A HOMEOMORPHISM INVARIANT FOR SUBSTITUTION TILING SPACES

by

Nicholas Ormes[1,2], Charles Radin[1] * and Lorenzo Sadun[1,3] **

[1] Mathematics Department, University of Texas at Austin

[2] Mathematics Department, University of Connecticut

[3] Physics Department, Technion – Israel Institute of Technology

Abstract

We derive a homeomorphism invariant for those tiling spaces which are made by rather general substitution rules on polygonal tiles, including those tilings, like the pinwheel, which contain tiles in infinitely many orientations. The invariant is a quotient of Čech cohomology, is easily computed directly from the substitution rule, and distinguishes many examples, including most pinwheel-like tiling spaces. We also introduce a module structure on cohomology which is very convenient as well as of intuitive value.

July 2000

Subject Classification: 37B50, 52C23, 52C20

---

* Research supported in part by Texas ARP Grants 003658-152 and 003658-158

** Research supported in part by Texas ARP Grants 003658-152 and 003658-158 and NSF Grant DMS-9626698

1. Introduction.

In this paper we study the topology of substitution tiling spaces. Specifically, we are interested in ways in which the geometry of the substitution is reflected in the topology.

We follow a history of work where one studies the topology of spaces created with an underlying dynamical system. An early reference to the circle of ideas is Parry and Tuncel [PaT; chap. 4]. They show that the first cohomology group of the standard suspension space $\widetilde{X}$ for a $\mathbb{Z}$-subshift (X, T) is a useful invariant for topological conjugacy of (X, T). Here, $\widetilde{X} = (X \times [0, 1])/\sim$, where $(x, 1) \sim (Tx, 0)$.

Later, Herman, Putnam and Skau [HPS] associated a C*-algebra with general minimal $\mathbb{Z}$-subshifts via the crossed product construction, in analogy with the way von Neumann algebras had been associated previously with measureable dynamics. For these systems the K_0-group of the C*-algebra coincides with the first cohomology of the suspension space [HPS]. Giordano, Putnam and Skau [GPS] used the (ordered) K-groups of the C*-algebra to characterize orbit equivalence for minimal $\mathbb{Z}$-subshifts.

In a similar vein Connes [Con] used Penrose tilings to motivate ideas on noncommutative topology. Specifically, he considered the quotient X of the space of Penrose tilings by the equivalence relation of translation. Using the hierarchical structure of the Penrose tilings there is a natural homeomorphism between X and a quotient Y of sequences modulo another equivalence relation. As is common for equivalence relations coming from dynamics, the quotient topology of Y is trivial. And so Connes considers a (noncommutative) C*-algebra associated with the quotient Y. The K-theory of this C*-algebra is his substitute for the (trivial) quotient topology.

This was pushed one step further by Kellendonk and Putnam [Ke1-4, KeP, AnP] in which examples of tiling dynamical systems such as the Penrose system are considered and their (unordered) K-theory is worked out, obtaining for instance cohomology groups for the spaces of tilings.

We consider substitution tiling systems in $\mathbb{R}^d$, such as the Penrose tilings in $\mathbb{R}^2$, as dynamical systems with $\mathbb{R}^d$ (translation) actions. These are natural generalizations of the suspensions of $\mathbb{Z}^d$ (substitution) subshifts. (We will restrict attention to $d = 2$; while most of our methods carry over to higher dimensions, there are complications which we wish to avoid.) As noted above, suspensions are sometimes introduced to get invariants to distinguish between subshifts, insofar as the spaces on which the subshifts act are all the same — Cantor sets. For tilings the natural objects already have $\mathbb{R}^d$ actions and different tiling spaces need not be homeomorphic. In particular, the cohomology groups of such spaces are invariants of topological conjugacy. For 1-dimensional tiling systems, Barge and Diamond [BD] worked out complete homeomorphism invariants.

Although in principle it should be possible to work out the cohomology groups for tiling spaces, in practice it can be difficult since the number of different tiles in interesting models (especially after collaring — see below), which gives the number of generators for the chain groups, is usually large, even infinite for examples such as pinwheel tilings. New in this work is analysis of two aspects of the cohomology of substitution tiling spaces that are both easy to compute and informative about the topology of the spaces.

The first aspect is an order structure for the top dimensional Čech cohomology. For some history in a "noncommutative" setting see [GPS], but the idea is much older, appearing for instance in [SuW]. Fundamental to this paper as well is a representation of our tiling spaces as inverse limits, as discussed in [AnP], and in particular we use this to define a positive cone in cohomology. More specifically, we construct a homeomorphism from the top Čech cohomology group to the real numbers, and define the positive cone to be the preimage of the non-negative reals. The image of the map is closely related to the additive group of $\mathbb{Z}[1/\lambda]$, where λ is the area stretching factor of the substitution. We show that the positive cone is a homeomorphism invariant, and that one can therefore extract invariant topological information from the easily computable stretching factor λ.

We prove results of this form for two classes of tiling spaces. For "fixed-orientation" tiling spaces (denoted X_x below), we need the additional assumption that tiles only appear in a finite number of distinct orientations in each tiling. For "all-orientation" tiling spaces (denoted X_ϕ below), this assumption is not needed. Our main result is

Theorem 1. *Let X and Y be substitution tiling systems whose substitutions have area stretching factors λ and λ', respectively. Suppose that either i) X and Y have finite relative orientation groups and the fixed-orientation spaces X_x and Y_y are homeomorphic, or ii) the all-orientation spaces X_ϕ and Y_ϕ are homeomorphic. Then*

(1) If λ is an integer so is λ', and $\mathbb{Z}[1/\lambda] = \mathbb{Z}[1/\lambda']$ as subsets of $\mathbb{R}$. In particular, λ and λ' have the same prime factors, although not necessarily with the same multiplicities.

(2) If λ is not an integer then λ and λ' are irrational, and $\mathbb{Q}[\lambda] = \mathbb{Q}[\lambda']$ as subsets of $\mathbb{R}$.

(We will see in examples below that this distinguishes topologically between tiling spaces, such as that of the pinwheel and the (2,3)-pinwheel, which had not previously been known to be distinguishable even in the stronger sense of topological conjugacy, that is, including the natural $\mathbb{R}^2$ action of translations on the spaces.)

To prove this theorem we show that the order structure, although defined in terms of the (not necessarily invariant) inverse limit structure, is in fact invariant. The relations between λ and λ' then follow by algebraic arguments.

To describe the second aspect of our analysis we must discuss the unusual rotational properties of substitution tilings such as those of Penrose. No Penrose tiling is invariant under rotation by $2\pi/10$ about any point, but every translation invariant Borel probability measure on the space of Penrose tilings (X_x or X_ϕ) is invariant under rotation by $2\pi/10$. We say $\mathbb{Z}_{10}$ is the "relative orientation group" of the Penrose tilings. This symmetry has been the source for most of the explosion of work on such tilings, in particular their use in modelling quasicrystals [Ra2]. (We have already investigated the use of rotational symmetry as a conjugacy invariant [RaS].)

As we will see below, each element of the cyclic rotation group $\mathbb{Z}_{10}$ gives an automorphism of the space X_x of Penrose tilings. As a result, the group $\mathbb{Z}_{10}$, and by extension its group ring $\mathbb{Z}[t]/(t^{10}-1)$, acts naturally on the (Čech) cohomology groups of X_x. This gives $H^*(X_x)$ the structure of a module over the group ring, and this module structure is quite useful. In particular, for any realization of X_x as the inverse limit of simplicial complexes, as in [AnP], the (co)chain groups are themselves modules, and the boundary and substitution maps are equivariant. This greatly simplifies the calculations. For Penrose tiles for example [AnP], there are 4 types of tiles, each appearing in 10 different orientations. Without using the rotational symmetry the action of substitution on 2-chains would be given by a 40×40 matrix. Using the rotational symmetry we express it instead as a 4×4 matrix with entries in the group ring of $\mathbb{Z}_{10}$, which can be analyzed one irreducible representation at a time. The substitutions on 0- and 1-chains, and the boundary maps, are handled similarly. The result is a streamlined calculation that yields information beyond that obtained by older methods. Anderson and Putnam [AnP] computed $H^0(X_x)=\mathbb{Z}$, $H^1(X_x)=\mathbb{Z}^5$ and $H^2(X_x)=\mathbb{Z}^8$. We note that, as modules, a finite-index additive subgroup of $H^1(X_x)$ equals $\mathbb{Z}[t]/(t-1)\oplus\mathbb{Z}[t]/(t^4-t^3+t^2-t+1)$ while a finite index subgroup of $H^2(X_x)$ equals $(\mathbb{Z}[t]/(t-1))^2 \oplus (\mathbb{Z}[t]/(t+1))^2 \oplus \mathbb{Z}[t]/(t^4-t^3+t^2-t+1)$, where rotation by $2\pi/10$ is multiplication by t.

2. Tiling spaces and their associated inverse limits.

Before defining substitution tiling spaces in general we present some examples whose topological properties we will later consider.

A "chair" tiling of the plane, Fig. 1, can be made as follows. Consider the L-shaped tile of Fig. 2. Divide this tile (also called a "tile of level 0") into four pieces as in Fig. 3 and rescale by a linear factor of 2 so that each piece is the same size as the original. This yields a collection of 4 tiles that we call a "tile of level 1". Subdividing each of these tiles and rescaling gives a collection of 16 tiles that we call a tile of level 2. Repeating the process n times gives a collection of 4^n tiles — a tile of level n. A "chair" tiling is a tiling of the plane with the property that every finite subcollection of tiles is congruent to a subset of a tile of some level. A chair tiling has only one type of tile, appearing in 4 different orientations. (For chair tilings, and in fact very generally for tilings made by such substitution rules, there are "matching rules" which provide a different method of construction [Goo, Ra2].)

Somewhat more complicated are the "Penrose" tilings, Fig. 4, which have 4 different types of triangular tiles, each appearing in 10 different orientations. The 4 tiles, and the rule for subdividing them, are shown in Fig. 5. (These tiles are not the familiar kites and darts. However, the space of substitution tilings developed from these 4 triangles is homeomorphic to the space of Penrose kite and dart tilings). Here the linear stretching factor is the golden mean $\tau=(1+\sqrt{5})/2$, and the area

