A CIRCLE PACKING ALGORITHM

CHARLES R. COLLINS AND KENNETH STEPHENSON *

Abstract. A circle packing is a configuration P of circles realizing a specified pattern of tangencies. Radii of packings in the euclidean and hyperbolic planes may be computed using an iterative process suggested by William Thurston. We describe an efficient implementation, discuss its performance, and illustrate recent applications. A central role is played by new and subtle monotonicity results for "flowers" of circles.

Key words. circle packing, conformal geometry, discrete Dirichlet problem

Introduction. A circle packing is a configuration P of circles realizing a specified pattern of tangencies. As such, it enjoys dual natures — *combinatoric* in the pattern of tangencies, encoded in an abstract "complex" K, and *geometric* in the radii of the circles, represented by a radius "label" R. As an early example, Figure 0.1 displays a simple complex K and a circle packing having its combinatorics. More substantial packings involve several hundred thousand circles.

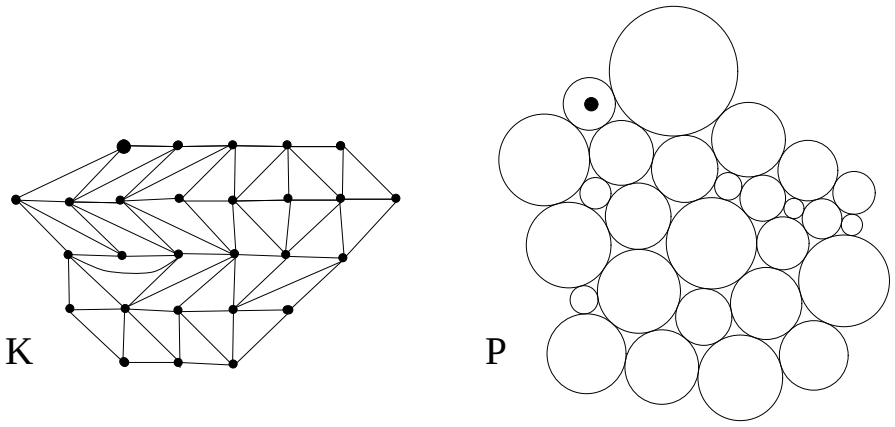

FIG. 0.1. *A simple packing example*

Our problem: *Given a complex K and appropriate "boundary conditions", compute the radii of the corresponding circle packing for K.*

There now exists a rather complete theory covering the existence and uniqueness of these radii. In this paper we describe an efficient algorithm for numerically approximating them in euclidean and hyperbolic geometry.

Though circle packings appear first with Koebe [23], they were rediscovered by William Thurston in [34]. (Important note: our circle packings are NOT those in the extensive "sphere packing" literature.) Thurston conjectured in 1985 [35] that maps between circle packings could be used in the approximation of classical conformal

*Department of Mathematics, University of Tennessee, Knoxville TN 37996-1300 (ccollins@math.utk.edu, kens@math.utk.edu). The second author gratefully acknowledges the support of the National Science Foundation and the Tennessee Science Alliance.

(analytic) mappings. His conjecture was confirmed by Rodin and Sullivan [29]. Since then, many additional uses of circle packings, both practical and theoretical have emerged: discrete conformal mapping [18, 33, 16, 17, 19], analytic function theory [13, 14, 15, 31, 27, 28], graph embedding [24, 25], discrete potential theory [32, 4], conformal tilings [9], and Riemann surface theory [10, 7, 6, 37, 38, 1]. There is a significant experimental component to circle packing, so both theory and applications benefit from an efficient implementation.

In computing packing labels R, one faces large, highly nonlinear, nonstructured, heterogeneous systems of equations. The underlying geometry plays a central role, with a mixture of local and global considerations which reflects the "discrete conformal" nature of circle packings; in particular, the system displays certain characteristics of classical discrete Laplace equations, including conserved geometric quantities. The global strategy in our packing algorithm, akin to "relaxation", was suggested by Thurston. Alternate approaches involving energy minimization and convexity [11] and random walks [32] have been suggested, but to our knowledge, not implemented on a significant scale.

The key implementation issues in our iterative approach are local in nature, depending heavily on the special properties of circles, and in particular, on a new "monotonicity" result of independent geometrical interest. We describe our algorithm in the context of the simplest Dirichlet-type problem; however, the implementation handles much more general situations and is now incorporated in the software package `CirclePack` developed by the second author.

In the next section we start with definitions and notation and describe the basic Dirichlet problem. In Section 2 we give the monotonicity properties of local circle patterns which are key both in theory and practice. The global iterative strategy is described in Section 3, with emphasis on the local/global interaction. In Section 4 we lay out our implementation and discuss rates of convergence, stability, and speed, and provide sample run data. The final section concerns more general circle packings, open questions, and selected applications.

1. Definitions and Notation. The principal objects of concern are circle packings P, their complexes K, and their associated labels (putative radii) R. In fact, the numerical manipulations involve only K and R: one solves for the label R satisfying desired boundary conditions and meeting numerical "packing" conditions which reflect local geometric compatibility. The circle packing itself results from a simple laying-out process and in particular, circle centers play a purely secondary role.

Geometries: Our algorithm applies in both the euclidean and hyperbolic settings. The euclidean plane is the familiar complex plane $\mathbf{C}$. The hyperbolic plane will be represented in the Poincarè disc model: that is, it consists of the open unit disc $\mathbf{D} = \{|z| < 1\}$ equipped with the Riemannian metric of constant curvature -1 having length element $\frac{2|dz|}{(1-|z|^2)}$. Note that hyperbolic circles in $\mathbf{D}$ are also euclidean circles, though with hyperbolic centers and radii. Horocycles, circles internally tangent to $\partial\mathbf{D}$, may be consistently treated as circles of infinite hyperbolic radius with centers at their points of tangency.

Complexes: Packing combinatorics are encoded in abstract simplicial 2-complexes K which triangulate oriented topological surfaces. *We restrict to the case in which K is a finite triangulation of a closed topological disc*, so we have a finite number of vertices (0-simplices), edges (1-simplices), and faces (oriented 2-simplices). (See the

concluding section for comments on the more general cases.)

The vertices of K are of two types, interior and boundary. If u and v are neighboring vertices (i.e., $\langle u, v\rangle$ is an edge of K) we write $u \sim v$. A vertex v and its neighbors form a (combinatorial) flower, $F_v = \{v; v_1, \cdots, v_k\}$: The petals v_j are listed in counterclockwise order about v with $v_{j+1} \sim v_j$; k is the degree of v, $\deg(v)$. When v is interior, the list of petals is closed; writing $v_{k+1} = v_1$, v belongs to the k faces $\{\langle v, v_j, v_{j+1}\rangle : j = 1, \cdots, k\}$. To avoid minor pathologies, we assume that the set of interior vertices of K is edge-connected and that every boundary vertex has an interior neighbor.

Packings: A configuration P of circles in the (euclidean or hyperbolic) plane is a *circle packing* for K if it has a circle c_v associated with each vertex v of K so that the following conditions hold: (1) if $\langle u, v\rangle$ is an edge of K, then c_u and c_v are (externally) tangent, and (2) if $\langle u, v, w\rangle$ is a positively oriented face of K, then $\langle c_u, c_v, c_w\rangle$ is a positively oriented triple of mutually tangent circles.

We emphasize that there is **no** *univalence* condition (as occurs in certain parts of the circle packing literature); that is, when vertices v and u are not neighbors, then there is no guarantee that their circles c_v and c_u have mutually disjoint interiors.

Labels: A label for K is a function $R : K^{(0)} \longrightarrow (0, \infty]$ assigning an (extended) positive value to each vertex of K; write $K(R)$ for the labeled complex. The archetype, of course, is the "radius label" taken from a packing P for K, wherein $R(v) = \text{radius}(c_v)$. In this case we write $K(R) \leftrightarrow P$ to indicate the association. (Note that the label ∞ is permitted only in the hyperbolic setting, and then only for boundary vertices.) The collection of all labels R for K will be denoted $\mathcal{R}$. Of course, in general a label represents only putative radii; it could not be associated with a coherent configuration of circles unless rigid compatibility conditions were satisfied.

Angle Sums: Those all-important compatibility conditions on labels are entirely local in our setting. Local compatibility at a vertex v involves the labels of the flower F_v and is expressed in terms of an angle sum θ.

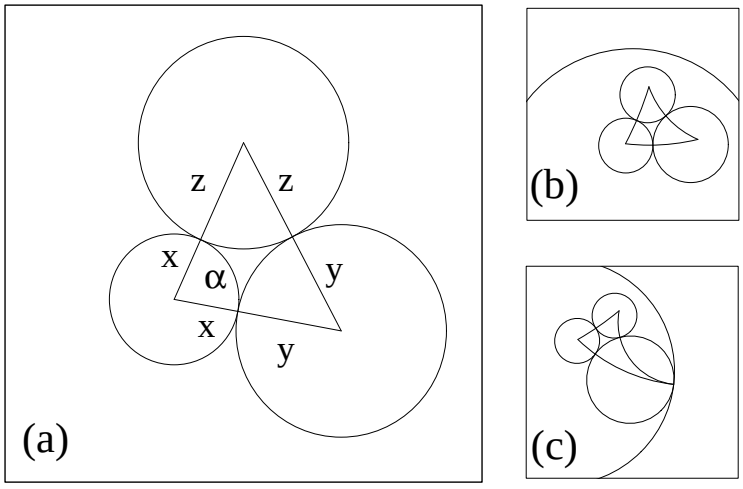

FIG. 1.1. *Triples in the euclidean and hyperbolic planes*

We describe the euclidean case first. Given labels $x, y, z \in (0, \infty)$, lay out a mutually tangent triple $\langle c_x, c_y, c_z \rangle$ of circles in the plane with radii x, y, z and connect the circle centers with geodesic segments to form a triangle T, as in Figure 1.1(a). The triangle T is unique up to rigid motions and the angle of T at the center of c_x, denoted by $\alpha(x; y, z)$, can be computed from the labels using the law of cosines:

$$\alpha(x; y, z) = \arccos\left[\frac{(x+y)^2 + (x+z)^2 - (y+z)^2}{2(x+y)(x+z)}\right]. \tag{1.1}$$

