Circle Packing: Experiments in Discrete Analytic Function Theory

Tomasz Dubejko and Kenneth Stephenson

CONTENTS

Circle packings are configurations of circles with specified patterns of tangency, and lend themselves naturally to computer experimentation and visualization. Maps between them display, with surprising faithfulness, many of the geometric properties associated with classical analytic functions. This paper introduces the fundamentals of an emerging "discrete analytic function theory" and investigates connections with the classical theory. It then describes several experiments, ranging from investigation of a conjectured discrete Koebe $\frac{1}{4}$ theorem to a multigrid method for computing discrete approximations of classical analytic functions. These experiments were performed using CirclePack, a software package described in the paper and available free of charge.

1. INTRODUCTION

The topic of "circle packing" is of relatively recent origin and is a natural one for computer experimentation and visualization. What may be surprising, however, are the deep connections it shares with classical complex analysis. These connections are the subject of our paper.

Circle packings were introduced by Thurston, first in the construction of hyperbolic polyhedra and only later in connection with complex analysis. In particular, Thurston [1985] conjectured that maps between certain packings could be used to approximate classical conformal mappings. This conjecture was proved in [Rodin and Sullivan 1987]; many additional connections with analytic functions have emerged since then. This is fertile new ground for computer experimentation in, among other topics, geometry, combinatorics, probability, numerical approximation, and discrete complex analysis. Our purpose here is to demonstrate

1058-6458/96 $0.50 per page

features of the emerging theory through several experiments, conducted using CirclePack, a software package developed by Stephenson (Section 6).

A circle packing is a configuration of circles having a prescribed pattern of tangencies. Our interest lies in maps from one circle packing to another that preserve tangency relationships. These maps exhibit two related but distinct connections to complex analysis: *approximation* and *analogy*. Avoiding formalities for now, let's begin with the seminal example of approximation.

Approximation

Thurston's conjecture concerned the approximation of a conformal map F from the unit disc $\mathbb{D}$ onto a domain Ω. (We want to have $F(0) = 0$ and $F'(0) > 0$, but we'll suppress the issue of normalizations.) Three pairs of circle packings giving successively better approximants are illustrated in Figure 1. In each case the packing P_j on the right is a portion of a regular hexagonal packing—namely, those circles whose centers lie in Ω—while Q_j on the left is a repacking of the same combinatoric pattern (that is, the circles have changed sizes and locations, but remain tangent to one another in the same pattern as P_j) and is extremal in the unit disc. The existence of Q_j follows from the Koebe–Andreev–Thurston Theorem quoted later.

Given any such pair Q_j, P_j, one defines a piecewise affine mapping f_j that identifies the center of each circle of Q_j with the center of the corresponding circle of P_j. The Rodin–Sullivan Theorem tells us that such maps f_j (appropriately normalized) approximate the classical conformal mapping from $\mathbb{D}$ onto Ω. In particular, as the radii of the circles used to form P_j decrease to zero, the piecewise affine maps $f_j : Q_j \to P_j$ converge uniformly on compact subsets of $\mathbb{D}$ to the classical conformal map $F : \mathbb{D} \to \Omega$.

Discrete Analogy

The analogy aspect of our topic is suggested when one looks at individual maps, such as the f_j's in Figure 1. One finds that each, in isolation, already displays many geometric properties traditionally associated with analyticity. That is, each individual map between circle packings seems to present us with an object that behaves geometrically like a discrete analytic function. This perspective, spelled out below, is at the heart of our

THE FUNDAMENTAL ANALOGY

Classical analytic functions may be treated geometrically as mappings between domains. Globally, they must be continuous and orientation-preserving; however, the key conditions are entirely local: At a generic point the mapping is locally one-to-one and "conformal". This means angle-preserving, but is typically recast as "maps infinitesimal circles to infinitesimal circles". At isolated critical points the mapping may be branched, meaning that it has the local mapping properties of $z \mapsto (z-a)^k$ for some integer $k > 1$.

In the discrete setting, the domains and ranges are circle packings. A map $f : Q \to P$ is said to be a *discrete analytic function* if it preserves tangency and orientation: that is, if tangency of c_1, c_2 in Q implies tangency of $f(c_1), f(c_2)$ in P. This is the discrete version of continuity. Any mutually tangent triple of circles in P gets mapped to such a triple in Q; orientation-preservation is simply the requirement (assumed in the sequel without further comment) that if the former triple is positively oriented in Q, then so is the latter in P. We find that the "local" mapping conditions are quite automatic: A flower in P consists of a circle c, the "center", and the closed chain $\{c_1, \dots, c_n\}$ of successively tangent neighbors, the "petals", and comparing a flower in Q to its image flower in P, one sees the local mapping behavior of f. At a branch point, for example, the petals of the image flower wrap more than once around its center.

Summarized rather loosely: *A classical analytic function is a continuous orientation-preserving map carrying one pattern of infinitesimal circles to another, while a discrete analytic function is a continuous orientation-preserving map carrying one pattern of real circles to another.*

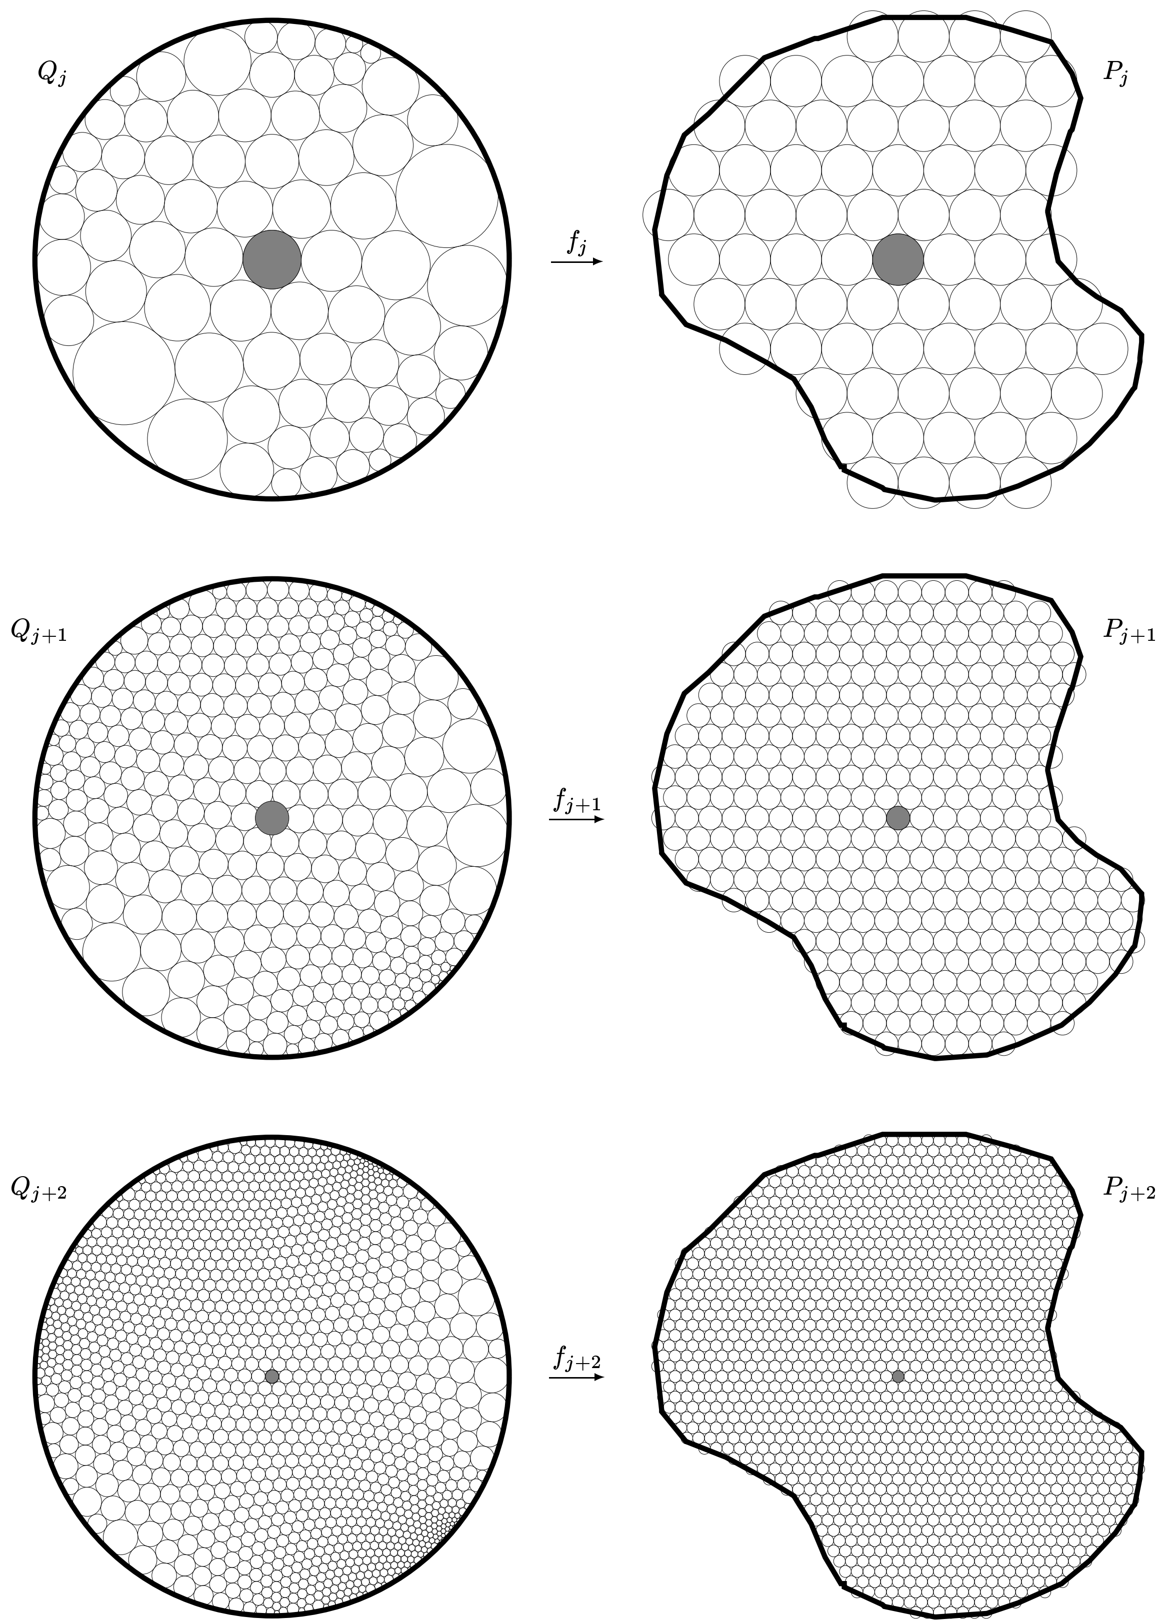

FIGURE 1. Approximations to a conformal mapping.

experiments, and we want to establish it firmly here at the beginning because it provides a familiar context for the entire paper.

In classical function theory, the Uniformization Theorem allows one to concentrate on functions defined on the three standard domains: the unit disc $\mathbb{D}$, the plane $\mathbb{C}$, and the Riemann sphere $\mathbb{S}^2$. Our discrete experiments can be arranged analogously: The Discrete Uniformization Theorem, discussed later, tells us that given an appropriate abstract pattern for a packing, encoded in an abstract complex K, there is an essentially unique extremal circle packing P_K, termed the "maximal packing" for K. Its circles have mutually disjoint interiors and pack one and only one of the spaces $\mathbb{D}$, $\mathbb{C}$, or $\mathbb{S}^2$, depending on K. (For instance, in Figure 1, the Q_j's are the maximal packings in $\mathbb{D}$ associated with the combinatorics of the P_j's.) Maximal packings serve as the standard domains for our maps.

Experimental Setup

Each of the discrete analytic functions f we consider will map a maximal packing P_K for some simply connected complex K to another circle packing P having the same combinatorics: $f : P_K \to P$. We adhere to the following standard normalization. A distinguished vertex of K is designated as α and the corresponding circle is centered at the origin in both P_K and P; in other words, $f(0) = 0$. Another distinguished vertex of K, designated γ, has its circle centered on the positive y-axis.

Examples

We present several examples, classified according to their classical models. When K triangulates a sphere, for example, P_K and P must both pack $\mathbb{S}^2$, so $f : P_K \to P$ is a discrete *rational* function. We will not be working with packings of $\mathbb{S}^2$ in this paper, but the reader may enjoy the illustration in Figure 2. The domain is the univalent maximal packing, while the range is a seven-sheeted branched cover of $\mathbb{S}^2$ with twelve branch points (the shaded circles). The combinatoric pattern of these packings happens to be dual to that of the buckminsterfullerene molecule C_{60}.

If P_K packs $\mathbb{C}$, then f is clearly a discrete *entire* or *entire meromorphic* function, depending on whether P is a planar or spherical packing. Liouville's Theorem holds, so P cannot lie in a bounded region; in fact, Callahan and Rodin [1993] have proved an analogue of Picard's Theorem (for the hexagonal case), stating essentially that P must cover the sphere, with the possible exception of at most two points. Examples are difficult to display, but Figure 3 demonstrates packings for a discrete

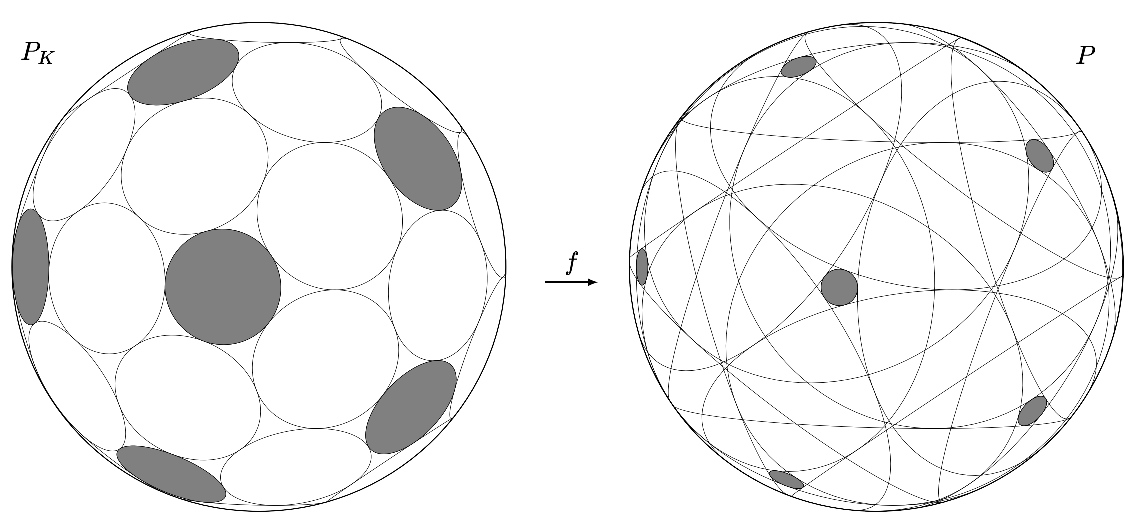

FIGURE 2. A discrete rational function.

analogue f of the complex sine function: the top diagram shows a portion Q of the domain, which is mapped by f to the upper half of the plane in the bottom (a regular hexagonal packing); note that $f(0) = 0$ and $f(\pi/2) = 1$. The full domain fills $\mathbb{C}$ and is generated by reflecting Q in the real axis and then repeatedly reflecting the result in its vertical sides. Simultaneously, f is extended by Schwarz reflection in the range. The reader can confirm that the mapping properties mimic those of $\sin z$.

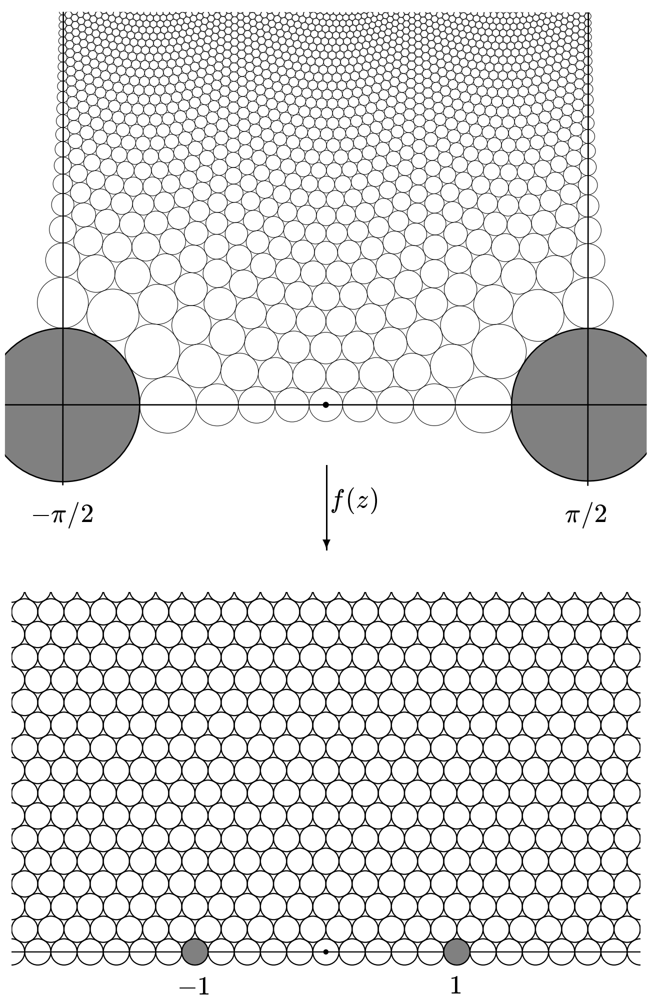

FIGURE 3. A discrete sine function.

One finds much more flexibility among discrete analytic functions on the disc—when P_K packs $\mathbb{D}$—and that is where most of our experiments take place. Figure 4 illustrates three examples, all with the same complex K (and therefore the common domain P_K). The first, f_1, has a range packing whose circles have mutually disjoint interiors; we call it *univalent*. The function f_2 has a range packing that twists back over itself, covering some parts of its range twice; we say it is a *locally univalent*, (globally) two-valent discrete analytic function. The third function f_3 is a branched two-valent function; the *branch point* is associated with a circle in its range packing P whose chain of neighbors wraps twice around it. This behavior is difficult to see in Figure 4, but we investigate it more closely in the next section.

These examples are intended to give the reader a basic intuition about discrete analytic functions. The emerging theory seems remarkably faithful to its classical counterpart: certain topological underpinnings are *pro forma*, but the rigidity imposed by the circles seems to force a geometric behavior tantamount to discrete analyticity. In addition to the Uniformization, Liouville, and Picard Theorems already mentioned, one finds in the literature geometrically precise discrete analogues of the Schwarz and Pick Lemmas, of Dirichlet-type boundary value problems and the Perron method, of distortion theorems, discrete versions of familiar classes of functions, polynomials, exponentials, Blaschke products, analogues of Brownian motion and harmonic measure, and so forth. Also, each discrete analytic function f induces a "ratio function" $f^{\#}$, defined later, whose behavior is remarkably like that of the modulus of a classical derivative.

Outline of the Paper

We investigate both the approximation and analogy facets of discrete analytic functions. The paper begins with the basics of circle packing: terminology, notation, existence and uniqueness results, and numerical algorithms. In Section 3 we formally define discrete analytic functions, state analogues of certain classical theorems, and give the fundamental approximation results.

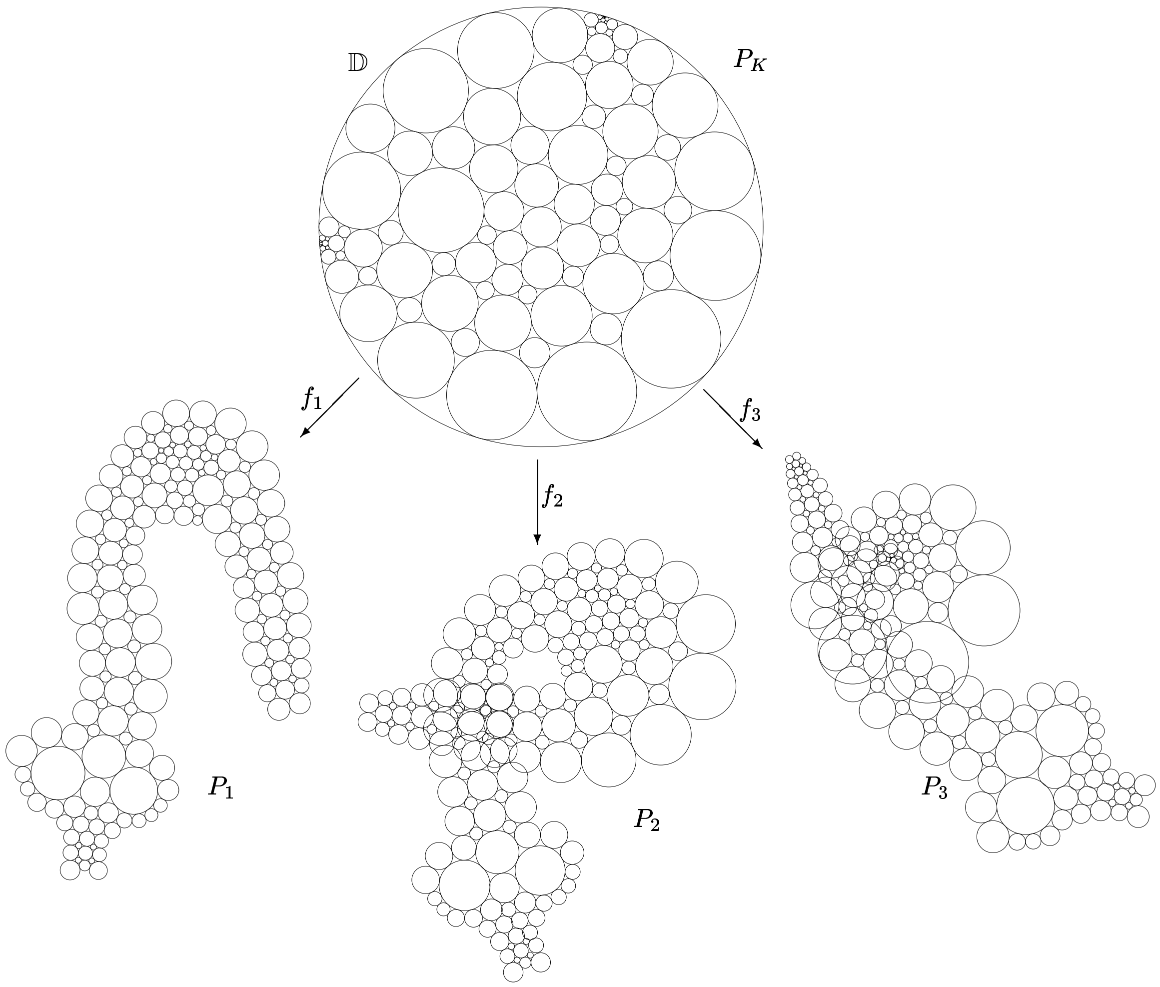

FIGURE 4. Discrete analytic functions on the disc.

