Computational Methods and Function Theory
Volume 3 (2003), No. 1, 325–347

Curvature Flow in Conformal Mapping

Charles R. Collins, Tobin A. Driscoll and Kenneth Stephenson

(Communicated by Lloyd N. Trefethen)

In memory of Dieter Gaier

Abstract. We use a simple example to introduce a notion of *curvature flow* in the conformal mapping of polyhedral surfaces. The inquiry was motivated by experiments with *discrete* conformal maps in the sense of circle packing. We describe the classical theory behind these flows and demonstrate how to modify the Schwarz-Christoffel method to obtain classical numerical confirmation. We close with some additional examples.

Keywords. Circle packing, conformal mapping, curvature, Schwarz-Christoffel.

2000 MSC. Primary: 52C26, 30C30, 53A30; Secondary: 65D15, 37E35, 53A05.

1. Introduction

A *polyhedral surface* is a piecewise affine 2D surface (defined either concretely or abstractly) whose faces are euclidean polygons. Polyhedral surfaces are becoming increasingly important in various parts of mathematics, science, and engineering; some examples are discussed in the concluding section of the paper. Polyhedral surfaces are known to carry natural conformal structures, yet it is only with the advent of circle packing methods that explicit computations with these structures have became readily accessible.

We will demonstrate conformal flattening using circle packing methods. As often happens, new machinery brings new insights and questions, and that is the case here. The cone points of a polyhedral surface represent point concentrations of gaussian curvature and that curvature must be redistributed or cancelled when the surface is mapped to the plane. Curvature redistribution is an integral feature in the numerics of circle packing, but in studying the packing algorithm we

Received June 30, 2003.
The authors gratefully acknowledge support of the National Science Foundation grants DMS-0101324, DMS-0104229.

ISSN 1617-9447/$ 2.50

stumbled upon remarkably persistent intrinsic curvature flows which demanded explanation.

This paper is intended as an introduction to these phenomena, so we have chosen to concentrate on a single example of a polyhedral surface $\mathfrak{S}$ and the conformal flattening of $\mathfrak{S}$ to a plane rectangle $\mathfrak{R}$. Circle packing methods permit us to formulate a continuous conformal deformation between $\mathfrak{S}$ and $\mathfrak{R}$ and to interpret that deformation in terms of curvature flow. We then establish a classical formulation based on modifications of the Schwarz-Christoffel method. At the conclusion we show some additional examples and discuss open questions and applications.

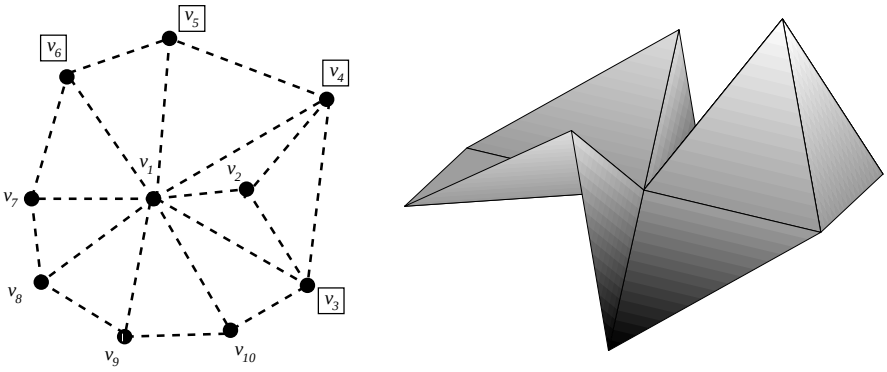

FIGURE 1. Combinatorics and embedded surface.

We thank David Minda, John Sullivan, and Kevin Pilgrim for interesting conversations on the material presented here.

2. The Equilateral surface $\mathfrak{S}$

The example surface $\mathfrak{S}$ we use throughout the paper is *equilateral*, that is, its faces are equilateral triangles. Its combinatorics are encoded in the simplicial complex $\mathcal{T}$ shown schematically on the left in Figure 1. $\mathfrak{S}$ is simply connected and has a set $\mathcal{V}$ of ten labelled points: v_1, v_2 are interior vertices, the remainder are boundary vertices, with v_3, v_4, v_5, v_6 designated as "corners".

The piecewise euclidean structure on $\mathfrak{S}$ is straightforward. Each face is an equilateral triangle, each edge a line segment. Each vertex v has a *link* of k faces which contain it, closed when v is an interior vertex. The point v is called a *cone point*, and since each face in its link has angle $\pi/3$ at v, the *cone angle* at v

is $k\pi/3$. For later use we record the cone angles for the points of $\mathcal{V}$:

$$\begin{aligned} \text{Interior: } & \Theta(v_1) = 3\pi,\ \Theta(v_2) = \pi, \\ \text{Corner: } & \Theta(v_3) = \Theta(v_4) = \pi,\ \ \Theta(v_5) = \Theta(v_6) = 2\pi/3, \\ \text{Edge: } & \Theta(v_7) = \Theta(v_8) = \Theta(v_9) = \Theta(v_{10}) = 2\pi/3. \end{aligned} \tag{1}$$

In a completely routine way the euclidean metrics on the faces and edges of $\mathfrak{S}$ determine a path metric on $\mathfrak{S}$ itself, the distance between any two points being the infimum of the lengths of locally rectifiable paths connecting them. We normalize our surface to have area 1, so the side lengths in $\mathfrak{S}$ are set at $\mathfrak{c} = 2\sqrt{3}/15$. Note that $\mathfrak{S}$ is not "flat" in this metric; the interior cone angles of 3π and π prevent $\mathfrak{S}$ from being locally isometrically immersed in $\mathbb{R}^2 = \mathbb{C}$. (In fact, our discussion applies to abstract polyhedral surfaces, which might even fail to be immersible in $\mathbb{R}^3$.)

A perhaps less well known companion to the piecewise euclidean structure on $\mathfrak{S}$ is the associated *conformal structure*. We may briefly describe its atlas $\{(U, \psi)\}$; see, e.g., Beardon [1] for details. Given any pair of faces sharing an edge, let U be the interior of their union. Each of the faces is isometrically identified with an equilateral triangle, and by sliding these triangles together in $\mathbb{C}$ so they intersect in this common edge one defines an isometry $\psi\colon U \to \mathbb{C}$, hence a chart (U, ψ). Using all such pairs provides charts for $\mathfrak{S}\backslash\mathcal{V}$. The charts at interior cone points require local *power* maps. Say $v \in \mathcal{V}$ is interior with cone angle $\Theta(v)$. The link of triangles containing v can be identified with a fan of equilateral triangles at the origin in $\mathbb{C}$ whose angles at the origin sum to $\Theta(v)$. Define the power map $\psi(z) = z^{2\pi/\Theta(v)}$; applied to a small disc U centered at the origin, ψ will close up the fan and provide a chart (U, ψ) compatible with those for the faces (see Figure 10). With the atlas thus defined, $\mathfrak{S}$ becomes a *Riemann surface*. (Similar considerations at boundary cone points provide half-neighborhoods, so $\mathfrak{S}$ is in fact a *bordered* Riemann surface.)

As a Riemann surface, $\mathfrak{S}$ serves as a domain for various conformal maps and analytic functions. Among these are maps to plane regions. We have chosen to treat $\mathfrak{S}$ as a *conformal rectangle* (recall the previously designated corners). It is well known classically that there exist a rectangle $\mathfrak{R}$ in $\mathbb{C}$ and a univalent conformal mapping $\phi\colon \mathfrak{S} \to \mathfrak{R}$ which maps the corners of $\mathfrak{S}$ to the corners of $\mathfrak{R}$; the rectangle is unique up to similarity and the map is then unique. Figure 2 shows the flattened image of $\mathfrak{S}$.

The reader will appreciate that there are various questions one can now address. What is the aspect ratio of $\mathfrak{R}$? (the *conformal modulus* of $\mathfrak{S}$?) What are the shapes of the images of the faces of $\mathfrak{S}$ under ϕ? Our interest, however, is in the transition from $\mathfrak{S}$ to $\mathfrak{R}$. Let us set the terminology.

The rectangle $\mathfrak{R}$ of Figure 2 is, like $\mathfrak{S}$, a piecewise euclidean surface. Its cone angles will be considered our "target" cone angles, and we record them here for

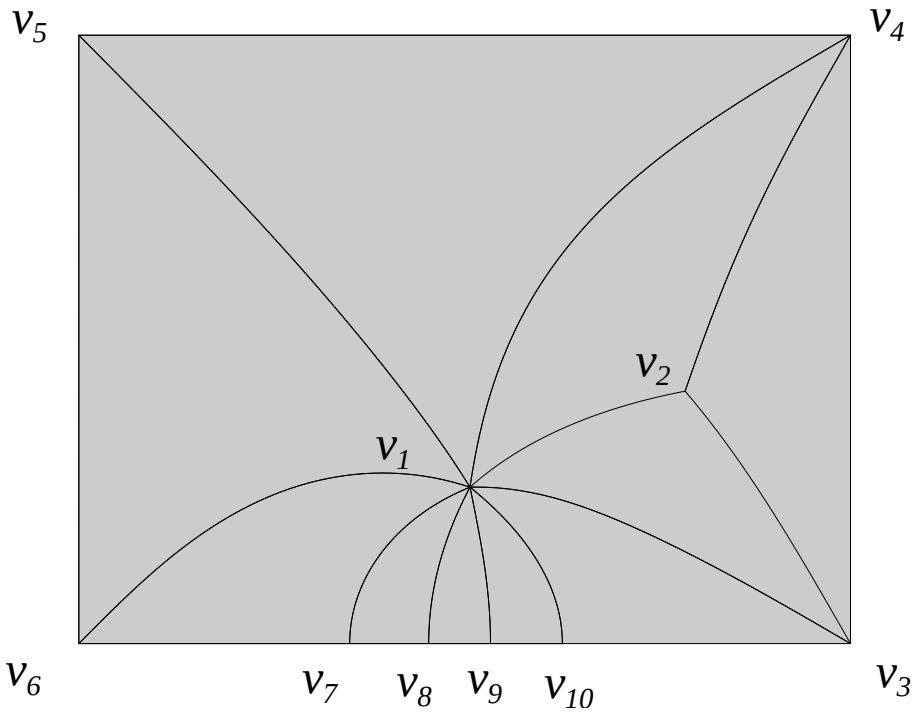

FIGURE 2. The conformally flattened $\mathfrak{S}$.

