

CONFORMAL TILINGS II: LOCAL ISOMORPHISM, HIERARCHY, AND CONFORMAL TYPE

PHILIP L. BOWERS AND KENNETH STEPHENSON

ABSTRACT. This is the second in a series of papers on conformal tilings. The overriding themes here are local isomorphisms, hierarchical structures, and the conformal “type” problem. Conformal tilings were introduced by the authors in 1997 with a conformally regular pentagonal tiling of the plane. This and even more intricate hierarchical patterns arise when tilings are repeatedly subdivided. Deploying a notion of expansion complexes, we build two-way infinite combinatorial hierarchies and then study the associated conformal tilings. For certain subdivision rules the combinatorial hierarchical properties are faithfully mirrored in their concrete conformal realizations. Examples illustrate the theory throughout the paper. In particular, we study parabolic conformal hierarchies that display periodicities realized by Möbius transformations, motivating higher level hierarchies that will emerge in the next paper of this series.

INTRODUCTION

We continue the development of conformal tilings begun in earnest in [7]. These form a new class of planar tilings that arise by placing a natural conformal structure on combinatorial decompositions of the plane. While still in its infancy, the theory clearly takes inspiration from the traditional theory of aperiodic hierarchical tilings, which has garnered much attention since at least the introduction of the Penrose tilings in the 1970s. Traditional tiling matured in the 1990s and has enjoyed intense development since; [2, 3, 13–15, 19, 20, 23, 24]. Much of this theory depends on combinatorial arguments, bereft of the influence of Euclidean geometry, and thus has become one of the two main influences upon our work.

The second main influence, well removed from the intricacies of the traditional aperiodic tiling theory, is Cannon, Floyd, and Parry’s articulation of finite subdivision rules in a series of papers over the last two decades. This targets the *Cannon Conjecture*, that every negatively curved group with 2-sphere Gromov boundary is, essentially, a cocompact Kleinian group. See, e.g., [8, 10, 11].

The authors have a long history of interaction with Cannon, Floyd, and Parry, beginning with our introduction of expansion complexes in [5]. The general construction was developed a bit by Cannon, Floyd, and Parry, but the theory remained largely unexplored until Maria Ramirez-Solano’s dissertation [21]. She brought sophisticated tools of traditional aperiodic tiling to dissect the pentagonal tiling of [5]. The present work can be thought of as a more general development inspired by Ramirez-Solano’s work. With our backgrounds in geometric topology, conformal

Received by the editors November 22, 2017, and, in revised form, August 7, 2018.

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

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

geometry, and function theory, we have chosen a unified development that we hope is accessible to readers uninitiated in the traditional theory and apologize for any oversights in credit to the tiling community.

The authors intend a series of papers that carries one from the basic definitions of individual conformal tilings through increasingly sophisticated layers of structure within families of conformal tilings.

(I) In the first paper, [7], the authors have laid out general foundations for the theory of *conformal tilings*, tilings by conformally regular curvilinear polygons with reflective structures, building on [5] and [6]. Existence and uniqueness are established, basic phenomena are illustrated with examples, and central themes, such as subdivision and conformal type, are discussed. Particularly relevant here are conditions under which a *combinatorial* subdivision of a tiling complex leads to a *conformal* subdivision of its conformal tiling. These and other results needed in the present work are reviewed and repackaged in the first section, so this paper is largely self-contained.

(II) In this, the second paper, we pursue links between conformal tilings and traditional aperiodic tilings that began to emerge for the pentagonal tiling in the dissertation of Maria Ramirez-Solano [21]. Our study centers on hierarchical families of tilings generated from subdivision rules. In traditional theory, largely euclidean, a tiling T determines a continuous hull in which every tiling locally isomorphic to T has natural residence, a canonical dynamical system on that hull, and a transversal for its action, which is a Cantor set slice providing local isomorphism classes of pointed tilings. In the conformal setting, a tiling \mathcal{T} is determined essentially uniquely by the pattern of its tiles, which is encoded in a planar polygonal complex K , so our work must start with combinatorics: We develop a metric space of polygonal complexes, a slice of local isomorphism classes to serve as transversal, and combinatorial versions of subdivision. Thus we are led to *combinatorial hierarchies*, bi-infinite sequences $\{K_n\}$ of complexes, each a subdivision of its predecessor. Reconnecting fully to geometry requires the special notion of *conformal* subdivisions developed in [7], and for these we can define *conformal hierarchies*, sequences $\{\mathcal{T}_n\}$ of conformal tilings, each an in situ subdivision of its predecessor. Among our main results are conditions for determining the conformal type, *parabolic* or *hyperbolic*, of conformal hierarchies and whether that type is consistent across the local isomorphism class. We illustrate with examples, culminating with special parabolic cases displaying a *periodic hierarchy*. This ties our work to its beginnings in [5] and to the Cannon, Floyd, and Parry efforts and serves as motivation for the third paper.

(III) The conformal hierarchies developed in the current paper are admittedly rare and special. When coupled with the notion of aggregate tiles introduced in **(I)**, however, they point the way to a new level of hierarchical structure. Our third paper [4] will pursue this with a study of *aggregate* conformal tilings and their limits. The results bring together in one setting conformal tiling hierarchies, traditional euclidean hierarchies, and new, perhaps fractal, tiling hierarchies. All arise spontaneously from pure combinatorics when endowed with appropriate conformal geometry.

Let us focus the lens a bit to preview our work in this paper. Conformal tilings \mathcal{T} are decompositions of the plane \mathbb{C} or the hyperbolic plane \mathbb{D} into curvilinear polygonal tiles which are conformally regular and have a reflective structure. Conformal tilings are determined essentially uniquely by their combinatorics, which we encode as planar polygonal complexes K . We are interested in families of tilings generated via subdivision rules τ , rules which have both combinatorial and geometric aspects. In the first section we review the fundamentals in combinatorics, conformal tiling, and subdivision on the way to establishing a key connection between combinatorics and geometry: dihedrally symmetric simple subdivision operators are conformal subdivision operators.

The second section concentrates on combinatorics with a view to the continuous hull. We define the metric space \mathbf{RC} of planar polygonal complexes and, for a given complex K , a local isomorphism class $\mathbf{LI}(K)$ of rooted complexes. This material is developed slightly beyond our immediate needs, but illuminates properties of K that mirror the traditional setting (cf. [3, Chp 5]). In particular, we will see the importance of two features: finite local complexity and repetitiveness. The broad ideas of this section promise to be old hat to the traditional tiling experts, but the particulars will be new, most certainly to those in discrete conformal geometry.

There seems to be no formal definition of the term *hierarchical tiling*, and we certainly refrain from suggesting one. Nonetheless, in the third section we define what it means for a planar polygonal complex K to exhibit a *combinatorial hierarchy*, a bi-infinite sequence $\{K_n\}$, each K_n a combinatorial subdivision of its predecessor. We are not aiming for full generality, so we quickly restrict attention to hierarchies based on subdivision rules τ . The key combinatorial notions of *expansion complexes* and *aggregate tilings* are developed. Expansion complexes, introduced in [5] and refined and exploited by Cannon, Floyd, and Parry in [10, 11], formalize the extraction of limit combinatorial tilings. They allow us to define, for appropriate rules τ , an aggregation operation which functions as τ^{-1} . The final theorem of the section, Theorem 3.3, guarantees a wealth of combinatorial hierarchies.

In the fourth section we bring in the additional combinatorial machinery, *expansiveness*, *repetitiveness*, and finiteness conditions, needed to capture the full sense of hierarchy for appropriate subdivision rules τ . Several examples illustrate. The section's main result, Theorem 4.4, tells us that the existence of combinatorial hierarchy for K can propagate to its whole local isomorphism class.

Conformal geometry enters in section five; when τ is a *conformal* subdivision rule, then a hierarchy $\{K_n\}$ of combinatorial tilings is realized as a *conformal hierarchy* $\{\mathcal{T}_n\}$ of conformal tilings. Here each \mathcal{T}_n is an in situ conformal subdivision of its predecessor \mathcal{T}_{n-1} , parallel to traditional euclidean substitution tilings. A principal theme of the paper is the *type problem*: does the conformal tiling \mathcal{T}_K for K live in the euclidean plane \mathbb{C} , *parabolic type*, or in the hyperbolic plane \mathbb{D} , *hyperbolic type*? The section concludes with the paper's Main Theorem, 5.5, establishing parabolicity for certain complexes K and subdivision rules τ —indeed, showing that parabolicity can propagate to the full local isomorphism class of K . This answers affirmatively the question of Maria Ramirez-Solano [21] which motivated this series of papers.

Concrete experiments have been central in the development of conformal tiling. In section six we explore in some detail the phenomenon of *periodic expansion complexes*, wherein the in situ subdivisions by τ are realized by Möbius transformations. We are able to identify and construct all such periodicities, which turn out to be

quite rare—a countable family among the local isomorphism class (K) . We describe and illustrate several examples.

Although periodicity itself is rare, the experiments and images of section six motivate the next stage in the topic. There is much work yet to do in regard to Cannon’s Conjecture, and the conformal hierarchies discussed here are suspended somewhere between the traditional euclidean hierarchical model and what is needed in the Cannon, Floyd, and Parry program. In our next paper [4] we introduce *conformal aggregate tilings*, which we hope will further the Cannon, Floyd, Parry program while also appealing to those who enjoy tiling for its own sake. We foreshadow those developments with some observations in the final section of this paper.

CONTENTS

Introduction	52
1. Conformal tilings	55
2. Complexes and their local isomorphism classes	63
3. Combinatorial hierarchies	70
4. Combinatorial hypotheses	78
5. Conformal hierarchies	85
6. Hierarchical periodicity	90
7. What next?	101
Acknowledgments	102
References	102

1. CONFORMAL TILINGS

In this section we review the essentials of conformal tiling and subdivision and establish notation. In particular, we provide enough background that the following theorem, proved at the end of this section and used later in the paper, is understandable to the reader.

Theorem 1.1. *Suppose τ is a dihedrally symmetric simple subdivision operator and K is a planar polygonal complex. Then the conformal tiling $\mathcal{T}_{\tau K}$ is a conformal subdivision of the conformal tiling \mathcal{T}_K .*

In [7] we laid out the combinatorial, equilateral, conformal, and concrete geometric structures of tilings. We abused notation by using the same symbol \mathcal{T} for all. Here we must take more care: K will denote the planar polygonal complex underlying a tiling, $|K^\beta|_{\text{eq}}$ its piecewise euclidean structure, S_K the associated Riemann surface, and \mathcal{T}_K the conformal tiling itself, occupying either the euclidean plane \mathbb{C} or the hyperbolic plane represented as the unit disc \mathbb{D} . In a sense, then, a tiling is a 4-tuple $\mathcal{T} = \{K, |K^\beta|_{\text{eq}}, S_K, \mathcal{T}_K\}$. (Note that in the hyperbolic case, \mathcal{T}_K will always refer to the so-called *maximal* tiling for K , so it fills \mathbb{D} and is unique up to conformal automorphisms (Möbius transformations) of \mathbb{D} .)

Let us begin with the combinatorics K . Informally, a *planar polygonal complex* is a decomposition of the plane into curvilinear polygons which meet one another at vertices or along full edges of their boundaries. More precisely, K is an oriented 2-dimensional regular CW-decomposition of the plane whose attaching maps are homeomorphisms. The 0-skeleton is a countable discrete collection of *vertices*; the

1-skeleton consists of *edges*, each spanning two distinct vertices; and the 2-cells are the *faces*, each a topological disc bounded by a closed cycle of two or more edges. When a face k has n edges ($n \geq 2$), we say that k is a *combinatorial n -gon* and its *polygonal type* is n . We find it convenient to think of the cells of K —the edges e and the faces k —as closed, and when we mean the corresponding open cell we will write e° or k° . We observe that these complexes may be defined from their 1-skeletons interpreted as “drawings” in the terminology of [7], but in this paper we do not allow the loops, dangling edges, and non-embedded polygons which can occur with drawings.

Our complexes K are locally finite, meaning that at most finitely many edges emanate from any given vertex. If there is $d < \infty$ so that each face has polygonal type at most d and at most d edges meet at any vertex, we say K has *bounded degree*. The smallest such d is the *degree* of K . Complexes K and K' are combinatorially equivalent, or CW-isomorphic, if there is an orientation-preserving cellular isomorphism $\phi : K \rightarrow K'$, that is, a homeomorphism which maps vertices to vertices, edges to edges, and faces to faces. Write $K \cong K'$. If F and L are CW-complexes, we write $F \hookrightarrow L$ to indicate that F isomorphically embeds in L . (Generally we will be dropping the adjective *orientation-preserving*, but we ask the reader to keep in mind that in this paper the complexes are oriented and maps between them are orientation-preserving.)

Geometric structures are imposed via more refined combinatorics. The *barycentric subdivision* of K , denoted as K^β , is defined by introducing a new vertex, a *face barycenter*, to each open face and a new vertex, an *edge barycenter*, to each open edge, and then adding new edges in each face connecting its barycenter to both its original vertices and to the barycenters added to its edges. This process subdivides each edge into two edges and each face k of K of polygonal type n into $2n$ triangles with the face barycenter as a common vertex, so the complex K^β is a simplicial decomposition of the plane into combinatorial triangles. A typical triangle in some face k of K has vertices a , b , and c , where a is a vertex of k , b is the edge barycenter for some edge of k , and c is the barycenter added to k° .

We now endow K with the *β -equilateral metric* using this triangulation K^β : identify each face of K^β with a unit-sided euclidean equilateral triangle, so each edge has unit length. Observe that two faces meeting along an edge e isometrically reflect to one another across e . The resulting metric space $|K^\beta|_{\text{eq}}$ is piecewise flat with cone type singularities at the vertices of K^β . There is a canonical maximal conformal atlas \mathcal{A} for $|K^\beta|_{\text{eq}}$, giving it what we call its *β -equilateral conformal structure*. The open faces are charts in \mathcal{A} , edge charts are defined by mapping two contiguous faces isometrically to the union of two equilateral triangles in the plane meeting along a common unit edge, and vertex charts are defined using local power mappings to resolve their cone type singularities. The resulting surface $S_K = (|K^\beta|_{\text{eq}}, \mathcal{A})$ is a non-compact, simply connected Riemann surface. The classical Uniformization Theorem implies that S_K is conformally equivalent to one of the plane \mathbb{C} or the disc \mathbb{D} . Throughout the paper we use \mathbb{G} to denote a generic plane, always either \mathbb{C} or \mathbb{D} . The conformal isomorphism mapping S_K to \mathbb{G} is unique up to conformal automorphisms of \mathbb{G} . We refer the reader to the first paper [7] for details of this procedure.

Definition. Given a planar polygonal complex K , the *conformal tiling* \mathcal{T}_K for K is the collection of *tiles* in \mathbb{G} determined by a uniformizing map of S_K onto \mathbb{G} ,

where the tiles are the images of the faces of K under the uniformizing map. The conformal tiling is unique up to conformal automorphisms of \mathbb{G} .

The tiles of \mathcal{T}_K are curvilinear polygons in \mathbb{G} that meet in the pattern of K . When S_K is *parabolic*, \mathcal{T}_K tiles the complex plane and is unique up to (orientation-preserving) euclidean similarities; when S_K is *hyperbolic*, \mathcal{T}_K tiles the Poincaré disc and is unique up to (orientation-preserving) hyperbolic isometries. The type problem now is manifest. Given K , is S_K parabolic or hyperbolic? Does \mathcal{T}_K tile the complex plane \mathbb{C} or the Poincaré disc \mathbb{D} ? In general, this is very difficult to resolve.

Two conformal properties characterize conformal tilings, *regularity* and *reflectivity*. Regularity means that if t is a tile of \mathcal{T}_K of polygonal type n , then t is homeomorphic to a regular euclidean n -gon by a homeomorphism that identifies vertices and is conformal on the interior of t . This is a direct consequence of the equilateral metric on $|K^\beta|_{\text{eq}}$ and it is in this sense that t is a *conformally regular* n -gon. However, there are uncountably many inequivalent ways to realize K by conformally regular polygons. This is where the additional reflective property of \mathcal{T}_K enters. Let us describe this property.

Let \mathcal{T} be a locally finite tiling by conformally regular polygons. The dihedral group D_{2n} of order $2n$ acts as a group of conformal/anticonformal isomorphisms on each tile t of polygonal type n . The action preserves and is transitive on vertices and has an interior fixed point, the *conformal center* of t . The elements of D_{2n} that preserve orientation act as a conformal rotation group of order n and those that reverse orientation act as anticonformal reflections whose fixed point sets—the *axes* of the action—are analytic arcs spanning across two vertices or two edges of t if n is even, and across vertex-edge pairs if n is odd. The intersection of an axis with an edge e is the *conformal center* of e determined by t . Slicing every tile of \mathcal{T} along its axes breaks it into conformal triangles f . The resulting triangles form a tiling we denote by \mathcal{T}^β .

Definition. A tiling \mathcal{T} by conformally regular polygons is said to be *reflective* if three conditions are satisfied whenever tiles t and t' meet along an edge e : (1) The conformal centers of e determined by t and t' coincide; (2) e is the fixed point set of an anticonformal reflection ψ that exchanges the two triangles f and f' meeting along e ; and (3) ψ exchanges the conformal centers of t and t' , which are vertices of f and f' , respectively.

Note that conformal tilings \mathcal{T}_K are clearly reflective: the conformal charts for edges in the atlas \mathcal{A} were chosen with precisely this in mind. On the other hand, suppose we start with a combinatorially equivalent tiling \mathcal{T} with conformally regular polygons. If \mathcal{T} is reflective, then \mathcal{T}^β has the combinatorics of the barycentric subdivision K^β . This is the pattern used in the construction of \mathcal{T}_K , leading to the following result.

Theorem 1.2. *Suppose \mathcal{T} is a tiling of \mathbb{G} by conformally regular polygons with the combinatorics of a polygonal complex K . If \mathcal{T} is reflective, then \mathcal{T} is an image of the conformal tiling \mathcal{T}_K under some conformal automorphism of \mathbb{G} .*

Proof. The proof is a consequence of a Schwarz reflection process described in the first paper [7]: Each triangular face \mathbf{k} of K^β is associated with a conformal triangle in each of \mathcal{T}_K and \mathcal{T} . A conformal homeomorphism ϕ between such triangles for a

given face \mathbf{k} extends analytically across each edge of \mathbf{k} by classical Schwarz reflection. Because \mathcal{T}_K and \mathcal{T} are both reflective, ϕ can be shown to extend analytically to the entirety of each triangle of K^β neighboring \mathbf{k} . The Schwarz reflection argument may now be repeated triangle-by-triangle throughout \mathcal{T}_K leading to a bijective conformal extension $\phi : \mathcal{T}_K \rightarrow \mathcal{T}$ which maps tiles to tiles. In particular, ϕ is necessarily a conformal homeomorphism of \mathbb{G} , an automorphism. \square

We will work with patches of tilings. A connected subcomplex $L \subset K$ for which every cell of L is contained in a closed face of L is termed a *combinatorial patch* of K . The corresponding set of tiles of \mathcal{T}_K will be denoted \mathcal{T}_L and termed a *patch* of \mathcal{T}_K . Patches \mathcal{T}_L and $\mathcal{T}_{L'}$, are conformally equivalent if there is a conformal isomorphism of $|\mathcal{T}_L| = \bigcup\{t : t \in \mathcal{T}_L\}$ onto $|\mathcal{T}_{L'}|$ that maps tiles to tiles. Obviously conformally equivalent patches have combinatorially equivalent subcomplexes L and L' ; this theorem provides the converse.

Theorem 1.3. *If L and L' are combinatorially equivalent patches of K and K' , respectively, then \mathcal{T}_L and $\mathcal{T}_{L'}$ are conformally equivalent patches of \mathcal{T}_K and $\mathcal{T}_{K'}$.*

The properties of conformal tilings, fully developed in [7], are taken for granted here. For example, note that all tile edges are analytic arcs and that if m faces meet at a vertex v , then each face has angle $2\pi/m$ there. Due to the β -equilateral structure, the conformal tiling \mathcal{T}_K brings with it companion conformal tilings: \mathcal{T}_K^β is a conformal tiling for K^β ; \mathcal{T}_K^\dagger is the conformal tiling for the dual complex K^\dagger with edges that are orthogonal bisectors of their dual edges in \mathcal{T}_K ; \mathcal{T}_K^\diamond is a 4-gon tiling for the quad complex K^\diamond . All these are formed by appropriate groupings of the conformal triangles of \mathcal{T}_K^β . Through repeated Schwarz reflection, a single tile of any of these serves to regenerate them all. The reader is referred to [7] for details; Figure 4 there, illustrating the interlocking tilings, is particularly recommended.

The final review topic for understanding the statement and proof of Theorem 1.1 is that of subdivision. A *subdivision* of the planar polygonal complex K is a planar polygonal complex K_s for which each open cell (vertex, edge, or face) of K_s is contained in an open cell of K . In particular, each closed face of K is the union of finitely many closed faces of K_s , and each closed edge of K is a union of finitely many closed edges of K_s . Write $K_s \leq K$. We are interested in *non-trivial* subdivisions in which each face of K is a union of at least two faces of K_s . Faces of K_s are *subfaces* of those of K , faces of K are *aggregates* of those of K_s , and K is an *aggregation* of K_s . Note that K plays two roles throughout this paper. For example, as a CW-complex, K may have its faces decomposed into faces forming K_s . On the other hand, as a combinatorial object, K could still be CW-isomorphic to K_s , $K \cong K_s$. We will encounter precisely this later, where both $K_s \leq K$ and $K \cong K_s$ hold. This should cause no confusion—the use of K depends on context.

Subdivisions of interest in our work are generated by various types of *subdivision rules*. Traditional examples, such as the *Penrose*, *pinwheel*, *sphinx*, and others, and the more nuanced ones introduced by Cannon, Floyd, and Parry using model complexes and subdivision maps, e.g., their *mixed* rule, were discussed at length in the first paper [7]. For use in this first section, however, we introduce the slightly more pedestrian *subdivision operators*.

Define a *seed* $\tau_n, n \geq 1$, to be a rotationally symmetric, oriented, regular CW-decomposition of the standard n -gon Δ_n . The seed is *non-trivial* if it breaks Δ_n into at least two pieces. The collection $\tau = \{\tau_n\}$ of seeds is *edge-compatible* if there

is an integer $N \geq 0$ such that each τ_n subdivides each edge of Δ_n into the same number N of subedges.

Definition. A *subdivision operator* τ is an edge-compatible collection $\tau = \{\tau_n : n = 1, 2, 3, \dots\}$ of rotationally symmetric, oriented, non-trivial seeds. Given any polygonal complex K , $\tau K \leq K$ is the subdivision obtained from K by subdividing each face of polygonal type n according to the pattern τ_n respecting orientation, with edge subdivisions matching on shared edges.

The finiteness, orientability, edge-compatibility, and rotational symmetry of the τ_n ensure that τK is determined unambiguously and is again a locally finite polygonal complex. The notation is fairly consistent with the more nuanced notion of subdivision *rules* that will be described in §3.1, though subdivision *operators* have the advantage of applying to any and all polygonal complexes. Note that if F is a subcomplex of K , τF has the obvious meaning as the subdivision of F induced by τ and is a subcomplex of τK .

We illustrate a variety of rotationally symmetric subdivision algorithms in Figures 1 and 2—the *star*, *barycentric*, *hex*, *pentagonal*, *twisted pentagonal*, and others. All but the twisted pentagonal are, in fact, dihedrally symmetric, a property that will be important in future developments. Each of these algorithms applies to n -gons for any positive n . Rotational symmetry means that one can encode the pattern in a triangular fundamental region for the action of the rotation group, as we have done in the central columns of the two figures: the bottom of the fundamental triangle represents an edge of the face to which the rule is applied and the apex of the triangle represents the fixed point of the action, the face barycenter. Vertices of the subdivision appear in yellow and edges appear as solid segments. By identifying the left and right sides, the pattern may also be encoded as a seed in a 1-gon, as in the left-hand columns of the figures (though our polygonal complexes do not admit 1-gons). In the right-hand columns of Figures 1 and 2 we illustrate seeds for 5-gons or 4-gons, respectively, but of course each algorithm leads to a seed τ_n for each n by identifying a cycle of n of these fundamental regions.

These and similar algorithms may be combined in various ways to describe subdivision operators; for example, we might specify that seed τ_4 comes from the pentagonal algorithm, and all other seeds τ_n , for $n = 2, 3$, and $n \geq 5$, come from the quad algorithm, and since each subdivides an edge into two edges, this describes an edge-compatible subdivision operator $\tau = \{\tau_n\}$. We have a special interest in operators which apply the same algorithm to every n -gon—all quad subdivisions on all faces regardless of polygonal type, or all pentagonal subdivisions, etc. A subdivision operator that applies the same rotationally symmetric subdivision algorithm to every tile is said to be *simple*.

