Sketch of the techniques in
"Matching rules and substitution tilings"

Chaim Goodman-Strauss

The following orginally appeared as part of the introduction to "Matching rules and substitution tilings", soon to appear in the Annals of Mathematics. Because of space limitations, this section was cut from the final draft. Nontheless, it may be helpful in understanding the paper.

In "Matching rules and substitution tilings", we set out to prove:

Theorem *Every substitution tiling of* $\mathbb{E}^d$, $d > 1$, *can be enforced with finite matching rules, subject to a mild condition:*

the tiles are required to admit a set of "hereditary edges" such that the substitution tiling is "sibling-edge-to-edge".

As an immediate corollary, infinite families of aperiodic sets of tiles are constructed.

The fundamental idea behind the construction is rather simple: in essence, we wish the tiles to organize themselves into larger and larger images of the inflated tiles (these images are "supertiles"). Each of these supertiles is to be associated with a small amount of information— its alleged position with respect to its parent supertile. This information resides on a "skeleton" of edges; these skeletons are designed so that every edge in the tiling belongs to only finitely many skeletons, and the skeletons of one generation of supertile connect to the skeletons of the previous generation. The proof that the construction succeeds is by induction: if every structure is well-formed for one generation, we show every structure must be well-formed for the previous generation as well.

Before continuing, we give a more detailed summary. We will define our terms precisely in Section 1, and in "Addressing and Substitution Tilings" (In preparation).

"Tilings" are to be coverings of d-dimensional Euclidean space $\mathbb{E}^d$ by congruences of a finite set of "prototiles"— marked compact subsets of $\mathbb{E}^d$. The images of the prototiles under congruences are to be called "tiles"; we require the tiles in a tiling to have disjoint interiors.

A "matching rule tiling" ($\mathcal{M}$, $\mathcal{T}'$) is the set of all tilings by prototiles $\mathcal{T}'$, that satisfy some local rules $\mathcal{M}$ that specify allowed bounded configurations.

Given a set of prototiles $\mathcal{T}$, a "substitution acting on the prototiles" is an expanding linear map σ (called an "inflation" or "similarity") on E^{d}, such that for each prototile A, $\sigma(\mathsf{A})$ is the union of a set A^+ of "daughter" tiles with disjoint interiors. Thus σ may also be thought of as an "inflate and subdivide" operation on configurations of tiles. We take care to require that σ can be iterated; any configuration congruent to some $\sigma^k(\mathsf{A})$, for some prototile A, we call a "supertile". The specific substitutions are encoded in a set $\mathcal{S}$ of images of the prototiles in the inflated prototiles.

A "substitution tiling" ($\mathcal{T}$, σ,$\mathcal{S}$) is the set of tilings by "polyhedral" prototiles $\mathcal{T}$ such that any bounded subset of the tiling appears in some supertile given by the substitution defined through σ and $\mathcal{S}$.

To define the "enforcement of substitution tiling by matching rules" we must define a "labeling" of a substitution tiling; essentially, this formally allows one to mark information concerning the hierarchy on the supertiles. Then a matching rule tiling "enforces" a substitution tiling if and only if it reproduces this labeling.

We begin with a given substitution tiling ($\mathcal{T}$,σ,$\mathcal{S}$). We intend for the tiles to organize themselves into larger and larger supertiles —inflations of the original tiles— further and further up the hierarchy. Each n-level supertile congruent to, say, $\sigma^{(n-1)}(\mathsf{A})$, $\mathsf{A} \in \mathcal{T}$ is to lie in a $(n+1)$-level supertile congruent to $\sigma^n(\mathsf{A}^-)$, where $A^- = \{\mathsf{B} \in \mathcal{T} \mid \mathsf{A} \in \mathsf{B}^+\}$. The essential information associated with each supertile is its own shape and position in the next level of the hierarchy.

Much of the construction given here is foreshadowed in "Tilings, substitution systems and dynamical systems generated by them", in which S. Mozes gives matching rules enforcing substitution tilings in which the tiles are all rectangular blocks. Mozes uses two key observations: each supertile needs only to know its ancestry only a finite number of generations back, and each supertile should be combinatorially active at only a few key sites.

This information needs to be consistent across the supertile, needs to be manifest at a few key points on the boundary of the supertile, and any neighborhood in the tiling must contain only a finite amount of information.