reference:

$$\begin{aligned} \text{Interior: } & A(v_1) = A(v_2) = 2\pi, \\ \text{Corner: } & A(v_3) = A(v_4) = A(v_5) = A(v_6) = \pi/2, \\ \text{Edge: } & A(v_7) = A(v_8) = A(v_9) = A(v_{10}) = \pi. \end{aligned} \tag{2}$$

In the piecewise euclidean setting, at an interior cone point v with cone angle θ, the difference $2\pi - \theta$ is commonly referred to as the *curvature* at v and acts as concentrated gaussian curvature. The terminology may be extended to boundary cone points v, where the difference $\pi - \theta$ is the *curvature* and acts as concentrated boundary curvature. In particular, the Gauss-Bonnet Theorem applies and the curvatures of our simply connected surfaces must sum to 2π. (The reader can check from (1) and (2) that, with 2 interior and 8 boundary vertices, the total curvature of $\mathfrak{S}$ is $12\pi - \sum_{j=1}^{10} \Theta(v_j) = 2\pi$ and the total curvature of $\mathfrak{R}$ is $12\pi - \sum_{j=1}^{10} A(v_j) = 2\pi$.)

Our intention is to parameterize the process of flattening $\mathfrak{S}$ to $\mathfrak{R}$. We keep all the curvature concentrated at $\mathcal{V}$, and as we progress through the intervening surfaces we will watch as that curvature redistributes itself. In the end, the only non-zero curvature remaining is the $\pi/2$ residing at each of the four corners.

3. Circle Packing

There is well established circle packing machinery for approximating conformal mappings such as $\phi\colon \mathfrak{S} \to \mathfrak{R}$, and if ϕ were our only goal, we would refer the reader to [8, 15] and be done. For the parameterization, however, we will have to tinker with that machinery and we have chosen the ten triangle example so that the reader can follow the geometry at an intuitive level. For those who wish further details on circle packing we suggest the surveys [15, 16] and the papers referenced therein.

A *circle packing* P is a configuration of circles with a specified pattern of tangencies. The pattern is encoded as a (abstract, oriented, simplicial) complex K; the configuration P has a circle c_v associated with each vertex $v \in K$; whenever $\langle u, v\rangle$ is an edge of K, then circles c_u, c_v are tangent; and whenever $\langle u, v, w\rangle$ is an oriented face of K, then c_u, c_v, c_w form an oriented triple of mutually tangent circles. The radii of the circles are given in a *label* for K, $R = \{R(v)\colon v \in K\}$.

Our approach, however, will start with the abstract complex K. We impose a geometry on K by choosing a label R of positive numbers. Given any oriented face $t = \langle u, v, w\rangle \in K$, an oriented triple of mutually tangent circles with radii $R(u), R(v), R(w)$ determine a euclidean triangle. These triangles can be attached to one another along shared edges and vertices to define an abstract polyhedral surface which is simplicially equivalent to K; we denote this metric surface by $K(R)$. The circles with radii from R form an *in situ* circle packing P in $K(R)$, and we may refer to $K(R)$ as the *carrier* of P, $\mathrm{carr}(P)$. The geometry of $K(R)$ shows up in its various cone angles (the term *angle sum* is used in the circle packing literature). The cone angle at v will be denoted by $\theta_R(v)$ and can be computed from the label R using the law of cosines.

In circle packing one can prescribe target cone angles, subject only to the Gauss-Bonnet Theorem and mild degree restrictions, and guarantee the existence of a label R, unique up to scaling, so that $K(R)$ has precisely those cone angles. Moreover, numerical methods are available for computing approximations to R — the process is called *repacking*. (All computations, layouts, and illustrations for circle packings in this paper were carried out using the software package `CirclePack`, available on Stephenson's web site.)

Let us now review the procedure for conformal flattening of $\mathfrak{S}$ to a rectangle. Write $K^{(0)}$ for the complex $\mathcal{T}$ triangulating $\mathfrak{S}$. As observed above, the labelled complex $K^{(0)}(R)$ with label $R \equiv \mathfrak{c}/2$ is isometrically isomorphic to $\mathfrak{S}$. The associated circle packing $Q^{(0)}$ is thus an *in situ* circle packing in $\mathfrak{S}$; you can visualize it by marking each equilateral face as in Figure 4(a). Repacking to meet the cone angle targets of (2) results in a label $R^{(0)}$ for $K^{(0)}$ which is flat at every vertex; laying out circles of these radii gives the univalent circle packing $P^{(0)}$ in Figure 3 with carrier a rectangle $\mathfrak{R}^{(0)}$.

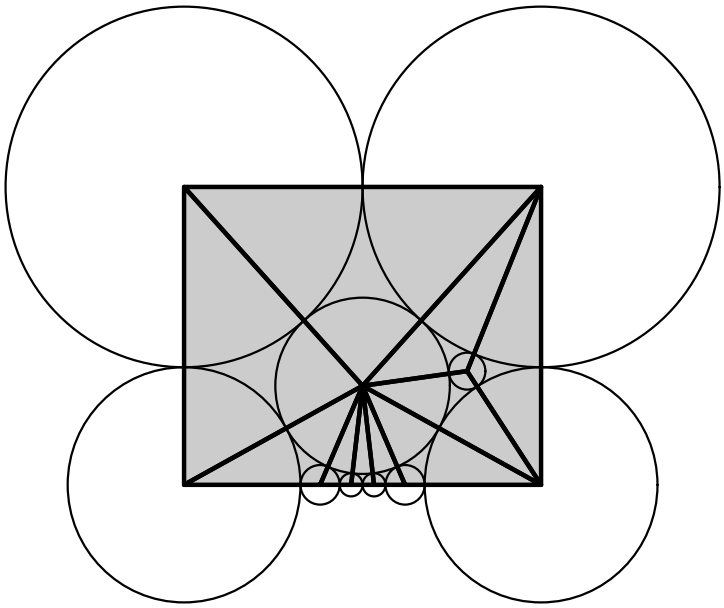

FIGURE 3. The coarse flat packing for $\mathcal{T}$.

A simplicial mapping can be defined from $\mathrm{carr}(Q^{(0)})$ (i.e., $\mathfrak{S}$) to $\mathrm{carr}(P^{(0)}) = \mathfrak{R}^{(0)}$; namely, identify the center of each circle of $Q^{(0)}$ with the center of the corresponding circle of $P^{(0)}$ and extend affinely (using barycentric coordinates) to the edges and faces. The resulting function $\phi^{(0)}$ from $\mathfrak{S}$ to $\mathbb{C}$ is called a *discrete conformal mapping*; with a slight abuse of notation we will write $\phi^{(0)}\colon Q^{(0)} \to P^{(0)}$.

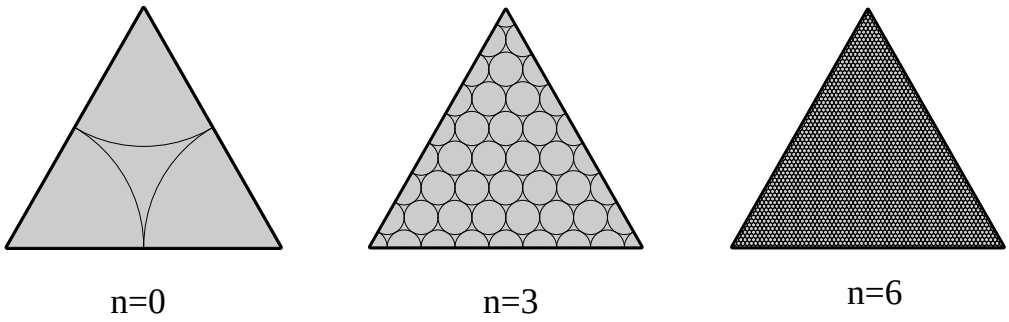

FIGURE 4. *In situ* faces.

It is clear that the "coarse" circle packing $Q^{(0)}$ cannot carry all the conformal information in $\mathfrak{S}$, so we deploy a hexagonal refinement scheme. An equilateral triangle is hex refined by connecting the midpoints of its edges to break it into

four smaller equilateral triangles. Figure 4 illustrates the circles for 3 and 6 stages of successive hex refinement. A simultaneous level n hex refinement of all the faces of an equilateral surface $\mathfrak{S}$ yields a complex $K^{(n)}$ and an associated *in situ* packing $Q^{(n)}$ for $\mathfrak{S}$ with constant label $\mathfrak{c}/2^{n+1}$. To compute the flat packing, we again specify the target cone angles in (2) for $\mathcal{V}$. In addition we must specify target cone angles at the vertices added during refinement; we define $A(v) = 2\pi$ at all new interior vertices, and $A(v) = \pi$ at all new boundary vertices. Repacking yields a label $R^{(n)}$ whose cone angles are precisely the specified targets, i.e., $\theta_{R^{(n)}}(v) = A(v), v \in K^{(n)}$. Laying out the associated circles gives a univalent circle packing $P^{(n)}$ which lies in the plane. We thus arrive at the level n discrete conformal mapping $\phi^{(n)}\colon Q^{(n)} \to P^{(n)}$. By way of normalization we may assume that $\mathfrak{R}^{(n)} = \mathrm{carr}(P^{(n)})$ has unit area and that the corners associated with v_6 and v_3 have been placed at $z = 0$ and on the positive x-axis, respectively.