There are two points of view on subdivisions. One is combinatorial and takes place at the level of the tiling complex K —as with the subdivision operators we just described. The other is geometric and takes place at the level of the geometric tiles themselves. In traditional tiling the combinatorial and geometric mesh perfectly: one has a finite collection of model “prototile” shapes and a subdivision (also called a *substitution*) process in which each of those tiles is broken into sub-tiles drawn from that same collection. In Conway’s “pinwheel” tiling, for example, a certain right triangle is partitioned into 5 congruent triangles, each similar to the original with one fifth the area: the subdivision is both combinatorial and concrete.

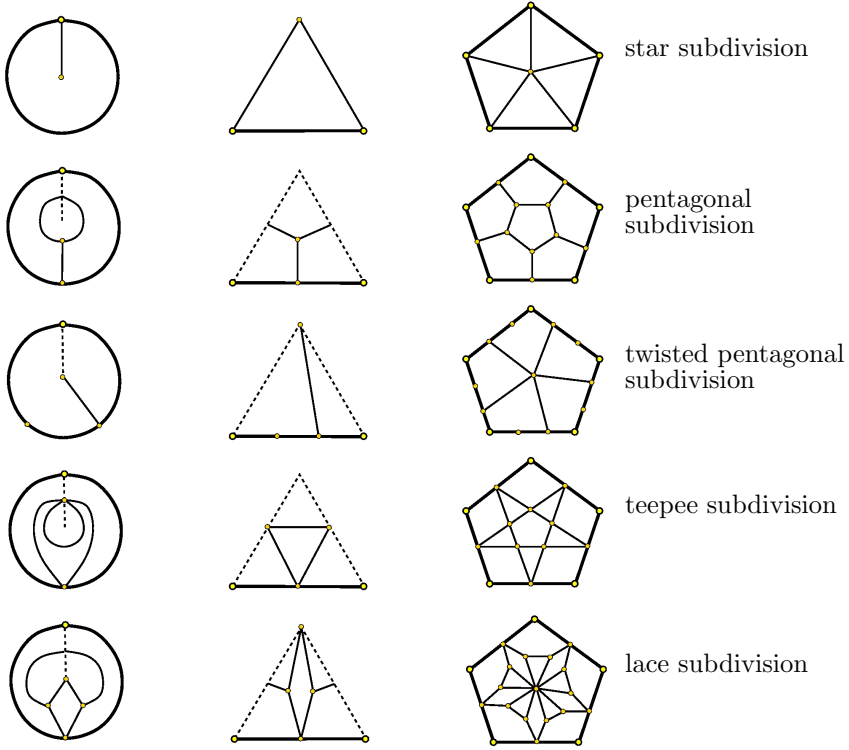


FIGURE 1. Rotationally symmetric subdivision algorithms for n -gons for any $n \geq 1$. Each pattern is encoded in a 1-gon (left-hand column) and in a fundamental triangle (central column), and is applied to a 5-gon (right-hand column).

This geometric compatibility accounts, in part, for the appeal of the relatively few combinatorial patterns that have been realized in traditional hierarchical tilings.

In the conformal setting, on the other hand, one seems at first to have complete freedom: any locally finite polygonal complex K has its conformal tiling \mathcal{T}_K , and after applying a combinatorial subdivision, the new polygonal complex K_s likewise has its conformal tiling \mathcal{T}_{K_s} . However, there is new subtlety here. For starters, there are no prototile shapes; nearly every tile in a non-trivial conformal tiling has a unique shape. The more important fact, however, is that generically the combinatorial and the geometric worlds do not mesh. If you aggregate the tiles from \mathcal{T}_{K_s} associated with some tile t of \mathcal{T}_K —that is, you “forget” the subdivision—the resulting aggregate shape will not in general be the same as the shape of t itself. The special cases in which the aggregate and original shapes agree are described in detail and with examples in [7, §3.7]; they are covered under this definition, adapted from there.

Definition. Let K_s be a combinatorial subdivision of K , $K_s \leq K$. The conformal tiling \mathcal{T}_{K_s} is said to be a *conformal* subdivision of the conformal tiling \mathcal{T}_K if the

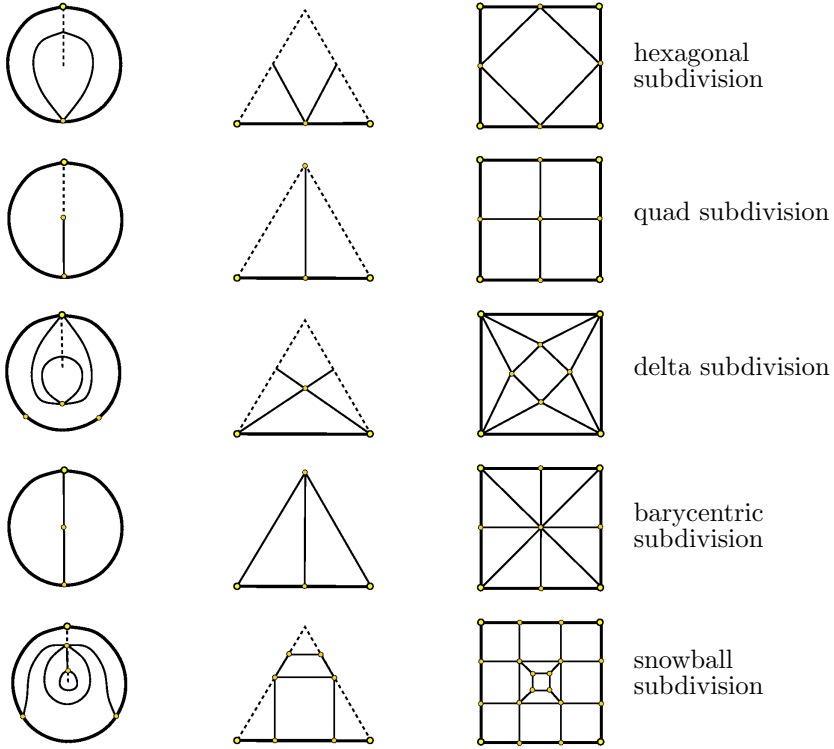


FIGURE 2. Dihedrally symmetric subdivision algorithms for n -gons for any $n \geq 1$. Each pattern is encoded in a 1-gon (left-hand column) and in a fundamental triangle (central column), and is applied to a 4-gon (right-hand column).

forgetful map between their Riemann surfaces, $\pi : S_{K_s} \rightarrow S_K$, is a conformal homeomorphism. A subdivision rule or operator τ is said to be *conformal* if $\mathcal{T}_{\tau K}$ is a conformal subdivision of \mathcal{T}_K whenever τ can be applied to K .

The upshot is that in the case of a *conformal* subdivision rule or operator τ , the conformal tiling $\mathcal{T}_{\tau K}$ (appropriately normalized) is obtained from \mathcal{T}_K by partitioning its tiles in situ, parallel to what happens in traditional tiling. In particular, when the conformal tiling \mathcal{T}_{K_s} is a conformal subdivision of the conformal tiling \mathcal{T}_K , each tile of \mathcal{T}_{K_s} is contained in a tile of \mathcal{T}_K , each tile of \mathcal{T}_K is a union of finitely many tiles from \mathcal{T}_{K_s} , and the tiles of both \mathcal{T}_K and \mathcal{T}_{K_s} are conformally regular and reflective across their respective edges. We may safely write $\mathcal{T}_{\tau K} \leq \mathcal{T}_K$, so we are in position to explore subdivision families for a fixed tiling, and in particular to study the type problem for its local isomorphism class.

We may now prove Theorem 1.1, whose content is reduced to the claim that every dihedrally symmetric simple subdivision operator τ is conformal.

Proof of Theorem 1.1. Suppose $\tau = \{\tau_n\}$ is a dihedrally symmetric, simple subdivision operator. Recall that for each n , τ_n is a combinatorial decomposition of a standard n -gon Δ_n . However, we can impose the β -equilateral conformal structure $|\tau_n^\beta|_{\text{eq}}$, making τ_n into a conformal n -gon. The rotational symmetry of τ_n further ensures that it is a conformally regular n -gon. Thus, for every n , we have at our disposal the conformally regular *model* n -gon τ_n subdivided into a conformally regular and reflective pattern of subtiles.

Consider a conformal tile $t \in \mathcal{T}_K$, say of polygonal type n . Since t and τ_n are both conformally regular, there is a homeomorphism $f_t : \tau_n \rightarrow t$ which is conformal on the interior and takes the n vertices of τ_n to the n -vertices of t . The image of the subtiles of τ_n under f_t now provides an in situ decomposition of t into conformally regular and reflective subtiles. Applying the map f_t for each $t \in \mathcal{T}_K$ results in a global tiling $\tau\mathcal{T}$ by conformally regular polygons with the combinatorics of τK . Note that edge compatibility is due to the dihedral symmetry in the τ_n .

It remains only to show that $\tau\mathcal{T}$ is also reflective. Suppose tiles $s, s' \in \tau\mathcal{T}$ share an edge e . If both are subtiles of the same aggregate tile $t \in \mathcal{T}_K$, then they are reflective across e by construction. Otherwise, s, s' are subtiles of $t, t' \in \mathcal{T}_K$, respectively, with t and t' sharing an edge E which has e as a subedge.

It is here that the *simple* and *dihedrally symmetric* properties of τ are needed. Let u, u' denote the unions in t, t' , respectively, of the two triangles of $|K^\beta|_{\text{eq}}$ neighboring E , so u and u' are conformal quadrangles. The reflective property described earlier now provides an anticonformal reflection $\psi : u \rightarrow u'$ with fixed point set E . Suppose t is an n -gon, t' an n' -gon. Because τ is a *simple* subdivision operator, τ_n and $\tau_{n'}$ share the same pattern on fundamental domains of the rotational group actions for t and t' . That is, the pattern of subtiles of t lying in u is identical to that of t' lying in u' . Here we recall that τ_n and $\tau_{n'}$ are in fact dihedrally symmetric, so even though $\psi : u \rightarrow u'$ is orientation reversing, it still maps subtiles to subtiles.

Now return to consider the subtiles s, s' , which share edge $e \subset E$. Due to the regularity and dihedral symmetry of the decompositions τ_n and $\tau_{n'}$, it is not difficult to check that the triangles of $|s^\beta|_{\text{eq}}$ and $|s'^\beta|_{\text{eq}}$ neighboring e lie entirely in u and u' , respectively. The restricted map $\psi|_{s \cap u}$ is thus an anticonformal map with fixed point set e and identifying these pairs of triangles and their vertices. This establishes the reflective structure for s and s' across e .

We conclude that the tile-by-tile decomposition of \mathcal{T}_K denoted by $\tau\mathcal{T}$ is conformally regular and reflective. Since $\tau\mathcal{T}$ and $\mathcal{T}_{\tau K}$ both have the combinatorics of τK , Theorem 1.2 implies that they are conformally equivalent. It follows that τ is a conformal subdivision operator and $\mathcal{T}_{\tau K}$ may be realized as a conformal subdivision of \mathcal{T}_K . In other words, $\mathcal{T}_{\tau K} = \tau\mathcal{T}_K$. \square

We close this section with an example. Though our interest is in infinite planar complexes, the methods apply more widely and we illustrate with a compact example introduced in our earlier paper.

Example 1.1. The example starts with a combinatorial cube K , that is, a topological sphere decomposed as six 4-gons. \mathcal{T}_K is a tiling of the Riemann sphere \mathbb{S}^2 . We apply successive stages of the “snowball” subdivision rule τ whose seed τ_4 (the only one we need) is shown at the bottom right in Figure 2. This seed subdivides each 4-gon into thirteen 4-gons. Since τ is dihedrally symmetric and simple, τ is *conformal* by Theorem 1.1. Figure 3 shows (numerical approximations of) four




FIGURE 3. This represents a conformally correct picture of the first four stages of the conformal “snowball” subdivision rule: Blue subdivides the original bold black-edged cube; the green then subdivides the blue; the red, the green; and the black, the red. In situ conformal subdivision may be continued ad infinitum.

stages of subdivision. Subdivision occurs in situ and can be continued ad infinitum without disturbing earlier stages. \checkmark

2. COMPLEXES AND THEIR LOCAL ISOMORPHISM CLASSES

In the traditional world of aperiodic euclidean tilings, an appropriate tiling determines a tiling space, the so-called *continuous hull* of the tiling, an action of the translation group on the hull, and a *canonical transversal* for the action that encodes pointed tilings locally isomorphic to the original. In an impressive doctoral thesis, Maria Ramirez-Solano [21] managed to generalize much of this to the conformal setting for certain pentagonal tilings. We aspire to follow her lead and develop the geometric structures which might support a more comprehensive parallel conformal theory. To provide the combinatorial basis for that effort, though perhaps beyond the immediate needs of this paper, we define various spaces of rooted polygonal complexes and identify necessary topological properties of the local isomorphism classes. Traditional tiling experts will see mirrored here several familiar structures, but only those that emerge from the combinatorics of tilings rather than their geometry.

2.1. Spaces of planar polygonal complexes. Our main space is \mathbf{RC} , the *space of rooted planar polygonal complexes*. The elements are the rooted isomorphism classes of pairs (K, f) , where K is an oriented planar polygonal complex and f is a fixed *root face* of K . In this setting, a rooted isomorphism of (K, f) to (L, g) is a cellular isomorphism of K to L that takes f to g (note: f and g must be of the same polygonal type). We will employ the notational use of sans serif letters to denote elements of \mathbf{RC} , so we may write $(K, f) \in \mathbf{k} \in \mathbf{RC}$. There is in this setting a map $*$: $\mathbf{RC} \rightarrow \mathbf{C}$ to the set of isomorphism classes of planar polygonal complexes \mathbf{C} that forgets the root face and that is countable-to-one. Thus, for example, $K \in \mathbf{k}^*$ whenever $(K, f) \in \mathbf{k}$. (As is usual practice, we do not distinguish between K and its equivalence class $[K]$, so, e.g., we write $K = \mathbf{k}^* = [K]$.) Recall, our complexes are always oriented and our maps orientation-preserving.

To define a metric on \mathbf{RC} , we first define a combinatorial distance within each complex. Given faces f and g of a planar polygonal complex K , the *fat path distance* from f to g (terminology from [10]) is the smallest integer k for which there exists a chain $\{f = f_0, f_1, \dots, f_k = g\}$ of faces such that f_j and f_{j+1} share a common edge in K for $j = 0, 1, \dots, k-1$. This distance gives a metric on the collection of faces of K and may also be described as the edge metric of the dual complex K^\dagger .

The metric we now impose on \mathbf{RC} is modeled after the so-called *big ball* metric in traditional tiling or the *local* metric in planar graph theory; however, completeness requires a slightly more complicated set up than usual—we must use *filled* balls. Given a face f of K and $n > 0$, define the (*combinatorial*) *n-neighborhood* $C_K(f, n)$ as the subcomplex of K formed by those faces g having fat path distance from f of at most n . Its complement in K has one unbounded component, but may also have bounded components. We therefore define the *filled n-neighborhood* $B_K(f, n)$ to be the smallest *simply connected* subcomplex of K containing $C_K(f, n)$. The combinatorial boundary of $B_K(f, n)$ is a simple closed edge path that separates f from the ideal boundary of K and is called the *outer sphere* of *radius* n centered on f . The complex K is the nested union

$$B_K(f, 1) \subset B_K(f, 2) \subset \dots \subset B_K(f, n) \subset B_K(f, n+1) \subset \dots$$

of filled neighborhoods. Note that $B_K(f, n)$ may not lie in the interior of $B_K(f, n+1)$, but it does lie in the interior of $B_K(f, N)$ for $N > n$ sufficiently large. Define $B_K(f, \infty)$ to be the whole complex K .

Given $\mathbf{k}, \mathbf{k}' \in \mathbf{RC}$ and rooted complexes $(K, f) \in \mathbf{k}$ and $(K', f') \in \mathbf{k}'$, let

$$\rho(\mathbf{k}, \mathbf{k}') = \rho((K, f), (K', f')) = e^{-n}$$

where, either n is the largest integer, or the symbol ∞ , for which the rooted complexes $(B_K(f, n), f)$ and $(B_{K'}(f', n), f')$ are rooted isomorphic, or $n = -1$ whenever f and f' have differing polygonal types. This defines a metric on \mathbf{RC} that satisfies the property that, whenever $\mathbf{j}, \mathbf{k}, \mathbf{l} \in \mathbf{RC}$, then either

$$\rho(\mathbf{j}, \mathbf{l}) \leq \rho(\mathbf{j}, \mathbf{k}) = \rho(\mathbf{k}, \mathbf{l})$$

or one of the permutations of this relation, holds. This is equivalent to

$$\rho(\mathbf{j}, \mathbf{l}) \leq \max \{ \rho(\mathbf{j}, \mathbf{k}), \rho(\mathbf{k}, \mathbf{l}) \},$$

which is the defining triangle inequality for an *ultrametric*, a topic of significant application in p -adic analysis. (See, e.g., [22, §2.1] for basic properties of ultrametries.)

Theorem 2.1. *The metric ρ on the space of rooted planar polygonal complexes \mathbf{RC} is bounded and complete, and the metric space (\mathbf{RC}, ρ) is totally disconnected.*

Proof. Obviously ρ is bounded since $n \geq -1$ in the definition of ρ so that, in fact, $\text{diam}_\rho \mathbf{RC} = e$. Let \mathbf{k}_i , $i = 1, 2, 3, \dots$, be a ρ -Cauchy sequence in \mathbf{RC} . By passing to a subsequence if necessary, we may assume that for each i , $\rho(\mathbf{k}_i, \mathbf{k}_{i+1}) \leq e^{-i}$. Choose representatives $(K_i, f_i) \in \mathbf{k}_i$, and observe that our definition of the metric ρ implies that for each i , the filled neighborhood $B_{K_i}(f_i, i)$ is rooted isomorphic to the filled neighborhood $B_{K_{i+1}}(f_{i+1}, i)$. Let $h_i : B_{K_i}(f_i, i) \rightarrow B_{K_{i+1}}(f_{i+1}, i)$ be a rooted isomorphism and enlarge the codomain of h_i to $B_{K_{i+1}}(f_{i+1}, i+1)$ to obtain the directed sequence

$$B_{K_1}(f_1, 1) \xrightarrow{h_1} B_{K_2}(f_2, 2) \xrightarrow{h_2} B_{K_3}(f_3, 3) \xrightarrow{h_3} \dots$$

Since the $B_{K_i}(f_i, i)$ are nested and simply connected, and since for each integer i , there exists an integer $N > i$ such that image of $B_{K_i}(f_i, i)$ under $h_{N-1} \circ \dots \circ h_i$ is contained in the interior of the disc $B_{K_N}(f_N, N)$, the direct limit complex $\varinjlim B_{K_i}(f_i, i)$ is a simply connected planar polygonal complex. Let \mathbf{k} denote the rooted isomorphism class of the rooted planar polygonal complex (K, f) , where $K = \varinjlim B_{K_i}(f_i, i)$ and $f \in K$ corresponds to f_1 . It is easy to show that $\mathbf{k}_i \rightarrow \mathbf{k}$ as $i \rightarrow \infty$, and so ρ is complete. If we do not distinguish between rooted isomorphic complexes, this direct sequence may be treated as a sequence of set containments and the direct limit is simply its union.

For any point $\mathbf{k} \in \mathbf{RC}$ and positive integer n , the metric ball $B_\rho(\mathbf{k}, e^{-n}/2)$ of radius $e^{-n}/2$ is both open and closed, since the metric takes on only the values in the countable set $\{0\} \cup \{e^{-k} : k = -1, 0, 1, \dots\}$. This implies that (\mathbf{RC}, ρ) is totally disconnected (as is the case with every ultrametric space). \square

For each integer $n \geq 3$, let \mathbf{RC}^n be the subspace of \mathbf{RC} of rooted isomorphism classes of planar polygonal complexes of degree bounded by n . ($\mathbf{RC}^3 = \emptyset$ is included for later convenience.) Let $\mathbf{RC}^\omega = \bigcup_{n=3}^\infty \mathbf{RC}^n$, the space of bounded degree, rooted planar polygonal complexes. The subsets \mathbf{C}^n and \mathbf{C}^ω of \mathbf{C} are defined in the obvious manners as the images of the respective spaces \mathbf{RC}^n and \mathbf{RC}^ω under the forgetful map $*$.

Theorem 2.2. *For each $n \geq 3$, the subset \mathbf{RC}^n is a closed, nowhere dense subspace of \mathbf{RC} .*

Proof. Let $\mathbf{k}_i \rightarrow \mathbf{k}$ where $\mathbf{k}_i \in \mathbf{RC}^n$ and $\mathbf{k} \in \mathbf{RC}$. An arbitrarily large finite filled neighborhood of the root face of \mathbf{k}^* is isomorphic to a filled neighborhood of the root face of \mathbf{k}_i^* for large enough i , so the degree of \mathbf{k}^* is at most n so that $\mathbf{k} \in \mathbf{RC}^n$.

\mathbf{RC}^n is nowhere dense in \mathbf{RC} since every element \mathbf{k} of \mathbf{RC}^n is the limit of a sequence \mathbf{k}_i , each term with at least one vertex of degree greater than or equal to $n+1$, as the reader may construct rather easily. \square

Corollary 2.3. *The space \mathbf{RC}^ω of bounded degree, rooted planar polygonal complexes is a dense F_σ subspace of \mathbf{RC} , and the space $\mathbf{RC} - \mathbf{RC}^\omega$ of unbounded degree rooted planar polygonal complexes is a dense G_δ subspace of \mathbf{RC} .*

The next property delineates special compact subsets of \mathbf{RC}^n . We say that a planar polygonal complex K satisfies a θ -isoperimetric inequality if every finite

combinatorial disc $D \subset K$ satisfies the θ -isoperimetric inequality, meaning that

$$\text{size}(D) < \theta(\text{size}(\partial D)).$$

Here, $\theta : \mathbb{N} \rightarrow \mathbb{R}$ is a positive function, $\text{size}(D)$ is the number of faces in D , and $\text{size}(\partial D)$ is the number of edges in the simple closed edge path ∂D forming the combinatorial boundary of D . For each integer $n \geq 3$ and positive function θ , let $\mathbf{RC}^{n,\theta}$ consist of those $\mathbf{k} \in \mathbf{RC}^n$ for which the complex \mathbf{k}^* satisfies a θ -isoperimetric inequality.

Theorem 2.4. *For each $n \geq 3$ and positive function θ , the subset $\mathbf{RC}^{n,\theta}$ is a compact, nowhere dense subset of \mathbf{RC} .*

Proof. Since \mathbf{RC}^n is nowhere dense in \mathbf{RC} , so too is $\mathbf{RC}^{n,\theta} \subset \mathbf{RC}^n$. That $\mathbf{RC}^{n,\theta}$ is closed in \mathbf{RC} is again easily seen by looking at large neighborhoods of the root.

To verify compactness, we show that $\mathbf{RC}^{n,\theta}$ is totally bounded in \mathbf{RC} , which for closed subsets of complete metric spaces is equivalent to compactness. Recall that a metric space is totally bounded if, for each $\varepsilon > 0$, there exists a finite ε -dense subset of the space. Let $\varepsilon > 0$ and choose a positive integer m so that $e^{-m} < \varepsilon$. Let $(H_1, f_1), \dots, (H_J, f_J)$ be a list of all the finite, connected planar polygonal CW-complexes of degree at most n that appear as the filled m -neighborhood of a face in some planar polygonal complex in $\mathbf{RC}^{n,\theta}$, up to isomorphism. This list is finite precisely because, for all complexes in $\mathbf{RC}^{n,\theta}$, the degree is bounded by n and the θ -isoperimetric inequality holds; indeed, the bounded degree condition guarantees that there are only finitely many isomorphism classes of m -neighborhoods $C_K(f, m)$ as K ranges over $\mathbf{RC}^{n,\theta}$ and f ranges over K , and for each such class, the θ -isoperimetric inequality guarantees that there are only finitely many ways to fill the holes of any $C_K(f, m)$ to obtain its filled m -neighborhood $B_K(f, m)$. The upshot is that these two conditions together place a bound $M > 0$ on the number of faces in any such filled neighborhood, and there are then only finitely many ways to arrange a set of at most M combinatorial polygons of polygonal type bounded by n to form a filled neighborhood. For each $i = 1, \dots, J$, let K_i be a planar polygonal complex that contains an isomorphic copy of H_i as the filled m -neighborhood $B_{K_i}(f_i, m)$. Letting $\mathbf{k}_i = (K_i, f_i)$, the set $\{\mathbf{k}_1, \dots, \mathbf{k}_J\}$ is a finite ε -dense subset of $\mathbf{RC}^{n,\theta}$. \square

Key features in this proof, the bound on degree and the θ -isoperimetric inequality, play the role in compactness that finite local complexity plays in the traditional setting. However, whereas the geometry attached to traditional tilings brings the isoperimetric condition for free, in the purely combinatorial setting we must impose it as a hypothesis. The m -neighborhoods $C_K(f, m)$ can have “holes”, bounded complementary components, that are filled in to get $B_K(f, m)$. Without an isoperimetric bound, these fillings could require arbitrarily many faces, independent of m . The isoperimetric condition is important for compactness in this paper, so we add it to the definition of finite local complexity from our earlier paper.

