A BRANCHED ANDREEV-THURSTON THEOREM FOR CIRCLE PACKINGS OF THE SPHERE

Phil Bowers and Kenneth Stephenson

Florida State University and the University of Tennessee

Introduction

Our interest is in the existence and uniqueness of collections of circles on the Riemann sphere with prescribed combinatoric patterns, prescribed angles of overlap, and prescribed branch sets. The central result in this topic is the following, commonly referred to now as the Andreev-Thurston Theorem:

Andreev-Thurston Theorem. *Let τ be a triangulation of S^2 which is not simplicially equivalent to a tetrahedron and let $\Phi : \tau^{(1)} \longrightarrow [0, \pi/2]$ be given, where $\tau^{(1)}$ denotes the set of edges of τ. Suppose that*

> *(A). If $e_1, e_2, e_3 \in \tau^{(1)}$ form a closed loop on S^2 and if $\sum_{i=1}^{3} \Phi(e_i) \geq \pi$, then (e_1, e_2, e_3) are the three edges of a two-simplex of τ.*
>
> *(B). If $e_1, e_2, e_3, e_4 \in \tau^{(1)}$ form a closed loop on S^2 and if $\sum_{i=1}^{4} \Phi(e_i) \geq 2\pi$, then (e_1, e_2, e_3, e_4) form the boundary of the union of two adjacent two-simplexes of τ.*

Then there is a realization of τ as a geodesic triangulation and a family P of circles centered at the vertices of this triangulation so that the two circles centered at the vertices of an edge $e \in \tau^{(1)}$ meet at angle $\Phi(e)$. This geodesic triangulation and family of circles is unique up to conformal automorphisms of S^2.

This statement provides the framework for our work: combinatorics of circle patterns are coded in triangulations τ of the sphere and "overlap angles" (or just "overlaps") are coded in functions Φ defined on $\tau^{(1)}$; we call the pair (τ, Φ) a **labeled triangulation**. Our interest is in finding **circle packings** for (τ, Φ); that is, collections $P = \{C_v : v \in \tau^{(0)}\}$ of

---

The second author gratefully acknowledges support of the National Science Foundation and the Tennessee Science Alliance.

Typeset by $\mathcal{AMS}$-TEX

circles in S^2 so that whenever $e = \langle v, w \rangle$ is an edge of τ, circles C_v and C_w overlap with angle $\Phi(e)$.

The circle packing P guaranteed by the Andreev-Thurston Theorem will be called the **maximal** circle packing for (τ, Φ). An example in which all overlaps are set to $\pi/2$ is shown in Figure 1. It might appear from the statement that P is (essentially) unique; however, this is contingent on the fact that it induces a geodesic triangulation of S^2. Removing this requirement, one can well expect additional circle packings P for (τ, Φ). In particular, there is nothing to prevent the neighbors of some circle C_v from wrapping more than once around C_v. We will introduce a natural notion of "branch point" to accomodate this and will extend the Andreev-Thurston Theorem to include existence and uniqueness of a certain class of branched circle packings.

FIGURE 1. A circle packing with overlaps $\pi/2$.

It may be helpful to the reader to see some early examples: We illustrate two circle packings in Figure 2. We have intentionally chosen these to share the same combinatorics τ of the packing in Figure 1, though this cannot be verified directly in the pictures. In each case a certain vertex $v_\infty \in \tau$ has been centered at the north pole, which is pointing directly at the viewer. Figure 2(a) is another maximal packing for τ, but this time with all overlaps zero (called a **tangency** packing). Figure 2(b) is a branched packing for τ and may be a little more difficult to understand: it's a tangency packing, but with two simple branch points, one at the center of the shaded circle and one at the center of the circle at the south pole (not seen). A careful look will suggest the meaning of branch point: Namely, follow the circles tangent to the shaded circle; you will find that they wrap twice around it. This circle packing induces a double covering of the sphere, locally two-to-one at the north and south poles and locally one-to-one elsewhere.

FIGURE 2. Maximal and branched tangency packings for τ.

Our aim in this paper is to give a new, complete, and constructive proof of the Andreev-Thurston Theorem, while simultaneously extending it to incorporate packings with "polynomial-like" branch structures (Main Theorem, §7). Moreover, the proof rests on solving a general "boundary value" problem for branched circle packings for triangulations of the disc (Proposition 1, §5); this allows extension of results (such as the Discrete Schwarz and Distortion Lemmas) which have played a central role in recent studies of tangency packings. We also illustrate a branched circle packing which is not polynomial-like and pose the existence and uniqueness question for general branched packings.

The authors thank Jeff van Eeuwen for valuable conversations on an earlier version of the material in this paper. The computer package *CirclePack* used to compute the packings illustrated in the paper is available from the second author.

1. Past and Future

The theory of circle packing has played roles in three main areas: in the construction of polyhedra in hyperbolic 3-space, in the approximation of analytic functions, and most recently, in development of a discrete analytic function theory.

History. The Andreev-Thurston Theorem above was proven by E. M. Andreev and used in the construction of hyperbolic polyhedra having preassigned dihedral angles, which were in turn glued together to form hyperbolic 3-manifolds ([A1], [A2]). The theorem was observed independently by Thurston in his study of hyperbolic orbifolds; he interpreted it in terms of patterns of circles on the boundary of hyperbolic 3-space, namely, the 2-sphere, S^2. For an early use of packings in the study of Schottky groups, see Brooks [Br].

Our statement of the Andreev-Thurston Theorem is taken from [Mo, Thm 15.1], where it appears without proof. It turns out that Paul Koebe had proved the result for tangent circles in 1936 [Kb]; see also Beardon/Stephenson [BSt], Bowers [Bo]. A proof by Marden and Rodin [RM] allows nonzero overlaps, but requires that the three overlaps for any one face sum to less than π. Igor Rivin's work on convex polyhedra [Ri] implies Andreev's Theorem in its original formulation (see especially Hodgson [H]) and even allows packings with overlaps up to π. The proofs applying to nontangency packings have all been nonconstructive, though in personal correspondence, Oliver Goodman reports a constructive method. None of the earlier work has addressed branched packings.

Connections with analytic functions began with the 1985 conjecture of Thurston [T2] that circle packings could be used to approximate conformal mappings; this was subsequently proven by Rodin and Sullivan [RS]. Thurston's conjecture also spawned development of a **discrete analytic function theory** based on maps between circle packings; see [BSt1], [BSt2], and references therein. In all of this work, the tangency case of the Andreev-Thurston Theorem (i.e., Koebe's Theorem) provided the initial motivation and remains as the fundamental rigidity result.

Future. In regard to connections with analytic functions, it seems likely that general overlap assignments will become increasingly important. A central open question concerns the extension of Thurston's 1985 conjecture on approximation to this setting. The methods in this paper permit extension of several of the ingredients which have been important in the tangency case — results such as the Schwarz and Distortion Lemmas and the Uniformization Theorem. (See Remark 4 in §5). In regard to discrete analytic function theory, the branched packings of this paper are new even in the tangency case and play a natural role as discrete rational functions (see §8).

There also remain questions of interest on the connection between circle packings on the surface of the sphere and hyperbolic polyhedra inside. For instance, branched circle packings seem to give "branched" hyperbolic polyhedra, which may be of independent theoretical interest.

Practical Matters: The constructive nature of our proof should not be overlooked: the use of supersolutions to the packing problem yields an iterative approximation scheme which converges monotonically to a solution, and error estimates can be related to hyperbolic area. See Remark 5 in §5; also a complexity analysis of Mohar [Mh].

Computations of circle packings have several potential uses beyond those related to the approximation of analytic functions: (a) For graph embedding, see [FH] and [MT]. (b) In the study of conformal geometry, compare the "squared rectangles" of Cannon, Floyd, and Parry [CFP] to the "circled rectangle" in Figure 3, obtained from an overlapping circle packing. Computational efficiency is a major issue in squaring rectangles. (c) In connection with hyperbolic polyhedra, specific circle packings may yield useful estimates of polyhedral volumes associated with orbifolds and manifolds.

FIGURE 3. A "circled" rectangle.

Overview: Our extension of the Andreev-Thurston Theorem is self-contained, with a proof which depends on elementary geometric ideas only and which shows that the packings are effectively computable. Despite their elementary nature, the details can get rather complicated nonetheless, and with the addition of branch considerations, we feel that an overview of the paper may be helpful.

Throughout the sequel, (τ, Φ) will be a labeled triangulation of S^2; we treat τ as an oriented simplicial 2-complex. In a circle packing P for (τ, Φ), the centers of the triple of circles associated with a face form a spherical triangle; we assume throughout that our packings give an orientation to these triangles which agrees with that of τ. The geometric complex which the triangles form is simplicially equivalent to τ and will be called the **carrier** of P, $\mathrm{carr}\,(P)$.

The circle packing guaranteed by the Andreev-Thurston Theorem is the maximal circle packing P (the term follows [BSt]) and is unique (up to conformal automorphisms) under the condition that its carrier triangulate S^2. However, removing this condition, other circle packings for (τ, Φ) may be possible. These necessarily involve some type of branching: specifically, if the neighbors of some circle C_v of P wrap $n+1$ times around C_v for some $n \geq 1$, we say that v is a **branch point** for P of **order** n. The branch set of P, denoted $\mathrm{br}\,(P)$, is the set of branch points, with each repeated according to its order. These more general packings P are what we call **branched circle packing** for (τ, Φ). By purely topological considerations, $\mathrm{carr}\,(P)$ is a branched covering of the sphere and its total branch order (the cardinality of $\mathrm{br}\,(P)$) is an even integer.

We must leave the statement of our Main Theorem for Section 7. It parallels the Andreev-Thurston Theorem, proving existence and uniqueness of circle packings; however, in addition to prescribed combinatorics and overlaps, it also involves certain prescribed branch points — namely, "polynomial" branch structures in which half the branching occurs at a single vertex. Our general approach involves converting to packings in the hyperbolic plane, solving a general Dirichlet-type (boundary value) problem, and converting back to the sphere. The proof in the hyperbolic setting uses the so-called "Perron" methods, principal tools in recent work in the tangency setting.

Suppose that we have a circle packing P for (τ, Φ). Normalize by a conformal automorphism to make the interior of the circle for some vertex v_∞ into the northern hemisphere of S^2; in the branched case, v_∞ must account for half the branching. Remove v_∞ and its star of faces from τ, leaving a complex K. The circles associated with K will intersect or lie within the southern hemisphere. *Via* the usual stereographic projection, the southern hemisphere is identified with the unit disc Δ, and when equipped with the hyperbolic (Poincaré) metric, Δ gives a model of the hyperbolic plane. In this standard way, then, we shift our attention to the study of packings for K in the hyperbolic plane.

Of course, we don't have a packing — we just have the data, K, Φ, and a prescribed branch set. Naively, one might start with a batch of circles, one for each vertex of K, and keep moving and resizing them in Δ until they fit together in the prescribed way. That's more or less what we do. Fortunately, spherical circles intersecting Δ turn out to be precisely the curves of constant curvature in the hyperbolic metric. We can describe this batch of circles, therefore, by specifying a "label" for K, that is, an assignment of curvature to each circle. This label (along with Φ) induces a hyperbolic structure on K, but with possible singularities at interior vertices, reflecting the fact that the corresponding circles may not fit properly with all their neighbors. The curvature label associated with the desired circle packing must be one which: (1) assigns values to boundary vertices so that their circles correctly overlap the northern hemisphere and (2) assigns compatible values to interior vertices so that each has neighbors which overlap it correctly while wrapping around it the desired number of times.

Finding the circle packing is reduced, then, to finding the correct label. We show its existence and uniqueness by solving a general boundary value problem for curvatures using a Perron method, introduced classically in the study of harmonic functions. (In circle packing, see [C], [BSt1], [Bo], and [DS].) Using several monotonicity results which tell how changes in labels effect the geometric structures they induce, we obtain our solution as the infimum of a certain family of labels. The main work in the generic case involves showing that this Perron argument gives a nondegenerate solution.

