to ∂D_0, map there are two cases. For concave vertices v (interior angle $\geq \pi$) choose a medial axis vertex disk D so that $v \in \partial D$ and apply the map ι_D to v. If v lies on the boundary of more than one such disk, each of the corresponding maps must agree at v, so it does not matter which disk we choose. Each convex vertex (interior angle $< \pi$) it the endpoint of a type 2 edge of the medial axis. The other endpoint is a vertex of the medial axis and we map v onto the boundary of the corresponding disk D by simply mapping it to the closest point of ∂D and then applying ι_D to map the point to ∂D_0.

We now consider each type of medial axis edge.

Case 1 (point-point bisector): If the two disks are $D_1 = D(z_1, r_1)$ and $D_2 = D(z_2, r_2)$ 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 $\partial D_2 \cap \overline{D_1}$ by an elliptic rotation τ of angle α, the interior angle of the crescent. If $r = |z_1 - z_2|$, law of cosines implies

$$\cos\alpha = \beta = \frac{r^2 - r_1^2 - r_2^2}{-2r_1 r_2},$$

and hence $e^{i\alpha} = \beta + i\sqrt{1-\beta^2}$. Similarly,

$$a = z_2 + \frac{r_2}{r} e^{i\theta}(z_1 - z_2), \tag{1}$$

$$b = z_2 + \frac{r_2}{r} e^{-i\theta}(z_1 - z_2), \tag{2}$$

where $\cos\theta = (r_1^2 - r^2 - r_2^2)/(-2rr_2)$. Thus τ is given by the formula $\tau(z) = \sigma^{-1} \circ R_\alpha \circ \sigma$, where $R_\alpha(z) = e^{i\alpha} z$ and $\sigma(z) = (z-b)/(z-a)$ sends a to ∞ and b to zero.

Case 2a (edge-edge bisector, parallel case): Consider Figure 11. Suppose the two endpoint disks have centers z_1, z_2 and common radius $r_1 = r_2$. We normalize so that these disks become $D_1 = D(0,1)$ and $D_2 = D(A,1)$ using the map

$$\eta(z) = (z - z_1) \cdot \frac{|z_2 - z_1|}{(z_2 - z_1) r_1}.$$

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.

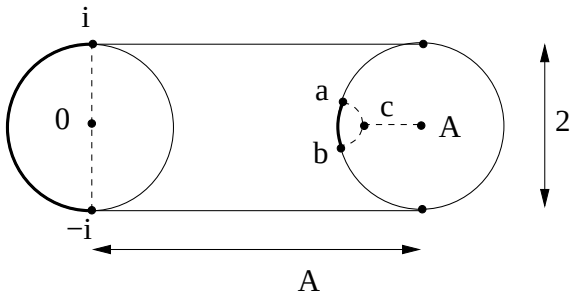

FIGURE 11. Case 2a: between parallel lines

With the points labeled as shown, i and $-i$ will map to a and b respectively and 0 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

$$w = \frac{1-e^{-A}}{1+e^{-A}}. \tag{3}$$

Thus

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

and this determines τ. Thus $\eta^{-1} \circ \tau \circ \eta$ is the desired map.

If we consider the case when $\Omega = \{x+iy : -1 < y < 1\}$ is an infinite strip, we can compute the hyperbolic metric exactly by mapping the strip to the right half-plane by the exponential map $z \to \exp(\pi z/2)$. In the half-plane the hyperbolic distance between two points $x < y$ on the real axis is $\log(y/x)$, so we see that the distance between 0 and A in the strip is $(\pi/2)A$. However, the points on the dome which project to 0 and A are hyperbolic distance A apart. Thus we should get an even

better approximation to the Riemann map map by replacing A by $(\pi/2)A$ in the definition above. We will refer to the corresponding map as the "modified ι map".

