

APERIODIC HIERARCHICAL TILINGS

CHAIM GOODMAN-STRAUSS

University of Arkansas

Fayetteville, Arkansas USA 72701

Abstract. A substitution tiling is a certain globally defined hierarchical structure in a geometric space. In [6] we show that for any substitution tiling in E^n , $n > 1$, subject to relatively mild conditions, one can construct local rules that force the desired global structure to emerge. As an immediate corollary, infinite collections of forced aperiodic tilings are constructed. Here we give an expository account of the construction. In particular, we discuss the use of hierarchical, algorithmic, geometrically sensitive coordinates—“addresses”, developed further in [9].




Figure 1. A substitution tiling

1. Introduction

Global structure arises constantly from local properties, in nature and in abstract: trees, people and crystals are good examples of the former; groups, automata and curvature provide good examples of the latter.




Figure 2. A matching rule tiling

Substitution rules generate global hierarchical structure in the plane, such as spaces of aperiodic hierarchical tilings. For example, in figure 1, we first see on the left a small L-tile and rules for “inflating and subdividing” the tile. As this process is iterated, larger and larger regions of the plane are tiled with L-tiles. On the right of figure 1, these L-tiles are arranged hierarchically into larger and larger L-shaped “supertiles”— the images of inflated and subdivided L-tiles. The thicker lines have been added to emphasize this hierarchy.

We can thus define a global structure—the “substitution tiling” induced by the inflation and division of the tiles. If a set of prototiles admits no tiling that is invariant under a translation (or more generally, no tiling invariant under any infinite cyclic group of isometries), the prototiles are *aperiodic*.

For example, the infinite L-tiling suggested at right on figure 1 *cannot* be translated since any translation will place some giant L-shaped supertile onto itself, and each tile is in only one L-shaped supertile of each size. But the unmarked L-tiles by themselves are not aperiodic since they can tile rectangles and hence can form periodic tilings.

By themselves L-tiles are not *forced* to tile aperiodically. They themselves cannot recreate the hierarchical structure of figure 1.

On the other hand, the tiles in figure 2 *must* make tilings in which the original hierarchical L-tiling is clearly visible, when the simple rule of matching black to black and white to white is followed [7] (The proof of this is by different techniques than that of the theorem below) . That is, local conditions force L-shaped supertiles of arbitrary size!

As the hierarchical structure of figure 1 is precisely reproduced, we say the original substitution tiling has been “enforced” by these marked “matching rule tiles”. The tiles themselves are “aperiodic” and “hierarchi-

cal” as they can only produce aperiodic hierarchical tilings. Precise definitions of these all of these terms are in [6] and elsewhere.

Beginning in 1964 with R. Berge’s insightful paper [1], a number of examples of aperiodic hierarchical tilings have appeared. S. Mozes produced a method for producing matching rule tiles enforcing a certain class of hierarchical tiling [14], but no truly general method had emerged until [6].

This theorem includes all examples known to the author:

Theorem [6] *Every substitution tiling of E^n , $n > 1$, can be enforced with finite matching rules, subject to a very mild condition:*

the tiles are required to admit a set of “hereditary edges” such that the substitution tiling is “sibling-edge-to-edge”.¹

The proof is somewhat technical but rests on simple principles. Here we attempt only to discuss a few general principles that may have wider application.

The primary structures are *addresses* generated by the substitution process; these addresses can serve tilings, configurations, tiles and points and encode a variety of information. In essence addresses are the geometric expression of regular languages, and form an extension of symmetry groups to semigroups [9].

In a sense, addresses are a geometric lexicon of the regular language given by the “substitution graph”. The task is to break the space of addresses into fundamental domains and give rules for correctly piecing them together. These addresses are hung upon *skeletons* of edges in the tiling, which effectively parse the addresses into digits and piece digits into addresses.

Importantly, these are new combinatorial, geometric structures that may be useful in describing natural phenomena.

