

CONFORMAL TILINGS I: FOUNDATIONS, THEORY, AND PRACTICE

PHILIP L. BOWERS AND KENNETH STEPHENSON

ABSTRACT. This paper opens a new chapter in the study of planar tilings by introducing *conformal* tilings. These are similar to traditional tilings in that they realize abstract patterns of combinatorial polygons as concrete patterns of geometric shapes, the tiles. In the conformal case, however, these geometric tiles carry prescribed conformal rather than prescribed euclidean structure. The authors develop the topic from the ground up: definitions and terminology, basic theory on existence, uniqueness and properties, numerous experiments and examples, comparisons to traditional tilings, patterns unique to conformal tiling, and details on computability through circle packing. Special attention is placed on aperiodic hierarchical tilings and on connections between abstract combinatorics on one hand and their geometric realizations on the other. Many of the motivations for studying tilings remain unchanged, not least being the pure beauty and intricacy of the patterns.

INTRODUCTION

A “tiling” (or “tessellation”) of a surface is a decomposition into a pattern of disjoint open sets whose closures, the *tiles*, cover the surface. Nature is rife with examples, as are various art traditions, from Greek mosaics, to the Moorish tilings of the Alhambra, to Escher’s angels and demons drawing, to textiles and quilts of myriad cultures. The mathematical fascination with such patterns traces back to the earliest geometers and continues unabated.

The direct line to our work begins with the famous Penrose tiling of the plane and the interest in “aperiodic hierarchical” tilings that it spawned. We call these “traditional” to distinguish them from the “conformal” tilings that we study here. Figure 1 displays four such traditional tilings. In each the tiles are euclidean polygons attached along edges, sometimes with coloring as a visual aid.

These patterns may have an immediate visual appeal, a comforting regularity. Spend some time with them, however, and you quickly find more subtle things going on. Each pattern has only a finite number of tile shapes: “kites” and “darts” in the Penrose, three different rhombi for dimers, a single right triangle for the pinwheel, and a single “chair”.

Despite the modest building material, these patterns are the antithesis of regular. In point of fact, all these examples are “aperiodic”, meaning that they have no global symmetries. Different regions are eerily reminiscent of one another, yet on closer inspection are never quite the same.

Received by the editors November 20, 2014 and, in revised form, August 26, 2015.

2010 *Mathematics Subject Classification*. Primary 52C23, 52C26; Secondary 52C45, 68U05.

The second author gratefully acknowledges support of a Simons Foundation Collaboration Grant.

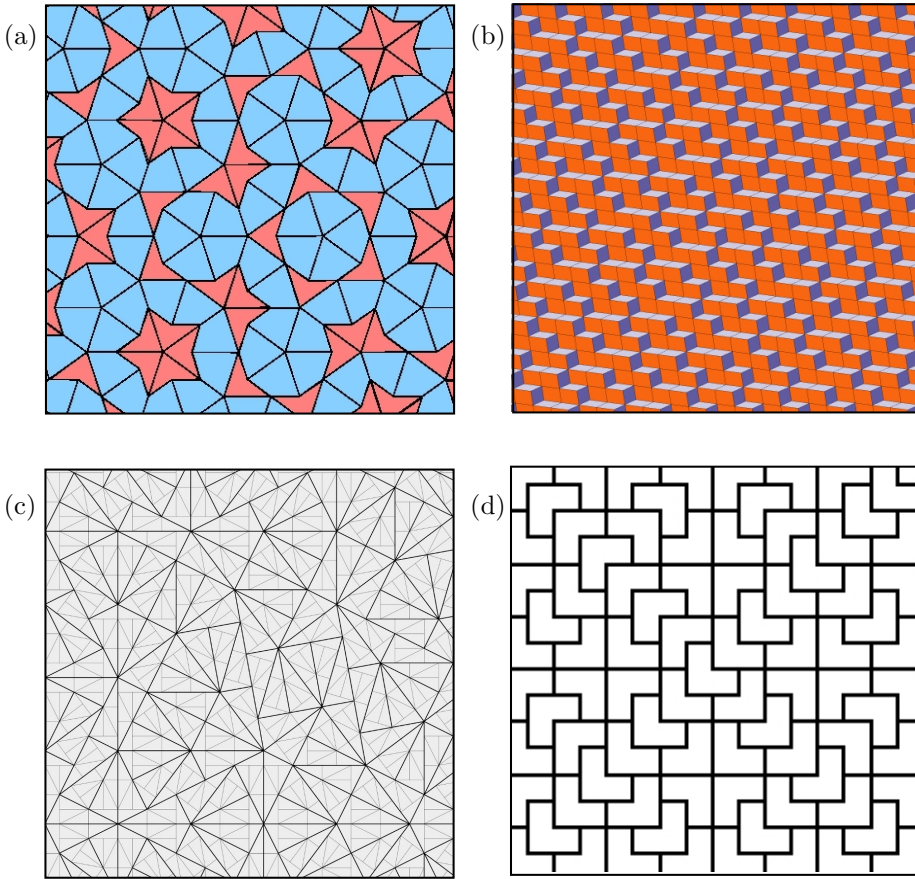


FIGURE 1. Traditional tiling examples: (a) Penrose kite and dart, (b) Dimer, (c) Pinwheel, and (d) Chair.

These patterns are also “hierarchical” in that they arise from subdivision rules. Look closely at the pinwheel image and you will see markings to break each tile into 5 tiles, all similar to the original. Form a new tiling from these smaller tiles, expand by a factor of $\sqrt{5}$ so each tile is congruent with the original, and you get an entirely new pinwheel pattern. Remarkably, the pinwheel and its progeny are guaranteed to be combinatorially distinct, yet any arbitrarily large finite piece of either will be found replicated infinitely many times within the other.

The variety of distinct tilings is quite staggering. Nonetheless, they can often be organized into tiling families, imbued with topologies and so forth. It is not the individual tilings — gorgeous though they may be — that are of prime interest, rather tiling families are the real target in this paper, its successor [5], and future work. Broadly speaking, *conformal tiling* represents a new chapter in a topic whose long history is a succession of new examples, new tools, themes, questions, and links to new topics.

The Example: Our progenitor example is the *regular pentagonal* tiling of Figure 2, introduced by the authors in [8]. This was motivated by work of Jim Cannon, Bill Floyd, and Walter Parry and realized through circle packing. Connections to traditional tilings emerged only recently with the doctoral dissertation and subsequent deep study of this example by Maria Ramirez-Solano (see [30–34]). The image in Figure 2 uses colors to highlight the underlying hierarchical structure.




FIGURE 2. The regular pentagonal tiling.

Tools: Conformal geometry is the new tool being introduced; this is the geometry of conformal (angle preserving) maps, analytic functions, and Riemann surfaces. Indeed, just two key properties fully characterize conformal tilings:

- First, the tiles are *conformally regular*, meaning that each tile is the image under a conformal homeomorphism of a regular euclidean polygon.
- Second, the tilings are *reflective*, meaning that whenever two tiles share an edge, they are anticonformal reflections of one another in that edge.




FIGURE 3. Conformal tiling examples: (a) Conformal Fibonacci, (b) Conformal dimer, (c) Conformal pinwheel, (d) Conformal chair.

Figure 3 shows four conformal tilings, three of which are conformal versions of tilings from Figure 1. Our “polygonal” tiles typically have sides that are curvilinear rather than straight. Not only *can* tile shapes differ, they *must* differ: By the reflective property, the shape of any single tile in a conformal tiling uniquely determines the full tiling — its combinatorics, every tile shape, every tile location. Conformal geometry is a rich and well studied area, and brings a wealth of tools. Even the reader not familiar with the details may see the geometric intuition bubble up from the many images we provide.

Themes: Tiling *patterns* take center stage. A pattern is called a *combinatorial tiling* and is simply an abstract cell complex. Surprising as it may seem, *every* combinatorial tiling can be realized as a concrete conformal tiling. Restricting ourselves largely to simply connected patterns, we will see examples in the euclidean plane \mathbb{C} , the hyperbolic plane \mathbb{D} , and on the sphere \mathbb{S}^2 .

The combinatorial theme is a major departure. Conjuring traditional tilings — the Penrose, chair, pinwheel, and so forth — is akin to an art: The results are beautiful, rare, and highly prized. The wealth of conformal

tilings, in contrast, raises different issues: Which patterns are worthy and why? Our interest is ultimately in tiling families, so we focus largely on aperiodic hierarchical tilings generated by subdivision rules. A central theme is always how the combinatorics and geometry are linked, whether studying conformal versions of traditional tilings, or new patterns which have no traditional counterparts. Among the latter are tilings generated by special *conformal subdivision rules* which recapture a key traditional feature in that the subdivisions occur *in situ*.

Methods: Conformal tiling practice is grounded firmly in the topic of *circle packing*. Circle packings, configurations of circles with specified patterns of tangency, were introduced by William Thurston along with a conjecture relating them to conformal mapping, [41]. Over the intervening years they have become the basis of a comprehensive theory of *discrete conformal geometry*; see [39] for an overview and [40] for details. Their association with tiling began with the regular pentagonal pattern and casual experiments requested by Bill Floyd to build an embedding. The experiments, images, and theory of conformal tiling have gone hand-in-hand with circle packing since this beginning.

It must be noted that our proofs for existence and properties of conformal tilings in this paper are formally independent of circle packing. Nonetheless, circle packing provides all the images, circle packing experiments have inspired most of the results, and *discrete conformal tilings*, the discrete parallel to the theory, often serve as a bridge between the combinatorial and the fully conformal settings.

For our purposes there are three important observations about circle packing: (1) Circle packings are conformal by nature. The discrete analogues of classical conformal objects that they provide, such as analytic functions, Riemann surfaces, and now conformal tilings, invariably manifest classical conformal phenomena in a faithful discrete form. (2) Under appropriate refinement, these discrete objects invariably (and provably) converge to their classical counterparts. (3) Circle packings are computable and with software `CirclePack` (see Appendix C) are easy to manipulate and render.

There are also core questions, the principal one having to do with “type”: Is a given infinite simply connected tiling *parabolic* or *hyperbolic* — that is, is it realized in the euclidean or the hyperbolic plane? For many individual tilings we resolve this issue here using circle packing results and quasiconformal mapping arguments. Nonetheless, for tiling families, type remains a key issue in our second paper.

Links: Conformal tiling links tiling theory with a range of other topics: the deep studies of subdivision rules of Cannon, Floyd, and Parry aimed at Cannon’s Conjecture, Grothendieck’s theory of *dessins d’Enfants* and its algebraic implications, dimer tilings in statistical mechanics, random triangulations and quadrangulations in probability, and “carpets” in metric geometry, to name a few. We comment on these in the paper as we discuss examples. We particularly target in this and our second paper a way to define tiling spaces and their dynamical systems. Conformal tiling brings with it an intrinsic metric which may be the missing link in this regard.

Studies in conformal tiling naturally have a different character than those in traditional tiling, yet much of the underlying appeal remains. Foremost perhaps — although not all will admit to this — is the pure beauty and intricacy of the patterns. We don't want the reader to overlook the central role that images, experiments, pattern recognition, and outright luck have played in this new topic. That is likely to continue, as nearly every image raises more questions than it answers — this is a fluid and target-rich area. There is also a note of caution: Images are finite and only approximate, they can take us only so far. We try to be careful here in describing what we have proven and cautious about conjecture. We are, nevertheless, optimistic that we provide a base for continued development, perhaps to notions of the tiling spaces, dynamical systems, and associated C^* algebras that have been developed in the traditional setting; see [3, Chp 5].

This paper lays the foundation of conformal tiling from the ground up — structures, terminology, main theoretical results, and numerous examples — with sufficient background on conformal geometry and the practical aspects of circle packing to be broadly accessible.

Here is a brief *Outline*: In Section 1 we carefully lay out the successive combinatorial, conformal, metric, and geometric structures behind conformal tilings. Section 2 describes the two key notions, *conformally regular* and *reflective*, and proves strong rigidity, limit, and normal families theorems. In Section 3 we move to examples, beginning with the broader landscape, but then focusing on conformal versions of traditional aperiodic hierarchical examples: the *chair*, *domino*, *sphinx*, *pinwheel*, and *Penrose* tilings of the euclidean plane and the *hyperbolic pentagonal* tiling in the hyperbolic plane. We describe subdivision rules, and their hierarchical natures with non-traditional *lace*, *mixed*, and *snowball* tilings. A new notion, that of *conformal* subdivision rules, is defined and illustrated. This is used extensively in our second paper, [5], and seems likely to be a core theme in conformal tiling.

In Section 4 we shift to the practical side of conformal tiling. Enough background is provided on circle packing for the reader to understand the machinery that makes the theory computable. Key is a faithful parallel discrete-conformal tiling world that both mimics and approximates the continuous. Several technical results are gathered in Section 5, including background on piecewise affine and conformal structures, results on the important “type” problem, and some technical lemmas needed for circle packing. The paper closes with appendices, including one on non-planar conformal tilings.

CONTENTS

Introduction	1
1. Structural stages	7
1.1. Combinatorial structure	7
1.2. Conformal structure	9
1.3. Metric structure	10
1.4. Geometric structure	11
2. The inner workings of conformality	14
2.1. Individual conformal tiles	14

2.2. Shape and reflection	16
2.3. Tile ensembles	17
2.4. Rigidity	19
2.5. Limits of tilings	20
2.6. Observations	22
3. Conformal tiling examples	23
3.1. Broad examples	23
3.2. The main examples	27
3.3. Finite subdivision rules	29
3.4. Traditional subdivision rules	30
3.5. Observations	35
3.6. Aggregate shapes	36
3.7. Conformal subdivision rules	38
4. Discretization and approximation	44
4.1. Tiling complexes	44
4.2. Circle packing	45
4.3. Discrete conformal tilings	46
4.4. Limit tilings	47
4.5. Observations	53
5. Technical results	54
5.1. Piecewise affine structures	54
5.2. Type problems	55
5.3. Circle packing lemmas	57
Appendix A. Alternate conformal structures	59
Appendix B. Multiply connected tilings	59
Appendix C. <code>CirclePack</code> software	61
References	61

1. STRUCTURAL STAGES

This section lays out the basic structures behind conformal tilings. The nearly unlimited flexibility of conformal methods requires structural discipline if one hopes for any coherent theory. We discuss in succession the combinatorial structure, conformal structure, metric structure, and lastly, our main target, concrete geometric structure. Our interest is in conformal realizations, and the final section explains the variety one encounters based on conformal mapping theory and the conformal type of surfaces.

1.1. Combinatorial structure. The *combinatorics* of a tiling refers to its abstract pattern, which we will encode as an abstract 2-dimensional cell decomposition. Let S be an oriented topological surface and D a connected graph embedded in S . We assume the following:

- D is allowed to have loops and/or multiple edges.
- S may be a bordered surface, in which case its border is part of D .
- Each component C of $S \setminus D$ is a topological disc whose closure is compact and meets at most finitely many edges of D .
- D is locally finite in S , meaning each point of S has a neighborhood meeting at most finitely many vertices and edges of D .

The graph D will be called a *drawing* in S , in deference to Grothendieck’s theory of *dessins d’Enfants*, “drawings by children”, for compact surfaces, though our drawings often have infinitely many vertices and edges. The vertices and open edges of D determine a CW-complex, the components of $S \setminus D$ being its open cells. (In our second paper [5], we will restrict to infinite drawings in the plane without loops or free edges, using the term “planar polygonal complex” and denote the combinatorics by K .)

Definition. A *combinatorial tiling* \mathcal{T} is an oriented topological surface S endowed with a cell-decomposition by an embedded drawing D as described above. The components of $S \setminus D$, the 2-cells, represent the tiles.

Each tile is a combinatorial n -gon for some $n \geq 1$, so it is simply connected, has n vertices and n edges, and meets at most finitely many neighboring tiles, each in one or more shared vertices and/or full edges. The surface S may be simply or multiply connected, compact, bordered, or open. If S is not compact, then \mathcal{T} will have infinitely many tiles. Nevertheless, our conditions guarantee that \mathcal{T} is always *locally finite*, meaning that the closure of any cell meets the closures of at most finitely many other cells.

The pattern of the tiling \mathcal{T} is conveniently encoded in its *tiling graph*; namely, the graph D^\dagger which embeds in S as the graph dual to D . This dual has a node for each tile $T \in \mathcal{T}$ and an edge e^\dagger for each edge e of D : If edge $e \in D$ is shared by tiles $T, T' \in \mathcal{T}$, then edge $e^\dagger \in D^\dagger$ connects the nodes for T and T' . Note that D^\dagger inherits an orientation from \mathcal{T} and that in the embeddings, edges e and e^\dagger cross at a point. We will routinely use combinatorial terminology associated with the tiling graph. A combinatorial tiling is *rooted* means that a node of D^\dagger has been designated as a root tile. A combinatorial *automorphism* of \mathcal{T} is associated with a (orientation preserving) graph automorphism of D^\dagger . Combinatorial distances are in the *path metric* in D^\dagger : Tiles T and T' are distance d apart means that d is the minimum count of edges for edge paths in D^\dagger between the nodes for T and T' . We will use a type of “big-ball” metric on rooted combinatorial tilings when we later discuss convergence.

There are two standard ways to triangulate cell complexes, and we use both. Though defined abstractly, these can be illustrated using our embeddings of D and D^\dagger in S . Begin with a *star* subdivision. Each node of D^\dagger is interior to a tile and is treated as the *tile barycenter*; draw edges from this barycenter to each of the tile’s vertices. These edges along with D itself give a triangulation of S denoted \mathcal{T}^* . (This is the *face barycenter triangulation* of [12].) A tile of \mathcal{T} that is an n -gon is realized as a union of n triangles in \mathcal{T}^* . Next, we decompose further by breaking each triangle of \mathcal{T}^* in two using the edges of D^\dagger . This defines an *edge barycenter* for each edge e of \mathcal{T} where its dual edge e^\dagger crosses. The result of these two subdivisions is the *barycentric subdivision* of \mathcal{T} , denoted $\beta\mathcal{T}$. Using the notation of cell complexes,

$$\beta\mathcal{T} \leq \mathcal{T}^* \leq \mathcal{T},$$

where the “ \leq ” indicates cellular embedding — that is, each cell on the subordinate side is embedded in a single cell on the superior side (using the identity map on S as the embedding map).

Figure 4 provides an example. Each vertex of \mathcal{T} is marked with \bullet , each edge barycenter with \times , and each tile barycenter with \square . The black edges form the original drawing D , the red edges are those added to get \mathcal{T}^* , and the blue edges are

those added to get $\beta\mathcal{T}$. Note that these are *abstract* triangulations: The embedding of the figure is merely for purposes of display (but will become our main concern shortly).

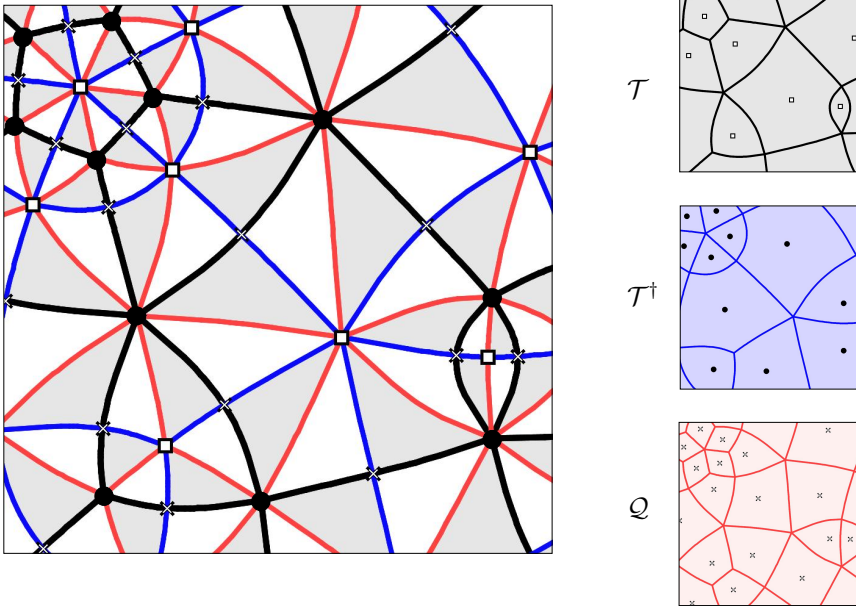


FIGURE 4. On the left is the barycentric subdivision $\beta\mathcal{T}$ for a tiling \mathcal{T} . On the right are the three tilings within $\beta\mathcal{T}$: the tiling \mathcal{T} itself, its dual tiling \mathcal{T}^\dagger , and its quad tiling \mathcal{Q} .

Since $\beta\mathcal{T}$ is the fundamental combinatorial object behind our work, some notation is in order. The sets of vertices, edges, and faces of $\beta\mathcal{T}$ will be denoted $\mathcal{V}, \mathcal{E}, \mathcal{F}$, with \mathcal{V} a disjoint union $\mathcal{V} = \mathcal{V}_\bullet \cup \mathcal{V}_\times \cup \mathcal{V}_\square$. Each face is a combinatorial triangle with one vertex from each of $\mathcal{V}_\bullet, \mathcal{V}_\times$ and \mathcal{V}_\square . The faces of $\beta\mathcal{T}$ fall into two equivalence classes based on vertex marking: triangles with $\{\bullet, \times, \square\}$ positively (i.e., counter-clockwise) oriented will be termed *white*, the others, *gray*. A tile of \mathcal{T} with n edges is broken into $2n$ faces, alternating white and gray, as seen in Figure 4.

The image of $\beta\mathcal{T}$ is rather busy, but also rich: In addition to \mathcal{T} , it encodes two additional tilings. These are parsed out on the right in Figure 4. The black edges give us the tiling \mathcal{T} itself, the blue edges give the dual tiling \mathcal{T}^\dagger (the tiling associated with the drawing D^\dagger), and the red edges give the *quad tiling* \mathcal{Q} (each of its tiles has four edges).

1.2. Conformal structure. We choose a canonical way to represent \mathcal{T} as a Riemann surface; namely, by defining it as an *equilateral* surface. This is a standard technique, used, for example, in [8, 9]. Given a triangulated topological surface S , one identifies each triangular face with a unit-sided euclidean equilateral triangle. From these identifications one constructs a conformal atlas, thereby realizing S as a Riemann surface. Equilateral surfaces are special cases of the more general *piecewise affine* surfaces, about which the reader will find additional detail in §5.1.

For a combinatorial tiling \mathcal{T} , we could choose to work with the barycentric triangulation $\beta\mathcal{T}$, giving it the β -*equilateral* structure, or with the coarser star

triangulation \mathcal{T}^* , giving it the *star equilateral* structure. The resulting conformal structures on \mathcal{T} are identical. For both theoretical and practical reasons, we work largely with the β -equilateral structure in this paper and with the star equilateral structure in our second paper.

Definition. For a combinatorial tiling \mathcal{T} , the *canonical conformal structure* on \mathcal{T} is the β -equilateral structure associated with $\beta\mathcal{T}$. With this conformal structure, \mathcal{T} is termed a *conformal tiling*. Its *conformal tiles* are the 2-cells, each with its inherited conformal structure. Other regroupings of triangles of $\beta\mathcal{T}$ give the conformal tilings \mathcal{T}^\dagger and \mathcal{Q} .