The strong parallels with classical theory motivate three experiments on the analogy side, described in Section 4. Experiment #1 is aimed at Dieudonné's extension of the Schwarz Lemma, #2 at the Koebe $\frac{1}{4}$ Theorem, and #3 at generation of a discrete entire function. The fact that discrete functions also numerically approximate their classical models motivates four experiments on the approximation side, described in Section 5. Experiments #4 and #5 approximate finite Blaschke products and complex polynomials; #6 approximates disc algebra functions with specified boundary curves; and #7 returns to the conformal mappings first proposed by Thurston, introducing a "multigrid" method that shortens packing computations considerably.

The paper's experiments were carried out and the images generated using CirclePack, a software package developed by the second author (see Section 6). The multigrid method is implemented using specialized routines by the first author, executed through CirclePack. Needless to say, the live experiments and animated sequences of images available using CirclePack are far more informative than the few isolated images we could present here. The reader is encouraged to acquire the software by anonymous ftp (see the section on Electronic Availability at the end). Accompanying scripts allow one to run through most of the experiments and images of the paper using nothing more than mouse clicks. For the more involved reader, just a handful of rudimentary commands in CirclePack will permit open-ended experimentation—and there is a lot of new, unexplored territory.

2. BASIC CIRCLE PACKING

The key to working with circle packings lies in recognizing their dual natures: combinatoric on one hand and geometric on the other. To illustrate in a simple case, take the euclidean circle packing of Figure 5 (left). Its (euclidean) *carrier*, denoted $\operatorname{carr} P$, is the geometric two-complex shown on the right, formed by centers of circles, edges joining the centers of tangent circles, and faces formed by mutually tangent triples of circles. We encode the combinatorics of the packing in the abstract simplicial two-complex K that is simplicially equivalent to $\operatorname{carr} P$. The geometric information resides in a vector $R = \{\rho_0, \rho_1, \dots\}$ that has a radius ρ_j for each vertex v_j of K (that is, for each circle of P). We call $K(R)$ a labeled complex and write $P \sim K(R)$ to indicate the association with P. The centers of the circles would also seem to be important in our bookkeeping; but in fact, as we will see, they are essentially determined by K and R.

The labeled complex is the central organizing mechanism for all our subsequent work. One typically does not expect things to be as straightforward as in Figure 5, however: Radii may be in one of three different geometries; packings may twist around, with carriers that overlap themselves; they may have branch circles; they might even have infinitely many circles; and so forth. One can see the potential variety from illustrations in Section 1. The best way for the reader to develop intuition is by hands-on manipulation using CirclePack—which, incidentally, also uses labeled complexes for its bookkeeping.

Two circle packings Q and P have the same combinatorics if they share the same complex K, and in this case there is a natural *circle packing map* between them that identifies corresponding circles—that is, circles associated with the same vertex of K. These maps are the subjects of our experiments. We work primarily in the euclidean plane $\mathbb{C}$ and the hyperbolic plane, represented here as the unit disc $\mathbb{D}$ with the Poincaré metric.

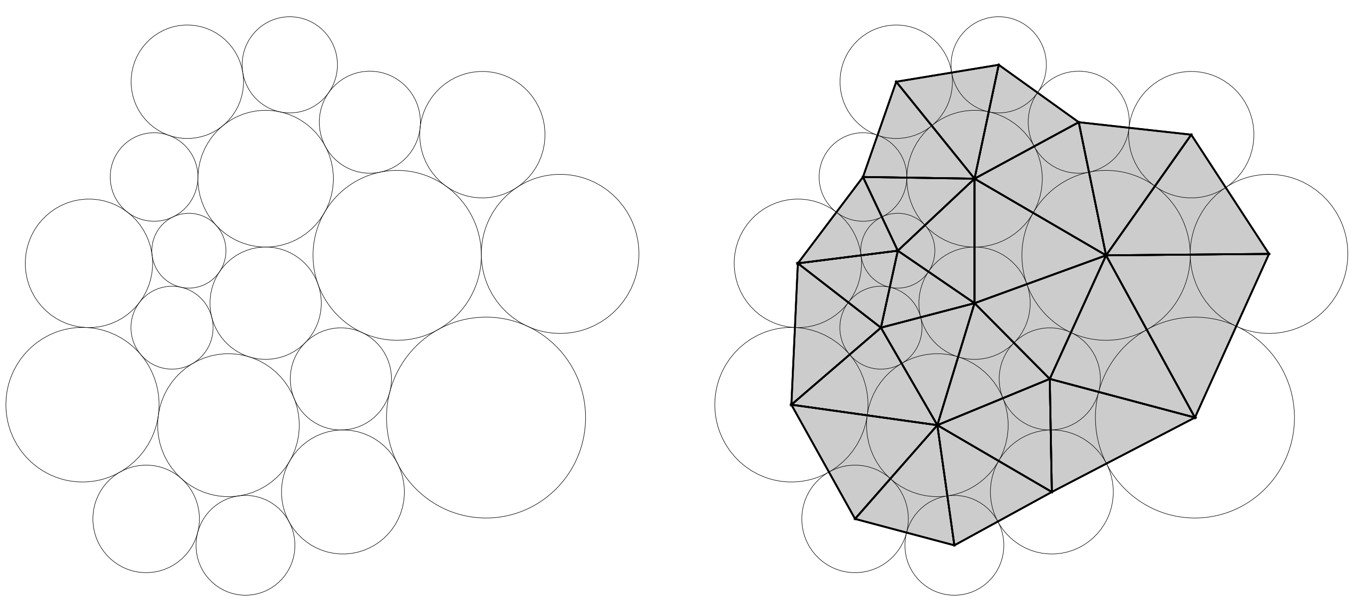

FIGURE 5. Left: A univalent euclidean packing. Right: The shaded object is the packing's carrier.

Combinatorics

The abstract pattern of tangencies in a packing—which circles are supposed to be tangent to which others—is referred to as the packing's "combinatorics" and is encoded in K. The appropriate conditions on K are most easily captured as follows: *The complex K is (simplicially equivalent to) a triangulation of an oriented topological surface.* Furthermore, we restrict attention here to complexes K that are simply connected and of bounded degree; the latter condition means that there is a finite bound on the number of vertices neighboring any one vertex. Two technical conditions are also assumed without further comment: namely, every boundary vertex of K must have at least one interior neighbor, and the set of interior vertices must be edge-connected.

When K is finite—as in any concrete experimental situation—it triangulates either a closed topological disc (the typical situation in this paper) or a sphere (if it has no boundary). Infinite complexes K that we encounter will be triangulations of open topological discs.

Geometry

The geometry of an actual circle packing for K resides in the radii and centers of its circles. In fact, the bookkeeping is best handled by concentrating on the radii, which may be treated as parameters. In particular, let D denote one of the geometric spaces $\mathbb{S}^2$, $\mathbb{C}$, or $\mathbb{D}$. A vector of positive numbers, denoted $R = \{\rho_0, \rho_1, \ldots\}$ with an entry $\rho_j = R(v_j)$ for each vertex v_j of K, is termed a *label* for K (in the geometry of D), and we refer to $K(R)$ as a *labeled complex.* Of course, most labels R will not be compatible collections of radii: circles of these sizes just wouldn't fit together in D according to the prescription of K.

Necessary and sufficient conditions for geometric compatibility are given in terms of angle sums: Fix a vertex v of K. For each face $\langle v, u, w\rangle \in K$, circles of radii $R(v)$, $R(u)$, $R(w)$ may be arranged in the prescribed geometry to form a mutually externally tangent triple. The appropriate law of cosines then gives an angle at v in the geometric face formed by the centers of these circles. Summing the angles over all faces of K containing v gives the *angle sum* at v, denoted $\theta(v) = \theta_R(v)$. If v is an interior vertex, circles having the radii associated with v and its immediate neighbors can be fit together in D to form a coherent flower if and only if

$$\theta_R(v) = 2\pi\eta \quad \text{for some integer } \eta = \eta_R(v) \geq 1. \tag{2.1}$$

This is known as the *packing condition* for R at v. The integer $\eta - 1$ is called the *order* for v; the neighbors of (the circle for) v wrap η times around it. When the packing condition (2.1) holds for all interior vertices v, we say that R is a *packing label* for K.

Theorem. *Suppose K is a simply connected complex and R is a label for K in the space D. Then R records the radii of some circle packing P for K in D if and only if R is a packing label. The packing $P \sim K(R)$ is uniquely determined up to automorphisms of D.*

Given a packing label R for K, an associated circle packing P is easily constructed. We describe the procedure as implemented in CirclePack: Recall that two distinguished vertices of K have been designated as α and γ. We build P by first centering a circle of radius $R(\alpha)$ at the origin. Next a vertex v neighboring α is chosen (the default in CirclePack is the first entry in the stored list of all neighbors of α) and a circle of radius $R(v)$ is drawn tangent to the circle for α with center on the positive x-axis. Now, remaining circles are added successively, each placed (with proper orientation) only after two contiguous neighbors have already been drawn. The conditions on our complexes ensure that all circles may eventually be placed and the packing condition (2.1) on interior radii guarantees that the order in which they are placed is irrelevant. A final normalization rotates the whole collection of circles so that the one associated with vertex γ is centered on the positive imaginary axis.

(This seems to work best for visual clues to packing symmetries.) The final configuration P is what we term *essentially unique*, that is, it is determined up to placement of the initial two circles (in other words, up to an automorphism of the underlying space D).

The natural simplicial map from K to $\operatorname{carr} P$ provides an immersion of K in the space D. If v is a vertex of K for which $\eta_R(v) = 1$, the immersion is locally univalent at v. If $\eta \geq 2$, the immersion has multiplicity η at v, and v is said to be a *branch point* (or *branch vertex* or *branch circle*) of order $\eta - 1$. As we will see, different packing labels R provide different immersions.

Maximal Packings

The above discussion begs the question of whether there exist any packing labels for a given K. The seminal result in the topic addresses this [Morgan 1984]:

Theorem (Koebe–Andreev–Thurston). *Let K be a triangulation of $\mathbb{S}^2$. There exists a circle packing P_K in $\mathbb{S}^2$ whose carrier is simplicially equivalent to K. Moreover, the circles of P_K have mutually disjoint interiors, and P_K is unique up to automorphisms and inversions of $\mathbb{S}^2$.*

Generalizations and extensions of this theorem say that for each (simply connected) complex K there exists an essentially unique canonical packing denoted P_K, which we call the *maximal packing* for K [Beardon and Stephenson 1990]. The appropriate space, $\mathbb{C}$, $\mathbb{D}$, or $\mathbb{S}^2$, is determined by the combinatorics of K, and we call K *parabolic*, *hyperbolic*, or *elliptic*, accordingly. The maximal packing label associated with P_K will be denoted by R_K. (Maximal packings may be defined using covering theory even for nonsimply connected complexes; see [Beardon and Stephenson 1990].)

Maximal packings enjoy the following properties:

- They are always univalent.
- When K triangulates the sphere, P_K packs $\mathbb{S}^2$ and K is elliptic.
- When K is finite but has boundary vertices—the situation that pertains in nearly all the packings of the paper—P_K is the so-called *Andreev packing* of $\mathbb{D}$, in which all the boundary circles are horocycles; in a completely natural way, a horocycle in $\mathbb{D}$ may be treated as a circle having infinite hyperbolic radius and center at the point of tangency with the unit circle $\mathbb{T}$. In particular, for each vertex $w \in \partial K$, we have $R_K(w) = \infty$.
- Infinite complexes must be either hyperbolic or parabolic. We will use only two parabolic complexes in the sequel: the hexagonal complex H, whose maximal packing is the familiar "penny packing" of Figure 6, left, and the "ball bearing

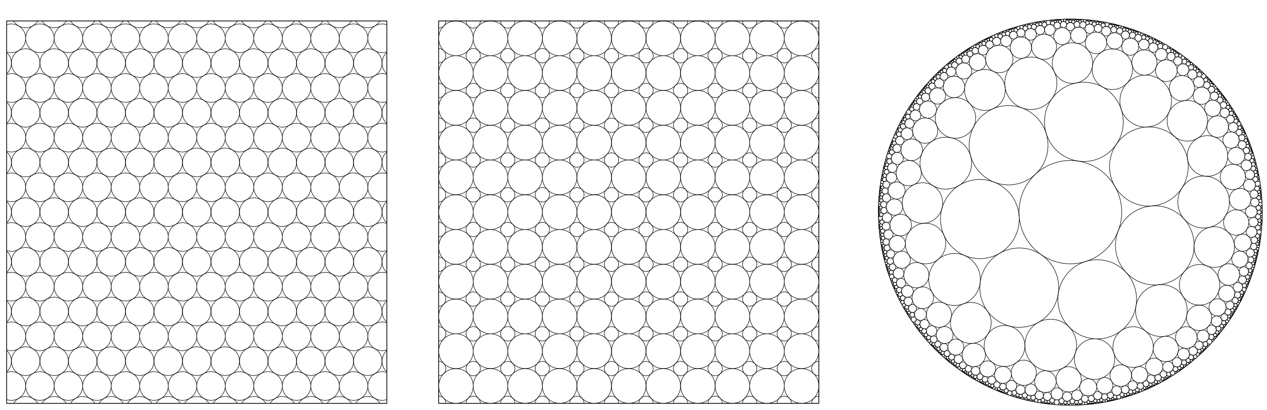

FIGURE 6. Left: The penny packing P_H. Middle: The maximal packing P_B of the ball-bearing complex. Right: The regular heptagonal packing (of the hyperbolic plane).

complex" B of Figure 6, middle. (The normalizations for P_H and P_B depend on the circumstances.) By way of contrast, the last part of Figure 6 illustrates the regular heptagonal circle packing, which is infinite but hyperbolic.

Results

Here we accumulate the basic theoretical results on circle packing that underlie our experiments. (They represent a consolidation of results from a variety of sources.) Assume that K is finite and simply connected, with boundary—that is, it triangulates a closed topological disc. Of first importance, of course, is the existence of an essentially unique maximal packing P_K, as noted previously; in this instance, P_K packs $\mathbb{D}$. Its "maximal" nature is reflected in the first result:

Discrete Schwarz Lemma. *If R is a hyperbolic packing label for K, then $R(v) \leq R_K(v)$ for every vertex v of K. Moreover, equality at a single interior vertex v implies $R \equiv R_K$.*

Global Monotonicity Lemma. *Let R_1 and R_2 be packing labels for the (hyperbolic or euclidean) complex K. If $R_1(w) \leq R_2(w)$ for every boundary vertex w, then $R_1(v) \leq R_2(v)$ for every vertex of K. Moreover, equality at a single interior vertex v implies $R_1 \equiv R_2$.*

Dirichlet Boundary Radii. *Let g be a positive function defined on the boundary vertices of K. Then there exists a unique euclidean packing label R_g for K with the property that $R_g(w) = g(w)$ for every boundary vertex w. The analogous result holds for hyperbolic packings, with the added feature that g may assume the value ∞.*

Dirichlet Boundary Angle Sums. *Let φ be a positive function defined on the vertices of ∂K, and assume that there is some euclidean packing $P \sim K(R)$ such that $\varphi(w) \leq \theta_R(w)$ for every $w \in \partial K$. Then there exists a unique packing label R' for K so that $\theta_{R'}(w) = \varphi(w)$ for all $w \in \partial K$. The analogous result holds in the hyperbolic setting, with the added feature that φ may assume the value 0.*

The Dirichlet results provide the theoretical basis for most of our packings; note, in fact, that a maximal packing in $\mathbb{D}$ is the solution of the hyperbolic Dirichlet problem with boundary radii specified to be infinity. Fortunately the solutions of these problems are effectively computable, meaning that one can compute approximations to any desired degree of accuracy. CirclePack employs an iterative algorithm first proposed by Thurston: For instance, to solve a typical boundary radius problem, one first assigns the specified radii to boundary vertices of K and arbitrary radii to the interior ones, giving an initial label R_I. Now, one repeatedly visits the interior vertices and upon each visit adjusts that component of the label so that the angle sum at that vertex is 2π. The succession of adjusted labels will converge to the desired packing label. (See [Bowers 1993], for instance.) The algorithm is an issue in Experiment #7; we also discuss CirclePack further in Section 6. Finally, we point out that our theoretical and computational results hold equally well when the packings share a common branch set, as described in the next section.

3. DISCRETE ANALYTIC FUNCTIONS

In spirit, a discrete analytic function f is a map between circle packings that preserves tangencies and orientation. Specifically, if Q and P are two circle packings with the same combinatorics, then $f : Q \to P$ is the map that identifies corresponding circles. This is slightly abstract for our purposes, however, so let's establish notation and give a more concrete definition.

Denote the vertices of K by $\{v_0, v_1, v_2, \ldots\}$, with $v_0 = \alpha$, the distinguished interior vertex used for normalizations. In the domain packing Q, denote the circle associated with v_j by C_j, the radius of C_j by r_j, and the center of C_j by z_j; in the range packing P denote the analogous quantities by c_j, ρ_j, and w_j. (Note that the radii and centers depend on the geometry; a packing in $\mathbb{D}$, for instance, might be treated as hyperbolic, euclidean, or spherical depending on the circumstances.)

Definition. A *discrete analytic function* f between two circle packings Q and P sharing the same complex K is a simplicial map

$$f : \operatorname{carr} Q \to \operatorname{carr} P$$

satisfying $f(z_j) = w_j$ for $j = 0, 1, 2, \ldots$. The associated *ratio function* $f^{\#}$ is the real function defined on the centers by

$$f^{\#}(z_j) = \frac{\rho_j}{r_j} = \frac{\operatorname{radius} c_j}{\operatorname{radius} C_j}, \quad \text{for } j = 0, 1, 2, \ldots.$$

The key condition on f is that it map centers to corresponding centers; from there, its simplicial nature (i.e., mapping edges to edges and faces to faces) can be arranged in various ways—for instance, using some type of barycentric coordinates. In any case, the resulting map f is continuous, open, discrete, and orientation-preserving, important mapping properties paralleling classical analytic functions. We will see that the ratio function $f^{\#}$, which after all reflects the factor by which a circle is stretched or shrunk under f, plays the role of the modulus of the derivative of f; it too could be extended in some ad hoc way to all of $\operatorname{carr} Q$, but we only use its values at the centers z_j.

As we have said, in spirit f is a map between circle packings, and so, despite the formal definitions, we will frequently abuse notation: In place of $f : \operatorname{carr} Q \to \operatorname{carr} P$, we generally just write $f : Q \to P$. We may well say $c_j = f(C_j)$, though technically this is not true of the circles as point sets under the map f between carriers. The vertex v_j of K, its circle C_j, and the center z_j of C_j are so strongly identified that we may refer to $f(v_j)$ when we actually mean $f(z_j)$, and to $f^{\#}(v_j)$ when we mean $f^{\#}(z_j)$. In discussing branching below, we attach the same meaning to the statement that "v_j is a branch vertex of K" or "c_j is a branch circle of P" as we do to the formally correct "z_j is a branch point of f".

We are now in a position to state various results in forms reminiscent of their classical models. For instance, the existence and essential uniqueness of the maximal packing P_K may be interpreted as a discrete version of the Uniformization Theorem: Within the combinatoric regime dictated by K, we may view P in the role of a general simply connected Riemann surface. The packing P_K is the "standard" domain for K—packing one of $\mathbb{S}^2$, $\mathbb{C}$, or $\mathbb{D}$, depending on K.

Discrete Uniformization Theorem. *Assume that P is a circle packing whose complex K is simply connected. There exists a discrete analytic function $f : P_K \to P$, where P_K is (that is, packs) the sphere, the plane, or the unit disc. The function f is unique up to automorphisms of P_K.*

In our experiments, the domain packing Q is always P_K, the maximal packing for K in its native geometry. The image packing P will be considered in a geometry appropriate to the circumstances. Our standard normalization places vertex α at the origin in both P_K and P, so $f(0) = 0$, and places the vertex γ on the positive y-axis in both domain and range.

The branch structure of $f : P_K \to P$ requires some explanation: Suppose v is an interior vertex of K and $\{u_1, u_2, \ldots, u_k\}$ is the closed chain of neighboring vertices. The associated circles of P_K form a necessarily univalent flower, with the angle sum $\theta_{R_K}(v) = 2\pi$. In the range packing $P \sim K(R)$, however, the corresponding chain of circles $\{c_1, c_2, \ldots, c_k\}$ may wrap some number η of times around v, meaning $\theta_R(v) = (2\pi)\eta$; thus v is associated with a branch point of order $\eta - 1$ for f. Figure 7 illustrates a flower from P_K with six petal circles and the corresponding branched flower from P; in the latter, the petals wrap twice around the shaded center ($\eta = 2$), for a simple branch point.