Outline: In Sections 2 and 3 we set up the necessary geometric structures and prove local monotonicity results. In Section 4 we define branch structures and labeled complexes. The key boundary value problem is Proposition 1 of Section 5, and in Section 6 we see how to convert between labels and packings. Our Main Theorem, the extension of Andreev-Thurston, is stated and proved in Section 7. Section 8 addresses non-polynomial branched packings, with a concrete example and open questions. We close in Section 9 by commenting on how the various hypotheses on packings come into play — the conditions are amazingly tight — and by illustrating a few missing and/or degenerate packings.

2. Geometric Structures

The unit disc $\Delta = \{z : |z| < 1\}$ with the hyperbolic metric $ds = 2\,dz/(1-|z|^2)$ serves as the hyperbolic plane. Conformal automorphisms of Δ are hyperbolic isometries — in fact they are precisely the Möbius transformations of S^2 which map Δ onto itself. If C is a spherical circle (i.e., an euclidean circle or straight line) and if $c = C \cap \Delta$ is not empty, then the curve c is called a **cycle** and C is called its **parent** circle. If C lies in Δ, $c = C$ is a **proper hyperbolic circle**, while if C intersects $\partial\Delta$, it is called a **hypercycle**. A hypercycle internally tangent to $\partial\Delta$ is called a **horocycle**.

In this section, we start with elementary properties of cycles, work our way successively through pairs and then triples of overlapping cycles, arriving finally at structures on K. At each stage we will need certain compatibility conditions on the data — particularly that for boundary vertices — and we will invoke the term **proper** for well-behaved data. We're careful with this because all requirements are sharp — exceptional cases could be easily overlooked.

Cycles: It is a pleasant fact that cycles are precisely the curves in Δ having constant curvature in the hyperbolic metric. This follows from the observation that the group of automorphisms of Δ mapping a cycle c to itself is transitive on c. The (geodetic) **curvature** of a cycle c, denoted $\operatorname{curv}(c)$, provides an especially convenient parameter for our purposes.

There is a normalized setting in which the curvature can easily be read off: Apply a conformal automorphism of Δ so that c contains the origin and its parent circle C is centered on the positive real axis. Define the **positive** direction on c to be associated with the counterclockwise direction on C. Let t, $0 < t \leq \infty$, denote the point where C intersects the positive real axis. (If c is a hyperbolic geodesic, C is the imaginary axis and $t = \infty$.) A computation shows that the curvature of c with respect to the hyperbolic metric is given by $\operatorname{curv}(v) = 1/t$. For hypercycles, $\operatorname{curv}(c) = \cos\alpha$, where $\alpha \in [0, \pi/2]$ is the acute angle c makes with the unit circle — that is, the overlap angle C makes in S^2 with the northern hemisphere. Of course, $\operatorname{curv}(c) = 0$ when c is a hyperbolic geodesic, and $\operatorname{curv}(c) = 1$ when c is a horocycle. For $\operatorname{curv}(c) \in (1, \infty)$, c is a proper hyperbolic circle, and the curvature is a strictly decreasing function of its hyperbolic radius r, namely, $\operatorname{curv}(c) = \coth(r)$. The following elementary observations about curvature will be helpful in the sequel.

Remark 1. *Let c be a cycle of positive curvature having as parent an euclidean circle C of radius r and center z. Then:*

(a). with r fixed, $\operatorname{curv}(c)$ is strictly decreasing as $|z|$ increases (until c becomes geodesic); and

(b). with z fixed, $\operatorname{curv}(c)$ is strictly decreasing as r increases (until c becomes geodesic).

Note that our cycles are always treated as oriented curves, with the orientation chosen to give nonnegative curvature. With this convention, we may define the **disc** bounded by c to be the portion of the hyperbolic plane on its left; the context will make clear which half space is the disc when c is a geodesic.

Pairs of Cycles: Two distinct cycles c_1, c_2 are said to **overlap** if they intersect in Δ; the **(angle of) overlap** ϕ is then the angle formed by and exterior to the discs which c_1 and c_2 bound. There is also one limit situation to include: distinct geodesics ending at a common point of $\partial\Delta$ and bounding disjoint discs are said to overlap with angle $\phi = 0$. Two details should be pointed out: If cycles overlap, then their parent circles overlap, but the converse can fail. Also, the discs bounded by two cycles may intersect without the cycles themselves overlapping.

Suppose data is given in the form of $\kappa_1, \kappa_2 \in [0, \infty)$ and $\phi \in [0, \pi/2]$. It is elementary to prove that there exists a pair of cycles c_1, c_2 with curvatures κ_1, κ_2, respectively, and overlapping with angle ϕ and that this pair is unique up to conformal automorphisms. We would expect an "edge" associated with these overlapping cycles, and this would clearly be the geodesic segment connecting their hyperbolic centers if both were proper hyperbolic circles. However, we're faced with several situations and it is convenient to distinguish two possibilities: (1) The intersection (one or two points) of the parent circles C_1, C_2 lies in Δ, or (2) One point of intersection of C_1, C_2 lies in the complement of Δ, possibly in $\partial\Delta$.

In situation (1), there is a unique geodesic γ orthogonal to c_1 and c_2. If c_1 is a proper hyperbolic circle, then hyperbolic symmetry in γ implies that γ passes through its hyperbolic center; likewise for c_2. We define the **edge** from c_1 to c_2, denoted $\text{edge}\,(c_1, c_2)$, as the geodesic segment in γ which passes through the point $c_1 \cap c_2$ if $\phi = 0$ or through the intersection of the discs bounded by c_1 and c_2 if $\phi > 0$ and which ends if it encounters a center of c_1 and/or c_2. Thus $\text{edge}\,(c_1, c_2)$ will be a finite, a one-way infinite, or a two-way infinite geodesic segment depending on the curvatures of the two cycles. In situation (2), a symmetry argument shows that there exists no common orthogonal geodesic to c_1 and c_2, so $\text{edge}\,(c_1, c_2)$ is explicitly defined as the empty set.

We say that the data κ_1, κ_2, ϕ for a pair of cycles is **proper** if the resulting edge is nonempty. A general observation is useful for finding practical criteria:

Remark 2. *Given a triple C_1, C_2, C_3 of overlapping circles in $\mathbf{C}$, let λ be the sum of their three overlap angles. If $\lambda < \pi$, then the circles form a (curvilinear) triangular interstice; if $\lambda = \pi$, then they pass through a common point; and if $\lambda > \pi$, then their interiors overlap in a triangular set having nonempty interior.*

Now, if $\kappa_1 > 1$ and/or $\kappa_2 > 1$, then the data is proper, since c_1 and/or c_2 will be a proper hyperbolic circle and $\text{edge}\,(c_1, c_2)$ will have a finite endpoint. However, when $\kappa_1, \kappa_2 \leq 1$, things are more complicated. If we apply the remark in S^2, with C_1, C_2 the parent circles and C_3 bounding the northern hemisphere, and we recall that $\arccos(\kappa_j)$ is the overlap of C_j with C_3, we obtain this criterion: *If $\kappa_1, \kappa_2 \leq 1$, then the data κ_1, κ_2, ϕ is proper if and only if* $\arccos(\kappa_1) + \arccos(\kappa_2) + \phi < \pi$.

Triples of Cycles: We next consider triples of overlapping cycles. The data will now consist of three curvatures, $\kappa_1, \kappa_2, \kappa_3 \in [0, \infty)$ and three overlaps $\phi_{12}, \phi_{13}, \phi_{23} \in [0, \pi/2]$. Each pair of curvatures and the associated overlap determine a pair of cycles. We will call the data for the triple **proper** if the data for at least one such pair is proper. Our fundamental existence result for triples is this:

Lemma 1. *Let curvatures $\kappa_1, \kappa_2, \kappa_3 \geq 0$ and overlap angles $\phi_{12}, \phi_{13}, \phi_{23} \in [0, \pi/2]$ be given. If this data is proper, then there exists a triple $\{c_1, c_2, c_3\}$ of cycles in Δ so that $\text{curv}(c_j) = \kappa_j$ and so that c_i and c_j overlap with angle ϕ_{ij}. The set of cycles is unique up to conformal automorphisms.*

Proof. In the *euclidean* setting, existence of triples of circles with specified radii and overlaps is proven in [T1, Lemma 13.7.2] and [MR, Lemma 2]. This is of use here (and later) because overlap angles are independent of the geometry, and we sketch the hyperbolic argument:

Pick a pair i, j and corresponding cycles c_i, c_j for which $\text{edge}(c_1, c_2) \neq \emptyset$. Apply a conformal automorphism putting $\text{edge}(c_1, c_2)$ along the real axis, with the disc bounded by c_1 on the left and that bounded by c_2 on the right, and with $c_1 \cap c_2$ on the imaginary axis. Fix parent circles C_1, C_2. Given any $r > 0$, there is a unique circle C_r of euclidean radius r which overlaps C_1 and C_2 with angles ϕ_{13} and ϕ_{23}, respectively, and has center z_r in the upper half plane. Moreover, C_r lies in the closure of the upper half plane and it's known that $|z_r|$ is strictly increasing in r (monotonicity is discussed in the next section).

Let $c_r = C_r \cap \Delta$. When r is small, $\text{curv}(c_r)$ is clearly large, and as r increases, $|z_r|$ increases, and by Remark 1, $\text{curv}(c_r)$ decreases. The intermediate value theorem gives an r for which $\text{curv}(c_r) = \kappa_3$. Set $c_3 = c_r$, and we have the desired triple. Uniqueness follows from euclidean uniqueness. $\square$

The triples of cycles we consider are oriented, positive orientation corresponding to counterclockwise placement (as with the choice of c_3 in the previous paragraph). The hypothesis that the data be "proper" is tailored to later uses. However, to see that some condition is needed in any case, observe that the data $\kappa_1 = 1, \kappa_2 = \kappa_3 = 0$, and $\phi_{12} = \phi_{13} = \pi/2$ are incompatible.

Boundary Labels: We are nearly ready to impose structures on K. Recall that (τ, Φ) is a labeled triangulation and that K is obtained from τ by deleting the star of v_∞. Write Φ_K for the restriction of Φ to the edges of K and treat (K, Φ_K) as a "labeled triangulation". The information on the overlaps of boundary vertices of K with v_∞ is encoded separately as an assignment g_Φ of curvatures to vertices of ∂K: Specifically

$$g_\Phi(w) = \cos(\Phi(\langle w, v_\infty \rangle)), \ \forall w \in \partial K.$$

Though Φ_K remains fixed throughout our work, we gain considerable flexibility by treating general boundary curvature assignments g in place of g_Φ. The hypotheses of the Andreev-Thurston Theorem suggest appropriate compatibility requirements for legitimate boundary labels:

Definition. *Let g be a nonnegative function defined on the boundary vertices of K. A* **critical edge** *of K is one which has assigned overlap $\pi/2$ (under Φ_K) and ends at a boundary vertex w with $g(w) = 0$. We will say that g is a* **proper boundary label** *for (K, Φ_K) if the following conditions hold:*

(a). No interior vertex v is the endpoint of more than two critical edges.

(b). If an interior vertex v is the endpoint of two critical edges, then those edges belong to a single face f and the overlap assigned to the edge of f opposite to v is positive.

(c). Every interior edge of K having both endpoints in ∂K has proper data.

