Numerical conformal mapping using cross-ratios and Delaunay triangulation

Tobin A. Driscoll* Stephen A. Vavasis†

Cornell ACRI Tech Report CTC96TR233
February 1996

Abstract

We propose a new algorithm for computing the Riemann mapping of the unit disk to a polygon, also known as the Schwarz-Christoffel transformation. The new algorithm, CRDT, is based on cross-ratios of the prevertices, and also on cross-ratios of quadrilaterals in a Delaunay triangulation of the polygon.

The CRDT algorithm produces an accurate representation of the Riemann mapping even in the presence of arbitrary long, thin regions in the polygon, unlike any previous conformal mapping algorithm. We believe that CRDT can never fail to converge to the correct Riemann mapping, but the correctness and convergence proof depend on conjectures that we have so far not been able to prove. We demonstrate convergence with computational experiments.

The Riemann mapping has applications to problems in two-dimensional potential theory and to finite-difference mesh generation. We use CRDT to produce a mapping and solve a boundary value problem on long, thin regions for which no other algorithm can solve these problems.

Keywords: numerical conformal mapping, Schwarz-Christoffel mapping, cross-ratios, Delaunay triangulation, crowding

AMS subject classifications: 30C30, 65E05

1 Introduction

Let P be an open, simply-connected, nonempty subset of the complex plane $\mathbf{C}$ that is not the entire plane. The celebrated Riemann Mapping Theorem [8] states that there is an analytic function f with a nowhere-vanishing derivative such that $f(D) = P$, where D denotes the open unit disk, and such that f is bijective on D. Furthermore, let z_0 be

*Center for Applied Mathematics, Cornell University, Ithaca, NY 14853, `driscoll@cam.cornell.edu`.
This work was funded in part by DOE Grant DE-FG02-94ER25199 to L. N. Trefethen.

†Department of Computer Science, Cornell University, Ithaca, NY 14853, `vavasis@cs.cornell.edu`.
This author's work was supported in part by an NSF Presidential Young Investigator grant with matching funds from Xerox, AT&T, Sun and Tektronix. This work was partially supported by Xerox Corporation, during a visit to the Xerox Palo Alto Research Center. This work was partially supported by the Research Institute for Advanced Computer Science (RIACS) during a visit to RIACS, sponsored by NASA under contract NAS 2-13721.

an arbitrarily specified point in P and let α be an arbitrary angle in $[0, 2\pi)$. Then f can be chosen so that $f(0) = z_0$ and $\arg f'(0) = \alpha$. With such a specification, f is uniquely determined. The point z_0 is called the *conformal center* of f.

This mapping f can be used to solve problems in potential theory posed on the original domain P. The classical example is Laplace's equation $\triangle u = 0$ with Dirichlet data. Because f preserves Laplace solutions, the original problem is reduced to solving Laplace's equation on a disk, which can be done via Fourier transforms.

A second application is finite-difference mesh generation. Because the conformal mapping f preserves angles, any grid on the disk with orthogonal grid-line intersections is mapped by f to a grid on P whose grid lines also meet orthogonally. Orthogonal grid lines simplify the task of discretizing a differential operator via finite difference approximations.

In the case that P is a simple polygon, the Riemann mapping can be written down almost in closed form. Let P be a finite polygon whose boundary is piecewise linear, with no interior angles equal to 0 (no cusps). Let the vertices of P in counterclockwise order be denoted $z_1, \ldots, z_n$. Let the interior angles at $z_1, \ldots, z_n$ be $\alpha_1, \ldots, \alpha_n$, and let $\beta_j = \alpha_j/\pi - 1$ for each j, so that $\beta_j \in (-1, 1]$ for each j. (If $\beta_j{=}1$, z_j is the tip of a slit, and the sides adjacent to z_j coincide at least partially.) Then any conformal mapping $f : D \to P$ has the form [8]

$$f(w) = A + B \int_0^w \prod_{j=1}^{n} \left(1 - \frac{\omega}{w_j}\right)^{\beta_j} d\omega, \tag{1}$$

where A, B are complex parameters ($B \neq 0$), $w_1, \ldots, w_n$ are points in counterclockwise order on the boundary of the unit disk, and the integral denotes a complex contour integral. The points $w_1, \ldots, w_n$ are called *prevertices*; they map to the points $z_1, \ldots, z_n$ under f. Formula (1) is known as the *Schwarz-Christoffel (S-C) formula.*

The Schwarz-Christoffel formula is not quite in closed form, because there is no explicit expression for the $n + 4$ real parameters $A, B, \theta_1, \ldots, \theta_n$, where $\theta_i = \arg w_i$. In practice, these parameters are determined by solving a system of nonlinear equations derived from geometric constraints. Any particular specification of the unknown parameters will yield some polygon (possibly covering parts of $\mathbf{C}$ more than once) whose side lengths and orientation can be measured and compared to the desired image polygon. By using ratios of sides, we can eliminate the affine scaling constants from the system, and by arbitrarily specifying the three degrees of freedom in the mapping, we can reduce the size of the square nonlinear system to $n - 3$ [18].

Actual software packages for S-C mapping like SCPACK [17] and its cousin, the SC Toolbox for MATLAB [7], solve such nonlinear systems numerically, after applying a transformation to the primitive variables that eliminates the need for explicit enforcement of the ordering constraints on the θ_j's. But two difficulties limit the generality of polygons these packages can map:

- The system of nonlinear equations does not have any special structure that lends itself to easy solution. In fact, the system can have local minima that can trap nonlinear solvers and prevent convergence entirely [9].

- More important, SCPACK and the SC Toolbox cannot generally handle *crowding*, a phenomenon of conformal mapping that occurs whenever the domain has any long, narrow channel [14]. The effect of such a channel is to make the prevertex positions badly skewed, to a degree exponential in the aspect ratio of the narrow region (see Fig. 5). The canonical example of crowding is a rectangle R of side lengths a and 1, where $a \gg 1$. If the conformal center is chosen at the center of R and z_1, z_2 are endpoints of one of the short edges of R, then $\theta_2 - \theta_1 = O(e^{-a\pi/2})$. Thus, if the aspect ratio is 20 to 1, then at least 14 significant digits are lost when computing a difference between two θ_j's.

One partial solution to the problem of crowding in the SC Toolbox is its provision for mapping elongated domains to rectangles. The elongation in the domain can be matched by a similar elongation in image rectangle, alleviating the crowding problem [9, 12]. While this technique can be generalized to certain classes of multiply elongated polygons [11], the technique becomes much more delicate, and the fundamental domain is no longer a disk.

We propose a new algorithm that remedies both of these deficiencies. There are two principal innovations in the new algorithm:

- Our algorithm uses as primitive variables certain *cross-ratios* of the w_j's. Cross-ratios are defined in Section 4. Because cross-ratios are invariant under fractional linear transformations, we can compute many different embeddings of the w_j's that are all conformally equivalent and hence yield the same polygon. In particular, when evaluating f for one part of the domain, we can recompute the w_j's so that no crowding occurs near the points where f needs to be evaluated. Thus, crowding is no longer a problem.
- Our system of equations enforces the constraints that certain absolute cross-ratios come out correctly in the image polygon (rather than enforcing conditions about side lengths and orientations as above). These cross-ratios in the image polygon appear to be strongly correlated with the corresponding primitive-variable cross-ratios. The resulting nonlinear system appears to have a monotonicity property that makes it much easier to solve than the other formulations.

The CRDT algorithm consists of the following steps:

1. Some of the edges of the polygon P are split, i.e., new vertices are introduced whose angles are π. Let the number of vertices in the split polygon be n.
2. A Delaunay triangulation of P with the new vertices is computed.
3. A solver is called for an implicitly specified $(n-3) \times (n-3)$ system of nonlinear equations. The variables of these equations are the $n-3$ cross-ratios of the w_j's associated with the Delaunay triangulation. Each of the $n-3$ equations enforces a constraint that a cross-ratio in the image polygon comes out correctly.

CRDT stands for "cross-ratios of the Delaunay triangulation."

The remainder of the paper is organized as follows. In Section 2 we provide the definition and basic properties of the constrained Delaunay triangulation. In Section 3 we describe the splitting step of our algorithm. In Section 4 we define cross-ratios and establish some of their properties. In particular, we show that $n-3$ cross-ratios uniquely determine the image polygon under (1) up to similarity transform. In Section 5 we present the whole algorithm and details on how to compute the forward mapping. In Section 6 we explain how CRDT circumvents crowding. In Section 7 we discuss solvers for the nonlinear system and report on experiments with various polygonal domains. In Section 8 we describe how to use CRDT in two applications. We know of no other algorithm that could duplicate the results of Section 8.

2 The Delaunay triangulation

Let P be a simple polygon. A *triangulation* of the polygon is a division of P into nondegenerate triangles such that (a) the intersection of any two triangles is either a common edge of the two, a common vertex, or empty; (b) the union of the triangles is P; (c) all three vertices of every triangle are also vertices of P; and conversely (d) every vertex of P is also a vertex of a triangle.

It is well known that for any n-vertex simple polygon P, there exists a triangulation of P, and furthermore, any triangulation of P has exactly $n-2$ triangles. A triangulation edge that is not a polygon edge is called a *diagonal*. It is also known that any triangulation of P has exactly $n-3$ distinct diagonals.

Let the *dual* of a triangulation be the following $(n-2)$-node graph. The graph has one node for each triangle, and an edge between two nodes if their corresponding triangles have a common diagonal. Thus, the dual graph has exactly $n-2$ vertices and exactly $n-3$ edges; each edge is in correspondence with a diagonal of the triangulation. It is well-known that the dual graph of a triangulation of a simple polygon is always a *tree*. See Fig. 1 for an example of a triangulation and its dual.

Among all triangulations of P, there is a distinguished triangulation known as the *constrained Delaunay triangulation* or just the *Delaunay triangulation* [1]. This is a triangulation with the following (defining) property: If d is a diagonal of the triangulation, let $Q(d)$ be the quadrilateral associated with d, that is, the union of the two triangles on either side of d. Then the sum of the two opposite interior angles of $Q(d)$ that are split by d is at least π.

It can be proved that such a triangulation always exists and, if no four points of P are cocircular, that it is unique. Moreover, there is a simple algorithm that converges to the Delaunay triangulation in $O(n^2)$ steps: First, compute an arbitrary triangulation of P. Find a diagonal d such that the condition in the last paragraph is violated. Then "flip" d—i.e., delete d and replace it with the other diagonal of $Q(d)$, and replace the two triangles formerly adjacent to d with the two triangles thus formed. Repeat until no flips are possible. This is the algorithm used in our computational tests, although a more efficient $O(n \log n)$ algorithm for constrained Delaunay triangulation was developed by Chew [3].

