

Fibbinary Zippers in a Monoid of Toroidal Hamiltonian Cycles that Generate Hilbert-Style Square-Filling Curves

Douglas M. McKenna*

*Mathemaesthetics, Inc.
 Boulder, Colorado USA
 Email: doug@mathemaesthetics.com

Received: August 13, 2021, **Accepted:** November 16, 2021, **Published:** November 26, 2021
 Released under the CC BY-NC-SA license (International 4.0)

ABSTRACT: Within the recursive subdivision of the $n \times n$ square, what characterizes a Hilbert-style space-filling curve motif of length n^2 when—under iterated, self-similar, pure edge-replacement—a sequence of always self-avoiding lattice paths results? How many motifs are there and what do they look like? Such motifs are composable elements of a monoid, where all such motifs map to a particular subset of Hamiltonian cycles on the $n \times n$ toroidal grid-graph. We prove that for any odd $n \geq 1$ each motif has a shape that falls into exactly one of $F_{(n-3)/2}$ boundary “zipping” modes, where F_i is the i th Fibonacci number; for even n no solution motifs exist. Each mode is governed by a special palindromic Fibbinary bit sequence (i.e., having no adjacent 1 bits). To varying degrees, each zipping mode emanates further combinatorial constraint inward from the square’s boundary, especially at the corners. The zipping mode whose Fibbinary bits have the most consecutive 0s freezes over half of the n^2 edges of an order- n motif into only one distinct (either left- or right-handed) configuration. Manual and machine enumeration for small n is significantly enhanced by these results. For $n = 1, 3, 5, 7, 9, 11$ there are 1, 0, 1, 7, 10101, 20305328 distinct, globally self-avoiding motifs, falling into $F_{(n-3)/2} = 1, 0, 1, 1, 2, 3$ zipping modes, respectively. For $n \geq 5$, each such motif, when infinitely exponentiated within its monoid, converges to an open-ended, square-filling, continuous curve.

Keywords: Fibonacci words; Fibbinary; Fractals; Monoids; Space-filling curves; Self-avoiding lattice paths; Toroidal Hamiltonian cycles

2020 Mathematics Subject Classification: 05C38; 05C30; 06D50; 11B39; 52C05

Introduction

Space-filling curves continuously and surjectively map every point in a line segment onto every point in a higher dimensional region. Two-dimensional, “Hilbert-style” curves fill a square by connecting one corner to an *adjacent* corner in an open-ended, infinitely long, continuous but non-differentiable curve. We will concern ourselves here not with the limiting space-filling curves *per se*, but rather with the combinatorial constraints governing the everywhere-repeated, edge-connected paths—called “motifs”—on the square lattice that generate increasingly accurate—yet always self-avoiding—paths which, in the limit, converge to the self-similar, square-filling curve.

These constructions rely on a self-similar tiling property of the square: it can be subdivided into an array of sub-squares, where each of the smaller sub-squares is geometrically similar to the whole. When these sub-squares are suitably ordered and rotated so as to guarantee “threaded” connectivity, they form a sequence in which all pairs of consecutive sub-squares are also adjacent—either by edge or only by corner—within the subdivision.

The Hilbert Curve [4], the Wunderlich Curve [22], and their usual order- n generalizations [13,15] are all based on pure edge-adjacency: every sequential pair of threaded sub-squares shares an edge.¹ Their approximations are typically drawn as piecewise-linear, Hamiltonian paths on the dual edge-adjacency graph (Fig. 1, left), whose vertices mark the centers of edge-adjacent sub-squares. Because it is Hamiltonian, a *dual path* is always self-avoiding, and therefore unambiguously specifies the underlying sequence of all n^2 sub-squares.

But when we permit corner-only adjacency, the underlying graph for an arbitrary Hamiltonian dual path is no longer a square grid-graph, as illustrated in Fig. 1, right. Diagonal line segments are now possible when drawing a picture of the dual path.

¹Pure edge adjacency is also true of the original Peano Curve [16], but it relies on a non-Hilbert-style recursive traversal from corner to diagonally opposite corner.

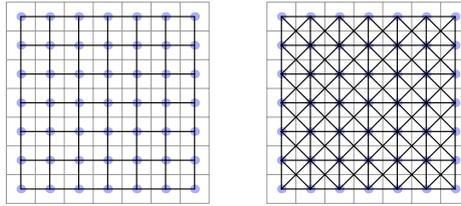


Figure 1: Left: Adjacency graph (black) of subdivided square (gray) for dual paths of pure edge-adjacent motifs. Right: Adjacency graph for dual paths where corner-only adjacency of consecutive sub-squares is permitted.

In contradistinction to the dual path, there is a second way, called the *tile path*, to specify and illustrate a motif's consecutive sub-squares. For Hilbert-style motifs, a tile path comprises a sequence of pairwise-connected, sub-square edges, each edge bijectively bound to exactly one sub-square. Fig. 2 illustrates the tile and dual paths for the well-known Hilbert Curve construction, for its first six iterations. Notice, though, that the edges of the successive, piecewise-connected, approximation tile paths in the top row are regularly self-contacting.

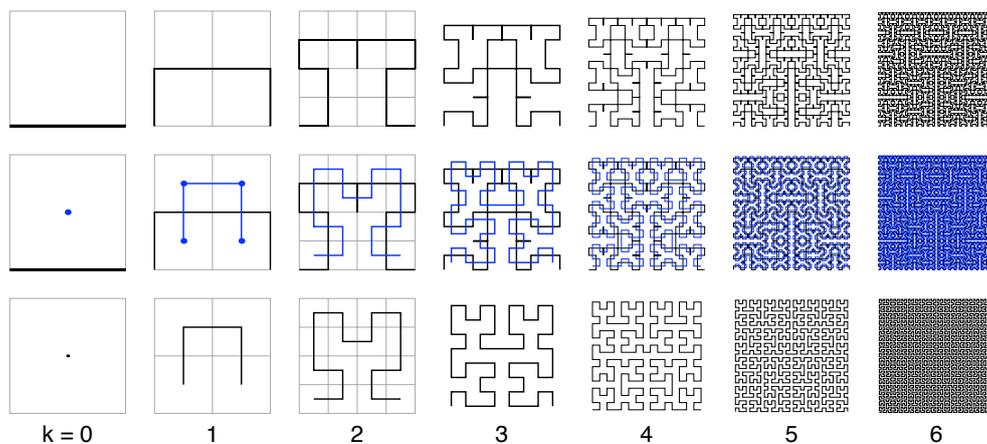


Figure 2: The Hilbert Curve construction, for stages $0 \leq k \leq 6$. Top: Approximating tile paths (self-contacting). Middle: Approximating dual paths superimposed on the tile path. Bottom: Just the (self-avoiding) dual path.

In this paper, within only the space of 2-D Hilbert-style constructions,² we focus on the shapes of motif tile paths³ for any $n \geq 1$ when we (a) prohibit any mirroring of replicated motifs, (b) allow consecutive sub-squares, at least once and as needed, to be adjacent only at a corner, and (c) require piecewise-connected tile paths on the square lattice at all finite recursive levels of space-filling curve approximation to be self-avoiding.

In Section 1, we define terms, including an edge-replacement composition operator that preserves salient inside/outside parity information, and we show that these space-filling curve motifs are elements of a monoid \mathcal{H} .

In Section 2, we show the results of hand and brute-force machine enumerations of all distinct motifs for small n , including the interesting experimental discovery that all order-9 solutions fall into two boundary modes.

Motivated by this discovery, in Section 3, we prove the Fibbinary Zipper Theorem, which tells us how many boundary modes there will be for a given n and what they look like.

In Section 4, we show how the Fibbinary zipping mode characterized by the most consecutive 0 bits emanates significant constraint into the interior of the order- n square subdivision, freezing over half the sub-tile edges.

In Section 5, we show the results of manual enumerations of order-5 (1 zipping mode) and order-7 (1 mode) distinct motifs, and machine enumeration for order-9 (all in the first mode, samples from second), samples of all order-11 motifs in each of 3 modes, and a sample solution motif found from partial machine enumerations within each of the 5 zipping modes for order-13.

Finally, in Section 6, we show that any one of these globally self-avoiding motifs can be composed with any other such motif, resulting in a new globally self-avoiding motif with its own zipping mode, thereby demonstrating that globally self-avoiding motifs are elements of their own submonoid $\mathcal{Z} \subset \mathcal{H}$.

²Peano-style curves continuously fill a square region from a corner to its diagonally opposite corner [11,16] and are constructible only for odd n . Their self-avoiding approximation paths are built, like the Hilbert Curve's, on the dual edge-adjacency graph. McKenna has also described a third, completely asymmetric, purely recursive traversal style (that relies on mirroring), for any odd or even n , that connects a square's corner with the center of an opposite side [9,10].

³In Hilbert-style constructions, we might call the tile path an *edge path*. But because there are other styles of square-filling curve construction where line segments mapped bijectively to tiles are not tile edges, *tile path* is a more general term.

1. Notation, Terminology, and Preliminary Observations

Definition 1.1. A tile T is a square with a bijection between itself and exactly one of its four edges, called its “distinguished edge” or (more generally) its “segment,” which imparts an orientation (bottom, right, top, left). To disambiguate which of two edge-sharing square tiles is associated with the distinguished edge of one of them, we illustrate the bijection using a small mark into the interior of its tile, see Fig. 3.

Definition 1.2. For each tile orientation, we also assign a direction to the tile’s segment or, equivalently, a “parity” to the tile itself. A counterclockwise segment imparts (arbitrarily) a “forward” parity to the tile; a clockwise segment imparts a “backward” parity. We denote these eight possibilities as BottomForward, RightForward, TopForward, LeftForward, BottomBackward, RightBackward, TopBackward, and LeftBackward; we also use the shorthand labels B, R, T, L, b, r, t, and l, respectively, see Fig. 3. Different shading also distinguishes parity, because sometimes a tile’s parity can be determined before we know which of its edges is distinguished.

	Bottom	Right	Top	Left	Partial
Forward					
Backward					

Figure 3: The eight orientation/parity combinations of a tile T .

Definition 1.3. When, due to nearby tile configurations or other constraints, a tile’s parity and orientation can be only one of the eight possibilities, we say that the tile is “forced.” When all tiles in a group are forced, we describe the group as “frozen.” If only the parity of a tile is known, the tile’s parity is “fixed.” If the orientation of a tile is known, whether we know its parity or not, we say the tile and its edge are “bound” to one another.

Definition 1.4. For $n \geq 1$, T_n denotes square T subdivided into an $n \times n$ array of equal-sized sub-tiles, $t_1 \dots t_{n^2}$, without regard to those sub-tiles’ orientations or parities. T_n is the “supertile” for each of its sub-tiles t_i .

Definition 1.5. A “thread” of length L within T_n is a sequence $t_1 \dots t_L$ of sub-tiles where (a) no sub-tile appears more than once in the sequence; (b) each pair $[t_i, t_{i+1}]$ of consecutive sub-tiles in the thread shares a corner at which the arrowhead of t_i ’s distinguished edge meets the arrowtail of t_{i+1} ’s distinguished edge, thereby creating a piecewise-connected path of coherently directed edges along some portion of the square lattice comprising all possible edges within T_n ; and (c) the sequence has no disconnected parts: every t_i in the thread can be reached following distinguished edges from t_1 , thereby ensuring piecewise-connectedness of all L sub-tiles.

Remark 1.1. In general, within T_n there are six ways that $[t_i, t_{i+1}]$ can be piecewise-connected, as shown in Fig. 4 for a BottomForward sub-tile t_i : one left turn, two straight ahead “turns,” two right turns, and one 180° U-turn. Two pairs are corner-only adjacent; four are edge adjacent. By symmetry, an equivalent set of possibilities exists for the other seven orientation/parity combinations that any given t_i might take on.

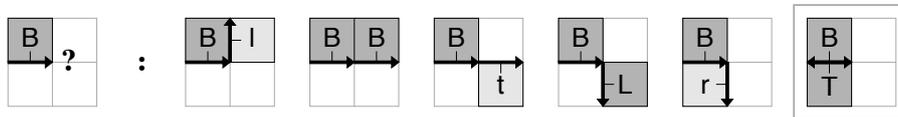


Figure 4: Six distinct ways to thread a second unique sub-tile in a thread’s consecutive pair.

In a thread, each of its t_i ’s occurs exactly once, but the distinguished edges of two sub-tiles (consecutive or not) can contact one another in the lattice of edges and even fall on top of one another, as the final boxed pair above in Fig. 4 shows, and as the Hilbert Curve’s tile path approximations in the top row of Fig. 2 illustrate. A thread can thus include edge sequences that self-contact, or even that cross each other to create parity-inverted loops.

Definition 1.6. When \mathbf{a} and \mathbf{b} are two adjacent corners of square tile T , a Hilbert-style “motif” within T_n is a thread of length n^2 where the arrowtail of t_1 begins at \mathbf{a} and the arrowhead of t_{n^2} ends at \mathbf{b} .

Remark 1.2. A motif in T_n implicitly specifies a distinguished edge $\mathbf{a} \rightarrow \mathbf{b}$ of T having one of four orientations and one of two parities. Hence a motif can also be characterized as BottomForward, RightForward, TopForward, LeftForward, BottomBackward, RightBackward, TopBackward, or LeftBackward with respect to T_n .

Definition 1.7. The “reversal” \overline{T} of a tile T flips its parity while maintaining its orientation. The “reversal” of a thread $t_1 \dots t_L$ is the thread $\overline{t_L} \dots \overline{t_1}$, which maintains the thread’s shape in T_n ’s edge lattice while reversing the direction of all the thread’s sub-tile arrows (i.e., all sub-tile parities). The “reversal” of a motif \mathbf{P} is denoted $\overline{\mathbf{P}}$.

Definition 1.8. A “composition” of an order- n motif \mathbf{J} with an order- m motif \mathbf{H} , denoted by $\mathbf{J} \circ \mathbf{H}$ or just \mathbf{JH} , replaces each of \mathbf{J} ’s sub-tiles j_i with a scaled-by- $\frac{1}{n}$ copy of \mathbf{H} having the same orientation/parity as j_i .

Remark 1.3. Motif composition accomplishes what is called “edge-replacement,” enhanced with directional coherence. Notice that composition is not commutative, as Fig. 5 shows, but is associative, as Fig. 6 shows.

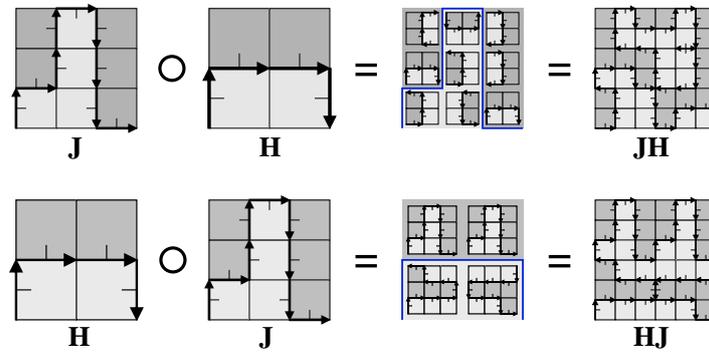


Figure 5: Composition of two different motifs is not commutative: $\mathbf{J} \circ \mathbf{H} \neq \mathbf{H} \circ \mathbf{J}$.

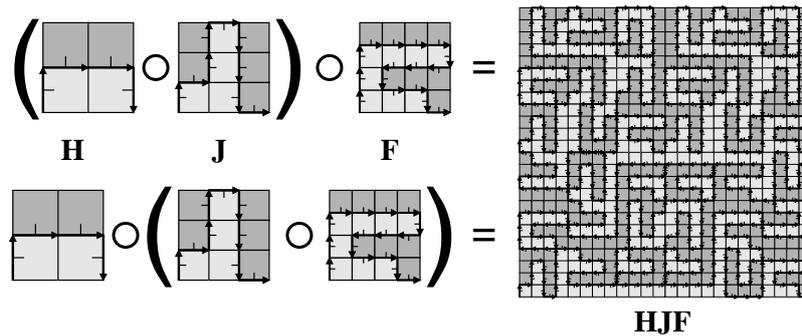


Figure 6: Composition is associative: $(\mathbf{HJ})\mathbf{F} = \mathbf{H}(\mathbf{JF}) = \mathbf{HJF}$ (result not drawn to scale).

Lemma 1.1. Under composition, Hilbert-style motifs are elements of a monoid \mathcal{H} .

Proof. A monoid is an algebraic structure $\langle \Sigma, \circ, \mathbf{I} \rangle$ comprising a set Σ of elements where any two elements can be composed using a binary operator \circ to result in another member of Σ , and where Σ contains an identity element \mathbf{I} . Here, because composition maintains the directional coherence of all arrows in the resulting tile path, an order- n motif composed with an order- m motif produces an order- nm motif, as Figs. 5, 6, and 54 illustrate. As Fig. 7 shows, a BottomForward tile is an order-1 motif that acts as the monoid \mathcal{H} ’s identity element \mathbf{I} . ■

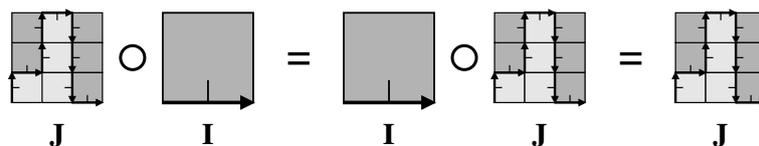


Figure 7: Composing the order-3 motif \mathbf{J} with the identity motif \mathbf{I} in either order results in that same motif \mathbf{J} .

Lemma 1.2. For any motif $\mathbf{P} \neq \mathbf{I}$, \mathbf{P}^k converges as $k \rightarrow \infty$ to a self-similar, square-filling, continuous curve.

Proof. This follows directly from a straightforward generalization (for any n) of the usual Hilbert Curve proof of continuity for any space-filling curve construction where recursively nested bounding boxes apply [18, 20]. ■

Definition 1.9. A “locally self-avoiding” motif has no two sub-tile segments coinciding in T_n ’s edge lattice, and no more than one arrowhead (arrowtail) arriving at (leaving from) a lattice vertex. \mathbf{H} , \mathbf{J} , and \mathbf{F} in Fig. 6 are each an example of a locally self-avoiding motif, whereas the result of composing any two of them is not.

Definition 1.10. A “globally self-avoiding” order- n motif \mathbf{G} is one for which \mathbf{G}^k is a locally self-avoiding order- n^k motif, for any integer $k \geq 0$. By convention, $\mathbf{G}^0 = \mathbf{I}$.