The collection β of branch vertices, each repeated according to its order, is called the *branch set* of f. Combinatoric necessary and sufficient conditions on branch sets are described in [Dubejko 1995]: β is a branch set for some f if and only if every simple closed edge-path σ in K has at least $2k + 3$ edges, where k is the number of points of β that σ encloses. This condition is easily verified when the vertices of β are not too crowded together, and no

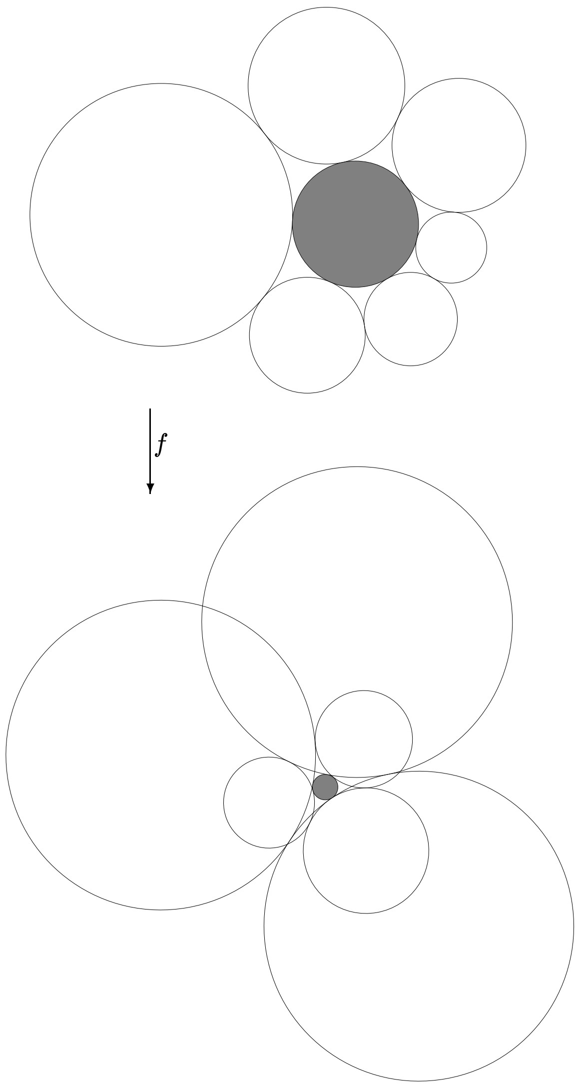

FIGURE 7. Local behavior at a branch point.

difficulties arise in arranging the branch points in our experiments.

The circle packing map f_3 of Figure 4 has one simple branch point; indeed, the flowers of Figure 7 are taken from its domain and range. The structure of P_3 is actually not so complicated if one views it as being comprised of two "sheets", as suggested in Figure 8: when these sheets are cross-connected along the indicated branch cut ending at the branch circle, they project to form P_3.

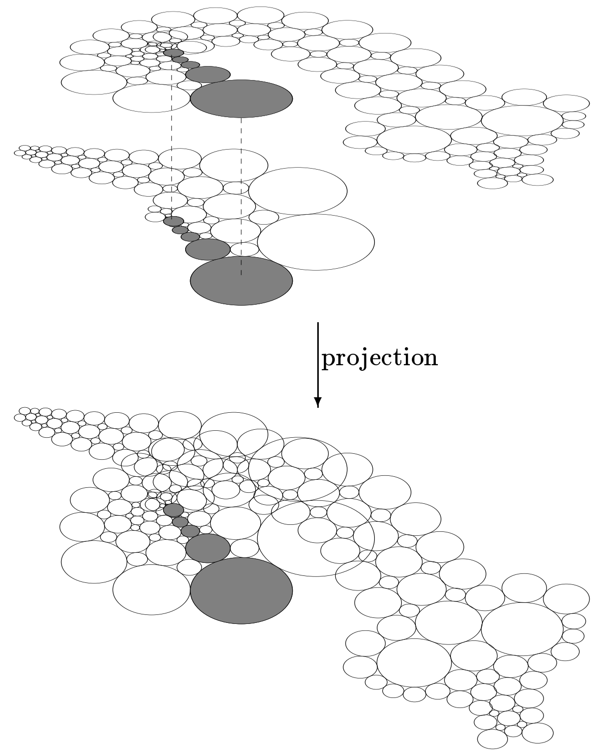

FIGURE 8. Sheets pasted to form a branched packing.

Discrete Analogy

We hope that the analogy with analytic functions is clear to the reader. It might help to restate (and extend) the Discrete Schwarz Lemma of the previous section in the new terminology.

Discrete Schwarz–Pick Lemma [Beardon and Stephenson 1990, Lemma 5]. *Let $f : P_K \to P$ be a discrete analytic function with $P \subset \mathbb{D}$ and $f(0) = 0$. Then $f^{\#}(0) \leq 1$, with equality if and only if f is an automorphism of $\mathbb{D}$. Moreover, if d denotes hyperbolic distance in $\mathbb{D}$, then*

$$d(f(z_j), f(z_k)) \leq d(z_j, z_k)$$

for any two interior centers z_j and z_k of P_K, with equality if and only if f is an automorphism of $\mathbb{D}$.

In Section 1, we mentioned several other specific parallels, though in less precise terms. The Discrete Liouville Theorem, for instance, would say that there is no nontrivial discrete analytic function defined on the plane (that is, with a parabolic complex) whose image lies in a bounded region of the plane. We leave formulations of other stated results to the reader. Experiments with further analogies are described in the next section.

Discrete Approximation

We are also in position, now, to quote the fundamental approximation result in the terminology of discrete analytic functions. Here Ω is a bounded simply connected domain in the plane with distinguished points a, b; the P_j, for $j = 1, 2, \ldots$, are univalent circle packings lying in Ω with complexes K_j. Write f_j for the univalent discrete analytic functions $f_j : P_{K_j} \to P_j$.

Theorem. *Using the notation established above, assume the following conditions:*

(1) *The radius of the largest circle of P_j goes to zero as j goes to infinity.*
(2) *The carriers of the P_j exhaust Ω.*
(3) *The degrees of the K_j are uniformly bounded above.*
(4) *$f_j(0) \to a$ as j goes to infinity.*
(5) *$f_j^{-1}(b) > 0$ for all j.*

Then the discrete analytic functions f_j converge uniformly on compact subsets of $\mathbb{D}$ to the classical conformal mapping $F : \mathbb{D} \to \Omega$ satisfying $F(0) = a$ and $F^{-1}(b) > 0$. Moreover, the ratio functions $f_j^{\#}$ converge uniformly on compact subsets of $\mathbb{D}$ to $|F'|$.

The convergence $f_j \to F$ was proven in [Rodin and Sullivan 1987] in the case that the K_j are hexagonal complexes, confirming Thurston's 1985 conjecture. It was extended to general complexes in [Stephenson 1990], though with an added assumption of a uniform bound on the ratios of the radii of any two circles from P_j. The result as stated is proven in [He and Rodin 1993].

Because analyticity is a local property, this original conformal case has wider implications. By relaxing global univalence and incorporating branch points, Dubejko has shown how to approximate, first finite Blaschke products, then polynomials, and thereby general analytic functions on $\mathbb{D}$ [Dubejko 1995; a]. Also, by results in [Dubejko and Stephenson 1995], convergence $f_j \to F$ in any of these settings implies convergence $f_j^{\#} \to |F'|$. Our experiments involving approximations are in Section 5.

4. DISCRETE ANALOGY

Our first three experiments address analogies between discrete and classical analytic functions. Basic topological parallels—open mapping, discreteness, orientation, and so forth—come free in the discrete theory because we use simplicial maps. It is only with the circles, however, that we acquire the geometry and rigidity so reminiscent of classical analyticity.

Experiments #1 and #2 investigate fundamental geometric constants. Both suggest not only that analogous constants exist in the discrete theory, but that they may be quite close to their classical values. (We display only a representative sample of our experimental trials.) Of course, experiments alone prove nothing. They can be valuable, nonetheless, particularly since in some sense it is "coarse" circle packings, those with small numbers of circles, that should be considered; by the approximation results, one expects the discrete analytic functions based on very "fine" circle packings to reflect the behavior of classical analytic functions.

Analogy is a two-way street, of course, and Experiment #3 attempts to use classical motivations to address an existence question in the discrete theory.

Experiment #1: The Dieudonné–Schwarz Lemma

According to the classical Schwarz Lemma, if $F : \mathbb{D} \to \mathbb{D}$ is analytic with $F(0) = 0$, then $|F'(0)| \le 1$, and we have seen the discrete version stated above. However, there is an extension of the classical result due to Dieudonné that shows that

$$|F'(z)| \le \begin{cases} 1 & \text{for } |z| \le \sqrt{2} - 1, \\ \dfrac{(1 + |z|^2)^2}{4|z|(1 - |z|^2)} & \text{for } \sqrt{2} - 1 < |z| < 1, \end{cases}$$

and these bounds are sharp. Does this also carry over to discrete analytic functions? In this section we consider the inequality for $|z| \leq \sqrt{2}-1$, and in Section 5 the inequality for $|z| > \sqrt{2}-1$. Since we are following an experimental philosophy here, we pose the question as a hypothesis:

Hypothesis. *There exists a universal constant* $\mathcal{C}_1 > 0$ *so that the following holds*: *If* K *is hyperbolic and* $f : P_K \to P$ *is a discrete analytic function with* $P \subset \mathbb{D}$ *and* $f(0) = 0$ *and if* z *is the* (*hyperbolic*) *center of a circle of* P_K *with* $|z| \leq \mathcal{C}_1$, *then* $f^{\#}(z) \leq 1$.

Figure 9 illustrates a simple prototype experiment. The maximal packing P_K is on the left, the range packing P on the right. CirclePack was used to compute the ratio of euclidean radii of circles in P to their counterparts in P_K and to shade in the circles in P_K for which this ratio exceeds 1. The Hypothesis concerns how closely the centers of these shaded circles can approach the origin. For purposes of reference, the classical "exclusion zone", the region wherein the modulus of the derivative cannot exceed 1, is indicated by a circle of radius $\sqrt{2}-1$ superimposed on P_K.

There is an endless supply of experiments to try: different complexes, choices of branch point(s), and choices of boundary radii. As a challenge, we decided to fix a complex—the one underlying Figure 9—and see which circles we could force to become shaded. Imagine this as a video game:

PackMan. The domain P_K on the left of your screen is fixed, and on the right is another hyperbolic packing P for K. With your controls you can change the boundary radii and/or branch structure of P. With each change, CirclePack immediately recomputes the new P and then shades any circle of P_K that happens to be (euclideanly) smaller than its new counterpart. Your goal is to manipulate P so as to shade as many circles of P_K as possible (thereby saving the planet).

Figure 10 displays the accumulated results—we did our best, but were not able to shade any circles having centers of modulus less than $\sqrt{2}-1$.

What's your best strategy? How can you force the shading closer to the origin? The game is quite intriguing: The Discrete Schwarz Lemma implies that a circle of P is *hyperbolically* smaller than its counterpart in P_K; thus the only hope that it be

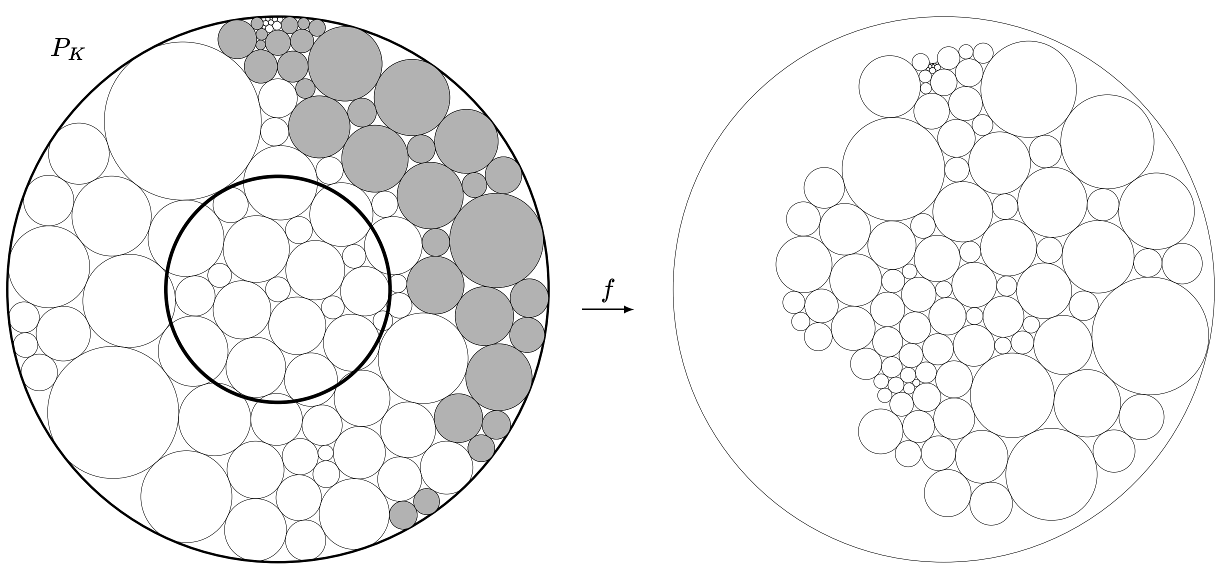

FIGURE 9. Prototype of Experiment 1. The inner circle in the domain represents the exclusion zone, of radius $\sqrt{2}-1$.

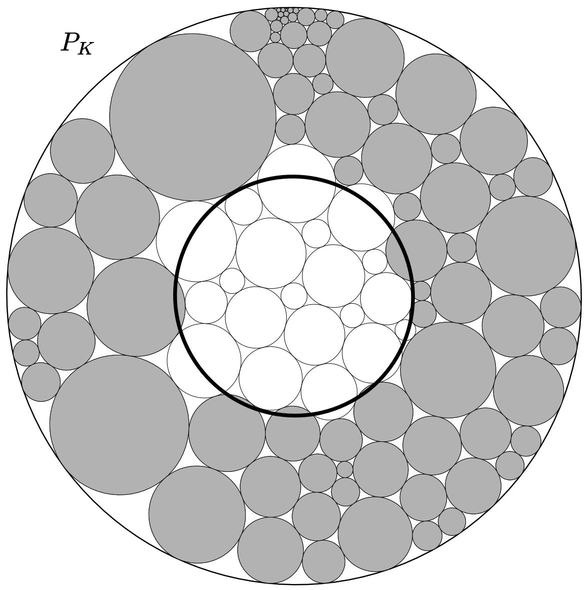

FIGURE 10. Accumulated PackMan trials with fixed K.

euclideanly larger is to force its center towards the origin in P. There are two ways you might try to arrange this: either by introducing branch points in P, thus sweeping some of its circles towards the origin, or by decreasing some boundary circles of P, in the hope of pulling others towards the origin.

The first approach has a classical precedent: The functions that show the sharpness of Dieudonné's classical constant are twofold Blaschke products. The discrete analogue would involve a packing P having a single branch point and infinite boundary radii. More about such functions in Section 5.

The second approach is a tricky balancing act: When you reduce boundary circles, the repacking computations (to form the new P) will bring circles towards the origin while simultaneously shrinking their hyperbolic radii. To see how delicate nature is in balancing these competing effects, consider the discrete analytic function f of Figure 11: The packings P_K and P are quite similar. Visualize P as being obtained from P_K by pushing in to make a small dent in $\mathbb{D}$ near $w = i$. As one might expect, reducing circles to make this dent while keeping the origin fixed forces a *stretching* behavior on the opposite side of the origin, and that is reflected in the shaded circles of P_K. Note that some of their centers lie perilously close to the classical exclusion zone.

All our experiments tend to support the existence of a discrete exculsion zone of some radius $\mathcal{C}_1$ about the origin, and only recently did we find a situation that cracks the classical barrier, showing that $\mathcal{C}_1$—if it exists—is necessarily less than $\sqrt{2}-1$. The example f is shown in Figure 12. The range packing P has a branch circle far out near

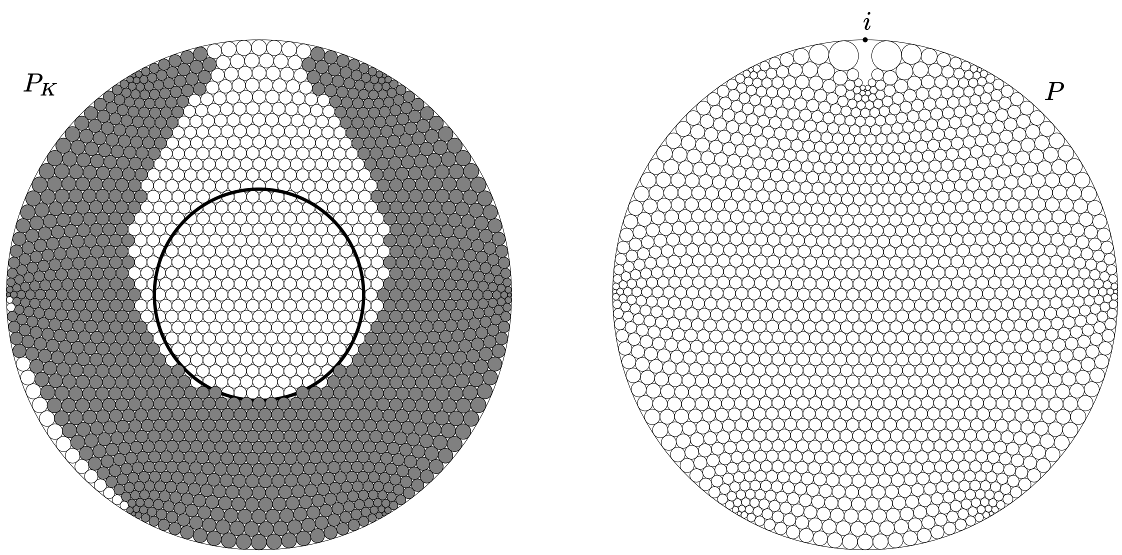

FIGURE 11. A delicate balance.

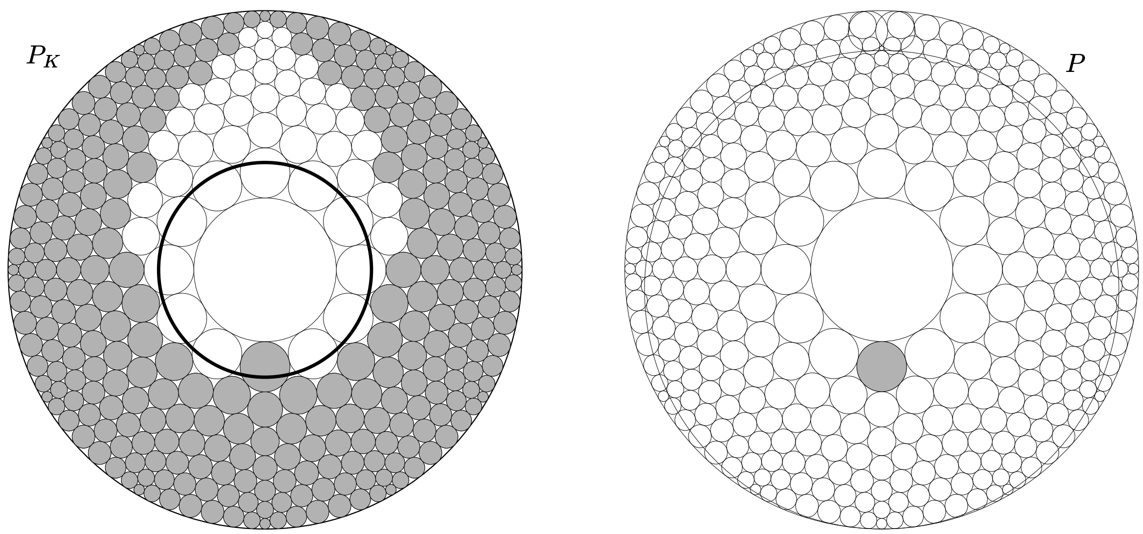

FIGURE 12. Penetrating the classical zone.

the boundary, which accounts for the large horocycle covering most of the disc. There is one vertex v whose circle is shaded and penetrates the exclusion zone, its center having modulus approximately 0.378; we point out, however, that it is a very near thing since $f^{\#}(v) = 1.0000112$ just barely exceeds 1 in this example. (Our thanks to Richard Rupp for running the experimental trials that suggested this function.)

In summary, there seems to be a discrete exclusion zone somewhat smaller than the classical one.

Experiment #2: Koebe's $\frac{1}{4}$ Theorem

The classical Koebe $\frac{1}{4}$ Theorem may be restated as follows: *Let F be analytic and univalent on the unit disc $\mathbb{D}$, with $F(0) = 0$. If its range $F(\mathbb{D})$ fails to contain the closed unit disc $\bar{\mathbb{D}}$, then $|F'(0)| \le 4$; this bound is sharp.* Our experiments probe the discrete analogue. For $\lambda > 0$, let $D(\lambda)$ denote the disc $\{|z| < \lambda\}$.

Hypothesis. *There exists a universal constant $\mathcal{C}_2$ so that the following holds: Let K be hyperbolic and let $f : P_K \to P$ be a univalent discrete analytic function with $f(0) = 0$. Suppose* $\operatorname{carr} P_K$ *contains the disc $D(\lambda)$. If* $\operatorname{carr} P$ *fails to contain the closed unit disc $\bar{\mathbb{D}}$, then $f^{\#}(0) \le \mathcal{C}_2/\lambda$.*

For the classical theorem (wherein λ is 1), the sharp bound is 4. Is the discrete result true with $\mathcal{C}_2$ equal to 4? Approximation results imply that $\mathcal{C}_2$—if such a constant exists—can be no smaller than 4.

This problem highlights some more-or-less typical *discretization* issues that can arise when mimicking a continuous theory. For instance, what do we mean by the range of f? We have chosen to interpret this as the carrier of P; in our runs, then, we can avoid covering $\bar{\mathbb{D}}$ by having some boundary circle of P centered at the point $w = i$. Likewise, we are forced to consider the carrier in the domain, which is why the disc $D(\lambda)$ is needed in the Hypothesis. Let $\lambda_K \le 1$ denote the supremum of those λ with $D(\lambda) \subset \operatorname{carr} P_K$. Then λ_K can be arbitrarily small, depending on K and v_0, and we must account for it in the bound on $f^{\#}(0)$—just consider that $D(\lambda_K)$ plays the role of domain played by the unit disc $D(1)$ in the classical setting.

