

AN APERIODIC MONOTILE

David Smith¹, Joseph Samuel Myers^{*2}, Craig S. Kaplan³, and Chaim Goodman-Strauss⁴

¹Yorkshire, UK

ds.orangery@gmail.com

²Cambridge, UK

jsm@polyomino.org.uk

³School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada

csk@uwaterloo.ca

⁴National Museum of Mathematics, New York, New York, U.S.A.

chaimgoodmanstrauss@gmail.com

Preprint: March 2023

© The authors.

Abstract. A longstanding open problem asks for an aperiodic monotile, also known as an “einstein”: a shape that admits tilings of the plane, but never periodic tilings. We answer this problem for topological disk tiles by exhibiting a continuum of combinatorially equivalent aperiodic polygons. We first show that a representative example, the “hat” polykite, can form clusters called “metatiles”, for which substitution rules can be defined. Because the metatiles admit tilings of the plane, so too does the hat. We then prove that generic members of our continuum of polygons are aperiodic, through a new kind of geometric incommensurability argument. Separately, we give a combinatorial, computer-assisted proof that the hat must form hierarchical—and hence aperiodic—tilings.

Keywords. Tilings, aperiodic order, polyforms

Mathematics Subject Classifications. 05B45, 52C20, 05B50

1. Introduction

Given a set of two-dimensional tiles, the nature of the planar tilings that they admit arises from a deep interaction between the local and the global. Constraints on the ways that pairs of tiles can be neighbours determine the structure of an infinite tiling, at all large scales. Constraints encoded in a set of tiles determine the structure of the space of the tilings that it admits, in subtle ways.

Aperiodic sets of tiles walk a fine line between order and disorder, admitting tilings, but only those without any translational symmetry, never permitting the simple repetition of periodic

*Development of software used in this work was supported in part by a Senior Rouse Ball Studentship for 2002–3 from Trinity College, Cambridge.




Figure 1.1: The gray “hat” polykite tile is an “einstein”, an aperiodic monotile. In other words, copies of this tile may be assembled into tilings of the plane (the tile “admits” tilings), yet copies of the tile cannot form periodic tilings, tilings that have translational symmetry. In fact, the tile admits uncountably many tilings. In Sections 2, 4, and 5 we describe how these tilings all arise from substitution rules, showing that they all have the same local structure.

tiling. Their study dates to Wang’s work [Wan61] on the then remaining open cases of Hilbert’s *Entscheidungsproblem*. Wang encoded logical fragments by what are now known as *Wang tiles*, congruent squares with coloured edges, to be tiled by translation only with colours matching on adjoining edges. He conjectured that every set of Wang tiles that admits a tiling (possibly using only a subset of the tiles) must also admit a periodic tiling, and showed that this would imply the decidability of the *tiling problem* (or *domino problem*): the question of whether a given set of Wang tiles admits any tilings at all. The algorithm would consist of enumerating all possible ways to cover larger and larger disks. Eventually one either will run out of ways to continue, and the tiles do not admit a tiling; or, if there is a fundamental domain for a periodic tiling, one will eventually discover it. If Wang’s conjecture held and aperiodic sets of tiles did not exist, this algorithm would always terminate.

Berger [Ber66] then showed that it was undecidable whether a set of Wang tiles admits a tiling of the plane. He constructed the first aperiodic set of 20426 Wang tiles, which he used as a kind of scaffolding for encoding finite but unbounded runs of arbitrary computation.

Subsequent decades have spawned a rich literature on aperiodic tiling, touching many different mathematical and scientific settings—we do not attempt a broad survey here. Yet there remain remarkably few really distinct methods of proving aperiodicity in the plane, despite or due to the underlying undecidability of the tiling problem.

Berger’s initial set comprised thousands of tiles, naturally prompting the question of how small a set of tiles could be while still forcing aperiodicity. Professional and amateur math-

ematicians produced successively smaller aperiodic sets, culminating in discoveries by Penrose [Pen78] and others of several consisting of just two tiles. Surveys of these sets appear in Chapters 10 and 11 of Grünbaum and Shephard [GS16] and in an account of the Trilobite and Cross tiles [GS99]. A recent table appears in the work of Greenfeld and Tao [GT21].

The obvious conclusion of this reduction in size would be to arrive at an “einstein”,¹ a single shape that tiles aperiodically. It has long been an open question whether such a tile exists. Can one tile embody enough complexity to forcibly disrupt periodic order at all scales?

1.1. The search for an einstein

Several candidate tiles have been proposed as einsteins, but they all challenge in some way the concepts of “tile”, “tiling”, or “aperiodic”.

Gummelt [Gum96] and Jeong and Steinhardt [SJ96, JS97] describe a single regular decagon that can cover the plane with copies that are allowed to overlap by prescribed rules, but only non-periodically, in a manner tightly coupled to the Penrose tiling. Senechal [Sen] similarly describes simple rules that allow copies of the Penrose dart to overlap and cover the plane, but never periodically. The result is an ingenious route to aperiodicity, but not a tiling in the usual sense.

The Taylor-Socolar tile [ST11] comes within a hair’s breadth of being aperiodic. Their hexagonal prototile tiles aperiodically, but only when additional matching conditions are enforced, preventing certain adjacencies that would otherwise be permitted. These matching conditions mandate relationships between non-adjacent tiles, making it impossible to encode the global aperiodicity in the shape of a two-dimensional topological disk. The matching conditions can be expressed purely geometrically, but doing so requires either a disconnected tile, a three-dimensional tile, or one with cutpoints [ST12].

The structure of the Taylor-Socolar tiling is closely related to Penrose’s $1+\epsilon+\epsilon^2$ tiling [Pen97, BGG12, Tay]. Like the Trilobite and Crab tiles [GS99], these can be adjusted so that an arbitrarily high fraction of the area lies in copies of just one kind of tile, but no matter how thin or small, the other tiles remain necessary.

Matching rules for how tiles may fit together have taken many different forms in the literature, and, loosely speaking, it is often possible to shift the complexity in a construction from the tiles to the matching rules or vice versa. For example, if we use a finite atlas of finite configurations as our allowed matching rules, even the lowly 2×1 rectangle is an aperiodic monotile!² «««| HEAD Watson and Whittaker recently produced a hexagon marked with an elegant and simple edge-to-edge “orientational” rules, that nonetheless cannot ===== Watson and Whittaker recently produced a hexagon marked with an elegant and simple “orientational parity” rule, that nonetheless cannot be »»»| 0bec59a29b004399ff47d3bd3da0658dad36882e encoded as a single unadorned shape [WW21]. The plainest possible rules are just that the tiles must fit together without gaps or overlapping interiors, and these remain the gold standard for an aperiodic

¹A pun from the German “ein stein”, roughly “one shape”, popularized by Danzer.

²Beginning with an aperiodic set of tiles with, say, geometric matching rules, pixelate pictures of the tiles and how they fit together, in some black and white bit-map. Take an atlas of these pictures, splitting black pixels vertically and white ones horizontally into identical rectangles. The rectangle is an aperiodic monotile with this atlas of matching rules.

monotile.

Moving to higher dimensional space permits richer forms of aperiodicity to arise. The Schmitt-Conway-Danzer tile [Sen96, Section 7.2] tiles \mathbb{R}^3 , with tilings that have a screw motion but not translations as symmetries; no periodic tiling by copies of this tile has a compact fundamental domain, and we call this tile *weakly aperiodic*. Weak aperiodicity is indeed weak, and readily appears in the hyperbolic plane and other non-amenable spaces—as early as 1974, Böröczky exhibited a weakly aperiodic monotile in the hyperbolic plane [Bör74], the elegantly simple basis of the “binary tilings” [BW92, GS09, MM98, Moz97].

Following Mozes [Moz97], we say a set of tiles is *strongly aperiodic* if it admits tilings but none with any infinite cyclic symmetry. In the plane, for sets of “normal” tiles, these concepts coincide [GS16, Theorem 3.7.1] and we do not need to make this distinction.

Recently, Greenfeld and Tao [GT22] showed that in sufficiently high dimensions, a single tile, tiling *only by translation*, can be aperiodic in \mathbb{Z}^n (and thus in \mathbb{R}^n). They also showed that it is undecidable whether a single tile, again tiling by translation, admits a tiling of a periodic subset of $\mathbb{Z}^2 \times G$ for some nonabelian group G [GT21]. By contrast, in \mathbb{R}^2 , Kenyon [Ken92, Ken93, Ken96], building on the work of Girault-Beauquier and Nivat [GBN91], showed that any topological disk that admits a tiling by translation also admits a periodic tiling, while Bhattacharya [Bha20] showed the same for any set in \mathbb{Z}^2 .

Little is known about limits on what sorts of shapes could potentially be aperiodic monotiles. Rao [Rao17] showed through a computer search that the list of 15 known families of convex pentagons that tile the plane is complete, thereby eliminating any remaining possibility that a convex polygon could be an einstein. Jeandel and Rao [JR21] showed that the smallest aperiodic set of Wang tiles is of size 11.

Even when a single tile admits periodic tilings, that periodicity may be more or less abstruse, in a way that offers tantalizing hints about aperiodicity. Here, the *isohedral number* of a tile is the minimum number of transitivity classes in any tiling it admits; a tile is *anisohedral* if its isohedral number is greater than one. The second part of Hilbert’s 18th problem [Hil02] asked whether there exist anisohedral polyhedra in \mathbb{R}^3 . Grünbaum and Shephard suggest [GS16, Section 9.6] that this question was asked in \mathbb{R}^3 because Hilbert assumed that no such tiles exist in the plane, yet Reinhardt [Rei28] found an example of such a polyhedron, and Heesch [Hee35] then gave an example of such a tile in the plane. Many anisohedral prototiles are known today. The computer enumeration by Myers [Mye19] furnished numerous anisohedral polyominoes, polyhexes, and polyiamonds, including a record-holding 16-hex that tiles with a minimum of ten transitivity classes. It is unknown whether there is an upper bound on isohedral numbers of monotiles.³

Related insights can be gleaned from the study of shapes that do not tile the plane. A shape’s *Heesch number* is the largest possible combinatorial radius of any patch formed by copies of the shape (or equivalently, the maximum number of complete concentric rings that can be constructed around it). A shape that tiles the plane is said to have a Heesch number of ∞ . Heesch first exhibited a shape with Heesch number 1, and a few isolated examples with Heesch numbers

³The problem of determining whether or not a given set of tiles admits a periodic tiling is also undecidable, at least for larger sets of tiles [GK72]. If we enumerate sets of tiles, and define $I(n)$ to be the isohedral number of the n th set, if it admits a periodic tiling, and let $I(n) = -1$ otherwise, then $I(n)$ cannot be bounded by any computable function. This defies our imagination.

up to three were discovered thereafter [Man04]. Mann and Thomas discovered marked polyforms with Heesch numbers up to 5 through a brute-force computer search [MT16]. Kaplan conducted a search on unmarked polyforms [Kap22], yielding examples with Heesch numbers up to 4. Bašić discovered the current record-holder, a shape with Heesch number 6 [Baš21]. *Heesch's problem* asks which positive integers can be Heesch numbers; beyond specific examples with Heesch numbers up to 6, nothing is known about the solution. An upper bound on finite Heesch numbers would imply the decidability of the tiling problem for monotiles. The algorithm would simply consist of generating all possible concentric rings around a central tile; eventually you will either fail (in which case the shape does not tile the plane) or exceed the upper bound on Heesch numbers (in which case it must tile the plane).

1.2. Outline




Figure 1.2: The high-level structure of the first proof of aperiodicity in this paper

In this paper, we prove the following:




Figure 1.3: The high-level structure of the second proof of aperiodicity in this paper

Theorem 1.1. *The shape shown shaded in Figure 1.1, a polykite that we call “the hat”, is an aperiodic monotile.*

No special qualifications or additional matching conditions are required: as shown, this shape tiles the plane, but never with any translational symmetries. The shape is almost mundane in its simplicity. It is a *polykite*: the union of eight kites in the Laves tiling [3.4.6.4], the dual to the (3.4.6.4) Archimedean tiling.

We provide two different proofs of aperiodicity, both with novel aspects. The first proof follows the structure shown in Figure 1.2, centred on a new approach in Section 3 for proving aperiodicity in the plane. We observe that any tiling by the hat corresponds to tilings by two different polyiamonds, one with two thirds the area of the other. If there were a strongly periodic tiling by the hat, the other two tilings would also be strongly periodic. We prove that if so, the lattices of translations in the polyiamond tilings would necessarily be related by a similarity; but no similarity between lattices of translations on the regular triangular tiling can have the scale factor $\sqrt{2}$ required by the ratio of the areas. This does not show that a tiling exists, and so must be combined with an explicit construction of a tiling (outlined in Section 2 and given in detail in Sections 4 and 5) to complete the proof of aperiodicity.

The second proof presented (but the first one found) follows the structure shown in Figure 1.3. Here we generally adhere to Berger’s approach, but we must begin with a novel step to get to the point where such a proof is possible. We first show that in any tiling by the hat polykite, every tile belongs uniquely to one of four distinct clusters (Section 4), and that those clusters fit together following certain matching rules. It is then the clusters that allow a more standard style of hierarchical construction. This remarkable behaviour has not been seen previously in the literature.

Defining matching rules on the boundaries of the clusters allows us to discard details of how the clusters are made up of hats, and instead consider simplified outlines we refer to as *metatiles*. Following Berger et al., we then give an inductive proof in Section 5 showing that any tile in any tiling by these four metatiles lies in a unique hierarchy of larger and larger *supertiles*, effectively combinatorial copies of the metatiles, at larger and larger scales. The proof is constructive: we first show the metatiles can only lie within larger clusters, the 1-level supertiles, uniquely, and that these clusters have the same combinatorial structure as the metatiles. In turn, in the same manner, the 1-level supertiles can only lie uniquely within 2-level supertiles, again clusters with the same combinatorics, and so on for subsequent levels. This construction proves that tilings by copies of the metatiles can only be non-periodic, because if there were a translational symmetry, these hierarchies of supertiles could not be unique. It also shows that the metatiles (and hence the hats) admit tilings of the plane, because we construct clusters of arbitrary size [GS16, Theorem 3.8.1].

Because of the combinatorial complexity of the hat polykite, a significant fraction of our second proof relies on exhaustive enumeration of cases, which we carried out and cross-checked with two independent software implementations developed by two of the authors in isolation. These calculations are necessarily ad hoc, and are essentially unenlightening. This case analysis is only needed to show that all tilings follow the substitution structure; it is not needed for showing that a tiling exists, and thus is not needed to show that the tile is aperiodic, given the proof in Section 3 that no periodic tiling exists.

We close this introduction with definitions of the essential terminology we will need for the rest of the article. In Section 2, we then present a compendium of provisional observations about this polykite, including an explicit construction of a tiling and aspects of its structure that deserve further study. Our two proofs of aperiodicity follow: we show that there are no periodic tilings (Section 3), then that tiles must group into clusters that define metatiles equipped with matching rules (Section 4), and finally that metatiles must compose into supertiles with combinatorially equivalent matching rules (Section 5). In Section 6, we offer additional remarks about the continuum of tiles that contains the hat polykite. As noted there, computer search shows that the hat is the smallest aperiodic polykite.

1.3. Terminology

Terminology used for tilings generally follows that of Grünbaum and Shephard [GS16].

A *tile* in a metric space is a closed set of points from that space. A *tiling* by a set of tiles is a collection of images of tiles from that set under isometries, the interiors of which are pairwise disjoint and the union of which is the whole space; we say a set of tiles *admits* the tiling, or in

the case of a single tile that it admits the tiling. For most purposes, it is convenient for tiles to be nonempty compact sets that are the closures of their interiors; the tiles considered here are polygons, or more generally closed topological disks. A tiling is *monohedral* if all its tiles are congruent. All tilings considered here are also *locally finite*: every point has some open neighbourhood that meets only finitely many tiles (all monohedral plane tilings by closed topological disks are locally finite).

In any locally finite tiling of the plane by closed topological disks, the connected components of the intersection of two or more tiles are isolated points, which are called *vertices* of the tiling, and Jordan arcs, which are called *edges* of the tiling, and the boundary of any tile is divided into finitely many edges, alternating with vertices. Each edge lies on the boundary of exactly two tiles, which we refer to as lying on opposite sides of the edge. Two distinct tiles are *neighbours* if they share any point of their boundaries, and *adjacents* if they share an edge.

When a (closed topological disk) tile has a polygonal boundary, we refer to it as having *sides* (maximal straight line segments lying on that boundary) and *corners* (between two sides), to distinguish from the edges and vertices of a tiling. We rely on context to distinguish the meanings of “side” as referring to sides of a polygon or the two sides of an edge of a tiling. A tiling by polygons is *edge-to-edge* if the corners and sides of the polygons coincide with the vertices and edges of the tiling.

A *patch* of tiles is a collection of non-overlapping tiles whose union is a topological disk. More specifically, a *0-patch* is a patch containing a single tile, and an $(n + 1)$ -*patch* is a patch formed from the union of an n -patch P and a set S of additional tiles, so that P lies in the interior of the patch and no proper subset of S yields a patch with P in its interior. (In other words, an n -patch is a tile surrounded by n concentric rings of tiles.) In a fixed tiling, every tile generates an n -patch for all finite n , by recursively constructing an $(n - 1)$ -patch and adjoining all its neighbours in the tiling, along with any other tiles required to fill in holes left by adding neighbours.