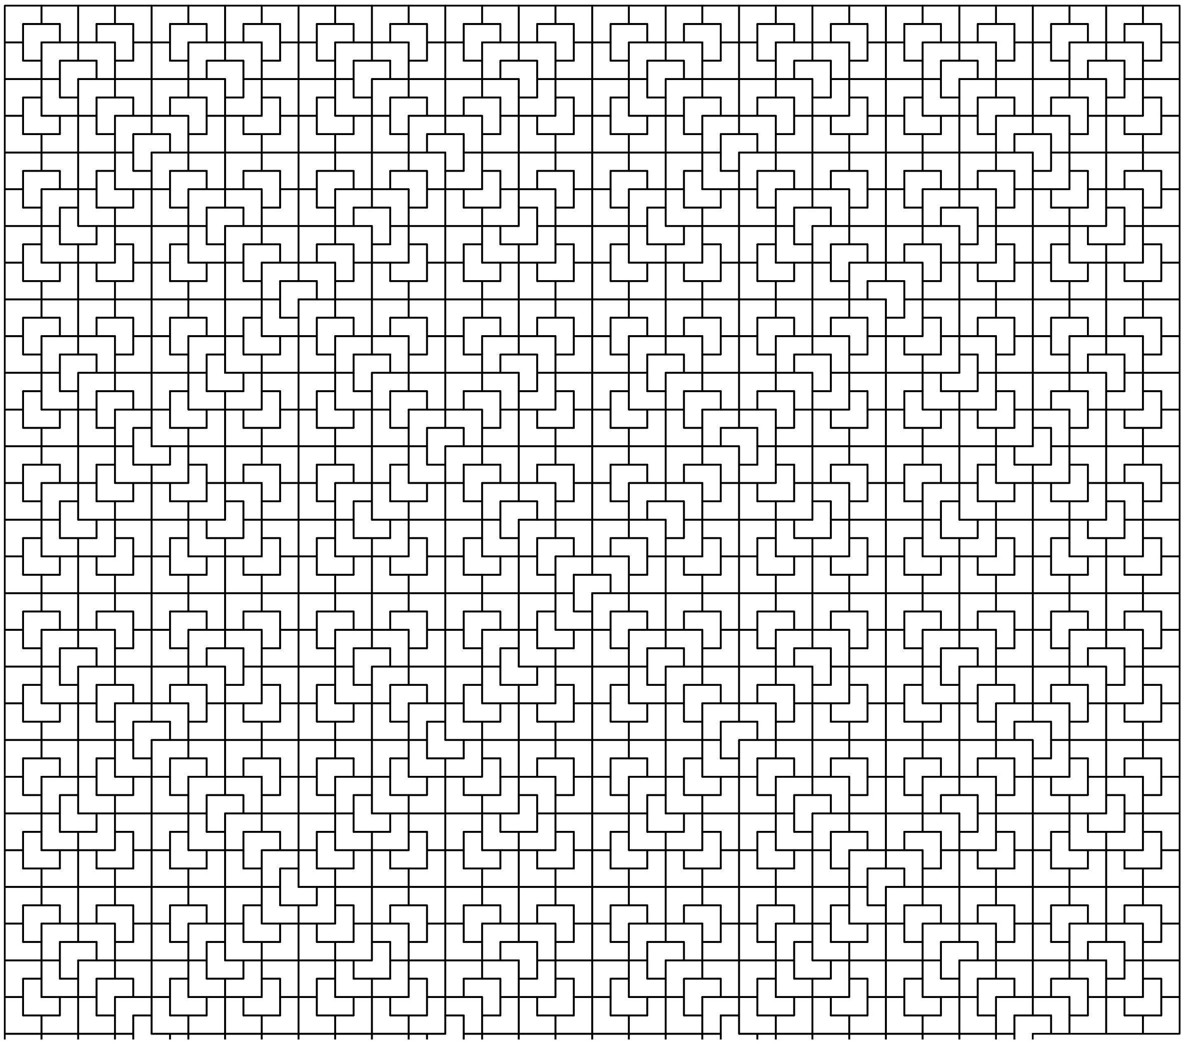

Figure 1. A chair tiling

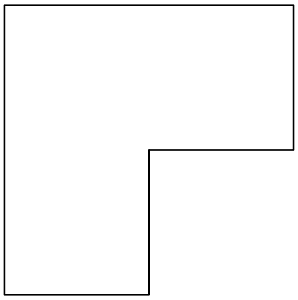

Figure 2. The chair tile

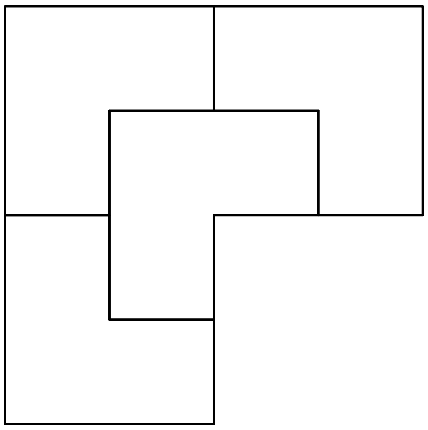

Figure 3. The chair substitution

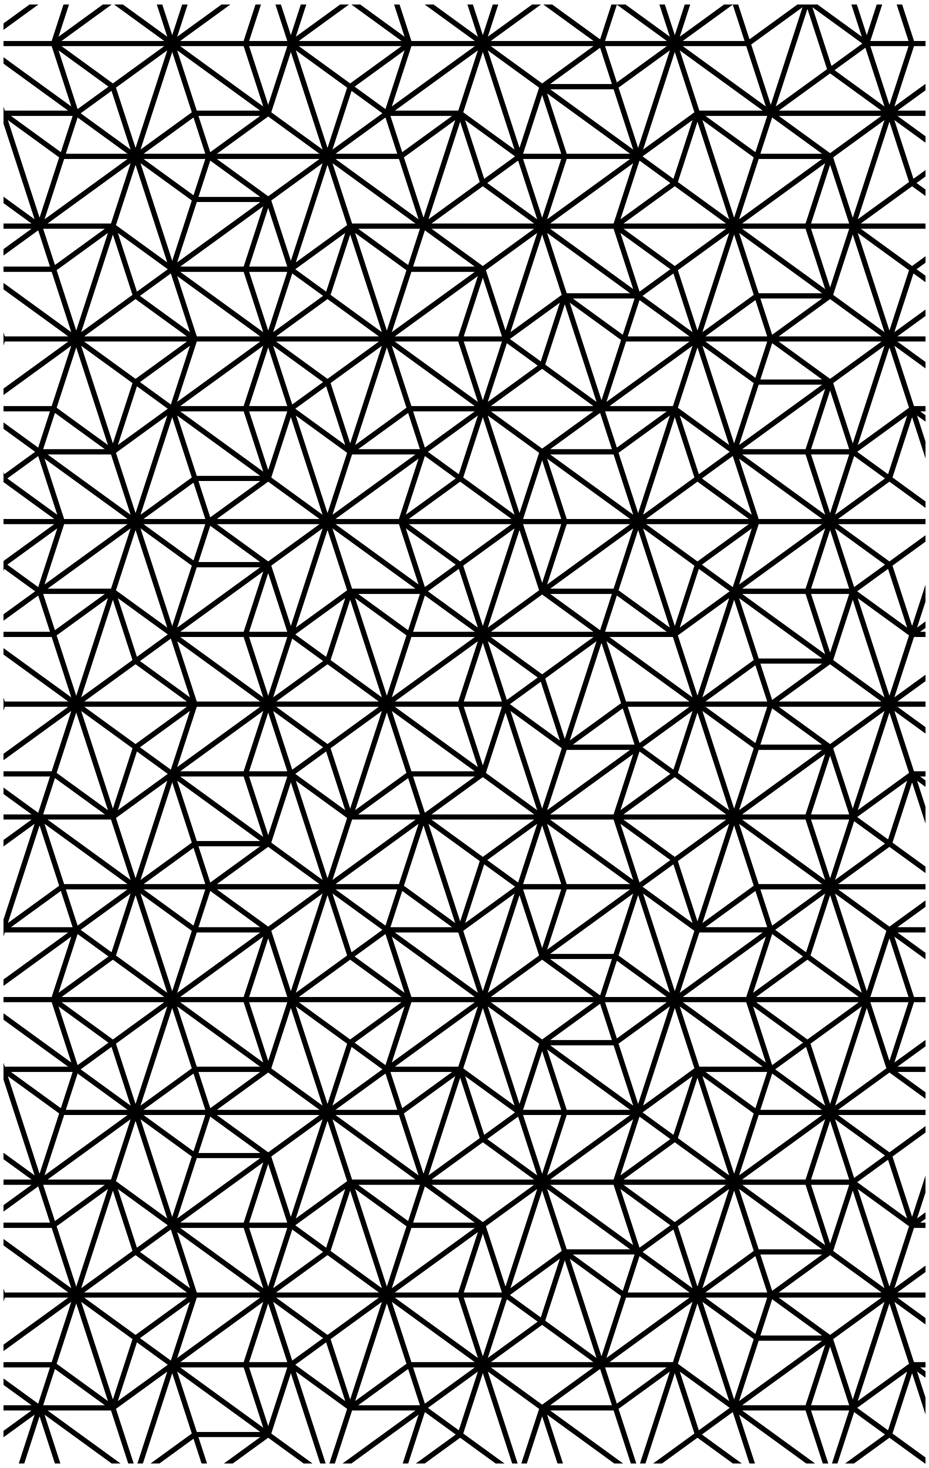

Figure 4. A Penrose tiling

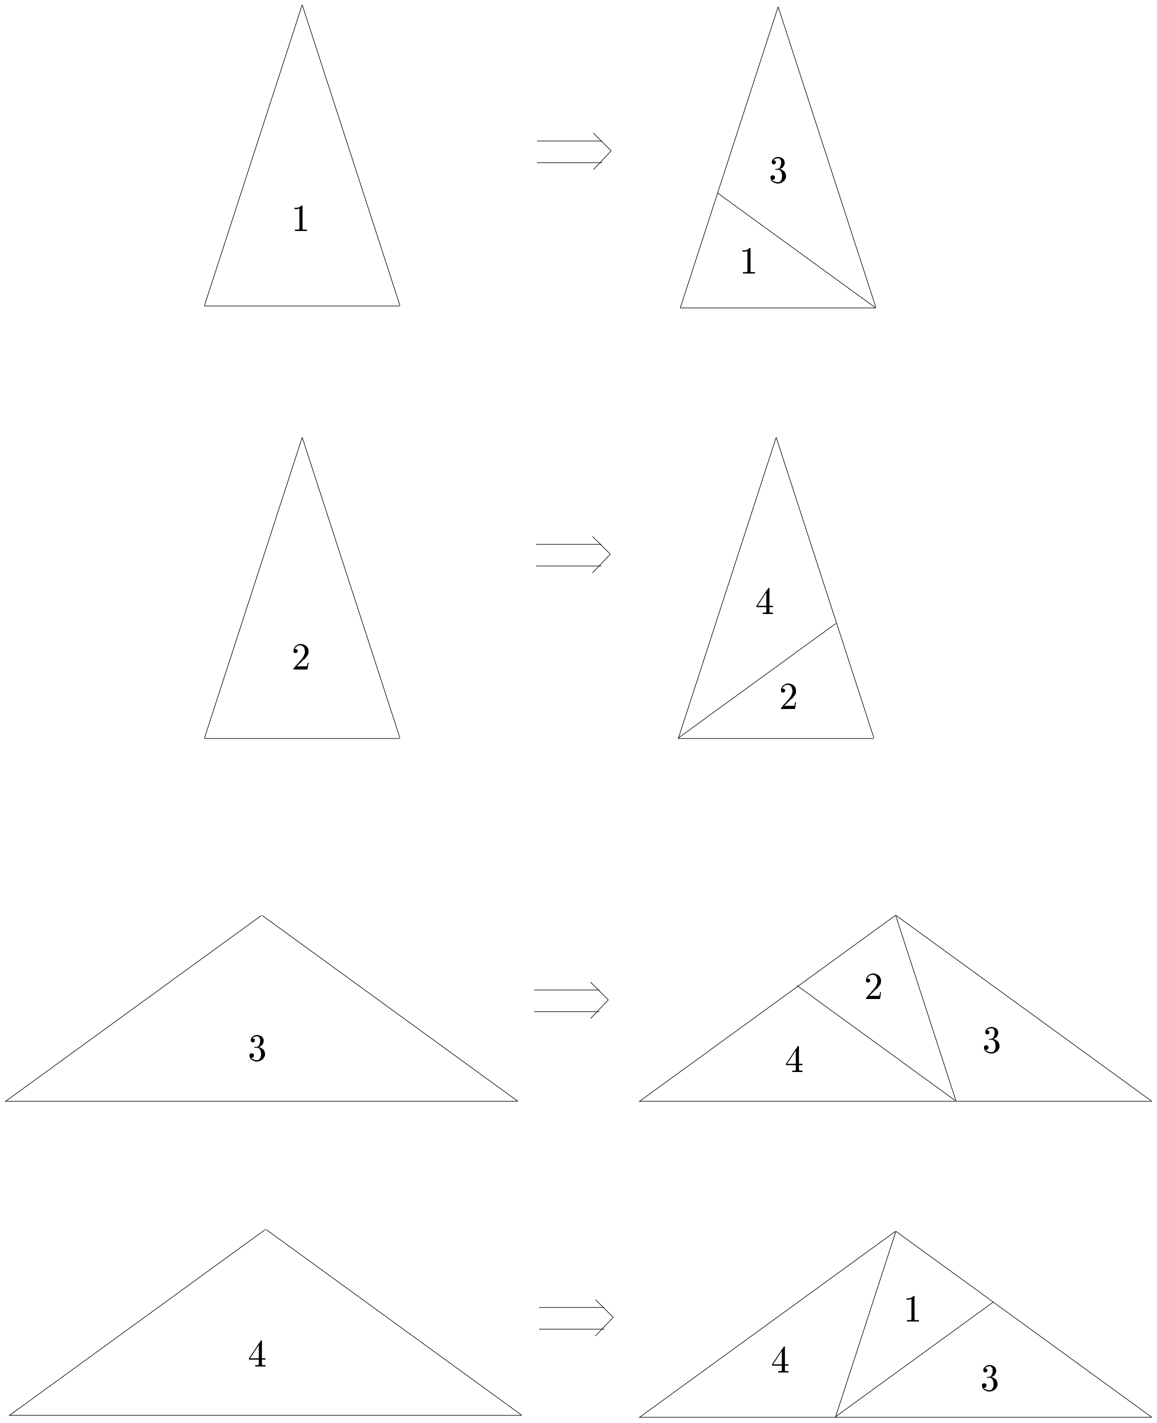

Figure 5. The Penrose substitution

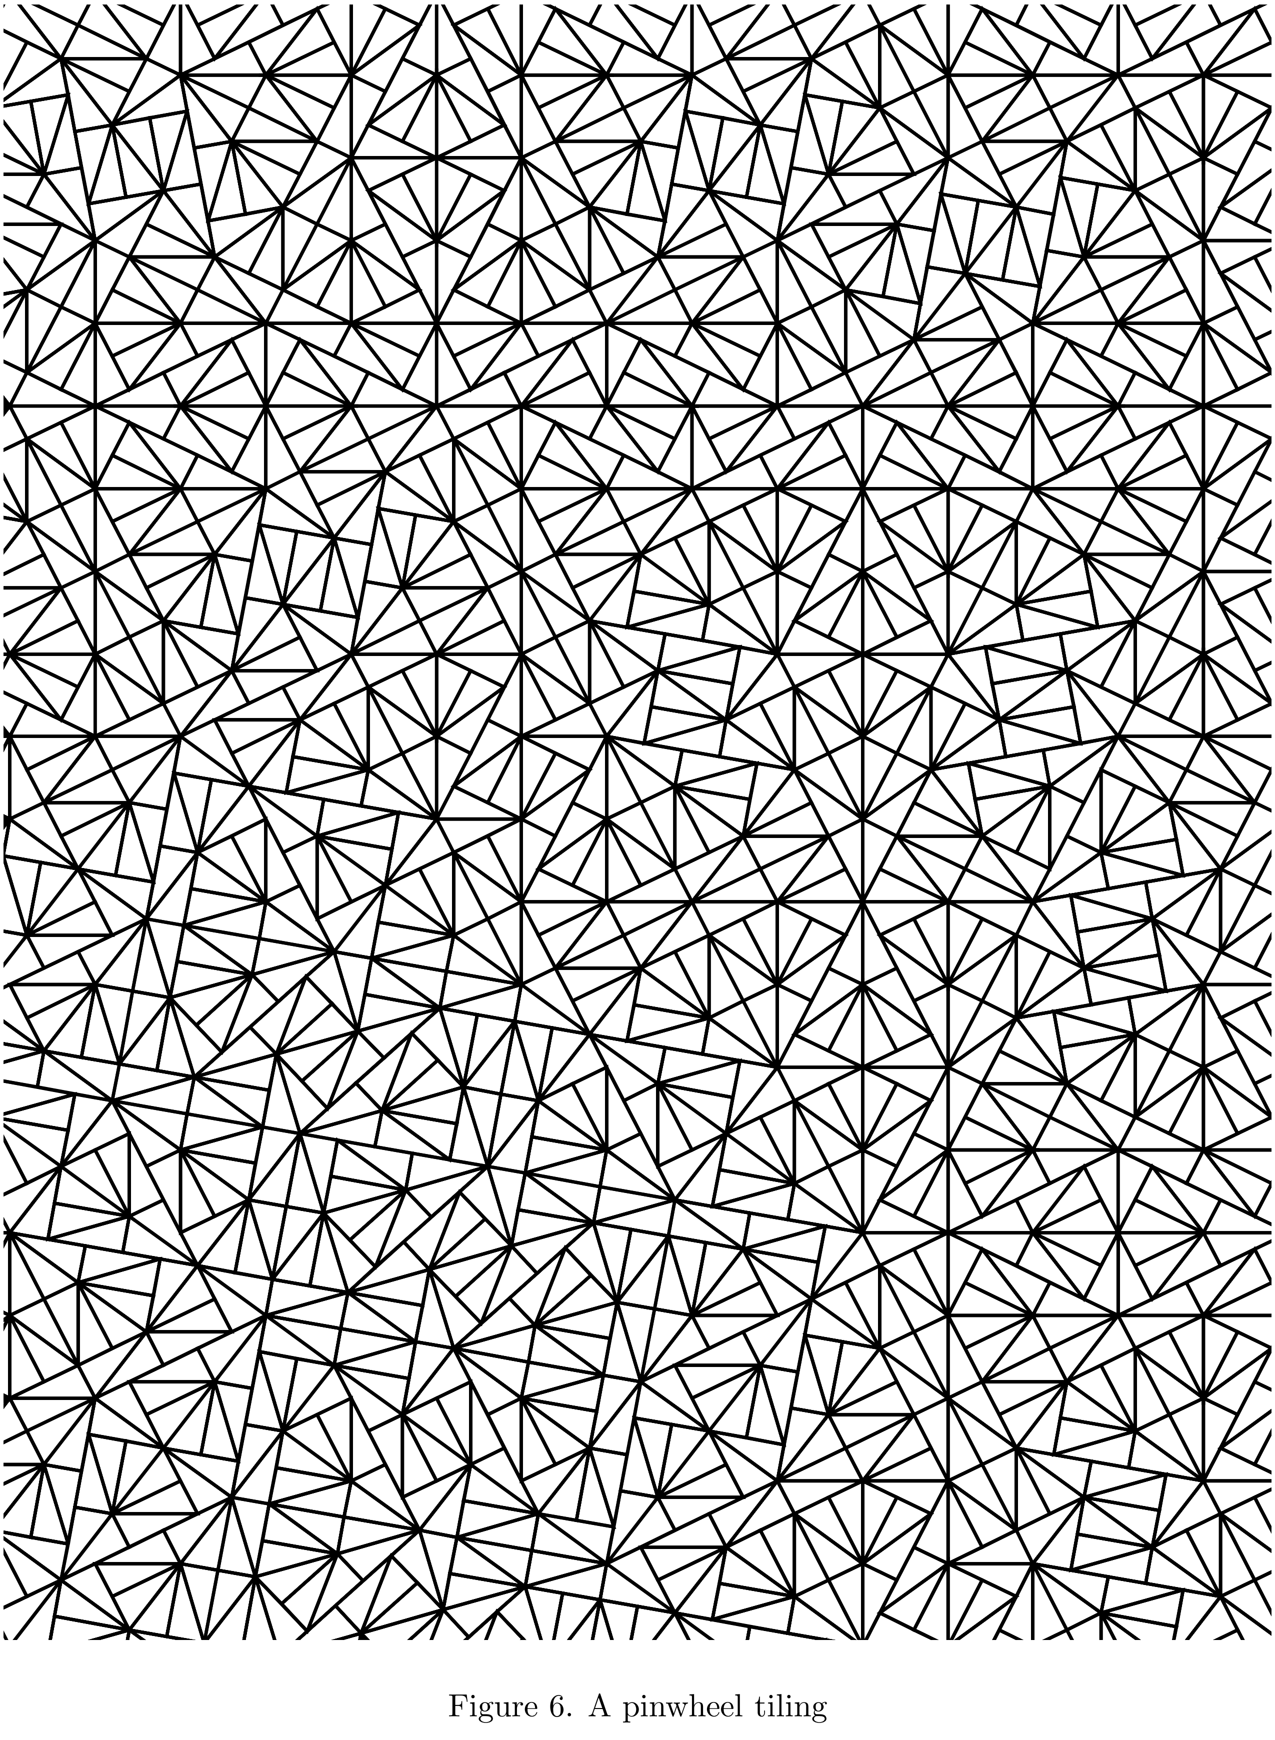

Figure 6. A pinwheel tiling

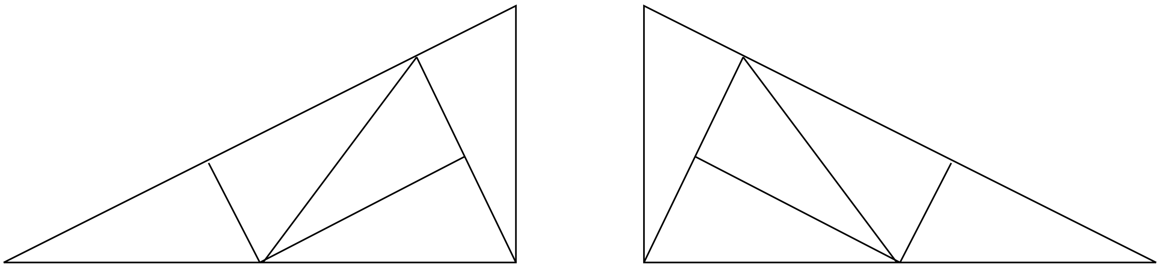

Figure 7. The substitution for pinwheel tilings

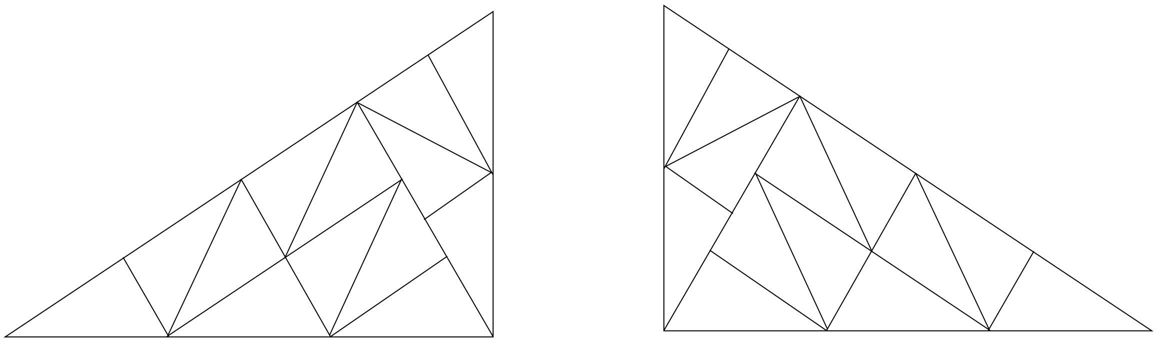

Figure 8. The substitution for (2,3) pinwheel tilings

stretching factor is $\tau^2 = (3+\sqrt{5})/2$.

A "pinwheel" tiling of the plane, Fig. 6 [Ra1], has tiles that appear in an infinite number of orientations. The two basic tiles, a 1-2-$\sqrt{5}$ right triangle and its mirror image, are shown in Fig. 7, with their substitution rule.

Notice that at the center of a tile of level 1 there is a tile of level 0 similar to the level 1 tile but rotated by an angle $\alpha = \tan^{-1}(1/2)$. Similarly, the center tile of a tile of level n is rotated by $n\alpha$ relative to the tile. Since α is an irrational multiple of π, we see, using the fact that within a tile of level 2 there is a tile of level 0 similar to the level 2 tile, that this rotation never stops, and each tiling contains tiles in an infinite number of distinct orientations.

Finally, we consider the "(2,3)-pinwheel" tilings defined by the substitution of Fig. 8, whose tiles are 2-3-$\sqrt{13}$ right triangles. (Such variants of the pinwheel are easily constructed for any integral legs $m < n$.) Like the ordinary pinwheel, variant pinwheel tilings also necessarily have tiles in an infinite number of distinct orientations. In fact the relative orientation groups for all pinwheel tilings are algebraically isomorphic. Theorem 1 shows that the tiling spaces for the pinwheel and (2,3)-pinwheel are not homeomorphic.

In all the above cases it is easy to construct explicit examples of tilings. Pick a tile to include the origin of the plane. Embed this tile in a tile of level 1 (there are several ways to do this). Embed that tile of level 1 in a tile of level 2, embed that in a tile of level 3, and so on. The union of these tiles will cover an infinite region, typically – though not necessarily – the entire plane.

We now give a general definition of substitution tiling systems. Let $\mathcal{A}$ be a nonempty finite collection of polygons in the plane. Let $X(\mathcal{A})$ be the set of all tilings of the plane by congruent copies, which we call tiles, of the elements of (the "alphabet") $\mathcal{A}$. We label the "types" of tiles by the elements of $\mathcal{A}$. We endow $X(\mathcal{A})$ with the metric