Consider a vertex v and its flower $F_v = \{v; v_1, \cdots, v_k\}$ in K. The sum of angles associated with v in the various faces $\langle v, u, w \rangle \in K$ is termed the *angle sum* at v for label R, denoted

$$\theta(v; R) = \sum_{\langle v,u,w \rangle} \alpha(R(v); R(u), R(w)),$$

where the sum is over faces $\langle v, u, w \rangle \in K$. If $\{r; r_1, \cdots, r_k\}$ denotes the labels from R for F_v, then the angle sum depends only on these labels. Assuming v is interior, it belongs to k faces and, abusing notation, we write

$$\theta(v; R) = \theta(r; r_1, \cdots, r_k) = \sum_{j=1}^{k} \alpha(r; r_j, r_{j+1}).$$

An elementary but crucial observation: *A set of circles $c_v, c_{v_1}, \cdots, c_{v_k}$ with the labels from R as radii will fit together coherently in the plane if and only if $\theta(v; R) = 2\pi n$ for some integer $n \geq 1$.* In this case, the petal circles will wrap precisely n times around c_v. A nine-petal flower is shown in Figure 1.2: in (a) the petals wrap once, $n = 1$; in (b), petals of the same radii wrap twice around the smaller center circle, $n = 2$. Angle sums are defined similarly at boundary vertices v, but since their petals are not required to form a closed chain, a coherent open flower exists irrespective of the angle sum.

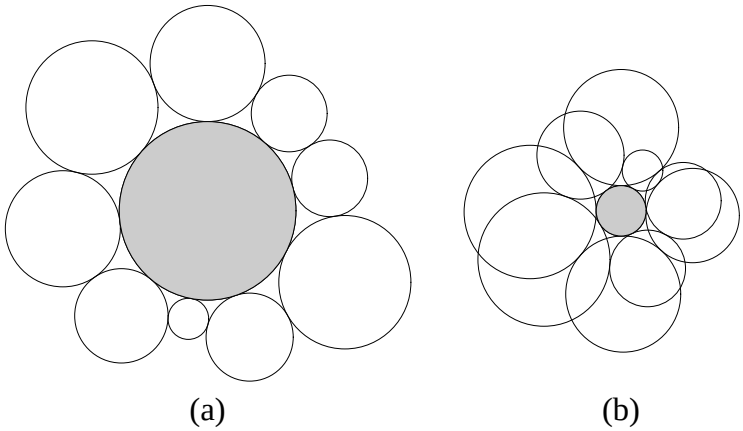

FIG. 1.2. *Nine-petal flowers*

The totality of angle sums for the vertices of a labeled complex $K(R)$ may be treated as a point in *angle space* $\mathcal{A}$. Thus θ is a map from label space to angle space, $\theta : \mathcal{R} \longrightarrow \mathcal{A}$.

Moving to hyperbolic geometry, similar local considerations apply; recall that in our disc model, hyperbolic geodesics are arcs of euclidean circles which intersect $\partial\mathbf{D}$ at right angles. See Figure 1.1(b) and (c) for sample triples, the latter having a vertex at the ideal boundary. As in the euclidean case, any triple of labels $x, y, z \in (0, \infty]$ determines a geodesic triangle T in the hyperbolic plane, unique up to rigid motions (i.e., Möbius transformations of $\mathbf{D}$), and it determines an angle $\alpha(x; y, z)$ at the center of the circle of radius x. The formula for α, to be given later, involves now the hyperbolic cosine law and must accommodate infinite radii. The angle sum $\theta(\cdot; R)$ is defined as before and has the same geometric implications.

Now for the result which motivates our computational effort:

DEFINITION 1.1. *Given a complex K, a label R is said to satisfy the* **packing condition** *at an interior vertex $v \in R$ if $\theta(v; R) = 2\pi n$ for some integer $n \geq 1$. The label R is said to be a* **packing label** *if the packing condition is satisfied at every interior vertex.*

The next theorem says that under our assumption that K triangulates a closed topological disc, the local compatibility conditions are enough to ensure a circle packing. For the proof of the theorem see [3]. (In more general multiply-connected cases, global compatibility conditions also enter.)

THEOREM 1.2. *Given a labeled complex $K(R)$, a necessary and sufficient condition for existence of a circle packing P with $P \leftrightarrow K(R)$ is that R be a packing label. In this case, P is uniquely determined up to rigid motions (isometries) of the euclidean or hyperbolic plane, as appropriate.*

The angle sum map θ is nonnegative (positive in the euclidean case) and its value at a vertex v is bounded by $\pi\deg(v)$. If v is interior, $\theta(v; R) = 2\pi n$, and $n \geq 2$, then the label is said to have a *branch point of order* $n-1$ at v; a packing label with one or more branch points is called a *branched* packing label. In computing packing labels, the branch structure (branch points and their orders) is specified in advance.

DEFINITION 1.3. *Given the complex K, an angle sum* **target** *function A assigns to each interior vertex v a value $A(v) = 2\pi n$, where $n - 1$ is the desired* **order** *of branching at vertex v. The* **default** *target is no branching, $n \equiv 1$.*

References to the target angle sums will generally be suppressed until they arise in actual computations. Dubejko [12] has established the following necessary and sufficient conditions for A: *The function A can be the target function for a circle packing of K if and only if, for any simple closed edge-path γ in K,*

$$\sum_{v \in \gamma^\circ} (A(v) - 2\pi) \leq (k - 3)\pi$$

where the sum is over vertices v interior to γ and k is the number of edges in γ. We assume henceforth that A satisfies these (purely combinatorial) conditions. Nothing will be lost if the reader assumes the default target, which is always legal.

The packings we intend to compute are guaranteed by the following fundamental existence and uniqueness result:

THEOREM 1.4. *(The Dirichlet Problem) Let K be a complex triangulating a closed topological disc, let A be an angle sum target function of K, and assume that $g : \partial K^{(0)} \longrightarrow (0, \infty)$ (resp. $(0, \infty]$) is a function defined on the boundary vertices of K. Then there exists a unique euclidean (resp. hyperbolic) packing label R for K with the property that $R(w_j) = g(w_j)$ for each boundary vertex w_j.*

More explicitly, the solution R satisfies the following nonlinear system of N equations, one for each interior vertex u_j.

$$\left\{ \sum_{\langle u_j, v, w \rangle} \alpha(R(u_j); R(v), R(w)) = A(u_j) : j = 1, \cdots, N \right\}.$$

We say that the solution label R "solves the Dirichlet problem", since the theorem statement and proof both parallel the classical Dirichlet problem for harmonic functions.

2. Local Geometry. Circles have been objects of study for well over two thousand years. The dynamics associated with small configurations of circles — triples and flowers — underly both the theoretical and practical solution of the Dirichlet problem. We refer to these lemmas collectively as "monotonicity" results.

LEMMA 2.1. *Let x, y, z denote euclidean or hyperbolic radii in the configurations of Figure 1.1. The angles α, β, and γ and Area(T) are differentiable functions of x for $0 < x < \infty$. Moreover,*

(a) α is decreasing in x,
(b) β and γ are increasing in x,
(c) Area(T) is increasing in x,
(d) $\lim_{x\to 0} \alpha(x; y, z) = \pi$, and
(e) $\lim_{x\to\infty} \alpha(x; y, z) = 0$.

Monotonicity is strict (except for the hyperbolic case of (b) when y (resp. z) is infinite).

LEMMA 2.2. *Let $F_v = \{v; v_1, \cdots v_k\}$ denote a closed flower, $\{r; r_1, \cdots, r_k\}$ the corresponding euclidean or hyperbolic labels, and $\theta(r; r_1, \cdots, r_k)$ the angle sum for v. Then θ is a differentiable function of its (finite) labels. Moreover,*

(a) θ is strictly decreasing in r,
(b) θ is strictly increasing in $r_j, j = 1, \cdots, k$,
(c) $\lim_{r\to 0} \theta = k\pi$, and
(d) $\lim_{r\to\infty} \theta = 0$.

In particular, given a with $0 < a < k\pi$, there exists a unique label $r = r_0$ so that $\theta(r_0; r_1, \cdots, r_k) = a$.

The previous results are standard in the circle packing literature (see [3]) and as we see in the next section, suggest the numerical approach to solving the Dirichlet problem. Later, Lemma 3.1 introduces a new, more subtle monotonicity, which largely accounts for the efficiency of our implementation.

3. The Packing Algorithm.

3.1. The Perron Method. The basis for packing algorithms lies with the Perron method; we describe the hyperbolic, "upper" version of Bowers [5]. We will say that a label R for K is a *superpacking label* for the boundary value problem in Theorem 1.4 if two conditions hold: (1) $R(w) \geq g(w)$ for every boundary vertex w, and (2) $\theta(v; R) \leq A(v)$, the target angle sum, for every interior vertex v. Together, these imply that the label R is too large.

The collection $\Phi \subset \mathcal{R}$ of all superpacking labels forms what is known as a *Perron family*. In particular, Φ is nonempty, since in hyperbolic geometry a label R_0 satisfying (1) and having sufficiently large interior labels will have small interior angle sums. By monotonicity, $R_1, R_2 \in \Phi \Longrightarrow \min\{R_1, R_2\} \in \Phi$. This suggests consideration of $\widehat{R} = \inf_\Phi\{R\}$. If $\widehat{R}$ is nonvanishing, continuity of angle sums with respect to

their labels easily implies that $\widehat{R}$ will lie in Φ. Monotonicity tells us that it must be a solution and elementary hyperbolic area computations give uniqueness. The argument that $\widehat{R}$ does not vanish requires a little more work, using hyperbolic areas, the Gauss-Bonet Theorem, the Euler characteristic of K, and the necessary conditions on the target function A. (The solution of euclidean boundary value problems may be inferred from the hyperbolic case because hyperbolic quantities are infinitesimally euclidean.)