Remark 1.4. The idempotent identity motif $\mathbf{I} = \mathbf{I}^k$ is both locally and globally self-avoiding: no vertex in the 1×1 lattice for \mathbf{I} has more than one arrowhead or arrowtail leaving or arriving at that vertex. Also notice that \mathbf{G}^k must itself be a globally self-avoiding motif. But for $\mathbf{G} \neq \mathbf{I}$, k must be finite, because when $k \rightarrow \infty$, the limiting space-filling curve (to which the sequence of tile paths converges) is surjective (a form of self-contact).

Definition 1.11. A “distinct” motif refers to a BottomForward motif \mathbf{P} and its seven other orientation/parity variants having the same shape, as well as \mathbf{P} ’s mirror image \mathbf{Q} and its seven orientation/parity variants having the same shape. To count distinct motif shapes is to ignore orientation, mirrors, and reversals.

Remark 1.5. The purpose of discounting mirror images is because in any machine enumeration, there is no point expending resources looking for mirror images of already-found motifs. As it turns out, there are no bilaterally symmetric, globally self-avoiding \mathbf{G} s, so a count of distinct motifs will always be half of the number of motifs otherwise blindly enumerated.

2. Enumerating Globally Self-Avoiding Motifs For Small n

Motif composition replicates its right operand’s shape everywhere that a tile’s distinguished edge appears in the left operand, so disregarding the self-avoiding shape of the left operand, self-avoidance of any composition’s result depends at the very least upon the right operand’s shape and how it interacts with adjacent copies of itself in variously oriented, multiple instances. The simplest form of composition is exponentiation of a single motif \mathbf{G} , so we focus on \mathbf{G}^k to better understand the constraints on self-avoidance under composition in general.

Because $\mathbf{G}^k = \mathbf{G}^{k-2}\mathbf{G}^2$ (for $k \geq 2$), if \mathbf{G}^2 is not locally self-avoiding, then \mathbf{G} cannot be globally self-avoiding. In other words, once self-contact muddies a tile path’s waters, iterated motif composition leaves those waters forever muddied with congruent replicas of that self-contacting tile path.

The simplest distinct motif for T_2 is the Hilbert Curve’s $\mathbf{H} = [\mathbf{l} \ \mathbf{B} \ \mathbf{B} \ \mathbf{r}]$, see Fig. 8; \mathbf{H} is locally self-avoiding, bilaterally symmetric, and unique, but it is not globally self-avoiding. And for T_3 there are just two locally self-avoiding, motifs: $\mathbf{J} = [\mathbf{l} \ \mathbf{B} \ \mathbf{l} \ \mathbf{R} \ \mathbf{t} \ \mathbf{L} \ \mathbf{L} \ \mathbf{r} \ \mathbf{B}]$ and its mirror $\mathbf{I} = [\mathbf{B} \ \mathbf{l} \ \mathbf{R} \ \mathbf{R} \ \mathbf{t} \ \mathbf{L} \ \mathbf{r} \ \mathbf{B} \ \mathbf{r}]$. So again there is only one distinct, locally self-avoiding motif, as shown in Fig. 9. But \mathbf{J}^k is, like \mathbf{H}^k , self-contacting when $k \geq 2$.

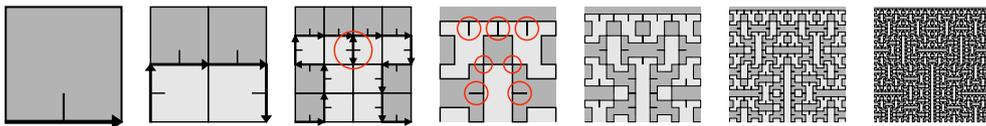


Figure 8: The Hilbert Curve construction \mathbf{H}^k for $0 \leq k \leq 6$, where $\mathbf{H} = [\mathbf{l} \ \mathbf{B} \ \mathbf{B} \ \mathbf{r}]$. Self-contacting tile paths (circled) occur on T_{2^k} for all $k \geq 2$.

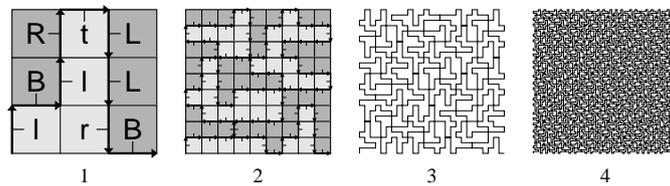


Figure 9: \mathbf{J}^1 , \mathbf{J}^2 , \mathbf{J}^3 , and \mathbf{J}^4 for the locally, but not globally, self-avoiding order-3 motif $\mathbf{J} = [\mathbf{l} \ \mathbf{B} \ \mathbf{l} \ \mathbf{R} \ \mathbf{t} \ \mathbf{L} \ \mathbf{L} \ \mathbf{r} \ \mathbf{B}]$.

An exhaustive, brute-force computer backtrack enumeration,⁴ for locally self-avoiding, distinct motifs for T_4 and T_5 finds exactly 7 and 138 solutions, respectively. Fig. 10 shows all seven for T_4 ; Fig. 11 shows all those for order-5. These are distinct solutions, so their otherwise-valid mirror images are not shown.

There are no globally self-avoiding motifs for T_4 , but for T_5 , McKenna [8, 17] found that exactly one of the 138 distinct, locally self-avoiding motifs, $\mathbf{E} = [\mathbf{l} \ \mathbf{l} \ \mathbf{B} \ \mathbf{L} \ \mathbf{t} \ \mathbf{l} \ \mathbf{R} \ \mathbf{b} \ \mathbf{R} \ \mathbf{R} \ \mathbf{t} \ \mathbf{L} \ \mathbf{t} \ \mathbf{B} \ \mathbf{B} \ \mathbf{r} \ \mathbf{T} \ \mathbf{b} \ \mathbf{r} \ \mathbf{B} \ \mathbf{L} \ \mathbf{b} \ \mathbf{r} \ \mathbf{B} \ \mathbf{B}]$ (Fig. 12),

⁴As related in an endnote in [8], McKenna first performed this search in 1978. It required an hour of CPU time on a PDP-20 mainframe computer.

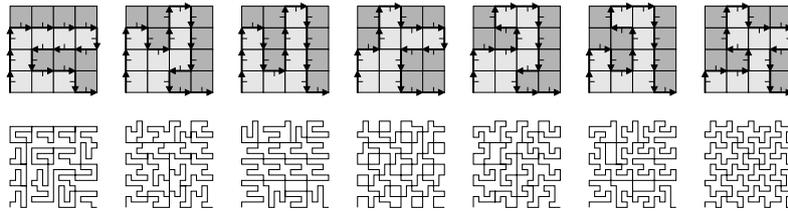


Figure 10: Top: the seven distinct locally self-avoiding motifs for T_4 . Bottom: One exponentiation for each shows that none is globally self-avoiding, although the rightmost, which has the fewest tile segments along T_4 's perimeter, comes tantalizingly close.

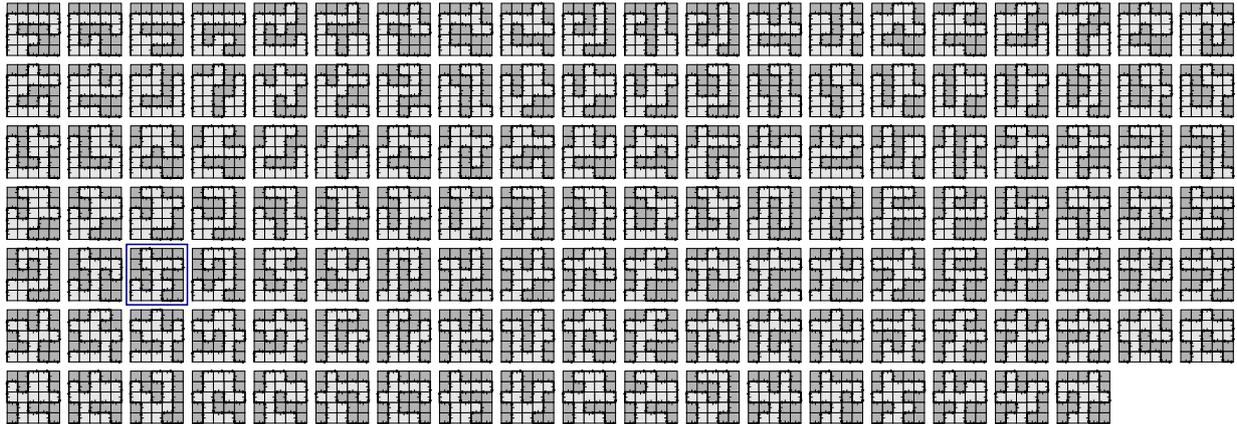


Figure 11: 138 distinct, locally self-avoiding motifs for T_5 (mirror images not shown). Exactly one of them (highlighted, third row/column from bottom/left) is globally self-avoiding (see also Fig. 9 of [8]).

is globally self-avoiding.⁵ The same computer search finds 5690 distinct, locally self-avoiding motifs on T_6 , but, perhaps unexpectedly, none is globally self-avoiding [8].⁶ This paper reports that based on exhaustive enumerations of T_7 and T_8 , there are respectively exactly 484,760 and 85,942,455 distinct, locally self-avoiding motifs. Of the 484,760 for T_7 , only seven are globally self-avoiding when exponentiated, which comports with an unpublished drawing of Kim [5] made during discussions with the author in 1983. There are no globally self-avoiding motifs on T_8 . By the next sections' results, we construct \mathbf{E} and prove it unique, and we construct the seven order-7 distinct motifs, all without resorting to computer enumeration (see Figs. 46 and 47 below).

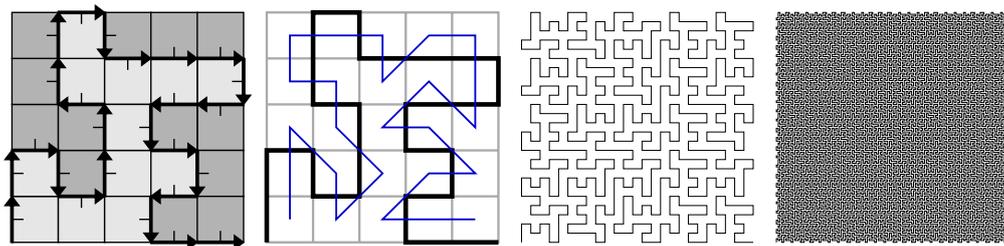


Figure 12: Left: The only distinct, globally self-avoiding motif \mathbf{E} for T_5 . Left middle: The tile path of \mathbf{E}^1 with its dual-path superimposed (in blue). Right: \mathbf{E}^2 and \mathbf{E}^3 , whose tile paths are self-avoiding on the lattice.

For T_9 , the number of distinct, locally self-avoiding motifs is not yet known, but a combinatorial explosion of distinct globally self-avoiding motifs begins in full force, with exactly 10,101 found using two different backtrack graph searching algorithms, neither of which relied on the Fibonacci Zipper Theorem proved below.

⁵Mandelbrot, with whom the author corresponded about this discovery in 1979, once proposed to call this construction the Peano-McKenna curve, although its mention did not make it into [7]. The \mathbf{E} motif (in mirror form) was later independently described by Dekking [2] as a "Gosper-style" construction (because Gosper's "Flowsnake" curve is a self-avoiding tile path based on the hexagonal tiling using pure edge replacement without mirroring). This resulted in an inadvertent misattribution via a musical composition's title [14].

⁶A typographic error in [8] reported the count as 5960. This enumeration, first performed by the author in 1980 on a 3081 mainframe computer at IBM Research, took approximately 50 hours of CPU time.

3. Fibbinary Zippers

A browsable visual display of all 10,101 distinct, globally self-avoiding motifs for T_9 revealed something unexpected: they fall into just two visual “rhythmic” modes, each strikingly different from the other. Fig. 13 illustrates an example of each mode. Motifs in the second mode are more visually “energetic.” They evince, for example, a much higher frequency of right or left tile path turns. Indeed, it is impossible, upon seeing this for the first time while scrolling through the entire solution space in enumerated lexicographic order, not to wonder whether one is witnessing some kind of “phase transition” in a discrete solution space, which is why the word *mode* is appropriate. Every member of a mode is characterized by exactly the same perimeter, or boundary, tiles as the others in that mode, a discovery that begged for further investigation of its “crystalline” nature. Additionally, one mode had an order of magnitude fewer solution motifs in it (660 vs. 9441), indicating—all else being equal—significantly less combinatorial freedom (i.e., the system is somehow more constrained) than does the second mode. And the absence of any solutions for $n = 2, 4, 6, 8$ naturally led to the conjecture that when n is even, no globally self-avoiding motifs exist, which is proved in Lemma 3.27 below.

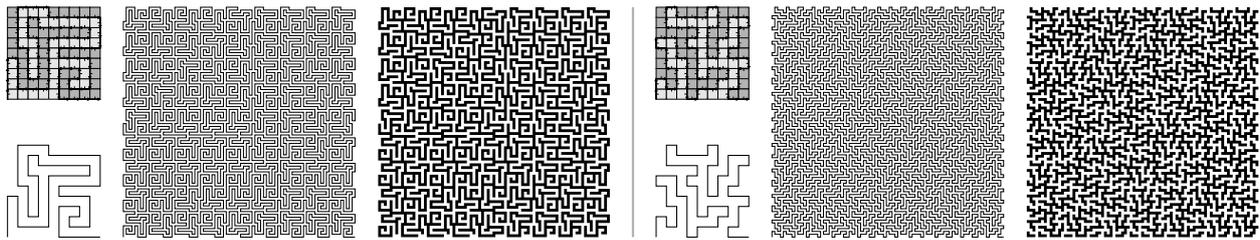


Figure 13: A sample globally self-avoiding G and G^2 (as path and filled in on one side, thereby showing parities) within T_9 , one for each of two modes into which all 10,101 distinct, globally self-avoiding motifs fall.

3.1 Edge contact combinations within motifs

When any motif P is composed with another motif Q , in PQ a reduced, oriented, sometimes reversed, copy of Q 's shape replaces each of P 's distinguished edges. Thus, two copies of Q 's tile path will contend for portions of the same shared supertile edge. So P 's global self-avoidance under exponentiation depends on how the segments in T_n 's outermost layer of $4(n - 1)$ perimeter sub-tiles are arranged with respect to other variously oriented copies of themselves along the shared edges and at the corners of supertiles. These perimeter tiles are where the motif's tile path has the opportunity to make contact with a neighboring copy of itself, but must not.

Given two edge-adjacent square tiles, each with four possible orientations (Bottom, Right, Top, or Left), there are $4^2 = 16$ ways one tile can share an edge with its neighbor. But as Fig. 14 shows, of these sixteen combinations only ten are distinct, because any pair's *combination* is invariant as the pair is rotated together through a multiple of 90° . For example, a Right-Right pair (upper right entry) creates the same contact combination as a Left-Left pair (lower left entry), regardless of the pair's 180° difference in rotation.⁷

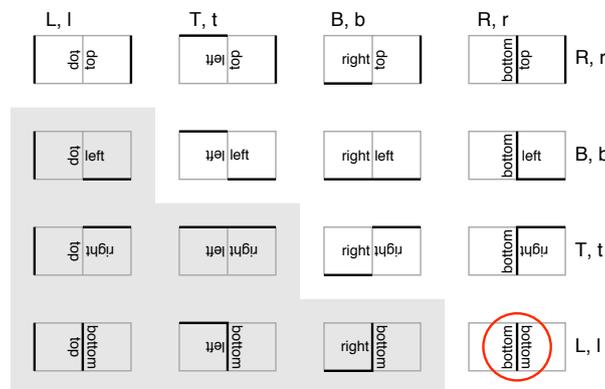


Figure 14: Of the sixteen possible ways two threaded—but not necessarily consecutive—tiles (regardless of parity) can be edge adjacent, six pairs (below diagonal, in gray) are invariant under 180° rotation to their counterparts reflected across the diagonal, and one (circled) represents a touching tile path.

⁷Compare with Fig. 2 in [12], re Hilbert-style curves with always edge-adjacent sub-tiles that require top:top combinations.

Definition 3.1. *The possible ways that a globally self-avoiding motif \mathbf{G} might interact with copies of itself in a self-avoiding manner along shared supertile edges are called “contact combinations.” They depend on \mathbf{G} ’s shape and orientation, but not on parity. Reading across and down the upper half of the 4×4 matrix, we denote these ten contact combinations as: top:top, left:top, right:top, bottom:top; left:left, right:left, bottom:left; right:right, bottom:right; and bottom:bottom. But bottom:bottom represents a self-contacting tile path, and is therefore not a possible edge-adjacency pair of two sub-tiles in any self-avoiding, much less globally self-avoiding, motif.*

With respect to globally self-avoiding motifs, the rub, so to speak, is that it is not *a priori* obvious which of the nine possible contact combinations occur in $\mathbf{G}^{k>0}$. Any given pair of edge-adjacent tiles in T_{n^k} are not necessarily consecutive nor even sequentially close to one another along the motif’s threaded tile path. Worse, under exponentiation, the placement of pairs of copies of \mathbf{G} is self-referentially governed by the shape of \mathbf{G} itself, so there might be different distributions of contact combinations based on myriad different motif shapes. And if a contact combination of supertiles were *never* to occur for some shape or under some condition, this might relax the perimeter boundary conditions in some way to permit some self-avoiding shape that otherwise might not be possible during composition. So a convincing proof of the following theorem must examine each of these nine contact combinations to determine which of them—*regardless of overall motif shape*—always (or never) occur, which are merely optional, and then explore the consequences of the mandatory constraints.

Because $\mathbf{G}^{k \geq 2} = \mathbf{G}^{k-1}\mathbf{G} = \mathbf{G}^{k-2}\mathbf{G}^2$, we can look for these contact combinations among pairs of tiles in either \mathbf{G} or \mathbf{G}^2 . If we can show that some combination always occurs somewhere in either \mathbf{G} or \mathbf{G}^2 , it must occur in all $\mathbf{G}^{k \geq 2}$. A sufficient “covering” of required contact combinations leads to our main result.

3.2 Fibbinary Zipper Theorem

Theorem 3.1. *For all odd $n \geq 1$, the tile path of any globally self-avoiding motif \mathbf{G} has a shape whose four-sided boundary, comprising $4(n - 1)$ sub-tiles of T_n , is frozen by exactly one of $F_{(n-3)/2}$ zipping modes, where F_i is the i th Fibonacci number ($F_{-1} = 1, F_0 = 0, F_1 = 1$, and $F_i = F_{i-1} + F_{i-2}$). Each mode is characterized by an n -bit, palindromic, Fibbinary bit pattern (i.e., in which no 1 bits are adjacent) whose central bit is a 1 and, when $n \geq 5$, whose first and last two bits are 0s. When n is even, there are no globally self-avoiding motifs.*

We prove this by defining or proving lemmas representing mostly prohibitive constraints.⁸ In the figures, a (red) circle indicates where an impossible or contradictory condition occurs.

