Computational Complexity of Flattening Fixed-Angle Orthogonal Chains

Erik D. Demaine
Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, USA

Hiro Ito
School of Informatics and Engineering, The University of Electro-Communications, Japan

Jayson Lynch
Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, USA

Ryuhei Uehara
School of Information Science, JAIST, Japan

Abstract

Planar/flat configurations of fixed-angle chains and trees are well studied in the context of polymer science, molecular biology, and puzzles. In this paper, we focus on a simple type of fixed-angle linkage: every edge has unit length (equilateral), and each joint has a fixed angle of 90° (orthogonal) or 180° (straight). When the linkage forms a path (open chain), it always has a planar configuration, the zig-zag with alternating 90° angles between left and right turns. But when the linkage forms a cycle (closed chain), or is forced to lie in a box of fixed size, we prove that the flattening problem — deciding whether there is a planar noncrossing configuration — is strongly NP-complete.

Back to open chains, we turn to the Hydrophobic–Hydrophilic (HP) model of protein folding, where each vertex is labeled H or P, and the goal is to find a folding that maximizes the number of H–H adjacencies. In the well-studied HP model, the joint angles are not fixed. We introduce and analyze the fixed-angle HP model, which is motivated by real-world proteins. We prove strong NP-completeness of finding a planar noncrossing configuration of a fixed-angle orthogonal equilateral open chain with the most H–H adjacencies, even if the chain has only two H vertices. (Effectively, this lets us force the chain to be closed.)

Keywords and phrases Computational origami, equilateral linkage, fixed-angle linkage, HP model, NP-completeness, orthogonal linkage

Digital Object Identifier 10.57717/cgt.v2i2.34

Related Version A preliminary version of this paper was presented at the 34th Canadian Conference on Computational Geometry.

Funding *Hiro Ito*: Partially supported by JSPS Kakenhi Grant Number JP20K11671.
Ryuhei Uehara: Partially supported by JSPS Kakenhi Grant Number JP18H04091, JP20H05961, JP20H05964, JP20K11673, and JP22H01423.

Acknowledgements This work was initiated at the 3rd Virtual Workshop on Computational Geometry held in March 2022. We thank the other participants of that workshop — in particular Martin Demaine, David Eppstein, Timothy Gomez, and Aaron Williams — for helpful discussions and for providing a fruitful collaborative environment.

1 Introduction

In this paper, we introduce and investigate a new model of protein folding. We are given an ***equilateral fixed-angle chain*** (“protein”), where each vertex is marked H or P and has a specified fixed angle, and edges all have unit length. The goal is to embed the chain into a given grid (e.g., 2D square, 3D cube, 2D triangular, or 2D hexagonal) while

1. respecting the fixed angles (but each angle is still free to be a left or right turn in 2D or spin in 3D);
2. avoiding self-crossing in the embedding; and
3. maximizing the number of H–H grid adjacencies.

This is a fixed-angle version of the well-studied HP model of protein folding (where the angles are normally free to take on any value), which is known to be NP-hard in the 2D square grid [4] and 3D cube grid [3]. Fixed angles are motivated by real-world proteins; see [7, Chapters 8–9] for a comprehensive survey. Originally, it is a graph embedding problem with unit-length edges and fixed angles that originated from the protein folding problem, however, such problems have a history of being widely studied as graph embedding and linkage embedding problems (see, e.g., [13, 11, 9, 12, 14]). In the 2D square grid or 3D cube grid studied here, we can restrict to ***orthogonal*** fixed-angle chains where all fixed angles are 90° or 180°. For example, the popular "Tangle" toy restricts further to all fixed angles being 90°; see [5].

In the 3D cube grid, NP-hardness of fixed-angle unit-length linkage folding follows from [1] which proves NP-hardness of embedding a fixed-angle orthogonal equilateral chain of n^3 vertices into an $n \times n \times n$ 3D cube grid. Precisely, the authors of [1] proved NP-hardness of the snakes cube puzzle which asks us to fold a given linkage of unit-length edges with angles fixed to 90° or 180° of length $n^3 - 1$ into a the grid cube of size $n \times n \times n$.