Given a tiling \mathcal{T} , a *poly- \mathcal{T} -tile* is a closed topological disk that is the union of finitely many tiles from \mathcal{T} . Poly- \mathcal{T} -tiles are also referred to generically as *polyforms*. Poly- \mathcal{T} -tiles may also be defined so that they are permitted to have holes. Because we are mainly concerned with tiles that admit monohedral tilings, it is not generally significant for the purposes of this paper whether shapes with holes are allowed or not.

The *symmetry group* of a tiling is the group of those isometries that act as a permutation on the tiles of the tiling. A tiling is *weakly periodic* if its symmetry group has an element of infinite order; in the plane, this means it includes a nonzero translation. A tiling is *strongly periodic* if the symmetry group has a discrete subgroup with cocompact action on the space tiled. In Euclidean space, all strongly periodic tilings are also weakly periodic. A set of tiles (or a single tile) is *weakly aperiodic* if it admits a tiling but does not admit a strongly periodic tiling, and *strongly aperiodic* if it admits a tiling but does not admit a weakly periodic tiling.

Any finite set of polygons in the plane that admits a weakly periodic edge-to-edge tiling also admits a strongly periodic tiling [GS16, Theorem 3.7.1], and a similar but simpler argument shows the same to be the case for a finite set of poly- \mathcal{T} -tiles where \mathcal{T} is itself a strongly periodic tiling and the weakly periodic tiling consists of copies of the tiles all aligned to the same underlying copy of \mathcal{T} , instead of being edge-to-edge. Thus in such contexts it is not necessary to

distinguish weak and strong aperiodicity and we refer to tiles and sets of tiles simply as *aperiodic*.

A *uniform tiling* [GS16, Section 2.1] is an edge-to-edge tiling by regular polygons with vertex-transitive symmetry group; notation such as (3.4.6.4), listing the sequence of regular polygons round each vertex, denotes a uniform tiling. A *Laves tiling* [GS16, Section 2.7] is an edge-to-edge monohedral tiling by convex polygons with regular vertices (all angles between consecutive edges at a vertex equal) and tile-transitive symmetry group; analogous notation such as [3.4.6.4] is used for Laves tilings, listing the sequence of vertex degrees round each tile, and in an appropriate sense Laves tilings are dual to uniform tilings.

2. The hat polykite and its tilings

Before proceeding to the full proof of aperiodicity, we first offer a less formal presentation of the hat, including an explicit construction of a tiling. This section fulfills three goals. First, it offers an abundance of visual intuition, which provides context for the technical machinery that will follow. Second, it gives some sense of our process of discovery and analysis, though it should not be interpreted as an ordered timeline. Third, it includes a few observations that will not be considered further in this article, but which might provide opportunities for future work by others.

One of the authors (Smith) began investigating the hat polykite as part of his open-ended visual exploration of shapes and their tiling properties. Working largely by hand, with the assistance of Scherphuis’s PolyForm Puzzle Solver software (www.jaapsch.net/puzzles/polysolver.htm), he could find no obvious barriers to the construction of large patches, and yet no clear cluster of tiles that filled the plane periodically.

Because the hat is a polyform, it was natural at this point to obtain an initial diagnosis of its tiling properties computationally. We modified Kaplan’s SAT-based Heesch number software [Kap22] to determine that if the hat does not tile the plane, then its Heesch number must be at least 16. Similarly, we modified Myers’ polyform tiling software [Mye19] to determine that if the hat admits periodic tilings, then its isohedral number must be at least 64. These two computations already establish that the hat is of extreme interest—if it had turned out not to be an einstein, then it would have shattered either the record for Heesch numbers or the record for isohedral numbers, in both cases by a wide margin!

Figure 2.1 (left) shows a computer-generated 10-patch (i.e., ten concentric rings of tiles around a shaded central tile, where each tile in a ring touches the ring it encloses in at least one point). It was constructed by allowing Kaplan’s software to work outward to that radius, and then stopping it manually. At first glance, it can be difficult to discern any structure at all in this patch. However, by colouring the tiles in different ways, clear “features” begin to emerge. Of course, we cannot infer any conclusive properties of infinite tilings from a finite computed patch. We must be particularly wary of tiles near the periphery of the patch, where features may break down under the extra freedom afforded by the proximity to empty space. However, for a sufficiently large patch, we might hope that tiles near the centre will be representative of configurations that arise in generic tilings.

The most important colouring for the purposes of this article is the one shown on the right in Figure 2.1. A single hat is asymmetric, and so in any patch we can distinguish between




Figure 2.1: A computer-generated 10-patch of 391 hats (left), arranged in ten concentric rings around a central shaded hat. The tiles can be coloured (right), showing that the reflected hats (dark blue) are sparsely distributed and each is surrounded by a congruent “shell” of three unrelected hats (light blue). A thickened outline shows the boundary of the maximal cluster of tiles that appears congruently around every reflected tile.




Figure 2.2: Long chains of similarly oriented tiles pass through reflected tiles in six directions (left). We can merge each reflected tile with one of its neighbours in its chain (centre), yielding a structure that can be placed into one-to-one correspondence with a patch of regular hexagons (right).

“unreflected” and “reflected” orientations of tiles. In the patches we computed, reflected tiles, shown in dark blue, are always distributed sparsely and evenly within a field of unreflected tiles. Furthermore, every reflected tile is contained within a congruent cluster of nine tiles, where the other eight tiles in the cluster are unreflected. One such cluster is outlined in bold in the illustration. The interior of the patch can be covered completely by overlapping copies of that cluster. Within the cluster, we are particularly interested in the “shell” of three light blue tiles adjacent to each reflected tile. Every reflected tile resides in a congruent, non-overlapping copy of this shell.

We have also observed that unreflected tiles tend to form long “chains” of like orientation, occasionally interrupted by reflected tiles. The chains contained in the example patch are shown coloured on the left in Figure 2.2. Because the hats are aligned with the underlying kite grid, unreflected tiles come in six orientations, all of which also appear as chain directions. Chains may end at reflected tiles or pass through them, but each reflected tile is a hub for at least two, and at most five spokes. Long segments of these chains have boundaries with halfturn symmetry. It is tempting to seek parallels between these chains and linear features in other aperiodic tilings, such as Ammann bars [GS16, Section 10.6] and Conway worms [GS16, Section 10.5]. Finally, we have noticed that these chains seem to impart a rough hexagonal arrangement to the hats, which is particularly clear in the triangular and parallelogram-shaped structures that are surrounded by chains. We have found that if we merge each reflected tile with its immediate neighbour as shown in Figure 2.2 (centre), then the tiles in any patch can be put into one-to-one correspondence with a patch of hexagons, as in Figure 2.2 (right). The hexagonal grid may provide a convenient domain in which to perform computations on the combinatorial structure of tilings by hats.




Figure 2.3: The two edge lengths in the hat polykite can be manipulated independently, producing a continuum of shapes. A selection of those shapes is shown here, normalized for scale. $\text{Tile}(0, 1)$, $\text{Tile}(1, 1)$, and $\text{Tile}(1, 0)$ admit periodic tilings; all others are aperiodic.

In the course of his explorations, the first author discovered a *second* polykite that did not seem to have a finite isohedral number or a finite Heesch number, this one a union of ten kites that we call the “turtle”. The idea of identifying two einsteins back-to-back seemed too good to be true! It was both a relief and a revelation when we determined that not only were the hat and the 10-kite related, they were in fact two points from a continuum of shapes that all tile the plane the same way. The hat is derived from the $[3, 4, 6, 4]$ grid, and therefore its edges come in two lengths, which we can take to be 1 and $\sqrt{3}$. Furthermore, these edges come in parallel pairs, allowing us to set the two lengths independently to any non-negative values. We use the notation $\text{Tile}(a, b)$ with a and b not both zero to refer to the shape produced when the length-1 and length- $\sqrt{3}$ edges are altered to have lengths a and b , respectively. Note that $\text{Tile}(a, b)$

is similar to $\text{Tile}(ka, kb)$ for any $k \neq 0$. In Figure 2.3 we show a selection of shapes along this continuum, normalized for scale. By this reckoning, the hat is $\text{Tile}(1, \sqrt{3})$ and the 10-kite is $\text{Tile}(\sqrt{3}, 1)$. We have also created an animation showing a continuous evolution of $\text{Tile}(a, 1-a)$ as a moves from 0 to 1 and back—see youtu.be/W-ECvtIA-5A. The tetriamond $\text{Tile}(0, 1)$, the octiamond $\text{Tile}(1, 0)$, and the equilateral $\text{Tile}(1, 1)$ admit simple periodic tilings; in Section 6, we will show that all other shapes in this continuum are aperiodic monotiles with combinatorially equivalent tilings. In Section 3, $\text{Tile}(0, 1)$ and $\text{Tile}(1, 0)$ will play a crucial role in establishing that the hat is aperiodic. Inspired by cut-and-project methods [DB81a, DB81b], we are also left wondering whether it would be productive to construct a closed path in four or six dimensions, which projects down to this family of tiles from a suitable set of directions.




Figure 2.4: A grouping of tiles into clusters in the example patch. In addition to four-tile clusters consisting of a reflected hat and its three-hat shell, we identify clusters consisting of a single tile, and parallelogram-shaped clusters consisting of pairs of tiles. The parallelograms come in two varieties: one separates two nearby shells, and the other joins up with two rotated copies to make a three-armed propeller shape called a *triskelion*. An isolated triskelion is shown shaded in grey in the lower right.

Given the colouring in Figure 2.2 showing non-overlapping clusters of reflected tiles and their shells, it is natural to wonder whether the remaining unaffiliated tiles in the patch reliably form clusters of other kinds. Figure 2.4 illustrates that we can account for all remaining tiles using two additional cluster types (shown separately on the right). First, where three shells meet they enclose a single isolated tile, which must be accepted as a cluster of size one. Then the remaining tiles group into congruent clusters of size two, which are roughly parallelogram-shaped. These appear in two varieties, depending on the local arrangement of clusters around them. In the first case, coloured white in the drawing, the parallelogram is adjacent to two shells along its long edges. In the second case, coloured grey, one end of the parallelogram is plugged into a local centre of threefold rotation, joining six hats into a three-armed propeller shape called

a *triskelion*. A triskelion is shown in isolation on the bottom right of Figure 2.4.

These clusters are the starting point for the definition of a substitution system, one that can be iterated to produce a patch of hats of arbitrary size. Interestingly, the substitution rules do not apply to the hats directly. Instead, we derive new *metatiles* from the clusters, and build a substitution system based on the metatiles. The underlying hats are simply brought along for the ride.




Figure 2.5: The H , T , P , and F metatiles (top), constructed by simplifying the boundaries of clusters of hats. We mark the H , T , and P metatiles with arrows when needed (bottom), to distinguish between otherwise symmetric orientations.

Figure 2.5 shows the shapes of the metatiles. Each one is constructed by simplifying the boundary of one of the clusters of hats in Figure 2.4. In order to ensure that the metatiles do not overlap, we must distinguish between the two varieties of two-hat parallelograms discussed above. Specifically, we remove a triangular notch from the parallelogram associated with each leg of a triskelion. Thus the three clusters yield four metatiles: an irregular hexagon (H), an equilateral triangle (T), a parallelogram (P), and a pentagonal triskelion leg (F). The original clusters can now be seen as endowing the metatiles with matching conditions along their edges; these matching conditions will be formalized in Section 4.

The H , T , and P metatiles have rotational symmetries. In the bottom row of Figure 2.5, we mark tiles with arrows showing their intended orientations. In each case, the arrow points to the (unique) side of the metatiles from which two adjacent kites protrude. The arrows suffice to distinguish symmetric rotations and our construction will not use reflections. (We will not need these arrows in later sections, as metatiles orientations will be implied by labels on their edges.)

We can now define a family of supertiles that are analogous to the metatiles, following the procedure illustrated in Figure 2.6. We first assemble the patch of oriented metatiles shown on the left. It can easily be checked that the elided hat polykites borne by these tiles fit together with no gaps and no overlaps. This patch is large enough to pick out one or more copies of each supertile, drawn in red in the central diagram. The supertile shapes are fully determined by two constraints: the red dots coincide with the centres of triskelions, and all interior angles of the hexagonal outlines are 120° . The diagram on the right shows the supercluster outlines in isolation, with their inherited orientation markings. Here, each arrow points to the unique supertile edge that passes through an outward-pointing P tile from the previous generation.




Figure 2.6: The construction of a family of supertiles from a patch of metatiles. The patch of metatiles on the left can be used to locate key vertices of the supertiles, marked with red dots in the central diagram. Those dots, together with constraints on angles, fully determine the shapes of the supertiles, which are not merely scaled-up copies of their progenitors. On the right, the supertiles are marked with arrows indicating their orientations.

At first glance, these supertiles appear to be scaled-up copies of the metatiles. If that were so, we could perhaps proceed to define a typical substitution tiling, where each scaled-up supertile is associated with a set of rigidly transformed tiles. However, with the obvious exception of the T , none of the supertiles is truly similar to its corresponding metatiles. Despite that discrepancy, the supertiles are fully compatible with the construction in Figure 2.6—they can be arranged in the same configuration shown on the left, and used as a scaffolding for deriving outlines of super-supertiles (implicitly yielding a much larger patch of hats along the way). Indeed, the construction can be iterated any number of times, with slightly different outlines in every generation. We are not aware of other substitution systems that use rules like these, where successive generations are combinatorially but not geometrically compatible. To see this construction in action, please try our interactive browser-based visualization tool at cs.uwaterloo.ca/~csk/hat/.

We know from Figure 2.5 that each of the four metatiles can be associated with a cluster of hats. The construction in Figure 2.6 can then be iterated any number of times to form ever-larger patches of metatiles, and hence of hats. We can, for example, consider the H supertiles formed through this process of iteration, and the patch of hats each one contains. These patches form a sequence that grows in radius without bound, each patch a subset of its successor. The Extension Theorem [GS16, Section 3.8] allows us to continue this iteration process “to infinity”, yielding the following result.

Theorem 2.1. *The hat polykite admits tilings of the plane.*

Of course, this theorem is not sufficient to establish aperiodicity on its own—we must also show that the hat does not also admit periodic tilings. In Section 5 we revisit this substitution process, tracking matching conditions on supertile edges after every step. There we show that *all* tilings by the hat necessarily obey the substitution rules given here (Theorem 5.1). The




Figure 2.7: The first four iterations of the H metatile and its supertiles. At each level, tiles partially overlap the boundary of their supertile. Overlaps are acceptable here, because the supertile will be met by neighbouring supertiles with the same configuration of smaller tiles on its boundary.

construction in that section also incidentally leads to a more detailed proof that the hat tiles the plane.

The shapes of each generation of supertiles are different from those of the generation before it. However, by normalizing the tiles for size, we have computed that they quickly converge on a fixed point, a set of tiles that truly do yield scaled copies of themselves under the construction in Figure 2.6. These converged tile shapes are particularly interesting because they can be used to define a geometric substitution system that operates via inflation and replacement. The converged tiles, together with their substitution rules, are shown in Figure 2.8. By virtue of its connection to the original metatiles in Figure 2.5, we know that this substitution tiling is aperiodic when the tiles are endowed with suitable matching conditions on their edges. We can also use this system as an alternative means of constructing patches of hats. We cannot simply associate a cluster of hats rigidly with each converged tile, but a patch of converged tiles is combinatorially equivalent to a corresponding patch of metatiles, which are equipped with hats.

If we rescale the converged tiles so that the short H edges have unit length, then all tile edges except the two F edges adjacent to a triskelion centre will have lengths in $\mathbb{Z}[\phi]$, where ϕ is the




Figure 2.8: A substitution system based on converged tile shapes. Scaling the tiles so that the short edges of the H tile have unit length, all tile edges except the two adjacent to triskelion centres have lengths in $\mathbb{Z}[\phi]$, where ϕ is the golden ratio. In each substitution rule, tiles shown with dashed boundaries can be omitted, leading to patches in which there are no duplicate tiles contributed by supertiles sharing an edge.

golden ratio. Furthermore, this substitution system has an inflation factor of ϕ^2 . The factor of ϕ^2 can also be derived algebraically, through an eigenvalue computation on the substitution matrix corresponding to the system presented in this section.

At first blush, it may be surprising to see ϕ arise in a tiling closely associated with the Laves tiling [3.4.6.4]; it appears more naturally in contexts such as Penrose tilings, which feature angles derived from the regular pentagon. The involvement of ϕ appears to be closely related to the appearance of $\sqrt{2}$ in the argument of Section 3; that number is also not expected on the regular triangular tiling or related contexts (where distances are the square roots of integers that can be expressed by the quadratic form $x^2 + xy + y^2$, so expected square roots are of 3 and primes of the form $6k + 1$). However, $1 + \phi^{-1} + \phi^{-2} = 2$, from which it follows that a triangle with a 120° angle between sides of lengths 1 and ϕ^{-1} has a third side of length $\sqrt{2}$ (Figure 2.9, left). And indeed, by aligning corresponding tiles of tilings using $\text{Tile}(1, 0)$ and $\text{Tile}(0, \sqrt{2})$, it appears that there is an angle $\tan^{-1} \sqrt{3/5}$, which is one of the angles of the triangle with sides 1, ϕ^{-1} , and $\sqrt{2}$, between the edges of the regular triangular tilings underlying $\text{Tile}(1, 0)$ and $\text{Tile}(0, \sqrt{2})$ (Figure 2.9, right).

