

Tinka Bruneau

School of Mathematics & Statistics
University of Glasgow
2580916b@student.gla.ac.uk

Michael F. Whittaker

School of Mathematics & Statistics
University of Glasgow
mike.whittaker@glasgow.ac.uk

The Solution

Planar aperiodic tile sets: from Wang tiles to the Hat and Spectre monotiles

A brief history of planar aperiodic tile sets is presented, starting from the Domino Problem proposed by Hao Wang in 1961. In this article, Tinka Bruneau and Mike Whittaker provide highlights that led to the discovery of the Taylor–Socolar aperiodic monotile in 2010 and the Hat and Spectre aperiodic monotiles in 2023. The Spectre tile is an amazingly simple monotile; a single tile whose translated and rotated copies tile the plane but only in a way that lacks any translational periodicity. They showcase this breakthrough discovery through the 60+ years that aperiodic tile sets have been considered.

In 1961, Hao Wang asked if it is possible to algorithmically determine whether translated copies of a finite set of marked square tiles can tile the plane. This became known as the Domino Problem. As a first attack on the Domino Problem, Wang posed his *Fundamental Conjecture*: a set of Wang tiles either tiles the plane with translational periodicity or does not tile. The other possibility is a set of tiles that can form a tiling of the plane but always without any translational periodicity, such a tile set is called *aperiodic*.

Robert Berger found a counterexample to Wang’s Fundamental Conjecture by finding an aperiodic tile set with 20426 tiles. This proved undecidability of the Domino Problem. After finding such a set of tiles, the search was on for the minimal size of such a set. Emmanuel Jeandel and Michaël Rao re-

cently proved that an aperiodic set of Wang tiles must have at least eleven tiles and gave an example of such a tile set. Note that we provide pictures of all tilings mentioned in the introduction throughout the article and the reader is invited to look ahead to get a feeling for these tile sets and their tilings.

Wang tiles are always squares, but what if we allow tiles to come in arbitrary shapes and allow tiles to be rotated and reflected? In the 1970s, both Sir Roger Penrose and Robert Ammann found aperiodic tile sets with just two tiles, now commonly referred to as the Penrose and Ammann–Beenker tiles, respectively. This left open the aperiodic monotile problem: Is there a single tile that can tile the plane but only in a way that forbids translational periodicity?

Two early attempts found monotiles that use overlaps to force aperiodicity. While

these are not our focus, they are very interesting both historically and from the standpoint of aperiodic order. The first of these was Gummelt’s covering of the Penrose tiling by a single decorated decagon [10]. Gummelt’s overlapping tile was really the first mono-cluster, and provided the first evidence that a monotile could be possible. The second was Penrose’s $(1 + \varepsilon + \varepsilon^2)$ -tiling [2, 14]. This was originally described as a single hexagon with a smaller hexagon and half-hexagons inscribed in its interior. The matching rules required overlapping hexagons to form the same type of hexagons as at the larger scale. An excellent discussion of both of these tilings can be found in [5].

In a breakthrough result of Joan Taylor in 2010 [20], the first connected and completely geometrically defined monotile was discovered, but the tile is not simply connected. That is, the tile is not a topological disc. Taylor joined forces with Joshua Socolar to introduce their tile to the mathematical community [18, 19]. The original formulation of the Taylor–Socolar tile was described as a decorated hexagon with certain matching rules between both

neighbouring and next to neighbouring tiles, and the tiling requires a reflected tile.

In March 2023, Dave Smith, Craig Kaplan, Joseph Myers and Chaim Goodman-Strauss found a simply connected monotile that they called the Hat [16]. Amazingly, the Hat is a simple polygonal shape and is entirely geometric, in the sense that no additional matching rules on how tiles are allowed to meet are required to enforce aperiodicity. What's more, in the article, a beautiful proof of aperiodicity was used to show that there are, in fact, an uncountable number of tiles in the Hat family that are also simply connected monotiles. The tiling requires a reflected version of the Hat. This left open the question of whether a simply connected and geometric solution to the monotile problem without reflections is possible.