Case 2b (edge-edge bisector, non-parallel case): The situation is shown in Figure 12. We will assume D_1 is smaller than D_2; the opposite case is handled similarly. Suppose the endpoint disks are $D(z_1, r_1)$ and $D(z_2, r_2)$ with $r_1 < r_2$ and normalize by a linear map η so that they become $D_1 = D(1, B)$ and $D_2 = D(y, yB)$ for some $B > 0, y > 1$. This can be done with the map

$$s = r_1|z_1 - z_2|/(r_2 - r_1)$$
$$z_3 = z_1 + s(z_1 - z_2)/|z_2 - z_1|$$
$$\eta(z) = (z - z_3)/(z_1 - z_3).$$

Thus $B = r_1/|z_1 - z_3|$, $y = |(z_2 - z_3)/(z_1 - z_3)|$. The points labeled a, b in Figure 12 are mapped to c, d respectively. The points on the dome above 1 and y are joined by a geodesic which projects to the straight line between 1 and y. Above a point $t \in [1, y]$ the dome has height tr so the length of this geodesic on the dome is

$$A = \int_1^y \frac{dt}{tB} = \frac{\log y}{B}.$$

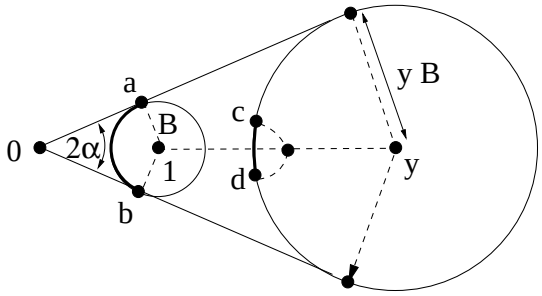

FIGURE 12. Case 2b: between non-parallel lines

Hence the image of 1 in D_2 is hyperbolic distance A from y and must be on the segment $[1, y]$. Thus by (3) 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

$$\tau(z) = yB\sigma(\frac{1}{B}(z - 1)) + y.$$

If $r_1 > r_2$ then the mapping is almost same, except that now $0 < y < 1$, so $A = \log y/r$ is negative. Thus we define $w = -(1 - e^A)/(1 + e^A)$ and proceed as before.

As in the case of parallel edges, we can compute the hyperbolic metric exactly for an infinite cone. Doing so leads shows that we should replace A by $A = \frac{\pi}{2\alpha} \log y$ to get the correct map for a cone. If we make this change we will call the resulting map the modified ι map, as before.

Case 3 (point-line bisector): This case is pictured in Figure 13. Here we have a parabolic edge of the medial axis which is the bisector of a line segment and a point (a vertex of the polygon). There are actually two cases depending on the orientation of the picture: as we traverse the medial axis edge towards the root, the vertex can either be on our left or on our right. We will assume it is one the right, as illustrated in Figure 13, and the other case is handled similarly.

Suppose a is the vertex in question (computed from D_1 and D_2 as in (1)) and L is the line. Then L is tangent to D_1 and D_2 at points c, d given by

$$c = z_1 + r_1 e^{i\theta} \frac{z_2 - z_1}{r}, \qquad d = z_2 + r_2 e^{i\theta} \frac{z_2 - z_1}{r},$$

where $\cos\theta = (r_1 - r_2)/r$ and $r = |z_2 - z_1|$, as before. Let a^* be the reflection of a across L. This is given by

$$a^* = c + (\bar{a} - \bar{c}) \cdot (d - c)/(\bar{d} - \bar{c}).$$

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. Applying each of these rotations in turn produces a sequence of disks $\{B_k\}$ intermediate between D_1 and D_2. Let $\tau_k : B_k \to B_{k+1}$ be the elliptic map that fixes the two points of intersection $\{a, b_k\} = \partial B_k \cap \partial B_{k+1}$. The generalized crescent being considered is the limit of the union of these disks and so $\tau_n \circ \cdots \circ \tau_1$ converges to the desired $\tau : D_1 \to D_2$ as $n \to \infty$.