The arrangement of three H tiles and a T tile inside of an H supertile mimics the arrangement of a reflected hat and its three unreflected neighbours in a single H metatile. We are naturally led to wonder whether the clusters of hats that make up the metatiles are primordial, or whether they are preceded by a set of “subclusters” that launch the substitution process one step earlier. A possible form for such subclusters is shown in Figure 2.10. The labels on the edges denote




Figure 2.9: A demonstration of how the golden ratio ϕ might arise in the context of tilings by hats. The triangle on the left has an angle of 120° between sides of lengths 1 and ϕ^{-1} ; from trigonometric identities we can compute that $x = \sqrt{2}$ and $\theta = \tan^{-1} \sqrt{3/5}$. On the right we show portions of tilings by $\text{Tile}(1, 0)$ (green) and $\text{Tile}(0, \sqrt{2})$ (blue), registered to the same centres of local threefold rotation. The angle between the edges of the triangle tilings underlying these two polyiamond tilings is approximately θ , and we believe this approximation converges as we register larger patches of the two tilings.

matching conditions that will be explained in detail in Section 4. Note that subclusters P_0 and F_0 have zero area; their boundaries are shown split into multiple parts to clarify the sequence of edges (in the case of F_0 , some of those edges intersect others); note also that the X^+ and X^- edges from Figure 4.1 have length zero in the subclusters and are not shown in the diagram. Defining what exactly it means to partition a tiling into subclusters following matching rules, when some subclusters have area zero and some edges have length zero but must still adjoin in the correct orientations, seems more awkward than the corresponding argument based on metatiles, so we do not pursue the subclusters further.

Finally, in Figure 2.11 we exhibit an alternative substitution system based on a different set of clusters. Here, each reflected hat is merged with a specific neighbour in its shell, as shown in the upper left, to form a larger two-hat compound. We can then define just two combinatorial substitution rules: one replaces a two-hat compound by a cluster of a compound and five hats (labelled H_7 in the figure), and the other replaces a single hat by a cluster of a compound and six hats (labelled H_8). As with the original metatiles, this process can be iterated to produce a patch of any size, after which each compound can be split back into a pair of hats. This substitution system is attractive for its minimality, though we believe it would be more cumbersome for proving aperiodicity. Although the tilings themselves are MLD (mutually locally derivable) [BSJ91] with those by the H , T , P , and F metatiles presented earlier, deriving those metatiles from the clusters shown here requires considering a radius larger than a single cluster. The two-hat compound's cluster is always equivalent to the union of an H metatile, a T metatile, and an F metatile, but the hat's cluster corresponds to three different combinations of H , P , and F . Alternatively, to establish an MLD system, we could define four congruent but combinatorially inequivalent copies of the substitution rule for a single hat. It is impossible to define the hat tiling through a single substitution rule—the implied substitution matrix would necessarily yield

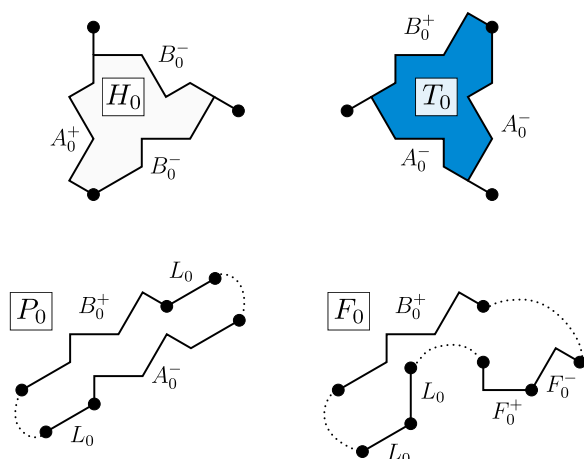


Figure 2.10: Four subclusters that may be thought of as preceding the clusters making up the metatiles in Figure 2.5. Edges are marked with the labels that will be used in Section 4. The P_0 and F_0 subclusters have zero area; dotted lines indicate vertices that should be regarded as coincident.




Figure 2.11: An alternative substitution system that yields the same tiling by hats as the system presented earlier. We merge a reflected hat with one of its neighbours (top left) to produce a two-hat compound. We can then define substitution rules that replace a single hat by a cluster labelled H_7 and a compound by a cluster labelled H_8 . Two additional iterations of the H_8 rule are shown.

a rational asymptotic increase in area after substitution, whereas we already know that the hat tiling inflates areas by a factor of ϕ^4 .




Figure 2.12: An excerpt from a very large patch generated using the substitution system presented in this section. In the right half of the drawing, hats belonging to F metatiles are coloured black, to highlight the interlocking tree structures formed by the triskelions and the other metatiles.

The ideas presented in this section are sufficient to show that the hat does in fact tile the plane. Figure 2.12 offers a final large patch of tiles as a demonstration. On the right side of the illustration we observe that the tiles belonging to triskelions form a connected tree structure that interlocks with a tree formed from the remaining tiles. This structure is reminiscent of those found in other aperiodic tilings, such as the Taylor-Socolar tiling and the $1 + \epsilon + \epsilon^2$ tiling.

However, exhibiting a tiling is usually the easy part of a proof of aperiodicity; it is also necessary to prove that none of the tilings admitted by the hat can be periodic. In the next section we present a novel geometric proof of aperiodicity. Then, in Sections 4 and 5 we turn to a more standard combinatorial argument that the matching rules implied by the substitution system shown earlier are forced in tilings by the hat.

3. Aperiodicity via coupling of polyiamond tilings

In this section, we prove the following result:

Theorem 3.1. *Let \mathcal{T} be a tiling by the hat polykite. Then \mathcal{T} is not strongly periodic.*




Figure 3.1: A patch from a hat tiling \mathcal{T} (left), which in this section, for eventual contradiction, is hypothesized to be strongly periodic. By contracting edges we can construct patches from corresponding \mathcal{T}_4 (centre) and \mathcal{T}_8 (right) tilings. Equivalent reference tiles are marked with asterisks in the three patches.

As noted in Section 1.3, a tile that does not admit strongly periodic tilings also cannot admit weakly periodic tilings. Therefore, together with the substitution system outlined in Section 2, and described in detail in Sections 4 and 5, this theorem establishes that the hat is an aperiodic monotile, thus proving Theorem 1.1. The latter two sections provide a detailed case analysis showing that all tilings by the hat polykite are given by the substitution system, and as a consequence giving an alternative proof that it is an aperiodic monotile. Deducing aperiodicity from the result of this section does not rely on that case analysis, only on the *existence* of tilings. We require only the description of the substitution system and the clusters of tiles used in that system, not the proof that all tilings necessarily obey that system.

We suppose throughout this section that there is a strongly periodic tiling \mathcal{T} by the hat polykite, and derive a contradiction. We also suppose this tiling is aligned to an underlying [3.4.6.4] Laves tiling; this supposition is justified by Lemma A.6, which shows that all tilings by the hat polykite are aligned to such an underlying Laves tiling.

Contracting the polykite sides of length 1 or 2 to length 0 produces a strongly periodic tiling \mathcal{T}_4 by tetriamonds; contracting the polykite sides of length $\sqrt{3}$ to length 0 produces a strongly periodic tiling \mathcal{T}_8 by octiamonds. (Because this contraction process is well-defined around any tile, edge or vertex, it yields a combinatorial tiling of the plane, and a combinatorial tiling corresponds to a geometrical tiling of the entire plane [GS09, Lemma 1.1].) We suppose that \mathcal{T}_4 and \mathcal{T}_8 use the same side lengths as the corresponding sides of the polykite, and have corresponding tiles in the same orientation. Figure 3.1 shows a patch from an example tiling \mathcal{T} , together with corresponding patches from \mathcal{T}_4 and \mathcal{T}_8 . This mapping to tiles of different side lengths is discussed in more detail in Section 6, where the tetriamond is described as $\text{Tile}(0, \sqrt{3})$ and the octiamond is described as $\text{Tile}(1, 0)$.

The tilings \mathcal{T} , \mathcal{T}_4 , and \mathcal{T}_8 are coupled, in the sense that there is a bijection between their tiles, with corresponding tiles in corresponding orientations and translation symmetries of any one

mapping directly to translation symmetries of the others. They also have close combinatorial relationships: any neighbours in the original polykite tiling are also neighbours in both polyiamond tilings. Furthermore, there must exist affine maps between the lattices of translations of the \mathcal{T}_4 and \mathcal{T}_8 tilings, such that each translation maps to a corresponding translation (one between corresponding pairs of tiles). The affine map sending the translations for \mathcal{T}_4 to those for \mathcal{T}_8 must scale areas by $2/3$.

That affine map cannot be a similarity. Consider the regular tiling by equilateral triangles, positioned to include a unit edge from $(0, 0)$ to $(1, 0)$. Every vertex of this tiling is given by $a(1, 0) + b(1/2, \sqrt{3}/2)$ for integers a and b , meaning that vectors joining vertices must have this form as well. We can then deduce that any distance d between two vertices must have d^2 of the form $a^2 + ab + b^2$ for integers a and b , and if $2^k \parallel a^2 + ab + b^2$ then k must be even. Therefore, a scale factor of $\sqrt{2}$ is not possible between translations on two triangular tilings with the same edge length, and a scale factor of $\sqrt{2/3}$ is not possible between translations on two triangular tilings with edge lengths $\sqrt{3}$ and 1.

Using the fact that the six translation classes of kites must appear with equal frequency in any aligned tiling by polykites, we now proceed to show that the affine map from \mathcal{T}_4 to \mathcal{T}_8 must in fact be a similarity, which gives the required contradiction.

In Figure 3.2, we show the hat polykite, with sides of length 1 or 2 shown as thin **black** line segments and sides of length $\sqrt{3}$ shown as thick **orange** line segments; the two orientations of kite of which it has two kites (it has exactly one kite of each other orientation); and the two polyiamonds, in orientations corresponding to those of the polykite.



Figure 3.2: The hat polykite, its imbalance in kites, and corresponding polyiamonds

The imbalance in kites of different orientations partitions the twelve orientations of the hat polykite (that can occur while aligned to an underlying [3.4.6.4] Laves tiling) into three sets of four, such that any tiling must have equal proportions of polykites from each set of orientations (meaning that in any patch with perimeter x , the imbalance between the numbers of polykites with orientations from any two of the sets is $O(x)$). The same applies to tilings by either polyiamond derived from tilings by the polykite. (The tetriamond is symmetric, so two orientations of the polykite in one of those sets can give rise to identical-looking tetriamonds. Those should still be considered as different orientations of the tetriamond, as if it were given an asymmetric marking.)

Note that given any two sides of polyiamonds in \mathcal{T}_4 , the corresponding vector in \mathcal{T}_8 between those two sides is well-defined: a side of a tile in \mathcal{T}_4 corresponds to a point on the boundary of the corresponding tile in \mathcal{T}_8 (and adjoining sides on adjacent tiles correspond to the same point on the boundaries of two neighbouring tiles in \mathcal{T}_8), so the vector is just the vector between those




Figure 3.3: A visualization of the three families of i -strips in each of \mathcal{T}'_4 and \mathcal{T}_4 , for the sample patch shown in Figure 3.1 (centre). The i -strips run along the channels between the heavy black boundaries.

corresponding points. It is also convenient for some purposes to divide the tetriamond into two congruent rhombi, resulting in a tiling \mathcal{T}'_4 . The four orientations of the tetriamond in one of the three sets of orientations all contain rhombi in the same two orientations (out of a possible three), so \mathcal{T}'_4 contains equal proportions of rhombi in each orientation.

The edges of the equilateral triangles (of side length $\sqrt{3}$) in the regular triangular tiling underlying \mathcal{T}_4 and \mathcal{T}'_4 lie in three sets of parallel lines; call those sets \mathcal{L}_1 , \mathcal{L}_2 , and \mathcal{L}_3 . For each $i \in \{1, 2, 3\}$, we now form an infinite graph \mathcal{G}_i . The vertices of \mathcal{G}_i are those rhombi of \mathcal{T}'_4 that include an edge lying in some line in \mathcal{L}_i , thereby encompassing rhombi in two of the three orientations. Two vertices in \mathcal{G}_i are connected by an edge exactly when the rhombi corresponding to those vertices are adjacent along an edge of \mathcal{T}'_4 lying on some line in \mathcal{L}_i . The connected components of this graph are two-way infinite paths, which we refer to as i -strips in \mathcal{T}'_4 . We then partition the tiles of \mathcal{T}_4 into a related set of i -strips by assigning each polyiamond to the same strip as one of its rhombi. If a polyiamond has both rhombi in an orientation that appears in i -strips, they must be in the same i -strip, because the line segment between those two rhombi lies in a line in \mathcal{L}_i . Note that this assignment must constitute a partition of the tiles in \mathcal{T}_4 . Clearly, the i -strips cannot cross, and any line parallel to those in \mathcal{L}_i passes through the i -strips

in \mathcal{T}'_4 in the same order as any other such line passes through them. Furthermore, any translation preserves both i -strips themselves and that ordering of i -strips. Figure 3.3 shows examples of the three families of i -strips for sample patches of \mathcal{T}'_4 and \mathcal{T}_4 .

Let \mathbf{v}_i be a vector between two consecutive lines in \mathcal{L}_i , orthogonal to those lines, chosen so the pairwise angles between those vectors are all 120° . Let \mathbf{v}'_i be a vector orthogonal to \mathbf{v}_i and with length $1/\sqrt{3}$ times that of \mathbf{v}_i , again chosen so the pairwise angles between those vectors are all 120° . Note that $\sum_i \mathbf{v}_i = 0$ and $\sum_i \mathbf{v}'_i = 0$. Considering the sides of rhombi in an i -strip in \mathcal{T}'_4 that lie in consecutive lines of \mathcal{L}_i , the vector between the midpoints of such sides is $\mathbf{v}_i \pm \mathbf{v}'_i$, where the sign depends on the orientation of the rhombus. Thus, if the vector between any two such midpoints is $a\mathbf{v}_i + b\mathbf{v}'_i$, then between those two sides there are $(a+b)/2$ rhombi of one orientation and $(a-b)/2$ of the other orientation that can appear in an i -strip in \mathcal{T}'_4 .

Considering as symmetries of \mathcal{T}'_4 only those translations that correspond to translation symmetries of the strongly periodic tiling \mathcal{T} , there are only finitely many orbits of rhombi under the action of the group of such translations, so in any i -strip \mathcal{S} there must be two rhombi in the same orbit. The translation mapping one to the other is a translation symmetry of the tiling, and therefore maps i -strips to i -strips. Because it maps \mathcal{S} to itself and preserves the ordering of i -strips, it must map every i -strip to itself. If that translation is by a vector $a\mathbf{v}_i + b\mathbf{v}'_i$, it follows that $b = 0$, because otherwise rhombi of the two orientations that appear in i -strips would appear in the tiling in different proportions.

Thus for each i we have some positive integer a_i , such that a translation by $a_i\mathbf{v}_i$ is a symmetry of \mathcal{T}'_4 and of \mathcal{T}_4 (corresponding to a translation symmetry of \mathcal{T} , and thus to one of \mathcal{T}_8) that sends each i -strip to itself. We may replace each a_i by their lowest common multiple, so we have a positive integer a such that for each i a translation by $a\mathbf{v}_i$ is a symmetry that sends each i -strip to itself.

We now examine the translation vectors in \mathcal{T}_8 corresponding to $a\mathbf{v}_i$ in \mathcal{T}_4 or \mathcal{T}'_4 . These may be calculated based on the tiles in any i -strip in \mathcal{T}_4 (between any two lines in \mathcal{L}_i related by a translation by that vector), and every i -strip (and choice of lines) must produce the same vector in \mathcal{T}_8 . Figure 3.4 shows the corresponding translations for three vectors between opposite pairs of parallel sides of the tetriamond. For each such pair, first the vector within the tetriamond is indicated, then the corresponding octiamond vector as a sequence of octiamond sides, then that vector decomposed into parts parallel to and orthogonal to the sides between which the vector is drawn. Rotating, reflecting or reversing the direction of the tetriamond vector has the same effect on the octiamond vector.




Figure 3.4: Corresponding translations for the two polyiamonds

Note that in the first case, the octiamond vector is parallel to the sides between which the tetriamond vector is drawn; the second and third cases have equal components orthogonal to those sides. For the orthogonal component of the corresponding translations in \mathcal{T}_8 to be equal for all i -strips, it follows that every i -strip must have the same proportion of the second and third cases relative to the first case. As the first case corresponds exactly to one of the three sets of orientations that occur in equal proportions in any tiling, the first case must thus be a third of the tetriamonds in any i -strip, while the second and third cases (which together correspond to the other two sets of orientations; however, each case does not correspond to a single set of orientations) in that figure must add to two thirds of the tetriamonds.

Since that proportion is independent of i , the orthogonal components of the translation vectors in \mathcal{T}_8 corresponding to av_i (which are $\frac{a}{2}\mathbf{v}_i$) add to zero. Since the av_i add to zero, the vectors in \mathcal{T}_8 also add to zero, and so their parallel components must also add to zero. But $\sum_i b_i \mathbf{v}'_i = 0$ if and only if all the b_i are equal; say they all equal b . That means the three translation vectors in \mathcal{T}_8 (which are $\frac{a}{2}\mathbf{v}_i + b\mathbf{v}'_i$) are at 120° angles to each other, and so the affine map between lattices of translations is a similarity, which is a contradiction as discussed above.