Definition. The planar polygonal complex K is said to have *finite local complexity*, abbreviated as FLC, if the degree of K is bounded and K satisfies a θ -isoperimetric inequality for some positive function $\theta : \mathbb{N} \rightarrow \mathbb{R}$.

To be rather clear on its intended use, we separate out as a corollary the important feature of the preceding proof as applied to a single complex K .

Corollary 2.5. *Suppose the planar polygonal complex K has FLC. Then, up to isomorphism, there are only finitely many filled m -neighborhoods in K of a given radius m .*

2.2. Local isomorphism classes of complexes. In studying conformal tilings, we are not interested in a single tiling \mathcal{T}_K as much as in the family of those tilings \mathcal{T}_L for which L is locally isomorphic to K . The notion of local isomorphism arises naturally.

For planar polygonal complexes K and L , we say that K *locally embeds* in L , written as $K \preceq L$, if every finite combinatorial patch $F \subset K$ isomorphically embeds in L , $F \hookrightarrow L$. We say K is *locally isomorphic* with L if $K \preceq L$ and $L \preceq K$, and in this case we write $K \sim L$. We let (K) denote the local isomorphism class of the planar polygonal complex K so that $(K) = \{L \in \mathbf{C} : K \sim L\}$. The relation \preceq is reflexive and transitive, and hence a pre-order on \mathbf{C} , and it defines a partial order on the set of local isomorphism classes of elements of \mathbf{C} .

We will see a dichotomy shortly, as (K) is either a singleton or is uncountably infinite. Let $\text{Li}(K)$ denote the pre-image of (K) under the forgetful map $*$ defined on \mathbf{RC} , so that $\text{Li}(K)$ is the set of rooted polygonal complexes $k \in \mathbf{RC}$ with $k^* \sim K$. The set $\text{Li}(K)$ is the combinatorial analogue of the canonical transversal $\Xi = \Xi_T$ for the action of the translation group on the continuous hull in the traditional theory of aperiodic tilings (cf. [3, Chp 5]).

Our interest is in uncovering the structure and properties of the local isomorphism class (K) and of the *rooted local isomorphism class* $\text{Li}(K)$. This working lemma is proven with a straightforward diagonalization argument.

Lemma 2.6. *Let K and L be planar polygonal complexes and let $L_1 \subset L_2 \subset \dots$ be a sequence of finite, connected subcomplexes exhausting L , meaning that $L = \bigcup_{n=1}^{\infty} L_n$. For each positive integer n , let $h_n : L_n \hookrightarrow K$ be an isomorphic embedding of complexes. If there are faces $g_0 \in L_1$ and $f_0 \in K$ such that $h_n(g_0) = f_0$ for all $n \geq 1$, then K is isomorphic to L , $K \cong L$.*

The first application is to confirm that planar polygonal complexes that are highly symmetric globally are uninteresting in terms of local isomorphism. Let $\text{Aut}(K)$ denote the set of orientation-preserving automorphisms of K .

Theorem 2.7. *If $\text{Aut}(K)$ acts cocompactly on the planar polygonal complex K and $L \preceq K$, then $L \cong K$. It follows that the local isomorphism class (K) is a singleton.*

Proof. Let F be a finite subcomplex of K that serves as a fundamental region for the action of $\text{Aut}(K)$ on K and write $L = \bigcup_{n=1}^{\infty} L_n$, where $L_1 \subset L_2 \subset \dots$ are finite, connected subcomplexes. Let g be a face of L_1 and let $(L_n, g) \cong (K_n, f_n)$, where K_n is a subcomplex of K isomorphic to L_n . Since $\text{Aut}(K)$ acts with fundamental region F , we may assume that for each n , f_n is a face of F . Since F is finite, infinitely many of the faces f_n must be the same face f , so we may assume without loss of generality that $f_n = f$ for all n . Apply the working Lemma 2.6. \square

Let H be a finite connected CW-complex. If K contains an isomorphic copy of H , then we say that the H is *represented* in K . If there are infinitely many disjoint isomorphic copies of H in K , we say that H is *infinitely represented* in K . If those copies are uniformly distributed, we say H is *quasi-dense* in K . More precisely, H is *quasi-dense* in K if there exists a positive integer n such that every vertex of

K is contained in an edge path of length at most n that meets a vertex in some embedded copy of H .

Definition. If every finite connected subcomplex H of K is quasi-dense in K , we say that K is *combinatorially repetitive*.

When $\text{Aut}(K)$ acts cocompactly on K , then every finite connected subcomplex is quasi-dense. However, our interests are in settings without such global symmetries.

Theorem 2.8. *If there is a finite connected complex H that is finitely represented in K , then the local isomorphism class (K) is a singleton. Alternately, if (K) is not a singleton, then every finite connected subcomplex of K is infinitely represented in K .*

Proof. The proof is similar to that of Theorem 2.7. Assume that L is a planar polygonal complex such that $L \sim K$. Let $L_1 \subset L$ be a finite, connected subcomplex that contains an isomorphic copy of H , which exists because $K \preceq L$. Let F be a finite, connected subcomplex of K that contains all subcomplexes of K isomorphic to L_1 , which exists since H , and therefore L_1 , is finitely represented in K . Write $L = \bigcup_{n=1}^{\infty} L_n$, where $L_1 \subset L_2 \subset \dots$ are finite, connected subcomplexes which exhaust L . Let g be a face of L_1 and for each n let $(L_n, g) \cong (K_n, f_n)$, where K_n is a subcomplex of K isomorphic to L_n . By our choices of F and L_1 , for each n , f_n is a face of F . Since F is finite, infinitely many of the faces f_n must be the same face f , and we may assume without loss of generality that $f_n = f$ for all n . Apply the working Lemma 2.6. \square

Symmetries of a complex K typically bring to mind global symmetries, elements of $\text{Aut}(K)$. These don't exist in traditional aperiodic tilings by definition. Yet if one looks at, say, the familiar Penrose tiling, the eye immediately picks out large finite patches that seem to repeat everywhere. Local isomorphism provides the refined notion of symmetry that the eye picks out: namely, numerous *local symmetries* between finite patches in differing regions of the tiling that do not extend to global symmetries.

In the combinatorial setting of this paper, Theorems 2.7 and 2.8 position complexes having singleton local isomorphism classes at both ends of the spectrum of symmetry. The middle ground, complexes K without global symmetry, but such that every finite patch reappears infinitely often, have rich local isomorphism classes (K) , the ones of most interest here.

We close this section with a detailed examination of the topology of the rooted local isomorphism class $\text{LI}(K)$. We find that the condition of combinatorial repetitiveness is the key ingredient for identifying when $\text{LI}(K)$ is compact.

Lemma 2.9. *The set $\text{LI}(K)$ is a dense G_δ subspace of the complete metric space $\overline{\text{LI}(K)}$, the closure of $\text{LI}(K)$ in \mathbf{RC} , and as such is a completely metrizable Baire space.*

Proof. Obviously $\text{LI}(K)$ is dense in the closure $\overline{\text{LI}(K)}$. Fix a face f of K and for each positive integer n let

$$U_n = \left\{ l \in \overline{\text{LI}(K)} : B_K(f, n) \hookrightarrow l^* \right\}.$$

To verify the theorem, we observe that each U_n is open in $\overline{\text{LI}(K)}$ and that $\text{LI}(K) = \bigcap_{n=1}^{\infty} U_n$.

U_n is open in $\overline{\text{LI}(K)}$: Let $\mathfrak{l} = (L, g) \in U_n$ and choose $N \geq n$ such that $B_K(f, n) \hookrightarrow B_L(g, N)$, which is possible since $B_K(f, n)$ embeds isomorphically in L . Then the set $B_\rho(\mathfrak{l}, e^{-N}) \cap \overline{\text{LI}(K)}$ is an open neighborhood of \mathfrak{l} in $\overline{\text{LI}(K)}$ that is contained in U_n , and U_n is open in $\overline{\text{LI}(K)}$.

$\text{LI}(K) = \bigcap_{n=1}^{\infty} U_n$: The “ \subset ” containment follows from the observation that $\text{LI}(K) \subset U_n$ for each n . For the containment “ \supset ”, suppose that $\mathfrak{l} = (L, g) \in \overline{\text{LI}(K)}$ is an element of U_n for all n . Then $B_K(f, n) \hookrightarrow \mathfrak{l}^* = L$ for all n , implying that every finite connected subcomplex of K embeds isomorphically in L , so that $K \preceq L$. For any positive integer n , choose $\mathfrak{e} = (E, h) \in \text{LI}(K)$ such that $\rho(\mathfrak{e}, \mathfrak{l}) < e^{-n}$, which is possible since \mathfrak{l} is in the closure of $\text{LI}(K)$. Then $B_L(g, n) \cong B_E(h, n) \hookrightarrow K$, the existence of the embedding following from $E \sim K$. Since n is arbitrary, this implies that every finite connected subcomplex of L embeds isomorphically in K , so that $L \preceq K$. We conclude that $L \sim K$, hence $\mathfrak{l} \in \text{LI}(K)$.

$\text{LI}(K)$ is completely metrizable since it is a G_δ subspace of a complete metric space (by Mazurkiewicz’s Theorem), and hence is a Baire space (by the Baire Category Theorem). \square

The following “invariance of domain” observation will be used several times later in the paper.

Lemma 2.10. *Let K and L be planar polygonal complexes, let f be a face of K , and let g be a face of L . Let H be a subcomplex of K containing f so that (H, f) is rootedly isomorphic to the combinatorial n -neighborhood $C_L(g, n)$ (resp., its filled n -neighborhood $B_L(g, n)$); then $H = C_K(f, n)$ (resp., $H = B_K(f, n)$).*

Theorem 2.11. *If K is a planar polygonal complex whose local isomorphism class (K) is not a singleton, then the space $\text{LI}(K)$ has no isolated points. It follows that $\text{LI}(K)$ is a completely metrizable, uncountably infinite, perfect Baire space. In particular, either (K) has one element or it has uncountably many elements.*

Proof. Let $\mathfrak{l} = (L, g) \in \text{LI}(K)$. By hypothesis there exists a planar polygonal complex J that is locally isomorphic to, but not isomorphic to L . Let n be a positive integer. Since $L \sim J$, there is an isomorphic embedding $h : B_L(g, n) \hookrightarrow J$ with, say, $h(g) = f$. By the preceding lemma, the image of the filled n -neighborhood $B_L(g, n)$ under h is precisely the filled n -neighborhood $B_J(f, n)$, and this implies that $0 \neq \rho(\mathfrak{j}, \mathfrak{l}) \leq e^{-n}$, where $\mathfrak{j} = (J, f)$. It follows that \mathfrak{l} is not an isolated point of $\text{LI}(K)$.

Lemma 2.9 guarantees that the metric space $\text{LI}(K)$ is a completely metrizable Baire space and the argument of this proof thus far guarantees that it is perfect. This implies that $\text{LI}(K)$ cannot be countable.

The fact that (K) is uncountably infinite when not a singleton is an immediate consequence of the fact that (K) is the image of the uncountable set $\text{LI}(K)$ under the countable-to-one forgetful map $*$. \square

The set $\text{LI}(K)$ generally fails to be closed in \mathbf{RC} . The next theorem demonstrates that its closure captures precisely those complexes that precede K under the pre-order \preceq .

Theorem 2.12. *For an arbitrary planar polygonal complex K , the closure of $\text{LI}(K)$ in the space \mathbf{RC} is*

$$\overline{\text{LI}(K)} = \{\mathfrak{l} \in \mathbf{RC} : \mathfrak{l}^* \preceq K\}.$$

Proof. (\subset): Let $(L, g) = l \in \overline{\text{LI}(K)}$ and let H be a combinatorial patch in L . Choose n so that $H \subset B_L(g, n)$ and let $(J, f) = j \in \text{LI}(K)$ such that $\rho(j, l) < e^{-n}$. Then the neighborhoods $B_L(g, n)$ and $B_J(f, n)$ are isomorphic and therefore $H \hookrightarrow J$. Since $J \sim K$, we conclude that $H \hookrightarrow K$, and therefore $l^* = L \preceq K$.

(\supset): Suppose that $(L, g) = l \in \mathbf{RC}$ with $L = l^* \preceq K$. Then for any n , there is an isomorphic embedding $h_n : B_L(g, n) \hookrightarrow K$. By Lemma 2.10, the image $h_n(B_L(g, n)) = B_K(h_n(g), n)$. Hence $\rho(l, k_n) \leq e^{-n}$, where $k_n = (K, h_n(g)) \in \text{LI}(K)$, and therefore $l \in \overline{\text{LI}(K)}$. \square

We now want to ask, when K is infinite, when the perfect, totally disconnected metric space $\text{LI}(K)$ is compact, and therefore a Cantor set. First note that if K has FLC, say with degree bounded by a positive integer d and satisfying a θ -isoperimetric inequality, then $\text{LI}(K)$ is a subspace of the compact set $\mathbf{RC}^{d, \theta}$ (Theorem 2.4), and the question becomes: When is the space $\text{LI}(K)$ closed in $\mathbf{RC}^{d, \theta}$?

Theorem 2.13. *If K is combinatorially repetitive, then $\text{LI}(K)$ is a closed subspace of \mathbf{RC} , and therefore complete in the metric ρ , and every complex in (K) is combinatorially repetitive. If in addition (K) is infinite and K has FLC, then $\text{LI}(K)$ is a Cantor set.*

Proof. By Theorem 2.12, to verify that $\text{LI}(K)$ is closed all we need show is that $l \in \text{LI}(K)$ whenever $(L, g) = l \in \mathbf{RC}$ satisfies $L \preceq K$. Let $L \preceq K$ and let H be a finite combinatorial patch in K . Since K is combinatorially repetitive, there exists an integer n such that an isomorphic copy of H appears in the combinatorial n -neighborhood of every face of K . Since $L \preceq K$, there is an isomorphic embedding $h : B_L(g, n) \hookrightarrow K$. By Lemma 2.10, $h(B_L(g, n)) = B_K(h(g), n)$ and our choice of n guarantees that $H \hookrightarrow B_K(h(g), n)$. It follows that $H \hookrightarrow B_L(g, n) \subset L$ and therefore $K \preceq L$. Therefore $K \sim L$ and $(L, g) = l \in \text{LI}(K)$.

That L is combinatorially repetitive whenever $K \sim L$ is left as an exercise. If K is of bounded degree d , satisfies the θ -perimetric inequality, and is combinatorially repetitive, then $\text{LI}(K)$ is a closed subspace of the compact set $\mathbf{RC}^{d, \theta}$. If in addition (K) is infinite, Theorem 2.11 implies that $\text{LI}(K)$ is perfect and, therefore, is a Cantor set. \square

Our last result is a consequence of the two preceding theorems.

Corollary 2.14. *If $\text{LI}(K)$ is a closed subspace of \mathbf{RC} , then $L \preceq K$ if and only if $L \sim K$. In particular, if K is combinatorially repetitive, then $L \preceq K$ if and only if $L \sim K$.*

3. COMBINATORIAL HIERARCHIES

For the rest of the paper, planar polygonal complexes K will be referred to as *combinatorial tilings*, their “tiles” \mathbf{k} being combinatorial polygons. Our interest lies with sequences of tilings which display *hierarchical structure*. This could take many meanings, but for us it has to do with subdivision (or substitution) processes, originally combinatorial but eventually conformal. This section develops the basics of combinatorial hierarchies and quite general subdivision rules τ . In the next section we make the conditions on τ progressively stronger, reaping progressively stronger features in their hierarchies. We are then prepared to move to the parallel conformal tiling story in §5.

Recall that $L \leq K$ indicates that L is a *combinatorial subdivision* of K . In this case, we also say K is a *combinatorial aggregation* of L .

Definition. The combinatorial tiling K *exhibits a combinatorial hierarchy* if there is a bi-infinite sequence $\{K_n : n \in \mathbb{Z}\}$, called a *combinatorial hierarchy* for K , of combinatorial tilings indexed by the integers such that the following three conditions hold:

- (1) $K_0 = K$;
- (2) $K_{n+1} \leq K_n \ \forall n \in \mathbb{Z}$ (K_{n+1} is a combinatorial subdivision of K_n);
- (3) $K_{n+1} \sim K_n \ \forall n \in \mathbb{Z}$ (K_{n+1} is locally isomorphic to K_n).

We bring condition (3) especially to the reader's attention: a "hierarchy" is not just a bi-infinite sequence of nested tilings since the members must belong to the same local isomorphism class.

When K exhibits a combinatorial hierarchy $\{K_n\}$, there is a subdivision process which we will denote by σ_n for each n so that $\sigma_n K_n = K_{n+1}$. The collection $\{\sigma_n\}$ is called the *subdivision sequence* for the hierarchy $\{K_n\}$. For $m \geq n$ write σ_n^m for the rule for K_n that applies $\sigma_n, \dots, \sigma_{m-1}$ successively to obtain the complex K_m . For now, the σ_n may be quite ad hoc processes. (In our notation, rule σ_m^m is the identity and $\sigma_n^{n+1} = \sigma_n$.)

We should point out here that there are many ad hoc ways to generate hierarchies. It is useful to observe, for example, that any bi-infinite subsequence of a hierarchy is itself a hierarchy. However, our primary concern is with hierarchies generated by subdivision rules τ , so let's look at some details.

3.1. Subdivision rules. Up to this point, our "tiles" have simply been combinatorial polygons. However, circumstances generally require more sophistication. From now on tiles will be classified as to "type". By default, the type simply refers to the number n of sides (and of vertices), but optionally type may include additional information, such as a "label", an orientation, and/or a designated *principal* vertex. In a sense, type serves as a substitute in the combinatorial setting for the way that shape is used to distinguish tiles in traditional tiling. In any case, *We assume henceforth that there is some fixed set \mathfrak{T} of tile types behind the tilings we are studying at the moment, with each tile corresponding to a unique member of \mathfrak{T} .* Often we need no details about the members of \mathfrak{T} , or only generic properties such as finiteness of \mathfrak{T} (which is typical). At the other extreme, for specific subdivision rules one must fully describe \mathfrak{T} ; for instance, the Cannon, Floyd, and Parry "mixed" rule, described in Figure 24 of [7], has six (6) tile types, four 4-gons, and two 3-gons.

Definition. A *subdivision rule* τ is, in fact, a collection of rules, one associated with each tile of a 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} . The rule τ is said to be *finite* if \mathfrak{T} is finite.

Some cautionary comments are in order on subdivision compatibility. Two tiles of τK must intersect only at isolated vertices and/or along full edges; thus a vertex of one tile cannot fall in the interior of an edge of a neighboring tile. If tiles \mathbf{k}_1 and \mathbf{k}_2 of K share a common edge e , then their subdivisions $\tau \mathbf{k}_1$ and $\tau \mathbf{k}_2$ must decompose e in identical fashion. We covered ourselves in the case of subdivision operators $\tau = \{\tau_n\}$ by the requirement of "edge compatibility" of the seeds τ_n

in §1, so τ can be applied to virtually any combinatorial tiling. Things are not so straightforward in general. Though a subdivision rule τ may apply to each individual tile type in a given combinatorial tiling K , it might not apply to K as a whole due to edge incompatibilities. Conditions may seem easy to state, but when designing a subdivision rule from scratch with the aim of making repeated application, compatibility can be devilishly difficult to guarantee.

Introducing tile types is rather innocuous when it comes to the mechanics, but it does require some adjustments in topics we have discussed earlier.

- In a given setting, we assume that \mathfrak{T} is *minimal*, meaning there are no extraneous tile types included. Thus, every type $\mathfrak{t} \in \mathfrak{T}$ is represented in K or in a subdivision $\tau^k K$ for some positive integer k . For example, a subdivision operator $\tau = \{\tau_n\}$ as defined in §1 is nominally infinite, but if the tiles of K and its subdivisions are m -gons for only some finite list $\{m_1, m_2, \dots, m_k\}$ of values for m , then \mathfrak{T} is finite and τ acts as a *finite* subdivision rule.
- A combinatorial tiling K having *finite local complexity*, FLC, now means that \mathfrak{T} is finite, that there is an upper bound on the number of tiles which come together at any vertex, and that there is a positive function θ so that K satisfies the θ -isoperimetric inequality.
- With \mathfrak{T} in the background, it is also understood that all maps between tilings or patches of tilings—embeddings, isomorphisms, local isomorphisms, and so forth—respect type. If $f : L \rightarrow K$ is a cellular isomorphism between tilings, then faces $\mathbf{k} \in L$ and $f(\mathbf{k}) \in K$ automatically have the same number of sides; but if the type of \mathbf{k} also involves, for example, a designated principal vertex, then f must respect this as well.
- Since maps respect types in \mathfrak{T} , the results of §2 apply with the understanding that all tilings have tiles with these types. Thus we can speak of the spaces $\mathbf{RC}_{\mathfrak{T}}$, $\mathbf{C}_{\mathfrak{T}}$, $\mathbf{RC}_{\mathfrak{T}}^{n,\theta}$, $\mathbf{LI}_{\mathfrak{T}}(K)$, and so forth. Since \mathfrak{T} may be assumed in a given setting, we will abuse notation and drop the subscript.
- If \mathcal{T}_K is the conformal tiling for K , then type designations from K carry over to \mathcal{T}_K as well.

It is immediate that, whenever $K \cong L$, we have $\tau K \cong \tau L$. This of course implies that τ induces a function $\hat{\tau} : \mathbf{C} \rightarrow \mathbf{C}$ in the obvious way, by sending the isomorphism class of K to the isomorphism class of τK . The hat symbol $\hat{}$ is useful here: τK refers to a specific combinatorial subdivision of K while $\hat{\tau} K$ refers to its isomorphism class in \mathbf{C} . The induced function $\hat{\tau}$ is also handy in another situation that will come up frequently later in the paper, namely, periodicity.

Definition. Suppose K is combinatorial tiling and τ is a subdivision rule applicable to K . If there is some $m \geq 1$ so that $\tau^m K \cong K$, then we say τ^m is a *combinatorial period map* for K and that K is τ -*periodic*. The least value of m for which $\tau^m K \cong K$ is called the τ -*period* of K .

We may rephrase this result by saying that the isomorphism class of K is a fixed point for $\hat{\tau}^m$. Clearly the collection of combinatorial period maps τ for K is closed under composition.

The connections between subdivision rules and combinatorial heirarchies now become our main focus, and we use the following terminology.

Definition. A subdivision rule τ is said to *manifest a combinatorial hierarchy* for the combinatorial tiling K if K exhibits a combinatorial hierarchy $\{K_n\}$ with $\tau K_n = K_{n+1}$ for all $n \in \mathbb{Z}$.

It is important to recall that this means τK_n is locally isomorphic to K_n for each n , a non-trivial requirement. Also note that by reindexing, the rule τ manifests a combinatorial hierarchy of every K_n . (In our generic notation, $\sigma_m = \tau$ for all $m \in \mathbb{Z}$, and $\sigma_m^n = \tau^{n-m}$ for all $m \leq n$.)

3.2. Expansion complexes. Our theory remains rather abstract until we can actually produce combinatorial hierarchies. From what we have said, the intuition is clear—build a hierarchy for K by repeated application of some subdivision rule τ . There are (at least) two hurdles: First, although the forward half of the sequence, $\{K_n = \tau^n K : n = 0, 1, 2, \dots\}$, seems straightforward, recall that we require that τK_n be locally isomorphic to K_n , which is not generally the case. Second, generating the backward half, $\{K_{-n} : n = 1, 2, \dots\}$, suggests a need for an *aggregation rule* τ^{-1} , which generally fails to exist.

We introduce expansion complexes to address these issues. The basic idea will be familiar to anyone in tiling theory: A subdivision rule τ defines a refinement process for tilings: starting with a single “initial” tile, one subdivides, then subdivides the result, and continues subdividing. Each stage is, of course, only a *finite* tiling, but diagonalization allows one to extract *infinite* tilings in the limit. By subdividing such a limit repeatedly with τ , one obtains *one-way* infinite sequences of infinite tilings. Expansion complexes formalize the diagonalization procedure typically used in these situations, though it avoids the global finite complexity hypotheses which are commonly required. Moreover, as the formulation is combinatorial, we can often extract *two-way* infinite sequences of infinite tilings, which are our main objects of study. We begin with a definition of expansion complexes applicable to quite arbitrary subdivision rules τ , though we will be more restrictive shortly.

In the following, we use a definition of “*interior*” for CW-complexes that comes from manifold theory rather than general topology. It will be used only for finite, 2-dimensional planar CW-complexes F and refers to the union of the open faces of F , the open edges that meet two closed faces of F , and the vertices of F that have an open neighborhood in $|F|$. We denote it by F° .

