

On substitution tilings of the plane with n -fold rotational symmetry

Gregory R. Maloney

Newcastle University, United Kingdom

received 13th Sep. 2014, revised 9th Mar. 2015, 27th Apr. 2015, accepted 12th May 2015.

A method is described for constructing, with computer assistance, planar substitution tilings that have n -fold rotational symmetry. This method uses as prototiles the set of rhombs with angles that are integer multiples of π/n , and includes various special cases that have already been constructed by hand for low values of n . An example constructed by this method for $n = 11$ is exhibited; this is the first substitution tiling with elevenfold symmetry appearing in the literature.

Keywords: Rhombs, Tiling, Algorithm

1 Introduction

Rotational symmetry is one of the most distinctive qualities that aperiodic planar tilings can have. The most famous aperiodic family of tilings is the family of Penrose tilings [16], which possess fivefold rotational symmetry. The Ammann–Beenker tilings [3] are arguably the second-most famous aperiodic tilings, and they possess eightfold rotational symmetry. Aperiodic tilings are often used as models of quasicrystals; when Dan Shechtman and his co-authors made the Nobel prize-winning discovery of quasicrystals in [19], the first sign by which they knew that they had found something special was the presence of tenfold rotational symmetry in the X-ray diffraction patterns.

Yet there is a lack of known examples of aperiodic planar tiling families with higher orders of rotational symmetry (see the Tilings Encyclopedia [6] for an extensive list of substitution tilings, including most of the ones cited here). It is true that there are projection tilings possessing arbitrary orders of rotational symmetry—see [2] for a method with $n = 5$ that generalises easily to arbitrary n , and also [22]—but the same cannot be said for substitution tilings. And under a certain assumption substitution tilings possess the desirable property of *repetitivity*, which implies in particular that any local configuration of tiles with rotational symmetry that appears in one such tiling must appear in every ball of sufficiently large radius in every such tiling. Therefore local patterns with rotational symmetry must appear quite regularly in substitution tilings if they appear at all, whereas in projection-method tilings heuristic arguments and empirical observation suggest that patterns with high orders of rotational symmetry must necessarily be sparse (see [1, Remark 7.12, p. 300f] and [14]).

The purpose of this work is to provide a method for finding substitution tilings with high orders of rotational symmetry. In practical terms, “high orders” of rotational symmetry are numbers in the range eleven

to nineteen; indeed, to my knowledge, the highest known order of rotational symmetry of any substitution tiling is twelve (see for instance the variant of the rule from [21] that appears in [6]), and there are no known substitution tilings in the literature with elevenfold symmetry (but see [17] for some unpublished examples similar to the ones presented here, including one with elevenfold symmetry). Therefore the substitution tiling space with elevenfold symmetry described in Section 5 of this paper is the first of its kind to be published.

Much previous work has been done on this topic, some of which claims to produce planar substitution tiling spaces with n -fold rotational symmetry for arbitrary n . My claim that there are no previously-known examples of substitution tilings with elevenfold rotational symmetry would seem to contradict this previous work, but in fact there is no contradiction; instead the confusion lies in differing interpretations of n -fold symmetry in a tiling space. Let us therefore be more explicit about what n -fold symmetry means in the current context.

Under mild conditions on the substitution rule, any two tilings that arise from it will be repetitive and locally indistinguishable from one another, meaning that any bounded patch that appears in one of them also appears in the other. Therefore it is natural to consider the space of all tilings arising from a given substitution rule, rather than just a single such tiling. Many results in the literature use the term n -fold symmetry to refer to tiling *spaces* that are n -fold symmetric; that is, a rotation of any tiling in the space by $2\pi/n$ yields another tiling in the space. This is not the same as saying that the tiling space contains an individual tiling that is invariant under n -fold rotation, which is precisely the property that is sought here.

A tiling that is invariant under n -fold rotation contains bounded rotation-invariant patches of arbitrary radius centred on its centre of rotation. Conversely, the existence of a non-trivial rotation invariant patch implies the existence of a rotation-invariant tiling in the following way. Place the patch with its centre of symmetry over the origin and apply the substitution rule to it repeatedly until two patches are found, one of which is contained in the other and both of which contain the origin in their interior. The smaller of these is the seed for a tiling that is a fixed point of some power of the substitution; this tiling will necessarily have global rotational symmetry.