4. Clustering of tiles

As discussed in Section 2, tilings by the hat polykite are composed of certain clusters of tiles. These clusters can be used to define simplified tile shapes that we call *metatiles*. The metatiles inherit matching conditions from boundaries of the hats that they contain. Furthermore, through a set of substitution rules they form larger, combinatorially equivalent supertiles that fit together following the same matching conditions. In this section, we give a precise definition of how tiles are assigned to clusters, and a computer-assisted proof by case analysis that this assignment does result in the clusters claimed, fitting together in accordance with the matching rules given.

The clusters and their associated metatiles are shown in Figure 4.1. Each metatile is a convex polyiamond outlined in lime; its hats are overlaid, and each is given a unique label. The union of the polykites in a cluster approximates the shape of its metatile, but with some indentations and protrusions along its boundary. At two corners of cluster T , and one of cluster P , an additional line is drawn from a corner of a polykite to a corner of the boundary of the polyiamond; this shows how an indentation to a corner of the polyiamond is considered to be associated with a particular side of the polyiamond.

The boundaries of the four metatiles are divided into labeled segments by marked points. The labels represent matching conditions to be obeyed in tilings by the metatiles. To satisfy the matching conditions, the four metatiles must form a tiling using copies that are only rotated and not reflected; edge segments marked A^+ and A^- must adjoin on adjacent tiles of the tiling; likewise, edge segments B^+ and B^- , X^+ and X^- , F^+ and F^- , and L and L must adjoin. We will show in Section 5 that any tiling by the metatiles has a substitution structure: the tiles may be grouped (after bisecting some tiles) into supertiles that satisfy combinatorially equivalent matching conditions. From this structure, it is deduced that no tiling by the metatiles is periodic; also, the substitution structure implies that the metatiles can tile arbitrarily large regions of the plane, and so the whole plane, implying that they form an aperiodic set.

We now show the following result:




Figure 4.1: The four clusters

Theorem 4.1. *Any tiling by the hat polykite can be divided into the clusters shown in Figure 4.1 (or reflections thereof, but not mixing reflected and non-reflected clusters), satisfying the given matching conditions, with the resulting tiling by metatiles having the same symmetries as the original tiling by polykites.*

Since inspection of the cluster shapes shows that, conversely, any tiling by metatiles induces one by the hat polykite (for example, A^+ and A^- are equal and opposite modifications to the shape of an edge and are consistent wherever they appear in the clusters), that suffices to show that the hat polykite is an aperiodic monotile.

This proof that any tiling can be divided into clusters, satisfying the matching conditions, is computer-assisted. We define rules (Section 4.1) for assigning labels to tiles in any tiling by the hat polykite, corresponding to the labels shown for the tiles in the four clusters. Those rules assign a label to a tile based only on its immediate neighbours, and because no arbitrary choices are involved in the rules, they preserve all symmetries of the tiling. It then remains to show that (a) the labels assigned do induce a division into the clusters shown, and (b) the clusters adjoin

other clusters in accordance with the matching rules.

Both (a) and (b) may be demonstrated by a computer case analysis of 2-patches of hats. This analysis is not very sensitive to the precise list of 2-patches used, as long as it includes all 2-patches that can actually occur in a tiling. It may also include some 2-patches that cannot occur in a tiling, since many such 2-patches do in fact also satisfy the conditions that need to be checked. For the purposes of this proof we worked with the 188 “surroundable 2-patches”, i.e., 2-patches that can be surrounded at least once more to form a 3-patch. To generate this set of 188 patches we modified Kaplan’s SAT-based software [Kap22] to enumerate all distinct 3-patches of hats, and extracted the unique 2-patches in their centres. We validated this list by creating an independent implementation based on brute-force search with backtracking; the source code for this implementation is available with our article. A more sophisticated case analysis can be carried out to show that at most 63 of the 188 surroundable 2-patches can actually appear in a tiling by hats. However, all 188 of them satisfy the conditions given in this section, allowing us to obtain the results we need with simpler and more transparent computations.

It is also possible to demonstrate both (a) and (b) by a shorter case analysis using only 1-patches. However, the analysis using only 1-patches is more complicated because the classification rules assume that all the neighbours of a tile are known. Those rules can therefore not be applied directly to the outer tiles in a 1-patch, making it necessary to work with partial information about which labels are consistent with such a tile. For more details of this alternative case analysis, see Appendix B.

The computer analyses of 2-patches and 1-patches depend on an assumption that it is only necessary to consider tilings where all polykites are aligned to the same underlying [3.4.6.4] Laves tiling. This assumption is not in fact obvious for tilings by polykites or other polyforms in general; it is justified in Appendix A.

To demonstrate (a), for each of the kinds of cluster, form a connected graph whose vertices are the tiles of that cluster, and such that, where the graph has an edge between two tiles, those tiles are neighbours in the cluster. (This does not uniquely determine the graph for cluster H ; below we use a spanning tree where H_1 has edges to each other tile.) For each ordered pair L_1 and L_2 of two labels in a cluster, with an edge between their tiles in the graph for that cluster, we check (Section 4.2) that, in every 2-patch centred on a tile with label L_1 , there is a tile with label L_2 in the expected position and orientation relative to that with label L_1 . If this relationship holds for all such pairs, for every 2-patch that might occur in a tiling and that has L_1 at its centre, it follows that the labels assigned in any tiling do indeed induce a division into clusters as desired.

In fact, the rules given do not distinguish between labels P_1 and F_1 ; they assign all such tiles the common label FP_1 . However, this does not affect the argument; it simply means that in a certain position and orientation relative to a tile FP_1 , there must be a tile that may be P_2 or F_2 , and, conversely, in a certain position and orientation relative to P_2 or F_2 , there must be a tile FP_1 . Given that this is true in all tilings, each FP_1 tile may then be relabeled as F_1 or P_1 based on whether its neighbour is F_2 or P_2 .

To demonstrate (b), for each marked edge segment (referred to here as E) of each kind of cluster (referred to here as C), we identify all the edge segments E' in all the kinds of cluster C' that could adjoin it, consistent with the matching rules. If any one tile in C' adjoining E' is in the correct position and orientation relative to any one tile in C that adjoins E , it follows as a result

of (a) that the entire edge segment properly matches between the two clusters. Furthermore, because the matching conditions on the boundaries of F_1 and P_1 are identical, it is not a problem that those are both handled as a single label FP_1 . So for each E we pick one tile in C , and for each choice of E' , we pick one tile in C' that would be a neighbour of the tile picked in C (Section 4.3). The computer analysis of 2-patches then verifies that, in each 2-patch whose central tile has the label of the tile picked in C , there is a neighbour in a position and orientation and with a label that matches the tile picked in C' for one choice of E' .

We now present details of the classification rules for tiles, and of the exact checks implemented that satisfy the requirements described above: the ordered pairs of tiles in a spanning tree for each cluster, and those associated with each edge segment of each cluster. The reference software implementation mentioned above also performs all of these checks on the 188 surroundable 2-patches.

4.1. Classification rules for the hat polykite

Each diagram in Figure 4.2 shows a central tile and some number of neighbours of that tile, and represents a classification rule. The order of the rules is significant; the first rule that matches determines the label on the central tile. If all the neighbours shown are present, and no previous rule matched, the tile has the label indicated. The last rule has no neighbours shown, so always matches if no previous rule matched. Thus every tile is assigned some label.




Figure 4.2: Classification rules (part 1)




Figure 4.2: Classification rules (part 2)

4.2. Within-cluster matching checks for the hat polykite

Each diagram in Figure 4.3 shows a central tile (shaded) and a neighbour, with labels on both. The central tile in every 2-patch that can occur in a tiling should be checked against all figures shown here with that central tile's label on the shaded tile; if, for all such 2-patches, the neighbour indicated is present with the correct label, then the labels assigned by the classification rules do induce a division into the clusters shown, as explained above.




Figure 4.3: Within-cluster matching checks (part 1)




Figure 4.3: Within-cluster matching checks (part 2)

4.3. Between-cluster matching checks for the hat polykite

Each diagram in Figure 4.4 shows a central tile (shaded) and a neighbour, with labels on both, and represents a tile on one side of a cluster edge and some options for a tile on the other side of that edge. In some cases, there are two alternatives listed for the same edge, with separate figures for each, marked in the form “(alternative k of 2)”. Also, in some cases there are multiple options for the labels on one or both tiles, shown in a single figure. The central tile in every 2-patch that can occur in a tiling should be checked against all figures shown here with that central tile’s label as one of the options for the shaded tile; if, for all such 2-patches, one of the alternatives listed for that edge is present with one of the labels indicated, then the clusters adjoin other clusters in accordance with the matching rules. (Where multiple alternatives are listed for the same edge, only one of those alternatives needs to pass the check.)




Figure 4.4: Between-cluster matching checks (part 1)




Figure 4.4: Between-cluster matching checks (part 2)

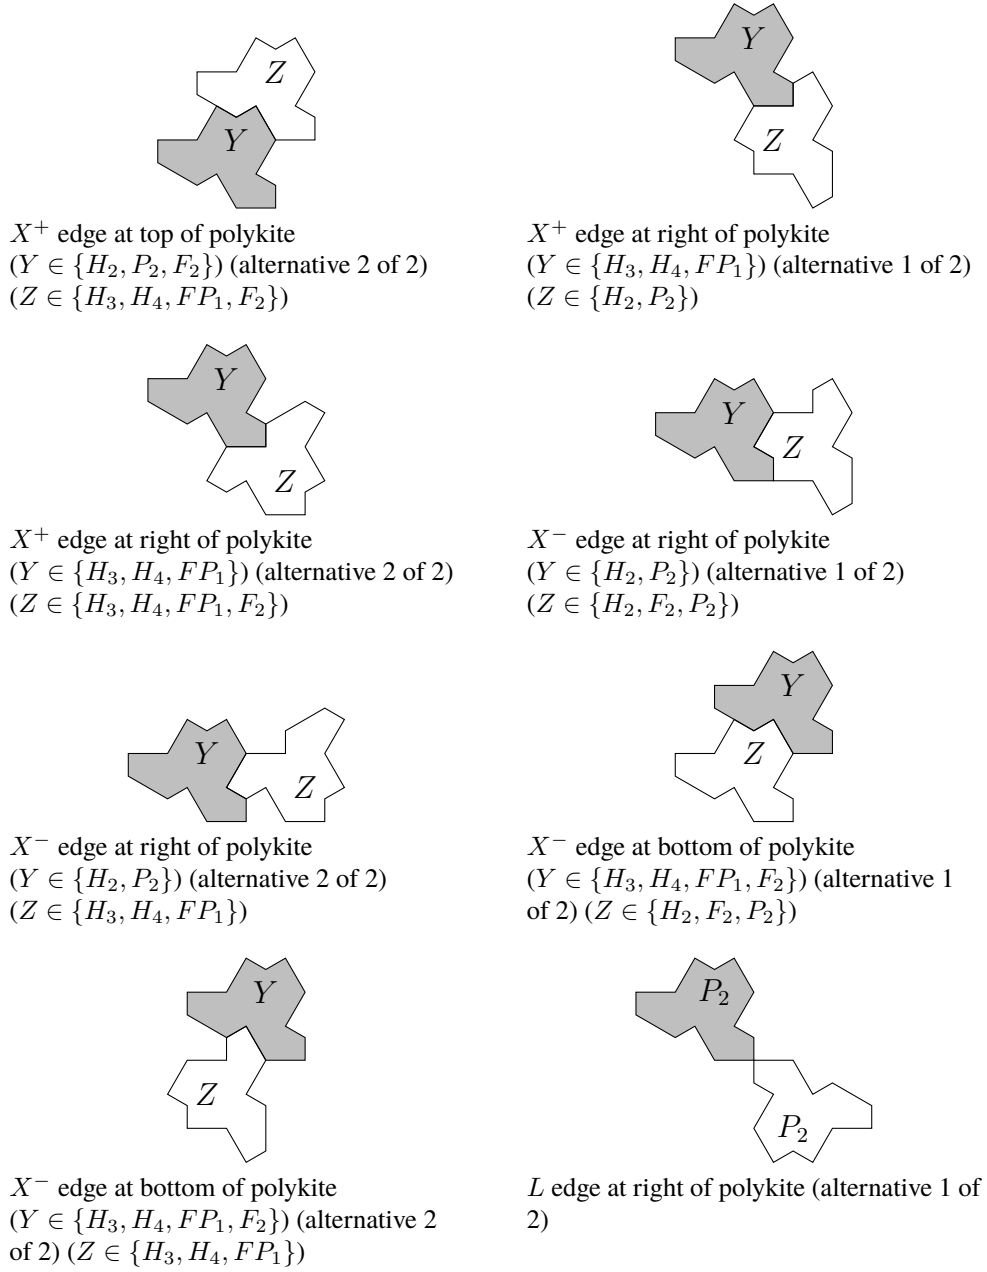


Figure 4.4: Between-cluster matching checks (part 3)





Figure 4.4: Between-cluster matching checks (part 4)

5. A four-tile substitution system

Consider the four metatiles, with matching conditions as in Figure 4.1, which are depicted in this section in the form shown in Figure 5.1. Edges A are **red**, B are **blue**, X are **green**, F are **pink**, and L are **gray**. Edges are marked with small geometrical decorations to indicate the signs (outward on the $^+$ side, inward on the $^-$ side): equilateral triangles for A , semicircles for B , orthogonal line segments for X , short slanted line segments for F . Note that the A and B on H are the opposite signs to those on T , P , and F . Also note that the tiles in this substitution system may not be reflected, only rotated.

Figure 5.1: Metatiles T , H , P , and F

Later in the argument it is convenient to bisect some tiles P and F , as shown in Figure 5.2. We refer to the edges resulting from the bisection of P as P^+ (in the sub-tile that has an edge B^+) and P^- , coloured **yellow** and decorated with a rectangle, and to the edges resulting from the bisection of F as G^+ (in the sub-tile that has an edge B^+) and G^- , coloured **violet** and decorated with an obtuse triangle. We also refer to the halves with a B^+ edge as the upper halves, and the other halves as the lower halves. We will show the following:

Theorem 5.1. *In any tiling by the four metatiles, after such bisection of P and F metatiles, the metatiles fit together to form larger, combinatorially equivalent supertiles, thereby forming a substitution system. The tiling by the supertiles has the same symmetries as the tiling by the metatiles.*




Figure 5.2: Bisection of tiles P and F

The bisection of tiles is not strictly necessary, in that the bisecting lines can be arbitrary curves—and, in particular, can go entirely along one side or other of the F or P tiles (keeping the same end points), effectively allocating an entire tile to one of two neighbouring supertiles. However, the bisected tiles are convenient for proving that the supertiles obey matching rules equivalent to those of the original tiles. In particular, bisection causes adjacencies between supertiles to be more clearly encoded in the boundaries of the supertiles themselves, without also relying on information about forced tiles that are not part of the supertiles. In some situations it may be more useful to assign whole tiles to supertiles at every level of substitution, with no bisection. For example, these whole tiles may be more convenient for analyzing sizes or growth rates of patches in the inflation process. If needed, we can define a symmetry-preserving bijection between the supertiles shown here and any alternative choice of supertiles that avoids bisection.

Diagrams for analysis of cases should be interpreted as follows. There are some unnumbered tiles that define the case being considered, then some numbered tiles that are forced in the sequence given by their numbers. If it is then necessary to split into multiple cases, the position at which multiple choices of tile must be considered is marked on the diagram with a filled circle, and there are then separate diagrams for each choice (on which the previous forced tiles are now unnumbered, but newly forced tiles are numbered).

The configuration in Figure 5.3, referred to as PP , often appears in the case analysis; the two adjoining copies of P in the same orientation force a contradiction because nothing fits at the marked point. Subsequently, when identifying forced tiles, as well as considering a tile as forced when it is the only one that would fit in a given place consistent with the matching conditions, we also consider a tile as forced when the only alternative consistent with the matching conditions would be to place a P tile in a way that produces this PP configuration.

5.1. Cases involving T

The two A^- edges of T must be adjacent to the A^+ edge of H , while the B^+ edge of T may be adjacent to either of the B^- edges of H . Thus we have two cases for the configuration around a T tile, which we refer to as $T1$ and $T2$ (Figure 5.4). As explained in the captions to this and subsequent figures, a sequence of deductions shows that any T in a tiling must occur in case $T1PF$ (Figure 5.9).




Figure 5.3: A common impossible configuration, referred to as PP




Figure 5.4: Cases $T1$ and $T2$. Consider the three marked places in each of $T1$ and $T2$. These can be filled with either P or H . On a side of the figure where there are two B^- edges, the marked place cannot be filled with H , because that would result in a 60° angle between two B^- edges, which cannot be filled. So both those sides must have P in the marked place, while the third side may have H (oriented to avoid such a 60° angle between two B^- edges; subsequently, when the same situation arises, we just consider the orientation of the H to be forced without further comment) or P . This results in four cases, which we call $T1P$ (Figure 5.5), $T2P$ (Figure 5.6), $T1H$ (Figure 5.7) and $T2H$ (Figure 5.8), and we proceed to draw further forced tiles in each of those cases.




