

A Computer Search for Planar Substitution Tilings with n -Fold Rotational Symmetry

Franz Gähler · Eugene E. Kwan ·
Gregory R. Maloney

Received: 21 April 2014 / Revised: 20 November 2014 / Accepted: 18 December 2014 /
Published online: 6 January 2015
© Springer Science+Business Media New York 2015

Abstract We describe a computer algorithm that searches for substitution rules on a set of triangles, the angles of which are all integer multiples of π/n . We find new substitution rules admitting 7-fold rotational symmetry at many different inflation factors.

Keywords Polygon · Substitution · Tiling · Computer

Mathematics Subject Classification 37B10 · 55N05 · 54H20 · 37B50 · 52C23

1 Introduction

1.1 Motivation

Planar tilings are of interest for their own inherent structure, for their roles as models of aperiodic solids called quasicrystals, and for aesthetic reasons. Among all planar tilings, those that exhibit global n -fold rotational symmetry are of particular interest.

F. Gähler
Fakultät für Mathematik, Universität Bielefeld, Universitätsstraße 25, D-33615 Bielefeld, Germany
e-mail: gaehler@math.uni-bielefeld.de

E. E. Kwan
Department of Chemistry and Chemical Biology, Harvard University,
12 Oxford Street, Cambridge, MA 02138, USA
e-mail: ekwan@fas.harvard.edu

G. R. Maloney
School of Mathematics & Statistics, Newcastle University, Newcastle upon Tyne NE1 7RU, UK
e-mail: gregory.maloney@ncl.ac.uk

This is because n -fold rotational symmetry is evident to the naked eye, it is an invariant of the tiling and the underlying tiling space (see [22] for a definition of a tiling space), and it can distinguish non-periodic tilings from periodic ones, as periodic planar tilings can only exhibit 2-, 3-, 4-, or 6-fold global rotational symmetry.

In [19] a method is described for constructing substitutions on the set of all triangles with angles that are integer multiples of π/n , subject to an appropriate normalization. This method gives rise to an infinite family of planar substitution tilings. (The terms “substitution” and “tiling” will be defined formally in Sect. 1.3.) The underlying tiling spaces of these substitutions are n -fold symmetric, but the tilings themselves do not appear to exhibit global n -fold rotational symmetry except in the special case $n = 9$.

Still there is another substitution described in [19, Fig. 12] that is defined on a proper subset of the triangles with angles that are integer multiples of $\pi/7$. This substitution does not arise from the general construction; indeed, the method of its discovery is not explained, yet it appears to give rise to a tiling with global 7-fold rotational symmetry. This turns out upon closer inspection to be false, as certain isosceles triangles appear in reflected positions, breaking the symmetry. The authors of [19] also observe that this last substitution is special in that it admits a local matching rule (see [9, 13]) whereas, in all of the cases that they checked, the substitutions arising from their general method do not.

The goal of this work is to search for other substitutions that are similar to this extra substitution in that they are defined on a proper subset of the triangles with angles that are integer multiples of π/n . In particular the intention is that, by selecting a minimal subset of these triangles, the substitutions found will produce at least one tiling possessing global n -fold rotational symmetry. Such substitutions will necessarily not arise from the general construction in [19], because the general construction uses all triangles with angles that are integer multiples of π/n , at least for n not divisible by 3.

It is particularly desirable to find multiple different substitution rules on the same set of prototiles that are compatible with one another, meaning that they can be combined to produce edge-to-edge tilings. There has been much recent work on different tiling spaces that arise from the combination of two or more substitution rules on the same prototile set. (Some terms mentioned here, such as “prototile,” will be defined formally in Sect. 1.3.) These fall into two classes: the multi-substitution tilings (see [4–6, 10, 11, 20]), that are obtained by choosing a substitution for each hierarchical level and applying it to all tiles at that level; and the random substitution tilings (see [2, 12, 17, 18]), that are obtained by making separate choices of substitution for each tile at each hierarchical level. While it is easy to find examples of such families of substitutions in one dimension, in two dimensions it is harder. Most known examples—with the noteworthy exception of [19]—are either constant length substitutions [7] or lack the edge-to-edge property [12].

The edge-to-edge property is obviously desirable from the point of view of modeling quasicrystals. But constant length substitutions have integer inflation factors, and so exhibit behavior markedly different from that of substitutions with non-integer inflation factors (see [3, 16]). Therefore there is a need for a collection of examples in two dimensions that are edge-to-edge and have non-integer inflation factors. This project addresses that need.

1.2 Background

Other projects have been undertaken with the goal of finding substitution rules admitting tilings with global n -fold symmetry, but with different approaches and different constraints. In [8], a family of substitution rules is introduced, one for each $n > 7$, that generalises the extra substitution in [19, Fig. 12]. This involves amalgamating adjacent triangles into quadrilaterals and pentagons to bypass a negative area obstruction. In [15], a family of substitutions on rhombic tiles is introduced, generalising a rule of Goodman-Strauss to orders of symmetry greater than 7. In the notation of Sect. 2.1, the inflation factor in [8] is $1 + a_2$ and the inflation factor in [15] is $2 + a_2$, where $a_2 = 2 \cos(\pi/n)$. The former is a unit in its ring of integers, but not a Pisot–Vijayaraghavan (PV) number, except in a few cases. The latter is neither a PV number nor a unit in its ring of integers. Representative pictures of both families of rules can be found in [9] under the names “cyclotomic trapezoids” and “Harriss’s 9-fold rhomb” respectively.

Both of these works successfully adapt a substitution rule that was originally defined for $n = 7$ to arbitrary n , therefore producing an infinite family of substitutions. In the process, both of them introduce extra prototiles, so that they no longer work with a minimal set. Also, the tilings that result from the substitution rules they describe do not exhibit local n -fold symmetry in the k -fold substituted image of any prototile, and hence the associated tiling spaces do not contain any tiling with global n -fold symmetry.

The approach in this work is exploratory rather than constructive. In [8, 15] the approach was to generalise to higher orders of symmetry n a substitution rule that has already been found at a low order of symmetry, thereby producing an infinite family of substitution rules, parametrized by n . Here we make no attempt to build on rules that already exist, but instead search exhaustively for all rules that fit a certain description. In this way we find many rules that do not appear in [8] or [15] because they do not fall into those families. The drawback is that we can only ever hope to find finitely many substitution rules by this method—although, as we shall see, this finite list is quite long, and can easily be made longer.

Since we impose no preconditions on the inflation factors that we search, accordingly we must search a large parameter space. To do this it is necessary to use a computer.

1.3 Definitions

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 only polygonal (in fact, triangular) tiles having some finite set S of representatives up to isometry. The elements of S are called *prototiles*, and we denote by $\mathcal{P}(S)$ the set of all patches consisting of tiles that are congruent to these prototiles. Let us also suppose that we

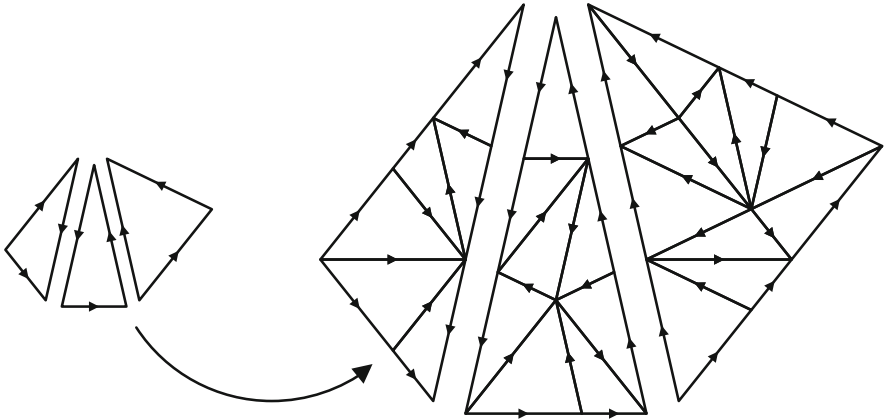


Fig. 1 A substitution on three prototiles, with edge orientations

have a *substitution*, that is, a map $\varphi : S \rightarrow \mathcal{P}(S)$ for which there is an *inflation factor* $\lambda > 1$ such that, for any $p \in S$, the support of $\varphi(p)$ is λp . Letting $S = \{p_1, \dots, p_k\}$, we define the *substitution matrix* of φ to be the $k \times k$ integer matrix, the entry of which at position (i, j) is the number of tiles isometric to p_i in $\varphi(p_j)$.