3.2. The Uniform Neighbor Model (UNM). One could implement the Perron method numerically. In fact, however, the geometric stability is such that more direct relaxation methods suffice. We now describe the basic model we use in our calculations.

Focusing on the flower for v, we treat the label r as a variable, and the petal labels $r_1, \cdots, r_k$ as fixed parameters. For a given value $r = r_0$, the associated "reference" label is the number $\hat{r}$ for which the following equality holds:

$$\theta(r_0; r_1, \cdots, r_k) = \theta(r_0; \overbrace{\hat{r}, \cdots, \hat{r}}^{k}) =: \widehat{\theta}(r_0; \hat{r}). \tag{3.1}$$

In other words, laying out a flower with petal circles of the uniform radius $\hat{r}$ would yield the same angle sum as with the original petal radii $r_1, \cdots, r_k$ when the center circle has radius r_0.

LEMMA 3.1. *Let $\theta(r) = \theta(r; r_1, \cdots, r_k)$ and $\widehat{\theta}(r) = \widehat{\theta}(r; \hat{r}) = \theta(r; \hat{r}, \cdots, \hat{r})$, as above, with $\hat{r}$ chosen so that $\theta(r_0) = \widehat{\theta}(r_0)$ for some $r_0 > 0$. Assuming the labels $r_1, \cdots, r_k$ are not all equal, then*

$$\frac{d\widehat{\theta}}{dr}(r_0) < \frac{d\theta}{dr}(r_0), \tag{3.2}$$

Moreover, $\theta(r) < \widehat{\theta}(r)$ for $0 < r < r_0$ and $\theta(r) > \widehat{\theta}(r)$ for $r > r_0$.

Proof. The last inequalities follow easily from (3.2), since θ and $\widehat{\theta}$ intersect at the reference label r_0 by definition.

The proof of (3.2) is complicated first by the presence of k parameters, but more subtly by the dependence of $\widehat{\theta}$ on $r_1, \cdots, r_k$ through $\hat{r}$, which is suppressed in the notation. Our strategy is to adjust the petal labels in pairs, moving the largest and smallest towards one another in such a way that the reference label $\hat{r}$ does not change, and watching the derivative of θ. We work in the euclidean setting. Circles in the Poincarè disc model of the hyperbolic plane are also euclidean circles; a hyperbolic flower with central circle at the origin is simultaneously a euclidean flower with the same angle sum, and the euclidean radii are monotone increasing in hyperbolic radii. In other words, the hyperbolic result follows from the euclidean.

Suppose that $S < L$, where S denotes the smallest of the petal labels $r_1, \cdots, r_k$ while L denotes the largest. By monotonicity $S < \hat{r} < L$. Fixing all remaining petal labels, θ is a function of L and S, $\theta = \theta(r; L, S)$. The condition that $\hat{r}$ remain fixed is expressed by

$$\theta(r_0, L, S) = \widehat{\theta}(r_0), \tag{3.3}$$

and this defines S as a function of L by the implicit function theorem. The slope of θ at r_0 becomes a function of L, and we will show that it is increasing; that is,

$$\frac{\partial^2 \theta}{\partial r \partial L}(r_0, L, S) > 0. \tag{3.4}$$

This means, of course, that decreasing L towards $\hat{r}$ (and hence increasing S towards $\hat{r}$) makes the slope of θ smaller (i.e., more negative). Strict inequality in (3.4) means that the current L and S may be adjusted until a new pair of parameters qualifies as largest and smallest, at which point one can shift to adjusting them in turn. It is an easy argument to show that (3.2) follows.

The verification of (3.4) is rather messy; we sketch the euclidean case and leave details to the interested reader. Since we adjust only L and S, we need only consider the contributions to the full angle sum of either three or four of the faces in the flower. Suppose, for instance, that x, y are the labels of the neighbors of L and z, w are the labels of the neighbors of S. The equation (3.3) reduces (see (1.1)) to

$$\alpha(r_0; L, x) + \alpha(r_0; L, y) + \alpha(r_0; S, w) + \alpha(r_0; S, z) = \text{constant}. \tag{3.5}$$

For notational convenience define the mixed partial derivative

$$\alpha_{1,2}(x, y, z) = \frac{\partial^2 \alpha(x; y, z)}{\partial x \partial y}. \tag{3.6}$$

Inequality (3.4) is equivalent to

$$\alpha_{1,2}(r_0, L, x) + \alpha_{1,2}(r_0, L, y) + \alpha_{1,2}(r_0, S, w)\frac{dS}{dL} + \alpha_{1,2}(r_0, S, z)\frac{dS}{dL} > 0. \tag{3.7}$$

The computations are messy, but the results are all rational expressions; equation (3.5) permits computation of $\frac{dS}{dL}$. After simplification, reorganization, and cancellation of clearly positive factors, one can identify a recurring subexpression in the left side of (3.7). In particular, define the auxiliary functions

$$f(r, a, b) := \frac{(a - r)(a + b) - 2r^2}{(r + a)(r + a + b)} \quad \text{and} \quad F(r, a, b, c, d) := f(r, a, b) - f(r, d, c).$$

A simple calculation confirms the following fact:

$$\textit{If } a > d > 0,\ r > 0, \textit{ and } b, c \in [d, a], \textit{ then } F(r, a, b, c, d) > 0. \tag{3.8}$$

Inequality (3.7) is equivalent, after further simplification and judicious pairing of subexpressions, to a linear combination with nonnegative coefficients of the following four expressions.

$$F(r_0, L, y, w, S), \qquad F(r_0, L, y, z, S), \qquad F(r_0, L, x, w, S), \qquad F(r_0, L, x, z, S).$$

The positivity of (3.8) implies inequality (3.7).

There are two other situations. If L and S share a common neighbor, then simply take $y = w$ (and/or $x = z$) in the above. On the other hand, if L and S are themselves neighbors, say x, L, S, z is the order of petals, then the expressions in (3.5) and (3.7) must be adjusted accordingly. The simplifications become slightly more involved, but the subexpression f recurs and the result is a nonnegative combination of these four expressions,

$$F(r_0, L, x, L, S), \qquad F(r_0, L, x, z, S), \qquad F(r_0, L, S, L, S), \qquad F(r_0, L, S, z, S).$$

Positivity again follows from (3.8). This covers all possibilities and completes the proof. □

See Figure 3.1 for a plot of θ and $\widehat{\theta}$ as functions of r for a sample 6-flower.

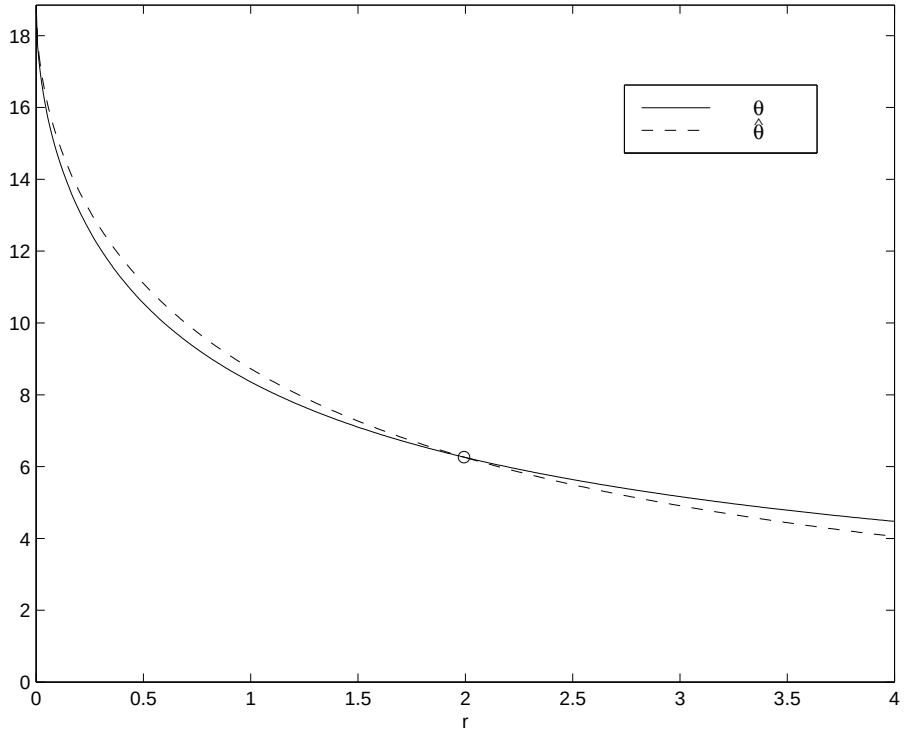

FIG. 3.1. *Angle sums of the original and reference flowers*

3.3. The Numerical Algorithm. Using the Uniform Neighbor Model, our basic algorithm generates a sequence of labels $\{R_j\}$ as follows:

1. Pick any initial label R_0, only requiring that $R(w) = g(w)$ for every $w \in \partial K$.
2. Given a label R_n, cycle through the list of interior vertices.
3. Given an interior vertex v, adjust $R_n(v)$ using the UNM by choosing $R_n(v)$ so that $\widehat{\theta}\big(R_n(v)\big) = A(v)$.
4. Denote the adjusted label by R_{n+1}; return to (2) until a prescribed accuracy in the angle sums is achieved.