Write $\{v_j\}$ for the set of vertices of K. A **proper label** *for (K, Φ_K) is an associated collection $\kappa = \{\kappa_j\}$ of nonnegative curvatures so that $\kappa_j > 1$ for every interior vertex v_j and so that the restriction of κ to ∂K is a proper boundary label. The resulting* **labeled complex** *will be denoted $K(\kappa)$.*

Condition (b) above is necessitated by an extraneous tangency pointed out by Thurston [T1, p. 13.52]: Elementary trigonometry shows that either of the subgraphs, Figure 4(a) or (b), occurring in (τ, Φ) (that is, two faces sharing an edge and having the indicated overlap angles) will lead to the circle pattern of Figure 4(c). In either case, two circles which do not share an edge in τ are forced nonetheless to be tangent; this is called an **extraneous tangency**. The subgraph 4(a) is a problem, since the circle for vertex v, an interior vertex of K, would be forced to be a horocycle. Condition (b) hypothesizes that no subgraph of this type is involved in our label. However, even if this subgraph were to appear, we could simply replace it by the subgraph of 4(b) — that is, replace the pair (τ, Φ) by a modified pair (τ', Φ'). It is not difficult to check that the new pair satifies the hypotheses of the Andreev-Thurston Theorem, and any circle packing for (τ', Φ') is simultaneously one for (τ, Φ). In other words: *Without loss of generality, one may assume that the subgraph of Figure 4(a) does not appear in (τ, Φ).* We will have occassion to use this fact when we define polynomial branch structures.

Structures on K: It is standard that $|K|$ denotes some geometric realization of K; our aim is to put a hyperbolic metric on $|K|$ based on the label κ. Typically the boundary faces will have vertices and edges in the ideal boundary at infinity, so we disregard the topological boundary and treat $|K|$ is an open topological disc. (The boundary edges and vertices remain as simplices of the abstract complex K, of course.)

Start with the faces: Let $f = \langle v_1, v_2, v_3 \rangle$ be a face of K, with order determined by the orientation on K, and let $\kappa_1, \kappa_2, \kappa_3$ be the corresponding curvatures from κ. Since τ is not

FIGURE 4. Subgraphs giving extraneous tangencies.

a tetrahedron, f must have at least one interior edge. By definition of a proper label, the data for f is therefore proper, and Lemma 1 gives an (essentially) unique triple $\{c_1, c_2, c_3\}$ of cycles in Δ with $\operatorname{curv}(c_i) = \kappa_i$ (and overlaps specified in Φ_K). To be compatible with the orientation of f, assume these are placed in counterclockwise order. Use this configuration to associate a **geometric face** F with f as follows: Define the set $G \subset \Delta$ by

$$G = \operatorname{edge}(c_1, c_2) \cup \operatorname{edge}(c_2, c_3) \cup \operatorname{edge}(c_3, c_1).$$

Since the data for f is proper, G is nonempty, and one easily verifies that the edges in G can intersect only at endpoints. Indeed, G separates Δ into two or more (possibly ideal) polygons; F is the one on the left as the edges are transited in the orientation inherited from K. Since we want $|K|$ to be an open topological disc, include as part of F, each geodesic segment $\operatorname{edge}(c_i, c_j)$ whose associated edge $\langle v_i, v_j\rangle$ is an interior edge of K, but exclude any geodesic boundary segments associated with boundary edges of K. Geometric faces will be important in the sequel, so we illustate some of the possibilities in Figure 5.

Define angles $\alpha_1, \alpha_2, \alpha_3$ associated with F as follows: If $\kappa_j > 1$, then c_j is a proper hyperbolic circle and its center is a finite vertex of F; we let α_j be the geometric angle in F formed there. Note from the orientation of $\{c_1, c_2, c_3\}$ and our definition of F that $\alpha_j < \pi$. If $\kappa_j \in [0, 1]$, then c_j is a hypercycle and gives rise to a point or arc of $\partial\Delta$ in ∂F. Since there is no geometric angle associated with c_j, we choose to define $\alpha_j = 0$.

We now lift the geometric structure of F to $f \subset |K|$: Let $h : \operatorname{int}(f) \longrightarrow \operatorname{int}(F)$ be a homeomorphism which extends continuously to map each edge $\langle v_i, v_j\rangle$ of f which is an interior edge of K onto the corresponding (necessarily nonempty) geodesic $\operatorname{edge}(c_i, c_j)$ of

FIGURE 5. Examples of geometric faces.

F. In particular, each interior vertex v_i is identified with the hyperbolic center of the cycle c_i. The hyperbolic structure on F induces, under h, a hyperbolic structure on f — that is, a Riemannian metric of constant curvature -1. The geometric properties of F, angles, distances, area, are thus lifted to f. In particular, each vertex v_i has an angle α_i associated with f as determined earlier. Also, the intersection of a cycle c_j with F is lifted to a curve in f which has constant curvature κ_j.

We place compatibility conditions on these face structures: Namely, consider two faces f and f' of K which share an edge, say $f = \langle v_1, v_2, v_3 \rangle$ and $f' = \langle v_1, v_3, v_4 \rangle$. The common edge $\langle v_1, v_3 \rangle$ is necessarily an interior edge of K. Applying a conformal automorphism, if necessary, corresponding geometric faces F and F' may be positioned to share edge (c_1, c_3) (this will be nonempty, will belong to $F \cap F'$, and F and F' will be on opposite sides of it by orientation). Assume in this setting that h and h' satisfy $h^{-1} \circ h' \equiv \mathrm{id}$ on $\langle v_1, v_3 \rangle$ so that the two metrics on that edge agree. This can be easily arranged with adjustments to h and h' in a neighborhood of the affected edge.

Working face-by-face, the labeled complex $K(\kappa)$ is thus provided with a hyperbolic structure, and as a metric space is denoted $|K(\kappa)|$. The metric has been made compatible across edges, so it has constant curvature -1 on all of $|K|$, except possibly at interior vertices. Our work comes in analyzing these "cone-type" singularities. The structure near a vertex v depends only on the data for its star, and the local curvature is captured with this

Definition. *Let κ be a proper label for (K, Φ_K). Given a vertex v, let $f_1, f_2, \cdots, f_n$ denote the faces of its star. In the hyperbolic structure of $|K(\kappa)|$, each face f_j has an angle α^j corresponding to v. Define the* **angle sum** *at v by $\theta_\kappa(v) = \sum_{j=1}^n \alpha^j$.*

Angle sums underly Proposition 2 of §6, which gives necessary and sufficient conditions for an abstract hyperbolic structure on K to correspond with one induced by an actual collection of circles.

Our final structural detail concerns a finite hyperbolic invariant which we may represent as area. Given a face f in $|K(\kappa)|$, its geometric face F may have infinite area, so we define a certain finite area polygon F_a within F called the **area-face** as follows: Let E be the union of those geodesic edges of F which have at least one end at some finite vertex of F. (Each interior vertex of f is associated with such a vertex of F, so E is not empty.) Define F_a as the (hyperbolically) convex hull of E. Because the angles α_j of F do not exceed π, F itself is hyperbolically convex, so F_a is always a subset of F. The faces differ if and only if one or more of the cycles c_j has curvature less than 1; see, for instance, Figure 6, where examples of F_a are darkly shaded.

FIGURE 6. Examples of area-faces.

The polygon F_a is 4-sided if precisely one of the curvatures κ_j is less than 1, and 3-sided otherwise. We may use the same angles α_j for F_a as were defined for F, since at any finite vertex, the geometric angle is the same for the two faces. In particular, the area formula for hyperbolic polygons gives

$$\text{Area}(F_a) = \begin{cases} \pi - \alpha_1 - \alpha_2 - \alpha_3, & \text{if } F_a \text{ is 3-sided} \\ 2\pi - \alpha_1 - \alpha_2 - \alpha_3, & \text{if } F_a \text{ is 4-sided.} \end{cases} \tag{1}$$

One can verify with elementary geometry that in both cases, Area$(F_a) \leq \pi$. Finally, observe that the definitions and all geometric quantities associated with F and F_a are invariant under conformal automorphisms and thus may be lifted to f.

3. Local Monotonicity Results

The monotonicities of interest are quite straightforward in the euclidean setting: Form a euclidean triangle T with the centers of three circles having specified overlap angles. If the radius of one circle is increased and its location is adjusted to maintain the specified overlaps, then the angle in T at that circle will decrease, the angles at the other two circles will increase, and the area of T will increase. (See [T1, Lemma 13.7.3] and [MR, Lemma 3].)

We need analogous monotonicity results for triples of cycles c_i in Δ, with radii replaced by curvatures and area replaced by hyperbolic area (of area-faces). Direct computations become unwieldy in the extreme. However, with some elementary hyperbolic geometry, appropriate normalizations by conformal automorphisms, and the euclidean results, one may prove the analogues fairly easily. We will indicate some arguments, but leave others for the interested reader. Recall that hyperbolic radius is inversely related to curvature.

FIGURE 7. Exceptional situations.

Two exceptional situations will haunt our monotonicity results, so we might as well name and describe them first. The main two are illustrated in Figure 7, with geometric and area-faces shaded. The cycle c_1 is a proper hyperbolic circle and has been centered at the origin; c_2, c_3 are placed counterclockwise. We use the usual notations for curvatures and overlaps.

ES1: $\kappa_1 > 1, \kappa_2 = \kappa_3 = 0, \phi_{12} = \phi_{13} = \pi/2, \phi_{23} > 0$. Here $\mathrm{Area}(F_a)$ and all angles are independent of κ_1.

ES2: $\kappa_2 = 0, \kappa_3 \geq 1, \phi_{12} = \phi_{23} = \pi/2$. Here $\mathrm{Area}(F_a)$ and all angles are independent of κ_1.

ES3: $\kappa_3 > 1, \phi_{13} = 0, \phi_{12} = \phi_{23} = \pi/2$. Here angle α_3 is independent of κ_1.

Note the existence in ES1 and ES2 of one or more critical edges; these account for certain of the restrictions placed on "proper" labels. The third exceptional situation comes up in the next lemma and is included for completeness.

Lemma 2. *Let c_j be cycles in Δ with curvatures $\kappa_j, j = 1, 2, 3$. Assume that $\kappa_1 > 1$ and $\kappa_2, \kappa_3 \geq 0$. Let ϕ_{ij} denote the usual overlap angles. Let F_a be the area-face, α_j the angles of F_a as defined in §2. The angles α_j and $\mathrm{Area}(F_a)$ are continuous functions of the curvatures (assuming that the overlaps maintain their assigned values), and the following monotonicity results hold. Except as noted in ES1, ES2, and ES3 above, monotonicity is strict.*

(a). α_1 is increasing in κ_1.

(b). α_2 and α_3 are decreasing in κ_1.

(c). $\mathrm{Area}(F_a)$ is decreasing in κ_1.

Moreover, we have the following limit results:

(d). $\lim_{\kappa_1 \to \infty} \alpha_1 = (\pi - \phi_{23})$ and, aside from ES1 and ES2, $\lim_{\kappa_1 \to 1} \alpha_1 = 0$.

(e). $\lim_{\kappa_2 \to \infty} \alpha_1 = U$, where $0 \leq U \leq \pi/2$ and $\lim_{\kappa_2, \kappa_3 \to \infty} \alpha_1 = 0$.

(f). Fix κ_2 and let $\kappa_1, \kappa_3 \to \infty$. If $\kappa_2 < 1$, then $\mathrm{Area}(F_a) \to \pi$; while if $\kappa_2 \geq 1$, then $\mathrm{Area}(F_a) \to 0$.

Proof. Using Remark 1, one can apply euclidean monotonicity results in several hyperbolic settings. For instance, one can establish this fact: The curvature of a cycle c which overlaps a fixed circle in Δ centered at the origin at a specified angle is strictly decreasing as a function of its parent's euclidean radius.

We illustrate how this can be used in confirming the monotonicity of α_1: The hypothesis $\kappa_1 > 1$ implies that c_1 is a proper hyperbolic circle. Normalize so it remains centered at the origin. Assume that $\kappa_2 > 0$ or that $\phi_{12} < \pi/2$, so that C_2 is a euclidean circle (i.e., c_2 is not a diameter). A specific increase in κ_1 can be accomplished by a euclidean scaling, $z \longrightarrow tz$, for an appropriate $t < 1$. However, this draws C_2 towards the origin and increases the curvature of c_2. The fact above tells us that we must increase the euclidean radius of C_2 and move it away from the origin in order to keep the correct overlap with (the new) c_1 while reestablishing the curvature of c_2. A similar adjustment is required of C_3. Because the euclidean radii of C_2 and C_3 increase while that of C_1 decreases and because C_1 is centered at the origin, the euclidean monotonicity results imply that the angle α_1 increases, as desired. Monotonicity of α_1 in other circumstances is left to the reader; this proves (a).