Two properties of this conformal structure are key to our work: (1) the tiles are *conformally regular*, meaning that each n -gon tile has a full dihedral symmetry group D_{2n} of conformal and anticonformal automorphisms which map corners to corners; (2) the structure is *reflective*, meaning that neighboring faces of $\beta\mathcal{T}$ are anticonformal reflections of one another in their common edge. We go into greater detail in §2.

1.3. Metric structure. We next endow \mathcal{T} with an intrinsic metric structure. This construction is attributed to Belyĭ in the case that \mathcal{T} is compact and is central to Grothendieck's *dessins d'Enfants* [17, 36]. It is also related to the construction of Klein's *j-invariant* on the upper half-plane. See [20] for background.

As a preliminary, break the Riemann sphere \mathbb{S}^2 into the two hemispheres bounded by the real axis. Each may be treated as a conformal triangle with vertices \bullet, \times, \square corresponding to boundary points $0, 1, \infty$, respectively. Treat the upper hemisphere as white, the lower hemisphere as gray. (Note: our convention is that stereographic projection of the plane to the sphere identifies 0 with the north pole of \mathbb{S}^2 , and ∞ with the south pole, so our hemispheres might well be called front and back.)

In the β -equilateral structure on \mathcal{T} , each face of $\beta\mathcal{T}$ is a conformal triangle with vertices \bullet, \times, \square , white or gray depending on orientation. Classical analytic function theory tells us that for every white (resp. gray) face f of $\beta\mathcal{T}$ there is a unique conformal mapping m_f of f onto the white (resp. gray) hemisphere of \mathbb{S}^2 for which $m(\bullet) = 0$, $m(\times) = 1$, and $m(\square) = \infty$. If f_1 and f_2 share an edge, $\bullet - \times$, $\times - \square$, or $\square - \bullet$, then $m_{f_1}(f_1)$ is, by the Schwarz Reflection Principle, the anticonformal reflection of $m_{f_2}(f_2)$ across the real axis through the segment $(0, 1)$, $(1, \infty)$, or $(\infty, 0)$, respectively. In consequence, the individual maps m_f for contiguous faces stitch together analytically, and we can define a locally one-to-one analytic function $m : \beta\mathcal{T} \setminus \mathcal{V} \rightarrow \mathbb{S}^2 \setminus \{0, 1, \infty\}$. As the points of \mathcal{V} are isolated singularities, m can be extended to a meromorphic function $M : \mathcal{T} \rightarrow \mathbb{S}^2$.

This meromorphic function, in turn, allows one to lift the spherical metric on \mathbb{S}^2 to a metric ρ on \mathcal{T} . If p, q are in the same white or gray face of $\beta\mathcal{T}$, then $\rho(p, q)$ is just the spherical distance from $M(p)$ to $M(q)$. For general points $p, q \in \mathcal{T}$, consider the family Γ of all paths γ in \mathcal{T} which connect p to q . Define $\rho(p, q) = \inf\{\text{length}(\gamma) : \gamma \in \Gamma\}$. It is straightforward to confirm that this defines a metric on the tiling; more precisely, a *complete length metric*, as defined, e.g., in [10, Chp 2].

Definition. Given \mathcal{T} , the meromorphic function constructed above is denoted $M_{\mathcal{T}}$ and the pull back of the spherical metric under $M_{\mathcal{T}}$, a *piecewise spherical* metric on \mathcal{T} , is denoted $\rho_{\mathcal{T}}$. Note that every tile of \mathcal{T} has $\rho_{\mathcal{T}}$ -diameter π .

We will not be using $M_{\mathcal{T}}$ or $\rho_{\mathcal{T}}$ often in this paper. However, meromorphic functions give a way to view tilings which may be exploited in the future, so we conclude this subsection by giving some context.

A meromorphic function m on a Riemann surface R has an *asymptotic value* $w \in \mathbb{S}^2$ if there exists a continuous path $\gamma : [0, \infty) \rightarrow R$ with $\gamma(t) \rightarrow \partial R$ as $t \rightarrow \infty$ so that $\lim_{t \rightarrow \infty} \gamma(t) = w$. The condition $\gamma(t) \rightarrow \partial R$ means that given any compact set $E \subset R$, there exists a t_E so that if $t > t_E$, then $\gamma(t) \notin E$.

Definition. A non-constant meromorphic function M on an open Riemann surface R will be termed a *Belyĭ map* under the following conditions:

- (1) M branches only at points of R lying over $\{0, 1, \infty\}$;
- (2) M has simple branching at all points over 1, and
- (3) M has no asymptotic values.

The pair (R, M) is then termed a *Belyĭ pair*.

This extends the terminology in §2.3 of [9] associated with Grothendieck *dessins*, since condition (3) is satisfied automatically when R is compact and condition (2) is necessary for the “clean” *dessins* of the theory. (One could also extend to allow bordered surfaces, but we omit the rather technical added conditions.)

Theorem 1.1. *Conformal tilings \mathcal{T} without borders are in one-to-one correspondence with Belyĭ pairs (R, M) ; namely, $\mathcal{T} \leftrightarrow (R, M)$ where the Riemann surface R is \mathcal{T} with the β -equilateral structure and M is the Belyĭ map $M_{\mathcal{T}}$.*

Proof. Given a conformal tiling \mathcal{T} of R , the function $M_{\mathcal{T}}$ defined earlier is easily seen to be a Belyĭ map on R . Conversely, suppose (R, M) is a Belyĭ pair. The preimages of the white and gray hemispheres of \mathbb{S}^2 under M decompose R into white and gray regions. One can use the maximum principle to see that each of these regions must be simply connected and then use the lack of asymptotic values to see that each has compact closure. The restriction of M to a region has no interior branch points, hence is a locally one-to-one proper map onto a hemisphere. In particular, each region is a curvilinear triangle with corner points mapping to $\{0, 1, \infty\}$ and by the Schwarz Reflection Principle is the conformal reflection of a neighboring region across each edge between these corners.

The local properties of analyticity guarantee that the set of white/gray regions is locally finite. The preimage of $[0, 1]$, the drawing $D = M^{-1}([0, 1])$, groups these regions into compact sets. These are the tiles that realize a conformal tiling \mathcal{T} of R . \square

1.4. Geometric structure. Geometry enters when our abstract structures on \mathcal{T} are realized in some concrete geometric setting. That setting will be denoted \mathbb{G} and in the simply connected cases we consider will always be one of the sphere, the plane, or the hyperbolic plane (\mathbb{S}^2, \mathbb{C} , or \mathbb{D}). Of course, these spaces all lie in \mathbb{S}^2 , but in specific cases the euclidean or hyperbolic geometry will be more appropriate.

Henceforth we work under the following **Assumption:** *We assume in the sequel that the tilings \mathcal{T} are simply connected.* We comment in Appendix B on how to leverage our work to handle general tilings.

Definition. Given combinatorial tiling \mathcal{T} , a *tiling map* is a continuous mapping $h : \mathcal{T} \rightarrow \mathbb{G}$ which is one-to-one on the interior of \mathcal{T} . Here, \mathbb{G} is one of \mathbb{S}^2, \mathbb{C} , or \mathbb{D} , and we do not assume the map is onto, though that will often be the case. The

images of the combinatorial tiles (i.e., the 2-cells) are concrete geometric tiles and their union is a geometric *realization* of \mathcal{T} in \mathbb{G} .

In the traditional setting, tiling maps are *ad hoc* — they have little significance since tile shapes are known in advance. In conformal tiling, on the other hand, the conformality of the map comes first: The map determines the tile shapes rather than the other way around.

Definition. A *conformal tiling* is the image of a combinatorial tiling under a conformal tiling map, $F : \mathcal{T} \rightarrow \mathbb{G}$, where \mathcal{T} has its β -equilateral structure and \mathbb{G} is one of the geometric spaces \mathbb{S}^2 , \mathbb{C} , or \mathbb{D} . The *tiles* are the images $F(T)$ of tiles $T \in \mathcal{T}$ — that is, they are concrete geometric regions in \mathbb{G} .

It is here in \mathbb{G} that we finally realize \mathcal{T} using concrete geometric shapes in parallel with traditional tilings. Note that with an abuse of notation we use \mathcal{T} to denote combinatorial, topological, β -equilateral, and geometric tilings depending on circumstances (and likewise for \mathcal{T}^\dagger and \mathcal{Q}). We want to study the concrete setting, so *tiles*, *vertices*, *edges*, *triangles*, and so forth will generally refer to their geometric realizations in \mathbb{G} .

This brings us to the question of the existence and variety of conformal tiling maps F available to us. Since we are considering the simply connected case, F can be obtained directly from the Uniformization Theorem: *For every simply connected Riemann surface R , there exists a conformal homeomorphism F from R onto one of \mathbb{S}^2 , \mathbb{C} , or \mathbb{D} . Moreover, F is essentially unique.* The term *essentially unique* will apply here and elsewhere when an object is unique up to conformal automorphisms (Möbius transformations) of \mathbb{G} . In any case, every simply connected combinatorial tiling can be realized.

Theorem 1.2. *Given any simply connected combinatorial tiling \mathcal{T} , there exists an essentially unique maximal conformal tiling with the combinatorics of \mathcal{T} in one of \mathbb{S}^2 , \mathbb{C} , or \mathbb{D} .*

The adjective “maximal” is needed only in the case $\mathbb{G} = \mathbb{D}$ as we will see shortly. Note that $\mathbb{G} = \mathbb{S}^2$ if and only if \mathcal{T} is finite and has no border, i.e., is a topological sphere. Two spherical examples are shown in Figure 5: one is a very regular tiling displaying soccerball-like dodecahedral symmetry, while the other is at the opposite extreme; namely, a randomly generated tiling. For display purposes (and aesthetics), tiles are color coded based on their number of edges.

If \mathcal{T} is not a sphere, then \mathbb{G} must be one of \mathbb{C} or \mathbb{D} . In the former case, \mathcal{T} is said to be *parabolic*, and in the latter case, *hyperbolic*. If \mathcal{T} has border edges, then \mathcal{T} is necessarily hyperbolic, and its uniformizing map F carries it to \mathbb{D} with any border edges mapped continuously to $\partial\mathbb{D}$. When \mathcal{T} is infinite and without border, determining whether \mathcal{T} is parabolic or hyperbolic is a version of the classical problem of “type”, an important topic to come.

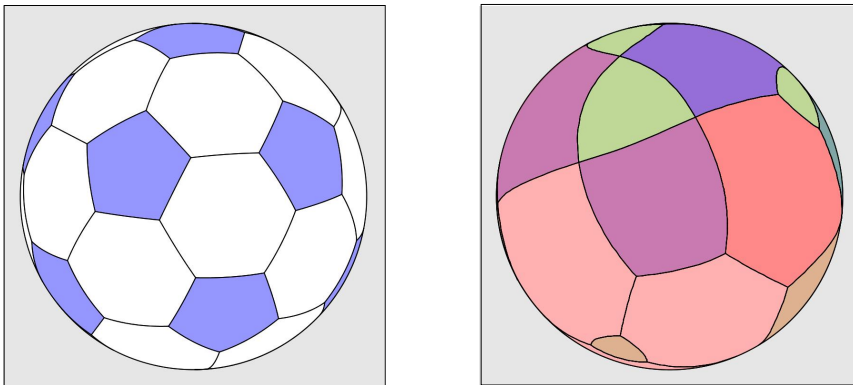


FIGURE 5. Tilings of the sphere, one very regular, one random.

Four apparently parabolic conformal tilings are illustrated in Figure 3 of the Introduction; parabolicity is established formally in Theorem 5.3. These are essentially unique (that is, up to scalings, rotations, and/or translations).

There is much greater variety when \mathcal{T} is hyperbolic. If Ω is any simply connected, proper, open subset of the plane, then by the Riemann Mapping Theorem there exists a conformal bijection $G : \mathbb{D} \rightarrow \Omega$. Composing with a uniformizing tiling map F , we get a new tiling map $G \circ F$ which realizes \mathcal{T} as a conformal tiling of Ω . If \mathcal{T} has non-empty border (and $\partial\Omega$ is sufficiently regular), $G \circ F$ extends continuously to the border, mapping it to $\partial\Omega$ (though not necessarily one-to-one). This flexibility is particularly useful with finite hyperbolic tilings, for by applying various “boundary conditions”, one can control the shape of the tiling’s realization. We provide a hyperbolic example \mathcal{T} in Figure 6. On the left is its maximal tiling, that obtained by a uniformizing map, while on the right it is realized as a square. (As with all examples, Figure 6 gives approximations only. Looking near the corners of the square, for example, you can see that this is, in fact, a finite tiling.)




FIGURE 6. Infinite hyperbolic conformal tilings sharing the same combinatorial pattern: the left is the associated “maximal” tiling, the right, a tiling filling a square.

2. THE INNER WORKINGS OF CONFORMALITY

When first encountered, conformal geometry presents an odd mixture of local rigidity within global flexibility. We must let go of the usual length and area considerations, straight lines, and corner angles, but we have angles (in the interior) and *Schwarz reflection*, we have notions such as *extremal length*, *harmonic measure*, and *brownian motion* as new hallmarks, and we have the powerful methods associated with *analytic functions* and *quasiconformal mappings*.

This section studies the workings and implications of conformality: the geometry of individual conformal tiles, conformal reflections and the meaning of shape, ensemble geometric effects, rigidity, and the roles played by our reflective structure. In the conformal setting, everything depends on limiting processes, and in Theorem 2.3 we provide a tool for handling the central geometric limit issues the topic encounters.

We assume here that \mathcal{T} has its canonical conformal structure and that $F : \mathcal{T} \rightarrow \mathbb{G}$ is a conformal tiling map. Since F is a conformal homeomorphism, most of what we say applies equally to the abstract conformal structure on $\beta\mathcal{T}$ and the conformal structure realized in \mathbb{G} .

2.1. Individual conformal tiles. First, a little background regarding conformality: Suppose $\psi : \Omega_1 \rightarrow \Omega_2$ is a homeomorphism between open subsets of Riemann surfaces. Then ψ is *conformal* if whenever smooth curves γ and σ intersect with angle θ at some point $z \in \Omega_1$, then $\psi \circ \gamma$ and $\psi \circ \sigma$ necessarily intersect with that same angle θ at $\psi(z) \in \Omega_2$. Likewise, ψ is *anticonformal* if it preserves the magnitude of such angles, but reverses their sense. Formulated in terms of complex analytic functions, ψ is conformal if and only if it can be represented in local coordinates as $w = g(z)$ where g is an analytic function (see, e.g., Ahlfors [1]), and anticonformal if and only if it can be represented in local coordinates as the complex conjugate $w = \bar{g}(z)$ of an analytic function.

A compact set P in a Riemann surface is termed a *conformal polygon* (a *conformal n -gon*) if there is a continuous map f from $\overline{\mathbb{D}}$ onto P whose restriction to \mathbb{D} is a conformal homeomorphism and maps finitely many points z_1, \dots, z_n in $\partial\mathbb{D}$ to the designated *vertices* (i.e., corners) of P . The images in P of the arcs $[z_j, z_{j+1}] \in \partial\mathbb{D}$ define the n *sides* of P . If the z_j are equally spaced around $\partial\mathbb{D}$, then P is said to be a *conformally regular n -gon*. Conformal n -gons P_1 and P_2 are said to be (*conformally*) *equivalent* if there is a homeomorphism $\psi : P_1 \rightarrow P_2$ whose restriction to the interior is conformal and which maps the vertices of P_1 onto the vertices of P_2 . **A strong cautionary note:** Conformal polygons realized in \mathbb{G} will often have well-defined corner angles where their sides meet. Conformality is about the interior, and these *boundary* angles are *not* necessarily preserved under conformal equivalence.

Let us now consider the β -equilateral surface \mathcal{T} . Triangles of $\beta\mathcal{T}$ got their conformal structures by being identifying with model unit-sided euclidean equilateral triangles. In particular, these are *conformal triangles* in this structure and their images under F are conformal triangles in \mathbb{G} .

These triangles do not live in isolation, however, and we must consider how they are attached. Suppose combinatorial triangles t_1, t_2 share an edge e . Placing their model euclidean triangles together in the plane as a quadrilateral with diagonal e , the two triangles are seen as normal euclidean reflections of one another in e .

The union $t_1 \cup t_2$ is mapped by F to the union of conformal triangles p_1, p_2 in \mathbb{G} sharing edge $F(e)$. Applying the classical Schwarz Reflection Principle to F , we conclude that p_1 and p_2 are anticonformal reflections of one another across $F(e)$. That is, there is a anticonformal involution of $p_1 \cup p_2$ which has fixed set $F(e)$ and interchanges the corners not adjacent to e . This implies that the interior of arc $F(e)$ is an *analytic arc*, the image of a real line segment under an analytic function with non-vanishing derivative. Henceforth, conformal triangles sharing an edge are always assumed to be anticonformal reflections of one another across that edge.

It is a fact of conformal mapping that every conformal triangle is conformally regular and hence any two conformal triangles are equivalent as conformal polygons. This applies as well to conformal 1-gons or 2-gons, but fails spectacularly for n -gons when $n \geq 4$, as these fall into a continuum of conformal equivalence classes for each n .

The reader will undoubtedly anticipate that every tile $T \in \mathcal{T}$ turns out to be a conformally regular n -gon for some n . We get additional information, however, by recognizing it as a chain of $2n$ conformal triangles, t_1, t_2, \dots, t_{2n} , all sharing the tile barycenter $v \in \mathcal{V}_\square$ as vertex. We can look at these in \mathbb{G} . Any one triangle will generate the full tile by Schwarz reflection as illustrated in Figure 7(a). Start with t_1 , for example. Reflections in the edges radiating from v , give successively t_2, t_3 , etc. Since the triangles result in accumulated angle 2π at v , it is clear that as the last triangle t_{2n} is placed, it will share an edge with t_1 , closing the chain. Figure 7(b) shows the analogous construction using only the 5 triangular faces for T in \mathcal{T}^* ; the conformal structures imposed on T are identical.




FIGURE 7. (a) Form a conformal n -gon of $\beta\mathcal{T}$ via reflection using $2n$ conformal triangles (10 in this case); (b) the identical conformal structure results in \mathcal{T}^* via reflection of n conformal triangles.

This representation of T highlights various conformal symmetries (leaving details to the reader). The n gray triangles are conformally equivalent (respecting the vertex markings), while the n white triangles are anticonformally equivalent to these, hence conformally equivalent among themselves. There is a conformal bijection $\gamma : T \rightarrow T$ of order n that rotates the faces by two notches around the barycenter (and hence maps the vertices of T to themselves). Just as with *regular* euclidean polygons, the existence of this n -fold rotational symmetry shows that T is a conformally regular n -gon. The barycenter $v \in T$ is fixed under the conformal rotations and is thus the *conformal center* of T .

Also, as in the euclidean case, there are orientation reversing symmetries. We can see these explicitly: Given any of the triangle edges radiating from the conformal center \square , the reflection across that edge which interchanges the neighboring triangles t_j, t_{j+1} clearly extends to an anticonformal bijection $\sigma : T \rightarrow T$. Its fixed points are an axis of T , an analytic arc formed by the union of two edges, either of type $\times - \square - \bullet$ if n is odd or type $\times - \square - \times$ or $\bullet - \square - \bullet$ if n is even.

Under composition, γ and σ generate a dihedral group D_{2n} of symmetries of T . Every map in this group respects the \bullet, \times , and \square labels: \square is the fixed conformal center, \bullet 's are the tile vertices, and \times 's are the conformal midpoints of the edges.

Similar considerations apply to construction of the dual tiles in \mathcal{T}^\dagger ; these are conformal polygons, but now with the \bullet 's as conformal centers and the \square 's as vertices. The tiles of the quad tiling \mathcal{Q} are only slightly different: Consider an interior quad. It is a *conformal square* (i.e., a conformally regular 4-gon) with an \times as conformal center. Since it has just four of our white/gray triangles, its symmetry group is $D_{2,4}$ and contains (generically) just two anticonformal reflections. One of these fixes the pair of edges of type $\bullet - \times - \bullet$. We knew these two were individually analytic arcs, but now see that their union is a single analytic arc. This $\bullet - \times - \bullet$ axis is actually the shared edge between two tiles of \mathcal{T} . Likewise, the other diagonal, of type $\square - \times - \square$, is a single analytic arc and the shared edge between two tiles of \mathcal{T}^\dagger . In particular, *each tile of \mathcal{T} is a conformally regular n -gon having edges which are analytic arcs. Likewise for tiles of \mathcal{T}^\dagger and \mathcal{Q} .*

We also note some explicit information about angles in \mathcal{T} . If k of its conformal tiles meet at a common interior vertex v , then $2k$ triangular faces of $\beta\mathcal{T}$ meet there. The conformal structure was defined in a neighborhood of v as a power function, $z \mapsto z^{6/2k}$. The k tile edges of \mathcal{T} radiating from v are therefore evenly spaced, separated by angle $2\pi/k$. Likewise, the $2n$ triangles forming an n -gon tile will all have angle π/n where they meet at the tile's conformal center. Since four triangles meet at each interior vertex of \mathcal{V}_\times , all have right angles there. In other words, each edge e of \mathcal{T} crosses its dual edge e^\dagger orthogonally. One could hardly imagine a more pleasing simultaneous embedding of a tiling and its dual.

Conformal reflections and symmetries are not as visually intuitive as their euclidean counterparts. Consider conformal "squares" for example. A Jordan region Ω in the plane with four distinguished boundary points is called a *conformal quadrilateral*; we can write $Q = \{\Omega; v_1, v_2, v_3, v_4\}$ where the corners v_i have counterclockwise order in ∂Q . Classical function theory provides a one-to-one conformal map g from Q to a euclidean rectangle R ; this map extends continuously to ∂Q and maps the corners to the corners of R . Then Q is a *conformal square* if and only if R is a euclidean square. Recognizing conformal squares among the conformal quadrilaterals is far from trivial. However, our construction of 4-gons in conformal tilings guarantees 4-fold rotational symmetry, so the quad tiles of \mathcal{Q} are automatically conformal squares.

2.2. Shape and reflection. In our discussion to this point it has been largely immaterial whether one works with tiles in the abstract Riemann surface \mathcal{T} or in a concrete realization of \mathcal{T} in \mathbb{G} . The latter will be quite striking, however, when the implications of analyticity are pointed out in §2.4.

Conformal tiles are rarely traditional polygons. Instead they are curvilinear regions with amazingly subtle shapes, as we shall see. Here, “shape” has a conformal meaning.

Definition. Tiles T and T' in \mathbb{G} have the same *shape* if they are *similar*; that is, if there is a conformal or anticonformal automorphism ϕ of \mathbb{G} so that $\phi(T) = T'$ and so that ϕ maps the vertices of T to the vertices of T' .

