

Chapter 5

Combinatorics Encoding Geometry: The Legacy of Bill Thurston in the Story of One Theorem



Philip L. Bowers

Abstract This chapter presents a whirlwind tour of some results surrounding the *Koebe–Andre’ev–Thurston Theorem*, Bill Thurston’s seminal circle packing theorem that appears in Chapter 13 of *The Geometry and Topology of Three-Manifolds*.

Keywords Circle packing · Uniformization · Circle domain · Discrete conformal geometry · Conformal type · Inversive distance · Polyhedron

AMS Codes 52C26, 57M20

5.1 Introduction

Bill Thurston was the most original and influential topologist of the last half-century. His impact on the discipline of geometric topology during that time is unsurpassed, and his insights in the topology and geometry of three-manifolds led finally to the resolution of the most celebrated problem of topology over the last century—The Poincaré Conjecture. He made fundamental contributions to many sub-disciplines within geometric topology, from the theory of foliations on manifolds to the combinatorial structure of rational maps on the two-sphere, and from geometric and automatic group theory to classical polyhedral geometry. Of course his foundational work on three-manifolds, first laid out in his courses at Princeton in the late 1970s, compiled initially as a Princeton paper-back monograph inscribed by Bill Floyd and available upon request as *The Geometry and Topology of Three-Manifolds* (GTTM) [73], and maturing as the famous *Thurston Geometrization Conjecture* of the early 1980s, was the driving force behind the development of geometric topology for the next thirty years. The final confirmation of the Geometrization Conjecture

P. L. Bowers (✉)

Department of Mathematics, Florida State University, Tallahassee, FL, USA

e-mail: bowers@math.fsu.edu

© Springer Nature Switzerland AG 2020

K. Ohshika, A. Papadopoulos (eds.), *In the Tradition of Thurston*,

https://doi.org/10.1007/978-3-030-55928-1_5

by Giorgi Perelman using the flow of Ricci curvature, following a program that had been introduced by Richard Hamilton, is one of the crown jewels of twentieth century mathematics.

Thurston marks a watershed in the short history of topology,¹ a signpost, demarcating topology before Thurston, and topology after Thurston. This is evidenced not only in the fabulous results he proved, explained, and inspired, but even more so in how he taught us to do mathematics. Topology before Thurston was dominated by the general and the abstract, entrapped in the rarified heights that captured the mathematical world in general, and topology in particular, in the period from the 1930s until the 1970s. Topology after Thurston is dominated by the particular and the geometric, a throwback to the nineteenth-century, having much in common with the highly geometric landscape that inspired Felix Klein and Max Dehn, who walked around and within Riemann surfaces, knew them intimately, and understood them in their particularity. Thurston's vision gave a generation of topologists permission to get their collective hands dirty by examining in great depth specific structures on specific examples.

One of the organizing principles that lies behind Thurston's vision is that geometry informs topology, and that the non-Euclidean geometry of Lobachevski, Bolyai, and Beltrami in particular is systemic to the study of topology. Hyperbolic geometry permeates topology after Thurston, and it is hyperbolic geometry that becomes the common thread of the present chapter. This will be seen in the interrelated studies presented here. All to varying degrees are due to the direct influence of Bill Thurston and his generalization of the earlier results of Koebe and Andre'ev. All involve hyperbolic geometry in some form or influence, and even further all illustrate how combinatorics encodes geometry, another of the principles that underlies Thurston's vision. To my mind, the proposition that *combinatorics encodes geometry, which in turn informs topology* has become a fundamental guiding motif for topology after Thurston. I offer this chapter as a celebration of Bill Thurston's vision and his immense influence over our discipline.

5.1.1 *An Introductory Overview*

The Koebe–Andre'ev–Thurston Theorem represents a rediscovery and broad generalization of a curiosity of Paul Koebe's from 1936, and has an interpretation that recovers a characterization of certain three-dimensional hyperbolic polyhedra due to E.M. Andre'ev in two papers from 1970. This theorem is the foundation stone of the discipline that has been dubbed as *discrete conformal geometry*, which itself has been developed extensively by many mathematicians in many different directions over the last thirty years. Discrete conformal geometry in its purest form is geometry born of combinatorics, but it has theoretical and practical applications.

¹I will use the term *topology* henceforth to mean *geometric topology*. By dropping the adjective *geometric* I certainly mean no slight of general, set-theoretic, or algebraic topology.

In the theoretical realm, it produces a discrete analytic function theory that is faithful to its continuous cousin, a quantum theory of complex analysis from which the classical theory emerges in the limit of large scales. In the realm of applications, it has been developed in a variety of directions, for practical applications in areas as diverse as biomedical imaging and 3D print head guidance. This rather large body of work flows from simple insights that Thurston presented in his lecture at Purdue University in 1985 on how to use the most elementary case of his circle packing theorem to provide a practical algorithm for approximating the Riemann mapping from a proper, simply-connected planar domain to the unit disk. A personal accounting of this development can be found in the author's own review [14] of the bible of circle packing theory, Ken Stephenson's *Introduction to Circle Packing: The Theory of Discrete Analytic Functions* [71].

A perusal of the section and subsection headings of this chapter will give the reader a clue as to where I am going in this survey. I primarily stick with the theoretical results for which there are fairly direct lines from the Koebe–Andre'ev–Thurston Theorem to those results. This means in particular that I almost totally ignore the really vast array of practical applications that circle packing has found, especially in the last two decades as discrete differential geometry has become of primary importance in so many applications among computer scientists and computational geometers. A survey of applications will have to wait as space constraints preclude a discussion that does justice to the topic.

5.1.2 Dedication and Appreciation

This chapter is dedicated to the memory of Bill Thurston and his student Oded Schramm, and to an appreciation of Jim Cannon and Ken Stephenson. I have spoken already of Bill Thurston's legacy. Oded Schramm was one of the first to press Thurston's ideas on circle packings to a high level of development and application, and his great originality in approaching these problems has bequeathed to us a treasure trove of beautiful gems of mathematics. Most of Oded's work on circle packing and discrete geometry was accomplished in the decade of the 1990s. As Bill is a demarcation point in the history of topology, Oded is one in the history of probability theory. In the late nineties, Oded became interested in some classical open problems in probability theory generated by physicists, in percolation theory and in random planar triangulations in particular. Physicists had much theoretical and computational evidence for the veracity of their conjectures, but little mathematical proof, or even mathematical tools to approach their verifications. In Oded's hands these venerable conjectures and problems began to yield to mathematical proof, using ingenious tools developed or refined by Oded and his collaborators, chief among which are SLE_κ , originally Stochastic-Loewner Evolution, now renamed as Schramm–Loewner Evolution, and UIPT's, or Uniform Infinite Planar Triangulations. For a wonderful biographical commentary on Oded's contributions to mathematics, see Steffen Rohde's article *Oded Schramm: From Circle Packing to SLE* in [58].

The two individuals who have had the greatest impact on my mathematical work are Jim Cannon and Ken Stephenson, the one a mathematical hero of mine, the other my stalwart collaborator for three decades. Jim's work has influenced mine significantly, and I greatly admire his mathematical tastes and contributions. Pre-Thurston, Jim had made a name for himself in geometric topology in the flavor of Bing and Milnor, having solved the famous double suspension problem and having made seminal contributions to cell-like decomposition theory and the characterization of manifolds. In the beginning of the Thurston era, his influential paper *The combinatorial structure of cocompact discrete hyperbolic groups* [24] anticipated many of the later developments of geometric group theory, presaging Gromov's thin triangle condition and, à la Dehn, the importance of negative curvature in solving the classical word and conjugacy problems of combinatorial group theory. He with Thurston invented automatic group theory and then Jim settled upon the conjecture that bears his name as the work that for three decades has consumed his attention. Ken has been a joy with whom to collaborate over the past three decades. He was inspired upon attending Thurston's Purdue lecture in 1985 to change his mathematical attentions from a successful career as a complex function theorist, to a geometer exploring this new idea of circle packing using both traditional mathematical proof and the power of computations for mathematical experimentation. I began my foray into Thurston-style geometry and topology by answering in [17] a question of Ken and Alan Beardon from one of the first research papers [7] to appear on circle packings after Rodin and Sullivan's 1987 paper [62] confirming the conjecture of Thurston from the Purdue lecture. Ken and I are co-authors on a number of research articles and his down-to-earth approach to the understanding of mathematics has been a constant check on my tendency toward flights of fancy. I have learned from him how to tell a good story of a mathematical topic. For Ken's warm friendship and collaboration I am grateful.

5.2 The Koebe–Andre'ev–Thurston Theorem, Part I

5.2.1 Koebe Uniformization and Circle Packing

In the early years of the twentieth century, rigorous proofs of the Riemann Mapping Theorem and the more general Uniformization Theorem were given by such eminent mathematicians as Osgood, Carathéodory, Poincaré, and Koebe, and refinements and re-workings would continue to be made by others, even up to the present.² The generalization of the Riemann Mapping Theorem to multiply-

²The author recommends rather highly the article *On the history of the Riemann mapping theorem* [38] by Jeremy Gray and the monograph *Uniformization of Riemann Surfaces: Revisiting a Hundred-Year-Old Theorem* [32]. These two works give insightful historical accountings of the discovery, articulation, understanding, and finally rigorous proofs of the Riemann Mapping

connected domains fell to the hands of Paul Koebe, who in 1920 in [49] proved that every finitely-connected domain in the Riemann sphere is conformally equivalent to a *circle domain*, a connected open set all of whose complementary components are points or closed round disks. Of course for a 1-connected, or simply-connected, domain this is nothing more than the Riemann Mapping Theorem. He proved also a rigidity result, that any conformal homeomorphism between any two circle domains with finitely many complementary components is in fact the restriction of a Möbius transformation.³ Koebe's real goal was what is known by its German name as *Koebe's Kreisnormierungsproblem* and by its English equivalent as *Koebe's Uniformization Conjecture*, which he posed in 1908.

Koebe Uniformization Conjecture ([48]) *Every domain in the Riemann sphere is conformally homeomorphic to a circle domain.*

This of course includes those domains with infinitely many, whether countably or uncountably many, complementary components. The general Koebe Uniformization Conjecture remains open to this day. More on this later.

In a paper of 1936, Koebe obtained the following circle packing theorem as a limiting case of his uniformization theorem of 1920. This went unnoticed by the circle packing community until sometime in the early 1990s.


Koebe Circle Packing Theorem ([50]) *Every oriented simplicial triangulation K of the 2-sphere \mathbb{S}^2 determines a univalent circle packing $K(\mathcal{C})$ for K , unique up to Möbius transformations of the sphere.*

Here the *circle packing* $K(\mathcal{C})$ is a collection $\mathcal{C} = \{C_v : v \in V(K)\}$ of circles C_v in the sphere \mathbb{S}^2 indexed by the vertex set $V(K)$ of K such that C_u and C_v are tangent whenever uv is an edge of K , and for which circles C_u , C_v , and C_w bound a positively oriented interstice whenever uvw is a positively oriented face of K . The circle packing is *univalent* if there is a collection $\mathcal{D} = \{D_v : v \in V(K)\}$ of disks with $C_v = \partial D_v$ whose interiors are pairwise disjoint.⁴ Connecting the centers of the adjacent circles by appropriate great circular arcs then produces a geodesic triangulation of \mathbb{S}^2 isomorphic to K . Figure 5.1 shows a circle packing of the sphere determined by an abstract triangulation K , and the realization of K as a geodesic triangulation. Of course the circle packings for a fixed K are Möbius equivalent, while the corresponding geodesic triangulations are not, simply because neither circle centers nor great circles are Möbius-invariant. I will look at a proof of the Koebe Circle Packing Theorem later, but first I'll present Thurston's generalization.

Theorem and the Uniformization Theorem. The narratives are at once engaging and perceptive, illustrating wonderfully the fact that mathematics is generally a common endeavor of a community of folks rather than the singular achievement of an enlightened few.

³Beware! This is not true in general. Two domains with uncountably many complementary components may be conformally equivalent yet fail to be Möbius equivalent.

⁴Without univalence, packings with branching would be allowed, where the sequence of circles tangent to a single circle C may wrap around C multiple times before closing up. See Sect. 5.3.3.



(a) The circle packing determined by a triangulation K of \mathbb{S}^2 .

(b) The corresponding geodesic triangulation of \mathbb{S}^2 .

Fig. 5.1 An abstract triangulation K of the 2-sphere determines (a) a circle packing, which in turn determines a realization of K as (b) a geodesic triangulation of the 2-sphere.

5.2.2 Koebe–Andre’ev–Thurston, or KAT for Short

In his Princeton course of 1978–79, Thurston greatly generalized the Koebe Circle Packing Theorem, though at the time he was unaware of Koebe’s result. He generalized in two ways, first by allowing adjacent circles to overlap and second by extending the theorem to arbitrary compact orientable surfaces. Thurston realized that his version of the theorem on the sphere \mathbb{S}^2 in fact encodes information about convex hyperbolic polyhedra, the connection of course through the fact that the sphere \mathbb{S}^2 serves as the space at infinity of the Beltrami–Klein and Poincaré ball versions of hyperbolic three-space \mathbb{H}^3 with circles on the sphere the ideal boundaries of hyperbolic planes in \mathbb{H}^3 . These polyhedra had been characterized in two papers of Andre’ev from 1970, whose results can be interpreted in terms of the existence and uniqueness of the circle packings Thurston examined in his generalization of Koebe. Thurston’s generalization to overlapping packings on the sphere is now known as the *Koebe–Andre’ev–Thurston Theorem*, honoring its three principle protagonists.

Koebe–Andre’ev–Thurston Theorem I (for the sphere) *Let K be an oriented simplicial triangulation of \mathbb{S}^2 , different from the tetrahedral triangulation, and let $\Phi : E(K) \rightarrow [0, \pi/2]$ be a map assigning angle values to each edge of K . Assume that the following two conditions hold.*

- (i) *If e_1, e_2, e_3 form a closed loop of edges from K with $\sum_{i=1}^3 \Phi(e_i) \geq \pi$, then $e_1, e_2,$ and e_3 form the boundary of a face of K .*
- (ii) *If e_1, e_2, e_3, e_4 form a closed loop of edges from K with $\sum_{i=1}^4 \Phi(e_i) = 2\pi$, then $e_1, e_2, e_3,$ and e_4 form the boundary of the union of two adjacent faces of K .*

Then there is a realization of K as a geodesic triangulation of \mathbb{S}^2 and a family $\mathcal{C} = \{C_v : v \in V(K)\}$ of circles centered at the vertices of the triangulation so that the two circles C_v and C_w meet at angle $\Phi(e)$ whenever $e = vw$ is an edge of K . The circle packing \mathcal{C} is unique up to Möbius transformations.

Now I want to point out that exactly what is called the *Koebe–Andre’ev–Thurston Theorem* is not at all settled. Some references use the term to mean the tangency case of the theorem ($\Phi \equiv 0$), which is nothing more than the Koebe Circle Packing Theorem, while others use the term to mean Thurston’s full generalization of the theorem to arbitrary closed surfaces that is presented in Sect. 5.3.2. Exactly what Thurston proved in GTTM also often is misreported. In fact my introduction to this section is a bit of a misreporting, so let me take a little time to say exactly what Thurston does in Chapter 13 of GTTM.

In terms of circle packings on the 2-sphere, Thurston does not allow overlaps of adjacent circles, only tangencies. His version of the tangency case appears as Corollary 13.6.2 in Chapter 13 of GTTM, and appears as a corollary of Theorem 13.6.1, which he attributes to Andre’ev. This theorem concerns hyperbolic structures on orbifolds and, as it was Thurston who invented the notion of orbifold in his course at Princeton during 1976–77 as recorded in the footnote on page 13.5 of Chapter 13 of GTTM itself, this theorem is an interpretation of Andre’ev’s in the context of orbifolds. Thurston does not give a proof of Theorem 13.6.1, but uses its result ensuring the existence of a hyperbolic structure on a suitable orbifold to prove Koebe’s Theorem of 1936, Corollary 13.6.2. He does this by using the triangulation K to define an associated polyhedron P by cutting off vertices by planes that pass through midpoints of edges. He then uses the Andre’ev result to realize P as a right-angled ideal polyhedron in \mathbb{H}^3 . The faces of this polyhedron then lie in planes whose ideal circular boundaries are the circles of the desired tangency packing complemented by the orthogonal circles through three mutually adjacent points of tangency. He then invokes Mostow rigidity for uniqueness.

It isn’t until he presents Theorem 13.7.1 that Thurston allows for adjacent circles to overlap with angle between zero and $\pi/2$, and that only for surfaces other than the sphere, those surfaces with nonpositive Euler characteristic. Thurston proves this by assigning polyhedral metrics with curvature concentrated at the vertices v_1, \dots, v_n by assigning a radius r_i at vertex v_i . Defining the mapping $c : \mathbb{R}^n \rightarrow \mathbb{R}^n$ that measures the curvature via $c(r)_i = 2\pi -$ (the angle sum at vertex v_i), he then argues in nine pages that the origin $\mathbf{0}$ is in the image of c , which implies the desired result. It is the case that the version Thurston presents on the sphere, Corollary 13.6.2, is Koebe’s result, and uses Andre’ev’s ideas for the proof. It is only with this positive genus version, Theorem 13.7.1, that Thurston puts forth new geometric ideas, fertile enough to spawn an industry dedicated to understanding polyhedral metrics on surfaces and their induced circle packings.

Thurston’s approach to circle packing is rather entwined with his overall concern, that of hyperbolic structures on three-dimensional manifolds and orbifolds. Since this work of the 1970s, Thurston’s circle packing results have spawned a rather extensive theory that is more combinatorial and geometric, and related more

to classical complex function theory and Riemann surfaces, and less to three-manifolds. It is related intimately to hyperbolic polyhedra and their generalizations, this the subject of Sect. 5.7, and has found several scientific applications. In the hands of folks like Ken Stephenson and his students and collaborators, it has spawned a discrete theory of complex analytic functions, laid out ever so elegantly in Stephenson's *Introduction to Circle Packing* [71]. It has yielded beautiful results on, for example, discrete minimal surfaces in the hands of the Berlin school of Bobenko, Hoffman, Springborn, Pinkall, and Polthier; see for example [11] and [56]. Though the theory now is rather mature, it continues to interact in new and interesting ways with new areas, for instance lying in the background in conformal tilings [21, 22], or in the foreground with its interaction with the classical rigidity theory of Euclidean frameworks [23]. There is an immense literature here, and so much of it owes a great debt of gratitude to the insights of Bill Thurston.

5.2.3 *A Proof of the Koebe Circle Packing Theorem*

Rather than proving the whole of KAT I, I will address the case where Φ is identically zero and prove Koebe's result. The proof presented here can be modified to give a complete proof of KAT I, which is done in [18] in proving a generalization.⁵ Koebe's original proof of his namesake theorem uses a limiting process on circle domains and classical analytic arguments on convergence of analytic families of maps, very much in the flavor of what we now teach as classical techniques in our complex analysis courses. There are now many proofs of the Koebe Circle Packing Theorem. To name a few, besides Koebe's, there is Thurston's in GTTM already outlined above based on Andre'ev's results on hyperbolic polyhedra, Al Marden and Burt Rodin's using piecewise flat polyhedral metrics, Alan Beardon and Ken Stephenson's [7] that adapts the classical Perron method for constructing harmonic maps as an upper envelope of subharmonic maps, Colin de Verdière's [33] based on a variational principle, Igor Rivin's hidden in his paper [60] on Euclidean structures on triangulated surfaces, the author's [12] that turns the Beardon–Stephenson proof upside down to address packings on punctured surfaces, and Alexandre Bobenko and Boris Springborn's [10] that uses a minimal principle on integrable systems. Here I present a geometric and combinatorial proof where hyperbolic geometry is the crucial ingredient. The proof is a twist on the Perron method used by Beardon and Stephenson in [7] and is specialized from a more general version that applies to arbitrary surfaces of finite conformal type that appears in [12]. We will see that it has the advantage of generalizing in interesting ways.

Proof of the Koebe Circle Packing Theorem By removing one vertex v_0 from K and its adjacent edges and faces, one obtains a triangulation T of a closed disk.

⁵See Sect. 5.3.3.

Place a piecewise hyperbolic metric on T as follows. For any positive function $r : V(T) \rightarrow (0, \infty)$, let $|T(r)|$ be the metric space obtained by identifying the face $v_1 v_2 v_3$ of T with the hyperbolic triangle of side lengths $r(v_i) + r(v_j)$ for $i \neq j \in \{1, 2, 3\}$. This places a piecewise hyperbolic metric on T with cone-like singularities at the interior vertices. This structure often now is called a piecewise hyperbolic *polyhedral metric*, and the function r is called variously a *radius vector* or *label*. For any vertex v , one can measure the angle sum $\theta_r(v)$ of the angles at v in all the faces incident to v . I will say that r is a *superpacking label* for T if the angle sums of all interior vertices are at most 2π , and a *packing label*⁶ if all are equal to 2π .

Now modify this a little by allowing r to take infinite values at the boundary vertices. This causes some ambiguity only if there is a separating edge in T that disconnects T when removed. This will be taken care of later, so for now assume no separating edge exists. The goal is to find a packing label τ with $\tau(w) = \infty$ whenever w is a boundary vertex. Assuming that such an τ exists, we may glue on hyperbolic half planes along the faces with two boundary vertices to give a complete hyperbolic metric on a topological disk, which must be isometric to the hyperbolic plane. This implies that the metric space $|T(\tau)|$ is isometric to an ideal polygon in the hyperbolic plane whose sides are hyperbolic lines connecting adjacent ideal vertices that correspond to the boundary vertices of T . Now placing hyperbolic circles of radii $\tau(v)$ centered at interior vertices v and horocycles centered at ideal vertices determined by the boundary vertices gives a univalent circle packing of the hyperbolic plane realized as, say, the Poincaré disk, the unit disk \mathbb{D} in the complex plane with Poincaré metric $ds = 2|dz|/(1 - |z|^2)$. The boundary circles are horocycles in the hyperbolic metric on the disk and are therefore circles internally tangent to the unit circle. Stereographic projection to the sphere \mathbb{S}^2 and addition of the equator as the circle corresponding to the vertex v_0 removed initially produces a univalent circle packing of the sphere in the pattern of K as desired. Uniqueness follows from uniqueness of the packing label τ with infinite boundary values, which follows from the construction of τ explained next.

Define the function τ as

$$\tau(v) = \inf \{r(v) : r \in \mathfrak{R}\} \tag{5.2.1}$$

where

$$\mathfrak{R} = \{r : V(T) \rightarrow (0, \infty] : r \text{ is a superpacking label for } T \text{ with infinite boundary values}\}.$$

The claim is that this is the desired packing label. The first observation is that $\mathfrak{R} \neq \emptyset$ so that we are not taking the infimum of the empty set. This is because one may choose label values so large on the interior vertices that all of the faces become hyperbolic triangles whose interior angles are no more than $2\pi/d$, where d is the

⁶For emphasis one sometimes calls this a *hyperbolic packing label* to distinguish it from *flat* or *Euclidean packing labels* that also find their use in this discipline.

maximum degree of all the vertices of T . It follows that τ is a non-negative function with infinite boundary values. To verify that τ is a packing label, I show that

- (i) τ cannot take a zero value on any interior vertex, which then implies that $\tau \in \mathfrak{R}$, and,
- (ii) the angle sum at any interior vertex is 2π , meaning further that τ is a packing label.

We need two preliminary observations.

- (iii) *Hyperbolic area is bounded away from zero.* The hyperbolic area of the singular hyperbolic surface $|T(r)|$ is $\geq \pi$ for all superpacking labels $r \in \mathfrak{R}$.
- (iv) *Monotonicity of angles.* For a face $f = v_0v_1v_2$ of T , let $\alpha_r(i)$, for $i = 0, 1, 2$, be the angle that the label $r \in \mathfrak{R}$ gives to f at vertex v_i . Then $\alpha_r(0) \uparrow \pi$, $\alpha_r(1) \downarrow 0$, and $\alpha_r(2) \downarrow 0$ monotonically as $r(v_0) \downarrow 0$ when $r(v_1)$ and $r(v_2)$ are held fixed.