Figure 5.5: Case $T1P$. The marked place can be filled with F or P , resulting in cases we call $T1PF$ (Figure 5.9) and $T1PP$ (Figure 5.10).




Figure 5.6: Case $T2P$, eliminated because PP occurs at the marked point.




Figure 5.7: Case $T1H$, eliminated because PP occurs at the marked point.




Figure 5.8: Case $T2H$, eliminated because PP occurs at the marked point.




Figure 5.9: Case $T1PF$. Any T in a tiling must occur in this case. Bisecting the P and F tiles in that case produces the configuration of Figure 5.16, which we call H' and which combinatorially acts like H (with the edge segments indicated marked for matching conditions) in a tiling along with configurations T' , P' , and F' .




Figure 5.10: Case $T1PP$, eliminated because PP occurs at the marked point.

5.2. Cases with H not adjacent to T

Any H not adjacent to a T tile must have a P tile adjacent to its A^+ edge, while the B^- edges may each be adjacent to P or F . This results in four cases, which we call HPP (Figure 5.11), HPF (Figure 5.12), HFP (Figure 5.13), and HFF (Figure 5.14), and we proceed to draw further forced tiles in each of those cases, with consequences explained in the captions to those figures.




Figure 5.11: Case HPP . Bisecting the P tiles and removing the forced F tiles produces the configuration of Figure 5.15, which we call T' and which combinatorially acts like T (with the edge segments indicated marked for matching conditions) in a tiling with the other supertiles. Although the forced F tiles are not included in T' , the fact that they are forced will be used in the proof that the supertiles must follow the matching conditions where they are adjacent to each other.




Figure 5.12: Case HPF . The second H cannot be adjacent to a T ; it must thus itself be in case HFP or HFF , and each of those cases turns out to force the H tile in case HPF .




Figure 5.13: Case HFP . Bisecting the P and F tiles produces the configuration of Figure 5.17, which we call P' and which combinatorially acts like P (with the edge segments indicated marked for matching conditions) in a tiling with the other supertiles.




Figure 5.14: Case HFF . Bisecting the P and F tiles produces the configuration of Figure 5.18, which we call F' and which combinatorially acts like F (with the edge segments indicated marked for matching conditions) in a tiling with the other supertiles.

5.3. The supertiles




Figure 5.15: Supertile T' , alongside corresponding T

The previous arguments have shown that every H or T tile appears in a configuration corresponding to the supertiles T' , H' , P' or T' . We now provide more detailed rules for allocating each H or T tile, and, after bisecting all P and F tiles, each half of such a tile, to groupings of tiles, such that each tile is allocated to exactly one grouping, the groupings all have the form of one of the supertiles, all symmetries of the original tiling are also symmetries of the tiling by supertiles (this property follows immediately from the form of the rules, which do not involve any arbitrary choices that could break symmetry), and the supertiles adjoin each other in accordance with the matching conditions indicated (A_2^+ adjoining A_2^- , B_2^+ adjoining B_2^- , X_2^+ adjoining X_2^- , F_2^+ adjoining F_2^- , and L_2 adjoining L_2).

- Each T tile is allocated to a H' supertile, along with all the H tiles adjacent to that T .
- Each H tile in case HPP is allocated to a T' supertile.
- Each H tile in case HFP is allocated to a P' supertile, along with the H tile in case HPF shown in Figure 5.13.




Figure 5.16: Supertile H' , alongside corresponding H




Figure 5.17: Supertile P' , alongside corresponding P




Figure 5.18: Supertile F' , alongside corresponding F

- Each H tile in case HFF is allocated to an F' supertile, along with the H tile in case HPP shown in Figure 5.14.
- Each H tile in case HPP was allocated to a supertile by exactly one of the previous two rules.
- Each half of a P tile, and the upper half of each F tile, is adjacent to exactly one H tile along its A^- or B^+ edge, and is allocated to the same supertile as that H tile. (Where both halves are allocated to the same supertile, the bisection isn't shown in the diagrams of the supertiles, to simplify those diagrams.)
- It remains to allocate the lower halves of F tiles. Each such lower half has a X^- edge between a L edge and a F^+ edge; it is allocated to the same supertile as the H tile adjacent to that X^- edge. For this allocation rule to be well defined, we need to show that this X^- edge is indeed adjacent to a H tile. The only other possibility not violating the matching rules would be adjacency to another F tile, as shown in Figure 5.19, but nothing fits at the marked point without violating the matching rules, so that cannot occur.

The X^- edge referenced in the last allocation rule cannot be adjacent to any of the exposed X^+ edges of H tiles in supertiles T' , P' or F' without violating the matching rules. Thus all lower halves of F tiles are the ones that appear on the diagrams of the supertiles, and we have shown that the tiling is partitioned into the supertiles.




Figure 5.19: Impossible adjacency of two F tiles

We now show that the supertiles adjoin each other in accordance with the matching conditions indicated. First, we examine P^+ edges (appearing in A_2^- and B_2^-) and P^- edges (appearing in A_2^+ and B_2^+). B_2^- and A_2^+ appear only in H' , where their P^+ and P^- edges cannot meet without tiles intersecting. So B_2^- can only join to B_2^+ and A_2^+ can only join to A_2^- .

Next we show that the converse holds: A_2^- can only join to A_2^+ and B_2^+ can only join to B_2^- . For a contradiction, suppose that the P^+ and P^- edges in some A_2^- and B_2^+ are joined. If the B_2^+ comes from a P' supertile, then that P^- bisects tile 5 in case HFP . Adjacent tiles 5 and 1 in that configuration both have B^+ edges, which must both be adjacent to H tiles; those H tiles are adjacent to each other, which can only occur in a supertile H' , which does not have an A_2^- edge. The same argument applies in the case of an F' supertile, considering tiles 6 and 2 in case HFF . If the P^+ in an A_2^- from P' is joined to a B_2^+ , a similar argument applies (considering tile 2 and an adjacent unnumbered tile in case HFP). So the only remaining case would be if both edges come from supertile T' , but that is inconsistent with the F tiles forced in case HPP .

Having shown that A_2^+ , A_2^- , B_2^+ , and B_2^- must meet the matching conditions for the supertiles, now note that the only X^+ and X^- edges on the boundaries of the supertiles that are not part of A_2^+ , A_2^- , B_2^+ or B_2^- are those forming part of F_2^+ and F_2^- . Thus it follows that F_2^+ and F_2^- must also adjoin each other. The only remaining G^+ and G^- edges of the tiles are then those that form X_2^- and X_2^+ edges of the supertiles, so those also match, and the remaining L edges form L_2 , so those also match.

To show that the supertiles are fully combinatorially equivalent to the original tiles, one more thing must be checked: that the same combinations of supertiles fit together at vertices as combinations of tiles fit together at vertices. Each supertile has been drawn with a copy of the corresponding tile alongside it, in a corresponding orientation. By inspection, if you take any class of edges of the tiles, including both sides of the edge (for example, A^+ and A^-), and take any line segment in the corresponding edges of the supertiles, the (directed) angle between the (directed) edges in the tile and in the supertile is consistent across all the diagrams.

This consistency of angles between edges of tiles and of supertiles means that the angles at vertices of supertiles around a point, each consecutive pair having matching edges, add up to the same amount as the corresponding angles for the corresponding tiles (an angle at a vertex of a supertile equals the angle at the corresponding vertex of the corresponding tile, plus the difference between the tile–supertile angles for the two edges, and those differences cancel when adding up around the point).

So the supertiles are combinatorially equivalent to the tiles, and thus the above arguments apply inductively to ensure that the composition of tiles into supertiles may be applied n times for all n . Since the radius of a ball contained in the supertiles goes to infinity with n (and this does not depend on how the bisection of supertiles is defined geometrically, although it may be easier to see if an alternative choice of supertiles is made at this stage that does not involve bisecting tiles), and the tilings by supertiles have all the symmetries of the original tiling, it follows that the original tiling cannot have a translation as a symmetry. Furthermore, the substitution structure implies that the metatiles tile arbitrarily large finite regions of the plane, and hence the whole plane.

6. A family of aperiodic monotiles

In the previous sections, we showed that the hat polykite is an aperiodic monotile. This polykite is formed of 8 kites from the [3.4.6.4] Laves tiling. Another small polykite, formed of 10 kites and shown in Figure 6.1, is also aperiodic. We have verified via a computer search that there are no other aperiodic n -kites for $n \leq 21$.

These two aperiodic polykites are two examples of a family of aperiodic monotiles, all of which have combinatorially equivalent sets of tilings, and which are determined by the choice of two side lengths.

The hat polykite has sides of lengths 1, 2, and $\sqrt{3}$; for the purposes of this section, we consider the side of length 2 as two consecutive sides of length 1 with a 180° angle between them. The tile of Figure 6.1 has the same angles, but with the side lengths of 1 and $\sqrt{3}$ swapped.

Let a and b be nonnegative reals, not both zero, and if $a \neq 0$ write $r = b/a$. Define $\text{Tile}(a, b)$ to be the polygon resulting from replacing the sides of length 1 in the hat polykite with sides of




Figure 6.1: An aperiodic 10-kite

length a (we refer to the resulting sides as 1 -sides) and replacing the sides of length $\sqrt{3}$ in the hat polykite with sides of length b (we refer to the resulting sides as r -sides). This process results in a closed curve (because the vectors of the 1 -sides add up to 0, as do those of the r -sides), and it is also straightforward to see that this curve is not self-intersecting; it is a 13-gon (or one with a smaller number of sides if a or b is zero), but considered as a 14-gon for the purposes of this section. This tile has area $\sqrt{3}(2a^2 + \sqrt{3}ab + b^2)$. The hat polykite is then $\text{Tile}(1, \sqrt{3})$ in this notation, while the tile of Figure 6.1 is $\text{Tile}(\sqrt{3}, 1)$.

For nonzero a , the value of r determines the tile up to similarity. In acknowledgment of these similarity classes, we write $\text{Tile}(r)$ as a shorthand for $\text{Tile}(1, r)$. We will show this tile is aperiodic for any positive $r \neq 1$. In fact, $\text{Tile}(1, k\sqrt{3})$ and $\text{Tile}(k\sqrt{3}, 1)$ are polykites for all odd positive integers k , so this continuum of aperiodic monotiles also contains a countably infinite family of aperiodic polykites.




Figure 6.2: Periodic tiling by $\text{Tile}(1, 1)$

We define a notion of combinatorial equivalence between tilings of these tiles, for two positive values of r , as follows: two tilings are combinatorially equivalent if there exists a bijection between their tiles, and a bijection between the maximal line segments in the unions of the boundaries of the tiles, such that corresponding tiles and line segments in the two tilings are in the same orientation, corresponding tiles adjoin corresponding line segments, on the same side of those line segments, in the two tilings, and corresponding tiles on the corresponding sides

of corresponding line segments appear in the same order along those segments. All the interior angles of the tile are at least 90° , and no two 90° angles appear consecutively, so any maximal line segment has at most two sides of tiles on each side of the line segment (and in particular is finite).

We now prove the following result:

Theorem 6.1. *Suppose $r \neq 1$ and $r' \neq 1$ are positive. Then there is a bijection between combinatorially equivalent tilings for $\text{Tile}(r)$ and $\text{Tile}(r')$, given by changing the lengths of all r -sides from r to r' , while preserving angles, orientations and adjacency to maximal line segments.*

Suppose first that r is irrational. If a maximal line segment in the union of the boundaries of the tiles has p 1-sides and q r -sides on one side of the line segment, it also has p 1-sides and q r -sides on the other side of the line segment. Because a maximal line segment has at most two sides of tiles on each side of the segment, the same argument also applies for any rational r except possibly $\frac{1}{2}$, 1, and 2.

If $r = 2$, there is the additional possibility that two 1-sides align with one r -side. When there are two consecutive 1-sides on one side of a line, with 90° corners of the two tiles between those two sides (or the 180° corner of a single tile), the other ends of those sides have corners with angles 120° or 240° . But for every r -side, one corner has angle 90° or 270° , and the angles of the tile do not permit 120° or 240° at the same vertex of a tiling as 90° or 270° .

Similarly, in the case $r = \frac{1}{2}$, the only additional possibility is that two r -sides align with one 1-side. The outer corners of the two r -sides have angles 120° or 240° , one corner of every 1-side has angle 90° , 180° or 270° , and those cannot appear at the same vertex.

Thus for any positive $r \neq 1$, we have shown that if a maximal line segment in the union of the boundaries of the tiles has p 1-sides and q r -sides on one side of the line segment, it also has p 1-sides and q r -sides on the other side of the line segment. We can now construct the required bijection. Because side vectors around any tile add up to zero, and the sides of tiles on both sides of a maximal line segment add up to the same length, the specified process converts a tiling by $\text{Tile}(r)$ into one by $\text{Tile}(r')$ that is combinatorially equivalent [GS09, Lemma 1.1].

(This argument relies on the fact that the plane is simply connected; a tiling by $\text{Tile}(r)$ of a region with a hole that cannot be filled with tiles might not convert to a tiling by $\text{Tile}(r')$ of a region with a combinatorially equivalent hole, and indeed the side vectors of the hole might be such that no combinatorially equivalent hole exists, if the vectors of the 1-sides among the sides of the hole do not add up to 0.)

As shown in Lemma A.6, all tilings by $\text{Tile}(\sqrt{3})$ are aligned to an underlying [3.4.6.4] Laves tiling, so in fact each maximal line segment is made up only of 1-sides or only of r -sides.

Finally, $\text{Tile}(1)$, or more generally $\text{Tile}(a, a)$, is not aperiodic, as shown by the periodic tiling in Figure 6.2. The polyiamonds $\text{Tile}(a, 0)$ and $\text{Tile}(0, b)$ are also not aperiodic. A tiling by $\text{Tile}(r)$ for positive $r \neq 1$ can still be mapped to a corresponding tiling by $\text{Tile}(a, a)$, $\text{Tile}(a, 0)$ or $\text{Tile}(0, b)$ following the process described above, but the map is not a bijection.

7. Conclusion

We have exhibited an einstein, the first topological disk that tiles aperiodically with no additional constraints or matching conditions. The hat polykite is in fact a member of a continuous family of aperiodic monotiles that admit combinatorially equivalent tilings. The hat forces tilings with hierarchical structure, as is the case for many aperiodic sets of tiles in the plane, but a new method introduced in Section 3 also suffices to show the lack of periodic tilings without needing that hierarchical structure, beyond demonstrating the existence of a tiling.

Our substitution system satisfies the relatively mild conditions needed to guarantee an uncountable infinity of combinatorially distinct tilings, all of which are hierarchical [Sen96, Section 7.6.2]. But not every tiling by hats is necessarily produced purely through substitution. As with Robinson’s aperiodic set of six shapes [Rob71], it is conceivable that hats could tile infinite sectors of the plane, which could then be combined into tilings with infinite “fault lines” that lie on the boundary of supertiles at all levels. Future work should examine the possibility of tilings with fault lines, as part of characterizing the full space of hat tilings. In particular, it should be determined whether every finite patch that appears in some hat tiling must appear infinitely often in all hat tilings, or whether there are patches that only appear on fault lines and not in the interior of a supertile.

The hat is a 13-sided non-convex polygon. A convex polygon cannot be an aperiodic monotile, and all non-convex quadrilaterals can easily be seen to tile periodically. Therefore, in terms of number of sides, the “simplest” aperiodic n -gon must have $5 \leq n \leq 13$. Subsequent research could chip away at this range, by finding aperiodic n -gons for $n < 13$ or ruling them out for $n \geq 5$.

Tilings by the hat necessarily include both reflected and unreflected tiles. We might therefore ask whether there exists an aperiodic monotile for which reflections are not needed, either because the tile has bilateral symmetry or because it covers the plane using only translations and rotations. We conjecture that there exists a reflection-free einstein.

Finding such a monotile pushes the boundaries of complexity known to be achievable by the tiling behaviour of a single closed topological disk. It does not, however, settle various other unresolved questions about that complexity. For example, all of the following questions remain open.

- Are Heesch numbers unbounded? That is, does there exist, for every positive integer n , a topological disk that does not tile the plane and has Heesch number at least n ? We conjecture that there is no bound on Heesch numbers.
- Are isohedral numbers unbounded? That is, does there exist, for every positive integer n , a topological disk that tiles the plane periodically but only admits tilings with at least n transitivity classes? Again, we conjecture that no bound exists. If the requirement of periodicity is omitted here, then the hat polykite requires infinitely many transitivity classes in any tiling. Socolar [Soc07] showed that if the tile is not required to be a closed topological disk, then tiles exist with every positive isohedral number.
- Is it undecidable whether a single closed topological disk (or indeed a more general single tile in the plane) admits a tiling? It would again be reasonable to conjecture yes, which

would also imply unbounded Heesch numbers. Greenfeld and Tao [GT21] demonstrated undecidability in a more general context. For sets of tiles in the plane, Ollinger [Oll09] proved undecidability for sets of five polyominoes.

- Is it undecidable whether a single closed topological disk (or indeed a more general single tile in the plane) admits a periodic tiling? It would again be reasonable to conjecture yes. Such an answer would imply unbounded isohedral numbers.