In the 2D grid, NP-hardness of fixed-angle HP protein folding follows from [10] which proves NP-hardness of embedding a fixed-angle orthogonal equilateral bicolored chain of n vertices into a given bicolored grid. That is, the placements of the colored vertices are strongly constrained by the colored grid, and we need linear number of H vertices and linear number of P vertices to prove hardness. In this paper, we prove that the fixed-angle HP protein folding problem is NP-hard in the 2D square grid, even if the chain has only two H vertices and those vertices are its endpoints. In other words, given a fixed-angle orthogonal equilateral HP chain, we prove it is strongly NP-hard to find any planar noncrossing embedding where the endpoints (the two H vertices) are adjacent. This result is tight in the sense that any fixed-angle orthogonal equilateral chain with fewer than two H vertices (and hence can have no H–H adjacencies) has a noncrossing embedding, given by zig-zagging the 90° angles to alternate between left and right turns.

Fixed-angle HP protein folding where only the two endpoints are H vertices is nearly equivalent to finding any planar noncrossing embedding of a ***closed*** fixed-angle chain (where the first and last vertex are identified, and vertices are no longer marked H or P). This is called the ***flattening problem*** for fixed-angle closed chains. The only difference is that, in the flattening problem, the first/last vertex has a fixed-angle constraint, whereas in the HP model, the two necessarily adjacent H vertices could form any angle.

Nonetheless, we show that the flattening problem for fixed-angle orthogonal equilateral closed chains is strongly NP-complete. Past work proved strong NP-hardness when this problem was generalized to fixed-angle orthogonal equilateral caterpillar tree (instead of a chain) or when we allow nonorthogonal fixed angles (and working off-grid) [6], but left this case open.

Finally our work also addresses two open problems from [1]. We solve one open problem by proving strong NP-completeness of deciding whether a given fixed-angle orthogonal equilateral chain can be packed into a 2D square (whereas [1] proved an analogous result for a 3D cube). We also prove that this problem remains NP-complete when the chain is only a constant factor longer than the side length of the square (and thus the square is sparsely filled), answering the 2D analog of a 3D question from [1].

2 Preliminaries

2.1 Linkages

A ***linkage*** consists of a ***graph*** $G = (V, E)$ and edge-length function $\ell : E \to \mathbb{R}^+$. A ***configuration*** of a linkage in 2D is a mapping $C : V \to \mathbb{R}^2$ satisfying the constraint $\ell(u, v) = \|C(u) - C(v)\|$ for each edge $\{u, v\} \in E$. Let $x(C(u))$ and $y(C(u))$ be the x- and y-coordinates of $C(u)$, respectively. A configuration is ***noncrossing*** if any two edges $e_1, e_2 \in E$ intersect only at a shared vertex $v \in e_1 \cap e_2$.

A linkage is ***equilateral*** if $\ell(e) = 1$ for every $e \in E$. A linkage with n vertices is an ***open chain*** if its structure graph G is a path $(v_0, v_1, \ldots, v_{n-1})$, and it is a ***closed chain*** if G is a cycle $(v_0, v_1, \ldots, v_{n-1}, v_n = v_0)$. A ***fixed-angle chain*** is a chain together with an angle function $\theta : V \to [0°, 180°]$, constraining configurations to have an angle of $\theta(v)$ at every vertex v, except for the two endpoints of an open chain. A fixed-angle chain is ***orthogonal*** if we have $\theta(v_i) \in \{90°, 180°\}$ (in clockwise or counterclockwise order) for every vertex v_i with $0 < i < n - 1$. A ***segment*** of a fixed-angle chain is a consecutive subchain $v_i, v_{i+1}, \ldots, v_k$ with intermediate flat angles $\theta(v_j) = 180°$ for $i < j < k$, which acts the same as a single edge of length equal to the sum ($k - i$ for equilateral chains).

The ***embedding problem*** asks to determine whether a given linkage has a noncrossing configuration in 2D. For general linkages, this problem is $\exists\mathbb{R}$-complete [2]. For fixed-angle orthogonal chains, the problem is in NP: given a binary choice of turning left or right at each vertex, we can construct an explicit embedding — placing the first vertex at the origin and the second vertex on the positive x axis, and adding and subtracting lengths to the x and y coordinates — and check for collisions and (for closed chains) closure. In fact, for fixed-angle orthogonal *open* chains, every instance is a "yes" instance:

▶ **Observation 1.** *Every fixed-angle orthogonal open chain has a noncrossing configuration.*