Figure 1 depicts a substitution on three prototiles, with the arrows on the edges of the tiles describing *edge orientations*, which are defined in Sect. 1.4 below. This substitution has the same prototile shapes and substitution matrix—that is,

$$\begin{bmatrix} 3 & 3 & 5 \\ 1 & 4 & 3 \\ 2 & 1 & 3 \end{bmatrix}$$

—as the substitution in [19, Fig. 12], although it is indeed a different substitution. In particular, the prototiles, although congruent to those of [19, Fig. 12], have different edge orientations, and so should be seen as different prototiles.

We can extend the definition of φ to all tiles isometric to the elements of S in the following way. If p is a prototile, τ is an orthogonal transformation of \mathbb{R}^2 , and $v \in \mathbb{R}^2$ is a translation vector, then $\varphi(\tau(p) + v) = \{\tau(t) + \lambda v \mid t \in \varphi(p)\}$. Once we have done this, we can extend the definition of φ further to include all patches in $\mathcal{P}(S)$ by declaring $\varphi(P) = \{\varphi(t) \mid t \in P\}$. The chief motivation for this is that, under certain easily-satisfied conditions, we can produce a tiling by starting with some prototile p and repeatedly applying φ [14]. Let us refer to such tilings as *substitution tilings*. Figure 2 depicts a substitution tiling constructed in this way from the substitution in Fig. 1. It clearly possesses local 7-fold rotational symmetry.

1.4 The Objective

Our purpose here is to find substitutions that obey a certain rule and use certain prototile sets. The rule is that the resulting substitution tilings must be *edge-to-edge* [14], which

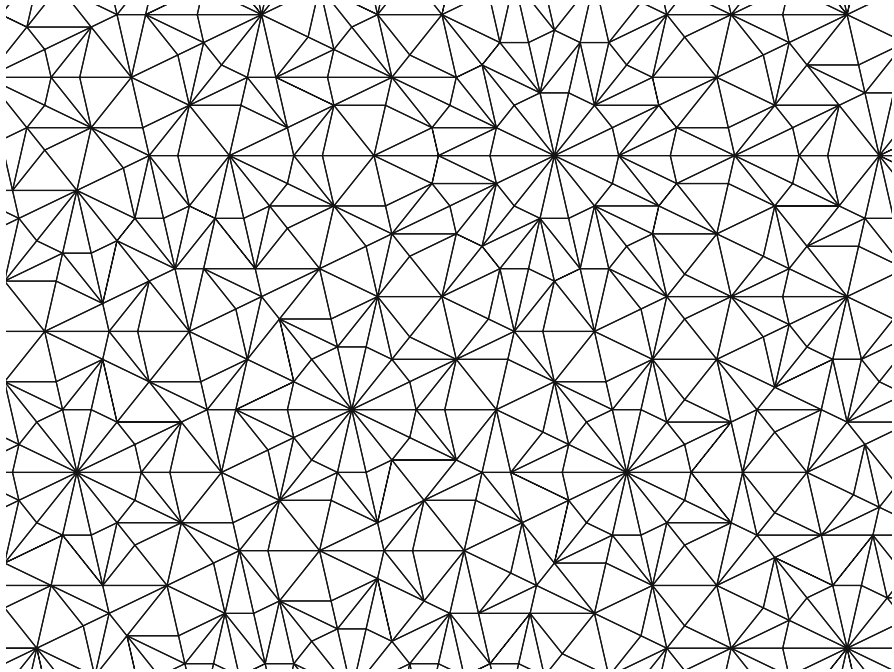


Fig. 2 A substitution tiling constructed from the substitution in Fig. 1

means that, if two tiles in a tiling intersect, their intersection is a face of both tiles. For polygonal tiles in two dimensions, this amounts to the property that, if the vertex of one tile touches another tile, then it is also a vertex of the other tile.

The edge-to-edge condition for substitution tilings arising from φ is equivalent to the condition that the patches $\varphi^m(p)$ be edge-to-edge for all $m \in \mathbb{N}$ and all prototiles p . Verifying this condition seems at first to require checking infinitely many patches, but in fact it is sufficient to check for all prototiles p that the patches $\varphi(p)$ are edge-to-edge and satisfy one additional condition, which is described below.

Given a tile t with an edge e having end points v_1 and v_2 , an *edge orientation* is a map f_t that assigns to e one of its vertices v_i . We can represent f_t graphically by drawing an arrow on e originating at $f_t(e)$ and terminating at the other vertex, as has been done in Fig. 1. Then we require that all prototiles p and tiles t have edge orientations on all of their edges satisfying the following conditions.

1. (Isometry equivariance) If $t = \tau(p) + v$, then, for any edge e of p , the corresponding edge in t has the same edge orientation; that is, $f_t(\tau(e) + v) = \tau(f_p(e)) + v$.
2. (Matching) If an edge e lies in two tiles t_1 and t_2 , then it receives the same edge orientation from both of them; that is, $f_{t_1}(e) = f_{t_2}(e)$.
3. (Preservation under φ) If two prototiles p and p' contain edges e and e' respectively such that $e = \tau(e') + v$ and $f_p(e) = \tau(f_{p'}(e')) + v$, then e and e' must have the same *edge breakdowns*, which means roughly that their inflated images must contain the same edges in the same order with the same orientations. More specifically,

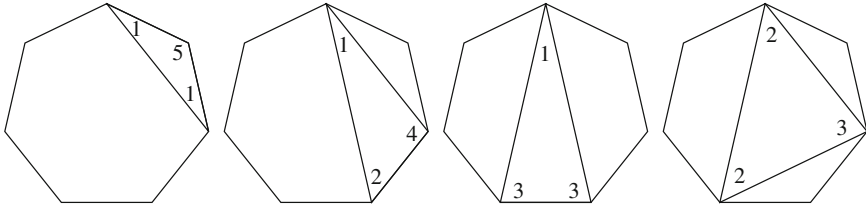


Fig. 3 The elements of $S(7)$ inscribed in regular heptagons

let $e_1 \subset t_1, \dots, e_k \subset t_k$ denote the edges in $\varphi(p)$ that are contained in λe , and let $e'_1 \subset t'_1, \dots, e'_m \subset t'_m$ denote the edges in $\varphi(p')$ that are contained in $\lambda e'$. Then $k = m$, $e_i = \tau(e'_i) + v$ for all $1 \leq i \leq k$, and $f_{t_i}(e_i) = \tau(f_{t'_i}(e'_i)) + v$ for all $1 \leq i \leq k$.

Now let us describe the special prototile sets that we will use. For a natural number $n \geq 3$, let $S(n)$ denote the set of isometry classes of triangles, the angles of which are integer multiples of π/n , normalized so that they can all be inscribed in circles of the same size. $S(7)$ is depicted in Fig. 3.