Although we have provided a description of tilings by the hat polykite and related tiles described here (all such tilings are given by the substitution system of Section 5, as applied to the clusters of tiles from Section 4, subject to the possibility mentioned above of tilings with fault lines, where each sector is produced by the substitution system), there are various informal observations in Section 2 that have not been fully explored or given a precise statement. Those observations could provide starting points for possible future investigation of the tiles described here and their tilings, the metatiles used in classifying tilings by the hat polykite, and other related substitution tilings. It is not clear which ideas from this work will be most promising for future work, so we have generally erred on the side of including observations that might be of use, rather than making the paper focus more narrowly on a single proof of a single main result.

We believe the approach of Section 3, of coupling two separate tilings to show that a third tiling cannot be periodic, is a new approach for proving such a result about plane tilings. It would be worth investigating whether it can be applied to other tiles and tilings. In particular, polykites (and, more generally, poly-[4.6.12]-tiles, a subset of the shapes known as *polydrafters*) may be unusually well-suited to this method of proof, because their edges can naturally be split into those falling in two sets of lines, with each set of lines forming a regular triangular tiling. It might also be applicable to some poly-[4.8.8]-tiles (a subset of the *polyaboloes*).⁴ This style of proof might help explain how small polykites proved to be aperiodic when polyominoes, polyiamonds and polyhexes up to high orders yielded no aperiodic monotiles. However, as noted in Section 6, searches of polykites have not found other aperiodic examples outside the family described in this paper.

Acknowledgements

Thanks to Ava Pun for her work on software for computing Heesch numbers and displaying patches. Thanks also to Jaap Scherphuis for creating a number of useful free software tools for exploring tilings. These tools all played a crucial role in the explorations that led to the discovery of our polykite.

A. Aligned and unaligned tilings of polyforms

The proof that the hat polykite is aperiodic involves a case analysis for ways of surrounding a copy of that tile, and that case analysis in turn involves considering possibilities for how an individual

⁴Note that if Lemma A.3 is applied to poly-[4.6.12]-tiles or poly-[4.8.8]-tiles, the conclusion is weaker than that of Lemma A.5 for polykites, so tilings may need to be considered that are only aligned in such a weaker sense.

kite in a copy of the hat polykite could fill a particular kite on an underlying [3.4.6.4] Laves tiling. By itself a proof founded on such a case analysis shows the absence of periodic tilings only when all tiles are *aligned* to the same underlying [3.4.6.4] Laves tiling. This argument leaves open the possibility that polykites might be able to tile periodically if they may be translated, rotated and reflected without regard to the underlying grid. Here we provide the necessary justification that it is sufficient to consider aligned tilings: for classes of polyforms that include polykites, if a tile admits a periodic tiling then it also admits one that is aligned (in a suitable precisely defined sense); the same is also shown for some other tiling properties such as “isohedral” in place of “periodic”, although for the purposes of the main result of the paper, this fact is needed only for the “periodic” case (in which case the result and proof are slightly simpler). We further show that the hat in particular does not admit any unaligned tilings.

In principle, the same issue arises for polyominoes and polyiamonds. However, the only unaligned tilings by congruent squares consist of parallel rows of squares slid relative to each other (Figure A.1), and much the same applies to tilings by congruent equilateral triangles (Figure A.2), and so it is clear that no interesting examples of unaligned tilings by polyominoes or polyiamonds can arise. However, kites can form nontrivial unaligned tilings such as that of Figure A.3, and so there is genuinely something to be proved here, less obvious than it is for polyominoes and polyiamonds.




Figure A.1: Sliding rows of squares




Figure A.2: Sliding columns of equilateral triangles




Figure A.3: Unaligned tiling of kites

In order to apply the results presented here to other classes of polyforms such as polyaboloes, we state the required conditions on the tiles in fairly general and technical form.

Let S be a set of real numbers that are linearly independent over \mathbb{Q} . Let \mathcal{P} be a finite set of closed topological disk polygonal tiles, such that all the angles of corners of tiles in \mathcal{P} are rational sub-multiples of π , all the lengths of sides of tiles in \mathcal{P} are integer multiples of elements of S , and such that, if a polygon in \mathcal{P} has two or more collinear sides, the lengths of those sides are integer multiples of the same element of S , as are the distances between their endpoints.

We now consider clusters of tiles built using copies of the polygons in \mathcal{P} . Let \mathcal{Q} be a nonempty set of tiles, each one congruent to one of the polygons in \mathcal{P} , with disjoint interiors. The set \mathcal{Q} may cover the entire plane, or just part of it. The union of the boundaries of the tiles in \mathcal{Q} decomposes into a set of maximal line segments, rays, and infinite lines, which we will refer to generically as *segments*. These segments are maximal in the sense that no segment is a subset of a longer line contained in the union of the tile boundaries.

Given one such maximal segment ℓ , and two tiles $A, B \in \mathcal{Q}$ (which may be identical), we say that A and B are ℓ -aligned if they both have sides that are subsets of ℓ , all sides of A or B that lie in ℓ have lengths that are integer multiples of the same $s \in S$, and the distance between any endpoint of one of those sides that lies in ℓ and any other such endpoint is also an integer multiple of s .

The set \mathcal{Q} naturally induces a graph whose vertices correspond to the tiles in the set. Two tiles A and B are connected by an edge in the graph if there is a maximal segment ℓ such that A and B are ℓ -aligned and intersect in a line segment of positive length that is a subset of ℓ . We say that \mathcal{Q} is *weakly aligned* if this graph is connected. We say that it is *strongly aligned* if it is weakly aligned and, for every maximal segment ℓ determined by \mathcal{Q} , and all tiles A and B that have sides lying in ℓ , A and B are ℓ -aligned. We say that \mathcal{P} has the *alignment property for side lengths S* if every weakly aligned set is strongly aligned. Here we drop the qualifiers “strongly”

and “weakly” and refer to \mathcal{Q} , given the combination of \mathcal{P} and S , simply as *aligned*.

Lemma A.1. *Any finite set of polykites, where the underlying kites have side lengths 1 and $\sqrt{3}$, has the alignment property for side lengths $\{1, \sqrt{3}\}$.*

Proof. In the Laves tiling [3.4.6.4], subdivide each kite into $30^\circ-60^\circ-90^\circ$ triangles as shown in Figure A.4, forming a [4.6.12] Laves tiling. Furthermore, if a kite congruent to one of those from the original [3.4.6.4] adjoins edge-to-edge a kite that is a union of triangles from that [4.6.12] tiling, then it too is such a union. Thus all polykites in any weakly aligned set are unions of tiles from the same [4.6.12] tiling. On any line in the union of the boundaries of the tiles from [4.6.12] that contains sides with rational length, sides of such kites can only be at integer offsets from each other, and on the other lines (containing sides with length a rational multiple of $\sqrt{3}$), sides of such kites can only be at offsets from each other that are integer multiples of $\sqrt{3}$ (both of these facts follow from consideration of which vertices have the correct angles to form a corner of such a kite). So every weakly aligned set is strongly aligned. \square




Figure A.4: Decomposition of a kite into 24 triangles

We now consider \mathcal{P} that has the alignment property for side lengths S , and proceed to showing that, in an appropriate sense, only aligned tilings need to be considered. Note that for polykites, at this point “aligned” means only that the kites adjoin edge-to-edge, which is weaker than all tiles coming from the same underlying [3.4.6.4] tiling; there will be further lemmas specific to polykites to show that we need only consider tilings where all tiles come from the same underlying [3.4.6.4].

The alignment property implies that the tiles of any tiling can be partitioned into strongly aligned sets such that for any maximal line segment ℓ in the union of the boundaries of the tiles, and any two tiles in different sets that have sides sharing a segment of positive length lying on ℓ , those two tiles are not ℓ -aligned. We refer to these as the *aligned components* of the tiling. Each aligned component is a connected set (possibly unbounded), with connected interior.

Suppose \mathcal{C} is an aligned component in a tiling, and \mathcal{D} is a connected component of the complement of \mathcal{C} (\mathcal{D} might be the interior of another aligned component, or might be the interior of the union of more than one aligned component). The boundary of \mathcal{D} consists of a single polygonal curve, either closed or infinite, and as with any other polygon we may speak of its

corners and sides. Furthermore, that curve cannot pass through the same point more than once; if it did, either \mathcal{D} (an open set) would not be connected, or \mathcal{C} would not have connected interior.

Consider traversing the boundary of \mathcal{D} ; note that \mathcal{D} always lies on the same side of the boundary during that traversal. When the traversal encounters a corner, say that corner is convex if an open line segment between two points on the curve sufficiently close to that corner but on opposite sides of it is entirely within \mathcal{D} .

When the boundary of \mathcal{D} is a closed curve, we must also initially allow for \mathcal{D} being inside that curve (a hole in \mathcal{C}) or outside (in which case \mathcal{C} is bounded). The following lemma shows that the first of those cases cannot occur, since if \mathcal{D} is inside the curve it must have at least three convex corners.

Lemma A.2. *The boundary of \mathcal{D} has no convex corners if it is a closed curve (so in that case \mathcal{C} must be a bounded convex set), and at most one convex corner if it is an infinite curve.*

Proof. If we consider any finite side of the boundary of \mathcal{D} , lying in some maximal line segment ℓ , all the tiles lying on the other side of the side from \mathcal{D} are not ℓ -aligned with any of those in \mathcal{D} , meaning that at least one of the two corners at the ends of that side lies in the middle of a side of such a tile. If v is a convex corner, and v_1, v_2, \dots are successive vertices traversing the boundary curve in one direction from v , then we conclude that v_1 lies in the middle of a side of a tile on the same line as vv_1 , then that v_2 lies in the middle of a side of a tile on the same line as v_1v_2 , and so on. This results in a contradiction if we encounter another convex corner (see Figure A.5 for an illustration), or encounter v again on a closed curve (see Figure A.6). \square




Figure A.5: Boundary curve with two convex corners




Figure A.6: Closed boundary curve with a convex corner

Now we need to list the tiling properties for which our argument says we do not need to consider unaligned tilings. Let H be one of the following predicates on a tiling \mathcal{T} ; here, k may be any positive integer.

- \mathcal{T} is a tiling (the trivial predicate).
- \mathcal{T} is a strongly periodic tiling.
- \mathcal{T} is a weakly periodic tiling.
- \mathcal{T} is a tiling with at most k orbits of tiles under the action of its symmetry group.
- \mathcal{T} is an isohedral tiling by 180° rotation.
- \mathcal{T} is an isohedral tiling by translation.

Lemma A.3. *If \mathcal{P} admits a tiling with property H , it admits an aligned tiling with property H .*

Proof. Consider a tiling \mathcal{T} with property H and look at the forms aligned components take in that tiling. By the previous lemma, such components must be simply connected; either bounded, or unbounded and with each boundary curve having at most one convex corner (in the sense defined above, i.e., convex considered as a corner of a connected component of the complement of the aligned component).

If such a component is the whole plane, the tiling is aligned and we are done. If it is a half-plane, form an aligned tiling of that component and its reflection, and that tiling has property H (which can only be the trivial property or “weakly periodic” in that case). If it is a strip infinite in both directions, with straight lines as its boundaries on both sides, repeat that strip by translation if there is such a tiling that is aligned, and otherwise repeat it by 180° rotation; by considering separately for each possible predicate H , the resulting aligned tiling has property H .

Otherwise, if there is any unbounded component, it does not have a translation as a symmetry and H is the trivial predicate. If an unbounded component contains balls of radius R for all R , there are aligned tilings of arbitrarily large regions of the plane, and so of the whole plane. The only way an unbounded component can avoid containing such balls (given that each boundary curve has at most one convex corner and all angles are rational sub-multiples of π , which implies that all boundary curves end in rays in finitely many directions) is for it to include a semi-infinite strip (bounded on either side by rays). But there are only finitely many ways for aligned tiles to cross the width of the strip at any point, so tiling a semi-infinite strip implies the existence of a periodic aligned tiling of an infinite strip, and thus a periodic aligned tiling of the whole plane.

It remains to consider the case where there is no unbounded component. If some component is a triangle or a quadrilateral, tiling that by 180° rotation yields an aligned tiling of the whole plane, which must have property H (if property H is ‘isohedral tiling by translation’, this case cannot occur; a component could be an infinite strip, but not a single parallelogram). If components contain unbounded balls, the tiling has no translation as a symmetry and aligned tilings of arbitrarily large regions of the plane imply aligned tilings of the whole plane. If components do not contain unbounded balls but also are not contained in bounded balls (i.e., they are of unbounded size in one direction only), they must have pairs of opposite parallel sides, of unbounded length but a bounded distance apart; the tiling is at most weakly periodic, and the same argument as for components including a semi-infinite strip applies since there are only finitely many possible distances between those opposite parallel sides.

Otherwise, all components are convex polygons of bounded size with at least five sides; we will show this case leads to a contradiction. Observe that every vertex of the induced tiling by these polygons lies in the middle of a side of one of the polygons and has degree exactly 3. If a vertex does not lie in the middle of a side, or has degree 4 or more, there is a vertex v of a polygon P , either not in the middle of a side or in the middle of a side that is not collinear with either of the sides vv_1 and vv_2 of P next to v , and the same argument that excluded convex corners on a closed curve earlier serves to exclude this possibility as well.

We now apply Euler's theorem for plane maps. Suppose, for some sufficiently large R , a ball of radius R contains t_k components that are k -gons (where $\sum_k t_k = \Omega(R^2)$). A vertex of the tiling is incident with two corners of tiles and a point in the middle of a side, so there are $\sum_k kt_k/2 + O(R)$ vertices in that ball. The number of sides of edges in the tiling in that ball (i.e., twice the number of edges) is $\sum_k \frac{3}{2}kt_k + O(R)$, since there are $\sum_k kt_k + O(R)$ sides of polygons, and each vertex is in the middle of a side so serves to increase the number of sides of edges by 1. But now $\sum_k t_k + \sum_k kt_k/2 = \sum_k \frac{3}{4}kt_k + O(R)$, so $\sum_k t_k = \sum_k kt_k/4 + O(R)$. Since all $k \geq 5$, we have $\sum_k kt_k/4 \geq \sum_k \frac{5}{4}t_k$, contradicting that equality. \square

Now we strengthen this lemma to a stricter notion of aligned tilings by polykites, by considering what edge-to-edge tilings by the monokite are possible.

Lemma A.4. *The only edge-to-edge tilings by the monokite are (a) the tiling resulting from a 180° rotation about the midpoint of each side (Figure A.7), and (b) tilings composed of rows of equilateral triangles each composed of three kites, where some of those rows may be slid relative to each other (by the length of the long side of the kite) so they are no longer aligned as in the Laves tiling.*




Figure A.7: 180° tiling by kites

Proof. There are exactly two possible vertex figures in an edge-to-edge tiling by the monokite that do not appear in the Laves tiling: one with angles of 90° , 120° , 90° , and 60° in that order (Figure A.8), and one with angles of 90° , 90° , 60° , 60° , and 60° in that order (Figure A.9). If the first one occurs in a tiling, successive surrounding tiles are forced (in the order numbered) that force all neighbouring vertices, and so all vertices, to have that vertex figure. If the other

one occurs in a tiling, successive surrounding tiles are forced (in the order numbered, taking into account that the first vertex figure cannot appear anywhere in the tiling) that force two neighbouring vertices to have that vertex figure, and thus force two rows of equilateral triangles, slid relative to each other, and then the only possibilities on either side of such a row are another such row in either of two positions. \square




Figure A.8: Vertex with angles of 90° , 120° , 90° , and 60° in that order




Figure A.9: Vertex with angles of 90° , 90° , 60° , 60° , and 60° in that order

Lemma A.5. *If \mathcal{P} is a finite set of closed topological disk polykites, all from the same underlying Laves tiling, and \mathcal{P} admits a tiling with property H , it admits a tiling with property H where all polykites in the tiling are aligned to the same underlying Laves tiling, except possibly when \mathcal{P} contains the monokite and H is “isohedral tiling by 180° rotation”.*

Proof. If the edge-to-edge tiling with property H (“aligned” in the more general sense) is the 180° tiling by the monokite, we are done because the monokite admits an isohedral tiling. Otherwise, taking a minimal block of consecutive rows of equilateral triangles filled exactly with tiles from \mathcal{P} and translating it so as to be aligned with the underlying Laves tiling produces a tiling with property H . \square

Lemma A.6. *All tilings by the hat polykite are aligned to an underlying [3.4.6.4] Laves tiling.*

Proof. Note that any maximal segment in the union of the boundaries of tiles in such a tiling can contain no more than two sides of tiles, since any 90° angle is adjacent on either side to angles

greater than 90° . In particular, there are no infinite rays contained in the union of the boundaries of tiles.

This constraint immediately excludes the case of aligned tilings decomposing into rows of equilateral triangles slid relative to each other, and no two adjacent kites in the polykite are consistent with the 180° rotation tiling by the monokite. So any tiling not aligned with an underlying [3.4.6.4] is also unaligned in the more general sense, and we consider aligned components. Because there are no infinite rays among the boundaries of tiles, such aligned components must be bounded convex sets. The corners of those sets must be corners of a single polykite (since any two angles of the polykite add to at least 180°). But no corner of the polykite can be a corner of a convex set tiled by the polykite: four have reflex angles, seven are adjacent to a reflex angle, and for the remaining two, extending one of the sides from that vertex cuts off a region too small to be filled by polykites (Figure A.10). \square




Figure A.10: Extending a side of the polykite from either vertex that is neither a reflex angle nor adjacent to one

B. Case analysis for 1-patches

We present here details of a computer-generated but human-verifiable case analysis, based on consideration of 1-patches rather than 2-patches. This analysis can be used to complete a variant of the proof in Section 4 that, when tiles in a tiling by the hat polykite are assigned labels following the rules given there, (a) the labels assigned do induce a division into the clusters shown, and (b) the clusters adjoin other clusters in accordance with the matching rules. As is justified in Appendix A, we only consider tilings where all tiles are aligned with an underlying [3.4.6.4] Laves tiling.