The term “reflection” in \mathbb{G} refers to the usual *anticonformal reflection* of complex analysis. However, it is convenient to extend the term to full tiles. As illustrated in Figure 7(a), a conformally regular n -gon is generated by repeated reflection of a white or gray subtriangle. We will call these “sector” triangles.

Definition. Suppose T and T' are conformally regular polygons with n and n' sides, respectively, which share an edge e . We will say that T' is a *reflection* of T across e if there is a sector triangle bordering e which generates T , and whose conformal reflection across e (in the usual sense) is a sector triangle generating T' .

This is the usual reflection when $n = n'$: In this case, there is, in fact, a global anticonformal homeomorphism $\phi : T \rightarrow T'$ fixing e . However, if $n \neq n'$, no such global map exists. Indeed, where the triangles of T meet at the conformal center, they form equal angles of π/n , whereas those meeting at the conformal center of T' form angles π/n' — thus, angle magnitudes are not preserved between the conformal centers.

2.3. Tile ensembles. As we move to tilings, we begin to see ensemble effects on tile shapes. Start with finite collections, which may be of interest in their own right or as stages in construction of infinite tilings. Figure 8 contrasts the traditional and conformal situations for finite “chair” tilings: Traditional tiles have fixed shapes and this limits the resulting global shapes. On the other hand, we can give a conformal tiling almost any global shape, but that affects the individual tile shapes. The first row shows three stages of the traditional construction, the second shows the corresponding maximal conformal tilings in $\overline{\mathbb{D}}$. The third and fourth rows are again conformal versions. In the third, the tilings fill rectangles, and in the fourth, they form euclidean chair global shapes. Conformality asserts itself: To get the rectangles and chair shapes, we have designated the angles of boundary vertices. But having done this, the edge lengths — the rectangle’s aspect ratio (extremal length) and the edge proportions of the chair — were forced on us by conformality.

In conformal tiling, individual and aggregate tile shapes continue to change as more generations of tiles are added. Before discussing these shapes, we would like to point out the importance of the “reflective” nature of our tilings, as distinct from their “conformal regularity.” For instance, it might be natural to impose a RPWA-structure on \mathcal{T} , so each combinatorial tile is identified with a regular, unit-sided, euclidean polygon (see Appendix A). This structure is conformally regular, but will be reflective if and only if all tiles have the same number of sides.

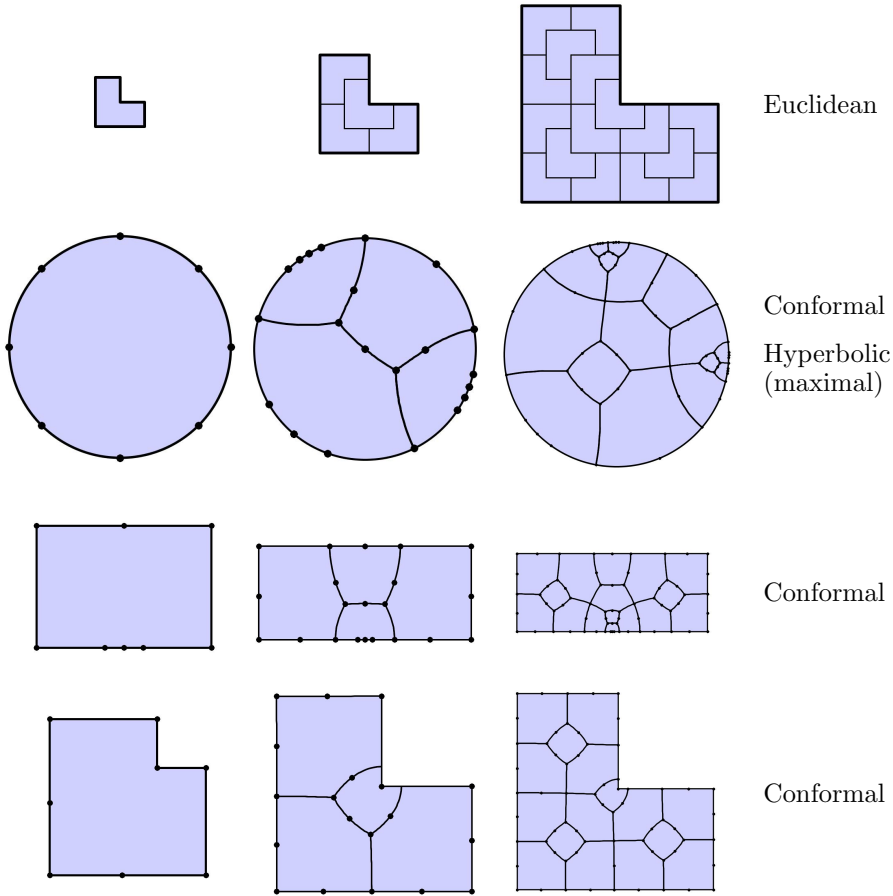


FIGURE 8. Three stages of a “chair” tiling in traditional euclidean and three conformal versions.

To illustrate, consider the most mundane of conformal tilings, the lattice of squares in \mathbb{C} with vertices on the Gaussian integer lattice. These are, of course, conformal squares reflective across the euclidean line segments as edges. Let the tiling be denoted \mathcal{T} and identify a root tile T_0 . Now, modify the combinatorics by using the diagonals of T_0 to break it into four 3-gons — leave the rest of the combinatorics of \mathcal{T} unchanged. Realize this new tiling conformally, then go in and reconsolidate the four 3-gons into a single tile. Denote this new tiling by \mathcal{T}' , and the reconstituted root as T'_0 .

Compare \mathcal{T} and \mathcal{T}' . They have identical combinatorics and the tiles of both are conformal squares — thus, they are **conformally regular** realizations of the same combinatorial tiling. However, \mathcal{T}' is, due to the single tile T'_0 , **not reflective**: If T'_1 shares an edge e' with T'_0 , the edge is an analytic arc, but T'_0 and T'_1 are not reflections of one another across e' . Figure 9 illustrates the gray and white sector triangles that give T'_0 and T'_1 their structures. Clearly, for example, reflection across the dashed edge does not identify the conformal centers of the two parent conformal squares.




FIGURE 9. An example of non-reflective but conformally square neighboring tiles.

The effect isolated in Figure 9 has global implications: Not a single edge of \mathcal{T}' will be a straight line segment — not a single tile of \mathcal{T}' will be a euclidean square, though they get closer and closer to being square as one moves away from the root. The important message here is this:

Fact. *Conformal realizations of tilings must be both conformally regular and reflective for the methods of this paper to apply.*

2.4. Rigidity. Among the key distinguishing features of conformal tilings — as we have defined them — is their global rigidity. This results entirely from the two local features we just discussed: (1) the individual tiles are conformally regular, and (2) tiles sharing an edge are reflections of one another (in the sense we have attached to “reflections” of tiles).

Theorem 2.1 (Shape Rigidity). *A conformal tiling \mathcal{T} is completely determined by the shape of any one of its individual tiles.*

In other words, if you have a single concrete tile which is known to be part of some conformal tiling \mathcal{T} , then that tile’s shape determines the rest of \mathcal{T} : the combinatorics, the geometric setting \mathbb{G} , be it \mathbb{C} , \mathbb{D} , or \mathbb{S}^2 , the shape and location of every other tile — all of \mathcal{T} .

Proof. Pick an arbitrary geometric tile T_0 from \mathcal{T} and suppose T_1 is a tile sharing an edge e with T_0 . The shape and location of T_1 is determined by the fact that it is a reflection (as a tile) of T_0 . Since every tile can be reached from T_0 through a finite chain of tiles, each sharing an edge with the previous, the shape and location of T_0 determine the shapes and locations of all tiles of \mathcal{T} . All of these, of course, determine the full tiling \mathcal{T} . \square

This result is in stark contrast to traditional tilings, where it is most often the case that every tile is similar to one of a finite number of *prototiles*, and where a tile of a given shape may occur in uncountably many distinct tilings. One of rigidity’s consequences is that similarities among tiles of a conformal tiling \mathcal{T} imply combinatorial symmetries in \mathcal{T} . The straightforward proof is left to the reader. A partial converse is given in Appendix B.

Corollary 2.2. *Let \mathcal{T} be a combinatorial tiling without boundary. Suppose that in some conformal realization of \mathcal{T} in \mathbb{G} two of its geometric tiles, T and T' , are*

similar. Then there exists a combinatorial automorphism of \mathcal{T} which identifies T and T' as combinatorial tiles.

2.5. Limits of tilings. Every stage in conformal tiling involves notions of convergence, from combinatorics, to tiles, to tilings, and finally to families of tilings — our long range interest. Tilings will be assumed to have designated roots, so convergence will be *rooted* convergence.

Combinatorial convergence uses a “big-ball” style metric, which is fairly standard in such situations. Each tiling \mathcal{T} is represented by its tiling graph D^\dagger . Write $B_{\mathcal{T}}(\eta, m) \subset D^\dagger$ for the (filled) ball of radius m centered at tile η . This is defined by starting with the subgraph spanned by nodes with graph distance at most m from η and then filling in any islands which this separates from the ideal boundary of D^\dagger . Because \mathcal{T} is simply connected, each such ball is simply connected. Fixing η , the balls $\{B_{\mathcal{T}}(\eta, m) : m = 1, 2, \dots\}$ are nested, finite, and exhaust D^\dagger .

We use these balls to define a metric ρ on the space \mathbf{RC} of rooted tiling graphs. If \mathcal{T} and \mathcal{T}' are combinatorial tilings with roots η and η' , then $B_{\mathcal{T}}(\eta, m)$ and $B_{\mathcal{T}'}(\eta', m)$ are isomorphic if there exists an orientation preserving graph isomorphism between them which identifies η and η' . Define the distance ρ from \mathcal{T} to \mathcal{T}' by

$$\rho(\mathcal{T}, \mathcal{T}') = e^{-m},$$

where m is either the largest integer for which $B_{\mathcal{T}}(\eta, m)$ and $B_{\mathcal{T}'}(\eta', m)$ are isomorphic or the symbol ∞ . This makes \mathbf{RC} into a complete metric space; details may be found in [5].

Definition. Combinatorial (rooted) tilings $\{\mathcal{T}^n\}$ converge to the combinatorial tiling \mathcal{T} if the distance $\rho(\mathcal{T}^n, \mathcal{T}) \rightarrow 0$ as $n \rightarrow \infty$. Write $\mathcal{T}^n \rightarrow \mathcal{T}$.

Note, in particular, that every combinatorial tiling \mathcal{T} is the limit of the finite subtilings associated with the nested balls $\{B_{\mathcal{T}}(\eta, m) : m = 1, 2, \dots\}$.

Moving now to convergence of associated conformal tilings and of individual tile shapes, conformality may appear to have a delicacy that would complicate things. However, by working in the setting of uniform convergence on compacta, the space of conformal tiling maps is, in fact, remarkably robust.

Theorem 2.3. *Suppose that $\mathcal{T}^n \rightarrow \mathcal{T}$ represents rooted combinatorial convergence of simply connected tilings to a simply connected limit tiling \mathcal{T} and that for each n , ϕ_n is a conformal tiling map for \mathcal{T}^n . Given any sequence of indices $\{n_k\}$, there is a further subsequence $\{n_j\} = \{n_{k_j}\}$ for which one of these holds:*

- (a) *The subsequence $\{\phi_{n_j}\}$ degenerates, meaning that it converges on \mathcal{T} to a constant function or to a function taking two values.*
- (b) *The subsequence $\{\phi_{n_j}\}$ converges uniformly on compacta of $\text{int}(\mathcal{T})$, the interior of \mathcal{T} , to a conformal tiling map ϕ for \mathcal{T} .*

In the latter case, $\phi(\mathcal{T})$ is a conformal realization of \mathcal{T} in one of \mathbb{C}, \mathbb{D} , or \mathbb{S}^2 .

Proof. Suppose first that \mathcal{T} is a sphere. Then the \mathcal{T}^n must be identical to \mathcal{T} for all sufficiently large n . As Riemann surfaces, the \mathcal{T}^n may be identified with \mathbb{S}^2 , meaning that the conformal tiling maps ϕ_n are conformal self-maps of \mathbb{S}^2 , that is, Möbius transformations. This is well known to be a normal family: Every

subsequence has a further subsequence that converges to a Möbius transformation or degenerates.

When \mathcal{T} is not a sphere, we can first simplify our situation. Each subsequence \mathcal{T}^{n_j} converges to \mathcal{T} in the big-ball metric, with root tiles $T_0^{n_j}$ identified with the root T_0 of \mathcal{T} . Choosing a further subsequence, if necessary, we may assume that the (filled, rooted) balls $B_{\mathcal{T}^{n_j}}(T_0^{n_j}, j) \subset \mathcal{T}^{n_j}$ are isomorphic to balls in \mathcal{T} . Cutting each \mathcal{T}^{n_j} down to the ball $B_{\mathcal{T}^{n_j}}(T_0^{n_j}, j)$ and redefining ϕ_{n_j} by restriction, we may therefore assume the following without loss of generality:

Assumption. *The sequence $\{\mathcal{T}^n\}$ consists of nested, finite, simply connected combinatorial tilings with common root T_0 which exhaust \mathcal{T} as n goes to infinity.*

This puts us in a normal families situation that is fairly typical. As a Riemann surface, if \mathcal{T} is not bordered, then it is conformally equivalent to \mathbb{C} or \mathbb{D} , while if it has a border, then $\text{int}(\mathcal{T})$ is conformally equivalent to \mathbb{D} . A note on terminology: To say that a sequence of functions $\{f_j\}$ on \mathcal{T} converges (or converges uniformly) on compacta means that for any compact set $E \subset \text{int}(\mathcal{T})$ there exists an open neighborhood Ω with $\bar{\Omega} \subset \text{int}(\mathcal{T})$ and $J > 0$ so that f_j is defined on Ω for all $j > J$ and $f_j(z) \rightarrow f(z)$ for all $z \in E$ (uniformly).

In preparation for later developments, we apply normal families from the quasiconformal mapping perspective, noting that conformal maps are κ -quasiconformal with $\kappa = 1$. We may treat the ϕ_n as maps to \mathbb{S}^2 . We have a number of successive subsequences to extract, and we will abuse notation by referring to each as $\{\phi_{n_j}\}$. Begin by labeling the conformal center of T_0 as the point α and choosing subsequence $\{\phi_{n_j}\}$ so that $\phi_{n_j}(\alpha) \rightarrow z_0$ for some point $z_0 \in \mathbb{S}^2$. If ψ is a fixed Möbius transformation with $\psi(z_0) = 0$, then because $\phi_{n_j}(\alpha) \rightarrow z_0$ we may choose a family ψ_j of Möbius transformations which converges uniformly on \mathbb{S}^2 to the identity and so that $\psi(\psi_j(\phi_{n_j}(b))) = 0$ for every j . It is clear that the conclusion of our theorem will follow if we prove it with the ϕ_{n_j} replaced by $\psi \circ \psi_j \circ \phi_{n_j}$. In other words, we may assume $\phi_{n_j}(\alpha) = 0$ for all j . It is well known in this situation that we can extract a further subsequence $\{\phi_{n_j}\}$ which either converges to a κ -quasiconformal map $\phi : \mathcal{T} \rightarrow \mathbb{S}^2$ or degenerates. If the limit does not degenerate, then since $\kappa = 1$, ϕ is conformal. Unwinding the auxiliary Möbius transformations, the same conclusion follows for the original maps. \square

The theory of quasiconformal mapping is crucial here. One feature may be hidden in the terminology; namely, that a non-degenerate limit function is necessarily one-to-one. Moreover, the images $\phi_n(\mathcal{T}^n)$ converge to $\phi(\mathcal{T})$ in a natural way *à la* the Carathéodory Kernel Theorem [24]. The possibility of degeneracy is also a feature, not a bug, as these examples show.

Example 2.1. The modes of degeneracy in families of quasiconformal mappings are modeled by these two Möbius families on the sphere: (a) Let $\phi_n : z \rightarrow z + n$, $n = 1, 2, \dots$. Each ϕ_n fixes ∞ , and $\phi_n(z) \rightarrow \infty$ uniformly on compacta of $\mathbb{S}^2 \setminus \{\infty\}$. Every orbit gets pulled into ∞ and the limit is constant, $\phi \equiv \infty$. (b) On the other hand, let $\phi_n : z \rightarrow nz$. These fix both 0 and ∞ , but $\phi_n(z) \rightarrow \infty$ uniformly on compacta of $\mathbb{S}^2 \setminus \{0\}$. Two-value degeneracy always occurs like this: Some point of the domain is mapped to one of the values, and the remaining points are swept to the other.

We avoid degeneracy by appropriate normalization. It is typical to require (for sufficiently large n) that the conformal center of the root tile, α , be mapped to the origin and that some other vertex, γ , be mapped, e.g., to $z = i$. In the case that \mathcal{T} is *spherical*, α, γ are mapped to 0 and 1, while the conformal center of a third tile is mapped to ∞ . In both situations, the ϕ_n then belong to a compact quasiconformal family and any limits we extract are non-constant and quasiconformal (conformal in our case).

Example 2.2. *The “type” problem:* Let \mathcal{T} be an infinite, simply connected combinatorial tiling without border with root T_0 . Define $\mathcal{T}^n = B_{\mathcal{T}}(T_0, n)$, $n = 1, 2, \dots$, so $\{\mathcal{T}^n\}$ is a nested sequence of finite simply connected combinatorial tilings, $\mathcal{T}^n \rightarrow \mathcal{T}$.

Let α denote the barycenter of T_0 , γ a designated corner. For each n let $\phi_n : \mathcal{T}^n \rightarrow \mathbb{D}$ be the uniformizing tiling map of \mathcal{T}^n (i.e., the maximal realization). Applying a conformal automorphism of \mathbb{D} we may assume that $\phi_n(\alpha) = 0$ and that $\phi_n(\gamma)$ is on the positive imaginary axis. Since the \mathcal{T}^n are nested, observe that if $m < n$, then the composition $\phi_m \circ \phi_n^{-1} : \mathbb{D} \rightarrow \mathbb{D}$ is analytic and fixes 0. The Schwarz Lemma from complex function theory implies that $|\phi_n(\gamma)| \leq |\phi_m(\gamma)|$. In other words, $|\phi_n(\gamma)|$ is decreasing in n . This gives us a fundamental dichotomy: either (1) $\phi_n(\gamma) \rightarrow 0$ as n grows or (2) there exists $\eta > 0$ so $|\phi_n(\gamma)| > \eta$ for all n . We know by the Riemann Mapping Theorem that there is a conformal tiling map ψ to one of \mathbb{C} or \mathbb{D} .

Case (1), \mathcal{T} is parabolic: Suppose ψ were to map to \mathbb{D} . We may assume $\psi(\alpha) = 0$. Defining ψ_n as the restriction of ψ to $B_{\mathcal{T}}(T_0, n)$, we can apply the Schwarz lemma to $\psi_n \circ \phi_n^{-1} : \mathbb{D} \rightarrow \mathbb{D}$ and conclude as above that $|\psi(\gamma)| \leq |\phi_n(\gamma)|$. Letting n grow, condition (1) implies $|\psi(\gamma) = \psi_n(\gamma) = 0|$, contradicting the one-to-one property of ψ . We conclude that \mathcal{T} is parabolic.

However, we can exploit the maps ϕ_n further by rescaling. Let $\lambda_n = 1/|\phi_n(\gamma)|$, noting that $\lambda_n \rightarrow \infty$. Replace each ϕ_n by $\lambda_n \phi_n$. Since these maps now omit ∞ and map α to 0 and γ to $z = i$, standard quasiconformal arguments imply that no subsequence can degenerate. The limit ψ of any convergent subsequence is conformal and, by Hurwitz’s theorem, one-to-one, so ψ is a conformal tiling map with $\psi(\alpha) = 0, \psi(\gamma) = i$. There is only one such map, so the full sequence of (rescaled) maps ϕ_n converges to ψ . This is, in fact, the standard approach to computing ψ .

Case (2), \mathcal{T} is hyperbolic: In this case, $\phi_n(\gamma) \rightarrow w$ for some $w \neq 0$ in \mathbb{D} . Again, no subsequence ϕ_{n_j} can degenerate. Any limit is one-to-one by Hurwitz’s theorem and maps onto \mathbb{D} by the Schwarz lemma. Therefore the full sequence $\{\phi_n\}$ converges to a maximal conformal tiling map ψ . In particular, \mathcal{T} is hyperbolic and we have an approach for computing ψ . Note that the bordered cases will always fall into case (2).

2.6. Observations. Creation of infinite tilings inevitably depends on some notion of convergence. In fact, absent symmetries, it is debatable whether one can actually conceive of infinite patterns — our minds must work through finite pieces and limits. Nevertheless, infinite traditional tilings have a feeling of concreteness. One may add tiles generation-by-generation, or apply successive stages of substitution and expansion, for example — but the growing finite patches are perceived as finished parts of a final infinite whole. Of course, this concreteness is illusory. Generically,

each finite patch belongs to *uncountably many* distinct final configurations — no matter how big a patch you look at, you can't be sure which global configuration you live in.

In conformal tiling, the situation is a bit reversed. A tile in \mathcal{T} doesn't even know its final shape until all of \mathcal{T} is in place. Yet by Theorem 2.1, the shape is not only (essentially) unique, but determines the rest of the tiles of \mathcal{T} as well. In practice, the shape evolves as finite subtilings of \mathcal{T} are realized (as in the previous example), and though it may seem less substantial during the process, convergence is surprisingly robust and fast. Even artificial boundary conditions along the way — as were illustrated, for example, with the chair in Figure 31 — fade quickly as a tile is surrounded and isolated from that boundary by more generations of tiles. Its “true” shape simply must emerge.

Ultimately, it is perhaps a matter of opinion if a traditional tiling or its conformal realization is more “concrete”.

3. CONFORMAL TILING EXAMPLES

As conformal tilings are determined entirely by their combinatorics, the supply of examples is inexhaustible. We begin this section with broad examples, an eclectic mix from extremely regular to completely random. We then turn to the main topic, tiling families parallel to those that have attracted so much interest in traditional tiling and in the work of Cannon, Floyd, and Parry. After recasting terminology and reviewing “subdivision rules”, we illustrate the conformal versions of six well-known traditional tilings. Particular attention is paid to the connections between the combinatorics and the geometry of conformal realizations and we summarized some of our observations.

This sets the stage for the examples of the most interest, subdivision tilings unique to the conformal setting. Key among these are tilings arising from “conformal subdivision rules”, rules under which conformal structures are invariant and which play a key role in [5].

3.1. Broad examples. These examples range from regular tilings which were known to the ancients to tilings appearing in recent research on dimers and random configurations.

Example 3.1. *(n, m)-Tilings:* We start by demonstrating that conformal tilings reproduce classical tilings in highly symmetric situations. The (n, m) -tilings are a convenient class of examples. An (n, m) tiling \mathcal{T} has regular n -gons as tiles, with m tiles meeting at each vertex. Figure 10 shows various spherical, euclidean, and hyperbolic examples. Each tile is a grouping of triangles from the $(n, m, 2)$ triangle group. These are traditional tilings in that the tiles are regular polygons with geodesic edges. In our approach, they are generated directly from their combinatorics.

