Fusion: A general framework for hierarchical tilings

Natalie Priebe Frank[1] Lorenzo Sadun[2]

[1]Vassar College

[2]University of Texas at Austin

Aperiodic Order, CMS Edmonton, June 3 2011

Fusion rules

A framework for general hierarchical systems, including

- Self-similar tilings and substitution sequences—our original motivation
- Cut-and-stack transformations
- Bratteli-Vershik systems
- Combinatorial substitutions
- Generalized substitutions
- S-adic transformations
- Random substitutions

Motivating results for substitution tiling dynamics

Acting by translation on the tiling space, we have

- Conditions for unique ergodicity
- There is an algorithm for finding eigenvalues
- Measurable eigenfunctions can be chosen continuous.
- The substitution is invertible if and only if the tilings are non-periodic
- No positive entropy or strong mixing
- Hierarchy can be enforced via local matching rules
- The spaces are either Cantor sets or Cantor set fiber bundles
- The spaces are inverse limits
- The cohomology and K-theory is computable

Today's talk will focus on invariant/ergodic measures.

Self-similar or substitution tilings

Self-similar tilings are "ideal" hierarchical tilings. Several ingredients:

- A group G of isometries, typically either $\mathbb{Z}^d$, $\mathbb{R}^d$ or the Euclidean group.
- A finite collection of shapes, called "prototiles". "Tiles" are prototiles moved around by group elements.
- An expansive linear tranformation $L : \mathbb{R}^d \to \mathbb{R}^d$, typically a pure dilation.
- A rule σ for replacing each tile t with a patch of tiles whose union is $L(t)$.

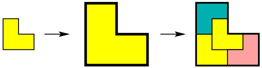

- Choose $G = \mathbb{R}^2$, moving tiles and tilings by translation,
- Thus we need four tile types, and
- The linear map L expands by a factor of 2.
- On all four tile types:

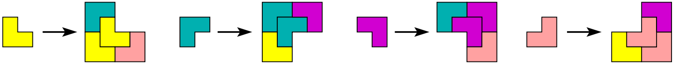

Iterating repeatedly creates "supertiles" that grow to cover the plane.