Lemma 3.1. OneToOneNess. *By Definition 1.1, in $\mathbf{G} = [t_1, \dots, t_{n^2}]$ there is a bijective correspondence between each sub-tile t_i and t_i ’s distinguished edge.*

Lemma 3.2. ArrowHeads. *By Definitions 1.5 and 1.6, each consecutive pair $[t_i, t_{i+1}]$ of tiles in \mathbf{G} has tile edges that are pairwise-connected, whose arrowheads connect only with arrowtails, and vice-versa, and with t_1 ’s arrowtail at \mathbf{a} and t_{n^2} ’s arrowhead at \mathbf{b} .*

Lemma 3.3. LongestDistance. *By Definition 1.6, \mathbf{G} ’s tile path within T_n must be n^2 in length.*

Lemma 3.4. SelfAvoidance. *By Definition 1.9, no tile’s corner has more than one tile segment arrowhead leading to it, nor more than one tile segment arrowtail leading away from it.*

Lemma 3.5. NoPlanarCycles. *By Definition 1.5, \mathbf{G} cannot contain a disconnected cycle in the plane.*

Lemma 3.6. InsideIsNotOutside. *Because \mathbf{G} is a globally self-avoiding motif, two edge-adjacent, though not necessarily consecutive, tiles in \mathbf{G}^k are of opposite parity if and only if exactly one of them has its distinguished edge along the two tiles’ shared edge in T_n ’s edge lattice.*

Proof. This follows directly from the polygonal version of the Jordan Curve Theorem [3], because \mathbf{G} is assumed a BottomForward self-avoiding piecewise-connected path, whose open ends at \mathbf{a} and \mathbf{b} can always be extended in the plane into a closed, oriented, self-avoiding loop. If a counterclockwise (clockwise) loop closure, all Forward (Backward) tiles will be on the outside of the loop, and all Backward (Forward) tiles will be on the inside. ■

⁸These constraints, whether by construction or proved, are numbered sequentially as lemmas as well as identified with descriptive phrase names in **CamelCase**. Because they are about to be referred to ubiquitously, later references forego using the words *definition*, *lemma*, or *constraint* in favor of providing the reference number as an innocuous subscript to a more memorable name. This concisely enhances readability and memorability while still allowing the reader to find the reference in this exposition.

Lemma 3.7. FourCorners. *The first tile t_1 and the last tile t_{n^2} in a motif \mathbf{G} can each be only one of two possible orientation/parity combinations: tile t_1 is either LeftBackward or BottomForward, and t_{n^2} is either BottomForward or RightBackward. Similarly, the upper left and upper right corner tiles in \mathbf{G} can each be only one of two possible orientation/parity combinations: the upper left corner tile can be either RightForward or BottomForward, and the upper right corner tile can be either BottomForward or LeftForward.*

Proof. As Fig. 15 shows, a segment can leave \mathbf{a} or arrive at \mathbf{b} in only two ways. A segment cannot arrive at an upper left or right corner of T_n and still continue the tile path without violating **OneToOneNess**_{3.1}. ■

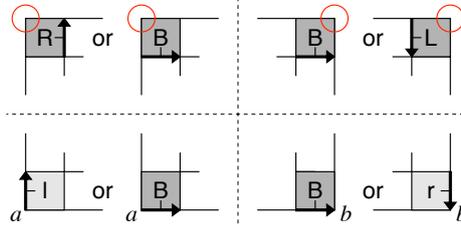


Figure 15: Each of \mathbf{G} 's four corner tiles can be in only one of two orientation/parity combinations. The upper left and right corners of T_n (circled) are not available to any tile path.

Remark 3.1. *For $\mathbf{G} = \mathbf{I}$ ($n = 1$), these four corner tiles all collapse into only one possible BottomForward tile.*

Lemma 3.8. PerimeterForces. *If in \mathbf{G} a sub-tile t_i 's segment travels along an edge of its supertile T_n , then t_i is forced into only one orientation and parity. Similarly, if t_i 's segment travels parallel to the perimeter but one sub-tile width into the interior, t_i is also forced. The same holds for $\overline{\mathbf{G}}$, but with all parities reversed.*

Proof. A segment on or near T_n 's perimeter must be bound to a perimeter sub-tile t_i of T_n . So t_i can have only one orientation, all shown in Fig. 16. But for each orientation, only one parity is possible. For each perimeter t_i , its reversal $\overline{t_i}$ precludes the tile path from reaching \mathbf{b} without violating **SelfAvoidance**_{3.4}. ■

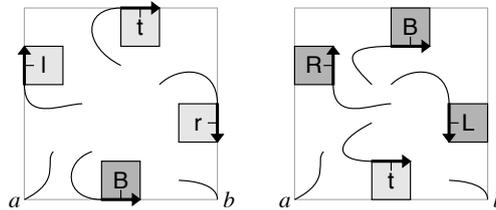


Figure 16: In order for \mathbf{b} to remain reachable from \mathbf{a} without violating **SelfAvoidance**_{3.4} or **ArrowHeads**_{3.2}, a sub-tile whose distinguished edge is on, or one sub-tile width away from, T_n 's perimeter is forced.

Lemma 3.9. PerimeterFlow. *If the arrowhead of a sub-tile's segment arrives at the perimeter of T_n from somewhere within its interior, the orientation/parity of that sub-tile is forced. If the arrowtail of a sub-tile's bound edge leaves from the perimeter of T_n into somewhere in T_n 's interior, that sub-tile is forced. For a BottomForward motif \mathbf{G} , the "flow" is thus always clockwise for the left, top, and right edges of T_n , and always counterclockwise along its bottom.*

Proof. See the left half of Fig. 17 for perimeter arrivals and the right half for departures. Consider a tile path with a segment that heads into the left side of T , somewhere above the corner \mathbf{a} (which is precluded from being the endpoint by **FourCorners**_{3.7}). There are two possible orientation/parity combinations for that tile segment: TopForward or BottomBackward. If the tile were BottomBackward (i.e., the tile labeled \mathbf{b}), the next segment of the path would not be able to make the right turn required by **PerimeterForces**_{3.8}. \mathbf{G} 's path from \mathbf{a} would therefore never be able to reach its destination at \mathbf{b} without violating **SelfAvoidance**_{3.4}. Recalling that by **FourCorners**_{3.7} \mathbf{b} cannot be an endpoint of a tile segment arriving from the interior of T , a similar argument holds for the other three sides of T . ■

Lemma 3.10. BottomToTopParity. *If two edge-adjacent sub-tiles form a bottom:top combination, then they must have opposite parity.*

Proof. This follows directly from **InsideIsNotOutside**_{3.6}. See Fig. 18. ■

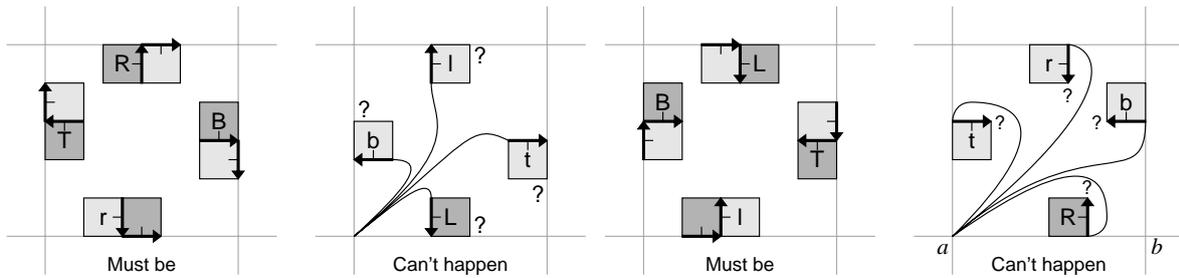


Figure 17: Tiles with segments that arrive at (left) or depart from (right) T_n 's perimeter are forced so that the tile path can continue on to eventually reach b from a .

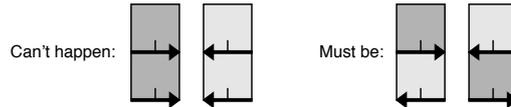


Figure 18: Every **bottom:top** contact combination requires opposite parities.

Lemma 3.11. PerimeterBottomToTops. *If a sub-tile's segment travels along the perimeter of T_n and not through a or b , then the edge-adjacent sub-tile on the other side of the perimeter in some adjacent copy of G or its reversal \bar{G} (in whatever orientation) in G^2 must have the same orientation but opposite parity, thereby creating a **bottom:top** contact combination.*

Proof. This follows from **BottomToTopParity**_{3.10} in G^2 . See Fig. 19. **FourCorners**_{3.7} precludes a tile segment from traveling along T_n 's edge into or away from either T_n 's upper left or upper right corner. Elsewhere along the perimeter, if the parities are the same, **BottomToTopParity**_{3.10} is violated in G^2 . Of the four remaining possible orientations for an edge-adjacent, cross-perimeter tile, three of them would create self-contact at the perimeter boundary. The only remaining possibility is the tile with the same orientation, but with necessarily opposite parity, so that the two tile segments are parallel, creating a **bottom:top** combination. ■

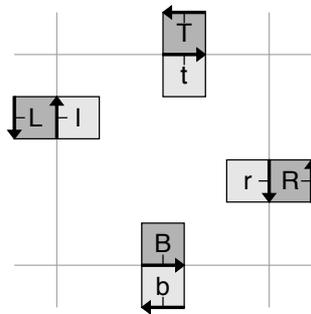


Figure 19: Tiles with segments on T_n 's perimeter (but not the lower left or lower right tiles) induce **bottom:top** combinations with their cross-perimeter neighbors in an adjacent copy of G .

Lemma 3.12. RightLeftContact. *Any G must have at least one **right:left** contact combination in some pair of edge-adjacent, necessarily consecutive, tiles in T_n .*

Proof. Ignoring $n = 1$ (remarked upon below), and recalling that there are no G s for $n = 2$ or $n = 3$, consider $n \geq 4$. Assume no **right:left** contact combination in G occurs. This is equivalent to asserting that no two consecutive threaded tiles have the same orientation and parity. So consider a motif's tile t_m (for some m) situated at the upper left corner of T_n , labeled D in Fig. 20. By **FourCorners**_{3.7}, there are two cases: D can be either **RightForward** or **BottomForward**. In both cases, under our assumption, a variety of later and earlier tiles in the thread become forced, such that in G^2 , self-contact occurs, violating **SelfAvoidance**_{3.4}. ■

Remark 3.2. *At first glance, this condition would not appear to apply to I for T_1 . But once we show below that globally self-avoiding motifs map to toroidal cycles, the left and right sides of T_1 become edge-adjacent.*

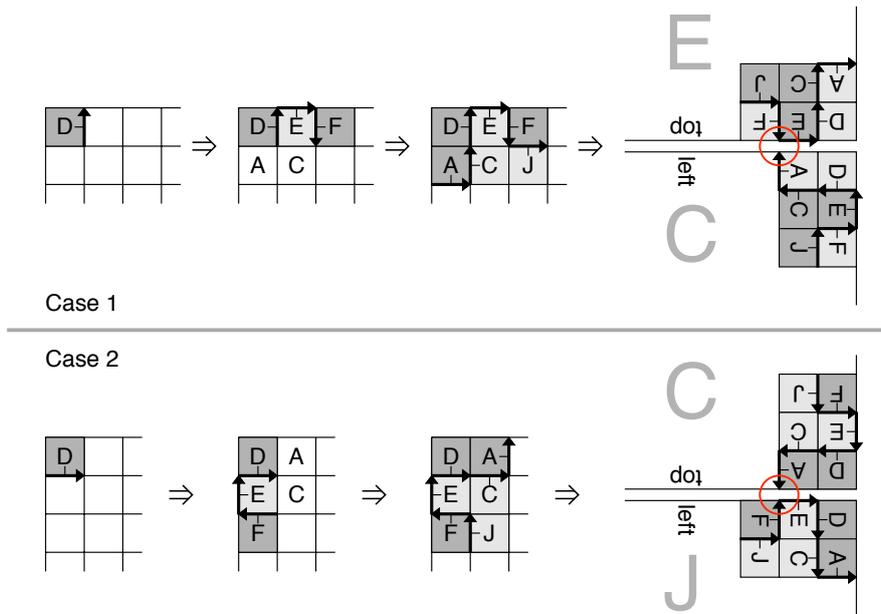


Figure 20: When assuming that no `right:left` contact combination occurs, both upper left corner tile possibilities (`RightForward` or `BottomForward`) create a group of frozen tiles that lead to self-contact in \mathbf{G}^2 .

Lemma 3.13. NoDoubleLeftRights. *The first tile t_1 of \mathbf{G} cannot be `LeftBackward` at the same time as the last tile t_{n^2} is also `RightBackward`.*

Proof. By `RightLeftContact`_{3.12}, both the first and last tile segments will coincide in \mathbf{G}^2 along the shared border of some `right:left` contact combination of supertiles, violating `SelfAvoidance`_{3.4} (see Fig. 21). ■

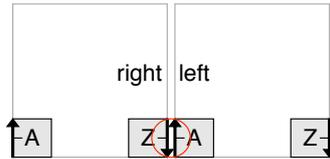


Figure 21: \mathbf{G} 's first and last tile cannot both have vertical tile segments or self-contact will occur in \mathbf{G}^2 .

Lemma 3.14. BottomTopContact. *Any \mathbf{G} must have at least one `bottom:top` contact combination in some pair of edge-adjacent tiles.*

Proof. Again, we need only consider $n \geq 4$. Assume \mathbf{G} has no `bottom:top` contact combinations among any edge-adjacent pair of tiles in T_n . By `FourCorners`_{3.7}, there are two cases: the m th tile t_m (for some m) in the upper left corner, labeled tile D in Fig. 22, is either (1) `RightForward` or (2) `BottomForward`.

In case (1), `PerimeterForces`_{3.8} requires E to be `TopBackward`. t_{m-1} leading to D must have a vertical distinguished edge, otherwise `PerimeterFlow`_{3.9} or our no `bottom:top` assumption is violated. So t_{m-1} is either A, or C (sub-case (a) or (b)). But in either sub-case, further tiles are frozen such that in sub-case (a) self-contact occurs in \mathbf{G}^2 between adjacent motifs at tiles E_E and C_A , or in sub-case (b) our assumption of no `bottom:top` contact is violated after going through the possible ways that various tiles become frozen.

In case (2) (bottom of Fig. 22) choose D as `BottomForward`. But a mirror of the case 1 argument used above applies to the upper *right* tile if that corner tile is assumed `LeftForward`. Therefore neither upper left nor upper right corner tiles can be vertical tile segments. Both must be `BottomForward`, which results in the tile path touching the right and left sides of T_n the same distance above its base, creating self-contact somewhere by `RightLeftContact`_{3.12}. Hence, our initial assumption must be false. ■

Remark 3.3. *Again, as in `RightLeftContact`_{3.12}, `BottomTopContact`_{3.14} would not appear to apply to \mathbf{I} . But it does because toroidally the top and bottom sides of T_1 are edge-adjacent. See `ToroidalCycles`_{3.17} below.*

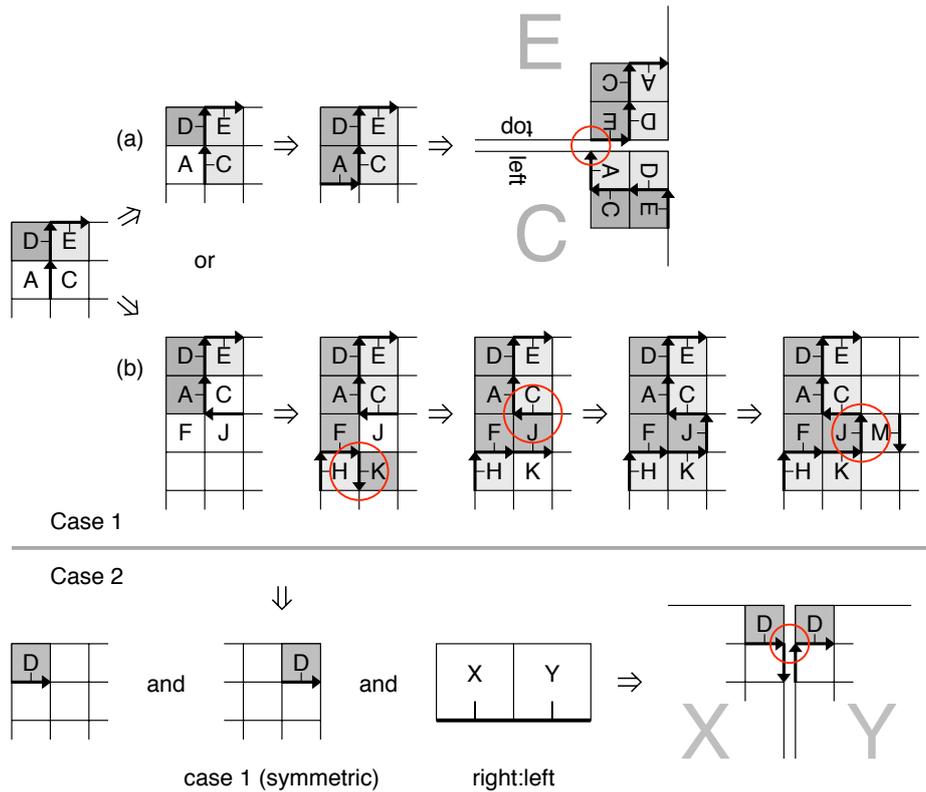


Figure 22: **bottom:top** must always occur. Otherwise, either self-contact in \mathbf{G}^2 or a contradiction occurs.

Lemma 3.15. NoDoubleBottoms. *When $\mathbf{G} \neq \mathbf{I}$, \mathbf{G} 's first tile t_1 cannot be BottomForward at the same time as the last tile t_{n_2} is also BottomForward.*

Proof. Assume that t_1 and t_{n_2} (A_D and Z_D in Fig. 23) are both BottomForward. By **BottomTopContact**_{3.14} and **SelfAvoidance**_{3.4}, the upper two corner sub-tiles (C_D and D_D) must be the same orientation and parity as A_D and Z_D . But by **RightLeftContact**_{3.12}, there will be self-contact at F_D and E_C . ■

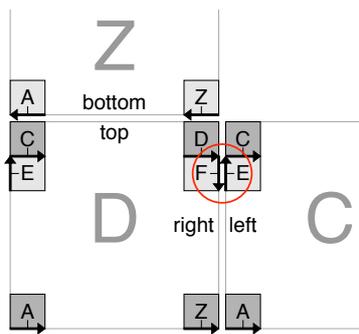


Figure 23: A motif \mathbf{G} 's first and last tile cannot both be BottomForward or self-contact occurs in \mathbf{G}^2 .