In calculating the hyperbolic area to confirm item (iii), let $V(T)$ and $F(T)$ be the respective vertex and face sets of T of respective cardinalities V and F . The sum of the angles of a face when given its metric by r is denoted $\alpha_r(f)$ so that its hyperbolic area is $A_r(f) = \pi - \alpha_r(f)$. Finally, with V_{int} and V_{bd} denoting the numbers of interior and boundary vertices of T so that $V = V_{\text{int}} + V_{\text{bd}}$, one has

$$\text{hyp-area}(|T(r)|) = \pi F - \sum_{f \in F(T)} \alpha_r(f) = \pi F - \sum_{v \in V(T)} \theta_r(v) \geq \pi(F - 2V_{\text{int}}), \tag{5.2.2}$$

since $\theta_r(v) \leq 2\pi$ at interior vertices and $\theta_r(v) = 0$ at boundary ones. An Euler characteristic exercise then shows that $F - 2V_{\text{int}} = V_{\text{bd}} - 2 \geq 1$, the inequality holding since K is simplicial. It follows that every superpacking label with infinite boundary values produces a metric on T with hyperbolic area at least π . Item (iv) is almost obvious from drawing examples, but can be given a rigorous proof using the hyperbolic law of cosines from hyperbolic trigonometry.

I now address item (i). First the claim is that the label τ cannot be identically zero on the set of interior vertices. Indeed, if τ is identically zero, one may choose a sequence of superpacking labels r_i with infinite boundary values such that, for each interior vertex v , $r_i(v) \rightarrow 0$ as $i \rightarrow \infty$. This latter fact in turn follows from the observation that the minimum label $\min\{r_1, r_2\}$ is in \mathfrak{R} whenever r_1 and r_2 are labels in \mathfrak{R} , which in turn is a consequence of the monotonicity of angles (iv). Recall that we are under the assumption that there are no separating edges so that at least one vertex of any face f of T is interior. Any such interior vertex has r_i -values converging to zero, and any boundary one is fixed at infinity, and with this it is easy to see that the hyperbolic area $A_{r_i}(f) \rightarrow 0$ as $i \rightarrow \infty$. But this implies that the hyperbolic area of $|T(r_i)|$ converges to zero as $i \rightarrow \infty$, which contradicts item (iii).

Now could it be that τ takes a zero value at some interior vertex, but not at all? The argument that this in fact does not happen is a generalization of what I have

argued thus far. I will but give an indication of how it goes, referring the reader to [12] for details. Let T' be the subcomplex of T determined by those faces of T that have a vertex in $\tau^{-1}(0)$. An argument using Euler characteristic similar to that already given implies that the hyperbolic area of $|T'(r)|$ is positive and bounded away from zero for every superpacking label r with fixed non-negative boundary values. But an argument as in the preceding paragraph shows that the hyperbolic areas of $|T'(r_i)|$ converge to zero for a sequence of superpacking labels with fixed boundary values and interior vertex values converging to zero. This contradiction implies that τ is a positive function on the interior vertex set, and continuity of angles of a triangle with respect to edge lengths implies that $\theta_\tau(v) = \lim_{i \rightarrow \infty} \theta_{r_i}(v) \leq 2\pi$ at any interior vertex, since $\theta_{r_i}(v) \leq 2\pi$ for all i . This shows that $\tau \in \mathfrak{R}$ and completes the verification of item (i).

Item (ii) follows quickly from item (iv). Indeed, if (ii) fails, then there is an interior vertex v of T such that $\theta_\tau(v) < 2\pi$. By the monotonicity properties (iv), varying τ by slightly decreasing its value at v without changing any other values increases $\theta_\tau(v)$ while decreasing $\theta_\tau(w)$ for any vertex w incident to v . By making that decrease of $\tau(v)$ small enough to keep the angle sum at v below 2π , we obtain a superpacking label r with infinite boundary values that satisfies $r(v) < \tau(v)$, which contradicts the definition of τ in Eq. 5.2.1.

At this point I have shown that τ is a packing label with infinite boundary values, and I now claim that it is the only one. Suppose there is a packing label r in \mathfrak{R} that differs from the infimum label τ defined in Eq. 5.2.1. Then $\tau(v) \leq r(v)$ for all vertices v , but there must be some interior vertex w with $\tau(w) < r(w)$. This implies that the hyperbolic area of the surface $|T(\tau)|$ is strictly less than that of $|T(r)|$. But this is impossible since τ and r are packing labels with infinite boundary values, and as argued above, both $|T(\tau)|$ and $|T(r)|$ are ideal hyperbolic polygons with V_{Bd} sides. An easy exercise shows that the hyperbolic area of any such hyperbolic polygon is equal to $(V_{\text{Bd}} - 2)\pi$.

This completes the proof modulo the assumption that T has no separating edge. This is handled by induction on the number of such edges. If there is one separating edge uv , cut T into T_1 and T_2 along that edge, circle pack each in the unit disk with horocyclic boundary circles, and then using Möbius transformations, place the T_1 packing in the upper half disk with the horocycles for u and v circles tangent at the origin and centered on the real axis, and place the T_2 packing in the lower half of the disk with those same horocyclic circles for u and v . This is possible since T is oriented, and this gives an appropriate packing label of T with infinite boundary values. □

5.2.4 Maximal Packings and the Boundary Value Problem

This proof actually proves the following extremely useful fact, which Beardon and Stephenson [7] exploited to give the first extension of the Koebe Circle Packing Theorem to infinite packings of the disk and the plane. The infinite theory is presented in Sect. 5.4.

Maximal Disk Packing Theorem *Every oriented simplicial triangulation T of a closed disk determines a univalent circle packing $T(\mathbb{C})$ for T in the unit disk \mathbb{D} in the complex plane \mathbb{C} , unique up to Möbius transformations of the disk, with the circles corresponding to boundary vertices of T internally tangent to the unit circle boundary $\partial\mathbb{D} = \mathbb{S}^1$. Moreover, when given its canonical hyperbolic metric making \mathbb{D} into the Poincaré disk model of the hyperbolic plane \mathbb{H}^2 , the circle radii of the packing are uniquely determined by T .*

The circle packing guaranteed by this theorem is called the *maximal packing* for T . This theorem is in fact a special case of the more general result of Beardon and Stephenson [8] that solves the discrete version of the classical Dirichlet boundary value problem of harmonic analysis. In that paper, the authors also prove a discrete version of the classical Schwarz-Pick Lemma of complex analysis. These two theorems finish up the present section.

Discrete Boundary Value Theorem (Beardon and Stephenson [8]) *Let T be an oriented simplicial triangulation of a closed disk and $f : V_{\text{Bd}}(T) \rightarrow (0, \infty]$ a function assigning positive or infinite values to the boundary vertices. Then there exists a unique hyperbolic packing label $\tau : V(T) \rightarrow (0, \infty]$ extending f . The resulting circle packing $T(\mathbb{C}_\tau)$ of the unit disk \mathbb{D} is unique up to Möbius transformations of \mathbb{D} .*

Proof The proof is a straightforward modification of that of the Koebe Circle Packing Theorem already presented. Again $\tau = \inf \mathfrak{R}$ is the desired packing label, provided that

$$\mathfrak{R} = \{r : r \text{ is a superpacking label for } T \text{ with } r(w) \geq f(w) \text{ when } w \in V_{\text{Bd}}(T)\}.$$

Of course, $f \equiv \infty$ gives the maximal packing of the preceding theorem. \square

This proof is a modification of the Beardon–Stephenson proof, which uses *subpacking* rather than superpacking labels. In a subpacking label, the interior angle sums are greater than or equal to 2π and one obtains the packing label as an upper envelope of subpacking labels, with the packing label given by $\tau = \sup \mathfrak{R}'$ where \mathfrak{R}' is the set of subpacking labels with boundary values given by f . The advantage of approaching the desired packing label τ from above using superpackings ($\inf \mathfrak{R}$) rather than from below using subpackings ($\sup \mathfrak{R}'$) is that this *upper Perron method* readily generalizes to include cusp type singularities and cone type singularities at interior vertices.⁷ This is presented in Sect. 5.3.4.

A word of warning here. When the boundary values are allowed to be finite, the resulting packing, though locally univalent, may not be globally univalent. This means that the disks bounded by the circles of the packing may overlap non-trivially, though ones neighboring the same interior vertex never do; this is the meaning of

⁷Another not insignificant advantage is that it is easy to show that $\mathfrak{R} \neq \emptyset$ while proving that $\mathfrak{R}' \neq \emptyset$ generally is difficult.




Fig. 5.2 A locally univalent circle packing that is not globally univalent

locally univalent. Figure 5.2 shows a locally univalent packing that is not globally univalent.

The second theorem of Beardon and Stephenson follows partly from the proof of the first (item (i)), and the rest of the theorem follows from a careful analysis of paths and angles in piecewise hyperbolic surfaces. The details of course appear in [8]. The theorem I state here is the generalization of what the reference [8] calls the Discrete Schwarz-Pick Lemma, which in turn is obtained by setting the boundary values of r identically to infinity.

Discrete Schwarz-Pick Lemma (Beardon and Stephenson [8]) *Let $r, r'; V(T) \rightarrow (0, \infty]$ be packing labels for the oriented simplicial triangulation T of a closed disk. Suppose that $r(w) \geq r'(w)$ at every boundary vertex $w \in V_{\text{Bd}}(T)$. Then*

- (i) $r \geq r'$; i.e., $r(v) \geq r'(v)$ at every vertex v of T ;
- (ii) $\rho_r(u, v) \geq \rho_{r'}(u, v)$ for any two vertices u and v , where ρ_r is the distance function on the metric surface $|T(r)|$, and similarly for $\rho_{r'}$;
- (iii) $A_r(f) \geq A_{r'}(f)$ for any face f of T . (Recall that $A_r(f)$ is the hyperbolic area of the face f .)

Moreover, if a single instance of finite equality occurs at an interior vertex in (i), or at vertices u and v at least one of which is interior in (ii), or at any face in (iii), then $r = r'$.

5.3 The Koebe–Andre’ev–Thurston Theorem, Part II

5.3.1 Circle Packings of Compact Surfaces

Thurston’s other avenue of generalization of Koebe, indeed the more far-reaching one, is his extension of KAT to arbitrary orientable closed surfaces. Here there are striking illustrations of how purely combinatorial information encodes precise geometry. I will start with Thurston’s tangency case of packings before presenting his version with overlaps.

Theorem 5.3.1 *Let K be an oriented simplicial triangulation of a closed surface S_g of positive genus. Then there is a metric of constant curvature, unique up to scaling, on S_g that supports a univalent, tangency circle packing $\mathcal{C} = \{C_v : v \in V(K)\}$ modeled on K . In particular, C_u and C_v are tangent whenever uv is an edge of K . The packing \mathcal{C} is unique up to isometries of S_g in this metric when $g \geq 2$, and up to scaling when $g = 1$. Connecting the centers of adjacent circles by geodesic shortest paths produces a geodesic triangulation of the surface in the pattern of K . The metric is locally Euclidean when $g = 1$ and locally hyperbolic otherwise.*

Just in case the reader blinked and missed it, I aim to emphasize the extent to which combinatorics determines geometry in this theorem. The simplicial complex K provides purely combinatorial data with topological overtones. Yet hidden inside of the combinatorics is precise geometry. For example in the hyperbolic case where $g > 1$, among the uncountably many possible pairwise distinct hyperbolic metrics of constant curvature -1 as tabulated in the $(6g - 6)$ -dimensional moduli space $\mathcal{M}(S_g) \cong \mathbb{R}^{6g-6}$, the complex K chooses exactly one of these metrics, and in that metric, determines a univalent circle packing unique up to isometry! For none of the other metrics that S_g supports is there a univalent tangency packing of circles in the pattern of K ! Since there are only countably many pairwise distinct simplicial triangulations of the fixed surface S_g , only countably many of the metrics parameterized by $\mathcal{M}(S_g)$ support any univalent tangency packing at all, though the set of metrics that do support such circle packings does form a dense subset of the moduli space.

I present a Proof of Theorem 5.3.1 based on the upper Perron method used to prove the Koebe Circle Packing Theorem.

Proof Let $\mathfrak{R} = \{r : V(K) \rightarrow (0, \infty) : \theta_r(v) \leq 2\pi \text{ for all } v \in V(K)\}$, the set of superpacking labels for K . Here again, exactly as in the proof of the Koebe Circle Packing Theorem, the label r determines a hyperbolic polyhedral metric surface $|K(r)|$. A unique packing label for which the angle sum at every vertex is equal to 2π would give all the claims of the theorem in the hyperbolic case. My claim is that when $g \geq 2$, the function $\tau = \inf \mathfrak{R}$ is the unique packing label for K , and when $g = 1$, then $\tau = \inf \mathfrak{R}$ is identically zero, but provides a way to place a flat polyhedral metric on K that meets the packing condition.

Exactly the calculation of Inequality 5.2.2 gives $\text{hyp-area}(|K(r)|) \geq (\mathbb{F} - 2V)\pi$ for any superpacking label $r \in \mathfrak{R}$, and an Euler characteristic argument gives

$$\mathbb{F} - 2V = -2\chi(S_g) = 4g - 4. \tag{5.3.1}$$

When $g \geq 2$ so that $\mathbb{F} - 2V$ is positive and hence $\text{hyp-area}(|K(r)|)$ is positive, the same argument used in the proof of the Koebe Circle Packing Theorem shows that items (i) and (ii) of that proof hold, so that τ is a packing label. Uniqueness follows exactly as in that proof.

The remaining case is when $g = 1$ so that S_g is a topological torus. Here are the steps in proving that S_1 supports a flat metric that supports a univalent circle packing in the pattern of K , both the packing and the metric unique up to scaling.

- (i) When $g = 1$, $\mathbb{F} - 2V = 0$ and this implies that $\tau = \inf \mathfrak{R}$ is identically zero on $V(K)$.
- (ii) Fix a vertex v^\dagger in K and let $\mathfrak{R}^\dagger = \{r^\dagger : r \in \mathfrak{R}\}$, where r^\dagger is the normalized label defined by $r^\dagger(v) = r(v)/r(v^\dagger)$.
- (iii) Show that $\tau^\dagger = \inf \mathfrak{R}^\dagger$ takes only positive values.
- (iv) Let $|K(\tau^\dagger)|_{\text{flat}}$ be the flat polyhedral surface with cone type singularities obtained by identifying a face $v_1v_2v_3$ with the Euclidean triangle of side-lengths $\tau^\dagger(v_i) + \tau^\dagger(v_j)$ for $i \neq j \in \{1, 2, 3\}$.
- (v) Show that $|K(\tau^\dagger)|_{\text{flat}}$ is non-singular; i.e., τ^\dagger is a flat packing label with Euclidean angle sums $\theta_{\tau^\dagger}^{\text{flat}}(v) = 2\pi$ at every vertex v .
- (vi) Show that τ^\dagger is the unique flat packing label with value 1 at v^\dagger .

The details of the argument appear in [7], but I will give an indication of why this outline works to prove the desired result. Let $A(r)$ be the hyperbolic area of the singular hyperbolic surface $|K(r)|$ when $r \in \mathfrak{R}$ and observe that

$$A(r) - s(r) = (\mathbb{F} - 2V)\pi, \quad \text{where} \quad s(r) = \sum_{v \in V(K)} (2\pi - \theta_r(v)). \tag{5.3.2}$$

Here $s(r)$ is the total *angle shortage*.⁸ In the genus 1 case, $\mathbb{F} - 2V = 0$ so $A(r) = s(r)$ for all superpacking labels $r \in \mathfrak{R}$. Now assuming that item (i) has been verified, any superpacking label r that is close to the infimum $\inf \mathfrak{R} = 0$ has area $A(r)$ close to zero and hence so too is the shortage $s(r)$ close to zero. In the limit as $r \rightarrow \inf \mathfrak{R} = 0$, the shortages $s(r) \rightarrow 0$ and this implies that the singular hyperbolic surfaces $|K(r)|$ have angle sums $\theta_r(v) \rightarrow 2\pi$ for every vertex v . Since Euclidean geometry is the small scale limit of hyperbolic geometry, this implies that the Euclidean angle sums $\theta_r^{\text{flat}}(v) \rightarrow 2\pi$ as $r \rightarrow 0$. Thus the collection $\{|K(r)|_{\text{flat}}\}_{r \in \mathfrak{R}}$ is a collection of singular flat surfaces whose singularities are removed in the limit as $r \rightarrow 0$. Of course there is no limiting surface since

⁸Also called the *discrete curvature*.

$r \rightarrow 0$. Whereas this cannot be remedied in hyperbolic geometry, it can be remedied in Euclidean geometry by rescaling the labels r as described in item (ii). With item (iii) confirmed so that the flat polyhedral surface $|K(\tau^\dagger)|_{\text{flat}}$ of item (iv) exists, since similarity transformations exist in Euclidean geometry, these rescalings preserve the Euclidean angles and imply that the limit surface $|K(\tau^\dagger)|_{\text{flat}}$ is non-singular. Items (v) and (vi) just state formally the result of making this imprecise but rather accurate discussion rigorous. \square

5.3.2 KAT for Compact Surfaces

Thurston's Theorem 13.7.1 of GTTM combines the introduction of surfaces of genus greater than zero in Theorem 5.3.1 with the overlap conditions of the KAT Circle Packing Theorem.

Koebe–Andre'ev–Thurston Theorem II (for compact surfaces) (Theorem 13.7.1, GTTM) *Let K be an oriented simplicial triangulation of a surface S_g of genus $g \geq 1$, and let $\Phi : E(K) \rightarrow [0, \pi/2]$ be a map assigning angle values to each edge of K . Assume that the following two conditions hold.*

- (i) *If e_1, e_2, e_3 form a closed loop of edges from K with $\sum_{i=1}^3 \Phi(e_i) \geq \pi$, then $e_1, e_2,$ and e_3 form the boundary of a face of K .*
- (ii) *If e_1, e_2, e_3, e_4 form a closed loop of edges from K with $\sum_{i=1}^4 \Phi(e_i) = 2\pi$, then $e_1, e_2, e_3,$ and e_4 form the boundary of the union of two adjacent faces of K .*

Then there is a metric of constant curvature on S_g , unique up to scaling, and a realization of K as a geodesic triangulation in that metric, as well as a family $\mathcal{C} = \{C_v : v \in V(K)\}$ of circles centered at the vertices of the triangulation so that the two circles C_v and C_w meet at angle $\Phi(e)$ whenever $e = vw$ is an edge of K . The circle packing \mathcal{C} is unique up to isometry.

I already have discussed the proof in GTTM. Let me say further that it was in this proof that Thurston introduced the idea of using labels, or radii assignments to vertices, to build a polyhedral surface with cone type singularities, and then to vary the labels until the packing condition is met. This is still the basic idea for proving many packing results, though the way in which one varies the labels and the choice of initial labels changes from researcher to researcher and from application to application. The Perron method used in this chapter is a modification of the method of Beardon and Stephenson [7]. This idea also led to a practical algorithm for producing the packing labels that was the starting point for Ken Stephenson's CirclePack. This sophisticated software package for computing circle packings has enjoyed extensive development over the past 30 years and is freely available at Ken's webpage.

Before I introduce infinite circle packings and their really interesting and novel features in Sect. 5.4, I'll discuss two generalizations of the KAT Theorems. The first

is presented in Sect. 5.3.3 and generalizes KAT I to certain branched packings of the 2-sphere where circles tangent to a given one wrap around that one more than once. These packings of course fail to be univalent, but provide a rich family of packings that model the behavior of polynomial mappings of the Riemann sphere. The ultimate goal is to model arbitrary rational mappings of the sphere, which would require the theory to extend to more general branch structures, this a topic of current research; see for example [5]. The second is presented in Sect. 5.3.4 and examines how to include both cusps with ideal vertices as well as prescribed discrete curvature at pre-chosen vertices.

5.3.3 A Branched KAT Theorem and Polynomial Branching

Ken Stephenson and I generalized KAT I by allowing for *polynomial branching* to occur in the circle packing. *Branching* means that we allow for the angle sums at predetermined vertices to be a positive integer multiple of 2π rather than just 2π , or stated differently, we allow the circles tangent to a given one to wrap around that given circle multiple times before closing up; see Fig. 5.3. *Polynomial* means that half the branching is concentrated at one vertex. The terminology comes from the classical theory of rational maps. Indeed, rational mappings may be thought of as




Fig. 5.3 Branching of multiplicity $m = 2$ or order $\sigma = 1$. Starting with the grey disk on the left and moving counterclockwise, four sequentially tangent grey disks wrap around the blue central disk nearly one full turn, at which point the sequentially tangent transparent (or white) disks take over to wrap around slightly more than one full turn to close up the flower of circles with angle sum $\theta = 4\pi$

branched self-mappings of the 2-sphere, and the polynomial mappings are precisely those in which there is an even amount of branching with half the branching occurring at a single point. Taken together, a circle packing promised by the next theorem mimics the behavior of a polynomial mapping of the Riemann sphere.

Our proof of the theorem as presented in [18] offers an independent proof of KAT I, which the branched version reduces to when the *branch structure* β is empty. In fact as far as I know, it was the first full direct proof of KAT I given that Thurston proves only the tangency case (the Koebe Circle Packing Theorem) and Marden-Rodin [54], though allowing overlapping circles, has more restrictive hypotheses. KAT I is implied by Igor Rivin's earlier work, which bears the same resemblance to KAT I as does Andre'ev's in that it is a result on the existence of hyperbolic polyhedra.

I state the result and then backtrack to fill in definitions and discuss the proof.

Polynomially Branched KAT Theorem (Bowers and Stephenson [18]) *Let K be an oriented simplicial triangulation of \mathbb{S}^2 , different from the tetrahedral triangulation, and let $\Phi : E(K) \rightarrow [0, \pi/2]$ be a map assigning angle values to each edge of K . Assume that the following two conditions hold.*

- (i) *If e_1, e_2, e_3 form a closed loop of edges from K with $\sum_{i=1}^3 \Phi(e_i) \geq \pi$, then e_1, e_2 , and e_3 form the boundary of a face of K .*
- (ii) *If e_1, e_2, e_3, e_4 form a closed loop of edges from K with $\sum_{i=1}^4 \Phi(e_i) = 2\pi$, then e_1, e_2, e_3 , and e_4 form the boundary of the union of two adjacent faces of K .*

If β is a polynomial branch structure for the edge-labeled triangulation (K, Φ) , then there exists a circle packing $\mathcal{C} = \{C_v : v \in V(K)\}$ for (K, Φ) , a family of circles in \mathbb{S}^2 so that the two circles C_v and C_w meet at angle $\Phi(e)$ whenever $e = vw$ is an edge of K , with $\text{br}(\mathcal{C}) = \beta$. The circle packing \mathcal{C} is unique up to Möbius transformations.

A branch structure essentially is a listing of some of the vertices of K , each paired with an integer ≥ 2 that indicates how many times the circles adjacent to the ones corresponding to the selected vertices wrap around before closing up. Before making this precise, let's observe that there must be further combinatorial conditions to ensure that a branched circle packing exists for the branch structure. Indeed, note that when there is no branching, the fact that K is a simplicial triangulation implies that the degree of each vertex is at least three, and this local condition guarantees that there are enough circles adjacent to a given circle to wrap around once, with angle sum 2π , at least in the tangency case. A moment's thought will show that if the desire is that there be branching of *multiplicity* $m \geq 2$ at a circle C_v , meaning that the circles adjacent to C_v wrap around m times before closing up, there had better be at least $1 + 2m$ adjacent ones to achieve the angle sum of $2\pi m$. This may not be sufficient but certainly is necessary, and the definition of a polynomial branch structure includes enough combinatorial conditions to ensure sufficiency.

To clothe this discussion in a bit of flesh, suppose that $\mathcal{C} = \{C_v : v \in V(K)\}$ is a circle packing for the pair (K, Φ) . For each vertex $v \in V(K)$, identify v with the

center of its corresponding circle C_v . Fixing a vertex v , let v_1, \dots, v_n be the list of neighbors of v forming the consecutive vertices in a walk around the boundary of the star $\text{st}(v)$ of v , and let α_i be the measure of the spherical angle $\angle v_i v v_{i+1}$. Then v is said to be a *branch point of order* $\sigma = m - 1$, or of *multiplicity* m , if $\theta(v) = 2\pi m$ for some integer $m \geq 2$, where $\theta(v) = \alpha_1 + \dots + \alpha_n$ is the *angle sum* at v ; again, see Fig. 5.3. The *branch set* $\text{br}(\mathcal{C})$ of the circle packing is the set of ordered pairs $(v, \sigma(v))$ as v ranges over the branch points and $\sigma(v)$ is the order of v . It is clear that the combinatorics of K as well as the values of Φ restrict the branch orders.

