

BOUNDS FOR THE CRDT CONFORMAL MAPPING ALGORITHM

CHRISTOPHER J. BISHOP

ABSTRACT. In [16] Driscoll and Vavasis described their CRDT algorithm for numerical computation of conformal maps. Although the method works well in practice, it is not proven to converge and they formulated a number of conjectures to quantify its behavior. In this paper we prove one of these conjectures: that the initial guess of the CRDT algorithm always lies within a uniformly bounded distance of the correct answer.

Date: February 11, 2010.

1991 Mathematics Subject Classification. Primary: 65E05 , Secondary: 30C30, 30C35, 30C62 .

Key words and phrases. numerical conformal mapping, CRDT algorithm, Schwarz-Christoffel formula, hyperbolic geometry, quasiconformal mappings, medial axis .

The author is partially supported by NSF Grant DMS 04-05578.

1. INTRODUCTION

The Schwarz-Christoffel formula says that a conformal map from the unit disk \mathbb{D} to the interior of a polygon P has the general form

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

where $\{\alpha_1\pi, \dots, \alpha_n\pi\}$ are the interior angles at the vertices $\mathbf{v} = \{v_1, \dots, v_n\}$ of P , and $\mathbf{z} = \{z_1, \dots, z_n\} = f^{-1}(\mathbf{v})$ are the preimages of the vertices (also known as the Schwarz-Christoffel parameters). The formula has been known since the 1860's ([11], [29]), but because of the difficulty of finding the parameters, it only became an effective tool for numerical conformal mapping in the 1980's. For a brief history of the formula and its applications see [15].

In [16] Driscoll and Vavasis introduced a new method for computing the Schwarz-Christoffel parameters: the CRDT algorithm (Cross Ratios and Delaunay Triangulations). Given a polygon P , CRDT proceeds in three steps: (1) add extra vertices to P so that the resulting polygon P' has edges that are “well separated”, (2) use the Delaunay triangulation of P' to construct an initial guess for the images of the vertices on the unit circle, \mathbb{T} , and (3) compute a conformal map using the Schwarz-Christoffel formula with the guessed prevertices (and known angles) and modify the guesses based on the result. See Section 7. The final step is iterated until the desired accuracy is achieved.

Driscoll and Vavasis formulated a number of specific questions regarding the behavior of the CRDT algorithm. In particular, they asked if steps (1) and (2) always give an initial guess that is within a bounded distance of the true Schwarz-Christoffel parameters when measured in terms of a metric derived from cross ratios. We will show that their specific formulation is false, but that a slight modification is true; one merely has to replace the cross ratio of four points by the conformal modulus of a related quadrilateral. This version follows immediately from a stronger result where the distance between n -tuples on $\mathbb{T} = \partial\mathbb{D}$ is measured in a quasiconformal sense, i.e., if $\mathbf{w} = \{w_1, \dots, w_n\}$ and $\mathbf{z} = \{z_1, \dots, z_n\}$ we define

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

The definitions and properties of quasiconformal maps are reviewed in Appendix B.

In this paper, P will denote a simple polygon and Ω its interior (so Ω is a simply connected, bounded, open set). We could include polygonal regions with slits, but this would require modifying some definitions and arguments, so for simplicity we will stick to simple polygons.

Theorem 1. *There is a $C < \infty$ (independent of P) so that the initial guess \mathbf{w} of the CRDT algorithm satisfies $d_{QC}(\mathbf{w}, \mathbf{z}) \leq C$, where \mathbf{z} are the true conformal prevertices.*

The theorem is illustrated in Figure 1 that shows a target polygon (on the left) and the Schwarz-Christoffel image using the CRDT initial guess as the parameters (in the center). On the right we show the Schwarz-Christoffel image if we simply take the parameters to be equidistributed on the circle (a common initial guess for some iterative methods).

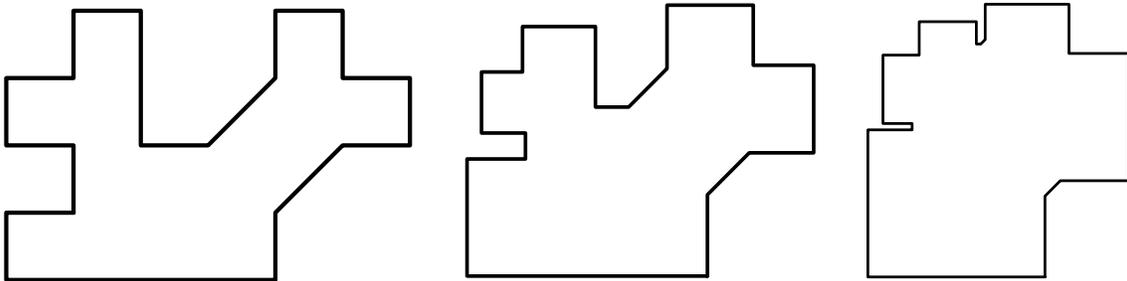


FIGURE 1. On the left is the target polygon. In the center is the Schwarz-Christoffel image using the initial guess of the CRDT algorithm. On the right is the image using equidistributed parameters. We quantify how distorted the polygons are using quasiconformal maps. See Section 7.

The proof of Theorem 1 shows that sometimes we get a good initial guess without adding the extra vertices in step (1) of CRDT. The following is a special case of a more general condition we will state later:

Theorem 2. *Let $\Gamma = \{\gamma_j\}$ denote the diagonals of the Delaunay triangulation of Ω and suppose that any two points $z, w \in \partial\Omega$ can be joined in $\partial\Omega \cup \Gamma$ by a path of length $\leq M|z - w|$. Then $d_{QC}(\mathbf{w}, \mathbf{z}) \leq C(M)$ (notation as in Theorem 1).*

By a triangulation we mean one using only the original vertices (no Steiner points added). Later we will talk about adding extra vertices to P , and will triangulate the

new polygon P' , but will not consider this a triangulation of P itself. The condition in Theorem 2 says the Delaunay triangulation is an M -spanner for P (a set Γ is a M -spanner for E if any two points on E can be connected in Γ length bounded by M times their Euclidean distance apart). It is known that the Delaunay triangulation is always a 2.42-spanner for the vertices of a polygon (see [14], [25]), but the condition in Theorem 2 requires it to be a spanner for all points on P (vertices and edges). Some examples where CRDT adds vertices unnecessarily are discussed in Section 7 (see Figures 27 and 28). CRDT also uses the extra vertices to control crowding (situations where distinct Schwarz-Christoffel parameters may be separated by less than machine precision), and they may still be needed for this reason even when Theorem 2 applies.

Given a generalized quadrilateral, i.e., a Jordan domain Ω and four distinct, ordered points $\mathbf{z} = \{z_1, z_2, z_3, z_4\}$ on the boundary, we can conformally map it to a rectangle $f : \Omega \rightarrow R$ with the four points mapping to the corners, and this rectangle is unique up to Euclidean similarities. We define $\text{Mod}_\Omega(\mathbf{z}) = |f(z_2) - f(z_1)|/|f(z_2) - f(z_3)|$ (the eccentricity of the rectangle). A K -quasiconformal map from one generalized quadrilateral to another can change the modulus by at most a multiplicative factor of K (see [1]). It follows immediately from this fact and Theorem 1 that if P is an n -gon, $\mathbf{z} \subset \mathbb{T}$ are conformal prevertices of P and $\mathbf{w} \in \mathbb{T}$ is the initial guess of the CRDT algorithm, then

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

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

Corollary 3. $-\log K \leq \log \text{Mod}_{\mathbb{D}}(\mathbf{z}') - \log \text{Mod}_{\mathbb{D}}(\mathbf{w}') \leq \log K$.

On page 1792 of [16] 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 show in Section 6 that this approach fails, so Corollary 3 is the appropriate modification of their conjecture. For 4-tuples on the circle, the conformal modulus is a function of the cross ratio and this function is discussed in Appendix A.

Suppose $f : \mathbb{D} \rightarrow \Omega$ is conformal and $g : \Omega \rightarrow \mathbb{D}$ is K -quasiconformal and maps the vertices of P to the CRDT parameter guesses on the circle. Then $g \circ f : \mathbb{D} \rightarrow \mathbb{D}$

is K -quasiconformal and maps the correct parameters to the CRDT guesses. Thus the proof of Theorem 1 reduces to showing that the CRDT initial guess (which is a map from the vertices of P into \mathbb{T}) can be extended to a quasiconformal map g from Ω , the interior of P , to the unit disk \mathbb{D} . To prove this we introduce a new region R and define maps $\Psi : \Omega \rightarrow R$ and $\Phi : R \rightarrow \mathbb{D}$. These maps extend to the boundary and the composition restricted to the vertices of P will agree with the initial guess of the CRDT algorithm. The map Ψ will be quasiconformal with a uniformly bounded constant. This means that it distorts the modulus by at most a bounded factor. The map Φ will not be quasiconformal, but it will be a quasi-isometry between the hyperbolic metrics on R and \mathbb{D} , again with uniformly bounded constants. This implies there is a quasiconformal map $\Omega \rightarrow \mathbb{D}$ with the same boundary values as Φ and a uniformly bounded quasiconformal constant and this is sufficient to prove Theorem 1. The connection between quasiconformal maps and quasi-isometries is reviewed in Appendix B.

A crescent is a domain bounded by two circular arcs meeting at two distinct points. If one of these arcs is a line segment, we call the crescent a “flat crescent”. We will construct R by adding a flat crescent to the outside of each edge of P . Suppose we have a triangulation of P . Each triangle T defines a unique circle passing through its three vertices (we call this the circumcircle corresponding to T and its interior the circumdisk). If e is an edge of P and T is the element of the triangulation that also has e as an edge, then e divides the circumdisk of T into two flat crescents. The one disjoint from T is called the “outer crescent”, and this is the one we attach to P .

The region R is essentially the union of P and the outer crescents of the triangulation. We say “essentially” because of a slight difficulty. If the collection of all outer crescents is pairwise disjoint, then R is exactly this union. It can happen, however, that some outer crescents overlap, e.g., see Figure 3. For a general n -gon, up to $\sim n$ outer crescents may overlap at a single point. When the outer crescents overlap outside P , we adjoin them to P , but consider them disjoint; this gives a Riemann surface R . See Figure 3.

First consider $\Psi : \Omega \rightarrow R$. The map Ψ fixes each vertex of P and simply “pushes” each edge of P out to the corresponding boundary arc of R . However, in order to do this with uniform quasiconformal bounds, we need to restrict the shapes of the

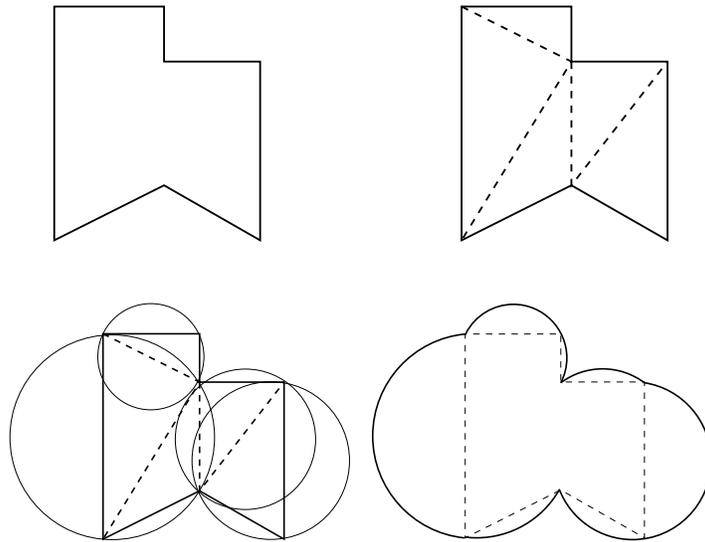


FIGURE 2. Triangulate P and attach the outer crescents to each edge of P . In this case the resulting region is planar, but sometimes the outer crescents may overlap, giving a Riemann surface.

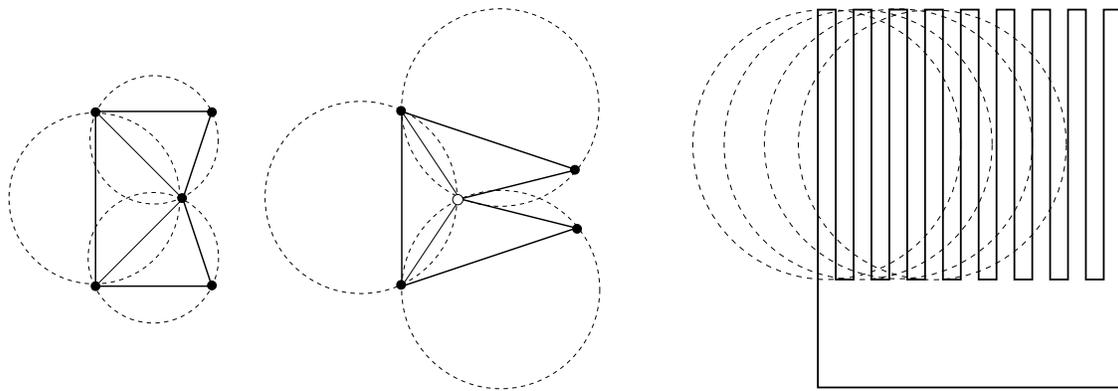


FIGURE 3. On the left the outer crescents are all disjoint; in the center some overlap; on the right is a domain where $\sim n$ outer crescents can overlap.

triangles. Driscoll and Vavasis showed that by adding extra vertices to P we can obtain a polygon P' whose Delaunay triangulation has three kinds of triangles T :

interior triangles: T has no edge on P' .

thin isosceles leaves: T is a leaf of the triangulation, is isosceles with its base in the interior of P and opposite angle $\leq \pi/4$.

well separated: if e is an edge of T that lies on P and v is the vertex of T opposite e , then $\text{dist}(v, e) \geq \frac{1}{2\sqrt{3}}\ell(e)$ (distance measured within P).