Next, what is the most appropriate interpretation of univalence? In the Hypothesis, we mean that the circles of P have mutually disjoint interiors. However, that turns out to be difficult to arrange computationally; there are as yet no univalence criteria in circle packing (no analogues, for example, of Nehari's condition) that might tell you from the radii alone whether a packing P will be

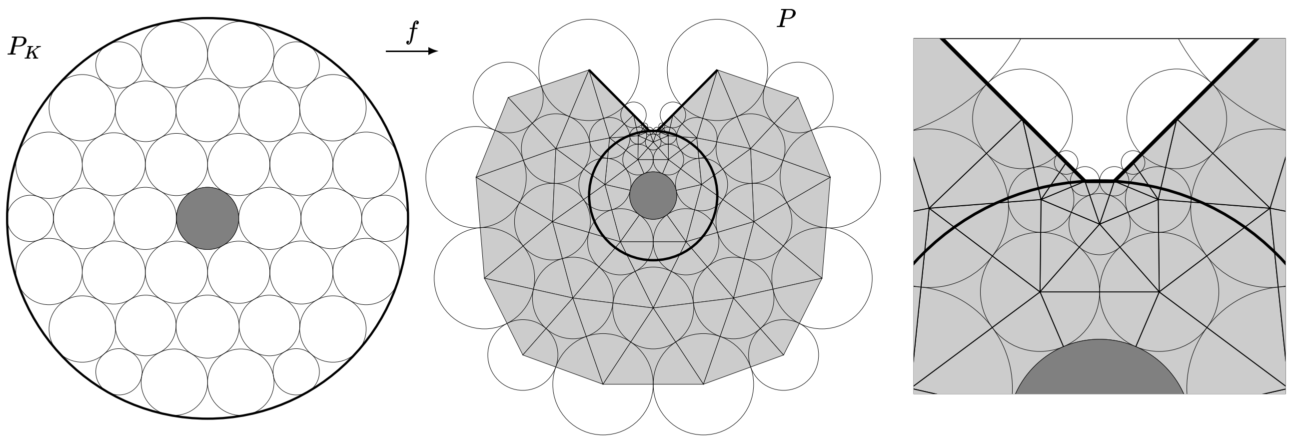

FIGURE 13. Prototype of Koebe Experiment. The thick circle in the range is the unit circle.

univalent when laid out. For practicality's sake, we substitute the weaker condition of *carrier univalence*, meaning that the carrier of P is univalently immersed in $\mathbb{C}$. This encompasses more functions, so any constant obtained provides an upper bound on the constant $\mathcal{C}_2$ of the Hypothesis.

In the classical theory, the extremal functions are derived from the Koebe function $k(z) = z/(1-z)^2$. Specifically, multiply this by 4 and rotate clockwise by $\pi/2$ to get $F(z) = -4ik(z)$, which is a classical univalent analytic function satisfying $F(0) = 0$ and with range $\mathbb{C}\backslash\{iy : y \in [1,\infty)\}$. Its image region, in essence, has had its boundary pushed off to infinity, apart from a portion that sneaks in to capture the point $w = i$, thereby failing to cover $\bar{\mathbb{D}}$. A computation shows that $|F'(0)| = 4$. We model our discrete experiments on F.

The prototype experiment is shown in Figure 13. Here $\lambda_K \approx 0.82$, so our aim is to compare $f^{\#}(0)$ to $4/\lambda_K \approx 4.88$. The carrier of P omits the imaginary axis from $w = i$ up; our normalization and the combinatoric symmetry of K were chosen to aid the reader in interpreting the pictures. In this instance, we forced P to form a deep notch by prescribing certain boundary angle sum conditions:

(a) The bottom of the notch, enlarged in the right-hand panel of Figure 13, was formed by two neighboring boundary circles with angle sums prescribed at $\frac{5}{4}\pi$.

(b) Each side of the notch was formed by two boundary circles having prescribed angle sums of π.

(c) The remaining boundary circles had no angle sum constraints, but rather had large prescribed radii.

This is a mixed Dirichlet problem. There exists an essentially unique circle packing satisfying these three conditions, and the approximate solution was computed with CirclePack; normalization places the proper circle at the origin and ensures that the bottom of the notch reaches $w \approx i$. The result is the packing P displayed in Figure 13. Care was taken in specifying the radii of condition (c) to maintain univalence (and in this case, symmetry).

The value of $f^{\#}(0)$ obtained is 2.36. It is clear, however, that we can improve on this. The idea is to further increase the circle at the origin by tightening up conditions (a), (b), and (c): make the notch sharper, include more circles in its sides, and enlarge the remaining free boundary circles, and do this while keeping the packing univalent. Two additional trials are shown in Figure 14. The values of $f^{\#}(0)$ were increased to 3.06 (top row) and 3.38 (bottom row).

We applied an additional technique, which was suggested by one of the referees of this paper, to get even more out of this same K. Namely, we augmented K by adding seven edges, one to connect each circle on the left side of the notch in Figure 14

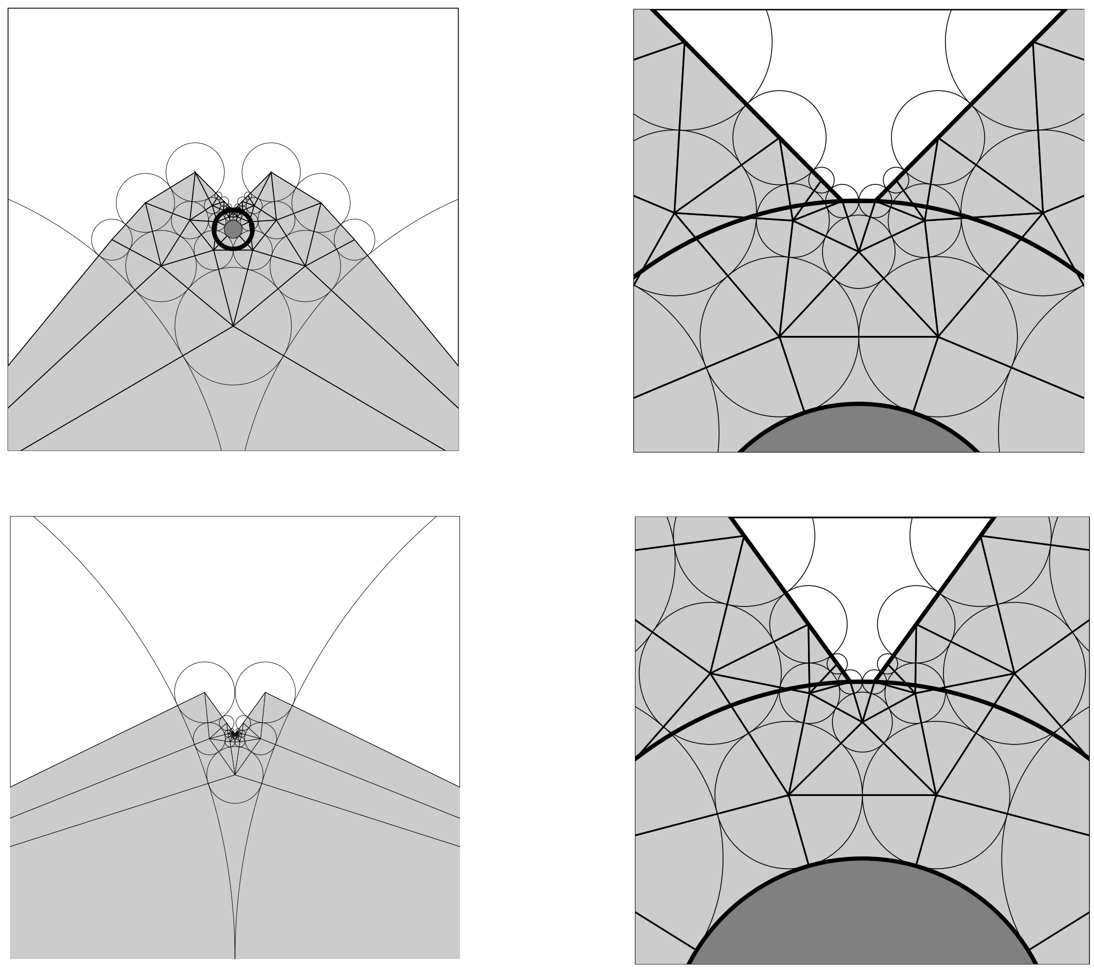

FIGURE 14. Two experiments to explore the bound in the circle-packing equivalent of Koebe's theorem. On the top, $f^{\#}(0) = 3.06$, and on the bottom, $f^{\#}(0) = 3.38$.

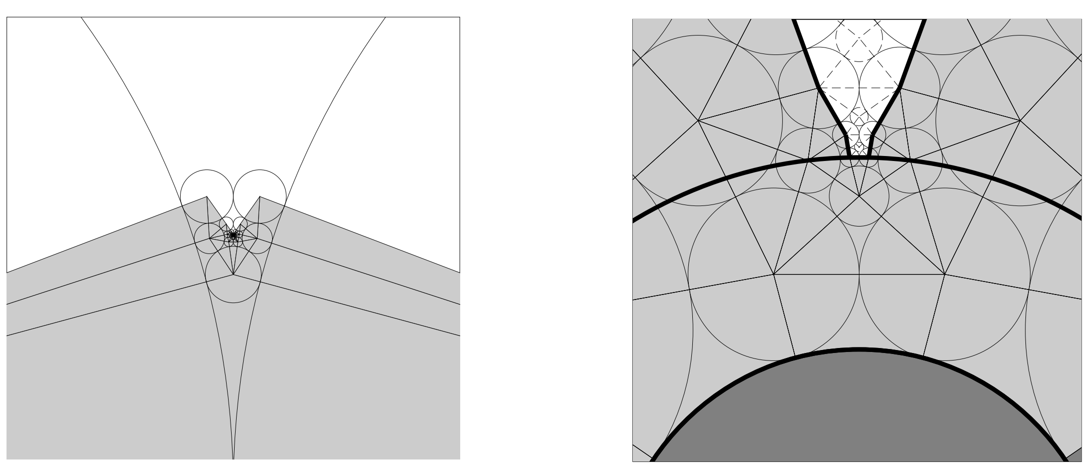

FIGURE 15. Augmenting the complex pushes the bound further, to 3.68.

(bottom) to its counterpart on the right; seven additional vertices are thrown in to triangulate the quadrilateral interstices thus formed. Any univalent packing of the augmented complex yields a univalent packing of K. The packing for this augmented complex is shown in Figure 15, with the new edges and circles appearing as dashed lines. Comparison with Figure 14 (bottom) shows that this technique draws the two sides of the notch much closer together while preserving univalence. By trying various boundary conditions, we have been able to push the value of $f^{\#}(0)$ to approximately 3.68.

The packing of Figure 16 is approaching the extreme situation for the prototype complex using our techniques: Some of its boundary circles overlap one another severely, but the packing remains carrier-univalent. The notch is deep and straight and the outer free boundary circles are huge. Nonetheless, the value of $f^{\#}(0)$ reaches only 3.84, still well within the working bound of $4/\lambda_K \approx 4.88$.

We have experimented with numerous complexes and packings, but have yet to find one for which the hypothesis fails with $\mathcal{C}_2$ equal to the classical constant 4. Large values for $f^{\#}(0)$ seem to arise from Koebe-like functions; attempts to vary from packings mimicking regions with single straight slits have led invariably to smaller values.

Experiment #3: The Error Function

It is evident that any attempt to study infinite circle packings would encounter experimental problems. However, there are also considerable theoretical hurdles. The fundamental issue for an infinite complex K concerns the existence and variety of packings. The entries of the associated infinite packing labels $R = \{\rho_0, \rho_1, \ldots\}$ must satisfy an infinite system of equations—namely, the angle sums of interior vertices must be multiples of 2π. There are currently very few methods for generating solutions. Indeed, the only examples so far are these: Maximal packings for infinite complexes; discrete polynomials created in [Dubejko a]; univalent packings filling simply connected plane domains [He and Schramm]; ad hoc constructions, such as the sine function described in Section 1 (see Figure 3); and "Doyle spirals".

We restrict attention to the hexagonal case, with complex H, where progress is most likely. If P is a circle packing for H, the discrete analytic function $f : P_H \to P$ would be an example of a *discrete entire function.* We know some things, a priori, about such packings P. For instance, by the discrete Picard Theorem [Callahan and Rodin 1993] (which can be extended to include branched packings), P must cover all of $\mathbb{C}$, with the possible exception of one point.

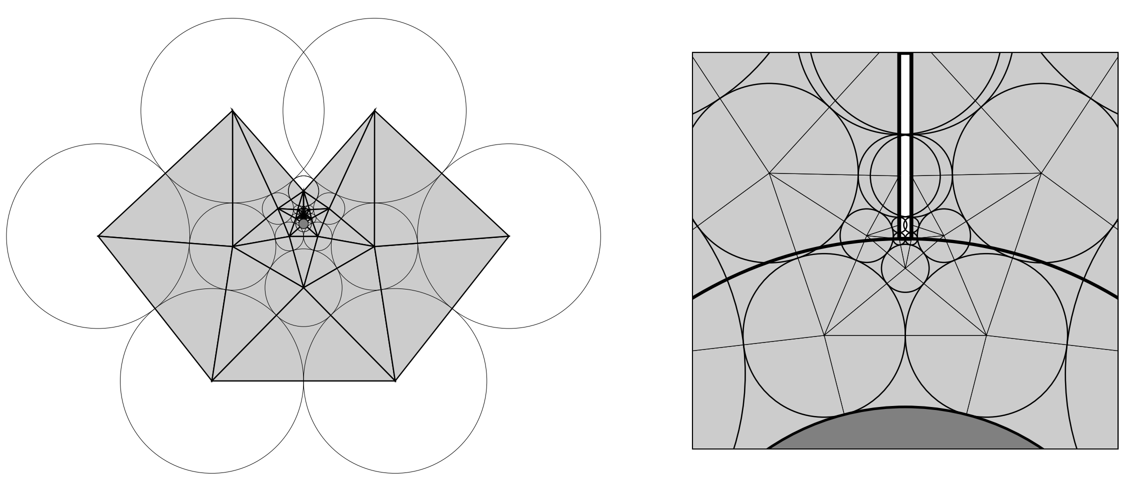

FIGURE 16. Approaching the extremal for the prototype.

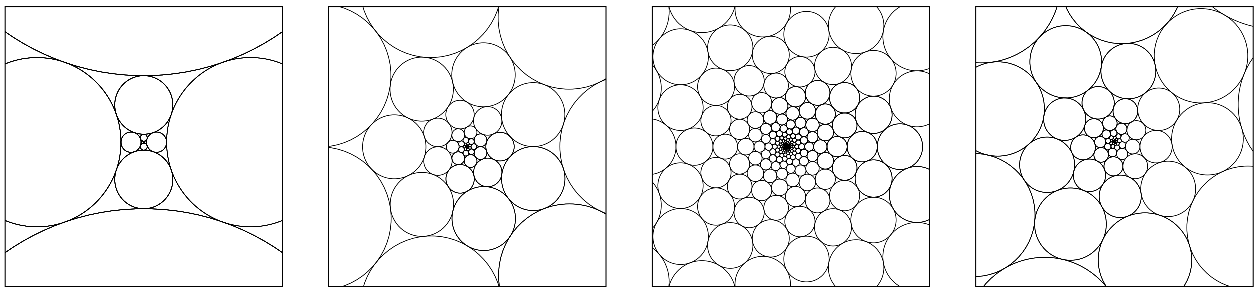

FIGURE 17. A selection of Doyle spirals.

How many packings P are there for H? By Sullivan uniqueness, any univalent P is a Möbius image of P_H; that is, f is a linear polynomial. A recent result [Dubejko 1996] yields the same conclusion if $f^\#$ is bounded. As for nonunivalent packings, Dubejko created finitely branched packings P whose carriers are finite valence (branched) covers of $\mathbb{C}$; these are *discrete polynomials*. An observation of Peter Doyle leads to a two-parameter family of "Doyle spirals" like those shown in Figure 17; the function f in each case behaves as a *discrete exponential* function. There are as yet no other packings known for H, and the search for further examples motivates the experiments in this section.

We begin by looking more closely at the spirals. First, note in Figure 17 that the packings are indeed hexagonal: every circle is tangent to six others. The circles spiral in towards a point, which we take to be $w = 0$, and out towards infinity. These spirals should be pictured as lying on the universal cover of $\mathbb{C}^* = \mathbb{C}\backslash\{0\}$; each circle is actually a projection to $\mathbb{C}^*$ of infinitely many circles on this universal cover. In order to get reasonable illustrations we have chosen pairs of parameters which cause the circles to line up with one another from one sheet to the next. (See [Beardon et al. 1994] for a complete analysis of Doyle spirals.)

Figure 18 shows the domain and range of one of these discrete exponentials f; with the help of a chain of shaded circles, one sees the familiar mapping properties of exponential functions, such as range, periodicity, and local univalence. Moreover, the growth of radii in P is precisely exponential, as one would expect—after all, if f is supposed to be a discrete exponential, then $f^\#$ should be exponential also.

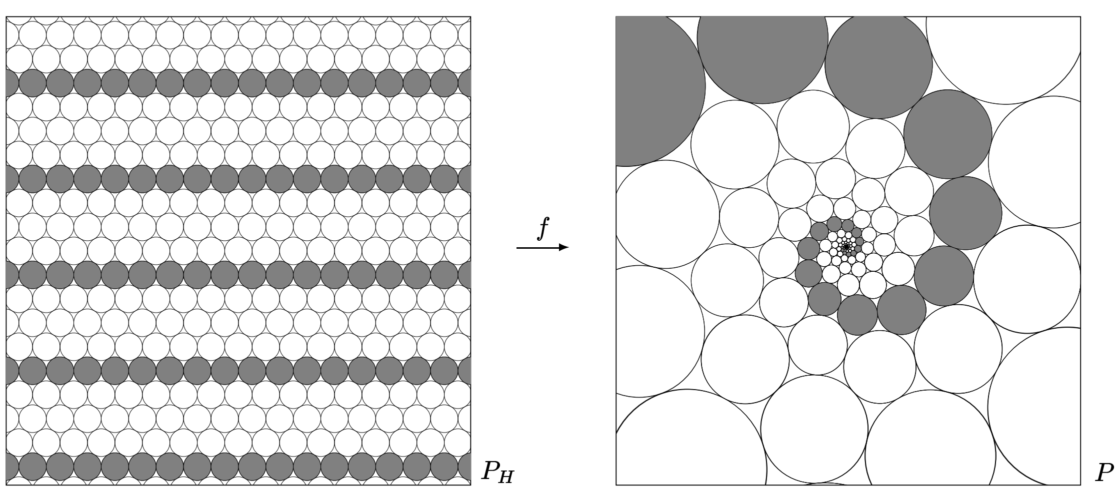

FIGURE 18. A discrete exponential function.

In a private communication, Peter Doyle posed the following question:

Question. *Do there exist any locally univalent circle packings of H other than Möbius images of P_H or Doyle spirals?*

One might be skeptical for reasons involving the distortions in flowers of the packing. Our experiments suggest a more optimistic view, though they are far from conclusive. We attempt to mimic the simplest classical locally univalent entire function after the exponential, namely, the *error function*

$$\mathrm{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2}\, dt.$$

For a classical description of the value distribution theory for this function, see [Nevanlinna 1970, p. 168], for example.

Our aim here is to develop a circle packing P for H whose associated discrete analytic function f mimics erf. An opening is provided by our knowledge of the derivative

$$\mathrm{erf}'(z) = \frac{2}{\sqrt{\pi}} e^{-z^2}.$$

This should tell us something about $f^{\#}$, allowing us to construct f by integration. Here is the experiment:

- For the domain, take P_H to have circles of radius $\varepsilon = 0.07$, with the circle for vertex α centered at the origin and that for γ centered on the positive y-axis.
- For each $n > 1$, let $H^{(n)}$ denote the subcomplex of H formed by vertices in generations $1, 2, \ldots, n$ from α. (The corresponding packings $P_H^{(n)}$ taken from P_H are shown in the first frame of Figure 19, with their boundaries marked, for $n = 5, 7, 9, 11$, and 13.)
- For fixed n, visit each vertex $v_j \in H^{(n)}$, find the center z_j of its circle in P_H (and $P_H^{(n)}$) and compute
$$\rho_j = \left|\mathrm{erf}'(z_j)\right| \varepsilon. \tag{4.1}$$
These values define a label R for $H^{(n)}$, and if they were used as radii for a new packing Q_n of $H^{(n)}$, then the function $f_n : P_H^{(n)} \to Q_n$ would satisfy $f_n^{\#} \equiv |\mathrm{erf}'|$ at all the points z_j, a strong sign that f_n is the restriction to $P_H^{(n)}$ of the function f we're looking for.
- Unfortunately, the label R will not be a packing label for $H^{(n)}$; that is, circles with these radii will not quite fit together. Therefore, we modify our procedure slightly: Compute the ρ_j by (4.1) only for the boundary vertices v_j of $H^{(n)}$. Then solve the Dirichlet problem for the unique interior radii that go with these to form a (unbranched) packing label; call it R_n.
- Let P_n be the resulting normalized circle packing of $H^{(n)}$. The images become complicated as n grows; more about that in a moment. (See Figure 19.)

Repeating the experiment for successive integers n, the hope is that the finite packings P_n will converge geometrically to an infinite packing P; equivalently, that the packing labels on the $H^{(n)}$ will converge to a packing label on H. Presumably, if P does not degenerate to a regular hexagonal packing, one could then show that P behaves like the error function—certainly it would not be a Doyle spiral.