Theorem 3.1. *Let $\phi^{(n)}\colon Q^{(n)} \to P^{(n)}$ be the discrete conformal mappings described above, $n = 0, 1, \ldots$. Then the image rectangles $\mathfrak{R}^{(n)} = \mathrm{carr}(P^{(n)})$ converge to a rectangle $\mathfrak{R}$ in the plane and the discrete conformal maps $\phi^{(n)}$ converge uniformly on $\mathfrak{S}$ to a conformal bijection $\phi\colon \mathfrak{S} \to \mathfrak{R}$ which maps the corner vertices of $\mathfrak{S}$ to the corners of $\mathfrak{R}$.*

This theorem follows in a straightforward way from methods in [2, 3]. The packing $P^{(6)}$ is shown in Figure 5, with the circles defining the ten faces drawn in. This is clearly the source of the image in Figure 2; our theorem justifies the claim that it approximates a conformally correct flattening of $\mathfrak{S}$.

4. Circle Packing Parameterization

For a fixed refined complex $K^{(n)}$ we can move from $Q^{(n)}$ to $P^{(n)}$ in incremental stages by manipulating cone angles. (We suppress the index n in the notations.)

Recall that the *target* cone angles (the cone angles of P) are denoted by $A(v)$, and let the *initial* cone angles (those of Q) be denoted by $\Theta(v)$. $\Theta(v)$ and $A(v)$ are identical at all vertices save those of $\mathcal{V}$; these values are given in (1) and (2), respectively. We set intermediate target cone angles $a(t, \cdot)$ defined by

$$a(t, v) = tA(v) + (1 - t)\Theta(v), \qquad v \in K,\ t \in [0, 1]. \tag{3}$$

There exists a unique label $R_t = \{R(t, v)\}$ so that $\theta_{R_t}(v) = a(t, v)$ for every vertex $v \in K$ and so that $\mathrm{area}(K(R_t)) = 1$; write P_t for the associated circle packing. In particular, $P_0 = Q$ and $P_1 = P$.

It is not difficult to show that $R(t, v)$ is a continuous function of t for each v. Figure 6 illustrates this with the face $\Delta = \langle v_1, v_6, v_7 \rangle$ at refinement level $n = 4$ for five values of t. On the left, $t = 0$, is the *in situ* packing in $\mathfrak{S}$. As t grows, the labels R_t on Δ change until when $t = 1$ we get the shape of Δ in the final flat packing P_1.

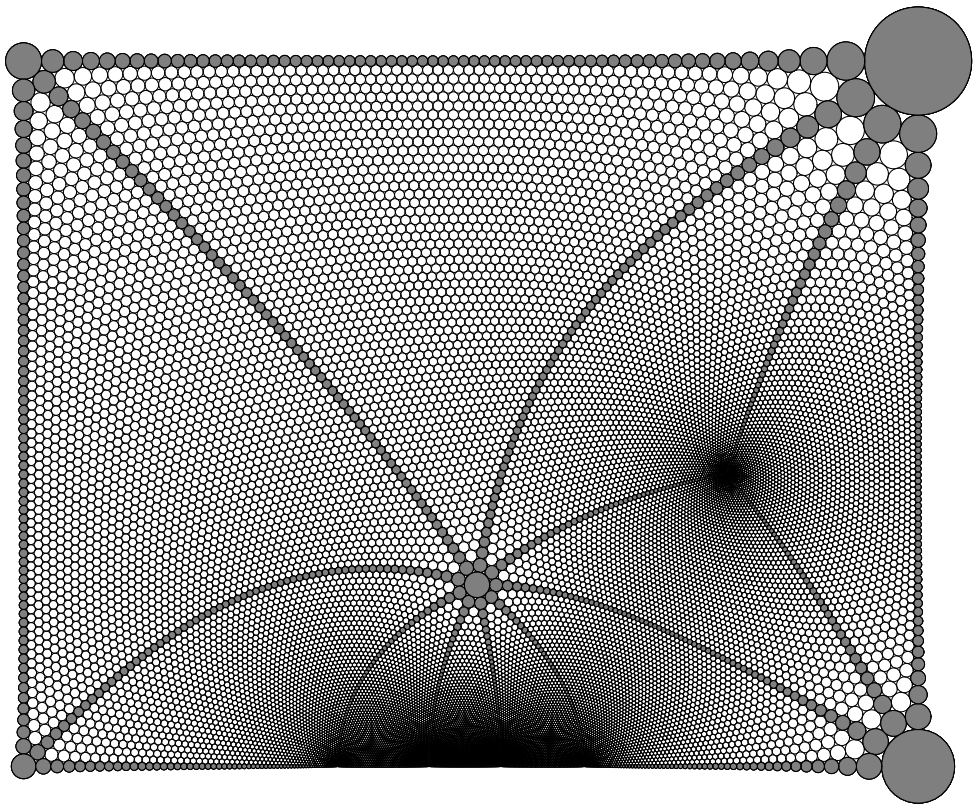

FIGURE 5. Level $n = 6$ flat packing for $\mathcal{T}$.

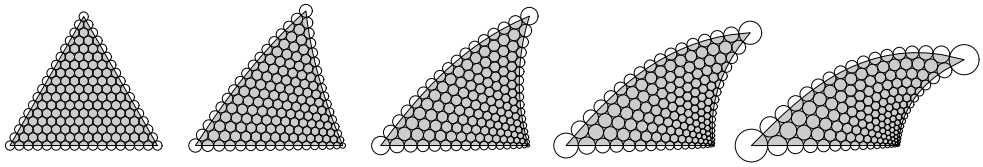

FIGURE 6. Face Δ at parameter values $t = 0.0, 0.25, 0.5, 0.75, 1.0$.

There are many ways one might formalize this parameterization; we choose to represent it in terms of metrics induced on $\mathfrak{S}$. Fix $t \in [0, 1]$. The simplicial map from $K(R_0)$ to $K(R_t)$ induces a piecewise affine metric d_t on $\mathfrak{S}$ which we will term a *discrete conformal metric*. This term requires some justification. Consider an individual face Δ of $\mathfrak{S}$. The corresponding portions of P_0 and of P_t are flat so they can be laid out in the plane. Define the simplicial map $h_t \colon P_0|_\Delta \to P_t|_\Delta$ (for example, the map between the face on the left and one of the other four

faces in Figure 6). This is a *discrete analytic function* in the sense of [16] and the associated *ratio function*, $h_t^{\#}(v) = R(t,v)/R(0,v)$, is a discrete analogue of the modulus of its derivative. For us $h^{\#}$ serves as a sort of "discrete density" for the induced metric d_t: the metric d_t on $\langle u, v, w\rangle$ is determined by the values $h_t^{\#}(u), h_t^{\#}(v), h_t^{\#}(w)$.

Also note that the maps h_t for various faces are not independent, but in fact respect the equilateral structure of $\mathfrak{S}$. If faces Δ_j, Δ_k share an edge e in $\mathfrak{S}$, then the portions of P_0 and P_t corresponding to $U = (\Delta_j \cup \Delta_k)$ are flat so they can be laid out together in the plane as in Figure 7. This means that we can define the maps h_{tj} and h_{tk} simultaneously for the two faces via the discrete version of Schwarz reflection across e. Also, note that by specifying the zero curvature at boundary vertices other than those of $\mathcal{V}$, the boundary segments in our faces are *straight* — isometrically isomorphic to euclidean line segments.

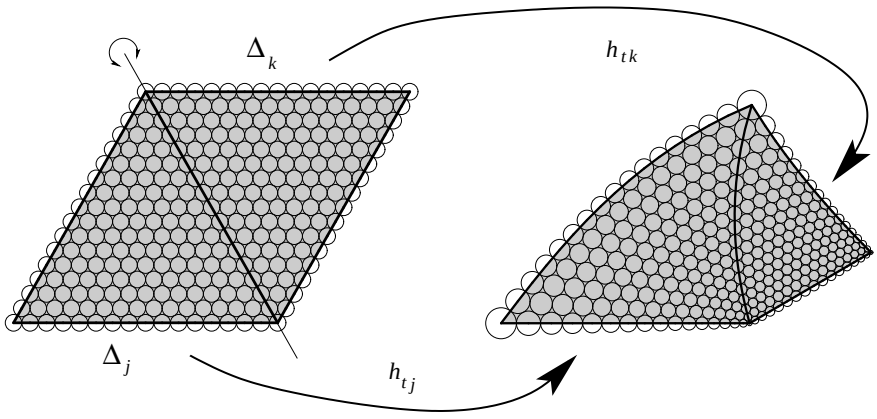

FIGURE 7. h_{tj} and h_{tk} related by Schwarz reflection.

The metrics d_t will be called discrete *PFR-metrics*, for "piecewise, flat, reflective". That is, the individual faces of $\mathfrak{S}$ are flat in d_t and contiguous faces are discrete Schwarz reflections of one another.