In a marvellous stroke of insight, Smith, Kaplan, Myers and Goodman-Strauss realised that a member of the Hat family could also be used to tile the plane without using a reflected copy of the tile [17]. The Spectre aperiodic monotile was discovered. The Spectre is in the family of Hat tilings, but was one of the singular members that allowed periodic tilings. The authors realised that forbidding the reflected tile still allowed tilings of the plane, but all such tilings lack translational periodicity. Thus, the Spectre provides a remarkable solution to the aperiodic monotile problem!

Tiles, tilings, and their properties

In this article, we restrict ourselves to two-dimensional tilings. We mostly follow the language and terminology defined in Baake and Grimm's *Aperiodic Order* [5], which is highly recommended to the interested reader. The second main reference on general tiling theory is the book of Grünbaum and Shephard [9], which contains an excellent introduction to Wang tiles and aperiodic tile sets.

We begin with the building blocks of a tiling. A *prototile* is a labelled subset of \mathbb{R}^2 that is equal to the closure of its interior. This just means that there are no dangling bits hanging from the tiles; in fact, one should probably just think of labelled polygons. In addition, we often allow *decorations* of the prototiles, like edge markings that must match, arrows to determine tile orientation, or lines that must continue across tile edges. Often, but not always, these decorations can be realised by puzzle




Figure 1 Tilings of the plane by a single square prototile. The prototile is on the left along with two possible tilings.

like bumps and dents in the tile edges. We'll see examples of this soon.

Let \mathcal{P} be a finite set of prototiles and G a subgroup of the isometry group of \mathbb{R}^2 . A *tiling* of \mathbb{R}^2 (the plane) is a countable collection of tiles $T = \{t_i \mid i \in \mathbb{N}\}$ such that:

- $t_i = \gamma \cdot p$ for some $p \in \mathcal{P}$ and γ in G ;
- $\bigcup_{i \in \mathbb{N}} t_i = \mathbb{R}^2$;
- $\text{interior}(t_i) \cap \text{interior}(t_j) = \emptyset$ if $i \neq j$.

Let's take a minute to understand the definition. The first bullet point specifies that all tiles in the tiling must be isometric copies of prototiles (i.e., some composition of a translation, rotation and/or reflection of a prototile), where the group G specifies the exact subset of isometries we allow for a given prototile set. Here we are thinking of the prototiles as actually sitting in the plane, so that we can move them about to form a tiling. In this article, G will always be the translation group, the direct isometry group (translations and rotations), or the full isometry group. Note that when we use the full isometry group G , we will include the reflected tile(s) in the prototile set for clarity. The second bullet point implies that the tiling covers the entire Euclidean plane and the third specifies that tiles never overlap except at their edges. Moreover, in the case that we have extra decorations, the tiling must also satisfy any extra rules specified by the decorations.

As Figure 1 depicts, it is easy to find tilings of the plane by a single square prototile of side length one. In fact, there are infinitely many possible tilings that arise from an unmarked square prototile. We note that it's impossible to pictorially represent a complete tiling of the plane, and the reader is meant to extrapolate how a complete tiling is produced from the small patch provided. For both examples in Figure 1, we can let G be the subgroup of translations. The tiling on the right reveals why there are infinitely many possible

square tilings, even taken up to translation, by choosing different relative shifts for rows of square tiles. In this case, the second row is shifted by $\frac{1}{2}$ with respect to the bottom row that's sitting on the x -axis, and the third is shifted by $\frac{1}{3}$ with respect to the bottom, and so on. We could use the same procedure to form rows below the bottom row to make a complete tiling of the plane.