Lemma 3.16. Handedness. *Left-handed \mathbf{G} s end with a BottomForward sub-tile; right-handed begin with a BottomForward sub-tile. The identity motif \mathbf{I} is thus ambidextrous, and $\mathbf{G} \neq \mathbf{I}$ cannot be bilaterally symmetric.*

Proof. This follows from **NoDoubleLeftRights**_{3.13} and **NoDoubleBottoms**_{3.15}, see Fig. 24. ■



Figure 24: Every globally self-avoiding, left-handed motif has a right-handed (chiral) mirror image.

Lemma 3.17. ToroidalCycles. \mathbf{G} 's tile path maps to a Hamiltonian cycle on the $n \times n$ toroidal grid-graph \mathbb{T}_n .

Proof. By **BottomTopContact**_{3.14} and **RightLeftContact**_{3.12}, tile vertices in the threaded tile path along any given side of \mathbb{T}_n are coincident at least once in \mathbf{G}^2 with their translationally respective vertices on the opposite side of \mathbb{T}_n . So a tile path in \mathbf{G}^2 can only comply with **SelfAvoidance**_{3.4} when exactly one or the other vertex—but not both—is included in the tile path. Hence we can map integer corner coordinates (i, j) of \mathbb{T}_n to $(i \bmod n, j \bmod n)$ on \mathbb{T}_n , the $n \times n$ toroidal grid graph, without violating **SelfAvoidance**_{3.4}.

By **FourCorners**_{3.7}, t_1 's distinguished edge can only leave from \mathbf{a} and t_{n^2} 's can only arrive at \mathbf{b} . But all four corner points of \mathbb{T} map toroidally to the origin $\mathbf{a} = \mathbf{b}$. Thus the mapped motif becomes a cycle among \mathbb{T}_n 's vertices with **ArrowHeads**_{3.2} applying at $\mathbf{a} = \mathbf{b}$ no differently from other lattice points: $[t_{n^2}, t_1]$ is now a consecutive, adjacent pair of threaded tiles in a single self-avoiding loop. Because \mathbb{T}_n has exactly n^2 vertices, and self-avoiding \mathbf{G} comprises exactly n^2 distinguished edges, \mathbf{G} is a Hamiltonian cycle on \mathbb{T}_n . ■

Definition 3.2. We call any 2×2 group of sub-tiles that are all the same parity a “nodule.” A nodule has a lattice point in the nodule’s center where four sub-tile corners meet.

Lemma 3.18. NoNodules. A nodule in a motif \mathbf{G} or in \mathbf{G}^2 prevents \mathbf{G} from being globally self-avoiding.

Proof. By **InsideIsNotOutside**_{3.6}, no distinguished edge in \mathbf{G} 's tile path can pass through a nodule’s central point when all sub-tiles are the same parity. This leaves only $n^2 - 1$ lattice points in the rest of \mathbb{T}_n available for \mathbf{G} 's tile path to pass through, which violates **ToroidalCycles**_{3.17}. ■

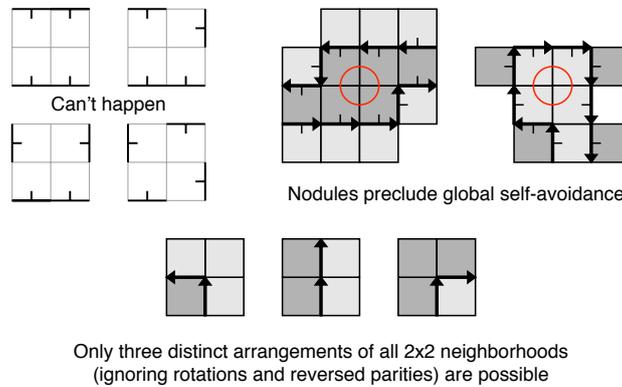


Figure 25: A globally self-avoiding motif must pass through every lattice point in a toroidal Hamiltonian cycle.

Remark 3.4. For any exponentiated motif (self-avoiding or not) of length n^2 , the number of vertices where self-contact occurs in the toroidal grid-graph equals the number of nodules. So the existence of at least one nodule implies self-contact elsewhere in a tile path. For instance, ignoring the highlighted **E** motif in Fig. 11, each of the remaining 137 distinct, non-globally-self-avoiding, order-5 motifs has—either directly or when exponentiated once—at least one nodule.

Lemma 3.19. TopEdgeTouched. For $\mathbf{G} \neq \mathbf{I}$, there must be a **TopBackward** sub-tile in \mathbb{T}_n 's top row.

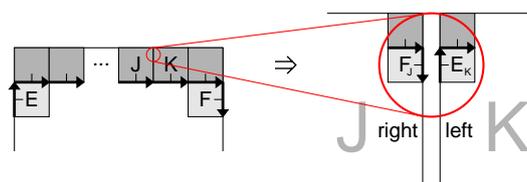


Figure 26: The assumption that a motif can avoid the top edge of \mathbb{T}_n leads to self-contact in \mathbf{G}^2 .

Proof. Assume that no sub-tile in the top row of \mathbb{T}_n has a **TopBackward** distinguished tile edge. All of those sub-tiles must then, by **PerimeterForces**_{3.8} and **PerimeterFlow**_{3.9}, be **BottomForward** (Fig. 26, top left), with E and F also forced. **RightLeftContact**_{3.12} then causes self-contact in \mathbf{G}^2 (Fig. 26, top right, circled). ■

Remark 3.5. This is one reason why $\mathbf{H}^{k \geq 2}$ (the Hilbert Curve’s exponentiated tile path) is self-contacting.

Lemma 3.20. BottomRightContact. Any $\mathbf{G} \neq \mathbf{I}$ must have a `bottom:right` contact combination in some pair of edge-adjacent tiles.

Proof. By **TopEdgeTouched**_{3.19} and **PerimeterForces**_{3.8}, at least one `TopBackward` tile segment must pass along the top edge of T_n (Fig. 27, left). By **PerimeterFlow**_{3.9}, the tile path’s necessary right turn creates a `bottom:right` contact combination. ■

Lemma 3.21. BottomLeftContact. Any $\mathbf{G} \neq \mathbf{I}$ must have a `bottom:left` contact combination in some pair of edge-adjacent tiles.

Proof. By **TopEdgeTouched**_{3.19} and **PerimeterForces**_{3.8}, at least one `TopBackward` segment must pass along the top edge of T_n (Fig. 27, right). By **PerimeterFlow**_{3.9}, the tile path’s turn away from the last `TopBackward` segment creates a `bottom:left` contact combination. ■

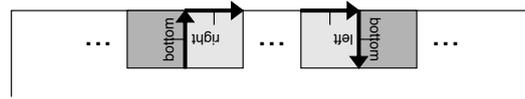


Figure 27: Arrival and departure for at least one required tile segment along the top edge of T_n induces `bottom:right` and `bottom:left` zipping, respectively.

Lemma 3.22. NoLowerCornerTurns. In any $\mathbf{G} \neq \mathbf{I}$, neither $[t_1, t_2]$ nor $[t_{n^2-1}, t_{n^2}]$ can have consecutive tile segments that form a turn, i.e., both $[t_1, t_2]$ and $[t_{n^2-1}, t_{n^2}]$ must form a `right:left` contact combination.

Proof. Again, we can assume $n \geq 4$. For each of the first and last consecutive tile pairs in \mathbf{G} , there is only one way to make a right turn. Consider first the lower left corner of T_n , and assume that \mathbf{G} ’s first two tiles, t_1 and t_2 (labeled A and E respectively in Fig. 28) make a right turn: $\mathbf{G} = [l B \dots]$. Tiles E and A thus form a `bottom:left` contact combination. Within \mathbf{G}^2 , whatever arrangement of tiles that have been composed into E must also occur in A, rotated 90° clockwise but with all parities (arrows) reversed because t_1 ’s parity is `Backward`, whereas t_2 ’s is `Forward`. But \mathbf{G}^2 is, by assumption, a globally self-avoiding motif itself, so by **ToroidalCycles**_{3.17} we can wrap \mathbf{G}^2 around the $n^2 \times n^2$ toroidal grid-graph. This guarantees that every supertile has all eight possible edge- and corner-adjacent neighbors.

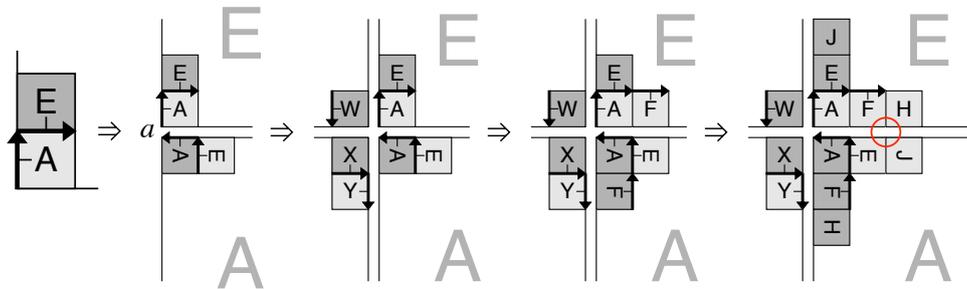


Figure 28: An initial right turn in the lower left corner of \mathbf{G} forces a nearby nodule in \mathbf{G}^2 .

Because threaded tiles A_A and A_E have tile segments that connect (at point **a**), tile W ’s tile segment must avoid A_E ’s, and so by **InsideIsNotOutside**_{3.6} tile W must be `LeftForward`. No segment can separate W from X (or it would collide at **a**), so **SelfAvoidance**_{3.4} forces X to be `BottomForward`. This in turn forces Y to be `RightBackward` along its supertile’s perimeter. **PerimeterBottomToTops**_{3.11} forces Y ’s neighbor tile F_A to be `RightForward`, which means that F_E must be `TopBackward`. Tile H_E cannot be a `Forward` tile, because for it to be `Forward`, **InsideIsNotOutside**_{3.6} requires there be a tile segment between it and F_E . The only possible choice for H_E would then be `LeftForward`, violating **PerimeterFlow**_{3.9}. For the exact same reason, tile J_A must also be a `Backward` tile. But this leaves the four tiles, F_E , H_E , J_A , and E_A forming a 2×2 nodule of `Backward` tiles that violates **NoNodules**_{3.18}, thereby guaranteeing self-contact somewhere else in \mathbf{G}^2 .

Next consider \mathbf{G} ’s lower right corner, and assume that \mathbf{G} ’s last two tiles, t_{n^2-1} and t_{n^2} form a right turn: $\mathbf{G} = [\dots B r]$. But this configuration is a symmetric right-left flip of the above: it mirrors the lower left corner (albeit with parities reversed), and hence a cross-perimeter nodule will again be created.

In the second case, we consider left turns. Again, for each pair of tiles, $[t_1, t_2]$ or $[t_{n^2-1}, t_{n^2}]$, there is only one way to make a left turn. Consider the lower left corner first, and assume $\mathbf{G} = [B l \dots]$. **BottomLeftContact**_{3.21} requires that there be a `bottom:left` contact combination in two edge-adjacent tiles somewhere along the top

row of tiles in T_n (see Fig. 27). Call these two tiles C and D, as shown in Fig. 29. As supertiles in \mathbf{G}^2 , the two left turns composed into them connect at the supertiles' shared corner, forming a configuration essentially the same (mirrored across the upper-left to bottom-right diagonal) as we found above for right turns: the four tiles, F_C , H_C , J_D , and E_D form a nodule.

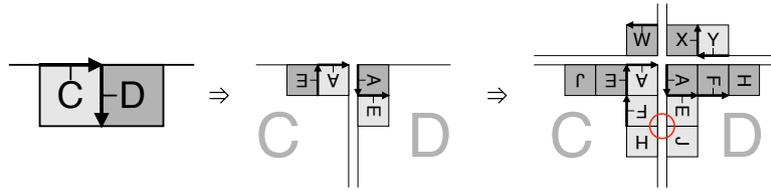


Figure 29: An initial left-right turn in \mathbf{G} creates a nodule in \mathbf{G}^2 as a result of the **bottom:left** zipping that must occur along the top of \mathbf{G} between tiles C and D.

Finally, as before, by reflective symmetry, the lower right corner cannot have a final left turn among the last two tiles. The same configuration will appear when C and D are chosen to be the two edge-adjacent tiles that form the required **bottom:right** contact combination along the top row of T_n . ■

Remark 3.6. *Composability makes the area near the toroidal origin $\mathbf{a} = \mathbf{b}$ so constrained that global self-avoidance under composition can't occur until the tile path has traveled at least two sub-tile widths along the perimeter of T_n , as the simplest **E** motif demonstrates (Fig. 12). In the terminology of [19], elements of this particular class of composable Hamiltonian cycles on the 2-D toroidal grid-graph can never be “bent” (having a turn at every vertex; this is also due to **RightLeftContact**_{3.12}). Perhaps even more counterintuitively, it turns out as a consequence of the Fibbinary Zipper Theorem (now under construction) that all solution motifs for order-7 must begin with three tile segments along the boundary. We will also later see that **NoLowerCornerTurns**_{3.22} is crucial to guarantee that any two globally self-avoiding motifs, one order- n and the other order- m , when composed result in an order- nm globally self-avoiding motif, see Theorem 6.1 below.*

Lemma 3.23. FrozenCorners. *In every left-handed $\mathbf{G} \neq \mathbf{I}$, there are at least 14 frozen tiles, 3 each in the upper left and right corners, and 4 each in the lower left and right corners, with 6 additional tiles of fixed parity.*

Proof. We know that there are no \mathbf{G} s for $n = 2, 3$, so assume $n \geq 4$.

By **Handedness**_{3.16} and **NoLowerCornerTurns**_{3.22}, a distinct \mathbf{G} begins with a pair of tiles, A and B, and ends with a pair of tiles, X and Z, each having a **right:left** contact combination (Fig. 30(a)). These two pairs toroidally surround the point $\mathbf{a} = \mathbf{b}$ at a supertile pair's connection point in \mathbf{G}^2 , where Z's tile segment meets A's (Fig. 30(b)). **PerimeterBottomToTops**_{3.11} and **SelfAvoidance**_{3.4} then force tiles E and F adjacent to X and Z, as well as tile Y above Z, into just one orientation and parity (Fig. 30(c)). **ArrowHeads**_{3.2} then forces tiles W and G (Fig. 30(d)), and **SelfAvoidance**_{3.4} and **ArrowHeads**_{3.2} in turn force tiles H, J, and K into their only possible orientations and parities (Fig. 30(e)). **PerimeterBottomToTops**_{3.11} forces tile C, and tile D is forced by **ArrowHeads**_{3.2} (Fig. 30(f)).

Finally, by **BottomToTopParity**_{3.10}, two tiles, each labeled Q in Fig. 30(f), must have parities opposite to their **bottom:top** forced neighbors E and J. Four more tiles (each labeled P) are then required to have the same parity as a forced edge-adjacent neighbor, because were they to be the opposite parity of that neighbor, by **InsideIsNotOutside**_{3.6} there would have to be a tile segment between the two. But in every case, such a segment would violate **SelfAvoidance**_{3.4} across a supertile edge.

We can then unmap the dissection from the torus back to the plane, recalling that the parities of all the tiles in the top two corners must be reversed, because supertiles in any **bottom:top** contact combination must by **PerimeterBottomToTops**_{3.11} have had opposite parities. The frozen and fixed corner areas are shown in Fig. 30(g). ■

Definition 3.3. *Along the shared edge of two edge-adjacent supertiles, a “gap” is a sequence of $w \geq 0$ pairs of tiles, one within each supertile, all having the same parity.*

Lemma 3.24. PerimeterGaps. *When \mathbf{G} 's tile path leaves from or arrives at the perimeter of T_n , those parts of the path in the cross-perimeter neighbor \mathbf{G} (regardless of orientation/parity) that arrive at or leave from the same perimeter create a single pair of cross-perimeter, edge-adjacent tiles having the same parity.*

Proof. Gaps are created when a motif turns into or away from T_n 's boundary. As can be inferred from Fig. 30(g) for **FrozenCorners**_{3.23}, at least one such turn must occur in a motif on each of the four sides of T_n , for $n \geq 5$. Consider the turn from tile E to tile F (Fig. 31, center left), or from tile X to tile Y (Fig. 31, center right). **SelfAvoidance**_{3.4}, **InsideIsNotOutside**_{3.6}, **PerimeterFlow**_{3.9}, **PerimeterForces**_{3.8}, and **NoNodules**_{3.18}

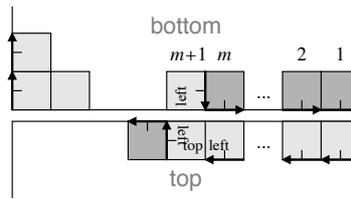


Figure 32: The rightmost gap in a **bottom:top** contact combination forms a **left:left** contact combination.

Lemma 3.27. NoEvenOrders. *There are no globally self-avoiding motifs when n is even.*

Proof. Assume n is even. **LeftLeftContact**_{3.25} requires that there be in \mathbf{G}^2 at least one pair of **left:left** neighbor tiles having either of the configurations shown on the left of Fig. 33. Without loss of generality, consider the top pair in the figure. In each of the two tiles' places, a **Forward** copy of \mathbf{G} will be composed in \mathbf{G}^2 , one upside down with respect to the other. Consider the central two tiles, E and F, along the left side of \mathbf{G} in T_n . Any distinguished edge for either of these tiles that passes along the perimeter must pass through the central corner point. But in an even order n dissection, by **LeftLeftContact**_{3.25}, this point will coincide with itself, and thus there is cross-perimeter self-contact in \mathbf{G}^2 . Hence, tile F can only be **BottomForward** or **RightForward**, and tile E can only be **RightForward** or **TopForward**. By **InsideIsNotOutside**_{3.6} and the parity rules of motif composition, however, there can be no parity difference in \mathbf{G}^2 between the two supertile neighbors from which the **left:left** combination is derived. Thus, the central four tiles violate **NoNodules**_{3.18}. A similar argument applies to **right:right** combinations had we chosen in **Handedness**_{3.16} to focus on right-handed motifs. ■

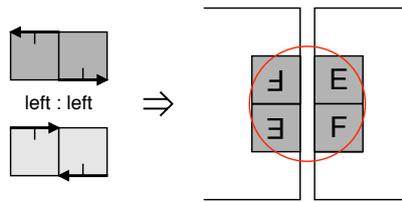


Figure 33: When n is even, some **left:left** pair results in a nodule in \mathbf{G}^2 at the center of \mathbf{G} 's left side.