My aim is to construct circle packings of \mathbb{S}^2 in the pattern of K with overlaps given by Φ with a given, predetermined branch set. Toward this end, I will define a branch structure on the complex $T = K \setminus \text{Int}[\text{st}(v_\infty)]$ that triangulates the closed disk one obtains by deleting one vertex, v_∞ , and its incident open cells from K . I will use Φ_T to mean the restriction of Φ to the vertices of T .

Definition (Branch Structure) A set $\beta = \{(v_1, \sigma_1), \dots, (v_\ell, \sigma_\ell)\}$, where v_1, \dots, v_ℓ is a pairwise distinct list of interior vertices of T and each σ_i is a positive integer, is a *branch structure* for the pair (T, Φ_T) if the following condition holds: for each simple closed edge path $\gamma = e_1 \cdots e_n$ in T that bounds a combinatorial disk D that contains at least one of the vertices v_i , the inequality

$$\sum_{i=1}^n [\pi - \Phi_T(e_i)] > 2\pi(\sigma(D) + 1) \tag{5.3.3}$$

holds, where $\sigma(D) = \sum \sigma_i$, the sum taken over all indices i for which $v_i \in \text{Int}(D)$.

We will see that this condition on the combinatorics of T and the values of Φ_T ensures that there are no local obstructions to the existence of a circle packing for (T, Φ_T) whose branch set is β , and in fact is enough to ensure that there are no global ones.

Definition (Polynomial Branch Structure) Let K be a simplicial triangulation of \mathbb{S}^2 with edge function $\Phi : V(K) \rightarrow [0, \pi/2]$. A collection


$$\beta = \{(v_\infty, \sigma_\infty), (v_1, \sigma_1), \dots, (v_\ell, \sigma_\ell)\}$$

is a *polynomial branch structure* for (K, Φ) if the following conditions prevail.

- (1) $\sigma_\infty = \sigma_1 + \dots + \sigma_\ell$.
- (2) The vertices v_1, \dots, v_ℓ are all interior vertices of the complex $T = K \setminus \text{Int}[\text{st}(v_\infty)]$.
- (3) $\beta_T = \{(v_1, \sigma_1), \dots, (v_\ell, \sigma_\ell)\}$ is a branch structure for (T, Φ_T) .
- (4) No Φ -edge labeled subgraph of the type given in Fig. 5.4 occurs in K where v is one of the branch vertices v_1, \dots, v_ℓ .

A few comments concerning this definition are in order. Item (1) says that there is an even amount of branching and half of it occurs at vertex v_∞ ; item (2) says that no branch vertex from the list v_1, \dots, v_ℓ is adjacent to the vertex v_∞ ; item (3) in

Fig. 5.4 A forbidden edge-labeled subgraph for a polynomial branch structure



particular says that Inequality Eq. 5.3.3 holds for (T, Φ_T) ; item (4) is a technical condition that avoids impossible configurations.

Discussion of Proof How do we put all of this together to prove the Polynomially Branched KAT Theorem? Letting $\beta = \{(v_\infty, \sigma_\infty), (v_1, \sigma_1), \dots, (v_\ell, \sigma_\ell)\}$ be a polynomial branch structure for (K, Φ) , we remove the vertex v_∞ and work with hyperbolic polyhedral metrics on the disk triangulation T as in the proof of the Koebe Circle Packing Theorem. The idea is the same as there in that we want to use vertex labels on T to describe hyperbolic triangles that then are identified with faces to form a singular hyperbolic surface, and then vary the labels to meet angle targets at the vertices. There are three new difficulties that appear.

- (i) Target overlap angles are given by Φ_T for adjacent circles rather than tangencies.
- (ii) Rather than 2π , the target angle sums at branch vertices are $2\pi m_i$ for integers $m_i = \sigma_i + 1 \geq 2$.
- (iii) As the boundary $\partial\mathbb{D}$ ultimately will serve as the circle corresponding to v_∞ in the desired circle packing, the overlaps of the boundary circles of the packing for T must intersect the unit circle at the angles demanded by Φ .

Now items (i) and (ii) are really no problem as superpacking labels can be described that allow for prescribed overlap angles for adjacent circles and target angles prescribed by the branch structure. The real difficulty is item (iii). If we use radius labels, the best we can do is, as in the proof of the Koebe Theorem, get boundary circles that meet the unit circle at single points with intersection angle zero. The hint for resolving this difficulty is found in thinking a bit more about the role of circles in hyperbolic geometry, and in particular in the Poincaré disk model where \mathbb{H}^2 is identified with the unit disk \mathbb{D} , and the ideal boundary of \mathbb{H}^2 is identified with $\partial\mathbb{D} = \mathbb{S}^1$. Euclidean circles that meet the Poincaré disk \mathbb{D} not only serve as hyperbolic circles, but also as horocycles and hypercycles. Those that lie entirely within \mathbb{D} are hyperbolic circles, those internally tangent to the ideal boundary \mathbb{S}^1 are

horocycles, and those that meet the boundary in two points a and b are hypercycles whose points in \mathbb{D} lie equidistant to the hyperbolic line with ideal endpoints a and b . This latter case includes the hyperbolic geodesic lines. What proves fruitful here is the fact that, when oriented, these Euclidean circles and circular arcs are precisely the curves of constant geodetic curvature in the hyperbolic plane. This is implied immediately by the fact that these are the flow lines of 1-parameter groups of hyperbolic isometries, the hyperbolic circles the flow lines of elliptic flows, horocycles of parabolic flows, and hypercycles of hyperbolic flows.

Here are the salient facts about the geodetic curvature κ of an arc of an oriented Euclidean circle that lies in the Poincaré disk \mathbb{D} . Call an arc $c = C \cap \mathbb{D}$, where C is a Euclidean circle that meets \mathbb{D} , a *cycle* with *parent circle* C . There is a normalized setting in which the curvature can be read off easily. Apply a conformal automorphism of the disk so that c passes through the origin and its parent circle C is centered on the positive real axis. Orient c counterclockwise and let $t, 0 < t \leq \infty$ denote the point of intersection of C with the interval $(0, \infty]$. Then the curvature satisfies $\kappa = \kappa(c) = 1/t$. In terms of intrinsic parameters, for counterclockwise-oriented hypercycles when $t > 1$, $\kappa(c) = \cos \alpha$ where α is the acute angle of intersection of C with the unit circle. This includes the case of a hyperbolic geodesic where $\alpha = \pi/2$ and $\kappa = 0$. Assuming still the counterclockwise orientation, when $t = 1$, c is a horocycle with $\kappa(c) = 1$, and when $t < 1$, $c = C$ is a hyperbolic circle of some hyperbolic radius r with $\kappa(c) = \coth r$.

For our purposes it is quite fortuitous that monotone curvature parameters for cycles can be used as vertex labels on T in place of radii labels to encode a singular hyperbolic metric on a disk that T triangulates. The curvature is inversely related to the radii, but the really important feature is that, unlike radii labels, the curvature label can be used to identify faces of T , not only with hyperbolic triangles with both finite and ideal vertices, but also triangles with “hyperideal vertices.”⁹ This means that when curvatures κ_1, κ_2 , and κ_3 label the vertices of the face f and values $\Phi_T(e_i)$ for $i = 1, 2, 3$ label the opposite edges, the face f may be identified with the region of the hyperbolic plane determined by cycles of curvatures κ_1, κ_2 , and κ_3 overlapping with angles $\Phi_T(e_i)$ for $i = 1, 2, 3$. This accomplishes two things. First, the overlaps of cycles are given by the edge function Φ_T . Second, and very importantly, if the vertex w of f is a boundary vertex and the value $\kappa = \cos \Phi_T(wv_\infty)$ is used for the curvature, then the boundary cycle corresponding to w overlaps with the unit circle by an angle of $\Phi_T(wv_\infty)$.

The important point is that the set \mathfrak{K} of curvature labels, ones whose boundary values are given by $g(w) = \cos \Phi(wv_\infty)$ for the boundary vertex w , and that produce superpackings where the angle sums at interior vertices are no more than 2π at non-branch points and no more than $2\pi m_i$ at branch point v_i , may be varied to obtain a β -packing label, this time as $\sup \mathfrak{K}$, the supremum instead of the infimum

⁹When the Klein disk is used as the model for the hyperbolic plane these are in fact Euclidean triangles that meet the disk, but whose vertices may lie within the disk, on the ideal boundary, or outside the closed disk. The hyperideal vertices are the latter ones.

since curvatures are inversely related to radii. Of course by β -packing label I mean that the angle sum at any interior vertex that is not a branch vertex is 2π , and at v_i is $2\pi m_i$. The argument is akin to that of the proof of the Koebe Circle Packing Theorem, but, though still elementary, is much more intricate and involved. The full detailed proof appears in [18] where the key proposition, stated below, generalizes the Discrete Boundary Value Theorem of Beardon and Stephenson. Setting up this result with appropriate definitions and analysis of hyperideal hyperbolic triangles, as well as the proof itself, takes up most of the content of the paper.

Proposition 5.3.2 (Bowers and Stephenson [18]) *Let g be a proper boundary label for T and β a branch structure for (T, Φ_T) . Then there exists a unique β -packing label \mathfrak{k} for (T, Φ_T) such that $\mathfrak{k}(w) = g(w)$ for every boundary vertex of T .*

This then is used to complete the proof of the Polynomially Branched KAT Theorem by using the circle packing produced by Proposition 5.3.2, augmented by the unit circle corresponding to the removed vertex v_∞ to define \mathcal{C} . Much of this becomes routine at this point, except that one still must confirm that half the branching occurs at v_∞ . This turns out to be nontrivial. Again the details are rather involved and can be found in [18]. □

5.3.4 Cusps and Cone Type Singularities

In this section I offer a generalization of KAT II where prescribed target angle sums at vertices are assigned, and necessary and sufficient conditions are sought to guarantee existence of such packings. This is the discrete version of the classical Schwarz-Picard problem of the existence of hyperbolic metrics on Riemann surfaces with prescribed cone type singularities. For simplicity I am going to restrict to the tangency case where Φ is identically zero.

To set up the problem, let K be a simplicial triangulation of a compact surface, possibly with boundary, with F faces, E edges, and V vertices. The vertex set $V(K)$ is partitioned into three sets: two disjoint subsets of interior vertices denoted as V_{Int} and V_{cusps} , and the set V_{Bd} of boundary vertices, with respective cardinalities V_{Int} , V_{cusps} , and V_{Bd} . Elements of V_{Int} are called *interior vertices* and of V_{cusps} are called *cusp vertices*. Two functions are given, the first $f : V_{\text{Bd}} \rightarrow (0, \infty]$ giving target radii for the boundary vertices and the second $\theta : V_{\text{Int}} \rightarrow (0, \infty)$ giving target angle sums at interior vertices. The target angle sums at the *cusp vertices* in V_{cusps} are zero. The task is to give necessary and sufficient conditions on K to guarantee the existence of a packing label $r : V(K) \rightarrow (0, \infty]$ for this data such that $r = f$ on V_{Bd} , $r = \infty$ on V_{cusps} , and $\theta_r(v) = \theta(v)$ for every interior vertex $v \in V_{\text{Int}}$.

To describe a solution to this problem, for any set V of vertices, let F_V denote the number of faces of K that meet V , and let $\theta(V) = \sum_{v \in V} \theta(v)$ denote the total angle sum of the vertices of V . Let

$$\mathfrak{R} = \{r : V(K) \rightarrow (0, \infty] : r = f \text{ on } V_{\text{Bd}}, r = \infty \text{ on } V_{\text{cusps}}, \theta_r(v) \leq \theta(v) \text{ for all } v \in V_{\text{Int}}\}.$$

This describes the set of *superpacking labels* for the data θ with boundary values given by f and cusp set V_{cusps} . A *packing label* for this data is a superpacking label where, in addition, the target angle sums given by θ are met, so that $\theta_r(v) = \theta(v)$ for all $v \in V_{\text{Int}}$. For any superpacking label r and vertex set V , let $\theta_r(V) = \sum_{v \in V} \theta_r(v)$. The next theorem gives necessary and sufficient conditions for a solution to the discrete Schwarz-Picard boundary value problem. The proof is a generalization of the proof presented herein for the Koebe Circle Packing Theorem. There the important invariant is $F - 2V_{\text{Int}}$. In the borderless case of Theorem 5.3.1, the important invariant is $F - 2V$. These arise from writing the hyperbolic area of the surface determined by a packing label, provided one exists, in terms of combinatorial invariants. The corresponding fact in this setting is that, for any packing label τ for the data f, θ , and V_{cusps} ,

$$\text{hyp-area}(K(\tau)) + \theta_\tau(V_{\text{Bd}}) = \pi F - \theta_\tau(V_{\text{Int}}) = \pi F - \theta(V_{\text{Int}}).$$

The right hand side of this equation is an invariant of K and θ and must be positive since the left hand side is positive. Also, for every interior vertex v ,

$$\theta(v) = \theta_\tau(v) < \pi \deg v$$

These give two necessary conditions for a desired label to exist, but these are not sufficient. Nonetheless, these two conditions are the extreme cases of the sufficient condition that appears as item (i) of the theorem.

Discrete Schwarz-Picard Boundary Value Theorem (Bowers [12]) *The following are equivalent.*

- (i) *For every edge-path connected set $V \subset V_{\text{Int}}$ of interior vertices, the invariant $\pi F_V - \theta(V)$ is positive.*
- (ii) *The function $\tau = \inf \mathfrak{R}$ does not take a zero value at any vertex.*
- (iii) *The function $\tau = \inf \mathfrak{R}$ is the unique packing label for K with data f, θ , and V_{cusps} .*
- (iv) *There exist a packing label for K for the data f, θ , and V_{cusps} .*

A word of caution is in order. Though this does solve the discrete Schwarz-Picard problem, the combinatorial condition of item (i), that $\pi F_V - \theta(V) > 0$ for every path connected subset V of interior vertices, is a very difficult condition to check once the size of K becomes in any way substantial. This pure mathematician has learnt to appreciate the difficulties our computational geometer cousins face when trying to make the elegant output of our theorems practical tools for performing geometric computations. This difficulty often is unrecognized or left unacknowledged by my pure mathematician siblings.

5.4 Infinite Packings of Non-compact Surfaces

I now turn our attention to infinite packings of non-compact surfaces. Here new and interesting phenomena arise, fraught with their own peculiar difficulties. To keep the conversation manageable, I am restricting attention to tangency circle packings of simply connected domains and will concentrate on one very interesting problem that arises in this setting—the *type problem*—and one great success in attacking the Koebe Uniformization Conjecture.

5.4.1 The Discrete Uniformization Theorem

Does every simplicial triangulation K of every topological surface S , compact or not, admit a circle packing in some geometric structure on S ? By passing to the universal covering surface \tilde{S} and lifting the triangulation to a triangulation \tilde{K} of \tilde{S} , the question may be approached by asking whether any G -invariant simplicial triangulation of a simply connected surface admits a G -invariant circle packing in some geometric structure, where G is a group of symmetries of the complex. There are only two simply connected surfaces up to homeomorphism, the sphere and the plane. The former case is addressed by the Koebe Circle Packing Theorem. In this section I will address the latter case.

Let \mathcal{T} be a *plane triangulation graph*, by which I mean that \mathcal{T} is the 1-skeleton of a simplicial triangulation K of the topological plane. There are precisely two inequivalent conformal structures on the plane, the one conformally equivalent to the complex plane \mathbb{C} and the other to the open unit disk \mathbb{D} . There are precisely two complete metrics of constant curvature up to scaling on the plane, the one isometric to Euclidean 2-space \mathbb{E}^2 and of constant zero curvature, the other isometric to the hyperbolic plane \mathbb{H}^2 and of constant negative curvature. Fortunately, the conformal and the geometric structures mesh nicely in that the complex plane \mathbb{C} is a conformal model of plane Euclidean geometry via its standard Euclidean metric $ds_{\mathbb{C}} = |dz|$, and the disk \mathbb{D} is a conformal model of plane hyperbolic geometry via the Poincaré metric $ds_{\mathbb{D}} = 2|dz|/(1 - |z|^2)$. Metric circles in these two geometries are precisely the Euclidean circles contained in their point sets, so circle packings in these geometric surfaces can be identified with Euclidean circle packings of \mathbb{C} and \mathbb{D} . I will use \mathbb{G} ¹⁰ to mean either \mathbb{C} or \mathbb{D} with the intrinsic Euclidean or hyperbolic geometry determined by either $ds_{\mathbb{C}}$ or $ds_{\mathbb{D}}$ when referring to geometric quantities like geodesics and angles, etc. Here is the foundational result in this setting.

Discrete Uniformization Theorem (Beardon and Stephenson [7], He and Schramm [42]) *Every plane triangulation graph \mathcal{T} can be realized as the contacts graph of a univalent circle packing $\mathcal{T}(\mathbb{C})$ that fills exactly one of the complex plane*

¹⁰ \mathbb{G} means \mathbb{G} eometry.

\mathbb{C} or the disk \mathbb{D} . The packing is unique up to conformal automorphisms of either \mathbb{C} or \mathbb{D} .

The *contacts* graph of a collection is a graph with a vertex for each element of the collection and an edge between two vertices if and only if the corresponding elements meet. The *carrier* of the circle packing \mathcal{C} in the geometry \mathbb{G} is the union of the geodesic triangles formed by connecting centers of triples of mutually adjacent circles with geodesic segments, and \mathcal{C} *fills* \mathbb{G} whenever its carrier is all of \mathbb{G} . When \mathcal{C} is univalent and fills \mathbb{G} , \mathcal{C} is said to be a *maximal packing* for \mathcal{T} or K , and K may be realized as a geodesic triangulation of \mathbb{G} whose vertices are the centers of the circles of \mathcal{C} with geodesic edges connecting adjacent centers.

Once this theorem is in place, the whole of the theory of tangency circle packings on non-compact surfaces comes into play. As already indicated, in a thoroughly classical way packing questions on surfaces can be transferred to questions of packings on simply connected surfaces, this by passing to covering spaces acted upon by groups of deck transformations. Any combinatorial symmetries of the complex K are realized as automorphic symmetries of \mathbb{G} , this from the uniqueness of the Discrete Uniformization Theorem, and this offers an alternate Proof of Theorem 5.3.1, and an extension of that theorem to triangulations of arbitrary, non-compact surfaces.

Beardon and Stephenson [7] proved the Discrete Uniformization Theorem when \mathcal{T} has *bounded degree*, a global bound on the degrees of all the vertices of \mathcal{T} . In this foundational paper as well as in their subsequent one [8], Beardon and Stephenson laid out a beautiful theory of circle packings on arbitrary surfaces, gave a blueprint for developing a theory of discrete analytic functions, and articulated one of the most interesting problems in the discipline, that of the *circle packing type problem* for non-compact surfaces, this latter the subject of the section following. The bounded degree assumption was needed both to verify that the packing fills \mathbb{G} and for the uniqueness, and He and Schramm [42] removed the bounded degree hypothesis and proved the general case where there is no global bound on the degrees of vertices. Earlier, Schramm [65] had proved a very general rigidity theorem for infinite packings of planar domains whose complementary domains are a countable collection of points, and He and Schramm [42] extended this to general countably connected domains.

Discussion of Proof The full proof is scattered throughout several articles published in the 1990s. In what constitutes a significant service to the discipline, Ken Stephenson has laid out a complete proof in roughly fifty pages of his wonderful text *Introduction to Circle Packing* [71]. I have not the space here to do justice to the argument, but I will make some comments.

Beardon and Stephenson's proof of existence relies on the Maximal Disk Packing Theorem and uses a diagonal argument on a sequence of finite subcomplexes of K that exhausts K . It does not depend on any bounded degree assumption and is quite straightforward. The proof of existence goes like this. Write $K = \bigcup_{i=1}^{\infty} K_i$ as a nested, increasing union of finite subcomplexes K_i , each a simplicial triangulation of a closed disk. Apply the Maximal Disk Packing Theorem to obtain a sequence

\mathcal{C}_i of univalent, maximal circle packings for the complexes K_i in the unit disk \mathbb{D} realized as the Poincaré disk model of hyperbolic geometry. Fix a base vertex v_0 of K_1 and let C_i be the circle of \mathcal{C}_i that corresponds to v_0 . By applying an automorphism of the disk if needed, assume that C_i is centered at the origin and of hyperbolic radius $\tau_i(v_0)$. Now the Discrete Schwarz-Pick Lemma implies that the sequence $\tau_i(v_0)$ of hyperbolic radii is non-increasing, hence has a limit, say $\tau(v_0) \geq 0$, as $i \rightarrow \infty$. There are two cases.

- (I) The limit radius $\tau(v_0) \neq 0$;
- (II) The limit radius $\tau(v_0) = 0$.

The first claim is that if v is any other vertex of K whose corresponding circle of \mathcal{C}_i , for large enough i , has hyperbolic radius $\tau_i(v)$, then $\lim_{i \rightarrow \infty} \tau_i(v)$ is not zero when case (I) occurs and is equal to zero when case (II) occurs. This means that the limit radius function $\tau : V(K) \rightarrow [0, \infty)$ never takes a zero value in case (I) and is identically zero in case (II). The proof of this claim uses the *Ring Lemma* of Burt Rodin and Dennis Sullivan that was crucial in [62] in their confirmation of Thurston’s outlined proof of the *Discrete Riemann Mapping Theorem* presented in his 1985 Purdue lecture; see Sect. 5.5.1. The Ring Lemma guarantees the existence of a sequence of positive constants c_d such that, when $d \geq 3$ disks form a cycle of sequentially tangent disks all tangent to a central disk of Euclidean radius R , and the disks have pairwise disjoint interiors, then the smallest disk has Euclidean radius $\geq c_d R$. The Ring Lemma is applied as follows. Let $v_0 v_1 \cdots v_n = v$ be a path of vertices in K from v_0 to v and choose N so large that this path of vertices is contained in the interior of K_i , for all $i \geq N$. The Ring Lemma applied sequentially to the chain of pairwise tangent circles in \mathcal{C}_i corresponding to the path $v_0 v_1 \cdots v_n = v$ implies that there is a positive constant c such that $R_i(v) \geq c R_i(v_0)$, where R_i is the Euclidean radius function on \mathcal{C}_i . This holds for all $i \geq N$ and the constant c is independent of i . As hyperbolic and Euclidean radii of circles in the disk are comparable in the small, this implies the claim.

Now order the vertex set $V(K)$ as v_0, v_1, \dots . In case (I), choose a subsequence i_j so that the hyperbolic centers of the circles of the sequence \mathcal{C}_{i_j} all corresponding to the vertex v_1 converge in the closed disk $\overline{\mathbb{D}}$ to a point c_1 . An application of item (ii) of the Discrete Schwarz Pick Lemma implies that c_1 is contained in the open disk \mathbb{D} . Repeat to find a subsequence of i_j for which the hyperbolic centers of the circles corresponding to v_2 converge to a point c_2 in \mathbb{D} . Iterating and applying a diagonal argument gives a subsequence of the sequence of circle packings \mathcal{C}_i for which the hyperbolic centers of the circles corresponding to the vertex v_n of K converges to a point c_n in \mathbb{D} for all positive integers n . Centering a circle of hyperbolic radius $\tau(v_n)$ at the point c_n produces a circle packing in the Poincaré disk in the pattern of K . In case (II) when τ is identically zero, a diagonal argument applied to the scaled packing $\frac{1}{R_i} \mathcal{C}_i$, where R_i is the Euclidean radius of \mathcal{C}_i , produces a circle packing in the plane \mathbb{C} in the pattern of K . Call the limit circle packing in either case \mathcal{C} .