We can translate a tiling T by a vector x in \mathbb{R}^2 via $T+x := \{t+x \mid t \in T\}$. Notice that the tiling in the middle of Figure 1 has the property that $T+(1,0) = T$. Indeed, if we translate all tiles in the tiling over by one unit to the left the sets $T+(1,0)$ and T are exactly the same! We say that a tiling T is *periodic* if there exists a nonzero translation x in \mathbb{R}^2 such that $T+x = T$, and we say that T is *nonperiodic* if $T+x = T$ implies $x = 0$. That is, if we took an infinite transparent photocopy of the tiling T , then the tiling is periodic if we can shift this photocopy by some non-trivial translation so that it perfectly matches T and nonperiodic if it only matches in exactly one place. The tiling on the right of Figure 1 is also periodic with $x = (1,0)$. We'll see examples of nonperiodic tilings in the next section.

There are many more concepts that are important in tiling theory. Again, the book of Baake and Grimm [5] contains a modern treatment and an exceptional foreword written by Roger Penrose. For example, there has been discussion on the *local indistinguishability classes* (LI-classes) for the Hat and Spectre monotiles, see [3, p. 2], where the notion of LI-classes can be found in [5, Section 5.1.1].

Wang tiles

A founder of modern tiling theory was the philosopher Hao Wang (1921–1995). A set of square prototiles with marked edges are called Wang tiles. Here we label the edges with a colour and also a number, but merely to aid legibility. Tiles must meet




Figure 2 Two sets of Wang tiles, the left can only tile periodically, the right can tile both periodically and nonperiodically.

along complete edges only when the symbol matches and we only allow G to be the group of translations; no rotations or reflections of prototiles allowed.

In 1961 Wang [22] proposed the Domino Problem: given a set of Wang tiles, is it possible to algorithmically determine whether the set tiles the plane? Wang tiles are theoretically important in logic, since the behaviour of any Turing machine can be mimicked using some particular set of Wang tiles, see [9, Section 11.4]. Wang realised that a set \mathcal{P} of Wang tiles has four possibilities:

1. The set \mathcal{P} does *not* tile the plane, such as a single tile with four distinct symbols on its edges.
2. The set \mathcal{P} can only tile the plane periodically, such as the Wang tile on the left of Figure 2.
3. The set \mathcal{P} can tile both periodically and nonperiodically, such as the Wang tile set on the right of Figure 2 (points for finding a method to fill the tiling out to the plane in a non-periodic way).
4. The set \mathcal{P} tiles the plane but only non-periodically.

Wang realised that the existence of a Wang tile set satisfying possibility 4 implies that the Domino Problem is undecidable [22]. This led to his *Fundamental Conjecture* that there are no Wang tile sets satisfying possibility 4. However, Wang’s student, Robert Berger, found the first such Wang tile set with 20426 tiles [7]! According to Grünbaum and Shephard [9, Chapter 11], Berger subsequently got the set down to only 104 tiles (actually 103 tiles, see [11, Section 1.2]). Shortly after, Donald Knuth (the inventor of TeX amongst other things) modified Berger’s set to only require 92 tiles. The hunt was on to find the smallest set of Wang tiles, with contributions from Hans Läuchli, Raphael Robinson, Roger Penrose, Robert Ammann and John Conway finding Wang tiles sets for

$n = 56, 52, 40, 35, 34, 32, 24,$ and 16, see [9, Section 1.11]. The following quote appears in [9, p.596]:

“The reduction in the number of Wang tiles in an aperiodic set from over 20000 to 16 has been a notable achievement. Perhaps the minimum possible number has now been reached. If, however, further reductions are possible then it seems certain that new ideas and methods will be required.”

Two further reductions were found, the Kari–Culik Wang tile set with only 13 tiles was found using number theoretic meth-

ods, see [5, Section 5.7.4]. The Jeandel–Rao Wang tile set in Figure 3 reduced this number to just eleven Wang tiles [11]! The new method in this case was an exhaustive computer search, so we now know that this is the minimal number, although we don’t know that the depicted Wang tile set is essentially unique as an aperiodic set of eleven Wang tiles. A patch of Jeandel–Rao tiles is in Figure 4.