But we must be careful what we mean by a “rotation-invariant patch.” It is a common situation to have two congruent tiles that are distinguished from each other by giving them different labels. Thus a rotation-invariant patch is collection of tiles for which there is a rotation that sends each tile to another tile in the collection that not only is congruent to the original tile, but also has the same label. Since the substitution behaves in the same way for two tiles with the same label, this guarantees that the image of the patch after several substitutions will also be rotation-invariant. This condition is not met in [15, Remark 6.3], in which there are patches consisting of fourteen isosceles triangles that appear to have fourteenfold rotational symmetry, but that in fact only have twofold symmetry. This is because some of the tiles are reflected copies of the others, and hence they have different labels.

The goal of this work is to find substitution tilings that contain patches of bounded size that are invariant under n -fold rotation in the sense just described. There are two types of tiles that are naturally suited to this purpose. The first is the collection of triangles with angles that are integer multiples of π/n ; the second is the set of rhombs with the same angles. These tiles have been used in substitution tilings before (see [15, 4, 8, 9, 7] for triangles and [16, 3, 10, 17] for rhombs), but except for a few specific examples with low values of n , it is always the tiling space itself that exhibits rotational symmetry, and not any individual tiling. The examples produced here do have n -fold symmetry; specifically, they contain patches consisting of stars of n congruent rhombs, each of which has angles $2\pi/n$ and $(n-2)\pi/n$.

Most previous work on this subject has focused on examples that can be created by hand. It is possible

to create substitutions with fivefold or sevenfold symmetry by hand, putting families of tiles together to form larger copies of themselves, but for higher orders of symmetry the scale becomes too large for this method to work. A common approach to this problem of scale has been to take a known substitution rule with fivefold or sevenfold symmetry and then to generalise it to higher n in some way. This has led to the discovery of various infinite families of tiling spaces, but in all cases the desirable property of having an individual tiling with n -fold symmetry is lost upon passage to higher n .

The approach here is not to use any specific fivefold or sevenfold example as a model, but rather to apply in a systematic way the ad hoc methods that have been used to discover such examples. The scale of this problem at higher values of n means that it is necessary to use a computer. In particular there is an algorithm of Kannan, Soroker, and Kenyon [11, 12] that will produce a substitution rule of a given size and shape, if it is possible to do so. Then a result of Kenyon says that any other substitution rule with that size and shape can be obtained from this one by a series of transformations, called rotations. Again with the help of a computer we can transform the initial substitution rule via rotations until it produces n -fold symmetric patches. This is not possible for every size and shape of substitution, but for large enough substitutions it works.

It should be mentioned that the method described here relies mostly on existing algorithms and ideas. The chief innovation is Rule 2.6 in Section 2. Nevertheless, it should not be dismissed just because its ingredients are not novel. Indeed, these ingredients have been known for many years, during which many researchers have tried without success to produce n -fold symmetric substitution tilings for large n ; the fact that this has been achieved here is evidence that the combination of these ingredients in this particular way is a significant development.

2 Definitions and notation

Let us suppose henceforth that an odd integer $n \geq 3$ is fixed. The case when n is even is addressed in Section 6.

Definition 2.1. A tile is a subset of \mathbb{R}^d homeomorphic to the closed unit disk. A patch is a collection of tiles, any two of which intersect only in their boundaries. The support of a patch is the union of the tiles that it contains. A tiling is a patch, the support of which is all of \mathbb{R}^d .

Let us restrict our attention to the case $d = 2$, and let us consider a restricted set of tiles consisting of all rhombs with sides of unit length and angles that are integer multiples of π/n . Let us select a representative for each isometry-equivalence class of such tiles: for each even integer $i < n$, let $r(i)$ denote the rhomb with vertices $(0, 0)$, $(1, 0)$, $(1 - \cos(i\pi/n), -\sin(i\pi/n))$, and $(-\cos(i\pi/n), -\sin(i\pi/n))$. Let $R(n) = \{r(i) \mid 2 \leq i < n, i \text{ is even}\}$. Then $R(n)$ contains $\lfloor \frac{n}{2} \rfloor$ elements if n is odd, where $\lfloor \cdot \rfloor$ denotes the floor function. The elements of $R(n)$ are called *prototiles*, and we denote by $\mathcal{P}(R(n))$ the set of all patches consisting entirely of tiles congruent to these prototiles. Figure 1 depicts the rhombs $R(7)$.

This work deals with tilings that arise from substitutions. A substitution is traditionally defined to be a map on a set of prototiles that assigns to each prototile p a patch, the union of which is λp for some $\lambda > 1$, called the *inflation factor*. However, in the present context we will need to relax this definition slightly; in particular, we will need to deform the edges of λp . For this it will be necessary to introduce some new notation and terminology.

For an integer k , let e_k denote the direction vector beginning at $(0, 0)$ and ending at $(\cos(k\pi/n), \sin(k\pi/n))$.