$$d(x,x') \equiv \sup_n \frac{1}{n} m_H[B_n(\partial x), B_n(\partial x')], \qquad 1)$$

where $B_n(\partial x)$ denotes the intersection of two sets: the closed disk B_n of radius n centered at the origin of the plane and the union ∂x of the boundaries of all tiles in x. The Hausdorff metric m_H is defined as follows. Given two compact subsets P and Q of $\mathbb{R}^2$, $m_H[P,Q] = \max\{\tilde{d}(P,Q), \tilde{d}(P,Q)\}$, where

$$\tilde{d}(P,Q) = \sup_{p\in P} \inf_{q\in Q} ||p-q||, \qquad 2)$$

with $||w||$ denoting the usual Euclidean norm of w.

Under this metric two tilings are close if they agree, up to a small Euclidean motion, on a large disk around the origin. The converse is also true for tiling systems whose tiles meet full edge to full edge (as we require) – closeness implies agreement, up to small Euclidean motion, on a large disk around the origin. Although the metric d depends on the location of the origin, the topology induced by d is translation invariant. A sequence of tilings converges in the metric d if and only if its restriction to every compact subset of $\mathbb{R}^2$ converges in m_H. It is not hard to show [RaW] that

$X(\mathcal{A})$ is compact and that the natural action of the connected Euclidean group $\mathcal{G}_E$ on $X(\mathcal{A})$, $(g, x) \in \mathcal{G}_E \times X(\mathcal{A}) \longrightarrow g[x] \in X(\mathcal{A})$, is continuous.

A "substitution tiling space" is a closed subset $X_\phi \subset X(\mathcal{A})$ satisfying some additional conditions. To understand these conditions we first need the notion of "patches". A patch is a (finite or infinite) subset of an element $x \in X(\mathcal{A})$; the set of all patches for a given alphabet will be denoted by W. Next we need, as for the above examples, an auxiliary "substitution function" ϕ, a map from W to W, with the following properties:

(1) There is some constant $c(\phi) > 1$ such that, for any $g \in \mathcal{G}_E$ and $x \in X$, $\phi(g[x]) = \phi(g)[x]$, where $\phi(g)$ is the conjugate of g by the similarity of Euclidean space consisting of stretching about the origin by $c(\phi)$.
(2) For each tile $T \in \mathcal{A}$ and for each $n \geq 1$, the union of the tiles in $\phi^n T$ is congruent to $[c(\phi)]^n T$, and these tiles meet full edge to full edge.
(3) For each tile $T \in \mathcal{A}$, ϕT contains at least one tile of each type.
(4) For each tile $T \in \mathcal{A}$ there is $n_T \geq 1$ such that $\phi^{n_T} a$ contains a tile of the same type as T and parallel to it.
(5) No tile $T \in \mathcal{A}$ has a nontrivial rotational symmetry.