Given a subdivision rule τ , let Δ denote a fixed combinatorial tile belonging to the collection \mathfrak{T} of tile types for τ . Define the sequence of subdivisions $\{\Delta_j\}$, wherein $\Delta_0 = \Delta$ and $\Delta_{j+1} = \tau \Delta_j$, $j = 0, 1, \dots$. We study nested patches within this sequence from which limit tilings may be extracted: Let

$$(\ddagger) \quad F_1 \hookrightarrow F_2 \hookrightarrow \dots \hookrightarrow F_m \hookrightarrow F_{m+1} \hookrightarrow \dots$$

be any sequence of isomorphic embeddings of CW-complexes that satisfies the following properties:

- (1) for each positive integer m , F_m is a connected subcomplex of Δ_{i_m} for some positive integer i_m , with $i_m < i_{m+1}$;
- (2) each map $F_m \hookrightarrow F_{m+1}$ is a cellular, orientation-preserving, isomorphic embedding of CW-complexes which respects tile types;
- (3) for each m , there exists a positive integer p for which the image of F_m under the composition $F_m \hookrightarrow \dots \hookrightarrow F_{m+p}$ is contained in the interior of F_{m+p} ;

- (4) for each m and each combinatorial simple closed edge path γ in F_m , there exists a positive integer q for which the image of γ under the composition $F_m \hookrightarrow \cdots \hookrightarrow F_{m+q}$ bounds a combinatorial disc in F_{m+q} .

Treating $\{F_m\}$ as a nested sequence of complexes, it is not difficult to see that it has a directed limit $K = \varinjlim F_m$. Property (3) guarantees that the topological space $|K|$ is a non-compact topological 2-manifold. Property (4) guarantees that $|K|$ is simply connected. It follows that $|K|$ is homeomorphic to the plane so that K is a CW-decomposition of the plane. Since each subdivision $\tau^{i_m} \Delta$ is regular, so too is the CW-decomposition K . It follows that K is a combinatorial tiling. We establish the following definition (omitting Δ since its significance generally fades in the limit).

Definition. Let the sequence $\{F_m : m = 1, 2, \dots\}$ be as in (\dagger) with the properties (1)–(4) given above. The CW-complex $K = \varinjlim F_m$, the direct limit of the system (\dagger) , is a planar polygonal complex (a combinatorial tiling) called an *expansion complex* for the subdivision rule τ .

Of course the tiles of K have the types of \mathfrak{T} and we may repeatedly apply τ to get the one-way infinite sequence $\{K_n = \tau^n K : n \geq 0\}$ of refinements, as anticipated. For the full two-way infinite sequence, we look to aggregation in the next subsection. First, remarks.

Remarks.

- We make no claim that these properties for the sequence $\{F_m\}$ are independent: under mild restrictions on τ , for example, property (4) is a consequence of properties (2) and (3). Property (3) guarantees that K is a CW-decomposition of a surface *without* border, playing the role that “forcing the border” plays in traditional tiling theory. Likewise, removing property (4) would allow for holes, perhaps even an infinite number. Bordered and multiconnected complexes would be interesting in their own right, but we avoid them here.
- In typical constructions, each F_m in (\dagger) will be a CW-decomposition of a closed topological disc whose image under $F_m \hookrightarrow F_{m+1}$ will lie in the interior of F_{m+1} . Properties (3) and (4) automatically adhere in this case. The slightly more complicated property (3) makes for easier proofs in what follows.
- On the other extreme, one can easily find subdivision rules having no expansion complex: a rule that breaks each n -gon into $n+1$ -gons, for example, will not allow the embeddings needed in (\dagger) .
- The initial tile Δ cannot be ignored in this definition, though it remains in the background. In Example 6.3 of §6 we will see how this subtlety might be overlooked.
- Lastly, we observe that expansion complexes as defined by Cannon, Floyd, and Parry in [10] fit within our definition, but are more restrictive; as with the pentagonal tiling of [5], their expansion complexes are associated with *expansion maps*, which fail to exist generically.

3.3. Aggregation. The reverse of subdivision is what we term “aggregation”, and the notion is precisely what one would expect.

Definition. A τ -aggregate of a combinatorial tiling K is a combinatorial tiling L for which $\tau L \cong K$.

One can easily find complexes K to which τ applies that have no τ -aggregates, and as will be seen in the cautionary Example 4.4, if a τ -aggregate exists it is not necessarily unique. Thus, in deploying the machinery for constructing τ -aggregates, we will place additional conditions on τ as circumstances demand. For the key theorem proved next and later uses, we modify the notion of *fat* paths used in §2. The *skinny path distance* $\delta_K(\mathbf{k}, \mathbf{h})$ is the smallest integer m such that there exists a chain $\{\mathbf{k} = \mathbf{k}_0, \mathbf{k}_1, \dots, \mathbf{k}_m = \mathbf{h}\}$ of faces of K with $\mathbf{k}_j \cap \mathbf{k}_{j+1} \neq \emptyset$, $j = 0, \dots, m-1$. Successive faces need not share edges, so one can step from face to face across vertices as well; in particular, the skinny path distance between faces never exceeds the fat path distance. We assume that τ is δ -bounded, meaning that there is a positive constant ℓ , called a δ -bound for τ , such that for any tile type in $\mathbf{t} \in \mathfrak{T}$, the skinny path diameter of its subdivision $\tau(\mathbf{t})$ is less than ℓ . Note, for instance, that all finite subdivision rules and all simple subdivision operators are δ -bounded.

Theorem 3.1. Any expansion complex $K = \varinjlim F_m$ for a δ -bounded subdivision rule τ has a τ -aggregate L which is itself an expansion complex for τ .

Proof. For each $m > 1$, let G_m be the smallest subcomplex of $\tau^{i_m-1}\Delta$ whose τ -subdivision τG_m , a subcomplex of $\tau^{i_m}\Delta$, contains F_m . Ideally, the complexes G_m would admit appropriate embeddings, $G_m \hookrightarrow G_{m+1}$, so that $\tau G_m \hookrightarrow \tau G_{m+1}$ extends the embedding $F_m \hookrightarrow F_{m+1}$. The directed limit $L = \varinjlim G_m$ would be an expansion complex for τ with $\tau L \cong K$. Unfortunately, in the face of non-uniqueness of aggregation, this direct approach may fail. There is, however, a finite list of candidates for each G_m , so we may use diagonalization. It may help the reader to recall that there is a sequence $\{i_m\}$ of indices given in property (1) for the complexes F_m ; our diagonalization involves choosing a subsequence $\{j_m\}$ of these indices.

First, by identifying each F_m with its canonical embedded copy in the direct limit K , we may write K as the increasing union $\bigcup_{m=1}^{\infty} F_m$. By passing to a subsequence if necessary, property (3) allows us to assume, for each positive integer m , that F_m is contained in the interior of F_{m+1} . In fact the use of property (3) allows us to assume, without loss of generality, that the skinny path distance from any face of F_m to any face in the complement of F_{m+1} is greater than $\ell + 1$, where ℓ is a δ -bound for τ .

Recall the embedding sequence (\ddagger) : for integers m and p , $1 < m < p$, we have the embedding $F_m \hookrightarrow F_{m+1} \hookrightarrow \Delta_{i_p} = \tau^{i_p}\Delta$. Thus there is a smallest connected complex $H_{m,p}$ of $\tau^{i_p-1}\Delta$ so that $F_m \hookrightarrow \tau H_{m,p} \hookrightarrow \Delta_{i_p}$. Since the skinny path distance from any face of F_m to the complement of F_{m+1} is greater than $\ell + 1$, we necessarily have $F_m \hookrightarrow \tau H_{m,p} \hookrightarrow F_{m+1}$. Each $H_{m,p}$ might play the role of G_m , but the $H_{m,p}$ need not be isomorphic for different p , and even if they were, the pairs $(\tau H_{m,p}, F_m)$ need not be isomorphic as pairs. Nonetheless, since for each $p > m$ the complex $\tau H_{m,p}$ is a subcomplex of the finite complex F_{m+1} , we may extract a subsequence p_j of indices for which the pairs $(\tau H_{m,p_j}, F_m)$ will be isomorphic. This is the basis for our diagonalization.

Applying the sequence extraction of the previous paragraph with $m = 1$, there is a subsequence we label $\{p_j^1\}$ so that the pairs $(\tau H_{1,p_j^1}, F_1)$ are isomorphic. Continue inductively: given sequence $\{p_j^m\}$ there's a subsequence $\{p_j^{m+1}\}$ of indices greater

than $m+1$ so that the pairs $(\tau H_{m+1, p_j^{m+1}}, F_{m+1})$ are all isomorphic. Having carried this out for each positive m , we arrive at the diagonal sequence $\{n_j\}$, where $n_j = p_j^j$, which we may assume is increasing.

For each $m > 0$, the pairs $(\tau H_{m, n_j}, F_m)$ are isomorphic for all indices $n_j > m$ so we define $G_m = H_{m, n_m}$ to represent this shared complex. In particular, G_m is isomorphic to a subcomplex of τ^{n_m-1} , so we also define the index $j_m = n_m - 1$. Observe that for an index $j > n_{m+1}$ both G_m and G_{m+1} are isomorphic to subcomplexes of $\tau^{i_j-1}\Delta$ with

$$(1) \quad F_m \hookrightarrow \tau G_m \hookrightarrow F_{m+1} \hookrightarrow \tau G_{m+1} \hookrightarrow F_{m+2}.$$

Since τG_{m+1} contains F_m , our definition of H_{m, i_j} tells us that $H_{m, i_j} \hookrightarrow G_{m+1}$; that is, $G_m \hookrightarrow G_{m+1}$.

As a result of this process we arrive at the sequence

$$(\ddagger) \quad G_1 \hookrightarrow G_2 \hookrightarrow \cdots \hookrightarrow G_m \hookrightarrow G_{m+1} \hookrightarrow \cdots.$$

The complexes G_m along with the indices j_m satisfy the properties (1)–(4), so $L = \bigcup_{m=1}^{\infty} G_m$ is an expansion complex for τ . Moreover, since $K = \bigcup_{m=1}^{\infty} F_m$, (1) implies that $\tau L = K$, and L is seen as a τ -aggregate of K . This completes the proof. \square

Using this theorem we can obtain the backward half $\{K_{-n} : n \geq 1\}$ of our nested sequences: Given an expansion complex K for τ , let $K_0 = K$ and inductively define K_{-n} to be a τ -aggregate of K_{-n+1} , $n = 1, 2, \dots$. The bi-infinite sequence $\{K_n\}$ now satisfies conditions (1) and (2) required of a combinatorial hierarchy for τ .

3.4. Primitive rules. Condition (3) in the definition of combinatorial hierarchy involves, in our situation, the local isometry condition, $\tau K_n \sim K_n$. We can only expect this with restrictions on τ which are common as well in traditional substitution tilings.

Definition. A subdivision rule τ (necessarily finite) is *strictly primitive* if for every tile type $\mathbf{t} \in \mathfrak{T}$, the subdivision $\tau \mathbf{t}$ contains subtiles representing all tile types in \mathfrak{T} . The rule τ is *primitive* if there is some positive integer s so that τ^s is strictly primitive.

As an example, look to the “Penrose” subdivision rule τ of [7, Fig. 17]. Though τ^2 is strictly primitive, τ itself is not.

Theorem 3.2. *Suppose τ is a primitive finite subdivision rule and that K is an expansion complex for τ .*

- (a) *K is combinatorially repetitive.*
- (b) *The local isomorphism class (K) is precisely the set of all isomorphism classes of expansion complexes for τ .*
- (c) *For all positive integers k , $\tau^k K$ is locally isomorphic to K .*

Proof. We first set notation. Suppose $K = \varinjlim F_m$ is an expansion complex for τ with initial polygon Δ of type $\mathbf{t} \in \mathfrak{T}$. We may identify each F_m with its canonical copy in K and write $K = \bigcup_{m=1}^{\infty} F_m$. Recall that for each F_m we have an integer i_m so that $F_m \hookrightarrow \tau^{i_m} \Delta$. Fix $s > 0$ so that τ^s is strictly primitive and let ℓ be a δ -bound for τ

(a) Let H be a finite connected subcomplex of K . Then for some positive integer m , $H \hookrightarrow F_m$, hence $H \hookrightarrow \tau^{i_m} \Delta$. Apply Theorem 3.1 to obtain a τ^{i_m+s} -aggregate tiling K' for K ; that is, so $\tau^{i_m+s} K' \cong K$. Let \mathbf{k} be any face of K and let \mathbf{k}' be the face of K' that contains \mathbf{k} . Note that the skinny path diameter of the complex $\tau^{i_m+s} \mathbf{k}'$, which contains \mathbf{k} as a face, is at most ℓ^{i_m+s} . Since τ^s is strictly primitive, $\tau^s \mathbf{k}'$ contains a tile \mathbf{g} of type \mathbf{t} . In particular, the subcomplex $\tau^{i_m} \mathbf{g}$ of $\tau^{i_m+s} \mathbf{k}'$ contains an isomorphic copy of H . It follows that the face \mathbf{k} of K is ℓ^{i_m+s} -close to an isomorphic copy of H in the skinny path metric δ_K , and this implies that each vertex of \mathbf{k} is contained in an edge path of K of length at most ℓ^{i_m+s} that also contains a vertex of an isomorphic copy of H . We conclude that H is quasi-dense in K , and thus that K to be combinatorially repetitive.

(b) For one direction, suppose $L = \lim_{\rightarrow} G_m$ is another expansion complex for τ with initial polygon Δ' of type $\mathbf{t}' \in \mathfrak{T}$. It suffices to verify that, for each positive integer n , the subdivided polygon $\tau^n \Delta'$ embeds isomorphically in K . From Theorem 3.1 we can obtain a tiling K' such that $\tau^{n+s} K' \cong K$. Since τ^s is strictly primitive, $\tau^s K'$ contains a tile \mathbf{k} of type \mathbf{t}' . Note that $\tau^n \mathbf{k} \subset \tau^{n+s} K'$, so $\tau^n \mathbf{k}$ embeds isomorphically in K . But $\tau^n \mathbf{k} \cong \tau^n \Delta'$, so we are done. Interchanging the roles of K and L , we conclude that any two expansion complexes for τ are locally isomorphic.

For the other direction, we need to show that if a combinatorial tiling L is locally isomorphic to the expansion complex K , then L is itself an expansion complex for τ . Let C_m for $m \geq 1$ be a sequence of pairwise disjoint simple closed edge paths in L such that C_{m+1} separates C_m from infinity and let D_m be the combinatorial disc bounded by C_m . Then D_m is contained in the interior of D_{m+1} . Since K is locally isomorphic to L , D_m is isomorphic to a subcomplex D'_m of K and there exists an index $k(m)$ such that D'_m is contained in $F_{k(m)}$. We may assume by choosing $k(m)$ sequentially that $k(m) < k(m+1)$, implying that $j_m = i_{k(m)} < i_{k(m+1)} = j_{m+1}$. We then have the isomorphic embedding $D_m \cong D'_m \hookrightarrow F_{k(m)} \hookrightarrow \tau^{j_m} \Delta$, and properties (1) through (4) in the definition of expansion complex are satisfied. We conclude that $L = \bigcup_{m=1}^{\infty} D_m \cong \lim_{\rightarrow} D'_m$ is an expansion complex for τ .

(c) Observe that τK is an expansion complex for the subdivision operator τ . Indeed, the sequence of embeddings defining an expansion complex for τK is merely $\tau F_m \hookrightarrow \tau F_{m+1}$, where τF_m is a subcomplex of $\tau \tau^{i_m} \Delta = \tau^{i_m+1} \Delta$. This and induction then imply that, for all positive integers k , $\tau^k K$ is also an expansion complex for τ , and an application of (b) finishes the proof. \square

3.5. Generating hierarchies. We finally find ourselves in position to generate a plethora of combinatorial hierarchies: indeed, for certain subdivision rules τ , every expansion complex exhibits a combinatorial hierarchy.

Theorem 3.3. *Every expansion complex K associated with a primitive finite subdivision rule τ exhibits a combinatorial hierarchy manifest by τ .*

Proof. Let $K = K_0$ be an expansion complex for a strictly primitive finite subdivision rule τ . For positive integers k , let $K_k = \tau^k K$ and define the combinatorial tiling K_{-k} inductively using Theorem 3.1 so that $\tau K_{-k} = K_{-k+1}$ for all k . Theorem 3.1 guarantees that each K_{-k} is an expansion complex associated to τ , while Theorem 3.2(c) implies that K_k is locally isomorphic to K_{k+1} for all integers k . It follows that $\{K_k\}$ is a combinatorial hierarchy manifested by τ . \square

Note that the converse of this theorem fails: a member K_n of a combinatorial hierarchy for τ need not be an expansion complex for τ . See the discussion in Example 6.3 for an illustration.

4. COMBINATORIAL HYPOTHESES

It turns out that combinatorial hierarchy alone is not enough to capture the sense of *hierarchical tiling* as used in the traditional tiling community. Other properties of tilings, such as repetitiveness and finiteness conditions, and properties of subdivision rules, such as expansiveness, must be brought to bear. In this section we introduce the needed combinatorial machinery. We start with the notion of engulfing.

4.1. Cores and engulfing. Suppose $\{K_n\}$ is a generic combinatorial hierarchy for $K = K_0$ with subdivision sequence $\{\sigma_n\}$. When F is a finite subcomplex of K and $m \geq 0$, we say that the subcomplex E of K_{-m} *engulfs* F if F is a subcomplex of $\sigma_{-m}^0 E$, itself a subcomplex of $K_0 = K$; that is, after E is subdivided m times, the resulting tiles of K_0 contain F . A *core* of K is a combinatorial patch that comes in one of three flavors—its faces consist either of a single face of K , known as a *face core*, two faces that meet along a common edge, an *edge core*, or all the faces that meet at a vertex, a *vertex core*.

Definition. The combinatorial hierarchy $\{K_n\}$ for K is *expansive* if, for every finite subcomplex F of K , there exists an integer $m \geq 0$ such that F is engulfed by a core of K_{-m} , and hence by a core of K_{-n} for any $n \geq m$.

This formalizes the idea that F may be thought of as a subset of a core once we have aggregated a sufficient number of times, but we will need a more quantitative statement. Recall that we have defined two combinatorial distances between faces \mathbf{k} and \mathbf{h} of K , the fat path distance using chains of faces sharing edges and the skinny path distance that allows faces to share just vertices. The latter, $\delta_K(\mathbf{k}, \mathbf{h})$ is more useful here; it cannot exceed the fat path distance, and, in fact, the two are comparable if K has finite degree.

If F is a finite subcomplex of K , the *skinny path diameter* of F is the largest skinny path distance between faces of K that meet F . Notice that if L is a subdivision of K and \mathbf{k}', \mathbf{h}' are faces of L contained in faces \mathbf{k}, \mathbf{h} , respectively, of K then skinny path distances satisfy $\delta_L(\mathbf{k}', \mathbf{h}') \geq \delta_K(\mathbf{k}, \mathbf{h})$, and we write $\delta_L \geq \delta_K$. Symbolically, $L \leq K$ implies $\delta_L \geq \delta_K$. In this sense we may say that the sequence of skinny path metrics δ_{K_n} for the combinatorial hierarchy $\{K_n\}$ increases distances as $n \rightarrow \infty$ and decreases distances as $n \rightarrow -\infty$. Notice that this increase and decrease of distance among faces in a hierarchy is not necessarily strict, as Example 4.3 below illustrates.

Definition. The combinatorial hierarchy $\{K_n\}$ for K is *strongly expansive* if there exists a function $\phi : \mathbb{N} \rightarrow \mathbb{R}^+$, $\phi(n) \rightarrow \infty$ as $n \rightarrow \infty$, and a positive constant M such that, for any integer $m > M$, any finite subcomplex F of K of δ_K -diameter at most $\phi(m)$ is engulfed by a core of K_{-m} .

This of course says that the number of aggregations for a core to engulf a finite subcomplex F of K depends only on the diameter of F : aggregating m times is sufficient to engulf any finite subcomplex having δ_K -diameter at most $\phi(m)$. For instance, if $\phi(m) = \lambda m$ for some $\lambda > 0$, then the hierarchy is *linearly expansive*,

while if $\phi(m) = \lambda^m$ for some $\lambda > 1$, then it is *exponentially expansive*. The latter will play an important role in our later work. We will see, for example, that the hierarchies for Z and P in Examples 4.1 and 4.2 below are exponentially expansive.

Here are three examples. The first two are strongly expansive, based on criteria we develop later in the section, while the third is non-expansive.

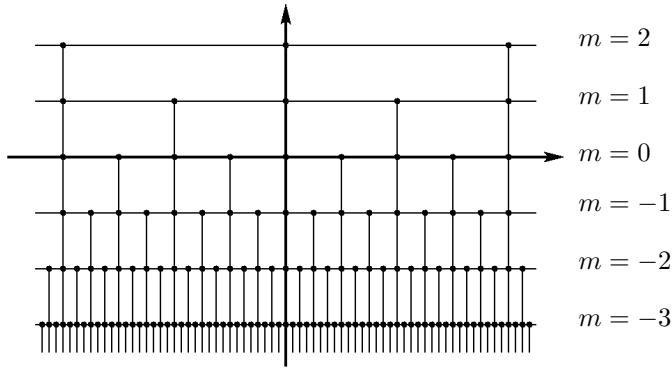
Example 4.1. The Gaussian complex. For an elementary, but nonetheless useful, example let Z denote the combinatorial tiling associated with the lattice of unit squares in \mathbb{C} having corners at the Gaussian integers \mathbb{Z}^2 . For each n let σ_n be the “quad” subdivision rule of Figure 2. Thus Z exhibits the combinatorial hierarchy $\{Z_n\}$ where each Z_n is isomorphic to Z itself. (In fact the conformal tiling \mathcal{Z} for Z is just the square tiling Z itself.) \checkmark

Example 4.2. The regular pentagonal complex. More interesting is the combinatorial *regular pentagonal* tiling P whose conformal tiling \mathcal{P} is the subject of [5]. If we let each σ_n be the “pentagonal” subdivision rule shown in Figure 1, then P exhibits a combinatorial hierarchy $\{P_n\}$. The key observation exploited in [5] was that $P_n \cong P_{n+1}$. However, (P) is uncountable, and we will see that every complex K that is locally isomorphic with P also exhibits a combinatorial hierarchy under the “pentagonal” subdivision rule, though generically, $K_n \not\cong K_{n+1}$. Examples of locally isomorphic conformal tilings for P are illustrated later in Figures 6 and 7. \checkmark

Example 4.3. A non-expansive combinatorial hierarchy. The combinatorial tiling H we describe here underlies the *Penrose hyperbolic tiling* familiar to the traditional tiling community, the *discrete hyperbolic plane* familiar to the geometric group theory and conformal geometry communities, and the Carleson (or Whitney) squares used in function theory. For example, the 1-skeleton of H occurs in the Cayley graph of a Baumslag-Solitar group. We described the subdivision rule τ and the resulting conformal tiling $\mathcal{H} = \mathcal{T}_H$ in [7, Example 3.11]; the reader can see half plane tilings there in Figure 21.

Our interest in H lies with its combinatorics, and we use the rather pedestrian schematic shown in Figure 4. The vertices of H lie along the horizontal lines in the complex plane \mathbb{C} with integer imaginary parts. At level $m \in \mathbb{Z}$, the vertices are $\{v_{m,k} = 2^m k + m\mathbf{i} : k \in \mathbb{Z}\}$ and the edges are the horizontal segments incident to $v_{m,k}$ and $v_{m,k+1}$ for all $k \in \mathbb{Z}$. The remaining edges are vertical segments incident to the vertices $v_{m,k}$ and $v_{m-1,2k}$ for all $m, k \in \mathbb{Z}$. The faces of H are pentagonal with cyclically ordered vertices $v_{m,k}, v_{m-1,2k}, v_{m-1,2k+1}, v_{m-1,2k+2}, v_{m,k+1}$ for the face $\mathbf{k}_{m,k}$, for $m, k \in \mathbb{Z}$. For any integer n , the n th tiling in the combinatorial hierarchy $\{H_n\}$ is just a copy of H translated vertically n units. Explicitly, $H_n = H + n\mathbf{i}$. Observe that the subdivision σ_0 subdivides each face $\mathbf{k}_{m,k}$ by bisecting each of its horizontal edges with a midpoint vertex, then adding a single vertical edge incident to the midpoint of the edge $\langle v_{m,k}, v_{m,k+1} \rangle$ and the vertex $v_{m-1,2k+1}$. This subdivides each of the pentagonal faces of H into two pentagonal faces and, easily, $\sigma_0 H = H + \mathbf{i}$.