Figure 1: The prototiles $R(7)$. Integers in the corners represent angle measures expressed as integer multiples of $\pi/7$.

Definition 2.2. An edge sequence is a finite sequence of integers $\sigma = (k_i)_{i=1}^m$. A standard edge sequence is one in which the number k occurs as often as the number $-k$, 0 appears arbitrarily many times, and $|k_i| < n/2$ for all $1 \leq i \leq m$. To an edge sequence σ and a point $x \in \mathbb{R}^2$ we associate a sequence $E_\sigma(x)$ of line segments, each of which begins at the end of the previous one and has direction vector e_{k_j} . The first line segment of E_σ begins at the point x .

Given an edge sequence σ , let $\sum \sigma$ denote the point $\sum_{i=1}^m e_{k_i}$. Then the endpoint of $E_\sigma(0)$ is $\sum \sigma$.

Given an integer j , define the rotation of σ by j to be the edge sequence $\sigma(j) := (k_i + j)_{i=1}^m$.

The conditions in the definition of a standard edge sequence σ imply that $\sum \sigma$ lies on the positive x -axis. A standard edge sequence should be thought of one that has the same number of steps up at a given angle of inclination as it has steps down at the corresponding supplementary angle.

Now we can use a standard edge sequence σ to distort the boundaries of the rhombs in $R(n)$.

Definition 2.3. Let σ be a standard edge sequence σ and $i < n$ an even integer. The σ -boundary of $r(i)$ is $\sigma r(i) := E_\sigma(0) \cup E_\sigma(\sum \sigma(-i)) \cup E_{\sigma(-i)}(0) \cup E_{\sigma(-i)}(\sum \sigma)$.

The end points of the various line segment sequences in the σ -boundary of $r(i)$ form the vertices of an inflated copy of $r(i)$. Figure 2 depicts the $(1, -1, 0)$ -boundaries of $r(6)$, $r(4)$, and $r(2)$ for $n = 7$. The outlines of the inflated copies of these rhombs are depicted as dashed lines.

The σ -boundary of $r(i)$ is the image of a closed curve, specifically the curve defined by traversing at a constant speed $E_\sigma(0)$ and $E_{\sigma(-i)}(\sum \sigma)$ in the positive direction, then $E_\sigma(\sum \sigma(-i))$ and $E_{\sigma(-i)}(0)$ in




Figure 2: σ -boundaries of $R(7)$ for $\sigma = (1, -1, 0)$. The outlines of the inflated copies of these rhombs are depicted as red dashed lines. The edges of $\sigma r(i)$ are depicted as solid black lines and the edges of $\sigma r(4)$ are labelled.

the negative direction. This curve is not necessarily simple—see for example $\sigma r(6)$ in Figure 2, which backtracks on itself near the bottom-left corner.

The goal here is to define a substitution rule on $R(n)$ that sends each prototile $r(i)$ to a patch, the support of which is the closure of the planar region inside of $\sigma r(i)$. This is only possible if it is understood what is meant by “the planar region inside of $r(i)$.” When $\sigma r(i)$ is a simple curve, the meaning of this is clear. It is also clear that $\sigma r(6)$ in Figure 2, which involves some backtracking, has a well-defined region inside of it. The notion of an inside makes sense for a more general class of curves, although they are somewhat difficult to describe.

Definition 2.4. Let E be a closed path of directed line segments of unit length. E is a good curve if it satisfies the following.

1. If any two line segments overlap at a point other than their endpoints, then both segments are the same with opposite orientations; and
2. it is possible to perform small homotopies to separate all pairs of overlapping oppositely-oriented line segments to obtain a planar-embedded graph for which, at any vertex where four edges meet, those four edges are alternating in- and out-edges of that vertex.

Note that condition (2) implies that all points contained in a connected component of a good curve’s interior (i.e. a bounded connected component of the complement of the curve) are wrapped around by the curve once, in the same direction.

The curve in Figure 3 is a good curve, as can be seen on the right side of the figure, in which homotopies fulfilling the requirements of Definition 2.4 are exhibited.

Definition 2.5. A substitution on $R(n)$ is a map $\varphi : R(n) \rightarrow \mathcal{P}(R(n))$ for which there is an edge sequence σ such that each $\sigma r(i)$ is a good curve, the closure of the inside of which is the support of $\varphi(r(i))$.




Figure 3: The $(0, 3, 1, 2, -1, -2, -3, 0)$ -boundary of $r(4)$. The picture on the right shows the images of the double lines under homotopies.




Figure 4: A substitution using the $(1, -1, 0)$ -boundaries of $R(7)$.