Condition (2) is significant. It is satisfied by the pinwheel tilings only if we add additional vertices at midpoints of the legs of length 2, creating boundaries of 4 edges. (A similar alteration is needed for the chair tilings.) The tile of level n, $\phi^n T$, will be said to be of "type" T.

Condition (5), by contrast, is technical, and does not significantly limit the scope of this work. If a tile does have n-fold rotational symmetry, we can recover condition (5) by breaking the tile into n congruent but asymmetric pieces, in a manner consistent with the subdivision rules.

Definition 1. For a given alphabet $\mathcal{A}$ of polygons and substitution function ϕ the "substitution tiling space" is the compact subspace $X_\phi \subset X(\mathcal{A})$ of those tilings x such that every finite subpatch of x is congruent to a subpatch of $\phi^n(T)$ for some $n > 0$ and $T \in \mathcal{A}$. We assume x can be decomposed in one and only one way by tiles of level n, for any fixed $n \geq 0$ [Sol].

The above definition gives spaces of tilings in which the tiles can appear in arbitrary orientations. Although any *fixed* chair tiling, for example, has tiles in only four orientations, the space of *all* chair tilings also contains rotated versions of that tiling, and so contains tilings in which chairs appear in any orientation.

In many cases, especially when working with tilings whose tiles have only a finite number of orientations per tiling, it is convenient (and customary) to allow tiles only a minimal set of orientations. Instead of using the metric d, we define a metric d' such that two tilings are close if (and only if) they agree on a large disk around the origin *up to a small translation*, rather than up to a small Euclidean motion. Specifically,

$$d'(x, x') = \inf\{\frac{R+S}{1+R+S} : B_{1+1/S}(g[\partial x]) = B_{1+1/S}(h[\partial x'])\} \qquad 2')$$

where $g, h \in \mathbb{R}^2$, $R > 0$, $S > 0$, $|g| < R$, $|h| < R$ and $d'(x, x') \equiv 1$ if the defining set is empty. We pick any one tiling $x \in X_\phi$, and define $X_x \subset X_\phi$ to be $\{x' \in X_\phi : d'(x, x') < 1\}$. For tiling spaces with only a finite number of orientations per tiling, X_x consists of those tilings in which tiles only appear parallel to tiles in x. We call X_x a "fixed-orientation" tiling space.

Finally, we consider the quotient of X_ϕ by rotations about the origin and denote this space X_0. For any fixed tiling x, X_0 is also the quotient of X_x by the "relative orientation group" G of the tiling x, defined as follows: G is generated by the rotations that take a tile in x and turn it parallel to another tile of the same type in x. In [RaS] it was shown that this group is the same for all tilings $x \in X_\phi$. The topologies of these three spaces are closely related. The space X_ϕ is a circle bundle over X_0 with some singular fibers; these fibers correspond to tilings with a discrete rotational symmetry about the origin.

Next we show how to give a substitution tiling space the structure of an inverse limit. We begin by doing this first for a fixed-orientation space, and we begin in particular with the example of the chair. Fix a chair tiling x, and in it a tile T. Each of the other tiles in x are rotated with respect to T by an element of the cyclic rotation group $G_{ch} = \mathbb{Z}_4$. We construct an (uncollared) chair complex Σ^{uc} as follows. We start with the topological disjoint sum Σ' of the elements of $\{g[T] : g \in G_{ch}\}$. We then identify those edges in Σ' which "meet" somewhere in x, defining Σ^{uc}. We consider a countable number of copies Σ^{uc}_j of Σ^{uc}, indexed by the non-negative integers. We think of the tiles in Σ^{uc}_j as referring to level j tiles. The substitution ϕ (whereby we think of each tile of level $j+1$ as the union of four tiles of level j) then defines a map from Σ^{uc}_{j+1} to Σ^{uc}_j for any $j \geq 0$, and allows us to define the inverse limit $\underset{\leftarrow}{\Sigma}^{uc} \equiv \underset{\longleftarrow}{\lim} \Sigma^{uc}_j$.

There is a map from the tiling space X_x to $\underset{\leftarrow}{\Sigma}^{uc}$ defined, on $x \in X_x$, by determining precisely how the origin in x sits in each successive tile of higher and higher level. For many but not all substitution systems, this map is neither 1-1 nor onto. A point in the inverse limit precisely describes the hierarchy of tilies containing the origin. In some instances this union of tiles may not cover the entire plane; if not, it may be that there is more than one way to extend the tiling to cover the entire plane (in which case the map is not 1-1), or it may be that there is no way to extend it (in which case the map is not onto). For example, in a chair tiling the tile containing the origin might sit in the upper left corner of the tile of level 1 of Fig. 2, which sits in the same position in a corresponding tile of level 2, and so on. The union of these tiles covers only a quadrant, and this can be extended to a complete chair tiling of the plane in more than one way.

This problem is absent in substitution systems (such as the Penrose tilings [AnP]) which "force the border", that is, for which there is some integer N such that, in every tiling, for every tile $\phi^N T$ of level N the (level 0) tiles which abut this tile are completely determined by the type T. From this it follows that, for any M, the tiles of level M that abut a given tile of level $N+M$ are also determined. An infinite-level tile, therefore, determines a tiling of the entire plane, and is consistent with a tiling of the entire plane, even if it does not itself cover the entire plane. The map from the space of tilings to the inverse limit is therefore 1-1 and onto.

It is easy to prove, but very useful, that for any substitution system we can extend the substitution in a simple way to a larger set of "collared" tiles producing a tiling space naturally homeomorphic to the original but now forcing the border. The new set of tiles consists of multiple marked versions of the original tiles, one for each way the original tile can be surrounded by tiles in a tiling. (For the chair one needs 14 versions of each of the original tiles; the pinwheel requires over 50 versions.) Any (fixed-orientation) substitution tiling space can thus be modelled as an inverse limit.

If the original tiles force the border, we let $\Sigma_x = \Sigma^{uc}$. If the original tiles do not force the border, let Σ_x denote the complex constructed from collared tiles. Either way, there is a natural substitution map from Σ_x to itself, and $X_x = \underleftarrow{\Sigma_x}$

The other two kinds of tiling spaces may also be constructed as inverse limits. To construct X_0 we consider the type *but not orientation* of the tiles of various level that contain the origin. In the complex Σ_0 the basic cells are tiles T, where each type appears in only one orientation. As before, if it is possible for a tile T_1 to meet a tile T_2 then their common edge is identified, and it may be necessary to consider collared tiles. As before the substitution ϕ maps Σ_0 to itself and we consider the inverse limit. A point in the inverse limit is a consistent instruction: the origin sits at such-and-such a point in such-and-such a tile, which sits in a particular way in a tile of level 1, which sits in a particular way in a tile of level 2, and so on. Since we are using collared tiles (or tiles that force the border), this prescription defines a unique tiling of the entire plane, up to an overall rotation, i.e. a point in X_0.

The construction of Σ_ϕ is similar, except that the basic cells are products $S^1 \times T$, where the T are as in the construction of Σ_0. If a tile T_1, rotated by an angle α (with respect to the standard orientation of $\mathcal{A}$), can meet a tile T_2, rotated by an angle β, along a common edge e, then we identify $(\alpha, e) \subset S^1 \times T_1$ with $(\beta, e) \subset S^1 \times T_2$. It then follows that for any angle α', $(\alpha', e) \subset S^1 \times T_1$ is identified with $(\alpha' + \beta - \alpha, e) \subset S^1 \times T_2$. Once again, substitution maps Σ_ϕ to itself, and the inverse limit corresponds to instructions on how to build a tiling around the origin. Only now the instructions include information on how to orient each tile, and so defines a tiling uniquely, i.e. a point in X_ϕ.

If the relative orientation group G is finite then Σ_ϕ, Σ_x and Σ_0 are all compact simplicial complexes. This is clear for Σ_x and Σ_0, as the cells, finite in number, are simplices that meet along common edges. The cells in Σ_ϕ are not simplices — they have the topology of S^1 — but can be divided into n contractible pieces, meeting along common edges, where n is the order of G.

If the relative orientation group is infinite then Σ_0 is a compact simplicial complex but Σ_x is not compact (as it contains an infinite number of cells). The space Σ_ϕ, while a compact CW complex, is not necessarily simplicial.

Finally, these complexes are closely related. Just as X_0 is the quotient of X_x by G and the quotient of X_ϕ by S^1, Σ_0 is the quotient of Σ_x by G and the quotient of Σ_ϕ by S^1.

3. Finite relative orientation group and P-positivity.

Here we define the "P-positive cone" of the top dimensional Čech cohomology group for substitution tiling spaces. In this section we assume our system has a finite relative orientation group, so the simplicial complex Σ_x is compact. We also restrict ourselves to fixed-orientation spaces X_x. Later, we will investigate the role of rotations, and consider order structures on the top cohomologies of X_ϕ and X_0 for tiling spaces with arbitrary relative orientation groups. (We continue to assume our tilings are 2-dimensional, although this construction applies equally well to other dimensions.)

As discussed above, the space X_x is an inverse limit of compact simplicial complexes Σ_x with the substitution (self-)map ϕ between them. One can therefore compute the Čech cohomology of X_x by taking a direct limit of the Čech = simplicial cohomology of Σ_x under the map ϕ^*.

Recall that the direct limit of an Abelian group under a map M is the disjoint union of an infinite number of copies of the group, indexed by the non-negative integers, modulo the equivalence relation

$$(g,k) \sim (Mg, k+1), \qquad 3)$$

for all g and all k, where (g,k) denotes the element g in the k-th copy of the group. The direct limit of the cohomology groups $H^2(\Sigma_x)$ can be obtained by taking the direct limit of the cochain groups $C^2(\Sigma_x)$ then moding out by the image of the coboundary map δ from the direct limit of $C^1(\Sigma_x)$. (The proof of this is simple diagram chasing.)

Since Σ_x is a finite simplicial complex, $C_2(\Sigma_x)$, the space of simplicial chains, is generated by the tiles themselves and is isomorphic to $\mathbb{Z}^n$ where n is the number of distinct tiles in Σ_x. To compute the cohomology, each tile must be given an orientation. We define this orientation in the following way: Select an orientation for $\mathbb{R}^2$ and select a tiling $x_0 \in X_x$. The translations comprise a free $\mathbb{R}^2$-action on x_0. Each tile that appears in x_0 thus inherits an orientation from the orientation on $\mathbb{R}^2$. All such orientations are consistent and are independent of x_0. The projection map from X_x to Σ_x defines an orientation on tiles in Σ_x. From here on, we refer to this as the "positive orientation" of the tiles in Σ_x.

If all tiles at all levels are given the positive orientation, the induced map $\phi_* : C_2(\Sigma_x) \to C_2(\Sigma_x)$ is an $n \times n$ matrix with only non-negative entries. (The i,j entry in the matrix is the number of copies of tile type i which appear in a level 1 tile of type j.) By assumption the matrix is primitive, that is, some power has all of its entries strictly positive. The cochain group $C^2(\Sigma_x)$ is generated by the duals to the tiles. Relative to this basis, the pullback map ϕ^* is described by the transpose of ϕ_*. Similar constructions apply to C^1, except that the elements of C_1 and C^1 are not naturally oriented, so the substitution matrix on C^1 may have negative entries.

Let $\mathcal{C}$ denote the direct limit of the cochains $C^2(\Sigma_x)$ under the map ϕ^*, and define $\mathcal{C}_+$ to be the semigroup of elements $(c,k) \in \mathcal{C}$ which have the property that $(\phi^*)^m c \in (\mathbb{Z}_+)^n$ for some m, where n is the number of distinct tiles in Σ_x. It is clear that $\mathcal{C}_+ - \mathcal{C}_+ = \mathcal{C}$. Since the matrix ϕ^* contains only non-negative entries, it is also clear that $\mathcal{C}_+ + \mathcal{C}_+ \subset \mathcal{C}_+$, $\mathcal{C}_+ \cap -\mathcal{C}_+ = \{0\}$ and thus $(\mathcal{C}, \mathcal{C}_+)$ is an ordered group (in fact, a dimension group [Ell]).

Let $H = H^2(X_x)$ denote the top Čech cohomology group for X_x. We wish to define a positive cone on H by

$$H_+ = \{[c] \mid c \in \mathcal{C}_+\} \qquad 4)$$

where elements of $\mathcal{C}$ are equivalent if their difference is in $\operatorname{Im}\delta$. To see that (H, H_+) is an ordered group, we introduce a function $\mu : \mathcal{C} \to \mathbb{R}$.

