

THE SPECTRE KNOT: ONE ENDLESS STRAND ON AN APERIODIC MONOTILE

JUSTIN ROUGHLEY

ABSTRACT. The Spectre of Smith, Myers, Kaplan and Goodman-Strauss tiles the plane, but only non-periodically. We decorate its straight-edged form $\text{Tile}(1, 1)$ with fourteen ribbon pieces, the same on every tile, so that the ribbons continue straight across every tile edge. We prove two theorems. In every tiling by copies of $\text{Tile}(1, 1)$ of one handedness whose supertile hierarchy covers the plane, no strand of ribbon is closed. In the *nest tiling*, grown outwards from a single tile through Δ -supertiles, each block sitting in the same place (child number 1) inside the next, all the pieces join into a single strand, a two-way infinite path through every tile; we do not claim this for other tilings. Every supertile connects the ribbon ends on its boundary in one and the same pattern, with no closed strand inside; ribbons zigzag along a boundary shared by two supertiles, so lengthening it by two edges changes nothing; and the side lengths at every level are reached from those at level 1 by such two-edge steps. So finitely many computations at level 1 settle every level. The combinatorial core is formalised in the Lean proof assistant. At the author's instruction, an AI system generated the proof, the programs, the formalisation, the figures and this text. This paper has not been peer reviewed.

1. INTRODUCTION

A *Celtic knot* is a pattern of interlaced ribbons [2, 6]; the number of strands in finite designs of this kind has been studied, for instance, by Fisher and Mellor [9]. A *knot tile* is a polygon carrying ribbon pieces that end on its edges, so that when copies of the tile are placed edge to edge, the pieces join up into longer ribbons; knot mosaics [23] assemble knot diagrams in this way from a set of square tiles. We ask for a knot tile whose ribbons, in a tiling of the whole plane, form one endless strand. In a periodic tiling this cannot happen. Suppose that the pieces formed one strand. Each piece has two ends, each joined to exactly one other piece, so the strand is a two-way infinite path $\dots, p_{-1}, p_0, p_1, \dots$, and a period τ of the decorated tiling (for a tile with no symmetry, such as $\text{Tile}(1, 1)$, every period of the tiling is one) maps it onto itself, keeping joins: it shifts it, $p_i \mapsto p_{i+k}$, or reverses it, $p_i \mapsto p_{c-i}$. A reversal is impossible, since τ^2 would then fix every piece, but τ^2 moves every tile; for the same reason $k \neq 0$. The pieces would then form $|k|$ orbits under $\langle \tau \rangle$, whereas the tiles form infinitely many, as $\mathbb{R}^2 / \langle \tau \rangle$ has infinite area and the tiles of one orbit cover at most the area of one tile there; and tiles in different orbits carry pieces in different orbits. So we look at a tile that tiles only non-periodically.

Socolar and Taylor [30] found a hexagonal tile that, with its mirror image, tiles only non-periodically when its markings are matched. In 2023 Smith, Myers, Kaplan and Goodman-Strauss described the *hat*, an aperiodic monotile [28], and then the *Spectre* [29]: a shape that tiles the plane, but only non-periodically and, even when reflections are allowed, only with all tiles of the same handedness [29, Thm. 2.2]. The Spectre is a curved-edge form of the polygon $\text{Tile}(1, 1)$, and tilings by Spectres correspond exactly to the tilings by $\text{Tile}(1, 1)$ with all tiles of one handedness (Figure 1, Section 2). For background on tilings and substitution see [16, 1].

Date: 22 October 2026.

2020 Mathematics Subject Classification. Primary 52C20; Secondary 05B45, 52C23, 57K10, 68V05, 68V20.

Key words and phrases. aperiodic monotile, Spectre, substitution tiling, Celtic knot, knot tile, computer-assisted proof, formal verification, Lean.




FIGURE 1. Left: the tile $\text{Tile}(1, 1)$ with its fourteen ribbon pieces. Right: the blocks N_1, N_2, N_3 of the nest tiling (Section 8), each drawn with the previous block, its child number 1, dashed. The pieces of the original tile N_0 (shaded) lie on 5 arcs of N_1 , 2 of N_2 and 1 of N_3 (Table 5, row $b = 0$), one colour per arc; an arc is a ribbon path through the block from boundary to boundary (Section 4).

We give a knot tile on the Spectre that is *mono-stranded*: in one explicit tiling the pieces of all the tiles form a single strand (Theorem B). Moreover, no tiling whose supertile hierarchy covers the plane (defined below) contains a closed strand (Theorem A). The decoration (Figure 1) puts one crossing point on each of the fourteen edges, with two ribbon ends at each, and joins these 28 ends in pairs inside the tile by fourteen pieces: twelve *hooks*, each turning back round one vertex, and two pieces that cross each other once. Across an edge, the ribbons continue straight through the crossing point.

Every tiling by $\text{Tile}(1, 1)$ with tiles of one handedness (a *homochiral* tiling) groups uniquely into a hierarchy of larger and larger *supertiles* [29, Thm. 2.2]. We say that the hierarchy *covers the plane* if the supertiles containing some tile exhaust the plane. This does not depend on the tile: if the supertiles containing a tile t exhaust the plane, one of them contains an interior point of any other tile t' , and so contains t' (supertiles are unions of tiles, and distinct tiles have disjoint interiors); from that level on, t and t' have the same supertiles. If the hierarchy does not cover the plane, the tiling splits into two or more *infinite supertiles*, unions of nested chains of supertiles, separated by a *fault* (Section 10).

Theorem A (The Daedalus Theorem). *In every homochiral tiling of the plane by $\text{Tile}(1, 1)$ whose supertile hierarchy covers the plane, with every tile carrying the decoration of Figure 1 (Definition 3.1), there is no closed strand: every strand is a two-way infinite path of ribbon pieces.*

Our second result is proved for one tiling, the *nest tiling* $\mathcal{N}$, grown from a single tile N_0 by repeatedly taking the Δ -supertile N_{b+1} that has N_b as its child number 1 (Section 8); the N_b are its *blocks*, and an *arc* of a block is a path of its ribbon pieces with both ends on its boundary (Lemma 4.1). In the language of substitution tilings, $\mathcal{N}$ is generated from one tile by iterating the substitution about it, and is a fixed point of the substitution in the combinatorial sense (compare [1]). In the tile addresses of [31, §2.8] (after [14]), N_0 has the address

$$\Delta \stackrel{1}{\leftarrow} \Delta \stackrel{1}{\leftarrow} \Delta \stackrel{1}{\leftarrow} \dots$$

Theorem B (The Ariadne Theorem). *In the nest tiling $\mathcal{N}$, with every tile carrying the decoration of Figure 1 (Definition 3.1), all the ribbon pieces of all the tiles form a single strand. It is a two-way infinite path that passes through every tile. More precisely, for every $b \geq 1$, all the pieces of the block N_b lie on a single arc of N_{b+3} if b is even, and of N_{b+4} if b is odd.*

We make no claim about the number of strands in any other tiling: computations indicate that in at least one other tiling whose hierarchy covers the plane, the pieces of a level-1 supertile lie on three strands. Tilings with faults are covered by neither theorem (Section 10).

The names recall the myth: Daedalus built the labyrinth, in which no strand closes ([Theorem A](#)); Ariadne gave the thread, which runs through every tile of $\mathcal{N}$ ([Theorem B](#)).

The method. The proof has three ideas.

- (1) *Black boxes.* Seen from outside, a supertile is just a list of which ribbon end on its boundary is joined to which. Every supertile of every kind and level $m \geq 1$ has the same *signature* as a single tile: a hook at every tile vertex on its boundary but two adjacent ones, the crossover corners A and B , where two ribbons cross over; and no closed strand inside ([Definition 4.3](#), [Theorem 7.6](#), [Figure 5](#)).
- (2) *Zigzags.* When two such boxes are glued along a common boundary path with no crossover corner inside it, the ribbons zigzag along it, and how they join the two ends of the path depends only on the parity of its length ([Lemma 5.1](#)). So lengthening a shared path by two edges changes nothing ([Lemma 7.1](#)).
- (3) *Fixed parities.* The supertiles are combinatorially hexagons and compose by the substitution of [[29](#), Fig. 5.1]. Their side lengths are given by Fibonacci numbers ([Proposition 6.2](#)), grow with the level and never change parity. So the side lengths at any level are reached from those at level 1 by lengthening, two edges at a time ([Lemma 7.2](#)). By induction on the level, a supertile is glued from children that already have the signature, so its box is that of a composition of boxes, fixed by its kind, the side lengths and where the children's corners A and B lie ([Section 7.1](#)); the 18 compositions at the level-1 lengths (nine kinds, B on either side of A), computed once, give the signature at every level.

A closed strand is finite, so when the hierarchy covers the plane it lies in one supertile, where the signature forbids it; this gives [Theorem A](#). The single strand ([Theorem B](#)) follows the same pattern. Unlike the signature, this uses how each block sits in the next (as child 1). In the next block up, every arc of a block lies on a hook (an arc with both ends at one boundary vertex), and which hook depends only on where the arc meets the boundary; for vertices well inside five of the six sides, only on the side ([Lemma 9.3](#), [Table 6](#)). This *junction map* is computed at level 1 and carried to every level by the same lengthening. Followed three or four levels up, it puts all the arcs of a block on one arc ([Theorem 9.1](#)). The finite computations (the level-1 data and [Computation 7.5](#); [Computation 7.3](#) on the 18 compositions at the level-1 lengths, and [Computation 9.2](#) on two of them) are exact; [Computations 7.3](#) and [9.2](#) are repeated in the kernel of the Lean proof assistant, and so is [Computation 7.5](#), on the level-1 supertiles of `exact.py`, with their corner A , taken as input. The agreement with [[29](#), Figs. 4.1, 5.1] is checked by scripts, against a render of Fig. 4.1, whose labels were transcribed from the figure, and against an encoding of Fig. 5.1 ([Appendix A](#)). [Figure 2](#) shows how the steps fit together, and which of them Lean checks.

Organisation. [Section 2](#) recalls the hierarchy of [[29](#)], in our numbering. [Section 3](#) defines the decoration and its strands, and [Section 4](#) introduces boxes and the signature. [Section 5](#) proves the Zigzag Lemma, [Section 6](#) finds the side lengths, and [Section 7](#) proves the Signature Theorem and [Theorem A](#). [Section 8](#) constructs the nest tiling and shows that it covers the plane, and [Section 9](#) proves [Theorem B](#). [Section 10](#) discusses scope and open questions; [Appendix A](#) describes the computations, and [Appendix B](#) the use of AI, the formal verification and the tools.

Related work. Curves drawn on tiles and joined across edges, each closed or two-way infinite as in [Lemma 3.4](#), have been studied for random Truchet tilings [[25](#), [17](#)] and for a deterministic family of them, in which a renormalisation, akin to our induction on supertiles, shows that for almost every parameter almost every curve is closed [[18](#)]. Mirror curves of Celtic, Tchokwe and Tamil designs are another case [[13](#), [19](#)]: a design drawn as a single line is called *monolinear*, and a rectangular grid of $m \times n$ cells with no inner mirrors carries $\gcd(m, n)$ curves [[13](#)]. So are the Conway–Penrose arcs on Penrose tiles, where at most two curves of each colour fail to close [[12](#), Ch. 1] (see also [[10](#)]), and girih interlace on Penrose tilings, whose strands close into loops and knots [[4](#), [5](#), [11](#)]. Single curves through




FIGURE 2. The route of the proof. Shaded: checked by the Lean formalisation in a combinatorial form, including the computations at level 1 ([Computation 7.5](#) on the level-1 supertiles of `exact.py`, with their corner A , taken as input). Dashed: the link with actual tilings, cited from [29] or argued in words. Shaded and dashed: Lean checks the combinatorial part, and the link with actual tilings is argued in words ([Appendix B](#)).

finite patches have been built as Hamiltonian cycles, through every vertex of arbitrarily large regions of Ammann–Beenker tilings [26] and through every tile of finite Penrose and Ammann–Beenker patches [22], and as space-filling curves from aperiodic substitution tilings [24]. Closest to this paper, an online project [3] decorates selected edges of $\text{Tile}(1, 1)$ and describes a single non-crossing curve through every tile of one nested tiling, with a machine check of the combinatorial core of its argument. Here every edge carries a crossing point, and two pieces cross inside each tile. In our drawing every crossing is a transversal double point ([Appendix A](#)), so over and under can alternate by the checkerboard rule [7], and the single strand becomes an alternating diagram of a line. Celtic knotwork is alternating in the same way, and up to topological equivalence its designs realise every alternating link projection [15]. We are not aware of earlier full knotwork decorations that form a single strand.

Use of AI and verification status. This work is an experiment in using AI for mathematical exploration. At the author’s instruction, an AI system (Claude, by Anthropic) generated the programs, the proofs, the Lean formalisation, the figures and this text; the author asked the question, steered the work, made the choices and decisions, corrected what was generated, and answers for the work. The author is responsible for the paper: for the correctness of its arguments and results, and for the completeness and accuracy of its citations of prior work. The AI system is not an author. Lean checks the combinatorial core of the proof; the link with actual tilings is argued in words. Most of the proof has not been reviewed by a human, and the paper has not been peer reviewed. [Appendix B](#) gives the full account: how the work was done, what the formalisation checks, and the tools used.

2. SPECTRE TILINGS

The proof climbs the supertile hierarchy of [29] one level at a time. This section fixes what it uses: the tile, its supertiles ([Theorem 2.1](#)), an orientation, a side numbering and a crossover corner for each supertile ([Section 2.3](#)), and the tables by which each level from level 2 up is assembled from the one below ([Section 2.4](#)). The decoration enters in [Section 3](#); the tables are first used in [Section 6](#), and a first reading may postpone their details until then.