Figure 4 depicts a substitution on $R(7)$ that uses the $(1, -1, 0)$ -boundaries; these boundaries appeared earlier in Figure 2. This substitution does not give rise to any tilings with global sevenfold rotational symmetry.

The patches in Figure 4 are ambiguous because they do not specify how the tiles are obtained as images of the prototiles under isometries. The prototiles in $R(7)$ are all self-symmetric under reflection and rotation by a half turn. Therefore each rhomb that appears in the figure could be the image of a prototile under any one of four isometries, two of which are orientation preserving and two of which are orientation reversing. In order to determine a substitution rule fully, we must specify which of these isometries to use for each tile in each patch.

The rule for specifying these isometries is simple but important; without this rule it would not be possible to iterate φ . Let us assign an orientation to each line segment L in the plane that is parallel to one of the direction vectors e_k .

Rule 2.6. If the line segment L is parallel to e_k for an even integer $0 \leq k < n$, then give L the same orientation as e_k ; otherwise give it the opposite orientation.

This gives orientations to the edges of all of the rhombs in $R(n)$, and also to the rhombs in the patches $\varphi(r(i))$. These orientations have the property that, if two line segments have a common end point, then if the angle between them is an even multiple of π/n the line segments are oriented either both toward or both away from their common end point; if the angle between them is an odd multiple of π/n then one of the line segments points toward the common end point and the other one points away from it.

These orientations determine, for each tile in a patch, a unique orientation-preserving isometry that carries that tile to its associated prototile. Figure 5 depicts the three prototiles in $R(7)$ and the patch




Figure 5: The prototiles $R(7)$ and the substituted image $\varphi(r(2))$, all with their edge orientations.

$\varphi(r(2))$, all with their edge orientations.

The reason for establishing these orientation rules is that it allows us to extend φ to arbitrary patches, which in turn means that we can iterate φ . This is a crucial procedure in the construction of substitution tilings (see [1, Chapter 6], where the less standard but more precise term *inflation tilings* is used).

First let us extend the domain of a substitution φ to include all tiles of the form $g(r(i)) + v$, where g is an orthogonal transformation of the plane and v is a translation vector. For this it will be necessary to generalise the notion of an *inflation factor*.

Definition 2.7. Let φ be a substitution with associated boundary sequence σ . Then the inflation factor of φ is the positive number $\lambda = \|\sum \sigma\|$.

Then given a substitution φ with inflation factor λ and a transformed prototile $g(r(i)) + v$, let us define

$$\varphi(g(r(i)) + v) := \{g(t) + \lambda v \mid t \in \varphi(r(i))\}.$$

Finally, for a patch P consisting of rhombs with edges parallel to the direction vectors e_k , let us define

$$\varphi(P) := \{\varphi(t) \mid t \in P\}.$$

The σ -boundaries of the prototiles in $R(n)$ are consistent with the orientations stipulated by Rule 2.6. This means that the image under φ of the patch P is also a patch—boundaries of two tiles that share an edge are distorted in the same way, so there is no overlap (of positive measure).

In particular every set $\varphi(r(i))$ is a patch consisting of rhombs with edges parallel to the direction vectors e_k , so we can define $\varphi^2(r(i)) := \varphi(\varphi(r(i)))$ and so on.

Definition 2.8. Let φ be a substitution. A substitution tiling arising from φ is a tiling, every bounded patch of which is contained in a translate of $\varphi^k(r(i))$ for some $k \in \mathbb{N}$ and some i .

It is a typical situation to have the same substitution give rise to infinitely many different tilings that are indistinguishable from one another by looking only at a bounded patch. Under the following mild condition on φ one can show that it gives rise to at least one substitution tiling.

Definition 2.9. The substitution φ is primitive if there is an integer k for which, given any ordered pair p_1, p_2 of prototiles, there is a translate of p_2 in $\varphi^k(p_1)$.

If φ is a primitive substitution then there is some translate t of a prototile p and a positive integer k such that the sequence $\varphi^{mk}(t)$ of patches is consistent in the sense that each one contains the previous one. With a bit more work, it can be arranged that the support of one contains the support of the previous one in its interior, so that these patches grow to cover the entire plane, and hence their union constitutes a tiling, which is in fact a substitution tiling arising from φ . This process is described in greater detail in [1, Chapter 6], where Figure 6.13 depicts the first two elements in such a sequence of nested patches that grows to cover the plane.

As described in Section 1, there are many known substitutions that give rise to substitution tilings with the property that the rotation of one such tiling by $2\pi/n$ is another such tiling. The objective here is to find substitutions that give rise to at least one substitution tiling that is invariant under rotation by $2\pi/n$.