Any core of H is sandwiched between horizontal lines at most distance 2 apart, and repeated subdivisions will not change this. In other words, if F is a finite subcomplex of H spanning, say, two lines distance 3 apart, then there is no core that can engulf F . In particular, the hierarchy $\{H_n\}$ is not expansive. It is useful to observe that while the Riemann surfaces S_Z and S_P for our two expansive examples above are parabolic, the Riemann surface S_H is hyperbolic. \checkmark

FIGURE 4. Geometric representation of the polygonal complex H .

4.2. Criterion for expansiveness. The hierarchies of interest to us are manifest by subdivision rules τ , such as those in the three examples above. We next present a criterion, simple to check, implying that any combinatorial hierarchy for τ must be strongly expansive. If v is a vertex of a combinatorial n -gon \mathbf{k} , then the set $\angle v = \{v\} \cup d^\circ \cup e^\circ \cup \mathbf{k}^\circ$, where d and e are the edges of \mathbf{k} incident with the vertex v , is called the *open angle* of \mathbf{k} at v . In this definition, \mathbf{k} is treated as a tile of a type to which τ applies, but only the fact that it is an n -gon is relevant.

Definition. We say that a subdivision rule τ is *strictly shrinking* if, for every n -gon \mathbf{k} to which it applies, every closed subtile of $\tau\mathbf{k}$ is contained in some open angle of \mathbf{k} . We say τ is *shrinking* if for some $n \geq 1$, the rule τ^n obtained from τ by iterating it n times is strictly shrinking.

Note in particular that if \mathbf{k} is a face of a combinatorial tiling K to which a strictly shrinking τ applies, and t is a subtile of $\tau\mathbf{k}$, then t cannot contain two vertices of \mathbf{k} nor can it meet more than two edges of \mathbf{k} non-trivially, and if it does meet two edges, those edges must be adjacent. If τ is (strictly) shrinking, observe that τ^m is (strictly) shrinking for all positive integers m . For examples, look to the subdivisions arising from the rules showcased in Figures 1 and 2. One can check that all but the star, delta, and hex rules are strictly shrinking. The hex rule is shrinking, as its second iterate is strictly shrinking.

Lemma 4.1. *Let τ be a shrinking subdivision rule with τ^s strictly shrinking for some integer $s \geq 1$. Then the skinny path metrics satisfy $\delta_{\tau^s K} \geq 2\delta_K - 1$ for any combinatorial tiling K to which τ applies. In particular, if $\{K_n\}$ is a combinatorial hierarchy for τ , then $\delta_{K_{n+s}} \geq 2\delta_{K_n} - 1$ for all integers n .*

The proof of this lemma, though slightly technical, is straightforward, so we leave it to the reader. With the help of this lemma we get the main result for this subsection.

Theorem 4.2. *Let τ be a subdivision rule that manifests a combinatorial hierarchy for the combinatorial tiling K . If τ is shrinking, then the hierarchy $\{K_n\}$ is exponentially expansive.*

Proof. By hypothesis, τ^s is strictly shrinking for some positive integer s . Note, however, that if one were to prove that the subhierarchy $\{K_{sn} : n \in \mathbb{Z}\}$ for τ^s is exponentially expansive, then one could conclude immediately that $\{K_n\}$ itself is exponentially expansive. It suffices, therefore, to assume that τ is strictly shrinking, i.e., $s = 1$.

Define $\phi : \mathbb{N} \rightarrow \mathbb{R}^+$ by $\phi(m) = 2^{m/2} = \alpha^m$, an increasing exponential function in the variable m with base $\alpha = \sqrt{2} > 1$. Let F be a finite subcomplex of K of δ_K -diameter at most α^m , for some positive integer m . We will verify that F is engulfed by a core of K_{-m} as long as $m > 4$. This will show that $\{K_n\}$ is ϕ -expansive and thus exponentially expansive.

First note that by the previous lemma with $s = 1$, we have the inequality $\delta_{K_{n+1}} \geq 2\delta_{K_n} - 1$ available. This implies that $\delta_K \geq 2^r \delta_{K_{-r}} - 2^r + 1$ for all positive integers r . Choose r so that $2^{r-1} < \alpha^m \leq 2^r$ and observe that this implies that $2(r-1) < m \leq 2r$. Fix a face \mathbf{k} of K that meets F . Since $\delta_K \geq 2^r \delta_{K_{-r}} - 2^r + 1$, if \mathbf{h} is another face of K that meets F with $\delta_K(\mathbf{k}, \mathbf{h}) > 1$, and \mathbf{a} and \mathbf{b} are faces of K_{-r} that contain \mathbf{k} and \mathbf{h} , respectively, then

$$\alpha^m \geq \delta_K(\mathbf{k}, \mathbf{h}) \geq 2^r \delta_{K_{-r}}(\mathbf{a}, \mathbf{b}) - 2^r + 1.$$

This implies, since $\alpha^m \leq 2^r$, that when $\delta_K(\mathbf{k}, \mathbf{h}) > 1$, then $\delta_{K_{-r}}(\mathbf{a}, \mathbf{b}) \leq 2 - 1/2^r < 2$, or $\delta_{K_{-r}}(\mathbf{a}, \mathbf{b}) \leq 1$. This says that after r aggregations, the subcomplex F is engulfed by the combinatorial patch $P_{K_{-r}}(\mathbf{a}, 1)$ of radius 1, the patch whose faces \mathbf{c} satisfy $\delta_{K_{-r}}(\mathbf{a}, \mathbf{c}) \leq 1$. Our claim is that after one more aggregation, the patch $P_{K_{-r}}(\mathbf{a}, 1)$, and therefore the subcomplex F , is engulfed by a core of $K_{-(r+1)}$. Since $2(r-1) < m$, it follows that

$$(\dagger) \quad (r+1) \leq 2(r-1) < m \quad \text{as long as} \quad r \geq 3.$$

If we set $m > 4$, then $\alpha^m = 2^{m/2} > 2^2$ implies that $2^{r-1} \geq 2^2$, so $r \geq 3$ and hence inequality (\dagger) holds. Since F is engulfed by a core of $K_{-(r+1)}$ and $m > (r+1)$, F is also engulfed by a core of K_{-m} . This implies that the hierarchy $\{K_n\}$ is exponentially expansive, and the proof is complete modulo the claim.

As for the claim, we need only show that when τ is strictly shrinking, the combinatorial patch $P_{\tau K}(\mathbf{a}, 1)$ of any face \mathbf{a} of τK is contained in a core of K . Let \mathbf{b} be the face of K that contains \mathbf{a} . Since τ is strictly shrinking, \mathbf{a} is contained in an open angle $\angle_{\mathbf{b}} v = \{v\} \cup d^\circ \cup e^\circ \cup \mathbf{b}^\circ$ determined by the vertex v of \mathbf{b} that is incident with the edges d and e of \mathbf{b} . Immediately then, the vertices of \mathbf{a} lie in this open angle and this implies that any face of τK that meets \mathbf{a} is contained in the vertex core $c(v)$ of K determined by v , and so the core $c(v)$ engulfs $P_{\tau K}(\mathbf{a}, 1)$. This completes the proof; however, we comment that $P_{\tau K}(\mathbf{a}, 1)$ is engulfed by the face core determined by \mathbf{b} if $\mathbf{a} \subset \mathbf{b}^\circ$ and that $P_{\tau K}(\mathbf{a}, 1)$ is engulfed by the edge core determined by d if \mathbf{a} does not meet the edge e . \square

4.3. Finiteness conditions. We gather various finiteness hypotheses and results needed later in the paper. Recall that a combinatorial tiling K has *bounded degree* if there exists a bound d both on vertex degree, the number of edges incident to a vertex, and on the number of sides of its polygonal tiles. There is an important subtlety here, for if K manifests a combinatorial hierarchy $\{K_n\}$, then the K_n , being locally isomorphic to K , must share this same bound. If τ is a finite subdivision rule based on the “star”, “delta”, or “barycentric” algorithms of Figures 1, 2, for example, then τ cannot manifest a combinatorial hierarchy of bounded degree

since each application of τ strictly increases some vertices degrees. In contrast, most of the finite subdivision rules we find of interest do not suffer this growth in vertex degree—look, for example, to rules based on “pentagonal”, “teepee”, “hexagonal”, or “snowball” algorithms. To distinguish these latter rules, we say that a subdivision rule τ has *bounded degree* if there is a positive integer β such that for each tile type $\mathbf{t} \in \mathfrak{T}$ and for every positive integer m , every vertex of $\tau^m \mathbf{t}$ is incident to at most β edge of $\tau^m \mathbf{t}$.

Lemma 4.3. *Suppose τ is a finite subdivision rule that manifests a combinatorial hierarchy for K . If τ is shrinking and K has bounded degree, then for any positive integer n there are at most finitely many isomorphically distinct filled n -neighborhoods $B_K(\mathbf{k}, n)$ in K .*

Proof. Fix $n \geq 1$ and fix any face $\mathbf{k} \in K$. The n -neighborhood $C_K(\mathbf{k}, n)$ has δ_K -diameter at most $2n$. By Theorem 4.2 the hierarchy is strongly expansive, so there is an associated expansion function $\phi : \mathbb{N} \rightarrow \mathbb{R}^+$. We may choose a positive integer m so large that $2n \leq \phi(m)$ and we then conclude that $C_K(\mathbf{k}, n)$ is engulfed by a core of K_{-m} . Assuming K has a bound d on its degree, then K_{-m} , being locally isomorphic to K , also enjoys this bound. Up to isomorphism, then, there are only finitely many distinct cores $c \in K_{-m}$, and since τ is finite, there are up to isomorphism only finitely many subdivided cores $\tau^m c \subset K$.

Suppose, then, that $C_K(\mathbf{k}, n)$ is engulfed after m subdivisions by a core $c \in K_{-m}$. If c were simply connected, its subdivision $\tau^m c$ would be simply connected as well, and since $B_K(\mathbf{k}, n)$ is obtained by filling in any holes in $C_K(\mathbf{k}, n)$, we would have the following containments:

$$C_K(\mathbf{k}, n) \subseteq B_K(\mathbf{k}, n) \subseteq \tau^m c.$$

However, a core c may fail to be simply connected, and an additional argument from Mayhook [18, Lem 5.8] comes into play. In this case, we may work with the augmented core c^+ in K_{-m} , which is defined as the union of c with its bounded complementary components. Since τ is shrinking, c^+ is engulfed by a core of $K_{-(m+m')}$ for some fixed $m' > 0$. In particular, up to isomorphism the number of augmented cores is finite. As these are simply connected, we may run our final counting argument with c^+ in place of c . \square

4.4. Local isomorphism. As noted at the beginning of this section, we need more than a *hierarchy* alone to capture hierarchical structure. The two other attributes we need, to wit, finiteness conditions and repetitiveness, will now be enforced to prove this theorem, the main result of the section. We find, in particular, that the existence of a combinatorial hierarchy is really a property not so much of a single tiling K , but of its whole local isomorphism class (K) .

Theorem 4.4. *Suppose τ is a finite subdivision rule that manifests a combinatorial hierarchy for K . If the hierarchy is strongly expansive and K has finite degree and is combinatorially repetitive, then τ manifests a combinatorial hierarchy for any combinatorial tiling L that is locally isomorphic to K .*

Proof. Let τ be a finite subdivision rule that manifests the combinatorial hierarchy $\{K_n\}$ for K . Since K is repetitive and of bounded degree, then any combinatorial tiling locally isomorphic to K has faces of the same tile types as K , is combinatorially repetitive, and has the same bounded degree d as K . This applies, in

particular, to all members of any combinatorial hierarchy manifest by τ for $L \sim K$. Also, finiteness of τ means we have only a finite number of tile types and the rule applied to each type gives only a finite number of subtiles, so there is an upper bound $\lambda < \infty$ on the number of faces into which τ subdivides any face to which it applies.

Let $L \sim K$ and define $L_0 = L$ and, for each positive integer n , $L_n = \tau^n L$ so that $L_n = \tau L_{n-1}$ when $n \geq 1$. Our first task is to define a τ -aggregate tiling L_{-1} , so $\tau L_{-1} = L_0$. We are starting with a hierarchy, not an expansion complex, so although Theorem 3.1 cannot be applied directly, the arguments are similar.

The data used to construct L_{-1} are the three complexes L_0 , K_0 , and K_{-1} , along with τ . Fix a face $\mathbf{k} \in L_0$ and write $L_0 = \bigcup_{\ell=1}^{\infty} B_{\ell}$, where $B_{\ell} = B_{L_0}(\mathbf{k}, \lambda(\ell+1))$ is the filled $\lambda(\ell+1)$ -neighborhood of \mathbf{k} . Using that $K_0 \sim L_0$ and Lemma 2.10, there are isomorphic embeddings $f_{\ell} : B_{\ell} \hookrightarrow K_0$ where the image of f_{ℓ} is $B_{K_0}(\mathbf{k}_{\ell}, \lambda(\ell+1))$, the filled $\lambda(\ell+1)$ -neighborhood of $\mathbf{k}_{\ell} = f_{\ell}(\mathbf{k})$ in K_0 . There is a unique face \mathbf{h}_{ℓ} of K_{-1} such that \mathbf{k}_{ℓ} is a face of the subdivision $\tau \mathbf{h}_{\ell}$. The definition of λ may be used to verify the first containment in the observation that

$$\tau C_{K_{-1}}(\mathbf{h}_{\ell}, \ell) \subset C_{K_0}(\mathbf{k}_{\ell}, \lambda(\ell+1)) \subset B_{K_0}(\mathbf{k}_{\ell}, \lambda(\ell+1)).$$

Since $B_{K_{-1}}(\mathbf{h}_{\ell}, \ell)$ is obtained by merely “filling in the holes” of $C_{K_{-1}}(\mathbf{h}_{\ell}, \ell)$, it follows from the containments above and the fact that $B_{K_0}(\mathbf{k}_{\ell}, \lambda(\ell+1))$ is simply connected, that $\tau B_{K_{-1}}(\mathbf{h}_{\ell}, \ell)$ is a subcomplex of $B_{K_0}(\mathbf{k}_{\ell}, \lambda(\ell+1))$.

We now claim that we may extract a subsequence $\{\mathbf{h}_{j_i}\}_i$ of the sequence of faces $\{\mathbf{h}_{\ell}\}_{\ell}$ with the following properties:

- (i) the finite complex $D_i = B_{K_{-1}}(\mathbf{h}_{j_i}, i)$ admits an isomorphic embedding e_i into the complex D_{i+1} with $e_i(\mathbf{h}_{j_i}) = \mathbf{h}_{j_{i+1}}$;
- (ii) the direct limit $L_{-1} = \varinjlim (D_i, e_i)$ is a combinatorial tiling that is locally isomorphic to L_0 ;
- (iii) $\tau L_{-1} \cong L_0$.

For item (i), we define the subsequence \mathbf{h}_{j_i} inductively as follows. By Lemma 4.3, there is a smallest subscript j_1 such that $D_1 = B_{K_{-1}}(\mathbf{h}_{j_1}, 1)$ is isomorphic to $B_{K_{-1}}(\mathbf{h}_{\ell}, 1)$ for infinitely many subscripts $\ell > j_1$. By Lemma 4.3 again, there is among these infinitely many subscripts a smallest subscript $j_2 > j_1$ such that $D_2 = B_{K_{-1}}(\mathbf{h}_{j_2}, 2)$ is isomorphic to $B_{K_{-1}}(\mathbf{h}_{\ell}, 2)$ for infinitely many subscripts $\ell > j_2$. Having chosen $\mathbf{h}_{j_1}, \dots, \mathbf{h}_{j_n}$ in this way so that $D_n = B_{K_{-1}}(\mathbf{h}_{j_n}, n)$ is isomorphic to $B_{K_{-1}}(\mathbf{h}_{\ell}, n)$ for infinitely many $\ell > j_n$, we again apply Lemma 4.3 to choose among these infinitely many subscripts the smallest subscript $j_{n+1} > j_n$ such that $D_{n+1} = B_{K_{-1}}(\mathbf{h}_{j_{n+1}}, n+1)$ is isomorphic to $B_{K_{-1}}(\mathbf{h}_{\ell}, n+1)$ for infinitely many subscripts $\ell > j_{n+1}$. This inductively defines the sequence \mathbf{h}_{j_i} for $i = 1, 2, \dots$, and for each i , the choice of j_{i+1} implies that there is an isomorphic embedding

$$e_i : D_i = B_{K_{-1}}(\mathbf{h}_{j_i}, i) \cong B_{K_{-1}}(\mathbf{h}_{j_{i+1}}, i) \subset B_{K_{-1}}(\mathbf{h}_{j_{i+1}}, i+1) = D_{i+1}$$

with necessarily $e_i(\mathbf{h}_{j_i}) = \mathbf{h}_{j_{i+1}}$. This confirms item (i).

Since the complexes D_i are filled i -neighborhoods and hence combinatorial discs with D_i contained in the interior of D_{i+1} , the direct limit complex L_{-1} is a CW-decomposition of the whole plane, and hence a combinatorial tiling. By construction, $L_{-1} \preceq K_{-1}$ and as $K_{-1} \sim K_0 \sim L_0$, we have $L_{-1} \preceq L_0$. We have yet to use

the hypothesis of combinatorial repetitiveness, but now it is invoked to prove that $L_0 \preceq L_{-1}$, implying that $L_{-1} \sim L_0$ and confirming item (ii). Let F be a connected subcomplex of L_0 . Since $L_0 \sim K_{-1}$, there exists an isomorphic copy F' of F in K_{-1} . Since K_{-1} is combinatorially repetitive, there exists an integer $k > 0$ such that every combinatorial k -neighborhood of any face in K_{-1} contains an embedded copy of F' , and therefore of F . In particular, $D_k = B_{K_{-1}}(\mathbf{h}_{j_k}, k)$ contains an embedded copy of F , and therefore so does L_{-1} . We conclude that $L_0 \preceq L_{-1}$, and this finishes the verification of item (ii).

For item (iii), we will use the fact derived above that for any positive integer ℓ , $\tau B_{K_{-1}}(\mathbf{h}_\ell, \ell)$ is a subcomplex of $B_{K_0}(\mathbf{k}_\ell, \lambda(\ell+1))$. Write $L_{-1} = \bigcup_{i=1}^{\infty} D'_i$, where D'_i is the canonical isomorphic copy of D_i in L_{-1} , and let \mathbf{h} be the face of the direct limit L_{-1} that corresponds to the faces \mathbf{h}_{j_i} of the factors. Then $\tau L_{-1} = \bigcup_{i=1}^{\infty} \tau D'_i$ so that $\tau D'_1 \subset \tau D'_2 \subset \tau D'_3 \subset \cdots$ is a sequence of finite subcomplexes of τL_{-1} that exhausts τL_{-1} , as in the hypothesis of the working Lemma 2.6. For each positive integer i , the mappings

$$\tau D'_i \cong \tau D_i = \tau B_{K_{-1}}(\mathbf{h}_{j_i}, i) \subset \tau B_{K_{-1}}(\mathbf{h}_{j_i}, j_i) \subset B_{K_0}(\mathbf{k}_{j_i}, \lambda(j_i+1)) \xrightarrow[f_{j_i}^{-1}]{\cong} B_{j_i} \subset L_0$$

define an isomorphic embedding $\tau D'_i \hookrightarrow L_0$. Notice that the image of one of the faces \mathbf{k}'_i of the subdivided face $\tau \mathbf{h} \subset \tau D'_i$ under this embedding is equal to the face \mathbf{k} of L_0 . Since there are only finitely many faces in the subdivided face $\tau \mathbf{h}$, by passing to a subsequence if necessary, we may assume without loss of generality that all the faces \mathbf{k}'_i are the same face \mathbf{k}' of τL_{-1} . An application of the working Lemma 2.6 now implies that $\tau L_{-1} \cong L_0$, and item (iii) is proved.

Having confirmed items (i)–(iii), we now may use the isomorphism $\tau L_{-1} \cong L_0$ to replace L_{-1} by an isomorphic copy that aggregates the faces of L_0 . That is, we may assume without loss of generality that $\tau L_{-1} = L_0$. Repeat the argument using the data L_{-1} , K_{-1} , and K_{-2} in place of L_0 , K_0 , and L_{-1} to construct a combinatorial tiling L_{-2} such that $\tau L_{-2} = L_{-1}$ with $L_{-2} \sim L_{-1}$. Iterating ad infinitum, this produces a sequence $\{L_n\}_{n < 0}$ such that $\tau L_n = L_{n+1}$ and $L_n \sim L_{n+1}$ for $n \leq -1$. Already we have defined L_n for $n \geq 0$ so that $\tau L_n = L_{n+1}$ and it follows from its definition that $L_n \sim K_n \sim K_{n+1} \sim L_{n+1}$, since $K \sim L$ implies that $\tau K \sim \tau L$.

All the ingredients are now in place as we have produced a bi-infinite sequence $\{L_n\}$, for $n \in \mathbb{Z}$, such that (1) $L_0 = L$, (2) $\tau L_n = L_{n+1}$, and (3) $L_n \sim L_{n+1}$, and this implies that the subdivision operator τ manifests a combinatorial hierarchy for L . \square

In the proof we have neither proved nor claimed that the backward sequence $\{L_n\}_{n < 0}$ is uniquely determined by $L = L_0$. There indeed may be differing ways to apply the aggregation operation τ^{-1} to L_n to produce a tiling L_{n-1} having $\tau L_{n-1} = L_n$ and $L_{n-1} \sim L_n$. This example illustrates.

Example 4.4. A cautionary example. Consider the lattice quad complex \mathbf{Z} and the quad subdivision rule τ of Example 4.1. Designate a fixed tile \mathbf{k}_0 of $\mathbf{Z} = \mathbf{Z}_0$. It is evident that we can aggregate the tiles of \mathbf{Z} to get $\mathbf{Z}_{-1} = \tau^{-1}\mathbf{Z}$ in four distinct ways: if \mathbf{k}' is the tile of \mathbf{Z}_{-1} that contains \mathbf{k}_0 , then \mathbf{k}_0 can occupy any one of the four quadrants of \mathbf{k}' . These aggregates are distinct tilings, but of course they are all isomorphic to \mathbf{Z} . $\sqrt{}$

5. CONFORMAL HIERARCHIES

We are approaching our main goal, meshing the combinatorics and geometry of hierarchical tilings. In this section we define conformal hierarchies and begin realizing geometric versions of the various combinatorial situations developed in the previous sections. We establish results on conformal type and conclude with the statement of our Main Theorem.

When K_s is a combinatorial subdivision of K we have written $K_s \leq K$. Suppose $\mathcal{T}_s = \mathcal{T}_{K_s}$ and $\mathcal{T} = \mathcal{T}_K$ are the associated conformal tilings. Recall that \mathcal{T}_s is a *conformal subdivision* of \mathcal{T} if \mathcal{T}_s and \mathcal{T} have compatible β -equilateral conformal structures. In other words, \mathcal{T}_s may be realized by subdividing the tiles of \mathcal{T} in situ. In this case we write $\mathcal{T}_s \leq \mathcal{T}$; we also say that \mathcal{T} is a *conformal aggregation* of \mathcal{T}_s .

Definition. The conformal tiling \mathcal{T} *exhibits a conformal hierarchy* if there is a bi-infinite sequence $\{\mathcal{T}_n : n \in \mathbb{Z}\}$, called a *conformal hierarchy* for \mathcal{T} , of conformal tilings indexed by the integers such that the following three conditions hold:

- (1) $\mathcal{T}_0 = \mathcal{T}$;
- (2) $\mathcal{T}_{n+1} \leq \mathcal{T}_n \ \forall n \in \mathbb{Z}$ (\mathcal{T}_{n+1} is a conformal subdivision of \mathcal{T}_n);
- (3) $\mathcal{T}_{n+1} \sim \mathcal{T}_n \ \forall n \in \mathbb{Z}$ ($K_{\mathcal{T}_{n+1}}$ is locally isomorphic to $K_{\mathcal{T}_n}$).

The tilings \mathcal{T}_n all lie in \mathbb{G} , one of \mathbb{C} or \mathbb{D} , and the hierarchy is said to be of *parabolic* or *hyperbolic* type, respectively.

The parallels to the combinatorial situation are self-evident. Indeed, lying behind any conformal hierarchy $\{\mathcal{T}_n\}$ is a combinatorial hierarchy, which by convention will be denoted $\{K_n = K_{\mathcal{T}_n}\}$. We say that $\{\mathcal{T}_n\}$ is a *conformal realization* of $\{K_n\}$, in the same sense that an individual conformal tiling \mathcal{T}_K is a conformal realization of K . As there, the \mathcal{T}_n are uniquely determined in \mathbb{G} up to a Möbius transformation (which is independent of n , since we have arranged for the \mathcal{T}_n to be nested).

