The Dissertation Committee for Betseygail Rand
certifies that this is the approved version of the following dissertation:

Pattern-equivariant Cohomology of Tiling Spaces With Rotations

Committee:

Lorenzo Sadun, Supervisor

Daniel Allcock

Rafael de la Llave

Michael Marder

Charles Radin

Robert Williams

Pattern-equivariant Cohomology of Tiling Spaces With Rotations

by

Betseygail Rand, BS

Dissertation

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

August 2006

This dissertation is dedicated to all the little tiles who made it possible. They were very still and did not wiggle.

Acknowledgments

First, I would like to thank the members of my committee. Without their time and consideration, this would not be possible.

Most important, though, I would like to express my deep gratitude and appreciation for my advisor, Lorenzo Sadun. For the past four years, he has made it a priority to guide me through learning to conduct mathematical research. Many times he expressed confidence in my work, where I had none. Without his substantial patience and support, I would not have gotten to this point. He has logged a tremendous amount of hours with me, often meeting more than once a week, merely because I dropped by his office with questions. I am indebted and very much thankful to him for being my advisor and mentor.

Betseygail Rand

The University of Texas at Austin
August 2006

Pattern-equivariant Cohomology of Tiling Spaces With Rotations

Publication No. ______________________

Betseygail Rand, Ph.D.
The University of Texas at Austin, 2006

Supervisor: Lorenzo Sadun

This paper develops a new cohomology theory on generalized tiling spaces. This theory incorporates both the rotational geometry of the tiling space and the local pattern geometry into the structure of the cohomology groups. Our use of the local pattern geometry is a generalization of pattern-equivariant cohomology, a theory developed by Ian Putnam and Johannes Kellendonk in 2003. It was defined for tilings whose tiles appear as translates. The most general setting in tiling theory is to work with tiling spaces, with an action of a subgroup of the Euclidean group. This paper defines a new, general pattern-equivariant cohomology for tiling spaces with finite rotation groups, and proves that it is preserved under homeomorphisms

which commute with the action of the group. It is conjectured here that this theory is not a topological invariant for tiling spaces with infinite rotation group.

Contents

List of Figures

Chapter 1

Introduction

Pattern-equivariant cohomology theory was developed by Ian Putnam and Johannes Kellendonk in 2003, [K], [KP], for tilings whose tiles appear in fixed orientations. It was an innovative development that linked the local geometry of a tiling to the global topology of the associated tiling space.

We set out to generalize this theory in two directions: first, as a theory on tiling spaces instead of individual tilings. Second, we want to include tiling spaces with tiles occurring in multiple orientations - possibly infinitely many. This adaptation is not a mere modification of definitions. Locally, a tiling space is a leaf $\times$ a Cantor set, where the leaf is locally isomorphic to G, a subgroup of rigid motions on $\mathbb{R}^2$. The original theory ties the Cantor structure to the topology; rotations are part of the leaf structure, and in this way the the rotational theory ties a portion of the leaf structure to the topology. We develop a large, flexible structure, in which a choice of representation of the rotation group serves to reveal the structural interplay between rotations and cohomology.

When [K] published his work, he proved topological invariance of the newly defined concept by showing that the pattern-equivariant cohomology of a tiling is isomorphic to the Čech cohomology of the associated space. In the case where the rotation group of the tiling space is finite, we prove from scratch that this cohomology is a topological invariant. When the rotation group is infinite, we conjecture that the cohomology groups are not in fact a topological invariant, and posit what a counter-example could look like.

Along the way, we prove an approximation theorem which is valuable outside of the context in which it is developed. Philosophically, it serves as the tiling space version of theorems which state that a stronger class of functions is dense within a weaker class of functions. For example, we have theorems approximating a continuous function by a smooth function, an L_1 function by a continuous function with compact support, or an integral by a Riemann sum. The approximation theorem holds in tiling spaces with both finite and infinite rotations.

Loosely speaking, we show that a function which is a topological conjugacy can be approximated arbitrarily closely by a function which preserves the local structure of the tiling space. This is important because it untangles the relationship between the topology and the local structure of tiling spaces. Pattern-equivariant cohomology is a spin-off of De Rham cohomology theory that places restrictions on forms on tiling spaces. These restrictions limit our attention to those forms which are local in nature. Thus in hindsight, this topic is ripe for linking topology of tiling spaces to their local structure, as done by the approximation theorem.

This paper will first introduce tilings and tiling spaces. We will next discuss the notions of equivalence, and what the approximation theorem achieves, but we will not formally state and prove the theorem in this context. We will then switch gears, and give the background to pattern-equivariant cohomology. First we describe an inverse limit structure to tiling spaces, which allows us to compute the cohomology of these spaces. We then lay out the framework of the original theory. Next we develop the theory of pattern-equivariant cohomology, including rotations, both the finite and the infinite situation. (The pattern-equivariant cohomology is still informative for a given space with infinite rotations; it just isn't known if it's a topological invariant.)

After this we thoroughly deconstruct the finite rotation case. We prove that pattern-equivariant cohomology in this case is a topological invariant, give examples, and classify the relationship between different choices of representation. In the course of proving the topological invariance, we prove the approximation theorem.

Lastly, we explore the case where there are infinitely many rotations in the tiling space. We clarify what complicates the picture and state our conjecture, that pattern-equivariant cohomology is not a topological invariant in this case.

Chapter 2

What is a tiling space?

2.1 Definitions

Intuitively, a tiling is pattern of polygons, covering the plane.

Definition 1 (tile) *A tile is a subset of* $\mathbb{R}^2$ *homeomorphic to a closed disk, with an associated label.*

We will assume our set of tiles is finite. Next, pick a group G of rigid motions on $\mathbb{R}^2$. (A priori, this work could be carried out in any dimension.) One immediate complication in higher dimensions is that the group of rotations on $\mathbb{R}^n$, for $n > 2$ is no longer abelian. Generally we pick G to be either the group of translations or translations along with rotations, some subgroup of $SO(2)$. G acts on tiles, so that we may speak of a given tile occurring anywhere in the plane.

Definition 2 (tiling) *Beginning with a finite set of tiles,* $\mathcal{A}$*, and a group* G*, a tiling of* $\mathbb{R}^2$ *is a covering of the plane by elements* $g \cdot A$*,* $g \in G$*,* $A \in \mathcal{A}$*, which intersect only on their boundaries.*

Our tiles may be wild; indeed there are circumstances where tiles may have fractal boundaries and other weird characteristics. However, there is a procedure for constructing a new tiling with very nice properties from any original tiling. (see [P]). The new tiling is called the derived Voronoi tiling, and has the feature that all its tiles are polygons which meet full edge to full edge. The dynamics of the tiling space associated to the derived Voronoi tiling are equivalent to the dynamics of the tiling space associated to the original tiling. Thus we may assume from the outset that our set of tiles is composed of polygons, and that they meet full edge to full edge.

2.2 Example - the half-hex tiling

We now introduce a tiling that we will trace throughout the rest of the paper. The half-hex tiling is built out of trapezoids occurring in the following orientations (see figure 2.1):

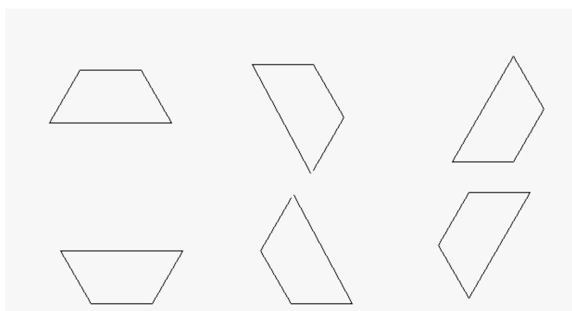

Figure 2.1: Tiles composing the half-hex tiling

and a portion of the full tiling looks like (see figure 2.2):

2.3 Requirements on tilings

Some notation

$B_r(x)$ is the open disk of radius r, with center at point x. A *patch* is a union of tiles. $[B_r(x)] =$ the patch composed of all tiles with non-trivial intersection with $B_r(x)$. If there is more than one tiling in the context, we may specify that the patch occurs within a tiling T by $[B_r(x)]_T$. We extend the action of G to patches.

We next pose some constraints on the tiling in order to focus on those with the most interesting dynamics. In order to eliminate tilings which are too simplistic,

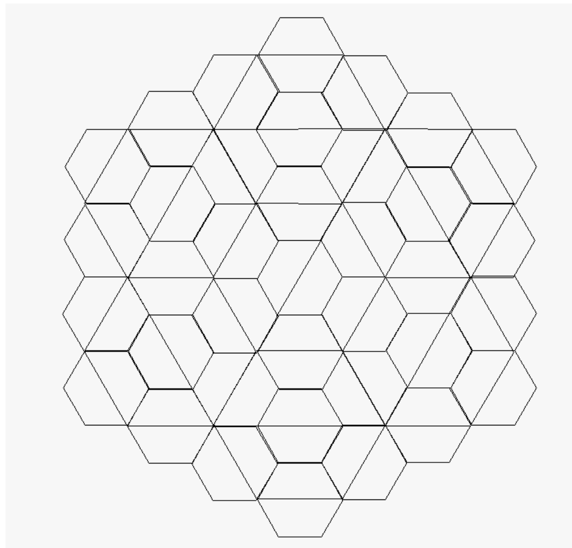

Figure 2.2: Sample patch of the half-hex tiling

we impose non-periodicity.

Definition 3 (non-periodicity) *A tiling T is non-periodic if $\nexists\ d \in \mathbb{R}^2$ such that $T + d = T$.*

In other words, we do not want regularly repeating patterns. The next two conditions keep our tilings from being too wild and unmanageable.

Definition 4 (finite local complexity) *A tiling T has the property of finite local complexity (FLC) if for every r there exist finitely many patches $P_1, \ldots, P_n$ such that for every $[B_r(x)]$ we have $g \cdot P_i = [B_r(x)]$ for some $g \in G, 1 \leq i \leq n$.*

FLC ensures that for a fixed radius, only finitely many patterns will be found of that size, up to the action of G. In the checkerboard pattern below,(see figure 2.3) squares on the right hand side have side length $\sqrt{2}$, while squares on the left hand side have unit side length. Along the fault line running up the y-axis, the squares occur at irrationally set periods with respect to each other. If we fix the radius r, then there will be infinitely many distinct $[B_r(p)]$, where p is a point on the y-axis, corresponding to the infinitely many irrational discrepancies between tiles on the left and on the right. Therefore this tiling lacks FLC.

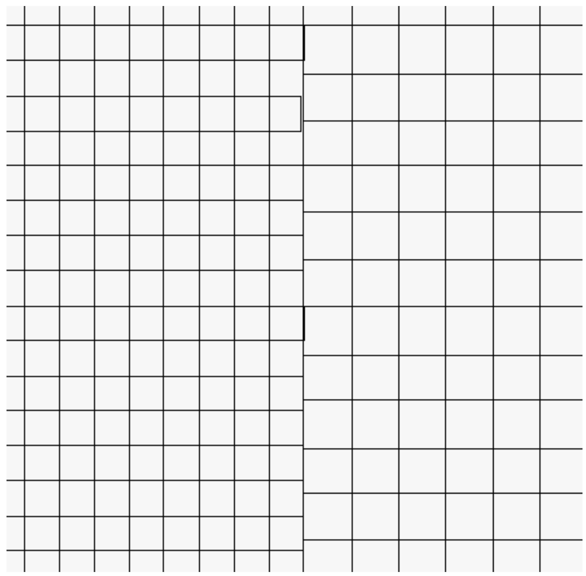

Figure 2.3: Because there are infinite variations along fault line, this tiling lacks finite local complexity

The last condition brings in a statistical orderliness.

Definition 5 (repetitivity) *A tiling T is repetitive if every patch P has an associated radius R such that every $[B_R(x)]$ contains a copy of P.*

In other words, repetitivity guarantees that any patch occurs on a regular basis; the distance between copies of a patch cannot grow arbitrarily large. The checkerboard pattern above also lacks repetitivity; the patches on the left do not occur anywhere on the right, and vice versa.

2.4 Tiling Spaces

From now on, all tilings will be assumed to have the properties from the previous sextion. Next we define the metric on two tilings.

Some notation: let $T|_r = B_r(0)$ and $[T|_r] = [B_r(0)]$. The notation $T|_r$ emphasizes a particular tiling T, useful in the context of a tiling space, whereas $[B_r(0)]$ is better when the specific tiling is understood, and we want to emphasize the location of a patch within that tiling.

Definition 6 *Let T_1 and T_2 be two tilings built out of the same set of labeled tiles, with the same group G. Then $d(T_1, T_2) = min(1, \sigma(T_1, T_2))$ where*

$$\sigma(T_1, T_2) = min\{|g| \text{ such that } (g \cdot T_1)|_{1/|g|} = T_2|_{1/|g|}\}.$$

That is, two tilings are ϵ-apart if they agree on a ball of $1/\epsilon$ about the origin, up to an ϵ-wiggle. Under this metric there are three ways for tilings to be close to each other. First, if we take a tiling and shift it by ϵ, the new tiling is ϵ away from the original, like one would desire. Second, consider $[B_r(0)]$ for large r. By repetitivity, copies of this large patch occur regularly throughout the tiling. By a judicious choice of translation, then, we can shift the tiling a large distance and wind up with a new tiling with the same patch at its origin. In this way very large actions can result in tilings that are arbitrarily close together. Finally, two tilings that are not in the same g-orbit can be close together, if they match on a large patch (give or take an ϵ-wiggle) and then mismatch beyond that point.

Definition 7 *Fix a group G and tiling T. The tiling space X_T is formed by taking the orbit of T under G, and then taking the closure under the metric above.*

Notation: often we will suppress the tiling T and simply write X.

Taking the closure under the metric above is a statement about local patches - the metric sees patches around the origin, but not the global structure of any particular tiling. A second tiling T' is in the space X if it is composed of all, and nothing more than the patches that appear in T.

Lemma 8 *$T' \in X_T$ iff $T \in X_{T'}$*

Case 1: $G = \mathbb{R}^2$. If $T' \in X_T$, then every patch in T' occurs somewhere in T. Specifically, each patch $[B_n(0)]_{T'}$, $n \in \mathbb{Z}$, occurs somewhere as $[B_n(x_n)]_T$ for some sequence of points x_n. The sequence of tilings $\{T - x_n\}$ converges to T'.

To show $T \in X_{T'}$, consider the sequence of patches $[B_n(0)]_T$, $n \in \mathbb{Z}$, which converges to T. Each $[B_n(0)]$ has an associated repetitivity radius R_n, so that any patch in T with radius R_n is guaranteed to contain a translate of the patch $[B_n(0)]$. Specifically then, for $m > R_n$, the tiling $T - x_m$ has a patch around its origin containing a copy of $[B_n(0)]$. Since $[B_m(x_m)]_T = [B_m(0)]_{T'}$, we have that the set of patches $[B_n(0)]_T$ all occur in T'. Therefore we can list a sequence of translates of T' that converge to T.

Case 2: $\mathbb{R}^2 < G \leq \mathbb{E}^2$. If the rotational portion of G is dense in S^1, then a given patch P in T', with orientation θ, may never occur in that precise orientation anywhere in T. However, there is a sequence of copies of P in T with orientations θ_n such that θ_n converges to θ. As before, our goal is to find a sequence patches that describe T within T'. We can take the sequence of patches $[B_n(0)]_T$, and let g_n be a Euclidean motion such that $g_n \cdot T'|_n = [B_n(0)]_T$. Then the sequence $g_n \cdot T'|_n$ will converge both in radius and rotation to the tiling T. Thus T is in the orbit closure of T'.

□

Note that patches in T' will automatically come equipped with the same repetitivity radii that they had under T. If P has repetitivity radius R in T, then by FLC we can list all possible patches of radius R in T. These must be the same possible patches of radius R in T', and so we know exactly that every patch of radius R in T' contains a copy of P.

Locally, a tiling space looks like $G\times$ a Cantor set, (see figure 2.4):

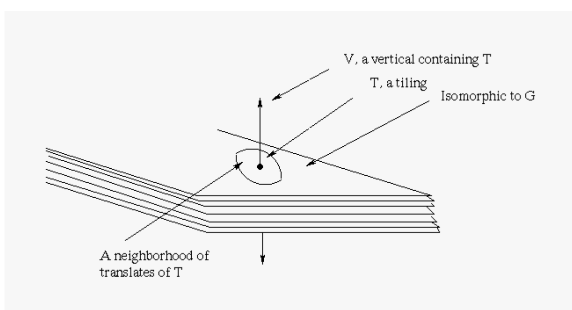

Figure 2.4: Schematic of a tiling space

A neighborhood of T in the leaf direction contains all tilings related to T by a small element of G. A tiling that is close to T in the Cantor direction matches T perfectly out to some fixed radius, and may mismatch beyond that radius.