Theorem 4. *Let P' be a polygon bounding a region Ω and suppose P' has a triangulation whose elements are each one of the three types described above. Let R be the surface constructed from P' as above. Then there is a quasiconformal map $\Psi : \Omega \rightarrow R$ that fixes each vertex of P' and has QC constant bounded independent of n and P' .*

The well separated condition is stronger than we need. Following the proof of Theorem 4 we will state a weaker condition that suffices. A special case of this is:

thick interior crescents: there is a $\phi > 0$ so that if e is an edge of T on P , then there is a flat crescent inside Ω with the flat side e and angle at the vertices ϕ .

Note that this is a crescent inside Ω , so is not one of the outer crescents we have discussed earlier. If we replace the well separated condition by the thick crescents condition, then Theorem 4 holds with a bound depending only on ϕ . If the thick crescents condition fails, then the condition in Theorem 2 also fails for $M = 1/\phi$. This implies Theorem 2.

Next consider the map $\Phi : R \rightarrow \mathbb{D}$. This is more interesting than $\Psi : \Omega \rightarrow R$ in the sense that Φ restricted to the vertices gives the CRDT initial guesses for the Schwarz-Christoffel parameters (whereas Ψ was the identity at these points). Driscoll and Vavasis define this map in terms of cross ratios, but we will give an equivalent, more geometrical description of Φ .

First note that elements of our triangulation form the vertices of a tree and this tree structure is shared by the circumdisks. Sometimes adjacent triangles will have the same circumdisk. If this happens, we collapse the tree of circumdisks by identifying the two corresponding vertices and removing the corresponding edge. In the resulting tree of disks, each distinct circumdisk is represented by a unique vertex and adjacent disks overlap in exactly two distinct points. See Figure 4.

Choose a root D_0 for the tree of disks. The parent, D^* , of any non-root disk D is the unique disk that is adjacent to D and closer to the root. There is an “obvious” map $\tau_D : D \rightarrow D^*$: the elliptic Möbius transformation whose fixed points are $\partial D \cap \partial D^*$ and that rotates D to D^* . (See Section 2 and Appendix A for the definitions and

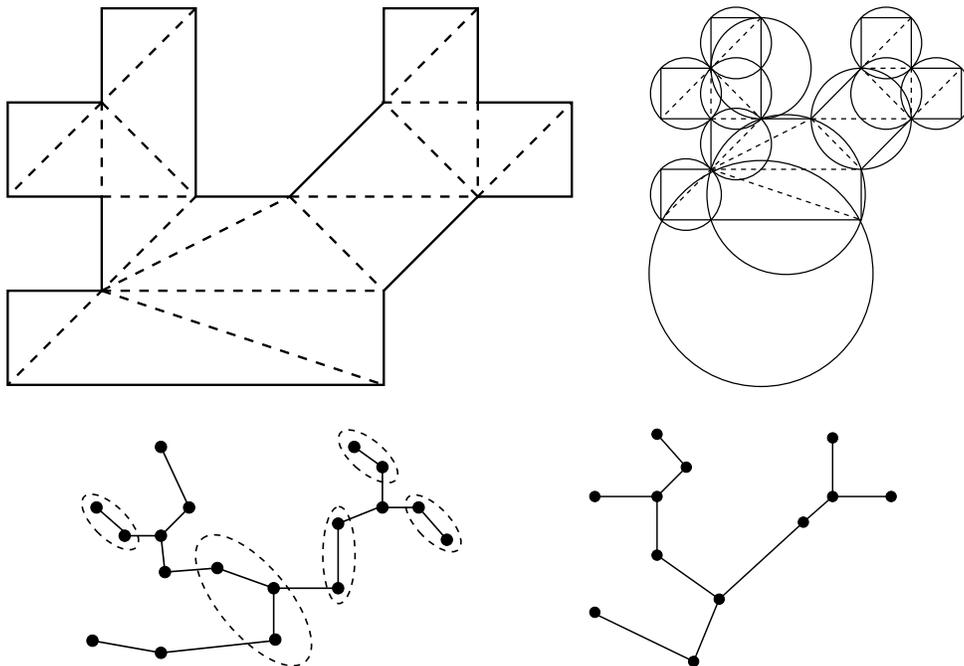


FIGURE 4. A triangulated polygon and the corresponding collection of circumdisks. Below is the triangulation tree and the tree of distinct circumdisks (vertices to be identified are grouped by the dashed lines).

properties of Möbius transformations.) If we set $\iota_{D_0} : D_0 \rightarrow D_0$ to be the identity, then we can inductively define maps $\iota_D = \iota_{D^*} \circ \tau_D : D \rightarrow D_0$ from any disk to the root. Any vertex v of P is on the boundary of some disk D and we set $\iota(v) = \iota_D(v)$ (if v is on the boundary of more than one disk, we can easily check that $\iota(v)$ is independent of the choice of D). We call this a “tree-of-disks” map. In Section 3 we will show that $\iota(v)$ agrees with the initial guess of the CRDT algorithm. Note that ι makes sense for any point on the boundary of R (other than the vertices of P , each point of ∂R is on the boundary of a unique disk D and we just take $\iota = \iota_D$ at these points). We will construct a map $\Phi : R \rightarrow \mathbb{D}$ that continuously extends these boundary values.

If z is an interior point of R that is in a unique disk D then we can define $\iota(z) = \iota_D(z) \in D_0$. However, if z is in more than one disk, the different possible definitions need not agree. We will divide the interior of R into subregions called gaps and crescents. (See Figure 9 for an example of this division of R .) Each gap is contained in some circumdisk D and Φ will agree with ι_D on the corresponding gap. Between

the gaps are the crescents and in these regions Φ will continuously interpolate its values on the gaps. The map Φ will collapse crescents onto arcs, hence it is not a homeomorphism, hence it cannot be quasiconformal. However, in Section 4 we will prove Φ is a quasi-isometry between the hyperbolic metrics on R and \mathbb{D} and then known results imply there exists a quasiconformal homeomorphism from R to \mathbb{D} with the same boundary values as Φ (see Appendix B). That is:

Theorem 5. *Suppose P is a planar polygon with vertices $\mathbf{v} = \{v_1, \dots, v_n\}$ and R is the Riemann surface constructed from a Delaunay triangulation of P as above. Then there is a continuous mapping of $\Phi : R \rightarrow \mathbb{D}$ so that $\Phi(\mathbf{v}) = \iota(\mathbf{v}) = \mathbf{w}$, are the initial guesses of CRDT algorithm and Φ is a quasi-isometry between the hyperbolic metrics on R and \mathbb{D} with constants that are independent of n and P .*

Note that this holds for all polygons; we do not assume that the extra vertices of step (1) of the CRDT algorithm algorithm have been added. The map ι can be constructed for any triangulation, but the uniform bounds can fail if the triangulation is not a Delaunay triangulation. See the second example in Section 6.

If R is a planar domain, ι is known to have a K -quasiconformal extension with $K < 8$; see [10]. Thus the novel aspect of this paper is to extend the planar result to the more general surfaces R constructed above using the CRDT extra vertices, and to describe geometrically when these extra vertices are needed. The ι map was originally invented in the context of hyperbolic 3-manifolds ([17], [31]), and its connection to planar conformal maps has been explored in a series of papers by Epstein, Marden, Markovic and the author (see [6], [7], [8], [9], [10], [18], [19], [20], [21]).

The first step of CRDT (which creates the well separated edges) may add an arbitrarily large number of extra vertices, so there is no time bound for the initial guess of CRDT depending only on the number of vertices. However, in [7] a method is given for computing a initial guess that has the same uniform distance bounds to the correct answer, but can be performed in time $O(n)$, with constant independent of the geometry. Moreover, [6] gives an algorithm that computes the correct prevertices to within d_{QC} distance ϵ in time $O(n)$ and constant depending only on ϵ (and no worse than $O(|\log \epsilon| \cdot \log |\log \epsilon|)$).

The rest of the paper is organized as follows:

Section 2: We define $\Phi : R \rightarrow \mathbb{D}$, the extension of ι .

Section 3: We show that the initial guess of CRDT agrees with the ι map.

Section 4: We prove Φ is a quasi-isometry. This proves Theorem 5.

Section 5: We prove Theorems 4 and 2.

Section 6: We show the Driscoll-Vavasis conjecture is false for cross ratios.

Section 7: We compute some examples of the CRDT initial guesses.

Appendix A: Definition of conformal modulus and its relation to cross ratios.

Appendix B: Definitions and properties of quasiconformal mappings.

Many thanks to Toby Driscoll and Stephen Vavasis for reading earlier versions of this paper and for many helpful comments about the CRDT algorithm. Also many thanks to editor Nick Papamichael and to the referees. Their careful reading of the manuscript and detailed comments improved the exposition and accessibility of the paper and I greatly appreciate all their suggestions.

2. THE DEFINITION OF Φ

In this section we define the continuous map $R \rightarrow \mathbb{D}$ described in Theorem 5. This map will not be a homeomorphism (hence is not quasiconformal), but we will prove in Section 4 that it is a quasi-isometry between the hyperbolic metrics on R and \mathbb{D} .

A triangulation of a simple polygon P is called Delaunay if, for every interior edge e of the triangulation, the two triangles T_1, T_2 sharing edge e have angles opposite e summing to at most π . Equivalently, if D_1, D_2 are the associated open circumdisks then $T_1 \not\subset D_2$ and $T_2 \not\subset D_1$. See Figures 5 and 6.

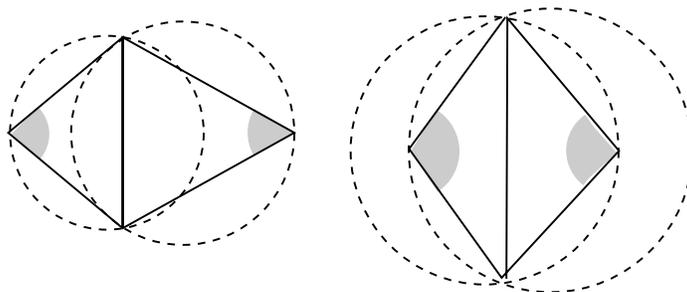


FIGURE 5. The triangles on the left satisfy the Delaunay condition, but the ones on the right do not.

Every polygon has a Delaunay triangulation. It is unique except when adjacent triangles define four vertices on a common circle: then we may “flip” the diagonal

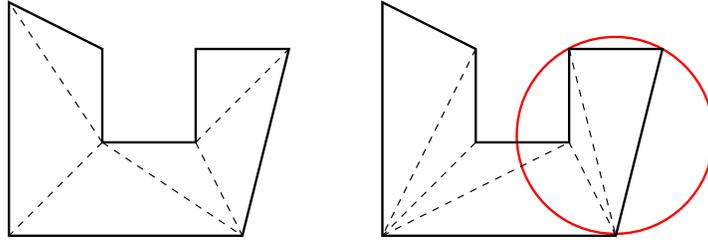


FIGURE 6. A Delaunay and non-Delaunay triangulation. We have drawn a circumcircle on the right that fails the Delaunay condition.

of the resulting quadrilateral and get another Delaunay triangulation. Thus the corresponding family of disks and the tree-of-disk map on the vertices does not depend on the choice of Delaunay triangulation. Delaunay triangulations have many nice properties and have been intensively studied. For example, they are dual to Voronoi diagrams; and they maximize the minimum angle in the triangulation [26]; they minimize the largest circumcircle [12], [28]. The basic facts can be found in various sources such as [2], [5], [22].

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

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

The formula can be easily derived by sending a and b to 0 and ∞ respectively by the map $w = \tau(z) = (z - a)/(z - b)$, then multiplying by $e^{i\theta}$ and then applying the inverse map $z = \tau^{-1}(w) = (bw - a)/(w - 1)$. Geometrically, a crescent is foliated by circular arcs orthogonal to both boundaries and this elliptic transformation identifies endpoints of leaves. See Figure 7.