3 Splitting edges

The first step of our algorithm is to split some edges of the polygon. Splitting an edge means replacing it by several smaller edges whose union is equal to the original edge. These new edges are joined by vertices whose angles are π. Notice that this operation does not affect

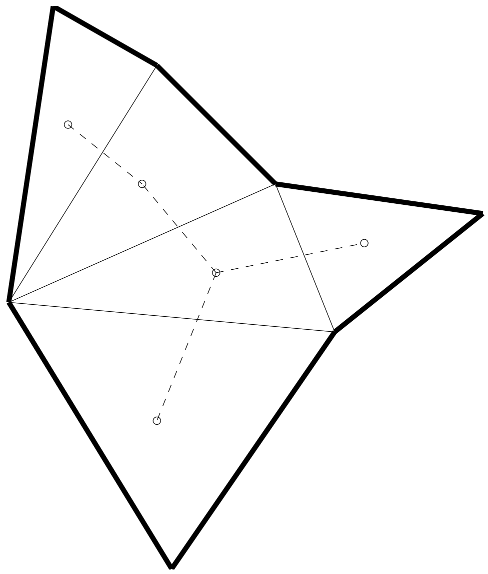

Figure 1: The Delaunay triangulation of a 7-sided polygon. The heavy solid segments are boundary edges, and the light solid segments are diagonals of the triangulation. The dual graph has five nodes (circles) and four edges (dashed lines). The dual is abstract; the geometry shown here is for convenience.

the Schwarz-Christoffel formula (1); a vertex whose interior angle is π has its β exponent equal to 0 in that formula.

The purpose of splitting is to make sure each individual quadrilateral in the Delaunay triangulation is well-conditioned. By "well-conditioned" we mean that the prevertices of the quadrilateral are not too crowded in some valid arrangement of the S-C prevertices. In particular, we want to avoid quadrilaterals that are long and narrow with the long edges equal to polygon edges, because the prevertices of such a quadrilateral will be crowded on the unit circle. (A long and narrow quadrilateral is acceptable provided that the polygon is "fat" around it. See Fig. 2.)

The splitting procedure has two phases. First, for every vertex v with an interior angle of $\pi/4$ or less, we chop off the corner at v as follows. Find the largest isoceles triangle T that can be formed by v and its two adjacent edges such that T is contained in P, and introduce new vertices along the two edges that are adjacent to v at the midpoints of the two sides of T. After this split, the two edges adjacent to v are said to be *protected*; that is, we do not allow them to be split during the second phase. Let P' denote the polygon obtained after this first part of the splitting procedure is complete.

The second phase of the edge-splitting procedure is iterative and generates a sequence of polygons, each of which is a subdivision of its predecessor, starting with P'. Let e be an unprotected edge of some polygon occurring in the iteration. Let $l(e)$ be its length. Let $d(e)$ be the smallest distance from e to any foreign vertex, where "foreign" means a vertex other than the endpoints of e, and distance is measured geodesically, i.e., along the shortest piecewise linear path that remains inside the polygon. (It turns out that $d(e)$ can

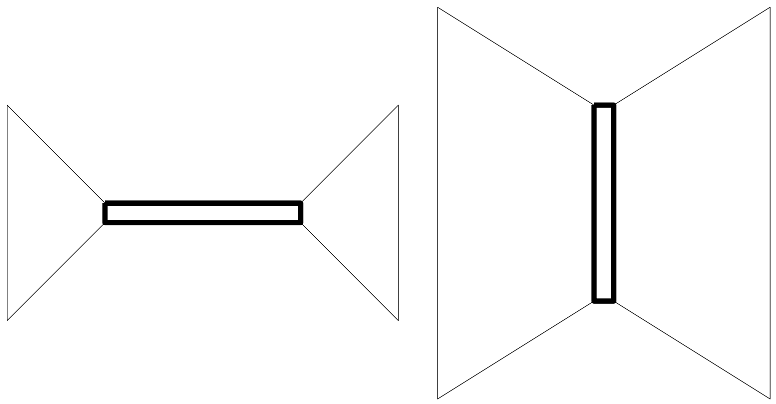

Figure 2: On the left is an octagon with a long, narrow quadrilateral in its triangulation (heavy outline). This quadrilateral would have to have its sides split because of crowding: in a mapping with conformal center at the center of the quadrilateral, the short edges of the quadrilateral are exponentially contracted in the preimage on the disk. In contrast, for the polygon on the right no splitting is necessary, because the polygon is "fat" around the quadrilateral. That is, the quadrilateral can be enclosed in a disk that is mostly interior to the polygon.

be determined efficiently given a triangulation of the polygon.) Then we say e is *ill-separated* if

$$d(e) < l(e)/(3\sqrt{2}). \tag{2}$$

At each iteration we identify all ill-separated edges and split them into three equal pieces. We repeat this until all edges are well-separated. See Fig. 3 for an example of both phases of the splitting process. The splitting of edges and protecting of sharp angles is reminiscent of techniques previously introduced in the finite-element mesh generation literature; see for example [2, 4, 15]. In the mesh generation literature, the purpose of these techniques is similar to our own purpose, namely, to prevent the occurrence of poorly shaped triangles that could arise in a triangulation of the original (unsplit) polygon. The main difference is that finite-element mesh generation subdivides the interior of the domain as well as its boundary and thus would avoid both kinds of long, skinny quadrilaterals illustrated in Fig. 2.

We do not try to prove that the splits computed by this procedure are "effective" for our algorithm, because we do not yet have an *a priori* characterization of well-conditioned quadrilaterals. We do prove, however, that the second phase of the splitting procedure described in this section always terminates after a finite number of steps. Let $r(e)$ be the geodesic distance from edge e to the closest foreign edge. (A foreign edge is one that is not adjacent to e.) Let r_0 be the minimum of $r(e)$ over all unprotected edges of P'. Notice that there can be no edge shorter than r_0 in P', for if there were an edge $e_0 = (v_1, v_2)$ of length shorter than r_0, let e_1 be the other edge whose endpoint is v_1 and let e_2 be the other edge whose endpoint is v_2. Then one checks that $\mathrm{dist}(e_1, e_2) \leq l(e_0) < r_0$, contradicting the choice of r_0.

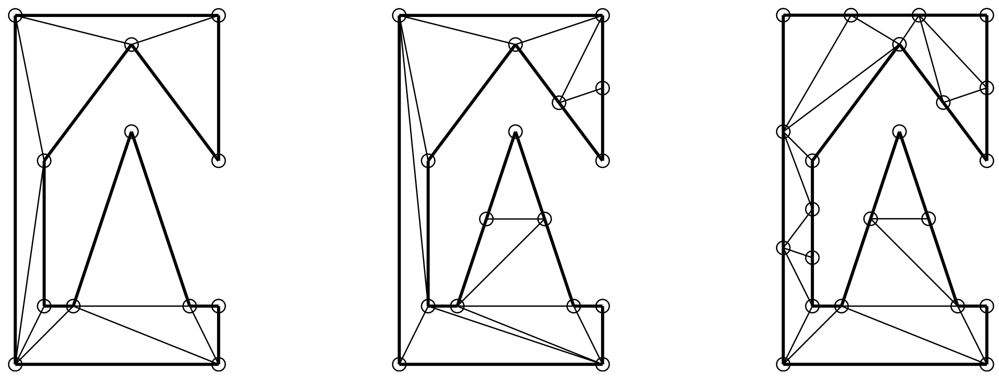

Figure 3: On the left is the Delaunay triangulation of a polygon. The middle shows the Delaunay triangulation after the sharp corners have been chopped in the first splitting phase. On the right is the subsequent result of the second phase, in which narrow regions are subdivided and the Delaunay triangulation is recomputed.

We claim that the splitting procedure above never produces an edge shorter than r_0. To see this, let $e = (v_1, v_2)$ be an unprotected edge of a polygon at some intermediate stage of the above algorithm whose length is less than $3r_0$. We must argue that we could never split e. By induction, let us assume that no edges up to now are shorter than r_0. Let e_0 denote the original edge of P' that contains e. Let v be the foreign vertex closest to e, i.e., $\mathrm{dist}(v, e) = d(e)$. There are three cases: v lies on an edge that was foreign to e_0 in P'; it lies on an edge that was adjacent to e_0; or it lies on e_0 itself.

In the first case, we know that $\mathrm{dist}(e_0, v) \geq r_0$ and hence $\mathrm{dist}(e, v) \geq r_0$. But $l(e) < 3r_0$, so (2) is not satisfied and e is not split.

In the second case, let e_0' be the original edge that contains v, so that e_0 and e_0' are adjacent; say their common point is v'. Since e_0 cannot be protected by assumption, the interior angle at v' greater than $\pi/4$. By assumption, the distance from v' to v is at least r_0. Therefore, some simple trigonometry shows that the distance from v to e_0 is greater than $r_0/\sqrt{2}$. Thus, $d(e) > r_0/\sqrt{2}$ whereas $l(e) < 3r_0$, so (2) is not satisfied.

In the third case v is collinear with e, and its distance from e again must be at least r_0, so the same reasoning shows that (2) is not satisfied.

4 Cross-ratios and embeddings

Let a, b, c, d be four distinct points in the complex plane such that the order $abcd$ forms a quadrilateral with counterclockwise vertex order and such that ac is an interior diagonal of the quadrilateral. We define the *cross-ratio* of these points to be

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

Note the identity $\rho(a, b, c, d) = \rho(c, d, a, b)$. Thus, the cross-ratio depends on the quadrilateral and the diagonal, but not on which endpoint of the diagonal we start at.

In general, the cross-ratio is a complex number, but there is an important special case when it is real.

Lemma 1 *Let a, b, c, d be four distinct points on a circle in counterclockwise order. Then $\rho(a, b, c, d)$ is a negative real number.*

Proof. The angle of the quadrilateral $abcd$ at a and the angle at c are inscribed in complementary arcs of the circle, so the sum of these angles must be π. A quick diagram shows that $(d-a)/(b-a)$ has its arg equal to the angle at a, and $(b-c)/(d-c)$ has its arg equal to the angle at c. Therefore, the arg of the cross-ratio, which is the sum of these args, is π. ∎

As mentioned in the introduction, the $n-3$ primitive real variables of the nonlinear system arise from $n-3$ cross-ratios of prevertices. The preceding lemma confirms that these variables are indeed real. However, cross-ratios are not quite suitable as variables, because we would have to impose side constraints that they be negative. Instead, our unconstrained primitive variables are logarithms of the negatives of the cross-ratios (see (3) below).