Lemma 3.28. NoPerimeterTopTops. *No two tile segments, one on each side of the shared edge of two edge-adjacent, supertile neighbors in \mathbf{G}^2 , can both travel parallel to the perimeter of T while both are one tile width away from the perimeter.*

Proof. Without loss of generality, consider two tiles, D and E, that are edge-adjacent neighbors in \mathbf{G}^2 across a supertile edge (Fig. 34, left). The constraints **ArrowHeads**_{3.2}, **FourCorners**_{3.7}, **InsideIsNotOutside**_{3.6}, and **PerimeterFlow**_{3.9} force tiles H and J to be the same parity as D and E, which violates **NoNodules**_{3.18}. ■

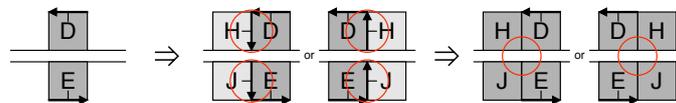


Figure 34: Two edge-adjacent neighbors in \mathbf{G}^2 whose edges are parallel to T 's perimeter force a nodule.

Lemma 3.29. CentralPerimeterGap. *The central tile on the left side of T_n (for odd n) forms a gap.*

Proof. By **LeftLeftContact**_{3.25}, the left side of \mathbf{G} must be self-avoiding when placed next to itself rotated by 180° . On an odd-order dissection, the central tile F is the only tile that is a cross-perimeter self-neighbor under 180° rotation, and thus any tile segment along the perimeter will coincide with itself, creating a **bottom:bottom** combination that violates **SelfAvoidance**_{3.4} (Fig. 35, left). Similarly, F can't have a vertical segment but not along the perimeter (e.g. **RightForward**) because **NoPerimeterTopTops**_{3.28} is then violated (Fig. 35, center). So the central tile must be half of a pair forming a gap, in one of two possible configurations. ■

Lemma 3.30. UpperLeftColumnForces. *All $(n - 1)/2$ tiles above the central gap tile of the left-most column of tiles in T_n , for odd n , are forced into exactly one configuration (in reverse order) by the arrangement of orientations and parities of the $(n - 1)/2$ tiles below the central gap tile in that column.*

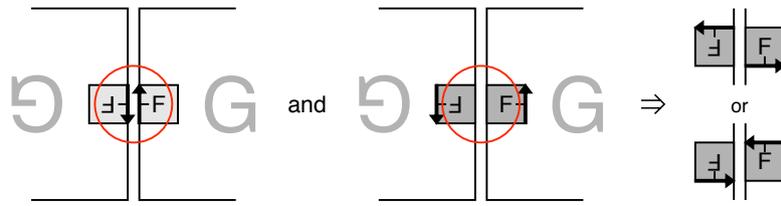


Figure 35: Central left side perimeter tile on odd order- n must be part of a perimeter gap.

Proof. Consider the column of n tiles up the left side of T_n (Fig. 36, left). By **PerimeterForces**_{3,8} and **PerimeterFlow**_{3,9}, there are only four possible tile orientations and parities, three Forward and one Backward, that any tile in the column can take on. These are labeled as tiles A, C, D, and J, having LeftBackward, BottomForward, RightForward, and TopForward tile segments, respectively. By **LeftLeftContact**_{3,25}, in G^2 these tiles will occur—in the same parity but oriented upside-down (Fig. 36, middle)—as cross-perimeter neighbors of other tiles (labeled W, X, Y, and Z) in the same column. By **PerimeterBottomToTops**_{3,11}, tiles Z and X are forced into only one possible orientation and parity by their respective cross-perimeter neighbors A and D. And tiles Y and W are also each forced, by **PerimeterGaps**_{3,24}, into just one orientation and parity by their respective cross-perimeter neighbors C and J (Fig. 36, upper right). For example, tile Y cannot be RightForward because were it RightForward its lower left corner would then be the center of a cross-perimeter nodule of Forward tiles. Hence, a tile along T_n 's shared left perimeter m tile widths below the central gap tile will force its counterpart tile along T_n 's left perimeter m tile widths above the central gap tile. ■

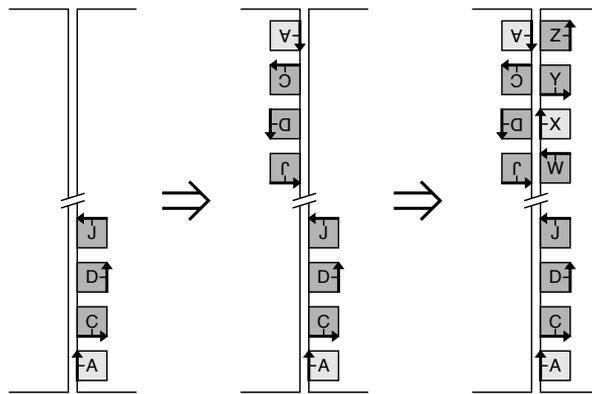


Figure 36: A required left:left contact combination means that all tile orientations and parities above the central perimeter gap are forced by the arrangement of tiles below the central gap tile.

Definition 3.4. A Fibonacci binary, or “Fibbinary,”⁹ bit sequence consists of $n \geq 0$ bits, each either 0 or 1, in which there occurs no two adjacent 1s.

Definition 3.5. An assignment of Fibbinary bits up T_n 's left sub-tile column is called the motif's “zipper bits.”

Lemma 3.31. LeftColumnZipperBits. For odd $n \geq 1$, the number of ways to arrange the n tiles in the left column of a left-handed G is $F_{(n-3)/2}$, where F_i is the i th Fibonacci number. Each zipper bits arrangement is determined by a palindromic, n -bit Fibbinary bit sequence whose central bit is a 1 and, when $n \geq 5$, whose first two and last two bits are 0s.

Proof. For each of the n edge-adjacent tiles, T_1, \dots, T_n , from the bottom to the top of the left column, assign a 1 to each tile with a horizontal tile segment, and a 0 to each remaining tile with a vertical tile segment. So, at least initially, for each 0 there are two possible ways to be vertical and for each 1 there are two ways to be horizontal. Consider the base cases for $n = 1, 3, 5, 7$, and then deal with the remaining general case for $n \geq 7$.

The exponentiated identity motif $I^k = I$ is a globally self-avoiding motif: although it doesn't converge to a space-filling curve (I is idempotent), it still conforms to **SelfAvoidance**_{3,4} for all k . So for $n=1$ there is only one distinct way to arrange the tiles in T_1 's “left column.” Its palindromic zipper bits are a single 1.

For $n=3$, **CentralPerimeterGap**_{3,29} requires the second (central) tile T_2 in the left column to create a gap, either BottomForward or TopForward. But the bottom two tiles must both be LeftForward along the perimeter,

⁹Marc LeBrun coined this portmanteau word in 2002, according to Sloane's A003714, see [21]; see also [6].

by **NoLowerCornerTurns**_{3.22}. Although the zipper bits 010 are Fibbinary, palindromic, and have a central 1 bit, the two constraints cannot simultaneously be satisfied for T_3 . So for $n = 3$ the number of ways to arrange the tiles in the left column is 0.

When $n \geq 5$, the first two required **LeftBackward** tiles at the bottom are 0s. The central gap tile (either **BottomForward** or **TopForward**) is a 1. Let (T_j, T_{j+1}) be two edge-adjacent—though not necessarily consecutively threaded—tiles in the column, where T_{j+1} is directly above T_j . As the left side of Fig. 36 shows, tile T_j can only be one of **LeftBackward**, **BottomForward**, **RightForward**, or **TopForward**. Consider each of these four cases:

- If T_j is **LeftBackward** (a 0), then by **ArrowHeads**_{3.2} tile T_{j+1} above it can only be either **LeftBackward** or **BottomForward** (a repeated 0 or a change to a 1).
- If T_j is **BottomForward** (a 1), then by **PerimeterGaps**_{3.24} T_{j+1} can only be **RightForward** (a change to 0).
- If T_j is **RightForward** (a 0), tile T_{j+1} can only be either **RightForward** or **TopForward** (a repeated 0 or a change to a 1) because **PerimeterFlow**_{3.9} or **InsideIsNotOutside**_{3.6} precludes any **Backward** tiles for T_{j+1} , because **ArrowHeads**_{3.2} precludes **BottomForward**, and **PerimeterForces**_{3.8} precludes **LeftForward**.
- And if T_j is **TopForward** (a 1), then by **PerimeterFlow**_{3.9} T_{j+1} can only be **LeftBackward** (a change to 0).

In other words, (a) any sequence of tiles up the column all with vertical (0) tile segments must all have the same orientation and parity, and therefore be consecutive (pairwise threaded); (b) no two edge-adjacent tiles in the column can both have horizontal (1) tile segments; (c) any single horizontal 1 tile is forced by the 0 tile below it into only one possible orientation and parity; and (d) the 0 tile directly above a 1 is forced into only one possible orientation and parity, depending on that of the 1 tile below it. The 1s (i.e., gaps) are thus self-avoiding, in that each 1 is always surrounded by a 0 on each side. In particular, this means that tile $T_{(n-1)/2}$ (just below the central gap tile) must be a 0 (either **LeftBackward** or **RightForward**). **NoLowerCornerTurns**_{3.22} requires both tiles T_1 and T_2 to be 0s, and tiles $T_{(n-1)/2}$ and $T_{(n+1)/2}$ must be 0 and 1, respectively, as shown in Fig. 37.

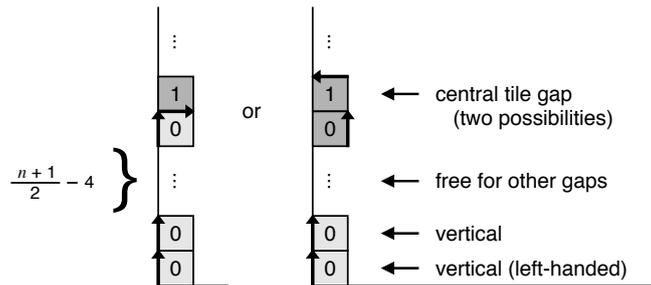


Figure 37: Tiles between the 00 pair at the bottom and 01 pair ending at the central gap are where all remaining combinatorial freedom in the left column can be found. The left or right possibility is dependent on n .

For $n = 5$, T_2 and $T_{(n-1)/2}$ collapse into the same **LeftBackward** (0) tile, and therefore these constraints can be satisfied in exactly one way: a left-handed **G** for T_5 begins as $\mathbf{G} = [1 \mid \mathbf{B} \dots]$, with zipper bits 00100.

Now consider the general case up the left column of left-handed T_n for odd $n \geq 7$. Between the two anchoring pairs of tiles, one of which pairs must be 00 at the bottom and the other 01 at mid-column, there are $m = (n + 1)/2 - 4 = (n - 7)/2$ tiles available to which a 0 or 1 might be assigned. Thus when $m = 0$, corresponding to $n = 7$, there are zero tiles between the two anchor pairs. **InsideIsNotOutside**_{3.6} precludes the second possibility shown in Fig. 37 on the right, so all left-handed, globally self-avoiding motifs for T_7 must begin in just one way: $\mathbf{G} = [1 \mid \mathbf{B} \dots]$, with palindromic zipper bits 0001000.

The number of ways to assign values to $m > 0$ bits is 2^m . But this includes values whose bit sequences contain adjacent 1s, corresponding to gaps wider than one tile width that would violate **PerimeterGaps**_{3.24} upon which the above cyclic constraints rely. Thus, we need to count how many m -bit integers from 0 to $2^m - 1$ have binary representations in which there occurs no pair of adjacent 1s, i.e., the Fibbinary integers (sequence A003714 in [21]). The number of $(m + 2)$ -bit Fibbinary integers is just the sum of the number for $(m + 1)$ -bit Fibbinary integers, where each integer’s bit sequence is prepended with a 0, plus the number of m -bit Fibbinary integers, where each integer’s bit sequence is prepended with a 10. Hence, for a given bit length m , the number of possible Fibbinary integers is F_{m+2} , where $F_0 = 0$, $F_1 = 1$, and $F_{m+2} = F_{m+1} + F_m$.

Finally, while initially there can be in general two states (two horizontal tile segments or two vertical tile segments) for each bit value, this is not true for the bottom two tiles in the left column, which are **LeftBackward**.¹⁰

¹⁰Recall that we are focusing on left-handed motifs, whereas for (indistinct) right-handed motifs, we would now be making the same analysis, mirrored, of T_n ’s *right* column, which would end in two **RightBackward** tiles.

Therefore, for any Fibbinary bit sequence of 0s and self-avoiding 1s, there is only one possible orientation and parity for tile T_{j+1} , given the state of T_j underneath it, determined by the constraints running up the left column from the bottom. If the number of 1s up to and including the central tile $T_{(n+1)/2}$ is odd, then the central tile on the left side of T_n is **BottomForward**; if the number of 1s is even, the central tile is **TopForward**.

Returning to our base cases, the number of arrangements for $n = 1, 3, 5, 7$ are 1, 0, 1, 1, respectively, so the formula nicely simplifies to $F_{(n-3)/2}$ for odd $n \geq 1$. This works even for the identity motif **I**, where $n = 1$ produces $F_{-1} = 1$. Figure 38 illustrates the number of ways to arrange the left-side column of tiles for T_1 through T_{17} . It also highlights those Fibbinary mode bits that embody the primary combinatorial freedom between the two anchoring pairs of tiles from the bottom 00 to the center 01 of the left column.

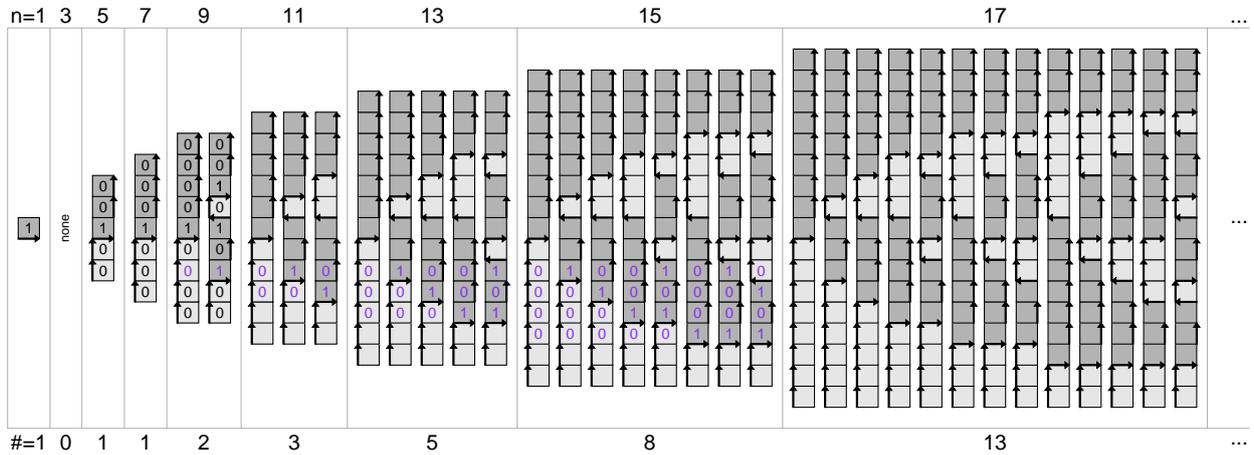


Figure 38: A left-handed **G**'s palindromic zipper bits arrange the left column of tiles of a globally self-avoiding motif for T_n , here for odd $n \leq 17$. Those mode bits that have any freedom to vary (starting at order-9) are highlighted in color. The 13 columns on the right omit mode bits to make it easier to see the Fibbinary patterns.

As the constraint cycle rises up past the central gap tile, there must be palindromic counts for all runs of 0 tiles, by **UpperLeftColumnForces**_{3,30}. So the zipper bits in the column form an n -bit palindromic Fibbinary bit sequence, whose central bit must be a 1, and, when $n \geq 5$, begins and (necessarily) ends with two 0s. ■

Lemma 3.32. ZipperModes. *The parities and orientations of all $4(n-1)$ outermost perimeter sub-tiles of T_n are frozen by **G**'s left column zipper bits.*

Proof. Not only must the left side of any left-handed **G** self-avoidingly zip with itself under 180° rotation, but by **RightLeftContact**_{3,12}, **BottomLeftContact**_{3,21}, and **LeftTopContact**_{3,26}, the left side must also zip at least once with each of the other three sides in some pair of neighbor tiles in **G** or \mathbf{G}^2 . As Fig. 39 illustrates for $n = 15$, the number and placement of all perimeter gaps must be such that **SelfAvoidance**_{3,4} is maintained. Hence our combined constraints literally zip each side of **G** together with its other three sides,¹¹ in a way that guarantees that Hamiltonian **NoNodules**_{3,18} is preserved everywhere along the wraparound boundary (for odd $n \geq 5$). So the pattern of tile segments and gaps along the outermost boundary layer of T_n 's $4(n-1)$ perimeter tiles is frozen. The same situation, in mirrored form, occurs in right-handed motifs. ■

Q.E.D. (Fibbinary Zipper Theorem)

Definition 3.6. *Because the zipper bits are palindromic (or perhaps, “ambidextrous”), they operate identically for both right- and left-handed motifs. For classification labeling purposes, we will ignore the upper $(n-1)/2$ bits, and use just the first $(n+1)/2$ of them, which we call its “mode bits,” as a subscript of the order n , e.g., 5_{001} or $13_{0010101}$. The central 1 bit thus becomes the low-order bit in the mode specification. Mode bits then correspond to the odd Fibbinary integers (sequence A022341 in [21]). Fig. 39 shows the result using the seventh Fibbinary mode of T_{15} , $15_{00100101}$, corresponding to zipper bits 00 1001 010 1001 00.*

3.3 Tidying Up

The foregoing proof of the Fibbinary Zipper Theorem depends upon showing, for left-handed motifs, that each of five of the nine possible contact combinations must occur either in **G** (**bottom:top**, **right:left**, **left:top**, **bottom:left**) or in \mathbf{G}^2 (**left:left**),¹² with appropriately mirrored (swapping right and left) combinations in

¹¹This is sometimes called *conic* wraparound.

¹²Based on examination of machine-enumerated motifs for small n , it is almost certainly true that every **G** (as opposed to \mathbf{G}^2) must have both a **left:left** and a **right:right** contact combination also. But that remains a conjecture.