Seeing the packings evolve graphically is fascinating in several regards:

(1) From the very first image one sees that the evolving mapping behavior of the f_n's closely mimics that of the error function. The image develops two logarithmic branch values at approximately $w = \pm 1$, the two asymptotic values of erf. As new generations of circles are added, the growing image flows out and around the singularities. The static images in Figure 19 quickly become too complicated to interpret; however, chains of circles that seem to "wrap" around one singularity begin to cover the other singularity (Picard's Theorem is not in danger). When running CirclePack, one can selectively highlight portions of the packings to

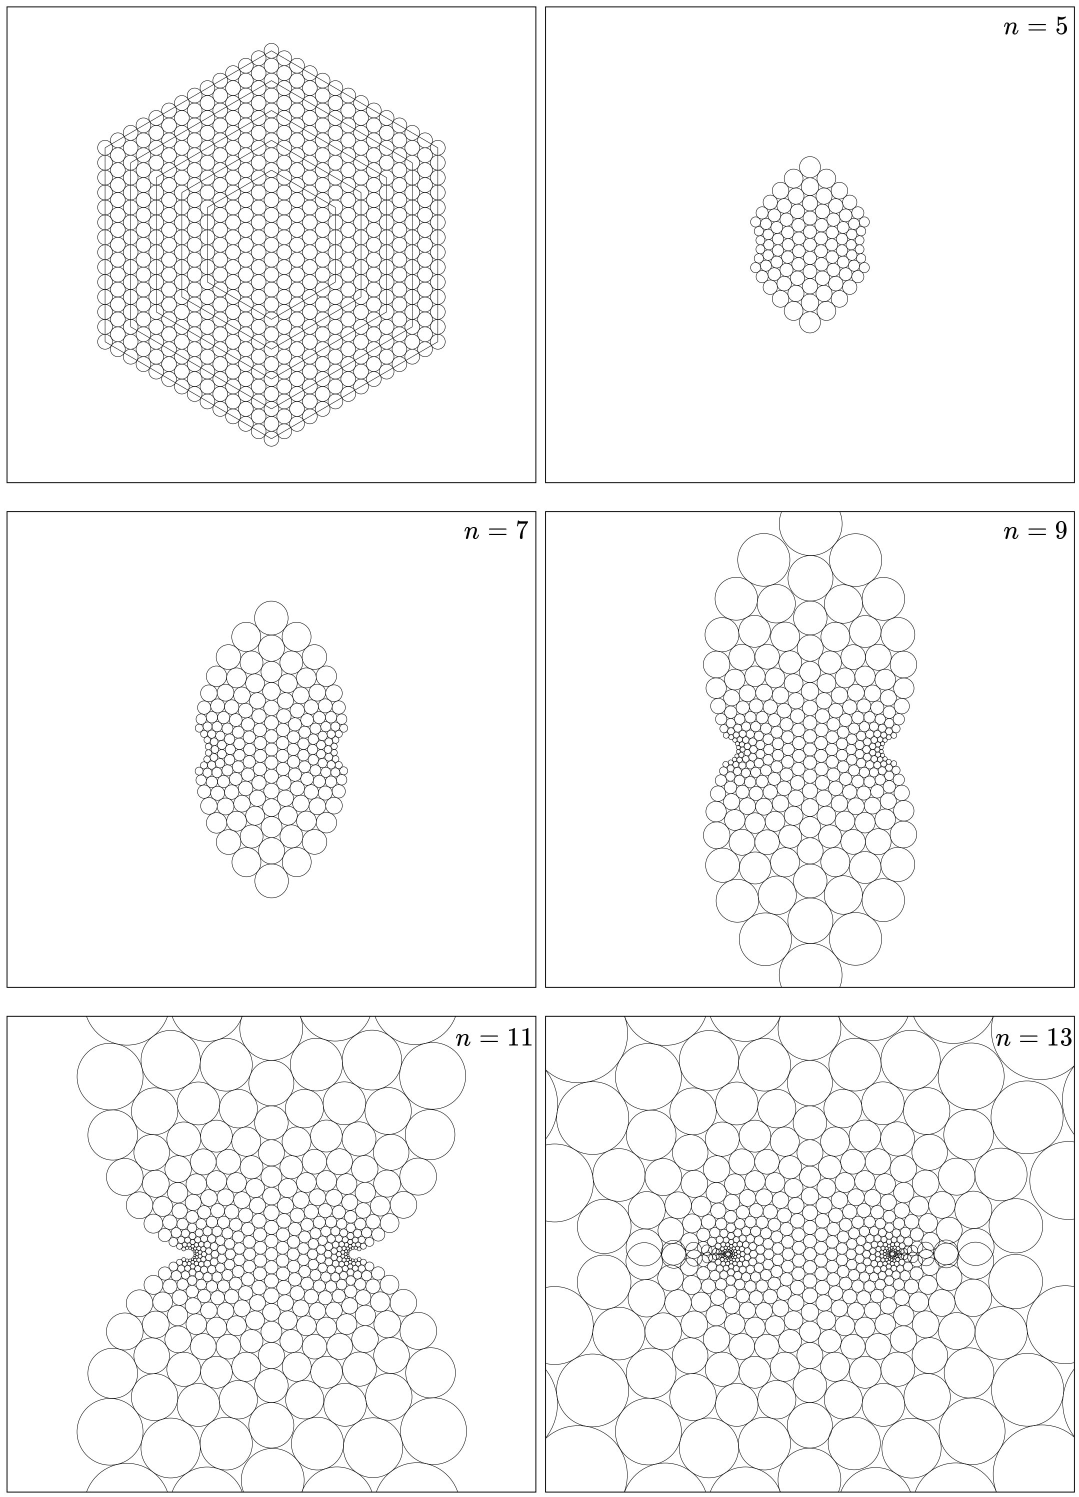

FIGURE 19. Towards a discrete error function. The first panel shows the packings $P_H^{(n)}$, for $n = 5, 7, 9, 11$, and 13; these packings are obtained by restricting P_H according to the boundaries shown. The remaining panels show P_n for the same values of n.

see the mapping behavior dynamically. It precisely models the well-known behavior of erf.

(2) The speed with which CirclePack computes successive packing labels is surprising. Recall that upon adding a new generation (going from $H^{(n)}$ to $H^{(n+1)}$), the new boundary radii are determined by (4.1), then associated interior radii are computed. When the second author first ran these experiments, the packing radii from P_n were retained as the initial guess for the interior radii of P_{n+1}. After adding the new boundary, it was found that no adjustments were needed in these interior radii: the radii from each stage seemed to be correct for the successive stage. This persisted for 25 new generations! The implication was exciting: perhaps the P_n are nested pieces of an infinite circle packing. On the 26th generation, however, the recomputation mechanism of CirclePack kicked in, showing that the earlier data had simply been within acceptable error tolerances.

(3) Nonetheless, the stability of the successive packings is impressive, as can be judged in these trials by watching one particular circle, say that at the origin, as n grows. Does its radius have a positive and finite limit? For large n, computational times and roundoff errors become problems. For instance, flowers of P_n can become quite distorted as one moves out in the generations from α because the ratios between the radii of neighbors will grow without bound. Figure 20 shows a central circle and its six neighbors; the radii of the petals vary enough to cause nonneighboring petals to overlap.

Aside from maximal packings and Doyle spirals, there are as yet no known locally univalent discrete entire functions based on any complex K, hexagonal or otherwise. Ultimately, experiments can only suggest whether an approach has possibilities and perhaps give clues to a rigorous solution. At a minimum, however, these trials mark the discrete error function as a frontrunner in addressing Doyle's question.

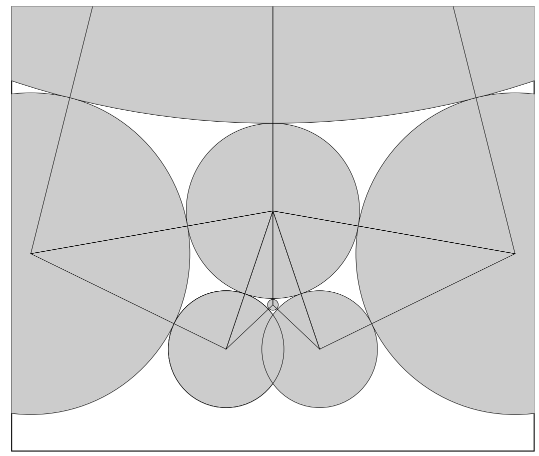

FIGURE 20. Distorted hex flower.

5. DISCRETE APPROXIMATION

The four experiments described in this section concern approximation in the vein of the Theorem at the end of Section 3; that is, we exploit the propensity of discrete analytic functions to converge to classical ones as their underlying circle packings become “finer” (contain more numerous and smaller circles).

To clarify the shift in emphasis, return for a moment to Experiment #3 on a discrete error function: We employed a parameter ε that remained fixed in those trials. Were we to make it successively smaller, the procedures there could produce a sequence of discrete analytic functions on finite subcomplexes of H that would converge uniformly on compacts of $\mathbb{C}$ to erf. Our goal before, however, was a single infinite packing whose associated map mimics erf, rather than a succession of finite packings whose maps approximate it.

Parallels will probably continue to emerge, but our aim now shifts to the generation of sequences of discrete analytic functions approximating classical ones. The next three experiments involve approximation of functions from important families: finite Blaschke products, polynomials, and disc algebra functions. The last experiment suggests an approach to speeding up circle packing computations.

Experiment #4: Discrete Blaschke Products

The analytic functions known as finite Blaschke products arise in several mathematical contexts, perhaps because they enjoy such a variety of characterizations: n-fold Blaschke products are precisely the n-valent proper (analytic) maps of $\mathbb{D}$ onto itself. Within the important class of "inner" functions in $\mathbb{D}$, these are the finite valence ones and the only ones that are continuous up to $\partial\mathbb{D}$. They play a role in iteration theory because they restrict to locally one-to-one maps of the unit circle to itself. They arise frequently in applications because, after a change of variables, they comprise precisely the n-degree rational mappings of the left half-plane onto itself.

In any serious discrete analytic function theory, one would certainly hope for analogues of finite Blaschke products. As it happens, circle packing versions are particularly easy to generate and are geometrically very faithful to their classical models. In particular, as in the classical case, they seem to provide the extremal functions for a variety of situations. Among these is the Dieudonné Schwarz Lemma from Section 4, which we will revisit in a moment.

The classical representation of a finite Blaschke product B depends on its zeros (counting multiplicities) and takes the form

$$B(z) = e^{i\theta} z^k \prod_{j=1}^{m} \frac{|a_j|}{a_j}\, \frac{a_j - z}{1 - \bar{a}_j z},$$

where $\theta \in \mathbb{R}$, k and m are nonnegative integers, and $\{a_1, \dots, a_m\}$ are nonzero points of $\mathbb{D}$. The function B is called an n-fold Blaschke product and the points a_j (along with the origin, if $k \geq 1$) are its zeros. Note that the zeros determine B, up to a rotation.

For the discrete setting, the product construction is not available—indeed, we have no complex arithmetic! However, there is an alternative: It is well known that B has $n-1$ branch points (counting multiplicities) and that these points determine B up to composition with an automorphism of $\mathbb{D}$. Branching is a geometric feature and, as we saw in Section 3, is available in the discrete setting. In fact, this provides a pleasing counterpoint to the classical theory, for we know of no practical classical methods for constructing Blaschke products from their branch sets.

Given K, a discrete n-fold finite Blaschke product is a discrete analytic function $b : P_K \to P$, wherein P is a packing of K in $\mathbb{D}$ having $n-1$ branch points (counting multiplicities) and having all boundary circles of infinite (hyperbolic) radius. The procedure for constructing such packings P (in our examples we use only simple branch points) goes as follows: Choose a set of $n-1$ distinct interior vertices $\beta = \{v_1, \dots, v_{n-1}\}$ forming a branch structure as described in Section 3. Set the "aims" (target angle sums) of the vertices v_j to 4π; the other interior vertices have their usual aims of 2π. Set the boundary radii of P to (hyperbolic) infinity. Repack using CirclePack. The existence and essential uniqueness of the resulting packing were established in [Dubejko 1995].

Our prototype experiments involve fourfold discrete Blaschke products based on hexagonal complexes K. We assume that the desired branch set consists of

$$z_1 = 0.15 - 0.25i,$$
$$z_2 = -0.15 - 0.25i,$$
$$z_3 = 0.5i.$$

The maximal packing P_K for a small hexagonal complex is displayed in Figure 21 (left). No circles are centered precisely at the z_j, so we have chosen three nearby ones as our branch circles, and shaded them darkly in Figure 21. The image packing P in the same figure, created using CirclePack, is difficult to interpret because its circles are four layers deep. Circles corresponding to boundary vertices of K are horocycles in both P_K and P: this is the discrete version of the requirement that Blaschke products map the unit circle to itself. With diligence, one may be able to trace these boundary circles in P as they wind four times about the inside of the unit circle. (While running CirclePack,

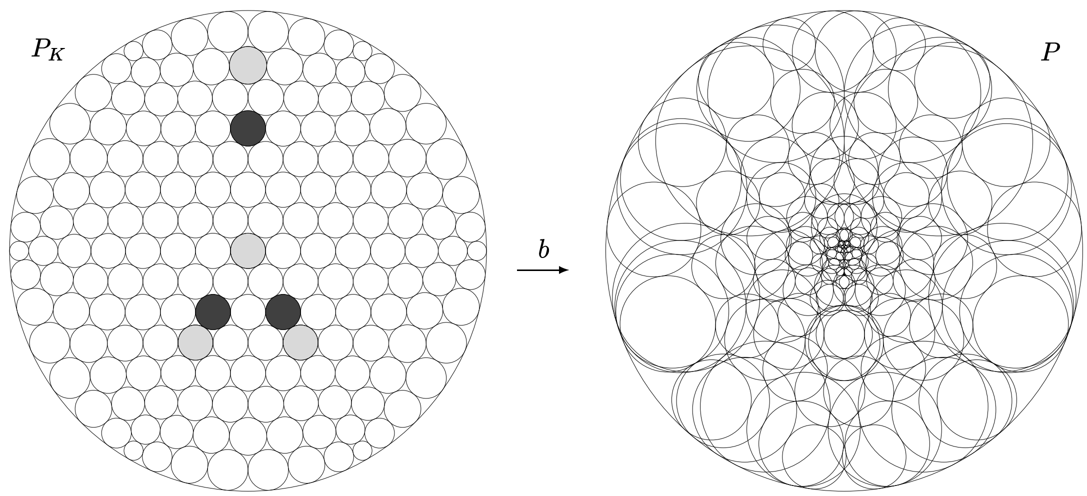

FIGURE 21. A fourfold Blaschke product. The dark circles in the domain are the branch circles, and the lightly shaded ones are those whose image in P contains 0 (they are easy to pick out using CirclePack).

one can graph the domain and range in various ways to better decipher the mapping properties.)

A larger example is illustrated in Figure 22. The packing P_K is much finer (1657 circles) and as before we have chosen three branch circles (darkly shaded) near z_1, z_2, and z_3. In the range, the circles are again four layers deep; we display only the boundary circles, which wrap four times around, and the three branch circles (the small points near the origin). It was recently shown [Dubejko 1995] that discrete Blaschke products approximate classical ones. To test this, we estimated (using euclidean barycentric coordinates) the zeros for the b of Figure 22, used them to construct a classical Blaschke product B, and computed the zeros of B'. We get

$$Z_1 \approx 0.138 - 0.245i,$$
$$Z_2 \approx -0.138 - 0.245i,$$
$$Z_3 \approx 0.000 + 0.488i,$$

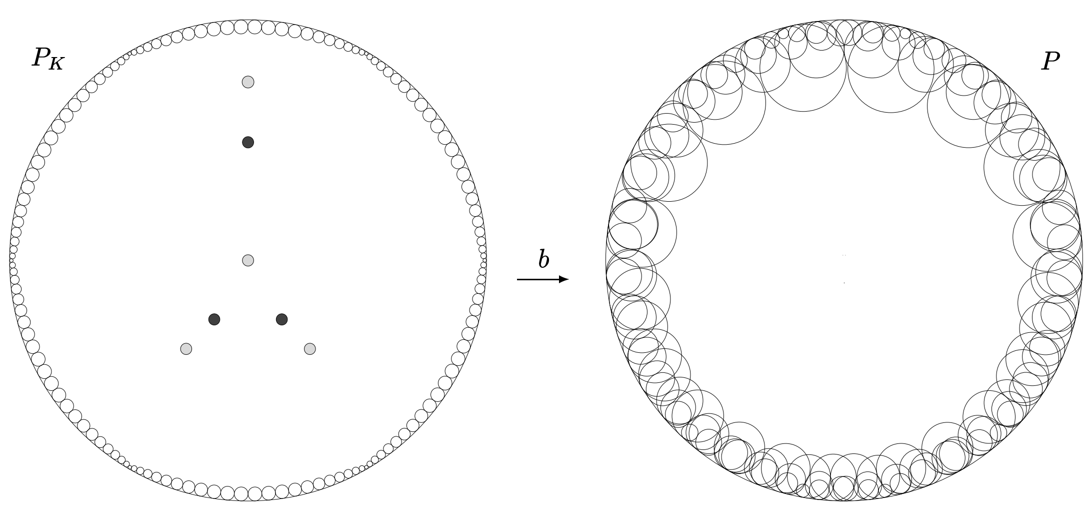

FIGURE 22. A finer fourfold Blaschke product. The shaded circles have the same meaning as in Figure 21, but otherwise only the boundary circles are shown.

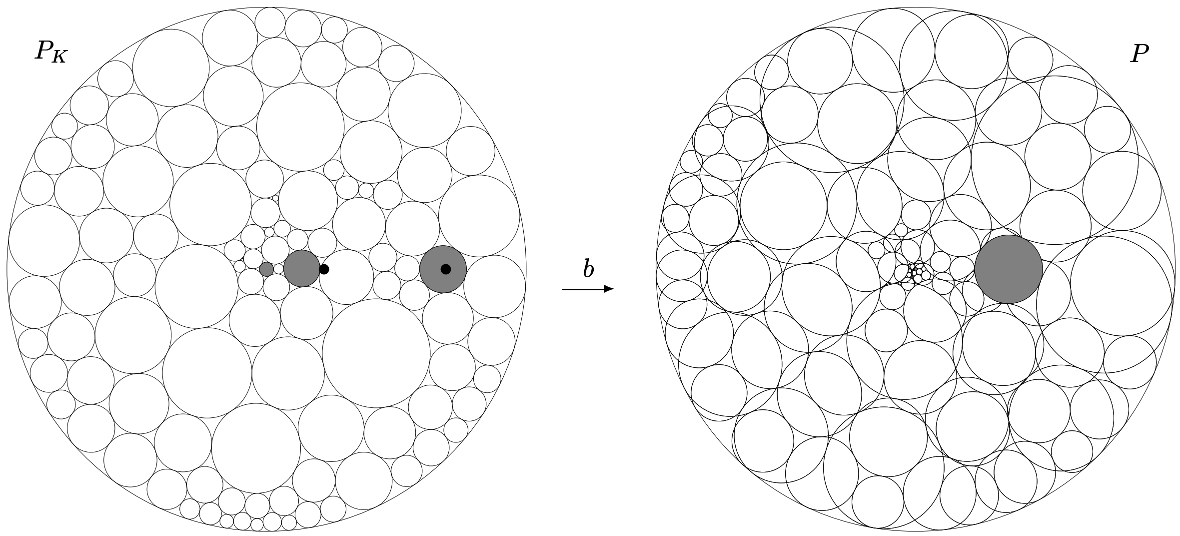

FIGURE 23. Extremal discrete Blaschke product. The dots in the domain represent x and y (nearer the origin).

values that are in fair agreement with the prescribed branch points z_1, z_2, and z_3 of b.

Let's apply finite Blaschke products in a further probe of the Dieudonné–Schwarz Lemma of Experiment #1 (where we already encountered a discrete finite Blaschke product in Figure 12). Now we are interested in that portion of Dieudonné's result that applies for $\sqrt{2}-1 < |z| < 1$. His extremal functions are twofold Blaschke products. In particular, suppose for convenience that $z = x$, a real number with $\sqrt{2}-1 < x < 1$. Then the Blaschke product

$$B(z) = z\frac{z+h}{1+hz} \quad \text{with} \quad h = \frac{1-3x^2}{x(1+x^2)}$$

will achieve precisely Dieudonné's bound, namely

$$B'(x) = M_x = \frac{(1+x^2)^2}{4x(1-x^2)}.$$

A computation [Carathéodory 1960, § 291] shows that the (unique) critical value of B is

$$y = \frac{\sqrt{1-h^2}-1}{h}.$$

A prototype experiment in the discrete setting is illustrated in Figure 23. We begin with a more or less random maximal packing P_K and proceed as follows: First, choose a circle of P_K centered at a real number x satisfying $\sqrt{2}-1 < x < 1$ (the outer dot in Figure 23). Observing that in this instance $x \approx 0.69032$, we see that Dieudonné's classical bound on derivatives at x is $M_x \approx 1.50834$. Computation also gives the auxiliary values $h \approx -0.42150$ and $y \approx 0.22105$ (the inner dot in Figure 23) associated with x.

We cannot precisely mimic the classical extremal function for this x because P_K has no circle centered at y. However, we can choose a nearby branch circle. The resulting map $b : P_K \to P$ is a discrete twofold Blaschke product. CirclePack tells us that $b^{\#}(x) \approx 1.4522$, which is within the classical bound M_x. Further experiments with this complex K suggest that b is indeed the discrete extremal for the ratio function at x—for instance, moving the branch point to various other circles results in smaller ratio functions—so Dieudonné's bound seems to hold there. Of course, there are other points x and other complexes to check, and we have uncovered instances in which Dieudonné's bound is violated by small margins. Nonetheless, the preliminary evidence is that the qualitative features in the discrete setting mimic the classical case and that there are likely to be discrete bounds on ratio functions that are close to the classical ones.

Experiment #5: Discrete Polynomials

Dubejko [1995; a] investigated infinite branched circle packings. In particular, he proved the existence of a full set of discrete polynomials and showed that any classical complex polynomial can be approximated by these discrete ones uniformly on compacts of $\mathbb{C}$. Indeed, this result implies that a wide variety of analytic functions on many domains are subject to such approximation. However, since discrete polynomials involve infinite circle packings, this has little practical value. In what follows we will describe a constructive method for approximation that uses finite branched hexagonal packings like those of Experiment #4 above.