We now explain *which* $n-3$ cross-ratios we use. Assume the vertices $z_1, \ldots, z_n$ of the polygon P are given in counterclockwise order. Let $d_1, \ldots, d_{n-3}$ and $Q(d_1), \ldots, Q(d_{n-3})$ be the $n-3$ diagonals and associated quadrilaterals of the Delaunay triangulation of P as defined in Section 2. Let the vertices of $Q(d_i)$, for $i = 1, \ldots, n-3$, be denoted by $(z_{\kappa(i,1)}, z_{\kappa(i,2)}, z_{\kappa(i,3)}, z_{\kappa(i,4)})$, so that for each i, $(\kappa(i,1), \kappa(i,2), \kappa(i,3), \kappa(i,4))$ is a four-tuple of distinct indices in $\{1, \ldots, n\}$.

Then the ith primitive variable σ_i is defined to be

$$\sigma_i = \ln(-\rho(w_{\kappa(i,1)}, w_{\kappa(i,2)}, w_{\kappa(i,3)}, w_{\kappa(i,4)})), \qquad i = 1, \ldots, n-3. \tag{3}$$

It is apparent that given a list of prevertices $w_1, \ldots, w_n$, the primitive variables $\sigma_1, \ldots, \sigma_{n-3}$ are easily computed from (3). For evaluating the nonlinear CRDT mapping, the process must be reversed. The remainder of this section explains how to find $w_1, \ldots, w_n$ on the unit circle to satisfying (3) given $\sigma_1, \ldots, \sigma_{n-3}$.

Notice that there are three degrees of freedom, because (3) imposes only $n-3$ real constraints on n real variables. We will use the flexibility afforded by these degrees of freedom to our advantage; indeed, they are the reason that the CRDT algorithm avoids problems with crowding, as discussed in Section 6. For now, let us fix these degrees of freedom by assuming that the three prevertices corresponding to a Delaunay triangle T_0 are arbitrarily placed on the unit circle in a manner preserving their ordering. (Later, we will show that the choice of T_0 and the three prevertex positions will not affect the S-C image. See Theorem 2 at the end of this section.)

We now show how to embed the remaining $n-3$ vertices using the cross-ratio information, starting with a lemma that tells us how to place a single prevertex.

Lemma 2 *Given distinct points a, b, c on the unit circle in counterclockwise order, and given a negative real number ρ_0, there exists a unique point d on the unit circle such that $\rho(a, b, c, d) = \rho_0$. Furthermore, this point is counterclockwise from c and clockwise from a.*

Proof. If we write out the formula and substitute, we get an explicit closed formula for d,

$$d = \frac{hc + a}{h + 1}, \tag{4}$$

where

$$h = \frac{\rho_0(b - a)}{(c - b)}.$$

We must first show that $h \neq -1$ so that the denominator in the formula for d is nonzero. But this is obvious, because $(b - a)/(c - b)$ must have a nonzero imaginary part (since a, b, c cannot be collinear), so h also has a nonzero imaginary part. Thus d is uniquely determined.

Next, we must show that d lies on the unit circle between c and a. Consider sliding a test point along the unit circle starting from very near a clockwise to c. It is easy to check that the cross-ratio, which is a negative real number by the earlier lemma, varies continuously from 0 to $-\infty$. Therefore, its value must be ρ_0 at some intermediate point. But the last paragraph shows that this intermediate point is unique. ∎

Now the first main theorem of this section tells us that, assuming the first three prevertices are determined, we can place the remaining $n - 3$ prevertices uniquely.

Theorem 1 *Let T_0 be any triangle in the Delaunay triangulation of P, and let its vertices in counterclockwise order be indexed as z_ϕ, z_ψ, z_χ. Suppose the prevertices w_ϕ, w_ψ, w_χ, are specified as distinct points on the unit circle in counterclockwise order. Then, given any real-number values for the primitive variables $\sigma_1, \ldots, \sigma_{n-3}$, there exists a unique solution to* (3)*, that is, a unique way to define the remaining $n-3$ w_i's on the unit circle satisfying* (3)*. The algorithm to find the w_i's satisfying* (3) *is linear time. Furthermore, for this solution to* (3)*, $w_1, \ldots, w_n$ will be in the correct counterclockwise order.*

We call such a placement of the prevertices an *embedding.*

Proof. This proof relies heavily on the fact that the dual of the Delaunay triangulation is a tree, as mentioned in Section 2. In the ensuing discussion, "nodes" and "edges" refer to nodes and edges of the tree, whereas "vertices" and "diagonals" refer to vertices and diagonals of the Delaunay triangulation.

In a tree, there is always a unique path between any pair of nodes. Let us root the tree at T_0. Every node in a rooted tree except the root has a parent, and every node except the leaves has children. Therefore, for each triangle in the triangulation except the root, we can identify the diagonal that separates it from its parent; we call this the *entry diagonal* of the triangle. The vertex opposite the entry diagonal is called the *new vertex.* Notice that choice of entry diagonal and new vertex depends not only on the triangle but on the choice of T_0 as well.

We next claim that the $n - 3$ new vertices of the $n - 3$ nonroot triangles are precisely the $n - 3$ vertices of the polygon whose prevertices are to be determined. It is clear that no new vertex can be a vertex of T_0. Furthermore, two distinct nonroot triangles cannot have a common new vertex. The justification for this claim is provided by Fig. 4. In particular, the figure shows that if two triangles had the same new vertex, then there would be a cycle in the dual, contradicting the fact that the dual is a tree. Let the nonroot triangles and

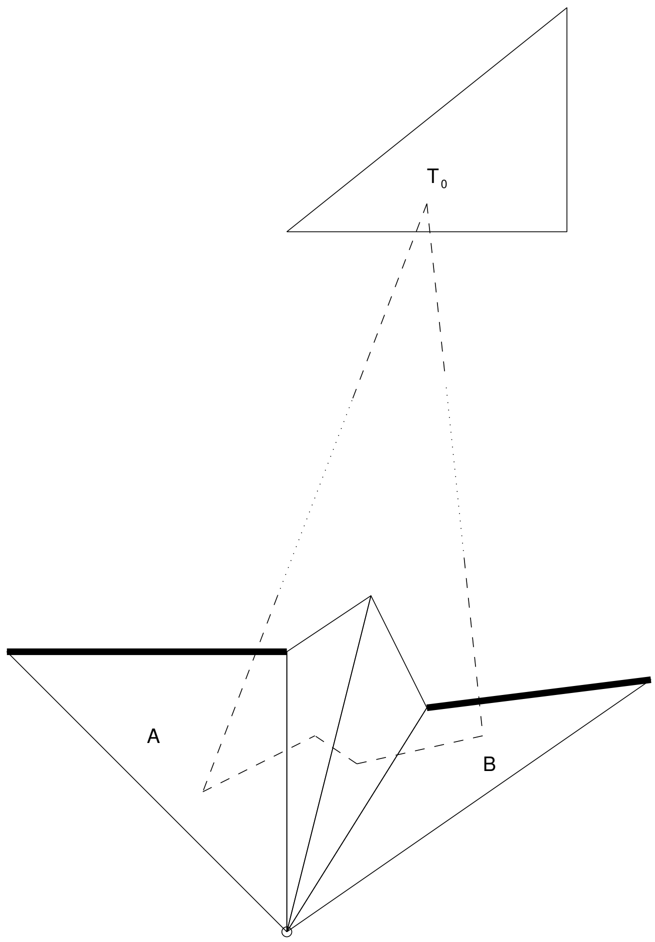

Figure 4: An illustration of the contradiction that would arise if two different triangles in the triangulation had the same new vertex. If triangles A and B, whose entry diagonals are shown as heavy lines, both shared this new vertex (marked by the circle), there would be a cycle in the dual graph (marked by dashed lines).

their new vertices be denoted by $T_1, \ldots, T_{n-3}$ and $z_{i_1}, \ldots, z_{i_{n-3}}$, respectively. Thus, the disjoint union of index sets $\{i_1, \ldots, i_{n-3}\}$ and $\{\phi, \psi, \chi\}$ is $\{1, \ldots, n\}$. Let the entry diagonal of nonroot triangle T_k be denoted by d_{j_k}, which is associated with quadrilateral $Q(d_{j_k})$. Each entry diagonal corresponds to exactly one of the primitive variables, because each diagonal (and its quadrilateral) comes up exactly once in the list for $k = 1, \ldots, n-3$.

We are now ready to describe our algorithm for embedding the $n-3$ remaining prevertices. Visit each node of the tree once. When visiting the node for triangle T_k, compute the position of prevertex w_{i_k}. Any search order that guarantees that parents are visited before their children (for instance, depth-first or breadth-first) is acceptable.

In more detail, we use (4) for computing w_{i_k} because we know that w_{i_k} corresponds to the quadrilateral $Q(d_{j_k})$, which is the union of T_k and its parent in the tree. The other three vertices of the quadrilateral are already embedded by previous steps in the search (because the parent is embedded before the child). Given the position of the three other vertices of this quadrilateral and the cross-ratio $-\exp(\sigma_{j_k})$, prevertex w_{i_k} is uniquely determined because of Lemma 2. But then an induction argument shows that the positions of all the w_{i_k}'s are uniquely determined, i.e., every step of the construction is forced. This shows uniqueness of the solution.

Furthermore, (3) is satisfied for this construction for $i = 1, \ldots, n-3$ because we used each σ_i exactly once in the preceding construction.

The only remaining claim is that $w_1, \ldots, w_n$ will end up in counterclockwise order. Again, this follows from a combination of Lemma 2 and the fact that the dual is a tree. Let T_k be the current triangle and z_{i_k} its new vertex. Observe that Lemma 2 ensures that when we place w_{i_k} on the unit circle, it will be on the arc between the endpoints (call them w_a and w_b) of the entry diagonal of T_k that is complementary to the arc that contains the parent triangle's prevertices. Therefore, w_{i_k} is placed correctly with respect to w_a and w_b. But notice that w_{i_k} must be the first prevertex placed between w_a and w_b because any other prevertices on this arc are the new vertices of children of T_k. Thus, the placement of w_{i_k} is correct with respect to all vertices placed before it. ∎

This theorem shows that given values for the primitive variables and the positions of the three prevertices corresponding to a Delaunay triangle, we can compute the positions of all of the prevertices. How should we choose the initial triangle and embed its prevertices? It turns out that *any* initialization is acceptable; all embeddings give the same image polygon, up to similarity transform.