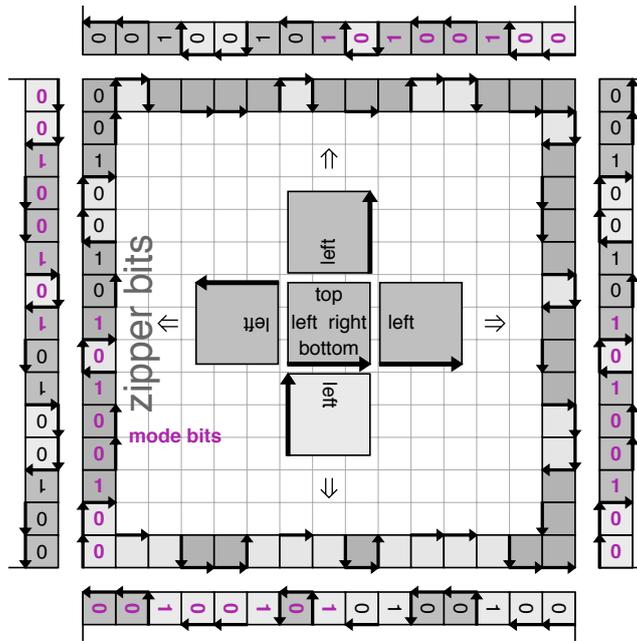


Figure 39: The mode bits (in purple/bold), when rotated 180° by a `left:left` contact combination, force all remaining zipper bits up the left side of T . The palindromic zipper bits then appear in neighbor tiles along the other three sides of G at least once in G or G^2 , leaving all perimeter tiles in G frozen into one configuration by the original mode bits. Here, in T_{15} , all $4(n - 1) = 56$ of G 's perimeter tiles are determined by T_{15} 's seventh (out of eight) Fibbinary mode, $15_{00100101}$.

right-handed motifs. The `bottom:right` contact combination was proved but not needed. For completeness, two more are also mandatory, but the one remaining contact combination, `top:top`, is not.

Lemma 3.33. RightTopRightRightContact. $G \neq I$ has at least one `right:top` contact combination and G^2 has at least one `right:right` contact combination, in some pair of neighboring tiles.

Proof. **FrozenCorners**_{3.23} requires that the lower left and upper left corner of every left-handed G must have certain forced tiles (Fig. 40, left). The tiles labeled A and C create a `right:top` contact combination (Fig. 40, middle), and **BottomTopContact**_{3.14} implies that `right:right` contact must occur also. ■

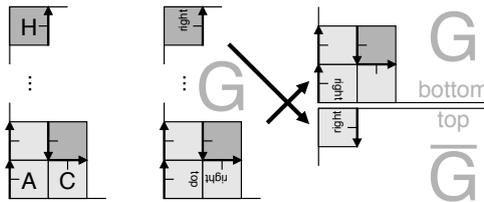


Figure 40: Both `right:top` and `right:right` contact combinations always occur in any G or G^2 .

Lemma 3.34. TopTopContactOptional. A motif G is not required to have a `top:top` contact combination in order to be globally self-avoiding.

Proof. By **NoPerimeterTopTops**_{3.28}, a `top:top` contact combination cannot occur along the shared edge of two edge-adjacent supertiles for globally self-avoiding motifs G . So we need only examine pairs of edge-adjacent tiles within G . To show that the `top:top` combination is optional, it suffices to find a counterexample G with no internal `top:top` neighbors. Orders 5 and 7 motifs offer no aid. But consider the two T_9 motifs shown in Fig. 41, one for each of the two possible zipping modes. By inspection, no two edge-adjacent sub-tiles form a `top:top` contact combination. ■

Remark 3.8. In motifs where all threaded sub-tiles are edge-adjacent, such as the Hilbert Curve and its pure edge-adjacent sub-tile generalizations, `top:top` combinations are not only required, they are ubiquitous (see [12], Fig. 2, or Ch. 7 of [13], describing “half-domino curves,” of which the Hilbert Curve is a degenerate case).

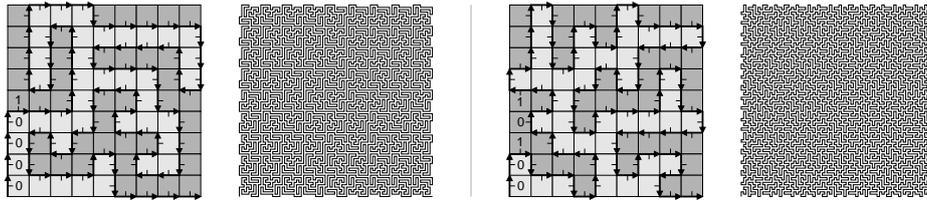


Figure 41: Left: A mode 9_{00001} motif \mathbf{G} and \mathbf{G}^2 containing no `top:top` neighbors. Right: A mode 9_{00101} motif \mathbf{G} and \mathbf{G}^2 containing no `top:top` neighbors.

4. The Highly Constrained First Fibbinary Zipping Mode

For any given odd n , the first Fibbinary zipping mode has zipper bits that are all 0s except for the central 1. For a given odd $n \geq 5$, we denote this family as $n_{0\dots01} = n_{0\bar{0}1}$, e.g., 5_{001} , 7_{0001} , 9_{00001} , 11_{000001} , etc.

When $n \geq 9$, there are at least two Fibbinary zipping modes governing possible motifs. For order 9, several brute force enumerations—made without any knowledge of the Fibbinary boundary conditions—found that only 660 of the 10,101 distinct motif solutions (all shown in Section 5 below) had boundaries governed by this first Fibbinary mode. In short, all else being equal, the 9_{00001} boundary creates significantly more constraint than does the 9_{00101} boundary. The same is true for a brute-force enumeration of all order 11 solutions. Only 58,936 of the 20,305,328 total are governed by the 11_{000001} boundary. The next several lemmas explain why.

As Fig. 38 illustrates, left-handed motifs subject to the first Fibbinary mode will have nothing but `LeftBackward` tiles below the central gap tile on the left side of T_n (for right-handed motifs, it would be `RightBackward` tiles along the right side). A mode $n_{0\dots01}$ motif, by virtue of having only one (central) perimeter gap per side, thus has a boundary that maximizes the number of tile segments in its tile path that travel along the edges of T_n . The resultant long lines of identical consecutive tiles “emanate” constraint into the interior of T_n , both directly and toroidally across wraparound boundaries.

Lemma 4.1. TripletForces. (a) *On the toroidal grid-graph, a thread containing a triplet of consecutive tiles $[t_{m-1}, t_m, t_{m+1}]$ that either (1) has colinear tile segments or (2) forms a 3×1 group of tiles, forces the edge-adjacent neighbor tile on the other side of the middle tile t_m ’s segment into the same orientation as t_m , but with opposite parity.* (b) *If the three tiles are all the same orientation and parity, the middle tile t_m ’s remaining edge-adjacent neighbor is also forced into the same orientation as t_m , but with opposite parity.*

Proof. The first case (a) follows from **SelfAvoidance**_{3.4} and **BottomToTopParity**_{3.10} (Fig. 42, left). To show (b), consider three connected `BottomForward` tiles (Fig. 42, right). If tile E is `Forward`, it cannot be `BottomForward`, for that violates **InsideIsNotOutside**_{3.6}. Nor can it be `LeftForward`, `TopForward`, or `RightForward`, because **NoNodules**_{3.18} requires tiles D and F to be `Backward`, one or the other of which then violates **OneToOneNess**_{3.1} or **InsideIsNotOutside**_{3.6}. Therefore, E must be `Backward`. But it can only be `BottomBackward` because any other orientation will also violate **InsideIsNotOutside**_{3.6}. A similar argument holds for all other orientations and parities that the original triplet might take on. ■

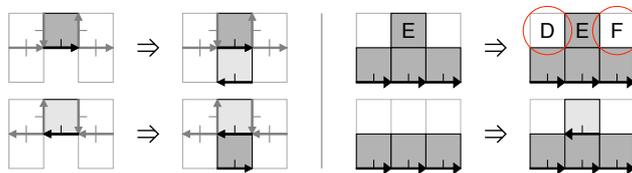


Figure 42: Left: Three consecutive tiles force the middle tile’s cross-tile segment neighbor into just one possible orientation and (opposite) parity, regardless of the various ways that tile segments can be arranged among the first and third tiles, and regardless of the middle tile’s parity. Right: If three consecutive tiles are all the same orientation and parity, the middle tile’s remaining edge-adjacent neighbor is forced into just one orientation and (opposite) parity.

Remark 4.1. *Like the next, this lemma is qualified as “on the toroidal grid-graph” to avoid considering boundary conditions. It is essentially true on the unwrapped T_n , except that we would have to include a caveat that it only holds for those triplets whose middle tiles’ neighbors exist within T_n . On \mathbb{T}_n , they always exist.*

Lemma 4.2. WeddingCakes. *On the toroidal grid-graph, a thread containing $R \geq 1$ consecutive tiles, all the same parity and orientation, forces a “wedding layer cake” pattern with $W_R = \lfloor (R+1)^2/4 \rfloor$ tiles in it, or, when counting layers on both sides, a diamond pattern with $D_R = \lfloor (R^2 + 1)/2 \rfloor$ tiles.*

Proof. The formulas hold for $R=1, 2$. For $R \geq 3$, consider the repeated application of **TripletForces**_{4,1}, one for each of $R-2$ (overlapping) triplets. The counts follow from simple summations. ■

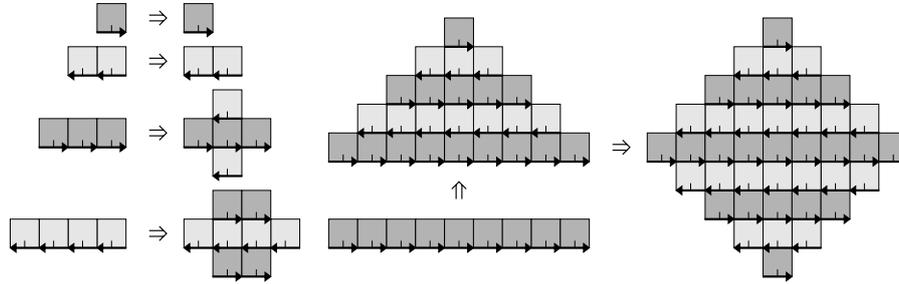


Figure 43: Consecutive tiles all having the same orientation and parity force successive shorter-by-2, adjacent layers of tiles, each layer alternating in parity. On the right, $R = 9$ tiles in a row create a frozen double-wedding cake of $D_9 = 41$ forced tiles.

Lemma 4.3. ModeOneHalfFrozen. *In every order- n \mathbf{G} whose boundary is governed by the first Fibbinary zipping mode $n_{0\dots 01}$, the number of initially frozen tiles is $(n^2 + 11)/2$.*

Proof. Distinct (left-handed) motifs governed by mode $n_{0\dots 01}$ each begin with a sequence of $\frac{n-1}{2}$ LeftBackward tiles. By **ZipperModes**_{3,32} and **WeddingCakes**_{4,2}, these, in turn, create eight wedding cakes growing inward from the perimeter, see Fig. 44. At each corner there are two wedding cakes, one for $k = \frac{n-1}{2}$ and the other for $k-1 = \frac{n-3}{2}$ tiles, that mesh together. It is a straightforward summation to determine that the number of forced tiles is $(n^2 + 11)/2$. Because n is only odd, this sum is always an integer. Thus, over half of the n^2 tiles in T_n are frozen when the boundary zipping mode is $n_{0\dots 01}$. ■

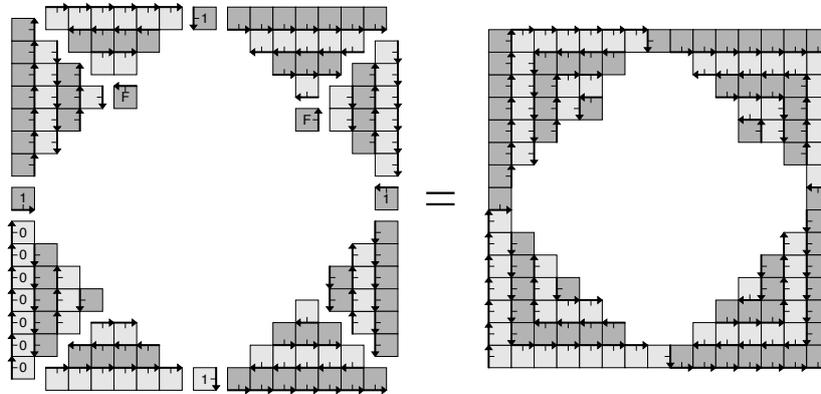


Figure 44: All initially forced tiles for zipping mode $n_{0\dots 01}$ for T_{15} can be counted as eight wedding cakes, along with four forced central gap tiles (labeled with 1s), as well as two additional forced-by-**NoNodules**_{3,18} tiles (labeled F), found in the upper half of T_n when $(n-1)/2$ is odd, and in the lower half when it is even.

Remark 4.2. *Fig. 44 shows this count for $n = 15$. After just one exponentiation of any mode $n_{0\dots 01}$ motif, each supertile lattice point must be surrounded—toroidally—by a rough-shaped diamond of frozen tiles (Fig. 45), where the composed Hamiltonian path passes through the super-lattice point in the center. It is as if constraint emanates from the lattice points, freezing those tiles within a “taxicab-geometry ball” of radius $r \approx n/2$.*

Remark 4.3. FrozenCorners_{3,23} *is essentially a special case of* **ModeOneHalfFrozen**_{4,3}. *The latter enables manual enumeration of all distinct, globally self-avoiding motifs for $n=5$ and $n=7$, confirming earlier machine results. A hand enumeration, especially for T_7 as below in* **SevenOrder7s**_{5,2}, *additionally gives a feel for how our various named constraints act locally to mold the asymmetric shapes of these handed motifs.*

Remark 4.4. *Any \mathbf{G}^2 cannot itself be a mode $n_{0\dots 01}$ motif.*

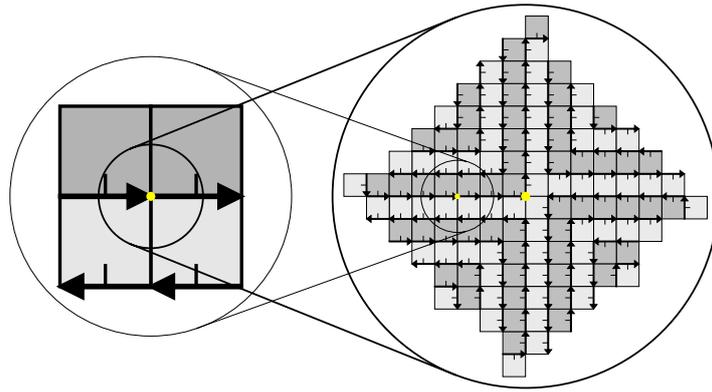


Figure 45: In a mode $n_0\dots_01$ motif \mathbf{G}^k (a portion shown here for $n = 15$), over half of all tiles surrounding every supertile lattice point (central dot) in \mathbf{G}^{k+1} are forced. Because this particular 2×2 sample (on left) toroidally surrounds a central point with each of the four “frozen wedding caked” corner areas of the motif, the number of forced sub-tiles in the right-hand circle for $n=15$ is $((15^2 + 11)/2 = 118$.

5. Globally Self-Avoiding Motifs for Orders 5 through 13

Lemma 5.1. ECurveUnique. *There is only one distinct globally self-avoiding motif for T_5 .*

Proof. There is only one zipping mode for T_5 , denoted 5_{001} . By **ModeOneHalfFrozen**_{4,3}, $(5^2 + 11)/2 = 18$ tiles are initially forced (Fig. 46, top left): the 16 perimeter tiles of T_5 as well as tiles H and K, with seven inner tiles remaining to be examined. Applying our various lemmas appropriately, all other tiles are forced.¹³ By **InsideIsNotOutside**_{3,6} (via **FrozenCorners**_{3,23}), we know that tiles A and C must both have **Backward** parities. **SelfAvoidance**_{3,4} requires tile J to be **LeftBackward** and further requires that there be a tile segment between tiles D and A whose arrow head connects with the path traveling up into the top left area of T_5 . Tile D must be **Forward** and **ArrowHeads**_{3,2} then requires A to be the tile to which that separating tile segment belongs (hence A must be **BottomBackward**). D must then be **RightForward**, as no other sides are available for its tile segment. **SelfAvoidance**_{3,4} then requires tile B to be **TopBackward**. Tile E must be **Backward** parity because of tile D being **Forward**. Tiles C, E, and F are then forced into their final orientations and parities by **LongestDistance**_{3,3} and **NoNodules**_{3,18}. Thus, all 25 tiles are completely forced by the initial boundary mode, and can be arranged in just one distinct (right- or left-handed) way. ■

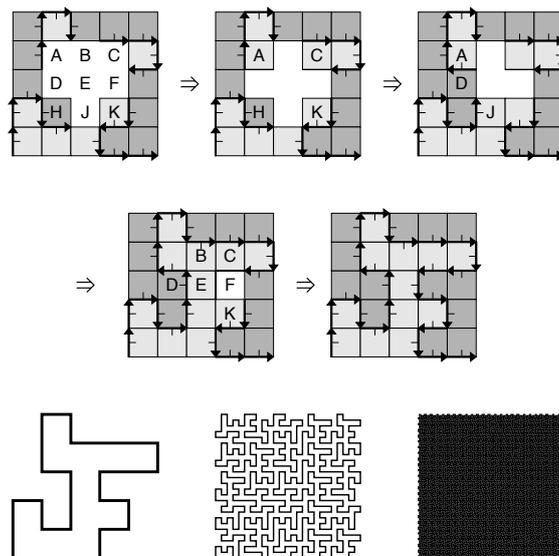


Figure 46: There is only one distinct (left-handed), globally self-avoiding motif for T_5 .

¹³McKenna [8] called this the “E” Curve because its left-handed motif resembles a block letter E.