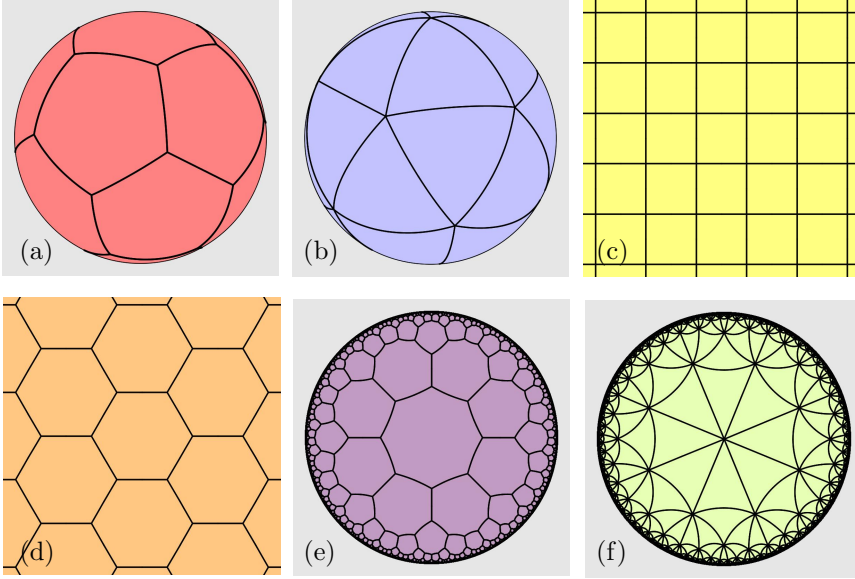


FIGURE 10. (a) The spherical dodecahedron, the $(5, 3)$ -tiling; (b) The spherical icosahedron, dual to (a); (c) The square or $(4, 4)$ -tiling in \mathbb{C} ; (d) The hexagon or $(6, 3)$ -tiling in \mathbb{C} ; (e) The octagonal or $(8, 3)$ -tiling in \mathbb{D} ; (f) The dual to (e).

We will be introducing discrete conformal tilings associated with circle packings in §4.3, but wish to note here that the (n, m) -tilings have so much symmetry that their coarse discrete conformal versions are already identical to the classical regular tilings.

Example 3.2. *Fusion Tiling:* We will shortly study subdivision tilings, but a related and powerful method for generating combinatorics is the notion of *fusion* tilings (see [14]). We illustrate with the Fibonacci tiling. Its fusion rule is illustrated in the transition from stage 0 to stage 1 in Figure 11; successive stage transitions take the same form. See [15, Figures 18 and 19].




FIGURE 11. The Fibonacci tiling fusion rule, along with discrete versions of the four tile types using parameters $3/1$ and $2/1$.

In the classical construction, side length compatibility in the pastings gives two degrees of freedom. We, of course, work only with combinatorics, and analogous parameters are available as ratios n/k and m/k of edge counts. The four tiles have dimensions: $a : n \times m$; $b : n \times k$; $c : k \times m$; and $d : k \times k$. Setting $n = 3$, $m = 2$, and $k = 1$ leads to the conformal Fibonacci tiling shown in Figure 3(a).

Example 3.3. *Dessins d’Enfants:* Alexander Grothendieck created the theory of *dessins d’Enfants*, drawings by children, to study algebraic number fields. The geometrically parallel discrete theory was developed by the authors in [9], initiating many of the themes we pursue here. In the theory of *dessins*, however, one works with finite drawings D on compact surfaces, and the surfaces may have positive genus. Two examples are illustrated in Figure 12. The left shows the tiling (and white/gray decomposition) for a genus 0 dessin, while the right shows the fundamental domain in \mathbb{D} for a genus 3 dessin. Each tiling is defined by dark lines in the figure; these are preimages $M_{\mathcal{T}}^{-1}[0, 1]$ under the meromorphic tiling map $M_{\mathcal{T}}$, as described in §1.3.




FIGURE 12. The left is a genus zero *dessin* on \mathbb{S}^2 , the right is a fundamental domain in \mathbb{D} for a *dessin* on a genus 3 surface.

The most remarkable feature of *dessins d’Enfants* is the algebraic nature of the surfaces involved. When tiling a compact surface of positive genus, for example, the β -equilateral structure picks out a particular Riemann surface — that is, a particular point in its moduli space. For this Riemann surface, there exists a defining polynomial $P(z, w)$ whose coefficients belong to some algebraic number field \mathbb{F} (a finite field extension of the rationals \mathbb{Q}). The converse also holds. Our genus 3 example in Figure 12 is a little anticlimactic. Shabat and Voevodsky [37, p. 217] have shown that this is the curve $y^3 = x^4 - 1$, a Picard surface, so \mathbb{F} is just \mathbb{Q} . Though the topic of conformal tiling subsumes *dessins d’Enfants*, whether the all-important algebraic implications can be generalized beyond the compact case has yet to be investigated.

Example 3.4. *Dimer Tilings:* In chemistry, a “dimer” is a polymer whose units consist of just two atoms. Mathematical models are typically parallelograms, dominoes, or lozenges; a sample dimer tiling is shown in Figure 1(b). It may take a moment to appreciate the subtleties of this pattern, but dimers turn out to be an impressively rich topic. (This particular example, due to Edmund Harriss, is dual to the 1D tiling $a \rightarrow ab, b \rightarrow c, c \rightarrow a$.) Dimer and related tilings are studied as polymer models in statistical mechanics, as perfect matchings in graph theory, as discrete integrable systems in differential geometry, and the list goes on. (See [21, 22].)

Conformally, dimers are simply quadrilaterals. As euclidean shapes, however, they come together in only a limited number of ways, so the local combinatorics are restricted. Conformal tiling perhaps suggests a new view of those combinatorics.

As every tiling \mathcal{T} has its associated quad tiling \mathcal{Q} , so goes the converse; that is, every conformal tiling by quadrilaterals is the “quad tiling” for some \mathcal{T} . Figure 13 displays a fragment of a euclidean dimer tiling, its conformal version, and the dual tiling. Note that the dual tiling contains n -gons for $n = 3, 4, 5$, and 6 only.




FIGURE 13. A euclidean dimer tiling, its conformal realization (colors correspond), and the dual tiling (colors based on degree).

Random tilings are a central theme in the study of dimers. Randomization is via “moves” which are a mixture of “fusion” and “subdivision”. Figure 14 illustrates a single dimer move. With the coloring, one naturally sees the tiling as a pile of blocks, and the move essentially takes one block off the pile. Below is the change in the dual tiling, where the move is more reminiscent of a move in the theory of braids.




FIGURE 14. A dimer “move” and its effect on the dual conformal tiling.

Example 3.5. *Random Tilings:* Random combinatorial constructs are now common in a range of pure and applied areas and they often have geometric associations. Though random *tilings* have not been studied explicitly, some active topics would qualify: random triangulations [2, 16], quadrangulations [23], and dimers [22], for example. In these and others, conformality is already a key feature. In [16] it is shown that a random triangulation of the plane is almost surely conformally equivalent to \mathbb{C} when given its equilateral structure, while in [18] it is shown that the associated circle packing is almost surely parabolic as well. (The possibility that these types were different is due to the lack of an upper bound on complexity in random situations.)

Random tilings, like that of Figure 5(b), may be created by mere aggregation in random triangulations. More interesting in the tiling world, however, might be

some notion of randomized subdivision processes. Since all these tilings can be realized conformally, they suggest a probabilistic approach to tiling spaces.

3.2. The main examples. A key motivation for conformal tiling has been pursuit of parallels to traditional tiling, in particular, to families of hierarchical aperiodic tilings and to their tiling spaces. We henceforth adhere largely to the framework in the traditional literature; namely, to simply connected infinite tilings of finite local complexity generated by subdivision rules. We begin by recasting terminology.

- **Tile Types:** In traditional tiling, each tile has a “type”, referring to its model prototile, perhaps with some distinguishing marks. In place of marks we use the general term “label”, which can be anything from a designated interior point, to boundary marks for enforcing attachment rules, to colors or abstract labels.

*In the conformal setting the **type** of a tile is the number n of its sides and perhaps an optional **label**. As part of its label, each tile has one corner designated as its **principal corner**.*

Labels are embellishments that do not directly affect tile geometry. They may be useful for normalization or pattern analysis, and necessary in some situations. For certain subdivision rules we investigate, “type” labels distinguish among tiles having the same number of sides (e.g., the “mixed” rule of Figure 24), and principal corners may be needed to orient the subdivision (e.g., the “domino” rule of Figure 15). Labels do not affect shapes for a given tiling, but may make all the difference under subdivision.

- **The big-ball metric:** Traditional tilings typically have designated “root” tiles, and part of the task of comparing two tilings involves drawing discs of some radius about their roots and comparing the patterns of tiles inside. Conformal tilings do not respect geometric distances, so the big-ball metric ρ , described in §2.5, is explicitly combinatorial. One consequence is that forming a ball about the root, $B_{\mathcal{T}}(\eta, m)$, may require filling in islands — there may be tiles enclosed by tiles which are combinatorially closer to the root. More subtly, there is no *a priori* isoperimetric inequality, even assuming finite local complexity: The number of tiles in a ball, though finite, is not necessarily bounded by some function of the number on the boundary. This will not affect our work here, but isoperimetric conditions play a role in [5].

- **Finite local complexity:** In traditional tilings, finite local complexity means not only that tiles have only finitely many types, but that tiles can be attached in only finitely many ways.

Conformal tiles intersect only at isolated vertices or along full edges or unions of full edges. So in the conformal setting, finite local complexity means simply that there is a uniform bound on the number of sides of any tile, on the number of tiles coming together at any vertex, and on the number of possible labels.

Finite local complexity is needed for compactness arguments and extends to larger ensembles of tiles. A finite set P of tiles taken from a tiling \mathcal{T} whose union has connected interior will be called a *patch*. Thus, finite local

complexity would imply, for example, that there are only finitely many patches having, say, one thousand tiles.

• **Repetitive:** A frequent hypothesis in traditional tiling is that any patch is repeated uniformly throughout the tiling. We phrase this in combinatorial form:

A combinatorial tiling \mathcal{T} is repetitive if given any finite patch P from \mathcal{T} there exists an m so that for every tile η the ball $B_{\mathcal{T}}(\eta, m)$ contains an isomorphic copy of P .

• **Aperiodic:** Maps between one tiling and another are generally intended as *tiling isomorphisms*; that is, as one-to-one maps that map full tiles onto full tiles, preserving type, and with corners going to corners. In the context of congruent shapes, a tiling is *aperiodic* if it has no translational symmetries, though other restrictions may be defined. In light of Theorem 2.2, however, automorphisms are the key now:

A conformal tiling \mathcal{T} is aperiodic if there are no infinite order conformal automorphisms of \mathbb{G} which define tiling isomorphisms of \mathcal{T} onto itself.

There is an additional nuance: A set of traditional prototiles is termed *aperiodic* if, in fact, there do exist tilings \mathcal{T} of the plane by tiles of these types, but that such tilings are necessarily aperiodic. Among our traditional aperiodic examples, both possibilities are represented: tilings whose prototiles are aperiodic and others whose prototiles are *not* aperiodic (though the tilings are). This nuance carries over to the conformal case; conformal realization of a periodic tiling will, by an easy application of our uniqueness results, be periodic as well. Existence of an aperiodic prototile set having a single (simply connected) prototile remains a major open question in traditional tiling, but is not relevant in conformal tiling.

• **Normalization:** In traditional tiling, normalization is *via* rigid motions, placing a designated tile in a standard position and orientation. In the conformal setting complications of scale, shape, and geometry enter.

A conformal tiling is normalized in a geometry \mathbb{G} by applying a conformal automorphism to place the conformal center of a designated tile at the origin and its principal corner on the positive imaginary axis.

Additional normalizations are typical in \mathbb{C} and \mathbb{S}^2 :

- (a) If $\mathbb{G} = \mathbb{C}$, then one scales by $t > 0$ so that the principal corner is at $z = i$.
- (b) If $\mathbb{G} = \mathbb{S}^2$, then one applies a Möbius transformation to \mathbb{S}^2 that puts the centroid of the conformal centers of the (necessarily finitely many) tiles at the origin in \mathbb{R}^3 . This tends to be visually pleasing.

From the normalized position, a conformal isometry of \mathbb{G} can move a specified point to the origin. In \mathbb{C} we will limit ourselves to translations. In \mathbb{D} and \mathbb{S}^2 , however, there are infinitely many conformal isometries moving that given point to the origin, so some further normalization may be needed. On the sphere there is also the possibility that \mathcal{T} has only one tile, which occurs when the associated drawing $D \in \mathbb{S}^2$ is a tree. In this case, one might, e.g., arrange

so the centroid of the vertices (there must be at least two) is at the origin in \mathbb{R}^3 .

Various other terminology carries over from the traditional setting in a straightforward way, though one must recall that combinatorics takes precedence over geometry and shape.

3.3. Finite subdivision rules. Tiling theory is largely about families of tilings, and these are typically generated *via* subdivision operators. The basic notion of subdivision is straightforward. One tiling is obtained from another by subdividing its tiles. The process is purely combinatorial, but it is convenient to use cell complexes.

Definition. Let \mathcal{T} and \mathcal{T}_s be combinatorial tilings of the same surface S . We say \mathcal{T}_s is a *subdivision* of \mathcal{T} if the identity map of S is a cellular map from \mathcal{T}_s to \mathcal{T} and we write $\mathcal{T}_s \leq \mathcal{T}$ or $\mathcal{T} \geq \mathcal{T}_s$. Tiles of \mathcal{T}_s are known as *subtiles* of the tiles of \mathcal{T} ; tiles of \mathcal{T} are *aggregates* of the tiles of \mathcal{T}_s .

In traditional tiling, subdivision tilings (also termed *substitution* tilings) involve rigid geometry. Each tile is similar to one of a given set of prototiles and is broken into subtiles, each again similar to one of these prototiles. Typically, there are only finitely many prototiles and subdivision rules, ensuring finite local complexity. Even in this setting, however, building subdivision tilings is an art — examples are prized. Conformal tiling, on the other hand, is driven by combinatorics, where subdivisions are (comparatively) easy to generate and tilings can always be realized.

Definition. A *finite subdivision rule* τ is, in fact, a finite collection of rules, one associated with each of a finite collection \mathfrak{T} of tile types. Each rule describes how its associated tile is to be subdivided into a finite pattern of subtiles, each again of a type in \mathfrak{T} .

Applying τ as an operator to an appropriate initial tiling \mathcal{T} breaks each tile into its subtiles, forming the new tiling $\tau\mathcal{T}$, $\mathcal{T} \geq \tau\mathcal{T}$. Repeated application leads to an infinite sequence

$$\mathcal{T} \geq \tau\mathcal{T} \geq \tau^2\mathcal{T} \geq \dots \tau^n\mathcal{T} \geq \dots$$

One now chooses root tiles in each combinatorial tiling and *via* the big-ball metric and diagonalization, extracts convergent subsequences. This generally leads to whole *tiling spaces*, and is the basis for many classical tiling families. We illustrate conformal versions of some familiar traditional examples here; natural issues about the rules and their tiling spaces are the topics of [5], and we will not go into detail now.

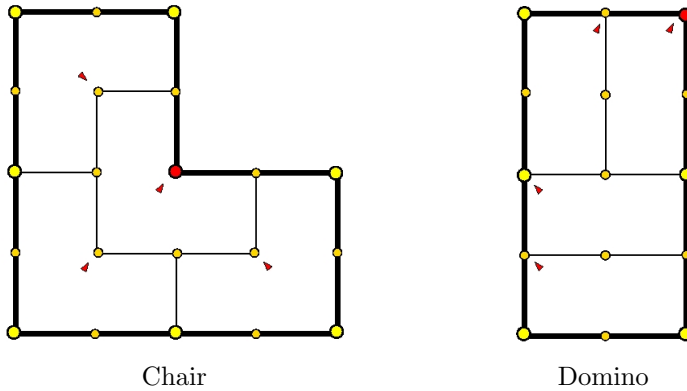


FIGURE 15. The “chair” and “domino” subdivision rules, each having only one tile type.

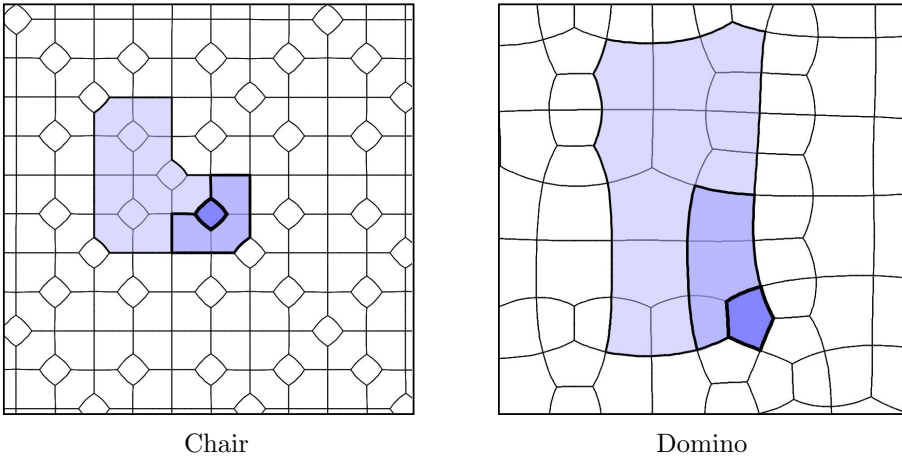


FIGURE 16. Tiles and aggregates in conformal “chair” and “domino” tilings.

3.4. Traditional subdivision rules. We consider conformal tilings associated with six traditional subdivision rules. Our attention attaches, in particular, to the shapes of individual and aggregate tiles, and ultimately to the emergence of geometric features which are somehow intrinsic to the subdivision rule. After the examples, we gather observations in §3.5.4 and later in §4.5.1.

We start with the well-known traditional subdivision rules shown in Figures 15 and 17. Tile vertices are the red and yellow dots, the principal corner being red; small yellow dots and thinner lines are vertices and edges introduced upon subdivision, red marks pointing to principal corners of the subtiles.

Example 3.6. *Conformal “chair”:* The single prototile is nominally a six-sided polygon. During subdivision, however, the corners of neighboring subtiles will meet

two of those sides, and since tiles must meet in full edges, this forces introduction of two additional vertices. In combinatorics, then, the single tile type is a combinatorial octagon.

A traditional chair tiling is shown in Figure 1(d), and its conformal tiling version in Figure 3(d). One may stare at the latter for some time before realizing that it is not the typical ceramic tiling one might buy for a countertop: Some tiles appear to be square, while others have 5 or 6 sides. In fact, none of the tiles is precisely square, there are no straight lines. None of the tiles appears to be octagons; this occurs because every tile shares contiguous edges with at least one neighbor, and these edges blend as parts of a common analytic arc, obscuring one or more “corners”.

Figure 16(a) shows the detail of the tiling and sets the pattern of later illustrations. An individual tile is highlighted along with aggregate tiles one and two levels above it. Optional additional vertices could be added to the chair subdivision rule. Edge compatibility issues force identical subdivision of each of the 8 edges so, for example, one could add a vertex to each of the chair’s 8 sides, making them combinatorial 16-gons. However, this combinatorial flexibility makes no difference in the conformal structure — 16-gons or 8-gons result in the same geometric tiles.

Example 3.7. *Conformal “domino”:* The domino rule of Figure 15 is one of many domino-type subdivision rules. Euclidean versions are regular square lattices with a pattern of horizontal and vertical cell unions. Each cell union is conformally a rectangle of aspect ratio 2:1, and the combinatorics of the cell unions affects the global shape.

We note that the domino prototile is *not* aperiodic (note the double negative) in tiling terminology, as it can be used to create an infinite tiling which is periodic. However, any *tiling* extracted by repeated subdivision is an aperiodic *tiling*. The same is true of the chair.

Next, we consider the subdivision rules of Figure 17 having multiple tile types, the most famous being that associated with the celebrated Penrose tiling. (The tile types are indicated; for historical reasons, type numbering starts at “4”.)

Example 3.8. *Conformal “sphinx”:* The sphinx tiling involves two tile types, one the reflection of the other. This example is instructive for those wishing to define their own subdivision rules, since ensuring consistency is a challenge. The prototiles have six (6) vertices; the need for the vertex on the base can be seen in Figure 17, where subtile #1 has a vertex falling in the base of subtile #3. It is very tempting (from personal experience!) to place that base vertex where the three subtiles come together. However, that leads to an inconsistency when a second subdivision is attempted. The placement chosen for the vertex (below the principal (red) vertex) leads to consistency, meaning one can subdivide *ad infinitum*. A tile and aggregates are shown in Figure 18.




FIGURE 17. The “sphinx” and “pinwheel” subdivision rules, with two tile types each, and the “Penrose” subdivision rule with four tile types.

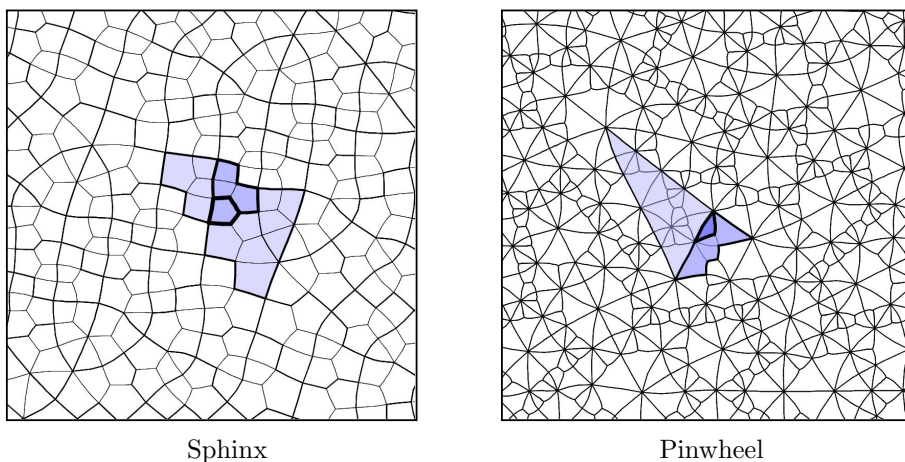


FIGURE 18. Tiles and aggregates in conformal “sphinx” and “pinwheel” tilings.

Example 3.9. *Conformal “pinwheel”*: The pinwheel is based on a triangle decomposition attributed to John H. Conway and is notable in traditional tiling for being the first to realize infinitely many tile orientations [29]. The subdivision requires a fourth vertex on the triangle’s longer leg, making these quadrangle tilings. A tile and its aggregates are shown in Figure 18.




FIGURE 19. For the conformal “Penrose” tiling, (a) shows a tiling with nested aggregates shown as usual. In (b), however, pairs of tiles forming kites are shaded blue, pairs forming darts are shaded red.