As to the behavior of the other angles, say we consider α_2. If $\kappa_2 \leq 1$, then $\alpha_2 = 0$ independent of κ_1. We may therefore assume that c_2 is a hyperbolic circle and center it at the origin. When κ_1 increases, the euclidean radius of C_1 decreases and the euclidean results imply that α_2 decreases. This proves (b).

The monotonicity of area requires some effort. The area formula (1) shows that the monotonicities for α_1 and for α_2, α_3 have competing effects on $A = \text{Area}\,(F_a)$; so the fact that A is decreasing in κ_1 is not immediate.

Suppose first that $\kappa_2, \kappa_3 > 1$. Center C_2 at the origin and C_3 on the positive real axis; by orientation, this puts C_1 in the upper half plane. Keep this normalization as κ_1 increases so that edge$\,(c_2, c_3)$ remains fixed. The angles α_2, α_3 strictly decrease, so the corresponding faces are properly nested, implying the area inequality. Next, suppose $\kappa_2, \kappa_3 \leq 1$. Then $\alpha_2 = \alpha_3 = 0$, so the inequality follows from (1) and the monotonicity of α_1.

This leaves a rather tricky case, that in which only one of κ_2, κ_3 exceeds 1. We leave this to the interested reader with a brief hint. Say that $\kappa_2 \leq 1$, $\kappa_3 > 1$, and consider two values $\kappa_1' > \kappa_1 > 1$ for the curvature of C_1. With some effort one can arrive at a generic situation normalized as in Figure 8; here the circle and and area face bounded by the dotted lines are those associated with κ_1'. It suffices to show that

$$\text{Area}\,(E_1) \geq \text{Area}\,(E_2).$$

A comparison of their angles would suggest equality, except that c_2 is not necessarily geodesic. Fortunately, c_2 has non-negative curvature κ_2 with respect to E_1, so the Gauss-Bonnet Theorem implies the inequality.

The first limit statement of (d) is evident if c_2 and c_3 intersect at $z \in \Delta$. If they intersect at $z \in \partial\Delta$, then necessarily $\kappa_2 = \kappa_3 = 0$ and $\phi_{23} = 0$. An appropriate normalization is

FIGURE 8. The area comparison.

then to center c_1 at the origin and put z at 1; it is clear that $\alpha_1 \to \pi = \pi - \phi_{23}$ as the hyperbolic radius of c_1 goes to zero. The second limit result is not true in ES1 and ES2, since α_1 is then a positive constant independent of κ_1. The generic case is left to the reader, as is (e).

Finally, (f) follows from (1): As κ_1 and κ_3 go to infinity, the sum $\alpha_1 + \alpha_3$ clearly goes to π. If $\kappa_2 < 1$, the area-face for our triple has four sides, while if $\kappa_2 \geq 1$, it has three. $\square$

Our last monotonicity results concern angle sums. We consider a labeled complex $K(\kappa)$, but concentrate on the star of a vertex.

Lemma 3. *Let κ be a proper label for (K, Φ_K). Suppose $M \subset K$ is the star of a vertex v_0 with neighboring vertices $v_1, v_2, \cdots, v_{n+1}$, faces $f_1, \cdots, f_n$, and area-faces $F_a^1, \cdots, F_a^n$. The angle sum $\theta = \theta_\kappa(v_0)$ depends only on the curvatures $\kappa_j = \kappa_j, j = 0, 1, \cdots, n+1$, (and on Φ_K) and we have the following monotonicities:*

> *(a). θ is a continuous function of $\kappa_j, j = 0, \cdots, n+1$, which is increasing in κ_0 and decreasing in κ_j for $j > 0$.*
>
> *(b). θ converges to zero when all $\kappa_j, j = 1, \cdots, n+1$, go to infinity.*

Moreover, suppose v_0 is an interior vertex (so $v_{n+1} = v_1$), ∂M is formed by the edge-path $\{e_1, \cdots, e_n\}$, and the curvatures $\kappa_j, j = 1, \cdots, n$, are fixed.

> *(c). θ is strictly increasing in κ_0, with $\lim_{\kappa_0 \to 1} \theta = \mu < 2\pi$ and*
>
> $$\lim_{\kappa_0 \to \infty} \theta = n\pi \sum_j \Phi(e_j).$$
>
> *(d). $A = \sum_{j=1}^n \text{Area}(F_a^j)$ is strictly decreasing in κ_0.*

Proof. (a) follows from Lemma 2, and (b) is elementary. For (c) and (d), we must contend with the possibility that one or more of the faces f_j fall into ES1 or ES2, for then the angle which these faces contribute to θ and the area they contribute to A are constant with respect to κ_0. However, each of ES1 and ES2 involves one or two critical edges, and the fact that κ is a proper label for (K, Φ_K) limits the number of critical edges from v_0. In particular, it is evident that M can have at most one ES1 face and/or at most two ES2 faces. Moreover, if M were to contain only these three exceptional faces, then the edges $\langle v_1, v_2\rangle, \langle v_2, v_3\rangle, \langle v_3, v_1\rangle$ would be a closed 3-edge loop whose overlaps sum to more than π, contradicting hypothesis (A) of the Andreev-Thurston Theorem. We conclude that M contains at least one nonexceptional face. By Lemma 2(a), the angle in that face at v_0 is strictly increasing in κ_0, proving the strict monotonicity of θ, and by Lemma 2(c), A is strictly decreasing.

For the limit statements in (c), we must again contend with ES1 and ES2: If a face $f = \langle v_0, v_j, v_{j+1}\rangle$ is in ES1, then the angle at v_0 in f is $\pi - \Phi(\langle v_j, v_{j+1}\rangle)$, while if in ES2, it is $\pi/2$, independent of κ_0. If μ denotes the portion of θ due to any such exceptional faces, then our counting above implies $0 \leq \mu < 2\pi$. By Lemma 2(d), the angle at v_0 in the other faces goes to zero as κ_0 decreases to 1, so $\lim_{\kappa_0 \to 1} \theta = \mu$.

In the edge-path forming ∂M, let e_j denote the edge opposite to v_0 in the face $f_j = \langle v_0, v_j, v_{j+1}\rangle$. By Lemma 2(d) for generic faces or by observation for the exceptional ones, the angle at v_0 in f_j goes to $\pi - \Phi(e_j)$ as κ_0 increases to infinity. In particular, the limit of the angle sum at v_0 is $\nu = \sum_j [\pi - \Phi(e_j)] = n\pi - \sum_j \Phi(e_j)$, as claimed. Note that $\nu > 2\pi > \mu$: This is clear for $n \geq 5$, since $\Phi(e_j) \leq \pi/2$, and follows for $n = 3, 4$ by hypotheses (A) and (B) of The Andreev-Thurston Theorem. $\square$

Remark 3. *When $\kappa_0 \in (1, \infty)$ is fixed, then as one or more of the κ_j goes to infinity, θ decreases to a nonnegative value. This allows us to define the angle sum at an interior vertex v_0 for any values of $\kappa_j \in [0, \infty], j = 1, 2, \cdots, n$, as long as the curvature for v_0 itself lies in $(1, \infty)$.*

4. Branch Structures and Packing Labels

When are the structures of abstract labeled complexes $K(\kappa)$ induced by circle packings? Necessary and sufficient conditions are purely local and depend on the "angle sums" defined in Section 2.

When v is an interior vertex, its star in the metric space $|K(\kappa)|$ consists of a closed chain $\{f_1, \cdots, f_n\}$ of contiguous faces. Of course, who's to say that geometric faces F_j having curvature and overlap data corresponding to these abstract faces will fit together coherently in the hyperbolic plane? The angle sum at the vertex v captures just this: *Geometric faces $\{F_1, \cdots, F_n\}$ may be chosen to fit together precisely, the final edge of F_n identified exactly with the initial edge of F_1, if and only if $\theta_\kappa(v) = 2\pi m$ for some positive integer m.* In this case, the geometric faces wraps m times around a common vertex corresponding to v. The associated circles form what is called a "flower", with the n neighbor circles as petals which wrap m times around the circle for v.

The metric of $|K(\kappa)|$, which is hyperbolic on the open faces and across interior edges, has singularities at interior vertices such as v. It is easy to check that the singularity at v is removable if and only if $m = 1$, and in particular, that $|K(\kappa)|$ is hyperbolic if and only if $\theta_\kappa(v) = 2\pi$ for every interior vertex v. However, as we wish to consider more general situations, we make the following definitions. These notions have been investigated in cases of tangent circle packings in [BSt2], [DS], and [D1]; our notation is based on that of Dubejko [D1].

Definition. *An interior vertex $v \in |K(\kappa)|$ is said to be a* **branch point** *of* **order** $m-1$ *(***multiplicity** m*) if $\theta_\kappa(v) = 2\pi m$ for an integer $m \geq 2$. The set of branch points, with each repeated according to its order, is denoted by $\mathrm{br}(K(\kappa))$.*

It is clear — for instance, by applying Lemma 3 to flowers — that the combinatorics of K along with the values of Φ_K restrict branch orders. We'll see that the following definition captures the global conditions precisely.

Definition. *A set $\beta = \{b_1, b_2, \cdots, b_m\}$ of interior vertices of K, possibly involving repetitions, is called a* **branch structure** *for (K, Φ_K) if the following condition holds: For each simple closed edge path $\gamma = \{e_1, e_2, \cdots, e_n\}$ in K which surrounds at least one vertex of K, the inequality*

$$\sum_{j=1}^{n} [\pi - \Phi_K(e_j)] > (2T + 2)\pi$$

holds, where T is the number of points of β inside γ, counting repetitions.

Our interest is in specifying a branch structure β and solving for a label κ so that $\beta = \mathrm{br}(K(\kappa))$. We proceed by studying supersolutions; the terminology is taken from the study of harmonic and superharmonic functions.

Definition. *Let β be a branch structure for (K, Φ_K). For interior vertices v of K, write $\beta(v)$ for the multiplicity of v; that is, $\beta(v) = k+1$, where k is the number of times v occurs in β. A proper label κ is said to be a* **β-packing label** *for (K, Φ_K) if $\theta_\kappa(v) = 2\pi\beta(v)$ for every interior vertex v of K. It is a* **β-subpacking label** *if $\theta_\kappa(v) \geq 2\pi\beta(v)$ and a* **β-superpacking label** *if $\theta_\kappa(v) \leq 2\pi\beta(v)$ for every interior vertex v.*

Note that $\beta = \emptyset$ always qualifies as a branch structure, and when this is the case we would drop the prefix β in these notations.

5. A Boundary Value Problem

We now solve a fairly general boundary value problem, generalizing to the overlapping case a result proven for tangent packings in [BSt1].

Proposition 1. *Assume that g is a proper boundary label and that β is a branch structure for (K, Φ_K). Then there exists a unique β-packing label $\widetilde{\kappa}$ for (K, Φ_K) so that $\widetilde{\kappa}(w) = g(w)$ for every boundary vertex w of K.*

Proof. We may assume that K has edge-connected interior and that every face has an interior vertex. (A set of vertices is edge-connected if any two are connected by a chain of edges with endpoints in the set.) Otherwise, one could restrict attention to an edge-component of interior vertices and its immediate neighbors, which are boundary vertices of K. After proving the result for each component, the conclusion would follow easily for the full complex.