The goal in this work is to find substitutions that use proper subsets of $S(n)$ as prototiles. The substitution depicted in Fig. 1 uses three of the four prototiles from $S(7)$.

Note that the general method for constructing substitutions that is described in [19] uses a bigger set of prototiles. In particular, that prototile set contains two copies of each scalene triangle in $S(n)$. These triangles have opposite edge orientations and therefore should be considered as different from one another. Let us not follow this convention here; for us, any two isometric tiles will always have the same edge orientations, and hence will be copies of the same prototile.

2 Lengths and Areas

We will need to know the areas of the triangles in $S(n)$, along with their various edge lengths. Denote by $T(k_1, k_2, k_3)$ the triangle with angles $k_1\pi/n, k_2\pi/n$, and $k_3\pi/n$.

2.1 Lengths

Since we have normalized the triangles in $S(n)$ so that they can all be inscribed in the same circle, in particular we know that their side lengths coincide with the lengths of the diagonals of a regular n -gon, as is depicted for $n = 7$ in Fig. 3.

Let the edge lengths of this regular n -gon be $a_1 = 1$. Then the first diagonal has length $a_2 := \sin((n - 2)\pi/n) / \sin(\pi/n) = 2 \cos(\pi/n)$. This is the length of the long edge of the triangle $T(1, 1, n - 2)$.

We can use similar triangles to produce a recursion relation for the length a_k of the k th diagonal of the n -gon. In particular, the triangles $T(1, k, n - k - 1)$ and $T(1, k - 1, n - k)$ can be placed against one another along their short edges to produce a triangle similar to $T(1, 1, n - 2)$, as has been done in Fig. 4. The edges of this triangle

are obtained from the edges of $T(1, 1, n - 2)$ by scaling by a factor of a_k . In particular, using the fact that the long edge of this new triangle has length $a_{k+1} + a_{k-1}$, we see that

$$a_{k-1} + a_{k+1} = a_2 a_k, \tag{1}$$

$$a_{k+1} = a_2 a_k - a_{k-1}. \tag{2}$$

Taking $a_0 = 0$, this recursion holds for all $k \geq 1$. If we view a_2 as a variable and a_k as a polynomial in a_2 , then a_k is the k th Chebyshev polynomial of type 2, subject to the reparametrization $a_2 = 2x$ [21]. If n is odd, then diagonal number $(n - 1)/2$ has the same length as diagonal number $(n - 1)/2 - 1$, so setting

$$a_{(n-1)/2} = a_{(n-1)/2-1} \tag{3}$$

yields a polynomial equation of degree $(n - 1)/2$ that a_2 satisfies.

If n is even, then the equation is

$$a_{n/2} = a_{n/2-1}.$$

Note that, by the symmetry of the regular n -gon, $a_k = a_{n-k}$.

Now let A denote the companion matrix of the minimal polynomial q_n of a_2 . If $b \in \mathbb{Q}(a_2)$, let $p_b \in \mathbb{Q}[x]$ denote the monic polynomial for which $b = p_b(a_2)$, and let $v(b) = (v_0(b), \dots, v_{\Phi(n)/2}(b))$ denote the vector of coefficients of this polynomial; i.e., the vector representation of b with respect to the basis $1, a_2, \dots, a_2^{\Phi(n)/2}$ of $\mathbb{Q}(a_2)$. Note that by (1) p_{a_i} has integer coefficients. Then $v(a_2 b)^t = A v(b)^t$ and, given an inflation factor $\lambda \in \mathbb{Q}(a_2)$, $v(\lambda b)^t = p_\lambda(A) v(b)^t$. This provides a means of representing the lengths λa_i as combinations of $a_0, \dots, a_{(n-1)/2}$. In particular, let $L_n := [v(a_0)|v(a_2)|\dots|v(a_{(n-1)/2})]$ denote the $\Phi(n)/2 \times (n - 1)/2$ matrix, the columns of which are $v(a_i)$; then the i th column of the solution X of the matrix equation

$$L_n X = p_\lambda(A) L_n \tag{4}$$

expresses λa_i as a combination of $a_0, \dots, a_{(n-1)/2}$. If n is prime, then L_n is square—the number $(n - 1)/2$ of edge lengths agrees with the degree of q_n —and upper triangular—eq. (1) expresses a_i as a polynomial of degree $i - 1$ in a_2 —and hence invertible. Then we can write

$$X = L_n^{-1} p_\lambda(A) L_n. \tag{5}$$

Let us use only inflation factors λ that are positive integer combinations of $a_0, \dots, a_{(n-1)/2}$. The reason for doing so is that, if $t \in S(n)$ contains an angle of measure π/n , then its shortest edge has length $a_1 = 1$, so the shortest edge of λt will have length λ , which must therefore be a sum of prototile edge lengths. This is not strictly necessary if we choose for our set of prototiles a subset of $S(n)$ that contains

Fig. 4 Two narrow triangles

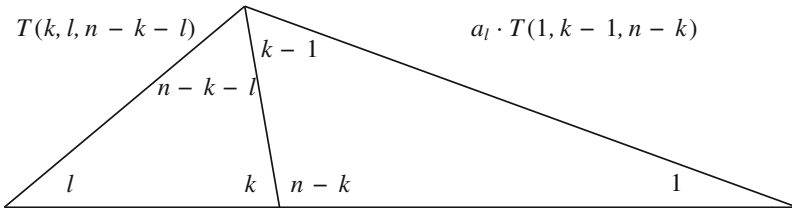
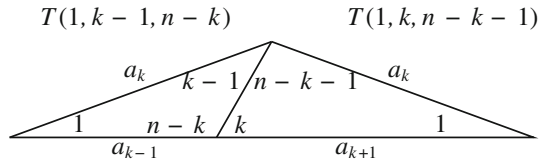


Fig. 5 A non-narrow triangle and an inflated narrow triangle

no triangle with an edge of length 1, but empirically such prototile sets do not work well, and we will soon focus our attention on a special set of prototiles that contains two triangles with minimal edge lengths—see Sect. 3.2.

2.2 Areas

Now we can calculate the areas of the triangles in $S(n)$. Given a triangle t , let $|t|$ denote its area. Then we can express the areas of all the triangles as elements of $\mathbb{Q}(a_2) \cdot |T(1, 1, n - 2)|$, where $a_2 = 2 \cos(\pi/n)$ as described in Sect. 2.1. Let us call a triangle in $S(n)$ a *narrow triangle* if it has an angle of π/n . Then we first calculate the areas of the narrow triangles recursively using the same similar triangles that we used to express a_k in terms of a_2 .

In particular, $T(1, k, n - k - 1)$ and $T(1, k - 1, n - k)$ fit together to form a triangle similar to $T(1, 1, n - 2)$, but inflated by a factor of a_k (see Fig. 4). Therefore

$$|T(1, k, n - k - 1)| = a_k^2 \cdot |T(1, 1, n - 2)| - |T(1, k - 1, n - k)|, \tag{6}$$

so, recursively, we obtain formulas expressing the areas of the narrow triangles in terms of $|T(1, 1, n - 2)|$. Note that each a_k can be written as an integer polynomial in a_2 , so this recursion formula expresses $|T(1, k, n - k - 1)|$ as an integer polynomial in a_2 .

To calculate the areas of the non-narrow triangles, note that every non-narrow triangle fits together with an inflated narrow triangle to form another inflated narrow triangle, as in Fig. 5. In particular, if $k < l < n - k - l$, then

$$|T(k, l, n - k - l)| = a_k^2 \cdot |T(1, l, n - l - 1)| - a_l^2 \cdot |T(1, k - 1, n - k)|, \tag{7}$$

and, since $|T(1, l, n - l - 1)|$ and $|T(1, k - 1, n - k)|$ are multiples of $|T(1, 1, n - 2)|$ by integer polynomials in a_2 , so is $|T(k, l, n - k - l)|$. Note that permuting the order

of the angles k, l , and $n - k - l$ gives us up to three different formulas for this area, any one of which will work.