3 The method

3.1 The Kannan–Soroker–Kenyon criterion

To find a substitution on the prototiles $R(n)$ that has n -fold rotational symmetry, the first step is to choose a standard edge sequence σ . This involves some work; let us discuss this further in Section 4.

The next step is to determine if there is at least one substitution φ associated to σ . This part also requires work, but fortunately Kannan and Soroker [11] and Kenyon [12] have developed a simple criterion for determining if such a substitution exists; this criterion can be checked with a computer. More specifically, the Kannan–Soroker–Kenyon (KSK) criterion is a test that determines if the region bounded by a curve consisting of unit-length line segments is tilable by the prototiles of $R(n)$; if each boundary $\sigma r(i)$ satisfies this criterion, then a substitution exists.

The KSK criterion can be applied to any polygonal boundary, but in the special case of the boundaries $\sigma r(i)$ it can be expressed in particularly simple terms. This involves *pseudoline arrangements* (see [5], where they are called *worms* following Conway). Specifically, if the region inside of $\sigma r(i)$ can be tiled by rhombs from $R(n)$, then it has an associated pseudoline arrangement.

Definition 3.1. A pseudoline is a smooth planar curve, and a pseudoline arrangement is a collection of pseudolines, any two of which, if they intersect, do so at most once, and cross each other at their point of intersection. Let us impose a further condition in the definition of a pseudoline arrangement that is not standard in the literature: no three pseudolines may share a common intersection point.

Then a pseudoline arrangement is associated to a tiling of the inside of the region bounded by $\sigma r(i)$ in the following way. Group the rhombs in the tiling into inclusion-maximal lines of linked rhombs with a common edge direction, each of which shares a parallel edge with the next. Every rhomb falls into exactly two such families. A pseudoline arrangement is obtained by drawing a pseudoline through each such family of rhombs.

Figure 6 depicts such a pseudoline arrangement for the patch $\varphi(r(2))$ from Figure 4. These pseudoline arrangements are considered equivalent up to diffeomorphisms that preserve the “no triple intersections” property, so the pseudoline arrangement on the left of that figure is equivalent to the one on the right, which has been deformed for clarity of drawing, and from which the background tiling has been removed.




Figure 6: Pseudoline arrangement for $\varphi(r(2))$ from Figure 4. The picture on the right is the same arrangement with the endpoints of the pseudolines shifted to lie on a square.

Figure 6 also has numbers at the ends of the pseudolines indicating the angles that the corresponding line segments of $\sigma r(2)$ make with the positive x -axis. Each pseudoline joins two edges with angles that differ by π , so labels of the end points of a single pseudoline always differ by n . This is true for any n and for any pseudoline arrangement arising from a tiling of the inside of $\sigma r(i)$.

Moreover, Definition 2.2 of a standard edge sequence implies that each angle $0 \leq j < 2n$ will only occur as a label of line segments on two of the four sides of $\sigma r(i)$, and those two sides cannot be opposite. Therefore, when reading the labels cyclically around the edge of $\sigma r(i)$, the labels j will not appear interspersed with the labels $n + j$. Also, two pseudolines with ends labelled j and $n + j$ cannot intersect one another in an arrangement arising from a tiling. These two facts together imply that there is a unique way to partition the line segments in $\sigma r(i)$ into pairs connected by a common pseudoline.

Whether a tiling of the inside of $\sigma r(i)$ exists or not, it is possible to label the edge segments with their angles modulo $2n$, and to determine which must be connected to which by pseudolines. It is also possible to determine which pairs of pseudolines must intersect each other; the only thing that cannot be determined just from looking at the boundary $\sigma r(i)$ is the order in which the pseudolines intersect. But it is not necessary to determine this in order to know whether a tiling of the inside of $\sigma r(i)$ exists or not.

If a tiling of the region bounded by $\sigma r(i)$ exists, then an intersection of two pseudolines in the associated pseudoline arrangement corresponds to a rhomb in the tiling. The angles in that rhomb can be calculated by subtracting the angles of the endpoint labels of the two pseudolines: pick an endpoint label of one of the lines and subtract it modulo $2n$ from the next label, in counterclockwise order, of the other line, to obtain one of the angles of the rhomb (expressed as an integer multiple of π/n). The KSK criterion simply says that a tiling of the region bounded by $\sigma r(i)$ exists if and only if the subtractions for all such pairs of crossing pseudolines yield angles less than π .