To show this, we begin by recalling some standard facts from complex analysis [16]. First, any conformal map from the unit disk to itself is a fractional linear transformation of the form

$$g(z) = e^{i\theta} \frac{z - r}{1 - \bar{r} z}, \tag{5}$$

where r is a complex number such that $|r| < 1$ and $\theta \in [0, 2\pi)$. As a consequence, there is a unique such mapping that takes any three distinct points on the unit circle to any other three points on the unit circle, preserving ordering [16]. We formalize another consequence in a lemma:

Lemma 3 *Let a, b, c, d be four points on the unit circle in counterclockwise order. Let g be* *a conformal mapping of the unit disk to itself. Then*

$$\rho(a, b, c, d) = \rho(g(a), g(b), g(c), g(d)).$$

Proof. Simple algebra verifies that cross-ratios are invariant under fractional linear transformations. See [16] for details. ∎

We now show that no matter which initial triangle T_0 we select in Theorem 1, and no matter how we choose the positions of its three prevertices, we end up with the same image polygon, up to similarity transform.

Theorem 2 *Let $\sigma_1, \ldots, \sigma_{n-3}$ be given. Let T_0, T_0' be two triangles in the Delaunay triangulation whose vertices are (z_ϕ, z_ψ, z_χ) and $(z_{\phi'}, z_{\psi'}, z_{\chi'})$ respectively. Let (w_ϕ, w_ψ, w_χ) and $(w'_{\phi'}, w'_{\psi'}, w'_{\chi'})$ be order-preserving embeddings of their prevertices in the unit disk as in Theorem 1. Use the construction of Theorem 1 to produce the two embeddings $(w_1, \ldots, w_n)$ and $(w'_1, \ldots, w'_n)$, respectively. Let $\tilde{P}$, $\tilde{P}'$ be the images of the unit disk under (1) using these two prevertex sets. Then $\tilde{P}$ and $\tilde{P}'$ are* similar, *i.e., they agree up to a translation, rotation, and scaling.*

Remark. We have not explained how to choose the affine constants A and B in (1). The theorem holds for any choice of these constants. Alternatively, the theorem asserts that, given the constants for one embedding, these constants can be chosen for the other embedding in such a way that the image polygons coincide.

Proof. Let $(w'_\phi, w'_\psi, w'_\chi)$ be the positions of the prevertices of T_0 in the second embedding. Let g be the conformal mapping of the unit disk to itself carrying the points (w_ϕ, w_ψ, w_χ) to $(w'_\phi, w'_\psi, w'_\chi)$. That is, g maps the prevertices of T_0 in one embedding to the prevertices of T_0 in the other.

We claim that $g(w_i) = w'_i$ for *all* i, not just $\{\phi, \psi, \chi\}$. Because g preserves cross-ratios, the embedding $(g(w_1), \ldots, g(w_n))$ has the same $n-3$ cross-ratios as $(w_1, \ldots, w_n)$, which by assumption has the same $n-3$ cross-ratios as $(w'_1, \ldots, w'_n)$. But since three entries in $(g(w_1), \ldots, g(w_n))$ are equal to the three corresponding entries in $(w'_1, \ldots, w'_n)$, the uniqueness part of Theorem 1 guarantees that $g(w_i) = w'_i$ for all i.

Now consider composing (1) with the conformal mapping g^{-1} of the disk to itself. The composition is a conformal mapping from the unit disk to $\tilde{P}$, so the S-C formula (1) must also hold when the prevertices are given by $(w'_1, \ldots, w'_n)$, for a suitable choice of the affine constants. So $\tilde{P}$ and $\tilde{P}'$ are similar. ∎

5 The CRDT algorithm

We are now prepared to specify the CRDT algorithm in more detail.

Step 1. Split the edges of the polygon as described in Section 3. We again use P to denote the polygon obtained after splitting. Let n be the number of its vertices.

Step 2. Compute the Delaunay triangulation of P. Now that the Delaunay triangulation is computed, we can fix a particular numbering of the diagonals and quadrilaterals in the triangulation. Recall the notation $\kappa(i,j)$ used in (3) to number the vertices of the quadrilaterals.

Let us define

$$c_i = \ln(|\rho(z_{\kappa(i,1)}, z_{\kappa(i,2)}, z_{\kappa(i,3)}, z_{\kappa(i,4)}|), \tag{6}$$

for $i = 1, \ldots, n-3$. Note that the cross-ratio in this formula in general will be a complex number, so the absolute value symbols denote magnitude.

Step 3. Solve the nonlinear system $F(\boldsymbol{\sigma}) = \mathbf{0}$. The map $F : \mathbf{R}^{n-3} \to \mathbf{R}^{n-3}$ is defined as follows. The input variables are the primitive variables $\sigma_1, \ldots, \sigma_{n-3}$ defined by (3). It was shown in Section 4 that there is a unique (up to similarity) Schwarz-Christoffel mapping that can be computed from these primitive variables. Let $\zeta_1, \ldots, \zeta_n$ be the vertices of the image of (1). For $i = 1, \ldots, n-3$, let

$$F_i(\sigma_1, \ldots, \sigma_{n-3}) = \ln(|\rho(\zeta_{\kappa(i,1)}, \zeta_{\kappa(i,2)}, \zeta_{\kappa(i,3)}, \zeta_{\kappa(i,4)})|) - c_i.$$

Observe that although the ζ's themselves are determined only up to similarity transform, the cross-ratio of four of them is invariant under similarity transform, so this definition makes sense.

We discuss nonlinear solvers in more detail in Section 7.

Note that at a solution to $F(\boldsymbol{\sigma}) = \mathbf{0}$, we have for $i = 1, \ldots, n-3$ that

$$|\rho(\zeta_{\kappa(i,1)}, \zeta_{\kappa(i,2)}, \zeta_{\kappa(i,3)}, \zeta_{\kappa(i,4)})| = |\rho(z_{\kappa(i,1)}, z_{\kappa(i,2)}, z_{\kappa(i,3)}, z_{\kappa(i,4)})|.$$

Furthermore, we know that all the interior angles of the polygon determined by $\zeta_1, \ldots, \zeta_n$ are correct because the angles are inherent in the Schwarz-Christoffel mapping. Is this polygon the correct one? We are not able to prove this, so we state it as a conjecture.

Conjecture 1 *Let P be an n-vertex triangulated simple polygon. Then P is uniquely determined up to similarity transform by the following data:*

- *the sequence of all interior angles of P at its vertices, and*
- *the list of $n-3$ absolute values of cross-ratios of the quadrilaterals determined by the triangulation of P.*

We have verified this conjecture analytically in the cases of $n = 4$ and $n = 5$. In practice, it is easy to check whether the right polygon has been computed. Our computational experiments support its validity in general: the CRDT algorithm has never failed to converge to the correct polygon.

If this conjecture turns out to be false, we can modify CRDT so that the $n-3$ equations enforce some condition about side lengths that guarantees that the obtained solution is correct. The advantage of equations that enforce cross-ratio conditions is that the system of nonlinear equations apparently has a desirable monotonicity property, described in Section 7.

We conclude this section with a description of our algorithm for computing F given a value of $\boldsymbol{\sigma}$. For each component F_i for some $i = 1, \ldots, n-3$, we need to know four image vertices $\zeta_{\kappa(i,1)}, \ldots, \zeta_{\kappa(i,4)}$. To find these, we construct a certain embedding E_i of the prevertices. Recall from Section 4 that given the vector $\boldsymbol{\sigma}$, we can arbitrarily embed three of the prevertices, such as $w_{\kappa(i,1)}, w_{\kappa(i,2)}, w_{\kappa(i,3)}$. We embed these three in such a way so that when $w_{\kappa(i,4)}$ gets placed by the algorithm in Section 4, these four prevertices on the unit disk will be arranged in a rectangle centered at the origin with the correct cross-ratio (i.e., cross-ratio $-\exp(\sigma_i)$). Then, as in Section 4, we position the rest of the prevertices to complete E_i. This embedding defines an S-C map f_i, given by (1) with $A = 0$ and $B = 1$, which we use to compute the relative positions of the four image vertices $\zeta_{\kappa(i,1)}, \ldots, \zeta_{\kappa(i,4)}$. The path of integration is a straight-line segment from the origin, and we use the compound Gauss-Jacobi quadrature rules described by Trefethen [17].

6 On the circumvention of crowding

In this section we explain the crux of our claim that CRDT is unaffected by crowding. If the domain P contains any long, narrow channel, then for any possible correct embedding of the prevertices, some of the prevertices will be extremely crowded. But the embedding E_i (described at the end of the last section) for computing the ith component of F guarantees that $w_{\kappa(i,1)}, \ldots, w_{\kappa(i,4)}$ will *not* be crowded, either against each other or against any other prevertex. Therefore, the crowding has no impact on the accuracy of the quadrature rule applied to these four prevertices, because the path of integration never passes close to crowded prevertices. Thus, it is the flexibility to re-embed the prevertices for each coordinate entry of F, along with the splitting of narrow channels, that allows us to circumvent crowding. See Fig. 5.

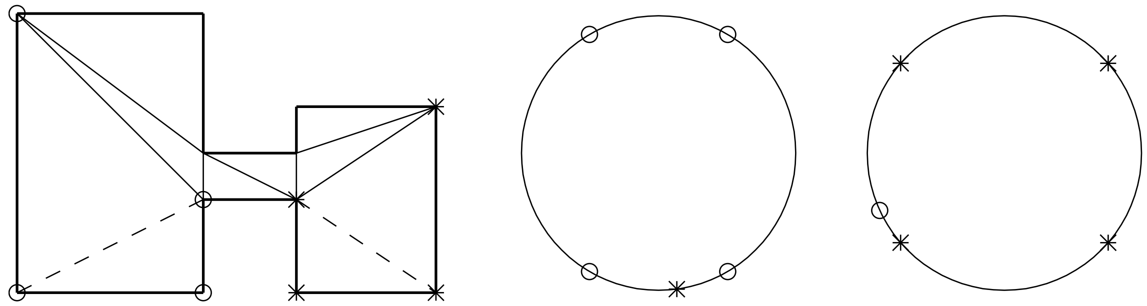

Figure 5: On the left is a triangulated polygon with two distinguished quadrilaterals whose diagonals are given as dashed lines. No embedding of the S-C prevertices will be universally uncrowded. In the middle picture we see an embedding that keeps the prevertices of the first quadrilateral (marked with circles) uncrowded, while the prevertices of the other (marked with stars) are crowded—too closely to be distinguished. However, in another embedding (far right), the crowding situation is reversed.