Relaxing the rules and the Penrose tiles

A set of prototiles, of any shape, that admits tilings of the plane but only nonperiodically is called an *aperiodic tile set*. During the flurry of activity to reduce the number of Wang tiles in the 1970s, Roger Penrose and Robert Ammann began considering aperiodic tile sets of polygons with specified markings. Both independently found an aperiodic tile set with just two tiles! However, we focus only on the Penrose tiles here. An excellent introduction to aperiodic tiles sets, with proofs that had not appeared previously in the literature, is [9, Chapter 10].




Figure 3 The smallest possible set of Wang tiles was found by Jeandel and Rao [11] consisting of just eleven tiles.




Figure 4 A patch of Jeandel–Rao Wang tiles shows the type of complexity required to reduce the number of tiles to eleven.




Figure 5 The two tile Penrose aperiodic tile set. The tiles must meet full-edge to full-edge and the red and blue lines must continue from one tile to the next.

The Penrose two tile aperiodic set is depicted in Figure 5 and a patch of tiles in Figure 7. The tiles are simple polygonal shapes that must match full-edge to full-edge with the blue and red lines continuing from one tile to the next. We note that the line matching rules can be encoded purely geometrically by puzzle like bumps and dents, see [5, p. 154]. These tiles were reduced from other versions of Penrose tile sets, with more prototiles, and have a ten-fold symmetry group. Roger Penrose gave a lecture at ‘Hatfest: Celebrating the Discovery of an Aperiodic Monotile’ at the University of Oxford in July 2023, where he showed how he constructed his two tile set in Figure 5. He suggested that he may have been indirectly inspired by a book of Kepler that was on his parents’ bookshelf. In fact, during his lecture he overlaid his tiles with an image produced by Kepler to an almost perfect match, see Figure 6. More information on this link can be found in Rodrigo Treviño’s Notices of the American Mathematical Society article [21].




Figure 6 Roger Penrose overlaying a patch of the Penrose tiling over Kepler's pentagon pattern.




Figure 7 A patch of a Penrose tiling.



Figure 8 Toilet paper embossed with Penrose tiles.

The success of the Penrose tiles in popular culture has been immense, to the point that his tiles were seemingly imprinted on Kleenex toilet paper, see Figure 8. Luckily, the tiles had been patented, and the following was issued by David Bradley, the Director of Pentaplex:

“So often we read of very large companies riding rough-shod over small businesses or individuals, but when it comes to the population of Great Britain being invited by a multi-national to wipe their bottoms on what appears to be the work of a Knight of the Realm without his permission, then a last stand must be made.”

The case was settled out of court and Kleenex stopped making the toilet paper. Squares are highly coveted by scientists studying aperiodic order, the second author has one framed on his office wall.

The Taylor–Socolar monotile

Given that two-tile aperiodic tile sets were found in the 1970s, it naturally begs the question of whether an aperiodic monotile is possible. That is, can one find a single tile and a set of local rules (decorations) that tile the plane but only nonperiodically. Here we have to be careful about what we mean by local rules. This is made very precise by Baake and Grimm in [5, Section 5.7]. For us, suffice it to say that a local rule is a decoration that can only ‘see’ a finite and predefined radius in \mathbb{R}^2 .

The Taylor–Socolar tile is an excellent example of a decoration that defines a local rule that is not edge-to-edge. The rules were discovered in 2010 by Joan Taylor, an amateur mathematician from Australia [20]. She made contact with tiling expert




Figure 9 The Taylor–Socolar monotile. On the left is the tile and its reflection. The image on the right demonstrates the (R2)-rule, flags that meet across a single hexagonal edge must point in the same direction.




Figure 10 A patch of a Taylor–Socolar tiling.

Joshua Socolar in order to verify the discovery and fully work out the details. In a crowning achievement, they introduced the first true aperiodic monotile with local rules that are defined only between neighbours and next nearest neighbours [18, 19]. The Taylor–Socolar tile consists of a hexagonal tile and its mirror image such that:

- (R1) The black lines must continue across tile edges.
- (R2) The purple flags at vertices of tiles that meet across a single hexagonal edge must point in the same direction.