Theorem 4.1. *For each $n > 0$ there exists a parameterized family $M(n) = \{\mathfrak{S}_t^{(n)} = (\mathfrak{S}, d_t^{(n)}) : t \in [0,1]\}$ of discrete PFR-metric spaces with the following properties:*

(a) *Boundary edges of $\mathfrak{S}_t^{(n)}$ are straight.*
(b) *$\mathfrak{S}_0^{(n)}$ is isometrically isomorphic to $\mathfrak{S}$ (i.e., with its original equilateral metric).*
(c) *$\mathfrak{S}_1^{(n)}$ is isometrically isomorphic to a euclidean rectangle $\mathfrak{R}^{(n)}$, corners identified with corners.*
(d) $\mathrm{Area}(\mathfrak{S}_t^{(n)}) = 1$.
(e) *For each $v \in \mathcal{V}$, the cone angle of $\mathfrak{S}_t^{(n)}$ is $tA(v) + (1-t)\Theta(v)$.*

This study originated in the attempt to understand changes in radii between the *in situ* packing Q in $\mathfrak{S}$ and the flat packing P of a rectangle. Fix the refinement level at $n = 6$ in the following. Using cone angles $a(t)$ of (3) to drive the process, we obtain a continuum of intermediate stages $K(R_t)$. A natural quantity to consider along this continuum is the relative rate of change:

$$\frac{R'(t,v)}{R(t,v)} \approx \frac{R(t+\delta,v) - R(t,v)}{\delta\, R(t,v)}, \qquad v \in K. \tag{4}$$

The intermediate packings do not lie flat in the plane, so we choose the final flattened packing $P = P_1$ as a common domain for our visualizations. Let us carry out a specific computation; fix $t \in [0,1]$ and some small δ. For each vertex $v \in K$, let z_v be the center of its circle in $\mathfrak{R} = \mathrm{carr}(P_1)$ and graph the associated value from (4). This plotting defines a triangulated surface in $\mathbb{R}^3$ with domain $\mathfrak{R}$, and we can get a sense of the function (4) from the (discrete) gradient field of this graph; We used MATLAB® to sprinkle drops of a simulated non-inertial fluid on the surface and capture their motion through the field. Take a moment to contemplate the result as shown in Figure 8. In comparing this to Figure 5, the reader may feel that the arrows are pointing the wrong way. This will be explained when we interpret this as a flow of gaussian curvature, and for now you may simply put a negative sign on gravity.

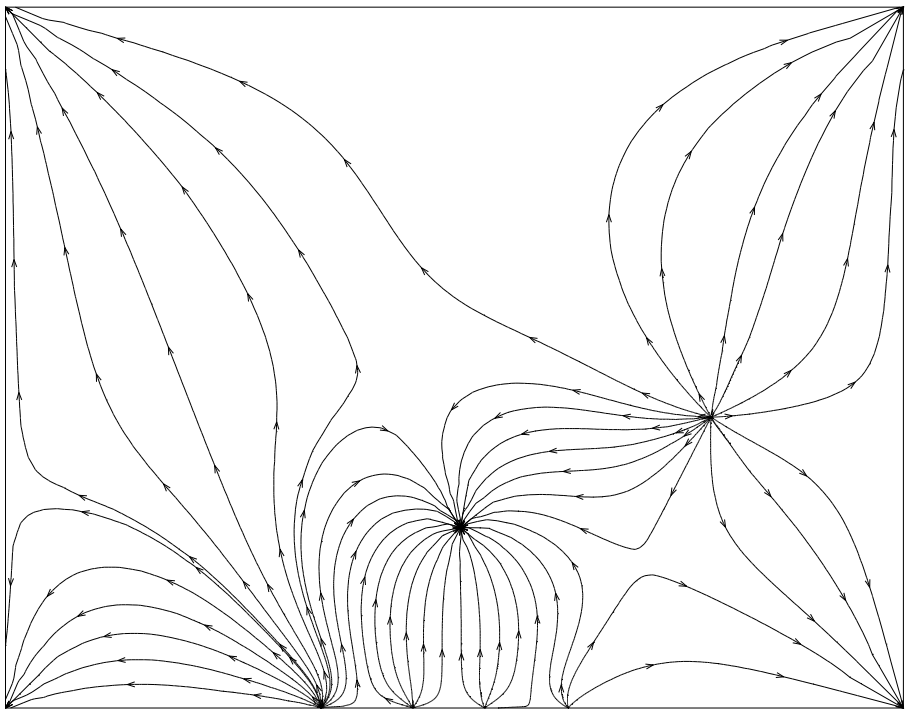

FIGURE 8. Simulated flow plotted on $\mathfrak{R}$.

A more interesting question: What was the value of t we used here? This brings us to the surprise which led to this paper: *This flow is essentially independent of t.* This is an experimental, not a formal mathematical statement. In our initial experiments, Collins and Stephenson computed the packings using 2% increments in t and were amazed to find that flows for successive increments were visually indistinguishable. We were well aware that the radius changes in moving from R_s to $R_{s+\delta}$ are tightly choreographed; remember that there is no non-zero curvature except at vertices of $\mathcal{V}$. The radii are all equal when $t = 0$ and change only slightly in moving to $t = 0.02$, whereas when $s = 0.98$, the radii have nearly reached their final values — they vary from one region to another by nearly 4 orders of magnitude. Nonetheless, the pictures suggested that the "impulse" for change at each stage was consistent. From the very beginning the circles seem to sense the interactions with their neighbors which would persist throughout the parameterization.

5. Classical Parameterization

There is a classical analogue to the parameterization of Theorem 4.1. If we let n grow to infinity, the limit behavior of the discrete simplicial maps between the various circle packings are quite well understood. The discrete conformal metrics $d_t^{(n)}$ will converge uniformly in n and t to classical conformal metrics d_t on $\mathfrak{S}$. The limit metrics are flat on each face of $\mathfrak{S}$, straight on each boundary edge, reflective between contiguous faces, and have cone angles $a(t, v)$ for $v \in \mathcal{V}$, hence we refer to them as *continuous* PFR-metrics.

Theorem 5.1. *There exists a parameterized family $M = \{\mathfrak{S}_t : t \in [0, 1]\}$, $\mathfrak{S}_t = (\mathfrak{S}, d_t)$, of continuous PFR-metric spaces, unique up to isometries, with the following properties:*

(a) *Boundary edges of $\mathfrak{S}_t$ are straight.*
(b) *$\mathfrak{S}_0$ is isometrically isomorphic to $\mathfrak{S}$.*
(c) *$\mathfrak{S}_1$ is isometrically isomorphic to a euclidean rectangle $\mathfrak{R}$, corners identified with corners.*
(d) $\text{Area}(\mathfrak{S}_t) = 1$.
(e) *For each $v \in \mathcal{V}$, the cone angle of $\mathfrak{S}_t$ is $tA(v) + (1 - t)\Theta(v)$.*

These results are not difficult once you see what happens on individual faces Δ of $\mathfrak{S}$. Fix t. The discrete analytic maps $h_t^{(n)}$, suitably normalized, converge uniformly in n to a classical analytic map $h_t \colon \Delta \to \Omega$ from Δ to some plane region Ω. (See [3, Chap 4].) The restriction of d_t to Δ is the conformal metric with density $|h_t'|$ (with respect to the equilateral metric on $\mathfrak{S}$). The map h_t is locally univalent, and though there is no guarantee that it will be univalent on Δ, we will refer to the image Ω as a euclidean "tile". Constructions of the $h_t^{(n)}$

guarantee that an edge of Δ in the boundary of $\mathfrak{S}$ will correspond to a euclideanly straight edge in its tile.

If faces of $\mathfrak{S}$ share an edge, then after rigid motions their tiles will fit together along the corresponding edges. The behavior of the discrete maps on contiguous faces (see Figure 7) implies that in the limit this edge must be an analytic arc and that the tiles will be Schwarz reflections of one another across that arc in the classical sense.

One can now describe the PFR-metric on $\mathfrak{S}$ very concretely: $\mathfrak{S}_t$ consists of ten euclidean tiles identified according to the combinatorics of $\mathcal{T}$ to form a cone space with cone angles $a(t, v)$. Figure 9 illustrates this for parameter values $t = 0.25, 0.5, 0.75$, and 1. For a given t, each face has three analytic arcs as boundary and contiguous faces could be moved rigidly to match along their shared edge. As t increases, the curvatures at the cone points are cancelling, leaving only curvature $\pi/2$ at the four corners when $t = 1$; only at that final stage can all these matchings be accomplished simultaneously in the plane.

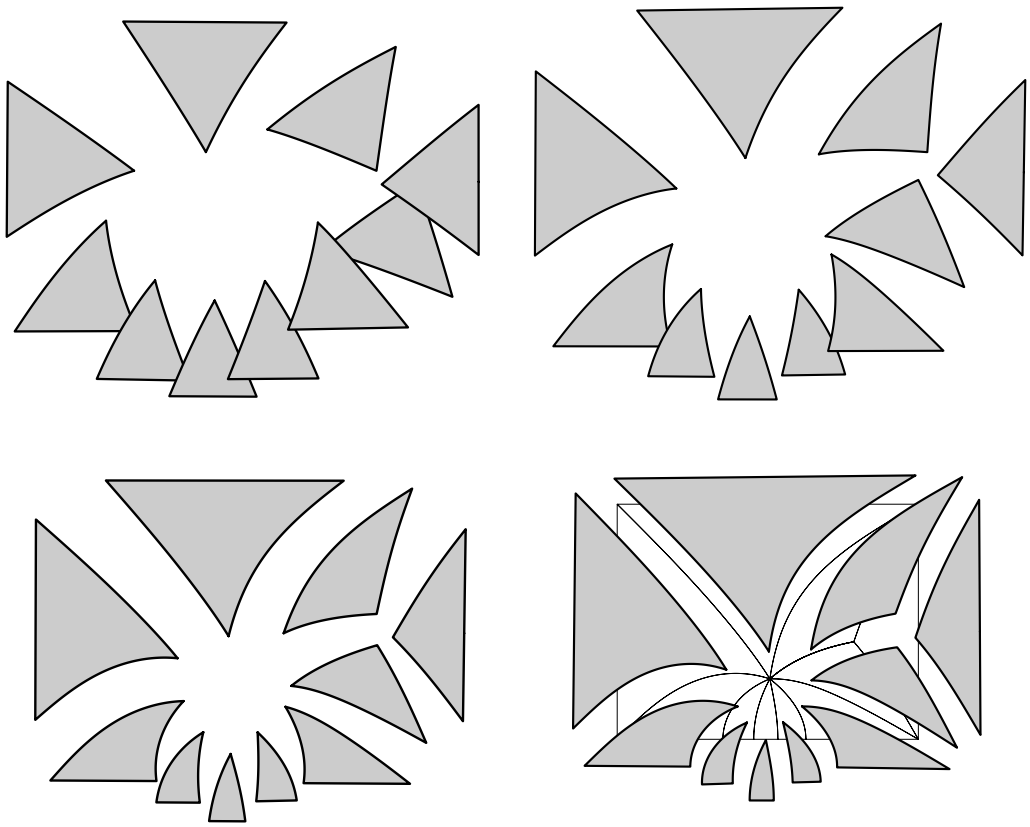

FIGURE 9. Banishing 25%, 50%, 75%, and 100% of misplaced curvature.