Lemma 5.2. SevenOrder7s. *There are exactly 7 distinct, globally self-avoiding motifs for T_7 .*

Proof. There is just one mode for the order-7 dissection: 7_{0001} . By **ModeOneHalfFrozen**_{4,3}, there are $(7^2 + 11)/2 = 30$ initially forced tiles, with 19 inner tiles remaining to be examined, see Fig. 47(a). Using our various lemma constraints as we did in **ECurveUnique**_{5,1}, the figure shows the several logical pathways to enumerating the seven solutions (and implicitly their right-handed mirror images). ■

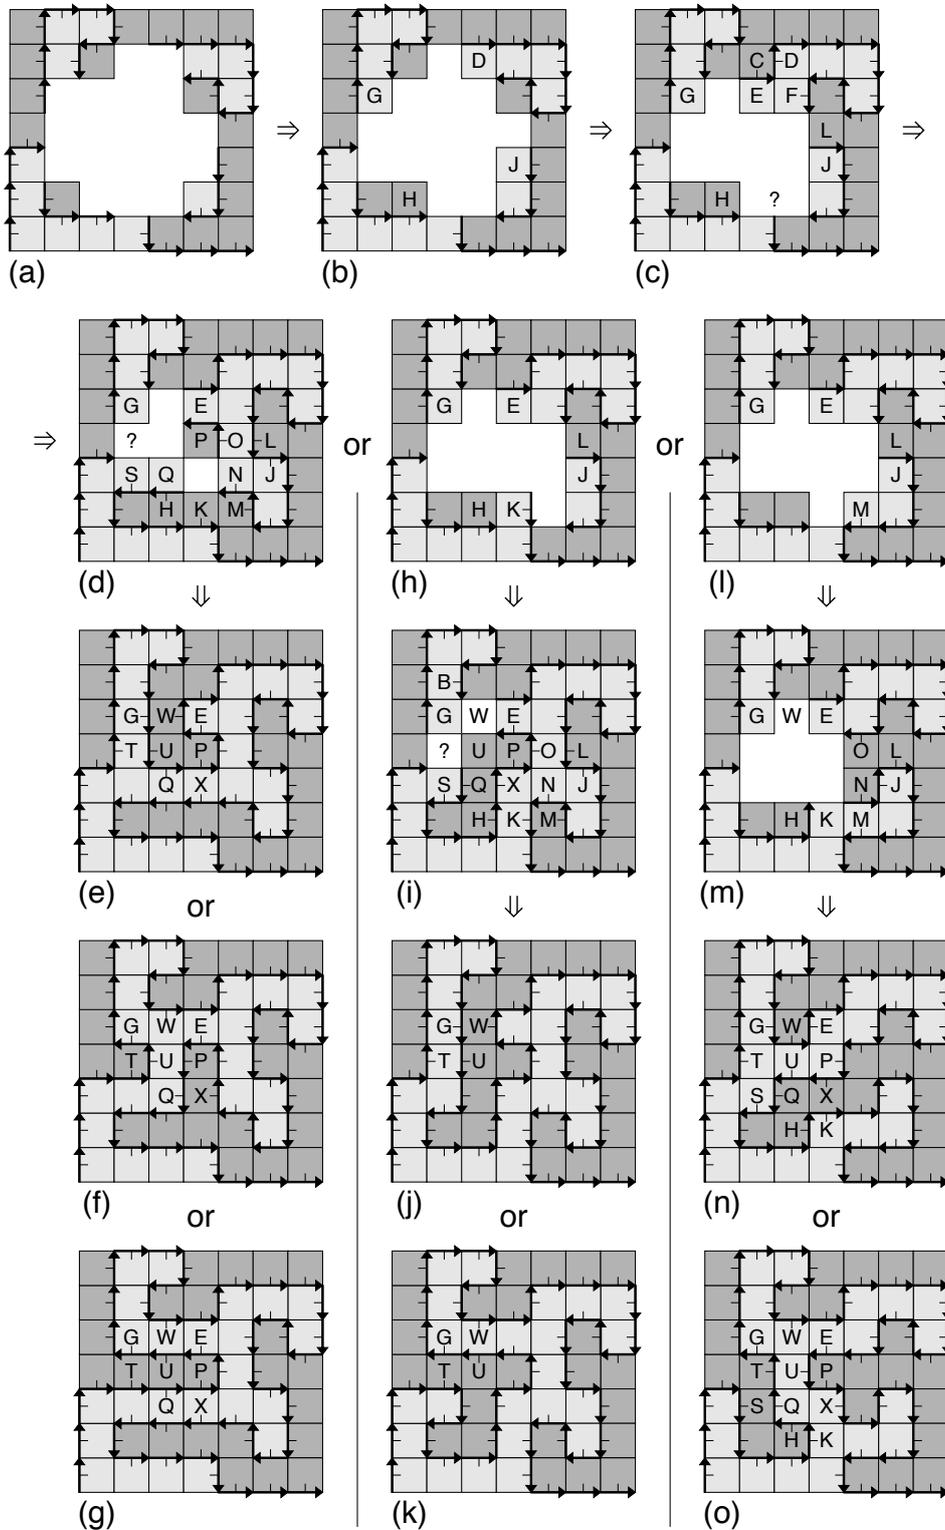


Figure 47: The seven distinct (left-handed), globally self-avoiding motifs for T_7 .

Remark 5.1. *The motif in the bottom center of Fig. 47 was manually found by McKenna in 1979 and later described as a “squarecurve” (square recursive curve) in [8] (see also mention of its right-handed mirror in [17]).*

On higher odd orders beginning with T_9 , hand enumeration is no longer feasible as the expected combinatorial explosion begins. Several computer enumerations of all distinct, globally self-avoiding motifs for T_9 each find exactly 10,101 motifs. These motifs fall into $F_{(9-3)/2} = F_3 = 2$ boundary modes, 9_{00001} and 9_{00101} . As Fig. 48 shows, each mode forces several more inner tiles into their final orientations and parities. Modes 9_{00001} and 9_{00101} yield 660 and 9441 distinct motifs, respectively. Fig. 49 shows all 660 mode 9_{00001} motifs in one particular lexicographic enumeration order. Fig. 50 shows a sampling of 280 of the 9441 motifs governed by mode 9_{00101} , also from top to bottom.

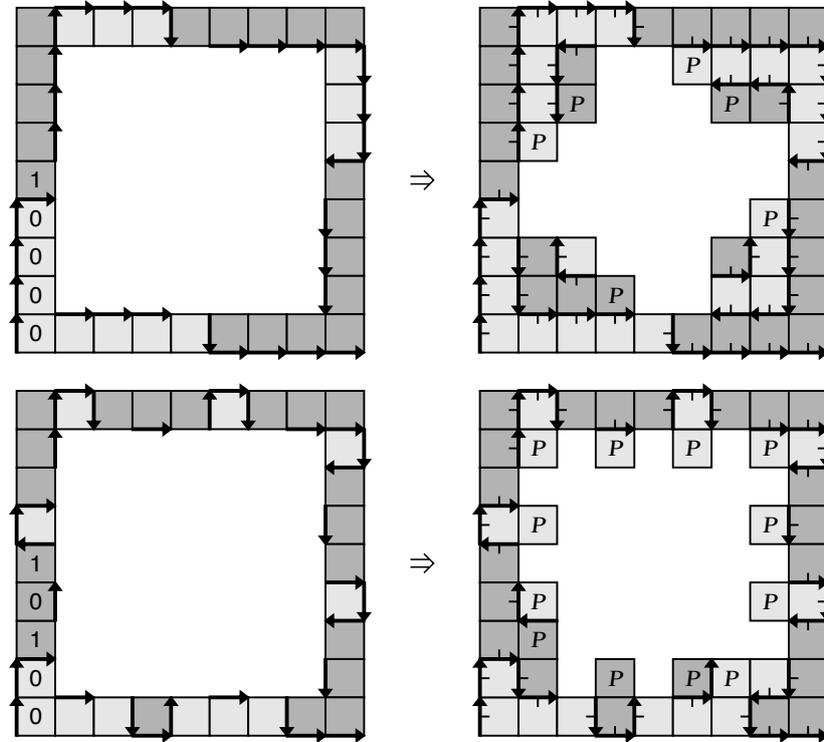


Figure 48: Left: Boundary tiles of left-handed motifs for each of T_9 's two Fibbinary modes, 9_{00001} and 9_{00101} . Right: Further forced or constrained tiles or path segments for each mode, where tiles labeled P have known parity, by **ModeOneHalfFrozen**_{4,3}, **InsideIsNotOutside**_{3,6}, and/or **NoPlanarCycles**_{3,5}.

For T_{11} , exhaustive machine enumeration¹⁴ found exactly 20,305,328 globally self-avoiding, distinct (left-handed) motifs. These fall into $F_{(11-3)/2} = F_4 = 3$ boundary modes, 11_{000001} , 11_{000101} and 11_{001001} , as shown in Fig. 51. For each of these boundary modes there are exactly 58,936; 19,854,452; and 391,940 motifs respectively. Fig. 52 shows an ordered sampling of motifs from each of the three boundary modes. And again, we see that the first boundary mode, with its much higher initial set of frozen tiles, is significantly more constrained than the others, leading to the lowest motif count.

For T_{13} there are $F_{(13-3)/2} = F_5 = 5$ boundary modes ($13_{0000001}$, $13_{0000101}$, $13_{0001001}$, $13_{0010001}$, $13_{0010101}$), whose initial boundary conditions are shown in the top row of Fig. 53. The next number in the series 0, 0, 1, 7, 10101, 20305328, ... is not currently known. But partial computer enumeration of each mode using standard backtracking methods was performed. The middle row of Fig. 53 shows a sample motif, plucked randomly out of the first few thousand found, for each of the five Fibbinary zipping modes in the top row. Different lexicographic search strategies were used for different modes, so as to be able to find an initial set of solutions quickly. The bottom row shows \mathbf{G}^2 for each motif, but filled by coloring Forward and Backward sub-tiles with white and black, respectively.

¹⁴This enumeration required 147 hours of CPU time on a circa-2004 450MHz personal computer to accomplish.

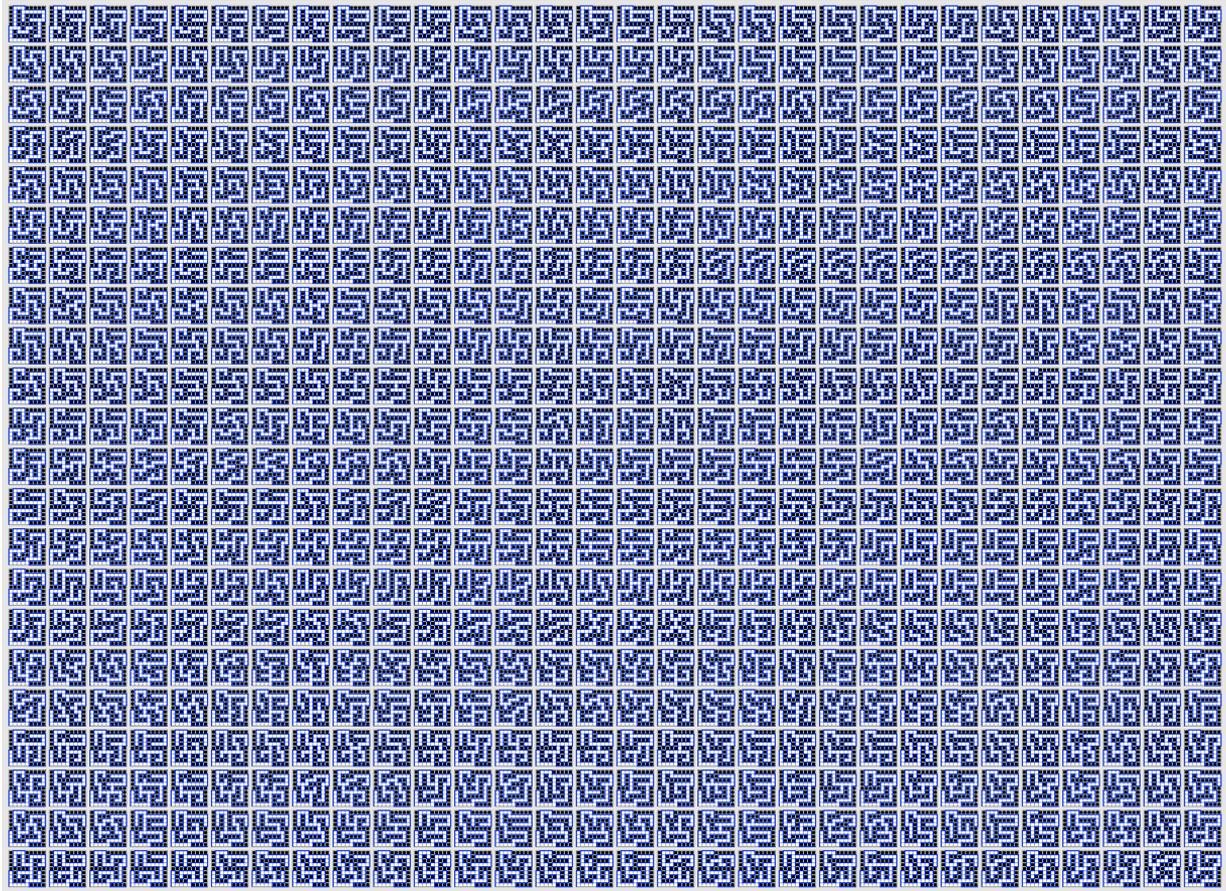


Figure 49: All 660 distinct globally self-avoiding motifs for mode 9_{00001} . Forward tiles are black, and Backward tiles are white. Each motif's tile path is drawn in blue. Bijection marks are suppressed.

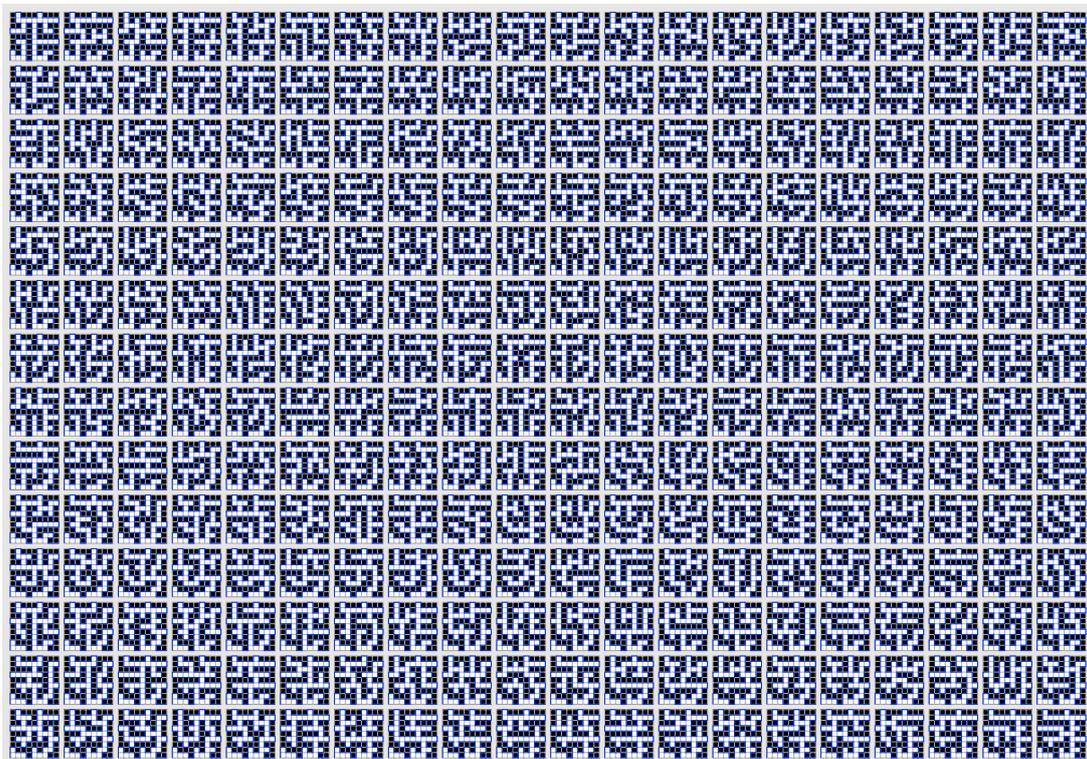


Figure 50: A sampling of the 9441 distinct globally self-avoiding motifs for mode 9_{00101} . Forward tiles are black, and Backward tiles are white. Each motif's tile path is drawn in blue. Bijection marks are suppressed.

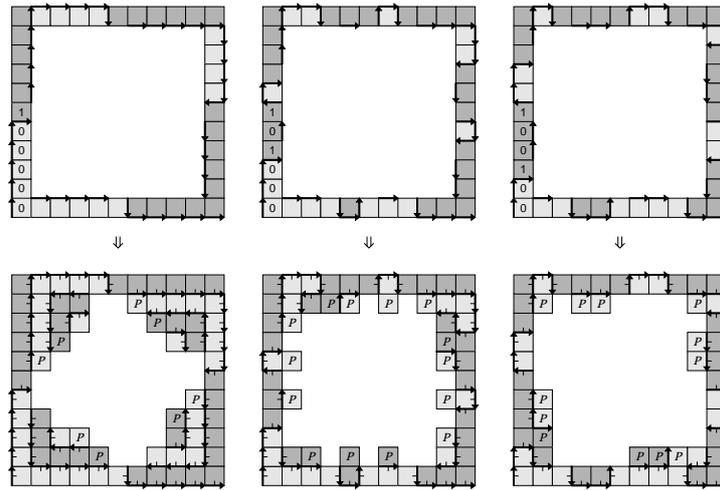


Figure 51: Top: Three boundary modes for T_{11} : 11000001 , 11000101 , and 11001001 . Bottom: Boundary tiles for each mode force a few more inner tiles and segments. Inner tiles with only known parity are labeled with a P .