There are three facts left to prove: first, that \mathcal{C} is univalent; second, that \mathcal{C} fills the disk in case (I) and the plane in case (II); third, that \mathcal{C} is unique up

to automorphisms. The first claim of univalence follows from the fact that each circle packing \mathcal{C}_i is univalent and the convergent subsequence of radii and centers described above essentially describes geometric convergence of circle packings. Beardon and Stephenson's original proof of the second claim that the packing fills \mathbb{G} relied critically on the bounded degree assumption. It was used to ensure that piecewise linear maps from the complexes K_i into the geometry \mathbb{G} defined using the convergent sequence of circle packings are uniformly quasiconformal so that the Carathéodory Kernel Theorem [27] applies to ensure that the image of the limit function is the kernel of the image sets, which is the whole of \mathbb{G} . The third claim of uniqueness in the hyperbolic case (I) follows from the uniqueness of the limiting radius function, but in the Euclidean case (II), uniqueness uses the bounded degree assumption. Later He and Schramm removed the bounded degree assumption. Their proof of uniqueness in case (II) is particularly elegant. It is a topological proof based on the winding numbers of mappings defined on the boundaries of corresponding interstitial regions in two circle packings for the same complex K , both of which fill \mathbb{C} . All of this is rather nicely laid out in Stephenson's *Introduction to Circle Packing* [71]. \square

5.4.2 Types of Type

The dichotomy between hyperbolic and Euclidean behavior is evident in the Discrete Uniformization Theorem. Indeed, the combinatorial complex K , or its 1-skeleton \mathcal{T} , determines uniquely its geometry in that the maximal circle packing $\mathcal{T}(\mathbb{C})$ fills either the disk \mathbb{D} or the complex plane \mathbb{C} , but forbids two packings where one fills the disk and the other the plane. This leads to the next definition.

Definition (CP-Type) A simplicial triangulation K of the plane, and its 1-skeleton plane triangulation graph $\mathcal{T} = K^{(1)}$, are said to *CP-parabolic* or *CP-hyperbolic* when the maximal circle packing $\mathcal{T}(\mathbb{C})$ fills respectively the complex plane \mathbb{C} or the disk \mathbb{D} . The *CP-type problem* is the problem of determining whether a given complex K or plane triangulation graph \mathcal{T} is CP-parabolic or CP-hyperbolic. One seeks conditions or invariants on the complex K or the graph \mathcal{T} , reasonably checked or computed, that can determine which of the two CP-types adheres. See Fig. 5.5.

This is a discrete version of the classical *conformal type problem*, or just *type problem* for short, that of determining whether, à la classical Uniformization Theorem, a given non-compact simply connected Riemann surface is *parabolic* and conformally equivalent to the complex plane \mathbb{C} , or *hyperbolic* and conformally equivalent to the disk \mathbb{D} .

Historically this is not the first discrete type problem. That honor probably goes to the problem of determining the *random walk type*, or *RW-type* for short, of an infinite graph. My aim in this section is to review this and several other species of discrete type problems and explore their interactions in the context of plane




Fig. 5.5 The CP-type of a plane triangulation graph is determined by the corresponding maximal circle packing and whether it fills the plane or the disk. **(a)** The penny packing, the maximal circle packing for the constant 6° plane triangulation graph \mathcal{G}_6 , fills the plane \mathbb{C} . The graph \mathcal{G}_6 is parabolic. **(b)** The maximal circle packing for the constant 7° plane triangulation graph \mathcal{G}_7 fills the disk \mathbb{D} . The graph \mathcal{G}_7 is hyperbolic, as are the graphs \mathcal{G}_d for all $d \geq 7$

triangulation graphs. In all I will examine six different species of discrete type that go under the abbreviations CP, RW, EL, EEL, VEL, and EQ-type.

Consider the standard simple random walk on a simple, connected, locally finite graph \mathcal{G} where the probability of walking across a particular edge uv from vertex u to vertex v is $1/\deg u$. The graph \mathcal{G} is said to be *random walk parabolic*, or *RW-parabolic*, if a walker almost surely returns to a fixed base vertex, and *random walk hyperbolic* or *RW-hyperbolic* otherwise. More common terminology is that the graph is *recurrent* when RW-parabolic and *transient* when RW-hyperbolic. In a transient graph, a random walker has a positive probability for escaping to infinity whereas in a recurrent one, the escape probability vanishes and, in fact, the walker almost surely returns to every vertex infinitely often. Woess [74] is a fantastic reference for the classical theory of random walks on graphs and Lawler and Limic [51] is an up-to-date reference with many recent results.

Early on in the development of circle packing theory, Ken Stephenson made a connection between the CP- and RW-type problems. The intuition for the connection arises from the close connection in classical complex function theory between the conformal type problem and brownian motion on a Riemann surface. Stephenson [70] proved that the CP- and RW-types of bounded degree plane triangulation graphs always coincide. Later in [43], He and Schramm gave an example of a plane triangulation graph, necessarily of unbounded degree, that is CP-parabolic but RW-hyperbolic. There the authors focused more sharply on the distinction between these two species of type and recalled Duffin's EEL-type from [35] and developed Cannon's VEL-type inspired by Cannon [25] in articulating the distinction.

Before continuing with the discussion of CP-type, let's review a bit of history. The story of discrete type really begins in the 1920s with Pólya's study [57] of the RW-type of the integer grid in \mathbb{R}^d where he proved that the integer grid in \mathbb{R}^2 is RW-parabolic while the grid in higher-dimensional Euclidean spaces is RW-hyperbolic. In 1959, Nash–Williams in [55] used a method of Lord Rayleigh to link the RW-type of a locally finite graph with its resistance to electric flow when each edge is thought of as a wire with a unit of electrical resistance, giving rise to EL-type. To be a bit less cryptic, when an infinite graph is thought of as an electric network with each edge representing a wire of unit resistance, the question is whether electricity will flow from a base vertex to infinity when a unit potential is applied to the base vertex and infinity is grounded. This is made a bit more precise by asking what the effective resistance is from the base vertex to infinity for the network. When the effective resistance to infinity is infinite, no current flows and the network is *EL-parabolic*, and when the resistance is finite so that current does flow, the network is *EL-hyperbolic*. In the beautiful 1984 Carus Mathematical Monograph [34] entitled *Random Walks and Electric networks*, Peter Doyle and J. Laurie Snell present an accessible proof that the RW- and EL-type of an infinite graph coincide. In 1962, Duffin [35] gave a combinatorial invariant of a graph, the *edge extremal length*, that characterizes the RW- or EL-type according to whether the edge extremal length of the set of transient edge-paths is infinite or not.

What is the edge extremal length of a path family? It is a discrete version of the classical conformal extremal length of a path family in a Riemann surface in quasiconformal analysis. For a graph \mathcal{G} , let Γ be any family of edge-paths, infinite or not. The edge extremal length is obtained by measuring the minimal length-squared of the curves in Γ divided by the area, this maximized over all metric assignments. This is the same as the classical definition, only what changes is how the admissible metrics are assigned. Here are the details. An *edge-path* in \mathcal{G} is a finite or infinite sequence $\mathbf{e} = e_1, e_2, \dots$ of directed edges of \mathcal{G} with the terminal vertex of e_i equal to the initial vertex of e_{i+1} . An *edge metric* on \mathcal{G} is a function $m : E(\mathcal{G}) \rightarrow [0, \infty]$ that assigns a non-negative value to each edge, and the *area of m* is defined as $\text{area}(m) = \sum_{e \in E(\mathcal{G})} m(e)^2$. An edge metric is *admissible* if its area is finite and I will let $M_E(\mathcal{G})$ denote the collection of admissible edge metrics. The *m -length* of the edge-path \mathbf{e} is $\ell_m(\mathbf{e}) = \sum_{i=1} m(e_i)$. Finally, the *edge extremal length* of the family Γ of edge-paths is

$$\text{EEL}(\Gamma) = \sup_{m \in M_E(\mathcal{G})} \frac{\inf_{\mathbf{e} \in \Gamma} \ell_m(\mathbf{e})^2}{\text{area}(m)}.$$

The notation $\text{EEL}(\mathcal{G})$ is reserved for the case where Γ is the set of paths to infinity that start at a given base vertex v_0 . These are called the *transient* edge-paths in \mathcal{G} based at v_0 , and any such transient edge-path $\mathbf{e} \in \Gamma$ has initial vertex v_0 at its first edge e_1 and is not contained in any finite collection of edges. One says that the graph \mathcal{G} is *EEL-parabolic* if $\text{EEL}(\mathcal{G}) = \infty$ and *EEL-hyperbolic* otherwise. It is an easy exercise to confirm that EEL-type does not depend on which base vertex is

chosen. Duffin’s result of [35] already mentioned is that both the RW- and EL-type of a graph coincides with the EEL-type. This was the state of the art in discrete type in the early 1990s when Stephenson connected CP-type with RW-type for bounded degree plane triangulation graphs.

In 1995, He and Schramm [43] in a remarkable article clarified the role of the bounded degree assumption. There, after constructing a plane triangulation graph that, though CP-parabolic, is RW-hyperbolic, they applied Cannon’s vertex extremal length to characterize CP-type combinatorially in the way that edge extremal length characterizes RW-type. Cannon [25] introduced the vertex extremal length of a discrete curve family made of shinglings and used it as a tool for assigning combinatorial moduli to ring domains in the space at infinity of a negatively curved group. He and Schramm adapted Cannon’s vertex extremal length to Duffin’s development of EEL-type to create VEL-type. The adjustment merely replaces edge-paths by vertex-paths and edge metrics by vertex metrics. The *vertex extremal length* of a family Δ of vertex paths is

$$\text{VEL}(\Delta) = \sup_{m \in M_V(\mathcal{G})} \frac{\inf_{\mathbf{v} \in \Delta} \ell_m(\mathbf{v})^2}{\text{area}(m)}. \tag{5.4.1}$$

Here, a *vertex-path* is a sequence $\mathbf{v} = v_1, v_2, \dots$ where each v_i is incident with its successor v_{i+1} , and a *vertex metric* is a non-negative function $m : V(\mathcal{G}) \rightarrow [0, \infty]$ with $\text{area}(m) = \sum_{v \in V(\mathcal{G})} m(v)^2$. The m -length of the vertex-path \mathbf{v} is $\ell_m(\mathbf{v}) = \sum_{i=1} m(v_i)$ and the set of *admissible metrics*, the ones of finite area, is denoted as $M_V(\mathcal{G})$. The VEL-type of \mathcal{G} now is defined analogously to EEL-type. Indeed, $\text{VEL}(\mathcal{G})$ means $\text{VEL}(\Delta)$, where Δ is the set of *transient* vertex-paths based at v_0 , those that meet infinitely many vertices. The graph \mathcal{G} is *VEL-parabolic* if $\text{VEL}(\mathcal{G}) = \infty$ and *VEL-hyperbolic* otherwise, and again it is an easy exercise to confirm that VEL-type does not depend on which base vertex is chosen. This seemingly innocuous adjustment to the definition of EEL-type turns out to be precisely the tool needed to characterize CP-type.

Though, easily, the EEL- and VEL-types of a bounded degree graph coincide, they may differ for a graph of unbounded degree. The relationships between the four types—RW, EL, EEL, VEL—are summarized in the next theorem.

Discrete Type Theorem for Graphs *Let \mathcal{G} be a connected, infinite, locally finite graph.*

- (i) [Nash–Williams [55], Duffin [35]] *The three types—RW, EL, EEL—coincide for \mathcal{G} .*
- (ii) [He–Schramm [43]] *If \mathcal{G} is EEL-parabolic then it is VEL-parabolic. If \mathcal{G} has bounded degree and is VEL-parabolic, then it is EEL-parabolic.*
- (iii) [He–Schramm [43]] *There is a VEL-parabolic plane triangulation graph that is EEL-hyperbolic, necessarily of unbounded degree.*

For a plane triangulation graph \mathcal{T} , all five types—RW, EL, EEL, VEL, CP—coincide provided \mathcal{T} has bounded degree. As stated above, it was Stephenson who

first proved this for RW- and CP-types. He and Schramm clarified the need for the bounded degree hypothesis, and the relationship between discrete types for plane triangulation graphs is summarized next.

Discrete Type Theorem for Plane Triangulation Graphs (He–Schramm [43])

Let \mathcal{T} be a plane triangulation graph. Then \mathcal{T} is CP-parabolic if and only if it is VEL-parabolic.

The proofs of these theorems are quite difficult and involved, though still elementary, and space forbids any sort of discussion of the proofs that would do justice to the subject. Suffice it to say that the interested reader can do no better than to consult the references cited in this section to fill in gaps in the desired detail of proofs.

The Discrete Type Theorem for Plane Triangulation Graphs reduces the very difficult problem of determining whether the maximal circle packing for \mathcal{T} is parabolic or hyperbolic to a combinatorial computation on the graph \mathcal{T} . The disappointment comes when one actually tries to do the computation of $\text{VEL}(\mathcal{T})$ from Eq. 5.4.1 for almost any given plane triangulation graph. One then finds out just how difficult it is to perform this computation; nonetheless, this development is useful for some theoretical considerations. For example, He and Schramm use the theorem to extend Stephenson’s result on RW- and CP-type. Here is an interesting result of the author that uses the computation of Eq. 5.4.1 for the proof of item (ii) of the theorem.

Theorem 5.4.1 (Bowers [13]) *Let \mathcal{G} be a connected, infinite, locally finite graph and \mathcal{T} a plane triangulation graph.*

- (i) *If \mathcal{G} is Gromov negatively curved and its Gromov boundary contains a nontrivial continuum, then \mathcal{G} is RW-hyperbolic.*
- (ii) *If \mathcal{T} is Gromov negatively curved, then \mathcal{T} is CP-parabolic if and only if its Gromov boundary is a singleton; alternately, it is CP-hyperbolic if and only if its Gromov boundary is a topological circle.*

I refer the reader to the appendix of the article [13] for definitions and basic theorems on Gromov negatively curved graphs and metric spaces. To show how the computation from Eq. 5.4.1 may proceed, I’ll prove the lemma used in [13] to prove the first assertion of item (ii) of Theorem 5.4.1.

Lemma 5.4.2 *Let v_0 be a vertex in the connected, infinite, locally finite graph \mathcal{G} and let $\{V_n\}$ be a sequence of pairwise disjoint sets of vertices, each of which separates v_0 from infinity. Suppose there exist positive constants C and ε such that, for $n \geq N$,*

$$\text{Card}(V_n) \leq Cn.$$

Then the graph \mathcal{G} is VEL-parabolic.

Proof Define the vertex metric m by $m(v) = 1/(n \log n)$ for any $v \in V_n$ when $n \geq N$, and $m(v) = 0$ otherwise. Then m is admissible since

$$\text{area}(m) = \sum_{n=N}^{\infty} \frac{\text{Card}(V_n)}{(n \log n)^2} \leq \sum_{n=N}^{\infty} \frac{C}{n(\log n)^2} < \infty.$$

For any transient vertex-path \mathbf{v} , the m -length satisfies $\ell_m(\mathbf{v}) \geq \sum_{n=N}^{\infty} 1/(n \log n) = \infty$, hence every transient vertex-path has infinite m -length. This implies that $\text{VEL}(\mathcal{G}) = \infty$ and \mathcal{G} is VEL-parabolic. \square

I'll end this section with a sixth version of discrete type that is of recent interest in several settings. It arose first for me when Ken Stephenson and I constructed expansion complexes of finite subdivision rules, for the first time in [19] when examining the pentagonal subdivision rule of Cannon, Floyd, and Parry [26]. More recently it arises in our examination of hierarchical conformal tilings [21, 22], and in Gill and Rohde's [37] examination of random planar maps. I name this version of discrete type *EQ-type* with *EQ* an abbreviation for *equilateral*. A plane triangulation graph $\mathcal{T} = K^{(1)}$ can be used to build a piecewise equilateral surface by setting each edge to unit length and isometrically gluing unit-sided equilateral triangles along their boundaries to the boundaries of the faces of K . This produces a piecewise flat surface $|\mathcal{T}|_{\text{eq}}$ that has a natural conformal atlas obtained as follows. Each edge e of \mathcal{T} indexes a chart map φ_e defined on the interior of the union of the faces incident with e . These have been identified with unit equilateral triangles and the chart map φ_e is an orientation-preserving isometry to the plane \mathbb{C} . Each vertex v also indexes a chart map φ_v defined on the metric neighborhood of v in $|\mathcal{T}|_{\text{eq}}$ of radius $1/2$, and uses an appropriate complex power mapping to flatten that neighborhood to a disk in the plane \mathbb{C} . The overlap maps are conformal homeomorphisms between the appropriate domains. The chart family $\mathcal{A} = \{\varphi_x : x \in V(\mathcal{T}) \cup E(\mathcal{T})\}$ forms a complex atlas making $|\mathcal{T}|_{\text{eq}}$ into a non-compact simply connected Riemann surface $S(\mathcal{T})$. The type problem now is manifest. Is $S(\mathcal{T})$ conformally the plane \mathbb{C} or the disk \mathbb{D} ? In the former case, \mathcal{T} and K are said to be *EQ-parabolic*, in the latter *EQ-hyperbolic*.

Notice that the question of the EQ-type of a plane triangulation graph is the classical question of the conformal type of a simply connected Riemann surface. It bares the moniker *discrete* because of how the surface is built—using discrete building blocks, the equilateral triangles, glued in a combinatorial pattern encoded in \mathcal{T} . The desire is for a combinatorial invariant of \mathcal{T} or K that will determine its EQ-type. So, what relationship exists between the discrete types already discussed and EQ-type? For plane triangulation graphs of bounded degree, easy arguments using quasiconformal mappings show that EQ-type coincides with CP-type—just map the equilateral triangle in $|\mathcal{T}|_{\text{eq}}$ at face f to the corresponding geodesic triangle in \mathbb{G} . When \mathcal{T} has bounded degree, this map is uniformly quasiconformal and so the EQ-type agrees with the conformal type of \mathbb{G} . For unbounded degree plane triangulation graphs, it remains an open question as to whether the EQ-type coincides with,

say, the EEL- or the VEL-type, or perhaps neither. I am bold enough to offer the following conjecture.

Conjecture 5.4.3 *For any plane triangulation graph, EQ-type coincides with VEL-type, and therefore with CP-type.*

A great reference for various expressions of discrete type and their stability under subdivision is Bill Wood's doctoral thesis [75] and the subsequent article [76]. I now turn our attention to Koebe's original inspiration for his circle packing theorem, his interest in circle domains, uniformization, and the Kreisnormierungsproblem.

5.4.3 *Koebe Uniformization for Countably-Connected Domains*

Zheng-Xu He and Oded Schramm's work on circle packing in the late 1980s and early 1990s led them to a study of Koebe's Uniformization Conjecture. Though the discrete circle packing tools they developed and used did not directly apply to Koebe's problem, the perspective they had gained turned out to be useful. By 1992–1993, they had made the greatest advance on Koebe's problem since its articulation and had proved a circle packing version that greatly generalized the Discrete Uniformization Theorem. Their work is detailed in the *Annals of Mathematics* article *Fixed points, Koebe uniformization, and circle packings*. The proofs are rather intricate and so I am content to state the two main results without any indication of the proofs, leaving it to the interested reader to peruse [42] for details.

He–Schramm Uniformization Theorem (He and Schramm [42], Schramm [67])
Every countably connected domain in the Riemann sphere is conformally homeomorphic to a circle domain. Moreover, the circle domain is unique up to Möbius transformations and every conformal automorphism of the circle domain is the restriction of a Möbius transformation.

A domain triangulation graph is the 1-skeleton of a simplicial triangulation of a planar domain.

He–Schramm Discrete Uniformization Theorem (He and Schramm [42])
Every domain triangulation graph with at most countably many ends has a univalent circle packing in the plane \mathbb{C} whose carrier is a circle domain. Moreover, the circle packing is unique up to Möbius transformations.

He and Schramm prove a theorem that generalizes their Uniformization Theorem to *generalized domains* and *generalized circle domains*. This more general uniformization theorem then is used to give a quick proof of their Discrete Uniformization Theorem.

I'll close this section by mentioning that Schramm in a 1995 paper [67] introduced the notion of *transboundary extremal length* that generalizes the classical extremal length of curve families. Transboundary extremal length is more suited to path families in multiply connected domains that allow for the curves of the family to pass through the complementary components of the domain. Using this tool, Schramm gives a short proof of Koebe uniformization of countably connected domains and generalizes it in two ways. First, he shows that circle domains as the target of uniformization may be replaced by more general domains, namely, those where the complementary components are what he calls τ -fat sets. Second, he shows that some domains with uncountably many complementary components may be uniformized to circle domains, namely those where the complementary components are uniformly fat. This includes for example domains whose boundary components are points and μ -quasircircles for a fixed constant $\mu \geq 1$.

5.5 Some Theoretical Applications

The theoretical work in circle packing has grown up hand-in-hand with various applications. In the past score of years, the needs of computer imaging have added a practical bent to the applications with the use of the theory for everything from medical imaging to 3D-printer head guidance. This has been one of the impetuses for the development of the discipline of discrete differential geometry with discrete conformal geometry as but one of its chapters. Circle packing theory à la Thurston as described in this chapter is one flavor of this, but several groups of computational geometers and computer scientists have developed discrete conformal geometry in a great variety of ways, with new techniques designed to solve both practical and theoretical problems. The discipline has grown to a vast enterprise too large and complicated for a review of this type. Rather than attempt a thorough discussion of these applications, I'll only mention a couple of the theoretical applications. The first stands as one of the linchpins of the discipline, and the second generalizes the first. I'll leave it for the interested reader to peruse the many resources available to learn of the state of the art today in practical applications.

5.5.1 Approximating the Riemann Mapping

The event that really got circle packing launched, piquing the interest of a small group of research mathematicians from as diverse fields as complex function theory, combinatorial and computational geometry, geometric topology, and the classical theory of polyhedra, was Bill Thurston's address entitled *The Finite Riemann Mapping Theorem* at Purdue University in 1985. He presented there an algorithm for computing discrete versions of the Riemann mapping of a fixed, proper, simply connected domain in the complex plane \mathbb{C} to the unit disk \mathbb{D} , with an indication

of why the discrete mappings should converge to a conformal homeomorphism of the domain onto \mathbb{D} . Burt Rodin and Dennis Sullivan published in [62] a proof of Thurston's claims in 1987, and this began a steady output of published research on circle packings that continues today. Here I review the content of Thurston's 1985 talk and explain the Rodin-Sullivan verification of Thurston's claims.

Thurston's algorithm is illustrated nicely in the graphics of Fig. 5.6. The scheme is rather simple. Overlay a domain D with a hexagonal circle packing \mathcal{H}_ε of constant circle radii ε , a 'penny packing.' Use the domain D as a cookie cutter to cut out a portion of the packing, say \mathcal{P}_ε , whose combinatorics are given by the simplicial complex T_ε . Apply the Maximal Disk Packing Theorem to obtain a maximal circle packing \mathcal{Q}_ε of the disk \mathbb{D} . Choosing two points x and y in the domain D , let u_ε and v_ε be the vertices of T_ε whose corresponding circles are closest to the respective points x and y . From the uniqueness of the Maximal Disk Packing Theorem, one may assume that the packings \mathcal{Q}_ε have been normalized so that the circle corresponding to u_ε is centered at the origin and the one corresponding to v_ε is centered on the positive real axis. Define the *discrete Riemann mapping* $f_\varepsilon : \text{carr}(\mathcal{P}_\varepsilon) \rightarrow \text{carr}(\mathcal{Q}_\varepsilon)$ as the piecewise linear mapping that takes centers of circles of \mathcal{P}_ε to corresponding centers of circles of \mathcal{Q}_ε . Thurston's claim of his 1985 lecture that Rodin and Sullivan verified in 1987 is the content of the next theorem.

Discrete Riemann Mapping Theorem (Rodin and Sullivan [62]) *The mappings f_ε converge as $\varepsilon \downarrow 0$, uniformly on compact subsets of D , to the Riemann mapping f of D onto \mathbb{D} with $f(x) = 0$ and $f(y) > 0$.*

Before I discuss the proof, I should say that there is nothing special about the hexagonal combinatorics. He and Schramm [44] verified that the particular combinatorics of the overlay packings are irrelevant as long as the maximum circle radii approach zero.

Sketch of Proof The proof applies classical tools from quasiconformal analysis to confirm convergence of the discrete mappings to the Riemann mapping. There are three parts. First, the Ring Lemma, already used on page 198 in the proof of the Discrete Uniformization Theorem, is used to observe that the discrete Riemann mappings f_ε for $\varepsilon > 0$ form a family of uniformly quasiconformal mappings with, say, dilatation of all maps bounded by $\mu \geq 1$. Second, standard results of quasiconformal analysis imply that the mappings converge to a μ -quasiconformal mapping f of D onto \mathbb{D} . Third, the limit mapping is proved to be 1-quasiconformal, or just conformal, so that it is a Riemann mapping of the domain D onto the disk \mathbb{D} . Allow me to fill in each of the three parts of the argument a bit.