Suppose D_1, D_2 are overlapping disks whose boundaries meet at two distinct points v_1, v_2 . Let γ_1 be the circular arc connecting v_1 and v_2 in D_1 that is perpendicular to ∂D_1 . This arc is the infinite geodesic connecting v_1 and v_2 in the hyperbolic geometry of D_1 (see Section 4), so we shall simply refer to γ_1 as the geodesic arc between v_1, v_2

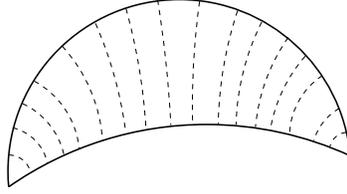


FIGURE 7. The orthogonal foliation of a crescent

in D_1 . Similarly, let γ_2 be the geodesic arc between v_1, v_2 in D_2 . Then $\gamma_1 \cup \gamma_2$ bounds a crescent C in $D_1 \cup D_2$, called the “normal crescent”. See Figure 8.

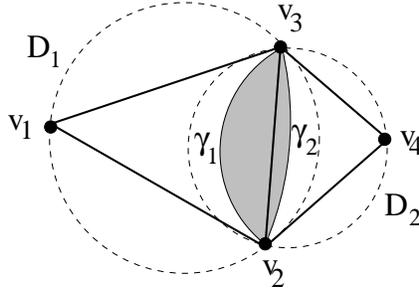


FIGURE 8. Given two overlapping disks D_1, D_2 we define the normal crescent (shaded region) to be bounded by the hyperbolic geodesics for these two disks. The Delaunay condition insures that γ_1 is to the left of γ_2 in this figure.

Note that $(D_1 \cup D_2) \setminus C$ consists of two components, each a crescent. If D_1, D_2 are the circumdisks of two adjacent triangles $T_1 = \{v_1, v_2, v_3\}$, $T_2 = \{v_4, v_3, v_2\}$, and these triangles satisfy the Delaunay condition, then the boundary of one of these components contains both γ_1 and v_1 . We call this the main component of D_1 with respect to D_2 . See Figure 8.

Given a crescent C and one of the two boundary arcs $\gamma \subset \partial C$, the collapsing map $\Gamma : C \rightarrow \gamma$ is the map that sends each leaf of the orthogonal foliation of C onto the endpoint of the leaf contained in γ . On γ , Γ is the identity. On the other boundary arc γ' , Γ agrees with the elliptic Möbius transformation that fixes each vertex of C and maps γ' to γ . Given an ordered pair $\{D_1, D_2\}$ of overlapping disks, define a map $\Gamma : D_1 \cup D_2 \rightarrow D_1$ as follows: on the main component of D_1 with respect to D_2 it is the identity, on the normal crescent it is the collapsing map of the crescent onto γ_1

and on the main component of D_2 with respect to D_1 it equals τ , the elliptic Möbius transformation that rotates γ_2 to γ_1 .

Suppose we have a tree of overlapping disks, coming from a Delaunay triangulation of a polygon. Given a disk D in the tree, the corresponding gap $G \subset D$ is the intersection of the main components of D with respect to all its neighbors in the tree. The gaps for distinct disks are disjoint and the gaps for adjacent disks are separated by a normal crescent. See Figure 9 for an example of this “gap/crescent” decomposition.

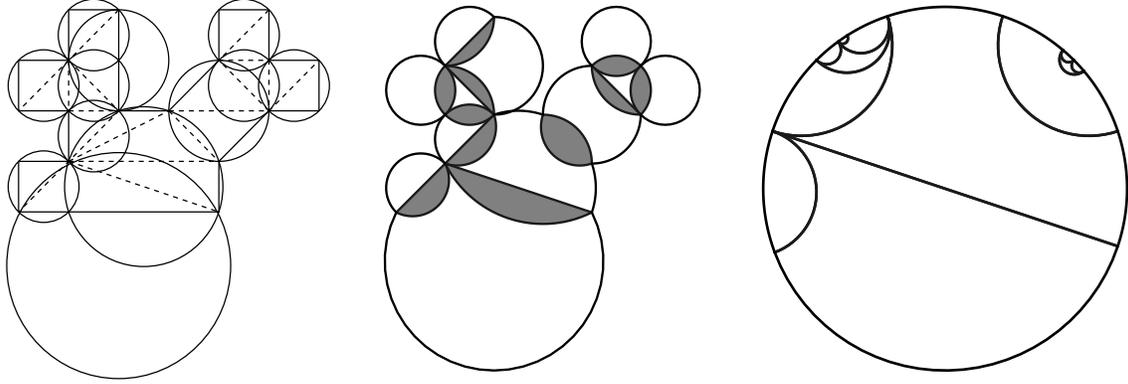


FIGURE 9. On the left is a tree of disks corresponding to a triangulated polygon (the same as in Figure 4). In the center is the corresponding gap/crescent decomposition. On the right is the Φ image of the decomposition. Every shaded crescent has been collapsed to a hyperbolic geodesic and each gap has been mapped into the disk by a Möbius transformation. The vertices in the rightmost picture were used as the parameters to draw the center picture of Figure 1.

Assume that the m disks in our tree are enumerated as $\{D_k\}$ so that D_0 is the root and so that any disk has a lower index than any of its children. Define a sequence of surfaces $\{R_k\}$ inductively by setting $R_0 = D_0$ and defining R_k by attaching the crescent $D_k \setminus D_k^*$ to R_{k-1} . Define a map $\Gamma_k : R_k \rightarrow R_{k-1}$ by extending the map $\Gamma : D_k^* \cup D_k \rightarrow D_k^*$ as the identity on the rest of R_k . Then $\Phi : \Gamma_1 \circ \cdots \circ \Gamma_m$ is a mapping from $R = R_m$ to $D_0 = R_0$. It is equal to the identity on the gap of D_0 , is equal to some Möbius transformation on the gap of every other disk and collapses each normal crescent to a hyperbolic geodesic of D_0 (these are called the bending geodesics, a term coming from hyperbolic geometry). See Figure 9. Note that Φ is

not a homeomorphism, since arcs in the normal crescents are collapsed to points, but at least it does map the interior of R to the interior of D_0 . We shall see in the next section that Φ is actually close to homeomorphism in a precise sense.

3. CRDT AND TREE-OF-DISKS MAPS

In this section we show that the map Φ , restricted to the vertices of P , agrees with the CRDT initial guess as defined by Driscoll and Vavasis in [16].

How is the second step of CRDT done (i.e., how does it create its initial guess from a triangulated polygon)? Suppose P is a simple n -gon and that $\mathcal{T} = \{T_1, \dots, T_{n-2}\}$ is a triangulation of P . (In this paper, triangulation means using only the original vertices of P ; no Steiner points are added.) Let D_k be the circumcircle associated to each triangle T_k , $k = 1, \dots, n-2$. Choose some root triangle for 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_1, v_2, v_3 and v_1, v_3, v_4 in counterclockwise order respectively). Also suppose that we have already defined $w_1, w_2, w_3 \in \mathbb{T}$. Then w_4 is uniquely determined by the condition

$$(2) \quad \text{cr}(w_1, w_2, w_3, w_4) = -|\text{cr}(v_1, v_2, v_3, v_4)|,$$

where $\text{cr}(a, b, c, d)$ denotes the cross ratio of four distinct complex numbers a, b, c, d ; See Appendix A. It is easy to see by induction that this uniquely determines the points \mathbf{w} up to a Möbius transformation of the circle. This is the CRDT initial guess.

Next, we will check that this map defined using cross ratios is the same as the tree-of-disks map ι defined in the introduction. Recall that we had a tree whose vertices were the circumdisks of the triangulation, that D_0 was the root of the tree. Suppose D_1, D_2 are the circumdisks associated to adjacent triangles T_1, T_2 and that D_1 is the parent of D_2 . Let $\tau_{D_2} : D_2 \rightarrow D_1$ be the elliptic Möbius transformation that fixes the two points $\partial D_1 \cap \partial D_2$. After conjugating by a Möbius transformation η that sends $v_1 \rightarrow 0$, $v_2 \rightarrow 1$ and $v_3 \rightarrow \infty$, the elliptic map τ_{D_2} is conjugated to the Euclidean rotation around 0 that sends the image to v_4 onto the negative real axis. Thus

$$\text{cr}(v_1, v_2, v_3, \tau_{D_2}(v_4)) = -|\text{cr}(v_1, v_2, v_3, v_4)|.$$

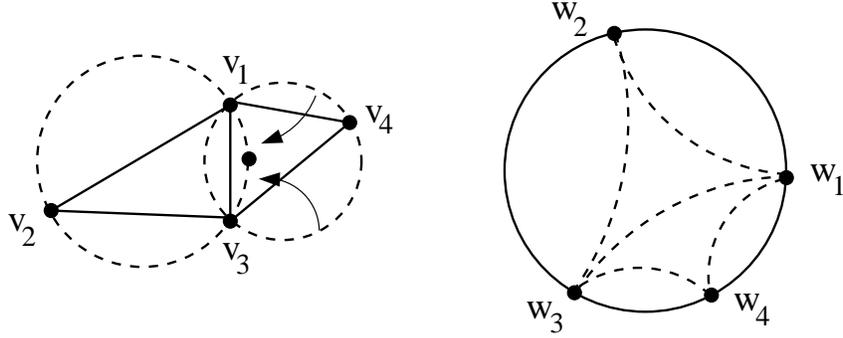


FIGURE 10. Given a quadrilateral $\{v_1, v_2, v_3, v_4\}$, and a mapping $\tau : \{v_1, v_2, v_3\} \rightarrow \{w_1, w_2, w_3\} \subset \mathbb{T}$, we map the fourth vertex $v_4 \rightarrow w_4$ by first using an elliptic rotation to map v_4 onto the circle defined by the first three points followed by the unique Möbius transformation determined by τ . This is the same as the map defined by Driscoll and Vavasis using cross ratios.

If $D_1 = D_0$ (i.e., D_2 is adjacent to the root), we define $\iota_{D_1} = \tau_{D_2}$. For other disks we inductively define $\iota_{D_2} = \iota_{D_1} \circ \tau_{D_2} : D_2 \rightarrow D_1 \rightarrow D_0$. When the Möbius transformation $\iota_{D_1} : D_1 \rightarrow D_0$ is applied to the four points $\{v_1, v_2, v_3, \tau_{D_2}(v_4)\} \in \partial D_1$ to give the points (see Figure 10),

$$\{w_1, w_2, w_3, w_4\} = \{\iota(v_1), \iota(v_2), \iota(v_3), \iota(v_4)\},$$

the cross ratio does not change. Thus

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

This shows that our “tree-of-disks” map sends the vertices of P to a n -tuple $\iota(\mathbf{v})$ on the circle that has the same set of cross ratios as the n -tuple \mathbf{w} constructed by the CRDT algorithm. Therefore \mathbf{w} and $\iota(\mathbf{v})$ are equivalent n -tuples under a Möbius transformation of the circle.

4. Φ IS A HYPERBOLIC QUASI-ISOMETRY

In this section we prove Φ is a quasi-isometry from (R, ρ_R) to (D_0, ρ_{D_0}) . Known results described in Appendix B then imply Theorem 5.

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. The orientation preserving isometries are exactly the Möbius transformations that preserve the disk. They all have the

form $z \rightarrow e^{i\theta}(z - a)/(1 - \bar{a}z)$, 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 for any hyperbolic metric when the domain is clear from context.

A map between metric spaces $f : (X, \rho_1) \rightarrow (Y, \rho_2)$ is called a quasi-isometry if there are constants A, B such that

$$\frac{1}{A}\rho_1(x, y) - B \leq \rho_2(f(x), f(y)) \leq A\rho_1(x, y) + B.$$

This says that the map is bi-Lipschitz at large scales and has bounded jumps at small scales.

A very useful fact is that on any simply connected planar domain, the hyperbolic metric is approximated by a more geometric quantity (e.g., page 20 of [23]),

$$(3) \quad \frac{1}{2}d\tilde{\rho} = \frac{|dz|}{2\text{dist}(z, \partial\Omega)} \leq d\rho_\Omega \leq \frac{2|dz|}{\text{dist}(z, \partial\Omega)} = 2d\tilde{\rho},$$

where $d\tilde{\rho}$ is called the quasi-hyperbolic metric. This inequality says that the metric spaces (Ω, ρ) and $(\Omega, \tilde{\rho})$ are bi-Lipschitz equivalent with each other (via the identity map). Thus to show Φ is a quasi-isometry on a planar domain, we may take either metric (at the cost of increasing the constants by a bounded factor). Our regions R need not be planar, but we shall prove that they satisfy an inequality ((4) below) similar to (3). The first step is to prove a simple geometric property of R that says that it is ‘‘almost planar’’.

Lemma 6. *Suppose R is a surface obtained by adding outer crescents to a simple polygon as described in the introduction. 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 of 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 its 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. \square

Lemma 7. *There is an $M < \infty$ so that the following holds. If R is the surface obtained from a Delaunay triangulation of a simple, planar polygon P , then*

$$(4) \quad \frac{|dz|}{M \operatorname{dist}(z, \partial R)} \leq d\rho_R \leq \frac{2|dz|}{\operatorname{dist}(z, \partial R)}.$$

Proof. The proof of the right hand inequality is exactly the same as the planar case: since R contains a disk D of radius $r = \operatorname{dist}(z, \partial R)$ around z , the monotonicity property of hyperbolic distance [27] (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 Euclidean 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 Γ connecting $D = D(z, \frac{1}{2}r) \subset R$ to ∂R is bounded away from zero. See Appendix A for the definitions and basic facts. Next we estimate this modulus in the usual way using the Cauchy-Schwarz inequality.