A few chair supertiles

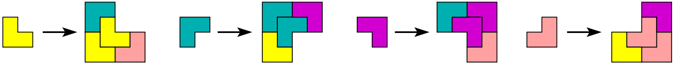

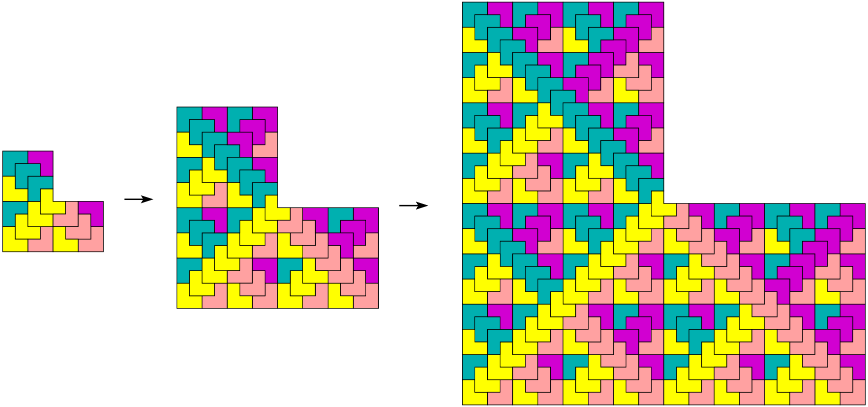

A tiling substitution without linear expansion L

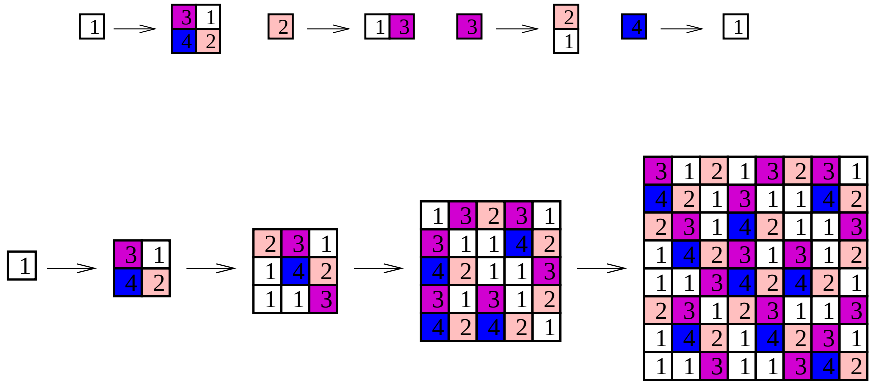

Definition An n-supertile is a tile t that has been substituted n times, i.e. $\sigma^n(t)$.

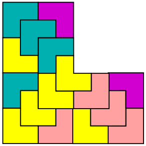

A chair 2-supertile (with t=yellow tile)

Two views of supertiles

1. $\sigma^n(t) = \sigma(\sigma^{n-1}(t))$ is what you get by expanding an $n-1$ supertile, replacing each tile with a 1-supertile:

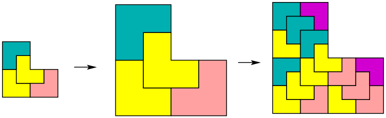

2. $\sigma^n(t) = \sigma^{n-1}(\sigma(t))$ is what you get by gluing several $n-1$ supertiles together, in a pattern determined by $\sigma(t)$.

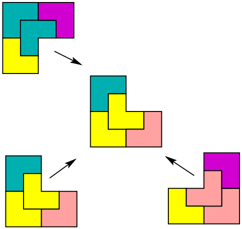

But why does gluing pattern have to be the same at each level?

Definition of a fusion rule $\mathcal{R}$

Definition. A *fusion* of a patch P_1 to another patch P_2 is a union of P_1 and P_2 that is connected and does not contain overlaps.

$$fusion = geometric\ concatenation$$

- *0-supertiles.* A finite collection $\mathcal{P}_0$ of tiles. These are "atoms".
- *1-supertiles.* A finite collection $\mathcal{P}_1$ of patches (fusions) of tiles from $\mathcal{P}_0$. These are "molecules".
- *2-supertiles.* A finite collection $\mathcal{P}_2$ of patches made by fusing together elements from $\mathcal{P}_1$.
- *n-supertiles.* For each $n > 0$, $\mathcal{P}_n$ is a finite set of patches that are fusions of $(n-1)$-supertiles.

How is this different from substitution?

For one thing, the combinatorics can change from level to level.

Example. Let $\mathcal{P}_0 = \{a, b\}$

$$P_n(a) = P_{n-1}(a)P_{n-1}(b),$$

$$P_n(b) = \begin{cases} P_{n-1}(a)P_{n-1}(b)P_{n-1}(b) & \text{if } n \text{ is prime} \\ P_{n-1}(b)P_{n-1}(b)P_{n-1}(a) & \text{if } n \text{ is not prime} \end{cases}$$

Our sets of supertiles are:

- $\mathcal{P}_1 = \{ab, bba\}$
- $\mathcal{P}_2 = \{abbba, abbbabba\}$
- $\mathcal{P}_3 = \{abbbaabbbabba, abbbaabbbabbaabbbabba\}$

Difference, part 2: The 10^n example.

$$\mathcal{P}_0 = \{a, b\}, \qquad \mathcal{P}_n = \{P_n(a), P_n(b)\},$$

where

$$P_1(a) = a^{10}b = aaaaaaaaaab \qquad P_1(b) = b^{10}a = bbbbbbbbbba$$

Let

$$P_2(a) = (P_1(a))^{100}P_1(b) \qquad P_2(b) = (P_1(b))^{100}P_1(a)$$

and in general

$$P_n(a) = (P_{n-1}(a))^{10^n}P_{n-1}(b) \qquad P_n(b) = (P_{n-1}(b))^{10^n}P_{n-1}(a)$$