Notation			
Tile(1, 1), v_i, e_i ; homochiral tiling $\mathcal{T}$	Section 2.1	rim ports $R(u_k), L(u_k)$	Section 4
kinds K , levels m , children, n_T	Section 2.2	π_P , arcs, boxes, closed strands of P	Lemma 4.1
rim, vertices, sides, rim vertices	Section 2.2	signature at (A, B)	Definition 4.3
own orientation, own order, ω_m	Definition 2.2	$\ell_m(K, j)$	Section 6
crossover corners $A, B = A + \omega_m$	Section 2.3	letters $\alpha, \dots, \theta$	Section 6
parent kind T , child kinds K_c	Definition 2.3	lengths $\alpha_m, \dots, \theta_m$	Proposition 6.1
child c , child side $c.j$, parent side J	Tables 2 and 3	d_m	Lemma 6.4
ports $R(v_i), L(v_i)$	Section 3	$C_T(L, d, \omega), A_c, B_c, c_\Sigma, A_\Sigma, A_\Gamma$	Section 7.1
pieces, hooks, crossover	Definition 3.1	moves, admissible pairs	Section 7.2
strands, closed strands, $\mathcal{G}(\mathcal{T})$	Definition 3.3	nest tiling $\mathcal{N}$, blocks N_b	Section 8
patch P , its rim ∂P , $E = E(P)$	Section 4	junctions q , hooks H_q	Section 9.2

2.1. The tile. Let $\rho = e^{i\pi/6}$. The polygon Tile(1, 1) has fourteen unit edges; we label its vertices $v_0, \dots, v_{13}$ anticlockwise (Figure 5, left) and its edges $e_i = v_i v_{i+1}$ (indices mod 14), with $v_0 = 0$ and

$$v_{i+1} - v_i = \rho^{k_i}, \quad (k_0, \dots, k_{13}) = (0, 10, 1, 3, 0, 2, 5, 7, 4, 6, 6, 8, 11, 9).$$

All vertices lie in $\mathbb{Z}[\rho]$. Tile(1, 1) is the equilateral member of the family Tile(a, b) [29, §§1–2], with its long side split into the collinear edges e_9, e_{10} . It is the mirror image of the polygon drawn in [29, Fig. 1.1] and has the handedness of the single Spectre and of the Mystic in [29, Fig. 2.1]. A tiling by Tile(1, 1) is *homochiral* if any two of its tiles are related by a *proper motion* (a rotation followed by a translation). *Spectres* are tiles whose tilings correspond exactly to the homochiral tilings by Tile(1, 1), each Spectre placed by the same isometry as the corresponding copy of Tile(1, 1) [29, §2, Lemma 2.1]; one is obtained by replacing every edge of Tile(1, 1) by a suitable curve. So what [29] proves for tilings by Spectres holds for homochiral tilings by Tile(1, 1) [29, p. 5]. Such tilings exist [29, Thm. 2.2], and in each, tiles meet full edge to full edge [29, proof of Lemma 2.1]. Reflecting the plane if necessary, we assume that every tile is the image of Tile(1, 1) under a rotation and a translation; a reflection carries the decoration of each tile (Definition 3.1) to that of its image, and so strands to strands. The tile has no non-trivial symmetry, proper or improper (its turning angles at $v_1, \dots, v_{13}, v_0$, in units of 30° , are $-2, 3, 2, -3, 2, 3, 2, -3, 2, 0, 2, 3, -2, 3$; the only 0 is at v_{10} , and the sequence is not symmetric about it), so an isometry between two tiles is unique.

2.2. Supertiles. We use the hierarchy of [29], with nine *kinds* $\Gamma, \Delta, \Theta, \Lambda, \Xi, \Pi, \Sigma, \Phi, \Psi$. A *level-0 supertile* of kind $K \neq \Gamma$ is a single tile; one of kind Γ is a *Mystic* [29, Fig. 2.1] (the darker pair of tiles in Figure 3): a pair of tiles $\{h(\text{Tile}(1, 1)), h(g(\text{Tile}(1, 1)))\}$, where $g(z) = \rho z + v_8$ and h is a proper motion, with *first* tile $h(\text{Tile}(1, 1))$. The first tile is determined by the pair, since the Mystic has no non-trivial proper symmetry: a symmetry exchanging its tiles would be an involution taking Tile(1, 1) to $g(\text{Tile}(1, 1))$, but the only proper motion doing so is g , a rotation by 30° . A level- $(m+1)$ supertile of kind T is the union of level- m supertiles, its n_T *children*, of the kinds listed in Table 1, where $n_\Gamma = 7$ and $n_T = 8$ for $T \neq \Gamma$. A level-1 supertile has the kind of its cluster in [29, Fig. 4.1]; the positions of its six vertices (Figure 3) determine that kind [29, proof of Thm. 4.1]; a level- $(m+1)$ supertile, $m \geq 1$, has kind T if its children form the panel T of [29, Fig. 5.1] (described below). The children of a level-1 supertile are numbered as in Figure 3, and those of a supertile of level ≥ 2 by their places in [29, Fig. 5.1], as listed in the last two rows of Table 1 (the match is unique: each panel has one Γ , one Δ and one Σ , and every other hexagon is told apart from those of its kind by which of these three it touches; Tables 1 and 2). A level-0 supertile that is a single tile gets the kind listed in Table 1 for its child number in its level-1 supertile (so a level-0 Δ is a child 1); beyond that, only the distinction between Γ and the other kinds is used at level 0. Child 1 of every supertile is a Δ and the last child a Γ , and a Δ of level $m \geq 1$ occurs only as child 1.

The level-1 supertiles are the Spectre clusters (9 tiles) and the Mystic cluster (8 tiles, kind Γ) of [29, Fig. 2.1], with the markings of [29, Fig. 4.1], up to reflection: [29, Fig. 2.1] draws

TABLE 1. The children of each kind of supertile. The last two rows give each child's place in the eight-hexagon layout of [29, Fig. 5.1]: L left, D the Δ , B bottom, R right, S the Σ , TR top right, TL top left, G the Γ .

parent T	child 0	1	2	3	4	5	6	7
Γ	Π	Δ	Θ	Σ	Ξ	Φ	Γ	
Δ	Ξ	Δ	Ξ	Φ	Σ	Π	Φ	Γ
Θ	Ψ	Δ	Π	Φ	Σ	Π	Φ	Γ
Λ	Ψ	Δ	Ξ	Φ	Σ	Π	Φ	Γ
Ξ	Ψ	Δ	Π	Φ	Σ	Ψ	Φ	Γ
Π	Ψ	Δ	Ξ	Φ	Σ	Ψ	Φ	Γ
Σ	Ξ	Δ	Ξ	Φ	Σ	Π	Λ	Γ
Φ	Ψ	Δ	Ψ	Φ	Σ	Π	Φ	Γ
Ψ	Ψ	Δ	Ψ	Φ	Σ	Ψ	Φ	Γ
place in [29, Fig. 5.1], $T \neq \Gamma$	L	D	B	R	S	TR	TL	G
place, $T = \Gamma$	L	D	R	S	TR	TL	G	




FIGURE 3. The nine level-1 supertiles, with tiles of the handedness of Tile(1, 1) ([29, Fig. 2.1] draws their mirror images). The Mystic is the darker pair of tiles. Grey numbers: the children (Table 1); black dots: the six vertices; blue numbers: the sides, numbered clockwise (the own orientation at level 1); open circle: the crossover corner A , which is vertex 0, except for Σ , where it lies inside side 5; filled square: the corner B (Section 2.3).

the clusters with tiles of the opposite handedness to its single Spectre and Mystic, and so to Tile(1, 1); Figure 3 draws them with that of Tile(1, 1). Explicitly, if the first tile of the Mystic (child 7) is Tile(1, 1), then child $c \leq 6$ of a level-1 supertile of kind $K \neq \Gamma$ is $f_c(\text{Tile}(1, 1))$, with $f_c(z) = \rho^{r_c} z + t_c$, $(r_0, \dots, r_6) = (8, 10, 10, 0, 2, 2, 4)$ and $(t_0, \dots, t_6) = (3 + \rho + \rho^3, 2 + \rho + 2\rho^2 + \rho^3, 4 + 2\rho + \rho^2 + 2\rho^3, 2 + 2\rho + 2\rho^2 + 2\rho^3, 1 + 2\rho + \rho^2 + 2\rho^3, 3\rho^2 + 3\rho^3, 1 + \rho^2 + 3\rho^3)$; the level-1 Γ -supertile is this cluster without child 2, with children $3, \dots, 7$ renumbered $2, \dots, 6$. In this frame the crossover corner A (Section 2.3) is 0, and the six vertices lie where Computation 7.5 records them: at the distances from A , in own order, that Table 4, ℓ_1 and d_1 give. Figure 3 draws the clusters up to a rotation and a translation.

Combinatorially, the level-1 supertiles are the marked hexagons of [29, Fig. 4.2], and the level- n supertiles that [29, §5] builds from these hexagons are our level- $(n + 1)$ supertiles. The edges of these hexagons carry the labels x^+ , x^- ($x \in \{\alpha, \beta, \gamma, \delta, \epsilon, \zeta, \theta\}$) and η of [29, Fig. 4.2] (Table 4); two hexagons sharing an edge carry x^+ and x^- , or η and η , on it (the matching rule of [29, §4]). For each kind T , the panel T of [29, Fig. 5.1] shows the eight (for

Γ , seven) hexagons, with their kinds, that make up a hexagon of kind T one level up, with that hexagon's six vertices and the labels of its edges.

The following theorem restates, in our language, results of [29].

Theorem 2.1 (Smith, Myers, Kaplan, Goodman-Strauss [29]). *Let $\mathcal{T}$ be a homochiral tiling by $\text{Tile}(1, 1)$.*

- (a) [29, Thms. 2.2, 3.1, 4.1, §5, Fig. 2.1] *The tiles of $\mathcal{T}$ group uniquely into level-1 supertiles, each the image under a rotation and a translation of one of the clusters of Figure 3 (those of [29, Fig. 2.1], reflected). Each contains one Mystic, in the position shown in [29, Fig. 2.1]; this Mystic and the other tiles, taken singly, are the level-0 supertiles inside it. For every $m \geq 1$ the level- m supertiles of $\mathcal{T}$ group uniquely into level- $(m+1)$ supertiles. Each tile lies in a unique nested sequence $X_0 \subset X_1 \subset \dots$ with X_m a level- m supertile.*
- (b) [29, Thms. 3.1, 4.1, 5.1, §5, Fig. 4.1] *For every $m \geq 1$, the tiling of the plane by the level- m supertiles of $\mathcal{T}$ is combinatorially equivalent (homeomorphic as a topological complex [29, §2.2]) to a tiling by regular hexagons. Each level- m supertile of kind K corresponds to a hexagon marked as the hexagon K of [29, Fig. 4.2] if m is even, and as its mirror image if m is odd; adjacent supertiles correspond to hexagons that share an edge, and the labels on that edge match. The grouping into level- $(m+1)$ supertiles corresponds, up to reflection, to the grouping of hexagons shown in [29, Fig. 5.1], and the six points where a level- $(m+1)$ supertile of kind T meets two others correspond to the six points marked on the outline of the panel T . At level 1, the six points where a supertile meets two others correspond to the dots marked on its cluster in [29, Fig. 4.1], the points by which the proof of [29, Thm. 4.1] tells the clusters apart.*

In (a), [29, Thm. 4.1] gives the clusters and [29, Thm. 2.2] the uniqueness at every level; we use uniqueness only to speak of the hierarchy of $\mathcal{T}$, and, in Section 8, to identify the hierarchy of the nest tiling. In (b), the choice of mirror image at odd m is fixed at $m = 1$, where our level-1 supertiles are the mirror images of the clusters in [29, Figs. 2.1, 4.1] (see the base case of $(*)$ in the proof of Proposition 2.4); the choice alternates with m by [29, Thm. 5.1].

By Theorem 2.1(b), every level- m supertile ($m \geq 1$) is a closed topological disc whose boundary, its *rim*, is a simple closed polygon made of tile edges. Its *vertices* are the six points of the rim where it meets two other level- m supertiles, and its *sides* are the six arcs of the rim between consecutive vertices. (The tile vertices on a rim are called *rim vertices*; for a supertile of level ≥ 1 , “vertex” alone always means one of the six.) If a vertex of one level- m supertile lies on the rim of another, it is a vertex of the other too, since exactly three level- m supertiles meet there. Each side is shared, as a path of tile edges, with exactly one neighbouring supertile.

2.3. Own orientation, side numbering and the crossover corner. By Theorem 2.1(b), the cyclic order of the labelled sides reverses at each level, so we orient supertiles alternately; then side j of a supertile of kind K carries the same label at every level (Table 4), and one composition pattern serves all levels.

Definition 2.2 (Own orientation). The *own orientation* of a level- m supertile is anticlockwise if m is even and clockwise if m is odd. We write $\omega_m = (-1)^m$. Its *own order* is the cyclic order in which this orientation runs round its rim.

Thus a child and its parent are always oriented oppositely. Each supertile of level $m \geq 1$ has a distinguished rim vertex A , its *crossover corner*, defined recursively: for a level-1 supertile, A is the vertex v_0 of the first tile of its Mystic; for a level- $(m+1)$ supertile, $m \geq 1$, A is the crossover corner of its Γ child. (That this point lies on the rim is computed at level 1 and follows from Proposition 2.4(d) at higher levels.) The second corner B is the rim vertex next to A anticlockwise (the sense in which v_1 follows v_0 on $\text{Tile}(1, 1)$); as own order alternates with m , this is $B = A + \omega_m$ in its own order. We call A and B the *crossover corners*, since




FIGURE 4. The composition pattern for Δ (left) and for Γ (right), in the layout of [29, Fig. 5.1]. Each hexagon is a child: kind, number (grey), sides 0–5 (blue); child 1, a Δ , is shaded light blue, and the Γ child darker. Thick black: the parent’s rim; dots: its vertices; vermillion: its sides; circled: its crossover corner, vertex 0 of the Γ child. Left, open diamonds: the junctions 9, 12, 16, 17, 21 of Section 9.2. Every parent kind but Γ has the children and gluing of the left panel, with the kinds of Table 1, but its own vertices and side numbers (Table 3). The children’s sides run anticlockwise and the parent’s clockwise, their own orientations when the children have even level; for odd level, take the mirror image.

the supertile’s two crossover arcs run between them (Theorem 7.6), as a tile’s two crossover pieces run between v_0 and v_1 (Definition 3.1). The vertices of a level- m supertile ($m \geq 1$) are numbered $0, \dots, 5$ in its own order, starting with the first vertex at or after A ; side j runs from vertex j to vertex $j + 1$ (Figure 3).

2.4. The composition pattern. By Theorem 2.1(b), every level from level 2 up is assembled from the one below in the same combinatorial way. We record this in our numbering.

Definition 2.3 (Composition pattern). The *composition pattern* gives, for each parent kind T , the kinds $K_0, \dots, K_{n_T-1}$ of the children (Table 1); for each side of each child, the sibling’s side across it, or ∂ if it lies on the parent’s rim (Table 2); and for each side of the parent, the children’s sides that make it up, in the parent’s own order (Table 3). Figure 4 draws it.

Proposition 2.4 (The pattern holds at every level). Let $\mathcal{T}$ be a homochiral tiling by $\text{Tile}(1, 1)$, let $m \geq 1$, and let Y be one of its level- $(m + 1)$ supertiles, of kind T , with children $X_0, \dots, X_{n_T-1}$. Then:

- (a) X_c has kind K_c (Table 1);
- (b) side j of X_c is the same path of tile edges as side j' of $X_{c'}$, traversed in the opposite direction, whenever Table 2 lists $c'.j'$, and it lies on the rim of Y when the table lists ∂ ;
- (c) side J of Y is the concatenation of the children’s sides listed in Table 3, and vertex J of Y is vertex $j + 1$ of the child c whose side $c.j$ is listed first for side J ;
- (d) the crossover corner A of Y is vertex 0 of its Γ child, which is vertex 0 of Y if $T \neq \Sigma$, and lies inside side 5 of Y if $T = \Sigma$.