To understand this composition better, we will conjugate by a Möbius transformation that converts the maps τ_k into linear maps. To do this, we have to map one of the fixed points to ∞. Consider the map

$$\eta(z) = (-i)(\frac{z - a^*}{z - a})(\frac{c - a}{c - a^*}).$$

This sends $a \to \infty$, $a^* \to 0$ and $c \to -i$. Moreover, the line L maps to the unit circle and each of the disks B_k is sent to a half-plane tangent to the unit circle. The maps τ_k must therefore be sent to rotations by angle θ/n around the points where the

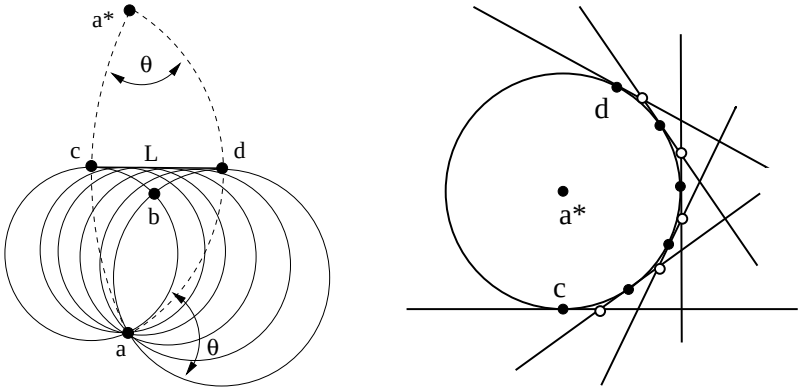

FIGURE 13. Case 3: equidistant from a point and a line. On the left is the family of circles passing through a and tangent to L. On the right we have conjugated a to ∞ and L to the unit circle. The converted τ_k maps are rotations around the white dots.

boundaries of the kth and $(k+1)$st half-planes meet. In the limit, the composition of these maps is simply "rolling" a line counterclockwise around the unit circle, i.e., is $\sigma(z) = e^{i\theta}(z-\theta)$. Thus $\tau = \eta^{-1} \circ \sigma \circ \eta$.

If we are in the other case, when the vertex is on the left, then we define the point a using (2) instead of (1), Similarly c, d are now defined using $-\theta$ in place of θ. Then define η as before. The σ map now rolls a line clockwise around the unit circle, i.e., $\sigma(z) = (z+\theta)e^{-i\theta}$. We have now completed the proof of Theorem 1.

Estimates of hyperbolic distance lead to a modification of ι in this case as well. Take $\sigma(z) = e^{\pm i\theta}(z \mp 2\theta)$ for the modified map.

5. Concluding remarks

Theorem 1 gives guesses for the Schwarz-Christoffel parameters of a polygon which are close to the correct parameters in a precise sense. If we put our guesses into the Schwarz-Christoffel formula we get a locally 1-1 map of the disk to a polygonal region (possibly overlapping itself) which has all the correct angles and which can be mapped to the desired region by a quasiconformal map with uniformly bounded constant and with vertices mapping to vertices. Figure 14 shows five such polygons and the corresponding regions using the ι map and modified ι map to guess the parameters. (The curved "S" shape is actually a 100-gon.)