Figure 7 depicts the $(0, 3, 1, 2, -1, -2, -3, 0)$ -boundary of $r(4)$ from Figure 3 along with an associated pseudoline arrangement. The region inside this boundary cannot be tiled by the rhombs in $R(7)$, as can




Figure 7: Pseudoline arrangement for the $(0, 3, 1, 2, -1, -2, -3, 0)$ -boundary of $r(4)$ from Figure 3. The red dot indicates the crossing pseudoline pair that violates the KSK criterion, and numbers indicate angles that line segments make with the positive x -axis.

be seen from the pseudoline arrangement, which contains a red dot indicating the crossing pseudoline pair that violates the KSK criterion. There are many pseudoline arrangements that we could draw for this boundary, and all of them would contain an intersecting pair of pseudolines connecting the same pairs of end points, so none of them is a pseudoline arrangement arising from a tiling by the elements of $R(7)$.

3.2 Rotations

Not only did Kannan, Soroker and Kenyon give a criterion to determine if there is a tiling of the inside of $\sigma r(i)$, they also provided an algorithm to produce such a tiling when it exists. The details of this algorithm are described in [11] and [12]. By applying this algorithm to each boundary $\sigma r(i)$, we can produce a substitution.

The final step of the method is to modify this substitution to produce a substitution that yields rotational symmetry. Kenyon [12] showed that every tiling of the region bounded by $\sigma r(i)$ using rhombs in $R(n)$ is related to every other such tiling by a sequence of simple changes, called rotations. Rotations are defined on patches that are triples of rhombs, any two of which share an edge. The support of such a triple is a hexagon with opposite edges parallel and of equal length. A rotation replaces such a triple with another triple of rhombs each of which is a translate of one of the original three rhombs by a unit vector that is a common edge of the other two rhombs. The new triple has the same support as the old one. Figure 8 depicts a rotation, with the corresponding change in the pseudoline arrangement underneath.

If we represent the tilings $\sigma r(i)$ using an appropriate data structure, then it is possible to implement rotations with a graphical interface, which makes it easy to find patches with the desired symmetry properties. In particular, it is possible to arrange for stars consisting of copies of $r(2)$ to appear inside of the patches $\varphi(r(i))$. Another strategy is to put copies of $r(2)$ in all the corners of each patch $\varphi(r(i))$ so that, after a second application of φ , stars appear in the corners of the first supertiles. This is what has been done in the example with elevenfold symmetry in Figure 9.

Not all edge sequences yield substitutions, and not all edge sequences that yield substitutions yield




Figure 8: A rotation. The corresponding change in the pseudoline arrangement is depicted underneath.

substitutions that produce n -fold rotational symmetry. Nevertheless, by passing to large enough edge sequences, it seems that it is always possible to find the desired symmetry.

3.3 Variants

Let us note here that some well-known substitution tilings arise from very similar methods, but with some important differences.

The Penrose rhomb substitution [16] uses the prototile set $R(5)$ and also involves tiling inflated prototiles with distorted boundaries, but pairs of opposite edges do not share the same distortion. Indeed, for each pair of opposite edges in a prototile, one edge is distorted with the edge sequence $(1, -1)$ and the other is distorted with $(2, -2, 0)$. This exploits the fact that, if σ_e and σ_o are sequences of all the even and all the odd numbers, respectively, between $-n/2$ and $n/2$, then the vectors $\sum \sigma_e$ and $\sum \sigma_o$ are the same. This follows easily from the fact that the sum of the edge vectors in a regular n -gon is the zero vector.

This approach works for $R(5)$, but it is very difficult to generalise because if opposite edges are not distorted in the same way, then Rule 2.6 is no longer sufficient to guarantee that the image of an arbitrary patch under φ is still a patch of non-overlapping tiles. Instead one must check very carefully to be sure that any two rhombs that meet in $\varphi(r(i))$ have the same distortion on their common edge.

Another well-known substitution that arises from a similar method is the binary substitution [13], which also uses the prototiles $R(5)$ with distorted boundaries. The difference in this case is that the boundaries are distorted using sequences of unit-length line segments that are parallel to the edges of a regular decagon, not just a regular pentagon. In other words, the edge sequence uses not only integers, but also half integers. By stretching Definition 2.2, one could say that the binary tiling uses the edge sequence $(1/2, -1/2)$.

Another class of examples that use half-integer edge sequences appears in [17]. These examples use a family of edge sequences parametrised by n , and, at least for values of n up to and including $n = 11$, they yield boundaries that satisfy the KSK criterion, and therefore give rise to substitution rules. Moreover, these substitution rules exhibit n -fold rotational symmetry, at least up to $n = 11$.