Let $v \in \partial R$ minimize $\operatorname{dist}(z, \partial R)$ 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 $\{w : |w - v| = t\}$ that hits D . By Lemma 6 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

$$\begin{aligned} M(\Gamma) = \iint_R \rho^2 dx dy &\geq \int_{r/2}^{3r/2} \int_{\gamma_t} \rho^2 t dt d\theta \\ &\geq \left[\int_{r/2}^{3r/2} \left(\int_{\gamma_t} t d\theta \right) dt \right]^{-1} \cdot \left[\int_{r/2}^{3r/2} \left(\int_{\gamma_t} \rho t d\theta \right) dt \right]^2 \\ &\geq \left[6\pi \int_{r/2}^{3r/2} dt \right]^{-1} \cdot \left[\int_{r/2}^{3r/2} 1 dt \right]^2 \\ &\geq \left[6\pi \frac{1}{2} \frac{3}{2} r \right]^{-2} r^2 \\ &= \frac{4}{81\pi}. \end{aligned}$$

Thus the modulus of Γ is bounded away from zero uniformly, as desired. This completes the proof of Lemma 7. \square

Lemma 8. *Suppose G is a gap in R and $z \in G$. Let D be the disk associated to G . Then (using Euclidean distances),*

$$\frac{1}{\sqrt{2}} \text{dist}(z, \partial R) \leq \text{dist}(z, \partial D) \leq \text{dist}(z, \partial R).$$

Proof. The right hand inequality is obvious since $D \subset R$. To prove the other inequality, note that if z is the center of D or if the radial projection of z is in ∂R then we have equality on the left. Otherwise, the radial segment through z hits a geodesic γ in D bounding a crescent. Replacing z with this point of intersection decreases $\text{dist}(z, \partial D)$ more than $\text{dist}(z, \partial R)$ and for a point on γ , the left-hand inequality holds by a simple computation. \square

Thus on a gap G corresponding to a disk D we have

$$(G, d\rho_R) \rightarrow (G, d\tilde{\rho}_R) \rightarrow (G, d\tilde{\rho}_D) \rightarrow (G, d\rho_D) \rightarrow (\mathbb{D}, d\rho),$$

where the first two maps are the identity and are bi-Lipschitz by Lemmas 7 and 8, and the third is bi-Lipschitz by (3). The last map is the restriction to G of a Möbius transformation from D to \mathbb{D} and hence is an isometry of the hyperbolic metrics on these two disks. Therefore Φ is hyperbolically bi-Lipschitz on each individual gap. To see what happens when points are in different gaps we will need the following result (which says that composing rotations with clustered fixed points is well approximated by rotating around a single point):

Lemma 9. *There is a $\epsilon_0 > 0$ and $C < \infty$ so that if $\epsilon < \epsilon_0$ then the following holds. Suppose $\{\tau_j\}_1^M$ is a finite collection of elliptic transformations such that each τ_j rotates by angle θ_j , that $\sum_j |\theta_j| = L \leq (4C\epsilon)^{-1}$ and that each τ_j has one fixed point in $D(0, \epsilon)$ and one fixed point outside $D(0, 1/\epsilon)$. Then $|\tau_1 \circ \dots \circ \tau_M(w) - w \exp(i \sum_j \theta_j)| \leq C\epsilon L$ for any w with $\frac{1}{2} \leq |w| \leq \frac{3}{2}$.*

Proof. Let $A_1 = \{w : \frac{1}{2} \leq |w| \leq \frac{3}{2}\}$ and $A_2 = \{w : \frac{1}{4} \leq |w| \leq \frac{7}{4}\}$. Each elliptic transformation preserves a family of circles and by the hypothesis on the fixed points, the circles in these families restricted to A_2 differ by at most an angle $C\epsilon$ from circles concentric with the origin. Thus if we apply one of these rotations of angle θ to a point z in A_2 the distance from A_1 increases by at most $C\epsilon\theta$. Thus if we start in A_1 and apply several maps $\{\tau_j\}$ with angles $\{\theta_j\}$, the total distance from A_1 will increase by at most $\sum C\epsilon\theta_j \leq C\epsilon L \leq 1/4$ so the image is still in A_2 .

Since τ_j differs from a rotation of angle θ_j around the origin by at most $C\epsilon\theta_j$ on A_2 , the desired result follows from summing. See Figure 11. \square

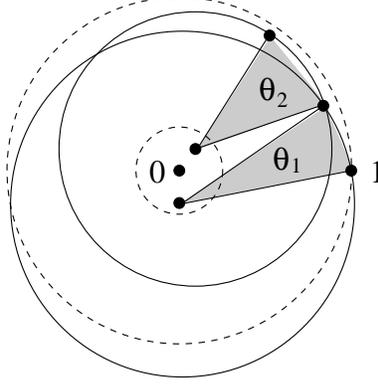


FIGURE 11. If the fixed points are ϵ -near 0 and ∞ the elliptic transformations look like Euclidean rotations up to error $C\epsilon$.

Lemma 10. $\Phi : R \rightarrow \mathbb{D}$ is a quasi-isometry between the hyperbolic metrics.

Proof. As noted before, Φ is bi-Lipschitz on each gap and is clearly Lipschitz on every crescent, so it is Lipschitz on all of R .

To prove Φ is a quasi-isometry, we need to prove the opposite direction

$$\rho_{\mathbb{D}}(\Phi(z), \Phi(w)) \geq A\rho_R(z, w) - B.$$

It suffices to show that for some $\epsilon > 0$ the preimage under Φ of a ball of hyperbolic radius ϵ in \mathbb{D} has uniformly bounded hyperbolic diameter in R . If this is true, then any two points distance $d > \epsilon$ apart in \mathbb{D} can be connected by a geodesic, and this geodesic can be cut into $\simeq d/\epsilon$ parts, each of which has preimage of diameter C . Thus the two original points have all preimages within $(C/\epsilon)d$ of each other, which is the desired lower bound.

Consider a ball B of hyperbolic radius ϵ in \mathbb{D} . Renormalizing by a Möbius transformation we may assume it is centered at the origin. If it is contained inside the image of one gap $G \subset D \subset R$, there is nothing to do. So assume B hits one or more bending geodesics (i.e., images of crescents under Φ). Let $\{\gamma_j\}$ be an enumeration of these and let θ_j be the angle measure of the corresponding crescents in R . Since each of these geodesics hits the small ball B , the endpoints of all the γ 's must be clustered in two balls of (Euclidean) radius 16ϵ , say, centered at ± 1 . See Figure 12.

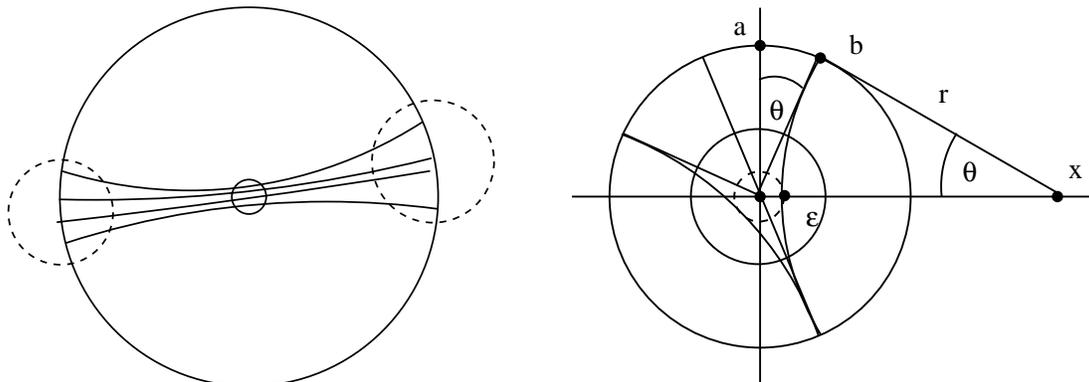


FIGURE 12. If many hyperbolic geodesics all hit a small ball at the origin, they must be almost parallel and all terminate at nearby points. The figure on the right shows why: if a geodesic passes at distance ϵ from the center and parallel to a diameter of the circle there, then the endpoints labeled a and b differ by at most θ where $\sin \theta = 1/x$, $\tan \theta = 1/r$, $\epsilon = x - r$ and so $|\theta| \leq 4\epsilon$. As shown, any other disjoint geodesic hitting the ϵ ball has its endpoints within 4θ of b , giving the claimed estimate.

By applying a Möbius transformation taking the unit disk to the upper half-plane we can instead assume these geodesics are in \mathbb{H} and each has one endpoint in $[0, \epsilon]$ and the other in $[C\epsilon, \infty]$. These geodesics also come with a natural left to right ordering and each has an associated elliptic Möbius transformation that rotates by angle θ_j around the endpoints of γ_j . If we apply these maps from rightmost to leftmost, letting the map act on everything to the right of the corresponding geodesic, we get the map Φ on the gaps.

Let B_2 be the ball of radius $1/2$ centered at i . Let $L = 7\pi$ and $\epsilon \leq 1/(7CL)$ where C is the constant from Lemma 9. If $\sum_j \theta_j = L \geq 7\pi$ then Lemma 9 implies the preimage of B_2 in R covers the unit circle more than 3 to 1 at some point. This is a contradiction, so we must have $L < 7\pi$. But this means that the preimage of B_2 is a connected set in R that can be covered by a uniformly bounded number of balls of radius $1/4$ that are more than distance $1/4$ from ∂R . Each element of this cover therefore has bounded hyperbolic radius in R and hence the preimage of B has uniformly bounded hyperbolic radius. \square

Thus Φ is a quasi-isometry with uniformly bounded constants and hence there is a K -quasiconformal map $R \rightarrow \mathbb{D}$ (with uniformly bounded K) that has the same boundary values (see the remarks in Appendix B). This proves Theorem 5.

5. PROOF OF THEOREM 4

Next we prove that there is a quasiconformal map $\Psi : \Omega \rightarrow R$ with uniformly bounded constant, assuming we have added extra vertices to P according to the first step of the CRDT algorithm. Following the proof we will observe that the same conclusion sometimes holds even without adding the extra points. To see that the extra points are needed in some cases, 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 $1/\log n$ (see Appendix A). However, if ι had a K -quasiconformal extension to P with K independent of P , these two moduli would be comparable to within a factor of K of each other.

Lemma 11. *Given any polygon P we can add vertices of angle π along the edges of P to form a new polygon P' so that every triangle T in the Delaunay triangulation of P' is one of the following types:*

- (1) *interior (i.e., T has no edge in P').*
- (2) *thin isosceles leaves (i.e., T is a leaf of the triangulation, is isosceles with base in Ω and opposite angle $\leq \pi/4$).*
- (3) *well separated (i.e., if e is an edge of T in P' and v is the vertex of T opposite e then $\text{dist}(v, e) \geq \frac{1}{2\sqrt{3}}\ell(e)$).*

Most of this is proven by Driscoll and Vavasis in [16]. Here we will sketch their proof and add a few remarks that give the precise statement above.

Their construction has two steps. In the first step, every vertex v of P with interior angle $\leq \pi/4$ is “chopped off” as follows. For such a v , let T be the largest isosceles triangle contained in the interior of P that has v as a vertex, and whose two equal length sides are subsets of the two edges of P adjacent to v . Add to P the midpoints of these two sides of T (thus the two edges of P adjacent to v are now replaced by four

edges occupying the same space). The two new edges adjacent to v are “protected”, i.e., they will not be further subdivided by the second part of the construction.

The second step is iterative. For each unprotected edge e in the current polygon compute its length L and the minimal distance D in P from e to any vertex that 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 13. The filled dots indicate the original vertices and the open dots the vertices added by the algorithm.

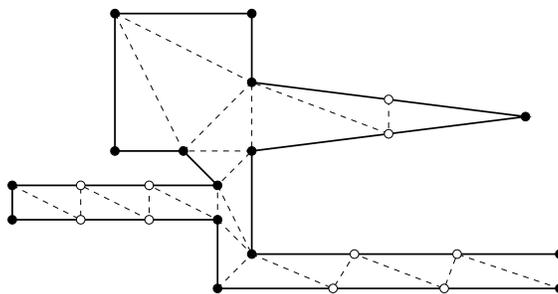


FIGURE 13. Adding extra (white) points to a polygon in the CRDT algorithm to remove small angles from the triangulation

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.

To prove Lemma 11, suppose we have a Delaunay triangulation on P' . If T is an element of the triangulation that is either interior, or only has unprotected boundary edges, then it clearly falls into cases (1) and (3) of the lemma. Otherwise, it has at least one protected edge. So to complete the proof of the lemma we need to show such a triangle is in case (2). That is, we need the following:

Lemma 12. *The “cut-off” triangles created in the first step of the construction are all in the Delaunay triangulation of P' .*

Proof. To prove the lemma, suppose it fails, i.e., suppose v is a vertex of P with interior angle $\theta \leq \pi/4$ so that the corresponding chopped off triangle $\triangle vab$ is not in the Delaunay triangulation of P' . Then there are a finite number of elements of

the Delaunay triangulation that have v as one vertex, say T_1, \dots, T_n where we may assume each triangle shares an edge with the next one, a is a vertex of T_1 and b is a vertex of T_n . See Figure 14.