One finds interesting new patterns in the conformal pinwheel. Extended smooth arcs might arise on those occasions when (aggregate) tiles of types 4 and 5 attach *via* reflection across their “longer” legs. One also sees, however, web-like filaments of smaller tiles. These seem to start when two type 4’s (or type 5’s) meet in the subdivision along their hypotenuses, attaching not by reflection but rather by

rotation. You see a jaggedness in the chains of analytic arcs, which seems to grow with subdivision. Patterns within patterns within patterns.

Example 3.10. *Conformal “Penrose”:* A common form of the famous Penrose tiling was shown in Figure 1. This involves two prototiles, the so-called “kite” and “dart”. To formulate it as a subdivision tiling, however, we use the four “Robinson” tiles of Figure 17: kite and dart are each broken into two halves, mirror images of one another. In illustrating the conformal version in Figure 19, the halves were combined in pairs to reconstitute kites and darts, and though the darts are not very dart-like, the overall pattern is easily reminiscent of Figure 1(a). We will have more to say about shapes in a moment.

Our last example puts us in hyperbolic geometry. Certainly the most familiar, even famous, hyperbolic tilings are those in M. C. Escher’s “Circle Limit” prints, including his “Angels and Devils”, [25]. These get their hyperbolic character due to underlying symmetry, but more general combinatorial tilings can easily turn out to be *hyperbolic* as well, meaning their maximal conformal tilings lie in the hyperbolic rather than the euclidean plane.

Talking about tilings obtained via subdivision, however, is another matter. The next example will be called the “hyperbolic pentagonal” tiling — in some sense, all traditional hyperbolic examples trace to the “doubling” feature of this rule.

Example 3.11. *Conformal “hyperbolic pentagonal”:* The subdivision rule for this tiling is given in Figure 20. Two conformal realizations are shown in Figure 21; on the left is one employing the minimal number of vertices, on the right, one using the added (green) vertices.

Though the subdivision rule has two tile types, a type 5 tile is its own subdivision. This can make it a challenge to understand the diagonalization argument that leads to the infinite example shown; all the tiles remaining in the limit are type 5, the type 4’s having receded to the ideal boundary (along the real axis). The tiling is unchanged under subdivision, and we do not show aggregates because each tile is its own aggregate.




FIGURE 20. The “hyperbolic pentagonal” subdivision rule.

The tilings of Figure 20 lie in the upper half-plane, but could equally well be realized in the unit disc. Traditional versions of this tiling can be realized using

isometric copies of a single hyperbolic pentagonal prototile (see, e.g., [27, 28]), and this may appear to be such a tiling. However, as usual, in the true conformal version, every single tile has a unique shape determined by its place in the global pattern. There are no combinatorial symmetries here.

We point out that the 1-skeleton of this tiling appears as a “slice” within the Cayley graph for the Baumslag-Solitar group $B(1, 2)$. Its fascinating and subtle combinatorics are investigated further in [5] and [26].

3.5. Observations. Each of the traditional cases above has its own character — one of the attractions in studying tilings. Below are some observations.

3.5.1. A cautionary note. A warning is in order regarding experimental images: All of ours have been created as described in §4, so they represent computational approximations to finite portions of infinite tilings. Though they stimulate the intuition, many features have yet to be proven. Take, for example, the kites or darts of the conformal Penrose tiling. Figure 19 may suggest that they have comparable area as one moves around the tiling. But this is not known, they may well grow, shrink, or oscillate in size as one moves away from the root. There are, in fact, uncountably many conformal Penrose tilings, and it is not even clear that the tiles would share the same behavior in this regard in all cases. Even the conformal chairs of Figure 16, visually so regular, may yet hold surprises. These are not issues one encounters with the traditional versions of these tilings but must now be taken into account. Note that for the “conformal” subdivision rules we introduce in the next subsection, there is some chance of provable regularity, as was found in analysis of the regular pentagonal tiling in [8, 34].

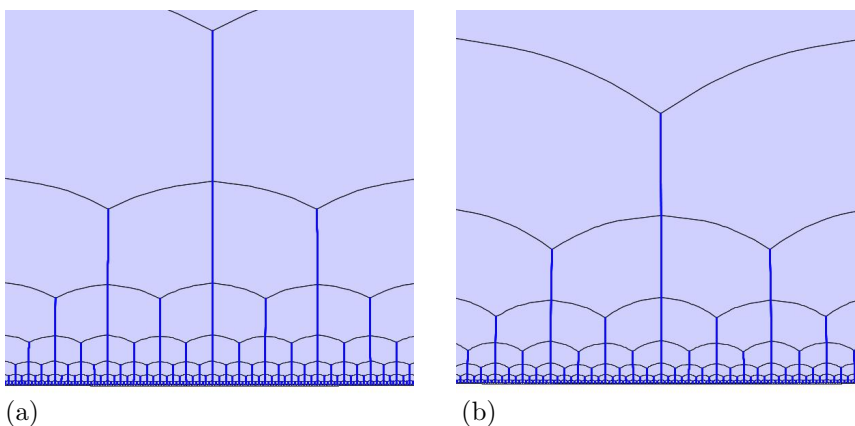


FIGURE 21. Two “hyperbolic pentagonal” tilings with identical combinatorics: (a) is realized when both tile types are pentagons and (b) when type 4 is a 10-gon and type 5 is a 14-gon (using the added green vertices of Figure 20).

3.5.2. *Combinatorial choice.* Conformal versions of tilings depend crucially on combinatorics, so the first task in mimicking a traditional tiling involves combinatorial choices. In the chair tiling, for instance, one could introduce an additional vertex in each edge, making the tiles 16-gons. In this case, the conformal structure is unchanged. In some of our rules we have illustrated optional vertices in green. It is natural, for example, to take the Robinson tiles used for the Penrose tiling in Figure 17 as 3-gons. However, a fourth vertex can be added to the type 4 and 5 tiles (the large green vertices and the small green ones that come with subdivision). Using these optional vertices gives a different tiling in this case, though the effects are nearly invisible. In contrast, for the hyperbolic pentagonal tilings in Figure 21, the optional vertices make a notable geometric difference.

3.5.3. *Type:* For all but the last of our traditional subdivision rules, the traditional realizations fill the euclidean plane. Our images suggest that the conformal versions do the same; that is, that they are parabolic. This is, in fact, the case as confirmed in Theorem 5.3. If a traditional tiling of finite local complexity fills the euclidean plane, then a conformal version will necessarily be parabolic. Type is a particularly important issue in conformal tiling. It is paramount in the seminal work on subdivision rules by Cannon, Floyd, and Perry, for example (see [12, 13]). We discuss it further below and as a central theme in [5].

3.5.4. *Aggregates.* The aggregate tilings in the examples we derived from traditional tilings are not themselves conformal tilings. This is clear from the various aggregate images we have displayed; that is, if you look at any of the aggregate tiles you can see that its sides (aggregated from sides of its subtiles) are not analytic arcs; neighboring aggregate tiles will not be reflections of one another. This is in sharp contrast with what we will see shortly in §3.7.

3.6. **Aggregate shapes.** Unlike with traditional tilings, conformal tiles take innumerable shapes — indeed, nearly every one in a tiling takes a different shape. The experiments described above have enabled us to look at samples of aggregate tiles, and one might speculate about the existence of limit shapes, though keeping in mind the experimental limitations that exist.

We have gathered a few examples in Figure 22. Every tile resulted from some subdivision operator τ applied to its parent tile. The reverse operation, aggregation, may be denoted τ^{-1} . Though existence of global aggregations is a subtlety we put off until our second paper, we may use the notation $\tau^{-n}\mathcal{T}$ for n stages of aggregation when they exist. In Figure 22 we show isolated tiles from our examples at varying levels of aggregation.

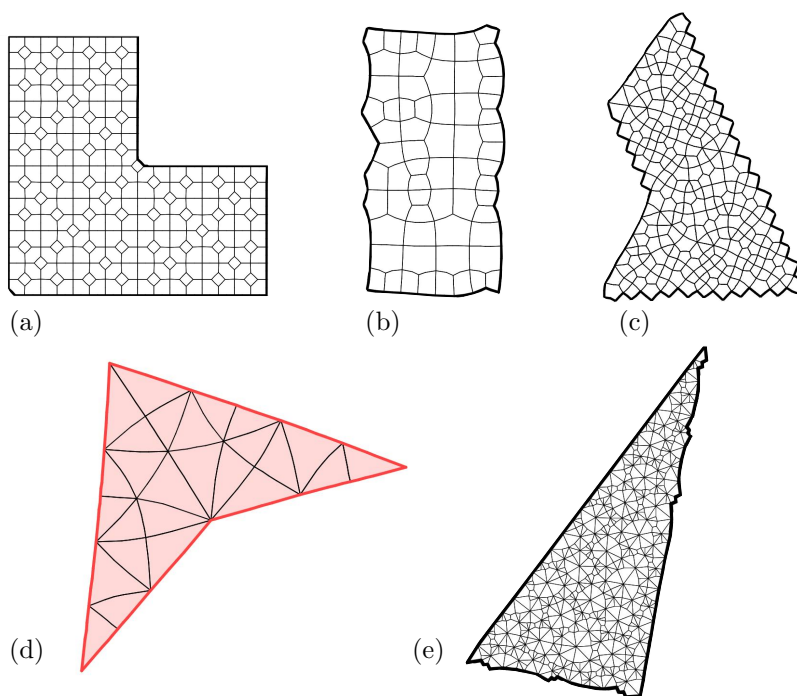


FIGURE 22. Aggregate tiles: (a) the “chair”, 4 levels; (b) the “domino”, 3 levels; (c) the “sphinx”, 4 levels; (d) A Penrose dart, 3 levels (union of two aggregate tiles); and (e) the “pinwheel”, 5 levels.

Let’s start with the Penrose “dart”, a union of two aggregate Robinson tiles. In the tiling itself, the darts are hardly dart-like — in fact, they seem to be convex. But already by the third aggregation shown here, the traditional euclidean shape seems to be emerging; the angle is measured at 2.545459 compared to $(4/5)\pi \sim 2.5133$ in the traditional prototile.

Likewise, conformal chair tiles are not chair-like, but by the level-5 aggregation, their only flaws appear to be corner cutoffs, which seem to get relatively smaller in each successive aggregation. The domino is looking more regular and closer in aspect to 2.0, the sphinx is looking more sphinx-like, with the jaggedness becoming finer as we aggregate.

The pinwheel tiling is another interesting case. Its prototiles are right triangles with side lengths in proportion $[1 : 2 : \sqrt{5}]$. A conformal pinwheel tiling, on the other hand, knows nothing of $\sqrt{5}$, it only knows combinatorics. Nevertheless, in the nested aggregates in Figure 22(e), the triples of tile corners form triangles with the side length proportions shown in this table, starting with level-7 (the base tile) and ending with the level-4 aggregate. One must ask: *Is $\sqrt{5}$ emerging from the combinatorics?*

TABLE 1. Edge proportions of triangular pinwheel tiles.

Tiling	Ratios		
Original Tile	1	:	0.90914 : 1.54389
Aggregate level-1	1	:	1.66243 : 1.86103
Aggregate level-2	1	:	2.20217 : 1.95415
Aggregate level-3	1	:	2.00517 : 2.23236
Traditional	1	:	2 : 2.23607

Question. *Is it the case that in any of the conformal versions of subdivision tilings we have considered — or perhaps more generally — the aggregate conformal tiles will converge in shape to their traditional prototiles?*

Here we would appropriately normalize the aggregate tiles and look for convergence, for example, in terms of the Hausdorff metric as the aggregation level n grows. If the answer were yes, then for that subdivision rule τ the combinatorics would have somehow encoded all the traditional geometric information, and one could likely extract a convergent subsequence from $\{\tau^{-n}\mathcal{T}\}$ whose limit is the traditional tiling. This would be a most pleasant surprise.

3.7. Conformal subdivision rules. We now leave traditional tiling examples to explore the wider landscape available in the conformal category. On one hand the combinatorial flexibility is liberating, while on the other, rather daunting. However, conformal methods originated in [8] and [9] to address concrete issues and new tiling themes are already emerging.

In departing from the classical motifs, we forgo the seminal “dodecahedral” tiling rule which set Cannon, Floyd, and Parry on their deep study of finite subdivision rules — and hence led to this paper. (For beautiful images, see [12, 13] and for the connection to Cannon’s Conjecture see [11].) Instead we start with two examples, the “lace” and “mixed” rules, intended to bring out a key distinction.

Example 3.12. *Conformal “lace”:* The lace rule is shown in Figure 23. A tile and its aggregates are shown in Figure 25. Though we give three tile types, these share an underlying elementary subfragment that could apply to any n -gon; this appears with others in Figure 27.




FIGURE 23. The “lace” subdivision rule with multiple tile types.

Example 3.13. *Conformal “mixed”:* The mixed finite subdivision rule of Figure 24 (see [13]) is a bit more complicated than we have encounter previously. Note, especially that there are multiple types having the same numbers of edges, so the principal vertices, marked with red, and the types must be taken into account for

subdivision consistency. A tiling with a highlighted tile and its aggregates is shown in Figure 25.

This tiling has been of particular interest to Cannon, Floyd, and Parry and is one we hope to address in future work on tiling spaces. The fragment in Figure 25 gives a sense of competition between the rule’s hyperbolic/parabolic tendencies. In some regions a “tripling” behavior, similar to the doubling in the hyperbolic pentagonal example, has a hyperbolic feel, yet in others there is an appearance of pseudo-regular grids which suggests parabolic.

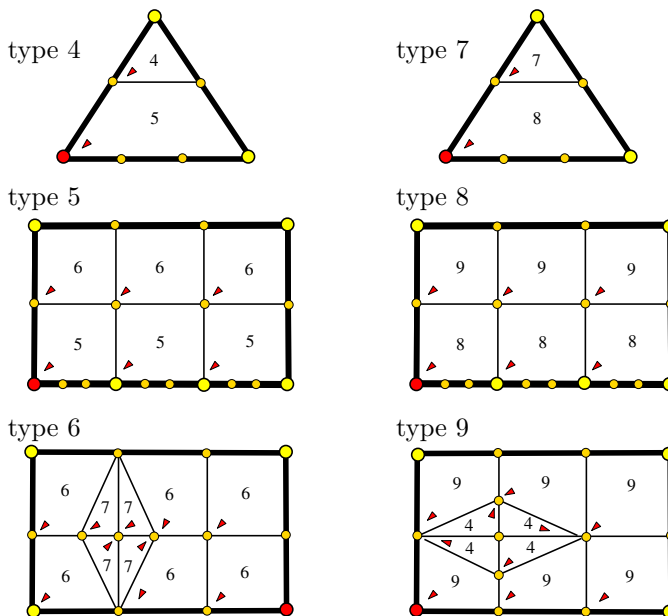


FIGURE 24. The “mixed” subdivision rule (due to Cannon, Floyd, and Perry) with multiple tile types.

We draw the reader’s attention to a fundamental contrast in behavior between the lace and mixed tilings of Figure 25. Notice the aggregates of the lace tile on the left in Figure 25: The conformal structure a tile gets as an aggregate appears to be precisely what it would get in its own tiling. In particular, for example, the aggregate edges are analytic arcs. Compare this to the aggregates of the mixed tile on the right. These aggregates clearly have edges which are not analytic (though, of course, they are piecewise analytic); the conformal structure a tile gets as an aggregate is not what it would get as a tile within its own generation. You see this as well in all the aggregates of Figure 22.

To codify this compatibility, suppose that \mathcal{T} is a combinatorial tiling and \mathcal{T}' is a combinatorial subdivision, $\mathcal{T}' \leq \mathcal{T}$ (obtained, for example, by applying a subdivision operator). Both \mathcal{T} and \mathcal{T}' are cell decompositions of a common topological surface S . We may define the *forgetful* homeomorphism, $\pi : \mathcal{T}' \rightarrow \mathcal{T}$, which is just the identity map from S to itself which forgets the subdivision. Each of \mathcal{T}' and \mathcal{T} imposes a β -equilateral structure on S . Are these the same?




FIGURE 25. Tiles and aggregates in the conformal “lace” and “mixed” tilings.

Definition. A subdivision $\mathcal{T}' \leq \mathcal{T}$ is said to be a *conformal* subdivision if the forgetful map $\pi : \mathcal{T}' \rightarrow \mathcal{T}$ is a conformal homeomorphism when \mathcal{T}' and \mathcal{T} are given their β -equilateral structures. A subdivision rule τ is said to be a *conformal subdivision rule* if $\tau\mathcal{T}$ is a conformal subdivision of \mathcal{T} whenever τ can be applied to \mathcal{T} .

Observation. With conformal subdivision operators τ we recapture a **key feature** of traditional tiling; namely, the subdivision $\tau\mathcal{T}$ is obtained by subdividing the conformal tiles of \mathcal{T} *in situ*.

Figure 26 may bring home this point (besides suggesting why we called this rule “lace”). It shows four levels of lace subdivision starting with a conformal square parent tile. If we were to subdivide yet again and impose the global conformal structure from this new tiling, the shapes of the aggregates at the first four levels would not change in the slightest. Indeed, one can conceive of an infinite succession of subdivisions, every point of the square being uniquely associated with some nested sequence of ever-smaller tiles.

(Note that Cannon’s Conjecture, [11], involves a combinatorial notion for “conformal subdivision”. This and our notion are conceptually related, but distinct; for example, the “twisted pentagonal” tiling illustrated in [8] is conformal in Cannon’s terminology but not in ours.)




FIGURE 26. Four levels of the “lace” subdivision.

We conclude this section by accumulating some sample conformal subdivision rules in Figure 27. All are dihedrally symmetric. They may be applied to arbitrary n -gons, so we show each rule applied to a sector and then to a representative n -gon. Those on the left are common in studying triangulations and the proof of their conformality is routine (see, e.g., [8]). The “barycentric” rule, denoted β as an operator, and the “hexagonal” rule, denoted \mathfrak{h} , are part of the standard machinery of conformal tiling. We used β in §1.1, converting \mathcal{T} into $\beta\mathcal{T}$; we use β again along with repeated applications of \mathfrak{h} in the coming discrete theory.

Regarding the rules on the right side of Figure 27, the proof that they are conformal and the study of associated tiling hierarchies are among the central topics of [5]. The first rule on the right generates the regular pentagonal tiling of Figure 2, the second is the lace rule (see Figure 26), and the last is the rule behind the snowball, our final example.




FIGURE 27. Conformal subdivision rules: Each applies to any n -gon, $n = 1, 2, 3, \dots$, so we apply it to a sector first, then to a representative n -gon.

Example 3.14. *Conformal “snowball”*: The snowball has been much studied for reasons unrelated to tiling. It began in metric geometry as an analogue in space of the von Koch snowflake in the plane. Start with a unit-sided cube: (1) break every square face into 9 equal square faces, (2) bump out the middle square to a cube, then repeat steps 1 and 2 *ad infinitum*.

The conformal version is shown in Figure 28 with four levels of subdivision from an initial spherical cube. As a consequence of conformality of this rule, these four subdivisions and the infinitely many that can follow all occur *in situ*. As a thought experiment, fix a point p of the sphere and zoom in. As you pass through the various subdivision levels you see a sequence of nested tilings. Extracting a subsequential limit leads to an infinite planar snowball tiling. You may think of this as living in the tangent plane at p . The various limit tilings as p ranges over the sphere exemplifies the hierarchical tiling families that we study in our second paper.

A final note and image regarding the snowball connects our work to that of the metric geometers. Figure 29 shows 4 levels of snowball subdivision of the unit square $U = [0, 1] \times [0, 1]$. On the right, the (interiors of) the center tiles at each level (those associated with the bumped out cube) have been omitted. Continuing to subdivide and discard the center tiles, one is left with a closed nowhere dense residual set Λ_s . For those familiar with it, this quickly evokes the classical Sierpinski “carpet” Λ , the two-dimensional version of the middle-thirds Cantor set; see [38].

The snow carpet Λ_s is a well-behaved example of more general situations studied by Mario Bonk in [4]. To put things in his terms, let $\mathcal{S} = \{S_i : i \in I\}$ denote the collection of Jordan curves bounding the omitted central tiles: (1) the curves S_i bound Jordan domains with *mutually disjoint interiors*, (2) they are *k-quasicircles* for a uniform constant k , and (3) they are *uniformly relatively separated* according to the definition of Bonk. These properties are easily established by the conformal regularity of our construction and standard tools, such as Koebe distortion. Details are left to the interested reader. Using Theorem 1.1, its Corollary 1.2, and §12 of [4], there exists a quasiconformal homeomorphism $f : \Lambda_s \rightarrow \Lambda$. Moreover, aside from its evident 4-fold rotational symmetry about the center of U , f is unique. A final comment: *apropos* of our coming discussion, it is interesting to note that in defining the map f , the intermediary in Bonk's Theorem 1.1 is a carpet whose defining Jordan curves are all round circles.




FIGURE 28. A tiling of S^2 : the conformal “snowball” of Example 3.14.




FIGURE 29. Define the conformal snow carpet Λ_s by omitting the middle tiles at every stage under the “snowball” subdivision.

4. DISCRETIZATION AND APPROXIMATION

In this section we discuss the mechanics of the creation, manipulation, analysis, and display of conformal tilings through the parallel and computable theory of *discrete conformal tilings*. The practical matters are not merely subordinate considerations: Circle packing technology lies at the heart of this new topic, from motivation, theory, and computation, to visualization and, perhaps most important, experimentation. Discrete conformal tilings were introduced by the authors in [8] and [9], and many of their key features are direct analogues of known results in circle packing theory. We develop or recall the necessary machinery in the next subsection. We use the adjective *continuous* for conformal notions when we need to distinguish them from their discrete versions.

An important note: due to the nature of circle packing, discrete tilings both *mimic* and *approximate* their continuous counterparts, giving two approaches to proofs. Mimicry is often so faithful that one can use simple quasiconformal maps to transfer properties. Failing that, approximation lets one transfer results by taking limits. We will see both approaches. We begin by establishing the discrete machinery and reviewing the basics of circle packing.

4.1. Tiling complexes. Discretization starts by generating refined triangulations. Let \mathcal{T} be a combinatorial tiling of a topological surface S and let $\beta\mathcal{T}$ denote its barycentric subdivision, as described in §1. For technical reasons involving possible 1-gons and 2-gons in \mathcal{T} , the barycentric subdivision operator β is applied once more, breaking each triangle of $\beta\mathcal{T}$ into 6 triangles. Successive applications of the hexagonal subdivision operator \mathfrak{h} (termed *hex refinement*) may then follow. (Subdivision operators β and \mathfrak{h} are described on the left in Figure 27.)

Definition. A *tiling complex* for \mathcal{T} is a triangulation of its surface S obtained by barycentrically subdividing $\beta\mathcal{T}$ and then applying some number $j \geq 0$ of stages

of hex refinement. The complex $K_0^{\mathcal{T}} = \beta(\beta\mathcal{T})$ is the *coarse tiling complex* and the j -stage tiling complex $K_j^{\mathcal{T}}$ is defined recursively by hex refinement: $K_j^{\mathcal{T}} = \mathfrak{h}(K_{j-1}^{\mathcal{T}}), j = 1, 2, \dots$.