A two-dimensional example

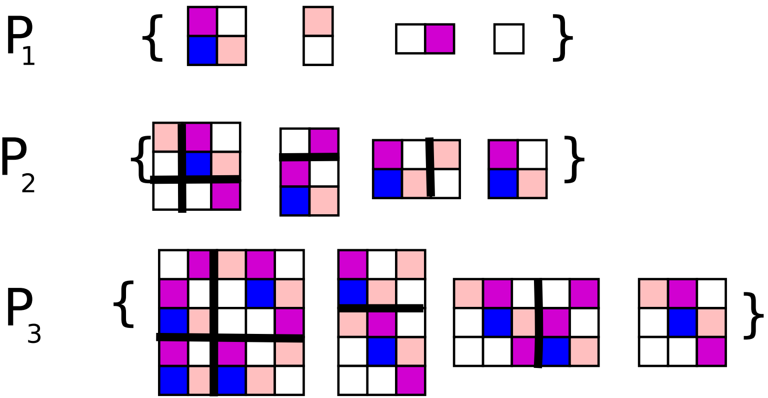

We could change the combinatorics at each level; we could change the number of *n*-supertiles at each stage too.

Tiling spaces

Consider a fixed substitution σ or fusion rule $\mathcal{R}$.

- A patch is *admissible* if it is found somewhere in a supertile.
- An infinite tiling $\mathbf{T}$ is *admissible* if every finite patch of tiles in $\mathbf{T}$ is admissible.
- The tiling space Ω_σ or $\Omega_{\mathcal{R}}$ consist of all admissible tilings $\mathbf{T}$.
- That is, each point in the tiling space is an infinite tiling.
- A tiling and its image under a rigid motion from G are, in the absence of symmetry, distinct points in the tiling space.

Note: While σ extends to a map from Ω_σ to itself, there is no self-map of $\Omega_{\mathcal{R}}$ induced by the fusion rule. *This is a major obstruction to proofs.*

Given a tiling space Ω constructed using isometry group G, we give Ω the “big ball topology”.

G acts continuously on Ω; our dynamical system is

$$(\Omega, G, \mu)$$

where μ is an invariant Borel probability measure.

For simplicity we often assume that our action is by continuous translations: $G = \mathbb{R}^d$.

What do we know about fusion systems?

- Without further assumptions, not a lot: every tiling system can be expressed as a fusion tiling system.
- With various assumptions, tons! We have versions of almost all of the motivating results listed at the beginning.

- Fusion allows for interesting constructions impossible for substitutions.
  - There can be measurable eigenfunctions/values that aren't continuous.
  - Strong mixing and entropy are possible.
  - Minimal systems can fail to be uniquely ergodic.

Essential technical stuff

- "Transition matrices" $M_{n,N}$ count how many n-supertiles make up each N-supertile:

$$M_{n,N}(i,j) = \#P_n(i) \text{ in } P_N(j)$$

  - For $n < m < N$ we have $M_{n,m}M_{m,N} = M_{n,N}$
- "Primitivity": for each n there exists an N such that all entries of $M_{n,N}$ are positive.
  - Implies $(\Omega_{\mathcal{R}}, G, \mu)$ is a minimal (the closure of one orbit).
- "Van Hove" (aka Følner): Supertile volumes grow large relative to their interiors.
  - Essential for frequency (and thus measure) calculations.
- "Recognizability": every tiling decomposes into n-supertiles unambiguously
  - Essential for anything involving Bratteli diagrams.

The Bratteli diagram model of $\Omega_{\mathcal{R}}$

Let the set of n-supertiles be denoted by
$\mathcal{P}_n = \{P_n(1), P_n(2), ..., P_n(j_n)\}$

- The vertex set V_n is the set of n-supertiles.
- Each copy of $P_n(i)$ in $P_{n+1}(j)$ gets an edge in E_n.
- If there is more than one way to fuse n-supertiles to get an $(n+1)$-supertile, fix a preferred one to use in this and all other computations.
- The transition matrix F_n is defined as the number of n-supertiles in each N-supertile. That is, $F_n = M_{n,n+1}$.