The tiles are depicted in Figure 9 and the arrows on the right point to the flag rule (R2). Figure 10 shows a patch and the reader should verify that (R1) and (R2) hold in any local patch in order to properly understand the local rules.

Theorem 1 [18, Theorem 1]. *The Taylor–Socolar monotile is aperiodic; that is, there are tilings formed by isometries of the Taylor–Socolar tile satisfying (R1) and (R2) in every local patch, and every such tiling is nonperiodic.*

The proof of Theorem 1 is very elegant and is recommended to the interested reader. Interestingly, the Taylor–Socolar tile can be realised by shape alone. However, the tile is not simply connected; that is, the tile is connected but is not a topo-

logical disc. See Figure 11, where distinct tiles are coloured differently so that one can see how tiles interconnect. Focussing on the single black tile in the centre of the image shows the interesting geometry of the tile.

For the experts, adding an extra vertex consistency rule, (R3), to the Taylor–Socolar tile forces a single LI-class, see [18, p. 2215]. This additional rule was included in the original discovery by Taylor [20]. In practice, this means that the tiling space of the Taylor–Socolar tiling satisfying (R1)–(R3) is minimal in the sense of dynamical systems. Moreover, the Taylor–Socolar tilings are model sets [2, 12], and thus are strongly related to mathematical



Figure 11 A geometric version of the Taylor–Socolar tile. Distinct tiles have different colours so that we can see where the connected tiles interpenetrate.

Image taken from [19, Figure 6]




Figure 12 An orientational monotile, where the local rule depends on the orientation of the arrow in the plusses and minuses. The top edges are labelled with positive $\mathcal{R}2$ -charges, which are oriented in each tile from left to right: anticlockwise, both clockwise and anticlockwise, and clockwise, respectively. The bottom edges are labelled with negative $\mathcal{R}2$ -charges, from left to right: clockwise, both clockwise and anticlockwise, and anticlockwise

quasicrystals. These properties imply that the Taylor–Socolar tile is immensely important from the perspective of aperiodic order.

Jamie Walton and the second author found a modification of the Taylor–Socolar tile that has edge-to-edge matching rules [23]. However, the $(\mathcal{R}2)$ rule below is somewhat unusual, it depends on orientation in the following way. Two tiles t_1 and t_2 are permitted to meet along a shared edge e only if:

- $(\mathcal{R}1)$ The black lines continue across e ,
- $(\mathcal{R}2)$ Whenever the two charges at e in t_1 and t_2 both have a clockwise orientation then they must be opposite in charge.

Theorem 2 [23, Theorem 1.1]. *The tile in Figure 13 is aperiodic; that is, there are tilings formed by isometries of the tile satisfying $(\mathcal{R}1)$ and $(\mathcal{R}2)$ in every local patch, and every such tiling is nonperiodic.*




Figure 13 A patch of the orientational monotile. Two neighbouring charges with clockwise orientation must be opposite.

In this case the proof of aperiodicity is surprisingly simple and does not require a lot of case-checking. Indeed, we show that $(\mathcal{R}1)$ -lines always lead to longer $(\mathcal{R}1)$ -lines. Thus, there are arbitrarily long $(\mathcal{R}1)$ -lines in any tiling. This precludes translational periodicity. For if there was a non-trivial translation $x \in \mathbb{R}^2$ such that $T+x=T$, then there must be an $(\mathcal{R}1)$ -line that is longer than $|x|$ and the structure of the $(\mathcal{R}1)$ -triangles forbids this translation.

The Hat and Spectre tiles

The Hat monotile [16] was discovered by David Smith, Craig S. Kaplan, Joseph Samuel Myers and Chaim Goodman-Strauss with the article appearing on the Mathematics arXiv in March 2023. The paper generated an immediate buzz, resulting in newspaper articles in both *The New York Times* and *The Guardian*. The Hat was originally discovered by David Smith in November 2022 and the authors worked furiously to understand whether the Hat is an aperiodic monotile.