We have not yet explained any "flow", but at least we now have the classical model of our initial discrete setting, with metrics defined by classical analytic functions. We can now exploit some classical machinery.

6. Schwarz-Christoffel Formulation

We explicitly construct the classical mappings for M using a modification of Schwarz-Christoffel (SC) methods for conformal mapping. See [7] for details on Schwarz-Christoffel mappings; our SC computations are based on Driscoll's `MATLAB` toolbox [6].

A typical application of SC involves a polygonal plane region Λ as target and the upper half plane $\mathbb{H}^+ = \{x + iy : y > 0\}$ as domain. One gets data from Λ in the form of corner locations $\{w_1, \ldots, w_n\}$ and interior angles $\{\alpha_1\pi, \ldots, \alpha_n\pi\}$. The conformal map $F\colon \mathbb{H}^+ \to \Lambda$ can be shown to have derivative of the form

$$F'(z) = \prod_{k=1}^{n-1} (z - z_k)^{\alpha_k - 1},$$

with z_n placed at infinity. The SC method depends on numerical integrals of F' in $\mathbb{H}^+$; adjustments are made in the locations of the pre-vertices z_j until numerical integration places their images sufficiently close to their intended images w_j. Once the pre-vertices are known, further numerical integrations can be used to find the images of any points under F or its inverse.

SC methods also work in nonplanar situations and permit interior branching compatible in the sense of the Gauss-Bonnet Theorem with the boundary specification; see §4.7 [7] for an example. A branch point of order m for F represents a cone angle of $(m+1)2\pi$ in the image. Our example involves more general interior cone angles, however. Each interior cone point requires a complex "pre-cone" point $b \in \mathbb{H}^+$ and an appropriate exponent. A cone angle of $\beta > 0$ requires that F behave locally at b like $(z - b)^{\beta/2\pi}$, so its derivative behaves like $(z - b)^{\beta/2\pi - 1}$; a companion factor for $\overline{b}$ makes the terms real on the real axis. Figure 10 illustrates the factors one would need in F' at interior cone points where various numbers k of equilateral triangles meet.

When β is not an integral multiple of 2π, then integration of F' around b would be multiple-valued. Therefore, for each such factor in F' we must prescribe a *cut*, γ, that is, a curve from b to $\partial\mathbb{H}^+$; integration of F' along curves which do not cross any of these cuts will remain single-valued.

Let us put our situation into this setting. Based on the cone angles (1), our map F requires real pre-vertices $z_5, \ldots, z_{10}$ on the x-axis (corner v_3 is put at ∞ and v_4 has target angle π so its factor is trivial) and complex pre-cones b_1, b_2

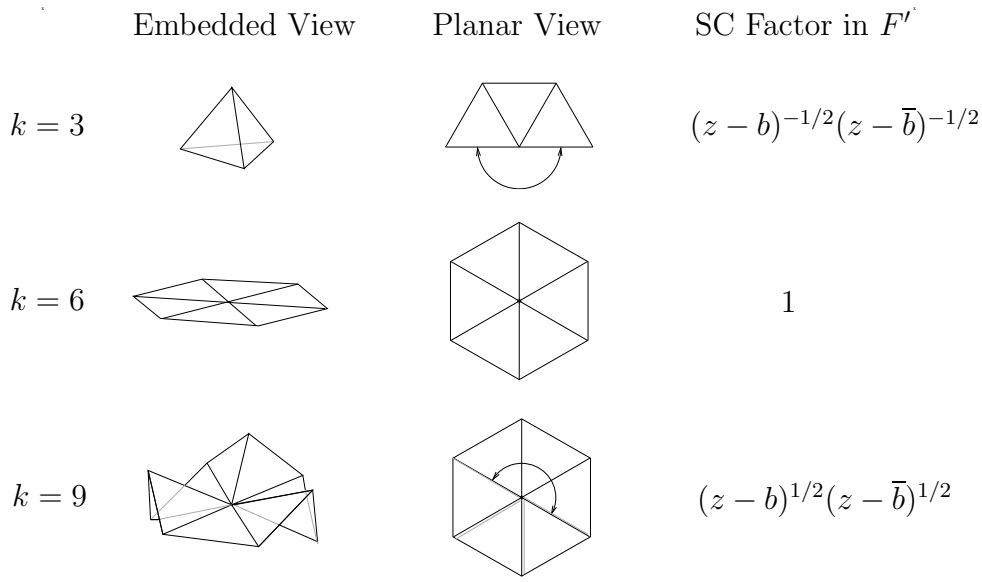

FIGURE 10. Factors in F' for various cone angles.

in $\mathbb{H}^+$. The expression for F' is

$$F'(z) = (z-b_1)^{1/2}(z-\overline{b}_1)^{1/2}(z-b_2)^{-1/2}(z-\overline{b}_2)^{-1/2}\prod_{k=5}^{10}(z-z_k)^{-1/3}. \tag{5}$$

The locations of the z_k and b_j are, of course, the missing ingredients. We would like to get these from circle packing, except that our flat circle packing lies in a rectangle, not in $\mathbb{H}^+$. To remedy that, note that $\phi\colon \mathfrak{S} \to \mathfrak{R}$ may be represented as a composition of mappings through $\mathbb{H}^+$. The diagram in Figure 11 commutes; all the maps are Riemann mappings, that is, one-to-one conformal maps.

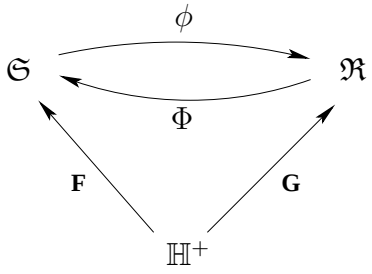

FIGURE 11. Factoring through $\mathbb{H}^+$.

Circle packing gives us an approximation of $\phi\colon \mathfrak{S} \to \mathfrak{R}$. A standard application of SC mapping provides $G\colon \mathbb{H}^+ \to \mathfrak{R}$. The preimage under G of the vertices and edges in the circle packing of Figure 2 yields the vertices and edges in $\mathbb{H}^+$ shown in Figure 12(a). The vertices there give us the presumptive pre-vertices and pre-cones in the expression (5). To check them we run a numerical integration of F'

on $\mathbb{H}^+$, avoiding the indicated cuts γ_1, γ_2, and display the images of the vertices and edges; the result is shown in Figure 12(b).

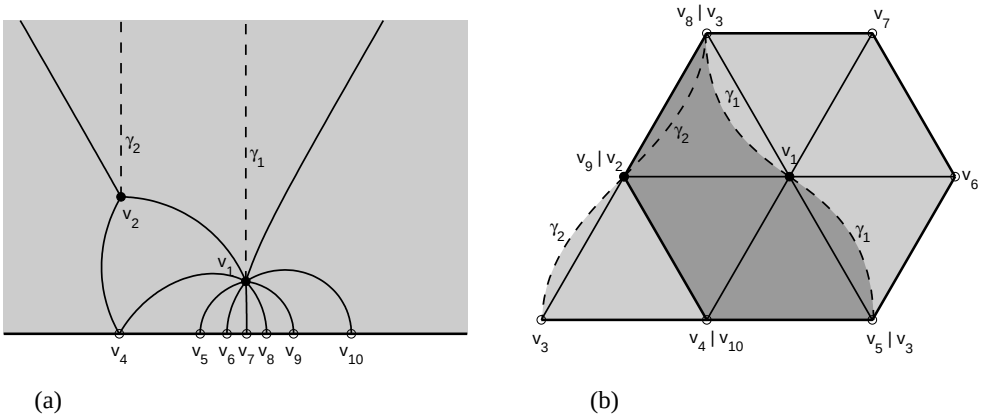

FIGURE 12. $\mathbb{H}^+$ and its computed image.

Figure 12(b) may be difficult to interpret. Note that the images of the vertices $v_1, \ldots, v_{10}$ are labelled, and using them one can check that the computed image of $\mathbb{H}^+$ under F consists of 10 *approximately* equilateral triangles identified in the pattern of the triangulation T, except that each of the two cuts has two edges (the dashed curves) separated by angle π. The computed function F differs from the exact function F only to the extent that these triangles differ from *exact* equilateral triangles. On a visual level one would have to say that the accuracy seems quite good.