Proof. By Theorem 2.1(b), the level- m supertiles inside Y are combinatorially the hexagons of the panel T of [29, Fig. 5.1], with hexagon edges as sides and hexagon vertices as vertices. We use the dictionary

- (*) side j of a level- m supertile of kind K carries the label of the edge of hexagon K in direction W, SW, SE, E, NE, NW for $j = 0, \dots, 5$ (Table 4; for odd m the supertile corresponds to the mirror image of hexagon K , and own order, then clockwise, meets

TABLE 2. Adjacency of the children's sides: the entry $c'.j'$ in row c , column j means that side j of child c is side j' of child c' ; ∂ means that it lies on the parent's rim.

every parent kind except Γ							parent kind Γ						
child	side 0	1	2	3	4	5	child	side 0	1	2	3	4	5
0	7.1	∂	∂	∂	∂	1.1	0	6.1	∂	∂	∂	∂	1.1
1	7.2	0.5	∂	2.0	3.0	4.0	1	6.2	0.5	∂	∂	2.0	3.0
2	1.3	∂	∂	∂	∂	3.1	2	1.4	∂	∂	∂	∂	3.1
3	1.4	2.5	∂	∂	∂	4.1	3	1.5	2.5	∂	4.0	5.0	6.3
4	1.5	3.5	∂	5.0	6.0	7.3	4	3.3	∂	∂	∂	∂	5.1
5	4.3	∂	∂	∂	∂	6.1	5	3.4	4.5	∂	∂	∂	6.4
6	4.4	5.5	∂	∂	∂	7.4	6	∂	0.0	1.0	3.5	5.5	∂
7	∂	0.0	1.0	4.5	6.5	∂							

TABLE 3. The parent's sides as concatenations of children's sides. An entry $c.j$ is side j of child c , which the parent's own orientation traverses from vertex $j + 1$ to vertex j of the child. Side J of the parent consists of the entries from its first position up to, but not including, the first position of side $J + 1$ (cyclically). For $T \neq \Sigma$ the parent's side 0 begins at vertex 0 of the Γ child, which is the parent's crossover corner. For $T = \Sigma$ the crossover corner lies inside side 5 (positions 19, 20, 21, 0, 1), after the children's sides 0.2, 0.1, 7.0.

Eight children (every $T \neq \Gamma$): the parent's rim, in the parent's own order from vertex 0 of child 7.

position	0	1	2	3	4	5	6	7	8	9	10
child side	7.5	6.4	6.3	6.2	5.4	5.3	5.2	5.1	4.2	3.4	3.3
position	11	12	13	14	15	16	17	18	19	20	21
child side	3.2	2.4	2.3	2.2	2.1	1.2	0.4	0.3	0.2	0.1	7.0

parent	first position of side						<i>Seven children ($T = \Gamma$), from vertex 0 of child 6; sides $0, \dots, 5$ begin at positions $0, 3, 6, 11, 16, 18$.</i>										
	0	1	2	3	4	5	position	0	1	2	3	4	5	6	7	8	9
Δ	0	5	7	11	14	19	child side	6.5	5.4	5.3	5.2	4.4	4.3	4.2	4.1	3.2	2.4
Θ	0	5	7	13	15	20	position	10	11	12	13	14	15	16	17	18	19
Λ	0	5	7	11	14	20	child side	2.3	2.2	2.1	1.3	1.2	0.4	0.3	0.2	0.1	6.0
Ξ	0	3	7	13	15	20											
Π	0	3	7	11	14	20											
Σ	2	5	7	11	14	19											
Φ	0	5	7	11	15	20											
Ψ	0	3	7	11	15	20											

its labels in the same order); its crossover corner A is vertex 0 if $K \neq \Sigma$; if $K = \Sigma$, A is not a vertex, and lies in the relative interior of side 5,

and show (*) by induction on m . For $m = 1$ it is a comparison of the level-1 supertiles, whose vertices and corner A are marked in Figure 3, with [29, Fig. 4.1]. Drawn with the edges in odd directions lengthened to $\sqrt{3}$, which turns the tiles into hats and turtles [29, §3], and in the handedness of that figure, they coincide with its nine marked clusters: the same outlines, with the vertices at the marked dots. In each cluster our vertex 0 is at the leftmost dot, and, read anticlockwise from it, the sides of cluster K carry the labels of row K of Table 4; A is that dot (for Σ , a point of the NW side, which ends there), on the light tile of the Mystic. A script checks this against a render of the figure, with the labels transcribed

TABLE 4. The label of [29, Fig. 4.2] carried by side j of a supertile of each kind: the label of the edge of the hexagon in direction W, SW, SE, E, NE, NW for $j = 0, \dots, 5$.

kind	side 0	side 1	side 2	side 3	side 4	side 5
Γ	α^-	α^+	γ^-	δ^-	β^+	β^-
Δ	γ^+	β^+	ϵ^-	α^+	γ^-	ζ^-
Θ	γ^+	β^+	θ^+	β^+	η	β^-
Λ	γ^+	β^+	ϵ^-	α^+	θ^-	β^-
Ξ	α^-	ϵ^+	θ^+	β^+	η	β^-
Π	α^-	ϵ^+	ϵ^-	α^+	θ^-	β^-
Σ	ζ^+	β^+	ϵ^-	α^+	γ^-	δ^+
Φ	γ^+	β^+	ϵ^-	ϵ^+	η	β^-
Ψ	α^-	ϵ^+	ϵ^-	ϵ^+	η	β^-

from the figure (Appendix A). Suppose (*) holds for the children of Y . The panel names the kind of each inner hexagon but does not show its edge labels; the matching rule fixes the rotation of each. By Table 4, δ occurs only on side 3 of Γ and side 5 of Σ , and ζ only on side 5 of Δ and side 0 of Σ . The Σ child has one side on the panel's rim, and the two sides next to it face siblings other than the Γ and the Δ child. Its adjacent sides 5 and 0 carry δ^+ and ζ^+ , so a sibling facing side 5 is the Γ child, and one facing side 0 the Δ child. Hence neither of them is the rim side, as the other would then be next to it and face some other sibling, and neither is next to the rim side; so they face the Γ and the Δ child, which fixes the rotations of these three. Every other child meets one of them along a side that the latter labels α^+ or γ^- , and no kind has two sides labelled α^- , or two labelled γ^+ , which fixes its rotation too (a script confirms that there is no other choice; Appendix A). With these rotations, Tables 1 to 3 transcribe the panel, including its six marked vertices; this gives (a)–(c). The crossover corner of Y is that of its Γ child, the child's vertex 0 by (*), which Table 3 places at the start of the parent's side 0 if $T \neq \Sigma$ and inside side 5 if $T = \Sigma$; this is (d). Finally, with the parent's sides numbered as in Section 2.3, the labels that the panel assigns to them are those of Table 4 for every T . So (*) holds for Y , and in particular A lies on the rim. $\square$

Remark 2.5 (Computer cross-checks of the pattern). Since the proof uses Tables 1 to 3 as transcriptions of [29, Fig. 5.1], they were also checked by computer on actual supertiles (evidence, not used in the proof). The patterns of the actual compositions $m \rightarrow m + 1$, $m \leq 7$, equal the tables (in exact arithmetic for $m \leq 3$; from level 5 on, the vertices are placed by Proposition 2.4(c), so that part is used rather than tested, though at levels 5 and 6 they agree with the vertices observed there); and Kaplan's implementation of the substitution [27] gives Table 1 and, tile for tile up to level 4, the same supertiles as `exact.py` (Appendix A).

Lemma 2.6 (Rigidity). *Any two level- m supertiles of the same kind, in possibly different homochiral tilings, are related by exactly one proper motion that maps descendants to descendants with the same numbers (children, their children, and so on, down to the tiles, and the first tile of a Mystic to the first tile). For $m \geq 1$ this motion maps the six vertices, the sides and the corner A of one supertile onto those of the other.*

Proof. Induction on m . For $m = 0$ this is clear, since a tile and a Mystic have no non-trivial proper symmetry. For $m = 1$, the level-1 supertiles are images under proper motions of the clusters of Figure 3 (Theorem 2.1(a)), with their vertices at the points marked in [29, Fig. 4.1] (Theorem 2.1(b)), their children numbered as in Figure 3, and A determined by the Mystic. Let $m \geq 1$ and let Y, Y' be level- $(m + 1)$ supertiles of kind T . By induction, child c of Y is mapped onto child c of Y' (both of kind K_c , by Proposition 2.4(a)), with its numbered vertices, by a unique proper motion g_c . If side j of child c is glued to side j' of child c' , then by Proposition 2.4(b), in Y and in Y' , vertices j and $j + 1$ of c are vertices $j' + 1$ and j' of c' ; so




FIGURE 5. Left: the decoration (centre lines of the ribbons of Figure 1). Each edge carries one crossing point (dot), where two pieces end; on e_i their ends are the ports $R(v_i)$, whose piece runs from the dot towards v_i , and $L(v_{i+1})$, whose piece runs from it towards v_{i+1} . Green: the twelve hooks $L(v_i) \leftrightarrow R(v_i)$; vermilion $R(v_0) \leftrightarrow R(v_1)$ and blue $L(v_0) \leftrightarrow L(v_1)$: the crossover at $A = v_0$, $B = v_1$. Right: the *signature* at (A, B) (Definition 4.3) on a rim of 14 vertices (dots), oriented anticlockwise as for the tile: a hook at every rim vertex but A and B , and the crossover arcs $R(A) \leftrightarrow R(B)$, $L(A) \leftrightarrow L(B)$. The tile has it (Lemma 4.4); so does every supertile of level $m \geq 1$, seen as a black box (Theorem 7.6).

g_c and $g_{c'}$ agree at two distinct points, and hence are equal. The children's adjacency graph is connected, so all the g_c are one motion g ; it maps descendants to descendants with their numbers and, by Proposition 2.4(c),(d), the vertices, sides and corner A of Y onto those of Y' . A proper motion that maps descendants to descendants maps the first tile of the Mystic at the bottom of the chain of last (Γ) children to its counterpart, so it equals g , since a tile has no symmetry. $\square$

3. THE DECORATION AND ITS STRANDS

3.1. Ports and pieces. Let t be a tile and e one of its edges, with endpoints x and y . The tile has two *ports* on e , the port *near* x and the port *near* y ; formally, they are the triples (t, e, x) and (t, e, y) . In the drawing (Figure 5), e carries one crossing point, where two ribbon ends of t meet; followed from there into t , one piece runs towards x and the other towards y . These ends are the two ports. For a vertex v_i of $\text{Tile}(1, 1)$, write

$$R(v_i) = \text{the port on } e_i \text{ near } v_i, \quad L(v_i) = \text{the port on } e_{i-1} \text{ near } v_i,$$

the ports just after and just before the vertex v_i in anticlockwise order. Each of the 28 ports is $R(v_i)$ or $L(v_i)$ for exactly one i .

Definition 3.1 (The decoration). The *decoration* pairs the 28 ports of $\text{Tile}(1, 1)$ by fourteen *pieces*:

$$\underbrace{L(v_i) \leftrightarrow R(v_i) \quad (2 \leq i \leq 13)}_{\text{twelve hooks}}, \quad \underbrace{R(v_0) \leftrightarrow R(v_1), \quad L(v_0) \leftrightarrow L(v_1)}_{\text{the crossover}}.$$

Every tile of a homochiral tiling carries the image of this decoration.

Each hook turns back round one vertex; the two crossover pieces cross once inside the tile.

3.2. The straight-through rule. Let $\mathcal{T}$ be a homochiral tiling, and let $e = xy$ be an edge shared by tiles t and t' .

Definition 3.2 (Straight-through rule). The port of t on e near x is *joined* to the port of t' on e near y , and the port of t near y to the port of t' near x .

This is what a ribbon does when it continues straight through the crossing point (Figure 6, $n = 1$). In our drawing the crossing point of e_i is at distance $0.5 + 0.15(-1)^i$ from v_i . Since e_i of one tile is always glued, traversed the other way, to some e_j of the other with $i + j$ odd (by the vertex colouring of [29, proof of Lemma 2.1]), the two tiles' crossing points on e coincide. Directly: by the turning angles in Section 2.1, the angle of $\text{Tile}(1, 1)$ at v_i is a multiple of 90° for even i and of 120° for odd i . At a vertex of a tile, every tile containing it has a vertex there, as tiles meet full edge to full edge, and their angles sum to 360° , which multiples of 90° and of 120° , some of each, never do. So the tile vertices at a point are all even or all odd, and as v_i of one tile is v_{j+1} of the other, $i + j$ is odd. The proofs use only Definitions 3.1 and 3.2, so any drawing that realises them may be used (Section 10).

3.3. Strands.

Definition 3.3 (Strand graph). Let $\mathcal{T}$ be a homochiral tiling. The *strand graph* $\mathcal{G}(\mathcal{T})$ is the multigraph with one vertex for each piece of each tile and, for each pair of ports joined by the straight-through rule, one edge between the two pieces that end at these ports. A *strand* is a connected component of $\mathcal{G}(\mathcal{T})$; a strand is *closed* if it is finite.

Lemma 3.4. Every vertex of $\mathcal{G}(\mathcal{T})$ has degree 2. Hence every strand is either a finite cycle (a closed strand) or a two-way infinite path.

Proof. A piece has two ends, each a port of its tile on some edge e . Tiles meet full edge to full edge (Section 2), so in a tiling of the plane e is an edge of exactly one other tile, and each end is joined to exactly one port, an end of one of that neighbour's pieces. Hence every vertex has degree exactly 2; joins always link pieces of different tiles, so $\mathcal{G}(\mathcal{T})$ has no self-loops. A connected, locally finite multigraph with all degrees 2 and no self-loops is a finite cycle (possibly of length 2) or a two-way infinite path. $\square$

Equivalently, in a tiling of the plane the pieces and the straight-through rule each pair off the ports of all the tiles, giving two fixed-point-free involutions; the ports of a strand form an orbit of the group they generate. This is how the formalisation describes strands.

4. BOXES AND THE SIGNATURE

4.1. Patches and boxes. A *patch* is a finite set P of tiles of a homochiral tiling whose union is a closed disc with a simple polygonal boundary, the *rim* ∂P . Every supertile of level ≥ 1 is a patch (Theorem 2.1), and so is a tile. Every rim edge belongs to exactly one tile of P . Every other edge of a tile of P is an *interior edge*, shared by two tiles of P : its relative interior lies inside the disc, so the tile across it is in P . Let $E = E(P)$ be the number of rim edges, and number the rim vertices $u_0, \dots, u_{E-1}$ in a chosen orientation of the rim, indices mod E ; for a supertile we use its own orientation unless we say otherwise. For a rim vertex u_k define the *rim ports*

$$R(u_k) = \text{the port on } u_k u_{k+1} \text{ near } u_k, \quad L(u_k) = \text{the port on } u_{k-1} u_k \text{ near } u_k,$$

where the port belongs to the tile of P containing that edge. There are $2E$ rim ports. For a single tile, in its own (anticlockwise) orientation, these are the ports $R(v_i), L(v_i)$ of Section 3. Reversing the orientation exchanges R and L at every rim vertex.

Lemma 4.1 (The box of a patch). Start at a rim port of P and follow the pieces of the tiles of P , crossing interior edges of P by the straight-through rule. The path ends at a different rim port. This defines a fixed-point-free involution π_P of the set of rim ports. The pieces of P that are not reached from any rim port form finitely many cycles, the closed strands of P .