In order to substantiate the claim in the last paragraph that none of the four prevertices are crowded against each other or their neighbors, we need a result stating that none of the cross-ratios $\rho_1, \ldots \rho_{n-3}$ of the prevertices is very large (close to $-\infty$) or very small (close to 0). Unfortunately, there cannot exist fixed (constant) upper or lower bounds on these cross-ratios that apply to all polygons, as the following example shows. Consider CRDT applied to the regular n-gon. Note that any triangulation of the regular n-gon is a Delaunay triangulation. Thus, CRDT might compute a triangulation that has a quadrilateral whose aspect ratio is $O(n)$. Since the Schwarz-Christoffel mapping of an n-gon is close to the identity mapping, there will also be four prevertices whose cross-ratio is $O(n^2)$, or $O(n^{-2})$. Thus, there is no *a priori* upper or lower bound possible on the ρ_i's, and therefore none on the σ_i's either. This growth of the σ_i's is very slow (logarithmic in n) and is thus not expected to have a significant impact on the accuracy of CRDT. But the absence of a constant upper bound or lower bound means, for example, that there is no *a priori* upper bound on how much adaptation is necessary in the compound Gauss-Jacobi integration used to evaluate (1).

The example in the last paragraph has bad triangles in the original Delaunay triangulation, so one could argue that the growth of the cross-ratios of the prevertices as $n \to \infty$ is unavoidable. Thus, a more plausible conjecture might be that the difference $\sigma_i - c_i$, (where c_i was defined by (6)) has constant upper and lower bounds. This result would also show that CRDT is not affected by crowding, in the sense that it never works with distances that are substantially shorter than edge-lengths in the original polygon. Note that both quantities σ_i and c_i are logarithms, so subtracting them is the right way to check how they differ. We do not know whether this conjecture is true, but we have not seen substantial divergence between σ_i and c_i in our computational experiments. In Fig. 6 we present a histogram of $\sigma_i - c_i$ for all polygons in our experiments, for all indices i, at the final solutions obtained by the CRDT algorithm. Notice that points in this histogram lie in a fairly narrow range. (The cross-ratios stay similarly bounded in intermediate steps of the iterations because of the monotonicity observed in our nonlinear equations, described in the next section.)

The representation of prevertices used by the SC Toolbox and SCPACK is also well-conditioned in the presence of crowding, because logarithms of distances between prevertices are used as primitive variables rather than the prevertex positions themselves. But the logarithmic distances cannot be used directly to compute S-C maps, so both packages pass (in an intermediate step) from these primitive variables to particular fixed prevertex positions. Hence forward evaluation of F is spoiled by cancellation in these packages because of this intermediate step. In CRDT, our representation is also unaffected by crowding, and, furthermore, when evaluating F we can work directly with the well-conditioned cross-ratio representation (by juggling the prevertices around so that the ones of interest are never crowded) to avoid cancellation.

Does this mean that CRDT "solves" all problems with crowding? If CRDT is applied to a problem to potential theory, and the application requires (as an intermediate step) a fixed embedding of the prevertices, then there is no improvement with CRDT compared to the Toolbox because crowding will presumably spoil the intermediate step. Therefore, if one wants to use CRDT as a subroutine, then one must devise an algorithm in which the cross-ratios produced by CRDT are used to shuffle between different embeddings of the prevertices. We give some examples of problems from potential theory that can be solved in this manner by CRDT in Section 8. In fact, for every reasonable problem in potential theory that we have considered, we have always been able to come up with a way to use the cross-ratios directly and avoid crowding. But in each case the technique we have devised is slightly different, so we are not able to state with certainty that CRDT can solve all

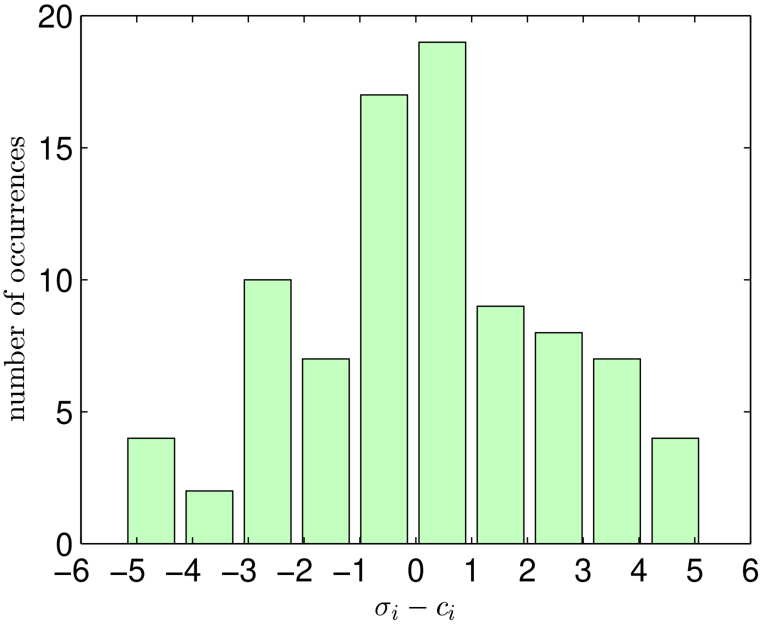

Figure 6: Evidence of controlled cross-ratio sizes. The histogram depicts the discrepancies $\sigma_i - c_i$ between the primitive cross-ratio variables and polygon cross-ratios for the experiments in Section 7. The variables are logarithmic. The clustering near zero and absence of large outliers supports the claim that CRDT circumvents crowding.

problems in potential theory posed on elongated polygons.

7 Computational experiments with CRDT

In this section we experiment with an implementation of CRDT in MATLAB. We present some evidence of the monotonicity of the nonlinear system and consider the matter of solving it numerically. We also compare the performance of the CRDT algorithm to the SC Toolbox for MATLAB [7] and find that CRDT is competitive for most regions. The principal exceptions are regions which cause the edge-splitting algorithm to add a great many extra vertices of angle π. While such vertices do not affect the amount of work in computing the S-C integral (1), they do affect the size of the nonlinear system to be solved.

The computational core of CRDT is solving the nonlinear system $F(\boldsymbol{\sigma}) = \mathbf{0}$. We consider nonlinear solvers that require only function evaluations. According to our experiments, a particular very simple linear iteration always converges at a linear rate. This iteration starts with $\boldsymbol{\sigma}^{(0)} = \mathbf{c}$, where c_i was defined above by (6). Then we iterate:

$$\boldsymbol{\sigma}^{(k+1)} = \boldsymbol{\sigma}^{(k)} - F(\boldsymbol{\sigma}^{(k)}). \tag{7}$$

Our conclusion from experiments is that this iteration always satisfies $||F(\boldsymbol{\sigma}^{(k+1)})||_2 \leq \alpha||F(\boldsymbol{\sigma}^{(k)})||_2$, for an α that is problem dependent but always satisfies $\alpha < 1$. We have not been able to come up with a convincing explanation for this behavior. The essential reason is apparently that the Jacobian F' approximates the identity, but none of the likely conditions on F' that would support this claim have been found to hold in experiments.

The fact that (7) converges linearly, and that F' appears to be like the identity, make us suspect that F has some strong monotonicity property. Expected consequences of this monotonicity are that $\|F\|$ has no local minima and that F is injective. In contrast, experiments indicate the nonlinear system used by SCPACK and the SC Toolbox does not have strong monotonicity properties and in fact is prone to local minima [9]. This is demonstrated by Fig. 7. If `dparam` of the SC Toolbox is used to solve for the polygon on the left in the figure, it arrives at the polygon on the right and terminates due to an apparent local minimum in the system. The minimum results from the fact that the two horizontal slits cannot be moved past each other without temporarily increasing the solution residual. (The fact that part of the plane is covered more than once is irrelevant.) This phenomenon was first described by Howell [9], who also points out that crowding often masks the effect. However, for the polygon of Fig. 7, the correct prevertices are pairwise separated by at least 10^{-10}, which is not fatally crowded in double precision, and `dparam` will find the solution if given a good enough starting guess. We do not know whether there exist polygons for which CRDT exhibits similar global convergence difficulties.

In practice, we do not use the simple iteration (7) for CRDT, because its convergence is too slow. Instead we use two variations of the nonlinear equation solver NESOLVE due to Behrens, which is based on a Gauss-Newton method with a Broyden update of F', as described in [6]. This is the same nonlinear system package used by the SC Toolbox. In one variant, Full CRDT, we use the standard finite-difference Jacobian to seed the Broyden update. In the other, Shortcut CRDT, we attempt to exploit the monotonicity by setting the initial $F' = I$. In Fig. 8 we show the convergence curves of the two NESOLVE variants

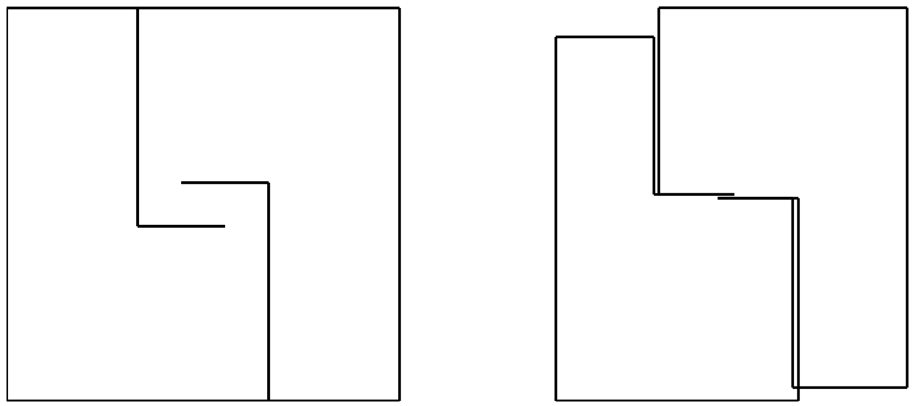

Figure 7: Lack of global convergence in the SC Toolbox. In solving for the polygon on the left, `dparam` arrives at the polygon on the right and terminates due to a nearby local minimum. The horizontal slits cannot move past each other without temporarily increasing the solution residual.