To check that accuracy, we will show how one can modify the SC methods to compute F independently — that is, without using the locations provided by the circle packing. We basically reverse the experiment above. If $\mathfrak{S}$ is cut along the edges $\langle v_1, v_3\rangle$ and $\langle v_2, v_3\rangle$, it lays out as a pattern of 10 equilateral triangles in the plane — we refer to this as a "generalized polygonal region" Λ. The region Λ is essentially that of Figure 12(b); if we assume that the point labelled "$v_4|v_{10}$" is at the origin and that labelled "$v_5|v_3$" is at $z = 1$, then we can infer precise locations and interior angles for all the vertices. For example, the image of interior cone point v_1 is $z = (1 + \sqrt{3})/2$, the image of the corner v_6 is $z = (3 + \sqrt{3})/2$; v_3 gets three locations due to the cuts, which both end there.

We now have from Λ the data that SC traditionally requires about the intended image. In addition, however, we must account for the cuts. These must be nonintersecting curves which are topologically equivalent to the cuts in $\mathfrak{S}$. The curves γ_1, γ_2 in Figure 12(a) would be examples, but of course, we don't know the locations of the pre-cones v_1, v_2. To get started, we need just a small bit of information about them — for example, even the coarsest packing, Figure 3, suggests (when transferred to $\mathbb{H}^+$) that v_1 will be to the left of v_2. Set initial locations for v_1, v_2 accordingly and assume the associated cuts extend vertically;

choose initial locations on the real line for $v_5, \ldots, v_{10}$, left to right. Now, SC mapping proceeds by incrementally shifting the pre-vertices and pre-cones — and updating their cuts — until numerical integration puts the images of the pre-vertices and pre-cones in their correct locations *vis-a-vis* Λ.

We anticipate that the final pre-vertices and pre-cones computed by the SC methods are accurate to 8 digits. The fact that they are within 2.3×10^{-5} of the locations predicted by circle packing is strong confirmation of circle packing accuracy. That is valuable information, for with surfaces even mildly more complicated than $\mathfrak{S}$, tracking and enforcing required cuts in SC will be a challenge.

7. Flow

Recall that the conformal mapping $\phi\colon \mathfrak{S} \to \mathfrak{R}$ is the inverse of the mapping $\Phi = F \circ G^{-1}\colon \mathfrak{R} \to \mathfrak{S}$. Figure 13 shows the flow lines (level curves of the harmonic conjugate) for $\log(|\Phi'|)$ in $\mathfrak{R}$. Compare this to the gradient field computed experimentally in Figure 8. Minor differences — as with the number and smoothness of the flow lines — are due to the pictures' very different origins: Figure 8 results from simulated water droplets flowing on a polyhedral surface, while Figure 13 is a standard numerical plot using the computed SC mapping. Our interest lies with the striking *similarities* between these two fields.

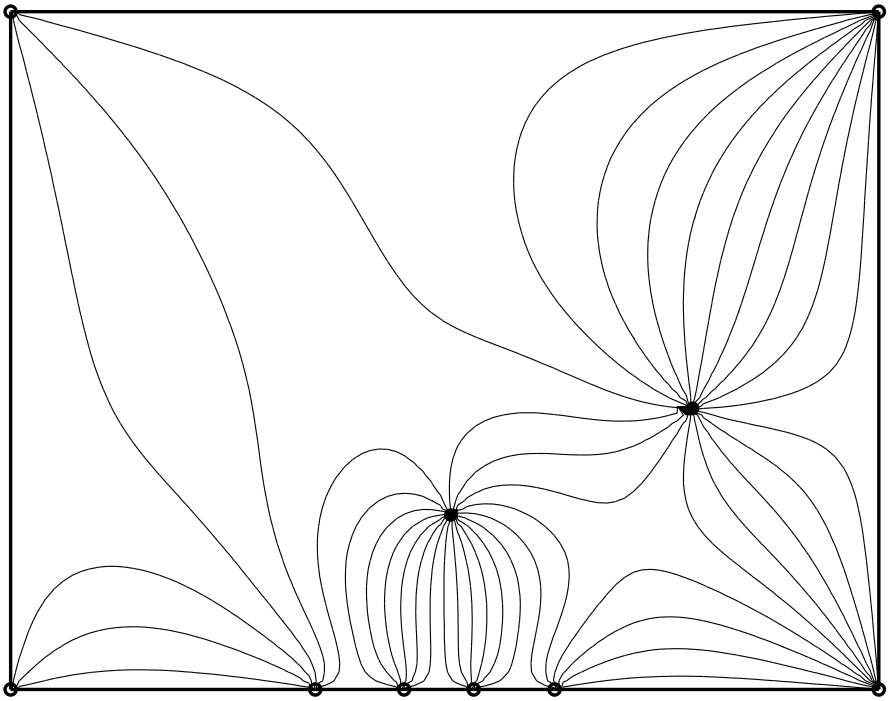

FIGURE 13. The gradient field for $\Phi\colon \mathfrak{R} \to \mathfrak{S}$.

The natural domain in working with Φ' is $\mathfrak{R}$, not $\mathfrak{S}$, so it takes a little work to recognize our parameterization. For $t \in [0,1]$, define Φ_t, up to a constant of integration, by

$$(\Phi_t)'(z) = (\Phi'(z))^{(1-t)} = \exp((1-t)\log(\Phi'(z))), \qquad z \in \mathfrak{R}.$$

Let m_t be the conformal metric on $\mathfrak{R}$ with density proportional to $|\Phi_t'|$ *vis-a-vis* the euclidean metric and area 1. One can confirm from the behavior of this density on the boundary and at the two interior cone points in $\mathfrak{R}$ that $(\mathfrak{R}, m_t)$ is isometrically isomorphic to $(\mathfrak{S}, d_t)$.

Fix a face Δ of $\mathfrak{S}$ and let $\Omega = \phi(\Delta)$ denote the corresponding (curvilinear) face in $\mathfrak{R}$. Fixing a point $z_0 \in \Omega$ for reference, define Φ_t on Ω so $\Phi_t(z_0) = z_0$, and define $h_t \equiv \Phi_t \circ \phi$ on Δ; then $|h_t'|$ is the density for the metric d_t on Δ. The images $\Omega_t = h_t(\Delta)$ morph smoothly from $\Omega_0 = \Delta$ to $\Omega_1 = \Omega$ as t goes from 0 to 1 in the continuous analogue of the progression suggested in Figure 6.

Now fix parameter $t \in [0,1]$ and small δ and consider the situation of Figure 14, where $u = h_t(z), w = \phi(z)$. Decompose the map $h_{t+\delta}$ as $h_{t+\delta} = g_\delta \circ h_t$. By the chain rule, $h_{t+\delta}'(z) = g_\delta'(u) \cdot h_t'(z)$. By definition,

$$h_t'(z) = \Phi_t'(w) \cdot \phi'(z) = \exp((1-t)\log(\Phi'(w))) \cdot \phi'(z),$$
$$h_{t+\delta}'(z) = \Phi_{t+\delta}'(w) \cdot \phi'(z) = \exp((1-t-\delta)\log(\Phi'(w))) \cdot \phi'(z).$$

Putting these facts together gives $g_\delta'(u) = \exp(-\delta \log(\Phi'(w)))$, implying

$$\log(|g_\delta'(u)|) = -\delta \log(|\Phi'(w)|).$$

If one chooses to plot the gradient field for $\log(|g_\delta'|)$ on Ω in the w-plane, it is simply a scalar multiple of that for $\log(|\Phi'|)$. In particular, the gradient field of $\log(|g_\delta'|)$, depicted for all the faces simultaneously in Figure 13, is independent of t and δ.

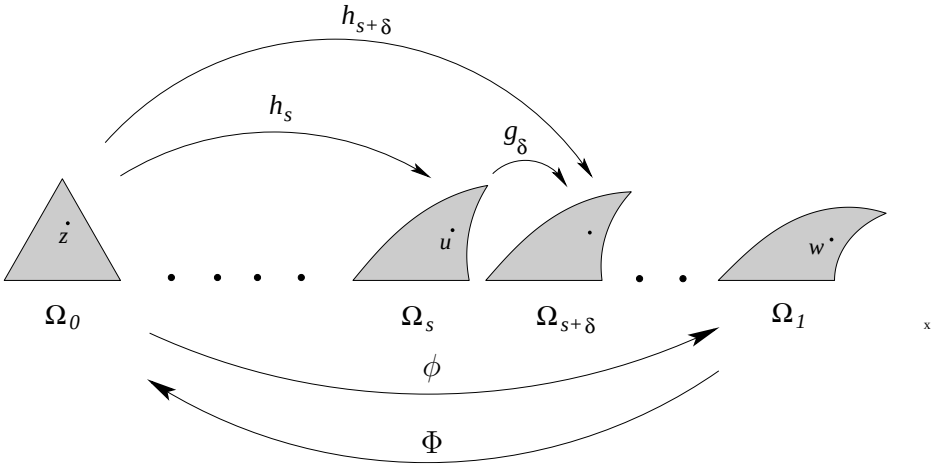

FIGURE 14. An incremental map g_δ.

Moving back to the discrete setting, take this same t and δ and work again at level $n = 6$. The discrete version of the incremental map g_δ is the discrete conformal map (the simplicial map) $\tilde{g}_\delta \colon P_t|_\Delta \to P_{t+\delta}|_\Delta$. Discrete analytic function theory tells us not only that $\tilde{g}_\delta$ approximates g_δ, but also that its *ratio function* $\tilde{g}_\delta^{\#}$ approximates $|g_\delta'|$. Recall that $\tilde{g}_\delta^{\#}$ is just the ratio of image radii to domain radii,

$$\tilde{g}_\delta^{\#}(v) = \frac{R(t+\delta, v)}{R(t, v)}.$$

Let v be a vertex of K, $v \in \Delta$, and parallel to our earlier notation let z, u, and w be the centers of the circles for v in P_0, P_t, and P_1, respectively. We have

$$-\delta \log(|\Phi'(w)|) = \log(|g_\delta'(u)|) \sim \log(\tilde{g}_\delta^{\#}(v)) = \log\left(\frac{R(t+\delta, v)}{R(t, v)}\right)$$
$$\sim \log\left(\frac{R(t, v) + \delta R'(t, v)}{R(t, v)}\right) = \log\left(1 + \delta \frac{R'(t, v)}{R(t, v)}\right).$$

We can now see why Figure 8 and Figure 13 are so similar: for small δ this last logarithm is roughly $\delta R'(t, v)/R(t, v)$, so we conclude

$$-\log(|\Phi'(w)|) \sim \frac{R'(t, v)}{R(t, v)}.$$

(The negative sign simply reflects the opposite mapping directions between the discrete and continuous settings.)