The Bratteli diagram model of $\Omega_{\mathcal{R}}$

μ-almost every tiling $\mathbf{T}$ in $\Omega_{\mathcal{R}}$ corresponds to an infinite path in the Bratteli diagram $\mathfrak{B}_{\mathcal{R}}$: the vertex in V_n corresponds to the n-supertile in $\mathbf{T}$ that contains the origin.

If $G = \mathbb{Z}$, the dynamics gives rise to an ordering and adic transformation on $\mathfrak{B}_{\mathcal{R}}$

but I don't know how to do this in general.

I do know that changing the ordering changes the combinatorics, and possibly the nature of invariant measures.

Moreover, a Bratteli diagram with fixed ordering can support invariant measures of different types.

Invariant measures and abstract patch frequency

$P =$ finite patch and $U =$ small subset of $\mathbb{R}^d$. The *cylinder set*

$$X_{P,U} = \{\mathbf{T} \in \Omega_{\mathcal{R}} | P - x \subset \mathbf{T} \text{ for some } x \in U\}$$

When μ is a translation-invariant Borel probability measure we have

$$\mu(X_{P,U}) = freq_{\mu}(P) Vol(U)$$

But we have to make sense of what this frequency means.

Invariant measures and abstract patch frequency

Fix a Van Hove sequence $\{A_n\}$ of subsets of $\mathbb{R}^d$: $Vol(A_n) \to \infty$ and the padded boundary of A_n goes to 0 relative to $Vol(A_n)$.

By the ergodic theorem for μ-almost every $\mathbf{T} \in \Omega_{\mathcal{R}}$,

$$freq_\mu(P) = \lim_{n\to\infty} \frac{\#(P \text{ in } \mathbf{T} \cap A_n)}{Vol(A_n)}$$

Note: Changing the Van Hove sequence will change the set of $\mathbf{T}$s for which the frequencies converge.

Sequences of supertile frequencies

Consider the sequence $\rho = \{\rho_n\}$ where each $\rho_n \in \mathbb{R}^{j_n}$ has all nonnegative entries.
Then $\rho_n(i)$ represents an abstract frequency of the supertile $P_n(i)$ if ρ is:

- *volume-normalized:* $\sum_{i=1}^{j_n} \rho_n(i) Vol(P_n(i)) = 1$ for all n. (This gives us a probability measure.)

- *transition-consistent:* $\rho_n = M_{n,N}\rho_N$ whenever $n < N$. (The frequency of an n-supertile consistent with the frequencies of the N-supertiles it lives inside)

The "10^n" example

$$P_n(a) = (P_{n-1}(a))^{10^n} P_{n-1}(b) \qquad P_n(b) = (P_{n-1}(b))^{10^n} P_{n-1}(a)$$

We have $M_{n,n+1} = \begin{pmatrix} 10^n & 1 \\ 1 & 10^n \end{pmatrix}$.
The system is minimal, but $P_n(a)$ is a-heavy and $P_n(b)$ is b-heavy.

Volume-normalization for n-supertiles: $l_n = |P_n(a)| = |P_n(b)|$

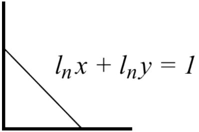

The frequencies of the n-supertiles satisfy $\rho_n(a) l_n + \rho_n(b) l_n = 1$

Invariant measures $\leftrightarrow$ well-defined supertile frequencies

Theorem. *Let $\mathcal{R}$ be a recognizable, primitive, Van Hove fusion rule. There is a one-to-one correspondence between the set of all invariant Borel probability measures on $(X_{\mathcal{R}}, \mathbb{R}^d)$ and the set of all sequences of well-defined supertile frequencies with the correspondence that, for all patches P,*

$$freq_{\mu}(P) = \lim_{n\to\infty} \sum_{i=1}^{j_n} \#\left(P \text{ in } P_n(i)\right) \rho_n(i) \qquad (1)$$

This means that finding invariant measures boils down to finding supertile frequencies. These are determined by the transition matrices.