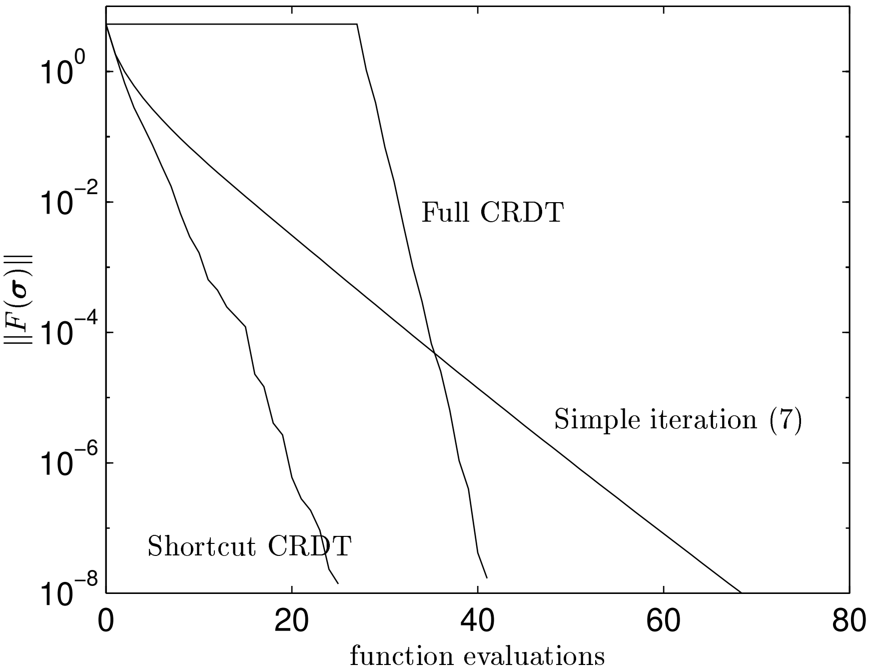

Figure 8: Convergence curves for numerical solutions of $F(\boldsymbol{\sigma}) = \mathbf{0}$. The iteration (7) converges linearly. The NESOLVE variations exhibit approximately linear convergence as well, but at much better rates.

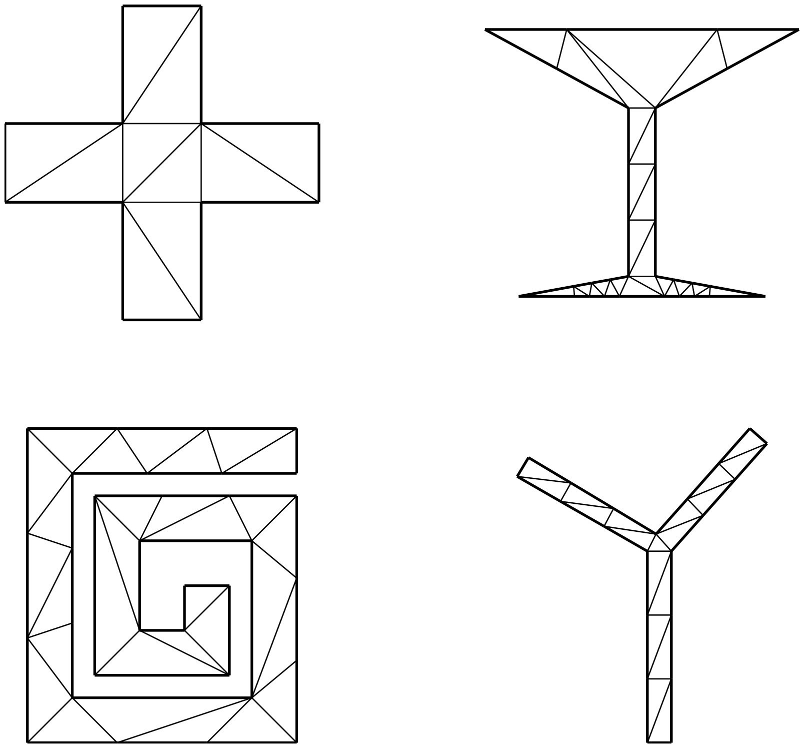

Figure 9: Four regions on which CRDT experiments were performed, after splitting and Delaunay triangulation. See text.

and the simple iteration for a typical case, the goblet shape in Fig. 9. The linear convergence of (7) is strikingly smooth. The convergence of the NESOLVE variations is more complex, but overall is approximately linear at much better rates than the simple iteration. Note that Shortcut CRDT is faster than Full CRDT because of the savings gained by not initializing F'.

We compare the performance of CRDT to the SC Toolbox functions `dparam` or `rparam` for maps from the disk or rectangle, respectively, depending on whether the target region is elongated. From each method we demand a nonlinear residual with maximum norm no larger than 10^{-8}. Since the CRDT solution is not known to produce the correct polygon in every case, the final CRDT solution is checked by applying the S-C formula and checking complex cross-ratios (not just absolute values). In every case the resulting error is within a factor of ten of the nonlinear residual. All the experiments reported here were performed on a SPARCstation-10.

Fig. 9 shows four experimental polygons with their Delaunay triangulations, after the edge splitting has been done. In Table 1 we present the number of function evaluations and total CPU time required by the nonlinear system solver for these regions.

Note that in every case the Shortcut CRDT variant indeed finishes more quickly than Full CRDT. Also observe that the individual SC Toolbox iterations are much faster than those for CRDT. This is because of the additional unknowns introduced by splitting and the

	SC Toolbox	Full CRDT	Shortcut CRDT
Cross	31/10.03	18/14.52	12/11.78
Goblet	78/14.35	42/142.3	26/112.3
Spiral	186/275.4	51/237.1	34/197.0
Y	$-/-$	35/76.37	25/70.38

Table 1: Performance of CRDT variants versus the SC Toolbox for the solution of the parameter problem for the regions in Fig. 9. The first number in each entry is the number of nonlinear function evaluations made by the nonlinear solver; the other number is the CPU time in seconds.

adaptive embedding used by CRDT to thwart crowding. However, the nonlinear systems posed by the SC Toolbox for these polygons are not as easily solved, and thus many more iterations may be required even though the systems are smaller.

The cross-shaped region (top left) is not elongated and both `dparam` and the CRDT variants converge rapidly. The SC Toolbox requires less setup effort, and the adaptive re-embedding of the CRDT algorithm is unnecessary, so `dparam` is slightly faster. The goblet region (top right) has 22 vertices added by the splitting algorithm to its original eight. These extra vertices greatly slow down the CRDT solvers, making `dparam` much faster even though it has some difficulty finding the solution. For the spiral (bottom left), relatively fewer vertices are added, and the solution is sufficiently difficult for `rparam` that CRDT is a little faster. Finally, for the Y-shaped region (bottom right), the SC Toolbox is unable to find a solution because of the doubly-elongated nature of the region. CRDT, however, finds the solution easily.

These examples demonstrate what we observe about CRDT in general. CRDT is least efficient when many extra vertices are added during splitting. This most often occurs near sharp corners and in narrow channels of the region, both of which are prominent in the goblet region. While it seems that there is no way to circumvent subdividing a channel because of crowding effects, we do not know if there is a more efficient way to produce well-conditioned quadrilaterals near sharp corners. On the other hand, CRDT handles multiply elongated regions with no difficulty, something which previously no general-purpose method for Schwarz-Christoffel mapping has been able to do.

8 Applications

In this section we describe two applications of CRDT that demonstrate its ability to handle crowding caused by elongation. While we cannot specify a recipe that will solve any conceivable problem of interest, we believe that the techniques of this section can be adapted to suit a variety of situations. A common thread in all the methods that we have explored is the careful use of local information. Any need for global information is typically obtained by taking a path through the polygon and compounding local effects.

Our first example is to compute the map from a rectangle to a generalized quadrilateral. For our purposes, a generalized quadrilateral is a polygon with four distinguished vertices, which map to the four corners of a rectangle. The fact that four, rather than three, of the

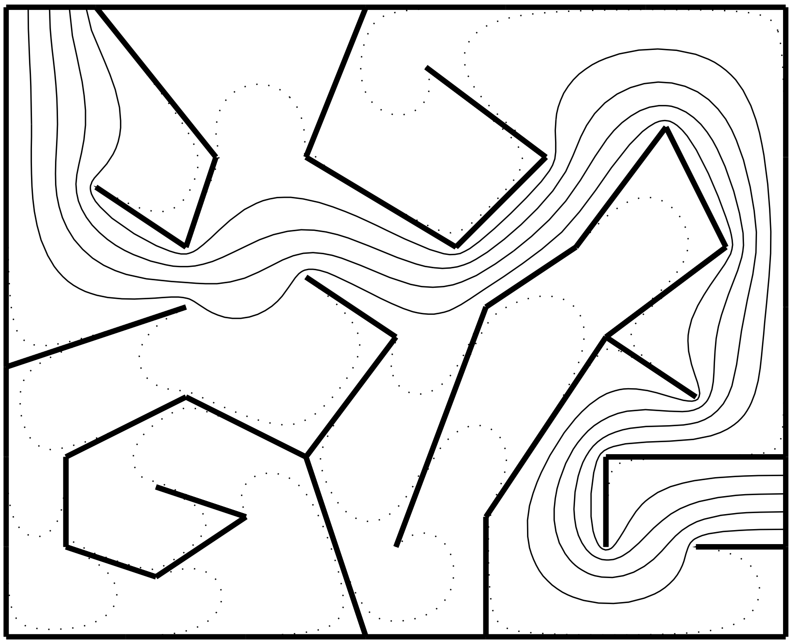

Figure 10: Rectangle map to a "maze." The exits of the maze map to the short rectangle sides, which are normalized to unit length. The maze has several branches that are very crowded on the long rectangle sides. The solid curves in the maze are images of lines with abscissae 0.2, 0.4, 0.6, and 0.8. The preimages of the dotted curves are separated from the long rectangle edges by 10^{-2}, 10^{-4}, ..., 10^{-16}. All computations were performed in double precision.

vertices are constrained is compensated by the fact that the aspect ratio of the rectangle must be a certain unknown value, known as the *conformal modulus*. Computation of this mapping is equivalent to solving Laplace's equation on the polygon with the Dirichlet values of 0 and 1 on two generalized sides separated by two generalized sides with homogeneous Neumann conditions.

For a multiply elongated region, branches other than the main channel will collapse into crowded clusters on the sides of the rectangle. Computing the map accurately in the vicinity of these clusters (into the collapsed branches) is therefore challenging.

In Fig. 10 we plot the images of straight lines in a rectangle to a certain quadrilateral. The rectangle has a width of 1 and a height of about 18.2, which is the conformal modulus. The solid curves are images of lines which are well separated from the long edges of the rectangle, so they stay out of the wrong turns, or "deadwaters." The dotted curves have preimages that are exponentially close to the long rectangle edges, and they closely follow the maze boundary into the deadwaters. The crowding of the spiral branch is comparable to machine precision, and it would pose no difficulty if it were much more crowded. As far as we know, no other conformal mapping algorithm can accurately compute these curves.