4 Finding boundaries that work

4.1 Permutations

It can be difficult to find an edge sequence σ that produces distorted boundaries $\sigma r(i)$ that all have tilable interiors. It would be good to have some general rules to help determine which sequences will work and which will not, but the approach used here is a simple brute-force computer search. This computer search begins with a standard edge sequence σ and permutes it, and for each permutation σ' applies the KSK criterion to see if all the σ' -boundaries enclose tilable regions.

4.2 Iterators

The edge sequence σ can be quite large—for instance, the elevenfold example in Section 5 uses a sequence of length 35 with many repetitions—so finding all of its permutations involves non-trivial work. Specifically, it is not practical to store a list of all permutations of such a large sequence in memory; a much better idea is to use an iterator that can produce the next permutation—in some specified order—given the current permutation. Such an iterator only holds a single permutation in memory at one time, which reduces memory use enough to make it possible to search the permutations of very large sequences.

The iterator used here is a variant of the cool-lex iterator [18] that has been adapted for iterating through permutations of multisets [23].

4.3 Simplifying the search

Even when an iterator makes it feasible in theory to check the KSK condition for every permutation of an edge sequence, it may still take a prohibitively long time. For such cases, it is sometimes useful to reduce the number of permutations checked by restricting attention to permutations that are obtained by concatenating certain selected subsequences.

For example, the elevenfold example in Section 5 uses the edge sequence

$$\sigma = (-1, 1, -3, 3, 0, 2, -2, -1, 1, 0, -5, 5, -3, 3, -1, 1, 4, -4, \\ 2, -2, 0, -1, 1, 2, -2, -3, 3, 0, 4, -4, -1, 1, 2, -2, 0).$$

This edge sequence was found using a brute-force search as described above, but rather than checking all permutations of the following multiset

$$\begin{array}{cccccc} 0 \times 5, & 1 \times 5, & 2 \times 4, & 3 \times 3, & 4 \times 2, & 5 \times 1, \\ -1 \times 5, & -2 \times 4, & -3 \times 3, & -4 \times 2, & -5 \times 1, & \end{array}$$

the search checks all sequences obtained by concatenating permutations of the following multiset of subsequences:

$$\begin{array}{ccc} (0) \times 5, & (-1, 1) \times 5, & (2, -2) \times 4, \\ (-3, 3) \times 3, & (4, -4) \times 2, & (-5, 5) \times 1. \end{array}$$

This drastically reduces the number of sequences that have to be checked.

5 An example with elevenfold symmetry

This section is devoted to a single example of a substitution rule that yields tilings with elevenfold rotational symmetry. To my knowledge, it is the first such example published.

The edge sequence is

$$\sigma = (-1, 1, -3, 3, 0, 2, -2, -1, 1, 0, -5, 5, -3, 3, -1, 1, 4, -4, \\ 2, -2, 0, -1, 1, 2, -2, -3, 3, 0, 4, -4, -1, 1, 2, -2, 0).$$

The associated inflation factor $\lambda = \|\sum \sigma\| \approx 27.2004$ has minimal polynomial $x^5 - 28x^4 + 21x^3 + 21x^2 - 17x + 1$ and is a Pisot number, which is a necessary condition for the existence of any non-trivial eigenvalue of the tiling dynamical system, that is, of any non-trivial discrete part in the spectrum (see [1, Chapters 6 and 7] and [20]).

The substitution rule appears in Figure 9, and a patch appears in Figure 10. This patch is generic and does not depict a centre of global elevenfold symmetry, but it does contain eleven-pointed stars. Then, as described in Section 1, by starting with an eleven-pointed star and repeatedly applying the substitution until two patches are found with one contained in the other, we can produce a tiling with global rotational symmetry, like the one depicted in Figure 11.




Figure 9: Substituted images of $R(11)$. From top to bottom: $\varphi(r(10))$, $\varphi(r(8))$, $\varphi(r(6))$, $\varphi(r(4))$, and $\varphi(r(2))$.



Figure 10: A patch generated using the substitution rule in Figure 9.




Figure 11: A centrally-symmetric patch generated using the substitution rule in Figure 9.

6 The case when n is even

The same methods can be applied to even n , but more care is necessary. Rule 2.6 no longer applies as it did for odd n ; indeed, if n is even then a line segment parallel to e_0 is also parallel to e_n , and these two vectors would give it two opposite orientations. This means it is not so easy to determine the isometries that carry tiles in a given patch onto the prototiles.