The first part, that the discrete Riemann mappings have quasiconformal distortion uniformly bounded, uses the fact that simplicial homeomorphisms are μ -quasiconformal with the distortion constant μ depending only on the shapes of the triangles involved. In particular, because the complexes T_ε have constant degree six on interior vertices, the Ring Lemma implies that there is a minimum possible angle $\omega > 0$ for any of the triangles in the Euclidean carrier $\text{carr}(\mathcal{Q}_\varepsilon)$, this independent of ε . This implies that the discrete maps f_ε are uniformly μ -quasiconformal since the

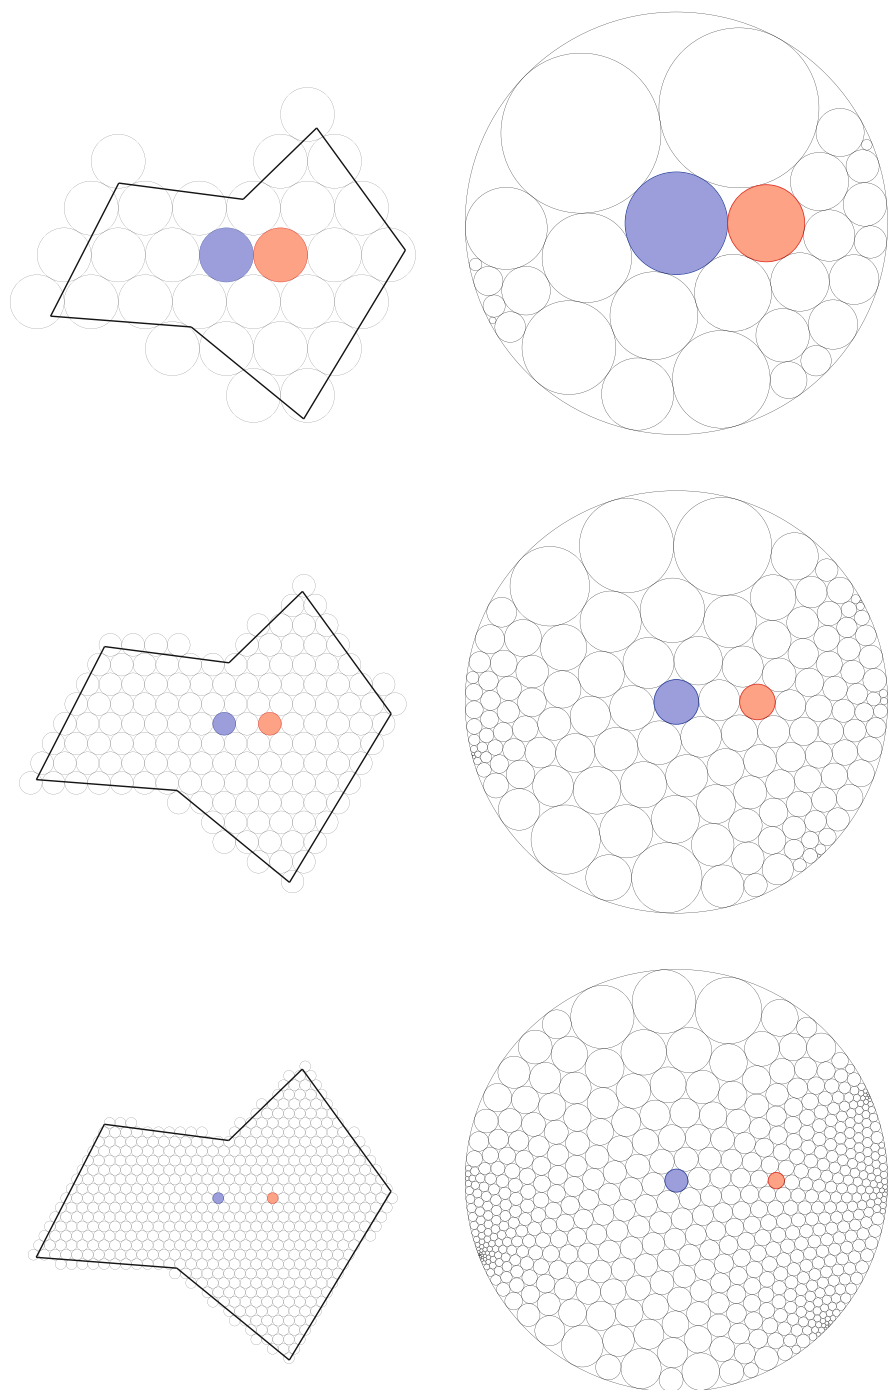


Fig. 5.6 Discrete Riemann mappings with finer and finer hexagonal circle packings

images of the equilateral triangles of $\text{carr}(\mathcal{P}_\varepsilon)$ are triangles of $\text{carr}(\mathcal{Q}_\varepsilon)$ of uniformly bounded distortion.

The second part now follows from standard tools of quasiconformal analysis. The uniformly quasiconformal maps f_ε are equicontinuous on compact subsets of D , as are the maps f_ε^{-1} on compact subsets of \mathbb{D} . It follows that the family $\{f_\varepsilon\}_{\varepsilon>0}$ is a normal family and any limit mapping f is bijective between D and \mathbb{D} . This latter claim uses the fact that any limit mapping is necessarily μ -quasiconformal, and the Carathéodory Kernel Theorem implies that f takes D onto \mathbb{D} .

Finally, that any limit mapping f is conformal follows from the *Hexagonal Packing Lemma*. This says that in a packing with hexagonal combinatorics, any two adjacent circles buried deeply within the packing have nearly equal radii. Here is the exact statement.

Hexagonal Packing Lemma (Rodin and Sullivan [62]) *There is a sequence c_n decreasing to zero as $n \rightarrow \infty$ such that in any packing with n generations of the regular hexagonal combinatorics surrounding circle C , the ratio of radii of C and any adjacent circle differs from unity by less than c_n .*

This lemma shows that as $\varepsilon \downarrow 0$, the mappings f_ε restricted to a fixed compact subset of D maps equilateral triangles to triangles of $\text{carr}(\mathcal{Q}_\varepsilon)$ that become arbitrarily close to equilateral, and this implies that any limit mapping is conformal.

This completes the proof modulo the proof of the Hexagonal Packing Lemma. This is proved as follows. Let H_n be any packing of circles in the plane with combinatorics given by greater than or equal to n generations of the hexagonal packing and whose central circle is the circle C_0 of unit radius centered at the origin. The Ring Lemma implies that the radii of the circles n generations removed from C_0 in the packings H_m for $m \geq n$ are bounded away from zero and infinity. A diagonal argument implies that there is a subsequence H_{n_i} that geometrically converges to a packing H , which necessarily has hexagonal combinatorics. But the uniqueness of the Discrete Uniformization Theorem implies that $H = \mathcal{H}_1$, the penny packing of unit radius. If the lemma were not true, one could choose the sequence H_n in such a way that the ratio of the center circle of H_n to at least one of its neighbors differs from unity by at least a fixed constant $\delta > 0$. This would imply that the limit packing H has a circle adjacent to C_0 of non-unit radius, contradicting uniqueness. \square

I should mention that Rodin and Sullivan did not have access to the Discrete Uniformization Theorem in 1987 as it was published only in 1990. They had to prove uniqueness of the penny packing of the plane, which they did by invoking results of Dennis Sullivan [72] extending the Mostow Rigidity Theorem to non-compact three-manifolds whose volumes grow slowly enough. This initiated an attempt to prove the Hexagonal Packing Lemma using only elementary means, which ultimately led to a better understanding of the rigidity of infinite circle packings over the next decade. This paper of Rodin and Sullivan was highly influential and can claim to be the genesis of the serious study of circle packings that now includes in its accomplishments hundreds of articles, thousands of citations, and a huge reservoir of applications in a great variety of different settings.

5.5.2 Uniformizing Equilateral Surfaces

I already have defined piecewise equilateral metrics determined by plane triangulation graphs in the context of the type problem. Of course there is nothing special about plane triangulation graphs. Any triangulation T of a surface may be endowed with a piecewise equilateral metric by identifying faces with unit equilateral triangles. Exactly as explained in Sect. 5.4.2, this endows the surface with a complex atlas of conformal charts indexed by the vertices and edges of the triangulation. Equilateral surfaces have become important in several different areas of mathematics. They arise for example in Grothendieck's theory of dessins d'enfants and their corresponding Belyı maps, see [20], in Angel and Schramm's theory of uniform infinite planar triangulations [4], in Gill and Rohde's study of random planar maps [37], in Bowers and Stephenson's theory of conformal tilings and especially those that arise from expansion complexes [21, 22], and in discrete conformal flattening of surfaces in \mathbb{R}^3 [16]. In this section I introduce a method of uniformizing these surfaces using the tools of Rodin-Sullivan [62] and basic surface theory.

Let T be a triangulation of the topological surface S . The notation $|T|_{\text{eq}}$ is used to denote the piecewise equilateral metric space determined by the triangulation T and \mathcal{S}_T to denote the Riemann surface determined by the atlas $\mathcal{A} = \{\varphi_x : x \in V(T) \cup E(T)\}$. Note that T need not be a simplicial triangulation for this to make sense. A face f of T first is identified as an equilateral triangle in $|T|_{\text{eq}}$ and then as a curvilinear triangle in the canonical metric of constant curvature on the surface \mathcal{S}_T . What is the shape of f in \mathcal{S}_T ? One fact about the shape of this curvilinear triangle is that the angle that two of its sides makes that emanate from the same vertex is $2\pi/d$, where d is the degree of the vertex. Another fact is that the sides are analytic arcs, and in fact any such arc is the fixed point set of an anti-conformal reflection that exchanges the two triangles incident with that arc. In the case \mathcal{S}_T is parabolic or hyperbolic, f can be lifted to the plane \mathbb{C} or the Poincaré disk \mathbb{D} and so this shape may be displayed as a curvilinear triangle in the plane. In case \mathcal{S}_T is elliptic, this shape may be stereographically projected from the 2-sphere to the plane. How does one get at this shape? The answer Ken Stephenson and I supplied in [20] is the content of this section.

For simplicity, let's restrict our attention to closed surfaces. The scheme for approximating a uniformizing map is to use the triangulation T as a pattern for a circle packing, and then refine iteratively using so-called *hex-refinement* to obtain a sequence \mathcal{P}_n of finer and finer circle packings, after an initial barycentric subdivision. Hex-refinement applied to a triangular face just adds a vertex to each existing edge and then connects the three new vertices on the three edges of the face by a 3-cycle of edges, thus subdividing the face into four smaller triangles. Thus barycentric subdivision followed by hex-refinement produces T_1 , and iteration of hex-refinement then produces the sequence T_n with \mathcal{P}_n the corresponding circle packing in the surface \mathcal{S}_n in the pattern of T_n . There is an added layer of difficulty here in that, unlike with the use of the hexagonal packing in the Discrete Riemann

Mapping Theorem, the circle packings in this setting do not occupy the same surface. The surfaces \mathcal{S}_n are determined by the triangulations T_n according to Theorem 5.3.1, and these need not be conformally equivalent to one another. Also, any face f of T with n th hex-subdivision f_n in T_n determines a sequence $\mathcal{P}_n(f)$ of circle packings, those circles in \mathcal{P}_n corresponding to the vertices of f_n .

Discrete Uniformization Theorem for Equilateral Surfaces (Bowers and Stephenson [20]) *The surfaces \mathcal{S}_n converge in moduli as $n \rightarrow \infty$ to a surface \mathcal{S} that is conformally homeomorphic to the surface \mathcal{S}_T , the Riemann surface determined by the equilateral surface $|T|_{\text{eq}}$. For any face f of T , the carriers of $\mathcal{P}_n(f)$ converge geometrically to the shape of f in \mathcal{S}_T when given its canonical constant curvature metric.*

The latter statement of the theorem may be understood to mean that when one lifts the carriers to the universal cover, the sphere \mathbb{S}^2 , the plane \mathbb{C} , or the disk \mathbb{D} , and normalizes appropriately, the carriers converge in the Hausdorff metric on compacta to the appropriate lift of f in \mathcal{S}_T .

Sketch of Proof Note that the realizations of the triangulation T in the metric surface $|T|_{\text{eq}}$ and in the Riemann surface \mathcal{S}_T are reflective, meaning that each edge e is the fixed point set of an anti-conformal reflection that exchanges the two faces contiguous to e .¹¹ Rather than the canonical constant curvature metric, I shall use the piecewise equilateral metric ρ_T on \mathcal{S}_T throughout the proof. Here is a key observation. Hex-subdivision may be performed metrically in \mathcal{S}_T by adding new vertices $v(e)$ as the mid-points of the edges $e \in E(T)$ and connecting $v(e)$ to $v(e')$ by a Euclidean straight line segment in the metric ρ_T in the face bounded by edges e, e' and e'' . This realizes the hex-refined triangulation T_1 as a reflective triangulation in \mathcal{S}_T .¹² Iterating, T_n may be realized as a reflective triangulation of \mathcal{S}_T that metrically hex-subdivides T_{n-1} .

Define homeomorphisms $h_n : \mathcal{S}_T \rightarrow \mathcal{S}_n$ so that the image of vertex v of T_n under h_n is the center of the circle that corresponds to v in the circle packing \mathcal{P}_n , extend linearly along edges and then with minimum quasiconformal distortion across faces. By the Ring Lemma, each mapping h_n is quasiconformal, and since hex-refinement does not increase degree, any bound ≥ 6 on the degrees of the vertices of T also bounds the degrees of the vertices of T_n , for all $n \geq 1$. This implies that the homeomorphisms h_n have uniformly bounded dilatations, and this implies that a subsequence of the surfaces \mathcal{S}_n converges in moduli to a Riemann surface \mathcal{S} .

My claim is that \mathcal{S} is conformally equivalent to \mathcal{S}_T . This would be confirmed were the maximum dilatations of the homeomorphisms h_n shown to limit to unity as $n \rightarrow \infty$, but unfortunately this does not occur. In fact these dilatations are bounded away from unity with large dilatations concentrated near the original vertices of T .

¹¹To be clear, the reflection is anti-conformal on the *interior* of the union of the two faces incident at e , but not at the vertices.

¹²Technically, this is after the initial barycentric subdivision, which also is performed in the metric ρ_T and yields a reflective triangulation.

To get around this, let D be a compact domain in \mathcal{S}_T disjoint from the vertex set $V(T)$. Note that the combinatorics of T_n away from the vertices of T is hexagonal, and this implies that as $n \rightarrow \infty$, the compact set D is surrounded by a number of generations of the hexagonal combinatorics that increases without bound. The Hexagonal Packing Lemma applies to confirm that the maximum dilatations of the restrictions of the homeomorphisms h_n to D converge to unity. This works for every compact domain that misses the vertex set $V(T)$, and this implies that the limit mapping $h : \mathcal{S}_T \rightarrow \mathcal{S}$ is conformal on the complement of the vertex set $V(T)$. Now the removability of isolated singularities comes into play and implies that the homeomorphism h is conformal at the vertices, and so is a conformal homeomorphism of \mathcal{S}_T onto \mathcal{S} . \square

Figure 5.7 shows an example of an approximation to a portion of an equilateral surface uniformized in the plane. In this figure each edge is the fixed point set of an anti-conformal reflection that exchanges the grey-white pair of triangles sharing that edge. This is an approximation of the conformally correct shapes of the equilateral triangles forming the equilateral surface being imaged.




Fig. 5.7 Conformal shapes of equilateral triangles in a planar equilateral surface approximated with the circle packing of the twice hex-refined barycentric subdivision of the original triangulation

5.6 Inversive Distance Circle Packings

Around 2001, Ken Stephenson and I began thinking about *inversive distance circle packings* and how they could be used to uniformize piecewise flat surfaces, those surfaces in which each face is identified with a flat Euclidean triangle, not necessarily equilateral. There is a tentative discussion of this in [20] and further discussion in [16] of the difficulties in proving convergence of discrete mappings to the uniformization mapping, though the method does seem to work well in practice; again see [16]. The first theoretical questions concern (1) the existence of circle packings with prescribed inversive distances between adjacent circles and (2) the rigidity and uniqueness of these packings.

When all inversive distances lie in the unit interval, adjacent circles overlap with specified angle $0 \leq \theta \leq \pi/2$. This is covered by the Koebe–Andre’ev–Thurston Theorems. When inversive distance is greater than unity, the circles do not overlap and the inversive distance is a Möbius-invariant measure of how separated the circles are. In this case Problems (1) and (2) seem much more difficult to approach. Problem (1) is especially difficult in that there are local assignments of inversive distances that must be avoided as there are no circle configurations that realize those distances. These are difficult to catalogue, but even if there are no local obstructions to the existence of a packing, it is not at all clear whether still there may be global obstructions. Little progress has been made on Problem (1), but the situation for Problem (2) has enjoyed some progress, initially in 2011 and more recently in the past couple of years. It is these recent successes in approaching Problem (2) that occupies this section. My contention is that a change of viewpoint can be effective in approaching inversive distance circle packings, and a hint as to how to proceed comes from the classical rigidity theory of bar-and-joint linkages. After a brief review of inversive distance, I will explore this new framework for circle packings and discuss some recent successes.

5.6.1 A Quick Introduction to Inversive Distance

There are a number of ways to define the inversive distance between two circles in the Riemann sphere. I will present several of these below, starting with the most mundane that gives a Euclidean formula for the inversive distance between two planar circles.¹³ Let C_1 and C_2 be distinct circles in the complex plane \mathbb{C} centered at the respective points p_1 and p_2 , of respective radii r_1 and r_2 , and bounding the respective *companion disks* D_1 and D_2 .

¹³This easily can be extended to the inversive distance between a circle and a line, or two lines. I will forgo this development since the next definition is completely general.

Definition (Inversive Distance in the Euclidean Metric) The *inversive distance* $\langle C_1, C_2 \rangle$ between C_1 and C_2 is

$$\langle C_1, C_2 \rangle = \frac{|p_1 - p_2|^2 - r_1^2 - r_2^2}{2r_1r_2}. \quad (5.6.1)$$

The *absolute inversive distance* between distinct circles is the absolute value of the inversive distance.


The absolute inversive distance is a Möbius invariant of the placement of two circles in the plane. This means that there is a Möbius transformation of \mathbb{C} taking one circle pair to another if and only if the absolute inversive distances of the two pairs agree. The important geometric facts that make the inversive distance useful in inversive geometry and circle packing are as follows. When $\langle C_1, C_2 \rangle > 1$, $D_1 \cap D_2 = \emptyset$ and $\langle C_1, C_2 \rangle = \cosh \delta$, where δ is the hyperbolic distance between the totally geodesic hyperbolic planes in the upper-half-space model $\mathbb{C} \times (0, \infty)$ of \mathbb{H}^3 whose ideal boundaries are C_1 and C_2 . When $\langle C_1, C_2 \rangle = 1$, D_1 and D_2 are tangent at their single point of intersection. When $1 > \langle C_1, C_2 \rangle \geq 0$, D_1 and D_2 overlap with angle $0 < \theta \leq \pi/2$ with $\langle C_1, C_2 \rangle = \cos \theta$. In particular, $\langle C_1, C_2 \rangle = 0$ precisely when $\theta = \pi/2$. When $\langle C_1, C_2 \rangle < 0$, then D_1 and D_2 overlap by an angle greater than $\pi/2$. This includes the case where one of D_1 or D_2 is contained in the other, this when $\langle C_1, C_2 \rangle \leq -1$. In fact, when $\langle C_1, C_2 \rangle < -1$ then $\langle C_1, C_2 \rangle = -\cosh \delta$ where δ has the same meaning as above, and when $\langle C_1, C_2 \rangle = -1$ then C_1 and C_2 are ‘internally’ tangent. When $-1 < \langle C_1, C_2 \rangle < 0$, then the overlap angle of D_1 and D_2 satisfies $\pi > \theta > \pi/2$ and again $\langle C_1, C_2 \rangle = \cos \theta$.

The more general definition measures the inversive distance between oriented circles. Note that an oriented circle determines a unique closed *companion* or *spanning disk* that the circle bounds. Indeed, assuming fixed orientations for \mathbb{S}^2 and $\widehat{\mathbb{C}}$ that are compatible via stereographic projection, the companion disk determined by the oriented circle C is the closed complementary disk D (of the two available) whose positively oriented boundary $\partial^+ D = C$, where of course the orientation of D is inherited from that of \mathbb{S}^2 or $\widehat{\mathbb{C}}$. This is described colloquially by saying that D lies to the left of C as one traverses C along the direction of its orientation.

Definition (General Inversive Distance) Let C_1 and C_2 be oriented circles in the extended plane $\widehat{\mathbb{C}}$ bounding their respective companion disks D_1 and D_2 , and let C be any oriented circle mutually orthogonal to C_1 and C_2 . Denote the points of intersection of C with C_1 as z_1, z_2 ordered so that the oriented sub-arc of C from z_1 to z_2 lies in the disk D_1 . Similarly denote the ordered points of intersection of C with D_2 as w_1, w_2 . The *general inversive distance* between C_1 and C_2 , denoted as $\langle C_1, C_2 \rangle$, is defined in terms of the cross ratio

$$[z_1, z_2; w_1, w_2] = \frac{(z_1 - w_1)(z_2 - w_2)}{(z_1 - z_2)(w_1 - w_2)}$$

Fig. 5.8 Inversive distances $d = \langle C_1, C_2 \rangle$. The shaded regions are the intersections $D_1 \cap D_2$, the points common to the spanning disks D_1 and D_2 for both circles C_1 and C_2



by

$$\langle C_1, C_2 \rangle = 2[z_1, z_2; w_1, w_2] - 1.$$

Subsequently, I'll drop the adjective *general* and refer to the inversive distance $\langle C_1, C_2 \rangle$ with its absolute value $|\langle C_1, C_2 \rangle|$ the *absolute inversive distance*.¹⁴

Recall that cross ratios of ordered 4-tuples of points in $\widehat{\mathbb{C}}$ are invariant under Möbius transformations and that there is a Möbius transformation taking an ordered set of four points of $\widehat{\mathbb{C}}$ to another ordered set of four if and only if the cross ratios of the sets agree. This implies that which circle C orthogonal to both C_1 and C_2 is used in the definition is irrelevant as a Möbius transformation that set-wise fixes C_1 and C_2 can be used to move any one orthogonal circle to another. Which one of the two orientations on the orthogonal circle C is used is irrelevant as the cross ratio satisfies $[z_1, z_2; w_1, w_2] = [z_2, z_1; w_2, w_1]$. This equation also shows that the inversive distance is preserved when the orientation of both circles is reversed so that it is only the relative orientation of the two circles that is important for the definition. In fact, the general inversive distance is a relative conformal measure of the placement of an oriented circle pair on the Riemann sphere. By this I mean that two oriented circle pairs are inversive equivalent if and only if their inversive distances agree. All of this should cause one to pause to develop some intuition about how companion disks may overlap with various values of inversive distances. See Fig. 5.8 for some

¹⁴The author first learned of defining inversive distance in this way from his student, Roger Vogeler. He has looked for this in the literature and, unable to find it, can only surmise that it is original with Prof. Vogeler. The definition appeared in [16] in 2003.

corrections to possible misconceptions. Finally, the inversive distance is symmetric with $\langle C_1, C_2 \rangle = \langle C_2, C_1 \rangle$ since $[z_1, z_2; w_1, w_2] = [w_1, w_2; z_1, z_2]$.

The inversive distance is real since the cross ratio of points lying on a common circle is real and, in fact, every real value is realized as the inversive distance of some oriented circle pair. Notice that if the orientation of only one member of a circle pair is reversed, the inversive distance merely changes sign. This follows from the immediate relation $[z_1, z_2; w_2, w_1] = 1 - [z_1, z_2; w_1, w_2]$. Despite its name, the inversive distance is not a metric as it fails to be non-negative and fails to satisfy the triangle inequality.¹⁵

The third definition is entirely in terms of the spherical metric.