Lemma 1. *There is a homomorphism $\mu : \mathcal{C} \to \mathbb{R}$ such that $\mathcal{C}_+ - \{0\} = \mu^{-1}(\mathbb{R}_+ - \{0\})$.*

Proof: Since the matrix ϕ^* is primitive it has a Perron eigenvalue $\lambda > 0$ (of multiplicity 1). This is the area stretching factor of the substitution. (To see this, note that a vector whose entries are the areas of the various tiles is an eigenvector of the substitution matrix. Since all entries of this vector are positive this is the Perron eigenvector, and the corresponding eigenvalue, which is the area stretching factor, must be the Perron eigenvalue. This eigenvalue is the square of the linear stretching factor $c(\phi)$.) Let r denote a left eigenvector of ϕ^*. Define the map $\mu : \mathcal{C} \to \mathbb{R}$ by $(v, k) \mapsto \lambda^{-k} r v$.

The map is well-defined since

$$\mu(\phi^* v, k+1) = \lambda^{-k-1} r \phi^* v = \lambda^{-k} r v. \qquad 5)$$

Suppose $(v, k) \in \mathcal{C}_+ - \{0\}$ then $(\phi^*)^m v \in (\mathbb{Z}_+)^n$ for some m. Since ϕ^* is primitive, we may choose m such that all entries of $(\phi^*)^m v$ are positive. Since all entries of r are non-negative,

$$\mu(v, k) = \lambda^{-k} r v = \lambda^{-k-m} r (\phi^*)^m v > 0. \qquad 6)$$

To prove the reverse direction, take a basis of $\mathbb{C}^n$ which includes w_0, the right eigenvector for λ, and $(n-1)$ vectors $\{w_1, w_2, \ldots, w_{n-1}\}$, each $w_i \in \operatorname{Ker}(I\lambda_i - \phi^*)^l$ for some l and $\lambda_i \neq \lambda$. We can and do choose w_0 with all its entries positive. Now $r w_i = 0$ since

$$0 = r(I\lambda_i - \phi^*)^l w_i = (\lambda_i - \lambda)^l r w_i. \qquad 7)$$

Therefore if $\mu(v, k) > 0$ then writing v in the basis $\{w_0, w_1, \ldots, w_{n-1}\}$ the coefficient c_0 of w_0 must be positive. Since λ exceeds the modulus of all other eigenvalues, as $l \to \infty$, we have $||\lambda^{-l}(\phi^*)^l v - c_0 w_0|| \to 0$. This implies $(\phi^*)^l v \in (\mathbb{Z}_+)^n$ for some l. ∎

Next we show that the map $\mu : \mathcal{C} \to \mathbb{R}$ induces a map from H to $\mathbb{R}$; since $H = \mathcal{C}/\text{Im } \delta$, the following will show that this map is well-defined.

Lemma 2. $\operatorname{Im}\delta \subseteq \operatorname{Ker}\mu$.

Proof: Let e be an edge and let $f_e \in C^1(\Sigma_x)$ be the 1-cochain which assigns a value 1 to the edge e and 0 to all other edges. Then $\delta(f_e)$ is a 2-cochain which acts on the tiles as follows: let T be an mth level tile. Then $\delta(f_e)(T)$ counts the number of occurences of e on the boundary T with a ± signs depending upon the orientation. In particular, $|\delta(f_e)(T)|$ is less than or equal to the number of occurences of e on the boundary of the level m tile T.

Now assume $\mu(\delta(f_e)) > 0$. Then there is an l such that for all for any level 0 tile T, $(\phi^*)^l\delta(f_e)(T) > 0$. Furthermore, for fixed T this quantity scales like λ^l for sufficiently large l. This means that for level l tiles there are at least $O(\lambda^l)$ occurences of e on the boundary. However, λ is the area stretching factor, so there can only be $O(\lambda^{l/2})$ edges on the boundary of a level l tile (since the perimeter scales like $\lambda^{l/2}$ and there is a minimum length to each edge). This is a contradiction.

A similar argument shows that $\mu(\delta(f_e))$ cannot be negative. Therefore the image of δ is a subgroup of the kernel of μ and $\mathcal{C}_+ \cap \operatorname{Im}\delta = \{0\}$. ■

Abusing notation, define $\mu : H \to \mathbb{R}$ by $\mu([x]) = \mu(x)$. Letting H_+ be the pre-image under μ of $\mathbb{R}_+$, (H, H_+) is an ordered group since

(1) $H_+ + H_+ \subset H_+$
(2) $H_+ \cap -H_+ = \{0\}$
(3) $H_+ - H_+ = H$ (this follows from the similar property of $\mathcal{C}_+$).

4. Some relations between X_x, X_0 and X_ϕ.

Before we prove the invariance of P-positivity and extend to arbitrary relative orientation groups, we need to understand the role of rotations in tilings. Let G be the relative orientation group of a tiling and let R be the associated group ring.

Finite relative orientation groups. We begin by assuming G is finite. As previously indicated, G, and therefore R, acts naturally on the simplicial (co)chain groups $C^i(\Sigma_x)$ and $C_i(\Sigma_x)$, giving them the structure of modules over R. A finite index subgroup of each module can be written as a direct sum of irreducible representations of G. The substitution maps are equivariant and so are module homomorphisms, mapping each irreducible representation to itself. The (co)boundary maps are also equivariant, and so map a representation in C^i (or C_i) to the corresponding representation in C^{i+1} (or C_{i-1}). Thus the integer (co)homology groups, and in particular the top cohomology, are themselves finite extensions of finite quotients of direct sums of irreducible representations.

The need for finite extensions and finite quotients comes about as follows. When a compact group acts on $\mathbb{R}^n$, $\mathbb{R}^n$ splits up as the direct sum of irreducible representations of that group. However, the same observation does not apply to $\mathbb{Z}^n$. For example, if $\mathbb{Z}_2$ acts on $\mathbb{Z}^2$ by permuting the two coordinates, the irreducible representations are all multiples of $(1,1)$ and all multiples of $(1,-1)$. The direct sum of these two representations gives all elements (a,b) with $a+b$ even — an index 2 subgroup of $\mathbb{Z}^2$. Similarly, finite index subgroups of the kernels and images of the (co)boundary maps are direct sums of representations of G. Computing (co)homology one representation at a time is equivalent to taking a finite index subgroup of the kernel of δ and modding out by a finite index subgroup of the image of δ. Correcting for our taking too small a numerator means doing a finite extension. Correcting for the denominator involves taking a finite quotient.

Since the substitution map is equivariant, rotating an eigenvector gives another eigenvector with the same eigenvalue. Since the Perron eigenvectors are unique, with purely positive entries, they must be rotationally invariant and therefore belong to the trivial representation. Since the left Perron eigenvector is invariant, μ of a cochain is also rotationally invariant. But that implies that μ acting on each nontrivial representation is identically zero.

For example in the uncollared chair complex there is just one tile, appearing in four different orientations, and $C^2 \sim R = \mathbb{Z}[t]/(t^4-1)$ where t represents rotation by $\pi/2$. The substitution map is given by the 1×1 matrix $2+t+t^3$. In the trivial representation we set $t=1$, so our matrix has Perron eigenvalue 4, with the Perron eigenvector being any multiple of an invariant element of C^2, *i.e.* of $(1+t+t^2+t^3)$.

This suggests a streamlined computation of the Perron eigenvalue and Perron eigenvector. Instead of working on all of Σ_x, work on the simplicial complex $\Sigma_0 = \Sigma_x/G$, whose (co)chain groups are precisely the trivial representations of G within the (co)chain groups of Σ_x. $C^2(\Sigma_x) \otimes_{\mathbb{Z}} \mathbb{R}$ equals $C^2(\Sigma_0) \otimes_{\mathbb{Z}} \mathbb{R}$ plus the other representations that appear in $C^2(\Sigma_x)$. The Perron eigenvalue on $C^2(\Sigma_x) \otimes \mathbb{R}$ is precisely the Perron eigenvalue on $C^2(\Sigma_0) \otimes \mathbb{R}$, and the (left and right) Perron eigenvectors on $C^2(\Sigma_x) \otimes \mathbb{R}$ come from Perron eigenvectors on $C^2(\Sigma_0) \otimes \mathbb{R}$.

Since the (co)chain groups of Σ_0 are isomorphic to the invariant parts of the (co)chain groups of Σ_x, and since the (co)boundary and substitution maps are equivariant, we have

Theorem 2. *Up to finite extensions, the (co)homology of Σ_0 is equal to the rotationally invariant part of the (co)homology of Σ_x, and the cohomology of X_0 is isomorphic to the rotationally invariant part of the cohomology of X_x.*

We cannot make any statements about the homologies of X_0 and X_x because homology, unlike cohomology, does not behave well under inverse limits.

Now X_ϕ is a circle fibration over X_0, with singular fibers corresponding to tilings that are symmetric about the origin. (These singular fibers have finite multiplicity, since any given tiling can only admit discrete rotational symmetry about the origin.) The existence of these singular fibers means that computing $\pi_1(X_\phi)$, $H^1(X_\phi)$ and $H^2(X_\phi)$ from corresponding data on X_0 can be complicated. However, computing the top cohomology $H^3(X_\phi)$ is quite easy. From the spectral sequence of the fibration, it follows immediately that $H^3(X_\phi)$ is isomorphic to $H^2(X_0)$ [BoT]. From this, and from the fact that Σ_ϕ is a circle fibration over Σ_0, we have

Corollary 1. *Up to finite extensions, the top cohomology of Σ_ϕ is equal to the rotationally invariant part of the top cohomology of Σ_x. Up to finite extensions, the top cohomology of X_ϕ is isomorphic to the rotationally invariant part of the top cohomology of X_x.*

Note also that the top-dimensional substitution matrix for Σ_ϕ is identical to the top-dimensional substitution matrix for Σ_0 — the S^1 factor just comes along for the ride. We may therefore compute our order structure on whichever of the three spaces Σ_x, Σ_0 or Σ_ϕ is most convenient.

To illustrate these principles, we compute the cohomology of the Penrose tilings. Here the relative orientation group is $\mathbb{Z}_{10}$, with group ring $R = \mathbb{Z}[t]/(t^{10}-1)$, where t represents rotation by $2\pi/10$. There are 4 irreducible representations of G, corresponding to the factorization $t^{10}-1 = (t-1)(t+1)(t^4+t^3+t^2+t+1)(t^4-t^3+t^2-t+1)$. The first two are one dimensional and t acts by multiplication by ± 1. The other two representations are 4 dimensional and correspond algebraically to $\mathbb{Z}[t]/(t^4+t^3+t^2+t+1)$ and $\mathbb{Z}[t]/(t^4-t^3+t^2-t+1)$.

In this tiling, there are 4 types of tiles, whose geometry and substitutions are shown in Fig. 9. Similarly, there are four kinds of edges, each in 10 orientations, and four kinds of vertices, which we label $\alpha, \beta, \gamma, \delta$, with the relations $\alpha = t\beta$, $\beta = t\alpha$, $\gamma = t\delta$ and $\delta = t\gamma$. The chain groups of Σ_x are therefore $C_2 = C_1 = R^4$ and $C_0 = (\mathbb{Z}[t]/(t^2-1))^2$. The boundary maps are given by the matrices

$$\partial_1 = \begin{pmatrix} 1-t & -1 & -t & -1 \\ 0 & 1 & 1 & t \end{pmatrix}; \qquad \partial_2 = \begin{pmatrix} -1 & t & t^4 & -t^7 \\ -1 & t^9 & -t & t^8 \\ 1 & -t^5 & 0 & 0 \\ 0 & 0 & 1 & -t^5 \end{pmatrix}, \tag*{8)}$$

where we have taken (α, γ) as our basis for C_0. The coboundary maps are given by the transposes of these matrices, with t replaced throughout by t^{-1}. The ranks of these matrices are easily computed by row reduction, one representation at a time:

(1) If $t = 1$, then ∂_2^* has rank 2 and ∂_1^* has rank 1, so $H^0(\Sigma_x)$ and $H^1(\Sigma_x)$ contain one copy of this representation while $H^2(\Sigma_x)$ contains two.

(2) If $t = -1$, then ∂_2^* and ∂_1^* have rank two, so $H^2(\Sigma_x)$ contains two copies of this representation while $H^1(\Sigma_x)$ and $H^0(\Sigma_x)$ contain none.

(3) If $t^4+t^3+t^2+t+1 = 0$, then ∂_1^* is zero, since C_0 (also C^0) does not contain this representation. ∂_2^* has rank 4, so all cohomologies are trivial.

(4) If $t^4-t^3+t^2-t+1 = 0$, then again ∂_1^* is trivial. Now ∂_2^* has rank 3, so $H^1(\Sigma_x)$ and $H^2(\Sigma_x)$ each contain this representation once.

Combining these results, we have that, up to finite extensions, $H^0(\Sigma_x) = \mathbb{Z}$, $H^1(\Sigma_x) = \mathbb{Z}[t]/(t-1) \oplus \mathbb{Z}[t]/(t^4-t^3+t^2-t+1)$ and $H^2(\Sigma_x) = (\mathbb{Z}[t]/(t-1))^2 \oplus (\mathbb{Z}[t]/(t+1))^2 \oplus \mathbb{Z}[t]/(t^4-t^3+t^2-t+1)$.