Therefore even after finding a tiling of the regions inside the σ -boundaries of the rhombs $R(n)$, one must still do more work to determine how the prototiles have been placed to obtain those tiles. I know of no way to automate this choice, but there are various ad hoc methods that have been successful [3, 10]. One such method is to use prototiles in which opposite edges do not have the same orientation. This is what was done in [10] and in the version of the Ammann–Beenker tilings that appears in the Tilings Encyclopedia [6].

7 Acknowledgements

I would like to thank the anonymous reviewers for their helpful comments, which significantly improved the content and presentation of this paper.

References

- [1] M. Baake and U. Grimm. *Aperiodic order. Vol. 1*, volume 149 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 2013. A mathematical invitation, With a foreword by Roger Penrose.
- [2] M. Baake, P. Kramer, M. Schlottmann, and D. Zeidler. Planar patterns with fivefold symmetry as sections of periodic structures in 4-space. *Internat. J. Modern Phys. B*, 4(15-16):2217–2268, 1990.
- [3] F. P. M. Beenker. Algebraic theory of non-periodic tilings of the plane by two simple building blocks: a square and a rhombus. TH-Report 82-WSK04, Eindhoven University of Technology, the Netherlands, 1982.
- [4] D. Frettlöh. Inflationäre Pflasterungen der Ebene mit D_{2m+1} -Symmetrie und minimaler Musterfamilie. Diploma Thesis, Universität Dortmund, 1998.
- [5] D. Frettlöh and E. Harriss. Parallelogram tilings, worms, and finite orientations. *Discrete Comput. Geom.*, 49(3):531–539, 2013.
- [6] D. Frettlöh and E. O. Harriss. The tilings encyclopedia. Accessed: 2014-08-29.
- [7] F. Gähler, E. E. Kwan, and G. R. Maloney. A computer search for planar substitution tilings with n -fold rotational symmetry. *Discrete Comput. Geom.*, 53(2):445–465, 2015.
- [8] J. García Escudero. Random tilings of spherical 3-manifolds. *J. Geom. Phys.*, 58(11):1451–1464, 2008.
- [9] J. García Escudero. Randomness and topological invariants in pentagonal tiling spaces. *Discrete Dyn. Nat. Soc.*, pages Art. ID 946913, 23, 2011.

- [10] E. O. Harriss. Non-periodic rhomb substitution tilings that admit order n rotational symmetry. *Discrete Comput. Geom.*, 34(3):523–536, 2005.
- [11] S. Kannan and D. Soroker. Tiling polygons with parallelograms. *Discrete Comput. Geom.*, 7(2):175–188, Mar. 1992.
- [12] R. Kenyon. Tiling a polygon with parallelograms. *Algorithmica*, 9(4):382–397, 1993.
- [13] F. Lançon and L. Billard. Two-dimensional system with a quasicrystalline ground state. *J. Physique*, 49(2):249–256, 1988.
- [14] J. Mikhael, M. Schmiedeberg, S. Rausch, J. Roth, H. Stark, and C. Bechinger. Proliferation of anomalous symmetries in colloidal monolayers subjected to quasiperiodic light fields. *Proceedings of the National Academy of Sciences*, 107(16):7214–7218, 2010.
- [15] K.-P. Nischke and L. Danzer. A construction of inflation rules based on n -fold symmetry. *Discrete Comput. Geom.*, 15(2):221–236, 1996.
- [16] R. Penrose. The rôle of aesthetics in pure and applied mathematical research. *The Institute of Mathematics and its Applications Bulletin*, 10(7/8):266–271, July 1974.
- [17] M. Rissanen. Sub rosa, a system of non-periodic rhombic substitution tilings. Unpublished, 2013.
- [18] F. Ruskey and A. Williams. The coolest way to generate combinations. *Discrete Math.*, 309(17):5305–5320, 2009.
- [19] D. Shechtman, I. Blech, D. Gratias, and J. W. Cahn. Metallic phase with long-range orientational order and no translational symmetry. *Phys. Rev. Lett.*, 53:1951–1953, Nov 1984.
- [20] B. Solomyak. Dynamics of self-similar tilings. *Ergodic Theory Dynam. Systems*, 17(3):695–738, 1997.
- [21] Y. Watanabe, T. Soma, and M. Ito. A new quasiperiodic tiling with dodecagonal symmetry. *Acta Crystallographica Section A*, 51(6):936–942, Nov 1995.
- [22] E. J. W. Whittaker and R. M. Whittaker. Some generalized Penrose patterns from projections of n -dimensional lattices. *Acta Cryst. Sect. A*, 44(2):105–112, 1988.
- [23] A. Williams. Loopless generation of multiset permutations using a constant number of variables by prefix shifts. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '09, pages 987–996, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics.