The main reason for computing triangle areas is to determine the substitution matrix of the substitution φ that we are trying to find. Choose a set of prototiles $S \subset S(n)$, the areas of which form a \mathbb{Q} -basis for $\mathbb{Q}(a_2) \cdot |T(1, 1, n - 2)|$. Enumerate the elements of this set t_1, \dots, t_m , and let $A_k := |t_k|/|T(1, 1, n - 2)|$. Let B_S denote the square matrix, the columns of which are the vectors $v(A_k)$ (defined in Sect. 2.1); that is, $B_S = [v(A_1) | \dots | v(A_m)]$. Our choice of S having areas that are a \mathbb{Q} -basis for $\mathbb{Q}(a_2) \cdot |T(1, 1, n - 2)|$ means that B_S is invertible.

Then, if there exists a substitution φ on S with inflation factor λ , it will have substitution matrix

$$M := B_S^{-1} p_\lambda(A)^2 B_S. \tag{8}$$

Sometimes an inspection of the matrix M is enough to prove the non-existence of a substitution on S with factor λ . For instance, if any of the entries of M are not integers or are negative, then no such rule can exist—although [8] describes a way of modifying S to address the problem of a negative entry.

If the entries of M are all positive integers, then we search the combination of S and λ for substitution rules.

3 The Program

Let us now describe the program used to search for substitution rules. The program is written in Java, although the discussion presented here contains no Java-specific details. We chose Java because of its native support for concurrency (see Sect. 3.5) and cross-platform compatibility.

3.1 Overview

Before describing the search algorithm in any detail, let us give an overview of how the program works.

The order of symmetry n is assumed to be fixed from the start. The program takes three ingredients as input: a set $S \subset S(n)$ of prototiles, an inflation factor λ that is a non-negative integer combination of $a_1, a_2, \dots, a_{(n-1)/2}$, and an inflated prototile λt_0 , where $t_0 \in S$. Then it tries to fill λt_0 with tiles congruent to prototiles from S in such a way that the tiles overlap at most in their boundaries, they meet edge-to-edge if at all, and their union is all of λt_0 . Let us refer to each such application of the program as a *search* of the triple (S, λ, t_0) . The output of a search is a (possibly empty) set of patches from $\mathcal{P}(S(n))$, each of which has support λt_0 . Let us call such a patch a *result*. A search will be considered successful if it returns at least one result.

In order to obtain a substitution rule on S with factor λ , we must run a search on (S, λ, t) for each prototile $t \in S$, and the search must be successful for each one. Then we can define a substitution rule φ as follows: for each $t \in S$, let $\varphi(t)$ be any one of the patches found in the search on (S, λ, t) . In order for the substitution rule so obtained

to be edge-to-edge, it must satisfy some additional conditions on its edge orientations and edge breakdowns, as described in Sect. 1.4. Isometry equivariance (condition (1)) and matching (condition (2)) are built into the search—so, in fact, the program will not find any patches that do not satisfy these conditions, although that restriction can be turned off.

Preservation under φ (condition (3)) is checked after the completion of all the searches for a given S and λ . More specifically, we assemble all of the results together and look for a tuple of results, one for each tile, that can be used to define a substitution that satisfies condition (3). We also check now that the tuple of results satisfies condition (1) as a whole, because the program only checks the patches individually during the search.

If such a tuple of results exists, then the substitution that it defines is edge-to-edge, and we have found what we set out to find for S and λ .

3.2 A Special Prototile Set

Of all subsets $S \subset S(n)$ for odd n , one works better than the others. This is the set consisting of all triangles that contain at least one edge of maximum length, that is, all triangles with at least one angle of measure $n/2 \pm 1/2$. There are $(n - 1)/2$ such triangles, and they can be combined along their short sides to produce triangles similar to the $(n - 1)/2$ isosceles triangles in $S(n)$.

3.3 Representation of Points

If a patch in $\mathcal{P}(S(n))$ contains a tile with a vertex at the origin, then all tile vertices in that patch must lie in the \mathbb{Z} -module \mathcal{V} generated by all of the vectors

$$\{a_k(\cos(i\pi/n), \sin(i\pi/n)) \mid 1 \leq k \leq (n - 1)/2, 0 \leq i < 2n\}. \tag{9}$$

If each $t \in S$ has a vertex at the origin, then so does each λt , and so we can restrict our attention to such patches. Therefore the program need not deal with arbitrary points in \mathbb{R}^2 , but only points in \mathcal{V} .

Internally, it represents these points as elements of \mathbb{Z}^{n-1} .

If n is prime, then \mathbb{Z}^{n-1} is isomorphic to \mathcal{V} . Specifically, multiplying on the left by the matrix

$$\begin{bmatrix} \cos \frac{0\pi}{n} & \cos \frac{1\pi}{n} & \cdots & \cos \frac{(n-2)\pi}{n} \\ \sin \frac{0\pi}{n} & \sin \frac{1\pi}{n} & \cdots & \sin \frac{(n-2)\pi}{n} \end{bmatrix} \tag{10}$$

takes the i th standard basis vector of \mathbb{Z}^{n-1} to the vector $(\cos(i\pi/n), \sin(i\pi/n)) \in \mathcal{V}$, which is an element of the set \mathcal{V} as $a_1 = 1$. To see that this map is onto, note that the images of the standard basis vectors are the direction vectors of $n - 1$ of the sides of a regular n -gon. This is depicted for the case $n = 7$ in Fig. 6. Then it is not hard to see that vectors with all different lengths $a_1, \dots, a_{(n-1)/2}$ can be obtained as sums of the images of the standard basis vectors—for instance,

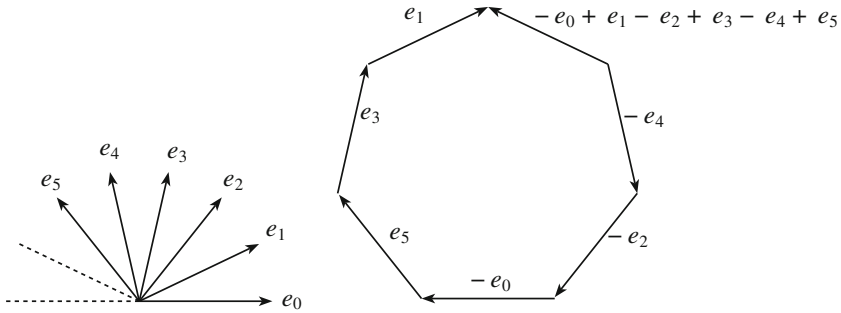


Fig. 6 The images of the standard basis vectors of \mathbb{Z}^6 , arranged to form the edges of a regular heptagon

$$a_2(\cos(1\pi/n), \sin(1\pi/n)) = (\cos(0\pi/n), \sin(0\pi/n)) + (\cos(2\pi/n), \sin(2\pi/n)).$$

If n is not prime then this map $\mathbb{Z}^{n-1} \rightarrow \mathcal{V}$ is onto but not one-to-one, as there are multiple non-trivial relations between the direction vectors of the edges of a regular n -gon.