Note that complexes $K_j^{\mathcal{T}}$ are purely combinatorial objects, as is \mathcal{T} itself initially. They are essential because circle packing provides a canonical way to impose geometry on triangulations. Figure 30 illustrates the white/gray triangles for the tiling \mathcal{T} behind Figures 4. We highlight the subcomplex $L_j^T \subset K_j^{\mathcal{T}}$ associated with a 5-gon tile T . The coarse level, $L_0^T \subset K_0^{\mathcal{T}}$, is on the left, the stage-2 refinement, $L_2^T \subset K_2^{\mathcal{T}}$, is on the right. The geometric differences are nearly imperceptible. We also reveal in each image the circles which clandestinely give shape to this tile as part of the full circle packings for $K_0^{\mathcal{T}}$ and $K_2^{\mathcal{T}}$. We review the circle packing machinery in the next subsection.

Subcomplexes L_j^T which are associated with individual tiles, as highlighted in Figure 30(a) and (b), will be called *discretely regular polygons* (n -gons). These are precise combinatorial analogues (at varying refinement levels) of *conformally* regular polygons. Comparing Figures 30 and 7, we recognize the conformal properties in combinatorial form — the reflective structures, the dihedral symmetries, corners, the tile barycenter, edge barycenters, the $2n$ white/gray triangles. (Note: one could equally well work in bottom-up fashion. A *tiling complex* is formed by a union of discretely regular polygons $\{L_j^T\}$ (all of the same refinement level j) identified combinatorially along tile edges.)

4.2. Circle packing. Geometry is imposed on tiling complexes *via* circles. A *circle packing* P is a configuration of circles realizing some prescribed pattern of tangencies, in our case, the pattern encoded in a particular tiling complex. We refer the reader to the literature for full details [40]. Manipulations of circle packings for this paper and its companion [5] have been carried out *via* the second author’s open software, `CirclePack`; see Appendix C. In particular, the `CirclePack` “scripts” for recreating many of the experiments behind this paper are available from the second author.




FIGURE 30. Triangulated subcomplexes for a discretely regular pentagonal tile T within a tiling: (a) the coarse level, $L_0^T \subset K_0^{\mathcal{T}}$, and (b) the stage-2 refinement, $L_2^T \subset K_2^{\mathcal{T}}$.

Each of our (circle) packings P will lie in one of the standard geometries $\mathbb{G} = \mathbb{S}^2, \mathbb{C},$ or \mathbb{D} . The pattern of tangencies behind P is encoded in a (simplicial) *complex*

K , which triangulates an oriented topological surface. The circles C_v in P are in one-to-one correspondence with the vertices v of K ; circles C_v and C_w are tangent if $\langle v, w \rangle$ is an edge of K , and $\{C_u, C_v, C_w\}$ form a positively oriented triple of mutually tangent circles if $\langle u, v, w \rangle$ is a positively oriented face of K . The geometric triangulation in \mathbb{G} formed by connecting the centers of C_u, C_v, C_w for every face $\langle u, v, w \rangle$ of K is called the *carrier* of P , $\text{carr}(P)$. All of our circle packings P will be *univalent*, meaning that the circles of P have mutually disjoint interiors, and thus $\text{carr}(P)$ provides an embedding of K . This complex K may be finite or infinite, with border or without border; the K we encounter here are all simply connected.

Every complex K has an essentially unique canonical univalent circle packing P_K known as its *maximal* circle packing. Since K is simply connected, this packing lies in \mathbb{G} , one of \mathbb{S}^2, \mathbb{C} , or \mathbb{D} as usual, and K is said to be *spherical*, *parabolic*, or *hyperbolic*, respectively. *Essentially unique* means that it is unique up to Möbius transformations of \mathbb{G} .

Suppose K is finite. If $\partial K = \emptyset$, then K is spherical and $\text{carr}(P_K) = \mathbb{S}^2$. On the other hand, if $\partial K \neq \emptyset$, then K triangulates a topological closed disc. In this case, K is hyperbolic: P_K lies in \mathbb{D} and the circles associated with boundary vertices of K are *horocycles*, meaning circles of \mathbb{D} which are internally tangent to $\partial\mathbb{D}$. Horocycles may be treated as circles having infinite hyperbolic radius and centers at their points of tangency with the unit circle.

Suppose, on the other hand, that K is infinite. If $\partial K \neq \emptyset$, then K is hyperbolic and again the circles of P_K associated with boundary vertices are horocycles. If ∂K is empty, however, we encounter a fundamental dichotomy of circle packing: K will be either hyperbolic or parabolic. If K is parabolic, then P_K fills the euclidean plane, $\text{carr}(P_K) = \mathbb{C}$, whereas if K is hyperbolic, then $\text{carr}(P_K) = \mathbb{D}$. Given K , the determination of whether P_K fills the euclidean or the hyperbolic plane is the circle packing type problem.

Key notions in our formulation of conformal tiling are motivated by circle packing. The type problem for tilings mirrors the circle packing type problem, and we will have much to say about the connection shortly. The term “maximal” for tilings derives from maximal circle packings and the implications are parallel. For instance, if K is hyperbolic, then there are many circle packings P for K in addition to the maximal packings. In this paper, these other packings are obtained by treating them as euclidean and applying some type of *boundary conditions*. This is the case, for example, for the circle packings behind Figure 8, as is evident in the detail of Figure 31. In practice, nearly all images in the paper are based on finite euclidean circle packings realizing prescribed *angle sums* along their boundaries (even though those boundaries may lie outside the displayed images).

4.3. Discrete conformal tilings. By realizing tiling complexes through circle packing, we get discretized versions of \mathcal{T} . Suppose P_j is a univalent circle packing for $K_j = K_j^T$. For each tile $T \in \mathcal{T}$ there is a discretely regular polygon $L_j^T \subset K_j$. The carrier of P_j restricted to L_j^T is termed a *discrete conformal tile*; its shape is imposed by the associated circles of the packing (as in Figure 30). Generically the tile is not “conformal” in the continuous sense. “Discrete conformal” is the relevant notion here, acknowledging the intrinsically conformal nature of circle packing objects.

Definition. If P is a univalent circle packing for a tiling complex K for \mathcal{T} , then $\text{carr}(P)$ is a union of the discrete conformal tiles. The resulting tiling is called a

discrete conformal tiling. We may designate it by \mathcal{T}_P , but typically we write \mathcal{T}_j , indicating that its packing is associated with some j -stage tiling complex $K = K_j^{\mathcal{T}}$.

Note that every tiling complex K for \mathcal{T} has at least one univalent circle packing P ; namely, its maximal packing P_K . If K is hyperbolic, there may be others as well. Each gives rise to a discrete conformal realization of \mathcal{T} .

Example 4.1. We illustrate the technology with creation of the “chair” tiling image in the bottom right of Figure 8. Each tile is a combinatorial octagon and neighboring tiles share one or two contiguous edges. In this example, \mathcal{T} represents two stages of the chair subdivision (see Figure 15) of a single initial chair tile. We describe the panels in turn:

- Figure 31(a) is a free-hand drawing of \mathcal{T} , emphasizing that we start with nothing but combinatorics.
- Figure 31(b) shows the coarse discrete tiling \mathcal{T}_0 . The underlying circle packing P_0 (833 circles) is shown and the edges and vertices of the discrete conformal tiles are highlighted. Here and in (c) and (d) we impose boundary conditions and a normalization so that the carrier of the packing takes a chair shape, reminiscent of (but not precisely equal to) the classical chair.
- Figure 31(c) shows \mathcal{T}_0 without the underlying circles, but rather with shading of the white/gray faces of $\beta\mathcal{T}$; the $\{0, 1, \infty\}$ points are marked with $\{\bullet, \times, \square\}$.
- Figure 31(d) is the stage-3 discrete tiling \mathcal{T}_3 , involving nearly 50,000 circles.
- Figure 31(e) and (f) zoom in on corresponding details in (c) and (d), respectively, to show the shape evolution under refinement. Note that each white/gray face is a 6-flower in (e), a circle for the barycenter vertex and its 6 petal flowers defining the boundary of the face; each is hex refined 3 times for the finer packing in (f).

The discrete conformal tilings in Figure 31 suggest the fundamentals of the convergence of discrete tilings \mathcal{T}_j to their continuous counterparts. For all intents and purposes, panel (d) shows the tiling and its 16 tiles with their true conformal shapes. Formalities of this convergence are the next issue.

4.4. Limit tilings. We face three parallel tiling worlds: combinatorial, discrete (at various refinement levels), and conformal. There are internal convergence issues in each along with transition issues from one to another.

At the combinatorial level, convergence of (rooted) tilings is in terms of the big-ball metric described in §2.5. With a combinatorial tiling in hand, transition to the discrete conformal world is via circle packing; this involves a web of convergence issues of its own, for which we defer to the circle packing literature. For our purposes, the result is the sequence of discrete conformal tilings \mathcal{T}_j associated with univalent packings at various refinement levels. The final transition, from *discrete conformal* to *conformal* involves the convergence $\mathcal{T}_j \rightarrow \mathcal{T}$. This could be formulated as the shape convergence seen in Figure 31, using, for example, the Hausdorff metric on sets. Instead, we will define discrete tiling maps $\phi_j : \mathcal{T} \rightarrow \mathcal{T}_j$ and prove a stronger result; namely, that the discrete tiling maps ϕ_j converge uniformly on compacta to the conformal tiling map ϕ .

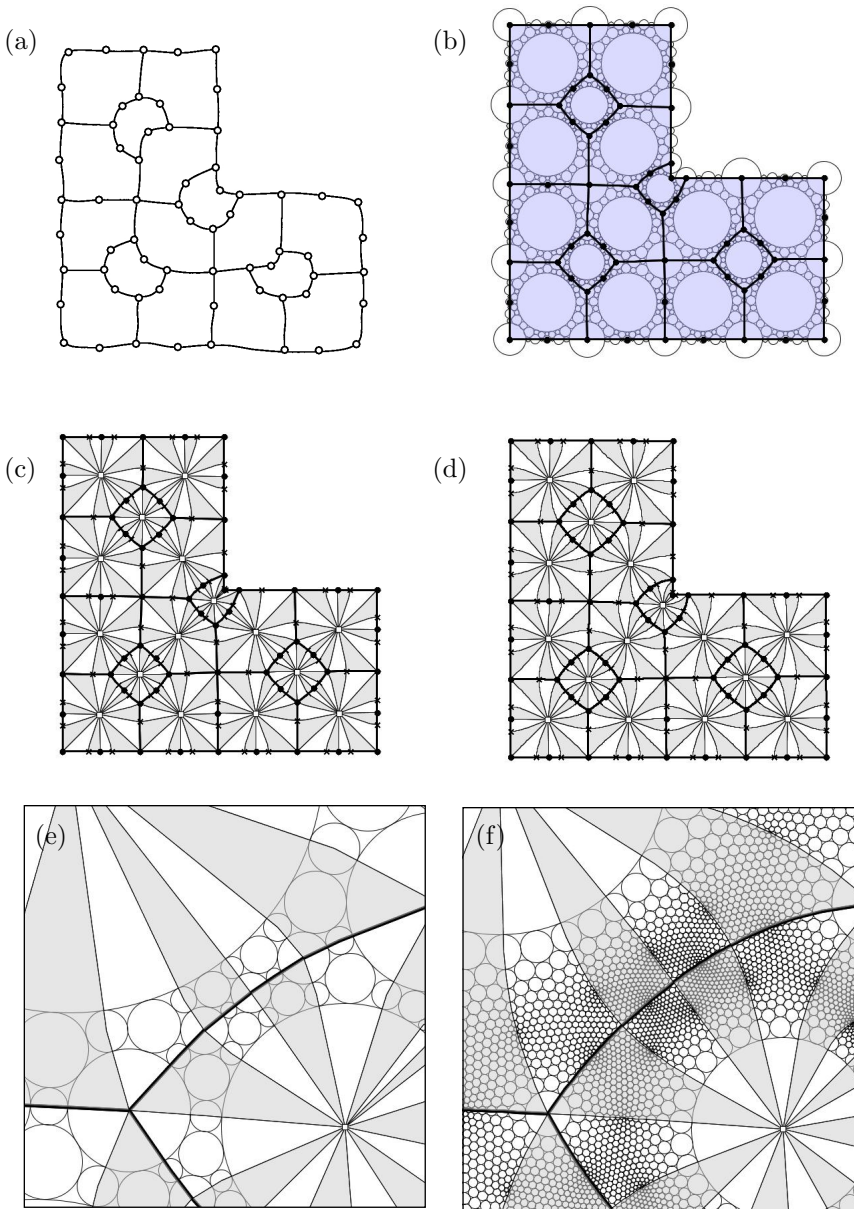


FIGURE 31. Two stages of the discrete conformal chair tiling.

Example 4.2 (Preview). An early example focusing on the mechanics may help. Suppose \mathcal{T} is an infinite, simply connected combinatorial tiling. By restricting to finite portions of \mathcal{T} , we get this computationally effective means for computing its maximal conformal tiling:

- (1) Choose a root tile T_0 . For each $j > 1$, let $\mathcal{T}^{(j)} \subset \mathcal{T}$ be the subtiling associated with the combinatorial ball $B_{\mathcal{T}}(T_0, j)$ (see §2.5).

- (2) Let K_j denote $K_j^{\mathcal{T}^{(j)}}$, the j -stage tiling complex for $\mathcal{T}^{(j)}$. For purposes of normalization, α and γ will be designated vertices of K_j , α being the center (i.e., barycenter) vertex of T_0 and γ the center for some other designated tile.
- (3) There are computer algorithms proven to effectively compute the maximal circle packing $P_j = P_{K_j} \subset \mathbb{D}$ and to put α at the origin and γ on the positive imaginary axis. (These are standard normalizations implemented in `CirclePack`.) This circle packing defines a discrete conformal tiling $\mathcal{T}^{(j)}$.
- (4) A normal families theorem will show that discrete tiling maps $\phi_j : \mathcal{T}^{(j)} \rightarrow \mathcal{T}^{(j)}$ will (perhaps after normalization, see Example 2.2) converge uniformly on compacta of $\text{int}(\mathcal{T})$ to a conformal tiling map $\phi : \mathcal{T} \rightarrow \mathbb{G}$, where \mathbb{G} is either \mathbb{D} or \mathbb{C} .
- (5) The image under the limit tiling map ϕ gives us a conformal realization \mathcal{T} with concrete tile shapes. In other words, the discrete conformal tilings $\mathcal{T}^{(j)}$ converge uniformly on compacta to the conformal tiling \mathcal{T} ; we will indicate this by $\mathcal{T}^{(j)} \rightarrow \mathcal{T}$.

We turn now to the full tiling \mathcal{T} . To consistently define discrete tiling maps on \mathcal{T} , we need a common domain; we use the β -equilateral structure on the barycentric subdivision $\beta\mathcal{T}$, described in §1.2 and §5.1. For each discrete tiling $\mathcal{T}^{(j)}$ we now show how to define $\phi_j : \mathcal{T} \rightarrow \mathcal{T}^{(j)}$.

Fixing $j \geq 0$, let $K_j = K_j^{\mathcal{T}}$ now be the j -stage refinement of $\beta\mathcal{T}$ and let P_j be a univalent circle packing for K_j in \mathbb{G} . In the RPWA-structure, each face f of $\beta\mathcal{T}$ is a unit-sided euclidean equilateral triangle with euclidean metric. In K_j , on the other hand, f is associated with a subcomplex L_j^f which is a j -stage hex refinement of a 6-flower. We can circle pack L_j^f so that its carrier is a unit-sided euclidean equilateral triangle and then identify this carrier isometrically with f in $\beta\mathcal{T}$. If we lift such packings (at the same refinement stage j) to every face of $\beta\mathcal{T}$, one can see that the circular arcs match up across edges and around vertices where the faces attach, yielding an honest-to-goodness circle packing Q_j for K in the $\rho_{\mathcal{T}}$ metric on $\beta\mathcal{T}$.

The mapping ϕ_j is a homeomorphism, $\phi_j : \text{carr}(Q_j) \rightarrow \text{carr}(P_j)$, defined face-by-face. We consider the planar and spherical settings separately.

First, suppose P_j lies in \mathbb{C} or \mathbb{D} . Treat P_j as euclidean. For each vertex $v \in K_j$, map the center of its circle in Q_j to the center of its circle in P_j . Each face of K_j is associated with a triple $\langle u, v, w \rangle$ of vertices, so the map defined on the centers can be extended *via* barycentric coordinates to map each face of $\text{carr}(Q_j)$ onto the corresponding face in $\text{carr}(P_j)$. In other words, ϕ_j is the piecewise affine map between the carriers.

If P_j lies in \mathbb{S}^2 , defining ϕ_j is slightly more complicated, as we need a normalization: We move the centers of three designated tiles (i.e., the centers of the circles associated with their barycenter vertices) *via* a Möbius transformation to lie at $0, 1, \infty$, respectively. Let \mathcal{C} denote the set of centers of circles of P_j . We get ϕ_j started as before by mapping circle centers in $\text{carr}(Q_j)$ to the corresponding centers in \mathcal{C} . Let H be the convex hull of \mathcal{C} in \mathbb{R}^3 . Each face in $\text{carr}(Q_j)$ corresponds to a planar triangular facet in ∂H , and we extend the map on centers to faces in two steps: First, extend *via* barycentric coordinates so that ϕ_j maps each face of

$\text{carr}(Q_j)$ affinely onto the corresponding facet of the hull H . Second, project each facet radially to \mathbb{S}^2 . Our normalization of P_j ensures that the origin is an interior point of the hull, so ϕ_j will be our homeomorphism in the spherical case.

Definition. Given a combinatorial tiling \mathcal{T} , non-negative integer j , and univalent circle packing P_j for $K_j^{\mathcal{T}}$, the homeomorphism $\phi_j : \text{carr}(Q_j) \rightarrow \text{carr}(P_j)$ as defined above is called a (*j-stage*) *discrete tiling map*. We write $\phi_j : \mathcal{T} \rightarrow \mathcal{T}(j)$ for brevity (though $\mathcal{T}(j)$ depends on P_j).

Note that as maps between circle packing carriers, the ϕ_j would be termed *discrete conformal mappings* in the lingo of circle packing; see [40, Chp 17]. With these mappings in hand, we next formulate a discrete convergence result parallel to Theorem 2.3 which can be tailored to handle the many convergence situations that arise in practice.

With inclusion of the intermediary discrete objects, the notation becomes a bit more difficult. We are assuming the combinatorial convergence $\mathcal{T}^n \rightarrow \mathcal{T}$. Associated with each \mathcal{T}^n is a discrete conformal tiling denoted $\mathcal{T}(n)$. In particular, this presumes we have: (1) a refinement level $j = \ell(n)$; (2) a univalent packing $P(n)$ for the j -stage tiling complex $K(n) = K_j^{\mathcal{T}^n}$; and (3) the associated discrete tiling map $\phi(n) : \mathcal{T}^n \rightarrow \mathcal{T}(n)$ (i.e., $\phi(n) : \mathcal{T}^n \rightarrow \text{carr}(P(n))$). It is convenient to view $P(n)$ as lying in \mathbb{S}^2 , though limits may end up in \mathbb{C} or \mathbb{D} .

Theorem 4.1. *Suppose that \mathcal{T} is a simply connected combinatorial tiling and that $\mathcal{T}^n \rightarrow \mathcal{T}$ represents rooted convergence of a sequence of simply connected combinatorial tilings to \mathcal{T} . Suppose that for each n we have a discrete conformal tiling $\mathcal{T}(n)$ in \mathbb{S}^2 and its associated tiling map $\phi_n : \mathcal{T}^n \rightarrow \mathcal{T}(n)$. Given any sequence of indices $\{n_k\}$ with the property that $\ell(n_k) \rightarrow \infty$ as $k \rightarrow \infty$, there is a further subsequence $\{n_j\} = \{n_{k_j}\}$ for which one of these conclusions holds:*

- (a) *The subsequence $\{\phi_{n_j}\}$ degenerates, meaning that it converges on \mathcal{T} to a constant function or to a function taking two values, or*
- (b) *The subsequence $\{\phi_{n_j}\}$ converges uniformly on compacta of $\text{int}(\mathcal{T})$ to a conformal tiling map ϕ for \mathcal{T} .*

In the latter case, $\phi(\mathcal{T})$ gives a conformal realization of \mathcal{T} .

Proof. This proof involves the usual concatenation of subsequence extractions and diagonalizations. We may start with an initial subsequence \mathcal{T}^{n_k} for which $\ell(n_k) \uparrow \infty$. (By the standard abuse of notation we reuse the same subsequence indices $\{n_k\}$ after each extraction.)

Note first that the maps $\phi(n_k)$ are defined on \mathcal{T}^{n_k} , not on \mathcal{T} , so the uniform convergence in (b) will need some explanation. Let $\Omega \subset \mathcal{T}$ be an open set whose closure is compact and lies in $\text{int}(\mathcal{T})$. There exists $M > 0$ so large that for $k \geq M$, the tilings \mathcal{T} and \mathcal{T}^{n_k} agree on some combinatorial ball B containing $\overline{\Omega}$; thus, all $\phi(n_k)$ are well defined on $\overline{\Omega}$ for sufficiently large k .

Since the refinement levels of the \mathcal{T}^{n_k} go to infinity as k grows, we can apply Proposition 4.2 below to extract from $\{\phi(n_k)\}$ a further subsequence so that $\phi(n_k)$ either converges uniformly on Ω to a conformal map or degenerates.

To grow the domains, choose $\{\Omega_m\}$ to be a sequence of simply connected open subsets of \mathcal{T} having nested compact closures which exhaust $\text{int}(\mathcal{T})$. Apply the

above argument to extract convergent subsequences successively from the Ω_m . Diagonalization yields a subsequence $\{\phi(n_k)\}$ which converges uniformly on compacta of $\text{int}(\mathcal{T})$. Degenerating on any Ω_m would imply degeneration throughout \mathcal{T} by uniqueness of analytic functions. Thus, if the limit ϕ is not degenerate, it must be a conformal map $\phi : \mathcal{T} \rightarrow \mathbb{S}^2$ (where, of course, we are treating the domain \mathcal{T} as a Riemann surface under its β -equilateral structure). This gives the limit conformal realization of \mathcal{T} and completes the proof, modulo the crucial proposition. \square

Proposition 4.2. *Let \mathcal{T} be a simply connected tiling with its β -equilateral structure and let Ω be an open subset of \mathcal{T} whose closure $\overline{\Omega}$ is compact and lies in the interior of \mathcal{T} . Suppose $\{\phi_n\}$ is a sequence of discrete tiling maps for \mathcal{T} with refinement levels $\ell(n)$ going to infinity. Then $\{\phi_n\}$ contains a subsequence that converges uniformly on Ω to a (one-to-one) conformal mapping $\phi : \Omega \rightarrow \mathbb{G}$ or a subsequence that degenerates on Ω to a constant or two constants.*

Proof. This concerns convergence of circle packings and is a standard line of argument in the packing literature. It is based on two early results of Rodin and Sullivan [35]. We use their Ring Lemma to show that the maps ϕ_n are uniformly κ -quasiconformal on Ω for some dilatation $\kappa \in [1, \infty)$. We then use their Hexagonal Packing Lemma to show that on compact subsets of faces of \mathcal{T} , the dilatations κ go to 1 as $n \rightarrow \infty$. The conclusion then follows by the normal families arguments traditional in quasiconformal mapping theory along with the fact that 1-quasiconformal maps are conformal.