We will write $\mathcal{V}$ for the vertex set of K, comprised of the interior vertices $\mathcal{V}_{\text{int}}$ and the boundary vertices $\mathcal{V}_\partial$. Write

$$\mathcal{L} = \{\kappa : \kappa \text{ a } \beta\text{-superpacking label for } (K, \Phi_K),\ \kappa(w) = g(w), \forall w \in \mathcal{V}_\partial\}.$$

It is necessary first to check that $\mathcal{L}$ is nonempty. Start by assigning curvature $g(w)$ to each boundary vertex and curvature 1 to each interior vertex. This is, of course, not a proper label, but we may modify it as follows: Suppose $v_1, v_2, \cdots, v_k$ is a listing of the interior vertices. By Lemma 3(c), the curvature at v_1 may be increased from 1 until the angle sum at v_1 is 2π. Similarly, the curvature at v_2 may be increased from 1 until the angle sum there is 2π, and so forth through v_k. Denote the resulting curvatures by κ^0 and notice that this is a proper label for (K, Φ_K) with boundary values equal to g. The key observation is that by Lemma 2(b), the angle sum at any v_j can only decrease due to the subsequent changes at other interior vertices, and hence is bounded above by 2π in the final label κ^0 (though it may well be less). This means that κ^0 is a β-superpacking label and hence that $\mathcal{L} \neq \emptyset$.

We may now define the function $\widetilde{\kappa}$ by $\widetilde{\kappa}(v) = \sup_{\kappa \in \mathcal{L}}\{\kappa(v)\}$. In particular, $\widetilde{\kappa} \equiv g$ on $\mathcal{V}_\partial$. We will prove that $\widetilde{\kappa}$ is a β-packing label for (K, Φ_K) and that it is the unique solution to our boundary value problem. This approach is based on the methods introduced by Perron to solve the classical Dirichlet problem for harmonic functions using superharmonic functions.

First, one easily verifies that $\mathcal{L}$ is closed under maxima: $\kappa_1, \kappa_2 \in \mathcal{L}$ implies $\kappa_3 \equiv \max\{\kappa_1, \kappa_2\} \in \mathcal{L}$.

Next, we consider a vertex $v \in \mathcal{V}_{\text{int}}$. If $\widetilde{\kappa}(v) < \infty$, then by Remark 3, the limiting angle sum $\theta_{\widetilde{\kappa}}(v)$ is well defined as $\kappa \to \widetilde{\kappa}$. We claim

$$\text{(2)} \qquad \text{If } v \in \mathcal{V}_{\text{int}} \text{ and } \widetilde{\kappa}(v) < \infty, \text{ then } \theta_\kappa(v) \longrightarrow 2\pi\beta(v) \text{ as } \kappa \to \widetilde{\kappa}.$$

One easily uses the finiteness of $\mathcal{V}$ and the fact that $\mathcal{L}$ is closed under maxima to find a sequence $\{\kappa_j\} \subset \mathcal{L}$ which increases monontonically to $\widetilde{\kappa}$. Suppose the degree of v is n and $\gamma = \{e_1, \cdots, e_n\}$ is the simple closed edge-path formed by its immediate neighbors. Because β is a branch structure, $\sum_j [\pi - \Phi(e_j)] > 2\pi\beta(v)$. By Lemma 3(c), we may assume that $\kappa_j(v)$ has been increased until $\theta_{\kappa_j}(v) = 2\pi\beta(v)$, and by Lemma 2(b) this change could only decrease the angle sums of its neighbors, preserving the β-superpacking property. In particular, the curvature $\kappa_j(v)$ will never exceed $\widetilde{\kappa}(v)$, which was assumed finite. Therefore, we may assume without loss of generality that $\theta_{\kappa_j}(v) = 2\pi\beta(v)$ for $j = 1, \cdots$. The continuity statement of Lemma 3(a) now establishes (2).

By this last observation, to conclude that $\widetilde{\kappa}$ solves our boundary value problem it suffices to show that $\widetilde{\kappa}(v) < \infty$, $\forall v \in \mathcal{V}_{\text{int}}$. This is where our main effort is required.

We begin with several notations: Let F, E, and V denote the number of faces, edges, and vertices of K, respectively, and write V_∂ and V_{int} for the number of boundary and interior vertices. Let κ denote a proper label for (K, Φ_K) with $\kappa \equiv g$ on $\mathcal{V}_\partial$. As we have seen, this label induces a metric structure on K which determines angles and angle sums. Define the **shortage** of κ by

$$\Sigma_\kappa = \sum_{v \in \mathcal{V}_{\text{int}}} (2\pi\beta(v) - \theta_\kappa(v)).$$

This reflects the shortage of interior angle sums for κ in comparison to a those for a β-packing label; note that $\Sigma_\kappa \geq 0$ if $\kappa \in \mathcal{L}$. If $\mathcal{H} \subseteq \mathcal{V}$, define $\theta_\kappa(\mathcal{H}) = \sum_{v \in \mathcal{H}} \theta_\kappa(v)$. Recall that κ also associates with each face $f \in K$ a geometric area-face, F_a. Write A_κ for the sum of the hyperbolic areas of these area-faces. Denote by N_g the number of faces which have one and only one vertex w satisfying $w \in \mathcal{V}_\partial$ and $g(w) < 1$. This appears in the formulas because it counts the number of faces whose associated area-faces have 4 rather than 3 edges. Write T for the total branch order of β, $T = \sum_{v \in \beta} [\beta(v) - 1]$.

Next we introduce the invariant combinatoric quantity $(F - 2V_{\text{int}})$, introduced in this context in [Bo]. Let κ be a proper label for (K, Φ_K) with $\kappa \equiv g$ on $\mathcal{V}_\partial$. Three angles $\alpha_j, j = 1, 2, 3$, were defined for each area-face associated with $K(\kappa)$. Using formula (1) and keeping count of those area-faces having 4 sides, we have $A_\kappa = \pi N_g + \sum_{j=1}^{F}(\pi - \alpha_1^j - \alpha_2^j - \alpha_3^j)$. Reorganizing by grouping the face angles into angle sums at the vertices of K yields

$$A_\kappa - \pi N_g = \pi F - \theta_\kappa(\mathcal{V}) = \pi F - (\theta_\kappa(\mathcal{V}_\partial) - \Sigma_\kappa + \sum_{v \in \mathcal{V}_{\text{int}}} 2\pi\beta(v)). \tag{3}$$

Because $|K|$ is an open topological disc, the Euler characteristic $\chi(|K|) = F - E + V = 1$. Also, $3F = 2E - V_\partial$, so $E = \frac{1}{2}(3F + V_\partial)$. Therefore, $F - 2V_{\text{int}} = V_\partial - 2$, and since $V_\partial \geq 3$, this is positive. Combining this with (3), we conclude that for any proper label κ with $\kappa \equiv g$ on $\mathcal{V}_\partial$,

$$A_\kappa - \pi N_g + \theta_\kappa(\mathcal{V}_\partial) - \Sigma_\kappa + 2\pi T = (F - 2V_{\text{int}})\pi > 0. \tag{4}$$

Denote by $\mathcal{W} \subseteq \mathcal{V}_{\text{int}}$ the vertices v with $\widetilde{\kappa}(v) = \infty$. We must show $\mathcal{W} = \emptyset$. It suffices to consider an edge-component $\mathcal{B}$ of $\mathcal{W}$. Let L be the subcomplex of K obtained by taking all (closed) faces which have a vertex in $\mathcal{B}$. If we prove that the cases $L = K$ and $\emptyset \subsetneqq L \subsetneqq K$ are impossible, then $L = \emptyset \Longrightarrow \mathcal{B} = \emptyset \Longrightarrow \mathcal{W} = \emptyset$, implying that $\widetilde{\kappa}$ is the desired β-packing label.

Case I: Suppose $L = K$: Because $\mathcal{B}$ is edge-connected, we must have $\mathcal{B} = \mathcal{W}$. We claim $\mathcal{W} = \mathcal{V}_{\text{int}}$. Suppose $v \in \mathcal{V}_{\text{int}} \backslash \mathcal{B}$. If all its neighbors are in $\mathcal{B}$, then by Lemma 3(b), the angle sum at v would go to zero as $\kappa \to \widetilde{\kappa}$, contradicting our earlier observation that the angle sum will converge to $2\pi\beta(v)$. If v belongs to an edge-component, say Q, of $\mathcal{V} \backslash \mathcal{B}$, then $L = K$ implies Q contains no closed loops with vertices in $\mathcal{V}_{\text{int}}$. In particular, Q will have an end, that is, a vertex having only one neighbor belonging to Q. Assume that q is such a vertex. All neighbors of q, save one, lie in $\mathcal{B}$. By Lemma 2(e), at most two of the angles forming the angle sum at q can possibly fail to go to zero, and these converge to values at most $\pi/2$. This means the angle sum at q has limit no larger than π as $\kappa \to \widetilde{\kappa}$, again contradicting (2). Thus, in Case I, all interior curvatures go to infinity.

We will consider areas of area-faces as $\kappa \to \widetilde{\kappa}$. Let $w_1, w_2, \cdots, w_n$ be the ordered list of boundary vertices of K, $e_1, e_2, \cdots, e_n$ the chain of boundary edges. Each boundary edge is in a face, and since it is assumed that each face has at least one interior vertex, the n boundary edges belong to n distinct faces $f_1, f_2, \cdots, f_n$. For each label κ let $A_1(\kappa), A_2(\kappa), \cdots, A_n(\kappa)$ be the areas of the area-faces corresponding to $f_1, \cdots, f_n$. Each f_j has two angles neighboring e_j, which we denote $\alpha_j(\kappa)$ and $\beta_j(\kappa)$, and an angle opposite to e_j and associated with an interior vertex, denoted $\gamma_j(\kappa)$.

As $\kappa \to \widetilde{\kappa}$, Lemma 2(e) implies that the angles forming the angle sum at a boundary vertex w_j go to zero, with the possible exception of the two angles neighboring the boundary edges emanating from w_j. In particular, adding over all boundary vertices, we have

$$\lim_{\kappa} \theta_\kappa(\mathcal{V}_\partial) = \lim_{\kappa}(\alpha_j(\kappa) + \beta_j(\kappa)).$$

By Lemma 2(f), the faces of $|K(\kappa)|$ having two or more interior vertices have areas going to zero or π as $\kappa \to \widetilde{\kappa}$; therefore, the only face angles playing a role in the limit are those of $f_1, \cdots, f_n$. We have

$$\lim_{\kappa} A_\kappa = \pi N_g + \sum_{j=1}^{n} \lim_{\kappa}(\pi - \alpha_j(\kappa) - \beta_j(\kappa) - \gamma_j(\kappa))$$

Rearranging and applying Lemma 2(d), we have

$$\lim_{\kappa}(A_\kappa + \theta_\kappa(\mathcal{V}_\partial) - \pi N_g) = \sum_{j=1}^{n} \lim_{\kappa}(\pi - \gamma_j(\kappa)) = \sum_{j=1}^{n} \Phi(e_j) \begin{cases} < \pi(V_\partial - 2T - 2), & V_\partial \geq 5, \\ < 2\pi, & V_\partial = 4, \\ < \pi, & V_\partial = 3. \end{cases}$$

The cases $V_\partial = 3$ and $V_\partial = 4$ follow from the hypotheses on Φ: namely, since K has nonempty interior, ∂K cannot bound a face or a pair of faces. The case $V_\partial \geq 5$ follows from the fact that β is a branch structure.

To see how one reaches a contradiction, consider the case $V_\partial \geq 5$. We have shown that

$$\lim_{\kappa \to \widetilde{\kappa}} [(F - 2V_{\text{int}})\pi - (A_\kappa + \theta_\kappa(\mathcal{V}_\partial) - \pi N_g)] = (F - 2V_{\text{int}})\pi - \sum_{j} \lim_{\kappa}(\pi - \gamma_j(\kappa))$$

$$> (V_\partial - 2)\pi - \pi(V_\partial - 2T - 2) > 2T\pi.$$

On the other hand, since the labels $\kappa \in \mathcal{L}$ are β-superpacking labels, $\Sigma_\kappa \geq 0$, so (4) implies

$$(F - 2V_{\text{int}})\pi - (A_\kappa + \theta_\kappa(\mathcal{V}_\partial) - \pi N_g) < 2T\pi.$$

This and the previous inequality contradict one another. Similar arguments apply to the cases $V_\partial = 3$ and $V_\partial = 4$. Therefore, Case I fails and $L \neq K$.