Proof. The pieces of P , joined across interior edges, form a finite graph in which a piece has one neighbour for each of its ends that is not a rim port. Its components are therefore cycles, and paths (possibly a single piece) with exactly two free ends, which are distinct rim ports. Every rim port is a free end of exactly one such path, so sending each rim port to the other free end of its path is a fixed-point-free involution, π_P . $\square$

We call the paths of [Lemma 4.1](#) the *arcs* of P ; there are exactly E of them. The rim of P with the pairing π_P is the *box* of P . All we need to know about the inside of P is its box and whether P has a closed strand; the next lemma computes both from smaller patches.

Lemma 4.2 (Gluing). *Let a patch P be the union of patches $P_1, \dots, P_n$ with pairwise disjoint interiors. Then the box of P is determined by the boxes of the P_i together with the identification of their shared rim edges. More precisely, π_P is computed by alternately applying the π_{P_i} and the straight-through rule across shared edges, until a rim port of P is reached. The closed strands of P are those of the P_i together with the cycles of this alternation that contain no rim port of P .*

Proof. Every arc of P ([Lemma 4.1](#)) is a concatenation of arcs of the P_i joined across edges shared by two of the P_i ; every closed strand of P is either a closed strand of some P_i or a cyclic concatenation of such arcs. The rim edges of P are the rim edges of the P_i that are not shared by two of them. $\square$

We use [Lemma 4.2](#) in two ways: to define the box of a composition, glued from boxes alone ([Section 7.1](#)), and to compute the box of a supertile from those of its children. For the first, a *box* is taken abstractly: a cyclically ordered set of rim vertices u , each with two ports $L(u), R(u)$, together with a fixed-point-free involution π of these ports; conditions (i) and (ii) of [Definition 4.3](#) make sense for any box.

4.2. The signature.

Definition 4.3 (Signature). Let A and B be adjacent rim vertices of a patch P . We say that P has the *signature* at (A, B) if

- (i) $\pi_P(L(u)) = R(u)$ for every rim vertex $u \notin \{A, B\}$ (a *hook* at u ; we also say that u is a hook of P , and call the arc of P from $L(u)$ to $R(u)$ its *hook arc* at u);
- (ii) $\pi_P(R(A)) = R(B)$ and $\pi_P(L(A)) = L(B)$ (the *crossover*);
- (iii) P has no closed strand: the ribbon followed from any port of P reaches the rim.

The definition is symmetric in A and B and does not depend on the orientation of the rim, since reversing it exchanges L and R . A patch with the signature has $E - 2$ hook arcs, one at each rim vertex other than A and B , and two crossover arcs. [Figure 5](#) (right) shows the signature.

Lemma 4.4. *A tile has the signature at (v_0, v_1) .*

Proof. A tile has no interior edge, so π is the pairing of [Definition 3.1](#), and every piece has both ends on the rim. $\square$

5. THE ZIGZAG LEMMA

In this section a patch P is the union of patches X_c (*children*) with disjoint interiors, each X_c having the signature at some (A_c, B_c) , its *crossover corners*. The children and P are all oriented in the same rotational sense, so two distinct children traverse a common rim path in opposite directions; $R_c(u), L_c(u)$ are the rim ports of X_c at u in this orientation. (When P is a supertile and the X_c are its children, the common orientation is the children's own orientation, which is opposite to the own orientation of P ([Definition 2.2](#)); this exchanges L and R at the rim vertices of P but does not affect whether P has the signature.) Along a shared path the ribbon zigzags between two children, and only the parity of the path's length matters ([Figure 6](#)).




FIGURE 6. The Zigzag Lemma. Along a shared path whose interior vertices are hooks of both children, the ribbons cross the path at every edge and turn back at every interior vertex, alternately in X_c (above) and $X_{c'}$ (below). The orange connection starts at $R_c(u_0)$ (top left) and the blue one at $L_{c'}(u_0)$ (bottom left). For odd n they end in the other child, at $R_{c'}(u_n)$ (bottom right) and $L_c(u_n)$ (top right); for even n they end in the child they started from, at $L_c(u_n)$ and $R_{c'}(u_n)$. Only the parity of n matters, so lengthening a shared path by two edges changes nothing (Lemma 7.1).

Lemma 5.1 (Zigzag Lemma). *Let $u_0u_1 \cdots u_n$ ($n \geq 1$) be a path of n edges lying on the rims of two distinct children X_c and $X_{c'}$, traversed by X_c from u_0 to u_n . Suppose that none of the interior vertices $u_1, \dots, u_{n-1}$ is a crossover corner of X_c or of $X_{c'}$. Then alternately applying the straight-through rule across the edges u_ku_{k+1} and the hooks of X_c and $X_{c'}$ at $u_1, \dots, u_{n-1}$ connects the ports of X_c and $X_{c'}$ near u_0 on u_0u_1 and near u_n on $u_{n-1}u_n$ in pairs as follows:*

$$\begin{aligned} n \text{ even: } & R_c(u_0) \leftrightarrow L_c(u_n), & L_{c'}(u_0) \leftrightarrow R_{c'}(u_n); \\ n \text{ odd: } & R_c(u_0) \leftrightarrow R_{c'}(u_n), & L_{c'}(u_0) \leftrightarrow L_c(u_n). \end{aligned}$$

These two connections use every port on the edges of the path.

For $n = 1$ the lemma is the straight-through rule (Definition 3.2) written in rim ports; Lemmas 7.1 and 9.3 use it for $n = 1$ and $n = 3$.

Proof. Start at $R_c(u_0)$, the port of X_c on u_0u_1 near u_0 (orange in Figure 6). The straight-through rule takes it to the port of $X_{c'}$ on u_0u_1 near u_1 . In the orientation of $X_{c'}$, which runs from u_1 to u_0 along this edge, that port is $R_{c'}(u_1)$. If $n = 1$ we have reached $R_{c'}(u_n)$. Otherwise u_1 is a hook of $X_{c'}$, which leads to $L_{c'}(u_1)$, the port of $X_{c'}$ on u_1u_2 near u_1 . Crossing u_1u_2 leads to the port of X_c on u_1u_2 near u_2 , which is $L_c(u_2)$. If $n = 2$ we stop. Otherwise the hook of X_c at u_2 leads to $R_c(u_2)$, and the pattern repeats with all indices shifted by two. So this connection visits the hooks of $X_{c'}$ at odd k and those of X_c at even k , and ends at $L_c(u_n)$ if n is even and at $R_{c'}(u_n)$ if n is odd. The second connection starts at $L_{c'}(u_0)$, the port of $X_{c'}$ on u_0u_1 near u_0 . By the same argument with $X_c, X_{c'}$ exchanged and L, R exchanged, it passes $L_c(u_1), R_c(u_1), R_{c'}(u_2), L_{c'}(u_2), \dots$, visits the hooks of X_c at odd k and those of $X_{c'}$ at even k , and ends at $R_{c'}(u_n)$ if n is even and at $L_c(u_n)$ if n is odd. Each connection crosses each of the n edges once, using two of its four ports, and the two connections use different ports, so together they use all of them. $\square$

6. SIDE LENGTHS

By the Zigzag Lemma, two boxes glued along a path with hooks at its interior vertices interact only through the parity of its length. This section shows that the side lengths of the supertiles are given by Fibonacci numbers (Proposition 6.2), so that they grow with the level, and never change parity (Corollary 6.3); it also places the crossover corner (Lemma 6.4). Section 7 uses these facts to reduce every level to level 1.

For $m \geq 1$, write $\ell_m(K, j)$ for the number of edges of side j of a level- m supertile of kind K ; by Lemma 2.6 it is well defined.

6.1. Letters. By (*) in the proof of Proposition 2.4, every side of a level- m supertile ($m \geq 1$) carries the edge label of [29, Fig. 4.2] that Table 4 lists. The *letter* of a side is its label without the sign; there are eight letters, $\alpha, \beta, \gamma, \delta, \epsilon, \zeta, \eta, \theta$. By the matching rule (Section 2.2), two

sides glued to each other have the same letter; for sides glued inside a supertile this can also be read off [Tables 1, 2 and 4](#).

Proposition 6.1 (Lengths by letter). *For every $m \geq 1$ the length $\ell_m(K, j)$ depends only on the letter of side j of K . Write x_m for the common length of the level- m sides with letter x . Then*

$$(1) \quad \begin{aligned} \alpha_{m+1} &= \beta_m + \epsilon_m + \eta_m, & \beta_{m+1} &= \alpha_m + \epsilon_m, \\ \gamma_{m+1} &= \beta_m + 2\epsilon_m + \eta_m + \theta_m, & \delta_{m+1} &= \alpha_m + \beta_m + \epsilon_m + 2\theta_m, \\ \epsilon_{m+1} &= 3\epsilon_m + \eta_m, & \zeta_{m+1} &= \alpha_m + \epsilon_m + \theta_m, \\ \eta_{m+1} &= 4\epsilon_m + \eta_m, & \theta_{m+1} &= 4\epsilon_m + \eta_m + \theta_m, \end{aligned}$$

with $(\alpha_1, \beta_1, \gamma_1, \delta_1, \epsilon_1, \zeta_1, \eta_1, \theta_1) = (7, 5, 10, 10, 8, 6, 10, 11)$.

Proof. For $m = 1$ this is a direct count on the level-1 supertiles ([Figure 3](#)), recorded in [Computation 7.5](#). By [Proposition 2.4\(c\)](#), $\ell_{m+1}(T, J)$ is the sum of the lengths of the children's sides listed in [Table 3](#). Inspection of the 54 parent sides in [Tables 1, 3 and 4](#) (also checked by a script; [Appendix A](#)) shows that the multiset of letters of these children's sides depends only on the letter of side J of T , and is the one in (1). For example, side 0 of Δ (letter γ) is made of 7.5, 6.4, 6.3, 6.2, 5.4, with letters $\beta, \eta, \epsilon, \epsilon, \theta$. The claim follows by induction on m . $\square$

The recurrence (1) has a closed form. Let F_n be the Fibonacci numbers, $F_1 = F_2 = 1$.

Proposition 6.2 (Fibonacci side lengths). *For every $m \geq 1$,*

$$\begin{aligned} \alpha_m &= F_{3m+2} + 2, & \beta_m &= F_{3m+1} + 2, & \gamma_m &= \delta_m = \eta_m = 2F_{3m+2}, \\ \epsilon_m &= F_{3m+3}, & \zeta_m &= 2F_{3m+1}, & \theta_m &= F_{3m+4} - 2. \end{aligned}$$

Proof. For $m = 1$, $(F_4, F_5, F_6, F_7) = (3, 5, 8, 13)$ gives the initial values. For the induction step put $a = F_{3m+1}$ and $b = F_{3m+2}$, so that $F_{3m+3} = a + b$, $F_{3m+4} = a + 2b$, $F_{3m+5} = 2a + 3b$, $F_{3m+6} = 3a + 5b$ and $F_{3m+7} = 5a + 8b$. By the induction hypothesis, $\alpha_m = b + 2$, $\beta_m = a + 2$, $\gamma_m = \delta_m = \eta_m = 2b$, $\epsilon_m = a + b$, $\zeta_m = 2a$ and $\theta_m = a + 2b - 2$. Substituting into (1):

$$\begin{aligned} \alpha_{m+1} &= (a + 2) + (a + b) + 2b = F_{3m+5} + 2, & \beta_{m+1} &= (b + 2) + (a + b) = F_{3m+4} + 2, \\ \gamma_{m+1} &= (a + 2) + 2(a + b) + 2b + (a + 2b - 2) = 4a + 6b = 2F_{3m+5}, \\ \delta_{m+1} &= (b + 2) + (a + 2) + (a + b) + 2(a + 2b - 2) = 4a + 6b = 2F_{3m+5}, \\ \eta_{m+1} &= 4(a + b) + 2b = 4a + 6b = 2F_{3m+5}, & \epsilon_{m+1} &= 3(a + b) + 2b = F_{3m+6}, \\ \zeta_{m+1} &= (b + 2) + (a + b) + (a + 2b - 2) = 2F_{3m+4}, \\ \theta_{m+1} &= 4(a + b) + 2b + (a + 2b - 2) = 5a + 8b - 2 = F_{3m+7} - 2. \end{aligned} \quad \square$$

Since F_n is even exactly when $3 \mid n$ (the parities of $F_1, F_2, \dots$ run odd, odd, even, and repeat), [Proposition 6.2](#) gives:

Corollary 6.3 (Parities are constant). *For every $m \geq 1$ the lengths $\alpha_m, \beta_m, \theta_m$ are odd and $\gamma_m, \delta_m, \epsilon_m, \zeta_m, \eta_m$ are even.*

6.2. The crossover corner.

Lemma 6.4 (Position of A). *Let $m \geq 1$. For $K \neq \Sigma$, the crossover corner A of a level- m K -supertile is its vertex 0. For $K = \Sigma$, A lies inside side 5 (letter δ), d_m edges along it from vertex 5 and $\delta_m - d_m = 2F_{3m}$ edges from vertex 0, where*

$$d_m = 2F_{3m+1}.$$

Proof. For $m = 1$ the lemma is part of [Computation 7.5](#): A is vertex 0 if $K \neq \Sigma$, and $d_1 = 6 = 2F_4$. For $m \geq 2$ and $K \neq \Sigma$ it is [Proposition 2.4\(d\)](#). For $K = \Sigma$: by [Table 3](#), which lists the children's sides in the parent's own order, side 5 of a level- $(m + 1)$ Σ -supertile begins, at its vertex 5, with the children's sides 0.2, 0.1, 7.0, and by [Proposition 2.4\(d\)](#) its crossover

corner comes right after them. Child 0 is a Ξ and child 7 a Γ , so these sides have letters θ, ϵ, α , and by Proposition 6.2 $d_{m+1} = \theta_m + \epsilon_m + \alpha_m = (F_{3m+4} - 2) + F_{3m+3} + (F_{3m+2} + 2) = 2F_{3m+4}$. Finally, $\delta_m - d_m = 2F_{3m+2} - 2F_{3m+1} = 2F_{3m}$ by Proposition 6.2. $\square$

Lemma 6.5 (Lower bounds). *For every $m \geq 1$: every side has length at least 5; $d_m \geq 6$; and $\delta_m - d_m \geq 4$.*

Proof. By Proposition 6.2 all letter lengths increase with m , and at $m = 1$ the smallest is $\beta_1 = 5$. Also $d_m \geq d_1 = 6$ and $\delta_m - d_m = 2F_{3m} \geq 2F_3 = 4$. $\square$

7. THE SIGNATURE AT EVERY LEVEL

We prove the Signature Theorem by induction on the level, reducing every step to a finite computation at level 1, and deduce Theorem A from it (Section 7.4).

7.1. Compositions as a function of side lengths. A box with the signature at (A, B) is determined by its rim and the places of A and B . So, once the children of a supertile have the signature, the box of the supertile depends, as we show below, only on the tables, the children's side lengths, the sign that places each B_c next to A_c , and the place of the corner A_Σ of the Σ child, the one child whose corner A is not a vertex (Lemma 6.4); d below records where on its side 5 it lies.