Geometric facts about angle "flow" explain why this sequence of corrected labels converges so well to the packing label. Consider the euclidean setting. For label R, define "excess" e at an interior vertex v and the "total error" E by

$$e(v) = \theta(v; R) - A(v), \qquad E = E(R) = \sum_{v \text{ interior}} |e(v)|.$$

Claim: *E is monotone decreasing with our label corrections.* Let F denote the number of faces of K. Each has three angles which sum to π; reorganizing these by vertex, the *total angle* is $\sum_{v \in K} \theta(v; R) = F\pi$, *independent* of R. Thus total angle is a conserved quantity; any adjustment of a label simply causes a *redistribution* of that angle among the vertices. Suppose, for instance, that $\theta(v; R)$ is too large at some interior v, so $e(v) > 0$; by Lemma 2.2 one can increase the label $R(v)$ until $e(v) = 0$. The excess angle at v is pushed to its immediate neighbors. At worst, E

remains unchanged. However, if $u \sim v$ is an interior with angle sum too *small* or is a boundary vertex (whose angle sum doesn't count in E), then the correction to $R(v)$ simultaneously *reduces* $|e(v)|$ and $|e(u)|$, and E decreases. Similar arguments apply when $e(v) < 0$. In any case, as long as the change made to $R(v)$ does not cause the angle sum at v to *overshoot* $A(v)$, E cannot increase. Since by Lemma 3.1 corrections obtained from the UNM are conservative — they do not overshoot —- the Claim is established. Considerations are slightly altered in the hyperbolic setting because area and angle are equivalent; this actually tends to improve the performance of the algorithm. (See [32] for the hyperbolic dynamics.)

Observations: The geometry of circle configurations makes the adjustment process so stable that almost any iterative procedure will succeed. This is tempered by the essentially arbitrary combinatorics permitted in K; the local geometry is variable and there is almost no *a priori* information on the global solution. Here are some observations regarding implementation:

- The process is insensitive to the initial label: one can generally set its values (for interior vertices) arbitrarily.
- There is no advantage in careful local computation, since results will be made obsolete by subsequent adjustments.
- The process is insensitive to the order in which local adjustments are made.
- Combinatorial variability (variable degrees, lack of symmetry, etc.) complicates data manipulation, storage, and vectorization.

4. Details of the Implementation. Given complex K, boundary function g, and angle sum target A, our task is to compute the associated packing label R, as guaranteed by Theorem 1.4. Index the vertices of K by $\{w_1, \cdots, w_M; u_1, \cdots, u_N\}$, with w_i denoting boundary vertices and u_i, interior vertices. The label entries which are subject to adjustment will be termed *free*; for the Dirichlet problem, these are the N interior labels.

Problem: *Find values* $\{r_1, \cdots, r_N\}$ *so that the label vector* R *satisfies the system* $G(R) = 0$, *where* $R = \{g(w_1), \cdots, g(w_M); r_1, \cdots, r_N\}$ *and*

$$\big\{G_j(R) = \theta(u_j; R) - A(u_j),\ j = 1, \cdots, N\big\}. \tag{4.1}$$

Some abuse of notation and label transformations will be highly advantageous in describing the algorithm.

> *Notational Convention:* The same letter to be used to denote both a vertex and its current label. Moreover, in hyperbolic geometry we use transformed labels; in all calculations (and without further comment), each hyperbolic label $h \in (0, \infty]$ will be replaced by the more convenient label $s = \exp\{-2h\} \in [0, 1)$.

Keeping these conventions in mind, the angle calculation associated with vertex v for face $\langle v, u, w\rangle$ (1.1) can be rewritten in a more efficient form as

$$\text{Euclidean:} \quad \alpha(v; u, w) = 2\sin^{-1}\left(\sqrt{\frac{u}{v+u}\cdot\frac{w}{v+w}}\right) \tag{4.2}$$

$$\text{Hyperbolic:} \quad \alpha(v; u, w) = 2\sin^{-1}\left(\sqrt{v\cdot\frac{1-u}{1-vu}\cdot\frac{1-w}{1-vw}}\right). \tag{4.3}$$

In computing the angle sum $\theta(v;R)$, only the labels for v and its petals are involved, so all our packing computations are "local". We will write $\theta(v;R) = \theta(v;\{v_j\})$, where $\{v_j\}$ is shorthand for the list of petal labels. The context should make our index usage clear.

4.1. Uniform Neighbor Calculation. Using the UNM requires two steps. First, given a value for v, determine $\hat{v}$ so that $\widehat{\theta}(v;\hat{v}) = \theta(v;\{v_j\})$. Second, solve for a new value for v (call it u) so that $\widehat{\theta}(u;\hat{v}) = A(v)$. The advantage of the UNM is that these equations can be solved explicitly as follows.

Let $\theta = \theta(v;\{v_j\})$ and $A = A(v)$. From these values, compute $\beta = \sin(\frac{\theta}{2k})$ and $\delta = \sin(\frac{A}{2k})$. For the euclidean case, we have $\widehat{\theta}(v) = k\alpha(v,\hat{v},\hat{v}) = \theta$. Using the formula for α (4.2), we get $\hat{v} = \frac{\beta}{1-\beta}\, v$. Then, solving $\widehat{\theta} = A$, we get $u = \frac{1-\delta}{\delta}\, \hat{v}$. Note that since $0 \le \theta < k\pi$ then $0 \le \beta < 1$ and so since $v > 0$ then $u > 0$.

For the hyperbolic case, the computations are slightly more complicated but proceed in a similar fashion. We get $\hat{v} = \frac{\beta - \sqrt{v}}{\beta\, v - \sqrt{v}}$. If $\hat{v} < 0$ we take $\hat{v} = 0$. Also $1 - \hat{v} = \frac{\beta\,(1-v)}{\sqrt{v}\,(1-\beta\sqrt{v})}$, thus since $\beta < 1$ and $v < 1$, $\hat{v} < 1$. We compute u from $u = t^2$ where

$$t = \frac{2\delta}{\sqrt{(1-\hat{v})^2 + 4\delta^2\hat{v}} + (1-\hat{v})}.$$

It is clear that $t > 0$. To see that $t < 1$, start with $\delta < 1$ to get $(1-\hat{v})^2 + 4\delta^2 < (2\delta - (1-\hat{v}))^2$. Thus

$$t < \frac{2\delta}{|2\delta - (1-\hat{v})| + (1-\hat{v})}.$$

If $2\delta > (1-\hat{v})$ this last equation reduces to $t < \frac{2\delta}{2\delta} = 1$. If, on the other hand, $2\delta < 1 - \hat{v}$, then it reduces to

$$t < \frac{2\delta}{2(1-\hat{v}) - 2\delta} < \frac{2\delta}{4\delta - 2\delta} = 1.$$

In any event, in the hyperbolic case $0 < u < 1$, as desired.

Let this process, in either geometry, be represented by $u = M(v,\{v_j\},A)$. Then one iteration of our algorithm loops through the N free labels and updates them by

$$u_i = M(u_i,\{u_j\},A(u_i)), \quad i = 1,\ldots N.$$

Note that since the values of adjacent labels may change during the iteration, later updates will effect the angle sums for prior circles — one expects this to be an iterative process. As we will soon show, this algorithm is locally linearly convergent. We will improve convergence overall by using this local convergence to create heuristics for global over-relaxation procedures (see §4.3).

4.2. Local Linear Convergence. Given petal labels $\{v_j\}$, let $\overline{v}$ denote the quantity of real interest; namely, the solution of $\theta(r;\{v_j\}) = A$. Lemma 3.1 implies that the computed value u lies between the current value v and $\overline{v}$:

$$\overline{v} \le u \le v \quad \text{or} \quad v \le u \le \overline{v}.$$

Thus, replacing label v by u is always a conservative improvement. We would like to see how much better $u = u(v)$ is than v, so we look at the ratio

$$\frac{u(v) - \overline{v}}{v - \overline{v}}.$$

In particular we are interested in the maximum value of this ratio over the admissible range for v and also the value as v approaches $\overline{v}$. From some simple calculations, we see that for the euclidean case, the maximum occurs at $v = 0$, and for the hyperbolic case, at $v = 1$ (hyperbolic radius $= 0$).

To understand the behavior of this ratio in the euclidean case, we need some information about the angle sum function. Consider θ as a function of v alone, with petal labels $\{v_j\}$ as fixed parameters. Then, for v near 0, we have

$$\theta(v) \approx k\pi - 2\sqrt{v}S, \quad \text{where} \quad S = \sum_{j=1}^{k} \sqrt{\tfrac{1}{v_j} + \tfrac{1}{v_{j+1}}}.$$

Also, $\theta'(v) \approx -\frac{1}{\sqrt{v}}S$. Next,

$$\lim_{v\to 0} u(v) = \frac{1-\delta}{\delta} \lim_{v\to 0} \frac{\beta v}{1-\beta} = -\frac{1-\delta}{\delta} \lim_{v\to 0} \frac{1}{\frac{d\beta}{dv}}$$

using L'Hôpital's rule (since $\beta_{v=0} = 1$). To evaluate this, we have

$$\frac{d\beta}{dv} = \frac{1}{2k}\cos(\tfrac{\theta}{2k})\theta'(v).$$

Thus for v near 0, we have

$$\frac{d\beta}{dv} \approx -\frac{1}{2k} \cdot \frac{1}{k}\sqrt{v}S \cdot \frac{1}{\sqrt{v}}S = -\frac{S^2}{2k^2} \Longrightarrow \lim_{v\to 0} u(v) = \frac{1-\delta}{\delta}\frac{2k^2}{S^2} = u_0.$$

And thus we have

$$\sup_{v>0} \frac{u(v) - \overline{v}}{v - \overline{v}} = \lim_{v\to 0^+} \frac{u(v) - \overline{v}}{v - \overline{v}} = 1 - \frac{u_0}{\bar{v}} < 1.$$

This follows from the above calculations and the fact that $0 < u_0 < \overline{v}$.

For the hyperbolic case, the results are similar in that there is an explicit expression for the maximum of this ratio in terms of $\{v_j\}$ and $\overline{v}$ and it is clearly less than 1. These results show that locally the convergence is at worst linear.