The Hat tile is unbelievably simple and elegant. It can be found by forming a hexagonal grid, dividing each hexagon into kites and then combining kites from three neighbouring hexagons into a tile. The kites are formed by cutting the hexagons with straight lines through the midpoints of opposite edges. For this reason the authors often refer to it as a polykite. See Figure 14 for the Hat tile, its mirror image, and a rendering into kites that combine to form regular hexagons.

Interestingly, the Hat was discovered by Smith when tinkering with polyforms to see what sorts of visually interesting tilings he could create. He got stuck building large patches of Hat tiles and sought help from Kaplan’s Heesch Number software. The *Heesch Number* of a tile is the largest number of concentric rings that form a patch around a single tile by isometric copies, where a ring consists of all tiles that touch the previous ring. The current record holder for a tile that does not tile the plane was found by Bašič [6] and has Heesch Number 6. Kaplan was able to show that the Hat has at least Heesch number 10 and then improved that to 16. So, it seemed pretty likely that the Hat would tile the plane, and in a way that didn’t seem to have any translational periodicity. This type of attack on the monotile problem was completely different from previous attempts, where authors typically start with a tiling of the plane and then add rules to try and enforce nonperiodicity. In some ways this approach goes back to a prototile set decidability problem, in the original sense of Wang. A patch of an infinite tiling formed from Hats is depicted in Figure 15. Notice that reflected Hats appear with a low frequency compared to Hat tiles.

Theorem 3 [16, Theorem 1.1]. *The Hat monotile is aperiodic; that is, there are tilings formed by isometries of the Hat tile, and every such tiling is nonperiodic.*

There are two proofs that the Hat tiles the plane and two proofs of aperiodicity. The authors show that Hat tilings arise from a (nonstone) substitution rule that is recognisable, in the sense that the hierarchy can be deduced in an infinite tiling. On one hand, this allows one to construct patches of arbitrarily large size that nest into each other in order to construct a tiling of the plane. On the other hand, recognisability implies that one can identify structure in the tiling of arbitrarily large size, and hence the tiling cannot be periodic. The second proof of existence uses a rather simple fusion system (see [8] for the definition of fusion) to define arbitrarily large patches of tiles that expand out from a fixed tile, see [16, Figure 2.11]. The second proof of aperiodicity is more involved but has already led to the discovery of the Spectre tile.




Figure 14 The Hat tile and its mirror image are an aperiodic tile set without any need for further decorations. On the right we see how the Hat is formed from kites that combine to form regular hexagons.

The idea behind the second proof of aperiodicity is to contract the edges of the hat tile to find other combinatorially equivalent tilings in the sense that the patterns formed by tiles are the same. Indeed, the two edge lengths in a Hat tile are at a ratio of $\sqrt{3}$ to one another and come in complementary (opposite) pairs. Thus, we can label the Hat tile as $\text{Tile}(\sqrt{3}, 1)$. The idea is now to consider $\text{Tile}(a, b)$ for $0 \leq a \leq \sqrt{3}$ and $0 \leq b \leq 1$. There is an excellent animation of this on YouTube titled ‘Aperiodic monotile animation’. At the two extremes are $\text{Tile}(0, 1)$, called the Comet and $\text{Tile}(\sqrt{3}, 0)$ called the Chevron, see [16, Figure 3.1]. Since these two tiles are combinatorially equivalent to the Hat tiling, the authors are free to use them to deduce properties of the Hat tiling. Thus, the authors suppose that the Comet tiling is periodic and aim to derive a contradiction, meaning that this hypothesis could not be correct and the tiling is nonperiodic. Since the Comet is assumed to be periodic

and forms the same combinatorial tiling as the Chevron, we can deduce that the Chevron is also periodic. Moreover, since the tilings are combinatorially equivalent there must be an affine map between the periodic lattice of the Comet and that of the Chevron. Using an argument, somewhat similar to an argument that $\sqrt{2}$ is irrational, they prove that such an affine map cannot exist. This implies that the Hat tiling must also be nonperiodic. See [16, Section 3] for further details.