3.4 The Search

The search works by recursion. It has two main objects that it updates with each recursion step, and it uses these objects to decide whether or not to terminate the recursion. The first object is a patch $P \in \mathcal{P}(S)$, the support of which is contained in λt_0 . The second is a list of non-negative integer multiplicities of prototiles. These multiplicities indicate how many prototiles of each type remain to be placed in order to obtain a patch with support equal to λt_0 ; it is initialized using the column of the matrix M from (8) that corresponds to λt_0 . We say that this list contains a tile t if the multiplicity of that tile is greater than 0. At each step in the recursion, the multiplicity of one prototile is decremented in this list and the corresponding tile is pushed on P , which is in fact a stack. Then the recursion is finished when all of the multiplicities are 0.

The main recursive procedure is called solve. Simplified pseudocode for this procedure appears in Algorithm 1.

Several of the statements in Algorithm 1 require further explanation. The place procedure on line 11 takes as input the prototile t and returns a tile t' that is congruent to t . Then the compatible procedure on line 12 returns true if $P \cup \{t'\}$ is a patch with support contained in λt_0 . (In fact, compatible also requires $P \cup \{t'\}$ to satisfy edge orientation conditions (1) and (2)—see Sects. 1.4 and 3.1.)

There are many conceivable ways of placing a copy of the prototile t in P , but we only consider a restricted range of possibilities. The edges of all of the tiles in P are divided into two lists, called open edges and closed edges. Closed edges are edges that are contained in two different tiles, or in one tile and in the boundary of λt_0 . Open edges are edges that are contained in only one tile, and not in the boundary of λt_0 .

Algorithm 1 The solve procedure

```

1:  $P$ : a patch with support contained in  $\lambda t_0$ 
2:  $t$ : the prototile we are currently trying to place in  $P$ 
3:  $l$ : a list of multiplicities telling us how many prototiles of each type remain to be placed
4: procedure SOLVE
5:   repeat
6:     if  $l$  is empty then
7:       store a deep copy of  $P$  somewhere ▷  $P$  is a result
8:       break
9:     end if
10:    if  $l$  contains  $t$  then
11:       $t' \leftarrow$  place  $t$ 
12:      if  $t'$  is compatible with  $P$  then
13:        push  $t'$  on  $P$ 
14:         $t \leftarrow$  first prototile
15:        solve ▷ recursion
16:         $t \leftarrow$  pop last tile from  $P$ , get associated prototile
17:      end if
18:    end if
19:     $t \leftarrow$  prototile after  $t$ 
20:  until back to first prototile
21: end procedure

```

In other words, open edges are edges against which another tile must be placed in order to obtain a completed patch.

The open edges and closed edges are stored in stacks, so in particular they have an order, with the open edge(s) coming from the most recently-placed tile being listed last. Then the rule is that we place the prototile t against the last open edge, if this is possible (t might not have an edge of the right length). This rule has two main implications.

The first implication is that any newly-placed tile t' has at least one edge in common with a tile that has already been placed—this edge becomes a closed edge. Therefore a newly-placed tile can contribute at most two new open edges. If it contributes one new open edge, or no new open edges, then there is no ambiguity as to which new open edge is listed last. But if it contributes two new open edges, then we have a systematic way of determining the order in which they are added to the stack; namely, they are pushed in the order in which they appear in a counterclockwise traversal of the edges of t' , starting with the closed edge.

The second implication is that, in order to call the place procedure, we need to have at least one open edge, even at the beginning when no tiles have been placed yet. This is a problem that we fix by placing an open edge, called a starter, on part of one of the edges of λt_0 (see Fig. 7). Then the first tile must be placed against the starter.

We need to choose a length for the starter edge, and some choices of edge length might preclude certain results. We deal with this by running several searches in parallel—one for each choice of starter edge length. This turns out to be a convenient method to divide up the work for multithreading—see Sect. 3.5. To produce a list of possible starter edge lengths, one option is to use all edge lengths $a_1, \dots, a_{(n-1)/2}$. An even better option is to select only those lengths for which the corresponding entry in the length substitution matrix X from (5) (Eq. 4 if n is not prime) is positive.

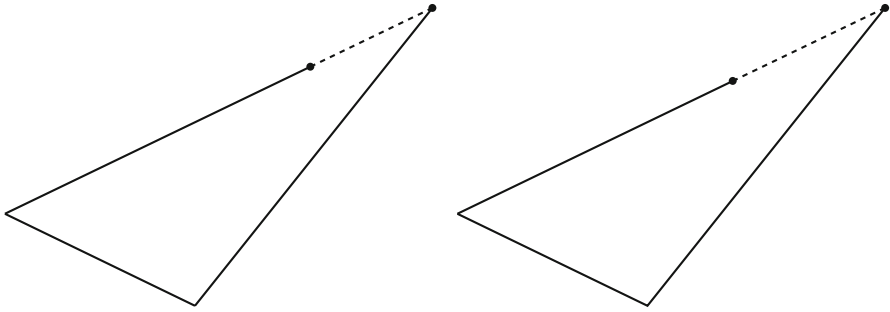


Fig. 7 Two copies of an inflated prototile with starter edges of different lengths shown as *dashed lines*

The pseudocode in Algorithm 1 has been simplified a bit for clarity, but we should note now that there might be several ways to place the prototile t against the last open edge. If t has an edge of the same length as the last open edge, then there are two possibilities: we can place t against the last open edge using a non-orientation preserving isometry—i.e., a reflection—or using an orientation preserving isometry. If t has two edges of that length, then there are four ways of placing it against the last open edge. (The case in which t is equilateral has not come up yet, because the program only works on prime n now, so in particular not on any n divisible by 3.)

To keep track of these possible placements, we store not only the current prototile t , but also two booleans *flip* and *second* that tell us respectively if we are trying to place a reflected copy of t and if we are trying to place t using the second of two edges of the same length. Therefore on line 11, we do not place t , but rather we place the triple $(t, \textit{flip}, \textit{second})$. On line 19, we get the next triple $(t, \textit{flip}, \textit{second})$, not just the next prototile t , and on lines 14 and 20 we refer to the first triple $(t_1, \textit{false}, \textit{false})$, not just the first prototile t_1 . On line 16, we pop the last tile t' from P and, using information from t' and P we infer which triple $(t, \textit{flip}, \textit{second})$ was input to the place procedure to produce t' .

This last point merits further discussion. The purpose of line 16 is to restore the states of $(t, \textit{flip}, \textit{second})$, and P to what they were on line 13. In principle we could store these states in memory, but note that we call the solve procedure recursively on line 15; this means that these states could change many times between lines 13 and 16. In practice it becomes unwieldy to store all of these changes—sometimes it even crashes the program—so we store none of them. Instead, when we remove t' from P on line 16, we use information about t' to restore everything to its state from line 13.

In particular, we remove any open edges coming from edges of t' , and if any closed edges come from edges of t' , then they are made open again. We also determine the prototile t to which t' is congruent. If t' is a reflected copy of t , then we set *flip* to *true*, and if t' has its second of two equal edges against the last open edge, then we set *second* to *true*. These calculations allow us to ascend and descend the recursion tree without having to use up the system memory storing states from previous levels of recursion.

3.5 Multithreading

The amount of time required to search an inflated prototile λt_0 scales up rapidly with its area, and it is worthwhile to take advantage of the multithreading capabilities of the Java language to reduce this time. The particular algorithm design that makes it possible to do this might be of independent interest, so let us describe it here.

The program holds a list of searches at varying stages of completion; at the same time, there are many threads, each working on a different search. When a thread finishes a search, it draws a new one from the list.

These searches are independent from one another, in the sense that no two can ever return the same result. This is because they all contain different configurations of open edges and tiles that have already been placed in their patches. So the initial searches are not actually empty; rather, they contain starter edges (see Sect. 3.4), which differ from one another in their lengths. Thus we can have up to $(n - 1)/2$ initial searches.