1.1. APERIODIC TILINGS, HIERARCHICAL TILINGS

The possibility of an aperiodic set of prototiles, in the plane or elsewhere, was essentially ignored until H. Wang and R. Berger began investigating connections between tilings and undecidability. In the early 1960’s Wang

¹The condition really is mild: we must be able to choose vertices and edges for our tiles such that the vertices/edges of the parent coincide with those of the children and such that the vertices/edges of sibling supertiles coincide.

encoded arbitrary Turing machines as tilings of the plane. However he required a somewhat awkward restriction—the placement of a “seed” tile to begin the computation. In 1964 Berger removed this restriction by constructing an underlying hierarchical structure; this structure provided arbitrarily large, enclosed domains for the run of the machine [1].

Berger thus could answer that the “Domino Question” is undecidable; that is, there is no algorithm to decide whether any given set of prototiles admit a tiling. Berger’s tiles admitted a tiling if and only if the underlying Turing machine did not halt: If every set of prototiles admitting some tiling admits a periodic tiling, one can produce an algorithm to check whether any given set of prototiles actually admits a tiling. Thus Berger produced an unexpected corollary—*there are sets of prototiles that do tile the plane, but cannot do so periodically!* He gave such a set, of some 20,000 tiles.

By 1971 R. Robinson had produced a much smaller set of six prototiles and streamlined Berger’s general result as well in an especially lovely and readable paper [20].




Figure 3. An aperiodic collection of square lattices and Robinson’s tiling

Robinson’s set of prototiles enforced aperiodicity by ensuring a specific

structure was formed: an infinite hierarchy of similar square lattices in the plane (figure 3). The collection as a whole is aperiodic, since no translation can leave every scaled lattice invariant.

R. Penrose found his celebrated tiles in 1972 [15]. The Penrose tiles have become canonical examples of aperiodic tiles, for they are both aperiodic hierarchical tiles but also are quasiperiodic (see below).




Figure 4. The penrose rhombs

Many other examples of aperiodic hierarchical tilings have been found since. In particular R. Amman, J. Socolar, and L. Danzer have constructed many elegant examples [10],[24],[4]. C. Radin has produced matching rules for the Conway pinwheel [17]. S. Mozes gave a general method for constructing specific class of aperiodic hierarchical tilings [14], and as described, a very general method is given in [6].

Only three other methods of constructing aperiodic tiles have emerged:

N.G. De Bruijn pointed the way to the second interpretation in 1981: aperiodic tilings could arise as projections of slices through a higher dimensional lattice [2]. That is, just as the penrose tiles *appear* to be projections of cubes, they indeed are, in a nice manner: the rhombs of the tiling are projections of the squares lying in a certain slab through a five dimensional cubic lattice. This “slice and project” method of constructing aperiodicity

has been very much studied. Recently Le T.T.Q. has given matching rules for a wide class of such tilings [12].

In 1988, Peter Schmitt gave a third class of aperiodic tiling— tilings arising from a single prototile in E^3 . These tilings are not isotropic and thus Schmitt introduced a new symmetry of space. This example was extended upon by Conway and Danzer [3].

In 1996 a fourth, completely new method of constructing aperiodic tilings has emerged: J. Kari has given a construction returning all the way to Wang for inspiration. Essentially he avoids the need for a seed tile by running infinite computations on infinite sequences [11]. Kari and K. Culik are beginning to extend this exciting work.

1.2. A VARIETY OF SUBSTITUTION TILINGS

We pause for a few examples of substitution tilings— heirarchical arrangements of tiles in the plane. A more lengthy discussion of certain technical issues can be found in Appendix A of [6]; a formal discussion is in Section 1.2 of [6].

The substitution tiling on the left of figure 5 is periodic and admits an infinity of hierarchies. Generically, however, if we use marked squares the resulting substitution tiling admits only one hierarchy! (That is, the tiling on the right has *unique decomposition*. Solmyak has proved that aperiodic hierarchical tilings for which a certain natural condition is satisfied have unique decomposition.)