We can use the combinatorial structure of our addressing scheme as the basis for a "labeling" of $(\mathcal{T}, \sigma, \mathcal{S})$: this labeling will encode "skeletons" and "wires" to compare and transport information across supertiles. We will define finite classes of "labels"— combinatorial encodings of these mechanisms; the labeling will consist of marking the tiling with these labels.

We will define the elements of our labeling in in Section 2. In Section 3 we derive tiles and matching rules from the local structure of the labeling, and in Section 4 show these force the hierarchy to organize.

Because we define our structures on inflated prototiles they will be available, scaled up, on every supertile.

Note we are *selecting* these structures. We do not assume the "nice" choices are being made. There is thus still ample room to find elegant constructions in more specific cases.

Each n-level supertile will consist of $(n-1)$-level supertiles held together by an n-level "skeleton", defined in Section 2.1, of edges for the parent supertile. The essential information defining this supertile is conveyed in a "packet" of labels along this skeleton. That is, our matching rules will assure that we have identified each supertile's intended position with respect to its parent, and perhaps with respect to a few recent ancestral supertiles.

A skeleton will be loose and floppy, a locally defined topological object, combinatorial in nature, on top of which is encoded information concerning the role of the supertile in the hierarchy. Supertiles are geometrically rigid, and combinatorially inert. Skeletons provide combinatorial cohesion; supertiles provide geometrical rigidity. Together they force the hierarchy to emerge.

Matching rules at its vertices ensure the skeleton is formed correctly; matching rules at certain "sites" ensure that an n-level skeleton correctly meets its descendant $(n-1)$-level skeletons and its parent $(n+1)$-level skeleton.

In the lower left of figure 1, the substitution for the pinwheel tiling ("The pinwheel tilings of the plane", C. Radin) is shown; above and to the right skeletons for three generations of supertile are shown. Note the sites, shown as half circles, connecting the skeletons of child to parent.

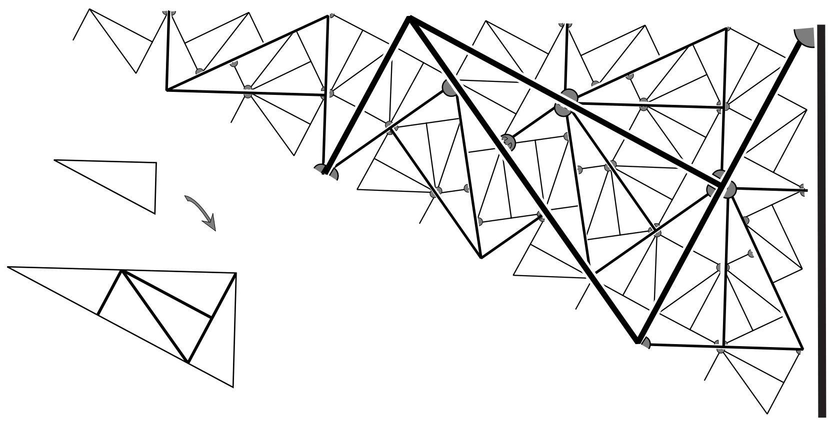

Figure 1: Skeletons

As a technical point, to ensure that skeletons are connected, that a supertile's skeleton meets each of the supertile's children, and that sites can be chosen, we allow an n-level skeleton to enter lower level supertiles (cf. figure 7). However, for any substitution tiling we find a constant κ so that all n-level skeletons include only edges of level at least $(n-\kappa)$ and less than n.(In fact, though, in most well known examples, we can take $\kappa = 1$. In this case, especially when $\mathsf{d} = 2$, the construction simplifies enormously.)

Thus skeletons might overlap, but only to a bounded depth.

Each n-level supertile, and skeleton, is associated with its "address" $\mathsf{X}_\kappa \mathsf{X}_{(\kappa-1)} ... \mathsf{X}_1$, relative to its $(n+\kappa-1)$-level ancestor. These digits are in $\mathcal{S}$. Each edge, vertex, and site in a skeleton carries the supertile's address and its own label relative to the skeleton, in classes defined throughout Section 2. As an edge or vertex might belong to many skeletons, we may encode many such pairs; however the total information at any point in a tiling will be bounded.