Compositions. Fix a parent kind T , with children $c = 0, \dots, n_T - 1$ of kinds K_c (Table 1), exactly one of which, c_Σ , is a Σ , and exactly one, the last, a Γ . Let L assign to each child c and side $j \in \{0, \dots, 5\}$ a positive integer $L(c, j)$, with equal values on sides glued by Table 2; let d be an integer with $2 \leq d \leq L(c_\Sigma, 5) - 2$; and let $\omega \in \{\pm 1\}$. The *composition* $C_T(L, d, \omega)$ is the following abstract patch of boxes. Child c has $E_c = \sum_j L(c, j)$ rim vertices, numbered $0, \dots, E_c - 1$ in its own order; its vertex j (one of the six) is rim vertex $V_c(j) = \sum_{i < j} L(c, i)$, and side j runs from $V_c(j)$ to $V_c(j + 1)$. Its crossover corner A_c is rim vertex 0, or, for $c = c_\Sigma$, rim vertex $V_c(5) + d$, at distance d from vertex 5 along side 5; $B_c = A_c + \omega$ (as $2 \leq d \leq L(c_\Sigma, 5) - 2$, the corners $A_\Sigma = A_{c_\Sigma}$ and B_{c_Σ} lie strictly inside side 5); and it is a box satisfying (i) and (ii) of Definition 4.3 at (A_c, B_c) . Sides are glued in pairs as in Table 2, reversing direction: the r -th edge of side j of c is the r -th edge from the end of side j' of c' . The sides marked ∂ form the rim of the composition, in the order of Table 3 reversed (the table follows the parent's own orientation; the composition, like its children, has the opposite one); the ports of a rim vertex are those of the children on its two rim edges. Following ribbons alternately through a child, by its pairing, and straight across a glued edge, as in Lemma 4.2, defines the box of the composition and its *closed strands*, the cycles of this alternation that contain no rim port of the composition. The Zigzag Lemma and its proof apply unchanged to compositions, since they use only the pairings of the children and the straight-through rule. We say that $C_T(L, d, \omega)$ *has the signature* if the crossover corners A_Γ, B_Γ of its Γ child are adjacent rim vertices of the composition and it has the signature at (A_Γ, B_Γ) (Definition 4.3, applied to its box and its closed strands).

Write ℓ_m also for the function $(c, j) \mapsto \ell_m(K_c, j)$; by Section 6 and Lemma 6.5, $C_T(\ell_m, d_m, \omega)$ is defined.

Reduction (A supertile is its composition). *Let $m \geq 1$, and let Y be a level- $(m + 1)$ supertile of kind T whose children X_c have the signature at $(A_c, A_c + \omega_m)$ in their own orientation. Then the children of Y , with their rims, sides, corners A_c, B_c and the gluing of their sides, and the rim of Y read in the children's order, match those of $C_T(\ell_m, d_m, \omega_m)$. So the box of Y is the box of the composition, and Y has a closed strand if and only if the composition has one.*

Indeed: (i) by Proposition 2.4(a),(b) and Proposition 6.1 the children have the kinds and side lengths of the composition, and side j of X_c is side j' of $X_{c'}$ traversed oppositely; (ii) two children meet only along the closed sides they share: at a point inside a side of a child, the rim edges there are shared with the partner across that side, so the child and its partner fill a neighbourhood; and at a vertex exactly three level- m supertiles meet, any two of them

sharing a side that ends there ([Theorem 2.1\(b\)](#)); (iii) so each side of a child is wholly shared with one sibling or wholly on the rim, as in [Table 2](#); (iv) A_c is placed by [Lemma 6.4](#), and $B_c = A_c + \omega_m$ by hypothesis; (v) the rim of the parent is given by [Table 3](#) ([Proposition 2.4\(c\)](#)). By [Lemma 4.2](#), the strands of the supertile are then those of the composition, and its box is the box of $C_T(\ell_m, d_m, \omega_m)$. Step (ii) is the one geometric step, argued here: it ensures that the children touch only along the gluings of (i), so that no other edges join their ribbons. With [Theorem 2.1](#), [Proposition 2.4](#) and [Lemma 2.6](#) it makes up part (M1) of the link with actual tilings, which the formalisation does not prove but takes as given ([Appendix B](#)). The compositions themselves are built from the tables alone, and the tables are consistent: [Table 3](#) lists exactly once each side that [Table 2](#) puts on the rim, and turning about the end vertex of each entry, across sides that [Table 2](#) glues, leads to the start of the next entry (a finite check, made in the formalisation). For instance, 7.5 ends at vertex 5 of child 7; side 4 of child 7 is side 5 of child 6, so this is vertex 5 of child 6, whose side 4 lies on the rim and is the next entry, 6.4.

7.2. Lengthening. Let (L, d) be as above. A *move* lengthens one side of one child by two edges, together with the side glued to it if there is one. In the composition, an edge e of that side is replaced by a path of three edges through two new vertices; e must be neither the first nor the last edge of its side (in each child that contains it), nor one of the two edges at A_Σ . This ensures that the edge between A_c and B_c is never cut, as the proof of [Lemma 7.1](#) needs: for a child other than the Σ child, A_c is vertex 0, the end of side 5 and the start of side 0, and B_c is next to it. The result is (L', d') : L' is L with the side, and its partner, two edges longer; $d' = d + 2$ if e lies on side 5 of the Σ child between its vertex 5 and A_Σ , and $d' = d$ otherwise.

Lemma 7.1 (Lengthening). *If $C_T(L, d, \omega)$ has the signature, then so does $C_T(L', d', \omega)$, for every move $(L, d) \rightarrow (L', d')$.*

Proof. The idea: on a shared edge, cutting e replaces one crossing by a zigzag along three edges, which joins the same ports ([Lemma 5.1](#)); on the rim it adds two hooks. Write $C = C_T(L, d, \omega)$ and $C' = C_T(L', d', \omega)$, and let $e = u_0u_1$. Then C' is C with e cut into three: in each child that contains it, e is replaced by the path $u_0w_1w_2u_1$, and every vertex of C has an image vertex in C' (the numbers of the rim vertices after e go up by two). This is only a relabelling: the side and its partner gain the two new edges, every other side keeps its length, and the r -th edge of a side is still glued to the r -th from the end of its partner. The conditions on e make the crossover corners of C' the images of those of C : for A_Σ by the choice of d' , and $B_c = A_c \pm 1$ stays next to A_c because the edge between them (the first edge of side 0, the last of side 5, or an edge at A_Σ) is never cut. In particular w_1 and w_2 are not crossover corners. Let φ map each port of C to the port of C' of the same child and the same hand (R or L) at the image vertex. Then φ commutes with the pairings inside the children, since the crossover corners correspond, and with the straight-through rule across every glued edge other than e . If e is shared by children c and c' , with c traversing it from u_0 to u_1 , crossing it joins $R_c(u_0)$ to $R_{c'}(u_1)$ and $L_c(u_1)$ to $L_{c'}(u_0)$ ([Lemma 5.1](#) with $n = 1$); in C' , by [Lemma 5.1](#) with $n = 3$, the ribbons run along the new path from $\varphi(R_c(u_0))$ to $\varphi(R_{c'}(u_1))$ and from $\varphi(L_c(u_1))$ to $\varphi(L_{c'}(u_0))$. If e lies on the rim, w_1 and w_2 are new rim vertices of C' , and hooks of the child, so hooks of C' .

Hence, by induction along the ribbon (each of its steps in C , through a child or across a glued edge, is matched in C' , across e by the path along $u_0w_1w_2u_1$), for every port x of C , the ribbon of C' from $\varphi(x)$ passes through the images of the ports that the ribbon of C from x meets, and reaches the rim at the image of the rim port reached in C . So C' has a hook at the image of every rim vertex of C other than A_Γ and B_Γ , and at each new rim vertex. As before, A_Γ and B_Γ are adjacent; they are still rim vertices, since the rim edges of the Γ child at them stay on the rim, and they carry the crossover. Finally, C' has no closed strand. Every port of C' at the image of a vertex of C is some $\varphi(x)$, whose ribbon reaches the rim.

A port at w_1 or w_2 lies on the hook there if e is on the rim, and otherwise on one of the two connections of [Lemma 5.1](#) along $u_0w_1w_2u_1$; these use every port on these edges and end at ports at u_0 and u_1 , so its ribbon too reaches the rim. $\square$

Call (L, d) *admissible* if sides glued by [Table 2](#) have the same length; $L(c, j) \geq \ell_1(c, j)$ and $L(c, j) \equiv \ell_1(c, j) \pmod{2}$ for all c and j ; $d \geq d_1$ and $d \equiv d_1 \pmod{2}$; and $L(c_\Sigma, 5) - d \geq \ell_1(c_\Sigma, 5) - d_1 = 4$. Then $d \geq 6$ and $L(c_\Sigma, 5) - d \geq 4$, so the composition $C_T(L, d, \omega)$ is defined. Every admissible pair is reached from (ℓ_1, d_1) by moves ([Lemma 7.2](#)), which only lengthen sides, two edges at a time. Admissible pairs need not be the side lengths of actual supertiles, since L may give different lengths to sides with the same letter.

Lemma 7.2 (Reachability). (a) For every $m \geq 1$, the pair (ℓ_m, d_m) is admissible. (b) Every admissible $(L, d) \neq (\ell_1, d_1)$ arises by a move from an admissible pair $(\bar{L}, \bar{d})$ with $\sum_{c,j} \bar{L}(c, j) < \sum_{c,j} L(c, j)$. The move can be chosen so that its edge e lies at least two edges from both ends of its side (at least two edges of the side precede e and at least two follow it), in each child that contains it, unless that side is side 5 of the Σ child or is glued to it. (This is used only for the junction map of [Section 9.2](#).)

Hence a property of pairs that holds at (ℓ_1, d_1) and is kept by the moves of (b) holds at every admissible pair, in particular at (ℓ_m, d_m) for every $m \geq 1$.

Proof. (a) Glued sides have the same letter ([Section 6](#)), so by [Proposition 6.1](#) the same length. By [Proposition 6.2](#) the letter lengths do not decrease with m , and by [Corollary 6.3](#) their parities do not change. By [Lemma 6.4](#), $d_m = 2F_{3m+1}$ is even, and by [Lemma 6.5](#), $d_m \geq 6 = d_1$ and $\ell_m(c_\Sigma, 5) - d_m = \delta_m - d_m \geq 4$.

(b) If $L = \ell_1$, then $L(c_\Sigma, 5) = \delta_1 = 10$ and $d_1 \leq d \leq 10 - 4 = d_1$, so $(L, d) = (\ell_1, d_1)$, which is excluded. Hence $L \neq \ell_1$, and by the parity condition some side j of some child c has $L(c, j) \geq \ell_1(c, j) + 2$; we undo a move on that side. If the side is neither side 5 of the Σ child nor glued to it, take $\bar{d} = d$ and the move whose edge is the third edge of the side (the third from the end in the partner): before the move the side has length at least $\ell_1(c, j) \geq 5$ ([Lemma 6.5](#)), so this edge lies at least two edges from both ends, in either child. Otherwise, measure distances along side 5 of the Σ child from its vertex 5. If $d \geq d_1 + 2$, take $\bar{d} = d - 2$ and the move whose edge is the second edge of the side, before A_Σ ; it does not touch A_Σ , since $\bar{d} \geq 6$. If $d = d_1$, take $\bar{d} = d$ and the move whose edge is the second edge after A_Σ : since $L(c_\Sigma, 5) \geq \ell_1(c_\Sigma, 5) + 2$ by the choice of the side, before the move it has $L(c_\Sigma, 5) - 2 - d \geq \ell_1(c_\Sigma, 5) - d_1 = 4$ edges after A_Σ , so this edge is at least two edges from the end of the side. In each case $(\bar{L}, \bar{d})$ is admissible: the side, and the side glued to it if there is one, are two edges shorter, so still equal, at least ℓ_1 and of its parity, as $L(c, j) \geq \ell_1(c, j) + 2$; $\bar{d} \geq d_1$ and $\bar{d} \equiv d_1 \pmod{2}$; and $\bar{L}(c_\Sigma, 5) - \bar{d}$ is $L(c_\Sigma, 5) - d \geq 4$ in the first two cases and at least 4 in the third, as shown. And $\sum \bar{L} = \sum L - 2$ or $\sum L - 4$. The final assertion of the lemma follows by induction on $\sum_{c,j} L(c, j)$. $\square$

Computation 7.3 (The compositions at the level-1 lengths). For each of the nine kinds T and both signs ω , the composition $C_T(\ell_1, d_1, \omega)$ has the signature.

This was computed by `hierarchy/compositions.py`, which builds the 18 compositions from [Tables 1 to 4](#) and ℓ_1, d_1 alone and follows the ribbon from every port; the Lean kernel repeats it ([Appendix A](#)). Both signs are needed, although a level-2 supertile is glued as $C_T(\ell_1, d_1, \omega_1)$ with $\omega_1 = -1$ ([Section 7.1](#)): moves keep ω , whereas the sign ω_m used at level $m + 1$ alternates with m .

Most of the signature of a composition is automatic. Only the Γ child has its crossover corners on the rim: its sides 5 and 0, which meet at A_Γ , lie on the rim ([Table 2](#)), while the crossover corners of every other child lie on sides it shares with siblings, away from the rim. So A_Γ, B_Γ carry the crossover of the Γ child, and every other rim vertex whose two rim edges belong to one child is a hook of that child. The computation decides the rest: the hooks at

the eight rim vertices (seven for $T = \Gamma$) where the rim passes from one child to another, and that no strand closes.

Proposition 7.4 (Compositions at every level). *For every $m \geq 1$, every kind T and both signs ω , the composition $C_T(\ell_m, d_m, \omega)$ has the signature.*

Proof. Fix T and ω , and consider the property of (L, d) that $C_T(L, d, \omega)$ has the signature. It holds at (ℓ_1, d_1) by [Computation 7.3](#), and every move keeps it by [Lemma 7.1](#); so by [Lemma 7.2](#) it holds at (ℓ_m, d_m) . $\square$

7.3. The induction.

Computation 7.5 (Base case). *Each of the nine level-1 supertiles of [Figure 3](#), and so by [Lemma 2.6](#) every level-1 supertile, has the side lengths $\ell_1(K, j) = x_1$, where x is the letter of side j ([Table 4](#)) and $(\alpha_1, \dots, \theta_1) = (7, 5, 10, 10, 8, 6, 10, 11)$; its crossover corner A is vertex 0 if its kind is not Σ , and for Σ lies inside side 5, $d_1 = 6$ edges along it from vertex 5; and it has the signature at (A, B) , where $B = A + \omega_1$ in its own order.*

This was computed from the tiles in exact arithmetic by the program `exact/exact.py`, which uses [Lemma 4.2](#) and checks each clause of the statement on the nine supertiles; `hierarchy/checks.js` checks the signature again in floating point. The kernel of Lean repeats the check on the nine level-1 supertiles computed by `exact.py`, taken as input, so it checks neither their geometry nor that A is v_0 of the first tile of its Mystic ([Appendices A](#) and [B](#)). The induction starts here, not at the tile ([Lemma 4.4](#)), because the composition pattern describes level $m + 1$ from level $m \geq 1$ only; level 1 is assembled from tiles and a Mystic.