Next we look at the substitution matrices, which in dimensions 2 and 1 are

$$\phi_2 = \begin{pmatrix} t^7 & 0 & 0 & t^4 \\ 0 & t^3 & t^6 & 0 \\ t^3 & 0 & t^4 & 1 \\ 0 & t^7 & 1 & t^6 \end{pmatrix}; \qquad \phi_1 = \begin{pmatrix} 0 & 0 & 0 & t^8 \\ t^4 & 0 & -t^7 & 0 \\ -t^7 & 0 & 0 & 0 \\ 0 & -t^3 & 0 & -t^5 \end{pmatrix} \tag*{9)}$$

These matrices (and their pullbacks) are easily seen to be invertible in all representations, and so induce isomorphisms in cohomology. Therefore $H^*(X_x) = H^*(\Sigma_x)$.

We can also compute the cohomology of Σ_0 and X_0. The complex Σ_0 has four faces, four edges and two vertices. The boundary maps can be read off from 8) by replacing t throughout by 1, and the substitution maps are similarly obtained from 9). In other words, the computation for Σ_0 and X_0 is precisely the same as the $t = 1$ part of the computation for Σ_x and X_x, with the result that $H^2(X_0) = H^2(\Sigma_0) = \mathbb{Z}^2$

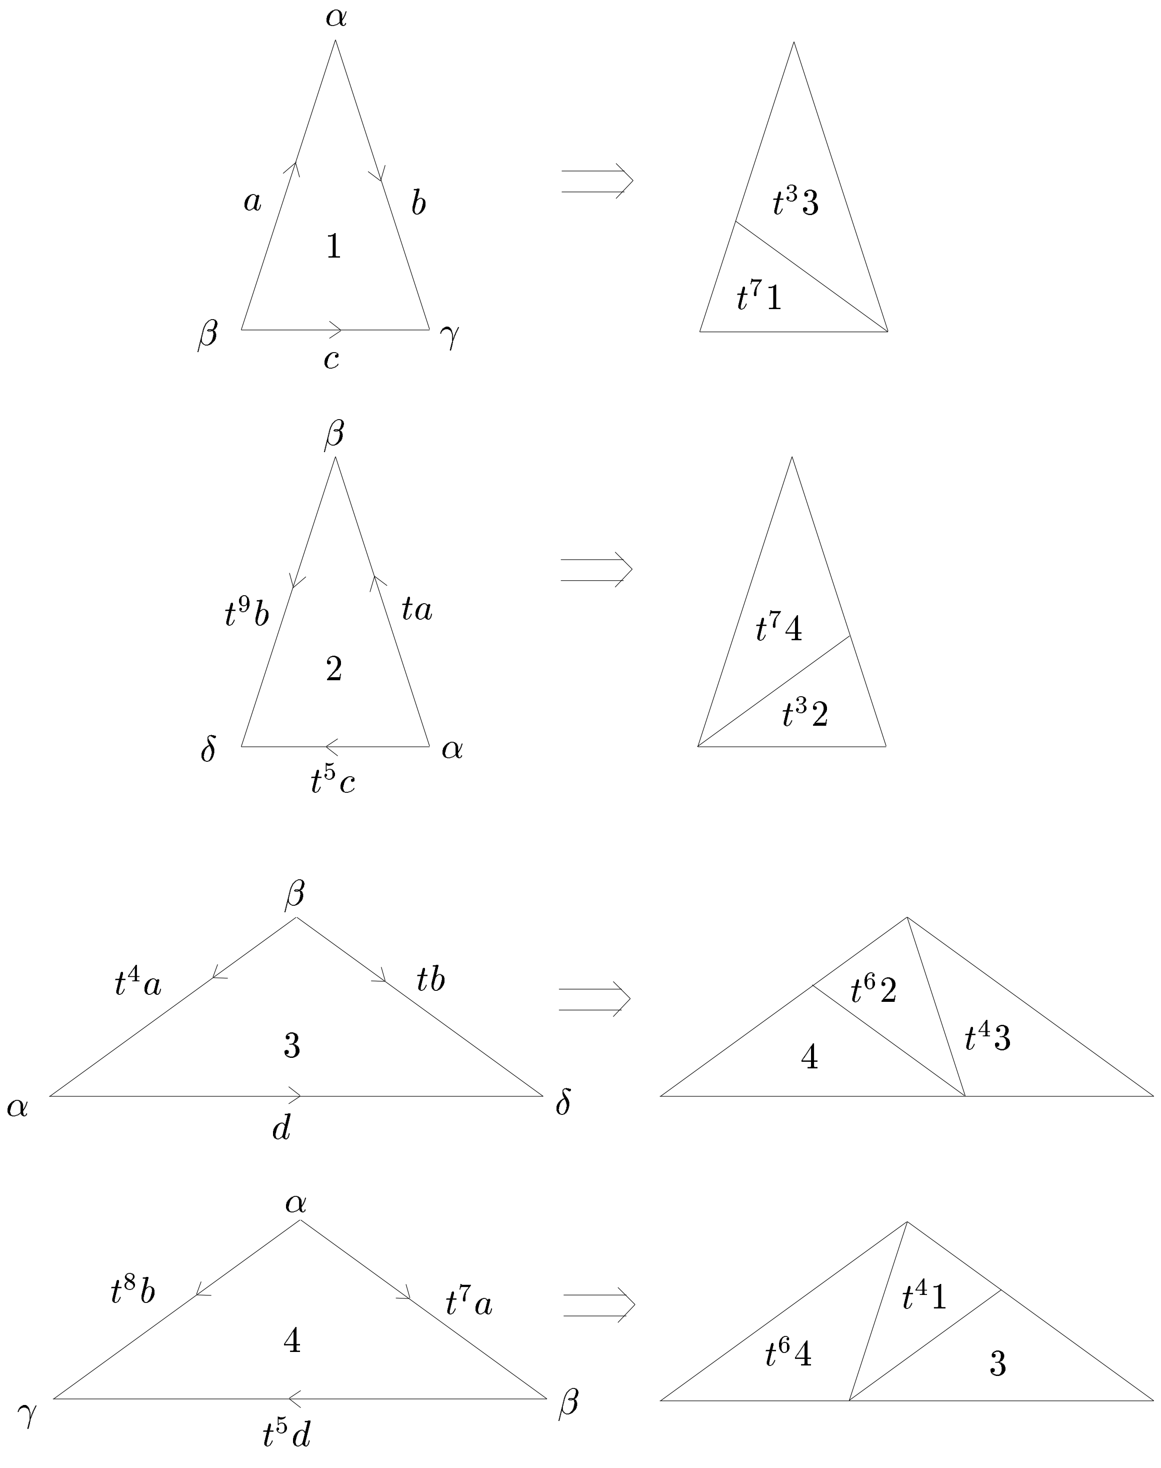

Figure 9. The Penrose substitution in detail

and $H^1(X_0) = H^0(X_0) = H^1(\Sigma_0) = H^0(\Sigma_0) = \mathbb{Z}$. We therefore also obtain $H^3(X_\phi) = \mathbb{Z}^2$.

Arbitrary relative orientation groups. We have seen that the spaces Σ_x and X_x are generally more complicated than the spaces Σ_0 and X_0, or Σ_ϕ and X_ϕ. In the case of finite relative orientation group this is a virtue, as the cohomology of X_x gives us information beyond that of X_ϕ. In the case of infinite relative orientation group, however, the topologies of Σ_x and X_x are too rich, and we shall see that the cohomologies (when nontrivial) are infinitely generated. This makes them too big to handle with the formalism of primitive matrices and Perron eigenvectors. Instead, we are forced to work on Σ_0 and X_0 (or Σ_ϕ and X_ϕ) where the cohomology groups are finitely generated and we can define an order structure. This structure then pulls back to an ordering of the rotationally-invariant elements of the top cohomology of X_x.

If a tiling space has an infinite relative orientation group, then the simplicial complex Σ_x is noncompact, but Σ_0 is still compact, and has a finite substitution matrix. In that case, the top homology of Σ_x, if nontrivial, is infinitely generated. This is seen as follows. Let C be a boundaryless 2-chain on Σ_x. Since C is a finite collection of tiles, only a finite number of orientations appear in C. By applying an appropriate rotation g, we obtain another boundaryless 2-chain $C' = g(C)$ whose tiles appear in a completely different set of orientations from those of C. Similarly, $C'' = g(C')$ is different from C'. If the element g has infinite order, we can generate an infinite collection of boundaryless chains, all linearly independent in $C_2(\Sigma_x)$. Thus $H_2(\Sigma_x)$ is infinitely generated. A similar argument, involving the dual to a tile in C, shows that H^2 is also infinitely generated. Furthermore, since the dual to a direct sum is a direct product, H^2 contains elements represented by infinite sums of duals of tiles.

Finding the left Perron eigenvector is also a problem. The left Perron eigenvector of ϕ^* is the same as a right Perron eigenvector of ϕ. This must be rotationally invariant, and is therefore a sum of an infinite number of tiles. Pairing this infinite sum with an element of C^2, which itself has an infinite number of terms, gives an infinite sum of real numbers that need not converge. Thus the μ map on $H^2(\Sigma_x)$, or on $C^2(\Sigma_x)$, is not well defined.

Working on Σ_0, however, we have no problem. The space Σ_0 is a compact simplicial complex, with finitely generated (co)chain complexes, and finitely generated (co)homology. The substitution matrix ϕ^* is finite and primitive, and has a Perron eigenvalue and Perron eigenvector exactly as before. On $H^2(\Sigma_0)$, and on the limit $H^2(X_0)$ ($= H^3(X_\phi)$) the order structure defined above works exactly as in the case of finite relative orientation group.

5. The invariance of positivity.

Having defined the P-positive cone in the top Čech cohomology group of our tiling spaces, we wish to show that this positive cone is invariant (up to an overall sign) under homeomorphisms. (This is not immediate, since the inverse limit structure that we used to define the positive cone is *not* invariant.)

In general, the Čech cohomology of a space is the direct limit of Čech cohomologies of open covers where the open covers generate the topology of the space (see [BoT] for details). In our situation we begin by considering a triangulation of Σ_n and the open cover of Σ_n given by "open stars" of the vertices of the triangulation. By the open star ($\mathrm{star}(v)$) of a vertex v we mean the collection of all simplices that touch v, but not including the opposite boundary simplex. We associate to this open cover a simplicial complex called the "nerve" of the open cover. The nerve of an open cover is constructed by including one vertex for each open set in the cover, then including a k-simplex between any k-vertices such that the intersection of the corresponding k sets in the open cover is nonempty. The Čech cohomology of the cover is defined as the simplicial cohomology of the nerve of that cover. Since every vertex of the triangulation lies in exactly one open star and every k-simplex lies in exactly k open stars, the nerve of the open star cover of Σ_n is identical to the triangulation of Σ_n as a simplicial complex. Thus the Čech cohomology of the open star cover of Σ_n is *canonically* identified with the simplicial cohomology of the triangulation of Σ_n. (This argument shows that simplicial and Čech cohomology agree for all simplicial complexes.)

Consider triangulations of Σ_n, generated recursively, so that each triangle in the triangulation of Σ_n maps by ϕ to a single triangle in the first barycentric subdivision of the triangulation of Σ_{n-1}. This triangulation of Σ_n corresponds to an open cover of Σ_n, which pulls back by the natural projection $\pi_n : X_x \to \Sigma_n$ to give an open cover $\mathcal{A}_n$ of X_x. That is, an open set in $\mathcal{A}_n$ is the set of tilings in which the origin sits in a specified neighborhood in a tile of level n. These neighborhoods, as defined by the triangulations of Σ_n, are increasingly fine as $n \to \infty$ (even after accounting for the expansion due to ϕ), so the sequence of open covers $\mathcal{A}_n$ generates the topology of X_x.

The Čech cohomology of X_x is the direct limit of the Čech cohomologies of the covers $\mathcal{A}_n$ (hence the direct limit of the simplicial cohomologies of Σ_n) with a bonding map defined as follows. From their construction, $\mathcal{A}_{n+1}$ refines $\mathcal{A}_n$. Each element of $\mathcal{A}_n$ (resp. $\mathcal{A}_{n+1}$) is a set of the form $\pi_n^{-1}(\mathrm{star}(v))$ where v is a vertex in the triangulation of Σ_n (resp. Σ_{n+1}). We define a vertex map from Σ_{n+1} to Σ_n by mapping each vertex u in the triangulation of Σ_{n+1} to a vertex v in the triangulation of Σ_n where $\pi_{n+1}^{-1}(\mathrm{star}(u)) \subset \pi_n^{-1}(\mathrm{star}(v))$. After extending this vertex map linearly to the interior of higher level simplices, the result is our bonding map $f_n : \Sigma_{n+1} \to \Sigma_n$. Although the vertex map, and therefore f_n, may not be uniquely defined, the induced chain map f_n^* on cohomology is always the same [BoT].

To define positivity recall that we used an orientation in $\mathbb{R}^2$ to define an orientation of all 2-simplices in Σ_n. Assume α and β are two such oriented 2-simplicies which both occur within the same tile $T \in \Sigma_n$. Then the duals to α and β are cohomologous. Any positive sum of these cohomology classes was declared to be P-positive in section 3.