We may have many more threads than initial searches, and, empirically, most of the initial searches seem to finish quickly with no results. Therefore we need a way to split up the few searches that remain in order to take advantage of the extra threads available. We do this by giving each search a kill switch that we trigger when the number of searches remaining in the list drops too low.

When the kill switch is triggered, the search algorithm changes: instead of making another solve call at line 15 of Algorithm 1, the search makes a deep copy of itself in its current state, and places this copy in the list of searches. A given search may have called solve recursively many times already when the kill switch is triggered, so it will pass line 15 many times before finishing, adding many searches to the list. These new searches represent work that the original search would have done had the kill switch not been triggered.

Figuratively, the behavior of the searches can be described as follows. Under ordinary conditions, a search climbs the recursion tree from bottom to top and searches all of the branches to their ends as they are encountered. After the kill switch has been triggered, the search no longer climbs up new branches that it encounters; instead, whenever it encounters a new branch on the way down the tree, it cuts it off at the base and throws it on a pile of branches to be searched. Then it reaches the bottom of the recursion tree and finishes.

4 Results

4.1 Fivefold Symmetry

To find substitutions using $S(5)$ it is not necessary to use a computer, but it certainly helps. For these prototiles, the smallest inflation factor is a_2 , the golden ratio. Substitutions with this inflation factor have been described elsewhere (see [11, 19]); these can be combined as multi-substitutions, but not as random substitutions—although see [11] for a description of a slightly different randomization strategy, called tile rearrangement.

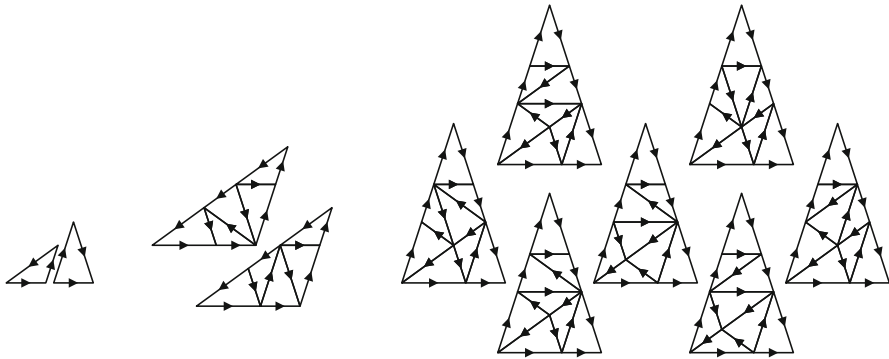


Fig. 8 Results for $S(5)$, $\lambda = 1 + a_2$

Already at the next inflation factor, $a_2^2 = 1 + a_2$, there are many substitutions to be found. Figure 8 depicts a family of these—two for $T(1, 1, 3)$ and seven for $T(1, 2, 2)$ —that can be used to produce random substitutions. The program also finds substitutions with different edge breakdowns that can be combined with these ones as multi-substitutions.

4.2 Sevenfold Symmetry

We have found many new substitution rules with sevenfold rotational symmetry using three different inflation factors: $\lambda_1 = 1 + a_2$, $\lambda_2 = a_2 + a_3$, and $\lambda_3 = 1 + a_2 + a_3$. The second and third of these are both PV numbers, while the first is not. A PV inflation factor is a necessary condition for the existence of any non-trivial eigenvalue of the dynamical system, that is, for any non-trivial discrete part in the spectrum [22]. All three inflation factors are units in the ring of integers of the number field generated by a_2 , which is necessary for the module generated by the eigenvalues to be finitely generated.

Of course, it is possible to search even larger inflation factors, and presumably to find more substitution rules, at the cost of more computing time and much greater memory usage.

The $T(1, 3, 3)$ and $T(2, 2, 3)$ triangles are both isosceles, so each one has two ways of placing it in any given position: an orientation preserving placement and an orientation reversing placement. If we choose one of these isosceles triangles and reflect each of its instances in a result P across its axis of symmetry, then this modified patch is still a result. We have processed the results obtained from the program to remove this redundancy and to select only those groups of results that have compatible edge breakdowns in the sense of Sect. 1.4, and that therefore can be used to produce tilings.

The results for λ_1 appear in Figs. 9 and 10. They fall into two classes, depicted in Figs. 9 and 10, according to the edge orientations of the prototiles, which are enclosed in boxes on the right sides of the figures. These should really be considered as two different sets of prototiles, because substitution rules coming from the two different

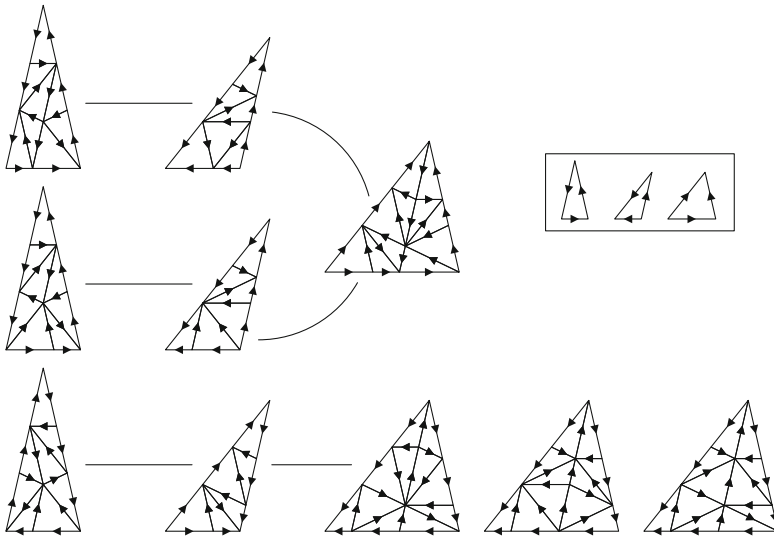


Fig. 9 The standard prototiles. These are compatible with results for λ_2 and λ_3 . Arcs between results indicate compatibility of edge breakdowns. The substitution depicted in Fig. 1 uses the results in the *middle row* of this figure

classes of results cannot be combined with one another even as multi-substitutions. On the other hand, two or more substitution rules coming from the same class of results can be combined as multi-substitutions.

An arc between two results in Fig. 9 or 10 indicates that they have the same edge breakdowns. Two or more substitution rules coming from results with the same edge breakdowns can be combined as random substitutions. So, for instance, we can make three different substitution rules by choosing, from the bottom row of Fig. 9, any one of the three results for $T(2, 2, 3)$, along with the results for $T(1, 3, 3)$ and $T(1, 2, 4)$; these three substitutions can be combined as random substitutions. This involves creating tilings by repeatedly substituting a patch, and for each $T(2, 2, 3)$ tile appearing in that patch, choosing at random which of the three rules to apply to it.

The results for λ_2 appear in Fig. 11. The substitutions built from the results in Figs. 9 and 11 use the same prototiles, and can be combined with one another to produce multi-substitutions. Any substitution built from results in Fig. 11 with the same edge breakdowns and edge orientations can be combined with one another to produce random substitutions. Of course, a substitution built from results in Fig. 11 cannot be combined with a substitution built from results in Fig. 9 or 10 to produce a random substitution because they have different inflation factors.

The results for λ_3 are too numerous to depict here, but a set of three of them, one for each prototile, appears in Fig. 12. These three all have compatible edge breakdowns and edge orientations, in the sense of 1.4. For these edge breakdowns and orientations there are 3 compatible $T(1, 2, 4)$ -results, 5 compatible $T(1, 3, 3)$ -results, and 36 compatible $T(2, 2, 3)$ -results, all of which can be combined with each other as random substitutions.