Definition (Inversive Distance in the Spherical Metric) In the 2-sphere \mathbb{S}^2 , the inversive distance may be expressed as

$$\langle C_1, C_2 \rangle = \frac{-\cos \sphericalangle(p_1, p_2) + \cos(r_1) \cos(r_2)}{\sin(r_1) \sin(r_2)} = \frac{-p_1 \cdot p_2 + \cos(r_1) \cos(r_2)}{\sin(r_1) \sin(r_2)}. \quad (5.6.2)$$

Here, $\sphericalangle(p_1, p_2) = \cos^{-1}(p_1 \cdot p_2)$ denotes the spherical distance between the centers, p_1 and p_2 , of the respective companion disks, $p_1 \cdot p_2$ the usual Euclidean inner product between the unit vectors p_1 and p_2 , and r_1 and r_2 the respective spherical radii of the companion disks. Note that $r_i = \cos^{-1}(p_i \cdot q_i)$ for any point q_i on the circle C_i , for $i = 1, 2$.

Verifying the equivalence of this with the general definition is an exercise in the use of trigonometric identities after a standard placement of C_1 and C_2 on \mathbb{S}^2 followed by stereographic projection. This standard placement is obtained by finding the unique great circle C orthogonal to both C_1 and C_2 and then rotating the sphere so that this great circle is the equator, which then stereographically projects to the unit circle in the complex plane. The details are left to the reader.

Here are two more quick descriptions of inversive distance. For those conversant with the representation of circles in \mathbb{S}^2 by vectors in de Sitter space, the inversive distance is the Minkowski inner product between the two points of de Sitter space that represent the two oriented circles. This is, perhaps, the most elegant formulation of the product. The final way I'll describe the inversive distance is a neat little curiosity. Let $C_1 = \partial\mathbb{D} = \mathbb{S}^1$ be the unit circle oriented clockwise and C_2 a circle oriented counterclockwise that meets the open unit disk non-trivially. Then, as explained on page 193, the intersection c_2 of C_2 with the open disk is a curve of constant geodesic curvature in the Poincaré disk $\mathbb{D} \cong \mathbb{H}^2$. The inversive distance is $\langle C_1, C_2 \rangle = \text{curv}(c_2)$, the geodesic curvature of the cycle c_2 in the Poincaré metric on \mathbb{D} . This includes all three cases for the cycle c_2 —a hyperbolic circle in \mathbb{D} , a

¹⁵Some authors, perhaps more aptly, call the inversive *distance* the inversive *product* of C_1 and C_2 .

horocycle that meets $\partial\mathbb{D}$ at a single point, or a hypercycle that meets $\partial\mathbb{D}$ at two points.¹⁶

5.6.2 Some Advances on the Rigidity Question

In [20], *inversive distance circle packings* were introduced. Rather than preassigned overlap angles labeling edges of a triangulation of a surface as in the Koebe–Andre’ev–Thurston Theorems, preassigned inversive distances label the edges. As stated already, questions of interest are of the existence and uniqueness of circle configurations in geometric structures on surfaces that realize the inversive distance data. Though the existence question is wide open, in 2011–2012 there were three advances on the uniqueness question for inversive distance packings. First, Guo [41] proved that inversive distance packings of closed surfaces of positive genus, ones supporting flat or hyperbolic metrics, are locally rigid whenever the inversive distances are non-negative. Shortly after that, Luo [52] improved this to global rigidity, or uniqueness of the packings in the cases considered by Guo. Then in a surprising result of the year following, Ma and Schlenker [53] produced a counterexample to global uniqueness for packings of the 2-sphere. They gave examples of pairs of circle packings of \mathbb{S}^2 in the pattern of the octahedral triangulation with six circles that satisfy the same inversive distance data, but that are not Möbius equivalent.

The ingredients of Ma and Schlenker’s example are Schönhardt’s twisted octahedron, which is an infinitesimally flexible polyhedron in Euclidean space \mathbb{E}^3 , embeddings in de Sitter space \mathbb{S}_1^3 , and special properties of the Pogorelov map between different geometries. In 2017, John Bowers and I [15] constructed a large family of Ma–Schlenker–like examples using only inversive geometry, producing many counterexamples to the uniqueness of inversive distance circle packings in the 2-sphere.

The Schönhardt octahedron is an example of a bar-and-joint linkage important in the rigidity theory of Euclidean frameworks, and its use in the Ma–Schlenker example hinted at a way forward in understanding the rigidity theory of inversive distance circle packings in the 2-sphere. This led to a fruitful change in viewpoint and a reformatting of the question of uniqueness of inversive distance circle packings to the question of the rigidity—local, global, and infinitesimal—of more general *circle frameworks*. These are analogues in Möbius geometry of the Euclidean frameworks in Euclidean geometry with point configurations in \mathbb{E}^3 replaced by circle configurations in \mathbb{S}^2 and the Euclidean metric replaced by the non-metric inversive distance. The analogy is not exact, but the theory of linkages in \mathbb{E}^3 has been found to be a good guide for understanding some of the rigidity theory of circle

¹⁶My student, Opal Graham, noticed, then proved this when I was lecturing on the curves of constant geodesic curvature in the hyperbolic plane.

frameworks. Part of why this works so well is because the space of circles in the 2-sphere is a three-dimensional incidence geometry that has much in common with the space of points in Euclidean 3-space. The lines of this geometry are coaxial circle families and the planes are what Carathéodory in [28] called bundles of circles. This allows one to define what is meant by a convex collection of circles, planar collections of circles, circle polyhedra, bounded circle configurations, etc. Space constraints in this chapter interfere with even a cursory account of these issues, so I am content with listing a couple of recent successes of the theory without all the definitions needed for a precise understanding, and then taking some time to set up the language of this change of viewpoint.

The two theorems following are the result, both the statements and the proofs, of an engagement between circle packing theory and the classical rigidity theory of Euclidean frameworks in \mathbb{E}^3 .

Theorem 5.6.1 (Bowers et al. [23]) *Let \mathcal{C} and \mathcal{C}' be two non-unitary, inversive distance circle packings with ortho-circles for the same oriented edge-labeled triangulation of the 2-sphere \mathbb{S}^2 . If \mathcal{C} and \mathcal{C}' are convex and proper, then there is a Möbius transformation $T : \mathbb{S}^2 \rightarrow \mathbb{S}^2$ such that $T(\mathcal{C}) = \mathcal{C}'$.*

The *edge-label* refers to prescribed inversive distances labeling each edge. *Non-unitary* means that the inversive distance between any pair of adjacent circles is not unity; in fact, these inversive distances are in the set $(-1, 1) \cup (1, \infty)$. Having *ortho-circles* means that each triple of mutually adjacent circles have an orthogonal circle. This generalizes to a global rigidity theorem about *circle polyhedra*, circle configurations in the pattern of three-dimensional polyhedra whose faces correspond to circle configurations that are planar in the incidence geometry of circle space; see [23] for details.

Theorem 5.6.2 (Bowers et al. [23]) *Any two convex and proper non-unitary circle polyhedra with Möbius-congruent faces that are based on the same oriented abstract spherical polyhedron and are consistently oriented are Möbius-congruent.*

Theorem 5.6.1 coupled with the Ma–Schlenker example of [53] and the examples of [15] show that the uniqueness of inversive distance circle packings, and more generally, of circle polyhedra is exactly analogous to that of Euclidean polyhedra—convex and bounded polyhedra in \mathbb{E}^3 are prescribed uniquely by their edge lengths and face angles whereas non-convex or unbounded polyhedra are not. The proof of this for convex and bounded Euclidean polyhedra is Cauchy’s celebrated rigidity theorem [29], which is reviewed in Sect. 5.7.5. The Proof of Theorem 5.6.2 follows Cauchy’s original argument, which splits the proof into two components—a combinatorial lemma and a geometric lemma. Cauchy’s combinatorial lemma deals with a certain labeling of the edges of any graph on a sphere, and applies to the present setting. The geometric lemma, known as *Cauchy’s Arm Lemma*, requires that a polygon with certain properties be defined for each vertex of the polyhedron, and fails to apply here. The main work of the proof is in describing and analyzing a family of hyperbolic polygons called *green-black polygons* that are defined for each vertex of a circle polyhedron in a Möbius-invariant manner. An analogue of

Cauchy's Arm Lemma for convex green-black polygons is developed and used to prove these theorems.

5.6.3 Circle Frameworks and Möbius Rigidity

I'll close out this section with a description of the change in viewpoint from circle packings to circle frameworks. This can be done using only absolute inversive distance, but I find it advantageous to remain as general as possible in setting up the viewpoint. The goal is to generalize the language of circle packings and patterns of triangulations and quadrangulations of the 2-sphere to that of circle realizations of oriented circle frameworks. Let G be a graph, by which I mean a set of vertices $V = V(G)$ and simple edges $E = E(G)$. Both loops and multiple edges are disallowed. An oriented edge incident to the initial vertex u and terminal vertex v is denoted as uv , and $-uv$ means the oppositely oriented edge vu . I will use the same notation, uv , to denote an un-oriented edge, context making the meaning clear. A *circle framework with adjacency graph* G , or *c-framework* for short, is a collection $\mathcal{C} = \{C_u : u \in V(G)\}$ of oriented circles in \mathbb{S}^2 indexed by the vertex set of G . This is denoted by $G(\mathcal{C})$. Two *c-frameworks* $G(\mathcal{C})$ and $G(\mathcal{C}')$ are *equivalent* if $\langle C_u, C_v \rangle = \langle C'_u, C'_v \rangle$ whenever uv is an edge of G . Let H be a subgroup of the inversive group $\text{Inv}(\mathbb{S}^2)$ of the 2-sphere. Two collections \mathcal{C} and \mathcal{C}' of oriented circles indexed by the same set are *H-equivalent* or *H-congruent* provided there is a mapping $T \in H$ such that $T(\mathcal{C}) = \mathcal{C}'$, respecting the common indexing and the orientations of the circles. When H is not so important they are *inversive-equivalent* or *inversive-congruent*, and when T can be chosen to be a Möbius transformation, they are *Möbius-equivalent* or *Möbius-congruent*. The global rigidity theory of *c-frameworks* concerns conditions on G or $G(\mathcal{C})$ that ensure that the equivalence of the *c-frameworks* $G(\mathcal{C})$ and $G(\mathcal{C}')$ guarantees their *H-equivalence*. Often one restricts attention to *c-frameworks* in a restricted collection \mathcal{F} of *c-frameworks*. In Theorem 5.6.2, \mathcal{F} is the collection of non-unitary, convex and proper *c-polyhedra* and the interest is in Möbius equivalence.

Definition (Labeled Graph and Circle Realization) An *edge-label* is a real-valued function $\beta : E(G) \rightarrow \mathbb{R}$ defined on the edge set of G , and G together with an edge-label β is denoted as G_β and called an *edge-labeled graph*. The *c-framework* $G(\mathcal{C})$ is a *circle realization* of the edge-labeled graph G_β provided $\langle C_u, C_v \rangle = \beta(uv)$ for every edge uv of G , which henceforth is denoted as $G_\beta(\mathcal{C})$. See Fig. 5.9.

Circle packings are circle realizations of edge-labeled graphs that arise as the 1-skeletons of oriented triangulations of the 2-sphere that also satisfy certain properties that ensure that the realizations of the triangular boundaries of faces respect orientation. The general definition allows for branch vertices and configurations of circles in which the open geodesic triangles cut out by connecting centers of adjacent circles overlap. There are subtleties in which I have no interest, so I am

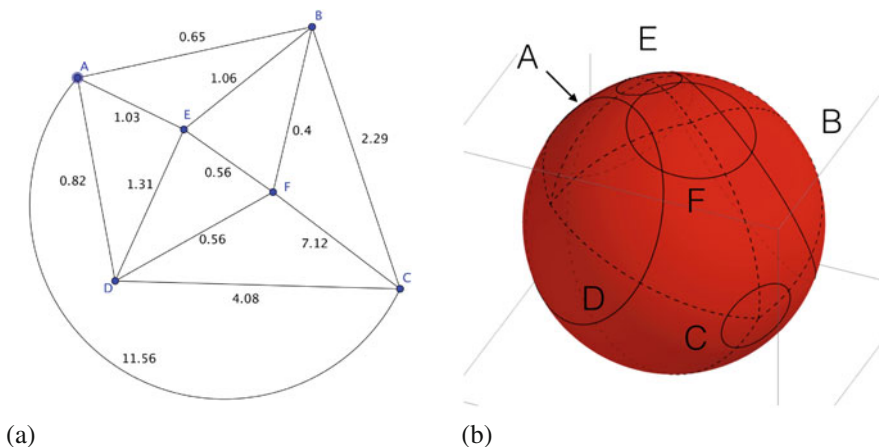


Fig. 5.9 An edge-labeled octahedral graph and its circle realization. The labels are proposed inversive distances between the circles corresponding to the vertices. (a) An edge-labeled octahedral graph \mathcal{O}_β . Labels < 1 imply overlapping circles, > 1 separated ones. (b) A corresponding c -framework realizing \mathcal{O}_β . Circle A is hidden on the back side of the sphere

going to adapt a restricted definition that corresponds to the circle packings that arise from spherical polyhedral metrics on triangulated surfaces. These are circle realizations of the edge-labeled 1-skeleton $G_\beta = K_\beta^{(1)}$ of an oriented triangulation K of \mathbb{S}^2 that produce oriented geodesic triangulations¹⁷ of the 2-sphere when adjacent circle centers are connected by geodesic arcs. The assumption here is that the centers of no two adjacent circles are antipodal, so that there is a unique geodesic arc connecting them, and that the centers of three circles corresponding to the vertices of a face of K do not lie on a great circle. Now this causes no particular problems when all adjacent circles overlap nontrivially, the traditional playing field of circle packing, but does cause some real concern when adjacent circles may have inversive distance greater than unity. For example, a circle realization may produce a geodesic triangulation of the sphere by connecting adjacent centers while its Möbius image may not. This is traced directly to the fact that neither circle centers nor radii, nor geodesic arcs, are Möbius invariants in the inversive geometry of the sphere. This behavior does not occur for inversive distance circle packings of the Euclidean or hyperbolic planes (and surfaces), precisely because circle centers and geodesics are invariant under automorphisms and radii are invariant up to scale in Euclidean geometry and invariant in hyperbolic geometry. My belief is that using centers and radii of circles in inversive geometry should be avoided except where these can be used to simplify computations (as in the use of the spherical definition of inversive distance). The shift then is from inversive distance circle packings to inversive

¹⁷By this I mean that the orientation of the geodesic triangulation determined by the packing is consistent with the orientation on K .

distance circle realizations. One is less concerned with possible underlying geodesic triangulations and more concerned with Möbius-invariant quantities. For example, rather than working with a geodesic face formed by connecting the centers of three mutually adjacent circles, one is more interested in the existence of an ortho-circle, a circle mutually orthogonal to the three, which is a Möbius invariant. Though the initial motivation was circle packing as reflected in Theorem 5.6.1, the real interest has evolved to circle realizations as reflected in the more general version represented by Theorem 5.6.2.

It turns out that Theorem 5.6.2 has implications for the rigidity of generalized hyperbolic polyhedra in \mathbb{H}^3 . Thurston was the first to exploit this connection between circle configurations on \mathbb{S}^2 and hyperbolic polyhedra in \mathbb{H}^3 in really significant ways, and his observations inspired several avenues of clarification and generalization. It is to this that I turn in the penultimate section of this chapter.

5.7 Polyhedra—From Steiner (1832) to Rivin (1996), and Beyond

In this section I survey the rich mathematical vein that has been mined in the geometric theory of polyhedra, particularly of three-dimensional hyperbolic polyhedra, that has its origins in Thurston's insights on using his circle packing theorem to characterize certain hyperbolic polyhedra. The initial observation of Thurston was that the study of polyhedra in hyperbolic three-space can be transferred to the study of overlapping circle packings in the two-sphere by realizing the Riemann sphere as the boundary of the Beltrami–Klein model of \mathbb{H}^3 sitting as the unit ball \mathbb{B}^3 in the real projective three-space. Theorems in one of these venues correspond to theorems in the other. Later Thurston's students, Oded Schramm and Igor Rivin, made great strides in the theory of both three-dimensional Euclidean and hyperbolic polyhedra, not so much using the techniques of circle packing but instead using very intricate and clever geometric arguments, often times in this classical setting of $\mathbb{H}^3 \cong \mathbb{B}^3 \subset \mathbb{E}^3 \subset \mathbb{RP}^3$. There is here a beautiful interplay among the classical geometries illustrating Arthur Cayley's aphorism that "All geometry is projective geometry." Here one sees the Beltrami–Klein model of hyperbolic three-space as a sub-geometry of the real projective three-space, with its orientation-preserving isometry group naturally identified with the Lorentz group of Minkowski space-time, which itself restricts to the two-sphere boundary of hyperbolic space as the group of circle-preserving transformations of the two-sphere, the group of Möbius transformations. This one geometry, the real projective geometry of dimension three, presents a playing field for studying three-dimensional polyhedra—classical Euclidean polyhedra, hyperbolic polyhedra of various types and generalizations, projective polyhedra, and circle polyhedra of Möbius geometry.

I will begin with an application of Thurston's circle packing theorem on using polyhedra to cage a sphere, and move then to Schramm's generalization. From there

I will discuss the characterization of certain hyperbolic polyhedra—compact by Hodgson and Rivin, ideal by Rivin, and hyper-ideal by Bao and Bonahon—and will finish with very recent work by Chen and Schlenker that characterizes those convex projective polyhedra all of whose vertices lie on the ideal boundary of hyperbolic space. I include a bonus final section on Cauchy’s 1813 Rigidity Theorem for the reader who is approaching this subject as a novice. This is the fundamental theorem of rigidity theory, and the techniques and tools Cauchy developed have been used time and again in proofs of rigidity in the past 200 years. Both Schramm and Rivin make use of Cauchy’s toolbox in their theorems on convex hyperbolic and Euclidean polyhedra, as do Bao and Bonahon as well as Bowers, Pratt and the author. Before these recent developments, previous generations of mathematicians who delved into the study of polyhedra made use of Cauchy’s toolbox—Dehn in his proof of infinitesimal rigidity, Aleksandrov in his rigidity results, Gluck in his examination of generic rigidity, and Connelly in various of his contributions.

5.7.1 *Caging Eggs—Thurston and Schramm*

In 1832, Jakob Steiner [69] asked

In which cases does a convex polyhedron have a combinatorial equivalent which is inscribed in, or circumscribed about, a sphere?

When a convex polyhedron P is *inscribed* in the sphere S so that its vertices lie on S , then its polar dual *circumscribes* the sphere S , so that each face of the dual P^* meets S in a single point. It wasn’t until 1928 that Ernst Steinitz found families of non-inscribable polyhedral types with the example of a cube truncated at one vertex being the simplest. Marcel Berger [9, p. 532] takes this long duration of time between Steiner and Steinitz as evidence that the theory of polyhedra in the years intervening had fallen into disrespect among mathematicians, being a subject of the old-fashioned mathematics of synthetic geometry.¹⁸ One would be hard pressed to say that the study of polyhedra in the time between Steinitz and Thurston was anything but a curiosity to many a mathematician schooled in the rarified heights of abstraction that had captured the mathematical mind of the time. The sort of “pedestrian geometry” offered by the study of polyhedra captured the imagination of a select few. There has been a healthy development of the rigidity theory of polyhedra, notably by Aleksandrov in the 1950s, and Gluck and Connelly in the 1970s. Aleksandrov’s work was largely ignored in the West until the 1980s. Coxeter had done truly foundational work in the combinatorial structure of polyhedra in the 1940s and 1950s, and Victor Klee and Branko Grünbaum began their foundational studies a bit later. Coxeter’s work in geometry was routinely dismissed by much of mainstream mathematics as old-fashioned nineteenth century

¹⁸Berger [9] uses the word *disdain* to describe the prevailing opinion of the study of polyhedra.

mathematics, uninteresting and pedestrian. Both Aleksandrov and Coxeter were “rehabilitated” by the larger community of geometers and topologists when their work of the forties and fifties—Aleksandrov’s on metric geometry and Coxeter’s on reflection groups—became important to the development of geometric group theory after Gromov’s publication of his hyperbolic groups essay [39] in 1987. With apologies to Aleksandrov, Coxeter, Klee, and Grünbaum, it has taken the attention of Thurston and his students Schramm and especially Rivin to resurrect more intense interest among topologists in this venerable old subject of classical geometry.¹⁹

Steinitz’s basic tool for attacking the Steiner question is the following observation. Suppose the polyhedron P circumscribes the sphere S . Let $e = uv$ be an edge of P with adjacent faces f and g . Since P circumscribes S , the face f is tangent to S at a point p and g is tangent at a point q . Then the angle $\angle upv = \angle uqv$ in measure and we let $\Theta(e)$ denote this common value. It is immediate that summing these edge labels for the edges of any face yields an angle sum of 2π . The reader might want to use this observation to see why a dodecahedron truncated at every vertex admits no inscribed sphere as there is no edge labeling Θ for this polyhedron that satisfies this property.

According to Steinitz then, the condition that an edge label $\Theta : E(P) \rightarrow (0, \pi)$ exists for the polyhedron P whose sum for the edges of each face is 2π is a necessary condition that P have a combinatorially equivalent realization that circumscribes a sphere, but it is not sufficient. It was not until Rivin’s study of hyperbolic polyhedra in the late 1980s and early—1990s that a characterization of polyhedra of *circumscribable type*, ones combinatorially equivalent to polyhedra that may circumscribe a sphere, was found. The definitive result is due to Rivin and reported in Hodgson, Rivin, and Smith [47], and follows from his characterization of ideal convex hyperbolic polyhedra that is presented in a later section.

Circumscribable Type Characterization (Rivin) *A polyhedron P is of circumscribable type if and only if there exists a label $\Theta : E(P) \rightarrow (0, \pi)$ such that the sum of the labels $\Theta(e)$ as e ranges over any circuit bounding a face is 2π , while the sum as e ranges over any simple circuit not bounding a face is strictly greater than 2π .*

A polyhedron is of *inscribable type* if it is combinatorially equivalent to one that may be inscribed in a sphere.

Inscribable Type Characterization (Rivin) *A polyhedron P is of inscribable type if and only if its dual P^* is of circumscribable type.*

The proofs will be discussed later, but first I want to generalize this discussion a bit. Inscription and circumscription are the respective cases, $m = 0$ and $m = d - 1$, of the question of whether a d -dimensional convex polytope has a realization in

¹⁹Grünbaum [40] addresses the disinterest of the mathematical community in the combinatorial theory of polytopes in the preface to his book.

\mathbb{E}^d each of whose m -dimensional faces meets a fixed $(d - 1)$ -dimensional sphere in a single point. One says that the polytope is (m, d) -scribable in this case. Egon Shulte [68] proved in the mid-1980s that when $0 \leq m < d$ and $d > 2$, then there are combinatorial types of d -dimensional polytopes that are not (m, d) -scribable, except for the single exceptional case when $(m, d) = (1, 3)$. The exceptional case then is when a convex polyhedron in \mathbb{E}^3 *midscribes* a sphere S , so that each edge of P is tangent to S , meeting S in exactly one point.

In light of Shulte’s result it perhaps is surprising that in his exceptional case, every convex polyhedron in \mathbb{E}^3 has a combinatorially equivalent realization that is midscribable about, say, the unit sphere \mathbb{S}^2 . Thurston in Chapter 13 of GTTM states that this is a consequence of Andre’ev’s theorems in [2, 3]. The proof I give merely applies the Koebe–Andre’ev–Thurston Theorem to an appropriately edge-labeled graph.

Midscribability of Convex Polyhedra (Thurston [73]) *Every convex polyhedron in \mathbb{E}^3 has a combinatorially equivalent realization that is midscribable about the unit sphere \mathbb{S}^2 . Considering $\mathbb{E}^3 \subset \mathbb{RP}^3$, any such realization is unique up to projective transformations of \mathbb{RP}^3 that set-wise fix the unit sphere \mathbb{S}^2 .*²⁰

Proof Let P be a convex polyhedron in \mathbb{E}^3 and let K be the simplicial 2-complex obtained by adding a vertex to each open face of P and starring to the vertices. Precisely, the vertices of K are those of P along with a new vertex v_f for each face f of P . The edges are the edges of P along with edges of the form vv_f , where v is a vertex of f . The faces are the 2-simplices of the form uvv_f where uv is an edge of f . Write the edge-set of K as $E(K) = E(P) \cup E'$, where E' are the new edges of the form vv_f . Define an angle map $\Phi : E(K) \rightarrow [0, \pi/2]$ by $\Phi(e) = 0$ when $e \in E(P)$ and $\Phi(e) = \pi/2$ when $e \in E'$. An application of the Koebe–Andre’ev–Thurston Theorem I produces a circle packing $K(\mathbb{C})$ on the 2-sphere \mathbb{S}^2 and a geodesic triangulation in the pattern of K with overlap angles of adjacent circles given by Φ . For each face f of P , let H_f be the half-space in \mathbb{E}^3 that meets all the circles of $K(\mathbb{C})$ and whose bounding plane ∂H_f contains C_{v_f} . My claim is that the convex polyhedron $Q = \bigcap_{f \in F(P)} H_f$ midscribes \mathbb{S}^2 and is combinatorially equivalent to P .