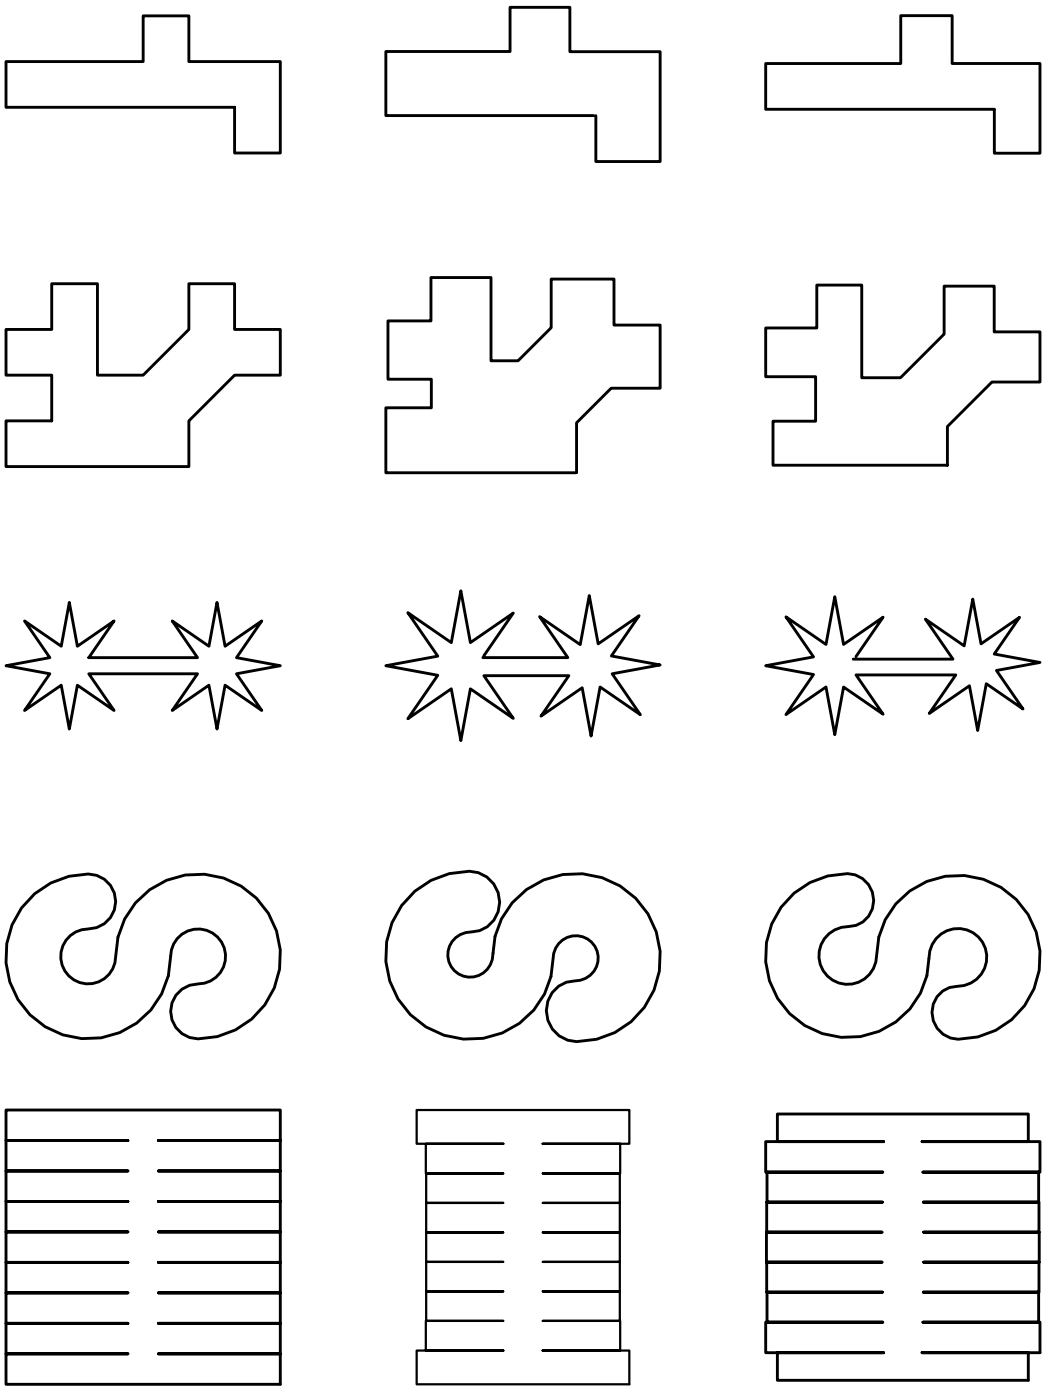

FIGURE 14.