Notation: let the vertical of cantor length ϵ containing T be

$$V_\epsilon(T) = \{T' \text{ such that } T|_{1/\epsilon} = T'|_{1/\epsilon}\}.$$

A vertical forms a Cantor set because it is composed of infinitely many tilings which are dense in the metric, yet also discrete set. It is known ([SW]) that when the group is the translation group R^2 that the global structure of the tiling space is that of a Cantor bundle over a torus. The global structure under general G is an open problem.

We now reformulate our earlier axioms on individual tilings to properties of tiling spaces. The first is aperiodicity. If that condition were relaxed, the translational orbit of a periodic tiling would be cylinder-shaped if it were periodic in one direction, and a torus if it were periodic in more than one direction. Combined with the property of repetitivity, we then get that tilings in the orbit closure must also all be periodic if any one tiling is periodic. The Cantor structure of a tiling space, under our axioms, implies and is implied by the two properties aperiodicity and repetitivity.

FLC implies that the tiling space is compact. For a sequence of tilings $\{T_i\}$, there are finitely many patches $[B_r(0)]$ for any radius r, up to a small action of our group. For a subsequence to converge under the metric means that the tilings agree on larger and larger radii around the origin, up to a smaller and smaller action of the group. Specifically, for every $n \in \mathbb{Z}$, there are finitely many patches $[B_n(0)]$, up to rotations and small translations. By FLC we can pick a subsequence of patches which are consistent insofar as $m > n$ implies that $[B_m(0)] = \theta[B_n(0)] + d$, where d is bounded. Since the group G is compact on regions of bounded translations, we can take a new subsequence of the first subsequence, which converges in the metric.

Finally, repetitivity implies that the tiling space is minimal - that there is one unique orbit closure. This follows from lemma 8. This property is often stated as follows: two tilings, T and T' are *patch-equivalent* if every patch in each one occurs in the other. As noted earlier, this implies that both will have the same radii of repetitivity, and $X_T = X_{T'}$.

Chapter 3

Differing notions of equivalence: topological equivalence compared to dynamical equivalence

Because of the G action on our space X, we may speak of the dynamics of a tiling space. Two spaces have intertwined dynamics if there is a map between them that commutes with the action of the group.

Definition 9 *Given two tiling spaces X and Y, with groups G_X and G_Y, these systems are* topologically conjugate *if there exists a homeomorphism $h : X \to Y$ such that for $g_X \in G_X, T_X \in X$, there exists $g_Y \in G_Y$ such that $g_X \cdot T_X = (h^{-1} \circ g_Y \circ h)(T_X)$.*

In particular, the homeomorphism h sends G_X-leaves to G_Y-leaves.

In our case, our tiling spaces X and Y will have the same group G acting upon them. Consider first that the rotational portion of the group G is finite. Then being *topologically conjugate* implies that the homeomorphism h sends translational leaves to translational leaves. Since the translational orbit of $T \in X$ gets sent to the translational orbit of $h(T) \in Y$, we have that $h(T)$ induces a homeomorphism of $\mathbb{R}^2$. In this way, the map h descends to T as an individual tiling identified with a covering of $\mathbb{R}^2$. In the case that G includes all of S^1, then being topologically conjugate does not imply that translational orbits get sent to translational orbits, and so $h(T)$ does not induce a homeomorphism of $\mathbb{R}^2$. When we discuss tiling spaces with rotation group S^1, we then sometimes require the extra condition that a homeomorphism also sends translational leaves to translational leaves.

Definition 10 *If $G_X = G_Y$, a map $h : X \to Y$ that commutes with the action of the group is called a* factor map*.*

Definition 11 *Given two tiling spaces X and Y, we say a map $h : X \to Y$* preserves verticals *if there are radii r and s such that $T_1|_r = T_2|_r \in X$ implies $h(T_1)|_s = h(T_2)|_s \in Y$. The map h is also called a local map.*

The approximation theorem in chapter 7.1 bridges these two notions. We show that a continuous map which sends orbits to orbits can be approximated by local map, with the property that the image of a tiling under the approximation map is a small translation of the image of the tiling under the original map.

Finally, there is a third equivalence, the intersection of being both locally and topologically conjugate.

Definition 12 *Given two tiling spaces X and Y, we say Y is locally derivable from X if there exists a map $\sigma : X \to Y$ which commutes with the action of the group, and a radius r such that for any T_1, $T_2 \in X$, $T_1|_r = T_2|_r$ implies that $\sigma(T_1)|_{\vec{0}} = \sigma(T_2)|_{\vec{0}}$. Two tiling spaces X and Y are mutually locally derivable, or MLD, if each is locally derivable from the other.*

In other words, two spaces are MLD if there are maps back and forth which preserve verticals and commutes with G. This is the strongest notion of equivalence, and informally this is what one means when one says two spaces are equivalent. Earlier we mentioned the derived Voronoi construction, producing a new tiling with polygonal tiles from an old tiling. The tiling spaces associated with the old and new tilings are MLD to each other.

For some time there was an open question: are there topologically conjugate spaces which are not MLD? This was resolved by [RS1], [Pe], and we give the following 1-dimensional example which illustrates this point (also found in [RS1]):

The Fibonacci substitution is given on two letters, a and b, by $\sigma(a) = b$ and $\sigma(b) = ab$. Repeatedly applying the substitution to one of the letters results in an infinite sequence. By "infinite sequence" we mean a right-infinite sequence. We can extend this to a bi-infinite sequence by beginning with a two letters, marking a decimal point between them, and repeatedly applying the substitution. For example, if we apply σ^2 repeatedly to the starting letters $b.a$, we get the:

$$b.a \to bab.ab \to bababbab.abbab \to bababbababbabbababbab.abbabbababbab$$

an so on. (By FLC, there is always a bi-infinite sequence which is fixed under some power of the substitution rule. We intentionally chose to apply σ^2 to $b.a$ in order to produce a fixed point of the substitution.)

As a bi-infinite sequence, we can form a *subshift space*, the analog of a tiling space. We first take the translational $\mathbb{Z}$-orbit of a sequence, and then take its closure under the metric that $\{a_n\}$ and $\{b_n\}$ are $1/m$ apart if they match on entries $[a_{-m}, a_m]$ and differ beyond, up to an m shift.

As with tiling spaces, there is a notion of topological conjugacy and of being MLD. In subshift spaces the convention is to use the term "sliding block code". (To be historically accurate, the concept of MLD was a generalization of a sliding block code, not the other way around.) Two subshift spaces are topologically conjugate if there is a continuous homeomorphism h that commutes with the $\mathbb{Z}$-action. The map h between two subshift spaces is a sliding block code if $\exists n$ such that if two sequences agree on the $2n$ entries around the origin, then their images agree on the entry at the origin. It turns out, (see [LM]), that for $\mathbb{Z}^d$ subshifts, every topological conjugacy is a sliding block code.

By assigning lengths to the tiles a and b, the bi-infinite sequence can be identified with a tiling of $\mathbb{R}$. Instead of a $\mathbb{Z}$-action, we now have an $\mathbb{R}$-action. The substitution $a \to b$ and $b \to ab$ has corresponding substitution matrix $\begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$, which has eigenvalues $\lambda = \frac{1 \pm \sqrt{5}}{2}$. [RS1] proves the following theorem:

Theorem 13 *Let X and Y be Fibonacci tiling systems, with tiles A_X, A_Y and B_X, B_Y. If $|A_X| + \lambda|B_X| = |A_Y| + \lambda|B_Y|$ then X and Y are topologically conjugate, but the conjugacy is not a sliding block code.*

As this result shows, MLD tiling spaces are finely separated from topologically conjugate tiling spaces. Our result is important in this area. We show that a continuous map which commutes with G can be approximated arbitrarily closely by a map which preserves verticals. This allows one to prove a property is a topological invariant by showing it is preserved by vertical-preserving maps.

Precisely:

Theorem 14 *If f is a continuous map from $X \to Y$, then for every ϵ sufficiently small, there exists $\hat{f}_\epsilon$ with the following properties:*

1. *$\hat{f}_\epsilon$ is a smooth, vertical-preserving map from $X \to Y$,*

2. *$f(T)$ and $\hat{f}_\epsilon(T)$ are in the same translational leaf, and*

$$\left| f(T) - \hat{f}_\epsilon(T) \right| < \epsilon,$$

3. *if f commutes with rotations, then $\hat{f}_\epsilon$ will also commute with rotations.*

In chapter 7.1 we will prove this theorem, in the pattern-equivariant cohomology context that it was developed.

Chapter 4

Tiling spaces are inverse limit spaces

There are many advantages to identifying the inverse limit structure within a tiling space. Primarily, we gain insight into the global structure, and specifically computations such as cohomology of the space become much more feasible to compute.

There are many constructions identifying an inverse limit structure within tiling spaces ([AP], [ORS], [S]); we choose the following for its flexibility. This particular construction does not appear in the literature.

Definition 15 (inverse limit space) *Let $\{K_r\}_{r\in I}$ be a set indexed by I, where I is a partially ordered, directed set. In other words, for any r and s $\exists$ t such that $t > s$ and $t > r$ are both defined. Where $r \geq s$ is defined, we have maps $\rho_{r,s} : K_r \to K_s$. These descending maps must be consistent:*

$$\rho_{r,r} = \text{ identity on } K_r,$$

and for $r > s > t$,

$$\rho_{s,t} \circ \rho_{r,s} = \rho_{r,t}.$$

Then the inverse limit space is the subset of the direct product of the K_r,

$$\varprojlim K_r = \{\{a_r\} \in \Pi K_r \text{ such that } \rho_{r,s}(a_r) = a_s\}.$$

The sets K_r are called approximants. Typically they are indexed by $I = \mathbb{R}$ or $I = \mathbb{Z}$ - we will use both cases.

We now construct a tiling space as an inverse limit space. A tiling can be represented as a sequence, $T = \{a_i\}_{i \in I}$, where:
1. I is the set of bounded regions of $\mathbb{R}^2$, partially ordered by inclusion,
2. a_i is a tiling of the region $i \in I$,
3. Where $s \subseteq r$ is defined, $\rho_{r,s} : a_r \to a_s$ by restriction of the partial tiling r to the region s.

The advantage of describing tiling spaces as inverse limit spaces is that it will ease computations. Under the construction above, it is unclear what the approximants K_r might be. Without understanding the approximants, we are not able to actually do any computations. However, there are choices of cofinal subsets of I which allow us to explicitly state what the approximants K_r will be.

Definition 16 *A subset B of A is cofinal if, for every a in A, $\exists$ b in B such that $a \leq b$.*

Lemma 17 *If B is cofinal in A, then $\varprojlim \{K_r\}_{r\in A} = \varprojlim \{K_r\}_{r\in B}$*

Proof: Let $\Phi : \varprojlim \{K_r\}_{r\in B} \to \varprojlim \{K_r\}_{r\in A}$ as follows: for $\{b_r\}$ in $\{K_r\}_{r\in B}$, we will form $\Phi(\{b_r\}) = \{a_s\}$. Starting with $s \in A$, and let $r \in B$ be such that $s \leq r$. Hence $\rho_{r,s}$ is defined and we can pick $a_s = \rho_{r,s}(b_r)$. Is this one-to-one? Let $\{b_r\} \neq \{b'_r\}$ in $\{K_r\}_{r\in B}$. Suppose they disagree at the m^{th} entry, ie $b_m \neq b'_m$. Pick $s \in A$ such that $s > m$. As before, this yields $t \in B$ such that $t \geq s$, and $a_s = \rho_{t,s}(b_t)$, and $a'_s = \rho_{t,s}(b'_t)$. Since $\rho_{s,m}(a_s) = b_m \neq b'_m = \rho_{s,m}(a'_s)$, we have that $a_s \neq a'_s$.

Let $\Psi : \varprojlim \{K_r\}_{r\in A} \to \varprojlim \{K_r\}_{r\in B}$ as follows: for $\{a_s\}$ in $\{K_r\}_{r\in A}$, form $\Psi\{a_s\} = \{b_r\}$ by just restricting to the subset of indices $B \subset A$. This too is one-to-one. Consider $\{a_s\} \neq \{a'_s\}$ in $\{K_r\}_{r\in A}$, and as before, assume they differ in the m^{th}

entry. Since B is cofinal in A, there exists $k \in B$ such that $m \leq k$. Since $a_k \neq a'_k$ (because $\rho_{k,m}(a_k) \neq \rho_{k,m}(a'_k)$), $\Psi(\{a_r\}) \neq \Psi(\{a'_r\})$.

$\square$

The index set I was the set of bounded regions in $\mathbb{R}^2$. Our cofinal set will be $\{[B_r(0)]\}$; we will reference these by their radius $r \in \mathbb{R}$. One can see that any bounded region in $\mathbb{R}^2$ will be contained in a sufficiently large $[B_r(0)]$.

A point a_r in K_r is a partial tiling $[B_r(0)]$. Schematically, the point $a \in K_r$ represents the union of tiles, with a fixed choice of origin (and orientation) (see figure 4.1). Then a point in the inverse limit space, $\{a_r\}$, is a consistent description of

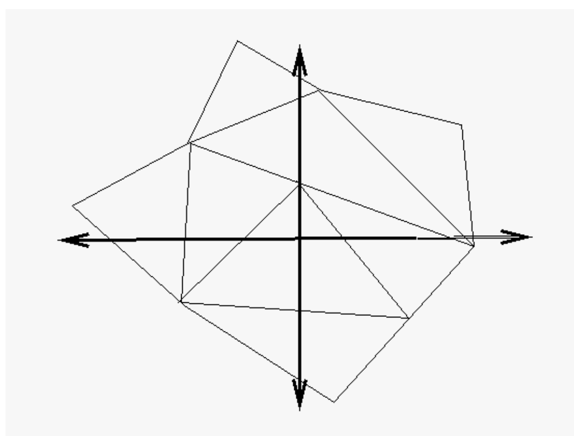

Figure 4.1: a_r in K_r is a partial tiling around the origin

the tiling by larger and larger choices of $[B_r(0)]$. The maps $\rho_{r,s}$ work by restricting the patch of larger radius to its internal patch of smaller radius.

We will see that the global structure of an approximant K_r is a branched manifold, composed of cells $\{N_i\}$ and branched edges where two or more cells meet.

First we form the cells, $\{N_i\}$. Consider the following: For a tiling T, $[T|_r]$ is the union of tiles with nontrivial intersection with $B_r(0)$. There is a set of tilings such that for any T' in this set, $[T|_r]$ and $[T'|_r]$ are the same patch, although possibly with slightly different placements of origins within that patch. A cell $\{N_i\}$ is the set of possible placements of origin that all yield the same patch.(See figure 4.2 for illustration.)

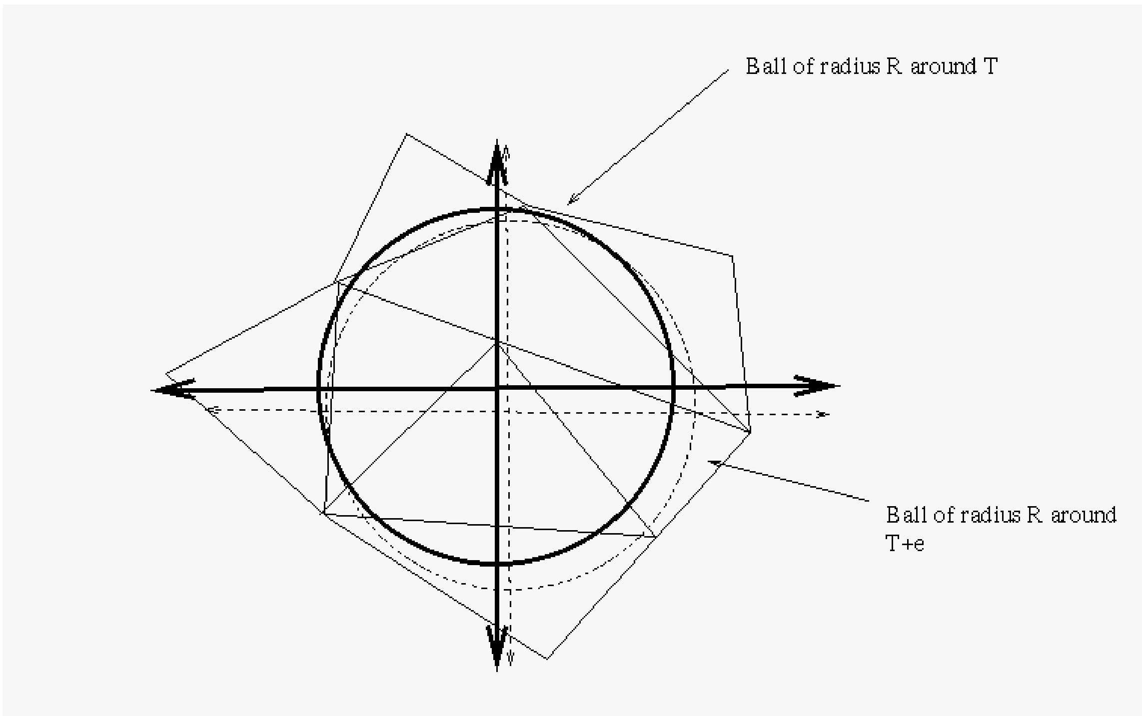

Figure 4.2: Small shift in origin may not affect $[B_r(0)]$.