Figure 5. Two quite distinct substitution tilings

In figure 6 are a number of examples, clockwise from upper-left: the “Sphinx”, discussed at length in [Godr] with matching rules provided by E.A. Robinson [RobE]; the “Pinwheel”, found by Conway with matching

rules provided by Radin [Rad]; a “Dimer” tiling, with matching rules found by E.A. Robinson; a triangle tiling found by the author; a triangle tiling found by Danzer and featured in Appendix B of [6]; and the “Half-hex”, with matching rules provided by Socolar [Soc]. These examples all generalize in a variety of ways.




Figure 6. The Sphinx, the Pinwheel, the Half-Hex and the Dimer, and two triangle tilings

Our last example, however points the way towards a far more general setting. Here we are no longer restricted to Euclidean congruences. We still might say that all the quadrilaterals of a given level in the hierarchy are congruent, and more or less leave the geometry at that. In this way,

substitution structures— addresses— may gain great utility, with widespread applications.




Figure 7. A substitution tiling?

2. Addresses

The following is developed more fully in [9] and [6]. The reader is asked to consider the useful properties of graph paper, or a ruler, as illustrated at the top of figure 8. First, each point on the paper or ruler is located to arbitrary precision through algorithmically produced strings of digits. These strings of digits have a hierarchical structure.

Although points are identified by infinite strings of digits, these strings are not necessarily unique. For example (in binary representation) $.01\bar{0} = .00\bar{1}$; nonetheless, there are simple algorithms that can identify when two strings refer to the same point (indeed, we learn such an algorithm as children).

Finally, the geometry of the space is reliable with respect to the addressing given by the graph paper or ruler. A certain geometry is closely bound to the particular coordinates imposed by the graph paper.

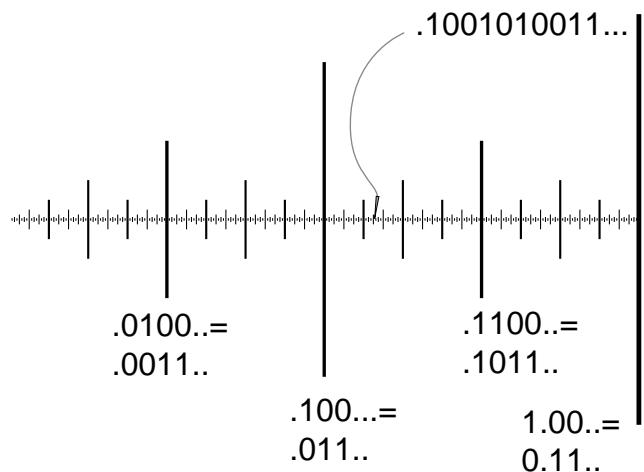


Figure 8. A ruler (base 2)

We will use the L-tiling to illustrate our addressing scheme. Consider the following:

First, label each image of the L-tile in the substitution rule. There are two games we might like to play: label points in an L-tile, and label L-tiles in the tiling.

To be useful, this labeling must have a few properties: We would like this labeling to be algorithmic, and algorithms to be available to tell us when two labels describe adjacent L-tiles, or given a label what the labels of adjacent L-tiles are. Similarly, we might allow points to be labeled in more than one way, but at least want an algorithm to determine when two labels describe the same point; and a way to find, given some label, the other labels describing the same point. Finally, we would like the addressing to be sensitive to the geometry of the substitution scheme.

So begin with addresses for inflations of a single tile. We start with a \bullet , a decimal point. If our original tile is, in say, the inside corner of an inflated L-tile, we label $\mathbf{a}\bullet$ its position with respect to the larger tile. If this larger tile is, in say, the outer corner of a twice inflated L-tile, the label of the smallest tile with respect to this larger tile is $\mathbf{ba}\bullet$. In this fashion we can describe positions of tiles with respect to arbitrarily large L-shaped supertiles through a string $\mathbf{x}_n \dots \mathbf{x}_1 \bullet$, $\mathbf{x}_i \in \{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}\}$.

In fact, these strings are exactly the words in the regular language given by the “substitution graph” (a finite state automaton) in figure 9. (The nodes of the graph correspond to the original tiles; the arrows to the images of the tiles in the substitution. Addresses are directed paths in the graph.