For any vertex v of P , let v^* be the apex of the cone in \mathbb{E}^3 that is tangent to \mathbb{S}^2 along the circle C_v , and when $e = uv$ is an edge of P , let $e^* = u^*v^*$ be the segment with endpoints u^* and v^* . Let f be a face of the polyhedron P with vertices v_1, \dots, v_n written in cyclic order. Since the circle C_{v_f} is orthogonal to the circles C_{v_i} , the apexes v_i^* all lie on the bounding plane ∂H_f , for $i = 1, \dots, n$. Let f^* denote the convex hull of the points v_1^*, \dots, v_n^* in ∂H_f . A moment’s thought should convince the reader that the convex polyhedron Q may be described as the convex hull of the set $V(Q) = \{v^* : v \in V(P)\}$. It follows that the vertex set of Q is $V(Q)$, edge set is $E(Q) = \{e^* : e \in E(P)\}$, and face set is $F(Q) = \{f^* : f \in F(P)\}$. This verifies that P and Q are combinatorially equivalent. Moreover, the edge $e^* = u^*v^*$

²⁰The projective transformations that fix \mathbb{S}^2 act as Möbius transformations on \mathbb{S}^2 .




Fig. 5.10 A midscribed polyhedron. Each face meets the sphere \mathbb{S}^2 in a red circle and each vertex is the cone point of a black circle. Each edge e^* meets \mathbb{S}^2 in exactly one point, at the intersection of the two red circles determined by the faces incident to e^* , or at the intersection of the two black circles determined by the endpoints of e^*

is tangent to the sphere \mathbb{S}^2 at the point of intersection of the circles C_u and C_v , which are tangent since $\Phi(e) = 0$. Hence Q midscribes the sphere \mathbb{S}^2 . See Fig. 5.10.

Uniqueness up to projective transformations that fix the unit sphere follows from the Möbius uniqueness of the circle packing $K(\mathcal{C})$ with edge angle data Φ and the fact that the Möbius group extends its action on \mathbb{S}^2 to a projective action of \mathbb{RP}^3 set-wise fixing \mathbb{S}^2 . □

Schulte introduced in [68] the question of whether the sphere can be replaced by other convex bodies. Schramm [64] proved that when the convex polyhedron P is simplicial, then for any smooth convex body S , a combinatorially equivalent polyhedron Q exists that midscribes S . Of course this means that each edge of Q is tangent to the boundary ∂S . Shortly thereafter, Schramm improved his result by removing the requirement that P be simplicial. A convex body S is *strictly convex* if its boundary contains no non-degenerate line segment, and is *smooth* if each point of the boundary has a unique supporting plane. This latter condition is equivalent to the boundary being C^1 -smooth. Schramm’s definitive result on midscription is the main theorem of his *Inventiones* article [66] whimsically entitled *How to cage an egg*.

Convex Body Midscription (Schramm [66]) *Let P be a convex polyhedron and S a smooth strictly convex body in \mathbb{E}^3 . Then there exists a convex polyhedron Q combinatorially equivalent to P that midscribes S .*

Discussion of Proof The proof is rather involved and so I am content to give the briefest of indication of its method. Schramm defines the *configuration space* $\mathcal{Z} = (\mathbb{E}^3)^{V(P)} \times G(2, 3)^{F(P)}$, where $V(P)$ and $F(P)$ are the respective sets of vertices and faces of P , and $G(2, 3)$ is the manifold of oriented affine planes in \mathbb{E}^3 . In this way P is identified with a single point of \mathcal{Z} , and the combinatorial type of P defines a submanifold \mathcal{Z}_P of \mathcal{Z} corresponding to various convex polyhedra in \mathbb{E}^3 that are combinatorially equivalent to P . Schramm then shows that there is a C^2 convex body S_0 with positively curved boundary that P midscribes. Let S_t , $0 \leq t \leq 1$, be a C^2 -path of convex bodies with positively curved boundaries with $S_1 = S$. The idea now is to flow S_0 to S_1 along this path and drag combinatorial realizations of P along as midscribing polyhedra. The proof relies on a fine analysis of the configuration space \mathcal{Z} and its submanifold \mathcal{Z}_P , and the method is to show that when S_t is midscribed by a realization of P , then so is $S_{t'}$ for all t' in an open interval about t . Then a delicate argument shows also that the set of parameter values for which S_t is midscribable by a realization of P is a closed set. Being open and closed, and nonempty since P midscribes S_0 , this set of parameter values must be the whole of the unit interval, hence $S = S_1$ is midscribed by a realization of P . \square

The remaining discussion on hyperbolic polyhedra has little to do, at least directly, with the Koebe–Andre’ev–Thurston Theorem. The arguments tend to be clever and technical, but ultimately involve the elementary geometry of hyperbolic space, often times realized as the unit ball in projective 3-space where the machinery of the Minkowski inner product and of de Sitter space is available. I include the discussion in order to complete for the reader the current state of affairs in the study of convex hyperbolic polyhedra, a study which I view as having been revitalized by Thurston’s articulation of KAT I and pushed forward into the broader mathematical consciousness by the seminal work of Thurston’s students, Oded Schramm and especially Igor Rivin.

5.7.2 Compact and Convex Hyperbolic Polyhedra—Hodgson and Rivin

In his doctoral thesis of 1986, Igor Rivin studied convex hyperbolic polyhedra. Therein he gave a characterization of compact, convex hyperbolic polyhedra that generalizes the Andre’ev results of [2],²¹ and in articles in the early 1990s, extended his characterization to ideal polyhedra, generalizing Andre’ev’s results in [3]. He

²¹See Roeder, Hubbard, and Dunbar’s paper [63] for a readable proof of Andre’ev’s classification of compact hyperbolic polyhedra with non-obtuse exterior dihedral angles.

used this latter generalization to answer definitively Steiner’s question of 1832 asking for a characterization of those polyhedra that circumscribe a sphere. This of course is the content of the Circumscribable Type Characterization Theorem of the preceding section. In this section, I present an overview of Rivin’s characterization of compact and convex hyperbolic polyhedra in terms of a generalized Gauss map. The overview embellishes Hodgson’s outline presented in [45] (and repeated in [46]). In the section following, I outline Rivin’s characterization of ideal polyhedra and make his observation that the Circumscribable Type Characterization Theorem is an immediate corollary of his characterization of ideal polyhedra.

To lay the groundwork, let’s review the *Gauss map* \mathbb{G} of a compact and convex Euclidean polyhedron P to the unit sphere \mathbb{S}^2 . This is a set-valued map from the 2-complex forming the boundary of P that assigns to the point p of ∂P the set of outward pointing unit normals to support planes to P at p . Thus when p is a point of an open face f , $\mathbb{G}(p) = \mathbb{G}(f)$ is a single point determined by the outward unit normal to f . When p is in the open edge e incident to faces f and g , $\mathbb{G}(p) = \mathbb{G}(e)$ is the great circular arc connecting $\mathbb{G}(f)$ to $\mathbb{G}(g)$ of length equal to the exterior dihedral angle between f and g . Finally, for a vertex p of P , $\mathbb{G}(p)$ is the convex spherical polygon bounded by the arcs $\mathbb{G}(e)$ for edges e incident with p . When edges e and e' of the face f are incident at p , the interior angle of the polygon $\mathbb{G}(p)$ at the vertex $\mathbb{G}(f)$ is $\pi - \alpha$, where α is the interior angle of the face f at p . In this way the Gauss map realizes the Poincaré dual P^* of P as a geodesic cellular decomposition of the 2-sphere \mathbb{S}^2 . Notice that the Gauss map does not encode all the information needed to reconstruct the polyhedron P . It encodes the interior angles of all the faces and the dihedral angles of all adjacent faces, but there is no encoding of side-lengths of the edges of P . For example, all rectangular boxes have the same image under the Gauss map, namely, a regular right-angled octahedral decomposition of the sphere \mathbb{S}^2 .

Another way to describe the convex spherical polygon $\mathbb{G}(p)$ for a vertex p of P is as the polar dual $\mathbb{L}^*(p)$ of the infinitesimal link $\mathbb{L}(p)$ of p in P .²² Note that $\mathbb{L}(p)$ is a convex spherical polygon with internal angles equal to the dihedral angles of the faces of P incident with p , and edge-lengths equal to the internal angles at the vertex p in the faces of P incident with p . Recall that an oriented great circle in \mathbb{S}^2 and its spherical center are polar duals of one another. The polar dual $\mathbb{L}^*(p)$ is obtained by replacing the edges of $\mathbb{L}(p)$ by the polar dual centers of their supporting great circles, and the vertices by appropriate arcs of the polar dual great circles. A nice exercise in spherical geometry verifies that $\mathbb{L}^*(p)$ is isometric to $\mathbb{G}(p)$. This gives an alternate construction of the Poincaré dual P^* as a geodesic, cellular decomposition of the 2-sphere—just isometrically glue the polar duals $\mathbb{L}^*(p)$ together as p ranges over the vertices of P along corresponding edges, $\mathbb{L}^*(p)$

²²For a Euclidean polyhedron, $\mathbb{L}(p)$ is the intersection of P with a small sphere centered at p , one whose radius is smaller than the lengths of edges incident with p , rescaled to unit radius, and is oriented so that its interior is “to the left” as one traverses the polygon in its positive direction.

glued to $L^*(q)$ whenever pq is an edge of P .²³ Obviously this gluing produces a 2-sphere, not only homeomorphic, but also isometric to the standard 2-sphere S^2 , and reproduces the cellular decomposition determined by the Gauss map.

It is this latter construction of the Poincaré dual P^* as a cellular decomposition of the 2-sphere that readily generalizes to convex and compact hyperbolic polyhedra. Indeed, let P now be a convex and compact hyperbolic polyhedron in \mathbb{H}^3 and for each vertex p , let $L^*(p)$ be the polar dual of the infinitesimal link $L(p)$ of p in P .²⁴ The link $L(p)$, as in the Euclidean case, is an oriented convex spherical polygon in S^2 with internal angles equal to the dihedral angles of the faces of P incident with p , and edge-lengths equal to the internal angles at the vertex p in the faces of P incident with p . The polar dual $L^*(p)$ then encodes the exterior dihedral angles at the edges of P incident with p as the lengths of its edges, and the interior angles α of the faces incident with p as its interior angles in the form $\pi - \alpha$. This construction acts as a local Gauss map in a small neighborhood of the vertex p . Now exactly as before, isometrically glue the polar duals $L^*(p)$ together as p ranges over the vertices of P along corresponding edges. The result is again a 2-sphere topologically, which is called the *Gaussian image* of P and denoted as $G(P)$, with a spherical metric of constant unit curvature, except at the vertices. The vertices have cone type singularities with concentrated negative curvature. Indeed, at the vertex corresponding to the face $f = p_1 \cdots p_n$ of P , the angle sum is $\theta(f) = n\pi - \sum_{i=1}^n \alpha_i$, where α_i is the internal angle of f at the vertex p_i . In the hyperbolic plane, the compact and convex polygon f always has interior angle sum strictly less than $(n - 2)\pi$ so that $\theta(f) > 2\pi$.

This brings us to Rivin’s characterization of compact and convex hyperbolic polyhedra.

Compact Convex Hyperbolic Polyhedra Characterization (Rivin) *A metric space (M, g) homeomorphic to S^2 can arise as the Gaussian image $G(P)$ of a compact and convex polyhedron P in \mathbb{H}^3 if and only if these three conditions adhere.*

- (i) *The metric g has constant curvature $+1$ except at a finite number of cone points.*
- (ii) *The cone angle at each cone point is greater than 2π .*
- (iii) *The lengths of the nontrivial closed geodesics of (M, g) are all strictly greater than 2π .*

Moreover, the metric g determines P uniquely up to hyperbolic congruence.

Recall that the Gauss map does not determine Euclidean polyhedra up to congruence since it contains no information about side lengths. In contrast, a

²³The edge pq determines respective vertices u and v of $L(p)$ and $L(q)$ whose respective polar edges u^* and v^* have the same lengths, namely the exterior dihedral angle of P at edge pq .

²⁴This is the link in the tangent space of \mathbb{H}^3 of the pre-image of the intersection of P with a small neighborhood of p under the exponential map.

hyperbolic polyhedron is determined up to a global hyperbolic isometry by its Gaussian image. The proof of this uniqueness uses Cauchy's toolbox that is reviewed in Addendum Sect. 5.7.5, wherein I recall the tools Cauchy used to prove his celebrated rigidity theorem of 1813. The necessity of items (i) and (ii) follows from the previous discussion and that of (iii) uses the fact that the total geodesic curvature of a non-trivial closed hyperbolic space curve is greater than 2π , a hyperbolic version of Fenchel's Theorem on Euclidean space curves. The proof of sufficiency is based on Aleksandrov's Invariance of Domain method used in his study of Euclidean polyhedra in [1].

Rivin also uses Cauchy's toolbox to prove this rather interesting theorem that illustrates again the enhanced rigidity of hyperbolic polyhedra vis-à-vis Euclidean ones.

Face Angle Rigidity (Rivin) *The face angles of a compact and convex polyhedron in \mathbb{H}^3 determine it up to congruence.*

The characterization of compact and convex hyperbolic polyhedra in terms of the Gaussian image surveyed here suffers from the same defect as Aleksandrov's characterization of compact and convex Euclidean polyhedra. Both characterizations posit a singular positively curved metric on a 2-sphere, but neither provides a way to decode from this metric space (M, g) the combinatorial type of the polyhedron P encoded in (M, g) . The proof is not constructive, but depends on a topological analysis within the space of admissible metrics on the 2-sphere satisfying the three conditions of the characterization and yields, finally, the abstract fact of existence of an appropriate polyhedron, without describing its combinatorial type.

5.7.3 Convex Ideal Hyperbolic Polyhedra—Rivin

Rivin turns his attention to convex ideal polyhedra in \mathbb{H}^3 in [61] where he gives a full characterization in terms of exterior dihedral angles. The characterization begins with an analysis of the exterior dihedral angles of such a polyhedron reported in [47] with details in [59] that goes as follows. Label each edge e^* of the polyhedron P^* dual to the ideal convex polyhedron P by the exterior dihedral angle $\theta(e^*)$ of the corresponding edge e of P . Rivin's argument that these labels satisfy the following conditions is reproduced in the next two theorems.

- (i) $0 < \theta(e^*) < \pi$ for all edges e of P .
- (ii) If the edges e_1^*, \dots, e_n^* are the edges bounding a face of P^* , then $\theta(e_1^*) + \dots + \theta(e_n^*) = 2\pi$.
- (iii) If e_1^*, \dots, e_n^* forms a simple nontrivial circuit that does not bound a face of P^* , then $\theta(e_1^*) + \dots + \theta(e_n^*) > 2\pi$.

Compare these conditions with the hypotheses of the Circumscribable Type Characterization on page 223. Now Condition (i) is a requirement of convexity and Condition (ii) is seen easily in the upper-half-space model by placing one of the

ideal vertices v of P at infinity and observing that the link of v is a convex Euclidean polygon. Indeed, the faces incident with v lie on vertical Euclidean planes whose intersections with the xy -plane cut out a convex Euclidean polygon $L(v)$, and quite easily the sum $\theta(e_1^*) + \dots + \theta(e_n^*)$ is precisely the sum of the turning angles of $L(v)$. Condition (iii) is a consequence of the following discrete, hyperbolic version of Fenchel’s Theorem, in this case for closed polygonal curves in \mathbb{H}^3 .

Discrete Total Curvature for Polygonal Hyperbolic Loops (Rivin [61]) *The total discrete geodesic curvature of a closed, polygonal, hyperbolic space curve is greater than 2π , unless the vertices are collinear, in which case the total curvature is 2π .*

Proof The total discrete geodesic curvature of the polygonal hyperbolic space curve γ with vertices $p_1, \dots, p_k, p_{k+1} = p_1$ is $\sum_{i=1}^k \alpha_i$, where α_i is the turning angle of γ at p_i . The angle α_i is just the exterior angle at p_i of the triangle $\tau_i = p_{i-1}p_i p_{i+1}$. For $2 \leq i \leq k - 1$, let T_i be the triangle $T_i = p_1 p_i p_{i+1}$ with internal angles $a_i, b_i,$ and c_i at the respective vertices $p_1, p_i,$ and p_{i+1} . Note that by considering the triangles τ_i, T_{i-1} and T_i with common vertex p_i , the spherical triangle inequality gives

$$c_{i-1} + b_i \geq \pi - \alpha_i \quad \text{for } 3 \leq i \leq k - 1,$$

and

$$b_2 = \pi - \alpha_2, \quad c_{k-1} = \pi - \alpha_k, \quad \text{and} \quad \sum_{i=2}^{k-1} a_i \geq \pi - \alpha_1.$$

Recalling that $\pi \geq a_i + b_i + c_i$ with equality only when $p_1, p_i,$ and p_{i+1} are collinear, and then summing, one has

$$(k - 2)\pi \geq \sum_{i=2}^{k-1} (a_i + b_i + c_i) \geq k\pi - \sum_{i=1}^k \alpha_i,$$

with equality only when p_1, \dots, p_k are collinear. □

Theorem 5.7.1 (Rivin [59]) *The edge label $\theta(e^*)$ of the polyhedron P^* dual to the ideal convex polyhedron P defined above satisfies Conditions (i)–(iii).*

Proof Conditions (i) and (ii) already are verified. For Condition (iii), the circuit e_1^*, \dots, e_n^* that does not bound a face of P^* corresponds to a chain of contiguous faces f_1, \dots, f_n in P with $f_i \cap f_{i+1} = e_i$. $F = \cup_{i=1}^n f_i$ is a hyperbolic surface with boundary and cusps, and can be completed by extending geodesically across the boundary components to a complete immersed surface \tilde{F} in \mathbb{H}^3 without boundary. The surface \tilde{F} is an immersed hyperbolic cylinder with both ends of infinite-area. This observation uses the fact that the circuit e_1^*, \dots, e_n^* does not bound a face of P^* . Let γ be the unique closed geodesic path on the surface \tilde{F} that is freely homotopic

to the meridian. The curve γ is immersed in \mathbb{H}^3 as a polygonal curve lying on \tilde{F} with turning angles at the edges e_i . But it is easy to see that the turning angle of γ at edge e_i is no more than the exterior dihedral angle of the faces f_i and f_{i+1} that meet along e_i . This implies that the sum, $\theta(e_1^*) + \dots + \theta(e_n^*)$, which is the sum of these dihedral angles, is at least as large as the discrete geodesic curvature of γ , which in turn is greater than 2π by an application of the preceding theorem. \square

Rivin was able to turn this around and prove a converse to the theorem, which gives the following characterization of convex, ideal hyperbolic polyhedra. The existence is proved in [61], uniqueness in [60], and necessity of the three conditions in [59].

Characterization of Convex Ideal Polyhedra (Rivin [61]) *Let P^* be an abstract polyhedron. Then for any label $\theta : E(P^*) \rightarrow (0, \pi)$ that satisfies Conditions (i)–(iii), there is a convex, ideal hyperbolic polyhedron P in \mathbb{H}^3 whose Poincaré dual is P^* , and whose exterior dihedral angles at edges e are given by the values $\theta(e^*)$. Moreover, P is unique up to hyperbolic congruence. Conversely, every such polyhedron P satisfies Conditions (i)–(iii) as shown in Theorem 5.7.1.*

This characterization also proves the Circumscribable and Inscriptible Type Characterizations, answering Steiner’s question of 1832. This is because a convex, ideal hyperbolic polyhedron in the Beltrami–Klein projective model of \mathbb{H}^3 is represented by a convex Euclidean polyhedron inscribed in the 2-sphere S^2 .

Since Rivin’s work of the 1990s, several topologists and geometers have taken up the mantle and continued to unearth these beautiful gems of discrete geometry. I’ll close this survey with the mention of two examples in the next section, the first from the first decade of the new century, and the second of very recent origin.

5.7.4 New Millennium Excavations

Space constraints forbid too much further development of the topic, but I would be remiss if I didn’t mention at least these two beautiful theorems, the first characterizing convex hyperideal hyperbolic polyhedra by Bao and Bonahon, and the second giving a complete answer to Steiner’s original question when interpreted as broadly as possible, this time by Chen and Schlenker. I develop just enough of these topics to state the main results, and leave the interested reader the task of perusing the original articles for details of the proofs.

5.7.4.1 Hyperideal Polyhedra—Bao and Bonahon

A *hyperideal* polyhedron in \mathbb{H}^3 is a non-compact polyhedron that may be described most easily in the Beltrami–Klein projective model $\mathbb{H}^3 = \mathbb{B}^3 \subset \mathbb{RP}^3$ as the

intersection with \mathbb{B}^3 of a projective polyhedron all of whose vertices lie outside of \mathbb{B}^3 while each edge meets \mathbb{B}^3 . Bao and Bonahon [6] classify hyperideal polyhedra up to hyperbolic congruence in terms of their dihedral angles and combinatorial type in much the same vein as Rivin’s classification of ideal hyperbolic polyhedra. Note that Bao and Bonahon do allow for the vertices to lie on the sphere $\mathbb{S}^2 = \partial\mathbb{B}^3$ and hence their characterization reduces to Rivin’s for ideal polyhedra.

I will state the characterization in terms of conditions on the 1-skeletal graph of the dual polyhedron using Steinitz’s famous characterization of those graphs that may serve as the dual graph of a convex polyhedron in \mathbb{E}^3 as precisely the planar, 3-connected graphs.

Characterization of Convex Hyperideal Polyhedra (Bao and Bonahon [6]) *Let \mathcal{G} be a 3-connected graph embedded in \mathbb{S}^2 and $\theta : E(\mathcal{G}) \rightarrow (0, \pi)$. There is a hyperideal polyhedron P in \mathbb{H}^3 with dual graph isomorphic with \mathcal{G} and exterior dihedral angles given by θ if and only if the following conditions are satisfied.*

- (i) *If e_1, \dots, e_n forms a simple nontrivial circuit of edges of \mathcal{G} , then $\theta(e_1) + \dots + \theta(e_n) \geq 2\pi$, with equality possible only if e_1, \dots, e_n bounds a component of $\mathbb{S}^2 - \mathcal{G}$.*
- (ii) *If $\gamma = e_1, \dots, e_n$ forms a simple path of edges of \mathcal{G} that connects two vertices of \mathcal{G} that lie in the closure of a component C of $\mathbb{S}^2 - \mathcal{G}$, but γ does not lie in the boundary of C , then $\theta(e_1) + \dots + \theta(e_n) > \pi$.*

Moreover if P' is the projective polyhedron with $P' \cap \mathbb{H}^3 = P$, a vertex v of P' is located on the sphere at infinity of \mathbb{H}^3 if and only if equality holds in Condition (i) for the boundary of the corresponding component of $\mathbb{S}^2 - \mathcal{G}$.

Finally, the hyperideal polyhedron P is unique up to hyperbolic congruence.

I should mention that Hodgson and Rivin’s [46] characterization of compact and convex hyperbolic polyhedra can be applied to appropriate truncated polyhedra associated with those hyperideal polyhedra for which no vertex lies on the sphere at infinity to characterize them.

Define a *strictly hyperideal* polyhedron to be the intersection of \mathbb{B}^3 with a projective polyhedron P all of whose vertices lie outside the closed unit ball $\overline{\mathbb{B}^3} = \mathbb{B}^3 \cup \mathbb{S}^2$, but all of whose faces meet \mathbb{B}^3 . Note that this definition allows that an edge of P may lie entirely outside the closed ball $\overline{\mathbb{B}^3}$. These are yet to be characterized, but I mention that the article [23] verifies the rigidity of these that are bounded and convex, as long as no edges are tangent to the unit sphere. The proof again uses Cauchy’s toolbox.