To be precise, let $P_1, \ldots, P_N$ be a list of all possible $[T|_r]$ for $T \in X$. Given a patch P_i, let N_i be the neighborhood of the origin such that $\epsilon \in N_i$ implies $[B_r(\epsilon)] = [B_r(0)] = P_i$. (See figure 4.3).

The approximant K_r is composed of these cells N_i. When do two cells share an edge? Pick a cell, N_i, and a tiling $T \in X$, and identify T as a covering of $\mathbb{R}^2$. Consider the set of regions of T which are g-actions of P_i. Loosely speaking, we want the union of all copies of P_i throughout T. In the center of each copy is a region congruent to N_i. Let the union of these be denoted $\{\hat{N}_{i_k}\}$.

We identify an edge between two cells N_i and N_j in K_r if any elements of

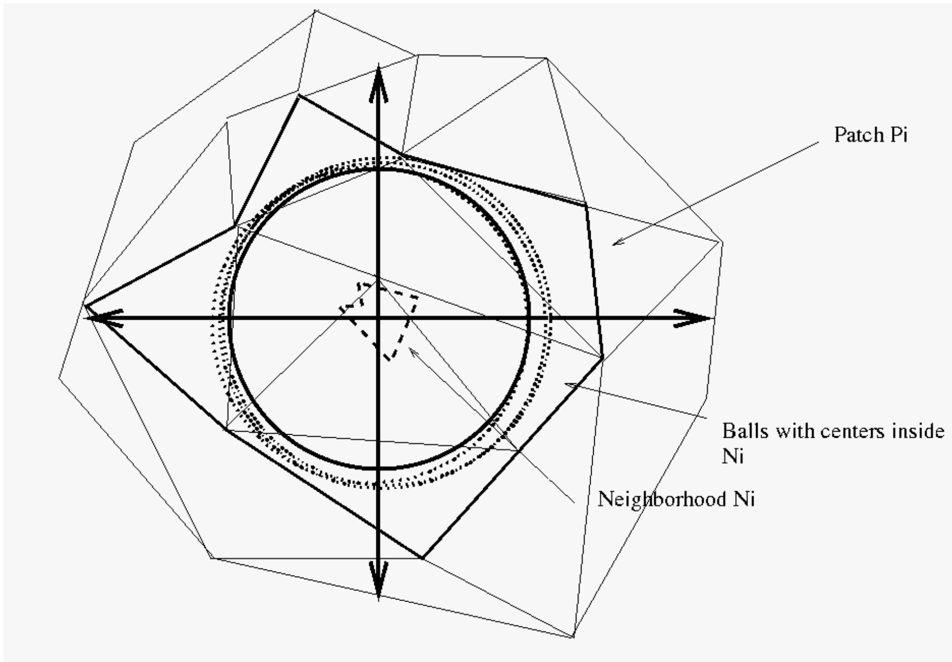

Figure 4.3: $[B_r(x)]$ is invariant for all x in the center cell N_i

the sets $\{\hat{N}_{i_k}\}$ and $\{\hat{N}_{j_l}\}$ share an edge as regions of $\mathbb{R}^2$. Notice that this implies that the corresponding patches P_i and P_j overlap to a very large degree. (See figure 4.4.)

In this way, the approximants form a branched manifold (see figure 4.5):

For points where the manifold branches, there is always a well defined tangent space. Suppose $x \in \partial N_i \cap \partial N_j$ in some K_n. Then there is some region of a tiling in X where two regions share an edge, one an element of $\{\hat{N}_{i_k}\}$, and the other an element of $\{\hat{N}_{j_l}\}$. These regions represent choices of origin such that $P_i = [B_n(x)]$ and $P_j = [B_n(y)]$ for $x \in \hat{N}_i, y \in \hat{N}_j$. Therefore P_i and P_j overlap up to their outermost tiles. Thus taking a chart of radius $< n - t$, where t is the maximum width of a single tile, will project onto the approximant K_n. Hence K_n has a well-defined tangent space at all points, lifted from the tiling space.

Half-hex tiling

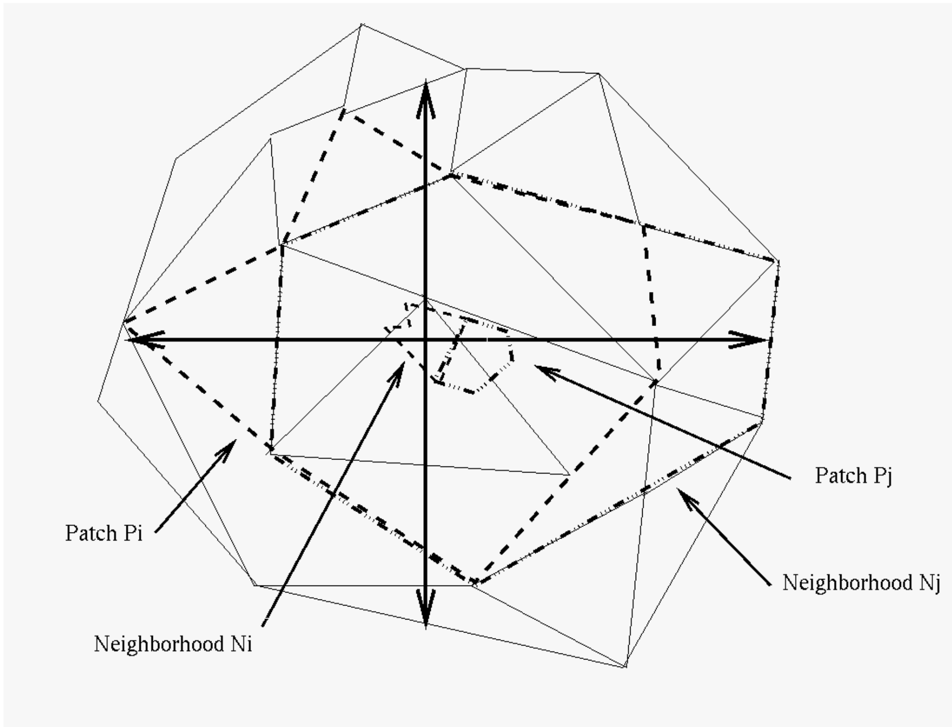

Figure 4.4: If two center cells are adjacent, their corresponding patches overlap quite a bit

What are the approximants in the case of the half-hex tiling space? This tiling space has much additional structure that cause the approximants to be particularly tidy to describe. We switch to a different cofinal set of the original $\{K_r\}_{r\subset\mathbb{R}^2}$, specific to the nature of the half-hex tiling. The half-hex tiling space is a substitution tiling space. This is a class of tiling spaces with significant hierarchical structure.

Definition 18 *$\mathcal{T}$ is a self-similar tiling if there exists a linear expansion $\phi : \mathbb{R}^n \to \mathbb{R}^n$ such that*

1. $\phi(t)$, the image of the support of any tile $t \in \mathcal{T}$, is composed of unions of tiles in $\mathcal{T}$.

2. If for two tiles in $\mathcal{T}$, $t_1 = g \cdot t_2$ for $g \in G$, then $\phi(t_1) = (\phi \circ g \circ \phi^{-1}) \cdot \phi(t_2)$ as collections of tiles in $\mathcal{T}$.

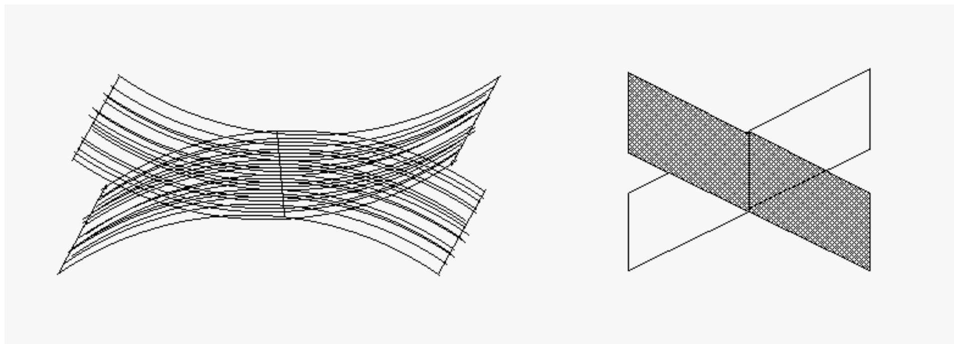

Figure 4.5: In the approximants, the surfaces branch smoothly, as shown on the left, as opposed to with corners, as shown on the right.

There is an associated map σ called the substitution map. Informally, it "inflates and chops": each tile t is associated with a patch of tiles, $\sigma(t)$ whose union is exactly the region $\phi(t)$. (See figure 4.7 for the substitution map specific to the half-hex tiling.) The substitution map σ can be applied to an entire tiling T. A self-similar tiling is a fixed point of a substitution map. A *substitution tiling space* is a tiling space that contains a self-similar tiling. (Recall the Fibonacci tiling example - the Fibonacci tiling space is a 1-dimensional substition tiling space, and the word generated by iterating σ^2 on $b.a$ is a self-similar tiling.)

Note that if T is self-similar, with corresponding space X_T, then we have $T' \in X_T$ implies $\sigma(T') \in X_T$. However, T' itself is not a self-similar tiling, because it is not a fixed point under the substitution map, whereas T does have the property that $\sigma(T) = T$. The map σ is invertible, a property known as *unique decomposition.* It follows that the supertile $\phi(t)$ is interchangeable with the patch $\sigma(t)$. We will exploit this fact in designing the approximants K_i.

Next, we abandon the construction of using radius to base our patches on in forming K_i. Any increasing region that will eventually cover the plane is also cofinal under the original definition of a tiling space inverse limit structure. This particular construction is due to Anderson and Putnam, [AP]. It is a general construction of the inverse limit structure of any substitution tiling. Let the tiles be labeled $\{t_1, \ldots, t_n\}$. Apply the inflation map ϕ to the basic tiles. We call the images

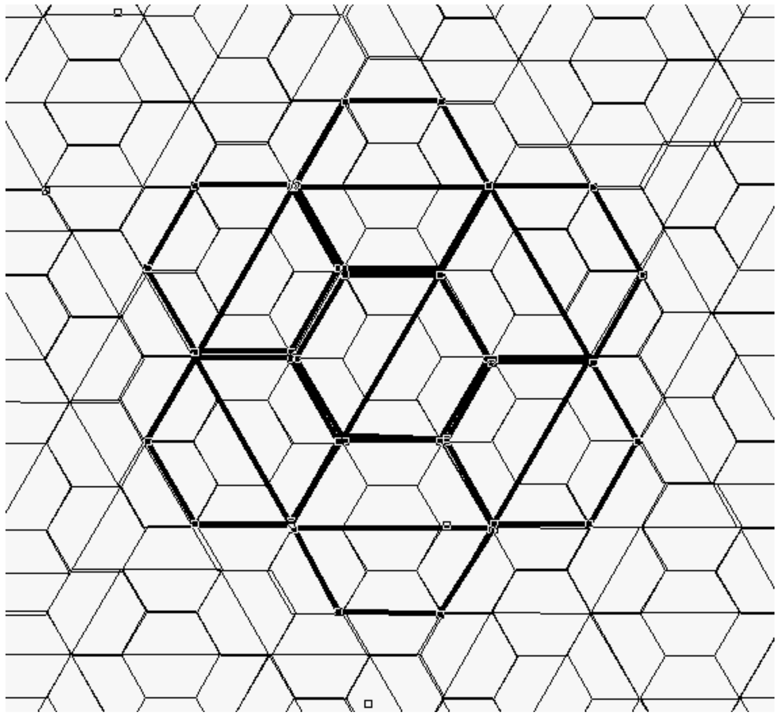

Figure 4.6: Under the inflation map, edges of the half-hex tiling line up with itself

under the linear expansion, $\phi(t_i)$ *supertiles*. Applying ϕ successively, we call the n^{th} image $\phi^n(t_i)$ an n^{th}-order supertile. Before, we determined a tiling by taking larger and larger regions $[B_r(0)]$. Here, we describe the tiling by the sequence of n^{th}-order supertiles containing the origin. Precisely, the n^{th} approximant is a description of where the origin sits in an n^{th}-order supertile, along with the supertile's rotational orientation.

As before with our center cells N_i, we identify an edge of $\phi(t_i)$ and $\phi(t_j)$ if the two patches meet anywhere in the tiling space. Our n^{th} approximant will be built out of $\sigma^n(t_i)$.

There is a hidden technicality: will applying the substitution map necessarily result in patches which cover $\mathbb{R}^2$? In other words, do the successive images of applying the substitution map necessarily produce a set of regions which is cofinal in our original index set I?

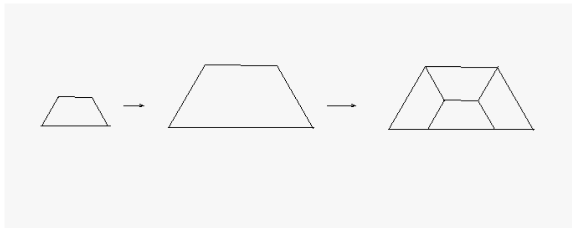

Figure 4.7: The map σ, applied to the trapezoid tile

The answer is no. Applying the substitution map can result in an infinite order supertile which does not cover the plane - for example, the region could be a quadrant. However, if all infinite order supertiles determine unique tilings of the whole plane, then we can determine the element of the inverse limit spaced uniquely from any infinite order supertile, and vice versa. The half-hexagon tiling space has this property, called "forcing the border". (There is a clever process developed by Anderson and Putnam, [AP], where the tiles are relabeled in a way that takes into account neighboring tiles, called "collaring" the tiles. This new tiling is MLD to the uncollared tiling, and so equivalent in the strongest sense. The collared tiling has the property that successive applications of σ will necessarily cover the plane. Thus this technicality never prevents any examples from going through.)

So, we have built K_n out of $\{\phi^n(t_i)\}$. As before, we identify edges of $\phi^n(t_i)$ and $\phi^n(t_j)$ if these two supertiles meet anywhere in the tiling space. Since $\phi^n(t_i)$ is just a linear expansion of $\phi^{n-1}(t_i)$, we have that $K_n = \phi(K_{n-1})$. This incredible simplicity in describing the approximants and the maps between them makes computations involving substitution tilings spaces some of the most studied and understood categories of tiling spaces. See figure 4.8 for the approximant K_n of the half-hex tiling.

Identifying the inverse limit structure of a tiling space allows us facility in computing cohomology that otherwise would be overly cumbersome. An approximant has a natural cell structure, namely the tiles; its cohomology is easy to compute. If the approximants are controlled in some way, such as having a substitution tiling space, then we may easily be able to compute the direct limit of the

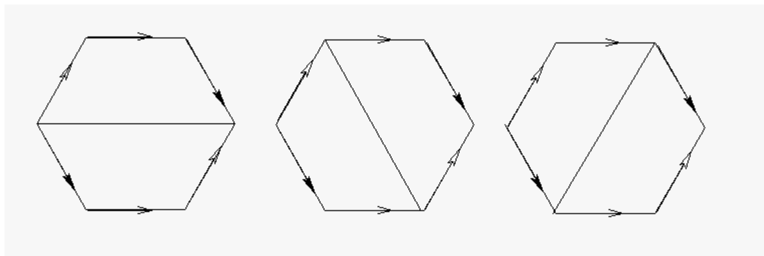

Figure 4.8: The approximant K_n of the half-hex tiling, where parallel edges with arrows are identified

cohomologies of the approximants.

Theorem 19 $\check{H}^*(\varprojlim K_r) = \varinjlim H^*(K_r)$

Proof: [AP]

The cohomology of the half-hex tiling space

We make the following computations of cohomology. Note at this point we are taking coefficients in $\mathbb{Z}$, since the approximants are themselves CW-complexes. Later, in the discussion of pattern-equivariant cohomology, we will switch to real coefficients because we will be using forms. For a discussion of pattern-equivariant cohomology with integer coefficients, see Sadun [S2].