One of the most interesting developments was again discovered by Smith and his coauthors, the Spectre tile [17]. Amazingly, this is $\text{Tile}(1, 1)$ from the previous paragraph, which can be used to tile the plane periodically if one allows reflection of the tile. However, the authors realised that it is still possible to tile the plane if reflections of the tile are not allowed to appear in a tiling, and more amazingly that all such tilings are nonperiodic!

$\text{Tile}(1, 1)$ appears on the left hand side of Figure 16 and a version of the Spectre tile appears in the centre. The Spectre is merely an edge modification of $\text{Tile}(1, 1)$ to curves, which eliminate the possibility of using both the tile and its reflection to tile the plane. We understand that Dave Smith proposed to call the tile the Spectre due to the image in the centre of Figure 16. The image on the right was constructed through a fusion system to build arbitrarily large patches [17], and we think that also looks spectre like.

Theorem 4 [17, Theorem 2.2]. *Tile(1, 1) and the Spectre monotile are aperiodic; that is, there are tilings formed by direct (orientation-preserving) isometries of each tile, and every such tiling is nonperiodic.*

Let us make a couple of remarks about Theorem 4. First, this is an absolutely incredible result that was completely unexpected, even given the recent Hat tile result. We note that $\text{Tile}(1, 1)$ is referred to by the authors of [17] as a weakly chiral monotile since it satisfies Theorem 4, but allowing a reflection results in a tile that can be used to construct a periodic tiling. A Spectre tile, see the centre of Figure 16, is referred to as a strictly chiral aperiodic monotile since it satisfies Theorem 4, but any prototile set containing both the




Figure 15 A patch of Hat tiles.




Figure 16 The image on the left is $\text{Tile}(1,1)$ and one must forbid reflections to obtain an aperiodic monotile. The tile in the centre is the Spectre tile that tiles the plane without allowing reflection, and does not tile the plane if there is at least one Spectre tile and at least one mirror image of the Spectre tile. The image on the right comes from a fusion system to construct arbitrarily large patches of Spectre tiles.




Figure 17 A patch of Spectre tiles.

Spectre and its mirror image nonredundantly does not tile the plane.

Theorem 4 was proved by showing that $\text{Tile}(1,1)$ and the Spectre arise from a (nonstone) inflation rule that is recognisable, similar to the first proof of Theorem 3 for the Hat tile.

For the experts, Baake, Gähler and Sadun have extended the 1-parameter family of tiles given by $\text{Tile}(a,b)$ to complex vari-

ables [3]. They showed that all the continuous hulls are topologically conjugate dynamical systems under these parameters, up to linear rescaling of the ambient space, and found a self-similar representative they call the CAP tiling. The name follows from their result that the tiling is a cut and project tiling, and hence forms a model set. They also compute the cohomology of the family and show that it has pure-

point dynamical spectrum. In current work, the same authors have also shown that the Spectre tile has similar properties [4].

Since the original Hat preprint appeared [16], Akiyama and Araki have provided yet another proof of existence and aperiodicity of a member of the Hat family [1]. They use the Golden Hex substitution to prove existence and Golden Ammann bars to prove aperiodicity.

Mathematicians are now discovering that the unique structure of the Spectre tiling lends it many other fascinating properties [15]. For example, the dimer model asks how many ways there are to colour the edges of the tiles so that each vertex meets exactly one coloured edge (dimer). Remarkably, this model can be exactly solved on the Spectre tiling. The number of dimer arrangements is $2^{N_{\text{Mystic}}+1}$, where N_{Mystic} is the number of *Mystic* tiles, see [17, p. 6] or [15, p. 2] for the definition of a *Mystic*. More remarkable still, the dimer model can also be exactly solved when quantum superpositions of dimer placements are allowed! Thus, Singh and Flicker have exactly solved the quantum dimer model for the first time in any setting. ❖

Acknowledgements

We'd like to thank Michael Baake, Kevin Brix, Felix Flicker, Robbert Fokink, Franz Gähler, Craig Kaplan, Jan Mazáč and Jamie Walton for excellent comments and suggestions on early drafts. This paper was written while the second author was a guest at the Fields Institute, and he thanks them for their hospitality and exceptional research environment.