4.3. Acceleration and Final Algorithm. Let $\bar{R}$ be the exact solution and let R_l and R_{l+1} be consecutive approximations. For large l computational experiments have shown that the local linear convergence discussed in the previous section is uniform, i.e.

$$R_{l+1} - \bar{R} \approx \lambda(R_l - \bar{R}),$$

holds element-by-element for some $\lambda < 1$. Assuming this result holds exactly, we get two heuristics for accelerating the convergence of this process. Taking this result for l and $l+1$, we can solve for R_{l+2} and $\bar{R}$ to get

$$R_{l+2} = R_{l+1} + \lambda(R_{l+1} - R_l) \tag{4.4}$$

and

$$\bar{R} = R_{l+1} + \frac{\lambda}{1-\lambda}(R_{l+1} - R_l). \tag{4.5}$$

When we use the equation for R_{l+2} (4.4) to replace R_{l+1} we call it *simple* acceleration. When we use the formula for $\bar{R}$ (4.5) to replace R_{l+1} we call it *super* acceleration.

Given:

- complex K
- boundary function g
- legal target function A
- initial label R
- tolerances $\epsilon > 0$ and $\delta > 0$.

Algorithm:

1. Set the boundary labels of R to their g values.
2. Initialize: $c = \epsilon + 1$, $\lambda = -1$, $flag = 0$
3. While $(c > \epsilon)$ do
   (a) $c_0 = c$, $\lambda_0 = \lambda$, $flag_0 = flag$, $R_0 = R$
   (b) For each free node u_j
      i. Calculate the angle sum θ_j
      ii. Update $u_j = M(u_j, \{u_i\}, A_j)$
      iii. Accumulate error estimate $c = c + (\theta_j - A_j)^2$
   (c) $c = \sqrt{c}$, $\lambda = c/c_0$, $flag = 1$
   (d) If $(flag_0 = 1)$ and $\lambda < 1$ then *perform super acceleration*
      i. $c = \lambda c$
      ii. If $|\lambda - \lambda_0| < \delta$ then $\lambda = \lambda/(1-\lambda)$
      iii. Determine largest λ^* s.t. $R + \lambda^*(R - R^0)$ in range
      iv. $\lambda = \min(\lambda, 0.5\lambda^*)$
      v. $R = R + \lambda(R - R^0)$
      vi. $flag = 0$

TABLE 4.1
Circle Packing Algorithm

To use these acceleration schemes we need an estimate of the error reduction factor λ. In practice, we use a ratio of values c_{l+1}/c_l where c_l is an approximation of $\|G(R_l)\|_2$. In addition, since we are using an estimate, we choose between the two different acceleration steps based on whether or not the ratio c_{l+1}/c_l is converging to a constant value. We also modify the size of the acceleration factor in order to assure that the new value is in a valid range for radius labels ($v_i > 0$ for euclidean and $0 < v_i < 1$ for hyperbolic). For this, we determine the largest λ_M for which all components of

$$R_{l+1} + \lambda_M(R_{l+1} - R_l)$$

remain valid. We use $\frac{1}{2}\lambda_M$ as an upper bound for the acceleration factor. Thus we are assured that we always stay in the proper range and when the situation presents itself, the iterates converge as quickly as possible. These modifications for acceleration and other features are incorporated into the Circle Packing Algorithm given in Table 4.1.

4.4. Results. Having considered certain theoretical properties of our algorithm, let us examine some numerical results.

Data File	eucl/hyp	N	Acceleration	Iters	Flops	Error
spiral78	eucl	50	none	16	48,040	2.8549E-5
spiral78	eucl	50	simple	14	43,590	2.4398E-5
spiral78	eucl	50	super	12	37,370	2.7130E-5
data237	eucl	223	none	297	3,924,600	4.6165E-5
data237	eucl	223	simple	199	2,740,200	4.5294E-5
data237	eucl	223	super	46	632,900	2.8921E-5
data237	hyp	223	none	220	4,919,200	4.4722E-5
data237	hyp	223	simple	147	3,377,400	4.5010E-5
data237	hyp	223	super	36	808,700	3.2906E-5

TABLE 4.2
Tests of acceleration

Acceleration: First, consider the effect of using the acceleration. In Table 4.2 we show the performance for sample runs with no acceleration, simple acceleration, and super-step acceleration. (Super-stepping occurs if the tests for step 3(d)(ii) in the meta-code given in the box are satisfied.) We report the number of iterations, the flop count (an estimate of the number of floating point operations), and the error of the final value. The error is computed as $\frac{\|G(R)\|_1}{3N}$, G as defined in (4.1).

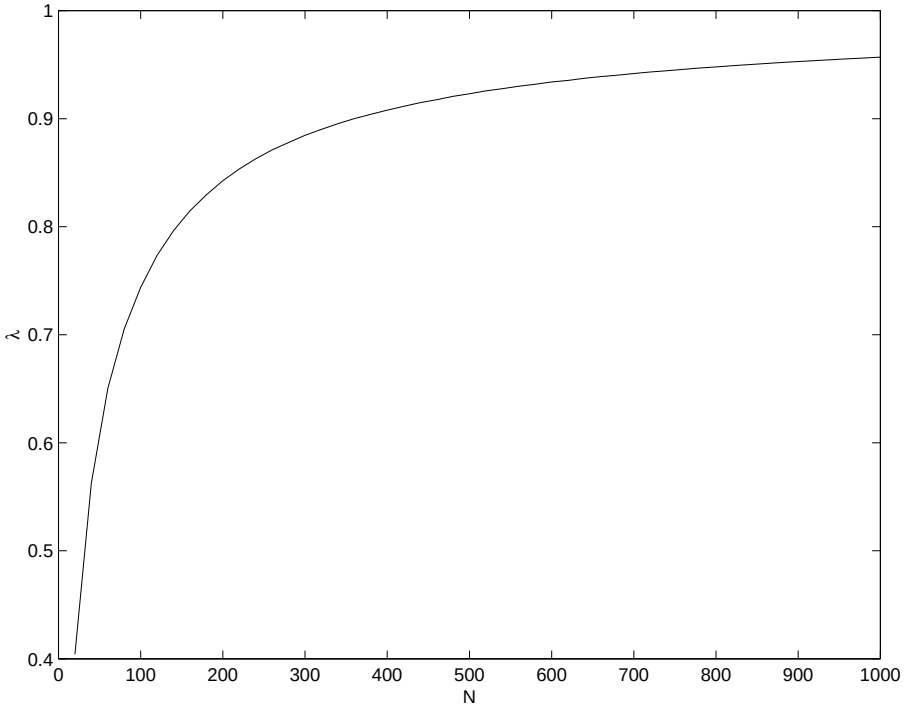

FIG. 4.1. *Experimental estimates of* λ

Size: The data suggest a clear advantage using acceleration, one which grows with size. Therefore, we next look at how the algorithm performs as N increases. We used a simple family of complexes which are built by spiraling out from the center point to include the desired number of interior circles. In the hyperbolic runs, we set the boundary labels to infinity and started the interior labels with small (< 0.1) values. Figure 4.1 is a graph of the asymptotic error reduction factor λ (without acceleration) plotted against the number of interior vertices N. This curve is approximately $\lambda = \frac{N}{N+C}$, for $C \approx 30$. Thus for large N the basic iteration converges slowly. However, with this same setup, we have compared the number of iterations needed to reach a tolerance of 10^{-8} using only the simple acceleration *versus* using the super acceleration.

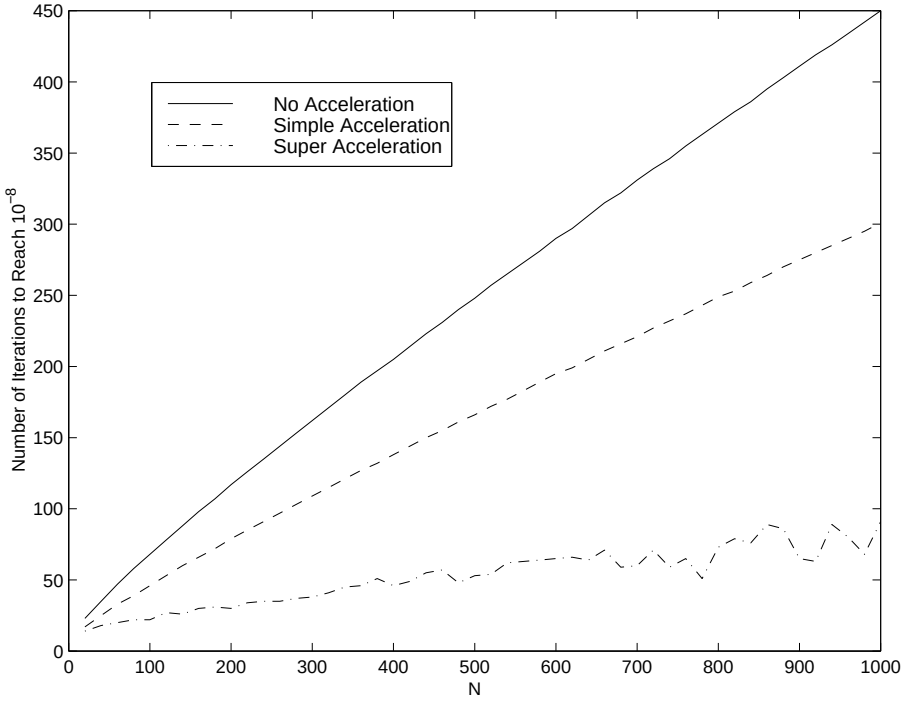

FIG. 4.2. *Effects of acceleration*

Ordering: To test the effect of the order in which free vertices $\{u_j\}$ are adjusted, we ran trials with a fixed hexagonal complex with 100 interior vertices. We set boundary labels to 2 and initial interior labels to 1; in the exact solution, all labels are 2. We randomized the indexing, recorded the number I of iterations it took to converge, noted the final error, and computed the approximate convergence factor λ by $\lambda^I =$(Error). In 41 random trials λ ranged from 0.5462 to 0.6094, I from 18 to 21. This was a limited test, but is in line with our experience that the ordering of vertices has a limited impact on the computations.