What is more noteworthy are the *geometric* parallels with traditional euclidean tilings based on the fact that conformal subdivisions occur in situ—a concrete geometric tile is a union of concrete geometric tiles of its subdivision. We first establish some terminology and notation associating combinatorial and geometric notions. A conformal hierarchy $\{\mathcal{T}_n\}$ is said to be (*strongly*) *expansive* if its combinatorial hierarchy $\{K_n\}$ is (strongly) expansive. If F is a subcomplex of a combinatorial tiling K , then the associated point set in the geometry \mathbb{G} of the conformal tiling \mathcal{T}_K will be denoted $|F|$. If F is a core c of K , then we call $|c|$ a *core* of \mathcal{T} . To say $|c|$ *engulfs* a compact set $D \subset \mathbb{G}$ simply means that $D \subset |c|$.

Lemma 5.1. *If $\{\mathcal{T}_n\}$ is an expansive conformal hierarchy in \mathbb{G} , then every compact set $D \subset \mathbb{G}$ is engulfed by a core of the tiling \mathcal{T}_{-m} for some $m \geq 0$.*

Proof. Since $\mathcal{T} = \mathcal{T}_0$ is a locally finite tiling of \mathbb{G} and D is compact, only finitely many tiles of \mathcal{T} meet D and there is a finite subcomplex F of $K = K_0$ such that $D \subset |F|$. Since the hierarchy is expansive, there is an integer $m \geq 0$ such that F is engulfed by a core of K_{-m} . If we write σ for the subdivision process σ_m^0 (so $\mathcal{T} = \sigma\mathcal{T}_{-m}$), then engulfing means that there is some core c of K_{-m} such that F is a subset of $\sigma c \subset K$. In particular, $|F| \subset |\sigma c|$ in \mathbb{G} . By item (2) in the definition of conformal hierarchy, \mathcal{T} is a conformal subdivision of \mathcal{T}_{-m} ; here we are using the fact that conformal subdivision of tilings is a transitive relation. Thus \mathcal{T}_{-m} is an in situ polygonal subdivision of \mathcal{T} , implying $|\sigma c| = |c|$. Thus, $D \subset |F| \subset |\sigma c| = |c|$, and D is engulfed by $|c|$. \square

5.1. Conformal type. Throughout this subsection we will be considering combinatorial tilings of bounded degree—that is, there’s a bound on the number of edges in any face and the number of faces meeting at any vertex. We establish conditions on conformal hierarchies ensuring that they will be parabolic, indeed, perhaps even constantly parabolic across their local isomorphism classes.

We begin with a review of our main tool, a classical criterion for parabolicity using ring domains. The reader is referred to the classic references Lehto and Virtanen [17] and Ahlfors [1] or the modern treatment given in Fletcher and Markovic [12] for details. A *ring domain* C is a doubly connected open set in the Riemann sphere \mathbb{S}^2 , i.e., a domain of \mathbb{S}^2 whose complement consists of two non-empty connected components. A classical theorem of Koebe implies that C is conformally equivalent to a *circle domain*, a domain for which each complementary component is either a closed circular disc or a point. It follows that any ring domain C is conformally equivalent to a *round annulus*, one of the form

$$A(r, R) = \{z \in \mathbb{C} : 0 \leq r < |z| < R \leq \infty\}.$$

The *modulus* of $A = A(r, R)$ is defined as $\text{Mod}(A) = \log(R/r)$, with the obvious interpretations that $\text{Mod}(A) = \infty$ if $r = 0$ and/or $R = \infty$. Two round annuli A and B with finite moduli are conformally equivalent if and only if $\text{Mod}(A) = \text{Mod}(B)$. The annuli with infinite moduli determine two conformal equivalence classes according to whether one or both of the complementary domains are points. For any ring domain C , then, we may define its *modulus* by $\text{Mod}(C) = \text{Mod}(A)$, where A is any round annulus conformally equivalent to C .

With these facts under our belts, we can give a tool for determining the type of a non-compact, simply connected Riemann surface S . Let D be a closed disc in S . By the Uniformization Theorem, the set $S \setminus D$ is conformally equivalent to a ring domain C in \mathbb{C} . Obviously, S is parabolic if and only if $\text{Mod}(C) = \infty$ and is hyperbolic if and only if $\text{Mod}(C) < \infty$. The theory of quasiconformal mapping provides the following characterization of parabolic surfaces.

Proposition 5.2 (Criterion for Parabolicity). *A non-compact, simply connected Riemann surface S is parabolic if and only if there is a constant $\mu > 0$ such that, for every compact subset E of S , there is a ring domain C separating E from the ideal boundary of S such that $\text{Mod}(C) \geq \mu$.*

This is proved by applying the superadditivity of the modulus, which avers that if C_1, C_2, \dots is a sequence of pairwise disjoint ring domains that are subdomains of the ring domain C , with every C_n separating the boundary components of C from one another, then

$$\sum_{n=1}^{\infty} \text{Mod}(C_n) \leq \text{Mod}(C).$$

Theorem 5.3. *If \mathcal{T} is a conformal tiling of bounded degree that exhibits an expansive conformal hierarchy, then \mathcal{T} is parabolic.*

Proof. Suppose $\{\mathcal{T}_n\}$ is an expansive conformal hierarchy for \mathcal{T} . All these tilings necessarily live in the same geometry \mathbb{G} .

We focus our attention first on the tiling $\mathcal{T} = \mathcal{T}_0$. For a core c of $K = K_{\mathcal{T}}$, we denote by $B(c)$ the subcomplex of K consisting of all faces in the unbounded complementary domain of c that meet c . The complex $B(c)$ will be called the *combinatorial collar* of c , and the corresponding patch of tiles in \mathcal{T} determines an open domain

$U(c)$ in \mathbb{G} , the interior of $|B(c)|$, that separates the core $|c|$ from infinity. Since $B(c)$ is a finite subcomplex of K , the domain $U(c)$ is finitely connected and by the Koebe Uniformization Theorem, there is a conformal isomorphism $f : U(c) \rightarrow V(c)$, where $V(c)$ is a circle domain, the complement of a finite number of closed round discs in the plane. By applying an appropriate Möbius transformation, we may assume two things: first, that $V(c)$ is the complement of a finite number of closed round discs inside the unit disc \mathbb{D} ; second, that the disc corresponding to the complementary domain of $U(c)$ that contains $|c|$, call it $D(c)$, is centered at the origin. There is a largest round annulus $A(c) \subset V(c)$ that shares one boundary with $D(c)$ and separates $D(c)$ from the rest of the boundary of $V(c)$. Let $R(c) = f^{-1}(A(c))$ be its image under the inverse isomorphism f^{-1} . Then $R(c) \subset U(c)$ is a ring domain that separates the core $|c|$ from infinity and is called the *standard collar* of the core $|c|$. As c ranges over all cores of K (hence all cores of any complex locally isomorphic to K) there are only finitely many different conformal isomorphism types of standard collars $R(c)$. This follows from the fact that \mathcal{T} has bounded degree and from Theorem 1.3. Indeed, there exists only finitely many combinatorial types of cores of K , and for each one of these combinatorial types of cores, only finitely many combinatorial types of collars. List combinatorial collars $B_1 = B(c_1), \dots, B_N = B(c_N)$ so that the combinatorial collar $B(c)$ of any core c of K is combinatorially equivalent to one from the list. By Theorem 1.3, the domain $U(c)$ of every core $|c|$ of \mathcal{T} is conformally equivalent to one of the domains $U(c_1), \dots, U(c_N)$. But this implies that the standard collar $R(c)$ is conformally equivalent to one of the standard collars from the list $R_1 = R(c_1), \dots, R_N = R(c_N)$.

Suppose $\mu > 0$ is the smallest conformal modulus of any ring domain in the list R_1, \dots, R_N . Let D be a compact subset of \mathbb{G} . By Lemma 5.1, there exists an integer $m \leq 0$ such that D is engulfed by a core $|c|$ of the tiling \mathcal{T}_m . It follows that the ring domain $R(c)$ separates D from infinity. By the observation of the preceding paragraph, $\text{Mod}(R(c)) \geq \mu$, and this verifies the Criterion for Parabolicity and completes the proof that $\mathbb{G} = \mathbb{C}$. \square

While this result guarantees that \mathcal{T} and the other tilings \mathcal{T}_n of its hierarchy are parabolic, it says nothing about the conformal type of their locally isomorphic cousins. Constancy of type across the whole local isomorphism class may be achieved by strengthening the expansive hypothesis.

Theorem 5.4. *Let $\mathcal{T} = \mathcal{T}_K$ be a conformal tiling of bounded degree that exhibits a strongly expansive conformal hierarchy; then \mathcal{T}_L is parabolic for any combinatorial tiling L which is locally isomorphic to K .*

Proof. Theorem 5.3 implies that \mathcal{T} is parabolic. Suppose \mathcal{T}_L is a conformal tiling of \mathbb{G} , either \mathbb{C} or \mathbb{D} , whose complex L is locally isomorphic but not isomorphic to $K = K_{\mathcal{T}}$. We will verify the Criterion for Parabolicity Proposition 5.2 for \mathbb{G} using the standard collars from the proof of Theorem 5.3.

Let D be a compact subset of \mathbb{G} and let F be a finite, connected subcomplex of L whose corresponding patch of conformal tiles within \mathcal{T}_L covers D , so that $D \subset |F|$. Since $L \sim K$ while L is not isomorphic to K , (K) is not a singleton, so Theorem 2.8 implies that the finite complex F is infinitely represented in K . Let \mathbf{k} be a face of F and let $(F_1, \mathbf{k}_1), (F_2, \mathbf{k}_2), \dots$ be a pairwise unequal listing of all the pairs where F_i is a finite subcomplex of K and \mathbf{k}_i is a face of F_i such that (F_i, \mathbf{k}_i) and (F, \mathbf{k}) are isomorphic as pairs.

Let $\{\mathcal{T}_n\}$ be a strongly expansive conformal hierarchy for $\mathcal{T} = \mathcal{T}_0$, so $\{K_n\}$ is a strongly expansive combinatorial hierarchy for $K = K_0$. Suppose $\phi : \mathbb{N} \rightarrow \mathbb{R}^+$ is the associated expansion function. Since F is finite and connected, there exists a positive integer m such that both the δ_L -diameter of F and the δ_K -diameter of F_i , for all positive integers i , are less than $\phi(m)$. Let σ denote the subdivision process σ_m^0 , so that $\sigma K_{-m} = K$. In particular, each of the subcomplexes F_i of K is engulfed by some core c_i of K_{-m} . It is precisely here that *strong* expansivity is used, so the m is independent of i . Let $B(c_i)$ be the combinatorial collar of c_i in K_{-m} as constructed in the proof of Theorem 5.3 and recall from that proof that, since $K_{-m} \sim K$, each collar $B(c_i)$ is combinatorially equivalent to one of a finite list, B_1, \dots, B_N , of such collars. Moreover, from this, as the proof of Theorem 5.3 avers, each standard collar $R(c_i)$ is conformally equivalent to a standard collar from the finite list, R_1, \dots, R_N , determined by the combinatorial collars B_1, \dots, B_N .

Since c_i engulfs F_i , \mathbf{k}_i is a face of the subdivided core σc_i . Choose a positive integer M so large that $\sigma(c_i \cup B(c_i))$ is a subcomplex of the combinatorial M -neighborhood $C_K(\mathbf{k}_i, M)$, and F is a subcomplex of $C_L(\mathbf{k}, M)$. Such an M exists since \mathcal{T} has bounded degree, there are only finitely many combinatorial types of cores in the list c_1, c_2, \dots and of collars in the list $B(c_1), B(c_2), \dots$, and there is an upper bound on the number of faces in the subdivision $\sigma \mathbf{h}$ of any face \mathbf{h} of K_{-m} . Consider the combinatorial M -neighborhood $C_L(\mathbf{k}, M)$ in the complex L . Since $L \sim K$, there is a subcomplex H of K that is isomorphic to $C_L(\mathbf{k}, M)$ via an isomorphism $\lambda : C_L(\mathbf{k}, M) \rightarrow H$. By Lemma 2.10, $H = C_K(\lambda(\mathbf{k}), M)$. Since F is a subcomplex of $C_L(\mathbf{k}, M)$, there exists a positive integer i such that $(F_i, \mathbf{k}_i) = (\lambda(F), \lambda(\mathbf{k}))$, a subcomplex of H . It follows that

$$F_i \subset \sigma c_i \subset \sigma(c_i \cup B(c_i)) \subset C_K(\mathbf{k}_i, M) = C_K(\lambda(\mathbf{k}), M) = H.$$

The first containment is just the statement that the core c_i of K_{-m} engulfs F_i , the second is trivial, and the third is by choice of M . By applying the inverse isomorphism λ^{-1} , there is a subcomplex J of $C_L(\mathbf{k}, M)$ isomorphic to $\sigma B(c_i)$ that contains F in one of its bounded complementary domains of L , implying that J separates F from infinity. By Theorem 1.3, \mathcal{T}_J is conformally equivalent to $\mathcal{T}_{\sigma B(c_i)}$. It follows that open domain $U(J)$, the interior of $|\mathcal{T}_J|$ in \mathbb{G} , is conformally equivalent to the open domain $U(\sigma B(c_i))$, the interior of $|\mathcal{T}_{\sigma B(c_i)}|$ in \mathbb{C} . Since the hierarchy $\{\mathcal{T}_n\}$ is conformal, $\sigma \mathcal{T}_{-m} = \mathcal{T}$ is a conformal subdivision of \mathcal{T}_{-m} , and therefore is a polygonal subdivision. It follows that the open domain $U(c_i)$, the interior of $|\mathcal{T}_{B(c_i)}|$ in \mathbb{C} , is precisely *equal* to the open domain $U(\sigma B(c_i))$. This implies that the standard collar $R(c_i)$ is contained in the open domain $U(\sigma B(c_i))$ and separates $|\mathcal{T}_{F_i}|$ from infinity in \mathbb{G} . From this, we conclude that $R = \lambda^{-1}(R(c_i))$ is a ring domain in \mathbb{G} that is conformally equivalent to $R(c_i)$ and separates $D \subset |\mathcal{T}_F|$ from infinity. But $R(c_i)$, and therefore R , is conformally equivalent to one of the standard collars from the list R_1, \dots, R_N , and therefore has conformal modulus at least μ , the smallest of the positive moduli of those ring domains. We have determined a positive constant μ such that, starting with an arbitrary compact subset of \mathbb{G} , there is a ring domain R of modulus $\geq \mu$ that separates D from infinity, and the Criterion for Parabolicity Proposition 5.2 applies to conclude that $\mathbb{G} = \mathbb{C}$, and hence that \mathcal{T}_L is parabolic. \square

5.2. Main Theorem. For the hierarchies discussed above, the associated subdivision sequences $\{\sigma_n\}$ were generic. However, the hierarchies of real interest to us

arise from subdivision rules. In particular, we bring in the special notion of *conformal subdivision rule* which was central to our work in [7] and was recalled earlier in §1.

Definition. If τ is a conformal subdivision rule that manifests a combinatorial hierarchy $\{K_n\}$ for K , then for every n the tiling $\mathcal{T}_{n+1} = \mathcal{T}_{K_{n+1}}$ can be so chosen that it is a conformal subdivision of \mathcal{T}_n , and in this case we say that τ *manifests a conformal hierarchy* $\{\mathcal{T}_n\}$ for \mathcal{T}_K .

Conformal subdivision occurs in situ, so we have

$$\mathcal{T}_{n+1} = \mathcal{T}_{K_{n+1}} = \mathcal{T}_{\tau K_n} = \tau \mathcal{T}_{K_n} = \tau \mathcal{T}_n \text{ for all } n \in \mathbb{Z}.$$

In other words, conformal hierarchies display the blend of combinatorics and geometry that we have been targeting, and we can state our Main Theorem.

Main Theorem 5.5. *Let τ be a shrinking conformal subdivision rule that manifests a combinatorial hierarchy for a combinatorial tiling K of bounded degree. Then τ manifests a conformal hierarchy for $\mathcal{T} = \mathcal{T}_K$ and \mathcal{T} is necessarily parabolic. Moreover, \mathcal{T}_L is parabolic for all combinatorial tilings L that are locally isomorphic to K .*

Proof. Let $\{\mathcal{T}_n\}$ and $\{K_n\}$ be the conformal and associated combinatorial hierarchies. Theorem 4.2 implies that $\{K_n\}$, and hence $\{\mathcal{T}_n\}$, is strongly expansive. By Theorem 5.4, type is constantly parabolic across the whole local isomorphism class of K . \square

The richest source we currently have for conformal subdivision rules is provided by Theorem 1.1—namely, dihedrally symmetric simple subdivision operators are conformal. Rules τ formed from compositions of these are likewise conformal. In any case, if τ satisfies the additional hypotheses of finiteness, primitivity, and bounded degree, Theorem 3.3 guarantees existence of a combinatorial hierarchy for any expansion complex for τ , and we may then apply the Main Theorem. A particular instance is the pentagonal tiling P discussed in Example 4.2. We are thus able to answer the question of Maria Ramirez-Solano that instigated this study; namely, we have proven the following.

Proposition 5.6. *All conformal tilings locally isomorphic to the pentagonal tiling P are parabolic.*

The Main Theorem applies to many other subdivision rules τ , and these typically manifest uncountably many different conformal hierarchies. We illustrate an example to motivate investigations in the next section.

Example 5.1. Let τ denote the “lace” subdivision rule, based on the lace algorithm of Figure 1. This is a simple, dihedrally symmetric, shrinking, and strictly primitive subdivision operator, hence conformal. Figure 5 shows four stages of τ -subdivision applied to an initial 4-gon. This fragment is—up to round-off error—conformally correct, but has been mapped to a euclidean square for display.

In this example you are asked to imagine the tiles in black as part of an infinite conformal tiling $\mathcal{T} = \mathcal{T}_0$, necessarily parabolic. Several stages of aggregation have been color-coded to visually highlight the fact that τ is conformal. The red-sided tiles approximate first level aggregates in \mathcal{T}_{-1} , the green-sided, second level aggregates in \mathcal{T}_{-2} , and the blue-sided, third level aggregates \mathcal{T}_{-3} in the conformal

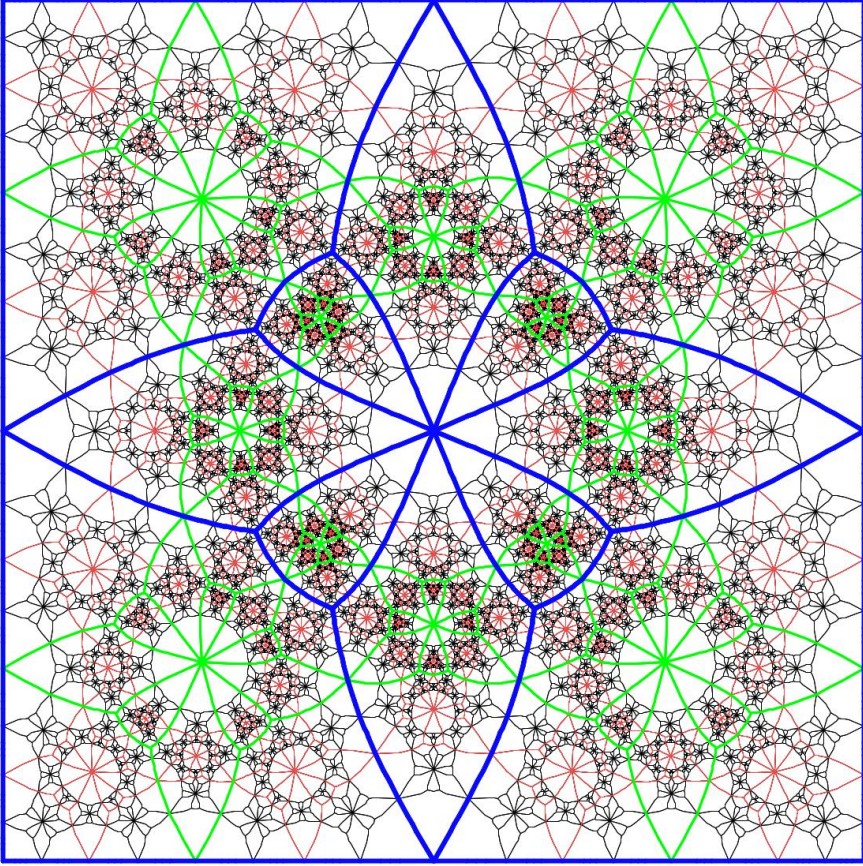


FIGURE 5. A patch of an infinite lace tiling \mathcal{T} , with colors encoding 4 levels of aggregation.

hierarchy. The surrounding blue-sided square represents a fourth level aggregate in \mathcal{T}_4 that we have isolated as a euclidean square. In the final section of this paper we refer back to this image as motivation for work in “Conformal tilings III” [4]. \checkmark

6. HIERARCHICAL PERIODICITY

In this, the penultimate section, we address an important but exceptional type of self-similarity which we call *hierarchical periodicity*. These symmetries occur in maps from a tiling to its own subdivisions within a hierarchy—in other words, these are symmetries at the level of *hierarchies* rather than of individual tilings.

Recall from §3 that τ is a combinatorial period map for K if $\tau^m K \cong K$. Our interest is in conformal realizations of this behavior. Let us start with a general observation: Suppose $\mathcal{T} = \mathcal{T}_K$ is a conformal tiling in $\mathbb{G} \in \{\mathbb{C}, \mathbb{D}\}$ (i.e., \mathcal{T} is conformally regular and reflective) and suppose μ is in $\text{Aut}(\mathbb{G})$, the group of conformal

automorphisms of \mathbb{G} . We write $\mu\mathcal{T}$ for the tiling

$$\mu\mathcal{T} = \{\mu(t) : t \in \mathcal{T}\}.$$

This is necessarily again a conformal tiling of \mathbb{G} and clearly $K_{\mu\mathcal{T}} \cong K_{\mathcal{T}}$. The situation of interest is captured in the following definition.

Definition. Let \mathcal{T} be a conformal tiling in \mathbb{G} , $\mu \in \text{Aut}(\mathbb{G})$. If every tile of \mathcal{T} is a finite union of tiles of $\mu\mathcal{T}$, then μ is called a (*conformal*) *period map* for \mathcal{T} .

Lemma 6.1. *Suppose \mathcal{T} is a conformal tiling of \mathbb{G} and μ and ν are two period maps for \mathcal{T} . Then $\lambda = \nu \circ \mu$ is also a period map for \mathcal{T} .*

Proof. Clearly, $\lambda\mathcal{T}$ is a conformal tiling for $K_{\mathcal{T}}$: the issue is whether the tiles of $\lambda\mathcal{T}$ decompose the tiles of \mathcal{T} in situ. By hypothesis, this does happen for μ , so as a preliminary, we apply μ to break each tile $t \in \mathcal{T}$, in situ, into a pattern of tiles from $\mu\mathcal{T}$ and we outline these subtile shapes in pencil on t . The fact that this pattern of subtiles is conformally regular and reflective persists in any image of t under a conformal mapping.

We next need to apply ν . Fix attention on a concrete tile $t_0 \in \mathcal{T}$. It is not $\nu(t_0)$ that we want to look at, but rather t_0 itself. Because ν is a period map, $t_0 = \bigcup_{j=1}^k \nu(t_j)$ for some tiles $t_j \in \mathcal{T}$. Each tile $\nu(t_j)$ carries with it our markings from the earlier application of μ . Applying ν to t_j carries these penciled shapes along into t_0 , and now, of course, the penciled shapes are images of tiles of \mathcal{T} under $\nu \circ \mu = \lambda$. The union of all of these penciled shapes, $j = 1, \dots, k$, gives a subdivision of t_0 . In other words, λ has decomposed t_0 in place as a finite union of tiles of $\lambda\mathcal{T}$. This happens for every concrete tile t_0 , so $\lambda\mathcal{T}$ is a conformal subdivision of \mathcal{T} . \square

For conformal subdivision rules, conformal and combinatorial period maps are in one-to-one correspondence. A conformal period map μ for $\mathcal{T} = \mathcal{T}_K$ induces a combinatorial subdivision rule σ_μ for K , namely, $\sigma_\mu K \cong K_{\mu\mathcal{T}}$. Because $K_{\mu\mathcal{T}} \cong K_{\mathcal{T}} = K$, we have $\sigma_\mu K \cong K$, so σ_μ is a combinatorial period map for K . In the other direction, a subdivision rule σ which is a combinatorial period map for K induces a conformal period map if and only if σ is a *conformal* subdivision rule. In that case, $\mathcal{T}_{\sigma K}$ can be realized as an in situ conformal subdivision of \mathcal{T}_K in \mathbb{G} . Because $\sigma K \cong K$, Theorem 1.2 then implies existence of a conformal automorphism μ so that $\mu\mathcal{T}_K = \mathcal{T}_{\sigma K}$. In other words, μ is a conformal period map for \mathcal{T}_K with $\sigma_\mu = \sigma$.