As expected, the figures show that the ι map tends to shorten long corridors and this is partially corrected by the modified map. We can measure this distortion by estimating the quasiconformal distance d_{QC} between our guesses and the actual prevertices as follows. This metric was defined by minimizing the quasiconformal constant over all self-maps of the disk which send the true prevertices to our guesses. Thus any such map gives an upper bound for the distance. Suppose Ω is our target polygon and Ω^* is the Schwarz-Christoffel image using our ι map guesses, and let f and f^* be the conformal maps from the disk onto these domains. If Φ is any quasiconformal map from Ω to Ω^* which maps vertices to vertices, then $\varphi = (f^*)^{-1} \circ \Phi \circ f$ is a quasiconformal map of the disk to the disk which sends the true prevertices to our guesses and has the same quasiconformal constant as Φ does (since composing by conformal maps does not change this). Moreover, there is a an obvious way to try to construct the desired map Φ: triangulate Ω and Ω^* in equivalent ways and use affine maps on each triangle.

Given two triangles in the plane with the same orientation, it is easy to compute the quasiconformal constant of the affine map that sends the three vertices of the first to the three vertices of the second. We can first use a conformal linear map to send each triangle to ones of the form $\{0, 1, a\}$ and $\{0, 1, b\}$. The affine map which fixes 0 and 1 and sends a to b is of the form $f(z) \to \alpha z + \beta \bar{z}$ where $\alpha + \beta = 1$ and $\beta = (b - a)/(a - \bar{a})$ and from this it is easy to compute that

$$\mu_f = \frac{f_{\bar{z}}}{f_z} = \frac{\beta}{\alpha} = \frac{b - a}{b - \bar{a}},$$

and so the quasiconformal constant is

$$K_f = \frac{1 + |\mu_f|}{1 - |\mu_f|}.$$

We compute this K for each element of our triangulation and take the maximum, which must be an upper bound for the d_{QC} distance between our guessed parameters and the true values. In general, this will be a strict upper bound, since the optimal map sending the true prevertices to our guesses will not correspond to a piecewise affine map. Moreover, a different triangulation might give a better estimate. Figure 15 shows triangulations for two of the polygons from Figure 14. Doing the computations above gives $K = 1.6988$ (for the usual ι map) and $K^* = 1.24475$ (for the modified ι map) for the polygon on the top row. The polygon on the bottom has

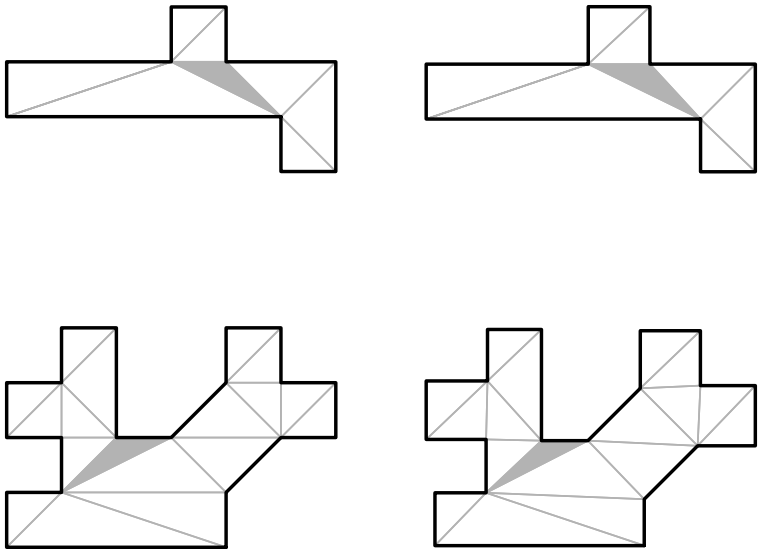

FIGURE 15. On left is the triangulation of the original polygon and on the right the triangulation of the Schwarz-Christoffel image using the modified ι map. The most distorted triangle is shaded.