We proceed to describe how we produced Fig. 10. Suppose the prevertices of the S-C map to polygon P are known. Define a new vector of turning angles $\tilde{\beta}$, where $\tilde{\beta}_j = -1/2$

if vertex j is distinguished and $\tilde{\beta}_j = 0$ otherwise. We call the S-C map defined by using $\tilde{\beta}$ in place of β in (1) and the same prevertices the *rectified map* for the polygon, because the image of the disk under this map is clearly a rectangle. Thus the composition of the inverse of the rectified map with the original map is therefore the desired rectangle map.

We must now consider the mechanics of forward and inverse mapping using CRDT's cross-ratio representation. Recall that each quadrilateral $Q(d_i)$ (Q_i for short) has an associated embedding, $E_i = (w_1^{(i)}, \ldots, w_n^{(i)})$, which depends on the primitive variable vector $\boldsymbol{\sigma}$. Each embedding in turn induces an S-C map f_i (having $A = 0$ and $B = 1$ in (1)) such that $f_i(D)$ is an affine transformation of the target polygon P, assuming $F(\boldsymbol{\sigma}) = 0$. An important feature of CRDT in the evaluation of F is that the prevertices of Q_i are well-separated in E_i.

In order to compute the map to P from embedding E_i, we must find the appropriate affine constants A_i and B_i that transform $f_i(D)$ to P. An obvious way to compute these constants would be to compute the S-C integral for the well-separated prevertices of Q_i, and solve for A_i and B_i by matching with vertices of P. However, there are two complications. First, in the case of the rectified map, the only vertices known initially are the ends of one side. Hence the other vertices, and the associated affine constants, must be found in a certain order. Second, even when the computation is ordered correctly, some of the vertices may be arbitrarily crowded, and the linear system that defines A_i and B_i might be ill-conditioned or even numerically singular.

The key to avoiding both of these difficulties is the idea of taking a path through P. Recall that the dual graph of P, which is a tree, has $n-3$ edges, one for each diagonal (hence quadrilateral) of the Delaunay triangulation. We define another tree, the *quadrilateral tree*, with a node for each quadrilateral. Quadrilaterals are adjacent if their edges in the dual graph share an endpoint, or, in terms of P, if they share three vertices.

Given a reference edge e of the target polygon P, we choose a quadrilateral Q_1 that has e as an edge. As in the CRDT iteration, we compute f_1 to find a scaled and translated version of Q_1. The affine constants for this embedding, A_1 and B_1, can be found using e as reference. Now we root the quadrilateral tree at Q_1 and visit the remaining quadrilaterals in any order which ensures parents are visited before children. For each quadrilateral Q_i, we apply f_i to the prevertices of Q_i. Since Q_i shares three vertices with its parent Q_j, we have sufficient reference for computing an affine transformation from the image plane of f_i to the image plane of f_j. But then by composition with the known affine constants A_j and B_j, we can find A_i and B_i. Even though the resulting scaling constant B_i may be exponentially small, it is found to high relative accuracy by multiplication. The important point is that for each embedding, reference is made to the raw, unscaled image of a neighboring embedding, because the transformation between the two is well-conditioned. If the scaled image were used instead, there could be a total loss of accuracy due to ill-conditioning in the presence of crowding.

Observe that once all of the A_i and B_i have been found, they can be used to compute the images of the prevertices on the rectangle. Even though some may be crowded, they will all be found with high accuracy relative to the overall size of the rectangle. Thus the conformal modulus can be found accurately as well.

To produce Fig. 10, we ran Shortcut CRDT on the polygon, which after splitting had 87 vertices. This found the solution vector $\boldsymbol{\sigma}$ to a residual tolerance of 10^{-8} after just 27 function evaluations. Then we used the procedure above to find the affine constants for the standard and rectified maps, and the rectangle prevertices. Note that the triangulation defined on P is still a triangulation of the rectangle, albeit with many degenerate triangles which lie on the long rectangle edges. Suppose the rectangle line we wish to map is separated from the rectangle wall by a distance h. In the image plane of f_i, that separation scales to $h/|B_i|$. If $h/|B_i| \ll 1$, the image of the line will be very close to the boundary of P in the vicinity of Q_i. If Q_i is degenerate, its raw image is a line segment, and it is also possible that $h/|B_i| \gg 1$. In fact, this means that Q_i is in a deadwater region and the image of the line will be far from Q_i. In sum, only for those embeddings in which h is comparable to $|B_i|$ do we need to track the image of the line. This is accomplished by choosing points on the portion of the line local to Q_i, inverting the rectified S-C map using Trefethen's technique [17], and computing the standard forward map from the same embedding.

We believe that this technique can be refined to perform grid generation for any polygon. The main obstacle is in defining the rectified map in general, when a rectangle is not a natural choice.

Our second example of applying CRDT involves the solution to a certain boundary value problem related to harmonic measure. Our BVP is Laplace's equation $\triangle u = 0$ with Dirichlet boundary conditions on a polygonal domain P. There is one edge s of the polygon, which we call the "forced edge," with boundary condition 1, and the solution on the rest of the boundary is zero. We want to find u only at a particular point x in the interior of P. Let the endpoints of s be denoted z_p and z_q, in counterclockwise order.

Mathematically the problem is easily solved. Let $f_0 : D \to P$ be a Schwarz-Christoffel map that has $f_0(0) = x$. The solution u of the original problem maps to a Laplace solution $\hat{u}$ on the unit disk via the formula $\hat{u}(w) = u(f_0(z))$. This Laplace problem has boundary condition one on the segment $f_0^{-1}(z_p)$ to $f_0^{-1}(z_q)$ and zero elsewhere. Because the solution to Laplace's equation at the center of a disk is identically the average of the boundary Dirichlet values, we have that $u(x)$ is precisely the arc-length distance θ measured in radians between $f_0^{-1}(z_p)$ and $f_0^{-1}(z_q)$, scaled by $1/(2\pi)$.

An obvious algorithm is to compute the prevertices of the disk map and read off the angular distance between prevertices p and q. In fact, it suffices to compute $h = |f_0^{-1}(z_p) - f_0^{-1}(z_q)|$. For, given the Euclidean distance h between two points on the unit circle, we can get the angular distance via

$$\theta = 2\sin^{-1}(h/2). \tag{8}$$

However, this direct method will fail when the forced edge s is separated from x by a long, thin region. In this case, the value $u(x)$ will be extremely close to zero, and cancellation error will dominate h. Yet we insist on computing $u(x)$ accurately in a *relative* sense. Note that it suffices to compute h with high relative accuracy, since (8) can be evaluated accurately when h is very small (indeed, (8) behaves like $\theta \approx h$ for small h).

If the domain P is singly elongated, a solution to the BVP is possible using S-C maps from a rectangle. A rectangle is transformed to the upper half-plane by the elliptic function $\mathrm{sn}\,(z|m)$, where m is obtained in the S-C solution. We can further construct a fractional linear transformation that maps the half-plane image of x to the center of the unit disk, and then accurately measure h. Of course, such a solution is not available when P is multiply elongated and the rectangle map fails.

We assume throughout that x is not too close to any edge of the polygon. If x is close to an edge, this creates a different problem with relative accuracy that is not addressed by the techniques in this section, though it could be addressed by other means.

We start by running CRDT on P and computing the affine constants as above. Let Q_1 be a quadrilateral having s as an edge. (We assume for now that s is an unsplit edge of the original polygon and deal with the case of split s below.) Let Q_m be a quadrilateral in P that contains the point x. Note that we can accurately invert the S-C map using embedding E_m to find $\xi = f_m^{-1}(x)$. Furthermore, ξ will not be close to the boundary of the disk.

Let $Q_1, Q_2, \ldots, Q_m$ be a path in the quadrilateral tree. In the embedding E_2, we can compute the distance $|w_p - w_q|$ accurately, because either w_p or w_q is a vertex of Q_2. Thus we can compute the cross-ratio ρ_2 of w_p, w_q, and two more prevertices of Q_2, including the prevertex that is not shared with Q_1. Now consider the embedding E_3. In this embedding, w_p and w_q may be crowded. However, we still know the cross-ratio ρ_2 of w_p, w_q, and two prevertices of Q_3. Hence of the four lengths that are factors in the formula for $|\rho_2|$, only one may be small, so we can compute it accurately using ρ_2 and the other lengths. By the same token, we can accurately compute ρ_3, the cross-ratio of w_p, w_q, and two other prevertices of Q_3, including the one not shared by Q_2. As in the previous example, we are computing a very small quantity (given by (8)) by repeated multiplication of small numbers rather than by subtracting two nearby complex numbers.

We can continue this procedure through ρ_m, a cross-ratio involving w_p, w_q, and two prevertices of Q_m. Finally, we find the image of these points under the transformation of the type (5) that moves ξ to the origin. Since ξ is not close to the boundary of the disk, the well-separated points will remain uncrowded. The invariance of ρ_m under fractional linear transformation allows us to recover $|w_p - w_q|$, which is now the desired h.

In the situation where segment s is split by CRDT, we simply add up the contributions to $u(x)$ separately from each subsegment using the preceding algorithm. A Laplace solution is linear in the boundary data, and there cannot be any cancellation at this step because each contribution is positive.

We use the method described above to solve the BVP on the T-shaped region of Fig. 11. The region is parameterized by L, the length of the long arm of the T, and the forced edge s lies at the end of the shortest arm. The point x lies along the centerline of the long arm and at a distance 0.1 from the end.

In Fig. 12 we plot the solution $u(x)$ for L up to 20. After a transient phase (shown in the inset), the behavior very quickly approaches $k \cdot e^{-\pi L}$ for some constant k. This is because as L grows, the configuration at the top of the T becomes irrelevant, and the solution is like that for a rectangle with forced edge at the top. The rectangle aspect ratio L is asymptotically related to the sn parameter m by

$$L = \frac{1}{\pi} \ln\left(\frac{4}{\sqrt{m}}\right) + O(m).$$

By further transformation to the unit disk, we find that to leading order, $u(x) \approx \sqrt{m} \approx e^{-\pi L}$.