Initial Label: To test the effects of the initial label on performance, we ran trials with the same complex, but in the hyperbolic setting. Boundary labels were random but fixed, ending up in the range $[0.68, 0.86]$. Table 4.3 summarizes the results

for various initial values for the interior labels. λ is the effective convergence rate, computed so that (Starting Error)$(\lambda)^I$ = (Final Error), where I is the number of iterations. (Here $U(a,b)$ indicates uniformly random variables in the ranges (a,b) and R is the exact solution). Note that even though the local method is conservative (i.e. always produces values on the same side of the correct answer as this initial guess), the super-step acceleration corrupts this property.

Initialization	Iters	Starting Error	Final Error	λ
$U(0,1)$	35	39.601	1.7274E-6	0.6162
$U(0,0.5)$	32	43.520	6.8814E-7	0.5705
$U(0.5,1)$	35	27.540	1.6074E-6	0.6213
$\approx R$	35	8.4378	1.2559E-6	0.6381
$> R$	29	5.9513	2.5086E-6	0.6028
$< R$	31	4.3359	1.6615E-6	0.6209
0.86	34	10.094	2.1669E-6	0.6366
0.68	31	5.9867	2.6891E-6	0.6241

TABLE 4.3
Effects of starting labels

5. General Packing Problems. We have described our algorithm for the simplest Dirichlet problem on simply connected complexes, but it applies much more broadly. Here we discuss more general settings, open questions, computationally intensive applications, and software.

Combinatorics: In fact, the only requirement on the complex K is that it triangulate an oriented topological surface. Thus, it may be finite or infinite, with or without boundary, planar or nonplanar, and simply or multiply-connected — we are guaranteed that there exists one or more circle packings for K. Paraphrasing the central existence result (see [2]):

Fact: *Given K, there exists a Riemann surface S and a univalent circle packing $\mathcal{P}_K$ in S with the combinatorics of K; S and $\mathcal{P}_K$ are unique up to conformal isometries.*

The packing $\mathcal{P}_K$ satisfies certain extremal conditions and is called the *maximal packing* for K. Note that K "chooses" the geometry in which its maximal packing must live. Figure 5.1 illustrates several examples. In computational terms:

(a) When K is a closed topological disc, the maximal packing lies in $\mathbf{D}$ and is computed by solving the Dirichlet problem (Theorem 1.4) with infinite boundary labels (and default target). The maximal packing for the complex of Figure 0.1(a) is shown in Figure 5.1(a).

(b) When K is the Riemann sphere, $\mathbf{S}^2$ (the unit sphere in $\mathbf{R}^3$), then one vertex is removed, the reduced complex is packed in $\mathbf{D}$ as in (a), the missing vertex is identified with the exterior of $\mathbf{D}$, and the results are projected back to $\mathbf{S}^2$, where a normalizing Möbius transformation may be applied. Figure 5.1(b) is an example with combinatorics dual to Buchminsterfullerene.

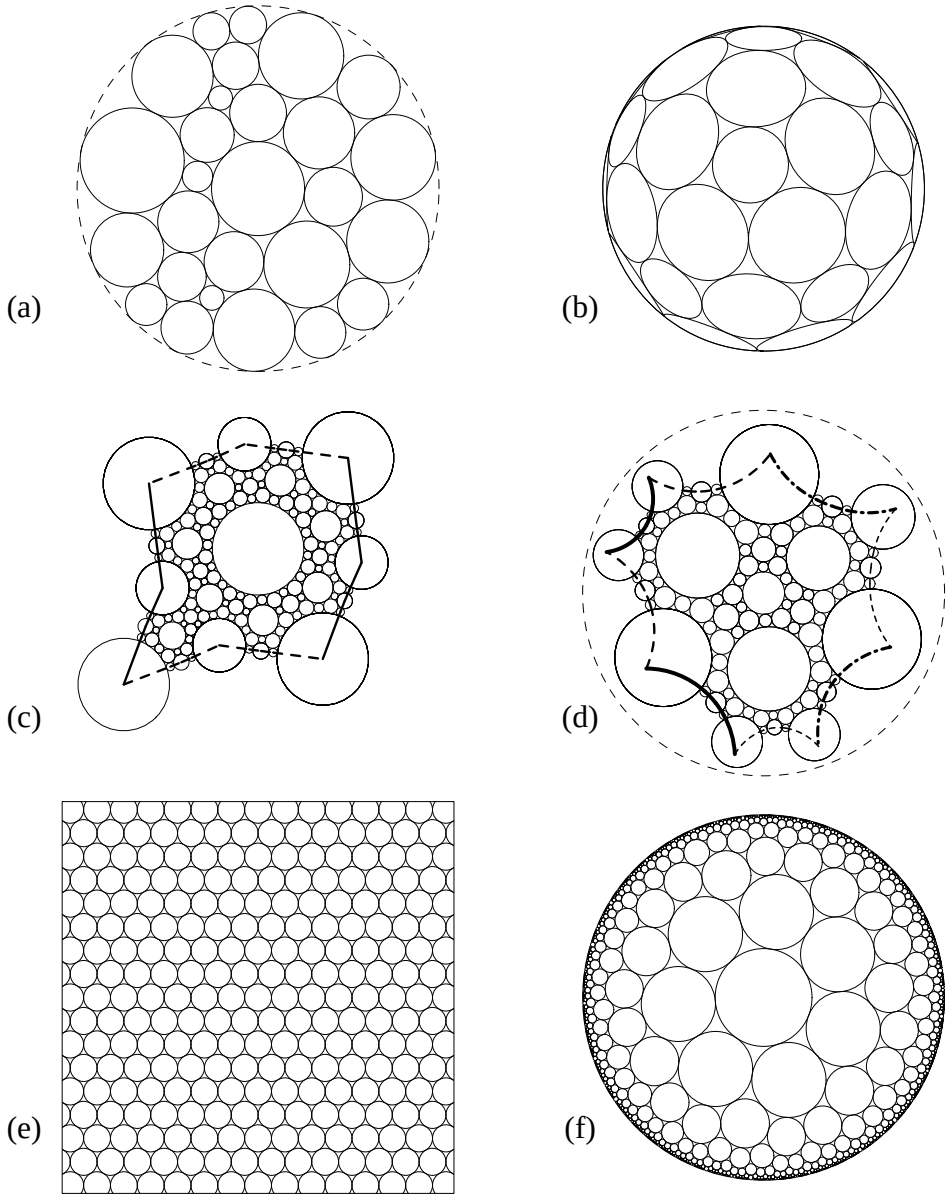

FIG. 5.1. *Maximal Packings*

(c) Suppose K triangulates a compact surface (hence is finite with no boundary) of positive genus g. Starting with any initial label, our algorithm generates labels converging (generally, quite rapidly) to a maximal packing label. When $g = 1$, the computations are necessarily euclidean and the result is unique up to scaling; when $g > 1$ the computations are hyperbolic and the result is unique. A fundamental domain for a covering of the packing can be displayed in $\mathbf{C}$ or $\mathbf{D}$, respectively. Figure 5.1(c) is a 1-torus and Figure 5.1(d) is a 2-torus; edges have been marked to show edge identifications.

(d) Suppose K is infinite and simply connected. Either K is *parabolic*, meaning $\mathcal{P}_K$ packs $\mathbf{C}$, as with the hexagonal "penny" packing of Figure 5.1(e), or K is *hyperbolic*, meaning $\mathcal{P}_K$ packs $\mathbf{D}$, as with the constant 7-degree packing of Figure 5.1(f). Computationally, these are approximated by appropriately normalized solutions of Dirichlet problems for finite, simply connected sub-complexes exhausting K. When given a nonsimply connected K, one works instead with its universal covering complex $\widetilde{K}$, which is infinite and simply connected.

Geometry: An important feature of circle packing is that the *combinatorics* of K largely determine the *geometry* in which its packings live. The fundamental dichotomy expressed in (d) above provides a striking example, since it can be shown that a *parabolic* complex has no packing in $\mathbf{D}$ (the Discrete Liouville Theorem, [15]). When K is multiply connected, the topology determines the appropriate geometry; namely, the intrinsic spherical, euclidean, or hyperbolic metric inherited from the universal covering surface. (Note in particular that all numerical computations involve the familiar metric quantities.) Thus, in (c) above, if $\text{genus}(K) = 1$, then the appropriate geometry is euclidean — the hyperbolic packing algorithm will diverge to zero. Conversely, when $\text{genus}(K) > 1$, hyperbolic geometry applies and the euclidean packing algorithm will fail.

The only influence one has on the geometry of a circle packing is through choice of boundary conditions (if there is a boundary) and through branching, which is tightly mediated by Gauss-Bonet, Riemann-Hurwitz, Euler characteristic, and other classical relations. In our experience, the packing algorithm always degenerates if the geometry being used is theoretically incompatible with the given complex K or with the boundary and branching conditions prescribed for K.

Boundary Angle Sums: Boundary angle sums can be specified in place of boundary labels. Tight compatibility conditions (involving geometry, combinatorics, and branching) have not yet been formulated, but our packing algorithm appears to work without change for legal prescriptions — simply set boundary angle sum targets, declare boundary labels as free, and run the algorithm. Figure 5.2(a) illustrates a mixed problem: some boundary vertices were given specified labels, while others were given specified angle sum targets.

Overlap Packings: The fundamental existence result for circle packings (see the Andreev-Thurston Theorem of [26]) actually applies to *overlapping packings*, of which our tangency patterns are a special case. One is allowed to specify an overlap angle $\phi(u,v) \in [0, \pi/2]$ for each edge $\langle u, v\rangle$ of K; in the associated packing P, the circles c_u and c_v will overlap (i.e., intersect) with angle $\phi(u,v)$ ($\phi(u,v) = 0$ means tangency).