B.1. Enumeration of neighbours

First we produce a list of possible neighbours of the hat polykite in a tiling. There are 58 possible neighbours when we only require such a neighbour not to intersect the original polykite; these are shown in Figure B.1, with that original polykite shaded. The first 41 of these neighbours remain in consideration for the enumeration of 1-patches. The final 17 are immediately eliminated (in the order shown) because they cannot be extended to a tiling: either there is no possible neighbour that can contain the shaded kite (without resulting in an intersection, or a pair of tiles that were previously eliminated as possible neighbours), or we eliminated Y as a neighbour of X and so can also eliminate X as a neighbour of Y .







Figure B.1: Possible neighbours (part 1)




Possible neighbour 13




Possible neighbour 14




Possible neighbour 15




Possible neighbour 16




Possible neighbour 17




Possible neighbour 18




Possible neighbour 19



Possible neighbour 20



Possible neighbour 21



Possible neighbour 22

Figure B.1: Possible neighbours (part 2)







Figure B.1: Possible neighbours (part 3)




Possible neighbour 33




Possible neighbour 34




Possible neighbour 35




Possible neighbour 36




Possible neighbour 37




Possible neighbour 38




Possible neighbour 39



Possible neighbour 40




Possible neighbour 41




Possible neighbour 42 (eliminated by considering neighbours containing the shaded kite)


Figure B.1: Possible neighbours (part 4)




Possible neighbour 43 (eliminated by considering neighbours containing the shaded kite)




Possible neighbour 44 (eliminated together with possible neighbour 43)




Possible neighbour 45 (eliminated by considering neighbours containing the shaded kite)




Possible neighbour 46 (eliminated by considering neighbours containing the shaded kite)




Possible neighbour 47 (eliminated together with possible neighbour 46)



Possible neighbour 48 (eliminated by considering neighbours containing the shaded kite)




Possible neighbour 49 (eliminated by considering neighbours containing the shaded kite)




Possible neighbour 50 (eliminated by considering neighbours containing the shaded kite)


Figure B.1: Possible neighbours (part 5)




Possible neighbour 51 (eliminated together with possible neighbour 50)




Possible neighbour 52 (eliminated by considering neighbours containing the shaded kite)




Possible neighbour 53 (eliminated together with possible neighbour 52)




Possible neighbour 54 (eliminated by considering neighbours containing the shaded kite)




Possible neighbour 55 (eliminated by considering neighbours containing the shaded kite)



Possible neighbour 56 (eliminated by considering neighbours containing the shaded kite)



Possible neighbour 57 (eliminated by considering neighbours containing the shaded kite)



Possible neighbour 58 (eliminated together with possible neighbour 57)

Figure B.1: Possible neighbours (part 6)

B.2. Enumeration of 1-patches

Having produced a list of possible neighbours, we now proceed to enumerating possible 1-patches. When we have a partial 1-patch (some number of neighbours for the original, shaded polykite), we pick some kite neighbouring that original polykite and enumerate the possible neighbouring polykites containing that kite, excluding any that would result in the patch containing two polykites that either intersect or form a pair of neighbours previously ruled out; the kite we use is chosen so that the number of choices for the neighbour added is minimal. This process results in 37 possible 1-patches; the partial patches from the search process are shown in Figure B.2 and the 1-patches are shown in Figures B.3.

Some of the 1-patches found can be immediately eliminated at this point, by identifying a tile in the 1-patch that cannot itself be surrounded by any of the 1-patches (that has not yet been eliminated) without resulting in either an intersection or a pair of neighbours that were previously ruled out. In the 12 cases implicated here, the tile that cannot be surrounded is shaded, and they are eliminated in the order shown, leaving 25 remaining 1-patches. For each of those remaining 1-patches, the classification of the central tile by the rules in Section 4 is shown.




Figure B.2: Partial patches (part 1)







Figure B.2: Partial patches (part 2)




Partial patch 15 (extends to partial patches 27–28)




Partial patch 16 (extends to partial patches 29–30)




Partial patch 17 (extends to partial patches 31–33)




Partial patch 18 (extends to partial patch 34)




Partial patch 19 (extends to partial patch 35)



Partial patch 20 (extends to partial patch 36)




Partial patch 21 (extends to partial patch 37)




Partial patch 22 (extends to partial patches 38–39)


Figure B.2: Partial patches (part 3)




Partial patch 23 (extends to partial patches 40–41)




Partial patch 24 (extends to partial patch 42)




Partial patch 25 (extends to partial patches 43–44)




Partial patch 26 (extends to partial patches 45–46)




Partial patch 27 (extends to partial patch 47)



Partial patch 28 (extends to partial patch 48)




Partial patch 29 (extends to partial patches 49–50)




Partial patch 30 (extends to partial patch 51)


Figure B.2: Partial patches (part 4)




Partial patch 31 (extends to partial patch 52)




Partial patch 32 (extends to partial patches 53–54)




Partial patch 33 (extends to partial patch 55)




Partial patch 34 (extends to partial patches 56–57)




Partial patch 35 (extends to partial patches 58–60)



Partial patch 36 (extends to partial patches 61–62)




Partial patch 37 (extends to partial patches 63–64)




Partial patch 38 (extends to partial patch 65)


Figure B.2: Partial patches (part 5)




Partial patch 39 (extends to partial patch 66)




Partial patch 40 (extends to partial patches 67–68)




Partial patch 41 (extends to partial patch 69, 1-patch 1)




Partial patch 42 (extends to partial patches 70–71)




Partial patch 43 (extends to partial patch 72)



Partial patch 44 (extends to partial patch 73)




Partial patch 45 (extends to 1-patch 2)




Partial patch 46 (extends to 1-patch 3)


Figure B.2: Partial patches (part 6)




Partial patch 47 (extends to partial patches 74–75)




Partial patch 48 (extends to partial patch 76)




Partial patch 49 (extends to partial patches 77–78)




Partial patch 50 (extends to partial patch 79, 1-patch 4)




Partial patch 51 (extends to partial patches 80–81)



Partial patch 52 (extends to partial patch 82)




Partial patch 53 (extends to partial patches 83–84)




Partial patch 54 (extends to partial patch 85, 1-patch 5)


Figure B.2: Partial patches (part 7)




Partial patch 55 (extends to partial patch 86, 1-patch 6)




Partial patch 56 (extends to 1-patches 7–8)




Partial patch 57 (extends to partial patches 87–88)



Partial patch 58 (no extensions)




Partial patch 59 (extends to partial patch 89)




Partial patch 60 (extends to partial patch 90)


Figure B.2: Partial patches (part 8)




Partial patch 61 (extends to partial patch 91)




Partial patch 62 (extends to partial patch 92)




Partial patch 63 (no extensions)




Partial patch 64 (extends to 1-patch 9)




Partial patch 65 (extends to 1-patches 10–11)



Partial patch 66 (no extensions)



Partial patch 67 (extends to 1-patch 12)



Partial patch 68 (extends to 1-patch 13)

Figure B.2: Partial patches (part 9)







Figure B.2: Partial patches (part 10)




Partial patch 75 (extends to 1-patch 21)




Partial patch 76 (extends to 1-patch 22)




Partial patch 77 (extends to 1-patch 23)



Partial patch 78 (extends to 1-patch 24)



Partial patch 79 (extends to 1-patch 25)



Partial patch 80 (extends to 1-patch 26)

Figure B.2: Partial patches (part 11)







Figure B.2: Partial patches (part 12)




Partial patch 87 (extends to 1-patch 33)




Partial patch 88 (extends to 1-patch 34)




Partial patch 89 (extends to 1-patch 35)



Partial patch 90 (extends to 1-patch 36)









Partial patch 91 (no extensions)








Partial patch 92 (extends to 1-patch 37)

Figure B.2: Partial patches (part 13)

1-patch 1 (central tile class F_2)1-patch 2 (central tile class H_3)1-patch 3 (central tile class H_3)1-patch 4 (central tile class F_2)1-patch 5 (central tile class P_2)

1-patch 6 (eliminated by trying to surround shaded tile)

Figure B.3: 1-patches (part 1)

1-patch 7 (central tile class H_2)1-patch 8 (central tile class H_2)1-patch 9 (central tile class H_1)1-patch 10 (central tile class H_4)

1-patch 11 (eliminated by trying to surround shaded tile)






1-patch 12 (central tile class FP_1)


Figure B.3: 1-patches (part 2)

1-patch 13 (central tile class FP_1)1-patch 14 (central tile class F_2)


1-patch 15 (eliminated by trying to surround shaded tile)



1-patch 16 (eliminated by trying to surround shaded tile)






1-patch 17 (eliminated by trying to surround shaded tile)




1-patch 18 (eliminated by trying to surround shaded tile)

Figure B.3: 1-patches (part 3)

1-patch 19 (central tile class H_4)1-patch 20 (central tile class H_4)

1-patch 21 (eliminated by trying to surround shaded tile)



1-patch 22 (eliminated by trying to surround shaded tile)



1-patch 23 (central tile class FP_1)1-patch 24 (central tile class FP_1)

Figure B.3: 1-patches (part 4)










Figure B.3: 1-patches (part 5)


1-patch 31 (central tile class P_2)1-patch 32 (central tile class P_2)1-patch 33 (central tile class H_2)1-patch 34 (central tile class H_2)

1-patch 35 (eliminated by trying to surround shaded tile)



1-patch 36 (eliminated by trying to surround shaded tile)

Figure B.3: 1-patches (part 6)



1-patch 37 (eliminated by trying to surround shaded tile)

Figure B.3: 1-patches (part 7)

B.3. Classification of outer tiles

For each of the possible neighbours that actually occurs in some of the remaining 1-patches, we can now list the possible classifications of a central tile that has such a neighbour; see Table B.1.

For each of the outer tiles in a 1-patch, we have some but not all of its neighbours, and can take the intersection of the sets from Table B.1 to produce a set of possible classes for that outer tile. Although this is not a single class, it can still be used for the within-cluster and between-cluster checks. In each case, it turns out that the set of possible classes for a neighbour appearing in one of those checks is a subset of the classes permitted by that check, and so we have a complete proof of the within-cluster and between-cluster matching properties that depends only on the enumeration of 1-patches presented here and not on a larger enumeration of 2-patches; the lists of checks and corresponding sets of classes appear below.

- 1-patch 1 (class F_2)
 - P_2 or F_2 neighbour FP_1 OK: $\{FP_1\} \subseteq \{FP_1\}$
 - F edge F^+ OK: $\{F_2\} \subseteq \{F_2\}$
 - F edge F^- OK: $\{F_2\} \subseteq \{F_2\}$
 - X^+ edge at top of polykite OK: $\{H_3\} \subseteq \{F_2, FP_1, H_3, H_4\}$
 - X^- edge at bottom of polykite OK: $\{H_2\} \subseteq \{F_2, H_2, P_2\}$
 - L edge at bottom of polykite OK: $\{P_2\} \subseteq \{P_2\}$
- 1-patch 2 (class H_3)
 - H_3 neighbour H_1 OK: $\{H_1\} \subseteq \{H_1\}$
 - H lower edge B^- OK: $\{FP_1, T_1\} \subseteq \{FP_1, T_1\}$
 - X^+ edge at right of polykite OK: $\{P_2\} \subseteq \{H_2, P_2\}$
 - X^- edge at bottom of polykite OK: $\{H_2, P_2\} \subseteq \{F_2, H_2, P_2\}$

Possible neighbour	Possible classes for central tile
2	$\{FP_1, F_2, H_4\}$
3	$\{H_2\}$
4	$\{H_4\}$
6	$\{H_3\}$
7	$\{H_1\}$
8	$\{H_3\}$
9	$\{FP_1, F_2, H_2, P_2, T_1\}$
11	$\{H_4\}$
12	$\{FP_1, T_1\}$
13	$\{FP_1, F_2, H_3, H_4\}$
14	$\{FP_1, F_2, H_4, P_2, T_1\}$
15	$\{H_2\}$
16	$\{H_1\}$
17	$\{F_2, H_2, P_2\}$
18	$\{FP_1, T_1\}$
19	$\{H_4\}$
20	$\{H_2\}$
22	$\{F_2, H_2, P_2\}$
23	$\{FP_1, T_1\}$
24	$\{H_1, H_3\}$
25	$\{FP_1, T_1\}$
26	$\{FP_1, F_2, H_4\}$
27	$\{P_2, T_1\}$
28	$\{H_1, H_2\}$
29	$\{FP_1, H_3, H_4, T_1\}$
30	$\{H_1\}$
32	$\{F_2, H_2, P_2\}$
33	$\{H_2\}$
34	$\{FP_1, F_2, H_3, H_4\}$
36	$\{H_2, P_2\}$
37	$\{FP_1, H_3, H_4\}$
38	$\{P_2, T_1\}$
39	$\{H_1\}$
40	$\{F_2, P_2\}$
41	$\{P_2\}$

Table B.1

- 1-patch 3 (class H_3)
 - H_3 neighbour H_1 OK: $\{H_1\} \subseteq \{H_1\}$
 - H lower edge B^- OK: $\{FP_1, T_1\} \subseteq \{FP_1, T_1\}$
 - X^+ edge at right of polykite OK: $\{F_2, FP_1, H_4\} \subseteq \{F_2, FP_1, H_3, H_4\}$
 - X^- edge at bottom of polykite OK: $\{F_2, P_2\} \subseteq \{F_2, H_2, P_2\}$
- 1-patch 4 (class F_2)
 - P_2 or F_2 neighbour FP_1 OK: $\{FP_1\} \subseteq \{FP_1\}$
 - F edge F^+ OK: $\{F_2\} \subseteq \{F_2\}$
 - F edge F^- OK: $\{F_2\} \subseteq \{F_2\}$
 - X^+ edge at top of polykite OK: $\{H_3\} \subseteq \{F_2, FP_1, H_3, H_4\}$
 - X^- edge at bottom of polykite OK: $\{FP_1, H_3, H_4\} \subseteq \{FP_1, H_3, H_4\}$
 - L edge at bottom of polykite OK: $\{F_2, FP_1\} \subseteq \{F_2, FP_1\}$
- 1-patch 5 (class P_2)
 - P_2 or F_2 neighbour FP_1 OK: $\{FP_1\} \subseteq \{FP_1\}$
 - T or P lower edge A^- OK: $\{H_2\} \subseteq \{H_2\}$
 - X^+ edge at top of polykite OK: $\{H_3\} \subseteq \{F_2, FP_1, H_3, H_4\}$
 - X^- edge at right of polykite OK: $\{FP_1, H_4\} \subseteq \{FP_1, H_3, H_4\}$
 - L edge at right of polykite OK: $\{F_2, FP_1\} \subseteq \{F_2, FP_1\}$
- 1-patch 7 (class H_2)
 - H_2 neighbour H_1 OK: $\{H_1\} \subseteq \{H_1\}$
 - H edge A^+ OK: $\{P_2, T_1\} \subseteq \{P_2, T_1\}$
 - X^+ edge at top of polykite OK: $\{F_2, FP_1, H_4\} \subseteq \{F_2, FP_1, H_3, H_4\}$
 - X^- edge at right of polykite OK: $\{F_2, P_2\} \subseteq \{F_2, H_2, P_2\}$
- 1-patch 8 (class H_2)
 - H_2 neighbour H_1 OK: $\{H_1\} \subseteq \{H_1\}$
 - H edge A^+ OK: $\{T_1\} \subseteq \{T_1\}$
 - X^+ edge at top of polykite OK: $\{H_3\} \subseteq \{F_2, FP_1, H_3, H_4\}$
 - X^- edge at right of polykite OK: $\{F_2, P_2\} \subseteq \{F_2, H_2, P_2\}$
- 1-patch 9 (class H_1)
 - H_1 neighbour H_2 OK: $\{H_2\} \subseteq \{H_2\}$