Non-spherical Case: Assume first that \mathcal{T} is not the sphere. Let K_n denote the complex $K_{\ell(n)}^{\mathcal{T}}$. Recall that ϕ_n maps the circle packing Q_n for K_n in the $\rho_{\mathcal{T}}$ metric on \mathcal{T} to a circle packing P_n for K_n in \mathbb{S}^2 . If \mathcal{T} has a boundary we might encounter some minor issues with the associated boundary circles in our later arguments. Since our interest lies with Ω and since $\overline{\Omega} \subset \text{int}(\mathcal{T})$, the easiest fix is to simply disregard the boundary circles. That is, we will assume henceforth that the packings Q_n and P_n are modified by removing the circles associated with any boundary vertices that K_n might have. In particular, this ensures that the circles of P_n have mutually disjoint interiors.

For full generality, we have assumed the P_n lie in \mathbb{S}^2 , but we would prefer they be in \mathbb{C} . Since each P_n is univalent, there is some point $z_n \in \mathbb{S}^2$ omitted by P_n (that is, which is neither in $\text{carr}(P_n)$ nor in any circle of P_n). We may choose a rigid motion ψ_n of \mathbb{S}^2 that moves z_n to ∞ . The circle packing $\psi_n(P_n)$ now omits ∞ , so we project it stereographically to a circle packing in \mathbb{C} . The sequence ψ_n has some subsequence converging uniformly to a Möbius transformation ψ , so at the expense this additional subsequence extraction, the fates of subsequences of ϕ_n and $\psi_n \circ \phi_n$ will be the same. In other words, *in the non-spherical case, we may assume, without loss of generality, that the packings P_n are euclidean.*

By compactness of $\overline{\Omega}$, at most finitely many vertices and barycenters of \mathcal{T} will lie in $\overline{\Omega}$, and all the other vertices of K_n have degree at most 6, so there is a uniform upper bound d on the degrees of the circles of P_n , independent of n . Note that in the full $\beta\mathcal{T}$, absent locally finite complexity, one may lose this bound on d ; it is precisely here that the restriction $\overline{\Omega} \subset \text{int}(\mathcal{T})$ is needed.

A uniform bound on quasiconformal dilatation of ϕ_n now follows *via* the Rodin/Sullivan Ring Lemma, which says that the radii of neighboring interior circles in univalent euclidean circle packings are comparable, with constant depending

only on degree. Apply this to Q_n first; the faces involve boundary circles, but as each is equilateral, it is clear that it could be embedded in a univalent euclidean packing with degrees bounded by 12. The Ring Lemma thus provides lower bound $\alpha_{12} > 0$ on all angles of all faces in $\text{carr}(Q_n)$. In P_n the circles of interest are also interior but now have degrees bounded by d . The Ring Lemma gives lower bound $\alpha_d > 0$ on the angles in the associated faces of $\text{carr}(P_n)$. Since ϕ_n carries faces of $\text{carr}(Q_n)$ to faces in $\text{carr}(P_n)$ affinely, the fact that the angles in domain and range are bounded away from zero by $\min\{\alpha_{12}, \alpha_d\}$ easily gives a uniform upper bound $\kappa < \infty$ on the quasiconformal dilatations of ϕ_n on the interior of each triangular face of $\beta\mathcal{T} \cap \bar{\Omega}$. The union of edges of these faces forms a removable set for quasiconformal maps on Ω . That is, ϕ_n is κ -quasiconformal on Ω for all sufficiently large n .

We may now apply standard normal families arguments from the theory of quasiconformal mapping to extract a subsequence ϕ_{n_k} that converges uniformly on Ω to a κ -quasiconformal mapping $\phi : \Omega \rightarrow \mathbb{C}$ or that degenerates. In the latter case we are done, so assume the subsequence does not degenerate.

We are now in position to improve our dilatation bound κ . Fix any face $f \in \beta\mathcal{T}$. The combinatorics of K_n restricted to f are purely hexagonal (degree 6) on the interior, and as n_k grows, the refinement level grows and the number of hex generations grows. In particular, fix any compact $E \subset \text{int}(f)$. The number of hex generations separating E from ∂f goes to ∞ as n_k grows. The Rodin/Sullivan Hexagonal Packing Lemma now applies. The faces of $\text{carr}(Q_{n_k})$ intersecting E have angles going to $\pi/3$ — that is, they look more and more like euclidean equilateral triangles as n_k grows. The same holds for the corresponding faces in $\text{carr}(P_{n_k})$, so the affine maps ϕ_{n_k} have dilatation on E which converges to 1. We conclude that the limit mapping ϕ is 1-quasiconformal at all points of $\text{int}(f)$. Restricting attention to Ω , we see that the limit function ϕ is 1-quasiconformal on the interiors of faces f intersecting Ω . Since the edges of these faces form a removable set for quasiconformal maps, ϕ is 1-quasiconformal on all of Ω . But 1-quasiconformal maps are conformal. That is, the limit map ϕ is conformal on Ω and we are finished in the non-spherical case.

Spherical Case: The arguments change little when \mathcal{T} is a topological sphere. We may assume $\Omega = \mathcal{T}$. The circle packings Q_n in \mathcal{T} remain piecewise euclidean as before, but the image packings P_n are now spherical. We start by assuming that \mathcal{T} has at least 3 tiles and we use the standard normalization on each P_n : The barycenter circles of three designated tiles are centered at $0, 1, \infty$, respectively. We comment on the general case at the end.

Recall that the maps ϕ_n were defined in two steps: first mapping each face of Q_n affinely onto its facet of the convex hull of P_n 's centers, followed by radial projection to \mathbb{S}^2 . In light of Lemma 5.6, the circles of P_n go to zero in radius as n grows, so the spherical versions of the Ring and Hexagonal Packing Lemmas apply in P_n . Facets associated with some compact set E in a face of $\beta\mathcal{T}$ are increasingly equilateral and go to zero in area. In particular, the affine maps to facets converges to 1-quasiconformal as before. Also, on nearly equilateral facets with vertices in \mathbb{S}^2 , the radial projection clearly converges to conformal as the facet area goes to zero. Due to our standard normalizations, no subsequence can degenerate. Every subsequence ϕ_{n_k} has a further subsequence that converges to a conformal map and appealing again to the normalization, there is a unique limit. We conclude that the

full sequence $\{\phi_n\}$ converges uniformly on all of \mathcal{T} to the unique conformal tiling map ϕ satisfying our normalization.

There are some last comments to make in the spherical case. First, if \mathcal{T} has only one or two tiles, then our arguments go through after adjusting the normalization; for example, normalizing so three *vertices* of \mathcal{T} map to $0, 1, \infty$. As for the normalization itself, how do we argue for general maps? Certainly something is needed in order to define reasonable discrete tilings maps. We cannot rely on Möbius transformations; they map circles to circles, but do not respect circle centers, and therefore do not respect our definition of discrete tiling maps. One can look instead to centroids. For each P_n , let $c_n \in \mathbb{R}^3$ be the centroid of its circle centers. If some subsequence c_{n_k} converges to a point of the sphere, then using any sensible definition of discrete tiling maps ϕ_{n_k} , the subsequence $\{\phi_{n_k}\}$ would necessarily degenerate. On the other hand, if all c_n remain within some distance $r < 1$ of the origin in \mathbb{R}^3 , then one can adjust our arguments to extract a subsequence converging to a conformal tiling map ϕ . By the essential uniqueness of conformal tilings in the spherical case, any limit conformal tilings, $\phi(\mathcal{T})$, will necessarily be Möbius images of one another, so in hindsight one might as well normalize at the beginning. \square

4.5. Observations. In these comments we assume the tilings are simply connected and of finite local complexity.

4.5.1. Theoretical. Various technical results established in §5 highlight the comprehensive nature of the discrete tiling world. A tiling \mathcal{T} is parabolic if it has a conformal realization filling the plane. According to Theorem 5.5, this is the case if and only if discrete conformal tilings for \mathcal{T} fill the plane as well. In other words, “parabolic” is an attribute of the combinatorics, regardless of whether you are working in the conformal tiling or discrete conformal tiling setting. Furthermore, Theorem 5.3 links the type in these settings to the type of traditional tilings. Numerous results on type are available in the theory of circle packing, connected with degree, random walks, vertex extremal length, and isoperimetric inequalities, for example. These now apply to tilings.

This paper lays the foundation for the future study of *families* of conformal tilings, and for that, Theorem 2.3 on limits will be of particular importance. But, by Proposition 5.7, any conformal tiling is the limit under refinement of the analogous discrete conformal tilings (even without finite local complexity). Therefore, we see that Theorem 4.1 implies Theorem 2.3. Whereas Theorem 2.3 stays within the category of conformal tilings, the discrete version is about conformal tiling in practice — it confirms the validity of the constructions and behaviors one sees directly in experiments.

4.5.2. Practical. Being able to transfer questions to the discrete tiling world has very real and practical implications. We observed in the proof of Theorem 2.3, for example, that in the combinatorial convergence $\mathcal{T}^n \rightarrow \mathcal{T}$, one may assume the \mathcal{T}^n are nested finite tilings exhausting \mathcal{T} . Since each \mathcal{T}^n is finite, its associated discrete conformal realizations are effectively computable. This is important in practice, since all experiments are necessarily finite (and, of course, only approximate), and it is the experiments that first motivated and continue to inform conformal tiling. We know, now, that we can effectively approximate tile shapes and visualize significant portions of global tiling patterns.

There remain practical issues, of course. Among these is “type”, particularly in regard to limits in families of tilings, as we discuss in [5]. We note that although one cannot reliably deduce type from finite stage experiments, those experiments can, nevertheless, be very suggestive. Borders are also an issue. We point the reader to the last part of the proof of the technical Lemma 5.7 below. This nicely illustrates the subtiles of theory and practice. In the practical approximation of the 16-tile chair Λ illustrated in Figure 31, discrete chair shapes Λ_j were created using boundary conditions on their packings P_j . None of these individual carriers is precisely equal to Λ . In this regard, the packings we appeal to in the proof of Lemma 5.7 will lie in Λ itself and so might seem preferable. However, packings lying in Λ are provided by a theorem of He and Schramm, which is non-constructive. We can use the more practical approach because boundary issues largely resolve themselves under refinement. Nonetheless, care must be taken.

5. TECHNICAL RESULTS

Several background technical results are gathered here.

5.1. Piecewise affine structures. A polygonal cell decomposition C of a topological surface S brings a *piecewise affine structure* (PWA) to S by identifying each n -sided cell of C with some (non-degenerate) euclidean n -sided polygon. These cells are then identified at shared vertices or isometrically along shared full edges to build the surface. (In particular, edges shared by two cells must have the same length in each.)

With the PWA structure, S becomes a metric space, the distance between points p and q being the infimum of the lengths of paths in S between p and q (always attained, though not necessarily uniquely). The PWA structure also provides a conformal atlas for S in a canonical way, using the euclidean structure for face interiors, the reflective structure for edges, and local power maps $z \mapsto z^\alpha$ at vertices. With this conformal atlas, S becomes a Riemann surface.

Definition. This construction defines a *PWA-structure* and C is termed a *PWA-surface*. In particular, C is both a Riemann surface and a metric surface with a piecewise euclidean metric. In case the model euclidean polygon for each cell of C is a unit-sided regular polygon, then the structure is a *RPWA-structure* (‘R’ for *regular*). Furthermore, if all cells of C are triangular, then the RPWA-structure is termed an *equilateral structure* and C is an *equilateral surface*.

Equilateral structures on tilings — meaning, our β -equilateral structures — are fundamental to our notion of conformal tiling, but more general piecewise affine structures enter as well. In particular, for use in the next subsection we state the following quasiconformal mapping lemmas. The straightforward proofs are left to the reader.

Lemma 5.1. *Let K and K' be combinatorially isomorphic triangulations of a surface S , each with a PWA-structure. Suppose that there is a uniform lower bound $a > 0$ on all angles of all model triangles for both K and K' . Then the homeomorphism $\phi : K \rightarrow K'$ which maps each face of K affinely onto the corresponding face of K' is κ -quasiconformal for some $\kappa \geq 1$ depending only on a .*

Suppose \mathcal{T} is a tiling of finite local complexity with the RPWA-structure, so each tile is identified with a unit-sided *regular* euclidean polygon. Break each n -gon

tile into $2n$ congruent euclidean triangles meeting at its center. Doing this for all tiles leads to a triangulation with the combinatorics of $\beta\mathcal{T}$; the inherited PWA-structure makes it into a PWA-surface which we denote by K . On the other hand, let K' denote $\beta\mathcal{T}$ with its equilateral structure (i.e., RPWA-structure). Applying the previous lemma, we get

Lemma 5.2. *Suppose \mathcal{T} has finite local complexity. Let K and K' denote $\beta\mathcal{T}$ with the two PWA-structures as noted above. Then the homeomorphism $\phi : K \rightarrow K'$ which maps each face of K affinely onto the corresponding face of K' is κ -quasiconformal for some $\kappa \geq 1$ that depends only on the complexity of \mathcal{T} .*

Note that in both of these lemmas the map ϕ is just the identity map on the common underlying topological surface S . Quasiconformality enters because S has distinct structures when treated as domain and as range.

5.2. Type problems. An infinite simply connected combinatorial tiling \mathcal{T} without boundary is either *parabolic*, living in \mathbb{C} , or *hyperbolic*, living in \mathbb{D} . Which is it?

Traditional tilings have generally been euclidean (though the hyperbolic pentagonal tiling is a much-studied departure). Since we have illustrated several conformal versions with suggestive images, it is incumbent upon us to ask if their type can be proven. This question will be taken up further in [5]. For comparisons in this paper to traditional examples, we establish the following result.

Theorem 5.3. *Suppose \mathcal{E} is a tiling of \mathbb{C} composed of euclidean polygonal tiles and displaying finite local complexity. Then the conformal tiling sharing the combinatorics of \mathcal{E} is parabolic, so it fills \mathbb{C} as well.*

Proof. Suppose that \mathcal{T} denotes the combinatorial tiling associated with \mathcal{E} and that \mathcal{C} denotes a conformal realization of \mathcal{T} . Assume \mathcal{C} lies in \mathbb{G} . We establish existence of a quasiconformal homeomorphism $\psi : \mathbb{C} \rightarrow \mathbb{G}$. Then by the quasiconformal version of Liouville's theorem, \mathbb{G} cannot be a proper subset of the plane, implying that $\mathbb{G} = \mathbb{C}$.

The map ψ is represented as a composition, $\psi = \phi_3 \circ \phi_2 \circ \phi_1$ of successive homeomorphisms. It suffices to show that ϕ_j is κ_j -quasiconformal for some $\kappa_j, j = 1, 2, 3$, for then ψ will be κ -quasiconformal for $\kappa = \kappa_3\kappa_2\kappa_1$.

Finite local complexity of \mathcal{E} implies that there are finitely many euclidean polygonal prototiles and that copies of these attach to one another in at most finitely many ways. In particular, by adding at most finitely many designated points on prototile boundaries as new corners, we may assume that any two tiles of \mathcal{E} that intersect will share isolated vertices and/or one or more fully closed edges. (See, for example, the chair prototile in Figure 15; two of the chair's original 6 sides were split in half to accommodate the subdivisions, and the tiles became octagons.)

Any (non-degenerate) euclidean polygon can be decomposed into a finite number of (non-degenerate) euclidean triangles in such a way that every edge of the polygon is an edge of one of the triangles. In particular, applying this to each of the prototiles we obtain a finite collection \mathcal{C}_1 of model euclidean triangles. Imprinting each tile of \mathcal{E} with the decomposition for its prototile leads to a triangulation K_1 of \mathbb{C} by triangles, each congruent to one of the triangles of \mathcal{C}_1 .

Our first map ϕ_1 carries \mathcal{E} to \mathcal{T}_{rpwa} , a copy of \mathcal{T} endowed with an RPWA-structure. Each tile T of \mathcal{T}_{rpwa} is a regular unit-sided euclidean n -gon P for some n and is identified with one of the euclidean prototiles p , vertices mapping to

vertices. Let graph G be the 1-skeleton for the decomposition of p . Applying Tutte's Embedding Theorem ([42]) we may realize a straight line embedding of G in P with the corner vertices from p positioned at the corresponding corners of P . This decomposes P into a certain set of euclidean triangles. Doing this for all prototiles leads to a second finite collection \mathcal{C}_2 of model euclidean triangles. Imprinting each tile of \mathcal{T}_{rpwa} with the decomposition transferred from its prototile leads to a triangulation K_2 of \mathcal{T}_{rpwa} , each tile congruent to one of the triangles of \mathcal{C}_2 .

Triangulations K_1 and K_2 are isomorphic, and since all triangles belong to the finite collection $\mathcal{C}_1 \cup \mathcal{C}_2$, there is a positive lower bound on all their angles. Lemma 5.1 now provides a κ_1 -quasiconformal map $\phi_1 : K_1 \rightarrow K_2$ which identifies corresponding triangles. By grouping these back into tiles we get $\phi_1 : \mathbb{C} \rightarrow \mathcal{T}_{rpwa}$.

Of course, the conformal surface of interest is not \mathcal{T}_{rpwa} but rather the β -equilateral surface, which we may denote \mathcal{T}_{eq} . Lemma 5.2 provides a κ_2 -quasiconformal map $\phi_2 : \mathcal{T}_{rpwa} \rightarrow \mathcal{T}_{eq}$.

Our final map is $\phi_3 = F$, the conformal map $F : \mathcal{T}_{eq} \rightarrow \mathbb{G}$. Being analytic, ϕ_3 is κ_3 -quasiconformal for $\kappa_3 = 1$, completing our construction. The composition

$$\mathbb{C} \rightarrow \mathcal{T}_{rpwa} \rightarrow \mathcal{T}_{eq} \rightarrow \mathbb{G}$$

is $\kappa_1 \kappa_2 \kappa_3$ -quasiconformal, implying by Liouville's Theorem that \mathbb{G} is the euclidean plane. \square

Certainly other conditions on the euclidean tiling \mathcal{E} might give the same conclusion with similar arguments, but we will not pursue that here. Another situation of interest follows with essentially the same arguments. Recall the notions associated with *finite subdivision rules* τ as described in §3.3. To apply τ to a tiling \mathcal{T} , it must have a rule for partitioning each tile type of \mathcal{T} into finitely many tiles. Though we have not gone into detail on the various properties and consistency requirements for building rules, this basic fact is all we want to point out here. The proof of the following is then left to the reader.

Theorem 5.4. *Let \mathcal{T} be a simply connected combinatorial tiling of finite local complexity and suppose τ is a finite subdivision rule which can be applied to \mathcal{T} . Then the tilings \mathcal{T} and $\tau\mathcal{T}$ are of the same type, be that spherical, parabolic, or hyperbolic.*

Of great practical interest is the role of our intermediaries, the *discrete* conformal tilings. The following result addresses their type.

Theorem 5.5. *Suppose \mathcal{T} is a simply connected combinatorial tiling with finite local complexity. Then \mathcal{T} and its associated discrete conformal tilings all have the same type.*

Proof. This is clear if \mathcal{T} is a topological sphere. Assume \mathcal{T} is not compact. Let $P_0 \subset \mathbb{G}$ be a univalent circle packing defining a discrete conformal tiling for the tiling complex $K_0^{\mathcal{T}}$. Treating P_0 as an euclidean packing (whether \mathbb{G} is \mathbb{C} or \mathbb{D}), the triangles of its carrier define a PWA-structure on $\beta\mathcal{T}$; denote this PWA-surface by \mathcal{T}_{pwa} . Finite local complexity implies that there is an upper bound d on the degree of the circles forming P_0 , i.e., on the number of circles tangent to any one circle. The Rodin/Sullivan Ring Lemma ([35]) implies that the angles of the triangles of \mathcal{T}_{pwa} are all greater than some lower bound $a > 0$. The angles are constantly $\pi/3$ in

\mathcal{T}_{eq} , of course, so by Lemma 5.1 there is a quasiconformal mapping $\phi : \mathcal{T}_{pwa} \rightarrow \mathcal{T}_{eq}$. Follow ϕ with the uniformizing conformal map $F : \mathcal{T}_{eq} \rightarrow \mathbb{G}'$, where \mathbb{G}' is either \mathbb{C} or \mathbb{D} . Then $F \circ \phi : \mathbb{G} \rightarrow \mathbb{G}'$ is quasiconformal. By Liouville's theorem, $\mathbb{G} = \mathbb{C}$ if and only if $\mathbb{G}' = \mathbb{C}$. So P_0 is parabolic if and only if \mathcal{T} is parabolic.

To finish, recall that circle packings P_j for hex refinements $K_j^{\mathcal{T}}$ provide additional discrete conformal tilings for \mathcal{T} . It has been shown by Bill Wood [43] that hex refinement does not change type, so any two P_j have the same type as one another, hence the same type as P_0 and \mathcal{T} . \square

Finally, note that we speak of the P_j as though they were concrete, yet as infinite objects they are also theoretical constructs. The computable packings one must actually work with are associated with finite portions of tilings.

5.3. Circle packing lemmas. Möbius transformations of the sphere map circle packings to circle packings, but do not respect the circle centers. One could, for example, force all but one of the centers to cluster in an arbitrarily small neighborhood of some point. This lemma states that our standard normalization avoids this so that as we refine, the convex hull of circle centers converges to the sphere. The arguments use basic results on conformal moduli under quasiconformal maps and are left to the reader.

Lemma 5.6. *Suppose L is a triangulation of the sphere with distinguished vertices v_0, v_1, v_∞ . For $j \geq 0$, let P_j be the univalent circle packing on \mathbb{S}^2 for the j -stage hex refinement of L normalized so the circles for v_0, v_1, v_∞ are centered at $0, 1, \infty$, respectively, and let r_j be the maximal spherical radius among circles of P_j . Then $r_j \rightarrow 0$ as $j \rightarrow \infty$.*

The next lemma shows the comprehensive nature of the discrete tiling theory. Any conformal tiling is the limit under refinement of analogous discrete versions.

Lemma 5.7. *Let \mathcal{T} be a simply connected combinatorial tiling, and $\phi : \mathcal{T} \rightarrow \mathbb{G}$ a conformal tiling map. There exist j -stage discrete tiling maps $\phi_j : \mathcal{T} \rightarrow \mathbb{G}$, $j = 0, 1, \dots$, so that $\phi_j \rightarrow \phi$ uniformly on compacta of $\text{int}(\mathcal{T})$ as $j \rightarrow \infty$.*

Proof. We first dispense with the two cases in which ϕ is essentially unique: when \mathcal{T} is the sphere or is of parabolic type. Here the discrete tiling maps ϕ_j are canonical, being associated with the maximal packings P_j of the tiling complexes $K_j = K_j^{\mathcal{T}}$. Assuming ϕ and the ϕ_j satisfy our standard normalizations, ϕ is the unique conformal tiling map for \mathcal{T} . According to Proposition 4.2, any subsequence $\{\phi_{j_n}\}$ will have a further subsequence which converges to a conformal tiling map of \mathcal{T} or will degenerate to one or two constants. The normalizations prevent degeneracy, so the limit must be ϕ . Since this applies to any subsequence, the full sequence $\{\phi_j\}$ must converge to ϕ . Note that in the parabolic case we do not need to know whether the circle packings P_j individually fill the plane, but at least in the case of finite local complexity, Theorem 5.5 shows that this is indeed the case.