Using the CW-complex structure, we compute:

$$H^k(K_r) = \begin{cases} \mathbb{Z} & \text{if } k = 0, \\ \mathbb{Z}^2 & \text{if } k = 1, \\ \mathbb{Z}^3 & \text{if } k = 2. \end{cases}$$

Next, we need to know how the generators of H^i pull back under the substitution map ρ: $\rho(K_r) =$

$$\rho^0 = \text{ identity on } H^0(K_r)$$

$$\rho^1 = \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix},$$

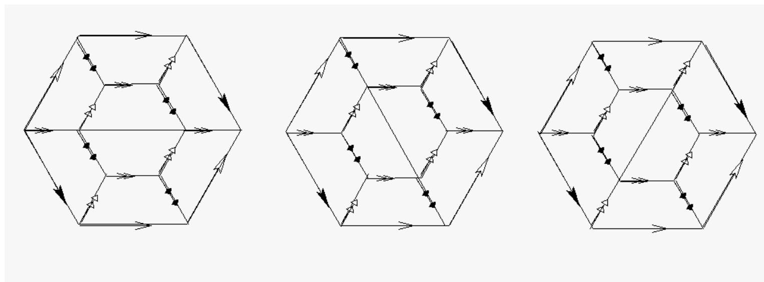

Figure 4.9: The approximant K_r after applying the substitution map ρ

$$\rho^2 = \begin{pmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{pmatrix},$$

and so, taking the direct limit we have:

$$\varinjlim H^*(K_r) = \begin{cases} \mathbb{Z} & \text{if } k = 0, \\ \mathbb{Z}[1/2]^2 & \text{if } k = 1, \\ \mathbb{Z}[1/4] \oplus \mathbb{Z}^2 & \text{if } k = 2. \end{cases}$$

4.1 Brief recap of introductory material

At this point we have developed three perspectives: that of an individual tiling, the tiling space, and the inverse-limit space structure. Research is done in all three settings, and with the appropriate change in terminology, theorems can be largely recast equivalently between any of them. We point out the key links between the three.

In the inverse limit space, a point in an approximant represents a patch, $[B_i(0)]$. In the tiling space, specifying a patch $[B_i(0)]$ is equivalent to singling out the vertical's worth of tilings in the Cantor direction which all feature the patch $B_i(0)$ around their origin. With a single tiling, we rephrase $[B_i(0)]$ as a patch, P_i, which may or may not be located at the origin. The vertical of the tiling space, or the point in the approximant of the inverse limit space, is equivalent to the set $\{[B_i(x)] : x \in \mathbb{R}^2,\ [B_i(x)] = g \cdot P_i,\ g \in G\}$.

Chapter 5

Translationally-equivariant cohomology

5.1 Basics

Pattern-equivariant cohomology was developed by Ian Putnam and Johannes Kellendonk in 2003 ([K], [KP]) Their model is of a single tiling T, with group $G = \mathbb{R}^2$, strictly translations.

Definition 20 (pattern-equivariant) *Let $f : \mathbb{R}^2 \to Y$. Define f to be pattern-equivariant with respect to T if there exists a radius r such that $[B_r(x)] = [B_r(y)] + d$ implies $f(x) = f(y)$.*

Pattern-equivariant functions are those which have a finite "sight" - there exists a radius r, such that we can list all patches of that radius, $P_1, \ldots P_N$, and knowing how the function evaluates on these patches is sufficient to tell the whole story. Representing a k-form on $\mathbb{R}^n$ as $\sum_{i_1,\ldots,i_k} f_{i_l,\ldots,i_k} dx_{i_1} \ldots dx_{i_k}$, define a form ω to be pattern-equivariant if the functions $f_{i_1,\ldots,i_k}$ are pattern-equivariant.

Let k-forms on $\mathbb{R}^2$ be denoted $\Omega^k(\mathbb{R}^2)$, and the subset of forms that are pattern-equivariant with respect to the tiling T be denoted $\Omega^k_T(\mathbb{R}^2)$. Note that if ω is pattern-equivariant, then $d\omega$ is also pattern-equivariant.

Thus this complex is well-defined:

$$\Omega^0_T \xrightarrow{d_0} \Omega^1_T \xrightarrow{d_1} \Omega^2_T.$$

(A priori this need not be limited to tilings of dimension 2, but the following analysis and interpretation is specific to this case.) Hence we have:

Definition 21 *The pattern-equivariant cohomology of T, $H^k(T) = (ker\ d_i/(im\ d_{i-1})$, where d is the standard exterior derivative map restricted to Ω^i_T.*

The major theorem established by Putnam and Kellendonk says that this cohomology is isomorphic to the Čech cohomology of the tiling space X containing T.

Theorem 22 $H^k(T) \simeq \check{H}^k(X)$.

We offer a different proof which, in its run, recasts pattern-equivariance in the tiling space context, which will be a nice foundation for later on.

Note that a tiling space X has a well-defined tangent space on the leaves, which are locally isomorphic to G ($= \mathbb{R}^2$ in this case). If T is a tiling in X, let $\mathcal{T}_X(T)$ be the tangent space to X at T. In the case where two tilings, T_1 and T_2 are in the same vertical, there is a canonical identification of the tangent spaces, $\mathcal{T}_X(T_1)$ and $\mathcal{T}_X(T_2)$ at those tilings. Since T_1 and T_2 agree out to some fixed radius r, then $(T_1 + \epsilon)|_{r-\epsilon} = (T_2 + \epsilon)|_{r-\epsilon}$ Each point in the neighborhood of T_1 can be identified with the vertical of Cantor length $r - \epsilon$ containing it. Thus a chart on a leaf of X of radius r can be extended vertically in the Cantor direction by at least $1/2r$. Any two tilings in the same vertical have canonically identified tangent spaces.

We can define a form on X to be a form on the leaves of X, and define pattern-equivariance as follows:

Definition 23 *A form ω on X is pattern-equivariant if there exists a radius r such that for $T_1, T_2 \in X$, $T_1|_r = T_2|_r$ implies $\omega_{T_1} = \omega_{T_2}$ as maps on $\mathcal{T}_X(T_1) = \mathcal{T}_X(T_2) \to \mathbb{R}$.*

Notation: Let the pattern-equivariant cohomology of X be denoted $H^i_{pe}(X)$.

Lemma 24 $H^k(T) = H^k_{pe}(X)$.

Since patches in a single tiling T are in one-to-one correspondence with verticals in the tiling space X containing T, we have that pattern-equivariant forms on T are in one-to-one correspondence with pattern-equivariant forms on X. In either context, there exists a radius r, and hence a finite list of patches $P_1, \ldots, P_N$. It is sufficient to know how a form evaluates on these patches to uniquely identify it with a form on the tiling or on the space. □

Consider the inverse limit structure. A point in the K_r^{th} approximant represents some $[B_r(0)]$. Knowing how a function evaluates on all patches of radius r is equivalent to having a function on the r^{th} approximant, K_r. (Recall that the approximants K_r are smooth branched manifolds, so this is well-defined.)

Thus pattern-equivariant functions are in one-to-one correspondence with functions on a finite approximant K_r. By the earlier theorem, $\check{H}^*(X) = H^*(\varprojlim K_r) = \varinjlim H^*(K_r)$. However, this theorem was stated in general terms, as a property of inverse limit spaces. Computing $H^*(K_r)$ requires that we have a well-defined Čech-DeRham theorem for branched manifolds. As clarified above, the branched manifolds in question have well-defined tangent spaces at the branch points. For a full proof see [S2].

What remains is to show that $\varinjlim H^*(K_r) = H_{pe}(X)$. Let $[\omega] \in H^*_{pe}(X)$, and ω a representative of $[\omega]$. Then $\exists r$ such that if $T_1|_r = T_2|_r$, then $\omega_{T_1} = \omega_{T_2}$. Therefore, for $s > r$, $T_1|_s = T_2|_s$ also implies that $\omega_{T_1} = \omega_{T_2}$. Thus, ω can be identified with a form on a finite approximant K_r, and also induces identification with forms on all approximants K_s, $s > r$. In particular, we identify ω with its pullbacks under the inclusion map $\rho_{s,r} : K_s \to K_r$. This is precisely the direct limit, $\varinjlim H^*(K_r)$. □

Chapter 6

Introduction to Rotations

There are several important motivations for incorporating rotations into the structure of pattern-equivariant cohomology. First, consider the half-hex tiling. If the goal is to understand the structural elegance of the tiling space, then we want to utilize its rotational nature. In other words, in figure 6.1 we have one tile in six orientations, not six coincidentally-shaped tiles.

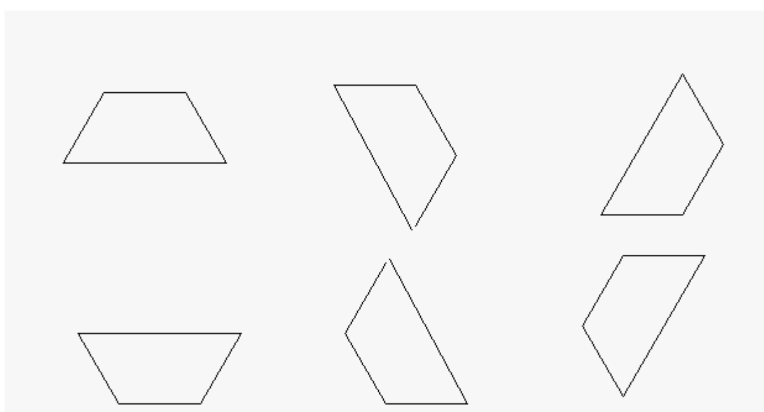

Figure 6.1: One tile in six orientations, as opposed to six independent tiles

The second motivation arises from considering tilings like the pinwheel tiling, (see figure 6.2). Like the half-hex tiling, the pinwheel tiling is also a substitution tiling. The substitution rule is shown in figure 6.3.

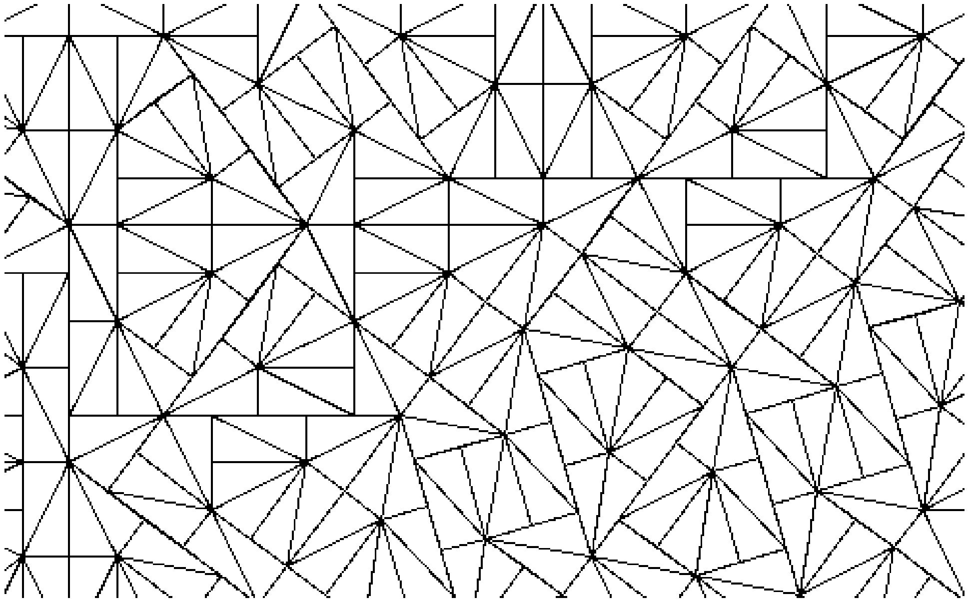

Figure 6.2: A portion of the pinwheel tiling

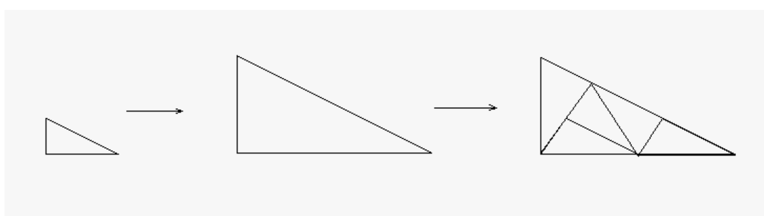

Figure 6.3: Substitution rule for the pinwheel tiling

Notice (in figure 6.3) that the interior-most triangle is rotated by the irrational angle $\tan^{-1}\sqrt{1/2}$ with respect to the original triangle. As the substitution is applied successively, every patch will occur in infinitely many orientations, the closure of which is S^1. Tilings like the pinwheel fail to have the property of finite local complexity if $G = \mathbb{R}^2$, and hence cannot be studied within the previous pattern-equivariant framework.

Tiling spaces with infinite rotations are very poorly understood. The cohomology of the pinwheel tiling has not been computed. There are no other examples of two-dimensional tilings with infinite rotations that are not substitution tilings. The few existing examples are variations on the pinwheel tiling [R], and substitution tilings with fractal boundaries [cite]. Since these are all substitution tilings, they all have substantial hierarchical structure, and so they are not typical.

The global structure of a tiling space with patches occuring in infinitely many rotations is not known. Any tiling $T \in X$ is a tiling of $\mathbb{R}^2$, but there are three continuous leaf directions: dx, dy and $d\theta$. Thus these tilings spaces are composed of two-dimensional tilings forming three-dimensional leaves. Above any tiling is an S^1 fiber of rotations of that tiling. This helps explain the gulf in complexity between tiling spaces without rotations and when $G' = S^1$.

By understanding rotations in the finite case as a property intrinsic to the tiling space, (instead of retyping them as translational tiling spaces) we hope to lay a groundwork which will guide future understanding of the infinite case.

Chapter 7

Pattern-equivariant Cohomology with Finite Rotations

The following definitions are the heart of the paper, as they forge new territory in an original way. Rotations are interesting because they are qualitatively different from translations. Translational equivariance is a statement about the Cantor direction. Rotations are built solely into the leaf direction. We will discuss the ways in which the vertical structure is tied unexpectedly to the local leaf structure.

That said, it is not canonical how we might include rotations. Do we want a pattern-equivariant form to be one that evaluates the same on all rotations of a tiling? In other words, if $T_1 = \theta T_2$, then $\omega_{T_1} = \omega_{T_2}$? If translationally-equivariant meant forms were constant on verticals in the cantor direction, perhaps we want to maintain that key distinction and disregard all leaf structure. In this case, ω evaluates independently on θT for different choices of θ.

We develop an innovative structure that is large and flexible enough to encompass all possibilities.

Definition 25 *Let the rotation group G' be the set of rotations on $\mathbb{R}^2$ such that $\theta \in G', T \in X$ implies $\theta T \in X$.*

Definition 26 *Let ω be a form on X with coefficients in the vector space Y. Pick a representation $\rho : G' \to End(Y)$. We say a ω is pattern-equivariant with respect to ρ if $\exists$ r such that $T_1|_r = \theta T_2|_r$ implies $\omega(T_1) = \rho(\theta)\omega(T_2)$.*

What do we mean by $\omega(T_1) = \rho(\theta)\omega(T_2)$? As maps on tangent spaces, $\omega(T_1)$ and $\omega(T_2)$ have as their domains $T_{T_1}(X)$ and $T_{T_2}(X)$ respectively. But we can canonically identify them using the rotation of θ, and so we can speak of $\omega(T_1)$ and $\omega(T_2)$ being equivalent maps.

Note this definition parallels our definition for translational PEQ cohomology on a tiling space, but breaks from the definition of translational PEQ cohomology on a single tiling $T \simeq \mathbb{R}^2$ as put forth by [K] and [KP]. In their original set-up, pattern-equivariance is defined on functions, and then a pattern-equivariant form is one in which the coefficient functions are pattern-equivariant. Our coefficient functions are not pattern-equivariant in a way that descends to functions on $\mathbb{R}^2$. This is because of how we are choosing to identify the tangent spaces of T and θT. We are not identifying the directional derivative dx_T with $dx_{\theta T}$, as one would if working with a patch P occuring in multiple rotations throughout $\mathbb{R}^2$. Instead we identify a directional derivative with respect to the pattern that occurs at that point - see figure 7.1 for illustration.

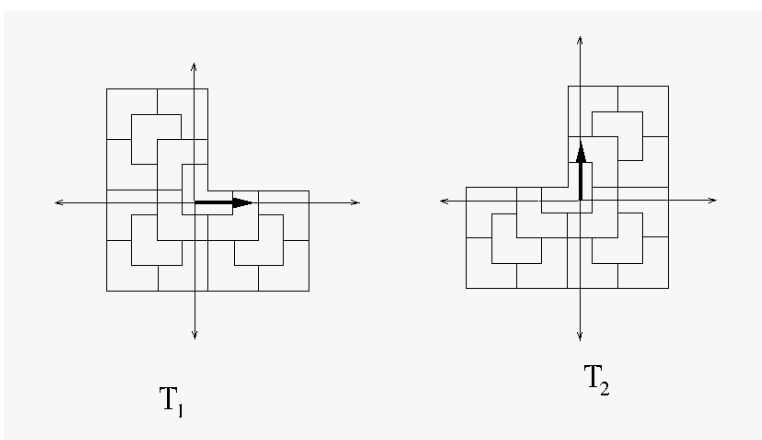

Figure 7.1: Directional derivatives of $T_T(X)$ and $T_{\frac{\pi}{2}T}(X)$ canonically identified with respect to the pattern

This choice of identification of $T_T(X)$ and $T_{\theta T}(X)$ more elegantly captures the rotational structure of X, and allows the results in chapter 7.3 to be phrased with ease and simplicity.