To show that this notion of P-positivity is invariant we must consider the isomorphism of Čech cohomology induced by a homeomorphism h. Let $\{\Sigma_n\}$ and

$\{\Sigma'_n\}$ be sequences of complexes corresponding to the tiling spaces and let $\{\mathcal{A}_n\}$ and $\{\mathcal{A}'_n\}$ be sequences of open covers, of X_x and X'_x, resp., corresponding to triangulations of the complexes. Let $\mathcal{B}_n = h^{-1}(\mathcal{A}'_n)$. From their construction, for every n there is an m such that $\mathcal{A}_m$ refines $\mathcal{B}_n$. Thus we may define the bonding map $f_{\mathcal{B}_n \mathcal{A}_m} : \Sigma_m \to \Sigma'_n$ by mapping each vertex u to a vertex v where $(\pi_m)^{-1}(\text{star}(u)) \subset h^{-1}({\pi'_n}^{-1}(\text{star}(v)))$. By taking limits we obtain an isomorphism between the Čech cohomology of X_x and the Čech cohomology of X'_x.

Theorem 3. *If X_x and X'_x are two fixed-orientation tiling spaces with finite relative orientation groups, and $h : X_x \to X'_x$ is a homeomorphism, then the isomorphism $h^* : H^2(X'_x) \to H^2(X_x)$ preserves the P-positive cones, up to an overall sign.*

Proof: To see that the P-positive cone is preserved by h^*, first consider whether the map $h : X_x \to X'_x$ is orientation preserving or reversing on each path-component. That is, fix a tiling $x_0 \in X_x$ and $h(x_0) \in X'_x$. There is a 1:1 map $j : \mathbb{R}^2 \to \mathbb{R}^2$ defined by $h(g(x_0)) = j(g)h(x_0)$ where $g \in \mathbb{R}^2$ acts by translation. We say h is orientation preserving or reversing based on whether j is orientation preserving or reversing as a homeomorphism of $\mathbb{R}^2$. Since the orbit of each point is dense, and since the action of $\mathbb{R}^2$ is continuous, whether h is orientation preserving or not is independent of the choice of x_0. Without loss of generality, assume that h is orientation preserving.

The goal is to show that every P-positive cohomology class on X' pulls back to a P-positive cohomology class on X. But the P-positive classes on X' are precisely the classes represented by positive linear combinations of basis cochains on Σ'_n, with n arbitrarily large. So we take a basic elementary cochain on Σ'_n (with n large), pull it back to X, and try to show that it's positive there. Without loss of generality we can assume that this cochain, which we call 1_Δ, is dual to a triangle Δ sitting in the middle of one of the tiles of Σ'_n. We then find an m large enough so that the open cover induced by the partition of Σ_m is a refinement of the open cover induced by the partition of Σ'_n. This refinement gives a vertex map ρ from the nerve of the cover of Σ_m (that is, Σ_m itself) to the nerve of the cover of Σ'_n (that is, Σ'_n itself). We must show that $\alpha = \rho^* 1_\Delta$ represents a P-positive cohomology class.

Now α itself is not necessarily a positive linear combination of elementary cochains. There will be some triangles T on Σ_m where α evaluates to $+1$ or -1, and others where it evaluates to 0. We call the triangles where it evaluates to ± 1 *essential*. Our task is to show that the positive essential triangles outnumber the negative ones in the following sense: there exists an $N > m$ such that every complete tile in Σ_N, mapped to Σ_m and paired with α, yields a positive number. Put another way, α evaluates positively on the image (in Σ_m) of every tile of level N in a particular tiling in X.

Here we are using the identification between a particular tiling in X (viewed as $\mathbb{R}^2$ with appropriate markers) and a path-component in X. In this identification, a point on a tiling corresponds to a translate of that tiling with that particular point situated at the origin. Call this space P (for path-connected) and let P' be the corresponding path-component of X'. The triangulation of Σ_m induces a triangulation of P, with each vertex, edge, or triangle of P mapping by π_m to a corresponding vertex, edge, or triangle of Σ_m. There are two metrics on P, namely the Euclidean metric and the metric inherited from X. For short-to-moderate distances, these metrics behave essentially identically:

Lemma 3. *There exist constants ϵ_1, ϵ_2 and M, with $0 < \epsilon_1 < M$ and $0 < \epsilon_2$, such that any two points in the same path component with Euclidean distance in (ϵ_1, M) have tiling-space distance greater than ϵ_2.*

Sketch of Proof: For our fixed set of polygonal tiles let K be a lower bound on the distance between the midpoint of any edge of any tile in any tiling and any other edge in that tiling, and let α be a lower bound of the interior angles of any of the tiles. Let N be such that any open ball of radius N in any tiling contains at least one tile and all the tiles touching it. The result then follows if $M < \min\{K/2, \sin(\alpha/2)\}$ and $\epsilon_2 < \epsilon_1/N$. ■

Corollary 2. *A path-connected set of tiling-space diameter ϵ_2 or less has Euclidean diameter ϵ_1 or less.*

We return to the proof of Theorem 3.

Pick $\epsilon_{1,2}$ for X as in the lemma. Since our homeomorphism h is between compact spaces, h^{-1} is uniformly continuous, so there is a δ for which all sets of (tiling-space) diamater δ or less in X' get mapped to sets of (tiling-space) diameter ϵ_2 or less in X. Then pick n large enough that we can choose our triangle Δ to have Euclidean diameter δ or less. This implies that each path-connected component (PCC) of $h^{-1}(\Delta)$ in X will have Euclidean diameter ϵ_1 or less.

Our proof of positivity now proceeds in four steps:

1. Show that preimages in P of essential triangles appear in disjoint and localized clusters, with each cluster associated to a specific PCC of $h^{-1}(\Delta)$.

2. Show that α applied to each complete cluster gives exactly $+1$.

3. Show that the number of complete clusters in a tile (of any level) is bounded from below by a multiple of the area of the tile.

4. Show that α, applied to triangles whose clusters are only partially in the tile, is bounded in magnitude from above by a multiple of the perimeter of the tile. (This contribution need not be positive.)

Since volume scales faster than perimeter, α then evaluates positively on all tiles of sufficiently high level.

Step 1. If T is an essential triangle, then its 3 vertices get sent by the bonding map to the 3 vertices of Δ. But this means that the open star of each of the vertices of T is contained in h^{-1} of the open star of the corresponding vertex of Δ. But then the intersection of these three open stars (namely T itself) is contained in h^{-1} of the intersection of the open stars of the vertices of Δ (namely Δ itself). Thus the open set T is contained in the open set $h^{-1}(\Delta)$, and every PCC of T must lie in a PCC of $h^{-1}(\Delta)$.

We define a cluster to be all the PCCs of essential triangles within a specific PCC of $h^{-1}(\Delta)$. The PCCs of Δ in P' are disjoint (since Δ was chosen to lie in the interior of a tile in Σ'_n) and have Euclidean diameter δ or less. This implies that the PCCs of $h^{-1}(\Delta)$ in P, and therefore the clusters, are disjoint and have Euclidean diameter ϵ_1 or less.

Step 2. Let D be a PCC of $h^{-1}(\Delta)$, and let Γ be the sum of all the triangles that intersect D. Γ is one of our clusters, plus a number of components of inessential triangles. We must show that $\alpha(\Gamma) = +1$. We do by showing that, as a simplicial chain, $\rho(\pi_m(\Gamma)) = \Delta$.

Let $\{v_i\}$ be the collection of all vertices in Γ. This is precisely the set of vertices whose stars contain points of D. If v_i is one of the vertices in this collection then $\rho(\pi_m(v_i))$ must be one of the vertices, d_1, d_2, d_3, of Δ. This follows since $(\text{star}(v_i))$ contains a point in D, and the image of this point lies in exactly three open sets in the open cover of Σ'_n, namely the stars of the three vertices of Δ. Since the cover induced by the triangulation of Σ_m is a refinement of the cover induced by the triangulation of Σ'_n, the open set corresponding to $\pi_m(v_i)$ is contained in one (or more) or these three open sets, and in no others. Thus if γ is a 2-simplex in Γ then $\rho(\pi_n(\gamma))$ is either $\pm\Delta$ or zero. Therefore, $\rho(\pi_m(\Gamma)) = k\Delta$ for some integer k. We will show that $k = 1$ by showing that the boundary of Γ maps to $+1$ times the boundary of Δ.

The boundary of Γ is a closed loop in X_x. Let x_1, x_2, x_3 be the three vertices of D (so that $\{h(x_i)\}$ are the 3 vertices of a PCC of Δ). We may choose 3 distinguished vertices v_1, v_2, v_3 on the boundary of Γ such that $\text{star}(v_i)$ contains d_i for $i = 1, 2, 3$. Because x_i is contained in a unique open set in $\mathcal{B}_n$ (namely that corresponding to the d_i), we must have $\rho(\pi_n(v_i)) = d_i$ for all i. Now let p_{12} denote the path from x_1 to x_2 along the boundary of D. There is a portion of the boundary of $\widetilde{\Gamma}$ which forms a path from v_1 to v_2 with the property that the star of each vertex along this path intersects p_{12}. Since the only elements of $\mathcal{B}_m$ that can contain points in p_{12} are the open sets corresponding to d_1 and d_2, the simplicial map $\rho \circ \pi_m$ applied to this portion of the boundary of Γ gives a path that begins at d_1, ends at d_2 and is completely contained in the edge between d_1 and d_2. The net output of the chain map $\rho \circ \pi_m$ applied to this path is therefore $+1$ times the edge from d_1 to d_2. Applying the same reasoning around the entire boundary of Γ and recalling that the map h is orientation-preserving, we see that $\rho \circ \pi_m$ applied to boundary of Γ is exactly $\partial\Delta$, as required.

Step 3. Since each preimage of an essential triangle has fixed Euclidean area, and since each PCC of $h^{-1}(\Delta)$ has diameter ϵ_1 or less, there is an upper bound to how many essential triangles can fit in a cluster. The number of clusters is then bounded from below by the number of PCCs of essential triangles, divided by the maximum number of triangles per cluster. But each essential triangle appears with prescribed positive density, so the number of clusters must scale with the area. Of these we must subtract off the partial clusters near the boundary, but this scales like the perimeter. See Step 4, below.

Step 4. Of the previously considered clusters, we must give special treatment to the ones within a distance ϵ_1 of the boundary of the tile. These may lie only partially within the tile, and the contribution of the essential triangles within the tile may not be positive. However, the area of this boundary region is at most ϵ_1 times the perimeter of the tile. Since the area of each triangle is bounded from below, the number of essential triangles in this region is bounded by a multiple of the perimeter. ■

6. Extension to X_ϕ for arbitrary relative orientation groups.

The previous development was for the cohomology of X_x in the case where the relative orientation group is finite. We now consider the Čech cohomology of X_ϕ with no restrictions on the relative orientation group.

In the case of arbitrary relative orientation group, P-positivity was defined for $H^2(X_0)$, using Perron eigenvectors and the substitution matrix, at the end of section 4. The natural isomorphisms between $H^3(X_\phi)$ and $H^2(X_0)$, and between $H^3(\Sigma_\phi)$ and $H^2(\Sigma_0)$, then give a consistent definition of P-positivity for $H^3(X_\phi)$. This notion of positivity can be rephrased directly in the language of Čech cohomology.

When working with X_ϕ, our complexes Σ_n are comprised of $S^1\times$ tiles. Triangulations of Σ_n then contain vertices, edges, triangles and tetrahedra. Correspondingly, our open covers, by stars of vertices, have single, double, triple and quadruple intersections. We orient the tetrahedra according to some global orientation on the 2-dimensional Euclidean group. The dual to any two 3-simplices within any $S^1\times$ tile are cohomologous. We declare that the positive sum of any of the duals to any of these 3-simplices is positive. A cohomology class on X_ϕ is positive if it can be represented as the pullback from some Σ_n of such a positive sum.

Theorem 4. *If X_ϕ and X'_ϕ are two all-orientation tiling spaces, and $h : X_\phi \to X'_\phi$ is a homeomorphism, then the isomorphism $h^* : H^3(X'_x) \to H^3(X_x)$ preserves the P-positive cones, up to an overall sign.*

Sketch of Proof: The proof of Theorem 3 carries over with only minor modification, as sketched below.

We first consider whether h preserves or reverses the orientations of path components of the tilings spaces – that is, whether it preserves or reverses the orientations of the action of the Euclidean group. Without loss of generality, assume that h is orientation-preserving.

If Δ is a 3-simplex arising from the triangulation of Σ'_n, with n large, we choose m such that the open cover induced by the triangulation of Σ_m refines the open cover induced by the triangulation of Σ'_n, pulled back to X_ϕ by h^{-1}. We then define the cochain $\alpha = \rho^* 1_\Delta$ on Σ_n, where $\rho : \Sigma_m \to \Sigma'_n$ is a vertex corresponding to the refinement. We call a tetrahedron on Σ *essential* if α applied to that tetrahedron is nonzero. We must show that every tile of sufficiently high level, crossed with S^1, contains more preimages of positive essential tetrahedra than of negative essential tetrahedra.