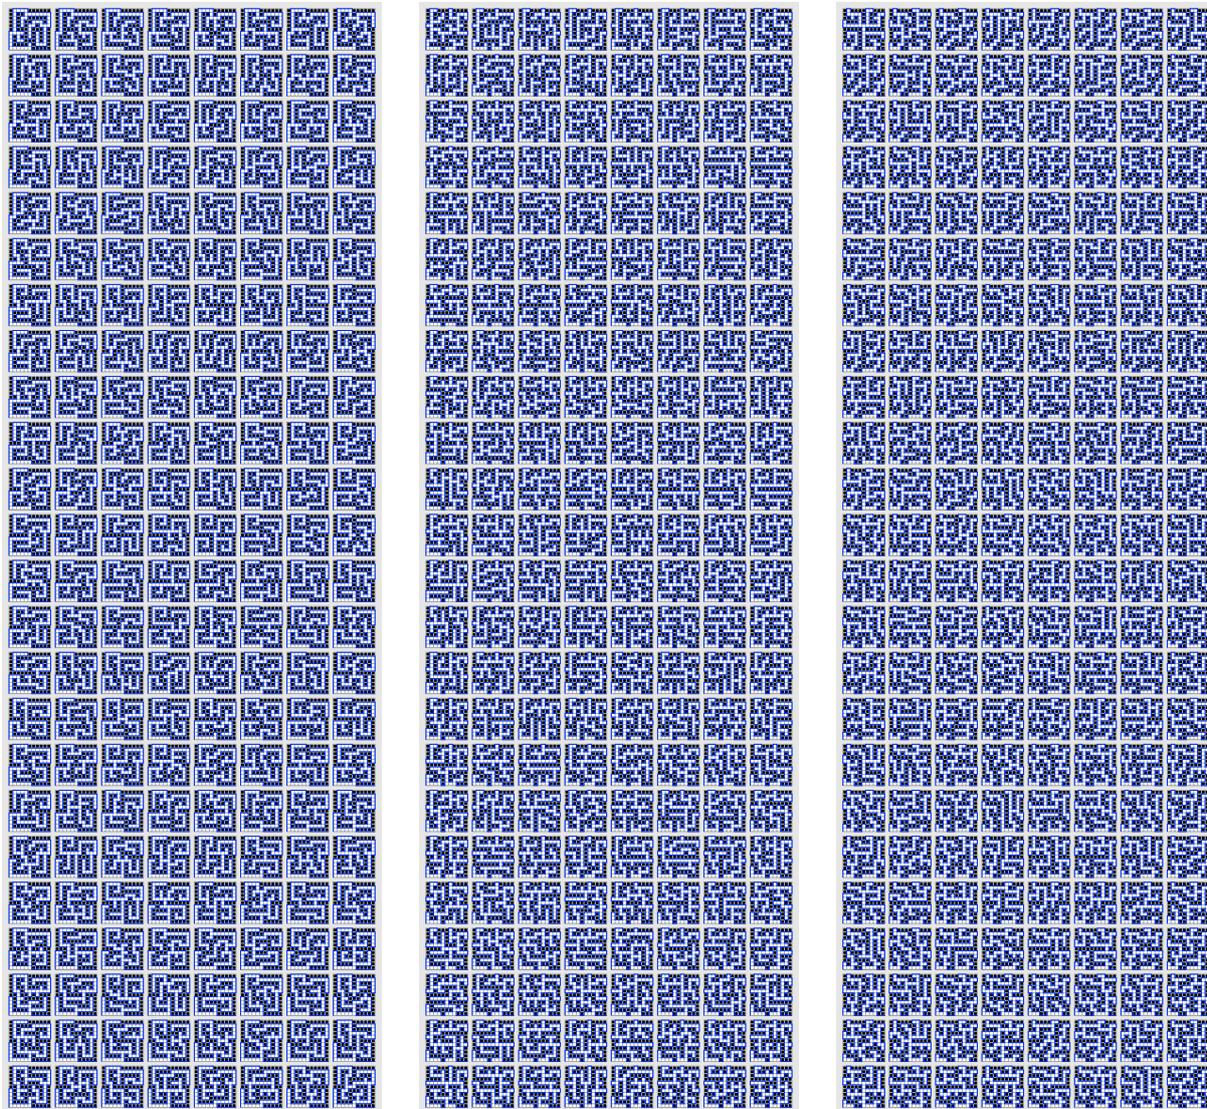


Figure 52: Left: A sampling of the 58,936 globally self-avoiding motifs for mode 11000001 . Middle: A sampling of the 19,854,452 globally self-avoiding motifs for mode 11000101 (between each adjacent motif are approximately 100,000 motifs in lexicographic order not shown). Right: A sampling of the 391,940 globally self-avoiding motifs for mode 11001001 . Bijection marks are suppressed.

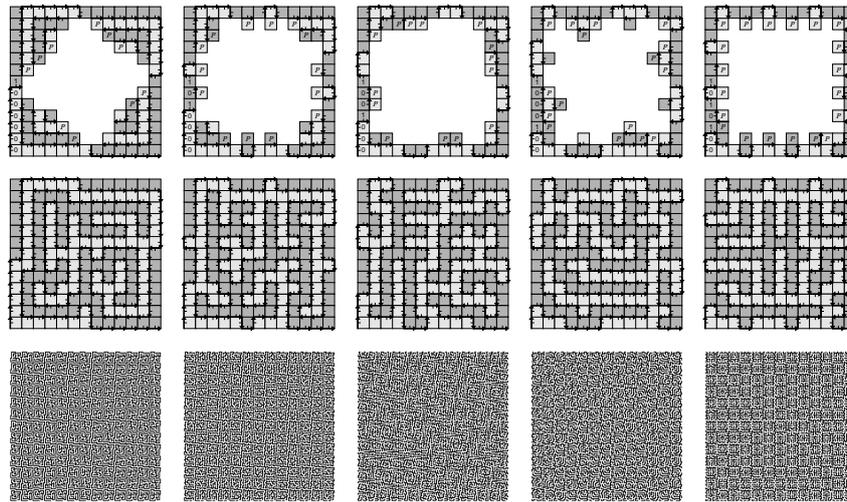


Figure 53: Top: The initially frozen tiles, and those tiles with fixed parity (marked with a P), for each of the five boundary modes governing globally self-avoiding motif boundaries for T_{13} . Middle: A sample motif \mathbf{G} for the mode above it. Bottom: \mathbf{G}^2 shown with Forward sub-tiles colored white, Backward sub-tiles black.

6. The Submonoid \mathcal{Z} of Globally Self-Avoiding Motifs

In the monoid \mathcal{H} of all Hilbert-style motifs, for every pair of motifs that each have a given property, if that property is preserved by the motif composition operator, then that property serves to distinguish a submonoid of motifs having that property, as long as the identity motif has it too. In particular:

Theorem 6.1. *Globally self-avoiding motifs form a submonoid $\mathcal{Z} \subset \mathcal{H}$ under the motif composition operator.*¹⁵

Proof. The identity motif \mathbf{I} is globally self-avoiding: no vertex of \mathbf{I}^k has more than one arrowhead arriving at, nor more than one arrowtail leaving from, that vertex.

Next, we must show that if \mathbf{F} on T_m and \mathbf{G} on T_n are each a globally self-avoiding motif, then $\mathbf{F} \circ \mathbf{G}$ is also a globally self-avoiding motif on T_{mn} . This is trivially true when either \mathbf{F} or \mathbf{G} is the identity, so we need only consider odd $m, n \geq 5$ (by the Fibbinary Zipper Theorem there are no globally self-avoiding motifs for even n, m or for $n = 3$ or $m = 3$).

Recall that composition is associative and that it replicates the shape of its right-hand operand in a variety of orientations according to the “instructions” embodied by the threaded tiles of its left-hand operand. In short,

$$(\mathbf{FG})^k = ((\mathbf{FG})^{k-1}\mathbf{F}) \circ \mathbf{G},$$

which is to say that $(\mathbf{FG})^k$ consists of nothing but $m^{2k} \cdot n^{2(k-1)}$ threaded, congruent copies of \mathbf{G} 's shape.¹⁶ But \mathbf{G} 's shape is globally self-avoiding, and zips with itself in all nine possible supertile contact combinations (eight required, one optional). Thus no matter what the orientations, each individual \mathbf{G} in (\mathbf{FG}) will zip with copies of itself, just as it would in \mathbf{G}^k . At consecutive supertile connection points when two connected copies of \mathbf{G} are placed in a pair of supertiles, by **Handedness**_{3.16}, **FourCorners**_{3.7}, and **ToroidalCycles**_{3.17}, the in-degree and out-degree of the two supertiles' shared point are invariant (they remain each 1) after composition with \mathbf{G} . This occurs regardless of which of the five possible ways (Fig. 4) that two consecutive supertiles can connect in a self-avoiding manner.

Additionally, consider the Fibbinary bit sequence of 0s and 1s along any side of the composition's result. That sequence starts its formation as just the concatenation of m copies of the Fibbinary bits of the right operand in the composition. Such a concatenation is guaranteed to remain Fibbinary because each set of bits always starts and ends with a 0. But motif composition doesn't just concatenate the palindromic bit patterns. Two consecutive copies of the motif can be in the same or in different orientations. Because of this the last bit of the zipper bits next to the first bit of a concatenated copy can be $..00..$, $..01..$, $..10..$, but not $..11..$, based on whether the composed tile path is leaving or arriving at the perimeter where that 1 appears. But the Fibbinary zipper bits always start or end with *two* 0 bits, so a change of just the first or last bit from 0 to 1 cannot introduce two adjacent 1 bits, thereby maintaining the Fibbinary zipper bits in the composed result. ■

Remark 6.1. **NoLowerCornerTurns**_{3.22} is thus the cornerstone of preserving Fibbinary bit self-avoidance during composition. Fig. 54 demonstrates the composition of two globally self-avoiding motifs, one with zipper

¹⁵We use \mathcal{Z} for “zipping.”

¹⁶Also recall that motif composition is not commutative, therefore in general $(\mathbf{FG})^k \neq \mathbf{F}^k\mathbf{G}^k$.

bits 001010100 and the other with 0001000, resulting in a composite globally self-avoiding motif on $T_{63=9 \times 7}$ with palindromic zipper bits 0001001 0001001 0001000 1001000 1001001 0001001 0001000 1001000 1001000 (the underlined 1s are where a first or last 0 in the nine initially concatenated order-7 zipper bits has changed to a 1).

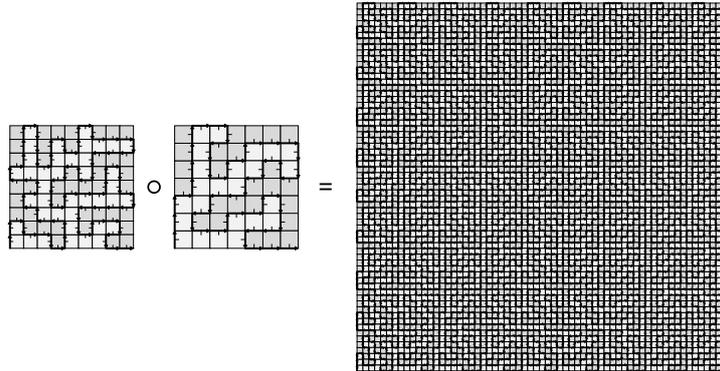


Figure 54: A composite globally self-avoiding motif on T_{63} composed from two globally self-avoiding motifs, one on T_9 and one on T_7 , each subject to a different mode (magnified result not drawn to scale).

7. Conclusions

The simplest, non-trivial, globally self-avoiding motif, **E**, for Hilbert-style space-filling curves using pure edge-replacement is a window into a fascinating combinatorial world: a monoid \mathcal{Z} of composable Hamiltonian cycles on the two-dimensional toroidal grid-graph. Imparting a slight asymmetry to a square by distinguishing one directed edge, and requiring continuity and self-avoidance during edge-replacement composition, ramifies into a space with interesting and elegant constraints that significantly affect how we might enumerate motifs.

The toroidal nature of the solution space is not an assumed convenience; it is induced by the requirements of self-avoidance, square-fillingness, and the one-to-one binding between a square and exactly one of its edges, under the monoid $\mathcal{Z} \subset \mathcal{H}$'s composition “growth” operator (especially self-similar growth). Interestingly, chirality and rotational symmetry emerge in the solution space while bilateral symmetry (à la the Hilbert Curve) disappears.

In this case, performing an exhaustive enumeration by machine in order to be able to better visualize an entire space of solutions led to an elementary, essentially visual, proof of an elegant boundary structure, thus demonstrating the utility of computers as a research/visualization tool in combinatorics.

A holy grail of any parameterized-by- n , constrained, combinatoric space is to find a closed formula $f(n)$, such as a generating function $g(n)$, that delivers or encodes the number of solutions for a given n . But the present result, for a very simple set of constraints, shows that a single formula is not likely, given that (a) each Fibbinary zipping mode creates differing amounts of constraint-based on a set of different boundary conditions, and (b) the number of such zipping modes increases without bound as $n \rightarrow \infty$.

Understanding and enumerating certain self-avoiding lattice paths can be enhanced by maintaining parity information, equivalent to knowing *locally* which side of an edge in a path is “outside” or “inside.”

With path self-avoidance as a fundamental constraint governing the problem, the appearance of a Fibonacci-related boundary condition is perhaps (in hindsight) not surprising (see, e.g., [1], [23]). But these constraints apply only to a subset of Hamiltonian cycles on the toroidal grid-graph: those that contact but do not cross the wraparound boundaries, i.e., when unwrapped, an order- n tile path remains bounded by a minimal $n \times n$ square. While self-avoidance characterizes all Hamiltonian paths and cycles, it is really the self-avoidance on both sides of a required, one-dimensional, shared, wraparound “cut” that induces the Fibbinary constraints.

Regardless, the submonoid \mathcal{Z} of a particular set of Hamiltonian cycles on the toroidal grid-graph—in which the cycle’s graph edges are bijectively bound to graph faces—embodies a beautiful combinatorial space that evinces an intriguing struggle between symmetry and asymmetry, and between order and disorder.

Acknowledgements

This paper reports the results of a private research project the author began in 1978, inspired by Bill Gosper’s seven-segment Flowsnake Curve construction based on edge-replacement in hexagonal tiles. Deeply appreciated are Oren Patashnik’s substantive criticism of and editing contributions to various longer and earlier drafts, which has taken shape over a good many years; what a long, strange thread it’s been. The University of Montana has my gratitude for letting me abuse their PDP-20 in 1978 for my first order-5 motif machine enumeration.

Discussions with Scott Kim encouraged and helped me, when the world was young, open the door into an interesting space of space-filling paths. Thanks also go to Don Knuth for his supportive comments when, over the years, our paths have crossed, or at least contacted toroidally. Also appreciated is the access to IBM Research's largest mainframe computer that Benoît Mandelbrot provided me in 1980, where my rudimentary brute-force search for order-6 solution motifs proved fruitless; had I only known then what I know now.

And, in memoriam, Malcolm C. McKenna, Donald C. McKenna, Jean Milliren, E. Butler Beaumont, Alan J. Perlis, William R. Bennett Jr., and Jef Raskin, for their support, teachings, and/or exemplary stick-to-it-iveness at various vertices of the author's distinctly non-Hamiltonian, non-toroidal, a piecewise-connected path through the graph of life.

References

- [1] A. T. Benjamin, *Self-Avoiding Walks and Fibonacci Numbers*, *Fibonacci Quart.* 44:4 (2006), 330–334.
- [2] F. M. Dekking, *Recurrent Sets*, *Advances in Math.* 44(1) (1982), 78–104 (p. 93).
- [3] T. Hales, *The Jordan Curve Theorem, Formally and Informally*, *Amer. Math. Monthly* 114:10 (2007), 882–894.
- [4] D. Hilbert, *Ueber Stetige Abbildung Einer Linie auf ein Flächenstück*, *Math. Annalen* 38 (1891), 359–460.
- [5] S. Kim, pers. corresp. with the author, 1983.
- [6] D. E. Knuth, *The Art of Computer Programming, Vol. 1, Fundamental Algorithms*, 2nd. Ed., Addison-Wesley (1973), p. 85, exercise 36.
- [7] B. B. Mandelbrot, *The Fractal Geometry of Nature*, W. H. Freeman (1982) (illustrated in part by D. M. McKenna, see p. 444).
- [8] D. M. McKenna, *SquaRecurves, E-Tours, Eddies, and Frenzies: Basic Families of Peano Curves on the Square Grid*, *The Lighter Side of Mathematics: Proc. Eugene Strens Memorial Conf. on Rec. Math. and its History*, 1986, *Math. Assoc. of America* (1994), 49–73.
- [9] D. M. McKenna, *Asymmetry vs. symmetry in a new class of space-filling curves*, *Conf. Proc. Bridges Math. Connections in Art, Music, and Science* (2006), 371–378, <http://archive.bridgesmathart.org/2006/bridges2006-371.html>.
- [10] D. M. McKenna, *The 7 Curve, Carpets, Quilts and Other Asymmetric Square-Filling Threaded Tile Designs*, *Conf. Proc. Bridges Math. Connections in Art, Music, and Science* (2007), 225–233, <http://archive.bridgesmathart.org/2007/bridges2007-225.html>.
- [11] D. M. McKenna, *Designing Symmetric Peano Curve Tiling Patterns with Escher-esque Fore-ground/Background Ambiguity*, *Conf. Proc. Bridges Math. Connections in Art, Music, and Science* (2008), 123–130, <http://archive.bridgesmathart.org/2008/bridges2008-123.html>.
- [12] D. M. McKenna, *Tendril Motifs for Space-Filling, Half-Domino Curves*, *Conf. Proc. Bridges Math. Connections in Art, Music, and Science* (2016), 119–126, <http://archive.bridgesmathart.org/2016/bridges2016-119.html>.
- [13] D. M. McKenna, *Hilbert Curves: Outside-In and Inside-Gone*, *Mathemaesthetics, Inc.* (2019), Ch. 3. (iPad/iPhone eBook available at <https://apps.apple.com/us/app/hilbert-curves/id1453611170>).
- [14] G. L. Nelson, *Real Time Transformation of Musical Material with Fractal Algorithms*, *Comput. Math. with Appl.* 32 (1996), 109–116 (discussing his composition entitled *Goss*).
- [15] A. Norton, *Eighty-eight Thousand, Four Hundred and Eighteen (More) Ways to Fill Space*, *Coll. Math. J.* 40:2 (2009), 108–112.
- [16] G. Peano, *Sur Une Courbe Qui Remplit Toute Une Aire Plaine*, *Math. Annalen* 36 (1890), 157–160.
- [17] P. Prusinkiewicz and A. Lindenmyer, *The Algorithmic Beauty of Plants*, Springer (1996), 12–14.
- [18] N. J. Rose, *Hilbert-Type Space-Filling Curves—In Memory of Hans Sagan (1928-2000)* (2000), N.C. State Univ., see, e.g., <http://www.dblab.ntua.gr/~gtsat/collection/SFCs/HilbertCurve.pdf>.
- [19] F. Ruskey and J. Sawada, *Bent Hamilton Cycles in d-Dimensional Grid Graphs*, *Elect. J. Combin.* 10(1) (2003), R1.
- [20] H. Sagan, *Space-Filling Curves*, Springer-Verlag (1994), ch. 1–3.
- [21] N. A. J. Sloane, *The On-Line Encyclopedia of Integer Sequences*, accessible electronically at <http://www.oeis.org>.
- [22] W. Wunderlich, *Über Peano-Kurven*, *Elem. Math.* 28 (1973), 1–10, see, e.g., <https://mathcurve.com/fractals/wunderlich/125.pdf>.
- [23] D. Zeilberger, *Self-avoiding walks, the language of science, and Fibonacci numbers*, *J. Statist. Plann. Inference* 54(1) (1996), 135–138.