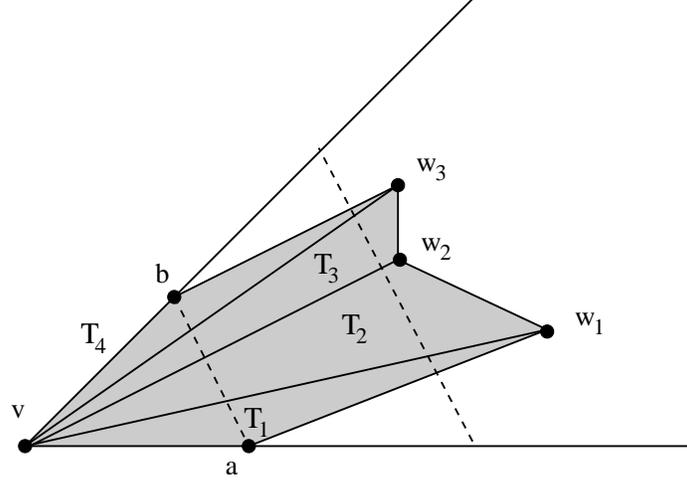


FIGURE 14. The sub-polygon W is shown shaded. We sum $\pi - \alpha_k$ over the interior angles of W , considering three cases: v , $\{w_k\}_1^n$ and $\{a, b\}$. The first is bounded by $\pi - \theta$, the second by 2θ and the third by 2θ .

Let W be the union of these triangles and let $\{\alpha_k\}$ be the interior angles of W . Note that $\sum_{k=1}^{n+1}(\pi - \alpha_k) = 2\pi$. Let $Q_k = T_k \cup T_{k+1}$ for $k = 1, \dots, n-1$ and let w_k be the vertex of Q_k opposite v . Each of these satisfies the Delaunay condition so if θ_k is the interior angle of Q_k at v and τ_k is the interior angle of Q_k at w_k , we have $\theta_k \geq \pi - \tau_k$. Therefore $\sum_k(\pi - \tau_k) \leq \sum_k \theta_k \leq 2\theta$ (the last inequality holds since we have counted the interior angle of each triangle at v at most twice). Therefore the sum of $\sum_k(\pi - \alpha_k)$ at the vertices w_k is at most 2θ and at v it is equal to $\pi - \theta$. The sum of these two cases is at most $\pi + \theta$.

The only two vertices of W that remain are at a and b . Suppose T is the triangle $\triangle vab$ and let $2T$ denote the triangle obtained by doubling the lengths of the sides va and vb as in Figure 14. Let S be the side of $2T$ opposite v . The CRDT algorithm chooses T so that the interior of $2T$ is contained in the interior of P and hence contains no vertices of P . Thus the sides of W adjacent to a, b and not connecting them to v have their other endpoints outside $2T$. These sides must cross S and by

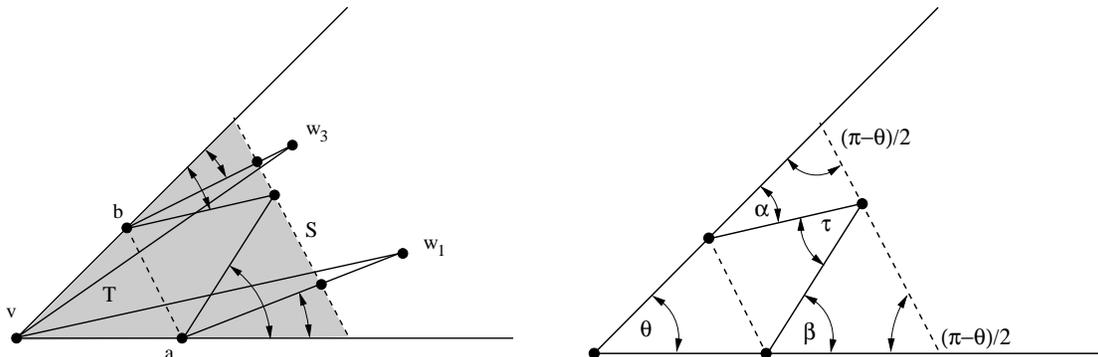


FIGURE 15. Summing the interior angles of W and a and b . First we replace the angles at a, b by larger angles α, β . The right-hand diagram shows $\alpha + \beta = \theta + \tau$ and it is a simple exercise to see $\tau \leq \theta$ with equality when it is the midpoint of S .

replacing them by two segments that meet at a common point p of S we can increase the sum $\sum(\pi - \alpha)$. See Figure 15.

It is an easy exercise to see that this sum is maximized when p is the midpoint of S in which case it is equal to $2\theta \leq \pi/2$. Thus $\sum(\pi - \alpha_k)$ over all vertices of W is at most $\pi + 3\theta < 2\pi$. However, since ∂W bounds a simply connected polygonal region, this sum must equal 2π , a contradiction. Thus the cut off triangle with vertex v must have been part of the Delaunay triangulation. (Note that the same argument would work as long as $\theta < \pi/3$.) \square

Now that we know the Delaunay triangulation of P' consists of three types of triangles, we can prove the following result by considering each case separately.

Lemma 13. *Suppose $P = \partial\Omega$ has a rooted triangulation whose elements satisfy (1)-(3) in Lemma 11. Then there is a K -quasiconformal map $f : \Omega \rightarrow R$ that fixes each of the vertices of P , and K is independent of P .*

Proof. On interior triangles we define f to be the identity. Moreover, on any triangle T , f will be the identity on any boundary edge of T not on P , so the definitions on adjacent triangles will fit together trivially.

To discuss the other two cases, we need to recall the relationship between the interior angles of a triangle and the angles formed between its edges and its circumcircle. See Figure 16. It illustrates a triangle $\triangle abc$, where the angles at a, b, c are denoted

α, β, γ respectively and the lengths of the opposite sides are denoted A, B, C . Let θ be the interior angle of the crescent formed by edge ab and the arc of the circumcircle between these points and similarly for τ and the edge ac . Note that $\alpha + \beta + \gamma = \pi$, $\tau + \alpha + \theta = \pi$ and $\tau + \gamma = \beta + \theta$. Combining these we see that $\theta = \gamma$. If, in addition, $\alpha = \beta$ then we deduce $\tau = \beta = \alpha$.

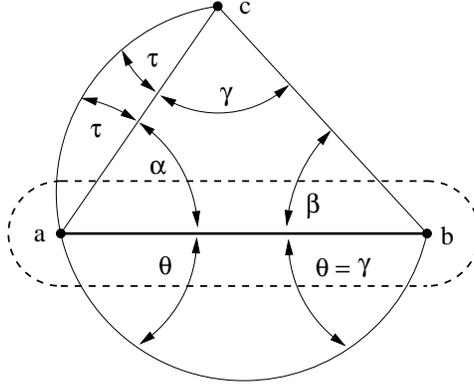


FIGURE 16. Diagram for proof of $\theta = O(\min(\alpha, \beta))$.

Also note that the Law of Sines implies

$$\sin \theta = \sin \gamma = \frac{C}{B} \sin \beta \leq \lambda \sin \beta,$$

if $\lambda = C/B$. In particular, if we know that $\text{dist}(c, ab) \geq \lambda C$ (as indicated by c being outside the dashed neighborhood of ab in Figure 16), then we have $\sin \theta = O(\sin \beta) = O(\beta)$ with a constant depending only on λ . Also, $\text{dist}(c, ab) \geq \lambda C$ implies that angle γ (and hence θ) is bounded away from π depending only on λ , so $\theta = O(\sin \theta)$ with a constant depending only on λ (in fact, $\gamma \leq \pi - 2 \arctan 2\lambda$). Thus $\theta = O(\beta)$.

Since the same argument works with β replaced by α we see that

$$(5) \quad \theta = O(\min(\alpha, \beta)),$$

with a constant bounded in terms of λ .

Now consider a triangle $T = \triangle abc$ so that the edge ab is on the boundary of the polygon and so that the opposite vertex c has distance at least $\lambda|ab|$ from ab . The angle bisectors of T split T into three subtriangles, and the subtriangle that contains the edge ab contains a crescent C_1 of angle $\nu = \frac{1}{2} \min(\alpha, \beta)$. Let C_2 denote a crescent of angle θ attached to the outside of T along ab . Then C_1 can be mapped to $C_1 \cup C_2$

by a quasiconformal map that is the identity on $\partial C_1 \setminus ab$ and the quasiconformal constant is bounded by $(\nu + \theta)/\nu$, which is uniformly bounded by (5).

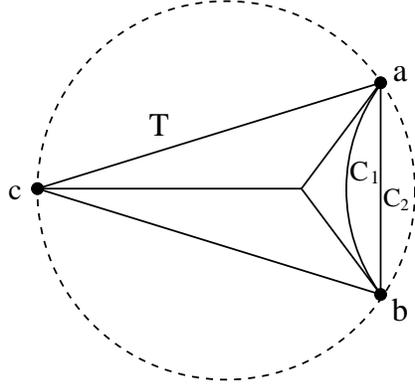


FIGURE 17. If the triangle is well separated, then it includes a crescent with angle comparable to the crescent attached to the boundary. The interior crescent can then be expanded by a uniformly quasiconformal map to also cover the attached crescent.

Finally suppose T is a thin isosceles leaf. By assumption T is attached to its parent along its base, and not adjacent to any other triangle. The vertex opposite the base has angle $\theta \leq \pi/4$ so that the two base angles are $\alpha = (\pi - \theta)/2 \geq 3\pi/8$. The two outer crescents attached to T both have angle α and their union with T is itself a crescent of angle 2α with one side being the base of T . Now we map this big crescent (with angle close to π) to a smaller crescent (with angle $\pi/2$) by an at most 2-quasiconformal map. See Figure 18. Then map the small crescent to the triangle as follows. Start with two conformal maps: an elliptic rotation to a half-plane minus a disk and the logarithmic map to a half-infinite strip. These two maps sent $[-1, 1]$ to $[0, \pi]$ by the map

$$x \rightarrow f(x) = \arcsin\left(\operatorname{Im}\left(\frac{z+i}{z-i}\right)\right) = \arcsin\left(\frac{1-x^2}{1+x^2}\right).$$

Then we apply the map of the form $(x, y) \rightarrow (f^{-1}(x), y)$ that makes the composition affine on the base and is quasiconformal since f' is bounded above and below. Finally we rescale the strip to have base $[-1, 1]$ and then apply a map of the form

$$(x, y) \rightarrow (e^{-y}x, C(1 - e^{-y})),$$

where C is the height of the triangle. A simple computation shows this is quasiconformal if C is bounded away from 0. See Figure 19. This proves Theorem 4 and hence Theorem 1.

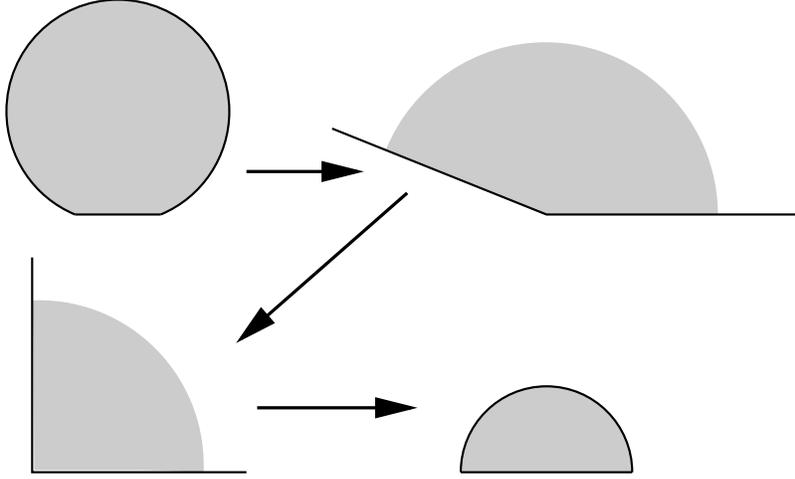


FIGURE 18. A crescent of angle α can be quasiconformally mapped to a crescent of angle β with $K = \beta/\alpha$. Just use a Möbius transformation to first map to a cone of angle α and then use a map of the form $re^{i\theta} \rightarrow re^{i\theta\beta/\alpha}$. This maps to a cone of angle β with the given QC bound. The map to the desired crescent by another Möbius transformation.

□

The well separated condition (condition (3) in Lemma 11) is stronger than we need. We only use it in the proof of Lemma 13 to establish the following property with a uniform $\eta > 0$:

interior crescents: There is a $\eta > 0$ so that for each edge e of T that lies on P , there is a crescent with base e and angle $\geq \eta\theta$, where θ is the angle in T opposite e .

The proof also needs that the crescents are disjoint, but we get this for free from the following simple result.