We can rephrase the definition in terms of forms on the approximants, as we did in the translation case. We can also apply this representation condition to cochains, and utilize the CW-structure of the approximants in computing cohomolgy.

Definition 27 *Let H^*_ρ be the cohomology found by restricting to forms in Ω_ρ.*

Typical choices for Y are $\mathbb{R}^d, \mathbb{C}^d$ and $\mathbb{Z}^d$. ($\mathbb{Z}^d$ is a module, not a vector space, but we can still speak of a map $\rho : G' \to End(\mathbb{Z}^d)$.) $\mathbb{Z}^d$ is convenient when considering the half-hex example because previous computations relied on the CW structure of the approximants. De Rham cohomology is the framework chosen by Putnam and Kellendonk, and so there are occasions to work over $\mathbb{R}^d$ or $\mathbb{C}^d$. Very often we will pass back and forth between cochains and forms, depending on whether we are emphasizing theoretical continuity with past work in De Rham cohomology, or ease of computations for the purposes of examples.

Recall that translationally-equivariant cohomology was shown to be isomorphic to Čech cohomology. Since Čech cohomology is a topological invariant, translationally-equivariant cohomology is invariant under a meaningful set of maps, and therefore worth studying. This new cohomology is clearly no longer Čech. Under what set of conditions is this cohomology invariant? In particular, is it a topological invariant?

We show that when the rotation group G' is finite, then H^*_ρ is invariant under homeomorphisms that commute with G'. Let X and Y be two tiling spaces with the same finite rotation group, G', and fix a representation ρ of G'.

Definition 28 *A continuous map $f : X \to Y$ is a* **factor map** *if f commutes with the action of G, the full group of rigid motions including both translations and rotations. A* **rotational-factor map** *commutes with only rotations.*

Definition 29 *A map f* **preserves verticals** *if there exists r, s such that $T_1|_r = T_2|_r$ implies $f(T_1)|_s = f(T_2)|_s$.*

Note that continuity applies both in the leaf and Cantor directions. If f is both continuous and a vertical-preserving map, then the given r and s are uniform bounds for the entire tiling space. By continuity in the Cantor direction, we also have that as $r \to \infty$, $s \to \infty$. The larger the ball on which T_1 and T_2 match, the larger the ball on which $f(T_1)$ and $f(T_2)$ will match. Because tiling spaces are compact, we can phrase it as uniform continuity in preserving verticals, (which will be useful): for every s $\exists$ r such that $T_1|_r = T_2|_r$ implies $f(T_1)|_s = f(T_2)|_s$.

Theorem 30 *$H *_\rho (X)$ is invariant under homeomorphisms that commute with rotations.*

The proof of this theorem is the heart of this paper. It is long and involved, and so we take a moment to outline the steps we will follow.

1. First, we show that rotationally pattern-equivariant forms pull back under smooth vertical-preserving rotational-factor maps.

2. If F is a smooth, vertical-preserving rotational-factor homotopy $F : [0,1] \times X \to Y$, then $F_0^*(H_\rho^k(Y) = F_1^*(H_\rho^k(Y)$.

3. If $f : X \to Y$ is a continuous map, then f can be approximated by a smooth, vertical-preserving map from X to Y that is homotopic to f. If f is a rotational-factor map, f_ϵ will also be a rotational-factor map.

4. If $F : [0,1] \times X \to Y$ is a continous rotational-factor homotopy, and F_o and F_1 preserve verticals, then $F_0^*(H_\rho^k(Y) = F_1^*(H_\rho^k(Y)$.

5. If $f : X \to Y$ is a continuous rotational-factor map, we define the pullback $f^*(H_\rho^k(Y))$ to be the pullback under f_ϵ, a smooth, vertical-preserving approximation which is homotopic to f.

6. This is well-defined: If f_ϵ and f_δ are two such approximations, then they are homotopic to each other, and hence $f_\epsilon^*(H_\rho^k(Y)) = f_\delta^*(H_\rho^k(Y))$.

7. If $h : X \to Y$ is a homeomorphism, then the pullback induced on cohomology is an isomorphism.

We take a moment to note that the proof that De Rham cohomology is a topological invariant follows the same template. First it is shown that forms pull back under smooth maps and behave nicely in smooth homotopies. Then using the fact that smooth functions are dense in continuous functions, it is shown that forms must behave nicely under continuous homotopies. The De Rham proof hinges on the approximation of continuous functions by smooth functions. This approximation theorem has much larger scope and impact in mathematics than this one application.

Likewise, step 3 approximates a continuous map on tiling spaces by a local map on tiling spaces. It follows in the tradition of these other theorems, like the one in the De Rham proof, or theorems approximating continuous functions by step functions, or L_1 functions by continuous functions with compact support. The importance of the approximation theorem in Step 3 reaches beyond pattern-equivariant cohomology.

We now prove Theorem 30.

7.1 Proof of Theorem 30

Lemma 31 *Let $f : X \to Y$ be a smooth vertical-preserving map. Suppose ω is PEQ with respect to ρ on Y. Then $f^*(\omega)$ is PEQ on X.*

Proof: Suppose ω has radius s on Y. Then for $T_1, T_2 \in Y$, $T_1|_s = \theta T_2|_s$ implies $\omega(T_1) = \rho(\theta)\omega(T_2)$. In X, there exists r such that if two tilings agree on a radius of r will map under f to tilings that agree out to s in Y. Finally, G' is discrete, so we can take r large enough to ensure that the distance between $d(f(T_2), f(\theta T_2)) < s < \frac{2\pi}{n}$, where $n = |G'|$, and hence the angle $= 0$.

So if $S_1|_s = \theta S_2|_s$ in X, then $f^*(\omega)(\theta S_2) = \omega(\theta f(S_1)) = \rho(\theta)\omega f(S_1) = \rho(\theta) f^*(\omega)(S_1)$. □

Lemma 32 *Suppose $F : [0,1] \times X \to Y$ be a homotopy such that all F_t are vertical-preserving maps, and F is smooth in t. For fixed ρ, $F_0^*(H_\rho^*(Y)) = F_1^*(H_\rho^*(Y))$.*

We want to show that $F_0^*([\omega]) - F_1^*([\omega]) = 0$, where ω is an equivalence class in $H_\rho^*(Y)$, or that $(F_0 - F_1)(\omega) \in$Im d, where $\omega \in [\omega]$. Let $\Phi(T) = \int_0^1 (\omega F_t(T))dt$. Then

$$d\Phi = d\int_0^1 (\omega F_t(T))dt = \int_0^1 d(\omega F_t(T))dt = \omega(F_1(T)) - \omega(F_0(T))$$

And so,

$$(F_1(T) - F_0(T))^*(\omega) = d\Phi(\omega) \in \text{ Im } d.$$

The question then is: does Φ qualify as a pattern-equivariant function? Let ω have pattern-equivariant radius R in Y, and suppose two tilings in X that match out to a distance S will be sent under all F_t to tilings that match out to at least R in Y. Suppose T_1 and θT_2 are two tilings in X that agree on a radius S. Then

$$\Phi(T_1) = \int_0^1 (\omega(F_t(T_1)))dt$$

$$= \int_0^1 (\omega(F_t(\theta T_2)))dt = \int_0^1 \rho(\theta)(\omega(F_t(T_2)))dt$$

$$= \rho(\theta)\int_0^1 (\omega(F_t(T_2)))dt = \rho(\theta)\Phi(T_2)$$

□

Theorem 33 *If f is a continuous map from $X \to Y$, then for every ϵ sufficiently small, there exists $\hat{f}_\epsilon$ with the following properties:*

1. $\hat{f}_\epsilon$ is a smooth, vertical-preserving map from $X \to Y$

2. $f(T)$ and $\hat{f}_\epsilon(T)$ are in the same translational leaf, and $\left|f(T) - \hat{f}_\epsilon(T)\right| < \epsilon$, where $\left|f(T) - \hat{f}_\epsilon(T)\right|$ denotes the translational gap within a leaf.

3 If f commutes with rotations, then $\hat{f}_\epsilon$ will also commute with rotations.

Since f is continuous and G' is discrete, f sends translational leaves to translational leaves. Under f, a vertical may get "sheared", (see figure 7.2):

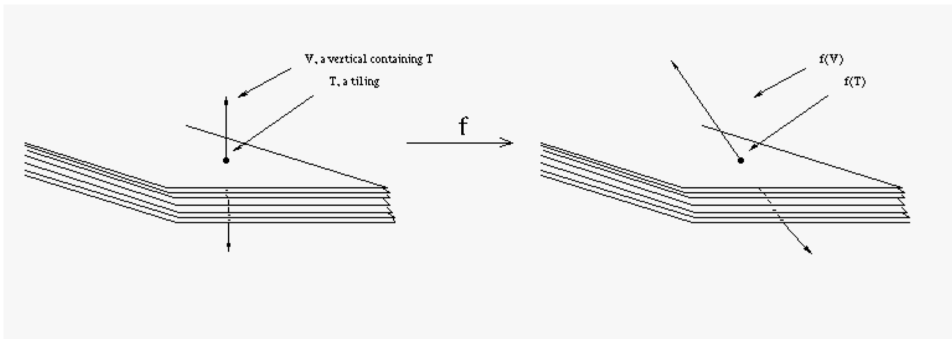

Figure 7.2: Verticals may be sheared under a continuous map

By continuity, if r is large and $T_1|_r = T_2|_r$, then $f(T_1)|_s = gf(T_2)|_s$, where $|g|$ is small. In other words, $f(T_1)$ and $gf(T_2)$ are in the same vertical. Our strategy will be to define $f_\epsilon(T_2) := gf(T_2)$, thereby forcing the image of verticals to be preserved. Loosely speaking, the plan is to comb the sheared images of verticals into nice straight verticals.

Proof:

Fix $\epsilon > 0$, and an anchoring tiling $T_o \in X$. Let $\delta > 0$ such that $d(T_1, T_2) < \delta$ implies $d(f(T_1), f(T_2)) < \epsilon$.

First we'll define $f_\epsilon(T_o) := f(T_o)$. Next we define B_o:
$B_o :=$ { a maximally connected set of translates of T_o such that no two tilings $\in B_o$ match on a radius of $1/\delta$ around the origin, independent of rotation. }

The set B_o is a maximally connected region of representatives of unique verticals. If T_o+d_1 and T_o+d_2 are in B_o, then $(T_o+d_1)|_{1/\delta} \neq \theta(T_o+d_2)|_{1/\delta}$ for any angle θ.

We then form

$$D_o := B_o \cup \{T = \theta T' : T' \in B_o\}.$$

Intuitively, D_o is the collection of tilings in B_o, times the rotation group. Note the following subtlety: a tiling T in B_o corresponds to a unique patch, $P = T|_r$. In forming D_o, we are not taking rotations of the patch P, but rotations of the original tiling $T \in B_o$. Therefore, unless T is a tiling of rotational symmetry, $\theta T \in D_o$ is again a unique representative of a new vertical. In the case that T is a tiling of rotational symmetry, we will not include duplicate copies of the same tiling. This is because we are forming D_o as a subset of X, not $B_o \times G'$.

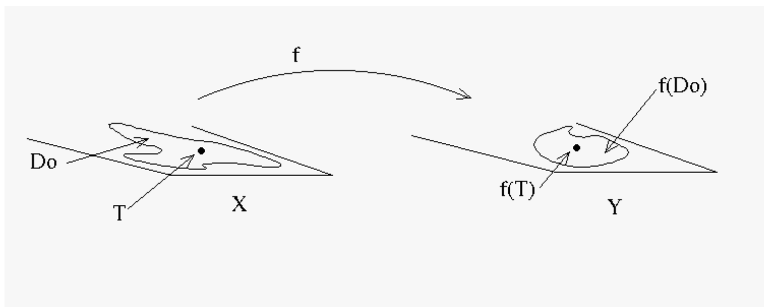

Figure 7.3: On D_o, the new function f_ϵ is defined by the original f

Define $f_\epsilon := f$ for any tiling in D_o, (see figure 7.3).

Let C_o be the cylinder with base D_o and height ϵ, defined as:

$$C_o := \{T' \in X : T'|_{1/\epsilon} = T|_{1/\epsilon}, \text{ for some } T \in D_o\}.$$

C_o is the collection of of tilings in the verticals of length ϵ with representatives in D_o. We will forcibly map C_o onto the cylinder with base $f(D_o)$, vertical to vertical.

Formally, we extend f_ϵ to C_o, as follows: Let $T' \in C_o$. Then T' is in the vertical of some element of D_o. By definition, every element of D_o can be written as $g \cdot T_o$, where T_o is our original anchor tiling. Therefore $T'|_{1/\delta} = g \cdot T_o|_{1/\delta}$. By continuity, $f(T')|_{1/\epsilon}$ matches $f(g \cdot T_o)|_{1/\epsilon}$ up to a ϵ-wiggle. We define $f_\epsilon(T')$ to be the tiling after "undoing" the ϵ-wiggle. Intuitively, we can think of $f_\epsilon(T') = \{$ leaf containing $f(T')$ $\} \cap \{$vertical containing $f(g \cdot T_o)\}$, but this is not quite precise, as we will see.

Definition 34 *On C_o, $f_\epsilon(T') := h \cdot f(T')$, where $h \in G$ is the unique element with $|h| < \epsilon$ such that $h \cdot f(T')|_{1/\epsilon} = f(g \cdot T_o)|_{1/\epsilon}$ for a unique $g \cdot T_o \in D_o$.*

By construction, the image of C_o under f_ϵ is a cylinder in Y with base $f_\epsilon(D_o)$ and cantor height ϵ. (See figure 7.4.)

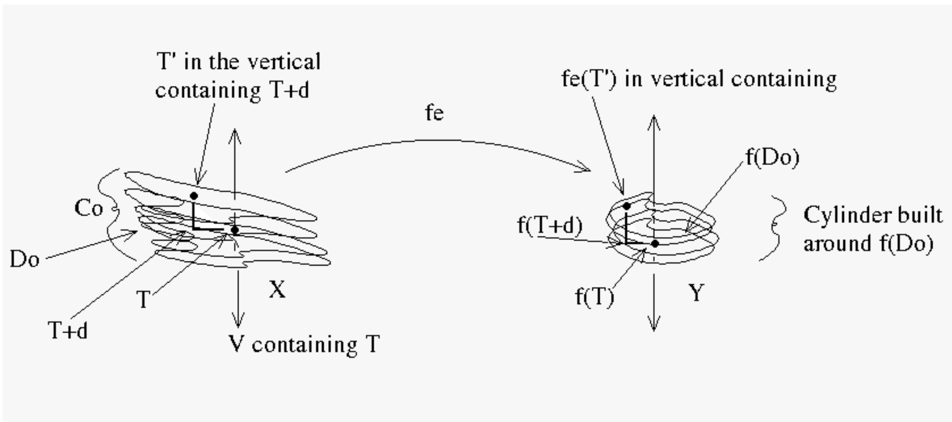

Figure 7.4: Extend f_ϵ to the cylinder C_o

Lemma 35 *f_ϵ is well defined on C_o.*

It is not obvious that the map f_ϵ is well-defined. In theory, h need not be unique. The vertical containing $f(g \cdot T_o)$ could intersect the leaf containing $f(T')$ multiple times in close proximity. This means that the patch $P = f(g \cdot T_o)|_{1/\epsilon}$ would appear more than once in the same leaf, both within a ϵ-wiggle of $f(T')$.

The patch P has radius $1/\epsilon$. Suppose two copies of P occur within 2ϵ of each other, P_1 and P_2. We look at this situation by case.

Case 1: P_1 and P_2 are related by a translation only, without any rotation. In other words, $P_1 = h + P_2$. Then they intersect on a region of radius $> 1/\epsilon - 2\epsilon$. $P_1 \cap P_2$ is periodic, with period h. If we take $2\epsilon <$ width of any tile in the space Y, then it is a contradiction for a patch to have period $h < 2\epsilon$.

Case 2: P_1 and P_2 are related by a rigid motion that involves both a rotation and a translation. (For example, suppose we are near a tiling with m-fold symmetry.

Tilings of symmetry get sent to tilings of symmetry under f, since f commutes with rotations. Therefore, while $g \cdot T_o$ is in D_o, it is in close proximity to m tilings which are rotations of $g \cdot T_o$ (and are not in D_o). Every time this large patch of symmetry occurs, there are m indistinguishable tilings around the pivot of symmetry.)

Therefore, $f(T')$ could be within ϵ of two indistinguisable patches P_1 and P_2. Suppose $\theta_1 f(T')|_{1/\epsilon} + h_1 = P_1$, and $\theta_2 f(T')|_{1/\epsilon} + h_2 = P_2$. Since the group of rotations is discrete, we know that $P_1 = \frac{2\pi}{m} P_2 + h$, $h < 2\epsilon$, for some integer m. Then $\theta_1 + \theta_2 \geq \frac{2\pi}{m}$. Since f is continuous, we can pick $\epsilon << \frac{2\pi}{m}$, and thus ensure that $f(T')$ is only within ϵ of one of the patches.