As before we work on a particular path-component P of X_ϕ. P is naturally identified, not with a single tiling, but with the unit tangent bundle of a single tiling as follows. A unit vector at a particular point in the tiling means that the tiling should be translated to put that point at the origin, and rotated so that the vector points in the positive x-direction. This unit tangent bundle has a natural Riemannian metric that plays the same role as that played by the Euclidean metric in the proof of Theorem 3. In particular, analogs to Lemma 3 and Corollary 2 are easily proven.

The proof then proceeds in 4 steps, exactly as before. We show that preimages of essential tetrahedra appear in clusters, that α applied to π_m of each complete cluster is $+1$, that the number of complete clusters in a tile of high level (crossed with S^1) scales as the area of the tile, and that the effect of the incomplete clusters is bounded by a multiple of the perimeter of the tile.

The only subtlety is in step 2, where previously we had made explicit use of the geometry of triangles. We define D and Γ as before, and show that each vertex in Γ must be mapped by $\rho \circ \pi_m$ to a vertex of Δ. Those vertices in Γ whose open stars contain a vertex x_i of D, must get mapped to the vertex d_i of Δ. Those whose open stars intersect the edge of D that runs from x_i to x_j must get mapped to d_i or d_j. Those that intersect the face whose corners are x_i, x_j and x_k must be mapped to d_i, d_j or d_k. These considerations, together with the fact that Γ is contractible, imply that $\partial\Gamma$, viewed as a chain, is mapped to $+1$ times $\partial\Delta$, and hence that Γ is mapped to exactly $+1$ times Δ. ∎

7. Algebraic consequences.

Below we examine some of the consequences of isomorphism of ordered groups obtained as direct limits. By Theorems 3 and 4 this will show that certain properties of the area stretching factor are homeomorphism invariants of X_x for the finite orientation case, and are homeomorphism invariants of X_ϕ in all cases.

Our setting is as follows. We assume we have Abelian groups H and H', obtained as direct limits of finitely generated free Abelian groups by primitive maps ϕ^* with Perron eigenvalues λ and λ'. An element of H will be denoted $[(v,k)]$, meaning the class of the element v in the k-th approximant to H. Of course, the application we have in mind is for H to be the top Čech cohomology of a tiling space, for the free Abelian groups to be the free part of the top cohomology of Σ_x or Σ_ϕ, and for the primitive map to be induced by substitution.

Lemma 4. *We can choose the left Perron eigenvector $r = (r_1, r_2, \ldots, r_n)$ of the primitive matrix ϕ^* so that each r_i is a polynomial in λ with integer coefficients.*

Proof: To find r, we wish to solve the equation

$$r(I\lambda - \phi^*) = 0. \qquad 11)$$

In order to solve this, we perform column operations until the matrix ϕ^* is in lower triangular form. Throughout this process, each entry of the matrix will be a quotient of integer polynomials in λ. Therefore there is a solution to the equation where each entry of r is a quotient of integer polynomials in λ. Thus by rescaling we may obtain a row eigenvector such that each entry is an integer polynomial in λ. ∎

Integer Case: In the case where λ is an integer, we will take the r from the previous lemma and scale by a factor of $1/g$ where g is the greatest common divisor of the entries of r.

Noninteger Case: In the case where λ is not an integer (and therefore irrational), take the r from the previous lemma and scale by a factor of $1/\lambda^p$ where p is the maximal power of λ that occurs in the entries of r. Now the row eigenvector is in $\mathbb{Z}[1/\lambda]$.

We define the map $\mu : H \to \mathbb{R}$ by $\mu[(v, k)] \mapsto \lambda^{-k} rv$ where r is fixed as above. In either the integer or noninteger case it is clear that the image of μ is a subset of $\mathbb{Z}[1/\lambda]$. Our positive cone H_+ is then the preimage of the positive real numbers.

Now we suppose that we have two such groups, H and H', with maps μ and μ' from each to $\mathbb{R}$, constructed as above, with Perron eigenvalues λ and λ', and suppose that $h : (H, H_+) \to (H', H'_+)$ is an order isomorphism. (We do not assume that h respects the direct limit structure, only that it respects the order structure.)

Lemma 5. *Let $x, y \in H$ with $\mu(x)$ and $\mu(y)$ both positive. Then*

$$\frac{\mu(x)}{\mu(y)} = \frac{\mu' h(x)}{\mu' h(y)}. \tag*{12)}$$

Proof: Let a, b be positive integers. Then

$$\begin{aligned} ax - by \in H_+ &\iff \mu(ax - by) \geq 0 \\ &\iff a\mu(x) - b\mu(y) \geq 0 \\ &\iff \frac{\mu(x)}{\mu(y)} \geq \frac{b}{a}. \end{aligned} \tag*{13)}$$

Since h is an order isomorphism,

$$\begin{aligned} ax - by \in H_+ &\iff ah(x) - bh(y) \in H'_+ \\ &\iff \frac{\mu' h(x)}{\mu' h(y)} \geq \frac{b}{a}. \end{aligned} \tag*{14)}$$

Therefore $b/a \leq \mu(x)/\mu(y)$ if and only if $b/a \leq \mu' h(x)/\mu' h(y)$. Since this is true for any positive integers a and b it must be the case that

$$\frac{\mu' h(x)}{\mu' h(y)} = \frac{\mu(x)}{\mu(y)}. \ \blacksquare \tag*{15)}$$

Lemma 6. *$\mathbb{Q}[\lambda]$ and $\mathbb{Q}[\lambda']$ are identical as subsets of $\mathbb{R}$.*

Proof: First we claim that there exist $p(x), q(x) \in \mathbb{Z}[x]$ such that

$$\lambda = \frac{p(\lambda')}{q(\lambda')}. \tag*{16)}$$

Choose $x = [(v, k)] \in H$ such that $\mu(x) > 0$. Let $y = [(v, k+1)] \in H$. Then

$$\frac{\mu' h(x)}{\mu' h(y)} = \frac{\mu(x)}{\mu(y)} = \lambda. \tag{17)}$$

Since the image of μ' is a subset of $\mathbb{Z}[1/\lambda']$, the claim follows.

Now let $f(x)$ be the minimal monic polynomial in $\mathbb{Z}[x]$ for the algebraic integer λ'. Since $q(\lambda') \neq 0$, the polynomial $f(x)$ does not divide $q(x)$. Since $f(x)$ is irreducible, the number 1 is a greatest common divisor of $f(x)$ and $q(x)$ in the ring $\mathbb{Q}[x]$.

Since $\mathbb{Q}[x]$ is a principal ideal domain, there exist polynomials $a(x), b(x) \in \mathbb{Q}[x]$ such that

$$a(x)q(x) + b(x)f(x) = 1. \tag{18)}$$

This implies

$$a(\lambda')q(\lambda') = 1. \tag{19)}$$

Therefore,

$$\lambda = a(\lambda')q(\lambda') \in \mathbb{Q}[\lambda']. \tag{20)}$$

which implies that as sets of real numbers,

$$\mathbb{Q}[\lambda] \subseteq \mathbb{Q}[\lambda']. \tag{21)}$$

The other inclusion follows similarly. ∎

(Of course, the previous lemma is trivial in the case where λ, λ' are integers.)

Lemma 7. *If λ is an integer then λ' is also an integer. Furthermore, λ and λ' have the same prime factors.*

Proof: Assume λ is an integer and recall $\mu[(v, k)]$ is defined as $\lambda^{-k} r v$ where r is an integer vector with relatively prime entries. Therefore there are integers $v_1, v_2, \ldots, v_n$ such that

$$\sum_{i=1}^{n} v_i r_i = 1. \tag{22)}$$

Define v to be the column vector $\lambda(v_1, v_2, \ldots, v_n)^{\mathrm{t}}$. Then $\mu[(v, 1)] = 1$ and the number 1 is in the image of μ.

From the previous lemma it is clear that λ' must also be an integer (it is a Perron eigenvalue which is either an integer or irrational). Thus for the same reason 1 is in the image of μ'.

Suppose $y \in H'$ with $\mu'(y) = 1$. Let $[(u, k)] = h^{-1}(y)$. Then

$$\begin{aligned} \frac{\mu[(u, k+1)]}{\mu[(u, k)]} &= \frac{\mu' h[(u, k+1)]}{\mu'(y)} \\ \frac{1}{\lambda} &= \mu' h[(u, k+1)]. \end{aligned} \tag{23)}$$

This implies $1/\lambda \in \mathbb{Z}[1/\lambda']$. Similarly, $1/\lambda' \in \mathbb{Z}[1/\lambda]$. Therefore, as subsets of $\mathbb{R}$,

$$\mathbb{Z}[1/\lambda] = \mathbb{Z}[1/\lambda']. \qquad 24)$$

This is the case if and only if λ and λ' have the same prime factors. ■

Combining these results with Theorems 3 and 4 then yields our main result, Theorem 1.

8. Conclusion.

We add here some remarks of a general nature. The basic objects of the above study are hierarchical tilings of the Euclidean plane, which are mainly of interest for their unusual symmetry (geometric) properties. Our analysis is an outgrowth of a traditional dynamical analysis, in which a space of tilings is equipped with a topological and Borel structure together with the action of Euclidean motions, the latter maintaining contact with geometry. It is notable that we have been ignoring the dynamical action, obtaining topological invariants of tiling spaces. This is a recent development, but not new to this paper, appearing as it has in noncommutative approaches. However a distinctive feature of this work is that our invariant is only sensible in the top (Čech) cohomology, which suggests to us that it might not be essentially topological; more specifically, it might be invariant under maps more general than homeomorphisms.

Acknowledgements. We gratefully acknowledge useful discussions with Mike Boyle, Gary Hamrick, Ian Putnam, Mike Starbird and Bob Williams.

References

[AnP] J. Anderson and I. Putnam, Topological invariants for substitution tilings and their associated C$^{\mathrm{star}}$-algebras, *Erg. Thy. Dyn. Syst.* **18** (1998), 509-537.

[BJV] M. Barge and B. Diamond, A complete invariant for the topology of substitution tiling spaces, preprint.

[BoT] R. Bott and L. Tu, *Differential forms in algebraic topology*, Graduate Texts in Mathematics, **82**, Springer-Verlag, New York, 1982.

[Con] A. Connes, *Noncommutative Geometry*, Academic Press, San Diego, 1994.

[Ell] G. Elliott, On the classification of inductive limits of sequences of semisimple finite dimensional algebras, *J. Algebra* **38** (1976), 29-44.

[Goo] C. Goodman-Strauss, Matching rules and substitution tilings, *Annals of Math.*, **147** (1998), 181-223.

[GPS] T. Giordano, I. Putnam and C. Skau, Topological orbit equivalence and C*-crossed products, *J. reine angew. Math.* **469** (1995), 51-111.

[HPS] R. Herman, I. Putnam and C. Skau, Ordered Bratteli diagrams, dimension groups, and topological dynamics, *Intern. J. Math.* **3** (1992), 827-864.

[Ke1] J. Kellendonk, Non-commutative geometry of tilings and gap labelling, *Rev. Math. Phys.*, **7** (1995), 1133-1180.

[Ke2] J. Kellendonk, Topological equivalence of tilings, *J. Math. Phys.* **38** (1997), 1823-1842.

[Ke3] J. Kellendonk, The local structure of tilings and their integer group of coinvariants, *Commun. Math. Phys.* **187** (1997), 115-157.

[Ke4] J. Kellendonk, On K$_0$-groups for substitution tilings, preprint.

[KeP] J. Kellendonk and I. Putnam, Tilings, C*-algebras and K-theory, preprint.

[PaT] W. Parry and S. Tuncel, *Classification Problems in Ergodic Theory*, Lecture Note Series, **67**, Cambridge University Press, Cambridge, 1982.

[Put] I. Putnam, The ordered K-theory of C*-algebras associated with substitution tilings, preprint.

[Ra1] C. Radin, The pinwheel tilings of the plane, *Annals of Math.* **139** (1994), 661-702.

[Ra2] C. Radin, *Miles of Tiles*, Student Math. Lib. **1**, Amer. Math. Soc., Providence, 1999.

[RaS] C. Radin and L. Sadun, An algebraic invariant for substitution tiling systems, *Geometriae Dedicata* **73** (1998), 21-37.

[RaW] C. Radin and M. Wolff, Space tilings and local isomorphism, *Geometriae Dedicata* 42 (1992), 355-360.

[Sol] B. Solomyak, Nonperiodicity implies unique composition for self-similar translationally finite tilings, *Discrete Comput. Geom.* **20** (1998), 265-279.

[SuW] D. Sullivan and R.F. Williams, The homology of attractors, *Topology* **15** (1976), 259-262.

Nicholas Ormes, Mathematics Department, University of Texas, Austin, TX 78712
Email address: `ormes@math.utexas.edu`

Charles Radin, Mathematics Department, University of Texas, Austin, TX 78712
Email address: `radin@math.utexas.edu`

Lorenzo Sadun, Mathematics Department, University of Texas, Austin, TX 78712
Email address: `sadun@math.utexas.edu`