Lemma 14. *Suppose I_1, I_2 are disjoint open line segments and C_1, C_2 are crescents with bases I_1, I_2 respectively. Let $C'_1 \subset C_1, C'_2 \subset C_2$ be crescents with bases I_1, I_2 , but*

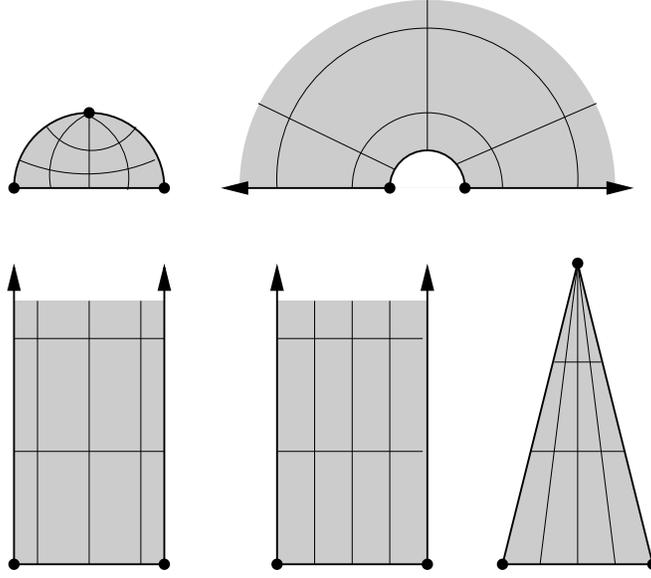


FIGURE 19. The half-disk is mapped quasiconformally to a triangle in a series of steps.

half the angle of C_1, C_2 respectively. If $C_1 \cap I_2 = C_2 \cap I_1 = \emptyset$, then C'_1 and C'_2 are disjoint.

Proof. We claim that the bisector of the segments I_1, I_2 separates the interiors of the reduced crescents C'_1, C'_2 . Suppose that this does not hold. Then there is a point z of the bisector contained in the interior of one of these crescents, say C'_1 . This point is the center of a disk D whose boundary hits both I_1 and I_2 and whose interior is disjoint from both of them. We claim this disk is contained in C_1 .

To prove this, note that $D \cap C'_1$ is a crescent with interior angle $> \pi$. If we map C'_1 to a cone by using a Möbius transformation to map its vertices to $0, \infty$, then the image of this crescent still has angle $> \pi$ (since such maps are conformal) and hence the image disk is centered in the image cone. It is then obvious that the image of D is contained in the cone with double the angle and hence D is contained in C_1 (actually is contained in a crescent with strictly smaller angle). Since ∂D hits I_2 , so does C_1 , contrary to assumption. See Figure 20. \square

Given the lemma, we convert our collection of interior crescents into a disjoint collection (dividing η by 2), and define a map that expands each interior crescent to the union of itself and the attached exterior crescent. This map is quasiconformal

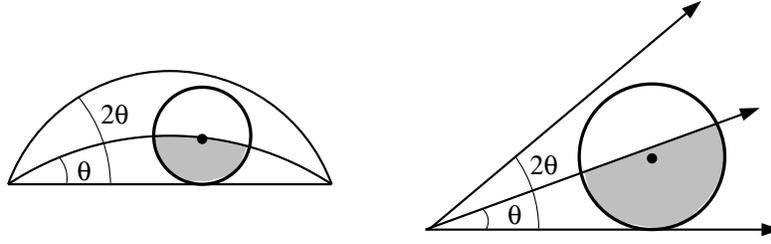


FIGURE 20. If a disk is centered in a crescent and misses one edge, it must be contained in the union of the crescent and its reflection across the other edge.

with constant depending only on η . We then apply the map that expands the thin leaves of the triangulation exactly as before. This shows that Theorem 4 holds if we replace the “well separated” condition by the “interior crescents” condition.

In the interior crescents condition, the opposite angles are bounded by π , so it suffices to require the crescents have base angles bounded away from 0, i.e., the thick crescents condition from the introduction. So, if the thick crescent condition holds with angle ϕ , then Theorem 4 holds with a constant depending only on ϕ .

6. TWO COUNTEREXAMPLES

First we shall show that Corollary 3 does not hold if the modules are replaced by cross ratios. Consider the domain Ω bounded by a polygon P in Figure 21. It is the square $[-1, 1] \times [-1, 1]$ with a 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

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

On the other hand, the conformal map f of Ω to the disk that 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

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

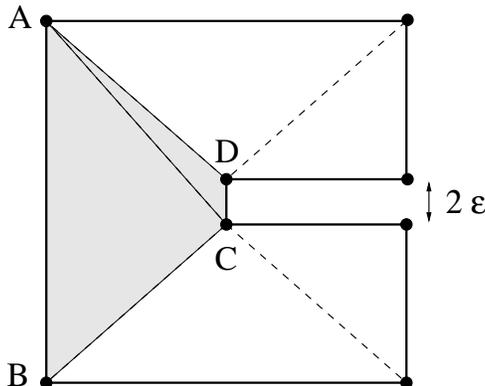


FIGURE 21. Cross ratios and modulus can scale at different powers of ϵ

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

$$\log cr_1 - \log cr_2 > \frac{1}{2} \log \frac{1}{\epsilon} - O(1),$$

which is as large as we wish if ϵ is small enough.

This polygon does not satisfy the 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.

Next we show that the uniform bounds in Theorem 5 can fail if we do not have a Delaunay triangulation. The example is given in Figure 22. Let C be a crescent of small angle θ that is symmetric with respect to the real axis and with vertices a, c at ± 1 . Let b, d denote the “highest” and “lowest” points of the crescent, as shown in Figure 22. Form a polygon by taking a, b, c, d and placing evenly spaced vertices along the boundary of the crescent, about distance $\epsilon \ll 1$ apart.

If we use the triangulation shown in Figure 22 (which is not Delaunay because of triangles bordering the single horizontal diagonal), then the corners labeled $\{a, b, c, d\}$ can be mapped to $\{1, i, -1, -i\}$ by the cross ratio map, and the vertex v will be mapped to a point about distance ϵ from 1. However, any conformal map that sends $\{a, b, c, d\}$ to these points on the circle, would map v to within $O(\epsilon^{1/\theta})$ of the point 1. Because of Mori’s theorem (which says a K -quasiconformal map is Hölder of order $1/K$), we see that the CRDT initial guess differs from the conformal map by a

quasiconformal map with constant $K \simeq 1/\theta$. Thus the bound is not independent of the domain.

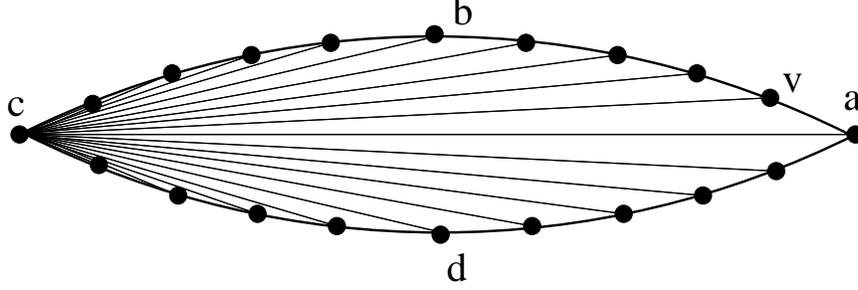


FIGURE 22. A non-Delaunay triangulation may give a bad initial guess for CRDT.

7. EXAMPLES AND OPEN QUESTIONS

Given a polygonal domain Ω , the CRDT initial guess produces a n -tuple on the unit circle that we have proven is uniformly close to the true parameters in a quasiconformal sense. In this section we will explicitly estimate how close in a few cases by plugging the CRDT guess into the Schwarz-Christoffel formula (using the known angles), and comparing the resulting image to the original polygon.

As part of the CRDT algorithm we computed the Delaunay triangulation of Ω . For each element T of this triangulation, consider the triangle T' corresponding to T in the guessed image. There is an obvious affine map between these triangles and we can easily compute its quasiconformal constant of this map as follows. First use a conformal linear map to send each triangle to one of the form $\{0, 1, a\}$ and $\{0, 1, b\}$. The affine map is then of the form $f(z) \rightarrow \alpha z + \beta \bar{z}$ where $\alpha + \beta = 1$ and $\beta = (b - a)/(a - \bar{a})$ and from this we see that

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

where

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

If the triangle T' is degenerate, or has the opposite orientation as T , we simply give ∞ as our QC bound K .

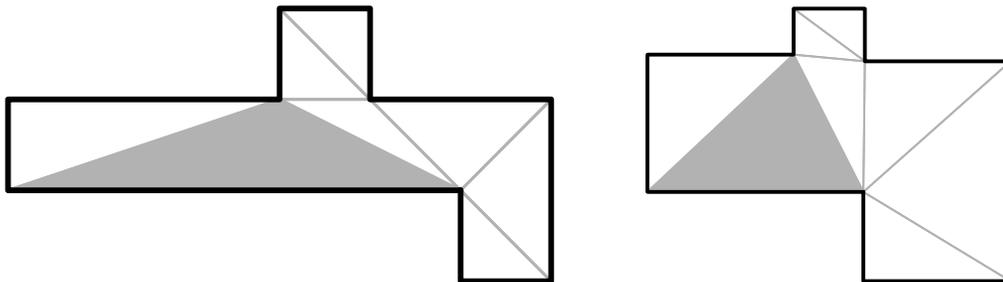


FIGURE 23. A polygon and Schwarz-Christoffel image using guessed parameters. The most distorted triangle is shaded.

The affine maps on each triangle define a QC map between the target polygon and the guessed image, so the maximum triangle distortion is an upper bound for the QC distance from the guessed parameters to the correct ones. This may not be an exact estimate, but at least it gives a rigorous upper bound for the QC distance to the correct parameters, without having to know what the correct parameters are. A few examples are worked out in Figure 24.

Next we describe the iterative step of CRDT. For fixed angles, evaluating the SC-formula defines a map S from n -tuples in \mathbb{T} to n -tuples in \mathbb{C} . If we change the n -tuple by a Möbius transformation of the circle, then the image polygon changes by a linear transformation. Thus S is a well defined map from \mathbb{T}_* (n -tuples of distinct points on \mathbb{T} modulo Möbius transformations) into \mathbb{C}_* (complex n -tuples modulo Euclidean similarities). Moreover, we can identify \mathbb{T}_* with \mathbb{R}^{n-3} as follows: fix a triangulation of the n points on \mathbb{T} , and for each pair of adjacent triangles record the logarithm of the cross ratio of the four vertices (the cross ratio is positive if we take the correct ordering of the four points). The original n -tuple can be recovered, up to Möbius transformations, from these $n - 3$ real values so $\mathbb{T}_*^n = \mathbb{R}^{n-3}$. This allows us to apply linear algebra to n -tuples on \mathbb{T} .

Suppose we have an explicit way of guessing the SC-parameters for a polygon, i.e., a map $G : \mathbb{C}_*^n \rightarrow \mathbb{T}_*^n$. Then $F = G \circ S$ gives a map $\mathbb{R}^{n-3} \rightarrow \mathbb{R}^{n-3}$ and the correct SC-parameters for Ω solve $F(\mathbf{z}) = \mathbf{z}_0$ (where $\mathbf{z}_0 = G(\mathbf{v}) \in \mathbb{T}_*^n$) and hence are fixed by

$$(6) \quad \mathbf{z}_{k+1} = \mathbf{z}_k - A(F(\mathbf{z}_k) - \mathbf{z}_0).$$

where A is any $(n-3) \times (n-3)$ matrix. The CRDT iteration is this iteration with the map G from polygons to n -tuples being the CRDT initial guess (the tree-of-disks maps