One final clarification: things do not go haywire in at the tiling of symmetry itself. As mentioned earlier, a symmetric tiling must be sent to another symmetric tiling.

□

Let's look more closely at the nature of the map f_ϵ on C_o:

1. Note that tilings on the boundary of C_o are also tilings on the perimeter of D_o: if $T' \in \partial D_o$, then there is some other tiling, T'' in D_o that matches T' exactly out to $1/\epsilon$. This is because the requirement on D_o is that it be any maximal connected set such that no two tilings match out to this radius, and so if $T' \notin D_o$ but it is in ∂D_o, this criterion must precisely fail. However, this means that T' is in the same vertical as T'' - they match precisely, and so have Cantor distance ϵ. Therefore $T' \in \partial C_o$.

2. D_o (and hence C_o) may not be convex; the translation necessary to take T_o to $T_o + D$ may pass through tilings excluded by D_o. This is not a problem - since $f|_{D_o} = f_\epsilon|_{D_o}$, we are only modifying tilings vertically.

3. Since the original map f was uniformly continuous, when we pick ϵ there is a δ for all of X such that $d(T, T') < \delta \Rightarrow d(f(T), f(T')) < \epsilon$. By definition of the distance metric, $f(T)|_{1/\epsilon} = h \cdot f(T')|_{1/\epsilon}$, $h \in G$ with $|f| < \epsilon$. Then $f_\epsilon(T') := hf(T')$, so the bound on h becomes a bound on how far apart $f(T')$ and $f_\epsilon(T')$ can be. Throughout this, we can control how much we are modifying the map f.

We now turn to the leaf containing T - we will construct a larger region containing D_o, and use it to build a larger cylinder which contains C_o. Consider how the cylinder C_o intersects the leaf containing T: (figure 7.5)

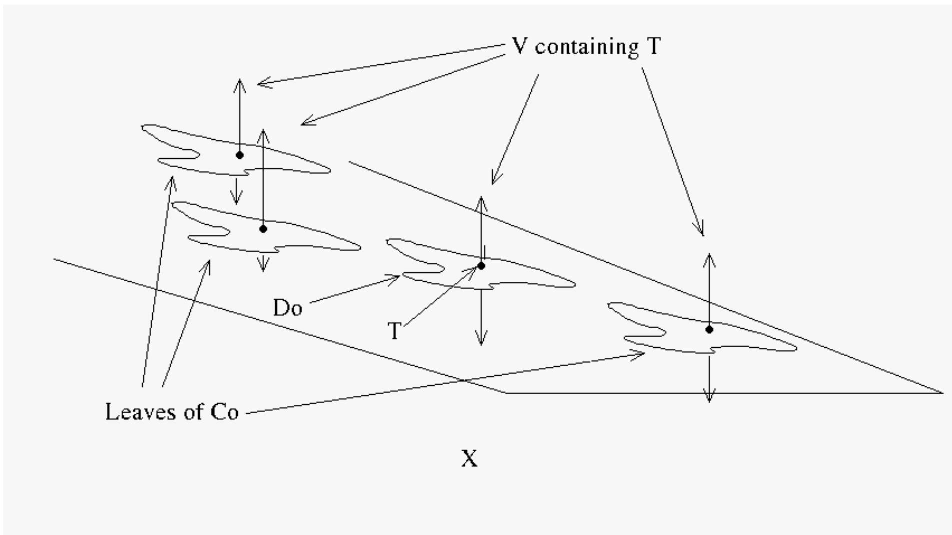

Figure 7.5: C_o intersects the leaf containing T

If a tiling is this leaf containing T and *not* in one of these regions of intersection with C_o, then this tiling must not match any tiling in D_o (or C_o) on a disk of $1/\epsilon$. In other words, this tiling is contained in a vertical on which we have not yet defined f_ϵ. We want to enlarge our base to include as many new verticals as possible.

Let $\tilde{D_o} = \{$ a connected component of C_o containing $D_o\}$.

Then let $B_1 = \{$ a maximally connected set of translates of T_o such that:

1. C_o, (and hence D_o and T_o,) has trivial intersection with B_1
2. The boundaries ∂B_1 and $\partial \tilde{D_o}$ have non-trivial intersection.
3. no two tilings in B_1 match on a disc of $1/\epsilon$ around the origin, up to any rotation. $\}$

Note that all tilings in B_1 correspond to verticals on which f_ϵ is not yet defined. As before, we then define $D_1 = B_1 \cup \{T = \theta T' : T' \in B_1\}$. D_1 is the base of our next cylinder. As before, define $f_\epsilon = f$ on D_1. (See figure 7.6.)

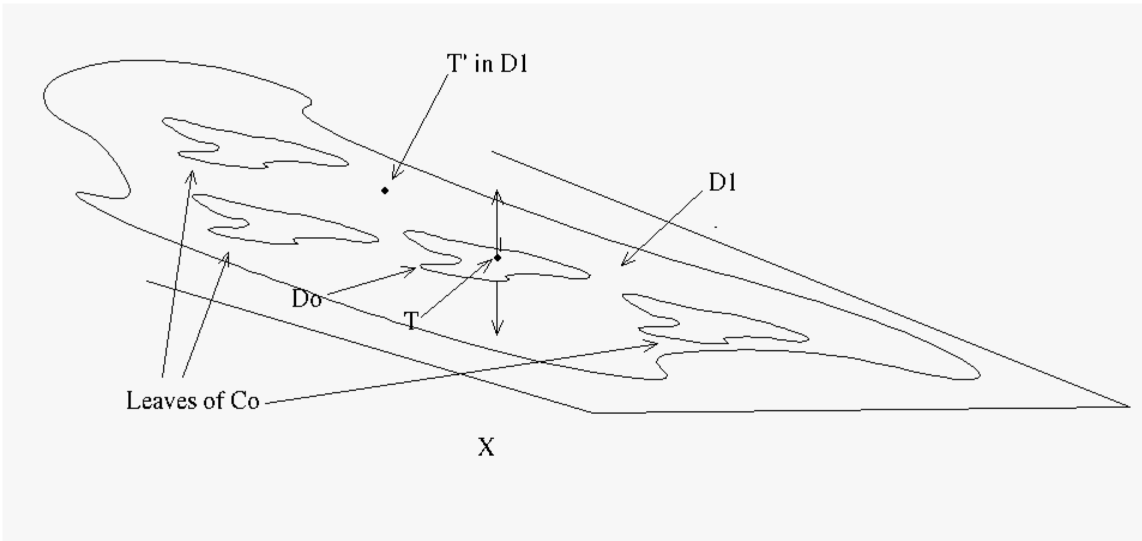

Figure 7.6: The base of the next cylinder, D_1

We now repeat the cylinder construction: Let C_1 be the cylinder of Cantor height ϵ and base D_1. If $T' \in C_1$, then $T'|_{1/\delta} = g \cdot T_o|_{1/\delta}$. Therefore, $f(g \cdot T_o)|_{1/\epsilon} = h \cdot f(T')|_{1/\epsilon}$ for some $h \in G$, $|h| < \epsilon$, and we define $f_\epsilon(T') := h \cdot f(T')$.

We now have $f_\epsilon : C_o \cup C_1 \to Y$ in a way that preserves verticals. Note f_ϵ is not continuous - we postpone this consideration for the time being.

Iterate this construction, using D_{i+1} as the base of the i^{th} cylinder. To form D_{i+1}, define as $\tilde{D}_i :=$ { a connected component of $\cup C_i$ containing T_o.} Then define: $B_{i+1} =$ { a maximally connected set of translates of T_o such that

1. $\cup C_i$ regions have trivial intersection with B_{i+1}
2. B_{i+1} has $\partial \tilde{D}_i$ as a boundary component
3. no two tilings in B_{i+1} match on a disc of $1/\epsilon$ around the origin, up to a rotation. }.

Define $D_{i+1} = B_{i+1} \cup \{T = \theta T' : T' \in B_{i+1}\}$, the union of G'-fibers over B_{i+1}. Form $C_{i+1} = \{T$ such that $T \in V_\epsilon(T')$ for $T' \in D_{i+1}$. Extend f_ϵ to C_{i+1} as before - mapping verticals onto verticals.

Theorem 36 *In a finite number of steps, f_ϵ will be defined on all of X.*

Proof: Recall from the discussion of inverse limit spaces in section 4 that we defined a neighborhood N in the center of a patch P of radius r to be the cell of possible choices of origin, such that for $x \in N$, $[B_r(x)] = P$.

We will show that if $T \in B_i$ with $[T|_r] = P$, then all verticals represented by all points of N are contained in C_i. There are finitely many patches $\{P_j\}$ of a given radius, and so finitely many corresponding cells $\{N_j\}$. However, the cell may have many components.

Lemma 37 *A cell N in an approximant K has finitely many components.*

Proof: N can have more than one connected component, as illustrated in figure 7.7. As above, let $[T|_r] = P$, and $N = \{x : [B_r(x)] = P\}$. Let $\{t_i\}$ be the

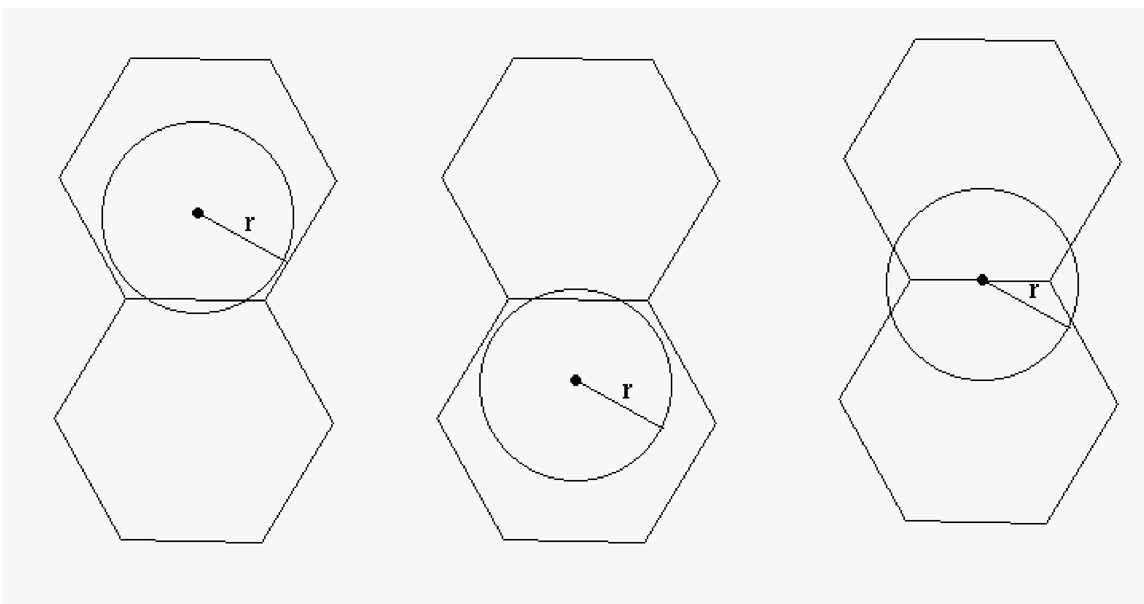

Figure 7.7: P = two stacked hexagons. The first two choices of $[B_r(x)]$ yield the union of the two hexagons, but the third choice of $[B_r(x)]$ will include adjacent cells. In this way, N has more than one component.

collection of tiles that share an external edge with P. For each of these tiles on the perimeter, define R_{t_i} and E_{t_i} as follows:

$$R_{t_i} = \{x \in P : B_r(x) \cap t_i \neq \emptyset\},$$

(See figure 7.8), and

$$E_{t_i} = \{x \in P : B_r(x) \cap (\partial t_i \cap \partial P)\}.$$

(See figure 7.9.)

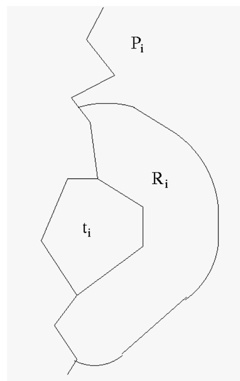

Figure 7.8: Region R_{t_i} composed of points whose $B_r(x)$ intersects t_i

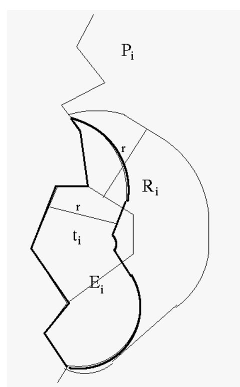

Figure 7.9: Region E_{t_i}, outlined in bold, is composed of points whose $B_r(x)$ intersects the outside edge of t_i

Claim: $R_{t_i} - E_{t_i}$ is a connected component. Notice that since t_i is a convex polygon, both R_{t_i} and E_{t_i} are composed of finitely many straight edges (corresponding to edges of t_i), and arcs of circles (corresponding to vertices of t_i.) R_{t_i} is a constant distance r from the internal edges of t_i, E_{t_i} is a constant distance r from the external edges of t_i, and the boundaries of R_{t_i} and E_{t_i} meet exactly on the boundary of P.

Form a region $R_{t_i} - E_{t_i}$ corresponding to each tile on the perimeter. The boundary of each region $R_{t_i} - E_{t_i}$ is composed of finitely many straight edges and arcs of circles, and therefore no two regions can intersect infinitely often. Since there are finitely many tiles on the perimeter of P, $\cap_i (R_{t_i} - E_{t_i})$ has finitely many components. $\Box$

Assume that N has one single component. Since we know N has finitely many components, we can repeat the following argument finitely many times if need be. We show that if a tiling representing N is in B_i, then all of N is represented in C_i.

We identify the cell N with its preimage in X. Let $\{\hat{N} =$connected regions of X containing tilings that represent $N\}$. Then $\{\hat{N}\}$ is a collection of regions congruent to N, the cell in the approximant K.

Case 1: $\{\hat{N}\} \cap B_i$ in exactly one region, and the patch P has no rotational symmetry. Since B_i is maximal, and no other copy of N appears in B_i, all of N must be included in B_i.

Case 2: $\{\hat{N}\} \cap B_i$ in exactly one region, and the patch P has rotational symmetry; $P = \theta P$. Then a maximal region will contain a θ-wedge of N. The rest of N is represented by verticals which agree up to rotation with a vertical in the θ-wedge contained in B_i. In forming D_i, we will take all rotations of the tilings in the θ-wedge of $\{\hat{N}\} \cap B_i$, which will be representatives of the rest of N.

Case 3: $P \cap B_i$ in more than one region, and P has no rotational symmetry. By maximality of B_i again, the union of these regions must represent all of N. And as in Case 2, if P has symmetry, N will be represented upon forming D_i.

Since there are finitely many cells in any approximant, and f_ϵ is defined on at least one entire cell in any iteration, f_ϵ will be defined on all of X in finitely many steps. $\Box$

As we have defined it, f_ϵ is not a continuous map. Our strategy will be to convolve with a bump function to smooth it in the leaf direction. Note that this will shorten the length of the verticals that are preserved, but still result in a vertical-preserving map. We will spell this out in full detail.

Let $h(x)$ be a bump function on $\mathbb{R}^2$ with compact support. Also let $h(x)$ be rotationally invariant: for $\theta \in G'$, let θ induce a map on $\mathbb{R}^2$, and $h \cdot \theta(x) = h(x)$. Let the support have radius which is sufficiently small, such that any ball of that radius intersects ∂D_i in at most one component. Let the radius of support be $|h|$. The actual final map we are interested in is not f_ϵ but

$$\hat{f}_\epsilon(T) = f_\epsilon(T) - \int_{\mathbb{R}^2} \left(f_\epsilon(T) - f_\epsilon(T - x)\right) h(x) dx,$$

where the subtraction $f_\epsilon(T) - f_\epsilon(T - x)$ denotes the translational gap. Note that $f_\epsilon(T) - f_\epsilon(T - x)$ and $f_\epsilon(\theta T) - f_\epsilon(\theta T - x)$ are both displacements on $\mathbb{R}^2$. We identify T with $\mathbb{R}^2$ and θT with $\theta\mathbb{R}^2$ such that $\int_{\mathbb{R}^2} \left(f_\epsilon(T) - f_\epsilon(T - x)\right) h(x) dx$ gives the same dispacement relative to $f_\epsilon(T)$ as $\int_{\mathbb{R}^2} \left(f_\epsilon(\theta T) - f_\epsilon(\theta T - x)\right) h(x) dx$ gives relative to $f_\epsilon(\theta T)$. In other words, the displacement behaves pattern-equivariantly.

We want to show that $\hat{f}_\epsilon$ satisfies the properties listed in Theorem 30:

1. $\hat{f}_\epsilon$ is a smooth, vertical-preserving map from $X \to Y$.