- H_1 neighbour H_3 OK: $\{H_3\} \subseteq \{H_3\}$
- H_1 neighbour H_4 OK: $\{H_4\} \subseteq \{H_4\}$
- H upper edge B^- OK: $\{FP_1, T_1\} \subseteq \{FP_1, T_1\}$
- 1-patch 10 (class H_4)
 - H_4 neighbour H_1 OK: $\{H_1\} \subseteq \{H_1\}$
 - X^+ edge at right of polykite OK: $\{P_2\} \subseteq \{H_2, P_2\}$
 - X^- edge at bottom of polykite OK: $\{H_2\} \subseteq \{F_2, H_2, P_2\}$
- 1-patch 12 (class FP_1)
 - FP_1 neighbour P_2 or F_2 OK: $\{F_2\} \subseteq \{F_2, P_2\}$
 - T, P or F edge B^+ OK: $\{H_4\} \subseteq \{H_3, H_4\}$
 - X^+ edge at right of polykite OK: $\{P_2\} \subseteq \{H_2, P_2\}$
 - X^- edge at bottom of polykite OK: $\{H_2\} \subseteq \{F_2, H_2, P_2\}$
 - L edge at bottom of polykite OK: $\{P_2\} \subseteq \{P_2\}$
- 1-patch 13 (class FP_1)
 - FP_1 neighbour P_2 or F_2 OK: $\{F_2\} \subseteq \{F_2, P_2\}$
 - T, P or F edge B^+ OK: $\{H_3\} \subseteq \{H_3, H_4\}$
 - X^+ edge at right of polykite OK: $\{P_2\} \subseteq \{H_2, P_2\}$
 - X^- edge at bottom of polykite OK: $\{H_2\} \subseteq \{F_2, H_2, P_2\}$
 - L edge at bottom of polykite OK: $\{P_2\} \subseteq \{P_2\}$
- 1-patch 14 (class F_2)
 - P_2 or F_2 neighbour FP_1 OK: $\{FP_1\} \subseteq \{FP_1\}$
 - F edge F^+ OK: $\{F_2\} \subseteq \{F_2\}$
 - F edge F^- OK: $\{F_2\} \subseteq \{F_2\}$
 - X^+ edge at top of polykite OK: $\{H_2\} \subseteq \{H_2, P_2\}$
 - X^- edge at bottom of polykite OK: $\{H_2\} \subseteq \{F_2, H_2, P_2\}$
 - L edge at bottom of polykite OK: $\{P_2\} \subseteq \{P_2\}$
- 1-patch 19 (class H_4)
 - H_4 neighbour H_1 OK: $\{H_1\} \subseteq \{H_1\}$
 - X^+ edge at right of polykite OK: $\{P_2\} \subseteq \{H_2, P_2\}$
 - X^- edge at bottom of polykite OK: $\{FP_1, H_3, H_4\} \subseteq \{FP_1, H_3, H_4\}$

- 1-patch 20 (class H_4)
 - H_4 neighbour H_1 OK: $\{H_1\} \subseteq \{H_1\}$
 - X^+ edge at right of polykite OK: $\{FP_1, H_4\} \subseteq \{F_2, FP_1, H_3, H_4\}$
 - X^- edge at bottom of polykite OK: $\{FP_1, H_4\} \subseteq \{FP_1, H_3, H_4\}$
- 1-patch 23 (class FP_1)
 - FP_1 neighbour P_2 or F_2 OK: $\{F_2\} \subseteq \{F_2, P_2\}$
 - T, P or F edge B^+ OK: $\{H_4\} \subseteq \{H_3, H_4\}$
 - X^+ edge at right of polykite OK: $\{P_2\} \subseteq \{H_2, P_2\}$
 - X^- edge at bottom of polykite OK: $\{FP_1, H_3, H_4\} \subseteq \{FP_1, H_3, H_4\}$
 - L edge at bottom of polykite OK: $\{F_2, FP_1\} \subseteq \{F_2, FP_1\}$
- 1-patch 24 (class FP_1)
 - FP_1 neighbour P_2 or F_2 OK: $\{F_2\} \subseteq \{F_2, P_2\}$
 - T, P or F edge B^+ OK: $\{H_3\} \subseteq \{H_3, H_4\}$
 - X^+ edge at right of polykite OK: $\{P_2\} \subseteq \{H_2, P_2\}$
 - X^- edge at bottom of polykite OK: $\{FP_1, H_3, H_4\} \subseteq \{FP_1, H_3, H_4\}$
 - L edge at bottom of polykite OK: $\{F_2, FP_1\} \subseteq \{F_2, FP_1\}$
- 1-patch 25 (class F_2)
 - P_2 or F_2 neighbour FP_1 OK: $\{FP_1\} \subseteq \{FP_1\}$
 - F edge F^+ OK: $\{F_2\} \subseteq \{F_2\}$
 - F edge F^- OK: $\{F_2\} \subseteq \{F_2\}$
 - X^+ edge at top of polykite OK: $\{H_2\} \subseteq \{H_2, P_2\}$
 - X^- edge at bottom of polykite OK: $\{FP_1, H_3, H_4\} \subseteq \{FP_1, H_3, H_4\}$
 - L edge at bottom of polykite OK: $\{F_2, FP_1\} \subseteq \{F_2, FP_1\}$
- 1-patch 26 (class FP_1)
 - FP_1 neighbour P_2 or F_2 OK: $\{F_2, P_2\} \subseteq \{F_2, P_2\}$
 - T, P or F edge B^+ OK: $\{H_4\} \subseteq \{H_3, H_4\}$
 - X^+ edge at right of polykite OK: $\{FP_1, H_4\} \subseteq \{F_2, FP_1, H_3, H_4\}$
 - X^- edge at bottom of polykite OK: $\{FP_1, H_4\} \subseteq \{FP_1, H_3, H_4\}$
 - L edge at bottom of polykite OK: $\{F_2, FP_1\} \subseteq \{F_2, FP_1\}$
- 1-patch 27 (class FP_1)

- FP_1 neighbour P_2 or F_2 OK: $\{F_2, P_2\} \subseteq \{F_2, P_2\}$
- T, P or F edge B^+ OK: $\{H_3\} \subseteq \{H_3, H_4\}$
- X^+ edge at right of polykite OK: $\{FP_1, H_4\} \subseteq \{F_2, FP_1, H_3, H_4\}$
- X^- edge at bottom of polykite OK: $\{FP_1, H_4\} \subseteq \{FP_1, H_3, H_4\}$
- L edge at bottom of polykite OK: $\{F_2, FP_1\} \subseteq \{F_2, FP_1\}$
- 1-patch 29 (class T_1)
 - T upper edge A^- OK: $\{H_2\} \subseteq \{H_2\}$
 - T or P lower edge A^- OK: $\{H_2\} \subseteq \{H_2\}$
 - T, P or F edge B^+ OK: $\{H_4\} \subseteq \{H_3, H_4\}$
- 1-patch 30 (class T_1)
 - T upper edge A^- OK: $\{H_2\} \subseteq \{H_2\}$
 - T or P lower edge A^- OK: $\{H_2\} \subseteq \{H_2\}$
 - T, P or F edge B^+ OK: $\{H_3\} \subseteq \{H_3, H_4\}$
- 1-patch 31 (class P_2)
 - P_2 or F_2 neighbour FP_1 OK: $\{FP_1\} \subseteq \{FP_1\}$
 - T or P lower edge A^- OK: $\{H_2\} \subseteq \{H_2\}$
 - X^+ edge at top of polykite OK: $\{H_2\} \subseteq \{H_2, P_2\}$
 - X^- edge at right of polykite OK: $\{FP_1, H_3, H_4\} \subseteq \{FP_1, H_3, H_4\}$
 - L edge at right of polykite OK: $\{F_2, FP_1\} \subseteq \{F_2, FP_1\}$
- 1-patch 32 (class P_2)
 - P_2 or F_2 neighbour FP_1 OK: $\{FP_1\} \subseteq \{FP_1\}$
 - T or P lower edge A^- OK: $\{H_2\} \subseteq \{H_2\}$
 - X^+ edge at top of polykite OK: $\{H_2\} \subseteq \{H_2, P_2\}$
 - X^- edge at right of polykite OK: $\{H_2\} \subseteq \{F_2, H_2, P_2\}$
 - L edge at right of polykite OK: $\{P_2\} \subseteq \{P_2\}$
- 1-patch 33 (class H_2)
 - H_2 neighbour H_1 OK: $\{H_1\} \subseteq \{H_1\}$
 - H edge A^+ OK: $\{P_2\} \subseteq \{P_2, T_1\}$
 - X^+ edge at top of polykite OK: $\{P_2\} \subseteq \{H_2, P_2\}$
 - X^- edge at right of polykite OK: $\{H_2, P_2\} \subseteq \{F_2, H_2, P_2\}$

- 1-patch 34 (class H_2)
 - H_2 neighbour H_1 OK: $\{H_1\} \subseteq \{H_1\}$
 - H edge A^+ OK: $\{T_1\} \subseteq \{T_1\}$
 - X^+ edge at top of polykite OK: $\{H_2\} \subseteq \{H_2, P_2\}$
 - X^- edge at right of polykite OK: $\{H_2, P_2\} \subseteq \{F_2, H_2, P_2\}$

References

- [Baš21] Bojan Bašić. A figure with Heesch number 6: pushing a two-decade-old boundary. *Math. Intelligencer*, 43(3):50–53, 2021. doi:10.1007/s00283-020-10034-w.
- [Ber66] Robert Berger. *The undecidability of the domino problem*. Number 66 in *Memoirs of the American Mathematical Society*. American Mathematical Soc., 1966. doi:10.1090/memo/0066.
- [BGG12] M. Baake, F. Gähler, and U. Grimm. Hexagonal inflation tilings and planar monotiles. *Symmetry*, 4(4):581–602, 2012. doi:10.3390/sym4040581.
- [Bha20] Siddhartha Bhattacharya. Periodicity and decidability of tilings of \mathbb{Z}^2 . *Amer. J. Math.*, 142(1):255–266, 2020. doi:10.1353/ajm.2020.0006.
- [Bör74] Károly Böröczky. Gömbkitöltések állandó görbületű terekben I. *Mat. Lapok*, 25(3–4):265–306, 1974. URL: http://real-j.mtak.hu/9373/1/MTA_MatematikaiLapok_1974.pdf#page=273.
- [BSJ91] M. Baake, M. Schlottmann, and P. D. Jarvis. Quasiperiodic tilings with tenfold symmetry and equivalence with respect to local derivability. *J. Phys. A*, 24(19):4637–4654, 1991. doi:10.1088/0305-4470/24/19/025.
- [BW92] Jonathan Block and Shmuel Weinberger. Aperiodic tilings, positive scalar curvature and amenability of spaces. *J. Amer. Math. Soc.*, 5(4):907–918, 1992. doi:10.2307/2152713.
- [DB81a] Nicolaas Govert De Bruijn. Algebraic theory of Penrose’s non-periodic tilings of the plane. I. *Kon. Nederl. Akad. Wetensch. Proc. Ser. A*, 43(84):39–52, 1981. doi:10.1016/1385-7258(81)90016-0.
- [DB81b] Nicolaas Govert De Bruijn. Algebraic theory of Penrose’s non-periodic tilings of the plane. II. *Kon. Nederl. Akad. Wetensch. Proc. Ser. A*, 43(84):53–66, 1981. doi:10.1016/1385-7258(81)90017-2.
- [GBN91] D. Girault-Beauquier and M. Nivat. Tiling the plane with one tile. In *Topology and category theory in computer science (Oxford, 1989)*, Oxford Sci. Publ., pages 291–333. Oxford Univ. Press, New York, 1991.
- [GK72] Yu. Sh. Gurevich and I. O. Koryakov. Remarks on Berger’s paper on the domino problem. *Sib. Math. J.*, 13(2):319–321, 1972. doi:10.1007/BF00971620.
- [GS99] Chaim Goodman-Strauss. A small aperiodic set of planar tiles. *European Journal of Combinatorics*, 20(5):375–384, July 1999. doi:10.1006/eujc.1998.0281.

- [GS09] Chaim Goodman-Strauss. Regular production systems and triangle tilings. *Theoret. Comput. Sci.*, 410(16):1534–1549, 2009. doi:10.1016/j.tcs.2008.12.012.
- [GS16] Branko Grünbaum and G.C. Shephard. *Tilings and Patterns*. Dover, second edition, 2016.
- [GT21] Rachel Greenfeld and Terence Tao. Undecidable translational tilings with only two tiles, or one nonabelian tile. 2021. arXiv:2108.07902.
- [GT22] Rachel Greenfeld and Terence Tao. A counterexample to the periodic tiling conjecture. 2022. arXiv:2211.15847.
- [Gum96] Petra Gummelt. Penrose tilings as coverings of congruent decagons. *Geom. Dedicata*, 62(1):1–17, 1996. doi:10.1007/BF00239998.
- [Hee35] H. Heesch. Aufbau der Ebene aus kongruenten Bereichen. *Nachr. Ges. Wiss. Göttingen, Math.-Phys. Kl. I, N. F.*, 1:115–117, 1935.
- [Hil02] David Hilbert. Mathematical problems. *Bull. Amer. Math. Soc.*, 8(10):437–479, 1902. doi:10.1090/S0002-9904-1902-00923-3.
- [JR21] Emmanuel Jeandel and Michaël Rao. An aperiodic set of 11 Wang tiles. *Adv. Comb.*, (1):1–37, 2021. doi:10.19086/aic.18614.
- [JS97] Hyeong-Chai Jeong and Paul J. Steinhardt. Constructing Penrose-like tilings from a single prototile and the implications for quasicrystals. *Phys. Rev. B*, 55:3520–3532, Feb 1997. doi:10.1103/PhysRevB.55.3520.
- [Kap22] Craig S. Kaplan. Heesch numbers of unmarked polyforms. *Contributions to Discrete Mathematics*, 17(2):150–171, 2022. URL: <https://cdm.ucalgary.ca/article/view/72886>.
- [Ken92] Richard Kenyon. Rigidity of planar tilings. *Invent. Math.*, 107(3):637–651, 1992. doi:10.1007/BF01231905.
- [Ken93] Richard Kenyon. Erratum: “Rigidity of planar tilings”. *Invent. Math.*, 112(1):223, 1993. doi:10.1007/BF01232432.
- [Ken96] Richard Kenyon. A group of paths in \mathbb{R}^2 . *Trans. Amer. Math. Soc.*, 348(8):3155–3172, 1996. doi:10.1090/S0002-9947-96-01562-0.
- [Man04] Casey Mann. Heesch’s tiling problem. *The American Mathematical Monthly*, 111(6):509–517, 2004. doi:10.1080/00029890.2004.11920105.
- [MM98] G. A. Margulis and S. Mozes. Aperiodic tilings of the hyperbolic plane by convex polygons. *Israel J. Math.*, 107:319–325, 1998. doi:10.1007/BF02764015.
- [Moz97] Shahar Mozes. Aperiodic tilings. *Invent. Math.*, 128(3):603–611, 1997. doi:10.1007/s002220050153.
- [MT16] Casey Mann and B. Charles Thomas. Heesch numbers of edge-marked polyforms. *Experimental Mathematics*, 25(3):281–294, 2016. doi:10.1080/10586458.2015.1096867.

- [Mye19] Joseph Myers. Polyomino, polyhex and polyiamond tiling, 2000–2019. Accessed: February 19th, 2023. URL: <https://www.polyomino.org.uk/mathematics/polyform-tiling/>.
- [Oll09] Nicolas Ollinger. Tiling the plane with a fixed number of polyominoes. In *Language and automata theory and applications*, volume 5457 of *Lecture Notes in Comput. Sci.*, pages 638–647. Springer, Berlin, 2009. doi:10.1007/978-3-642-00982-2_54.
- [Pen78] Roger Penrose. Pentaplexity. *Eureka*, 39:16–22, 1978. URL: <https://www.archim.org.uk/eureka/archive/Eureka-39.pdf#page=19>.
- [Pen97] Roger Penrose. Remarks on tiling: details of a $(1 + \epsilon + \epsilon^2)$ -aperiodic set. In *The mathematics of long-range aperiodic order (Waterloo, ON, 1995)*, volume 489 of *NATO Adv. Sci. Inst. Ser. C: Math. Phys. Sci.*, pages 467–497. Kluwer Acad. Publ., Dordrecht, 1997.
- [Rao17] Michael Rao. Exhaustive search of convex pentagons which tile the plane. 2017. arXiv:1708.00274.
- [Rei28] Karl Reinhardt. Zur Zerlegung der euklidischen Räume in kongruente Polytope. *Sitzungsber. Preuß. Akad. Wiss., Phys.-Math. Kl.*, pages 150–155, 1928.
- [Rob71] Raphael M. Robinson. Undecidability and nonperiodicity for tilings of the plane. *Invent. Math.*, 12:177–209, 1971. doi:10.1007/BF01418780.
- [Sen] Marjorie Senechal. Personal communication.
- [Sen96] Marjorie Senechal. *Quasicrystals and geometry*. Cambridge University Press, 1996.
- [SJ96] Paul J. Steinhardt and Hyeong-Chai Jeong. A simpler approach to Penrose tiling with implications for quasicrystal formation. *Nature*, 382:431–433, 1996. doi:10.1038/382431a0.
- [Soc07] Joshua E. S. Socolar. The hexagonal parquet tiling: k -isohedral monotiles with arbitrarily large k . *Math. Intelligencer*, 29(2):33–38, 2007. arXiv:0708.2663, doi:10.1007/BF02986203.
- [ST11] Joshua E. S. Socolar and Joan M. Taylor. An aperiodic hexagonal tile. *J. Combin. Theory Ser. A*, 118(8):2207–2231, 2011. doi:10.1016/j.jcta.2011.05.001.
- [ST12] Joshua E. S. Socolar and Joan M. Taylor. Forcing nonperiodicity with a single tile. *Math. Intelligencer*, 34(1):18–28, 2012. doi:10.1007/s00283-011-9255-y.
- [Tay] J.M. Taylor. Aperiodicity of a functional monotile. <https://sfb701.math.uni-bielefeld.de/preprints/sfb10015.pdf>.
- [Wan61] Hao Wang. Proving theorems by pattern recognition – II. *The Bell System Technical Journal*, 40(1):1–41, 1961. doi:10.1002/j.1538-7305.1961.tb03975.x.
- [WW21] James J. Walton and Michael F. Whittaker. An aperiodic tile with edge-to-edge orientational matching rules. *Journal of the Institute of Mathematics of Jussieu*, pages 1–29, 2021.