We need a supertile to "know" where certain vertices–"terminals"– are; these vertices are endpoints of the supertile's parent's edges. That is, if the supertile is of level n, the terminals are endpoints of n-level edges in the boundary of the supertile. We can hook the terminals into the supertile's skeleton, if they meet lower level edges in the interior of the supertile (such terminals are "endovertices"). Alas, this is not often the case and we must introduce another device– we link certain terminals ("mesovertices") to the skeleton through a series of lower level supertiles. Such a series is a "vertex wire". A supertile may thus carry, for certain of its vertices, certain information associated with some higher level supertile.

In figure 2 vertex wires are shown for three vertices on the pinwheel prototile (the vertex in the middle of the large edge is not really necessary in the actual enforcement of the pinwheel tiling but gives a more interesting vertex wire to illustrate). To the left the three vertices are illustrated; in the middle of the figure the wires are drawn as they would appear in the substitution tiling– a nested sequence of supertiles converging to the vertices; on the right the wires are drawn as they are abstractly represented– sequences of tiles. The vertex on the left of the prototile does not need a vertex wire; it is incident to the tile's skeleton.

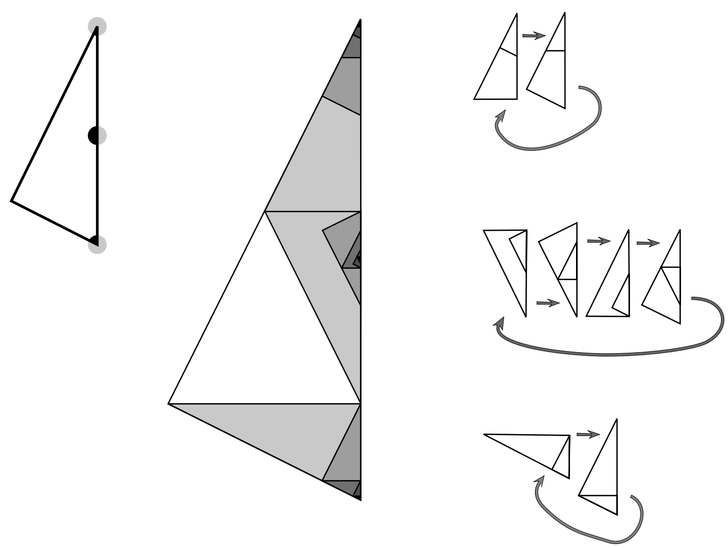

Figure 2: Vertex wires

A schematic of all the various structures we will exploit is shown in figure 3. The bulk of the construction will be an algorithm to encode these structures in our sets of labels, our packets and tiles.

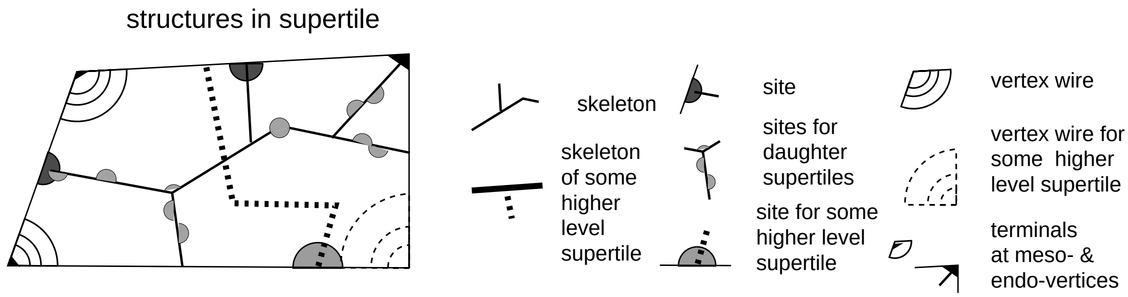

Figure 3: Structures in a supertile

Once these mechanisms are set up, the actual proof that they succeed in enforcing the hierarchical structure is relatively simple. This is necessarily so, since we cannot rely on combinatorial arguments specific to a given set of tiles. Here is a very quick sketch of the proof of the matching rules:

We derive a labeling of ($\mathcal{T}$, σ, $\mathcal{S}$), and from this construct a matching rule tiling ($\mathcal{M}$, $\mathcal{T}'$). A "well-formed supertile" is a configuration of tiles in ($\mathcal{M}$, $\mathcal{T}'$) that is essentially a labeled supertile in ($\mathcal{T}$, σ, $\mathcal{S}$).