Theorem 7.6 (Signature Theorem). *In every homochiral tiling by $\text{Tile}(1, 1)$, for every $m \geq 1$, every level- m supertile of every kind has the signature at (A, B) , where A is its crossover corner ([Section 2.3](#)), placed as in [Lemma 6.4](#), and $B = A + \omega_m$ in its own order. In particular no such supertile contains a closed strand.*

Proof. For $m = 1$ this is [Computation 7.5](#). Assume the theorem for level $m \geq 1$ and let Y be a level- $(m + 1)$ supertile of kind T . By the induction hypothesis its children have the signature at $(A_c, A_c + \omega_m)$. So, by the Reduction ([Section 7.1](#), with [Lemma 4.2](#)), the box of Y is that of $C_T(\ell_m, d_m, \omega_m)$, and Y has a closed strand only if this composition has one; the composition has the signature by [Proposition 7.4](#). Hence Y has no closed strand, and it has a hook at every rim vertex except the adjacent crossover corners A_Γ, B_Γ of its Γ child, which carry the crossover. By definition A_Γ is the crossover corner of Y , placed as in [Lemma 6.4](#). Finally $B_\Gamma = A_\Gamma + \omega_m$ in the child's own order, which is opposite to that of Y ; so $B = A - \omega_m = A + \omega_{m+1}$ in the own order of Y . $\square$

7.4. No closed strands. The hypothesis that the hierarchy covers the plane is used only to place a closed strand inside one supertile, where [Theorem 7.6](#) excludes it.

Proof of [Theorem A](#). Let $\mathcal{T}$ be a homochiral tiling whose hierarchy covers the plane: for some tile t , the supertiles $X_0(t) \subset X_1(t) \subset \dots$ containing t have union $\mathbb{R}^2$. Write $X_m = X_m(t)$.

Step 1 (in the plane). Every tile t' lies in some X_m , and every edge of t' is shared by two tiles of some X_m . Indeed, a point x of the interior of t' lies in X_m for some m , hence in some tile s of X_m . Every neighbourhood of x meets the interior of s , so the interiors of s and t' meet, and $s = t'$. An edge of t' is shared with exactly one other tile (tiles meet full edge to full edge, [Section 2](#); see the proof of [Lemma 3.4](#)), and since the X_m are nested, both tiles lie in one X_m .

Step 2 (combinatorial). Let $X_1 \subset X_2 \subset \dots$ be any supertiles, each a child of the next, such that every edge of every tile of their union is shared by two tiles of some X_m ; we show that every strand of the union is a two-way infinite path. A closed strand has finitely many pieces and is joined across finitely many edges, so there is an m such that X_m contains all its

tiles and each of these edges is shared by two tiles of X_m . It is then a closed strand of X_m , contradicting [Theorem 7.6](#). So there is no closed strand, and by the argument of [Lemma 3.4](#) every strand of the union is a two-way infinite path.

By Step 1, $X_1(t) \subset X_2(t) \subset \dots$ is such a chain, and its union is $\mathcal{T}$; this proves [Theorem A](#). Step 2 is [Theorem A](#) as the formalisation states it ([Appendix B](#)). $\square$

8. THE NEST TILING

This section builds the nest tiling $\mathcal{N}$ of [Theorem B](#) from Δ -supertiles $N_0 \subset N_1 \subset \dots$, each child number 1 of the next (child 1 of every supertile is a Δ , [Table 1](#)), and shows that it is a tiling of the plane ([Proposition 8.4](#)) with no closed strand. So as not to presuppose a tiling that contains the whole chain, each N_b is built in some tiling and moved into place ([Construction 8.1](#) and [Proposition 8.2](#)).

Construction 8.1 (The nest). Let $\mathcal{T}$ be a homochiral tiling (one exists [[29](#), Thm. 2.2]). For each $h \geq 0$, let $Y^{(h)}$ be a level- h Δ -supertile of $\mathcal{T}$; one exists, for instance child 1 of any level- $(h+1)$ supertile. Let $Y^{(h)} = Y_h^{(h)} \supset Y_{h-1}^{(h)} \supset \dots \supset Y_0^{(h)}$ be its chain of descendants, $Y_b^{(h)}$ being child 1 of $Y_{b+1}^{(h)}$ for $b < h$. Let g_h be the proper motion taking the tile $Y_0^{(h)}$ onto $\text{Tile}(1, 1)$ (whose vertex v_0 is the origin), and put $N_b^{(h)} = g_h(Y_b^{(h)})$ for $b \leq h$.

Proposition 8.2. $N_b^{(h)}$ depends neither on $h \geq b$ nor on the choices made. Write N_b for it; then $N_0 = \{\text{Tile}(1, 1)\}$, and N_b is child 1 of N_{b+1} for every b .

Proof. Let $h, h' \geq b$, with $Y^{(h)}, Y^{(h')}, g_h, g_{h'}$ as in [Construction 8.1](#). By [Lemma 2.6](#) there is a proper motion f with $f(Y_b^{(h)}) = Y_b^{(h')}$ that maps descendants to descendants with their numbers; hence $f(Y_{b'}^{(h)}) = Y_{b'}^{(h')}$ for all $b' \leq b$, in particular $f(Y_0^{(h)}) = Y_0^{(h')}$. So $g_{h'}f$ and g_h both take the tile $Y_0^{(h)}$ onto $\text{Tile}(1, 1)$; since a tile has no non-trivial symmetry, $g_{h'}f = g_h$. Thus $N_b^{(h')} = g_{h'}f(Y_b^{(h)}) = N_b^{(h)}$. As f maps descendants onto descendants with their numbers, vertices and corners A ([Lemma 2.6](#)), the children, vertices and corner A of N_b , and of its descendants, do not depend on h either. Taking $h > b$, $N_b = g_h(Y_b^{(h)})$ is child 1 of $g_h(Y_{b+1}^{(h)}) = N_{b+1}$, and $N_0 = g_h(Y_0^{(h)}) = \{\text{Tile}(1, 1)\}$. $\square$

The nest tiling is $\mathcal{N} = \bigcup_{b \geq 0} N_b$, a set of tiles; we call the N_b its *blocks*.

Lemma 8.3 (Grandparent Lemma). Let $b \geq 1$. (a) The block N_b does not meet the rim of N_{b+2} . (b) Every edge of a tile of N_b is shared by two tiles of N_{b+2} .

Proof. (a) Work in the tiling $g_h(\mathcal{T})$ for some $h \geq b+2$, of which N_b, N_{b+1} and N_{b+2} are supertiles ([Construction 8.1](#) and [Proposition 8.2](#)). N_b is child 1 of N_{b+1} ([Proposition 8.2](#)). By [Table 2](#), sides 0, 1, 3, 4, 5 of child 1 face siblings 7, 0, 2, 3, 4 respectively, and side 2 lies on the parent's rim. We show first that every point of N_b off its closed side 2 is an interior point of N_{b+1} . This is clear for interior points of N_b , and for points in the relative interior of the other sides, where N_b and one sibling together fill a neighbourhood, as in step (ii) of the Reduction ([Section 7.1](#)). At each of the vertices 0, 1, 4, 5 of N_b exactly three level- b supertiles meet: N_b and the siblings across its two sides there, which are distinct; together they fill a neighbourhood. It remains to treat the closed side 2 of N_b . By [Table 3](#) (the Δ row: positions 14 to 18), side 4 of the parent is the concatenation of 2.2, 2.1, 1.2, 0.4, 0.3. Since 1.2 is neither first nor last, the closed side 2 of N_b lies in the relative interior of side 4 of N_{b+1} . In N_{b+2} , side 4 of its child N_{b+1} faces sibling 3, so, by the same argument one level up, every point in the relative interior of that side is an interior point of N_{b+2} . Hence every point of N_b is an interior point of N_{b+2} . [Figure 4](#) (left) shows both facts: side 2 of child 1 lies inside the parent's side 4, and side 4 of child 1 faces child 3.

(b) An edge of a tile of N_b that is not on the rim of N_b is shared by two tiles of N_b . A rim edge on a side $j \neq 2$ of N_b is shared with a tile of the sibling facing that side inside N_{b+1} . A

rim edge on side 2 lies on side 4 of N_{b+1} (position 16 of Table 3), which faces sibling 3 inside N_{b+2} (Table 2), so it is shared with a tile of that sibling. $\square$

Proposition 8.4. *$\mathcal{N}$ is a homochiral tiling of the plane, and any finitely many of its tiles lie in a common block N_b with $b \geq 1$.*

Proof. The blocks are nested (Proposition 8.2), so any finitely many tiles of $\mathcal{N}$ lie in a common block; in particular any two have disjoint interiors and are related by a proper motion, being tiles of the homochiral tiling $g_h(\mathcal{T})$ for large h . We show that the tiles cover the plane. Fix h and work in the homochiral tiling $\mathcal{T}_h = g_h(\mathcal{T})$ of Construction 8.1, which contains $N_0, \dots, N_h$ (Proposition 8.2). Call two tiles of $\mathcal{T}_h$ neighbours if they share a point, and measure distance in the resulting graph. We claim that every tile at distance at most k from N_1 lies in N_{1+2k} , provided $1 + 2k \leq h$. For $k = 0$ there is nothing to prove. A tile t at distance $k + 1$ shares a point p with a tile at distance k , which lies in N_{1+2k} by induction. By Lemma 8.3(a), p is an interior point of N_{3+2k} , so the interior of t meets the interior of a tile of N_{3+2k} (the points of t near p lie in the finite union N_{3+2k}); hence t is that tile. Now let $x \in \mathbb{R}^2$ and let D be the disc of radius $r = |x| + 1$ about the origin. Every tile has diameter $3 + \sqrt{3} < 5$ and area $3 + 3\sqrt{3} > 8$, so the tiles of $\mathcal{T}_h$ meeting D lie in the disc of radius $r + 5$, and there are at most $k_0 = \lfloor \pi(r + 5)^2/8 \rfloor$ of them. They cover the connected set D , so their neighbour graph is connected: two groups with no common point would split D into two disjoint non-empty closed sets. One of them is the tile $\text{Tile}(1, 1)$ of $N_0 \subset N_1$, since the origin is its vertex v_0 . So every tile of $\mathcal{T}_h$ meeting D is joined to $\text{Tile}(1, 1)$ by a path of fewer than k_0 steps in this graph, and lies in N_{1+2k_0} as soon as $h \geq 1 + 2k_0$; in particular, so does a tile containing x . Since N_{1+2k_0} does not depend on h , x is covered by $\mathcal{N}$. $\square$

For every $h \geq b$, the block N_b is a level- b Δ -supertile of the homochiral tiling $g_h(\mathcal{T})$ (Construction 8.1 and Proposition 8.2), so Theorem 7.6 applies to it if $b \geq 1$. By Propositions 8.2 and 8.4 and Lemma 8.3(b), the chain $N_1 \subset N_2 \subset \dots$, whose union is the tiling $\mathcal{N}$, satisfies the hypothesis of Step 2 of the proof of Theorem A; so $\mathcal{N}$ has no closed strand. Moreover, for each $m \geq 1$ the level- m descendants of the blocks N_b , $b \geq m$, partition the tiles of $\mathcal{N}$ (they agree for different b , since N_b is child 1 of N_{b+1}), and each level is grouped into the next as in $g_h(\mathcal{T})$ for large h . So this grouping has the properties of Theorem 2.1(a): its level-1 groups are clusters of Figure 3, each with its Mystic, and each group of level $m + 1$ is a supertile of some $g_h(\mathcal{T})$, made of groups of level m . By the uniqueness in Theorem 2.1(a), these are the supertiles of $\mathcal{N}$; those containing N_0 are the blocks, whose union is the plane (Proposition 8.4). So the hierarchy of $\mathcal{N}$ covers the plane: $\mathcal{N}$ has no fault, and Theorem A applies to it.

9. MERGING AND THE SINGLE STRAND

This section proves Theorem B by way of Theorem 9.1. By the junction map (Table 6), every arc of N_m lies on a hook of N_{m+1} fixed by where the arc leaves N_m , for most of its rim by the side alone. As in Section 7, the map is computed at the level-1 lengths (Computation 9.2) and carried to every level by the moves of Lemma 7.2 (Lemma 9.3); the hooks it reaches lie well inside sides of N_{m+1} (Lemma 9.4), so it applies again one level up, and after three or four levels all the arcs of N_b lie on one arc.

9.1. Merging. For $b \geq 0$ and $k \geq 0$, every arc of N_b lies on a unique arc of N_{b+k} , because N_{b+k} has no closed strand. Let $\mathcal{P}_k(b)$ be the partition of the arcs of N_b in which two arcs are in the same class if they lie on the same arc of N_{b+k} . The classes only merge as k increases; Table 5 gives their number.

Theorem 9.1 (Merging). *For every $b \geq 1$, $\mathcal{P}_4(b)$ has a single class; if b is even, so does $\mathcal{P}_3(b)$. These bounds are exact: $\mathcal{P}_2(b)$ has two classes, and so does $\mathcal{P}_3(b)$ if b is odd.*

TABLE 5. The number of arcs of N_{b+k} on which the arcs of N_b lie, that is, the number of classes of $\mathcal{P}_k(b)$ (row $b = 0$: the 14 pieces of the single tile, computed on the drawing up to level 5). Rows $1 \leq b \leq 6$: computed by following every ribbon, up to level 9 (in floating point; for $b \leq 3$ also in exact arithmetic, up to level 7); rows 0–6 are evidence, not used in the proof. Last two rows: the proof of [Theorem 9.1](#), for every $b \geq 1$ of each parity. For $b \geq 1$ the column $k = 1$ also follows from [Lemma 9.3](#): it is $\epsilon_b + 3$ for even b and $\epsilon_b + 2$ for odd b (the $\epsilon_b - 1$ hooks strictly inside side 2 of N_b , which has letter ϵ , and the four or three hooks H_q that the column of [Table 6](#) for the parity of b names, each reached, since no row of that table is empty, and none at a vertex strictly inside side 2). A dash: not computed (it needs level 10).

b	tiles in N_b	arcs of N_b	arcs of N_{b+k}			
			$k = 1$	$k = 2$	$k = 3$	$k = 4$
0	1	14	5	2	1	1
1	9	46	10	2	2	1
2	71	182	37	2	1	1
3	559	758	146	2	2	1
4	4,401	3,198	613	2	1	1
5	34,649	13,534	2,586	2	2	1
6	272,791	57,318	10,949	2	1	–
even $b \geq 2$				2	1	1
odd $b \geq 1$				2	2	1

9.2. Junctions. Let $m \geq 1$. The rim of N_{m+1} , read in its own order from its corner A , is the concatenation of the 22 sides of children at positions $q = 0, \dots, 21$ in the eight-children part of [Table 3](#), and the Δ row gives the first position of each side of N_{m+1} . The *junction* q is the rim vertex of N_{m+1} at which the side at position q starts ([Figure 4](#), left). Junctions are where the rim passes from one of these sides to the next. Every other rim vertex of N_{m+1} lies strictly inside one child’s side, and by [Theorem 7.6](#) an arc of N_{m+1} with an end there is an arc of that child alone (its hook there or, at B , a crossover arc of the Γ child, the only child with crossover corners on the rim, as noted after [Computation 7.3](#)); so an arc of N_{m+1} that runs through more than one child, such as one continuing an arc of N_m into its siblings, has both its ends at junctions. The junctions 9, 12, 16, 17, 21 lie inside sides 2, 3, 4, 4, 5 of N_{m+1} , at least two edges from their ends ([Lemma 9.4](#) below, which uses only [Table 3](#) and [Lemma 6.5](#)), so they are hooks of N_{m+1} ([Theorem 7.6](#)); we write H_q for the hook arc of N_{m+1} at junction q . The side at position 16 is 1.2, side 2 of N_m , so junctions 16 and 17 are its two ends, vertices 3 and 2 of N_m respectively. The rim vertices of N_m are indexed $0, \dots, E - 1$, where $E = E(N_m)$ is their number, in its own order from its corner A , so that its vertices $0, \dots, 5$ have indices $0 = V_0 < V_1 < \dots < V_5 < E$; its corner $B = A + \omega_m$ has index 1 if m is even and $E - 1$ if m is odd.