5.7.4.2 Weakly Inscribed Polyhedra—Chen and Schlenker

Recall Steiner’s question of which polyhedra inscribe or circumscribe a sphere that Rivin answered. A more faithful translation of Steiner’s question from the German is “Does every polyhedron have a combinatorially equivalent realization that is

inscribed or circumscribed to a sphere, or to another quadratic surface? If not, which polyhedra have such realizations?” He includes the definition that “A polyhedron P is *inscribed to* a quadratic surface S if all the vertices of P lie on S ,” and further defines that P is *circumscribed to* S if all of its facets are tangent to S . As before I will concentrate on inscription since polarity relates circumscription to inscription. In the very recent preprint [30], Chen and Schlenker point out that the apparent grammar mistake—inscribed *to* instead of *in* S —makes a significant distinction.

Generally Steiner’s question has been interpreted to ask about inscription of the polyhedron P to a quadratic surface S in Euclidean space \mathbb{E}^3 , and in this setting P is contained in the bounded component of the complement of S , i.e., P is “inside” S , hence the change from inscribed “to” to “in”. But Steiner’s question makes sense in projective space as well, and in this setting a polyhedron may be inscribed to a surface without being inscribed in the surface. To be a bit more illustrative, consider the unit sphere \mathbb{S}^2 sitting in $\mathbb{E}^3 \subset \mathbb{R}\mathbb{P}^3$. Now \mathbb{S}^2 usually is thought of as the boundary of the open unit ball \mathbb{B}^3 that serves as the projective model of hyperbolic space, and this is what Rivin exploited in his characterization of those polyhedra inscribable in \mathbb{S}^2 . But \mathbb{S}^2 is also the boundary of the complement $\mathbb{R}\mathbb{P}^3 - \mathbb{B}^3$, which has a complete metric making it into a model of de Sitter space $d\mathbb{S}^3$. In this setting a projective polyhedron may have its vertices on the sphere \mathbb{S}^2 and yet not lie entirely in the ball \mathbb{B}^3 so that it is inscribed to \mathbb{S}^2 , but not inscribed in \mathbb{S}^2 in the usual meaning. Following Chen and Schlenker, I will revise Steiner’s terminology to emphasize the difference between *inscribed in* and *inscribed to but not in*.

Definition (Strong and Weak Inscription) In the real projective space $\mathbb{R}\mathbb{P}^3$, a polyhedron P inscribed to a quadratic surface S is *strongly inscribed in* S if the interior of P is disjoint from S , and *weakly inscribed to* S otherwise.

Before presenting a characterization of those polyhedra weakly inscribed to a sphere in $\mathbb{R}\mathbb{P}^3$, allow a word about polyhedra inscribed to other quadratic surfaces. This topic has been neglected until rather recently. There are only three quadratic surfaces in $\mathbb{R}\mathbb{P}^3$ up to projective transformations, and these are the sphere, the one-sheeted hyperboloid, and the cylinder. Danciger, Maloni, and Schlenker in [31] characterized the combinatorial types of polyhedra that are strongly inscribable in a one-sheeted hyperboloid or in a cylinder, and of course Rivin takes care of those strongly inscribable in a sphere. Chen and Schlenker’s work reported here characterizes those polyhedra weakly inscribable to a sphere, and the characterization of those weakly inscribable to the remaining two quadratic surfaces is the subject of current research by Chen and Schlenker.

Weak Inscription Characterization (Chen and Schlenker [30]) A 3-connected planar graph Γ is the 1-skeleton of a polyhedron $P \subset \mathbb{R}\mathbb{P}^3$ weakly inscribed to a sphere if and only if Γ admits a vertex-disjoint cycle cover by two cycles C_1 and C_2 with the following property. Color edge uv red if u and v both belong to C_1 or both belong to C_2 , and color it blue otherwise. Then there is a weight function $w : E(\Gamma) \rightarrow \mathbb{R}$ such that

- (i) $w > 0$ on red edges and $w < 0$ on blue ones;
- (ii) w sums to 0 over the edges adjacent to a vertex v , unless v is the only vertex on C_1 or C_2 (trivial cycle), in which case w sums to -2π over the edges adjacent to v .

I end this survey of progress in the characterization of polyhedra since Thurston's observation that every polyhedron type in \mathbb{E}^3 has a realization that midscribes a sphere with a description of the original rigidity theorem of Cauchy that is so instrumental in many of the proofs of the results surveyed here.

5.7.5 Addendum: Cauchy's Toolbox

In this bonus section I review Cauchy's celebrated rigidity theorem [29] of 1813 on the uniqueness of convex, bounded polyhedra in \mathbb{E}^3 . The theorem concerns two convex polyhedra with equivalent combinatorics and with corresponding faces congruent. Cauchy's Rigidity Theorem states that the two polyhedra must be congruent globally, meaning that there is a Euclidean isometry of the whole of \mathbb{E}^3 mapping one to the other. Like many of the great theorems of mathematics, the proof is of more importance than the theorem itself. As stated earlier in the introduction to this section, the toolbox Cauchy developed has been instrumental in the past 200 year development of the theory of polyhedra, especially in its rigidity theory. The proof, though at places clever and even subtle, overall is rather straightforward with a simplicity that belies its importance.

Cauchy's proof has two components—the one geometric and the other combinatorial. The geometric component is the Discrete Four Vertex Lemma, which follows from an application of Cauchy's Arm Lemma. Denote a convex planar or spherical polygon P merely by listing its vertices in cyclic order, say as $P = p_1 \dots p_n$. The Euclidean or spherical length of the side $p_i p_{i+1}$ is denoted as $|p_i p_{i+1}|$ and the interior angle at p_i is denoted as $\angle p_i$.

Cauchy Arm Lemma *Let $P = p_1 \dots p_n$ and $P' = p'_1 \dots p'_n$ be two convex planar or spherical polygons such that, for $1 \leq i < n$, $|p_i p_{i+1}| = |p'_i p'_{i+1}|$, and for $1 \leq i < n - 1$, $\angle p_{i+1} \leq \angle p'_{i+1}$. Then $|p_n p_1| \leq |p'_n p'_1|$ with equality if and only if $\angle p_{i+1} = \angle p'_{i+1}$ for all $1 \leq i < n - 1$.*

Cauchy's original proof of the lemma had a gap that subsequently was filled by Ernst Steinitz. A straightforward inductive proof, such as the one in [36], relies on the law of cosines and the triangle inequality.

Now let P and P' be convex planar or spherical polygons with the same number of sides whose corresponding sides have equal length. Label each vertex of P with a plus sign $+$ or a minus sign $-$ by comparing its angle with the corresponding angle in P' : if the angle at p_i is larger than that at p'_i , label it with a $+$, if smaller, a $-$, and if equal, no label at all. Using the Cauchy Arm Lemma, the proof of the following lemma is straightforward.

Discrete Four Vertex Lemma *Let P and P' be as in the preceding paragraph and label the vertices of P as described. Then either P and P' are congruent, or a walk around P encounters at least four sign changes, from $-$ to $+$ or from $+$ to $-$.*

Proof First note that because a polygon is a cycle, the number of sign changes must be even. If no vertex is labeled, then the two polygons are congruent. Assume then that some of the vertices are labeled, but all with the same label. Then Cauchy's Arm Lemma implies that there exists a pair of corresponding edges in P and P' with different lengths, a contradiction.

Assume now that there are exactly two sign changes of the labels of P . Select two edges $p_i p_{i+1}$ and $p_j p_{j+1}$ (oriented counter-clockwise) of P such that all of the $+$ signs are along the subchain from p_{i+1} to p_j and all of the $-$ signs are along the subchain from p_{j+1} back to p_i . Subdivide both edges in two by adding a vertex at the respective midpoints X and Y of $p_i p_{i+1}$ and $p_j p_{j+1}$. Similarly, subdivide the corresponding edges $p'_i p'_{i+1}$ and $p'_j p'_{j+1}$ in P' at midpoints X' and Y' . Denote the subchain of P from X to Y by P_+ and the subchain from Y back to X by P_- . Similarly for P'_+ and P'_- in P' . Applying the arm lemma to P_+ and P'_+ implies that $|XY| > |X'Y'|$, and, similarly, an application to P_- and P'_- implies that $|XY| < |X'Y'|$, a contradiction. \square

This brings us to the combinatorial component of Cauchy's proof. A nice proof of the following lemma appears in [36] and follows from an argument based on the Euler characteristic of a sphere.

Cauchy Combinatorial Lemma *Let P be an abstract spherical polyhedron. Then for any labeling of any non-empty subset of the edges of P with $+$ and $-$ signs, there exists a vertex v that is incident to an edge labeled with a $+$ or a $-$ sign for which one encounters at most two sign changes in labels on the edges adjacent to v as one walks around the vertex.*

Cauchy Rigidity Theorem *If two bounded, combinatorially equivalent, convex polyhedra in \mathbb{E}^3 have congruent corresponding faces, then they are congruent by a Euclidean isometry of \mathbb{E}^3 .*

Proof Assume that bounded, convex polyhedra P and P' have the same combinatorics and congruent corresponding faces. For each edge of P , label its dihedral angle with a $+$ or a $-$ depending on whether it is larger or smaller than the corresponding dihedral angle in P' . If P and P' are not congruent, Cauchy's Combinatorial Lemma provides a vertex v that is incident to an edge labeled with a $+$ or a $-$ sign, and around which there are at most two sign changes. Intersect P with a small sphere centered at v (one that contains no other vertex of P on its interior) to obtain a convex spherical polygon, and intersect P' with a sphere centered at the corresponding vertex v' and of the same radius. By construction both spherical polygons have the same edge lengths, and the angles between edges are given by the dihedral angles between faces at v and v' . An application of the Four Vertex Lemma implies that there are at least four sign changes, contradicting that there are at most two. It follows that P and P' are congruent. \square

Both the bounded and convex requirements are necessary. For example, a polyhedron \hat{H} in the shape of a cubical house with a shallow pyramidal roof has a cousin \check{H} obtained by inverting the roof. \hat{H} is not congruent to \check{H} , though these are combinatorially equivalent with congruent corresponding faces.

5.8 In Closing, an Open Invitation

This has been a whirlwind tour through the four decade history of the influence of one theorem brought to prominence by the mathematician we celebrate in this volume. Any result that has spawned such a great body of significant work leaves in its wake a bounty of open questions, problems, conjectures, and possible applications that await the right insights for resolution and explanation. What of the Koebe Uniformization Conjecture, of the question of where EQ-type sits among EEL- and VEL-type, of circle packings that mimic rational functions with arbitrary branching, of the existence and rigidity of inversive distance circle packings, of characterizations of projective polyhedra up to Möbius equivalence generalizing Bao-Bonahon, or of combinatorial rather than metric characterizations of hyperbolic polyhedra of various stripes? I have not covered in this survey the myriad of applications that circle packing has spawned, particularly in the realm of computer graphics and imaging, where each month sees more and more new and original publications. And so I close this tribute to the influence of this one theorem of Bill Thurston with an invitation to any reader who has been captured by the beauty and elegance of the results outlined in this survey to explore further on his or her own the wider discipline of Discrete Conformal Geometry, in both its theoretical and practical bents, and perhaps to add to our understanding and appreciation of this beautiful landscape opened up by the imagination of Bill Thurston.

Acknowledgments I thank Prof. Athanase Papadopoulos for inviting me to write on a favorite theme of mine to honor Bill Thurston and his legacy. It has been a pleasure for me to review the impact derived from this one beautiful little gem of Thurston. I thank Ken Stephenson for permission to use the graphics of Figs. 5.1, 5.2, 5.5, and 5.7, and John Bowers for generating the graphics for Figs. 5.3, 5.4, 5.6, 5.8, 5.9, and 5.10.

References

1. A.D. Aleksandrov, *Convex Polyhedra* (Translation of 1950 Russian ed.). Springer Monographs in Mathematics (Springer, Berlin, 2005)
2. E.M. Andre'ev, On convex polyhedra in Lobachevski spaces. *Mat. Sbornik* **81**(123), 445–478 (1970)
3. E.M. Andre'ev, On convex polyhedra of finite volume in Lobachevski spaces. *Mat. Sbornik* **83**(125), 256–260 (1970)
4. O. Angel, O. Schramm, Uniform infinite planar triangulations, *Commun. Math. Phys.* **241**, 191–213 (2003)

5. J. Ashe, E. Crane, K. Stephenson, Circle packings with generalized branching (2016). arXiv160703404A
6. X. Bao, F. Bonahon, Hyperideal polyhedra in hyperbolic 3-space. *Bull. Soc. Math. France* **130**(3), 457–491 (2002)
7. A.F. Beardon, K. Stephenson, The uniformization theorem for circle packings. *Indiana Univ. Math. J.* **39**, 1383–1425 (1990)
8. A.F. Beardon, K. Stephenson, The Schwarz-Pick lemma for circle packings, III. *J. Math.* **35**, 577–606 (1991)
9. M. Berger, *Geometry Revealed: A Jacob's Ladder to Modern Higher Geometry* (Springer, Berlin, 2010)
10. A.I. Bobenko, B.A. Springborn, Variational principles for circle patterns and Koebe's theorem. *Trans. Amer. Math. Soc.* **356**(2), 659–689 (2004)
11. A.I. Bobenko, T. Hoffmann, B.A. Springborn, Minimal surfaces from circle patterns: Geometry from combinatorics. *Ann. Math.* **164**(1), 231–264 (2006)
12. P.L. Bowers, The upper Perron method for labelled complexes with applications to circle packings. *Proc. Camb. Phil. Soc.* **114**, 321–345 (1993)
13. P.L. Bowers, Negatively curved graph and planar metrics with applications to type. *Mich. Math. J.* **45**, 31–53 (1998)
14. P.L. Bowers, Introduction to circle packing: the theory of discrete analytic functions [book review]. *Bull. Amer. Math. Soc.* **46**(3), 511–525 (2009)
15. J.C. Bowers, P.L. Bowers, Ma-schlenker c-octahedra in the 2-sphere. *Discrete Comput. Geometry* **60**, 9–26 (2017)
16. P.L. Bowers, M.K. Hurdal, Planar conformal mappings of piecewise flat surfaces, in *Visualization and Mathematics III*, chap. 1, (Springer, Berlin, 2003), pp. 3–34
17. P.L. Bowers, K. Stephenson, The set of circle packing points in the Teichmüller space of a surface of finite conformal type is dense. *Math. Proc. Camb. Phil. Soc.* **111**, 487–513 (1992)
18. P.L. Bowers, K. Stephenson, A branched Andreev–Thurston theorem for circle packings of the sphere. *Proc. London Math. Soc.* **73**(3), 185–215 (1996)
19. P.L. Bowers, K. Stephenson, A “regular” pentagonal tiling of the plane. *Conform. Geom. Dyn.* **1**, 58–86 (1997)
20. P.L. Bowers, K. Stephenson, *Uniformizing Dessins and Belyĭ Maps via Circle Packing*. *Memoirs of the AMS*, vol. 170, no. 805 (American Mathematical Society, Providence, 2004)
21. P.L. Bowers, K. Stephenson, Conformal tilings I: foundations, theory, and practice. *Conform. Geom. Dyn.* **21**(1), 1–63 (2017)
22. P.L. Bowers, K. Stephenson, Conformal tilings II: local isomorphism, hierarchy, and conformal type *Conform. Geom. Dyn.* **23**, 60 (2018)
23. J.C. Bowers, P.L. Bowers, K. Pratt, Rigidity of circle polyhedra in the 2-sphere and of hyperideal polyhedra in hyperbolic 3-space. *Trans. Amer. Math. Soc.* **371**, 4215–4249 (2018)
24. J.W. Cannon, The combinatorial structure of cocompact discrete hyperbolic groups. *Geometriae Dedicata* **16**(2), 123–148 (1984)
25. J.W. Cannon, The combinatorial Riemann mapping theorem. *Acta Math.* **173**, 155–234 (1994)
26. J.W. Cannon, W.J. Floyd, W.R. Parry, Finite subdivision rules. *Conform. Geom. Dyn.* **5**, 153–196 (2001)
27. C. Carathéodory, Untersuchungen über die konformen abbildungen von festen und veränderlichen gebieten. *Math. Annal.* **72**(1), 107–144 (1912)
28. C. Carathéodory, *Conformal Representation*. *Cambridge Tracts in Mathematics and Mathematical Physics*, vol. 28, Reprint of 1952 edn. (Cambridge University Press, Cambridge, 2008)
29. A. Cauchy, Sur les polygones et les polyèdres. *J. Ecole Polytechnique XVIe Cahier* **IX**, 87–98 (1813)
30. H. Chen, J.-M. Schlenker, Weakly inscribed polyhedra (2017). arXiv170910389C, to appear in *Trans. Amer. Math. Soc.*, Series B
31. J. Danciger, S. Maloni, J.-M. Schlenker, Polyhedra inscribed in a quadric. *Invent. Math.* **221**(1), 237–300 (2020)

32. H.P. de Saint-Gervais, *Uniformization of Riemann Surfaces: Revisiting a Hundred-Year-Old Theorem*. Heritage of European Mathematics, vol. 11 (European Mathematical Society, Zürich, 2016)
33. Y.C. de Verdière, Une principe variationnel pour les empilements de cercles. *Invent. Math.* **104**, 655–669 (1991)
34. P.G. Doyle, J. Laurie Snell, *Random Walks and Electric Networks*, vol. 22, 1st edn. (Mathematical Association of America, Washington, 1984)
35. R.J. Duffin, The extremal length of a network. *J. Math. Anal. Appl.* **5**, 200–215 (1962)
36. D. Fuchs, S. Tabachnikov, *Mathematical Omnibus: Thirty Lectures on Classic Mathematics* (American Mathematical Society, Providence, 2007)
37. J.T. Gill, S. Rohde, On the Riemann surface type of random planar maps. *Rev. Math. Iber.* **29**, 1071–1090 (2013)
38. J. Gray, On the history of the Riemann mapping theorem. *Rendiconti del Cir. Math. di Palermo* **II(34)**, 47/94 (1994)
39. M. Gromov, *Hyperbolic groups*, in *Essays in Group Theory*, ed. by S.M. Gersten (Springer, New York, 1987), pp. 75–263
40. B. Grünbaum, *Convex Polytopes*. Graduate Texts in Mathematics, vol. 221, 2nd edn. (Springer, Berlin, 2003)
41. R. Guo, Local rigidity of inversive distance circle packing. *Trans Amer. Math. Soc.* **363(9)**, 4757–4776 (2011)
42. Z.-X. He, O. Schramm, Fixed points, Koebe uniformization and circle packings. *Ann. Math.* **137**, 369–406 (1993)
43. Z.-X. He, O. Schramm, Hyperbolic and parabolic packings. *Discrete Comput. Geom.* **14**, 123–149 (1995)
44. Z.-X. He, O. Schramm, On the convergence of circle packings to the Riemann map. *Invent. Math.* **125(2)**, 285–305 (1996)
45. C.D. Hodgson, Deduction of Andreev’s theorem from Rivin’s characterization of convex hyperbolic polyhedra, in *Topology 90. Proceedings of the Research Semester in Low Dimensional Topology at O.S.U* (1993)
46. C.D. Hodgson, I. Rivin, A characterization of compact convex polyhedra in hyperbolic 3-space. *Invent. Math.* **111**, 77–111 (1993)
47. C.D. Hodgson, I. Rivin, W.D. Smith, A characterization of convex hyperbolic polyhedra and of convex polyhedra inscribed in the sphere. *Bull. Amer. Math. Soc.* **27**, 246–251 (1992)
48. P. Koebe, Über die uniformisierung beliebiger analytischer Kurven, III. *Nach. Ges. Wiss. Gott.* 337–358 (1908)
49. P. Koebe, *Abhandlungen zur theorie der konformen abbildung: VI. abbildung mehrfach zusammenhängender Bereiche auf Kreisbereiche*, etc. *Math. Z.* **7**, 235–301 (1920)
50. P. Koebe, *Kontaktprobleme der konformen abbildung*. *Ber. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Kl.* **88**, 141–164 (1936)
51. G.F. Lawler, V. Limic, *Random Walk: A Modern Introduction*. Cambridge Studies in Advanced Mathematics, no. 123 (Cambridge University Press, Cambridge, 2010)
52. F. Luo, Rigidity of polyhedral surfaces, III. *Geom. Topol.* **15(4)**, 2299–2319 (2011)
53. J. Ma, J.-M. Schlenker, Non-rigidity of spherical inversive distance circle packings. *Discrete Comput. Geom.* **47(3)**, 610–617 (2012)
54. A. Marden, B. Rodin, On Thurston’s formulation and proof of Andreev’s theorem, in *Proceedings of a Conference, held in Valparaíso Computational Methods and Function Theory, 1989*. *Lecture Notes in Mathematics*, vol. 1435 (Springer, Berlin, 1990), pp. 103–115
55. C. St. J.A. Nash–Williams, Random walk and electric currents in networks. *Proc. Camb. Phil. Soc.* **55**, 181–195 (1959)
56. U. Pinkall, K. Polthier, Computing discrete minimal surfaces and their conjugates. *Exp. Math.* **2**, 15–36 (2012)
57. G. Pólya, Über eine Aufgabe der Wahrscheinlichkeitsrechnung betreffend die irrfahrt im Strassennetz. *Math. Annal.* **84(1–2)**, 149–160 (1921)

58. S. Rhode, Oded Schramm: From circle packing to SLE, in *Selected Works of Oded Schramm*. Selected Works in Probability and Statistics, vol. 1 (Springer, Berlin, 2011), pp. 3–45
59. I. Rivin, On geometry of convex ideal polyhedra in hyperbolic 3-space. *Topology* **32**(1), 87–92 (1993)
60. I. Rivin, Euclidean structures on simplicial surfaces and hyperbolic volume. *Ann. Math.* **139**(3), 553 (1994)
61. I. Rivin, A characterization of ideal polyhedra in hyperbolic 3-space. *Ann. Math.* **143**, 51–70 (1996)
62. B. Rodin, D. Sullivan, The convergence of circle packings to the Riemann mapping. *J. Diff. Geom.* **26**, 349–360 (1987)
63. R.K.W. Roeder, J.H. Hubbard, W.D. Dunbar, Andreev’s theorem on hyperbolic polyhedra. *Ann. de l’Inst. Fourier* **57**(3), 825–882 (2007)
64. O. Schramm, Existence and uniqueness of packings with specified combinatorics. *Isr. J. Math.* **73**(3), 321–341 (1991)
65. O. Schramm, Rigidity of infinite (circle) packings. *J. Amer. Math. Soc.* **4**, 127–149 (1991)
66. O. Schramm, How to cage an egg. *Invent. Math.* **107**(1), 543–560 (1992)
67. O. Schramm, Transboundary extremal length. *J. d’Analyse Math.* **66**(1), 307–329 (1995)
68. E. Shulte, Analogues of Steinitz’s theorem about noninscribable polytopes, in *Proceedings of the Intuitive geometry, Siófok 1985*. Colloquia Mathematica Societatis János Bolyai, vol. 48 (1985), pp. 503–516
69. J. Steiner, *Systematische Entwicklung der Abhängigkeit Geometrischer Gestalten von Einander* (Fincke, Berlin, 1832)
70. K. Stephenson, A probabilistic proof of Thurston’s conjecture on circle packings. *Rendiconti del Seminario Mate. e Fisico di Milano* **LXVI**, 201–291 (1996)
71. K. Stephenson, *Introduction to Circle Packing: the Theory of Discrete Analytic Functions* (Cambridge University Press, New York, 2005). (ISBN 0-521-82356-0, QA640.7.S74)
72. D. Sullivan, On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions, in *Proceedings of the 1978 Stony Brook Conference on Riemann Surfaces and Related Topics* (Princeton University Press, Princeton, 1981), pp. 465–496
73. W.P. Thurston, *The Geometry and Topology of 3-Manifolds*. Lecture Notes (Princeton University, Princeton, 1980)
74. W. Woess, *Random Walks on Infinite Graphs and Groups*. Cambridge Tracts in Mathematics, no. 138 (Cambridge University Press, Cambridge, 2000)
75. W.E. Wood, Combinatorial type problems for triangulation graphs, Ph.D. Thesis, FSU, Tallahassee, Advisor Philip L. Bowers, 2006
76. W.E. Wood, Bounded outdegree and extremal length on discrete Riemann surfaces. *Conform. Geom. Dyn.* **14**, 194–201 (2010)