Case II: Suppose $\emptyset \subsetneqq L \subsetneqq K$: There are two possibilities: Suppose L is simply connected; that is, its boundary edges form one simple closed curve. Write Φ_L for the restriction of Φ_K to L. We may apply the reasoning of Case I to L as follows: L triangulates a close topological disc. Each face of L contains an interior vertex — namely a vertex of $\mathcal{B}$; and because $\mathcal{B}$ is an edge-component of $\mathcal{W}$, no vertex of ∂L is in $\mathcal{W}$. Thus, $\widetilde{\kappa}(v) < \infty$, $\forall v \in \partial L$. Indeed, the restriction of $\widetilde{\kappa}$ to ∂L is a proper boundary label *vis-a-vis* L and Φ_L. Consider the family $\mathcal{L}'$ of all β-superpacking labels for (L, Φ_L) which agree with $\widetilde{\kappa}$ on ∂L. This essentially contains the restrictions to L of the labels in $\mathcal{L}$. Suppose, for instance, that $\kappa \in \mathcal{L}$. Then its restriction to L will be a β-superpacking label there. Adjust its values at the boundary vertices of L upward to match $\widetilde{\kappa}$; by Lemma 2(b), the resulting label remains a β-superpacking label for (L, Φ_L), so we get a label in $\mathcal{L}'$. Using these labels, we can conclude that $\sup_{\kappa \in \mathcal{L}'}\{\kappa(v)\} = 0$ for all $v \in \mathcal{B}$. This puts us precisely in the situation covered in Case I above, namely, the supremum over $\mathcal{L}'$ is zero for all interior vertices of L. As we have seen, this results in a contradiction.

If L is not simply connected, then there clearly exists at least one simple closed edge path in ∂L which does not separate L from ∂K. We will call such a path a **loop**, and we may assume without loss of generality that its interior lies in $\mathcal{V} \backslash \mathcal{W}$ (otherwise it would surround an edge-component of $\mathcal{W}$ which we could consider in place of $\mathcal{B}$). Let $\mathcal{H}$ be an edge-component of vertices from $L \backslash \mathcal{B}$ containing a loop. Note that $\mathcal{H}$ can have no end by the same argument used in Case I. Hence, $\mathcal{H}$ will consist of one or more loops, say $l_1, \cdots, l_n$, joined by edges or vertices. The fact that $\mathcal{B}$ is edge-connected implies that $\mathcal{H}$ has at least one loop l_j which we might call an end loop, that is, it has at most one vertex with a neighbor in $\mathcal{H} \backslash l_j$. Assume this loop is l_1, and if it has a vertex neighboring a point of $\mathcal{H} \backslash l_1$, let that vertex be denoted v_1. In particular, notice that $l_1 \backslash \{v_1\} \subset \mathcal{V}_{\text{int}}$ and every neighbor of $v \in l_1, v \neq v_1$, which lies in $L \backslash l_1$ is in $\mathcal{B}$.

Let M be the complex formed by the vertices that are in l_1, the edges connecting them, and all the edges, vertices, and faces of K which lie in the interior of l_1. We reach a contradiction in Case II by analyzing the restriction of $\widetilde{\kappa}$ (and β and Φ_K) to M. Note that every vertex v of M, except possibly v_1, satisfies $1 < \widetilde{\kappa}(v) < \infty$, so (the restriction of) $\widetilde{\kappa}$ is a β-packing label for (M, Φ_M). The total angle sum of M with this label ($\sum_{v \in V^M} \theta_{\widetilde{\kappa}}(v)$) is less than πF^M, where F^M is the number of faces of M. Hence,

$$\sum_{v \in \mathcal{V}^M_{\text{int}}} 2\pi\beta(v) + \theta_{\widetilde{\kappa}}(\mathcal{V}^M_\partial) \leq \pi F^M. \tag{5}$$

Each boundary edge of E belongs to one face outside M, and the third vertex of that face belongs to $\mathcal{B}$ (otherwise the face would be included in M). Refer to the two angles of that face next to M as the **contiguous** angles. By Lemma 2(e), these sum to at most $\pi/2$. Thus, the sum of all contiguous angles is at most $\frac{\pi}{2} E^M_\partial = \frac{\pi}{2} V^M_\partial$. We may consider this sum in another way: Assuming $v \neq v_1$, any neighboring vertex $w \notin M$ lies in $\mathcal{B}$, so $\widetilde{\kappa}(w) = \infty$, while $1 < \widetilde{\kappa}(v) < \infty$. This means that the angles at v in faces not in M will go to zero, with the possible exception of the two contiguous ones, so the angle sum at v relative to M is 2π minus the sum of the contiguous angles. Summing the contiguous angles by vertex gives

$$\begin{aligned}\theta_{\widetilde{\kappa}}(\mathcal{V}_\partial^M) &\geq \theta_{\widetilde{\kappa}}(\mathcal{V}_\partial^M \backslash \{v_1\}) \\ &\geq \sum_{v \in (\mathcal{V}_\partial^M \backslash \{v_1\})} [2\pi\beta(v) - \frac{\pi}{2}] \\ &\geq (\sum_{v \in (\mathcal{V}_\partial^M \backslash \{v_1\})} 2\end{aligned}$$

$$pi\beta(v)) - \frac{\pi}{2} V_\partial^M$$

$$\geq 2\pi(V_\partial^M - 1) - \frac{\pi}{2} V_\partial^M.$$

With (5) this implies

$$F^M \geq 2(V^M) - \frac{1}{2} V_\partial^M - 2.$$

The fact that $\chi(|M|) = 1$, $3F^M = 2E^M - E_\partial^M$, and $E_\partial^M = V_\partial^M$, gives $V^M = 1 + \frac{1}{3}(E^M + E_\partial^M)$. With the inequality above this leads to

$$1 = V^M - E^M + F^M \geq V^M - E^M + 2V^M - \frac{1}{2} V_\partial^M - 2 = 1 + \frac{1}{2} V_\partial^M.$$

This contradiction proves that Case II fails.

We have shown that $L = K$ and $\emptyset \subsetneqq L \subsetneqq K$ lead to contradictions, so $L = \emptyset$. This implies $\mathcal{B} = \emptyset$, so $\mathcal{W} = \emptyset$. Thus $\widetilde{\kappa}(v) < \infty$ for all vertices $v \in K$, and we have proved that $\widetilde{\kappa}$ is a β-packing label for (K, Φ_K), as intended.

To establish uniqueness, suppose $\widehat{\kappa}$ is another β-packing label for (K, Φ_K) with $\widehat{\kappa}(w) = \widetilde{\kappa}(w) = g(w)$, $\forall w \in \partial K$. The shortages $\Sigma_{\widetilde{\kappa}} = \Sigma_{\widehat{\kappa}} = 0$, so (4) implies

$$A_{\widetilde{\kappa}} + \theta_{\widetilde{\kappa}}(\mathcal{V}_\partial) - \pi N_g = (F - 2V_{\text{int}})\pi = A_{\widehat{\kappa}} + \theta_{\widehat{\kappa}}(\mathcal{V}_\partial) - \pi N_g. \tag{6}$$

Since $\widehat{\kappa} \in \mathcal{L}$, we have $\widehat{\kappa} \leq \widetilde{\kappa}$: By Lemma 3(a), this implies $\theta_{\widehat{\kappa}}(\mathcal{V}_\partial) \leq \theta_{\widetilde{\kappa}}(\mathcal{V}_\partial)$, while by Lemma 3(d) it implies $A_{\widehat{\kappa}} \geq A_{\widetilde{\kappa}}$. Together with (6), we conclude that $A_{\widehat{\kappa}} = A_{\widetilde{\kappa}}$. By the strict monotonicity in Lemma 3(d), this equality implies that $\widehat{\kappa}(v) = \widetilde{\kappa}(v)$ for every interior vertex v. This proves uniqueness and concludes the proof of Proposition 1. $\square$

Remark 4: Some very useful global monotonicity results about circle packings follow in a straightforward way from arguments such as those use above in conjunction with the local monotonicities from §3. We won't need these here, but believe that an example might be of interest to the reader:

Global Monotonicity Lemma. *Suppose that $\beta_1 \subseteq \beta_2$ are two branch structures for a labeled triangulation (K, Φ_K). If κ_1 is a β_1-superpacking label for (K, Φ_K), κ_2 is a β_2-subpacking label for (K, Φ_K), and $\kappa_1(w) \leq \kappa_2(w)$ at all boundary vertices $w \in K$, then $\kappa_1(v) \leq \kappa_2(v)$ for all vertices of K.*

In studies of tangency circle packings, boundary value problems and monotonicity have played a central role. With this lemma and Propostion 1, many results may now be extended to more general overlapping and/or branched packings — for instance, the Discrete Schwarz Lemma [BSt1], [vE], the Discrete Distortion Lemma [DS], the Uniformization Theorem and "type" conditions [BSt2], and so forth. Details will appear elsewhere.

Remark 5: The proof here suggests a convenient numerical algorithm for approximating the solution label $\widetilde{\kappa}$; this is a based on Thurston's iterative algorithm and is similar to what's implemented in the software used to generate the pictures for this paper. It's a simple "relaxation" method:

> Suppose that K has N interior vertices $v_0, \cdots, v_{N-1}$. Start with any β-superpacking label κ_0 which is equal to g on the boundary vertices. (In fact, one can start with $\kappa_0(v_j) \equiv 1, j = 1, \cdots, N$.) Inductively define a sequence of β-superpacking labels as follows: Assuming κ_n has been chosen, define κ_{n+1} by adjusting the curvature at v_j, where $j = n (\text{mod } N)$. In particular, adjust the curvature so that the angle sum at v_j is $2\pi\beta(v_j)$. The labels κ_n converge to $\widetilde{\kappa}$ as $n \to \infty$.

In fact, monotonicity gives a great computational advantage here: the κ_n increase monotonically at each vertex to $\widetilde{\kappa}$ and the area A_{κ_n} of $K(\kappa_n)$ decreases monotonically to the area $A_{\widetilde{\kappa}}$. Since one can compute $A_{\widetilde{\kappa}}$ directly from (K, Φ_K), it provides a measure of complexity in the computation, allowing, e.g., estimates of the rate of convergence. For more details, see comments of Bowers [Application 4, Bo] on the tangency case.

6. Circle Packings

We must now see how one moves between labeled complexes and circle packings. First, a circle packing Q on the sphere is said to **lie in** the unit disc if its interior circles lie in the open disc and its boundary circles intersect the disc in cycles having nonnegative curvature. The corresponding label, denoted κ_Q, is defined by $\kappa_Q(v) = \text{curv}(c_v)$.

Proposition 2. *A label κ is a β-packing label for (K, Φ_K) if and only if $\kappa = \kappa_Q$ for some circle packing Q for (K, Φ_K) which lies in the unit disc and has branch set β.*

Proof. Necessity: Given Q, the faces formed, *in situ*, by overlapping triples of its cycles comprise a complete set of the geometric faces which induce the hyperbolic structure of $|K(\kappa_Q)|$. In particular, it is immediate that the angle sums meet the conditions for a β-packing label, $\beta = \mathrm{br}\,(Q)$. Also note that κ_Q is unchanged if Q is replaced by its image under a conformal automorphism of Δ.