Suppose T_1 and T_2 agree out to $1/\delta + d$, where d is big enough that $f_\epsilon(T_1)$ and $f_\epsilon(T_2)$ agree out to $1/\epsilon + |h|$, where $|h|$ is the radius of support of the bump function for the convolution. We know such a d exists by the comment following definition 29.

Then for $|x| < |h|$,

$$f_\epsilon(T_1) - f_\epsilon(T_1 - x) = f_\epsilon(T_2) - f_\epsilon(T_2 - x).$$

Therefore, in the convolution

$$\int_{\mathbb{R}^2} (f_\epsilon(T_1) - f_\epsilon(T_1 - x))\, h(x)dx = \int_{\mathbb{R}^2} (f_\epsilon(T_2) - f_\epsilon(T_2 - x))\, h(x)dx.$$

So we have that $\hat{f}_\epsilon$ preserves verticals. Note that $\hat{f}_\epsilon$ no longer preserves verticals that are as long as those preserved under f_ϵ. The function f_ϵ sent verticals of length δ (in the cylinders C_i) to verticals of length ϵ. The function $\hat{f}_\epsilon$ sends verticals of length $\frac{\delta}{1+d\delta}$ to verticals of length $\frac{\epsilon}{1+|h|\epsilon}$.

2. $f(T)$ and $\hat{f}_\epsilon(T)$ are in the same translational leaf, and

$$\left| f(T) - \hat{f}_\epsilon(T) \right| < \epsilon.$$

Proof:

$$\hat{f}_\epsilon - f_\epsilon(T) = \int_{\mathbb{R}^2} (f_\epsilon(T) - f_\epsilon(T - x))\, h(x)dx,$$

which can be bounded by taking the support of $h(x)$ arbitrarily small.
$f_\epsilon(T) - f(T) < \epsilon$ by construction. Thus

$$\left| f(T) - \hat{f}_\epsilon(T) \right| < \left| \hat{f}_\epsilon - f_\epsilon(T) \right| + |f_\epsilon(T) - f(T)|$$

can be bounded arbitrarily small-ly.

3. If f commutes with rotations, then $\hat{f}_\epsilon$ will also commute with rotations.

First we show that f_ϵ commutes with rotations. Suppose f commutes with rotations, and $|G'| = m$. Taking $\epsilon << \frac{2\pi}{m}$ ensures that when we define $f_\epsilon(T') := h \cdot f(T')$, the rotational portion of h will have angle less than ϵ. (See discussion in Case 2, under definition 34.)

Suppose T' is in the vertical of $g \cdot T_o \in D_i$. Then $\theta T'$ is in the vertical of $\theta \cdot g \cdot T_o \in D_i$, so T' and $\theta T'$ are defined in the same iteration of f_ϵ. By construction, $f_\epsilon(T)$ is defined to be in the vertical of $f(g \cdot T_o)$, and $f_\epsilon(\theta T)$ is defined to be in the vertical of $f(\theta \cdot g \cdot T_o) = \theta f(g \cdot T_o)$. Since f_ϵ is well defined, there is only one available choice for $f_\epsilon(\theta T')$. Since $\theta f_\epsilon(T')$ is available, we must have that $f_\epsilon(\theta T') = \theta f_\epsilon(T')$.

Next consider $\hat{f}_\epsilon$:

$$\hat{f}_\epsilon(\theta T) = f_\epsilon(\theta T) - \int_{\mathbb{R}^2} \left(f_\epsilon(\theta T) - f_\epsilon(\theta(T-x))\right) h(x)dx$$

$= \theta f_\epsilon(T) -$ displacement appropriately rotated, by design $= \theta \hat{f}_\epsilon(T)$. $\square$

Furthermore, f_ϵ has the following property concerning a patch appearing in multiple orientations in the same leaf, which we set aside in a lemma:

Lemma 38 *There exists $r > 1/\delta$ such that if $(T+x)|_r = \theta(T+y)|_r$, then $\hat{f}_\epsilon(T+x)|_s = \theta\hat{f}_\epsilon(T+y)|_s$.*

In other words, $\hat{f}_\epsilon$ commutes with rotations on sufficiently large patches within a translational leaf.

Proof: Since $\hat{f}_\epsilon(\theta T) = \theta\hat{f}_\epsilon(T)$, and $\hat{f}_\epsilon$ is continuous, there is some radius r such that $(T+x)|_r = \theta(T+y)|_r$ implies $d(\hat{f}_\epsilon(T+x), \hat{f}_\epsilon(\theta T+y)) < \pi/m$ where $m = |G'|$. Therefore $\hat{f}_\epsilon(T+x)|_s = \theta\hat{f}_\epsilon(T+y)|_s$.

$\square$

This theorem is important, because this is what fails to hold in the case of infinite rotations. It is *not* true that commuting with rotations of tilings implies that $\hat{f}_\epsilon$ commutes with rotations on sufficiently large patches.

Theorem 39 *Suppose that $F : X \times [0,1] \to Y$ is a homotopy such that each F_t are continuous maps on X, and F_0, F_1 preserve verticals. Then $H^*_\rho(F^*_0(Y)) = H^*_\rho(F^*_1(Y))$.*

Proof: Let's say F_0, F_1 preserve verticals of length ϵ_0, ϵ_1 respectively. Let $\epsilon := \min\{\epsilon_0, \epsilon_1\}$. Approximate each F_t with $\hat{F}_t$, a vertical-preserving map. Note this homotopy, $\hat{F}_t$, is not smooth (or even continuous) in t.

As before, we will convolve with a bump function $g(t)$, with compact support of radius 1. Let

$$\mathcal{F}(T,t) = \hat{F}_t - \left(t(1-t)\int_{\mathbb{R}}[F_t(T) - F_{t-x}(T)]g(x)dx\right).$$

(For $(t-x) < 0$ or > 1, define $F_{t-x} := F_0$ or $:= F_1$, appropriately.) Note that $F_t(T)$ and $F_{t-x}(T)$ are both tilings in Y, and that they are translates of each other. This is because originally F_t was smooth in the variable t, and the approximation procedure to preserve tilings only translates any particular tiling. Thus $[F_t(T) - F_{t-x}(T)]$ is a well-defined displacement vector. The factor of $t(1-t)$ out in front of the integral ensures that $\mathcal{F}(T,0) = F_0(T)$ and $\mathcal{F}(T,1) = F_1(T)$.

Now $\mathcal{F}$ satisfies the criterion of the earlier theorem, and hence $\mathcal{F}_0 = F_0$ and $\mathcal{F}_1 = F_1$ all induce the same pullback on cohomology.
□

Definition 40 $f^*(H^*_\rho(Y)) := f^*_\epsilon(H^*_\rho(Y))$.

Theorem 41 H^*_ρ *is a topological invariant under homeomorphisms that commute with* G'.

Proof: Let $h : X \to Y$ be a homeomorphism. Then $h^{-1} \circ h$ and $h \circ h^{-1}$ are homotopic to the identity on X and Y respectively. Approximate h and h^{-1} by vertical preserving maps f and g respectively.

Claim: $f \circ g$ and $g \circ f$ approximate $h^{-1} \circ h$ and $h \circ h^{-1}$, respectively. First we show that for $T \in X$, $\left|(h^{-1} \circ h)(T) - (f \circ g)(T)\right| < \epsilon$. This is because $|h(T) - f(T)|$ can be made arbitrarily small, and $\left|h^{-1}(h(T)) - g(f(T))\right| < \left|h^{-1}(h(T)) - g(h(T))\right| + \left|h^{-1}(f(T)) - g(f(T))\right|$, and so also can be made arbitrarily small. Therefore $f \circ g$ preserves translational leaves and moves no tiling more than ϵ from its image under $h^{-1} \circ h$.

Second, the composition of vertical preserving maps is again vertical preserving. Let $\epsilon > 0$. Then $\exists \delta_1 > 0$ such that g sends tilings that match out to $1/\delta_1$ in Y to tilings that match out to $1/\epsilon$ in X. Then there also exists δ_2 such that f sends tilings that match out to $1/\delta_2$ in X to tilings that match out to $1/\delta_1$ in Y. Hence $f \circ g$ is vertical preserving. Finally, smoothness and commuting with rotations are preserved under composition.

Since $f \circ g$ approximates $h^{-1} \circ h$, the maps induced on cohomology by $h^{-1} \circ h$ is defined to be that induced by $f \circ g$. Since is $f \circ g$ and $g \circ f$ are vertical-preserving maps homotopic to the identity, H^*_ρ is invariant under both compositions. Therefore both f and g, and hence h and h^{-1}, induce isomorphisms on cohomology. $\square$

To reconnect with our earlier point, the most important aspect of this proof is the approximation theorem, which has ramifications beyond the scope of this application. It is known that there are tiling spaces which are topologically conjugate but not related by a local factor map. Therefore, a homotopy which preserves the topology cannot possibly guarantee that the dynamics will be preserved. Approximating a factor map with one that preserves verticals is precisely straddling this divide.

Secondly, this proof was presented in the global tiling space context, not the inverse limit space context. Generally the inverse limit space is easier computationally to work with - why not consider it as the setting for this proof? The reason is the precise nature of what we are trying to do. We can phrase our task in terms of inverse limit spaces: we are trying to approximate a function by it's behavior on approximants. But when we restrict to approximants in the inverse limit space, we no longer have access to the leaf structure, and this construction crucially hinges on the G direction of the tiling space.

7.2 Example: half-hex tiling

Recall that a form ω (or cochain) is pattern-equivariant with respect to a representation ρ of the rotation group G' if there exists an r such that $T_1|_r = \theta T_2|_r$ implies that $\rho(\theta)\omega(T_1) = \omega(T_2)$. By utilizing the inverse limit structure, we can streamline this definition: a form ω is pattern-equivariant with respect to ρ if there is an approximant K_n such that $x, \theta x \in K_n$ implies $\rho(\theta)\omega(x) = \omega(\theta x)$.

The rotation group of the half-hex tiling is $G' = \mathbb{Z}_6$. The corresponding equation $\theta^6 - 1 = 0$ factors into irreducible components depending on the choice of coefficients. Over $\mathbb{R}$ (or $\mathbb{Z}$) there are four irreducible representations of G', corresponding to the factorization $\theta^6 - 1 = (\theta - 1)(\theta + 1)(\theta^2 + \theta + 1)(\theta^2 - \theta + 1)$.

The first two factors correspond to 1-dimensional representations, either $\rho(\alpha) = 1$ or $\rho(\alpha) = -1$, where α is a generator of $\mathbb{Z}_6$. The second two factors correspond to 2-dimensional representations. For example, $\rho(\alpha) = \begin{pmatrix} -1 & 1 \\ -1 & 0 \end{pmatrix}$ satisfies $\theta^2 + \theta + 1$. For $x \in K_n$, we have

$$\begin{pmatrix} dx_{\theta x} \\ dy_{\theta x} \end{pmatrix} = P^{-1} \begin{pmatrix} -1 & 1 \\ -1 & 0 \end{pmatrix} P \begin{pmatrix} dx_x \\ dy_y \end{pmatrix}$$

as maps on the tangent spaces to θx and x respectively. The matrix P is the change of basis matrix between the basis $[1, 0]$ and $[0, 1]$ to the basis $[1, 0]$ and $[\cos \frac{2\pi}{3}, \sin \frac{2\pi}{3}]$. Note that $P^{-1} \begin{pmatrix} -1 & 1 \\ -1 & 0 \end{pmatrix} P$ also satisfies $\theta^2 + \theta + 1$.

Over $\mathbb{C}$, we get six one-dimensional representations of $\mathbb{Z}_6$, corresponding to the factorization of $\theta^6 - 1$ into roots of unity.

Note that we are sticking with forms on a tiling space. If we pass to the form induced on a specific tiling T, we can speak of forms on $\mathbb{R}^2$ satisfying a pattern-linked requirement, as originally developed by [K]. However, note that the 1-forms dx and dy are not rotationally pattern-equivariant except with respect to the trivial representation, (although they together satisfy a 2-dimensional representation). Thus speaking of a pattern-equivariant function with a basis dx, dy would require less-elegant trigonometry in order to describe a form that twists as patches rotate. When we speak of a form on a tiling space, we can canonically identify the tangent spaces at T and θT so that dx_T, in terms of the patch at the origin of T, is naturally identified with $dx_{\theta T}$, relative to the rotation of the patch at the origin of θT.

Recall that we computed the cohomology of the half-hex tiling space as:

$$H^i(X) = \begin{cases} \mathbb{Z} & \text{if } i = 0, \\ \mathbb{Z}[1/2]^2 & \text{if } i = 1, \\ \mathbb{Z}[1/4] \oplus \mathbb{Z}^2 & \text{if } i = 2, \end{cases}$$

and that this was computed by finding the cohomology of each approximant K_n and then taking a direct limit, given by the substitution map. For each representation, we will compute the pattern-equivariant cohomology of an approximant, and then again take the direct limit to get the pattern-equivariant cohomology of the tiling space for a given representation.

Also recall that the cohomology of each of the approximants K_n is

$$H^i(K_n) = \begin{cases} \mathbb{Z} & \text{if } i = 0, \\ \mathbb{Z}^2 & \text{if } i = 1, \\ \mathbb{Z}^3 & \text{if } i = 2, \end{cases}$$

and was computed by using the CW structure of the approximants K_n. For the time being, we continue working over $\mathbb{Z}$.

In the top-dimensional cohomology, we take the trapezoid in each of the six orientations to be our generators; hence the relative orientation group ring is $\mathbb{Z}[\theta]/(\theta^6 - 1)$. Under $\mathbb{Z}$, $\theta = -1$ and $\theta = 1$ are each associated to one copy of $\mathbb{Z}$, and $\theta^2 + \theta + 1$ and $\theta^2 - \theta + 1$ are each associated with $\mathbb{Z}^2$.

In $H^1(K_n)$ we take a short edge in three orientations, and a long edge in three orientations, so that the relative orientation group ring of 1-chains is $\mathbb{Z}[\theta]/(\theta^3 - 1) \oplus \mathbb{Z}[\theta]/(\theta^3 - 1)$. Then $\theta^2 + \theta + 1$ is associated to $\mathbb{Z}^2 \oplus \mathbb{Z}^2$, and $\theta - 1$ is associated to $\mathbb{Z} \oplus \mathbb{Z}$.

For $H^0(K_0)$ we have one generator which appears in two orientations. Thus the orientation group ring is $\mathbb{Z}[\theta]/(\theta^2 - 1)$. Only $\theta = 1$ and -1 show up, each with a copy of $\mathbb{Z}$.

The following table illustrates the exact sequences by choice of representation:

Representation:	$\theta - 1$	$\theta + 1$	$\theta^2 + \theta + 1$	$\theta^2 - \theta + 1$
2 dimensions:	$\mathbb{Z}$	$\mathbb{Z}$	$\mathbb{Z}^2$	$\mathbb{Z}^2$
	$\uparrow$	$\uparrow$	$\uparrow$	$\uparrow$
1 dimension:	$\mathbb{Z} \oplus \mathbb{Z}$	0	$\mathbb{Z}^2 \oplus \mathbb{Z}^2$	0
	$\uparrow$	$\uparrow$	$\uparrow$	$\uparrow$
0 dimension:	$\mathbb{Z}$	$\mathbb{Z}$	0	0

Without closer scrutiny to the coboundary maps, we can compute the cohomology according to the representations $\theta + 1$ and $\theta^2 - \theta + 1$ as follows:

Cochains pattern-equivariant with respect to $\rho_1 = \theta + 1$ have the exact sequence:

$$\mathbb{Z} \leftarrow 0 \leftarrow \mathbb{Z}.$$

Hence

$$H^n_{\rho_1}(K_i) = \begin{cases} \mathbb{Z} & \text{if } k = 2, \\ 0 & \text{if } k = 1, \\ \mathbb{Z} & \text{if } k = 0. \end{cases}$$

Cochains pattern-equivariant with respect to $\rho_2 = \theta^2 - \theta + 1$ have the exact sequence:

$$\mathbb{Z}^2 \leftarrow 0 \leftarrow 0$$

and so the cohomology is

$$H^n_{\rho_2}(K_i) = \begin{cases} \mathbb{Z}^2 & \text{if } k = 2, \\ 0 & \text{if } k = 1, \\ 0 & \text{if } k = 0. \end{cases}$$

Cochains pattern-equivariant with respect to $\rho_3 = \theta^2 + \theta + 1$ have

$$\mathbb{Z}^2 \leftarrow \mathbb{Z}^2 \oplus \mathbb{Z}^2 \leftarrow 0$$

and so here we must understand the boundary maps to compute the cohomology.

Note that the boundary of a trapezoid in standard position is one long edge, and four short edges, in orientations θ^2, θ^3, and θ^4. Thus for choice of θ, we have $\partial(\theta\text{trap}) = \theta(\text{long}) + (\theta^2+\theta^3+\theta^4)(\text{short})$. The boundary of any edge is two vertices, one in each orientation. So $\partial(\theta\text{edge}) = (1+\theta)(\text{vertex})$.