Our algorithm requires only an adjustment in the computation of α. To illustrate in the euclidean case, suppose triangle T has labels x, y, z and overlaps ϕ_x, ϕ_y, ϕ_z (for opposite edges). Defining parameters $\eta_x, \eta_y, \eta_z \in [0,1]$, where $\eta_\cdot = \cos(\phi_\cdot)$, formula (1.1) becomes

$$\alpha(v; u, w, \eta_v, \eta_u, \eta_w) = \arccos\left(\frac{(x^2+y^2+2xy\eta_z)+(x^2+z^2+2xz\eta_y)-(y^2+z^2+2yz\eta_x)}{2\sqrt{x^2+y^2+2xy\eta_z}\sqrt{x^2+z^2+2xz\eta_y}}\right). \tag{5.1}$$

The monotonicity results of §2 continue to hold (see, e.g., [36]), so Thurston's itera-

tive algorithm yields packing labels as in the tangency case. In fact, the process is sufficiently robust that the uniform neighbor model works despite the fact that it is no longer strictly applicable.

Yet more general overlap situations are of theoretical and practical interest, but open the door to incompatibilities. See [20] for the most general existence and uniqueness statement for overlaps up to angle π. In another direction, "imaginary" values for overlap angles correspond geometrically to *inversive distances*, a classical, conformally invariant way to measure the "distance" between pairs of separated circles. The continuum of situations — from overlaps of angle $\pi/2$, through tangency, out to inversive distances approaching infinity — is accommodated in formula (5.1) by letting the η-parameters vary over $[0, \infty)$. The four packings of Figure 5.2 satisfy the same mixed boundary label/angle sum conditions but show a variety of "overlap" prescriptions. Figure 5.2(a) is the familiar tangency case; Figure 5.2(b) has constant inversive distances, all η set to 2.0; Figure 5.2(c) has all overlaps set to $\pi/3$; and Figure 5.2(d) involves a mixture of overlaps and inversive distances. There is yet little theoretical work on inversive distance packings, but our algorithm handled these without complaint. Improper specifications tend to show up in labels that degenerate during repacking.

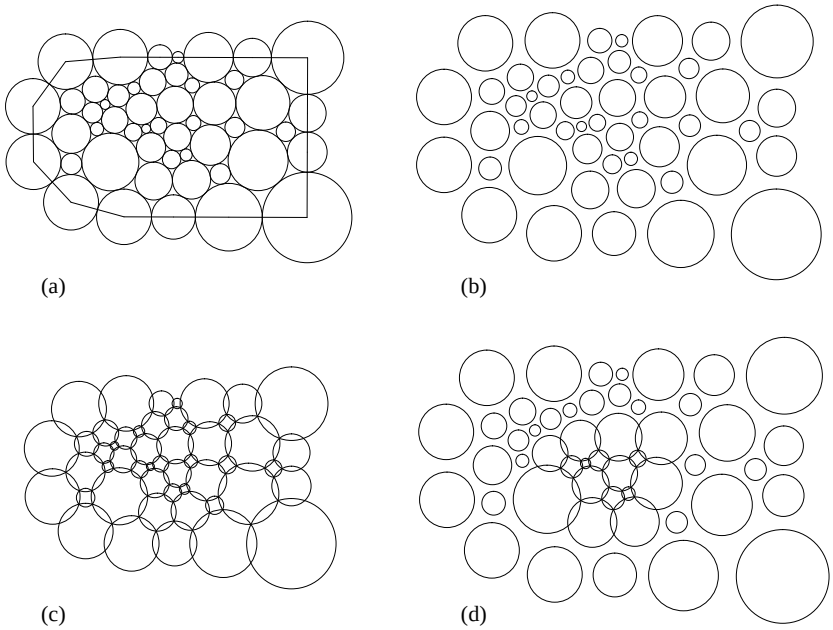

FIG. 5.2. *Mixed boundary value, overlap, and inversive distance examples*

Spherical Geometry: The sphere is the most rigid and difficult classical setting. To the authors' knowledge, no packing algorithm intrinsic to the geometry has been found; spherical packings are typically obtained by stereographically projecting from the disc. However, fundamental existence and uniqueness results for branched packings (i.e., discrete rational functions, see [8]) and even simple Dirichlet problems cannot be handled by projection and remain open.

Applications: In the numerical conformal mapping of plane regions, it is unlikely that circle packing can ever compete in speed or accuracy with classical numerical methods such as Schwarz-Christoffel. However, circle packing techniques are finding new applications in a number of more general conformal situations; see the survey [33]. We illustrate three for which no other methods are known. These are computationally intensive and happen to be of interest to mathematicians, physicists, and neuroscientists; they provided much of the motivation for our algorithm improvements. We do not mention other potentially valuable applications, such as graph embedding.

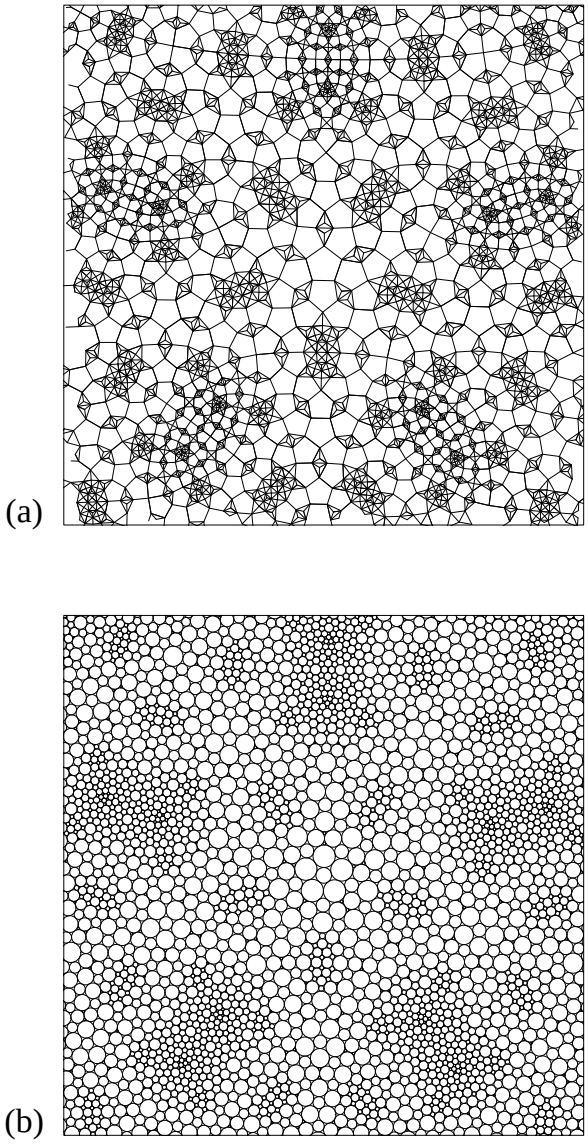

FIG. 5.3. *A sample conformal tiling*

Tilings: Figure 5.3(a) approximates a (finite piece of) a "conformal" tiling of **C**. In this theory, the tile shapes are determined purely by the abstract adjacency graph of the global pattern. Such a graph can be augmented to give a complex K which is then circle packed to provide approximations of the tiling; a refinement process and associated packings lead to more accurate shapes. We refer the reader to Bowers and Stephenson [6] for details. The circle pattern underlying the tiling is Figure 5.3(b). This example (indeed, this whole topic) was motivated by work of Cannon, Floyd, and Parry; see [22]. Our thanks to Bill Floyd, whose software created the underlying complex as input for our packing routine.

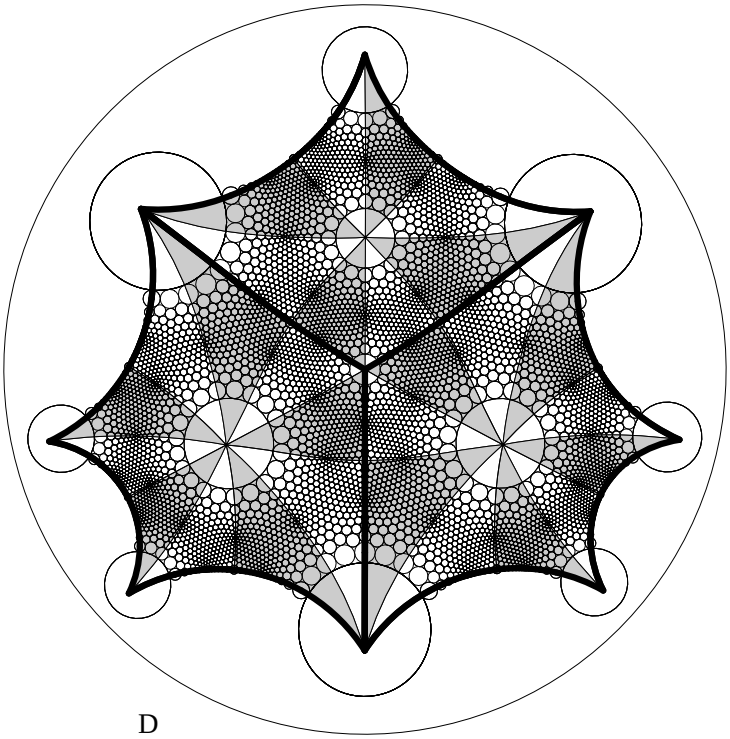

FIG. 5.4. *A stage-3, genus 2 dessin*

Dessins: In the theory of *dessins d'enfants* of Grothendieck, drawings on surfaces lead to algebraic number fields *via* triangulations and associated conformal structures (see [30]). In [6], Bowers and Stephenson develop circle packings techniques to provide both discrete parallels and approximations to these conformal structures. Figure 5.1(d) comes from a genus 2 dessin at a "coarse" stage. Figure 5.4 is a more accurate stage-3 refinement for the same dessin obtained with a "hex" refinement process. Each refinement stage roughly triples the number of vertices, so the ability to handle large packings becomes important quickly. In this setting of triangulated surfaces, particularly, there appears to be a potential for significant vectorization and parallelization of our algorithm.