bounds of $K = 2.80955, K^* = 1.23762$ for the two verions of the ι map. In both cases, the constant K is smaller than the theoretical upper bound of 7.82 and the modified map is better the unmodified map. (If we modify ι by only altering the maps for the edge-edge bisectors, but not for the edge-point bisectors, then the bounds for the two polygons become $K^* = 1.176455$ and $K^* = 2.36412$ respectively; this is better for the first but worse for the second. Thus it is not always clear how much modifying to do.)

Can we use these bounds to define an iteration which will converge to the correct parameters (i.e. decrease K to 1)? Note that in the definition of K_f we use only the modulus of μ_f, which measures the degree of stretching between the triangles. The argument of μ_f describes the direction of this stretching; can this information be used to modify our parameter guesses more accurately?

References

[1] A. Aggarwal, L. J. Guibas, J. Saxe, and P. W. Shor. A linear-time algorithm for computing the Voronoĭ diagram of a convex polygon. *Discrete Comput. Geom.*, 4(6):591–604, 1989.

[2] 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.

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

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

[5] C.J. Bishop. Bounds for the CRDT algorithm. preprint, 2008.

[6] C.J. Bishop. Conformal mapping in linear time. preprint, 2008.

[7] C.J. Bishop. An explicit constant for Sullivan's convex hull theorem. In *In the tradition of Ahlfors and Bers, III*, volume 355 of *Contemp. Math.*, pages 41–69. Amer. Math. Soc., Providence, RI, 2004.

[8] 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.

[9] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete Comput. Geom.*, 6(5):485–524, 1991.

[10] 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.

[11] 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.

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

[13] 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.

[14] D. B. A. Epstein and A. Marden. Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces [mr0903852]. In *Fundamentals of hyperbolic geometry: selected expositions*, volume 328 of *London Math. Soc. Lecture Note Ser.*, pages 117–266. Cambridge Univ. Press, Cambridge, 2006.

[15] D. B. A. Epstein, A. Marden, and V. Markovic. Complex angle scaling. In *Kleinian groups and hyperbolic 3-manifolds (Warwick, 2001)*, volume 299 of *London Math. Soc. Lecture Note Ser.*, pages 343–362. Cambridge Univ. Press, Cambridge, 2003.

[16] D. B. A. Epstein, A. Marden, and V. Markovic. Quasiconformal homeomorphisms and the convex hull boundary. *Ann. of Math. (2)*, 159(1):305–336, 2004.

[17] D. B. A. Epstein, A. Marden, and V. Markovic. Complex earthquakes and deformations of the unit disk. *J. Differential Geom.*, 73(1):119–166, 2006.

[18] D. B. A. Epstein and V. Markovic. The logarithmic spiral: a counterexample to the $K = 2$ conjecture. *Ann. of Math. (2)*, 161(2):925–957, 2005.

[19] P. Erdös. On the Hausdorff dimension of some sets in Euclidean space. *Bull. Amer. Math. Soc.*, 52:107–109, 1946.

[20] N. M. Patrikalakis and T. Maekawa. *Shape Interrogation for Computer Aided Design and Manufacturing*. Springer Verlag, 2002.

[21] 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.

[22] C. Rourke. Convex ruled surfaces. In *Analytical and geometric aspects of hyperbolic space (Coventry/Durham, 1984)*, pages 255–272. Cambridge Univ. Press, Cambridge, 1987.

[23] 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.

[24] 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.

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

[26] 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.

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

[28] 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.

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

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