Sufficiency: Reversing this process takes a little more work. However, given a β-packing label κ, the following straightforward construction process using geometric faces yields Q:

Each face f_j of $|K(\tau)|$ is isometrically isomorphic to a geometric face F_j in Δ (see §2), and whenever a face f_i shares an edge with f_j, one may (by applying an automorphism, if necessary) ensure that F_i shares the corresponding edge with F_j. Now, one obtains a (branched) immersion of $\xi : |K(\kappa)| \to \Delta$ by fixing one initial face, say F_0, and then laying out successively contiguous faces until all have been placed. The "packing" condition — that angle sums for all interior vertices are integral multiples of 2π — is necessary and sufficient to ensure that this yields a well-defined immersion. (This is a discrete monodromy argument, an easy extension of the proof for Theorem 1 of [BSt].) The immersion is locally isometric except at the branch points and is unique up to placement of F_0, i.e., up to conformal automorphisms of Δ.

As for the circles, each geometric face contains arcs of the three cycles which determined it; these have the curvatures specified in κ and intersect with the angles specified in Φ_K. It is elementary to verify that these arcs in the various immersed faces will piece together to form curves in Δ having the appropriate constant curvatures of κ. In other words, these piece together to form cycles, one for each vertex of K. Their parent circles on S^2 comprise the desired circle packing Q, and it is evident that $\kappa_Q = \kappa$. $\square$

Remark 6: The conditions we have laid down for a circle packing are local: they are met if each circle C_v intersects (with correct overlaps) the circles C_u for all vertices u of K neighboring v. Typically (and always in the case of branched packings), C_v will also intersect (or overlap, enclose, etc.) circles C_w for vertices w *not* neighboring v. This is perfectly fine, though it may make a picture of the packing difficult to interpret.

Remark 7: We have made no assumptions about the boundary curvatures coded in κ. If, for instance, κ is greater than 1 at all boundary vertices, then Q will lie entirely within Δ — no circles reach the northern hemisphere. On the other hand, any boundary vertex w with $\kappa(w) \leq 1$ will lead to a hypercycle in Q overlapping the northern hemisphere with angle $\arccos(\kappa(w))$. The latter situation, of course, is what we will face in the next section, and it leads to faces of K which, when laid out in the disc, intersect the unit circle.

7. The Extended Andreev-Thurston Theorem

As observed in the Introduction, one might expect a variety of conformally inequivalent circle packings P for a given labeled triangulation (τ, Φ) if the requirement that carr(P) triangulate S^2 has been removed. Each P, induces a natural simplicial map from τ to carr(P); and as a discrete, open mapping of a topological sphere into S^2, this is necessarily a branched covering with total branch order an even integer. It is this opportunity for branching which accounts for the added variety, and in order to classify the possible packings, we need necessary and sufficient conditions on the branch sets and an appropriate uniqueness statement. We don't get the complete answer here, but our methods yield an important special class. (The terminology will be justified in the next section.)

Definition. *Given a labeled triangulation (τ, Φ) of S^2, a set $\beta \subset \tau^{(0)}$ will be called a* **polynomial branch structure** *if: (1) β contains an even number $2n$ of points, counting repetitions; (2) one vertex v_∞ of τ occurs n times in β; (3) the restriction β' of β to $K = \tau \backslash star(v_\infty)$ contains only interior vertices and forms a branch structure for (K, Φ_K); and (4) no subgraph of the type given in Figure 4(a) occurs in (τ, Φ) wherein v belongs to β.*

We extend the Andreev-Thurston Theorem to include prescribed polynomial branch structures; the original theorem is the case $\beta = \emptyset$.

Main Theorem. *Let τ be a triangulation of S^2 which is not simplicially equivalent to a tetrahedron, let $\Phi : \tau^{(1)} \longrightarrow [0, \pi/2]$ be given, and suppose that*

> *(A). If $e_1, e_2, e_3 \in \tau^{(1)}$ form a closed loop on S^2 and if $\sum_{i=1}^{3} \Phi(e_i) \geq \pi$, then (e_1, e_2, e_3) are the three edges of a two-simplex of τ.*
>
> *(B). If $e_1, e_2, e_3, e_4 \in \tau^{(1)}$ form a closed loop on S^2 and if $\sum_{i=1}^{4} \Phi(e_i) \geq 2\pi$, then (e_1, e_2, e_3, e_4) form the boundary of the union of two adjacent two-simplexes of τ.*

If β is a polynomial branch structure for the labeled triangulation (τ, Φ), then there exists a circle packing P for (τ, Φ) for which $br(P) = \beta$. Moreover, P is unique up to conformal automorphisms of S^2.

Proof. Of course, the heart of the proof is the boundary value problem we solved in §5. Write v_∞ for the vertex at which half of the branching of β occurs, write $K = \tau \backslash \text{star}(v_\infty)$, let β' be the restriction of β to K, and let g_Φ be the boundary label on K obtained from

Φ. Before applying Proposition 1, we must verify that g_Φ satisfies the conditions (a), (b), and (c) required of "proper" boundary labels in §2.

Condition (a): If an interior vertex v of K has three critical edges, then hypothesis (B) of the Andreev-Thurston Theorem implies that τ is a double tetrahedron. Moreover, the packing for (τ, Φ) turns out to be degenerate: namely, the same circle is identified with both v_∞ and v. We describe these packings in §9, so for our proof here we may assume that condition (a) is met.

Condition (b): Suppose an interior vertex v has two critical edges. If they were non-contiguous, there would be a closed 4-edge loop through v and v_∞ in τ which does not bound two faces, contradicting hypothesis (B). Therefore, the critical edges must belong to a single face. If the overlap for the remaining edge is not positive, then we have an instance of the "extraneous tangency" noted in §2 with subgraph Figure 4(a). As pointed out there, this subgraph may be switched with that of Figure 4(b) without effect: that is, the associated four circles would be equally well described by either subgraph. By switching to the subgraph of Figure 4(b), we are in effect replacing (τ, Φ) by an altered labeled complex, and vertex v in the figure becomes a boundary vertex of the new K. One easily verifies that the altered data continue to satisfy the hypotheses of our extended Andreev-Thurston Theorem. Therefore, we assume henceforth that any such extraneous boundary tangencies have been removed by such a "switch" — that is, we may assume that condition (b) is met.

Condition (c): Suppose e is an interior edge of K without proper data. That means that its endpoints w_1, w_2 lie in ∂K and that the closed 3-edge loop through w_1, w_2, and v_∞ has overlaps summing to π or more. But because e is interior, this loop would not bound a face, contradicting hypothesis (A). Therefore, condition (c) is met.

We conclude that g_Φ is a proper boundary label, and we apply Proposition 1 to get a β'-packing label $\widetilde{\kappa}$ with the boundary values of g_Φ. This is associated by Propostion 2 with a circle packing Q for (K, Φ_K) lying in Δ. We augment this by adding the equator as the circle C_∞ for v_∞ (boundary of the northern hemisphere), giving the new circle packing $P = Q \cup \{C_\infty\}$ on S^2. It remains to verify that P is indeed a circle packing for (τ, Φ). This reduces to two questions:

(1). Does the circle C_w, for each $w \in \partial K$, intersect the equator with the correct overlap, namely, with overlap $\Phi(\langle w, v_\infty \rangle)$? This is automatically true: $\widetilde{\kappa}$ is a packing label with boundary values g_Φ, meaning that $C_w \cap \Delta$ has curvature $g_\Phi(w) = \arccos \Phi(\langle w, v_\infty \rangle)$, meaning that C_w overlaps the northern hemisphere with angle $\Phi(\langle w, v_\infty \rangle)$.

(2). Does $\operatorname{br}(P) = \beta$? This reduces to checking that v_∞ is a branch point of order n. For this we revisit the (branched) immersion $\xi : |K(\widetilde{\kappa})| \to \Delta$ constructed in §6. Recall that this is defined on the interior of K, where it is locally isometric (aside from isolated branch points). We need to extend ξ to ∂K and show that its restriction there maps ∂K in an $n+1$-to-1 oriented fashion onto $\partial\Delta$, the unit circle. This could be done abstractly, verifying completeness of $|K(\widetilde{\kappa})|$ and applying the Killing/Hopf Theorem, suitably extended to account for the finite set of branch points. However, our considerations are basically topological, and we can get by with the obvious constructive extension of ξ.

Namely, the vertices of ∂K are assigned curvatures in $[0, 1]$ by $\widetilde{\kappa}$. Therefore, if f is the face of K associated with an edge $e \in \partial K$, then in a geometric face F constructed using the data for f (see §2 and Figure 5), the edge corresponding to e will be an arc of the unit circle. It is clear that in the laying out of faces in §6 to define ξ, we may as well include these boundary edges; ξ maps each in an oriented way to an arc of the unit circle. At the end of the day, ξ is a continuous, orientation preserving map of ∂K to the unit circle.

Now, this extended map ξ, treated as a topological map from K to the closed unit disc, is continuous, open, discrete, and proper. Since ξ has n branch points on the interior of K, the topological argument principle implies that ξ maps ∂K $n+1$-to-1 onto $\partial\Delta$, as we intended to show. It is now evident that the circles of P associated the boundary vertices of K wrap $n+1$ times about C_∞, so C_∞ is a branch point of order n. This proves that $\operatorname{br}(P) = \beta$.

We have now shown existence of a circle packing P for (τ, Φ) with branch set β. The final issue is uniqueness, and we again refer to §6. Suppose P' is another such packing. Normalizing so that v_∞ is associated with the northern hemisphere, the remaining circles comprise a packing Q' for (K, Φ_K) intersecting Δ in a family of cycles $\{c_v : v \in K\}$. Because we have already eliminated any chance of extraneous boundary tangencies, Q' "lies in" Δ as described in §6 and provides a proper label κ for K with $\kappa(w) = \cos(\Phi(\langle w, v_\infty \rangle)) = g_\Phi$ for all $w \in \partial K$. By Proposition 1, $\kappa = \widetilde{\kappa}$, and the arguments in §6 imply that Q' and Q are equivalent up to conformal automorphisms of Δ. Throwing the northern hemisphere into each, we conclude that P' and P are equivalent up to conformal automorphisms of S^2. This gives uniqueness, and the proof is complete. $\square$

The hypotheses of this theorem are very tight — essentially necessary and sufficient. We discuss the exceptions and additional details in §9.

8. General Branching

Suppose $P_{(\tau,\Phi)}$ denotes the maximal circle packing and P is another circle packing for a labeled triangulation (τ, Φ) of the sphere. The simplicial mapping

$$f : \text{carr}\,(P_{(\tau,\Phi)}) \longrightarrow \text{carr}\,(P) \tag{7}$$

carrying the center of each circle of $P_{(\tau,\Phi)}$ to the center of the corresponding circle of P is in many ways a discrete analogue of a classical rational function from S^2 to S^2 — in particular, it will have the topological properties and branch structure of a rational map. (As an example, the map from Figure 2(a) to 2(b) would be a discrete analogue of $z \mapsto z^2$.) Polynomials may be characterized among the rationals (up to conformal automorphisms of domain and/or range) by noting that precisely half of their branching occurs at a single point. This is what motivated our earlier definition of "polynomial" branch structures. With the extension of the Andreev-Thurston Theorem, therefore, we have established existence and uniqueness for polynomial-like discrete rational maps. However, for general branching, both existence and uniqueness remain open:

Question. *Given a labeled triangulation (τ, Φ) of the sphere, what are necessary and sufficient conditions on a subset $\beta \subset \tau^{(0)}$ so that there exists a circle packing P for (τ, Φ) with $\text{br}\,(P) = \beta$? If a branched packing exists, is it unique up to conformal automorphisms?*

Of course, card (β) must be even, and certain of the previous necessary conditions will evidently carry over to this case. The real crux of the matter, however, is how to construct the packings. Here is a nontrivial example which will help frame the issues:

Example: We construct a branched, tangency circle packing P with a total of 42 circles and 12 simple branch points. It is easiest to start by describing the associated maximal packing $P_{(\tau,\Phi)}$, displayed in Figure 9. The triangulation τ, also shown there, is the dual triangulation to the familiar soccer-ball (a.k.a, Buckminsterfullerene, carbon C_{60}) cell decomposition of S^2. This is a tangency pattern, meaning that we have prescribed $\Phi \equiv 0$.