It is not difficult to take the next step in hierarchical periodicity. To be explicit, let $\mathcal{T}_0 = \mathcal{T}$ and for each integer n define \mathcal{T}_n by

$$(2) \quad \mathcal{T}_n = \mu^n \mathcal{T} \equiv \{\mu^n(t) : t \in \mathcal{T}\}.$$

Since μ is a conformal automorphism, each tiling \mathcal{T}_n is a reflective tiling by conformally regular polygons, hence is conformal. Since $\mu\mathcal{T}$ conformally subdivides \mathcal{T} , a moment's consideration will convince the reader that, for each integer n , the tiling \mathcal{T}_{n+1} conformally subdivides \mathcal{T}_n and, in fact, $\mathcal{T}_{n+1} = \sigma\mathcal{T}_n$, where $\sigma = \sigma_\mu$.

Theorem 6.2. *Let K be a combinatorial tiling of bounded degree and let $\mathcal{T} = \mathcal{T}_K$. Let τ be a shrinking conformal subdivision rule and suppose K is τ -periodic with τ -period m . Then \mathcal{T} tiles the complex plane \mathbb{C} and there exists a Möbius transformation μ that realizes the subdivision rule τ^m in the sense that $\mu\mathcal{T} = \{\mu(t) : t \in \mathcal{T}\} = \tau^m\mathcal{T}$. The transformation μ fixes ∞ and has a single attracting fixed point in \mathbb{C} .*

Proof. Define the subdivision rule $\sigma = \tau^m$, so K has σ -period 1. The existence of μ was discussed above. Define the conformal hierarchy $\{\mathcal{T}_n\}$ as in (2), so that for all integers n , $\mathcal{T}_{n+1} = \mu\mathcal{T}_n$. Since τ is shrinking, σ is shrinking, hence strongly expansive, and hence, by Theorem 5.4, \mathcal{T} tiles the complex plane. Therefore, μ has the form $\mu(z) = az + b$ for some complex constants $a \neq 0$ and b . Since $\sigma\mathcal{T}$ is a non-trivial subdivision, μ is not the identity, and we will argue in turn that μ is not a translation ($a \neq 1$) and not a rotation ($|a| \neq 1$), but rather is contractive ($0 < |a| < 1$).

Suppose μ were a translation, so $\mu(z) = z + b$ for some non-zero complex constant b . Let d be a positive integer strictly larger than the number of faces that meet at any vertex of the complex $K_{\mathcal{T}}$ and, given an arbitrary tile t of $\mathcal{T} = \mathcal{T}_0$, let D be the union of the tiles $\mu^j(t) = t + jb$ for $j = 1, \dots, d$. Since the conformal hierarchy $\{\mathcal{T}_n\}$ is expansive, there is, for some integer $n < 0$, a core $|c|$ of the tiling \mathcal{T}_n that engulfs D . Let p be a point interior to the tile t . Since the core $|c|$ is the union of fewer than d tiles, there is a tile s of \mathcal{T}_n that contains at least two of the points from the list $p + b, \dots, p + db$. Let k be the smallest integer in the list $1, \dots, d$ for which $p + kb \in s$, and let ℓ be the largest positive integer for which $p + \ell b \in s$. Note that $k < \ell$ so that $j = \ell - k > 0$ and note that it may well be that $\ell > d$. Consider the tile $t_k = \mu^k(t) = t + kb \in \mathcal{T}_k$. Since p is interior to t , $p + kb$ is interior to t_k , and so the tile s in \mathcal{T}_n meets the interior of the tile t_k of $\mathcal{T}_k = \sigma^{k-n}\mathcal{T}_n$. Since \mathcal{T}_k is a polygonal subdivision of \mathcal{T}_n , this implies that $t_k \subset s$, and since $p + kb$ is interior to t_k , $p + kb$ is interior to s . The same argument applied to $t_\ell = t + \ell b$ implies that $t_\ell \subset s$ and that $p + \ell b$ is interior to s . Choose $\varepsilon > 0$ so that the disc neighborhoods $D_k = \{|z - (p + kb)| < \varepsilon\}$ and $D_\ell = \{|z - (p + \ell b)| < \varepsilon\}$ are contained in the tile s , and observe that $\mu^j(D_k) = D_\ell$. It follows that the tile $\mu^j(s)$ of \mathcal{T}_{n+j} contains the open disc D_ℓ , as does s , and so the tile s and the tile $\mu^j(s)$ meet in an interior point of both. Since \mathcal{T}_{n+j} subdivides \mathcal{T}_n , we conclude that $\mu^j(s) \subset s$. This implies, since $p + \ell b \in s$, that $p + (\ell + j)b = \mu^j(p + \ell b) \in s$. Since j is positive, this contradicts the choice of ℓ as the largest positive integer for which $p + \ell b \in s$. We conclude that μ cannot be a translation, so $a \neq 1$.

Since $a \neq 1$, μ has a unique fixed point in \mathbb{C} , namely the fixed point $z_0 = b/(1-a)$. Let t be a tile of \mathcal{T} that contains z_0 and that has maximum area among all such tiles of \mathcal{T} . Then $\mu(t)$ is a tile of $\sigma\mathcal{T}$ containing $z_0 = \mu(z_0)$ and since $\sigma\mathcal{T}$ is a conformal subdivision of \mathcal{T} , $z_0 \in \mu(t) \subset t^*$ for some tile t^* of \mathcal{T} . Since σ is non-trivial, this means $\mu(t) \subsetneq t$ and $\text{area}(\mu(t)) < \text{area}(t^*)$. By our choice of t , $\text{area}(t^*) < \text{area}(t)$, implying $\text{area}(\mu(t)) < \text{area}(t)$. Thus μ must be contracting, so that $0 < |a| < 1$, and z_0 must be an attracting fixed point for μ . \square

The connections among μ , τ , and $\sigma = \tau^m$ can be subtle, as the authors know keenly from experience. Therefore, some further observations may be in order:

- 1 Just because K is τ -periodic, not every conformal hierarchy $\{\mathcal{T}_n\}$ for K is necessarily τ -periodic. We saw with the quad tiling (see Example 4.4) that

- aggregation, as an inverse to subdivision, can be ambiguous. That is to say, $\mu\mathcal{T}' = \mathcal{T}_{n+1} = \tau^m\mathcal{T}_n$ does not necessarily imply that $\mathcal{T}' = \mathcal{T}_n$.
- 2 On the other hand, the particular hierarchy $\{\mathcal{T}_n\}$ defined in (2) is σ -periodic by construction, with σ -period 1; so for all integers n , $\mathcal{T}_n = \mu^n\mathcal{T}_0$. That is, within this particular hierarchy, subdivision is realized by μ and a particular aggregation—one of potentially many—is realized by μ^{-1} .
 - 3 Suppose K is τ -periodic with τ -period $m > 1$ and $\mu\mathcal{T} = \tau^m\mathcal{T}$. Although the hierarchy constructed in (2) is for $\sigma = \tau^m$, there is a canonical extension to a conformal hierarchy $\{\mathcal{T}'_n\}$ for τ itself: For each integer n , write $n = sm + t$ for integers s and t , $0 \leq t < m$, and define $\mathcal{T}'_k = \tau^t\mathcal{T}_s$. This inserts $m - 1$ successive subdivisions by τ between each pair of successive subdivisions by σ . One easily verifies that for all integers n , $\mu\mathcal{T}_n = \mathcal{T}_{n+m}$. The associated combinatorial tilings $K'_n = K_{\mathcal{T}'_n}$ break into m distinct conjugacy classes which are cyclical under subdivision by τ .

Conformal period maps are rare but much sought after. Observed periodicity in the regular pentagonal tiling was key to the founding of this topic, and related *quasiconformal* period maps are the holy grail in Cannon, Floyd, and Parry's approach to Cannon's Conjecture. Moreover, in §6.2 we see how to easily construct concrete examples.

Nonetheless, periodicity is the exception. We will see shortly that in the first instance, for K to be τ -periodic is rare. Even then, Observation (1) is cautionary: a generic conformal hierarchy manifest by τ will not be periodic. Observations (2) and (3) suggest a loosening of the requirements of periodicity that will serve us well in future developments of [4].

6.1. Periodic expansion and the action of $\hat{\tau}$. If τ is a non-trivial subdivision rule, then τK is never equal to K , but it may be *isomorphic* to K . This is where $\hat{\tau}$, introduced on page 72, comes in: while τK may be a specific combinatorial tiling that subdivides K , $\hat{\tau}K$ is the isomorphism class of τK in \mathbf{C} . Thus we always have $\hat{\tau} : (K) \rightarrow (\tau K)$. When τK is locally isomorphic to K , then we have $\hat{\tau} : (K) \rightarrow (K)$, and this is the situation we address in this subsection. To avoid trivialities, we assume that (K) is not a singleton.

Let K be an expansion complex for a strictly primitive finite subdivision rule τ . By Theorems 3.1 and 3.2, $\hat{\tau} : (K) \rightarrow (K)$ is surjective. (It would seem that $\hat{\tau}$ is also *injective*. But this concerns the uniqueness of aggregation and is a challenging open question. See [25] for a proof of injectivity for certain traditional subdivision rules.) If $L \in (K)$, then L is τ -periodic if and only if L lies in a finite orbit in (K) under the forward iteration of $\hat{\tau}$. The τ -period of L is the number of tilings in that orbit. Our main goal in this section is to construct τ -periodic expansion complexes, but we first prove that these are the exception rather than the rule.

Theorem 6.3. *Let τ be a shrinking, primitive finite subdivision rule of bounded degree. Then for any $N > 0$ there are, up to isomorphism, at most finitely many τ -periodic expansion complexes K having τ -period bounded by N .*

Proof. Note first that an expansion complex K for τ is automatically an expansion complex for any positive power τ^s and vice versa. Moreover, if K is τ -periodic, then it is also τ^s -periodic, and the orbit under τ is no bigger than that under τ^s . It suffices, then, to prove our result with τ replaced by τ^s ; that is, we may assume

without loss of generality that τ is *strictly* shrinking and *strictly* primitive. The proof is divided into three parts.

PART 1: CLASSIFYING FIXED POINTS. Let J be an expansion complex for τ that is a fixed point of the action of $\hat{\tau}$. Then J is a CW-decomposition of the plane and τJ is a CW-complex that subdivides J with $\hat{\tau}J = J$. This means that J is isomorphic to τJ and we let $\lambda_J : \mathbb{C} \rightarrow \mathbb{C}$ be a homeomorphism of the complex plane that is a cellular isomorphism of J onto τJ . (Note that λ_J and all other cellular mappings in this proof respect the tile types of \mathfrak{T} .)

In Part 1, we work under this

Fixture Assumption: *There is a closed face or edge cell b_J of J for which $\lambda_J(b_J) \subset b_J^\circ$, or a vertex cell b_J of J for which $\lambda_J(b_J) = b_J$.*

We will call b_J a *fixture* for λ_J , and the goal of this first part of the proof is to characterize combinatorially those fixed points J of $\hat{\tau}$ with isomorphisms λ_J having fixtures, and thereby to put a limit on their number.

Let C_J be the core of J determined by the cell b_J and let c_J be the core of τJ determined by the cell $\lambda_J(b_J)$. C_J and c_J are both vertex, edge, or face cores depending on whether the cell b_J is, respectively, a vertex, edge, or face of J , and, in fact, it is easy to see that $\lambda_J(C_J) = c_J$. A straightforward argument using the fact that τ is strictly shrinking shows that the core $\lambda(c_J)$ of $\tau^2 J$ is contained in the interior $(\tau^2 C_J)^\circ$ of the τ^2 -subdivided core C_J .

Let J and L be expansion complexes that serve as fixed points of the action of $\hat{\tau}$, both of which satisfy the fixture assumption. Define J and L to be λ -equivalent if there exists a homeomorphism $h : |C_J| \rightarrow |C_L|$ that is a cellular isomorphism of the cell complex τC_J onto the cell complex τC_L such that $h(c_J) = c_L$ and for which $h^{-1} \circ \lambda_L \circ h = \lambda_J|_{|C_J|}$. Our aim is to show that if J and L are λ -equivalent fixed points of $\hat{\tau}$, then $J \cong L$ so that $J = L$ in (K) . To prove this, we first show how to recover J from the triple $(C_J, c_J, \lambda_J|_{C_J})$. Indeed, our claim is that

$$J = \bigcup_{i=1}^{\infty} \lambda_J^{-i}(\tau^i C_J),$$

where $\lambda_J^i : J \rightarrow \tau^i J$ is the isomorphism of complexes gotten from iterating the mapping λ_J i times. To verify the claim, note that, since $c_J = \lambda_J(C_J)$ is a subcomplex of τC_J , the core $C_J = \lambda_J^{-1}(c_J)$ is a subcomplex of $\lambda_J^{-i}(\tau^i C_J)$ for every non-negative integer i . It follows that the union $J' = \bigcup_{i=1}^{\infty} \lambda_J^{-i}(\tau^i C_J)$ is a connected subcomplex of J . To see that $J' = J$, it suffices to verify that, for each non-negative integer i , the complex $\lambda_J^{-i}(\tau^i C_J)$ is contained in the interior of $\lambda_J^{-i-2}(\tau^{i+2} C_J)$. But this follows from induction with the basis established from the fact that the complex $\lambda(c_J)$ is contained in the interior of $\tau^2 C_J$, implying after an application of λ_J^{-2} that C_J is a subcomplex of J contained in the interior of $\lambda_J^{-2}(\tau^2 C_J)$. Now, assuming that the expansion complex L is λ -equivalent to J , as with J write $L = \bigcup_{i=1}^{\infty} \lambda_L^{-i}(\tau^i C_L)$ and let $h : |C_J| \rightarrow |C_L|$ be a homeomorphism with the properties described in the definition of λ -equivalence. This means that h is a homeomorphism of $|C_J|$ onto $|C_L|$ that is a cellular isomorphism of τC_J onto τC_L , and this implies by the obvious τ -aggregation that h is also a cellular isomorphism of the complex C_J onto the complex C_L . From $h^{-1} \circ \lambda_L \circ h = \lambda_J|_{|C_J|}$ we may infer that $h = \lambda_L^{-1} \circ h \circ \lambda_J|_{|C_J|}$. For each positive integer i , this allows us to extend the cellular isomorphism h of C_J onto C_L to a cellular isomorphism $h_i = \lambda_L^{-i} \circ h \circ \lambda_J^i|_{|J_i|}$ of $J_i = \lambda_J^{-i}(\tau^i C_J)$ onto $L_i = \lambda_L^{-i}(\tau^i C_L)$ and we obtain the diagram of commuting cellular containments

and cellular isomorphisms h_i :

$$\begin{array}{ccccccc}
 C_J & \hookrightarrow & \lambda_J^{-1}(\tau C_J) & \hookrightarrow & \lambda_J^{-2}(\tau^2 C_J) & \hookrightarrow & \cdots \xrightarrow{\lim} J \\
 \downarrow h \cong & & \downarrow h_1 \cong & & \downarrow h_2 \cong & & \\
 C_L & \hookrightarrow & \lambda_L^{-1}(\tau C_L) & \hookrightarrow & \lambda_L^{-2}(\tau^2 C_L) & \hookrightarrow & \cdots \xrightarrow{\lim} L.
 \end{array}$$

This induces an isomorphism of CW-complexes $J \cong L$ and implies that $J = L$ in (K) .

Armed with the observation of the preceding paragraph, we can verify that there are at most finitely many fixed points of the action of $\hat{\tau}$ on the local isomorphism class (K) that satisfy the fixture assumption. Indeed, we have shown that each fixed point J of the action of $\hat{\tau}$ that satisfies the fixture assumption identifies a CW-pair $(\tau C_J, c_J)$, where C_J is a core of J and c_J is a core of τJ , and that $J \cong L$ whenever L is a fixed point of $\hat{\tau}$ satisfying the fixture assumption that is λ -equivalent to J . Since τ has bounded degree, the expansion complex K has bounded degree and, up to isomorphism, there are only finitely many pairs $(\tau C, c)$ where C is a core of K and c is a core of τK . Since the expansion complex J is locally isomorphic to K , the pair $(\tau C_J, c_J)$ is, up to isomorphism, one of these finitely many pairs $(\tau C, c)$ from K . Moreover, for any fixed pair (C, c) , there are, up to cellular isotopy, only finitely many orientation-preserving cellular isomorphisms λ from the cell complex C onto the cell complex c . This implies that each triple $(C_J, c_J, \lambda_J|_{C_J})$ is represented among the finitely many distinct triples (C, c, λ) , and this implies that there are only finitely many λ -equivalence classes of expansion complexes in (K) that satisfy the fixture assumption, and therefore only finitely many fixed points of $\hat{\tau}$ that satisfy the fixture assumption. This completes Part 1.

PART 2: EXISTENCE OF A FIXTURE FOR A POWER OF λ_J . In this second part of the proof, we would like to show that any cellular isomorphism λ_J of any fixed point J of $\hat{\tau}$ onto its subdivision τJ has a fixture; unfortunately this fails to be true. At the conclusion of the proof, we will give an example. Our aim in Part 2, then, is to prove that the strictly shrinking subdivision operator τ determines a positive integer M such that, for every fixed point J of $\hat{\tau}$, there exists a positive integer $m \leq M$ such that the cellular isomorphism λ_J^m of J to $\tau^m J$ has a fixture. In fact, we will show that $M = 6\beta$ works, where β is a degree bound for τ .

The verification of the existence of a fixture for a positive power of λ_J depends strongly on the fact that τ is expansive. Let J be any fixed point of the action of $\hat{\tau}$. We know from Theorem 3.3 that J exhibits a combinatorial hierarchy manifested by τ , but we can say more in this case. Indeed, forward and backward iteration of the mapping λ_J builds such a hierarchy $\{J_k\}$, where for each integer k , $J_k = \lambda_J^k(J)$. Since τ is shrinking, Theorem 4.2 implies that the hierarchy $\{J_k\}$ is exponentially expansive. Our first goal is to show that there exist positive integers p and q and a polygonal face \mathbf{h} of J_{-p} such that $\lambda_J^q(\mathbf{h}) \subset \mathbf{h}$. Toward this goal, let \mathbf{k}_0 be any face of J and, for any positive integer m , let \mathbf{k}_m be that unique face of J that contains the face $\lambda_J^m(\mathbf{k}_0)$ of J_m . Since the hierarchy is expansive, there is a positive integer p and a core c of J_{-p} that engulfs $\mathbf{k}_0 \cup \mathbf{k}_1 \cup \cdots \cup \mathbf{k}_\beta$, where β is a degree bound for τ , and therefore a bound on the number of faces that meet at any vertex of any J_k of the hierarchy. Since the subcomplex c of J_{-p} is a core, it has at most β faces

and it follows that, since each of the faces $\mathbf{k}_0, \mathbf{k}_1, \dots, \mathbf{k}_\beta$ is a face of the subcomplex $\tau^p c$, there is at least one face, say \mathbf{h} , of c that contains two of the faces from the list $\mathbf{k}_0, \mathbf{k}_1, \dots, \mathbf{k}_\beta$. Assume that the two faces \mathbf{k}_i and \mathbf{k}_j satisfy $\mathbf{k}_i \cup \mathbf{k}_j \subset \mathbf{h}$, where $0 \leq i < j \leq \beta$ are chosen so that $q = j - i$ is as small as possible. By the definition of the faces \mathbf{k}_m , we have $\lambda_J^i(\mathbf{k}_0) \cup \lambda_J^j(\mathbf{k}_0) \subset \mathbf{h}$ with $\lambda_J^q(\lambda_J^i(\mathbf{k}_0)) = \lambda_J^j(\mathbf{k}_0)$. It follows that the face $\lambda_J^q(\mathbf{h})$ of $\lambda_J^q(J_{-p}) = J_{-p+q}$ contains the face $\lambda_J^j(\mathbf{k}_0)$ of its subdivision J_j , as does the face \mathbf{h} of J_{-p} . This shows that the face $\lambda_J^q(\mathbf{h})$ of J_{-p+q} meets the open face \mathbf{h}° since $\lambda_J^j(\mathbf{k}_0)^\circ \subset \lambda_J^q(\mathbf{h}) \cap \mathbf{h}^\circ$. Since the face $\lambda_J^q(\mathbf{h})$ of the complex J_{-p+q} meets the open face \mathbf{h}° of J_{-p} , and since J_{-p+q} subdivides J_{-p} , we conclude that $\lambda_J^q(\mathbf{h}) \subset \mathbf{h}$.

Since \mathbf{h} is a face of J_{-p} , the combinatorial n -cell $\mathbf{a} = \lambda_J^p(\mathbf{h})$ is a face of J with $\lambda_J^q(\mathbf{a}) = \lambda_J^{p+q}(\mathbf{h}) \subset \lambda_J^p(\mathbf{h}) = \mathbf{a}$. If λ_J^q fixes a vertex v of \mathbf{a} , then $b_J = \{v\}$ is a fixture of λ_J^q , or if $\lambda_J^q(e) \subset e^\circ$ for an edge e of \mathbf{a} , then $b_J = e$ is a fixture of λ_J^q . Assume that λ_J^q neither fixes a vertex of \mathbf{a} nor maps an edge of \mathbf{a} into the corresponding open edge. Since τ is strictly shrinking, so too is τ^q . This implies that $\lambda_J^q(\mathbf{a})$, a face of $\tau^q J$ and a subset of \mathbf{a} , is a subset of an open angle $\angle_a v$ for a vertex v of \mathbf{a} . If $\lambda_J^q(\mathbf{a}) \subset a^\circ$, then $b_J = \mathbf{a}$ is a fixture of λ_J^q ; otherwise, $\lambda_J^q(\mathbf{a})$ meets one or both of the half-open edges $\{v\} \cup d^\circ$ and $\{v\} \cup e^\circ$, where d and e are the edges of \mathbf{a} incident to v . The remainder of the argument rests on where v goes under the action of λ_J^q . There are four possibilities: $\lambda_J^q(v) \in \angle_a v = \{v\} \cup d^\circ \cup e^\circ \cup \mathbf{a}^\circ$, and as this is a disjoint union, $\lambda_J^q(v)$ lies in exactly one of the open cells $\{v\}$, d° , e° , or \mathbf{a}° . We have assumed though that λ_J^q fixes no vertex of \mathbf{a} , so the first possibility is ruled out. The last possibility, that $\lambda_J^q(v) \in \mathbf{a}^\circ$, implies that the image of the open angle $\angle_a v$ under λ_J^q is contained in the open cell \mathbf{a}° , and from this we have $\lambda_J^{2q}(\mathbf{a}) \subset \mathbf{a}^\circ$, so that $b_J = \mathbf{a}$ is a fixture of λ_J^{2q} . The remaining two possibilities are symmetric, so we assume that $\lambda_J^q(v) \in e^\circ$. This implies, since λ_J^q is a cellular isomorphism of J to a subdivision J_q , that $\lambda_J^q(e^\circ)$ is contained in either e° or \mathbf{a}° . In the former case, since $\lambda_J^q(e)$ also is contained in the open angle $\angle_v \mathbf{a}$, then $\lambda_J^q(e) \subset \{v\} \cup e^\circ$. But this implies that $\lambda_J^{2q}(e) \subset e^\circ$ and therefore $b_J = e$ is a fixture of λ_J^{2q} . In the latter case, $\lambda_J^{2q}(v) \in \mathbf{a}^\circ$, implying that the image of the open angle $\angle_a v$ under λ_J^{2q} is contained in the open cell \mathbf{a}° . From this we have $\lambda_J^{3q}(\mathbf{a}) \subset \mathbf{a}^\circ$, so that $b_J = \mathbf{a}$ is a fixture of λ_J^{3q} . This paragraph's discussion verifies that at least one of the mappings λ_J^q , λ_J^{2q} , or λ_J^{3q} has a fixture. A quick inductive argument proves that, for any positive integer s , a fixture for λ_J^s is a fixture for λ_J^{st} for all positive integers t . We conclude that the mapping λ_J^{6q} has a fixture, and since $1 \leq q \leq \beta$, we may set $M = 6\beta$. This concludes the second part of the proof.

PART 3: THE GENERAL CASE. Our aim is to bound the number of τ -periodic expansion complexes in (K) having τ -periods bounded by some positive integer k . However, Part 2 gives us one more reduction of our problem; for every fixed point J of $\widehat{\tau}$, there exists a positive integer $m \leq 6\beta$ such that λ_J^m has a fixture, where λ_J is a cellular isomorphism of J onto τJ . Observe that $\lambda_J^m = \lambda_{\tau^m J}$, that is, λ_J^m is a cellular isomorphism from J to $\tau^m J$, and that every fixed point of $\widehat{\tau}^m$ with order at most k is a fixed point of τ with order at most mk . Without loss of generality, therefore, we may replace τ by τ^m and assume henceforth that λ_J has a fixture.