Figure 9. Addresses

A more interesting example is in figure 7 of [6]). Hence these labels are algorithmic.

Indeed we can describe positions in “infinitely large supertiles” $\dots x_1 \bullet$. These infinitely large supertiles might be tilings of the whole plane, but could be a tiling of just a portion of the plane; again it is straightforward to construct an automaton to test which case the address describes. For example, it is clear that any infinite-to-the-left string $\dots x_i \dots \bullet$, for which there is some N with $x_i = a$ for $i > N$, describes a tile in an infinitely large supertile covering three-quarters of the plane, whereas a “generic” string most likely corresponds to a tile in an infinitely large supertile that covers the entire plane.

Paradoxically, there are uncountably many distinct infinitely large supertiles, even if we allow equivalence up to isometry: there are simply uncountably many infinite-to-the-left-strings; moreover, two strings describe equivalent (up to a isometry) infinitely large supertiles if and only if the strings are equivalent to the left of some digit.

In [6] this labeling of hierarchies is precisely what we show is recreated

by matching rules.

Many interesting tricks can be performed [9]; for example, how can one detect that when strings describe adjacent L-tiles in some infinitely large supertile? In figure 10 a portion of an automaton to carry out such a check is illustrated. (A machine of this sort is a *Mealy* machine). Strings of digits simply are paths that follow the arrows; infinite-to-the-left strings are thus paths that have no beginning but do terminate. Two strings $\dots x_i \dots \bullet$ and $\dots y_i \dots \bullet$ describe adjacent L-tiles if there is a path with each arrow labeled x_i/y_i . The portion shown here illustrates the adjacency of the strings in boldface in figure 9. The reader is encouraged to complete this finite state automaton (there is one state in the automaton for each kind of edge that appears as we inflate the original L-tile).²




Figure 10. A Mealy machine describing adjacencies along the indicated edge

Similarly we can describe points in a tile as infinite-to-the right-strings $\bullet x_0 \dots$, and of course points in infinitely large L-tiles as bi-infinite strings. Again, points may be described by more than one string. Strikingly, the same Mealy machine that describes adjacency of tiles also describes when strings label the same point, and can be used for identifying pairs of infinite supertiles that can be matched to make a substitution tiling of the full plane.

3. Keys

We now turn a technique for unambiguously encoding the hierarchy: local and regional keys.

The goal here is to encode the labelings of all tiles in a substitution tiling *only using a finite amount of information in each tile*. For now we will assume that somehow finite amounts of information can be compared

²Such an automaton can always be found up to some condition on the substitution tiling; the author believes that requiring vertex-to-vertex is more or less sharp.

between neighboring supertiles and between a supertile and its parent and children.

The author imagines an infinite community piecing together a family tree through the exchange of brief postcards.

We use a simple idea: certain tiles—*keys*—are to determine large amounts of the tiling. Every tile is labeled with the last digit of its address: that is, its intended position with respect to its parent. Keys, however, will be labeled with one additional digit, giving the position of some ancestral supertile with respect to its parent.

We choose all the tiles labeled, say, \mathbf{a} to be *local keys* and all the tiles labeled, say \mathbf{c} to be *regional keys*. Thus a label \mathbf{ax} indicates the tile is in position \mathbf{a} with respect to its parent, and its parent is in position \mathbf{x} with respect to *its* parent. Now if \mathbf{x} is itself a key, this parent is supposed to carry an additional label. This label is tucked into the regional key adjacent to our original tile. That is, suppose the parent supertile is supposed to be labeled \mathbf{xy} . Then the local key within this parent is \mathbf{ax} and the regional key within this parent is \mathbf{cy} . Of course if the parent is not itself a key, \mathbf{y} is vacant.

In this fashion, the entire hierarchy can be encoded with only one or two digits in each tile — if information can be transmitted over arbitrary distances from regional keys to the boundary of the supertile the key labels. In figure 11 a portion of an L-tiling is encoded. The labels \mathbf{b} and \mathbf{d} , and unfilled keys \mathbf{c} have been left out of the illustration for clarity.