Returning to the case that $\rho_3 = \theta^2+\theta+1$, we have that $\partial(\theta\text{trap}) = \theta(\text{long}) - 2\theta(\text{short})$. In particular, the boundary map has no kernel, and image $\mathbb{Z}\oplus\mathbb{Z}$. We do not need to concern ourselves with the boundary map from dimension 1 to 0, since under this representation there is no 0^{th} cohomology. So we have:

$$H^n_{\rho_3}(K_i) = \begin{cases} 0 & \text{if } k=2, \\ \mathbb{Z}\oplus\mathbb{Z} & \text{if } k=1, \\ 0 & \text{if } k=0. \end{cases}$$

Lastly, cochains pattern-equivariant with respect to $\rho_4 = \theta - 1$ have exact sequence

$$\mathbb{Z} \leftarrow \mathbb{Z}\oplus\mathbb{Z} \leftarrow \mathbb{Z}.$$

Using the boundary maps above, when $\theta = 1$ we have $\partial(\text{trap}) = (\text{long}) + 3(\text{short})$, with no kernel and image $\mathbb{Z}$. So d has kernel $\mathbb{Z}$ and image $\mathbb{Z}$. From dimension 1 to 0, we have the map $\partial(\text{edge}) = 2(\text{vertex})$. The kernel is $\mathbb{Z}$ and the image is $\mathbb{Z}$, and so for d we have kernel $= 0$, image $= \mathbb{Z}$. Thus:

$$H^n_{\rho_4}(K_i) = \begin{cases} 0 \text{ if } k=2, \\ 0 \text{ if } k=1, \\ 0 \text{ if } k=0. \end{cases}$$

Notice that $H^n(K_i) = H^n_{\rho_1} \oplus H^n_{\rho_2} \oplus H^n_{\rho_3} \oplus H^n_{\rho_4}$. We have decomposed the cohomology into its rotationally invariant subspaces.

To get the pattern-equivariant cohomology of the tiling space, we must now take the direct limit of the cohomologies of each approximant, for each representation. Recall that the direct limit maps induced by the substitution are: $b_0 =$ identity on $H^0(K_i)$,

$b_1 = \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix}$ on $H^1(K_i)$, and

$b_2 = \begin{pmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{pmatrix}$ on $H^2(K_i)$, and so we get the following decomposition for the cohomology of the half-hex tiling space by representation:

Representation:	$\theta - 1$	$\theta + 1$	$\theta^2 + \theta + 1$	$\theta^2 - \theta + 1$
$H^2(X)$	0	$\mathbb{Z}[1/4]$	0	$\mathbb{Z}^2$
$H^1(X)$	0	0	$\mathbb{Z}[1/2] \oplus \mathbb{Z}[1/2]$	0
$H^0(X)$	0	$\mathbb{Z}$	0	0

This completes the example.

7.3 Classification by representation

Definition 42 *Let the canonical representation, $\tilde{\rho}$, be the representation of G' given by the factor $\theta^n - 1$.*

All other representations (besides the trivial representation) are factors of the canonical representation. The canonical representation transforms tautologically over rotations. It is the n-dimensional representation where $\rho(\theta)$ is the permutation matrix corresponding to the action of θ on G'. In other words, if θ is a generator of of G', and ω is the form, then for $p \in K_n$ we would have

$$\begin{pmatrix} \omega_{1_{\theta p}} \\ \omega_{2_{\theta p}} \\ \\ \vdots \\ \omega_{n_{\theta p}} \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ 0 & & & \ddots & 0 \\ 0 & & & & 1 \\ 1 & 0 & 0 & \dots & 0 \end{pmatrix} \begin{pmatrix} \omega_{1_p} \\ \omega_{2_p} \\ \\ \vdots \\ \omega_{n_p} \end{pmatrix}$$

Theorem 43 *If ρ_m, ρ_p, and ρ_q are three representations of G', with degrees $m = p + q$, such that the corresponding factors of equation $\theta^n - 1$ satisfy $\rho_m = \rho_p \cdot \rho_q$, then $H^*_{\rho_m} = H^*_{\rho_p} \oplus H^*_{\rho_q}$.*

Proof: Let ω_i be a collection of generators for H_{ρ_m}. Then for $t \in K_n$, $\omega(\theta t) = \rho_m(\theta)\omega(t)$, where $\rho_m(\theta)$ is a matrix that satisfies the factor of $\theta^n - 1$. (The factor is also denoted by ρ_m.) Since the term ρ_m factors into $\rho_p \cdot \rho_q$, the matrix $\rho_m(\theta)$ either satisfies one or the other, and hence ω is pattern equivariant with respect to ρ_p or ρ_q. Likewise, if ω satisfies ρ_p or ρ_q, it is also pattern-equivariant with respect to ρ_m. □

The proof does not require heavy machinery. Its simplicity illustrates why working by representation naturally takes advantage of the rotational structure of a tiling space.

Theorem 44 $H_{\tilde{\rho}} = \oplus_i H_{\rho_i}$, *where* $\{\rho_i\}$ *is the set of irreducible representations.*

Because $\tilde{\rho} = \Pi \rho_i$, this follows from the previous theorem.

Theorem 45 *If* ω *is a translationally pattern-equivariant form, and the rotation group* G' *has order* n*, then* $n\omega$ *is rotationally pattern-equivariant with respect to the canonical representation.*

A typical translationally-equivariant form ω is rotationally-equivariant with respect to the trivial representation, and evaluates independently on different rotations of a tiling, $T, \theta T, \ldots, \theta^{n-1}T$. Let θ act on forms by $\theta(\omega)(T) := \omega(\theta T)$. We will write $n\omega$ as a linear combination of $\omega, \theta\omega, \ldots, \theta^{n-1}\omega$ in a way that satisfies the canonical representation. Let $\hat{\omega} = \begin{pmatrix} \sum_i \eta_1^i \theta^i \omega \\ \vdots \\ \sum_i \eta_n^i \theta^i \omega \end{pmatrix}$ where $\eta_j = e^{\frac{2\pi i}{j}}$. By construction this is a form that fits $\hat{\omega}(\theta T) = \tilde{\rho}(\theta)\hat{\omega}(T)$.

Chapter 8

The infinite case

The case where the rotation group is S^1 is a wealth to be further explored and better understood. We outline what we do know and why it will be worth understanding completely. These are spaces composed of tilings whose patches occur in infinitely many orientations, a dense subset of S^1, and hence lack FLC with respect to translation. See figure 6.2 for the pinwheel example. These tiling spaces are composed of two-dimensional tilings, but the tangent space - and hence the cohomology - is three-dimensional. The cohomology has not yet been computed of any of these tiling spaces.

The definitions of being pattern-equivariant with respect to a representation still makes sense when $G' = S^1$. The problem is it does not seem to be a topological invariant anymore.

This may mean that the rotational structure of a tiling space is a finer category than homeomorphism. We do have invariance under the following intermediate situation, which we state and discuss without formal proof:

Statement *Let X and Y be homeomorphic tiling spaces, both with rotation group S^1, and suppose $f : X \to Y$ meets the following conditions:*
1. f commutes with rotations
2. f sends translational leaves to translational leaves
3. f satisfies Lemma 38, which states that for sufficiently large r, $(T + x)|_r =$

$\theta(T+y)|_r$ *implies* $f(T+x)|_s = \theta f(T+y)|_s$.

Under these criterion, f can be approximated by f_ϵ as in the finite case. Hence forms pattern-equivariant with respect to ρ pull back to forms pattern-equivariant with respect to ρ, and all follows as in the finite case.

Discussion:

First, if f fails to send translational leaves to translational leaves, then $f : T \to f(T)$ does not induce a map on $\mathbb{R}^2$. Thus our original identification of forms on $\mathbb{R}^2$ which are pattern-equivariant with respect to a tiling T, and forms on the tiling space X no longer holds. The pullback of a form under such a map f would be on the tangent space to a leaf of X, which is the tangent space of G, not $\mathbb{R}^2$. And on G, rotations are homotopic to the identity. Therefore this is inadequate to understand the rotational structure of the tiling space. (Other approaches to generalizing pattern-equivariant cohomology have relaxed the condition that translational leaves get sent to translational leaves. Johannes Kellendonk (in progress) is developing a version of forms on the Lie group G, instead of forms on the approximants that satisfy the representation condition. Again, because rotations disappear here, this approach does not clarify the role of rotations within a tiling space.)

An example of a map that fails to send translational leaves to translational leaves is the following: Let f be a map on the Pinwheel tiling space, X. . Let g be a smooth map from the tiles $\{t_i\}$ of the pinwheel tiling to the real numbers, such that for $x \in \partial t_i$, $g(x) = 0$ - the support of any tile is contained in its interior. Finally, define $f : X \to X$ by $f(T)$ is T, rotated by $g(x)$, where x is the location of the origin of T. This function preserves verticals - if T_1 and T_2 agree on their choice of origin, then they will be rotated by the same amount, and hence $f(T_1)$ and $f(T_2)$ will be in the same vertical. This map does *not* send translational leaves to translational leaves, and hence does not descend to a map on $\mathbb{R}^2$.

This map does however satisfy Lemma 38. For a while, the author was guided by this counter-example, and it appeared that a theorem requiring conditions 1 and 2 would be full generality: that rotationally-equivariant cohomology would be invariant under maps which commute with rotations and send translational leaves to translational leaves.

However, there is a subtle problem, which leads to the conjecture that pattern-equivariant cohomology is *not* invariant under homeomorphisms satisfying conditions one and two.

Conjecture 46 *There exists homeomorphisms $h : X \to Y$ such that h commutes with rotations but fail to satisfy condition 3.*

We sketch what a counter-example might look like. Let $h : X \to Y$. Suppose there is a series of patches $\{P_i\}$, whose radius $\to \infty$. For every i, suppose two copies of the same patch occur in two orientations, in the same leaf, ie $(T+x)|_r = \theta(T+y)|_r$. Then since f is continuous, $d(f(T+x), f(T+y)) <$ small, and so $f(T+x)|_s = \theta' f(T+x)|_s + d$ where d is a small translation, and $0 < \theta' - \theta <$ small.

Figure 8.1 illustrates what this might look like. For simplicity, this depiction is given on a single tiling, with patches sent to patches, instead of tilings in a space. The key is that angles are distorted in small amounts.

As $i \to \infty$, the patches $\{P_i\}$ must come in pairs whose angle gets closer and closer to being preserved, because in the limit, two tilings which differ by θ are sent to tilings which differ by θ. We conjecture that a map like this exists and that under our algorithm, cannot be approximated by a vertical-preserving map. If this is true, it would mean that rotationally-equivariant cohomology distinguishes between these two homeomorphic spaces, based on the geometry of the rotation group within each of them.

Outside of pattern-equivariant cohomologies, infinite rotation tiling spaces have been studied as S^1 fiber spaces over the quotient-space X/S^1. There are singular fibers corresponding to tilings with rotational symmetry. From work done by (cite), these singular fibers are the key elements to understanding the space. One can see that the representation ρ of S^1 chosen to form H^*_ρ will act on the fiber, and therefore determine which forms are pullbacks of forms on the base space. But how does this fiber structure deeply relate to the Cantor structure? What does

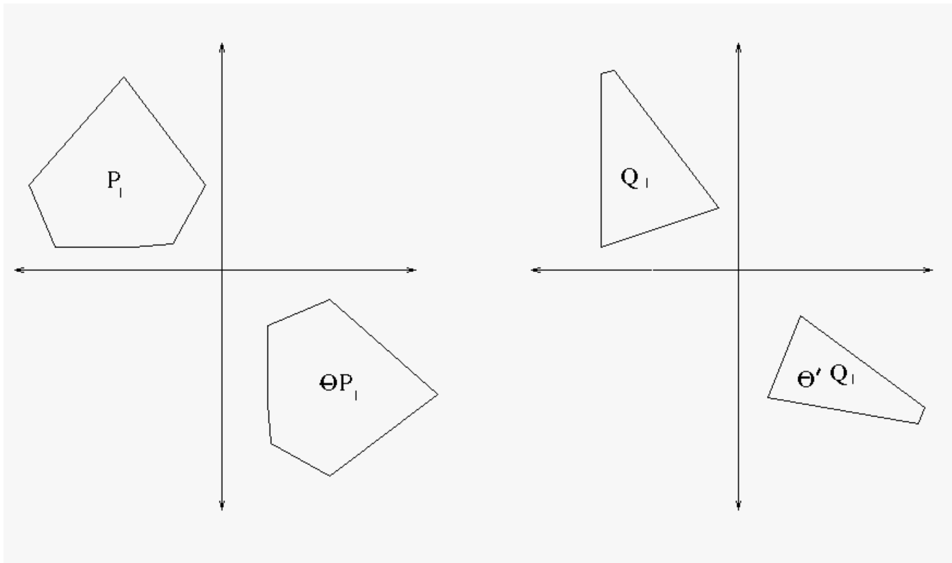

Figure 8.1: As a map induced on a tiling T, P and θP get sent to Q and $\theta' Q$

rotationally-equivariant cohomology really extract about these spaces? These questions are open and enticing.

Chapter 9

Conclusion

We have laid out the framework for a flexible structure to incorporate rotations into pattern-equivariant cohomology, by linking it to a choice of representation of the rotation group. The significance of this in decrypting the structure of tiling spaces will aid our understanding of them. Along the way, we develop a potentially far-reaching approximation theorem, which both clarifies and bridges tiling spaces which are MLD with those that are topologically conjugate spaces. There is still much to be understood. For example, intuitively a pattern-equivariant form is preserved under scaling, but our approach avoids dealing with this, since scaling is a topological conjugacy, and we reduce our cases to MLD pullbacks. What is the broadest category that pattern-equivariant forms pull back under? What does the pattern-equivariant structure inform you about the space? What will happen in higher dimensions? Can this be meaningfully generalized to any inverse limit space? Can this structure of a pattern-equivariant form transforming with respect to a group representation be generalized to Lie groups? To general dynamical systems? These questions will motivate much in the future.

Bibliography

[AP] J. Anderson and I. Putnam, Topological invariants for substitution tilings and their C*-algebras , *Ergodic Th. and Dynam. Sys.* **18** (1998), 509-537.

[HRS] C. Holton, C. Radin, and L.Sadun, Conjugacies for tiling dynamical systems, *Comm. Math. Phys.*, **254** (2005) 343-359.

[K] J. Kellendonk, Pattern-equivariant functions and cohomology, *J. Phys. A* **36** (2003) 5765-5772.

[KP] J. Kellendonk and I. Putnam, The Ruelle-Sullivan map for $\mathbb{R}^n$ actions, preprint 2005. http://wwww.math.uvic.ca/faculty/putnam/r/0206_final.pdf, to appear in *Math. Ann.*

[LM] D. Lind and B. Marcus, *An Introduction to Symbolic Dynamics and Coding*, Cambridge University Press, Cambridge, 1995.

[ORS] N. Ormes, C. Radin, L. Sadun, A Homeomorphism Invariant for Substitution Tiling Spaces, *Geometriae Dedicata* **90** (2002), 153-182.

[Pe] K. Petersen, Factor maps between tiling dynamical systems *Forum Math.* **11** (1999), 503-512.

[P] N. Priebe, Towards a Characterization of self-similar tilings in terms of derived Voronoi tessellations, *Geom. Dedicata* **79** (2000),239-265.

[PS] N. Priebe and B. Solomyak, Characterization of planar pseudo-self-similar tilings. *Discrete Comput. Geom.* **26** (2001), 289-306.

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

[RS1] C. Radin and L. Sadun, Isomorphism of hierarchical structures, *Ergodic Theory Dynam. Systems*, **21** (2001) 1239-1248.

[S] Tiling Spaces are Inverse Limits, *Journal of Mathematical Physics* **44** (2003) 5410-5414.

[S2] Pattern-Equivariant Cohomology with Integer Coefficients. (In submission)

[SW] L. Sadun and R. F. Williams, Tiling Spaces are Cantor Set Fiber Bundles , *Ergodic Theory and Dynamical Systems* **23** (2003) 307-316.

Vita

Betseygail Rand was born in Gainesville, Florida, on February 3^{rd}, 1978, to proud parents Kenneth and Colleen Rand. She graduated from Eastside High School in 1995 with both a high school diploma and an International Bacchalaureate degree. She graduated from the University of Michigan in 1999 with a BS in Mathematics. In the fall of 2000 she entered the Graduate School at the University of Texas.

Permanent Address: 216 Riverside Dr
San Marcos, Texas, 78666

This dissertation was typeset with LaTeX 2_ε[1] by the author.

[1] LaTeX 2_ε is an extension of LaTeX. LaTeX is a collection of macros for TeX. TeX is a trademark of the American Mathematical Society. The macros used in formatting this dissertation were written by Dinesh Das, Department of Computer Sciences, The University of Texas at Austin, and extended by Bert Kay, James A. Bednar, and Ayman El-Khashab.