Let $F : \mathbb{C} \to \mathbb{C}$ be a complex polynomial. We assume that the critical set of F, which we denote $\operatorname{br} F = \{t_1, \dots, t_m\}$ (with repeats to account for multiplicities), is known. We further assume, for convenience, that $F(0) = 0$ and $F(1) = 1$. We construct discrete approximants f_n of F that are defined on finite portions of the hexagonal complex H. As we go along, we will illustrate with a simple prototype experiment:

$$\text{Prototype: } F(z) = \frac{z^3 - 3z}{-2} \text{ with } \operatorname{br} F = \{1, -1\}.$$

First, let's establish notation: Let P_H denote the maximal packing for H, normalized to give all circles radius $\frac{1}{2}$ and to center circles at $z = 0$ and at $z = 1$. For positive integers n, define $D(n) = \{|z| < n\}$; scale P_H by the factor $1/n$ and let Q_n consist of the circles of P_H/n lying inside $D(n)$. Write K_n for the complex of Q_n. Figure 24 illustrates the prototype case $n = 3$.

For n sufficiently large, $\operatorname{br} F$ will lie in $D(n)$ and we may choose a set $V_n = \{v_1(n), \dots, v_m(n)\}$ of distinct vertices of K_n that forms a branch structure for K_n. Our strategy is to choose V_n in such a way that the centers $z_j(n)$ of the corresponding circles of Q_n approximate $\operatorname{br} F$. More precisely, we want

$$z_j(n) \to t_j \text{ as } n \to \infty \text{ for each } j = 1, 2, \dots, m.$$

(The specific conditions on branch structures were described earlier; these conditions and the one just given are very easily arranged if n is sufficiently large.)

For each (sufficiently large) n, construct a new packing P_n for K_n as was done in the previous experiment. That is, assign angle sum 4π to the vertices of the set $V_n \subset K_n$, assign infinite hyperbolic radii to the boundary vertices of K_n, and

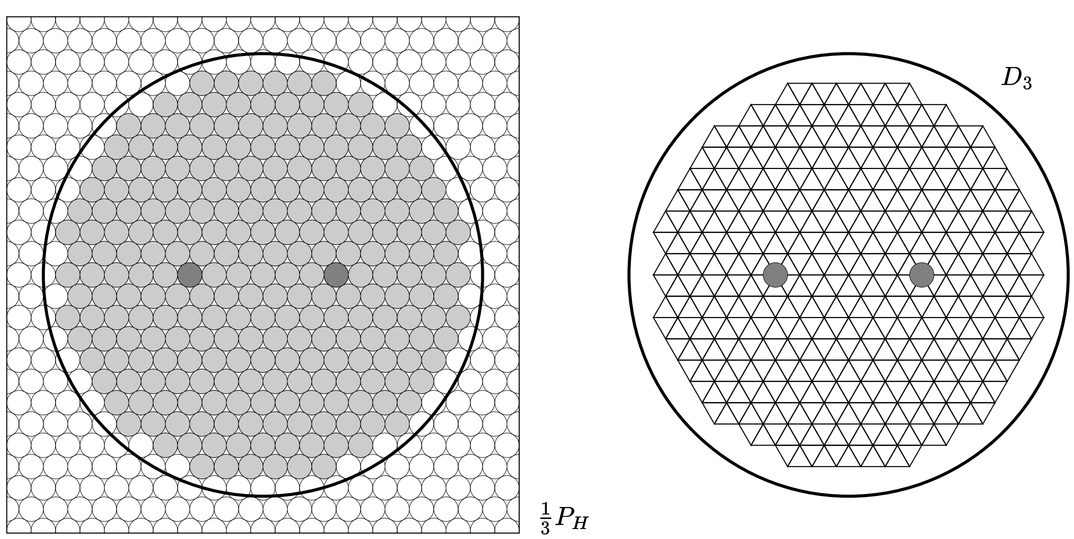

FIGURE 24. Q_3 (left, shaded region) and its complex K_3 (right).

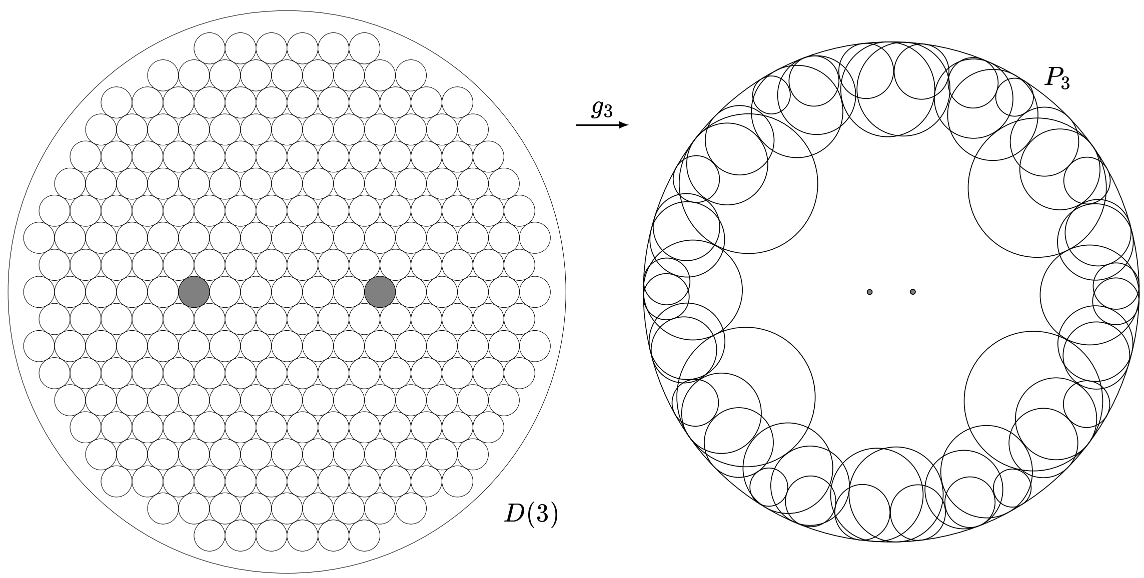

FIGURE 25. The intermediate function g_3. The range packing P_3 is 3 circles deep, so only the boundary circles and the (extremely small) branch circles are shown.

compute the resulting branched packing $P_n \subset \mathbb{D}$. Let $g_n : Q_n \to P_n$ denote the discrete analytic function; the prototype g_3 is illustrated in Figure 25. The branch vertices $v_1(3)$ and $v_2(3)$ are associated with the shaded circles in the domain; these are the ones centered nearest (in fact, at) ± 1.

The construction of P_n is carried out in the unit disc because the Dirichlet problem is conceptually and numerically easy to solve in hyperbolic geometry. However, the function f_n approximating F requires one final normalization: Define $f_n = \lambda g_n$, where the complex number λ is chosen so that $f_n(1) = 1$. In other words, $f_n : Q_n \to \lambda P_n$.

As n goes to infinity, the domains of the f_n exhaust $\mathbb{C}$ and the branch points of the f_n converge to the branch points of F. All the while, $f_n(0) = 0$ and $f_n(1) = 1$. In separate work, the second author has proved what the experiments suggest; namely, that the f_n will converge uniformly on compacts of $\mathbb{C}$ to F. Figure 26 is obtained from f_3 and f_4 in our prototype case: The domains Q_3 and Q_4 are shown with the branch circles shaded; the image packings are too confusing for static display, but we have shown the images under f_3 and f_4 of a representative closed curve σ, a regular hexagon. One can compare these images to $F(\sigma)$ (the dashed curves in Figure 26) to get some feel for the accuracy of the approximations.

The constructions of f_n for even larger values of n are relatively easy in CirclePack and the behavior of F emerges fairly rapidly. It is clear, however, that one would not choose this method to approximate polynomials in any practical situation. Nonetheless, it does work, and in slightly altered circumstances, circle packing might bring something new to the table. The next experiment may be a case in point.

Experiment #6: Prescribed Boundary Curves

Throughout this discussion, γ will denote a closed plane curve that intersects itself only transversely, if at all, and that has no triple points. We say that γ is admissible if it is the image of the unit circle under a function $F = F_\gamma$ in the disc algebra $A(\mathbb{D})$, that is, a function F that is analytic in $\mathbb{D}$ and continuous on $\bar{\mathbb{D}}$. Marx [1974] characterized admissible curves by their winding numbers and crossing behavior. In this section we describe a circle packing method for approximating F. This is an extension of the scheme of Thurston, which

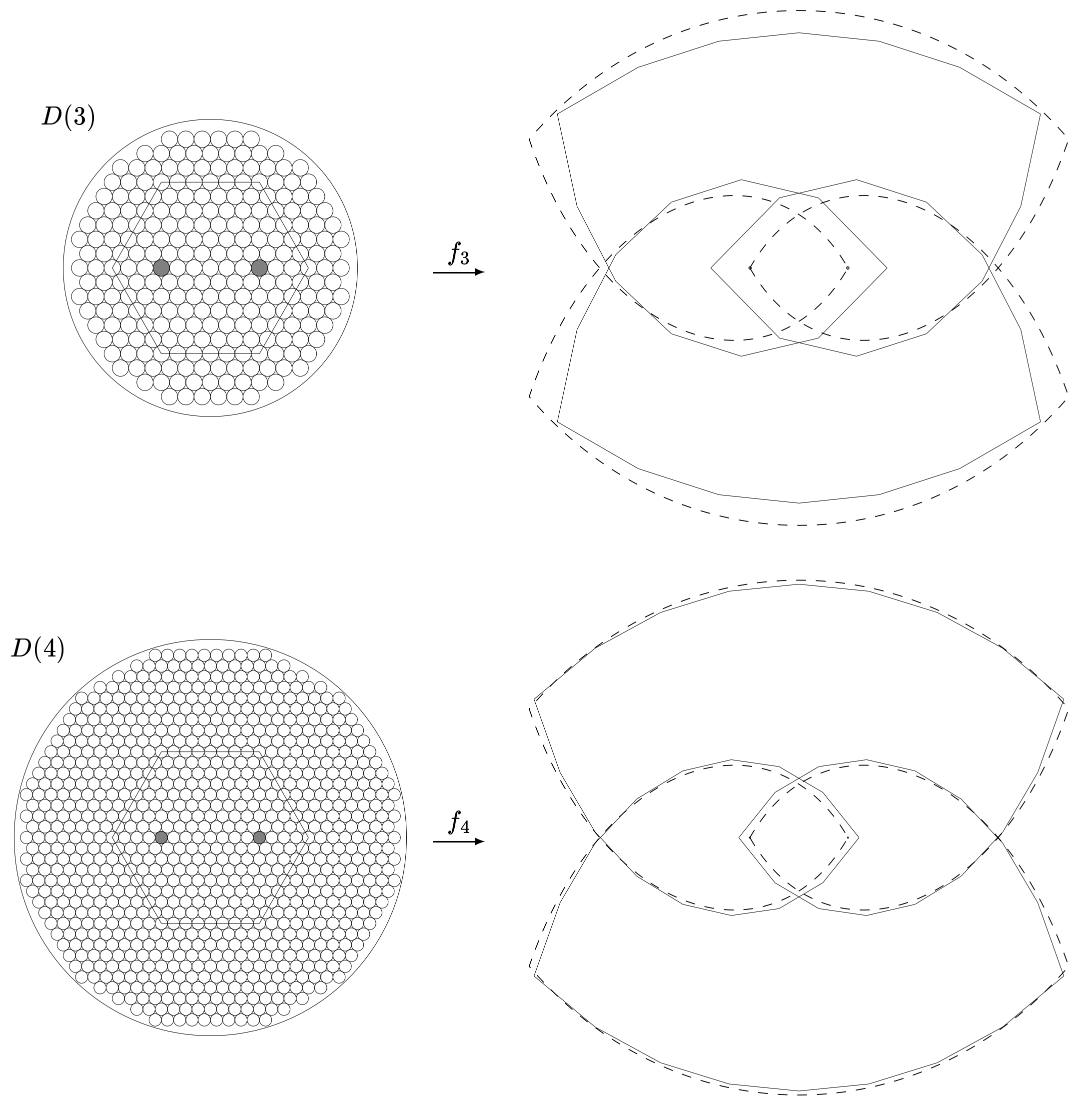

FIGURE 26. The approximating functions f_3 and f_4.

applies when γ is a Jordan curve (and hence automatically admissible). However, in general the function F will necessarily have branch points; our method permits specification of the branch sets so that we may approximate, uniformly on compacts of $\mathbb{D}$, any desired classical solution.

For our prototype experiment, let γ be the oriented curve of Figure 27 (left). This was chosen because it is evidently admissible—in fact, the image Riemann surfaces of the solutions in $A(\mathbb{D})$ are fairly easy to visualize. One of the legal branch sets consists of a simple branch point with image w_1 and a double branch point with image w_2, and this is the particular solution F_γ we choose to approximate. For the sake of variety, we will use the infinite ball-bearing complex B in place of hexagonal combinatorics.

Write Q_ε for the maximal packing of B scaled so that the larger circles have radius ε. Overlay γ with Q_ε as shown in Figure 27. Let a_1 and a_2 designate two smaller circles nearest w_1 and w_2. We build a range circle packing $P = P_\varepsilon$ whose carrier

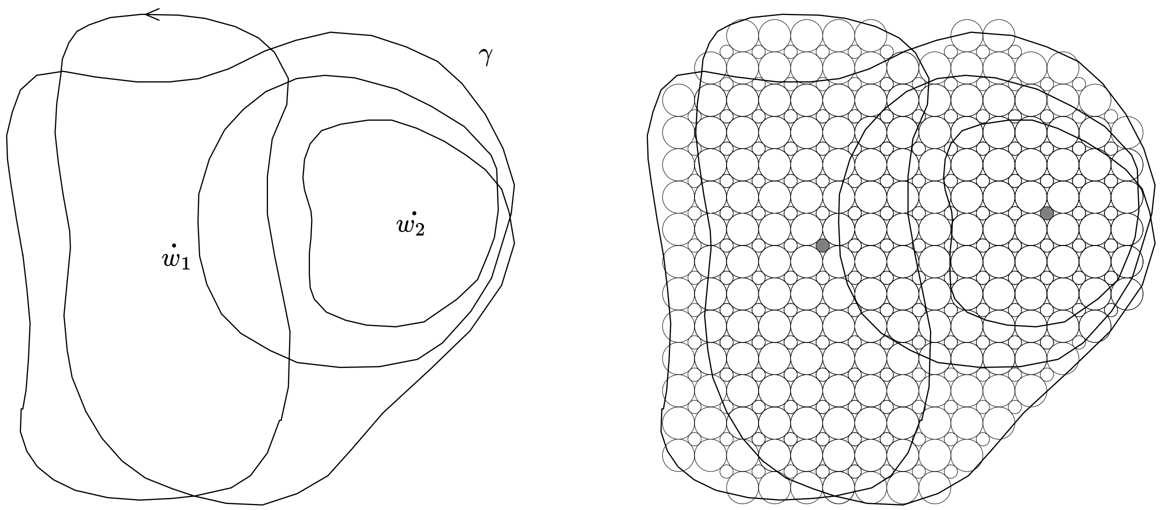

FIGURE 27. An admissible curve γ, and the portion of the ball-bearing packing it encloses. The shaded circles a_1 and a_2 are those nearest the points w_1 and w_2.

approximates the image Riemann surface of f by pasting together various pieces of Q_ε. The carrier pieces and their edge identifications are displayed in Figure 28. We omit the details because the analogous constructions are quite standard in geometric function theory. We point out, however, that the resulting geometric complex G is a simply connected branched covering of a portion of the plane.

The circles associated with the faces of G form P, which may be pictured as a branched covering over a portion of Q_ε (some circles have as many as three circles of P lying over them). Note that although P has only one circle lying over each of a_1 and a_2, the pastings give eight petals for the flower of a_1 and twelve for that of a_2; this is where the branching occurs.

The complex K for P is simplicially equivalent to G, hence is simply connected. Let the maximal packing P_K in $\mathbb{D}$ be normalized in the usual way. The map $f = f_\varepsilon : P_K \to P$ is the discrete analytic function that approximates F. In Figure 29, the circles of P_K and P associated with the branching are shaded for reference.

As ε goes to zero, one can show that the simplicial maps associated with the f_ε converge uniformly on compacts of $\mathbb{D}$ to an analytic function F mapping the unit circle to γ and having the

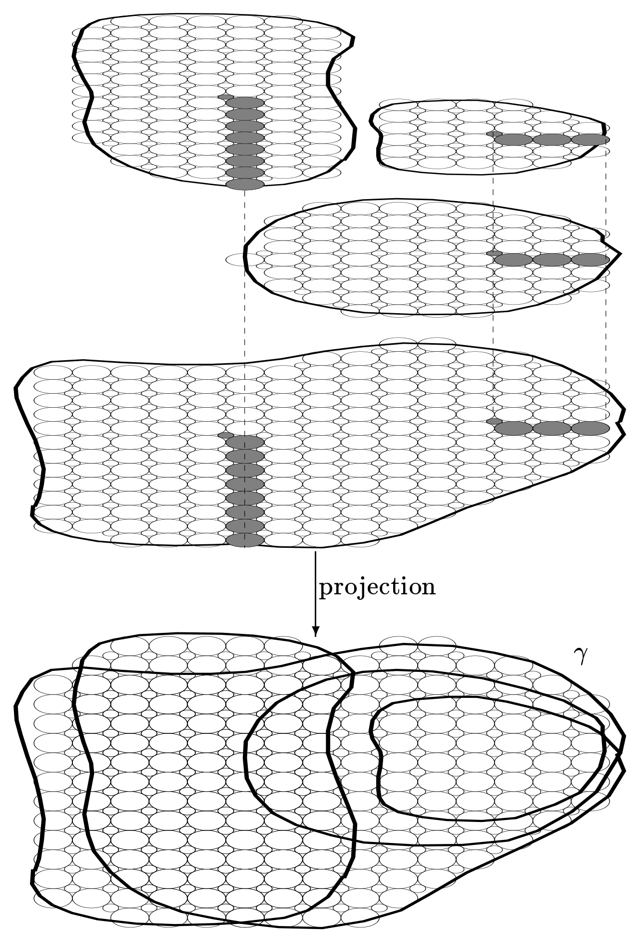

FIGURE 28. Constructing the range packing.

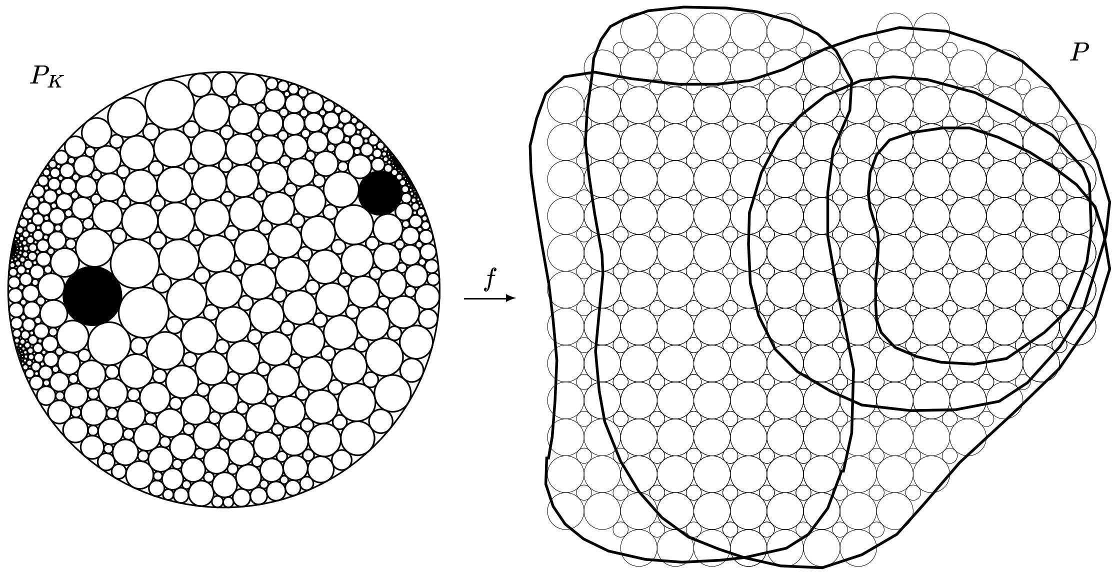

FIGURE 29. The approximating map.

specified branch values w_1, w_2. This, of course, is the desired map F_γ. The proof techniques are basically those of the conformal setting, though one must work outside small neighborhoods of the branch points and fill those in as removable singularities at the last [Dubejko 1995].

This construction process could be completely automated; Chad Sprouse, a 1993 Summer REU student of the second author, developed a very clever scheme for carrying out the cut/paste operations. This has not yet been implemented, but the eventual goal would be to generate the f_ε by providing nothing more than a parameterization of γ and a choice of branch points.

Experiment #7: Multigrid Packing Algorithms

Rodin and Sullivan's proof [1987] of Thurston's conjecture immediately raised questions about the potential of circle packing as a practical numerical technique. It is not yet competitive in terms of accuracy and speed with more classical numerical methods, such as implementations of Schwarz–Christoffel or more recent zipper methods of [Marshall]. Nonetheless, circle packing may ultimately have advantages on certain parallel computer architectures, such as massively parallel machines; for multiply connected regions; for regions undergoing dynamical change; or even as a preprocessor feeding into a classical method. Also, as our earlier experiments show, circle packing extends to branched and multivalent situations, and even perhaps to maps between Riemannian surfaces. (Certainly, in an increasingly "quantum" world, there is some place for quantum complex analysis!)

The algorithm currently used in CirclePack for computing (approximate) packing labels is the iterative method of Thurston described at the end of Section 2—it is roughly analogous to the "method of relaxation" (or Gauss–Seidel) for the Dirichlet problem in classical discrete potential theory. But, whereas this latter problem permits much more efficient algebraic methods, circle packing problems are highly nonlinear and no replacement for relaxation has been found.

A potential advantage in circle packing, however, lies in the faithful way in which discrete analytic functions on even coarse packings mimic classical functions. We exploit that here to implement a multigrid method for computing approximate conformal mappings. We describe it by means of a relatively simple prototype experiment and then give some results in more substantial trials. The same technique can be applied in nonunivalent and branched setting, such as those discussed in Experiments #4–#5.

Our goal in this section is to approximate the classical conformal mapping F from $\mathbb{D}$ onto a given simply connected domain Ω in the plane, $F : \mathbb{D} \to \Omega$. The approach was outlined in Section 1 and illustrated there by Figure 1: one fills Ω with a portion P of some regular hexagonal circle packing and computes the associated maximal packing P_K; the desired approximant is the piecewise affine map associated with $f : P_K \to P$. Using successively finer packings improves the approximation.