This simple mechanism is, forgive the author, the “key” to the whole construction (although it is deeply disguised and virtually unidentifiable in [6]).

4. Mechanisms

We prove the Theorem as follows: we first construct a new set of marked tiles and matching rules by examining structures in supertiles of some bounded size. We define “well formed supertiles” as configurations of the new tiles that are in some essential way equivalent to supertiles in the original tiling. Then we proceed by induction:

We assume that every tile in the matching rule tiling lies in a well-formed supertile of level n , and show this implies every tile lies in supertile of level $n + 1$.

The key, of course, is selecting the right structures to encode in the new marked tiles. Consider what properties we need from supertiles to force them to organize into the next larger size supertile:



a is local key
 c is regional key

 Thus ax cy indicate
 next level up is
 xy

 labels b, d and
 unfilled c_
 have been re-
 moved for clarity

Figure 11. Encoding addresses through local and regional keys

- Supertiles should be the correct shape and size.
- Supertiles should be labeled with their intended position with respect to their parent.
- Supertiles should be “combinatorially inert” except at a few special points. (Since every neighborhood in the tiling is to contain only a finite amount of information—and so can only serve a finite number of supertiles— whereas a given point may be on the boundary of infinitely many supertiles.)
- At these combinatorially active sites, sibling supertiles are to compare labels, fix their relative positions and orientations, and transmit information to the parent. “Sites” and “terminals” play this role in the

construction.

- Finally there must be some structure to transmit information from the sites connecting children to parent to the sites connecting parent to grandparent. Moreover, every neighborhood in the tiling should only lie in finitely many of these networks. (As a bonus, however, such a network is the ideal place to store key labels). “Skeletons” and “wires” play this role in the construction.

However, it is not hard to always construct such structures in any given substitution tiling. The seemingly endless technical details in [6] arise because of numerous special situations.

We will only discuss one of the structures we exploit: the *skeleton*, a self-organizing structure that simultaneously contains and transmits information about each supertile.

In the lower left of figure 12, the substitution for the pinwheel tiling [Rad] is shown; above and to the right skeletons for three generations of supertile are shown. Note that the skeletons are each connected sets, are connected to each other at special points (sites), include all the “highest level” edges in the supertile, every edge is in only finitely many (i.e. one) skeletons, and that the skeletons are all similar.




Figure 12. Skeletons

Given the children supertiles are formed correctly, in order for the parent’s skeleton to form correctly, we must know the locations of the endpoints (“terminals”) of each of its edges. If the endpoints lie on the skeletons of lower level supertiles, no further work is needed. However, we sometimes need an additional structure: we link certain terminals (“mesovertrices”) 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.

5. On the utility of substitution tilings

The title of this section admittedly overreaches its contents. But it is the author's belief that, ultimately, this will make a fine title for a much longer work. At the moment, this is only conjecture.

We will give two, simple, “theoretical applications.”

1) Grids are used for any number and variety numerical simulations. There is a seeming trade-off that must be faced by the designer of the simulation:

One can choose a highly regular grid, in which one can easily address locations, adjacencies, etc. without the use of direct pointers between neighboring cells. But such a grid may not easily conform to the geometry of the simulation, or worse, the geometry of the grid may introduce extraneous results. For example, suppose we simulate the geometry of the circle on a very fine square grid. Certain properties, such as area, can be calculated to arbitrary precision by subdividing the grid. But others, such as perimeter *will never* be approximated, regardless of how fine the grid is. In fact, even the length of a generic line segment is very badly estimated by a square grid, no matter how fine.

To remedy this one can choose a grid that appears irregular, that is transparent to the geometry of the simulation. But in such a grid, it is difficult to “look up” a given location; the grid is constructed only through local information: which cells are adjacent to which others. This also carries a fairly large memory cost, since all adjacencies must be encoded.