Once everything is prepared we induct: if every tile in our matching rule tiling ($\mathcal{M}$, $\mathcal{T}'$) lies in a well-formed supertile of level n, we show every tile must lie in a well-formed supertile of level $(n+1)$. A well-formed supertile has only a few relevant properties: it is clearly marked with its address $\mathsf{X}_\kappa \mathsf{X}_{(\kappa-1)} \ldots \mathsf{X}_1$ at the sites on its boundary and at its terminals, it is actually congruent to $\sigma^n(\mathsf{X}_1)$, and it is already anticipating any information it must carry for higher level supertiles (that is, it has clearly marked channels in which higher level skeletons or vertex wires will run).

Then, as sketched in figure 6:

(i) The skeleton of the parent emerges at the sites, propagating in a uncontrolled, but locally well-formed manner along the boundary of the supertile, until

(ii) terminating at the terminals. As edges meet, information concerning the parent supertile begins to be corroborated. Because the tiling is sibling edge-to-edge, the terminals must be incident to edges that seem to be edges for neighboring sibling supertiles.

(iii) As these neighboring edges propagate sites for sibling supertiles must have been present, although initially it may not be clear that the sibling supertiles are the right size or in the right position.

(iv) Again the vertex wires come to the rescue; that the tiling is sibling edge-to-edge ensures that the wires of siblings must meet, and so fixing the siblings adjacent to our original supertile. We can then fix siblings adjacent to these, and so forth, until the geometry of the entire parent supertile has been fixed.

Finally, because the parent supertile's skeleton is connected, all information concerning the parent's alleged role in the hierarchy is consistently represented across the entire parent supertile. Thus, every tile in every tiling in $(\mathcal{M},\mathcal{T}')$ lies in a well-formed supertile of level $(n+1)$.

The full proof is not much longer; however precisely defining the structures we need requires extreme care.

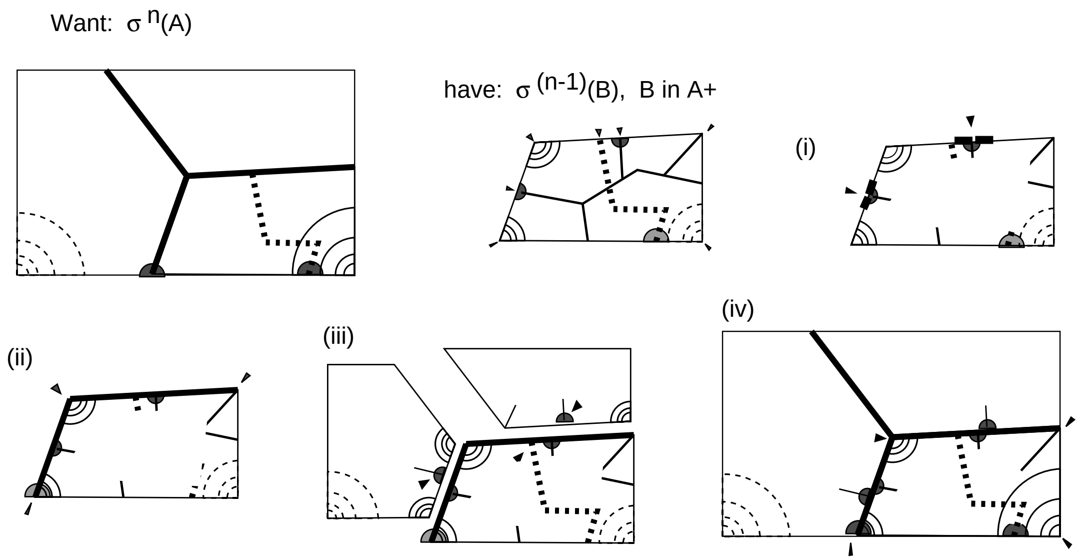

Figure 4: A sketch of the proof of that matching rules succeed

Chaim Goodman-Strauss
Dept. of Mathematics
University of Arkansas
Fayetteville, AR 72701

cgstraus@comp.uark.edu