At each stage, the packing P and its complex K are easily determined; the real work enters in computing the maximal packing label R_K, which easily yields P_K. We assume hereafter that Ω lies in $\mathbb{D}$ and carry out all computations in hyperbolic geometry. Thurston's algorithm proceeds from an initial guess R_I to R_K using iterative relaxation.

The direct method generates R_I by simply assigning to interior vertices the current (hyperbolic) radii of their circles in $P \subset \Omega \subset \mathbb{D}$ and assigning infinite radii to boundary vertices, since we know that they will be horocycles in P_K. CirclePack then cycles repeatedly through the list of interior vertices, making adjustments until it reaches an approximation to R_K within assigned tolerance. Of course, most radii change considerably from their values in R_I, and the direct method may converge relatively slowly. Our aim is to improve efficiency by making a more informed choice of the initial label R_I.

As its name suggests, the multigrid method attempts to use information gleaned from coarser circle packings to shorten computations on finer ones. Two stages of the prototype experiment are illustrated in Figure 30.

Assuming that $f_1 : P_{K_1} \to P_1$ is known, we want to compute $f_2 : P_{K_2} \to P_2$. In other words, P_{K_1}, P_1, and P_2 are known, and we seek P_{K_2}, the maximal packing at the finer stage. Our opening is provided by the ratio functions: Since f_1 approximates F (albeit not very well), its ratio function $f_1^\#$ approximates $|F'|$. As $f_2^\#$ will also approximate $|F'|$, we see that $f_1^\#$ gives some preliminary information about the (as yet unknown) f_2.

Strategy. Use the ratio function for a known approximation f_1 to set the initial radii for computation of a finer approximation f_2.

Here's the central idea: Suppose v is an interior vertex of K_2 whose associated circles are $C \in P_{K_2}$ (unknown) and $c \in P_2$. Choose a vertex $v' \in K_1$ with associated circles $C' \in P_{K_1}$ and $c' \in P_1$ so that c' is as close as possible to c in Ω; see Figure 30. The radius of C, $\operatorname{rad} C$, is unknown. We expect C to be near C' in $\mathbb{D}$, but more importantly, we expect $f_2^\#(v)$ to approximate $f_1^\#(v')$:

$$f_2^\#(v) \approx f_1^\#(v') \Longrightarrow \frac{\operatorname{rad} c}{\operatorname{rad} C} \approx \frac{\operatorname{rad} c'}{\operatorname{rad} C'}.$$

Note that since we are working in $\mathbb{D}$, the appropriate radii and ratio functions are all *hyperbolic*. We conclude from the implication above that

$$\operatorname{rad} C \approx r = \frac{\operatorname{rad} c}{(\operatorname{rad} c')/(\operatorname{rad} C')} = \frac{\operatorname{rad} c}{f_1^\#(v')}. \tag{5.1}$$

We take r to be our initial guess for the label at v: that is, we set $R_I(v) = r$. If we do likewise for all interior vertices v of K_2 and set the label to ∞ for all boundary vertices, we have our initial label R_I. CirclePack then takes over to compute R_{K_2}, and thereby P_{K_2}.

In our practical implementation of this idea, we employ hexagonal packings, arranged so that at each successive stage the radii are cut in half. For instance, Figure 31 superimposes the packings P_1 and P_2 of Figure 30. Carr P_2 is actually a refinement of carr P_1, and each vertex of K_1 may be identified with one from K_2.

Here, then, is the specific procedure for choosing R_I for K_2:

- If v is a boundary vertex of K_2, then obviously set $R_I(v) = \infty$.

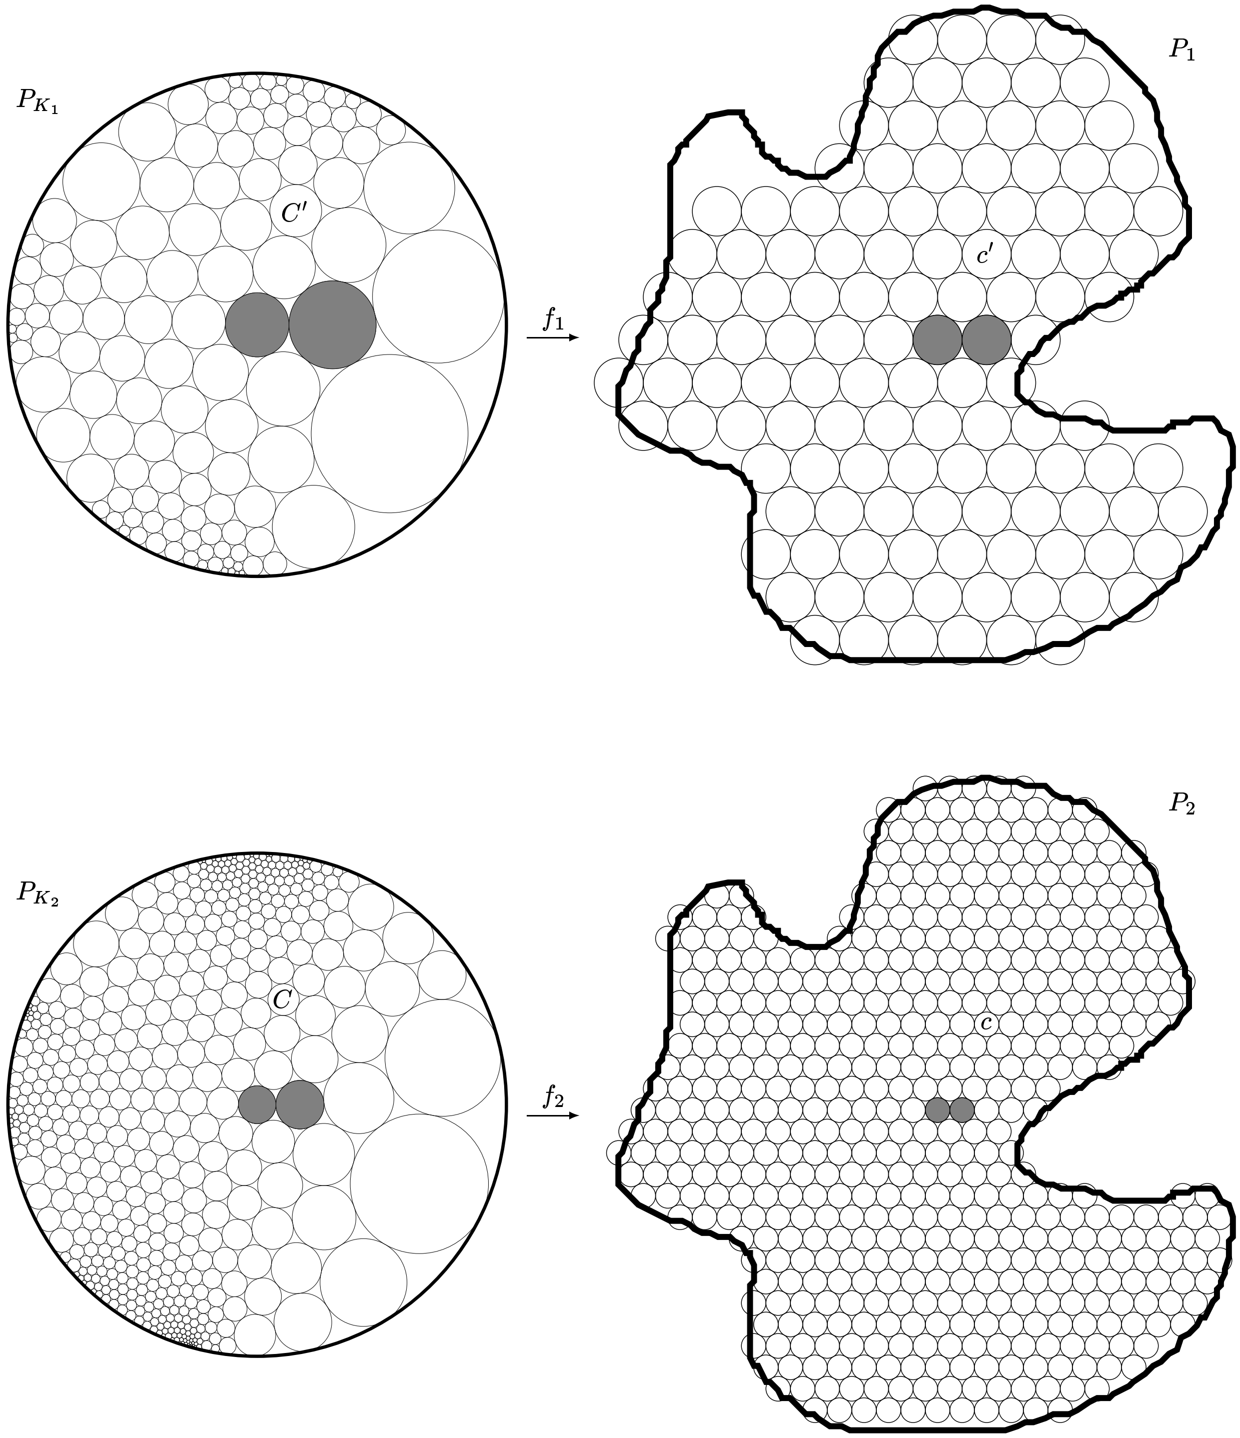

FIGURE 30. Two stages of the multigrid method.

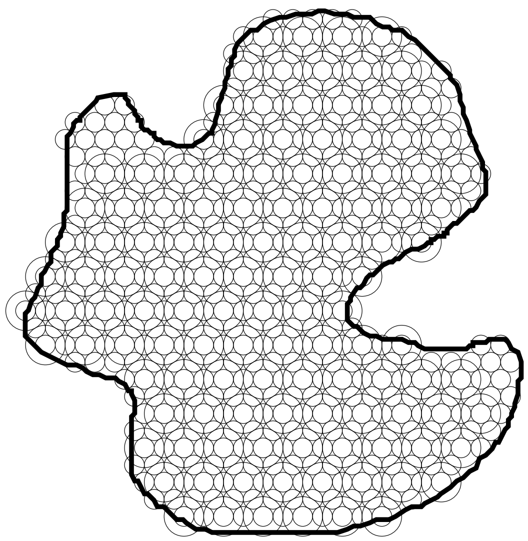

FIGURE 31. P_1 and P_2 superimposed.

- If v is an interior vertex of K_2 that is identified with an interior vertex v' of K_1, then in accordance with (5.1), set

$$R_I(v) = \frac{R_2(v)}{f_1^{\#}(v')},$$

  where R_2 is the hyperbolic label for P_2.
- If v is an interior vertex of K_2 that is not identified with a vertex in K_1 but that is the midpoint of some segment vw, where v' and w' are neighboring interior vertices of K_1, choose $R_I(v) = \frac{1}{2}(R_I(v) + R_I(w))$.
- If v is any other vertex of K_2, choose $R_I(v) = \ln(2)$ (which happens to be the maximal hyperbolic radius of the central circle in any univalent flower having six petals).

Of course, this whole process may now be repeated through a cascade of successively finer packings. The first author developed code to do this, using a link to CirclePack for the computations at each stage. Stages 2, 3 and 4 for the prototype experiment are illustrated in Figure 32; they involve 500, 1992, and 7990 circles. We computed a fifth stage also, but it has more circles (32001) than can reasonably be displayed here.

How efficient is this multigrid approach? We have accumulated rudimentary data in Table 1 concerning the discrete maps of $\mathbb{D}$ onto an ellipse Ω. The three middle stages of the cascade are shown in Figure 33. Roughly speaking, the multigrid approach cut by a factor of five the time needed to compute the finest maximal packing label.

Remarks

There are numerous open, and largely unexplored, issues in the numerics of circle packing, having to do, for example, with algorithms, strategies, accuracy, and rates of convergence.

Most intriguing is the computation of packing labels. The equations they must satisfy are horribly nonlinear. There may well be efficient, indirect approaches to solving them, such as vector-valued Newton's method [Carter 1989] or the minimization approach of [Colin de Verdière 1991]. However, for us the more interesting questions arise when the "labels" are associated directly with the geometry—that is, when they're thought of as radii subject to adjustment—and when features of discrete analyticity can be brought to bear. In this vein, our multigrid method, coupled with the iterative algorithm of Thurston, raises many fascinating questions. (See [Stephenson 1990], for instance, where radius adjustments, tied to the movement of hyperbolic area, are modeled by means of Markov processes.) Moreover, the iterative algorithm and multigrid method both fit quite beautifully with the massively parallel architectures of certain computers. Numerical efficiency has not been a high priority in CirclePack, but we are currently porting the algorithms to a MasPar MP-2, and we anticipate a powerful packing engine that can be linked to CirclePack, allowing a closer study of packing algorithms and the associated dynamics of circle packings.

This is not to say that circle packing will soon reach a practical stage. In regard to speed, a cursory glance at the Table shows that circle packing is slow: the CPU time needed for the finest multigrid approximation, for instance, was over two hours.

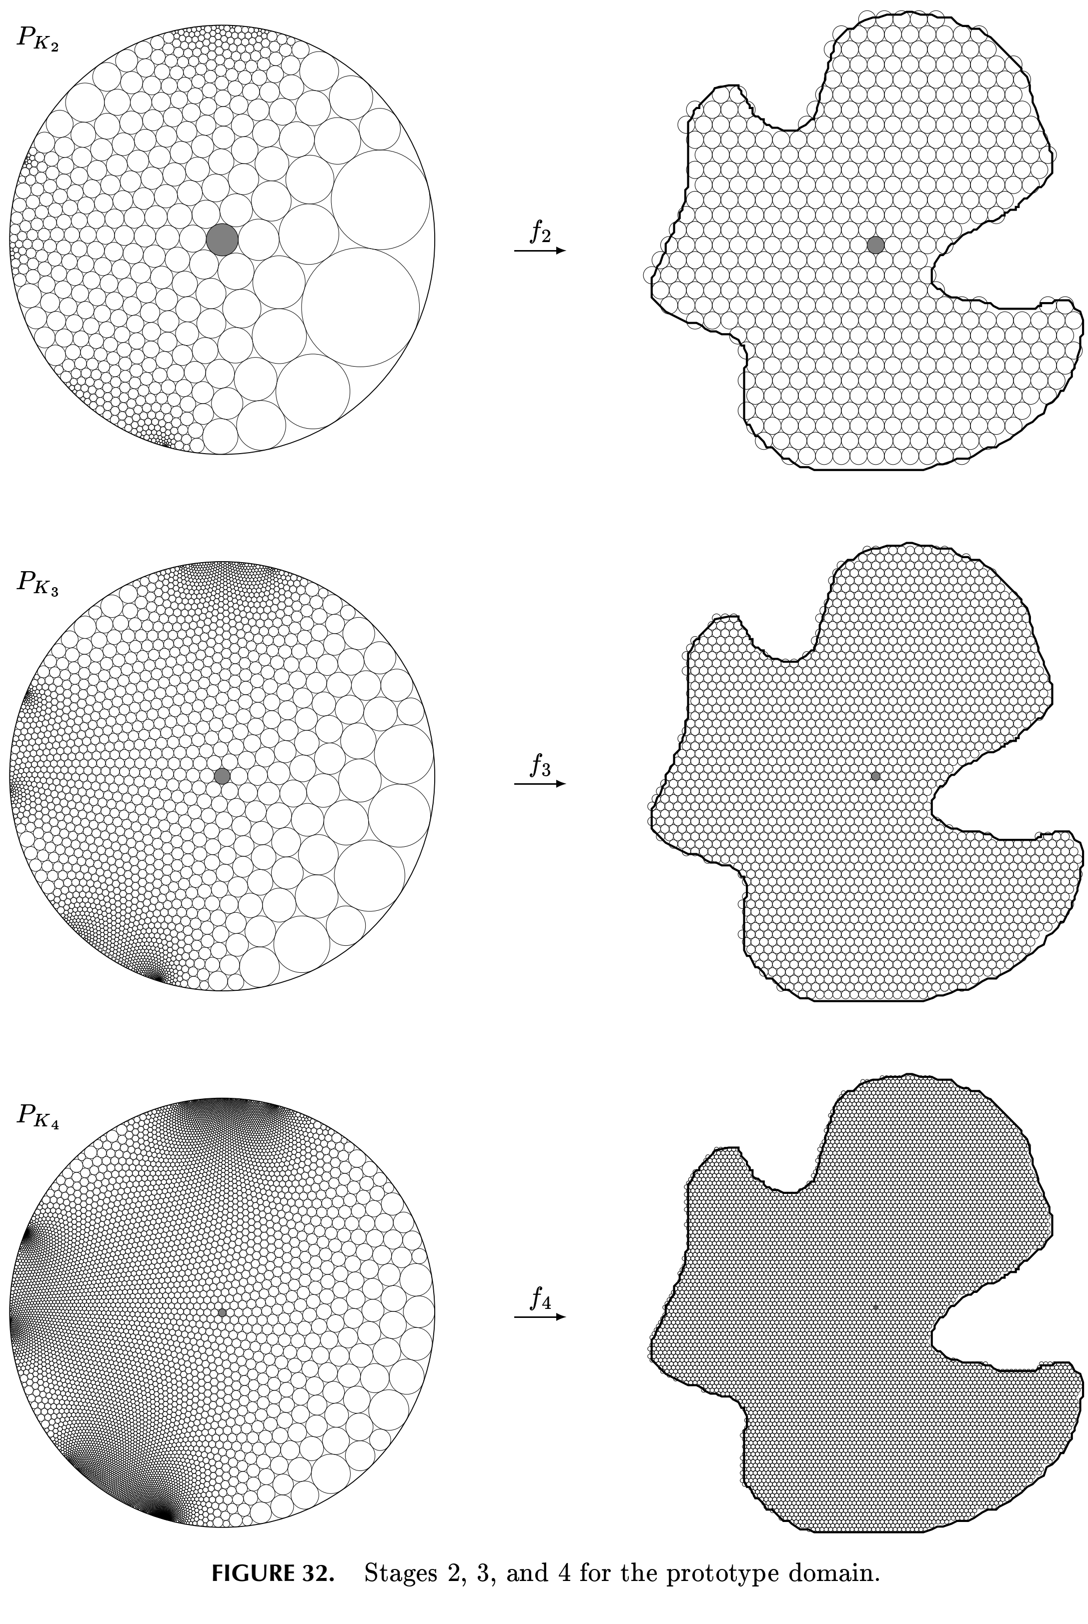

FIGURE 32. Stages 2, 3, and 4 for the prototype domain.

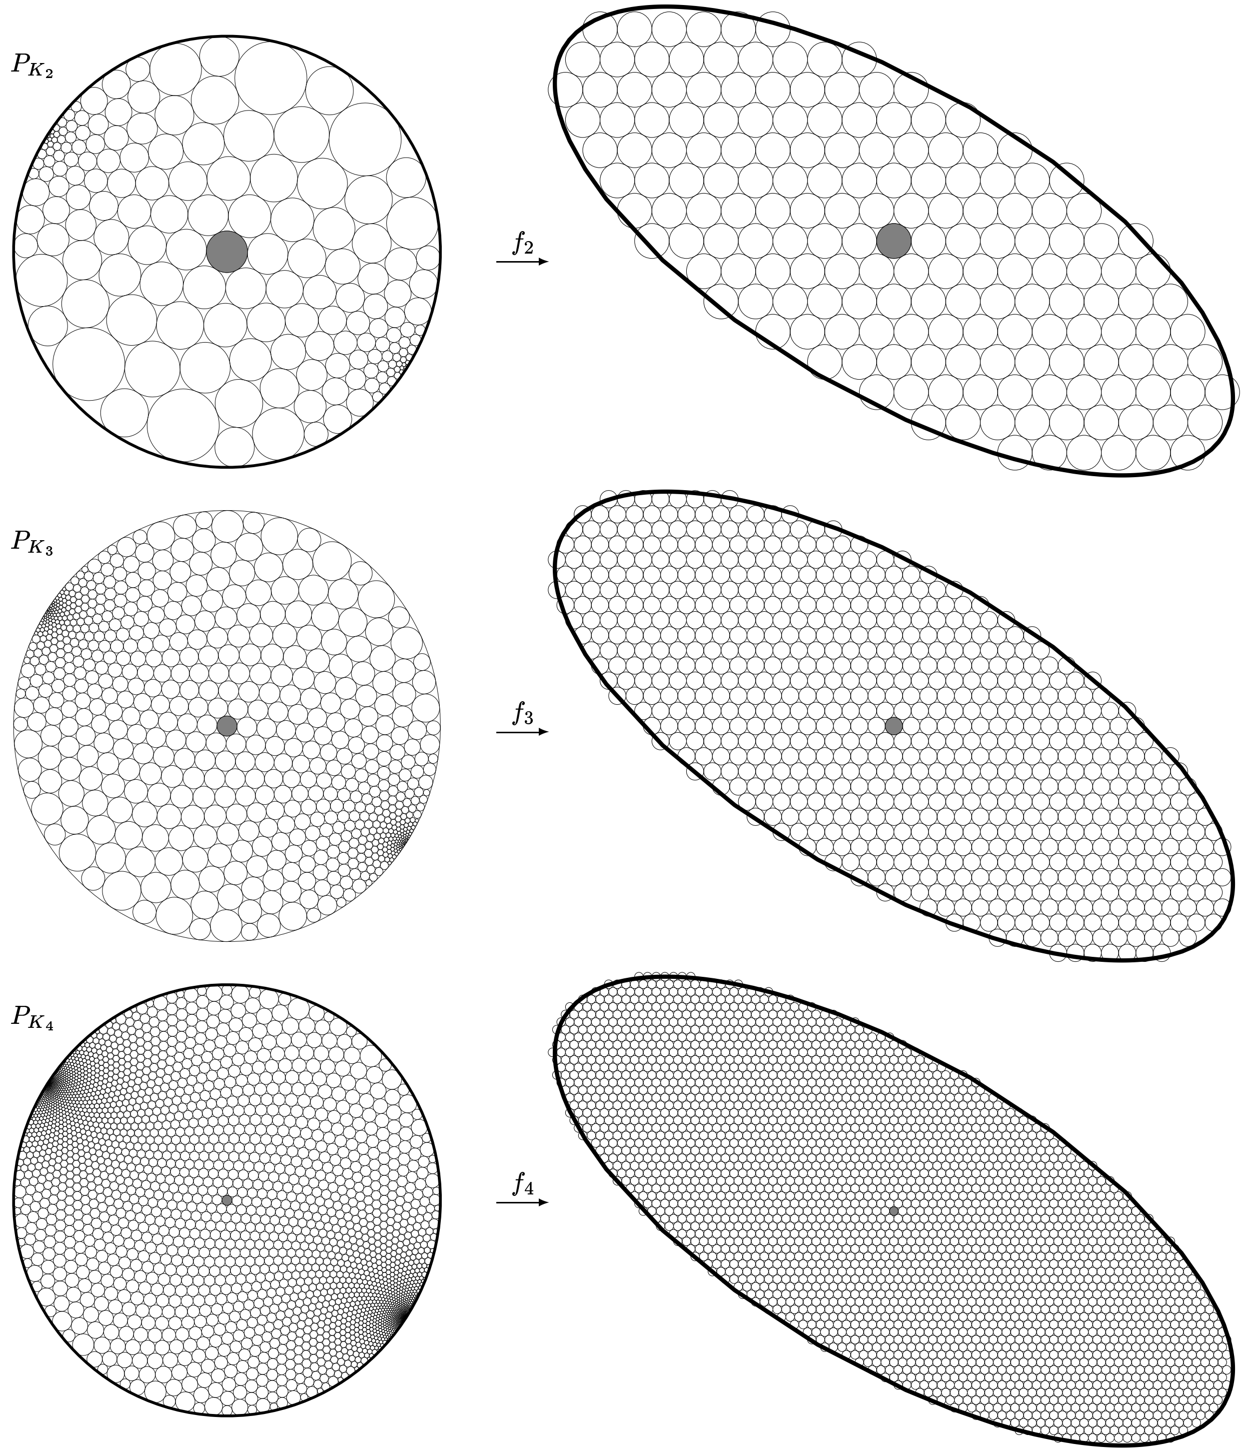

FIGURE 33. Stages 2, 3, and 4 for an ellipse.

			multigrid method		direct method		
stage	circles	tolerance	time	riffles	time	riffles	factor
1	173	10^{-4}	0.69	1445	1.28	2719	0.54
2	709	10^{-4}	6.76	19615	13.10	36313	0.62
3	2814	10^{-5}	69.28	220766	216.62	672822	0.36
4	11277	10^{-6}	1473.94	5440000	3808.87	13387000	0.41
5	45062	10^{-6}	6441.00	23234000	49014.00	163130000	0.16
Total			7991.67	28915826			

TABLE 1. Running times for the determination of the discrete maps from a disk into an ellipse (Figure 33), with and without the multigrid approach. The computations were carried out using CirclePack on a Sun SPARCstation 10, model 40. Tolerance indicates by how much angle sums of interior vertices were allowed to differ from 2π. Time refers to CPU user time required, in seconds. Riffles is the number of circle adjustments required during the iterative computation. Factor gives, for each stage j, the ratio between the time for the multigrid method (accumulated from stages 1 through j) and the time for the direct method at Stage j.