One issue that arises immediately is verification of the computed solutions. We do this by noting that the BVP is essentially singly elongated, even when P is not; the elongations not containing x nor the forced edge are largely irrelevant to the solution. For the T-shaped

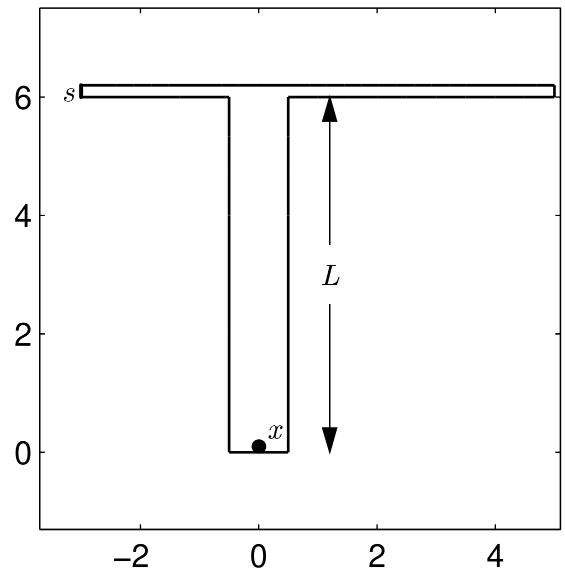

Figure 11: The BVP of Section 8 is solved on this T-shaped region. The length L is allowed to vary, and the forced edge is denoted s. The point x at which the solution is sought is at $0.1i$. The dimensions of the horizontal bar are 8×0.2.

polygon in Fig. 11, for example, we can shrink the right branch of the crossbar, decreasing the BVP solution by an exponentially small amount. Once this branch is sufficiently small, we can compute a rectangle map and solve the approximate BVP as described above. We have done this for several values of L and verified the CRDT solutions to at least 10 digits. We have also confirmed 12 digits for BVPs on other regions. The advantage of the CRDT method is that the solution is found for the original region directly without introducting any approximations. Furthermore, the CRDT method is fully automatic. In contrast, the continuation routines in the SC Toolbox can require a fair amount of user intervention: the user must select the branches of the polygon to shrink and the amount of shrinkage. In several experiments, we were always able to accurately solve this class of exponentially small harmonic-measure conformal mapping problems using the continuation routines. But convergence was attained only with prohibitively laborious hand-tuning that would not be feasible for novice users of the Toolbox.

The two applications presented in this section share the idea of composing a chain of operations via a path in the quadrilateral tree. By following such a path, we can take advantage of the overlap between neighboring quadrilaterals to ensure that each link in the chain is fairly well-conditioned, even though the global results of following the chain may vary over many orders of magnitude. We believe that this technique is central to the application of CRDT to potential theory problems.

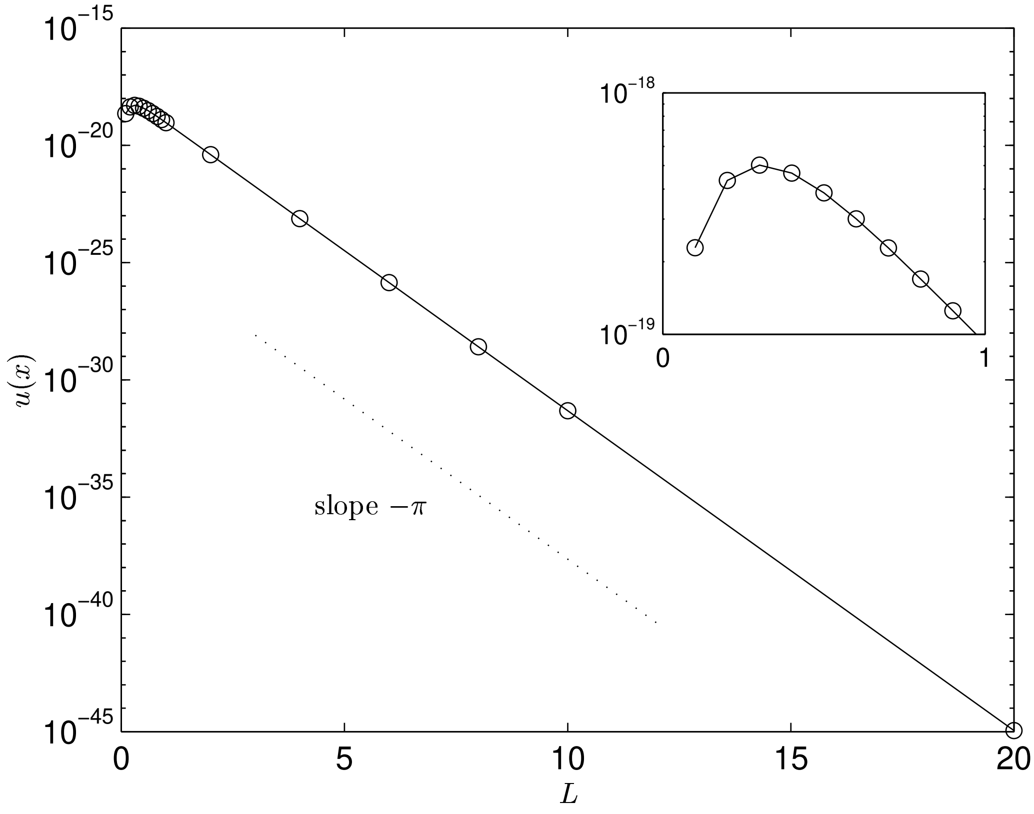

Figure 12: Solution of the BVP for the region in Fig. 11. The inset shows the solution at points for $L < 1$. As L grows, the curve quickly approaches $k \cdot e^{-\pi L}$, corresponding to the leading-order behavior of the solution for a rectangle.

9 Conclusions

We have introduced CRDT, a new algorithm for finding the Schwarz-Christoffel prevertices on the unit circle for arbitrary bounded polygonal regions. The classical crowding problem is avoided through conformally equivalent re-embeddings of the prevertices so that the numerical mapping is always locally accurate. In addition, the nonlinear system chosen for numerical solution apparently has a monotonicity property that makes it easier to solve numerically than previously posed systems for the S-C parameter problem, although we cannot verify this is so for all polygons. While we conjecture that the CRDT algorithm will always converge to the correct solution, we have been unable to prove this is so.

The polygon is first split so that long, narrow regions can be represented piecewise by well-conditioned triangles. A Delaunay triangulation of the resulting n-vertex polygon is computed and used to define $n-3$ quadrilaterals, whose diagonals appear as internal sides in the triangulation. The primitive variables of the nonlinear system are logarithms of the cross-ratios of the prevertices of those $n-3$ quadrilaterals. These cross-ratios define an infinite set of conformally equivalent configurations of the prevertices, each of which produces an S-C map to the same image polygon. The imposed constraints are on the magnitudes of the cross-ratios of the quadrilaterals in the image polygon.

The CRDT algorithm generally compares favorably with the SC Toolbox for MATLAB in numerical experiments. The principal exceptions are those regions which require a great many extra vertices to be added in the splitting phase of the algorithm. We do not know if there is a more effective splitting procedure. On the other hand, CRDT has no problem finding the prevertices for arbitrarily elongated polygons, something which no previous algorithm can claim.

We demonstrate the use of CRDT in applications. Fig. 10 shows the rectangle map to a multiply elongated polygon for which we believe no other algorithm would work. We also illustrate how to use CRDT to solve a particular elliptic boundary value problem. In each case, the key is to follow a path in the quadrilateral tree of the polygon, dual of the triangulation, the distance separating prevertices can be computed without cancellation, even when the distance is extraordinarily small. This technique can be used to solve a certain Dirichlet problem or compute a conformal modulus to an accuracy not achievable by other methods on most regions.

Besides the unresolved matters already introduced in this work, there are open questions about possible extensions of CRDT to work with polygons with infinite vertices, circular-arc polygons [10], or multiply-connected regions [5, 13]. Future work will include the application of CRDT to grid generation and the incorporation of CRDT into the SC Toolbox.

10 Acknowledgments

We would like to thank Nick Trefethen at Cornell, Marshall Bern and John Gilbert at Xerox, and David Eppstein at the University of California at Irvine for their suggestions and comments.

References

[1] M. Bern and D. Eppstein, *Mesh generation and optimal triangulation*, in Computing in Euclidean Geometry, D. Z. Du and F. Hwang, eds., World Scientific, 1992. Also Xerox Palo Alto Research Center Technical report CSL-92-1.

[2] M. Bern, D. Eppstein, and J. Gilbert, *Provably good mesh generation*, in Proc. 31st IEEE Symp. Foundations of Computer Science, 1990, pp. 231–241.

[3] L. P. Chew, *Constrained Delaunay triangulations*, Algorithmica, 4 (1989), pp. 97–108.

[4] L. P. Chew, *Guaranteed-quality triangular meshes*, Tech. Rep. 89–983, Department of Computer Science, Cornell University, 1989.

[5] H. D. Däppen, *Die Schwarz-Christoffel-Abbildung für zweifach zusammenhängende Gebiete mit Anwendungen*, PhD thesis, ETH Zürich, 1988.

[6] J. E. Dennis and R. B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Prentice-Hall, 1983.

[7] T. A. Driscoll, *A* Matlab *toolbox for Schwarz-Christoffel mapping*, ACM Trans. Math. Soft., (to appear, 1996).

[8] P. Henrici, *Applied and Computational Complex Analysis*, vol. 1, Wiley, 1974.

[9] L. H. Howell, *Computation of Conformal Maps by Modified Schwarz-Christoffel Transformations*, PhD thesis, MIT, 1990.

[10] ——, *Numerical conformal mapping of circular arc polygons*, J. Comput. Appl. Math., 46 (1993), pp. 7–28.

[11] ——, *Schwarz-Christoffel methods for multiply-elongated regions*, in Proc. of the 14th IMACS World Congress on Computation and Applied Mathematics, 1994.

[12] L. H. Howell and L. N. Trefethen, *A modified Schwarz-Christoffel transformation for elongated regions*, SIAM J. Sci. Stat. Comput., 11 (1990), pp. 928–949.

[13] C. Hu, *User's guide to DSCPACK*, Nat. Inst. Aviation Res. 95-1, Wichita State Univ., 1995.

[14] R. Menikoff and C. Zemach, *Methods for numerical conformal mapping*, J. Comp. Phys., 36 (1980), pp. 366–410.

[15] S. A. Mitchell and S. A. Vavasis, *Quality mesh generation in three dimensions*, in Proc. 8th ACM Symp. on Computational Geometry, 1992, pp. 212–221.

[16] Z. Nehari, *Conformal Mapping*, Dover, 1952.

[17] L. N. Trefethen, *Numerical computation of the Schwarz-Christoffel transformation*, SIAM J. Sci. Stat. Comput., 1 (1980), pp. 82–102.

[18] ——, *Numerical construction of conformal maps.* Appendix to *Fundamentals of Complex Analysis for Mathematics, Science, and Engineering*, by E. B. Saff and A. D. Snider, Prentice Hall, 1993.