We are now reduced to finding a bound on the fixed points of $\widehat{\tau}$ with order bounded by k . But if J is a fixed point of $\widehat{\tau}$, then it has a fixture, and we proved in

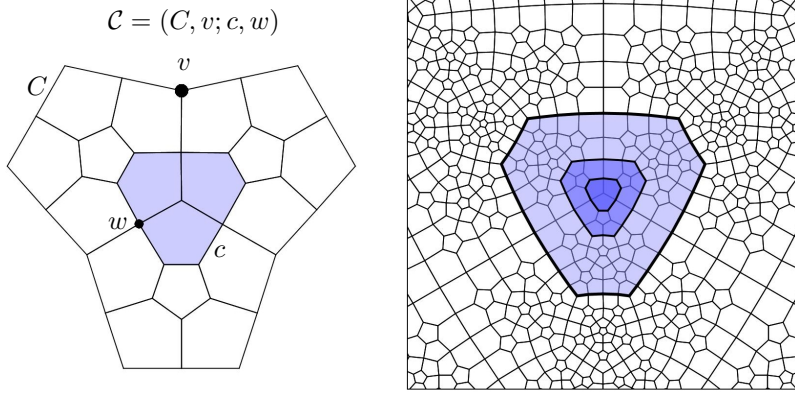


FIGURE 6. A τ -periodic pentagonal tiling based on a vertex core and having τ -period 1; see Example 6.3 for the description. Note in particular that the periodicity of interest is *not* the rotational symmetry that the eye picks out.

Part 1 that the set of such expansion complexes forms a finite subset of the local equivalence class (K) .

Finally, recall the integer s chosen in the opening paragraph of this proof and the bound 6β on the m of Part 2 and let $\tau' = \tau^{6s\beta}$. Given $N > 0$, we have actually shown that the number of τ' -periodic expansion complexes $J \in (K)$ having τ' -period bounded by N is finite. But if J has τ -period bounded by N , then it automatically has τ' -period bounded by N . The number of these is thus finite and the proof is complete. \square

Example 6.1. Consider the hyperbolic tiling \mathbf{H} of Example 4.3. The homeomorphism $z \mapsto z + \mathbf{i}$ of \mathbb{C} (see Figure 4) provides a fixture-free cellular isomorphism of \mathbf{H} to $\sigma_0\mathbf{H}$, where σ_0 is the non-shrinking subdivision rule of the example that divides each pentagonal face into two pentagons and for which $\widehat{\sigma}_0\mathbf{H} = \mathbf{H}$. All powers of σ_0 are also fixture-free. The difficulty here is that hierarchy constructed in Example 4.3 is not expansive. In fact, it is not difficult to see that the mapping $\widehat{\sigma}_0$ induced on the local isomorphism class (\mathbf{H}) is the identity, so that every element of (\mathbf{H}) is a fixed point even though (\mathbf{H}) is uncountably infinite.¹

Example 6.2. The Gaussian complex \mathbf{Z} of Example 4.1 and its “quad” subdivision rule τ provide an example of a cellular isomorphism that fails to have a fixture, even though τ is a strictly shrinking, rotationally symmetric, and of bounded degree. Let $R_{\pi/2}$ be the counterclockwise rotation of $\pi/2$ radians about $w = (1/2, 1/2)$, the center of the unit square face \mathbf{k} of \mathbf{Z} whose lower left-hand vertex is the origin $\mathcal{O} = (0, 0)$, and let $M_{1/2}$ be the map that multiplies each complex number by $1/2$, and so dilates the complex plane toward the origin \mathcal{O} . Then $\lambda_{\mathbf{Z}} = M_{1/2} \circ R_{\pi/2}$ is a cellular isomorphism of \mathbf{Z} onto $\tau\mathbf{Z}$ with no fixture. Though $\lambda_{\mathbf{Z}}^2$ also fails to have a fixture, a moment’s consideration shows that the face \mathbf{k} serves as a fixture of $\lambda_{\mathbf{Z}}^3$. $\sqrt{}$

¹Dane Mayhook’s doctoral thesis explores the local isomorphism class (\mathbf{H}) of this hyperbolic complex and gives a constructive description of all elements of (\mathbf{H}) .

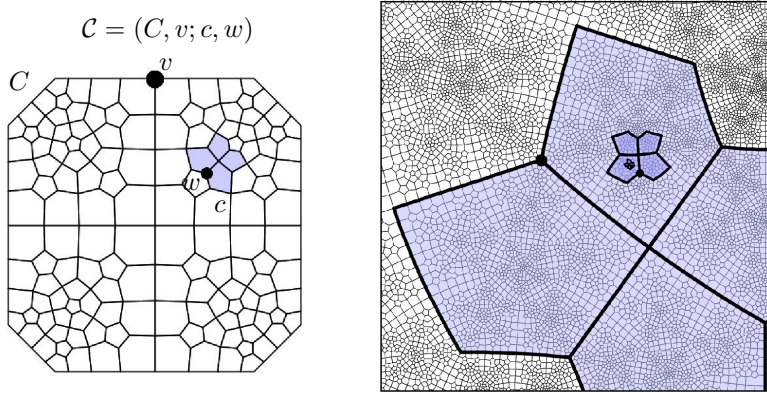


FIGURE 7. A pentagonal tiling having τ -period 2. There are no global symmetries in the tiling itself, the symmetry of interest lies within the hierarchy. See Example 6.4 for details.

6.2. Conformal period maps. Reverse engineering the proof of Part 1 of Theorem 6.3 opens the way to construction of conformal period maps. We describe the process supposing that τ is a shrinking, primitive, finite, subdivision rule with degree bounded by β . Indeed, for ease of construction, we assume that τ is *strictly* shrinking and *strictly* primitive. We intersperse images of concrete conformal examples to help the reader, though their details will only be provided at the end of this section.

It is the “fixtures” of Part 1 that we need to reverse engineer. Here is our mechanism.

Definition. Let Δ be a tile with type in the collection \mathfrak{T} for τ and fix a positive integer $k > 0$. A *combinatorial footprint* for τ is a 4-tuple $\mathcal{C} = (C, v; c, w)$ satisfying the following properties:

- (1) C is a core contained in the interior of the complex $\tau^t \Delta$ for some positive integer t ;
- (2) c is a core of $\tau^k C$ that is contained in the interior of a combinatorial disc in C ;
- (3) v is a vertex in the boundary of C ;
- (4) w is a vertex in the boundary of c ;
- (5) there exists an orientation-preserving cellular isomorphism $\lambda_0 : C \rightarrow c$ with $\lambda_0(v) = w$. (As usual, this must respect tile types in \mathfrak{T} .)

Let $\mu_0 : c \hookrightarrow \tau^k C$ be the cellular inclusion and, for each positive integer m , let $\mu_m : \tau^{km} c \hookrightarrow \tau^{k(m+1)} C$ be the cellular inclusion induced from μ_0 , and let $\lambda_m : \tau^{km} C \rightarrow \tau^{km} c$ be the cellular isomorphism induced from λ_0 . Setting $F_m = \tau^{km} C$, for each non-negative integer m we have the isomorphic embedding $\Xi_m = \mu_m \circ \lambda_m : F_m \hookrightarrow F_{m+1}$ and a quick check verifies that this sequence of embeddings satisfies properties (1)–(4) for the definition of an expansion complex for τ . Let $F = \varinjlim F_m$ be the expansion complex defined by this sequence of embeddings. The idea here is that the isomorphic embedding Ξ_m shrinks C via λ to c and then includes c back into C via μ , and the subscript m merely tells one the level at

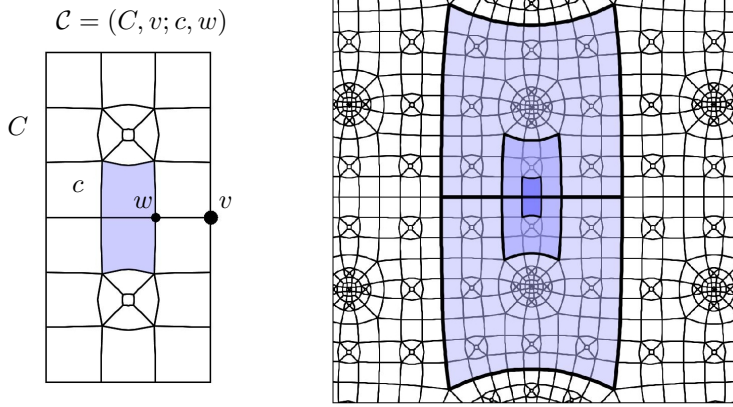


FIGURE 8. A conformal tiling for the “snowball” subdivision rule of period 1 based on an edge core. See Example 6.5 for details.

which Ξ is cellular, namely, at the level of the τ^{km} -subdivision. Now set $f_m = \tau^{km}c$ and let $\xi_m = \lambda_{m+1} \circ \mu_m : f_m \hookrightarrow f_{m+1}$, an isomorphic embedding of f_m into the interior of f_{m+1} . Again, a quick check verifies that the sequence of embeddings $\xi_m : f_m \hookrightarrow f_{m+1}$ satisfies the four properties that define an expansion complex for τ , and we let $f = \varinjlim f_m$.

The following diagram commutes, and this has interesting implications for the expansion complexes F and f :

$$\begin{array}{ccccccccccc}
 F_0 & \xrightarrow{\Xi_0} & F_1 & \xrightarrow{\Xi_1} & \cdots & \xrightarrow{\Xi_{m-1}} & F_m & \xrightarrow{\Xi_m} & F_{m+1} & \xrightarrow{\Xi_{m+1}} & \cdots & \xrightarrow{\lim} & F \\
 \lambda_0 \downarrow & \nearrow \mu_0 & \downarrow \lambda_1 & \nearrow \mu_1 & & \nearrow \lambda_m & \downarrow \mu_m & \nearrow \lambda_{m+1} & \downarrow \mu_{m+1} & & & & \\
 f_0 & \xrightarrow{\xi_0} & f_1 & \xrightarrow{\xi_1} & \cdots & \xrightarrow{\xi_{m-1}} & f_m & \xrightarrow{\xi_m} & f_{m+1} & \xrightarrow{\xi_{m+1}} & \cdots & \xrightarrow{\lim} & f.
 \end{array}$$

Since the vertical arrows are cellular isomorphisms, they induce a cellular isomorphism $\lambda : F \rightarrow f$; since the diagonal arrows are cellular inclusions, they induce a cellular inclusion $\mu : f \rightarrow \tau^k F$. Since the image of each μ_m lies interior to F_{m+1} , the induced inclusion μ must be onto, and this implies that μ is a cellular isomorphism of f onto $\tau^k F$. That the diagram commutes further implies that the composition $\lambda_F = \mu \circ \lambda : F \rightarrow \tau^k F$ is a cellular isomorphism, and hence $\hat{\tau}^k F = F$ in (K) . It follows that F generates an ℓ -orbit of the action of $\hat{\tau}$ on (K) for some positive integer $\ell \leq k$. If τ is not strictly shrinking and/or not strictly primitive, then there is some positive integer j so that τ^j has these properties. We can carry out the construction above for τ^j ; the only change is that the resulting ℓ satisfies $\ell \leq jk$.

All our concrete examples are conformal tilings. The pentagonal tilings in Figures 6 and 7 were constructed based on vertex cores, the “snowball” tiling of Figure 8, based on an edge core, and the “lace” tiling of Figure 9, based on a face core. Each figure has a schematic of its combinatorial footprint on the left. Creation of these footprints is straightforward: If C is a face core, then C has some type $t \in \mathfrak{T}$ and c must be a tile of the same type among the faces of $\tau^k C$ and interior

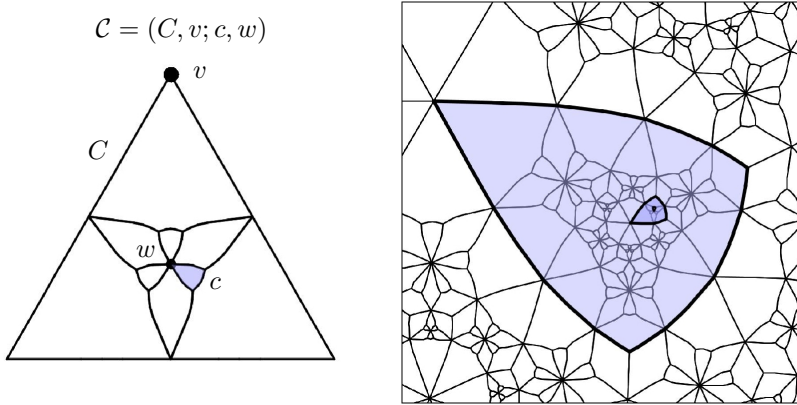


FIGURE 9. A face core for the “lace” subdivision rule with peroid 1.

to C . It remains only to choose corresponding vertices v of C and w of c to get the orientation-preserving isomorphism λ_0 . The face C , identified naturally as the base face of F , is a fixture of the isomorphism λ_F since, under this identification, $\lambda_F(C) = c \subset C^\circ$. The situation for edge and vertex cores is a bit more complicated. In these cases, the mapping λ_F may have no fixture, but of course the proof of Theorem 6.3 shows that a power $m \leq 6\beta$ of λ_F will have one.

The footprints we used for our examples are rather elementary, but computational complexity becomes quite daunting otherwise. In each case we have been able to compute a significant fragment of the conformal tiling. The periodicity is realized geometrically by applying an associated Möbius transformation μ of the plane. The shading in each tiling image highlights nesting of the conformal core under three applications of μ and gives some sense of the geometric action. After looking at these you might consider the following Question, which was answered affirmatively by Cannon, Floyd, Kenyon, and Parry [9] for scaling in the regular pentagonal tiling.

Question. Are the Möbius period maps illustrated here necessarily algebraic? For instance, in Figure 8 the period map involves a scaling factor $\lambda < 1$. Is λ necessarily an algebraic real number?

Here are some further details regarding our concrete tiling examples.

Example 6.3. Figure 6 displays a periodic conformal tiling of period 1 for the “pentagonal” subdivision operator τ . The construction is based on a vertex core, and in this case that vertex is fixed under the cellular map carrying C to c , so the associated point in the tiling is a fixed point for the conformal period map μ ; we have placed it at the origin. There is an evident 3-fold symmetry in the global tiling, but this is incidental: period maps are not symmetries of an individual tiling, but rather symmetries in the tiling’s hierarchy. In this case, μ is loxodromic, a contraction by some factor $\lambda < 1$ followed by a counterclockwise rotation of $2\pi/3$.

We take a moment here to point out a subtlety within the definition of *expansion complex*. Let $K = K_\tau$ be the combinatorial tiling associated with Figure 6. However, now imagine the footprint schematic on the left of Figure 6, built with, say, five pentagons coming together on the central vertex instead of the current

three. Subdivide each and denote the resulting complex by F_2 , then define subsequent complexes F_n inductively by $F_{n+1} = \tau F_n, n \geq 2$. It is not difficult to see that the directed limit $\lim_{\rightarrow} F_m$ is a combinatorial tiling K' much like K except that the vertex at the center has degree 5 instead of degree 3. As with K , K' is a fixed point for $\hat{\tau}$. Thus we can build a combinatorial hierarchy $\{K'_n\}$ with all $K'_n \cong K'$ and the associated conformal tilings $\{\mathcal{T}_n = \mathcal{T}_{K'_n}\}$ form a conformal hierarchy. However—and here is our point— K' is not an expansion complex for τ . The definition posits an initial face $\Delta \in \mathfrak{T}$, but there is no face core which, for example, engulfs F_2 , since subdivision by τ will never produce an interior degree 5 vertex. The converse to Theorem 3.3 fails.

Example 6.4. Figure 7 is also a pentagonal tiling, this time of period 2. This tiling is again based on a vertex core, but the cellular map does not fix that vertex. In \mathcal{T} the eye readily picks out the approximate location of the fixed point of the conformal period map μ , and that point is not associated with any tile vertex. Note that there are no global symmetries in this tiling.

Example 6.5. Figure 8 is constructed from an edge core for the “snowball” subdivision rule. Reflective symmetry across the edge core in the footprint leads to a line of global reflective symmetry for the tiling \mathcal{T} . The fixed point for μ has been placed at the origin; μ has period 1 and is a pure contraction $\mu(z) = \lambda z$. Note that, other than the line of symmetry, there is not a single straight edge among the conformal tiles.

7. WHAT NEXT?

This final short section foreshadows developments in the next paper in this series, [4]. With regard to conformal hierarchies, we first will be directing the reader’s attention to images like that of Figure 5 for motivation. This image represents a conformal hierarchy manifest by the “lace” subdivision rule. The finest of the tiles there (those in black) are part of the conformal tiling $\mathcal{T} = \mathcal{T}_0$, while several aggregate tiling levels are shown in various colors. However, our interest attaches to the infinite sequence of forward subdivisions $\{\tau^n \mathcal{T} : n = 1, 2, \dots\}$, which have necessarily been left to the reader’s imagination.

As these forward subdivisions occur in situ, the tilings become progressively finer-grained. One can show that individual tiles go to zero in diameter, meaning that every point $z \in \mathbb{C}$ can be identified with a bi-infinite sequence $\{t_n\}$ of nested conformal tiles:

$$(\star) \quad z \longleftrightarrow \{\dots, t_{-2}, t_{-1}, t_0, t_1, t_2, \dots\}, \quad t_n \in \mathcal{T}_n, \quad t_{n+1} \subset t_n, \quad \{z\} = \bigcap_{n \in \mathbb{Z}} t_n.$$

This identification is essentially unique, ambiguity arising only for points z landing on an edge of some tiling \mathcal{T}_n , and these points form a set of Lebesgue measure zero.

We can now move considerations to the combinatorial side: each conformal tile t_n is associated with a combinatorial tile $\mathbf{k}_n \in K_n$, and thus with a rooted tiling (K_n, \mathbf{k}_n) in \mathbf{RC} . The result is a pseudo-geometric/combinatorial identification

$$(\star\star) \quad z \longleftrightarrow \{(K_n, \mathbf{k}_n)\} \subset \mathbf{RC}, \quad \mathbf{k}_n \in K_n, \quad \mathbf{k}_{n+1} \subset \mathbf{k}_n.$$

A handy example of these identifications is provided by the periodic conformal tilings of the previous section. Each footprint is a prescription for building the sequence $\{t_n\}$, and the associated point z is just the fixed point of the conformal

period map μ . We mentioned that these periodic cases are “exceptional”, and that is the case. However a slightly modified notion might well be considered “ubiquitous”. In studying (\star) and $(\star\star)$ in [4] we will introduce associated limiting processes, simultaneously *pointwise*, *combinatorial*, and *analytic*. Among other things, we will find a countable dense set of points $z \in \mathbb{C}$ whose limits are precisely our τ -periodic tilings.

The attraction of the identifications (\star) and $(\star\star)$ is that they link objects in three categories: points z in the plane, rooted conformal tilings, and rooted combinatorial tilings. Each setting has its own internal life: topologies, limits, accumulation points, morphisms, and so forth. With these, one begins to glimpse the interlocking structures that are so prominent (and attractive) in traditional tiling theory.

There are limitations, however, and we also begin to address those. The basic observation is that whenever τ manifests a combinatorial hierarchy $\{K_n\}$, there is an associated bi-infinite sequence $\{\mathcal{T}_n = \mathcal{T}_{K_n} : n \in \mathbb{Z}\}$ of conformal tilings. True, this qualifies as a “hierarchy” only when τ is conformal, and we have rightly extolled the virtues of this case, but we will uncover fascinating opportunities in other cases as well. For example, suppose τ is the much-studied “Penrose” rule. Since τ is not conformal, we lose (\star) , and thereby $(\star\star)$ as well, even though Penrose tilings already enjoy euclidean versions of such identifications. In a bid to reestablish these by our methods, we will push in [4] beyond *conformal* hierarchies, using *conformal aggregate* tilings and their limits to define yet another level of hierarchy. Due to results of [16], our program succeeds for the Penrose and other traditional euclidean substitution rules. But it also scoops in various non-conformal, non-traditional rules τ . We subsume, for example, the “fractal hierarchies” of Mayhook ([18] associated with the “twisted pentagonal” and other non-conformal subdivision rules.

As to whether we can formulate the notions mentioned at the beginning of §2, the *continuous hull* of a tiling, the action of the “translation” group, its *canonical transversal*—all seem to be tantalizingly close. We also inch closer to the beautiful themes in the work of Cannon, Floyd, and Parry. Stay tuned.

ACKNOWLEDGMENTS

We reiterate our gratitude to those acknowledged in [7]: Maria Ramirez-Solano and Jean Savinien for introducing us to traditional tiling, Natalie Frank, Chaim Goodman-Strauss, and Lorenzo Sadun, for valuable conversations, and Jim Cannon and Bill Floyd for sharing their perspective with us over many years. Especially Maria, thanks. Thanks also to an anonymous referee on an earlier version of this paper. *Last Comment:* The experiments behind this paper as well as its images are due to the software package **CirclePack**, available from the second author.

REFERENCES

- [1] Lars V. Ahlfors, *Lectures on quasiconformal mappings*, 2nd ed., University Lecture Series, vol. 38, American Mathematical Society, Providence, RI, 2006. With supplemental chapters by C. J. Earle, I. Kra, M. Shishikura and J. H. Hubbard. MR2241787
- [2] Jared E. Anderson and Ian F. Putnam, *Topological invariants for substitution tilings and their associated C^* -algebras*, Ergodic Theory Dynam. Systems **18** (1998), no. 3, 509–537, DOI 10.1017/S0143385798100457. MR1631708
- [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] Philip L. Bowers and Kenneth Stephenson, *Conformal tilings III: Pushing the limits*, in preparation.
- [5] Philip L. Bowers and Kenneth Stephenson, A “regular” pentagonal tiling of the plane, *Conformal Geometry and Dynamics* **1** (1997), 58–86.
- [6] Philip L. Bowers and Kenneth Stephenson, *Uniformizing dessins and Belyĭ maps via circle packing*, *Mem. Amer. Math. Soc.* **170** (2004), no. 805, xii+97, DOI 10.1090/memo/0805. MR2053391
- [7] Philip L. Bowers and Kenneth Stephenson, *Conformal tilings I: foundations, theory, and practice*, *Conform. Geom. Dyn.* **21** (2017), 1–63, DOI 10.1090/ecgd/304. MR3594281
- [8] James W. Cannon, *The combinatorial Riemann mapping theorem*, *Acta Math.* **173** (1994), no. 2, 155–234, DOI 10.1007/BF02398434. MR1301392
- [9] J. W. Cannon, W. J. Floyd, R. Kenyon, and W. R. Parry, *Constructing rational maps from subdivision rules*, *Conform. Geom. Dyn.* **7** (2003), 76–102, DOI 10.1090/S1088-4173-03-00082-1. MR1992038
- [10] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Expansion complexes for finite subdivision rules. I*, *Conform. Geom. Dyn.* **10** (2006), 63–99, DOI 10.1090/S1088-4173-06-00126-3. MR2218641
- [11] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Expansion complexes for finite subdivision rules. II*, *Conform. Geom. Dyn.* **10** (2006), 326–354, DOI 10.1090/S1088-4173-06-00127-5. MR2268483
- [12] A. Fletcher and V. Markovic, *Quasiconformal maps and Teichmüller theory*, Oxford Graduate Texts in Mathematics, vol. 11, Oxford University Press, Oxford, 2007. MR2269887
- [13] Natalie Priebe Frank and Lorenzo Sadun, *Fusion: a general framework for hierarchical tilings of \mathbb{R}^d* , *Geom. Dedicata* **171** (2014), 149–186, DOI 10.1007/s10711-013-9893-7. MR3226791
- [14] 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
- [15] Chaim Goodman-Strauss, *Matching rules and substitution tilings*, *Ann. of Math. (2)* **147** (1998), no. 1, 181–223, DOI 10.2307/120988. MR1609510
- [16] Richard Kenyon and Kenneth Stephenson, *Shape convergence for aggregate tiles in conformal tilings*, *Proc. Amer. Math. Soc.*, to appear.
- [17] 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
- [18] Dane Mayhook, *Conformal tilings & type*, ProQuest LLC, Ann Arbor, MI, 2016. Thesis (Ph.D.)–The Florida State University. MR3597681
- [19] 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
- [20] Charles Radin, *The pinwheel tilings of the plane*, *Ann. of Math. (2)* **139** (1994), no. 3, 661–702, DOI 10.2307/2118575. MR1283873
- [21] 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.
- [22] Alain M. Robert, *A course in p -adic analysis*, Graduate Texts in Mathematics, vol. 198, Springer-Verlag, New York, 2000. MR1760253
- [23] Lorenzo Sadun, *Topology of tiling spaces*, University Lecture Series, vol. 46, American Mathematical Society, Providence, RI, 2008. MR2446623
- [24] Boris Solomyak, *Dynamics of self-similar tilings*, *Ergodic Theory Dynam. Systems* **17** (1997), no. 3, 695–738, DOI 10.1017/S0143385797084988. MR1452190
- [25] B. Solomyak, *Nonperiodicity implies unique composition for self-similar translationally finite tilings*, *Discrete Comput. Geom.* **20** (1998), no. 2, 265–279, DOI 10.1007/PL00009386. MR1637896

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

Email address: `bowers@math.fsu.edu`

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

Email address: `kstephe2@utk.edu`