Don Marshall's Zipper program, which uses traditional algorithms to compute conformal maps, handles this ellipse in less than a minute. Classical numerical methods seem likely to retain a speed advantage in situations where they apply. However, with modest improvements in packing efficiency, circle packing may have a role in approximating more general functions, as suggested in our earlier experiments. There are also potential applications in graph embedding (see [Mohar 1993; Miller and Thurston 1990], for example) and in grid generation.

Roundoff is another issue: as the number of circles in a packing grows, errors can begin to affect the packing's integrity. A detail of one small portion of the stage 5 packing for the region Ω of Figure 32, highlighted in Figure 34, shows this. The global properties of the maximal packings remain remarkably stable despite these errors in the outer reaches, plus there are some randomization techniques that should ameliorate this particular problem. Nonetheless, questions on the effects of accumulated errors as well as on the ultimate accuracy of the discrete approximations remain open.

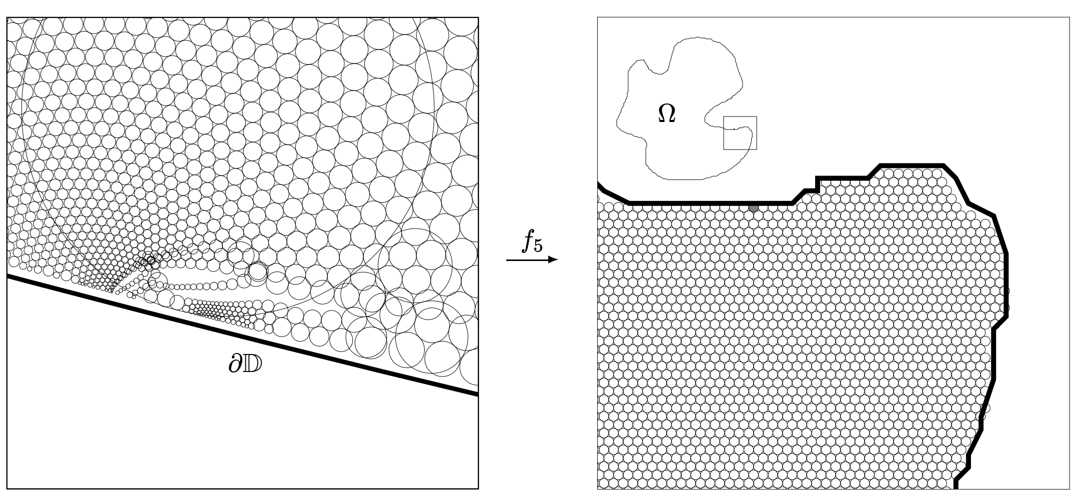

FIGURE 34. Left: Detail of the stage-5 maximal packing calculated for the region Ω of Figure 32; accumulated errors have led to fissures and misplaced circles. Right: Corresponding portion of the range packing in Ω.

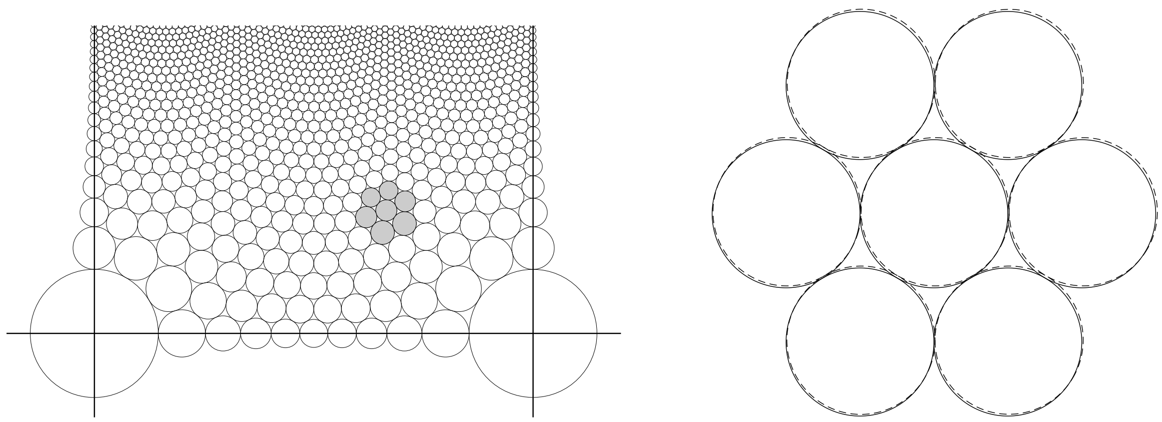

FIGURE 35. Comparison of images for the sine map from Figure 3. Left: We choose one flower in the domain. Right: This flower's image under the sine function (solid curves) is compared with the corresponding circles (dashed) from the regular hex packing.

Some sense of the faithfulness of circle packing maps can be gained with direct comparisons to their classical analytic counterparts. Suppose, for instance, that $f : Q \to P$ is a discrete analytic function that mimics a classical function F defined on some open set Ω containing Q. Each circle C of Q has, as a closed curve in Ω, a closed image curve $F(C)$. Of course, $F(C)$ will in general not be a circle. How close is it to a circle? In particular, how close is it to the corresponding actual circle $f(C)$? The paper provides many opportunities for such comparisons, but we illustrate with just two examples and leave additional comparisons to the reader.

The classical sine function is so familiar that we decided to run our first comparison with the packings of Figure 3. Applying the function sin to the circles in the upper half of Figure 3, we intended to obtain a set of image curves that we could juxtapose with the circles of the lower half. However, we were surprised to find that the two images are indistinguishable at the resolution of the figure. Even when we concentrate on one flower and magnify (Figure 35), the image under the true sine function differs from the image under the discrete analogue by only about $1/500$, and the images under sin are indistinguishable from circles. Considering the rather coarse nature of the domain packing here, the similarity of these range images seems remarkable.

For our second comparison (Figure 36), we chose the discrete conformal mapping f_2 of Figure 30. We obtained our conformal map $F : \mathbb{D} \to \Omega$ (also an approximation, of course) by using Don Marshall's Zipper conformal mapping package. Again, the images $F^{-1}(C)$ appear indistinguishable from circles. Inaccuracies in the locations are not surprising, and their pattern makes sense if one refers to Figure 30. Two large boundary horocycles of the maximal packing are a mere three generations from the circle at the origin. They can't provide much sensitivity for an important (in the sense of harmonic measure) section of $\partial\Omega$, and they seem to have pushed the other circles toward -1 with respect to their conformally "correct" locations.

6. ABOUT CIRCLEPACK

For readers who would like to experiment with circle packings on their own, here are some technical details about CirclePack, the software used for this paper. CirclePack allows the creation, manipulation, color display, printing, and storage of circle packings in any of the three standard geometries (the spherical routines are in an early stage of development). In addition to the tangency packings

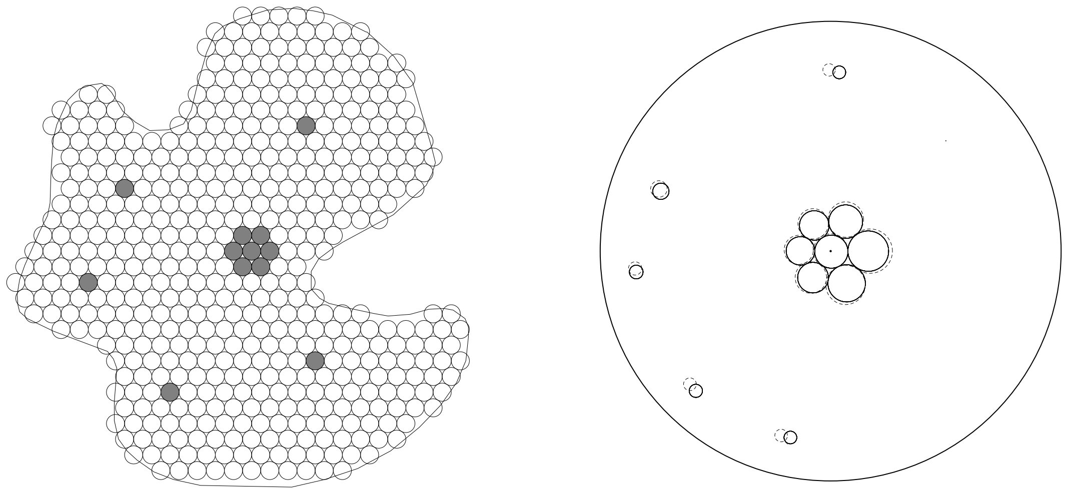

FIGURE 36. Comparison of images for the discrete conformal map f_2 of Figure 30. Left: Twelve circles were chosen from the packing of the region Ω. Right: The corresponding circles in the maximal packing (dashed curves), and inverse images of the shaded circles under the conformal map taking the disk to Ω (solid curves).

discussed in this article, CirclePack is capable of handling packings with overlap angles specified individually for edges of K. The complexes K underlying the packings may be simply connected, like those in this article, or multiply connected; planar or nonplanar; with or without boundary. Packings of up to 100,000 circles are possible, depending on machine resources. Parameters for the program's operation are in a configuration file, allowing various defaults and features to be tailored to the user's needs. User-developed remote routines may be called from within CirclePack; some examples, along with basic code segments and header files, are provided in the distribution. All data files have ASCII formats. Postscript is used for printable output.

CirclePack is written in C, runs under X-Windows, and is built on the XView graphical user interface. User commands are implemented through command-line arguments, mouse-based actions, or with prepared scripts. Figure 37 illustrates a typical screen. A detailed help file, accessible from within CirclePack, provides context-sensitive help.

The lines from a typical script are shown below. The first line activates Pack 0 (each packing being identified by a number); it then reads in a set of circle packing data, sets the screen size, and displays the packing. The next line sends a copy of the data to Pack 2, activates Pack 2, converts its packing to hyperbolic geometry, and sets two of its vertices' "aims" to 4.0π (to create branch points). The last line sets the boundary radii of Pack 2 to infinity, tells the program to repack using up to 20,000 iterations of Thurston's algorithm; recomputes the circle centers, and displays the packing. All these commands and many others can be issued individually by typing them on a command line or by using mouse operations, but it is particularly helpful to run through sample scripts to get a feeling for how things work.

```
[1]:= act 0; infile_read Q3.p; set_screen -h 7; disp -w -c;
[2]:= copy 2; act 2; geom_to_h;  set_aim -d; set_aim 4.0 26 33;
[3]:= set_rad -.1 b; repack 20000; fix; disp -w -u -c b -cf 26;
```

Sample CirclePack script

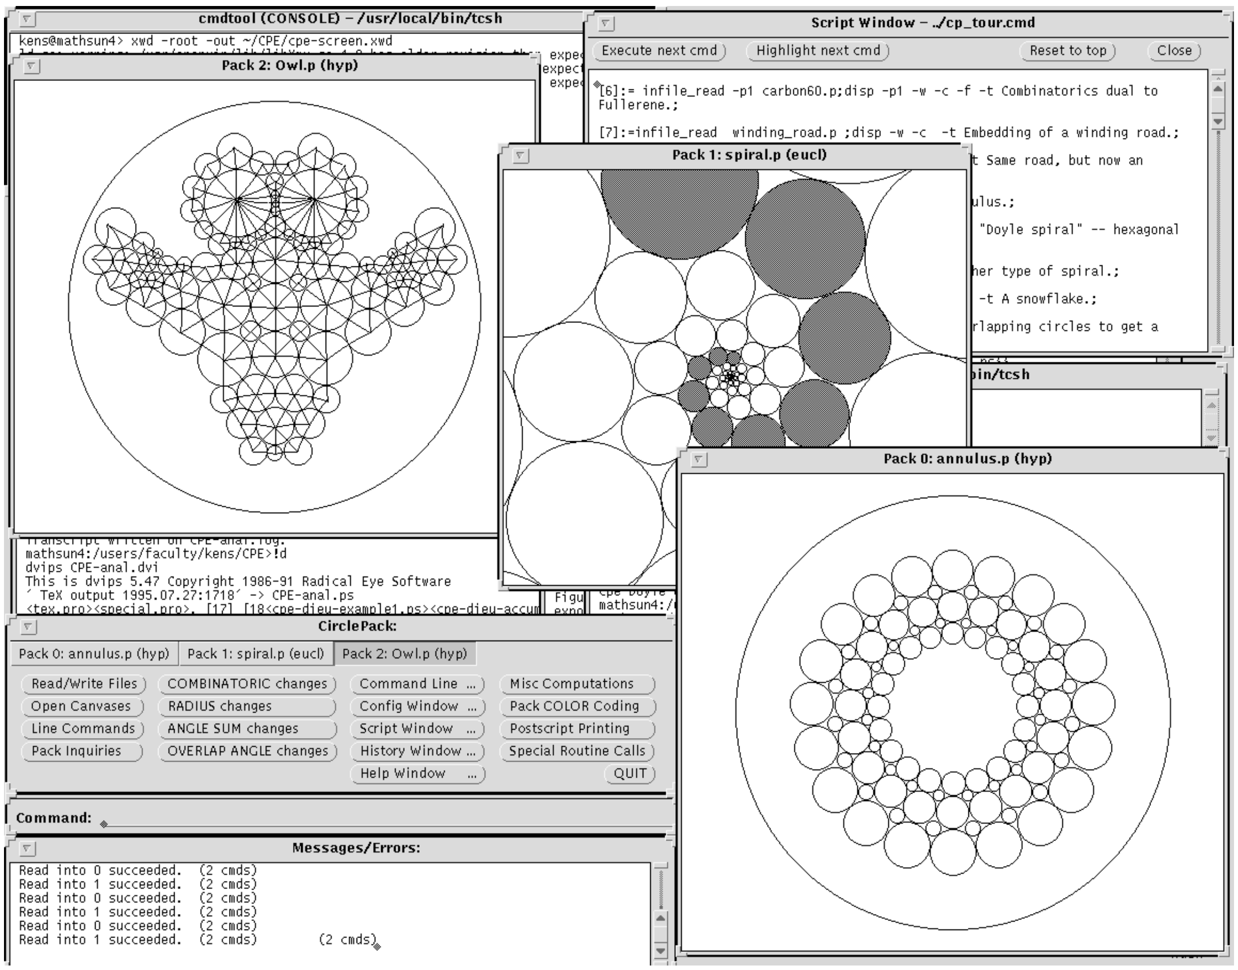

FIGURE 37. Sample CirclePack screen under X-Windows.

CirclePack was developed by Stephenson on a Sun SPARCstation with equipment and research support provided, in part, by the National Science Foundation and the Tennessee Science Alliance.

Retrieval information is given on the next page.

REFERENCES

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

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

[Beardon et al. 1994] Alan F. Beardon, Tomasz Dubejko, and Kenneth Stephenson, "Spiral hexagonal circle packings in the plane", *Geometriae Dedicata* **49** (1994), 39–70.

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

[Callahan and Rodin 1993] Kevin Callahan and Burt Rodin, "Circle packing immersions from regularly exhaustible surfaces", *Complex Variables* **21** (1993), 171–177.

[Carathéodory 1960] C. Carathéodory, *Theory of Functions of a Complex Variable II*, Chelsea, New York, 1960.

[Carter 1989] Ithiel Carter, "Circle packing and conformal mapping", Ph.D. Thesis, University of California, San Diego, 1989.

[Colin de Verdière 1991] Yves Colin de Verdière, "Un principe variationnel pour les empilements de cercles", *Inventiones Mathematicae* **104** (1991), 655–669.

[Dubejko 1995] Tomasz Dubejko, "Branched circle packings and discrete Blaschke products", *Trans. Amer. Math. Soc.* **347** (1995), 4073–4103.

[Dubejko 1996] Tomasz Dubejko, "Recurrent random walks, Liouville's theorem, and circle packings", to appear in *Proc. Cambridge Phil. Soc.* (1996). Also available at http://www.msri.org/MSRI-preprints/online/1995-040.html.

[Dubejko a] Tomasz Dubejko, "Infinite branched packings and discrete complex polynomials" to appear in *Journal of Lond. Math. Soc.*

[Dubejko and Stephenson 1995] Tomasz Dubejko and Kenneth Stephenson, "The branched Schwarz lemma: a classical result via circle packing", *Mich. Math. J.* **42** (1995), 211–234.

[He and Rodin 1993] Zheng-Xu He and Burt Rodin, "Convergence of circle packings of finite valence to Riemann mappings", *Comm. in Analysis and Geometry* **1** (1993), 41–41.

[He and Schramm] Zheng-Xu He and Oded Schramm, "Koebe uniformization for "almost circle domains", *Amer. J. Math.* **117** (1995), 653–667.

[Marshall] Donald E. Marshall, "Zipper" software package, available at http://www.math.washington.edu/~marshall/personal.html.

[Marx 1974] Morris L. Marx , "Extensions of normal immersions of S^1 into $\mathbf{R}^2$", *Trans. Amer. Math. Soc.* **187** (1974), 309–326.

[Miller and Thurston 1990] Gary L. Miller and William Thurston, "Separators in two and three dimensions", pp. 300–309 in *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing*, edited by H. Ortiz, ACM, Baltimore, 1990.

[Mohar 1993] Bojan Mohar, "A polynomial time circle packing algorithm", *Discrete Math.* **117** (1993), 257–263.

[Morgan 1984] John W. Morgan, "On Thurston's uniformization theorem for three-dimensional manifolds", pp. 37–126 in *The Smith Conjecture* (edited by Hyman Bass and John W. Morgan), Academic Press, New York, 1984.

[Nevanlinna 1970] Rolf Nevanlinna, *Analytic Functions*, Springer, New York, 1970.

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

[Stephenson 1990] Kenneth Stephenson, "Circle packings in the approximation of conformal mappings", *Bull. Amer. Math. Soc.* **23** (1990), 407–415.

[Thurston 1985] William Thurston, "The finite Riemann mapping theorem", Invited talk at the international symposium at Purdue University on the occasion of the proof of the Bieberbach conjecture, March 1985.

Electronic Availability

CirclePack is available by anonymous ftp from the host archives.math.utk.edu (Mathematics Archives), in the directory software/multi-platform/circle.packing. It occupies approximately 10 Meg when installed. The distribution contains all source code, executables for Sun SPARCstations, help files, numerous packing data sets, and prepared command scripts. A makefile is provided and should permit recompilation on other platforms, though help from someone having systems experience is recommended.

The distribution also contains, in subdirectory CPE, scripts that run through several of the experiments described in this paper. For instance, to generate the packings of Figure 26, run CirclePack using the script CPE-5-poly.cmd. The "script window" (seen in the upper right hand corner of Figure 37) will appear, and you can run through the sequence of prepared commands by simply clicking on the button "Execute next cmd".

Tomasz Dubejko, Department of Mathematics, Northwestern University, Evanston IL 60208
(tdubejko@math.nwu.edu, http://www.math.nwu.edu/~tdubejko)

Kenneth Stephenson, Department of Mathematics, University of Tennessee, Knoxville, TN 37920-1200
(kens@math.utk.edu; http://www.math.utk.edu/~kens)

Received January 5, 1995; accepted in revised form August 17, 1995