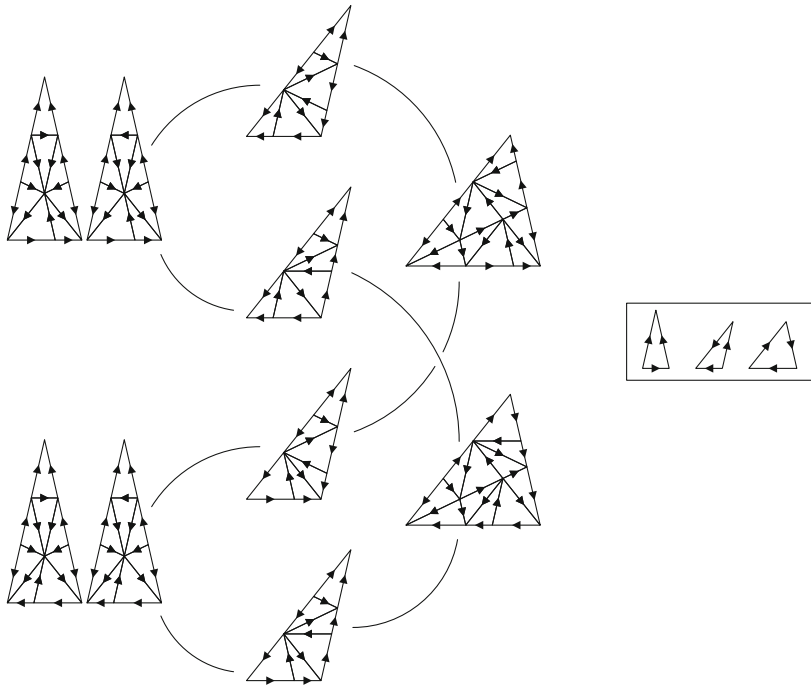


Fig. 10 The non-standard prototiles. The Danzer–Nischke example [19, Fig. 12] uses the result for $T(1, 2, 4)$ depicted at the top. Arcs between results indicate compatibility of edge breakdowns

Moreover, this is one out of ninety edge breakdown/edge orientation combinations; the other eighty-nine also have many results associated to them (some of which are repeated, since, for example, the $T(2, 2, 3)$ result appearing in Fig. 12 is compatible with any $T(1, 2, 4)$ result having the same edge breakdown on its medium and long edges, regardless of the edge breakdown of the short edge). A substitution built using results from any of these ninety edge breakdown/edge orientation combinations can be combined with any other such substitution, or any substitution coming from Fig. 9 or 11.

4.3 Observations

Let us now make a few miscellaneous observations on the results presented in Sects. 4.1 and 4.2.

1. The non-standard prototiles in Fig. 10 appear to be a phenomenon unique to the factor λ_1 , in that there are no results for the factors λ_2 or λ_3 with those prototiles.
2. The prototiles in Figs. 8, 9, 11, and 12 have the property that, if two edges meet at a vertex at an angle that is an even integer multiple of π/n , then they are either both oriented towards that vertex or both oriented away from it; if the angle is an odd integer multiple of π/n , then their orientations are opposite. This has the consequence that, in any tiling made with these prototiles, all edges with the same

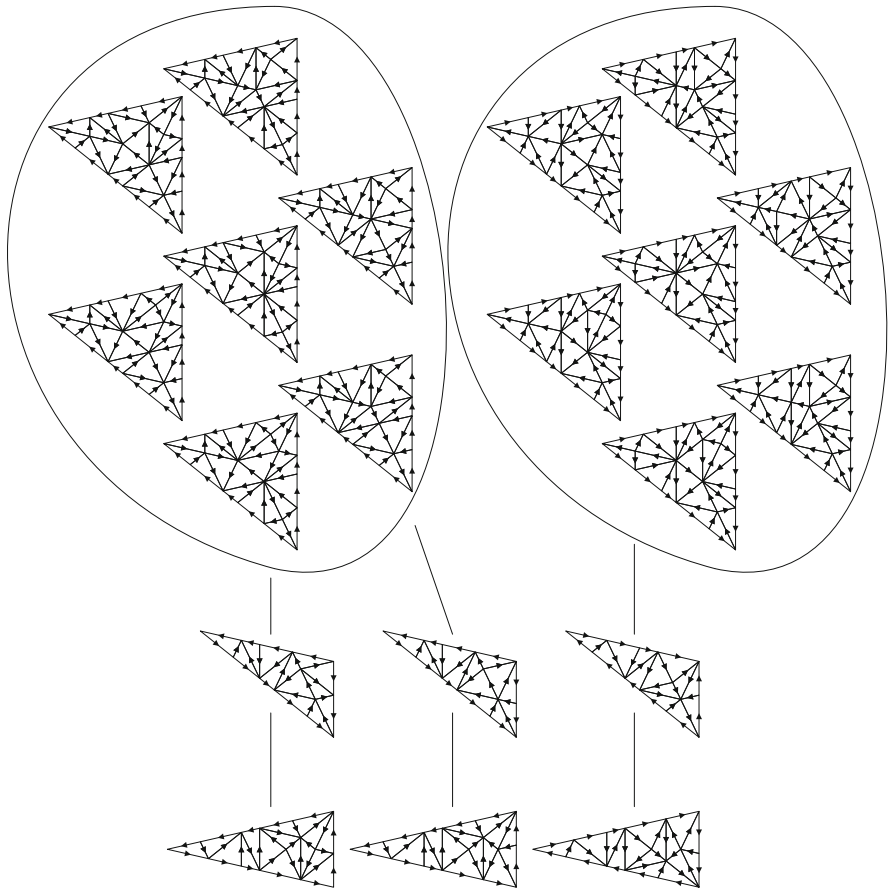


Fig. 11 The results for λ_2 . Arcs between results or groups of results indicate compatibility of edge breakdowns. The two groups of results for $T(2, 2, 3)$ are reflections of one another

angle will have the same orientation. It also means that these tilings are only n -fold symmetric, not $2n$ -fold symmetric, as is often the case when one tries to construct something n -fold symmetric. This is also true of substitutions arising from the general method described in [19], although it is not true of the special substitution in that work, nor of the substitutions in Fig. 10.

3. Some of the results in Figs. 9, 11, and 12 reverse the directions of the tile edges in the super-tile edges, while others preserve them.
4. The program produces many more results than are recorded here. Some of these results are obtained by reflecting all instances of an isosceles tile, as described in Sect. 4.2, and so in that sense are redundant. Other results are only partial in the sense that they cannot be combined with results for the other prototiles to produce a substitution rule. For example, there exist results for $T(2, 2, 3)$ that are compatible with other results for $T(1, 2, 4)$, but not with any result for $T(1, 3, 3)$. In principle such results could still be used in random substitutions, provided that they were

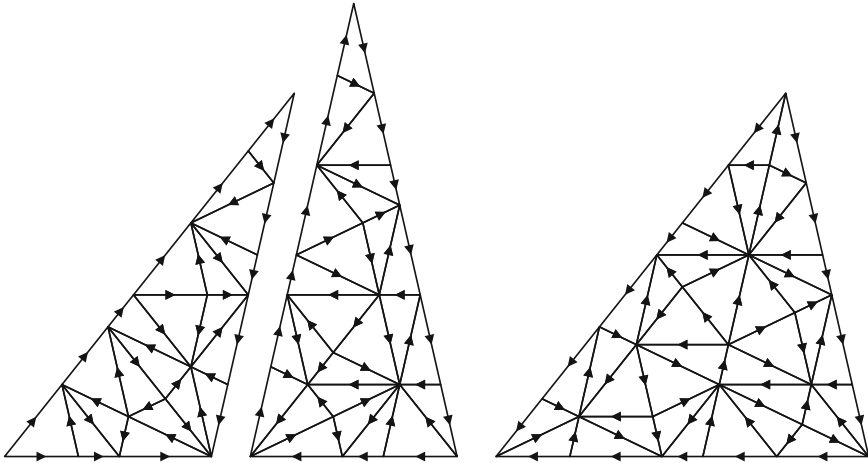


Fig. 12 Results for λ_3

not applied to any instance of $T(2, 2, 3)$ that shared an edge with an instance of $T(1, 3, 3)$.