Parameterization of invariant measures

- The jth column of $M_{n,N}$ represents the number of times each n-supertile appears in the N-supertile of type j.
- Normalizing this column by volume gives us the relative volume each n-supertile occupies in $P_N(j)$.
- Let $\Delta_{n,N}$ = volume-normalized vectors in positive cone spanned by columns of $M_{n,N}$.
- In the 10^n example we have 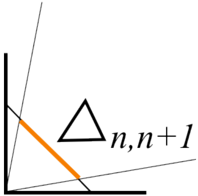

Parameterization of invariant measures

- Note: $\Delta_{n,N+1} \subset \Delta_{n,N}$, since $M_{n,N+1} = M_{n,N} M_{N,N+1}$.
- In the 10^n example we get 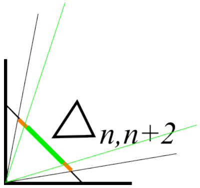

- Let $\Delta_n = \cap_N \Delta_{n,N} = \{$all possible vectors $\rho_n\}$.

Parameterization of invariant measures

Thus Δ_n is the set of possible n-supertile frequencies. A sequence of well-defined supertile frequencies $\rho = \{\rho_n\}$ must also be transition-consistent.

Take the inverse limit under $M_{n,n+1} : \Delta_{n+1} \to \Delta_n$
$\Delta_\infty = \varprojlim \Delta_n = \{$ invariant measures $\}$

Corollary. *Let $(X_{\mathcal{R}}, \mathbb{R}^d)$ be the dynamical system of a recognizable, primitive, Van Hove fusion rule. The set of all invariant Borel probability measures is parameterized by Δ_∞.*

Unique ergodicity

We have unique ergodicity if the Δ_ns collapse to points.

Corollary. *Let* $\delta_n = \min_{i,j,k} \frac{M_{n,n+1}(i,k)}{M_{n,n+1}(j,k)}$. *If* $\sum \delta_n = \infty$, *the fusion is primitive and* $\Omega_{\mathcal{R}}$ *is uniquely ergodic.*

Corollary. *If the transition matrix is fixed and primitive,* $\Omega_{\mathcal{R}}$ *is uniquely ergodic.*

The 10^n system has exactly two ergodic measures.

$M_{n-1,n} = \begin{pmatrix} 10^n & 1 \\ 1 & 10^n \end{pmatrix}$ has eigenvalues $10^n - 1$ and $10^n + 1$.

Δ_n is an interval for every value of n (defined by the limits of the first and second columns of $M_{n,N}$).

Likewise, Δ_∞ is an interval, whose endpoints μ_a and μ_b can be obtained from the supertile sequences $\kappa = (a, a, a, a, \ldots)$ and $\kappa = (b, b, b, b, \ldots)$.

The invariant measure $\mu = (\mu_a + \mu_b)/2$ corresponds to Lebesgue measure when the system is seen as a cut-and-stack transformation.

Other examples

- A minimal fusion rule with two ergodic measures, one having pure discrete spectrum and the other having some continuous spectrum.

- A fusion rule that is measurably pure discrete spectrum but topologically weakly mixing.
  - The Bratteli diagram is identical to one that is topologically pure point spectrum; the difference is determined by the order on the edges.

Conclusion

- Fusion gives a unified framework for lots of hierarchical structures.
- Some properties of substitutions carry over, others don't. Interesting counterexamples.
- For primitive Van Hove fusions, finding ergodic measures boils down to linear algebra. The *spectral type* of those measures doesn't, though.
- Spectral theory can be handled via return vectors in $\mathcal{V}_n$.
- Topological properties, like inverse limit structures, resemble those of substitution tilings.