Substitution tilings provide a middle road. They are clearly algorithmic and have all the advantages of a rigid hierarchical structure. On the other-hand, they admit an extraordinarily rich family of possible geometries. To give one example, substitution tilings *generically* are isoparametric— that is, one can “usually” approximate arbitrary smooth curves to arbitrary precision.

2) Substitution tilings provide models of cell-division and growth, often to great effect [16]. Whether matching rules or addressing are actually useful is beyond this author's expertise. But it can be said that matching rules encode local “well-formedness”. Our main theorem can be interpreted as: every sufficiently regular hierarchical structure can be endowed with local conditions such that if the structure is not faithfully reproduced, the local conditions are not satisfied. One might imagine a growing embryo; local conditions might exist which can detect a global defect as soon as it arises. (Whether such a mechanism actually exists is of course quite far beyond the scope of this discussion)

Finally, the proof of the theorem is an exercise in self-organization, a fundamentally interesting phenomenon.

References

1. R. Berger, *The undecidability of the domino problem*, Memoirs Am. Math. Soc. **66** (1966).
2. N.G. de Bruijn, *Algebraic theory of Penrose's non-periodic tilings*, Nederl. Akad. Wentensch Proc. Ser. A **84** (1981), 39-66.
3. L. Danzer, *A family of 3D-spacefillers not permitting any periodic or quasiperiodic tiling*, preprint.
4. L. Danzer, personal communication.
5. C. Godreche, *The sphinx: a limit periodic tiling of the plane*, J. Phys. A: Math. Gen. **22**1989, L1163-L1166.
6. C. Goodman-Strauss, *Matching rules and substitution tilings*, to appear in Annals of Math.
7. C. Goodman-Strauss, *A small aperiodic set of tiles*, preprint.
8. C. Goodman-Strauss, *An aperiodic tiling of E^n for all $n > 1$* , preprint.
9. C Goodman-Strauss, *Addresses and substitution tilings*, in preparation.
10. B. Grunbaum and G.C. Shepherd, *Tilings and patterns*, W.H. Freeman and Co. (1989).
11. J. Kari, *A small aperiodic set of Wang tiles*, preprint.
12. , T.T.Q. Le, *Local rules for quasiperiodic tilings*, Proceedings of the Fields Institute (1995).
13. J.M. Luck, *Aperiodic structures: geometry, diffraction spectra, and physical properties*, to appear in Fund. Prob. in Stat. Mech VIII.
14. S. Mozes, *Tilings, substitution systems and dynamical systems generated by them*, J. D'Analyse Math. **53** (1989), 139-186.
15. R. Penrose, *The role of aesthetics in pure and applied mathematical research*, Bull. Inst. of Math. and its Appl. **10** (1974) 266-71.
16. P. Prusinkiewicz and A. Lindenmeyer, *The Algorithmic Beauty of Plants*, Springer-Verlag (1990).
17. C. Radin, *The pinwheel tilings of the plane*, Annals of Math. **139** (1994), 661-702.
18. C. Radin, personal communication.
19. E.A. Robinson, personal communication.
20. R. Robinson, *Undecidability and nonperiodicity of tilings in the plane*, Inv. Math. **12** (1971), 177.
21. L. Sadun, *Some generalizations of the pinwheel tiling*, preprint.
22. M. Senechal, *Quasicrystals and geometry*, Cambridge University Press (1995).
23. J.E.S Socolar, personal communication.
24. J.E.S. Socolar and P.J. Steinhardt , *Quasicrystals II. Unit cell configurations*, Physical Review B **34** (1986), 617-647.
25. B. Solomyak, *Non-periodicity implies unique composition for self-similar translationally-finite tilings*, to appear in J.Disc. and Comp. Geometry.
26. P.J. Steinhardt and S. Ostlund, *The physics of quasicrystals*, World Scientific (1987).
27. W. Thurston, *Groups, tilings and finite state automata: Summer 1989 AMS colloquim lectures*, GCG 1, Geometry Center.
28. P. Walters, editor, *Symbolic dynamics and its applications*, Contemporary Mathematics **135** (1992).