Are we seeing curvature flow? In the very concrete world of circle packing, this seems to be a fair interpretation. Our parameterization $\{R_t : t \in [0, 1]\}$ is formulated to keep all non-zero curvature at the 10 cone points of $\mathcal{V}$ (the faces are euclideanly flat and the boundary edges between points of $\mathcal{V}$ are straight). As $t \to 1$ that curvature is redistributed in a continuous fashion by driving the cone angles at $\mathcal{V}$ to their target values in (2).

But let us look at the details. Say we fix level $n = 6$ and consider parameters $t < t'$ in $[0, 1]$. Adjusting the label R to move from the surface $K(R_t)$ to $K(R_{t'})$ is an iterative process (at least, in the algorithm used in `CirclePack`). On visiting a vertex v, whose intermediate cone angle target is $a(t', v)$, the difference $a(t', v) - \theta_R(v)$ represents unwanted curvature, so an adjustment is made in $R(v)$ to distribute this curvature to v's neighbors. These changes in R quite literally *push curvature around the surface.* Our experiments were attempts to better understand this process with the aim of improving computational efficiency. This flow among $20,767$ vertices is too complex, however, so we decided to repack incrementally: prescribe, say, a 2% reduction in the curvature at the vertices of $\mathcal{V}$, but let the repacking computation run to achieve zero curvature at all other vertices. In this way we captured the movement of this curvature just among the vertices of $\mathcal{V}$. The arrows in Figure 6 represent the direction of curvature flow: from the excess at v_2 to the shortage at v_1, for example; out of v_4 and into v_5; and so forth.

As mentioned earlier, the first surprise in the experiments was the consistency: the flows from t to $t + 0.02$ were virtually indistinguishable for various t. We also found that this flow organizes itself much earlier in the refinement stages: compare the flow for $n = 3$ shown in Figure 15, involving only 353 circles, with that for $n = 6$ in Figure 8.

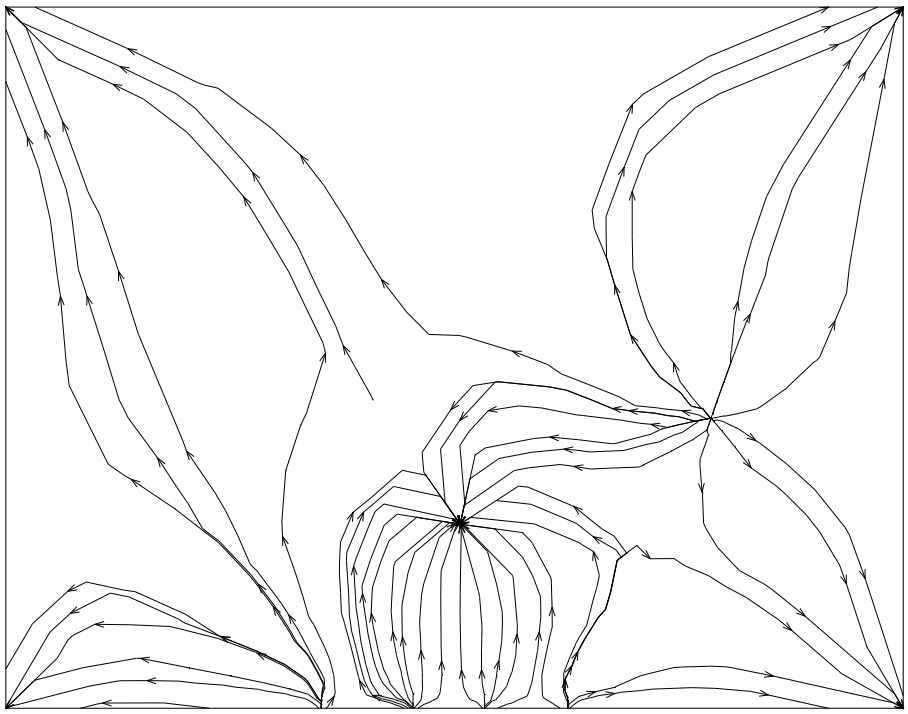

FIGURE 15. Simulated curvature flow, level $n = 3$.

There are also other ways to "drive" the curvature flows, many functions $a(t, v)$ that respect the Gauss-Bonnet Theorem and could be used in place of (3). For example, it happens by chance that the two interior curvatures of $\mathfrak{S}$ sum to zero; therefore, one can flatten the interior while holding the boundary curvatures fixed, and then drive the boundary to a rectangle while keeping the interior flat. The curvature flows within these two phases are shown for level $n = 6$ in Figure 16.

Not only does each flow persist throughout its parameter range, but in fact the two flows are indistinguishable from those obtained when we run the boundary flow *first* and the interior flow *second*. This is, of course, perfectly consistent with (5): $\log |\Phi'|$ is a sum of the logarithms of the boundary terms and the interior terms, and the flows in Figure 16 reproduce their gradient fields. Remember, however, that our images result from actual circle packing runs, so the fidelity is surprising nonetheless.

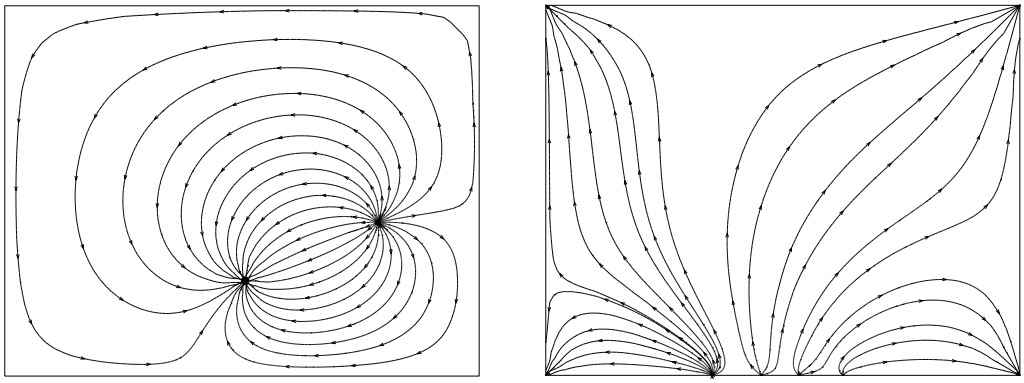

FIGURE 16. Splitting the flow: interior, boundary.

The phenomena we have seen in flattening $\mathfrak{S}$ also occur in more traditional SC mapping situations, that is, without interior cone points. Suppose we flatten the interior of $\mathfrak{S}$ as described above; the result is a flat polygonal region Λ in the plane with corner angles $\Theta(v_3), \ldots, \Theta(v_{10})$ (i.e., the boundary cone angles of (1)). Taking this as a starting point, we can prescribe a linear change in boundary angles to morph Λ into $\mathfrak{R}$ by setting the targets

$$a(t, v_j) = tA(v_j) + (1-t)\Theta(v_j), \qquad j = 3, \ldots, 10, \ t \in [0, 1].$$

For each t, the corresponding surface is a polygonal region Λ_t in the plane. Figure 17 shows Λ_t for $t = 0.0, 0.33, 0.85$, and 1.0; $\Lambda_0 = \Lambda, \Lambda_1 = \mathfrak{R}$.

The flow shown in Λ_1 is that on the right of Figure 16 and it is transported to the other polygons by the associated conformal maps. This may at first seem artificial, but it is not at all.

Circle packing is very much a *Lagrangian* world, with the action following the vertices. Label computations are associated with vertices of the *abstract* complex K, so patterns of coherence which one finds in these computations are, technically speaking, patterns on K. When a label satisfies the packing condition and a corresponding circle packing P is laid out in the plane, the vertices get concrete geometric locations; it is quite natural to transfer the coherence pattern from K to its geometric realization in P.

What is **not** natural is to expect that transferred pattern to have some intrinsic geometric significance! For each t, our flow reflects the behavior of R_t'/R_t within K, yet when transferred to Λ_t it is found to be the gradient field of a classical harmonic function. This reflects the intimate connection between circle packing and conformal geometry first noted in Thurston's Conjecture [17], the Rodin/Sullivan Theorem [12], and reconfirmed in subsequent developments [15, 16].