4.4 A Surprising Topological Property

Some of these new examples produce substitution tiling spaces with a surprising topological property.

These scaling factors are units in the relevant ring, which means that the induced substitution on the edge lengths (and the first cohomology of the Anderson–Putnam complex [1]) is unimodular, and so the first cohomology is finitely generated, and so is the module of eigenvalues in the dynamical spectrum (provided the factor is PV, otherwise there are no non-trivial eigenvalues). Despite all that, and unlike all examples found in the literature so far, (some of) these substitutions have a non-unimodular action on the second cohomology, which is thus not finitely generated.

All of the substitutions coming from Fig. 10 have this property, as does the substitution from the bottom row of Fig. 9 that uses the leftmost result for $T(2, 2, 3)$. Some of the substitutions from Fig. 11 have this property and others do not.

It is also easy to find examples with this property using the fivefold prototiles $S(5)$.

4.5 Tilings with Global Sevenfold Rotational Symmetry

Some of the tilings arising from the substitutions in Figs. 9, 10, and 11 possess global rotational sevenfold symmetry. Whether or not a substitution gives rise to such tilings can be determined by looking at the collection of vertex stars of the tilings—i.e., the translation equivalence classes of patches consisting of tiles that share a common vertex, which lies in the interior of the patch. There is a simple algorithm for obtaining all such patches: Choose a prototile as a seed and apply the substitution to it enough

times that the resulting patch contains a vertex in its interior (and hence contains vertex stars). Put these vertex stars in a list. Then apply the substitution to these vertex stars and add to the list any new vertex stars that appear in the resulting patches. Repeat this process until no more vertex stars are found.

Once all the vertex stars have been found, one can check visually whether or not any of them is sevenfold symmetric. If none of them is, then there is no globally sevenfold symmetric tiling in the substitution tiling space. If, on the other hand, there is at least one sevenfold symmetric vertex star, then there must be some sevenfold symmetric vertex star that lies at the center of its own image under some power of the substitution. This vertex star is then a seed for a sevenfold symmetric tiling that is a fixed point of a power of the substitution.

Of the substitutions in Fig. 9, the ones in the first two rows both produce sevenfold symmetric vertex stars, but in the third row only the substitution that uses the second result for $T(2, 2, 3)$ produces such stars. The substitutions in Fig. 10 do not produce any sevenfold symmetric vertex stars. Of the substitutions in Fig. 11, all of the ones using the $T(2, 2, 3)$ results from the collection on the right produce sevenfold symmetric vertex stars. Of the ones using results on the left, some do produce such stars and some do not.

4.6 Tilings with Global Elevenfold Rotational Symmetry

One of the original goals of this work was to discover new substitution tilings with elevenfold rotational symmetry. We were unable to find any; indeed, we found no substitutions on minimal subsets of $S(11)$. This is because the difficulty of the search increases rapidly with the size of the inflation factor λ .

We searched all small inflation factors for $\{T(1, 5, 5), T(1, 4, 6), T(2, 3, 6), T(2, 4, 5), T(3, 3, 5)\}$ and found no substitutions, and the larger inflation factors are too large to be searched within reasonable time. This is not due to a lack of memory—the discussion at the end of Sect. 3.4 describes a method of reducing memory usage to reasonable levels—but is rather due to CPU limitations.

Acknowledgments The first and third authors were partly supported by the German Research Council (DFG), CRC 701. The third author was also partly supported by the Fields Institute during a research visit.

References

1. Anderson, J.E., Putnam, I.F.: Topological invariants for substitution tilings and their associated C^* -algebras. *Ergod. Theory Dyn. Syst.* **18**(3), 509–537 (1998)
2. Baake, M., Moll, M.: Random noble means substitutions. In: Schmid, S., Withers, R.L., Lifshitz, R. (eds.) *Aperiodic Crystals*, pp. 19–27. Springer, Dordrecht (2013)
3. Dekking, F.M.: The spectrum of dynamical systems arising from substitutions of constant length. *Z. Wahrscheinlichkeitstheorie Verw. Gebiete* **41**(3), 221–239 (1977/1978)
4. Durand, F.: Linearly recurrent subshifts have a finite number of non-periodic subshift factors. *Ergod. Theory Dyn. Syst.* **20**, 1061–1078 (2000).
5. Ferenczi, S.: Rank and symbolic complexity. *Ergod. Theory Dyn. Syst.* **16**, 663–682 (1996)
6. Frank, N., Sadun, L.: Fusion: a general framework for hierarchical tilings of \mathbb{R}^d . *Geom. Dedicata* **171**, 149–186 (2013)

7. Frank, N.P.: Multidimensional constant-length substitution sequences. *Topol. Appl.* **152**(1–2), 44–69 (2005)
8. Frettlöh, D.: Inflationäre Pflasterungen der Ebene mit D_{2m+1} -Symmetrie und minimaler Musterfamilie. Diploma Thesis, Universität Dortmund (1998). <http://www.math.uni-bielefeld.de/~frettløe/papers/diplom>
9. Frettlöh, D., Harriss, E.O.: The tilings encyclopedia. <http://tilings.math.uni-bielefeld.de/>. Accessed 01 Apr 2014
10. Gähler, F., Maloney, G.R.: Cohomology of one-dimensional mixed substitution tiling spaces. *Topol. Appl.* **160**(5), 703–719 (2013)
11. García Escudero, J.: Randomness and topological invariants in pentagonal tiling spaces. *Discrete Dyn. Natl. Soc.*, 1–23 (2011)
12. Godrèche, C., Luck, J.M.: Quasiperiodicity and randomness in tilings of the plane. *J. Stat. Phys.* **55**(1–2), 1–28 (1989)
13. Goodman-Strauss, C.: Matching rules and substitution tilings. *Ann. Math.* **147**(1), 181–223 (1998)
14. Grünbaum, B., Shephard, G.C.: *Tilings and Patterns*. W. H. Freeman, New York (1986)
15. Harriss, E.O.: Non-periodic rhomb substitution tilings that admit order n rotational symmetry. *Discrete Comput. Geom.* **34**(3), 523–536 (2005)
16. Kamae, T.: A topological invariant of substitution minimal sets. *J. Math. Soc. Jpn.* **24**, 285–306 (1972)
17. Nilsson, J.: On the entropy of a family of random substitutions. *Monatsh. Math.* **168**(3–4), 563–577 (2012)
18. Nilsson, J.: On the entropy of a two step random Fibonacci substitution. *Entropy* **15**(9), 3312–3324 (2013)
19. Nischke, K.P., Danzer, L.: A construction of inflation rules based on n -fold symmetry. *Discrete Comput. Geom.* **15**(2), 221–236 (1996)
20. Pacheco, R., Vilarinho, H.: Statistical stability for multi-substitution tiling spaces. *Discrete Contin. Dyn. Syst.* **33**(10), 4579–4594 (2013)
21. Rivlin, T.J.: *Pure and Applied Mathematics. Chebyshev Polynomials*, 2nd edn. Wiley, New York (1990)
22. Solomyak, B.: Dynamics of self-similar tilings. *Ergod. Theory Dyn. Syst.* **17**(3), 695–738 (1997)