Computation 9.2 (The junction map at the level-1 lengths). *Let $\omega = \pm 1$. In $C_\Delta(\ell_1, d_1, \omega)$, put child 1 in the place of N_m , and define indices and junctions as above. Then the ribbon from every port p of child 1, except the ports at rim vertices strictly inside its side 2, followed from p through child 1 and on, reaches the rim of the composition at junction q , where q is given by [Table 6](#) with $\omega_m = \omega$.*

Like [Computation 7.3](#), this was computed by `hierarchy/compositions.py` and in the kernel of Lean ([Appendix A](#)).

From the signatures of the children and [Lemma 5.1](#), two rows of [Table 6](#) also follow by hand (this is not used in the proof). Let $V_3 < u < V_4$. Side 3 of N_m is glued to side 0 of child 2, whose crossover corners are its vertex 0, which is V_4 , and a rim vertex next to it; so no crossover corner of N_m or of child 2 lies strictly between V_3 and u . By [Lemma 5.1](#), the

TABLE 6. The junction map (Lemma 9.3): the junction q such that an arc of N_m lies on the hook H_q of N_{m+1} . Hooks are listed by the index u of their rim vertex, which is never A or B .

arc of N_m	m even ($\omega_m = +1$)	m odd ($\omega_m = -1$)
crossover arc through $R(A)$ and $R(B)$	17	21
crossover arc through $L(A)$ and $L(B)$	9	16
hook at u , $0 < u \leq V_1$	17	21
hook at u , $V_1 < u \leq V_2$	17	17
hook at u , $V_2 < u < V_3$	a hook of N_{m+1} itself	
hook at u , $V_3 \leq u < V_4$	16	16
hook at u , $V_4 \leq u < V_5$	12	16
hook at u , $V_5 \leq u < E$	9	16

ribbon leaving the hook at u through $L(u)$ reaches V_3 at $R(V_3)$ of N_m or at $L(V_3)$ of child 2, and the hook there takes it onto the rim of N_{m+1} at V_3 , junction 16 (V_3 is a crossover corner of neither child, as the crossover corners of each of these children are its vertex 0 and a rim vertex next to it, and sides have at least five edges). For $V_1 < u < V_2$ it is the same with child 0, the port $R(u)$ and junction 17: side 1 of N_m is side 5 of child 0, whose vertex 0 is V_1 , so no crossover corner lies strictly between u and V_2 , and V_2 is a crossover corner of neither child. At $u = V_3$ and $u = V_2$ the hook itself ends on the rim. As junctions 16 and 17 are hooks of N_{m+1} , these arcs of N_{m+1} are H_{16} and H_{17} .

Lemma 9.3 (The junction map). *Let $m \geq 1$. In N_{m+1} , the hook arc of N_m at each rim vertex strictly inside side 2 of N_m is a hook arc of N_{m+1} , and every other arc of N_m lies on the hook H_q of N_{m+1} , where q is given by Table 6.*

Proof. The first part needs no computation. Side 2 of N_m is the side at position 16 of Table 3, so it lies on the rim of N_{m+1} , and a rim vertex u of N_m strictly inside it is a rim vertex of N_{m+1} , with the same two rim ports. By Lemma 9.4, u lies inside side 4 of N_{m+1} , so it is not a crossover corner, and by Theorem 7.6 it is a hook of N_{m+1} ; this hook is the hook of N_m at u .

For the rest, say that (L, d) has the *junction property* if the statement of Computation 9.2 holds for $C_\Delta(L, d, \omega)$ for both ω . By Theorem 7.6 and the Reduction of Section 7.1, N_{m+1} is glued from its children as $C_\Delta(\ell_m, d_m, \omega_m)$ is, with N_m as child 1. In this gluing, the ribbon from a port p of child 1 runs along the arc of N_m that ends at p and then on along the arc of N_{m+1} that contains that arc, to one of its ends. The two ends of an arc of N_m lie in the same row of Table 6 (a hook arc has both its ends at one rim vertex, and each crossover arc has a row), so if the ribbons from both reach the rim at junction q , both ends of that arc of N_{m+1} are at q , and it is the hook H_q . So the rest of the lemma follows once (ℓ_m, d_m) has the junction property. By Computation 9.2, (ℓ_1, d_1) has it, so by Lemma 7.2 it suffices to show that the moves of Lemma 7.2(b) keep it. Since child 1 is a Δ , and the side glued to side 5 of the Σ child (child 4) belongs to child 7 (Table 2), such a move, when its edge lies on the rim of child 1, has that edge at least two edges from both ends of its side (Lemma 7.2(b)).

Let C, C' and φ be as in the proof of Lemma 7.1. Junctions and the vertices V_j of child 1 are vertices of children, so those of C' are the images of those of C , and φ keeps the order of the rim vertices of child 1. For a port x of child 1 at an old vertex, the ribbon from $\varphi(x)$ reaches the rim at the image of the rim port reached from x , and $\varphi(x)$ lies in the same row of Table 6 as x . It remains to consider the new vertices w_1, w_2 when the edge $e = u_0u_1$ lies on the rim of child 1. Then w_1, w_2, u_0, u_1 lie strictly inside the same side of child 1, at least two edges from its ends, so in the same row of Table 6, and none of them is a crossover corner. If that side is side 2, the property says nothing about w_1 and w_2 . Otherwise e is shared with a sibling; let child 1 run from u_0 to u_1 , with rim ports R_1, L_1 . By Lemma 5.1 with $n = 3$, as in the proof of Lemma 7.1, every port of child 1 at w_1 or w_2 lies on the arc of C' through $\varphi(R_1(u_0))$ or on the arc through $\varphi(L_1(u_1))$. The first arc also passes through

$\varphi(L_1(u_0))$, by the hook of child 1 at u_0 , and the ribbons from these two ports run along it in opposite directions; so both its ends are at the junction that Table 6 gives for the row of u_0 . Likewise both ends of the second arc are at the junction for the row of u_1 , which is the row of w_1 and w_2 , so the property holds at w_1 and w_2 too. $\square$

Lemma 9.4 (Where the junctions lie). *Let $m \geq 1$. The junctions 9, 12, 16, 17 and 21 of N_{m+1} lie on its sides 2, 3, 4, 4 and 5 respectively, at least two edges from both ends of the side. Every rim vertex strictly inside side 2 of N_m lies on side 4 of N_{m+1} , at least two edges from both ends.*

Proof. By Table 3, side 2 of a Δ is made of the children's sides at positions 7 to 10, side 3 of those at 11 to 13, side 4 of those at 14 to 18, and side 5 of those at 19 to 21. Junction q lies between the sides at positions $q - 1$ and q , which lie in the same side of N_{m+1} for $q = 9, 12, 16, 17, 21$; and side 2 of N_m is the side at position 16. Every side of a supertile of level ≥ 1 has at least five edges (Lemma 6.5), so each of these junctions is at least five edges from both ends of its side of N_{m+1} , and so is every point of side 2 of N_m , which lies between the sides at positions 15 and 17. $\square$

Proof of Theorem 9.1. In this proof a rim vertex, and the hook at it, lie *inside* a side if the vertex is at least two edges from both ends of that side. A rim vertex of N_m inside its side j is not a crossover corner, and by Lemma 9.3 its hook lies on the hook of N_{m+1} that Table 6 names for side j : for $j = 0, 1, 3, 4, 5$, the hook H_q with $q = 17, 17, 16, 12, 9$ if m is even and $q = 21, 17, 16, 16, 16$ if m is odd; for $j = 2$, its hook is itself a hook of N_{m+1} . By Lemma 9.4, the junctions 9, 12, 16, 17, 21 lie inside sides 2, 3, 4, 4, 5 of N_{m+1} , and the rim vertices strictly inside side 2 of N_m inside side 4. So each of these hooks of N_{m+1} lies inside a side of N_{m+1} , and the same rule applies to it one level up; for $j = 0, \dots, 5$ that side is 4, 4, 4, 4, 3, 2 if m is even and 5, 4, 4, 4, 4, 4 if m is odd.

Let $b \geq 1$ be even, so that $\omega_b = +1$. By Lemma 9.3 every arc of N_b lies on H_{17}, H_{16}, H_{12} or H_9 of N_{b+1} , or on a hook of N_{b+1} strictly inside side 2 of N_b ; by Lemma 9.4 these are hooks inside sides 4, 4, 3, 2 and 4 of N_{b+1} . Since $b + 1$ is odd, the hooks inside sides 3 and 4 lie on H_{16} of N_{b+2} , and those inside side 2 are hooks of N_{b+2} inside its side 4. All are inside side 4 of N_{b+2} , and since $b + 2$ is even they lie on H_{12} of N_{b+3} . So $\mathcal{P}_3(b)$ has a single class.

Let $b \geq 1$ be odd. By Lemmas 9.3 and 9.4, every arc of N_b lies on H_{21}, H_{17} or H_{16} of N_{b+1} , or on a hook of N_{b+1} strictly inside side 2 of N_b : hooks inside sides 5, 4, 4 and 4 of N_{b+1} . Since $b + 1$ is even, those inside side 4 lie on H_{12} of N_{b+2} , inside its side 3, and the one inside side 5 lies on H_9 of N_{b+2} , inside its side 2. Since $b + 2$ is odd, this H_{12} lies on H_{16} of N_{b+3} , and this H_9 is a hook of N_{b+3} inside its side 4. Both are inside side 4 of N_{b+3} , and since $b + 3$ is even they lie on H_{12} of N_{b+4} . So $\mathcal{P}_4(b)$ has a single class. In both cases, once $\mathcal{P}_k(b)$ has a single class, so does $\mathcal{P}_{k'}(b)$ for every $k' \geq k$.

The same steps give the exactness. For even b , the arcs of N_b lie on two arcs of N_{b+2} : H_{16} , and the hook of N_{b+2} at junction 9 of N_{b+1} . That vertex lies strictly inside side 2 of N_{b+1} , and junction 16 of N_{b+2} is an end of that side, so the two hooks are at different rim vertices. For odd b , they lie on two arcs of N_{b+2} , H_{12} and H_9 , and on two arcs of N_{b+3} : H_{16} , and the hook of N_{b+3} at junction 9 of N_{b+2} , a vertex that for the same reason is not junction 16 of N_{b+3} . In each case both arcs are reached (for instance from vertices inside sides 4 and 5 of N_b if b is even, and sides 4 and 0 if b is odd), and hooks at distinct rim vertices are distinct arcs. $\square$

9.3. One strand.

Proof of Theorem B. By Lemma 3.4, every component of the strand graph $\mathcal{G}(\mathcal{N})$ is a cycle or a two-way infinite path, and there are no cycles, as shown in the paragraph after Proposition 8.4. It remains to show that $\mathcal{G}(\mathcal{N})$ is connected. Let p, p' be pieces of tiles t, t' of $\mathcal{N}$. By Proposition 8.4 both tiles lie in some N_b with $b \geq 1$. Since N_b has no closed strand, p and p' lie on arcs of N_b . By Theorem 9.1 these arcs lie on one arc of N_{b+4} , so p and p' are on the same strand. The strand passes through every tile, since every tile has pieces. Finally, every

piece of N_b lies on an arc of N_b , and by [Theorem 9.1](#) these arcs lie on one arc of N_{b+3} if b is even, and of N_{b+4} if b is odd. $\square$

10. SCOPE AND OPEN QUESTIONS

The theorems use only the pairing of ports ([Definition 3.1](#)) and the straight-through rule ([Definition 3.2](#)), so they hold for any drawing that realises these, including drawings on the curved Spectre, whose tilings correspond to the homochiral tilings by $\text{Tile}(1, 1)$ ([Section 2](#)). For a decoration with its crossover at another pair of adjacent vertices, the same finite computations can be run; their outcome may differ.

In a tiling with a fault, the union of the supertiles containing a tile, an infinite supertile [[31](#), [§5](#)], is only part of the plane, and two or more infinite supertiles meet along the fault; Tatham [[31](#), [§7.3](#)] gives an example. A closed strand inside one infinite supertile is ruled out by [Theorem 7.6](#): its finitely many tiles lie in one supertile, and every edge it crosses is shared by two of them, so it is a closed strand of that supertile. The argument does not reach a closed strand that crosses the fault, and whether one can occur is open. [Theorem B](#) is not affected, since the nest tiling has no fault ([Section 8](#)).

The nest tiling follows the chain $\Delta \subset \Delta \subset \dots$, with child number 1 at every step. It is not the only chain of constant kind and child number, since Γ , Σ , Φ and Ψ each have a child of their own kind ([Table 1](#)); the others are not treated here. In any tiling whose hierarchy covers the plane, the supertiles $X_1 \subset X_2 \subset \dots$ containing a tile give a chain $(K_m, c_m)_{m \geq 1}$: X_m has kind K_m and is child c_m of X_{m+1} . The induction of [Section 7](#) does not depend on the chain, but the merging of [Section 9](#) does, and computations, not used in the proofs, indicate that the number of strands does too. Take $(K_m, c_m) = (\Theta, 2)$ for odd m and $(\Gamma, 7)$ for even m . This chain is built as in [Construction 8.1](#) and covers the plane by the argument of [Lemma 8.3](#) and [Proposition 8.4](#): by [Tables 2](#) and [3](#), a Θ that is child 2 of a Γ meets the rim of its parent only at the parent's vertex 3 and inside the parent's sides 2 and 3, and a Γ that is child 7 of a Θ only at the parent's vertex 0 and inside the parent's sides 5 and 0; one level up, these sides face siblings. Along this chain, the arcs of X_1 lie on exactly three arcs of X_m for $3 \leq m \leq 8$ (exact arithmetic), and still on three at $m = 20$ (a symbolic count); we have not proved that this persists. In all but five of 90,000 chains chosen at random (30,000 for each $b = 1, 2, 3$; X_b of a random kind, and each parent drawn uniformly from the pairs (T, c) of [Table 1](#) whose child c has the current kind), the arcs of X_b lie on a single arc of X_m for some $m \leq 20$ (folder `paper/code/experiments/` of the repository in [Appendix A](#)). It is open which tilings whose hierarchy covers the plane carry a single strand, and which numbers of strands occur.