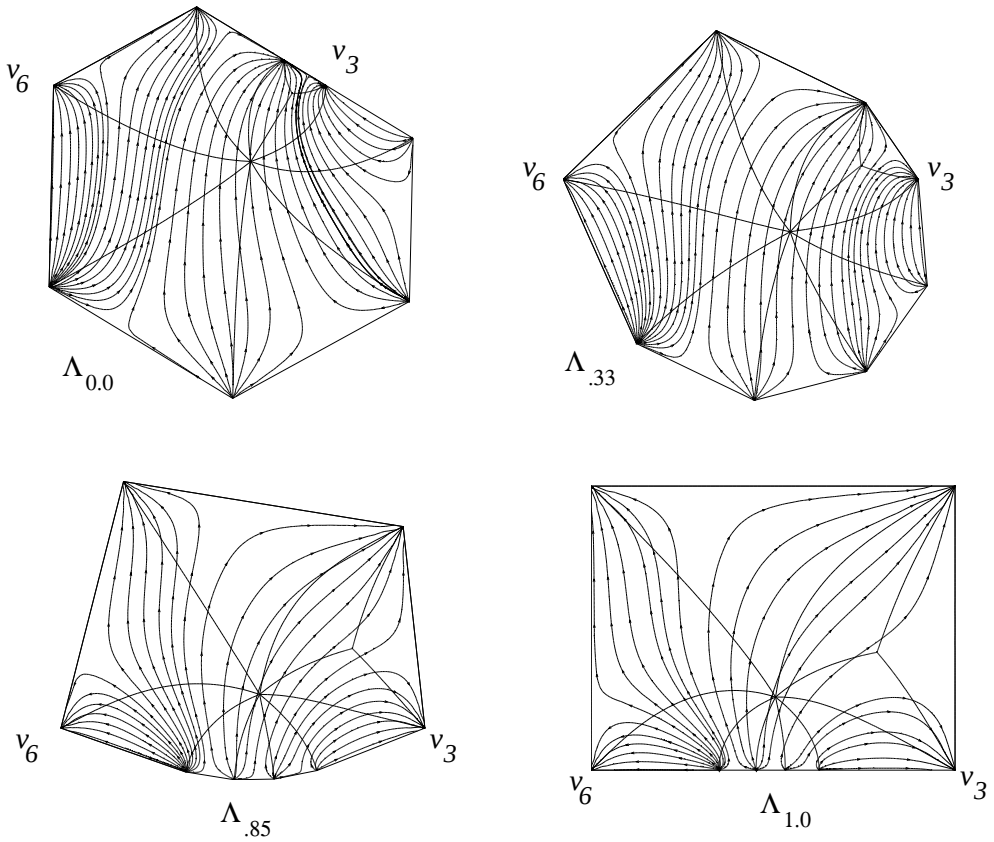

FIGURE 17. A planar flow.

8. Wrap Up

This paper has illustrated some experimentally observed phenomena and considered their classical counterparts. There are plausible connections in several other directions that remain to be investigated. We have driven our sample flows so that curvature *changes* are proportional to curvature, and in that sense these are perhaps *Ricci*-type flows (though not in the same sense as [4]). On the other hand, they are reminiscent of the *Löwner chains* of conformal maps onto slit regions, which are driven by Löwner's differential equation (see [9] and [13]). Is there some sensible notion of "energy" associated with these flows? Would $\mathfrak{S}$ naturally *relax* to $\mathfrak{R}$ using our linear flow? Is there a brownian motion formulation whose travellers have curvature strapped to their backs? (See the Markov model of [14]).) There are clear connections to cone spaces, and some similarity to quadradic differentials and Teichmüller theory, though our surfaces here have all been conformal discs.

In the other direction, we have not yet learned how to exploit these phenomena in our repacking algorithms. Starting from label R_0, for example, the computation of, say, $R_{.02}$ takes approximately the same time as would the computation of the

final label R_1. Nonetheless, it's clear that the circles learn quickly how they must change, and we are certainly closer to understanding the self-organization which we have already exploited numerically to some extent via the *superstep* methods of [5].

So, are these flows "explained" now that we have identified them with classical analytic functions and potentials? Perhaps in hindsight for the simple surface $\mathfrak{S}$. But in more complicated settings, circle packing may be the only access one has to the classical objects. These flows occur, for example, in the conformal mapping of Riemann surfaces, as with the equilateral surfaces at the heart of Grothendieck's *dessins d'Enfants*; see [3]. They can be seen, also, in more general analytic functions. For instance, we could drive the curvature in $\mathfrak{R}$ so that vertex v_1 becomes a simple branch point (cone angle 4π) by drawing the extra π curvature from v_2. These flows occur in flattening more general abstract polyhedral surfaces, for example in the topic of *conformal tiling* [2]. In all these cases, the Schwarz-Christoffel evidence we have shown for the accuracy of the circle packing maps is certainly encouraging. There is still much to learn, however: in recent "brain mapping" applications, the typical cortical surface from an fMRI scan will have $200,000$ vertices rather than 10 (see [10])!

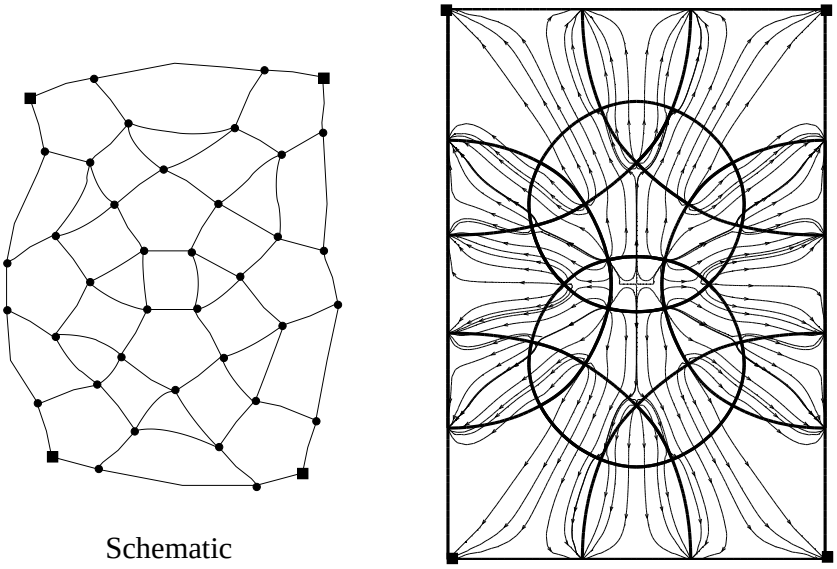

FIGURE 18. Curvature flow for a dodecahedral tiling.

Let us finish in the spirit of experimentation with the dodecahedral tiling pattern of Cannon, Floyd, and Parry [11]. On the left in Figure 18 is an abstract schematic with cells which are pentagons, rectangles, and triangles and four identified corners. This can be made into an abstract polyhedral surface by treating each face as a regular euclidean polygon, and if we put aside some subtleties concerning the edge pastings, this can be flattened to a rectangle. Figure 18 shows

an approximation to the conformally correct image at refinement level $n = 4$, and superimposed on that is the associated flow.

References

1. A. F. Beardon, *A primer on Riemann surfaces*, London Math. Soc. Lecture Note Series, vol. 78, Cambridge Univ. Press, Cambridge, 1984.
2. _____, A "regular" pentagonal tiling of the plane, *Conformal Geometry and Dynamics* **1** (1997), 58–86.
3. P. L. Bowers and K. Stephenson, Uniformizing dessins and Belyĭ maps via circle packing, *Amer. Math. Soc.*, to appear.
4. B. Chow and F. Luo, Combinatorial Ricci flows on surfaces, Preprint, 2002.
5. Ch. R. Collins and K. Stephenson, A circle packing algorithm, *Comput. Geom.* **25** (2003), 233–256.
6. T. A. Driscoll, *A MATLAB toolbox for Schwarz-Christoffel mapping*, 1996.
7. T. A. Driscoll and L. N. Trefethen, *Schwarz-Christoffel Mapping*, vol. 8, Cambridge Univ. Press, Cambridge, New York, 2002.
8. T. Dubejko and K. Stephenson, Circle packing: Experiments in discrete analytic function theory, *Exp. Math.* **4** (1995) no.4, 307–348.
9. P. L. Duren, *Univalent Functions*, vol. 259, Springer-Verlag, New York, 1983.
10. M. K. Hurdal, P. L. Bowers, K. Stephenson, D. W. L. Sumners, K. Rehm, K. Schaper, and D. A. Rottenberg, Quasi-conformally flat mapping the human cerebellum, in: C. Taylor and A. Colchester (eds.), *Medical Image Computing and Computer-Assisted Intervention* - MICCAI'99, vol. 1679, Springer, Berlin, 1999, 279–286.
11. J. W. Cannon, W. J. Floyd, and W. Parry, Finite subdivision rules, *Conform. Geom. Dyn.* **5** (2001), 153–196.
12. B. Rodin and D. Sullivan, The convergence of circle packings to the Riemann mapping, *J. Differ. Geom.* **26** (1987), 349–360.
13. O. Schramm, Scaling limits of loop-erased random walks and uniform spanning trees, *Israel J. Math.* **118** (2000), 221–288.
14. K. Stephenson, A probabilistic proof of Thurston's conjecture on circle packings, *Rend. Semin. Mat. Fis. Milano* **66** (1996), 201–291.
15. _____, Approximation of conformal structures via circle packing, in: N. Papamichael, St. Ruscheweyh, and E. B. Saff (eds.), Computational Methods and Function Theory 1997, Proceedings of the Third CMFT Conference, vol. 11, World Scientific, 1999, 551–582.
16. _____, Circle packing and discrete analytic function theory, in: R. Kühnau (ed.), *Handbook of Complex Analysis, Vol. 1: Geometric Function Theory*, Elsevier, 2002.
17. W. Thurston, *The finite Riemann mapping theorem*, 1985, Invited talk, An International Symposium at Purdue University on the occasion of the proof of the Bieberbach conjecture, March 1985.

Charles R. Collins E-MAIL: ccollins@math.utk.edu
ADDRESS: *Mathematics, University of Tennessee, Knoxville, TN 37996-1300, U.S.A.*

Tobin A. Driscoll E-MAIL: driscoll@math.udel.edu
ADDRESS: *Mathematical Sciences, University of Delaware, Newark DE 19716, U.S.A.*

Kenneth Stephenson E-MAIL: kens@math.utk.edu
ADDRESS: *Mathematics, University of Tennessee, Knoxville, TN 37996-1300, U.S.A.*
URL: www.math.utk.edu/∼kens