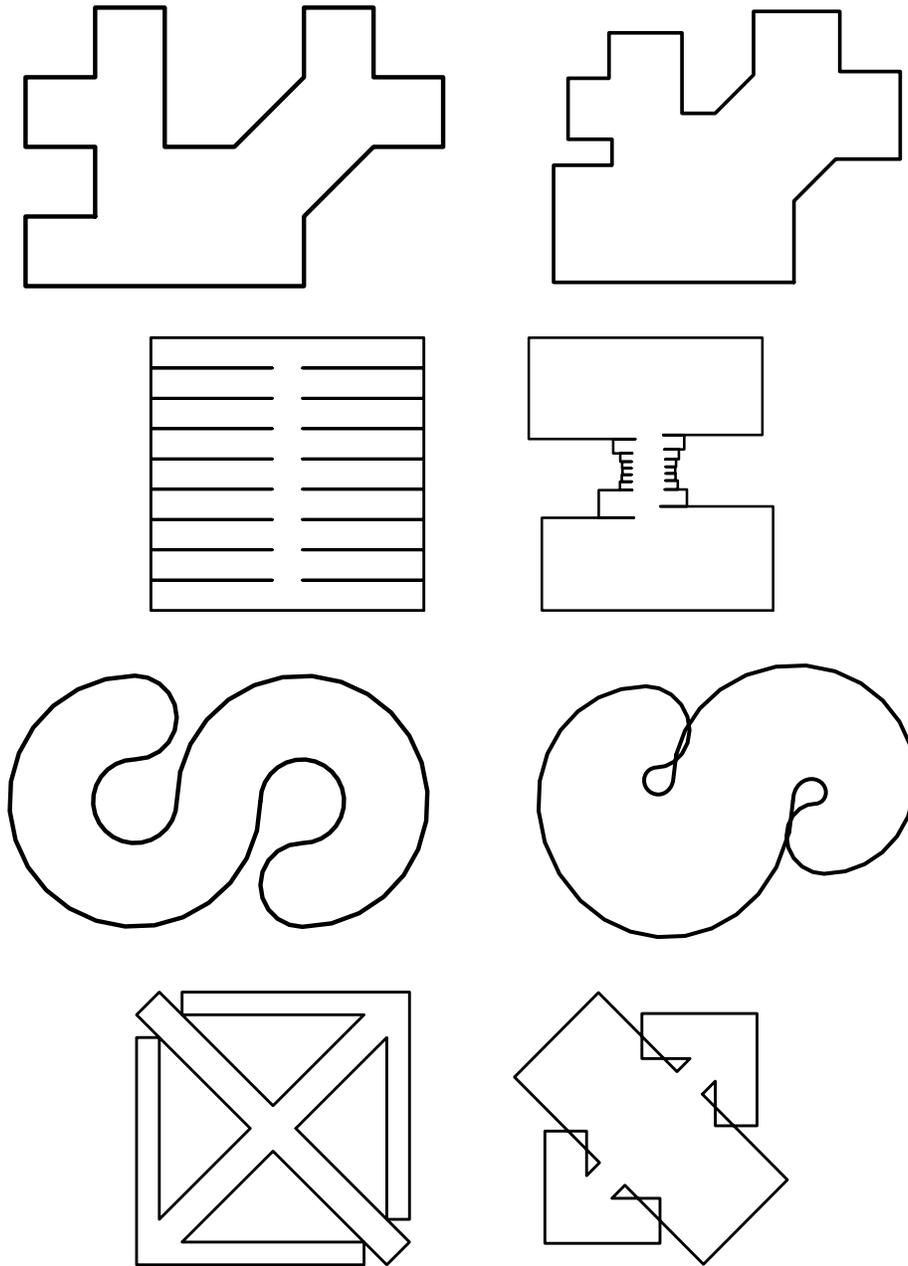


FIGURE 24. Target polygons (left) and Schwarz-Christoffel images using initial CRDT guesses (right). We have not added extra vertices. The QC distortions are bounded by 3.12157, 3.8577, 2.61153, 3.62306 respectively, using the triangle method described in the text. Note that CRDT tends to shorten long corridors.

from the introduction). Full CRDT consists of letting $A = D^{-1}$, where $D = DF$ is the derivative matrix of the function F , computed by discrete approximation. Simple CRDT assumes A is the identity matrix. Shortcut CRDT starts by assuming A is the identity, but uses Broyden updates at each iteration to make A approximate D^{-1} using the information gathered at each step (see Chapter 8 of [13]). In general, shortcut CRDT works best for problems of reasonable size. See Figures 25 and 26 for some examples of how simple and shortcut CRDT work.

Next we illustrate Theorem 2 and our remarks following Theorem 4. As pointed out in the introduction, the extra points added in the first step of CRDT are not always needed to guarantee that the initial guess of the Schwarz-Christoffel parameters are within a bounded distance of the correct ones. Figures 27 and 28 show two examples of polygons to which CRDT will add vertices, but in which the original polygon already has a guaranteed good initial guess. In Figure 27 we have two rectangles joined by a small gap. The Delaunay triangulation has a diagonal that bridges this gap, so it satisfies Theorem 2 and has a good initial guess. CRDT will subdivide the four edges adjacent to the gap (recall each slit has two sides) and then will further subdivide the closest edges, converting a 10-gon to a 26-gon.

At each step we compute the polygon corresponding to the current parameter guess and compare it to the target using triangulations as described above. We have plotted $-\log(K - 1)$ at each step for each of our two polygons. The CRDT iteration step is quadratic in the number of vertices, because in evaluating the Schwarz-Christoffel integrals, we evaluate n integrals, each of an n -fold product although this can be avoided by other methods, e.g., [4]. Also, the vertices of angle π added by CRDT do not contribute to the product in the SC formula. For the example in Figure 27 the running time increases by almost a factor of 10, with no gain in accuracy.

In Figure 28, the triangulation does not satisfy Theorem 2, but it does satisfy the isosceles leaf and thick crescent conditions of the introduction, so the initial guess will be good. CRDT however, cuts off the two narrow triangles, and then adds extra vertices to the remaining narrow channels, converting an 8-gon to a 50-gon. The running time increased by a factor of 86, without any gain in accuracy, as shown in Figure 29.

Finally, we note that the following questions remain unanswered.

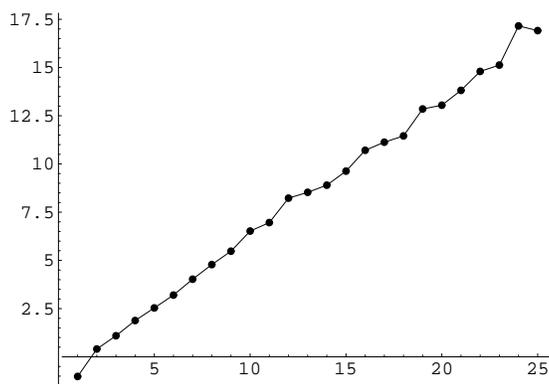
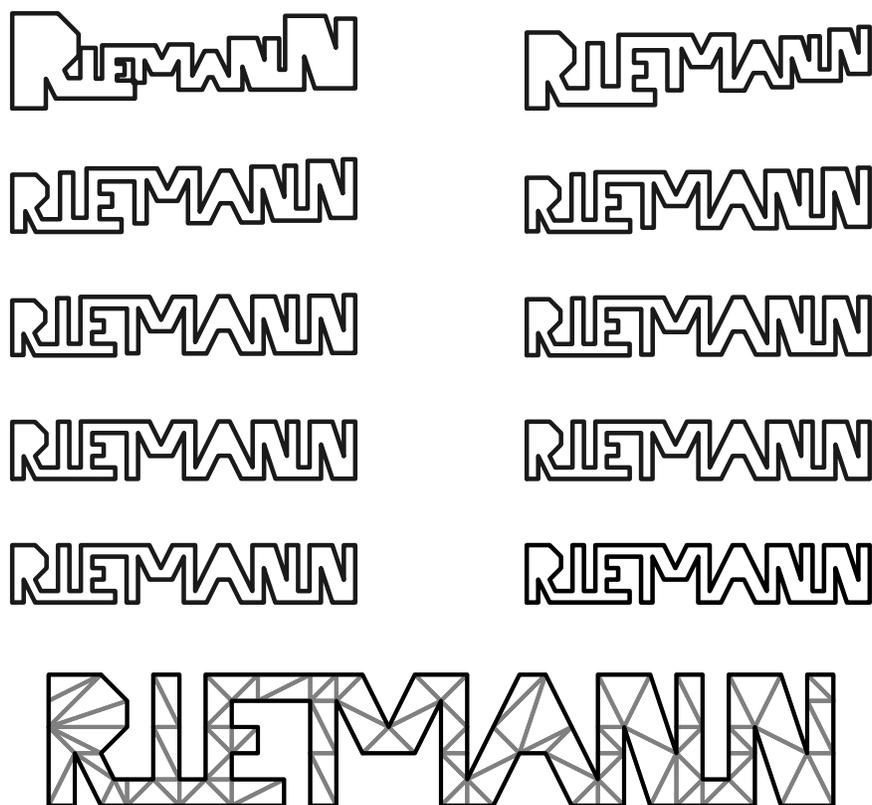


FIGURE 25. Shortcut CRDT applied to a 98-gon. The top left shows the CRDT initial guess, followed by 9 iterations of shortcut CRDT. The center figure shows the triangulation used. The bottom shows the gain in accuracy over 25 iterations. I have graphed $-\log(K - 1)$ where K is the upper bound for the QC error derived from triangulations

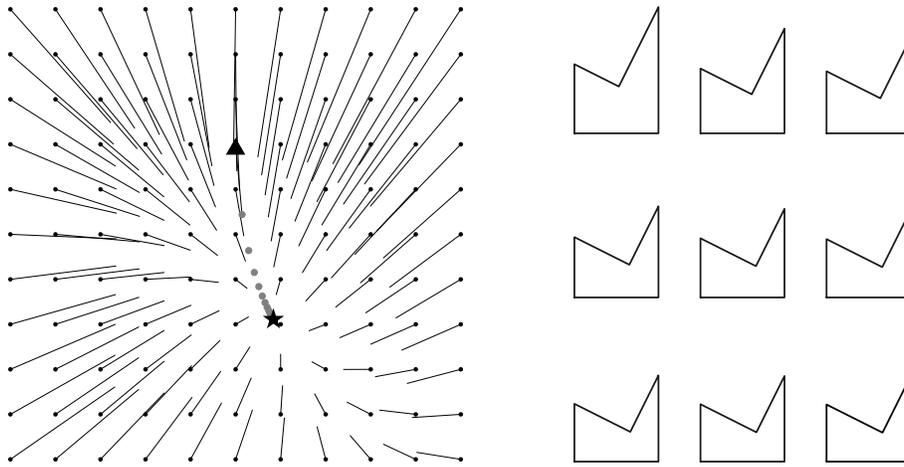


FIGURE 26. For a pentagon, CRDT iterations acts on \mathbb{R}^2 (5-tuples on the circle modulo Möbius transformations). On the left we show the map $\mathbf{z} \rightarrow \mathbf{z} - (F(\mathbf{z}) - \mathbf{z}_0)$ corresponding to simple CRDT for a particular pentagon. Convergence of simple CRDT is the claim that iterating the initial CRDT guess z_0 (the triangle) under this map will converge to the fixed point of the vector field (the star). In this picture convergence seems obvious from any starting point; what can we prove in general? The iterates shown on the left (the gray dots) correspond to the pentagons shown on the right.

- (1) How can we modify the first step of CRDT to only add enough points so that the interior crescents condition is satisfied? How many points do we need to add to make the Delaunay triangulation a M -spanner for all of P ?
- (2) Can we use the computation of quasiconformal maps between triangulations to choose better steps for the iteration? For example, what happens if we make extra evaluations in directions corresponding to the most distorted triangles and use Broyden updates to incorporate this information into A ?
- (3) Can the estimates in this paper be refined to prove that CRDT converges?
- (4) An alternative to the CRDT map from polygons to n -tuples is proposed in [7] using the medial axis of P . Like the CRDT map, this map has uniform QC bounds (as in Theorem 1) but in various examples it seems to give a better approximation than the CRDT guess gives. Can we prove this? Does replacing the CRDT map by the medial axis map give a better conformal mapping algorithm?

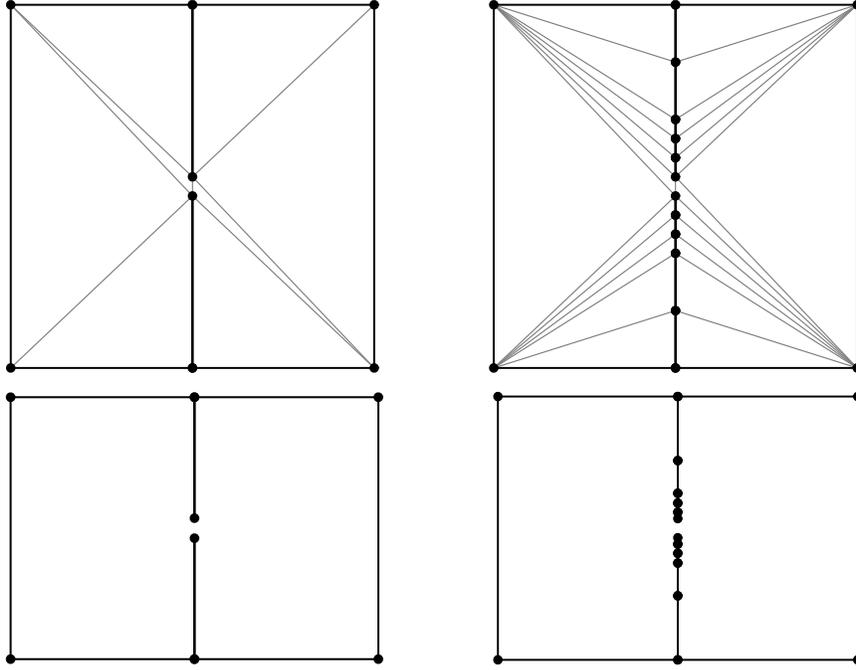


FIGURE 27. The upper left is a polygon satisfying Theorem 2, and the upper right is the result when CRDT adds extra vertices. The Delaunay triangulations are shown in gray. Below are the Schwarz-Christoffel images using the CRDT initial guesses for these triangulations. The QC errors are bounded above by 1.44394 and 3.37951 respectively. Thus adding the extra vertices actually gives a worse initial guess.

- (5) Is the CRDT iteration monotone with respect to the QC distance? (Figures 25 and 29 show that our upper bound estimates are not monotone everywhere, but appear “roughly” monotone.)
- (6) Is there an alternate metric that is more appropriate for studying the convergence of CRDT?