APPENDIX A. THE COMPUTATIONS

All programs (under the MIT licence) and their logs are in the folder `paper/code/` of the public repository https://github.com/Cortado-J/SPECTRE_KNOT_PROOF: there `./run_all.sh` re-runs everything the proof needs, and `README.md` maps each computation in this paper to the script that does it. The Lean formalisation is in the folder `lean/` ([Appendix B](#)). The version that goes with this paper is tagged in the repository with its date of publication. The proof site, https://cortado-j.github.io/SPECTRE_KNOT_PROOF/, gives each result with its proof, what it cites and the Lean code that formalises it.

What the proof computes. The geometric computations on actual supertiles (the level-1 data and [Computation 7.5](#)) are done in exact arithmetic by `exact/exact.py`, with every vertex in $\mathbb{Z}[\rho]$ and every isometry $z \mapsto \rho^k z + t$ or $z \mapsto \rho^k \bar{z} + t$; floating point is used only to read off the sign of a polygon's area, which gives the orientation of its rim (every such area exceeds 8 in absolute value, so the sign is not in doubt). The computations on compositions ([Computations 7.3](#) and [9.2](#)) use only [Tables 1](#) to [4](#) and the level-1 lengths ℓ_1 and d_1 : `hierarchy/compositions.py` builds each composition, follows the ribbon from every port and checks each clause of the statement. Lean's kernel repeats [Computations 7.3](#) and [9.2](#), and [Computation 7.5](#) on the nine level-1 supertiles of `exact.py`, taken as input. Two

negative controls, run as tests both in Lean and by the programs, show that these checks can fail: with each child's corner B_c placed two vertices from A_c instead of one, all 18 checks of [Computation 7.3](#) fail; with junctions 12 and 16 exchanged in [Table 6](#), [Computation 9.2](#) fails for both signs of ω .

For [Proposition 2.4](#), one script fits our level-1 supertiles to a rendered image of [\[29, Fig. 4.1\]](#). For each cluster, a proper similarity takes our six vertices to its six dots with residual below 1% of an edge, while any other matching of the dots, or the vertices of any other kind, fits at least 40 times worse; our rim lies on the printed outline, and the first tile of our Mystic on the light tile of the printed Mystic; and the labels, transcribed from the figure (not read by the script), agree with [Table 4](#). Another script reads an encoding of the nine panels of [\[29, Fig. 5.1\]](#), finds the unique consistent rotation of each inner hexagon, and checks the encoding against [Tables 1 to 3](#) and the labels on each panel's outline against [Table 4](#). On the same pattern, `hierarchy/letters.py` repeats the inspection of the 54 parent sides in the proof of [Proposition 6.1](#).

Evidence, not proof. Besides the cross-checks of the pattern on actual supertiles ([Remark 2.5](#)), separate programs, mostly in JavaScript, sharing no code with `exact.py` and working in floating point (we do not bound their rounding error), agree with the proof: the actual supertiles up to level 9 have the signature, and following every ribbon through the actual blocks up to level 9 gives the numbers of [Table 5](#). In exact arithmetic, on blocks built by `exact.py`, following every ribbon up to level 7 confirms the rows $1 \leq b \leq 3$ of [Table 5](#), and [Table 6](#) holds on $N_1, \dots, N_4$. Second routes, through skeletons of the compositions (only the parities of the shared runs) and a symbolic computation with certified linear inequalities, give [Theorems 7.6](#) and [9.1](#) again. The drawn ribbons realise the pairing of [Definition 3.1](#) on all supertiles up to level 3; each piece stays inside its tile, the two crossover pieces cross once, transversally, and the pieces meet each edge at $\pm 45^\circ$, so the curves cross only transversally and two at a time.

APPENDIX B. USE OF AI, FORMAL VERIFICATION AND TOOLS

How the work was done. The author came up with the original idea of looking for a mono-stranded aperiodic knot, and then used an AI system (Claude, by Anthropic) to search for it; from the 34 candidates found, the author chose design 0.1:01 (the decoration of [Definition 3.1](#)), drew its ribbons by hand and chose where its crossing points sit on the edges. From then on, at the author's instruction, the AI system generated the programs, the proofs, the Lean formalisation, the text of this paper and its figures; the author steered the work, made the choices and decisions, corrected what was generated, and answers for the work. The author has read parts of this paper and spot-checked figures and counts; no professional mathematician has reviewed it. A program, check or route called *separate*, *second* or *independent* here is another implementation, generated by the same AI system. At the author's instruction, the AI system also reviewed the argument in several passes and searched the literature for related work. AI systems can reproduce published ideas without naming their source, so this paper may use an idea whose source it does not cite. A fuller account: https://cortado-j.github.io/SPECTRE_KNOT/made.html.

Formal verification. The combinatorial part of the proof is formalised in the Lean 4 proof assistant [\[8\]](#), using its core library only (folder `lean/`). The formalisation describes a patch, or a whole tiling, combinatorially: by which edge of which tile is glued to which. In this form Lean's kernel has checked [Theorem 7.6](#) for every level $m \geq 1$; [Theorem A](#) for the union of any chain of supertiles, each a child of the next, in which every edge is eventually glued (Step 2 of its proof), and that the nest is such a chain; [Theorem B](#) for the nest tiling built as the union of the blocks N_b , with [Theorem 9.1](#), including its exactness; and [Lemma 8.3](#), with part (a) only at the vertices of tiles. The formal statements rest on [Tables 1 to 3](#) and on the nine level-1 supertiles as glued patches of tiles, which `exact/exact.py` computes in exact arithmetic and the formalisation takes as data. The formalisation has no Mystic, so

the position of the crossover corner of a level-1 supertile (Lemma 6.4 for $m = 1$) comes from that program, not from Lean.

The link with tilings of the plane is not formalised but assumed, in three parts, each argued in this paper: (M1) the supertiles of an actual tiling are glued as the tables say (Theorem 2.1, Proposition 2.4, Lemma 2.6 and step (ii) of Section 7.1); (M2) a tiling whose hierarchy covers the plane is the union of a chain as above, in which every edge is eventually glued (Step 1 of the proof of Theorem A); (M3) the nest tiling is the union of the blocks N_b , each child 1 of the next, and covers the plane (Propositions 8.2 and 8.4), with the step of Lemma 8.3(a) from the vertices of tiles to the points of the plane. The formal proofs follow Sections 3 to 9 step by step, apart from what (M1)–(M3) assume (lean/README.md maps each result to its formal counterpart). They differ only in bookkeeping: strands are orbits of ports (Section 3.3); the signature (Definition 4.3) is stated by following ribbons rather than through π_P ; and, inside the proofs of Lemma 7.1 and Proposition 7.4, the signature of a composition is carried in a form that cutting an edge leaves unchanged: each rim vertex is found by turning about a vertex of a child, and the positions of Table 3 and the adjacency of A and B are recovered at the end. The formalisation also proves a step the paper only sketches, that a move gives the old composition with an edge cut into three (proof of Lemma 7.1), and facts the paper takes from the geometry: every supertile is a disc near its rim (turning about a rim vertex leads from the tile on one rim edge to the tile on the other), and its gluing is symmetric and glues no tile to itself.

The proofs use only Lean’s standard axioms. The main formal theorems include signatureTheorem (Theorem 7.6), theoremA_strand (Theorem A), theoremB_strand (Theorem B), theoremB_sharp and theoremB_exact (Theorem 9.1), and grandparent (Lemma 8.3(a)), all in the namespace SpectreKnot.Statement. Like the proofs, the programs and this text, the formalisation was generated by the AI system at the author’s instruction. Its statement file (lean/SpectreKnot/Statement.lean), with the three files it imports (a few basic definitions; as data, Tables 1 to 3 and the nine level-1 supertiles, and Table 4, ℓ_1 and d_1 , which only the proofs use), holds every definition on which the formal theorems depend, with comments in plain English. No human has checked that these formal statements say what the theorems of this paper say.

Tool and computational resource disclosure. *AI system:* Claude, by Anthropic (various Claude models, primarily Opus), used in the Claude app and through Claude Code from August to October 2026; no other AI or machine-learning system was used. *Proof assistant:* Lean 4, version 4.34.1, core library only; leancore [21] replays every declaration through the kernel, and comparator [20] checks that the main theorems, as proved, have exactly the statements in Statement.lean. *Other software:* Python 3 with SymPy (for Section 6), NumPy and SciPy (to propose the certificates of the symbolic computation of Appendix A, which are then checked in exact arithmetic, and, with Matplotlib and Pillow, for the comparisons with [29]); Node.js; Poppler; and L^AT_EX. Except for L^AT_EX, their versions are recorded in paper/code/logs/A00_environment.log. One cross-check (Remark 2.5) downloads Kaplan’s spectre.js, which is not included in the repository. *Computing:* the computations the proof needs take about twelve minutes on a four-core machine, the corroborating ones of Appendix A up to level 9 about an hour (up to 6.5 GB of memory), and building the formalisation about ten minutes; none of the checks the proof needs uses the AI system or proprietary software. The AI system ran on Anthropic’s servers; the computing resources it used are not known to the author.

REFERENCES

1. Michael Baake and Uwe Grimm, *Aperiodic order. Volume 1: A mathematical invitation*, Encyclopedia of Mathematics and its Applications, vol. 149, Cambridge University Press, Cambridge, 2013.
2. George Bain, *Celtic art: The methods of construction*, William MacLellan, Glasgow, 1951.
3. bohemian-miser, *Spectre: Playing with the spectre tiling*, <https://github.com/bohemian-miser/Spectre>, file docs/FASS_THEOREM.md at commit f1a946c (accessed 8 October 2026), 2026.

4. Joseph I. Cline, *Interlace patterns emerging in a Penrose-type Islamic design*, Proceedings of Bridges 2024: Mathematics, Art, Music, Architecture, Culture, 2024, pp. 163–170.
5. ———, *Knotted strapwork strands in a Penrose-type girih pattern*, Proceedings of Bridges 2025: Mathematics, Art, Music, Architecture, Culture, 2025, pp. 61–68.
6. Peter R. Cromwell, *Celtic knotwork: Mathematical art*, The Mathematical Intelligencer 15 (1993), no. 1, 36–47.
7. ———, *Knots and links*, Cambridge University Press, Cambridge, 2004.
8. Leonardo de Moura and Sebastian Ullrich, *The Lean 4 theorem prover and programming language*, Automated Deduction – CADE 28, Lecture Notes in Computer Science, vol. 12699, Springer, 2021, https://doi.org/10.1007/978-3-030-79876-5_37, pp. 625–635.
9. Gwen Fisher and Blake Mellor, *On the topology of Celtic knot designs*, Bridges: Mathematical Connections in Art, Music, and Science (Winfield, Kansas), 2004, pp. 37–44.
10. Felix Flicker, Steven H. Simon, and S. A. Parameswaran, *Classical dimers on Penrose tilings*, Physical Review X 10 (2020), 011005, <https://doi.org/10.1103/PhysRevX.10.011005>.
11. Uli Gaenshirt, *Ammann grid and knot structure of a quasiperiodic girih pattern*, Proceedings of Bridges 2024: Mathematics, Art, Music, Architecture, Culture, 2024, pp. 155–162.
12. Martin Gardner, *Penrose tiles to trapdoor ciphers*, W. H. Freeman, New York, 1989.
13. Paulus Gerdes, *Lunda geometry: Mirror curves, designs, knots, polyominoes, patterns, symmetries*, second ed., Lulu, Morrisville, NC, 2007.
14. Chaim Goodman-Strauss, *Aperiodic hierarchical tilings*, Foams and Emulsions (J. F. Sadoc and N. Rivier, eds.), NATO Science Series E, vol. 354, Kluwer Academic Publishers, Dordrecht, 1999, https://doi.org/10.1007/978-94-015-9157-7_28, pp. 481–496.
15. Jonathan L. Gross and Thomas W. Tucker, *A Celtic framework for knots and links*, Discrete & Computational Geometry 46 (2011), 86–99, <https://doi.org/10.1007/s00454-010-9257-0>.
16. Branko Grünbaum and G. C. Shephard, *Tilings and patterns*, second ed., Dover, 2016.
17. W. Patrick Hooper, *Truchet tilings and renormalization*, arXiv:1103.6046, 2011.
18. ———, *Renormalization of polygon exchange maps arising from corner percolation*, Inventiones Mathematicae 191 (2013), no. 2, 255–320, <https://doi.org/10.1007/s00222-012-0393-4>; arXiv:1105.6137.
19. Slavik Jablan, Ljiljana Radović, Radmila Sazdanović, and Ana Zeković, *Knots in art*, Symmetry 4 (2012), no. 2, 302–328, <https://doi.org/10.3390/sym4020302>.
20. Lean FRO, *Comparator: A trustworthy judge for Lean proofs*, <https://github.com/leanprover/comparator>, tag v4.34.0.
21. ———, *leanchecker*, part of the Lean 4 toolchain from version 4.28.0, <https://github.com/leanprover/lean4>; formerly <https://github.com/leanprover/lean4checker>.
22. Lucy Liu, Gary P. T. Choi, and L. Mahadevan, *Quasicrystal kirigami*, Physical Review Research 4 (2022), no. 3, 033114, <https://doi.org/10.1103/PhysRevResearch.4.033114>; arXiv:2104.13399.
23. Samuel J. Lomonaco and Louis H. Kauffman, *Quantum knots and mosaics*, Quantum Information Processing 7 (2008), 85–115, <https://doi.org/10.1007/s11128-008-0076-7>; arXiv:0805.0339.
24. Stefan Pautze, *Space-filling, self-similar curves of regular pentagons, heptagons and other n-gons*, Proceedings of Bridges 2021: Mathematics, Art, Music, Architecture, Culture, 2021, pp. 157–164.
25. Gábor Pete, *Corner percolation on $\mathbb{Z}^2$ and the square root of 17*, Annals of Probability 36 (2008), no. 5, 1711–1747, <https://doi.org/10.1214/07-AOP373>.
26. Shobhna Singh, Jerome Lloyd, and Felix Flicker, *Hamiltonian cycles on Ammann–Beenker tilings*, Physical Review X 14 (2024), 031005, <https://doi.org/10.1103/PhysRevX.14.031005>.
27. David Smith, Joseph Samuel Myers, Craig S. Kaplan, and Chaim Goodman-Strauss, *A chiral aperiodic monotile*, Project page with the interactive implementation spectre.js, <https://cs.uwaterloo.ca/~csk/spectre/> (accessed 15 October 2026), 2023.
28. ———, *An aperiodic monotile*, Combinatorial Theory 4 (2024), no. 1, Paper No. 6, <https://doi.org/10.5070/C64163843>; arXiv:2303.10798.
29. ———, *A chiral aperiodic monotile*, Combinatorial Theory 4 (2024), no. 2, Paper No. 13, <https://doi.org/10.5070/C64264241>; arXiv:2305.17743. Theorem, figure, section and page numbers refer to the published version.
30. Joshua E. S. Socolar and Joan M. Taylor, *An aperiodic hexagonal tile*, Journal of Combinatorial Theory, Series A 118 (2011), no. 8, 2207–2231, <https://doi.org/10.1016/j.jcta.2011.05.001>; arXiv:1003.4279.
31. Simon Tatham, *Finite-state transducers for substitution tilings*, arXiv:2512.16595v2, 2025.

INDEPENDENT RESEARCHER

Email address: cortado.just@gmail.com