References

- S. Akiyama and Y. Araki, An alternative proof for an aperiodic monotile (2023), arXiv: 2307.12322.
- M. Baake, F. Gähler and U. Grimm, Hexagonal inflation tilings and planar monotiles, *Symmetry* 4 (2012), 581–602.
- M. Baake, F. Gähler and L. Sadun, Dynamics and topology of the hat family of tilings (2023), arXiv:2305.05639.
- M. Baake, F. Gähler and L. Sadun, Dynamics and topology of the Spectre tiling, work in progress.
- M. Baake and U. Grimm, *Aperiodic Order. Volume 1: A Mathematical Invitation*, Cambridge University Press, 2013.
- B. Bašić, A gure with Heesch number 6: pushing a two-decade-old boundary, *Math. Intellig.* 43 (2021), 50–53.
- R. Berger, *The Undecidability of the Domino Problem*, Memoirs of the American Mathematical Society, Vol. 66, American Mathematical Society, 1966.
- N. P. Frank and L. Sadun, Fusion: a general framework for hierarchical tilings of \mathbb{R}^d , *Geom. Dedicata* 171 (2014), 149–186.
- B. Grünbaum and G. C. Shephard, *Tilings and Patterns*, W. H. Freeman, 1987.
- P. Gummelt, Penrose tilings as coverings of congruent decagons, *Geom. Dedicata* 62 (1996), 1–17.
- E. Jeandel and M. Rao, An aperiodic set of 11 Wang tiles, *Advances in Combinatorics* 1 (2021).
- J.-Y. Lee and R. V. Moody, Taylor–Socolar hexagonal tilings as model sets, *Symmetry* 5 (2013), 1–46.
- R. Penrose, Pentaplexity: A class of non-periodic tilings of the plane, *Math. Intellig.* 2 (1979), 32–37.
- R. Penrose, Remarks on tiling: Details of a $(1+\varepsilon+\varepsilon^2)$ -aperiodic set. in *The Mathematics of Long-Range Aperiodic Order*, NATO ASI Series C, Vol. 489, Kluwer, 1997, pp. 467–497.
- S. Singh and F. Flicker, Exact solution to the quantum and classical dimer models on the Spectre aperiodic monotiling (2023), arXiv: 2309.14447.
- D. Smith, J. S. Myers, C. S. Kaplan and Chaim Goodman-Strauss, An aperiodic monotile (2023), arXiv:2303.10798.
- D. Smith, J. S. Myers, C. S. Kaplan and Chaim Goodman-Strauss, A chiral aperiodic monotile (2023), arXiv:2305.17743.
- J. E. S. Socolar and J. M. Taylor, An aperiodic hexagonal tile, *J. Comb. Th. A* 118 (2011), 2207–2231.
- J. E. S. Socolar and J. M. Taylor, Forcing non-periodicity with a single tile, *Math. Intel.* 34 (2012), 18–28.
- J. M. Taylor, Aperiodicity of a functional monotile, preprint (2010); available from <http://www.math.unibielefeld.de/sfb701/preprints/view/420>.
- R. Treviño, Aperiodic tilings, order, and randomness, *Notices of the AMS* (September 2023).
- H. Wang, Proving theorems by pattern recognition, II, *Bell Sys. Tech. J.* 40 (1961), 1–41.
- J. J. Walton and M. F. Whittaker, An aperiodic tile with edge-to-edge orientational matching rules, *J. Inst. Math. Jussieu* 22 (2023), 1727–1755.