Brain-mapping: A use for circle packings which is just emerging and placing new demands for speed and flexibility concerns the "flattening" of images of the human brain for use in neuroscience research. The cortical surfaces of the brain, the cerebellum and the hemispheres of the cerebrum, are essentially highly convoluted topological 2-spheres embedded in 3-space. Flat representations are needed in structural and functional studies of the cortex for purposes such as registration, visualization, and statistical data collection; conformal flattening is emerging as the preferred method because it preserves valuable geometric information. Various medical imaging technologies, such as PET, MRI, and fMRI, provide 3-dimensional representations from which the cortical surfaces can be extracted as triangulated topological spheres or discs. Circle packing is then a means for approximating the conformal maps of these surfaces and manipulating the resulting images. See [21] and references therein; our thanks to the authors for the examples shown in Figure 5.5 and Figure 5.6. These are grayscale images of color coded flat maps in the hyperbolic and spherical settings. These circle packings each involve roughly 50,000 circles; the spherical packing is computed in hyperbolic geometry and projected stereographically to the sphere.

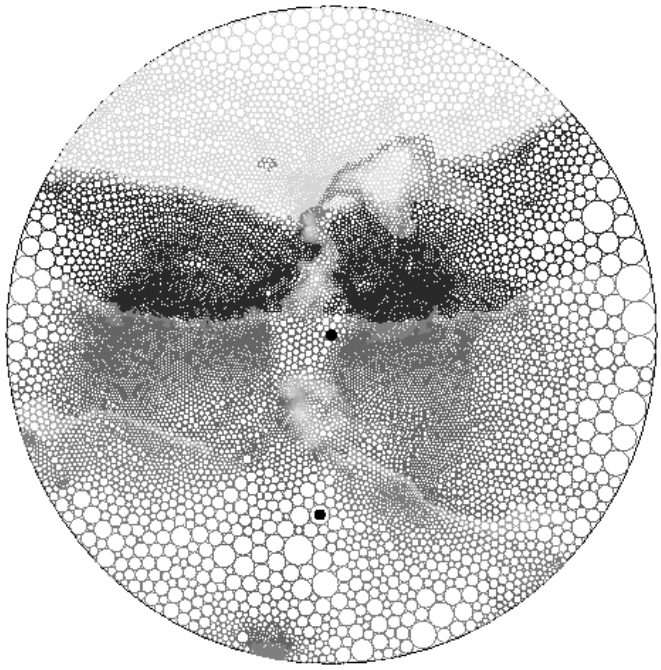

FIG. 5.5. *A hyperbolic flat map*

Software: Our circle packing algorithm is implemented in C in the standalone program `RePack` and as the compute engine behind the second author's graphical software package `CirclePack`. The software is available at `www.math.utk.edu/~kens`. With `CirclePack` the user can create, manipulate, display, and print circle packings. Functionality is provided for all the operations we have discussed — Dirichlet, boundary angle sums, overlaps, compact complexes — plus many more.

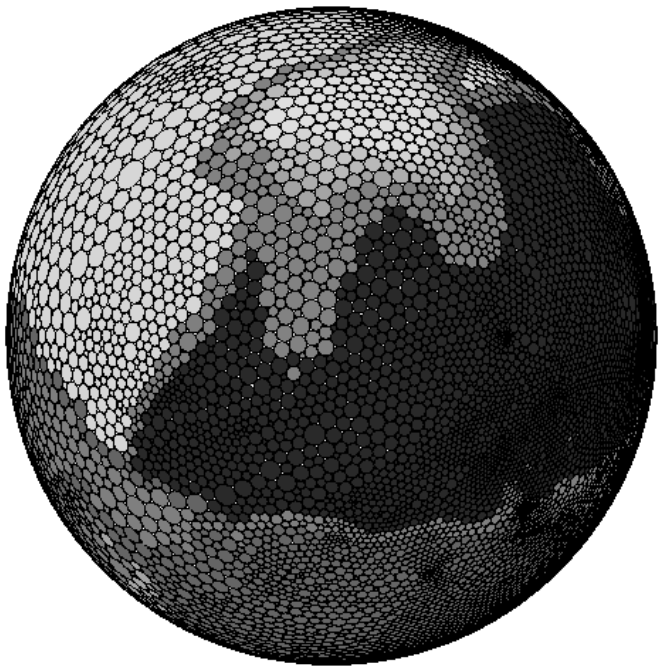

FIG. 5.6. *A spherical flat map*

REFERENCES

[1] R. BARNARD AND G. B. WILLIAMS, *Combinatorial excursions in moduli space*, Pacific J. Math. to appear.

[2] A. F. BEARDON AND K. STEPHENSON, *The uniformization theorem for circle packings*, Indiana Univ. Math. J., 39 (1990), pp. 1383–1425.

[3] ———, *The Schwarz-Pick lemma for circle packings*, Ill. J. Math., 35 (1991), pp. 577–606.

[4] I. BENJAMINI AND O. SCHRAMM, *Harmonic functions on planar and almost planar graphs and manifolds, via circle packings*, Invent. Math., 126 (1996), pp. 565–587.

[5] P. L. BOWERS, *The upper Perron method for labelled complexes with applications to circle packings*, Proc. Camb. Phil. Soc., 114 (1993), pp. 321–345.

[6] P. L. BOWERS AND K. STEPHENSON, *Uniformizing dessins and Belyĭ maps via circle packing*, preprint.

[7] ———, *Circle packings in surfaces of finite type: An in situ approach with applications to moduli*, Topology, 32 (1993), pp. 157–183.

[8] ———, *A branched Andreev-Thurston theorem for circle packings of the sphere*, Proc. London Math. Soc. (3), 73 (1996), pp. 185–215.

[9] ———, *A "regular" pentagonal tiling of the plane*, Conformal Geometry and Dynamics, 1 (1997), pp. 58–86.

[10] R. BROOKS, *Circle packings and co-compact extensions of Kleinian groups*, Inventiones Mathematicae, 86 (1986), pp. 461–469.

[11] Y. C. DE VERDIÈRE, *Une principe variationnel pour les empilements de cercles*, Inventiones Mathematicae, 104 (1991), pp. 655–669.

[12] T. DUBEJKO, *Branched circle packings, discrete complex polynomials, and the approximation of analytic functions*, PhD thesis, University of Tennessee (advisor: Ken Stephenson), 1993.

[13] ———, *Branched circle packings and discrete Blaschke products*, Trans. Amer. Math. Soc., 347 (1995), pp. 4073–4103.

[14] T. DUBEJKO AND K. STEPHENSON, *The branched Schwarz lemma: a classical result via circle packing*, Mich. Math. J., 42 (1995), pp. 211–234.

[15] ———, *Circle packing: Experiments in discrete analytic function theory*, Experimental Mathematics, 4 (1995), pp. 307–348.

[16] Z.-X. He and B. Rodin, *Convergence of circle packings of finite valence to Riemann mappings*, Comm. in Analysis and Geometry, 1 (1993), pp. 31–41.
[17] Z.-X. He and O. Schramm, *Inverse Riemann mapping theorem for relative circle domains*, Pac. J. Math, 171 (1995), pp. 157–165.
[18] ———, *On the convergence of circle packings to the Riemann map*, Invent. Math., 125 (1996), pp. 285–305.
[19] ———, *The c^∞-convergence of hexagonal disk packings to the Riemann map*, Acta Math., 180 (1998), pp. 219–245.
[20] C. D. Hodgson and I. Rivin, *A characterization of compact convex polyhedra in hyperbolic 3-space*, Invent. Math., 111 (1993), pp. 77–111.
[21] 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 Medical Image Computing and Computer-Assisted Intervention - MICCAI'99, C. Taylor and A. Colchester, eds., vol. 1679, Springer, Berlin, 1999, pp. 279–286.
[22] J. W. Cannon, W. J. Floyd, and Walter Parry, *Finite subdivision rules*, Conformal Geometry and Dynamics, 5 (2001), pp. 153–196.
[23] Koebe, *Kontaktprobleme der Konformen Abbildung*, Ber. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Kl., 88 (1936), pp. 141–164.
[24] G. L. Miller and W. Thurston, *Separators in two and three dimensions*, in Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, ACM, Baltimore, May 1990, pp. 300–309.
[25] B. Mohar, *A polynomial time circle packing algorithm*, Discrete Math., 117 (1993), pp. 257–263.
[26] J. W. Morgan, *On Thurston's uniformization theorem for three-dimensional manifolds*, in The Smith Conjecture, H. Bass and J. W. Morgan, eds., Academic Press, New York, 1984, pp. 37–126.
[27] B. Rodin, *Schwarz's lemma for circle packings*, Invent. Math., 89 (1987), pp. 271–289.
[28] ———, *Schwarz's lemma for circle packings II*, J. Differential Geometry, 30 (1989), pp. 539–554.
[29] B. Rodin and D. Sullivan, *The convergence of circle packings to the Riemann mapping*, J. Differential Geometry, 26 (1987), pp. 349–360.
[30] L. Schneps, ed., *The Grothendieck Theory of Dessins d'Enfants*, vol. 200 of London Math. Soc. Lecture Note Series, Cambridge Univ. Press, Cambridge, 1994.
[31] K. Stephenson, *Circle packing and discrete analytic function theory*, in Handbook of Complex Analysis, Vol. 1, R. Kühnau, ed., Elsevier. to appear.
[32] ———, *A probabilistic proof of Thurston's conjecture on circle packings*, Rendiconti del Seminario Mate. e Fisico di Milano, LXVI (1996), pp. 201–291.
[33] ———, *Approximation of conformal structures via circle packing*, in Computational Methods and Function Theory 1997, Proceedings of the Third CMFT Conference, N. Papamichael, S. Ruscheweyh, and E. B. Saff, eds., vol. 11, World Scientific, 1999, pp. 551–582.
[34] W. Thurston, *The Geometry and Topology of 3-Manifolds*, Princeton University Notes. preprint.
[35] ———, *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.
[36] J. van Eeuwen, *The discrete Schwarz-Pick lemma for overlapping circles*, PAMS, 121 (1994), pp. 1087–1091.
[37] G. B. Williams, *Discrete conformal welding*, PhD thesis, University of Tennessee, Knoxville, 1999. (advisor: Ken Stephenson).
[38] ———, *Earthquakes and circle packings*, J. Anal. Math., 85 (2001), pp. 371–396.