This pattern really originated with the regular dodecahedron, and we may construct $P_{(\tau,\Phi)}$ quite explicitly using Schwarz triangles (see [Cox]): Figure 10(a) is a (2,3,5)-Schwarz triangle t (i.e., spherical angles are $\pi/2, \pi/3, \pi/5$). Note that we have marked t with circular arcs centered at two of its vertices. Apply the classical Schwarz reflection process, which generates a surface by repeatedly reflecting t through its edges and attaching the reflected copies. One can show that for this particular t the process results in 120 copies of t having mutually disjoint interiors and tiling the sphere. Indeed, ten copies of t fit together as in

FIGURE 9. The maximal packing $P_{(\tau,\Phi)}$.

Figure 10(b) to form one petagonal face of a regular (spherical) dodecahedron, and 12 of these faces paste together to complete the tiling of S^2. The circular arcs carried along with the copies of t will piece together to give $P_{(\tau,\Phi)}$.

FIGURE 10. The Schwarz triangle t.

Constructing our branched packing is only slightly more difficult. The branch set β we choose to prescribe consists of the 12 vertices of τ having degree 5, each occuring once (i.e., simple branch points) — note that this is *not* a polynomial branch structure. The requisite branched packing P can again be constructed using Schwarz triangles: Consider the (2,3,5/2)-Schwarz triangle T pictured in Figure 11(a), with associated circular arcs. The Schwarz reflection process is applied to T. In particular, 4 reflected copies are required about vertex A and 6 about vertex B to complete their neighborhoods. However, 10 reflections are required about vertex C; this is where the branching occurs, since these 10 wrap twice around C. The reflection process continues until any further reflected copies are repeats; one can show that this yields a closed surface $\mathcal{S}$ consisting of 120 copies of T.

FIGURE 11. The Schwarz triangle T.

And where is P? Well, the circular arcs on the copies of T again piece together to form 42 complete circles, giving P. We leave it to the reader to verify that the combinatorics of P are correct. (Hint: note that by mapping the arcs α_j of t to η_j of T, $j = 1, 2$, respectively, and propagating this to all 120 triangles in domain and range *via* reflection, one obtains a tangency preserving map of $P_{(\mathcal{T},\Phi)}$ onto P.)

We have displayed P in Figure 12(a) with the branch circles shaded. We realize that the picture is rather difficult to interpret, primarily because $\operatorname{carr}(P)$ is a 7-sheeted covering of S^2. This follows using Euler characteristics and the Riemann-Hurwitz relation or by noting the decomposition of T into 7 copies of t shown in Figure 11(b). The branch values of P occur at the centers of the 12 pentagons of a regular dodecahedron, and we show one of these branched flowers in Figure 12(b). By studying the symmetry and geometry of this packing, Jack Quine has shown that the associated simplicial map f of (7) mimicks a classical 7-degree rational map of the form

$$R(z) = \frac{z^2(az^5 + b)}{(cz^5 + d)}.$$

FIGURE 12. The branched packing P.

This example is, of course, very special. Nevertheless, it illustrates the problem with non-polynomial branching — namely, projecting to the hyperbolic setting doesn't work. If $\mathrm{br}\,(P)$ has total order $2n$, then $\mathrm{carr}\,(P)$ is an n-fold branched covering of S^2. When the branching is polynomial, a circle with branch order n is simultaneously on all n "sheets", and by making it into the northern hemisphere, the other circles are forced into Δ. However, in the Example above, no circle is on more than 2 sheets. Any projection to Δ will omit some circles and/or lead to some negatively oriented cycles. What then?

Of course, aesthetically one would like to avoid moving to hyperbolic geometry at all: one would hope for intrinsic arguments which use the spherical metric and are $PSL(2, \mathbf{C})$-invariant. By-and-large, iterative packing algorithms in fact seem to work in spherical geometry. However, to our knowledge, no one knows how to prove convergence, even in the polynomial case. The difficulty seems to lie with the loss of monotonicity of angle sums with respect to radii in the spherical metric.

Putting existence aside, it does seem likely that a branched packing of (τ, Φ) with a given branch set is unique (up to conformal automorphisms), but again, there is as yet no proof in the non-polynomial case.

9. Closing Comments and Examples

The hypotheses of the Andreev-Thurston Theorem are amazingly tight, which accounts for the details we needed to follow in our proof. Here are some of the reasons:

(1). τ a tetrahedron: This case was omitted because certain overlap assignments Φ satisfying hypotheses (A) cannot be realized by circle packings (hypothesis (B) is vacuous). To illustrate, if all overlaps specified in Φ are equal to a constant angle ψ, then a circle packing for (τ, Φ) exists iff $0 \leq \psi \leq \arccos(1/3)$. (See [Thm 2, M].)

(2). The bound $\pi/2$ on overlaps: Without this, Lemma 1 fails and one can easily construe specifications of curvatures and overlaps which cannot be realized. The technique of varying labeled structures to reach a solution is lost, "extraneous overlaps" may be created, and the connection between the circles of a packing and the faces of a polyhedron can be broken. Rivin [R] has proved existence of circle packings with overlaps up to π, but the hypotheses involve existence of appropriate "dual weights" to Φ.

(3). Hypotheses (A) and (B): As noted in Remark 2, a triple of circles whose overlaps sum to less than π will form two (curvilinear) triangular interstices on S^2 — additional circles of a packing may lie in one or both of these. If, on the other hand, the overlaps sum to π or more, there is only one interstice to accomodate additional circles; the triple itself must therefore form a face, yielding hypothesis (A). Similar considerations explain hypothesis (B).

(4). Non-triangulations: If τ is a finite graph on S^2, not necessarily a triangulation, and Φ satisfies hypotheses (A) and (B) of the theorem, then following Thurston, one may complete τ to a triangulation (specifying, e.g., overlaps zero for the added edges), obtain a circle packing for the triangulation, and then discard those circles not associated with τ. Of course, uniqueness is lost.

(5). Degenerate Packings: There is a class of degenerate circle packings we'll call **equatorial**: Suppose τ is the "double suspension" of a simple closed edge-chain $\langle w_1, \cdots, w_n \rangle$; that is, there are distinguished vertices v_N and v_S, each sharing an edge with every w_j. If $\Phi(\langle v_N, w_j \rangle) \equiv \Phi(\langle v_S, w_j \rangle) \equiv \pi/2, \;\; j = 1, \cdots, n$, then, after an appropriate normalization, v_N will correspond to the northern hemisphere, v_S to the southern, and the circles for w_1, w_2, and w_3 will form a closed chain centered on the equator. The degeneracy occurs because the equator is the circle for both v_N and v_S.

Equitorial packings for $n = 3$ do occur as degenerate packings in the Andreev-Thurston setting. In this case τ is a double tetrahedron, and if the overlaps among the w_j satisfy

$$\Phi(\langle w_1, w_2 \rangle) + \Phi(\langle w_2, w_3 \rangle) + \Phi(\langle w_3, w_1 \rangle) < \pi,$$

then (τ, Φ) satisfies the hypotheses of the Andreev-Thurston Theorem. We note that this degeneracy was explicitly eliminated in Andreev's original statement, since it is associated with a degenerate polyhedron; see §2 of [H]. Equitorial packings for $n > 3$ are not at issue, since they fail hypothesis (B). Note that they also fail to be unique.

References

[A1] E. M. Andreev, *Convex polyhedra of finite volume in Lobacevskii space*, Mat. Sbornik (Russian) **83** (1970), 256–260; Math. USSR Sbornik (English) **12** (1970), 255–259.

[A2] E. M. Andreev, *Convex polyhedra in Lobacevskii space*, Mat. Sbornik (Russian) **81 (123)** (1970); Math. USSR Sbornik (English) **10** (1970), 413–440.

[BSt1] Alan F. Beardon and Kenneth Stephenson, *The Schwarz-Pick lemma for circle packings*, Ill. J. Math. **141** (1991), 577–606.

[BSt2] Alan F. Beardon and Kenneth Stephenson, *The uniformization theorem for circle packings*, Indiana Univ. Math. J. **39** (1990), 1383–1425.

[Bo] Philip L. Bowers, *The upper Perron method for labeled complexes*, Proc. Camb. Phil. Soc. **114** (1993), 321–345.

[BoS] Philip L. Bowers and Kenneth Stephenson, *Circle packings in surfaces of finite type: An in situ approach with applications to moduli*, to appear, Topology.

[Br] Robert Brooks, *On the deformation theory of classical Schottky groups*, Duke Math. J. **52** (1985), 1009–1024.

[C] Ithiel Carter, *PhD Thesis, UCSD* (1989).

[CFP] J. W. Cannon, W. J. Floyd, and Walter Parry, *Squaring rectangles: the finite Riemann Mapping Theorem*, Geometry Center preprint (1993).

[Cox] H. S. M. Coxeter, *Regular Polytopes, 3rd edition*, Dover, New York, 1973.

[D1] Tomasz Dubejko, *Branched circle packings, discrete complex polynomials, and the approximation of analytic functions*, PhD Thesis, University of Tennessee, August 1993.

[D2] Tomasz Dubejko, *Branched circle packings and discrete Blaschke products*, preprint 11/93.

[DS] Tomasz Dubejko and Kenneth Stephenson, *The branched Schwarz lemma: a classical result via circle packing*, preprint 1/93.

[H] Craig D. Hodgson, *Deduction of Andreev's theorem from Rivin's characterization of convex hyperbolic polyhedra*, Topology 90, Proceedings of Research Semester in Low Dimensional Topology at O. S. U., de Gruyter Verlag, New York.

[Kb] Koebe, *Kontaktprobleme der Konformen Abbildung*, Ber. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Kl. **88** (1936), 141–164.

[FH] Formann, Hagerup, Haralambides, Kaufmann, Leighton, Simvonis, Welzl, and Woeginger, *Drawing graphs in the plane with high resolution*, 31st Symp. Found. of Comp. Sci. (1990), 86–95.

[MR] Al Marden and Burt Rodin, *On Thurston's formulation and proof of Andreev's theorem*, Computational Methods and Function Theory, Proceeding, Valparaiso 1989, Lecture Notes in Mathematics, vol. 1435, Springer-Verlag, Berlin, Heidelberg, New York, Tokyo, pp. 103–115.

[M] J. G. Mauldon, *Sets of equally inclined spheres*, Canadian J. Math. **14** (1962), 509–516.

[MT] Gary L. Miller and William Thurston, *Separators in two and three dimensions*, Baltimore, May 1990, Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, ACM, pp. 300-309.

[Mo] John W. Morgan, *On Thurston's uniformization theorem for three-dimensional manifolds*, The Smith Conjecture, Edited by Hyman Bass and John W. Morgan, Academic Press, New York, pp. 37–126.

[Mh] Bojan Mohar, *Circle packings of maps in polynomial time*, preprint.

[R] I. Rivin, *On geometry of convex polyhedra in hyperbolic 3-space*, Ph.D. Thesis, Princeton (1986).

[RS] Burt Rodin and Dennis Sullivan, *The convergence of circle packings to the Riemann mapping*, J. Differential Geometry **26** (1987), 349–360.

[T1] William Thurston, *The Geometry and Topology of 3-Manifolds*, preprint, Princeton University Notes.

[T2] William Thurston, *The finite Riemann mapping theorem*, Invited talk, An International Symposium at Purdue University on the occasion of the proof of the

Bieberbach conjecture, March 1985.

[vE] Jeff Van Eeuwen, *The discrete Schwarz-Pick lemma for overlapping circles*, PAMS (to appear).

Phil Bower, Department of Mathematics, Florida State University, Tallahassee, FL 32306 bowers@math.fsu.edu

Kenneth Stephenson, Department of Mathematics, University of Tennessee, Knoxville, TN 37996-1300 kens@math.utk.edu

Version: 17 November 1994