APPENDIX A. 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

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

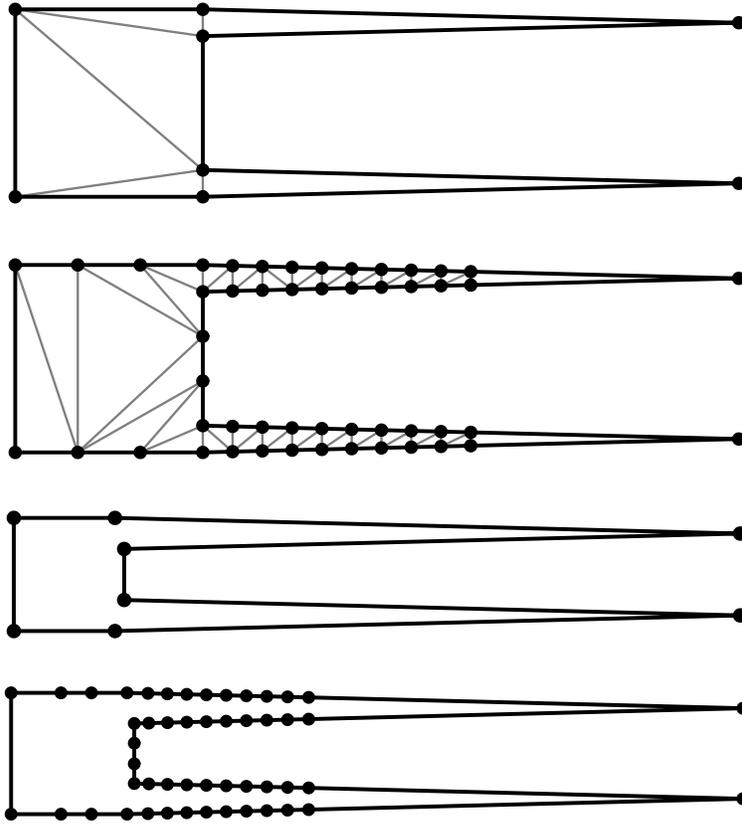


FIGURE 28. A polygon and the same polygon with the CRDT extra vertices. Below are the Schwarz-Christoffel images using the initial guesses generated from these polygons. The QC upper bounds are 2.38016 and 2.40062 respectively, so adding the extra vertices makes almost no difference.

and define the modulus of Γ as

$$\text{Mod}(\Gamma) = \inf_{\Omega} \int_M \rho^2 dx dy,$$

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]. One property that we use in this paper is that if $K \subset \mathbb{D}$ is a continuum, then its hyperbolic diameter is bounded iff the modulus of the path family connecting K to the $\partial\mathbb{D}$ is bounded (see Lemma 7).

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

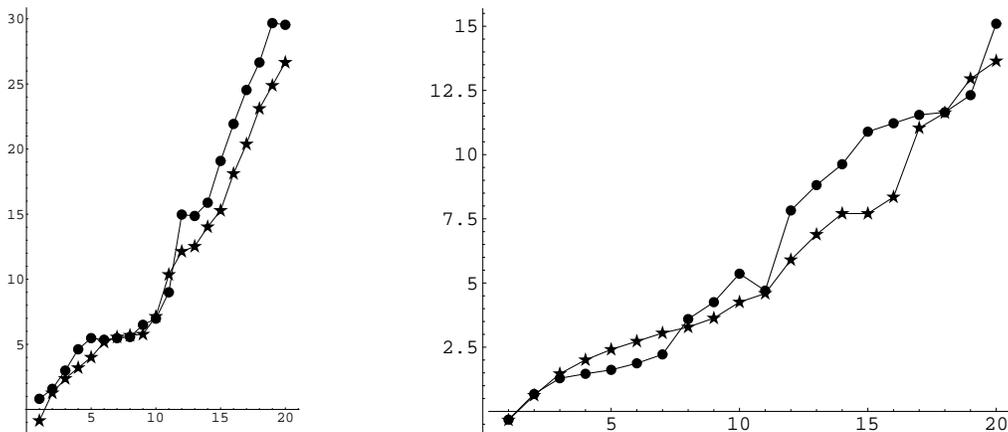


FIGURE 29. These two graphs compare the shortcut CRDT iteration for the polygons in Figures 27 and 28. We have plotted $-\log(K - 1)$ where K is the QC error upper bound at each step. On the left is the comparison between the polygons in Figure 27, with dots being the original polygon and stars the one with extra vertices. The results are similar, but the second graph took about ten times as long to compute. On the right is the corresponding graph for the polygons in Figure 28. Again, the results are similar, but the polygon with extra vertices took much longer to compute.

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 that 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 $\text{Mod}(Q') = 1/\text{Mod}(Q)$.

A linear fractional (or Möbius) transformation is a map of the form $z \rightarrow (az + b)/(cz + d)$. This is a 1-1, onto, holomorphic map of the Riemann sphere $\mathbb{S} = \mathbb{C} \cup \{\infty\}$ to itself. Such maps form a group under composition and are well known to map circles to circles (if we count straight lines as circles that pass through ∞). Möbius transforms are conformal, so they preserve angles.

The non-identity Möbius transformations are divided into three classes. Parabolic transformations have a single fixed point on \mathbb{S} and are conjugate to the translation map $z \rightarrow z + 1$. Elliptic maps have two fixed points and are conjugate to the rotation $z \rightarrow \lambda z$ for some $|\lambda| = 1$. The loxodromic transformations also have two fixed points and are conjugate to $z \rightarrow \lambda z$ for some $|\lambda| < 1$. If, in addition, λ is real, then the map is called hyperbolic.

The cross ratio of four distinct complex numbers a, b, c, d is

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

Note that $z \rightarrow \text{cr}(a, b, c, z)$ is the unique Möbius transformation that sends a to 0, b to 1 and c to ∞ . Thus the cross ratio is clearly Möbius invariant. Moreover, $\text{cr}(a, b, c, d)$ is real valued iff the four points lie on a circle; and is negative iff the points are labeled in counterclockwise order on the circle.

Since the cross ratio and the modulus of four points on the unit circle are both invariant under Möbius transformations of the disk to itself, each must be a function of the other. 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 $-P = \text{cr}(a, b, c, d) \in (-\infty, 0)$. 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 that separate $[-P, 0]$ from $[1, \infty]$. We will denote this modulus by $M(P)$. The transformation $z \rightarrow (z - 1)/(z + P)$ sends $0, 1, \infty, -P$ to $-\frac{1}{P}, 0, 1, \infty$, so by the Möbius invariance of modulus and the fact that conjugate quadrilaterals have reciprocal moduli, we see that

$$(7) \quad M_{\mathbb{H}}\left(\frac{1}{P}\right) = \frac{1}{M_{\mathbb{H}}(P)},$$

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

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

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

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

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

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 (7) to deduce

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

which implies

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

In other words,

$$\begin{aligned} M &\simeq \frac{1}{2\pi} \log P, & P \gg 1, \\ M &\simeq \frac{\pi}{2|\log P|}, & P \ll 1, \end{aligned}$$

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

$$\begin{aligned} M_{\mathbb{D}}(\mathbf{x}) &\simeq \frac{1}{\pi} \log |\text{cr}(\mathbf{x})|, & |\text{cr}(\mathbf{x})| \gg 1, \\ M_{\mathbb{D}}(\mathbf{x}) &\simeq \frac{\pi}{|\log |\text{cr}(\mathbf{x})||}, & |\text{cr}(\mathbf{x})| \ll 1. \end{aligned}$$

Another elegant connection between modulus and cross ratios is given by Bagby in [3], where it is shown that conformal modulus for a ring domain is given by minimizing an integral involving logarithms of cross ratios.

APPENDIX B. QUASICONFORMAL MAPPINGS

Quasiconformal mappings are a generalization of conformal mappings that 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 $f : \Omega \rightarrow \Omega'$ is a homeomorphism.

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

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 $k = (K - 1)/(K + 1)$. Here $f_{\bar{z}} = (1/2i)(f_x + if_y)$ and $f_z = (1/2)(f_x - if_y)$.

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

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

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 [24] 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 d_{QC} from the introduction satisfies the triangle inequality.

Even though they do not have 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 of the unit disk to itself generalize hyperbolic bi-Lipschitz maps, i.e., maps that satisfy

$$\frac{1}{K} \leq \frac{\rho(f(x), f(y))}{\rho(x, y)} \leq K.$$

Although a quasiconformal map $f : \mathbb{D} \rightarrow \mathbb{D}$ need not be bi-Lipschitz, it is a quasi-isometry of the disk with its hyperbolic metric ρ , i.e., there are constants A, B such that

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

In [19] Epstein, Marden and Markovic show that any K -quasiconformal selfmap of the disk is a quasi-isometry with 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 $h : \mathbb{D} \rightarrow \mathbb{D}$ is a quasi-isometry then a result of Vaisala [32] implies that h has a continuous extension to the boundary and that this extension is quasimetric, 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 that a circle homeomorphism has a quasiconformal extension to the disk, it suffices to prove it has a quasi-isometric extension. More generally, if $\Phi : \Omega \rightarrow \mathbb{D}$ is a quasi-isometry of the hyperbolic metrics and $f : \mathbb{D} \rightarrow \Omega$ is conformal, then $\Phi \circ f : \mathbb{D} \rightarrow \mathbb{D}$ is quasi-isometric and hence there is a quasiconformal map $\Psi : \mathbb{D} \rightarrow \mathbb{D}$ with the same boundary values. Thus $\Psi \circ f^{-1} : \Omega \rightarrow \mathbb{D}$ is a quasiconformal map with the same boundary values of Φ . This is the situation encountered in this paper.

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] F. Aurenhammer. Voronoi diagrams - a survey of a fundamental geometric data structure. *ACM Comp. Surveys*, 23:345–405, 1991.
- [3] T. Bagby. The modulus of a plane condenser. *J. Math. Mech.*, 17:315–329, 1967.
- [4] L. Banjai and L.N. Trefethen. A multipole method for Schwarz-Christoffel mapping of polygons with thousands of sides. *SIAM J. Sci. Comput.*, 25(3):1042–1065 (electronic), 2003.
- [5] 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.
- [6] C.J. Bishop. Conformal mapping in linear time. To appear in *Discrete and Computational Geometry*.
- [7] C.J. Bishop. A fast QC mapping theorem for polygons. Preprint, 2008.
- [8] C.J. Bishop. Divergence groups have the Bowen property. *Ann. of Math. (2)*, 154(1):205–217, 2001.
- [9] C.J. Bishop. Quasiconformal Lipschitz maps, Sullivan’s convex hull theorem and Brennan’s conjecture. *Ark. Mat.*, 40(1):1–26, 2002.
- [10] 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.
- [11] E.B. Christoffel. Sul problema della temperatura stazionale e la rappresentazione di una data superficie. *Ann. Mat. Pura Appl. Serie II*, pages 89–103, 1867.
- [12] E. F. D’Azevedo and R. B. Simpson. On optimal interpolation triangle incidences. *SIAM J. Sci. Statist. Comput.*, 10(6):1063–1075, 1989.
- [13] J. E. Dennis, Jr. and R. B. Schnabel. *Numerical methods for unconstrained optimization and nonlinear equations*. Prentice Hall Series in Computational Mathematics. Prentice Hall Inc., Englewood Cliffs, NJ, 1983.

- [14] D.P. Dobkin, S.J. Friedman, and K. J. Supowit. Delaunay graphs are almost as good as complete graphs. *Discrete Comput. Geom.*, 5(4):399–407, 1990.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] 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.
- [19] 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.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] J. B. Garnett and D. E. Marshall. *Harmonic measure*, volume 2 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2005.
- [24] J. Heinonen and P. Koskela. Quasiconformal maps in metric spaces with controlled geometry. *Acta Math.*, 181(1):1–61, 1998.
- [25] J.M. Keil and C.A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete Comput. Geom.*, 7(1):13–28, 1992.
- [26] 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.
- [27] Ch. Pommerenke. *Boundary behavior of conformal maps*. Grundlehren der Math. Wissenschaften, 299. Springer-Verlag, 1992.
- [28] V. T. Rajan. Optimality of the Delaunay triangulation in \mathbf{R}^d . *Discrete Comput. Geom.*, 12(2):189–202, 1994.
- [29] H.A. Schwarz. Conforme abbildung der oberfläche eines tetraeders auf die oberfläche einer kugel. *J. Reine Ange. Math.*, pages 121–136, 1869. Also in collected works, [30], pp. 84-101.
- [30] H.A. Schwarz. *Gesammelte Mathematische Abhandlungen*. Springer, Berlin, 1890.
- [31] D. Sullivan. Travaux de Thurston sur les groupes quasi-fuchsien et les variétés hyperboliques de dimension 3 fibrées sur S^1 . In *Bourbaki Seminar, Vol. 1979/80*, pages 196–214. Springer, Berlin, 1981.
- [32] J. Väisälä. Free quasiconformality in Banach spaces. II. *Ann. Acad. Sci. Fenn. Ser. A I Math.*, 16(2):255–310, 1991.

C.J. BISHOP, MATHEMATICS DEPARTMENT, SUNY AT STONY BROOK, STONY BROOK, NY 11794-3651

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