This brings us to the case that \mathcal{T} is of hyperbolic type. For now, put aside the finite and bordered tilings, and assume \mathcal{T} is infinite without border. The hurdle in building discrete analogues of ϕ is the lack of essential uniqueness in ϕ itself: Its image, which we denote $\Lambda = \phi(\mathcal{T})$, could be absolutely any simply connected proper open subset of the plane. This is where beautiful and deep circle packing results of He and Schramm come into play.

For normalization purposes, assume T_0 is a root tile for \mathcal{T} , α is its conformal center, and γ is a corner. Let $z_\alpha, z_\gamma \in \Lambda$ be the images of these two points under ϕ . Each tiling complex K_j has vertices associated with these points, and abusing notation we refer to these as α and γ as well. Given $K_j = K_j^{\mathcal{T}}$, Theorem 1.2 of [19] implies existence of a univalent circle packing P_j with $\text{carr}(P_j) = \Lambda$. Moreover, it says we may specify that the circle for α in P_j is centered at z_α and that the circle for γ is on the ray from z_α through z_γ . Define ϕ_j to be the discrete tiling map $\phi_j : \mathcal{T} \rightarrow \text{carr}(P_j)$.

The maps ϕ_j are κ -quasiconformal for some common $\kappa < \infty$ due to the Ring Lemma, and hence they form a normal family. Standard arguments with extremal ring domains let one show that $|\phi_j(\alpha) - \phi_j(\gamma)|$ is bounded below independent of j . In particular, there exists some subsequence $\{\phi_{j_k}\}$ so that $\phi_{j_k} \rightarrow \widehat{\phi}$. Here $\widehat{\phi}$ is a conformal tiling map for \mathcal{T} , and our task is to show $\widehat{\phi} \equiv \phi$. By the Carathéodory Kernel Theorem, $\widehat{\phi}(\mathcal{T}) = \Lambda$. Conjugating by the Riemann mapping from Λ to the disc \mathbb{D} and applying Theorem 1.2, we see that $\widehat{\phi} = g \circ \phi$ for some conformal automorphism g of Λ . Since $\widehat{\phi}(\alpha) = \phi(\alpha) = z_\alpha$, g fixes z_α , so it is a conformal rotation of Λ . By our normalization of the maps ϕ_j , note that $\widehat{\phi}(\gamma)$ is on the ray from z_α through $\phi(\gamma) = z_\gamma$. We may generally conclude that g is the identity, giving the desired convergence, $\phi_{j_k} \rightarrow \widehat{\phi} \equiv \phi$ on \mathcal{T} . There is one final detail: if there is a non-trivial g which maps $\widehat{\gamma}$ to the same ray from z_α , one simply needs to replace the corner γ used in this argument by a point of T_0 closer to α .

In summary, when \mathcal{T} is of hyperbolic type, infinite without border, and ϕ is a given conformal tiling map, we are able to define discrete tiling maps $\{\phi_j\}$ so that $\phi_j \rightarrow \phi$ uniformly on compact subsets of \mathcal{T} .

The final situation to consider is \mathcal{T} with border. The given ϕ maps $\text{int}(\mathcal{T})$ to a simply connected open set Λ and extends continuously to its border, $\phi : \partial\mathcal{T} \rightarrow \partial\Lambda$. Our previous arguments require adjustments due to the presence of boundary vertices in the tiling complexes K_j . We consider the finite case, as it highlights a pretty interplay between the discrete and conformal worlds. Consider the conformal tiling \mathcal{T} with 16 tiles we purport to show in Figure 31. Below is the clear target mapping.

Target Mapping: *A conformal tiling map $\phi : \mathcal{T} \rightarrow \Lambda$, where Λ is a euclidean chair-shaped polygonal region with designated corners of \mathcal{T} mapping to the corners of Λ .*

In fact, ϕ exists, and conditions on the corners will completely determine its image Λ (up to similarity). The difficulty is that it is essentially impossible to know the side lengths of Λ ahead of time. Instead, the process behind Figure 31 effectively teased out the shape of Λ *through* the process of discrete approximation. That turns our current proof on its head — at least in any practical sense.

Nonetheless, pretending we know Λ is our game here. This brings a new difficulty: the discrete approximations shown in Figure 31 have circles centered at points around the boundary. These are **not** the packings that our argument is using; the packings created *via* the He/Schramm result have all circles *interior* to Λ . This implies that $\text{carr}(P_j)$ will be a proper (in fact, compact in this finite case) subset of Λ . This introduces another feature in the convergence: As j increases, the Ring

Lemma will guarantee that the boundary circles go to zero in radius, so $\text{carr}(P_j)$ exhausts $\text{int}(\Lambda)$ as j grows. (Note one subtlety: Because the circle packings P_j are univalent, the Ring Lemma applies to boundary as well as interior circles.) It is clear that we still get subsequential convergence to limit conformal tiling maps $\hat{\phi}$ which map $\text{int}(\mathcal{T})$ onto $\text{int}(\Lambda)$. The reader can confirm that, as earlier, we can adjust normalizations to get $\hat{\phi} \equiv \phi$. These arguments extend with little additional effort to the case that \mathcal{T} is infinite with border as well. Therefore, our proof is complete. \square

In closing, note that while the `CirclePack` methods behind Figure 31 are limited in practice to tiling maps ϕ whose ranges are \mathbb{D} or polygonal regions, they have the advantage of being computable. In the chair case, for example, ϕ carries designated boundary vertices of \mathcal{T} to the corners of Λ ; ϕ , the shape of Λ , and the location of z_α are all computable. The He/Schramm approach, though more flexible, is not computable, even with *a priori* information on z_α and the shape of Λ .

APPENDIX A. ALTERNATE CONFORMAL STRUCTURES

Let \mathcal{T} be a combinatorial tiling of a surface S . For both practical and aesthetic reasons, we have chosen the canonical conformal structure on S to be the *conformally regular* and *reflective* β -equilateral structure associated with $\beta\mathcal{T}$. A natural competing choice would be the RPWA-structure on S as described in §5.2. We refer to these tilings as “equilateral” and “RPWA” and take a moment to compare them. Unless all tiles have the same number of sides, these are distinct — there is no conformal homeomorphism between them that maps tiles to tiles. In both tilings the tile edges are analytic arcs. In the equilateral tiling, if m tiles come together at a common corner, each tile has corner angle $2\pi/m$. In the RPWA tiling, the corner angles may differ, but we are assuming finite local complexity, so these angles are bounded away from zero. The upshot is that there is a *quasiconformal* homeomorphism between these two structures that maps tiles to tiles. In particular, they have the same type, parabolic or hyperbolic. When type is the key concern, either structure could be used.

In practice, the RPWA-structures face major difficulties (see [40, §22.4]), giving the equilateral structure the upper hand: Associated circle packings are easily defined, refined, computed and displayed, even for packings with several million circles. Also appealing is the fact that the tilings \mathcal{T} , \mathcal{T}^* , \mathcal{T}^\dagger , and \mathcal{Q} are realized simultaneously. Moreover, all of these move seamlessly into the discrete conformal world, *a la* §4.3. As with other uses of circle packing in function theory and conformal geometry, this discrete world is remarkably faithful to its continuous counterpart, so even at coarse refinement stages one can see the mathematics in action.

APPENDIX B. MULTIPLY CONNECTED TILINGS

Though we have concentrated on simply connected examples in this paper, combinatorial tilings \mathcal{T} , as defined in 1.1, make sense for arbitrary surfaces. The barycentric subdivision $\beta\mathcal{T}$, its β -equilateral structure, the tiling function $M : \mathcal{T} \rightarrow \mathbb{S}^2$, the piecewise spherical metric, the realizations *via* conformal mappings — these are all perfectly well defined.

There are some things to note, of course, and perhaps the most challenging is new subtlety in the “type” problem. When S is multiply connected, there is generically a moduli space, denoted $\text{Mod}(S)$, parameterizing the uncountably many distinct conformal structures that can be defined on S . A conformal tiling \mathcal{T} of S picks out precisely one point of $\text{Mod}(S)$; that is, there is precisely one Riemann surface R in which one can realize a conformal tiling of \mathcal{T} . To illustrate, we have chosen a combinatorial tiling \mathcal{T} having 11 tiles that form a torus. The conformal realization is shown in Figure 32. A fundamental domain is displayed on the left with the usual edge-gluing conventions (i.e., like-colored edges are identified). The side-pairing translations are $\lambda_1 : z \rightarrow z + w_1$ and $\lambda_2 : z \rightarrow z + w_2$ (the third side-pairing being a linear combination of these).




FIGURE 32. A combinatorial tiling \mathcal{T} of a topological torus S converts it to a conformal torus of conformal modulus $\tau = w_2/w_1$.

On the right in Figure 32 is the lifting of \mathcal{T} to a tiling of \mathbb{C} , the universal covering space of the torus. The lifted tiling, call it $\widehat{\mathcal{T}}$, is a geometric tiling like others we have considered, the only difference being that it is invariant under the covering group $\Lambda = \langle \lambda_1, \lambda_2 \rangle$. The moduli space for the torus, $\text{Mod}(S)$, is the strip $\{\tau \in \mathbb{C} : |\tau| \geq 1, -1/2 \leq \Re(\tau) \leq 1/2\}$ with sides identified under the maps $z \mapsto z + 1$ and $z \mapsto -1/z$. For our example, the conformal modulus is the ratio $w_2/w_1 \sim .1874 + 1.396i$.

Through the use of covering theory, the theory of tilings for general surfaces is quite straightforward. A Riemann surface R may be represented as $R = \mathbb{G}/\Lambda$ where \mathbb{G} is one of \mathbb{D} , \mathbb{C} , or \mathbb{S}^2 and Λ is a properly discontinuous group of automorphism of \mathbb{G} . (Note that Λ must be trivial if $\mathbb{G} = \mathbb{S}^2$, so we get nothing new in this case.)

Definition. A conformal tiling \mathcal{T} of a Riemann surface R represented as $R = \mathbb{G}/\Lambda$ may be defined geometrically as $\widehat{\mathcal{T}}/\Lambda$, where $\widehat{\mathcal{T}}$ is a conformal tiling of \mathbb{G} which is invariant under Λ . \mathcal{T} (and the corresponding combinatorial tiling) is said to be *hyperbolic*, *parabolic*, or *spherical* as \mathbb{G} is \mathbb{D} , \mathbb{C} or \mathbb{S}^2 .

Development of this theory follows that in circle packing, which the reader will find in [40, Chp 9]. Circle packing, in turn, provides *discrete conformal tilings* (see §4.3) along with all the usual machinery for computations, approximations, and visualization (such as Figure 32). There is one restriction in the hyperbolic setting, as noted already in the simply connected case: The canonical hyperbolic tiling must be *maximal*, meaning that its interior fills R . The complete story, then, is this, extending Theorem 1.2 and Corollary 2.2.

Theorem B.1. *Given a combinatorial tiling \mathcal{T} of a topological surface S , there exists a unique point of $\text{Mod}(S)$ whose associated Riemann surface R supports a maximal conformal realization of \mathcal{T} . Moreover, this realization is unique up to conformal automorphisms of R .*

Corollary B.2. *Suppose \mathcal{T} is a maximal conformal tiling in a Riemann surface R . Then \mathcal{T} contains two tiles T and T' which are similar in R if and only if there is a combinatorial automorphism (perhaps orientation reversing) of the tiling \mathcal{T} mapping T to T' .*

We close by observing that we face a new “type” problem, in the spirit of Theorem 5.4, but even more challenging. Let S be a multiply connected topological surface, \mathfrak{T} the set of all combinatorial tilings of S . By the previous theorem we have a map $\pi : \mathfrak{T} \rightarrow \text{Mod}(S)$. There is some information about this map. For compact S , the range of π is countable. Indeed, by the Grothendieck’s theory of *dessins d’Enfants* a Riemann surface of genus $g > 0$ is realized if and only if it has a defining polynomial $P(z, w) = 0$ with coefficients in some algebraic number field (see [9, 36]). Arguments paralleling those of [6, 7] should show the range of π is dense in standard metrics.

Regarding our subdivision theme, this may bridge the issues at the heart of Cannon’s Conjecture and Bonk’s work on carpets, [4]. Thus, suppose τ is a subdivision rule which can be applied to a tiling $\mathcal{T} \in \mathfrak{T}$. Then $\tau^n(\mathcal{T}) \in \mathfrak{T}, n = 0, 1, 2, \dots$. If τ is a *conformal* subdivision rule, then subdivision occurs *in situ* — it does not change the conformal structure, so τ commutes with π , $\tau(\pi(\mathcal{T})) = \pi(\tau(\mathcal{T}))$. In general, however, τ will be an action on the moduli space. What are the possibilities when τ is not conformal?

APPENDIX C. CirclePack SOFTWARE

The experiments leading to this paper and all the paper’s computations and images were done using the software `CirclePack`, developed by the second author. The software is available through his web site under a *GNU General Public License*. In addition, scripts are available from the second author which allow the user to replicate and extend many of the experiments discussed in this paper.

Acknowledgments. Our thanks to Jim Cannon and Bill Floyd for fascinating conversations about subdivision rules over many years, and also for “rules” files used in our experiments. Thanks also to Chaim Goodman-Strauss, Natalie Frank, and Lorenzo Sadun for valuable insights and for specific examples used here. Lastly, a special thanks to Maria Ramirez-Solano for digging so deeply into the regular pentagonal tiling and to her and Jean Savinien for arranging an exciting workshop on *Non-standard hierarchical tilings* at University of Copenhagen, 2012, which led more or less directly to this paper.

REFERENCES

- [1] Lars V. Ahlfors, *Conformal invariants: topics in geometric function theory*, McGraw-Hill Book Co., New York-Düsseldorf-Johannesburg, 1973. McGraw-Hill Series in Higher Mathematics. MR0357743
- [2] Omer Angel and Oded Schramm, *Uniform infinite planar triangulations*, Comm. Math. Phys. **241** (2003), no. 2-3, 191–213, DOI 10.1007/978-1-4419-9675-6-16. MR2013797

- [3] Michael Baake and Uwe Grimm, *Aperiodic order. Vol. 1*, Encyclopedia of Mathematics and its Applications, vol. 149, Cambridge University Press, Cambridge, 2013. A mathematical invitation; With a foreword by Roger Penrose. MR3136260
- [4] Mario Bonk, *Uniformization of Sierpiński carpets in the plane*, Invent. Math. **186** (2011), no. 3, 559–665, DOI 10.1007/s00222-011-0325-8. MR2854086
- [5] Philip L. Bowers and Kenneth Stephenson, *Conformal tilings II: Local isomorphism, hierarchy, and conformal type*.
- [6] Philip L. Bowers and Kenneth Stephenson, *The set of circle packing points in the Teichmüller space of a surface of finite conformal type is dense*, Math. Proc. Cambridge Philos. Soc. **111** (1992), no. 3, 487–513, DOI 10.1017/S0305004100075575. MR1151326
- [7] Philip L. Bowers and Kenneth Stephenson, *Circle packings in surfaces of finite type: an in situ approach with applications to moduli*, Topology **32** (1993), no. 1, 157–183, DOI 10.1016/0040-9383(93)90044-V. MR1204413
- [8] Philip L. Bowers and Kenneth Stephenson, *A “regular” pentagonal tiling of the plane*, Conform. Geom. Dyn. **1** (1997), 58–68 (electronic), DOI 10.1090/S1088-4173-97-00014-3. MR1479069
- [9] Philip L. Bowers and Kenneth Stephenson, *Uniformizing dessins and Belyi maps via circle packing*, Mem. Amer. Math. Soc. **170** (2004), no. 805, xii+97, DOI 10.1090/memo/0805. MR2053391
- [10] Dmitri Burago, Yuri Burago, and Sergei Ivanov, *A course in metric geometry*, Graduate Studies in Mathematics, vol. 33, American Mathematical Society, Providence, RI, 2001. MR1835418
- [11] James W. Cannon, *The combinatorial Riemann mapping theorem*, Acta Math. **173** (1994), no. 2, 155–234, DOI 10.1007/BF02398434. MR1301392
- [12] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Expansion complexes for finite subdivision rules. I*, Conform. Geom. Dyn. **10** (2006), 63–99 (electronic), DOI 10.1090/S1088-4173-06-00126-3. MR2218641
- [13] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Expansion complexes for finite subdivision rules. II*, Conform. Geom. Dyn. **10** (2006), 326–354 (electronic), DOI 10.1090/S1088-4173-06-00127-5. MR2268483
- [14] Natalie Frank, *Fusion: a general framework for hierarchical tilings*.
- [15] Natalie Priebe Frank, *A primer of substitution tilings of the Euclidean plane*, Expo. Math. **26** (2008), no. 4, 295–326, DOI 10.1016/j.exmath.2008.02.001. MR2462439
- [16] James T. Gill and Steffen Rohde, *On the Riemann surface type of random planar maps*, Rev. Mat. Iberoam. **29** (2013), no. 3, 1071–1090, DOI 10.4171/RMI/749. MR3090146
- [17] A. Grothendieck, *Esquisse d’un programme*, (1985), Preprint: introduced Dessins.
- [18] Ori Gurel-Gurevich and Asaf Nachmias, *Recurrence of planar graph limits*, Ann. of Math. (2) **177** (2013), no. 2, 761–781, DOI 10.4007/annals.2013.177.2.10. MR3010812
- [19] Zheng-Xu He and Oded Schramm, *The inverse Riemann mapping theorem for relative circle domains*, Pacific J. Math. **171** (1995), no. 1, 157–165. MR1362982
- [20] Gareth A. Jones and David Singerman, *Complex functions*, Cambridge University Press, Cambridge, 1987. An algebraic and geometric viewpoint. MR890746
- [21] Richard Kenyon and Andrei Okounkov, *What is ... a dimer?*, Notices Amer. Math. Soc. **52** (2005), no. 3, 342–343. MR2125269
- [22] Richard Kenyon and Andrei Okounkov, *Limit shapes and the complex Burgers equation*, Acta Math. **199** (2007), no. 2, 263–302, DOI 10.1007/s11511-007-0021-0. MR2358053
- [23] Grégory Miermont, *The Brownian map is the scaling limit of uniform random plane quadrangulations*, Acta Math. **210** (2013), no. no. 2, 319–401. MR3070569
- [24] O. Lehto and K. I. Virtanen, *Quasiconformal mappings in the plane*, 2nd ed., Springer-Verlag, New York-Heidelberg, 1973. Translated from the German by K. W. Lucas; Die Grundlehren der mathematischen Wissenschaften, Band 126. MR0344463
- [25] M. C. Esher, *The magic of M. C. Esher*, Thames and Hudson, 2013, ISBN-13: 9780500290736.
- [26] Dane Mayhook, *Conformal tilings and type*, Ph.D. thesis, 2016; advisor Phil Bowers, FSU.
- [27] Hervé Oyono-Oyono and Samuel Petite, *C*-algebras of Penrose hyperbolic tilings*, J. Geom. Phys. **61** (2011), no. 2, 400–424, DOI 10.1016/j.geomphys.2010.09.019. MR2746126
- [28] R. Penrose, *Pentaplexity: a class of nonperiodic tilings of the plane*, Math. Intelligencer **2** (1979/80), no. 1, 32–37, DOI 10.1007/BF03024384. MR558670

- [29] Charles Radin, *The pinwheel tilings of the plane*, Ann. of Math. (2) **139** (1994), no. 3, 661–702, DOI 10.2307/2118575. MR1283873
- [30] M. Ramirez-Solano, *Cohomology of the continuous hull of a combinatorial pentagonal tiling*, ArXiv e-prints (2013).
- [31] ———, *Construction of the continuous hull for the combinatorics of a regular pentagonal tiling of the plane*, ArXiv e-prints (2013).
- [32] ———, *Construction of the discrete hull for the combinatorics of a regular pentagonal tiling of the plane*, ArXiv e-prints (2013).
- [33] ———, *Continuous hull of a combinatorial pentagonal tiling as an inverse limit*, ArXiv e-prints (2013).
- [34] Maria Ramirez-Solano, *Non-commutative geometrical aspects and topological invariants of a conformally regular pentagonal tiling of the plane*, Ph.D. thesis, Department of Mathematical Sciences, University of Copenhagen, 2013.
- [35] Burt Rodin and Dennis Sullivan, *The convergence of circle packings to the Riemann mapping*, J. Differential Geom. **26** (1987), no. 2, 349–360. MR906396
- [36] Leila Schneps, *Dessins d’enfants on the Riemann sphere*, The Grothendieck theory of dessins d’enfants (Luminy, 1993), London Math. Soc. Lecture Note Ser., vol. 200, Cambridge Univ. Press, Cambridge, 1994, pp. 47–77. MR1305393
- [37] G. B. Shabat and V. A. Voevodsky, *Drawing curves over number fields*, The Grothendieck Festschrift, Vol. III, Progr. Math., vol. 88, Birkhäuser Boston, Boston, MA, 1990, pp. 199–227, DOI 10.1007/978-0-8176-4576-2_8. MR1106916
- [38] Warclaw Sierpinski, *Sur une courbe contoriennne qui contient une image biunivoque et continue de toute courbe donnée*, C. r. hebd. Seanc. Acad. Sci, Paris **162** (1916), 629–632.
- [39] Kenneth Stephenson, *Circle packing: a mathematical tale*, Notices Amer. Math. Soc. **50** (2003), no. 11, 1376–1388. MR2011604
- [40] Kenneth Stephenson, *Introduction to circle packing: the theory of discrete analytic functions*, Cambridge University Press, Cambridge, 2005. MR2131318
- [41] William Thurston, *The finite Riemann mapping theorem*, 1985, Invited talk, An International Symposium at Purdue University in celebration of de Branges’ proof of the Bieberbach conjecture, March 1985.
- [42] W. T. Tutte, *Convex representations of graphs*, Proc. London Math. Soc. (3) **10** (1960), 304–320. MR0114774
- [43] William E. Wood, *Bounded outdegree and extremal length on discrete Riemann surfaces*, Conform. Geom. Dyn. **14** (2010), 194–201, DOI 10.1090/S1088-4173-2010-00210-9. MR2672225

DEPARTMENT OF MATHEMATICS, THE FLORIDA STATE UNIVERSITY, TALLAHASSEE, FLORIDA 32306

E-mail address: `bowers@math.fsu.edu`

DEPARTMENT OF MATHEMATICS, THE UNIVERSITY OF TENNESSEE, KNOXVILLE, TENNESSEE 37996

E-mail address: `kstephe2@utk.edu`