Proof. Intuitively, we embed the chain C in a zig-zag. Precisely, let $P = (v_0, v_1, \ldots, v_{n-1})$ be the path structure graph. First we put v_0 at $(0, 0)$, and v_1 at $(1, 0)$. For each $i = 2, 3, \ldots, n - 1$, we define $x(C(v_i))$ and $y(C(v_i))$ as follows. When $\theta(v_{i-1}) = 180°$, we have no choice: $x(C(v_i)) = x(C(v_{i-1})) + (x(C(v_{i-1})) - x(C(v_{i-2})))$ and $y(C(v_i)) = y(C(v_{i-1})) + (y(C(v_{i-1})) - y(C(v_{i-2})))$. When $\theta(v_{i-1}) = 90°$ and $\overrightarrow{C(v_{i-2})C(v_{i-1})}$ is horizontal, we define $x(C(v_i)) = x(C(v_{i-1}))$ and $y(C(v_i)) = y(C(v_{i-1})) + 1$. If it is vertical, we define $x(C(v_i)) = x(C(v_{i-1})) + 1$ and $y(C(v_i)) = y(C(v_{i-1}))$. The obtained configuration is noncrossing because it proceeds monotonically in x and y, with strict increase in one of the coordinates in each step. ◀

We note that Observation 1 holds for any fixed-angle orthogonal open chain which is not necessarily equilateral.

In the ***HP model***, the structure graph $G = (V, E)$ has its vertices ***bicolored*** by a color function $\omega : V \to \{H, P\}$. For a configuration C of an equilateral orthogonal linkage, a pair (u, v) of vertices forms an ***H–H contact*** if $\omega(u) = \omega(v) = H$, $\|C(u) - C(v)\| = 1$, and $\{u, v\} \notin E$. The ***HP optimal folding problem*** of a bicolored fixed-angle orthogonal equilateral chain asks to find a noncrossing configuration of the linkage in 2D or 3D that maximizes the number of H–H contacts.

2.2 Linked planar 3SAT

Our reductions are from a variant of 3SAT, which we define and sketch the known NP-hardness proof in this section.

In the standard ***3SAT*** problem, we are given a Boolean formula ϕ over a set V of n binary variables, where ϕ is a conjunction of a set C of m ***clauses***, where each clause in C is a disjunction of at most three ***literals***, where each literal is of the form x or $\neg x$ for some variable $x \in V$. The 3SAT problem asks us to determine if ϕ has a truth assignment or not. In ***planar 3SAT***, we form the graph $G_\phi = (C \cup V, E)$ with a vertex for each variable in V and each clause in C, and edges between variables and the clauses that contain them, and require that G_ϕ has a planar embedding.

We mimic a variant of planar 3SAT with additional planarity restrictions: if we add edges to form a Hamiltonian cycle κ of $C \cup V$ that first visits all clauses in C in some order $c_1, c_2, \ldots, c_m$, and then visits all variables in V in some order $v_1, v_2, \ldots, v_n$, then the resulting graph $G'_\phi = G_\phi \cup \kappa$ must also be planar, as in Figure 1. The ***linked planar 3SAT problem*** asks, given ϕ, G_ϕ, κ, and a planar embedding of G'_ϕ, whether ϕ is satisfiable. Pilz [8] proved this problem NP-complete. Our proof follows the same high-level structure of this proof, so we briefly review it now:

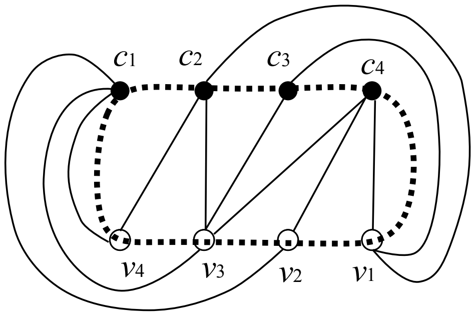

Figure 1 An example instance of linked planar 3SAT, where $c_1 = (\neg v_2 \vee \neg v_3 \vee \neg v_4)$, $c_2 = (v_4 \vee v_3 \vee \neg v_1)$, $c_3 = (\neg v_3 \vee v_1)$, and $c_4 = (v_1 \vee v_2 \vee v_3)$. Hamiltonian cycle κ (drawn dotted) visits $c_1, c_2, c_3, c_4, v_1, v_2, v_3, v_4$ in cyclic order.

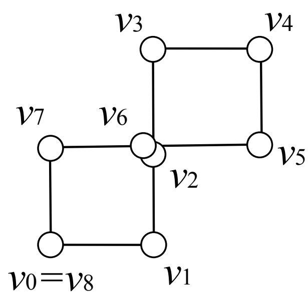

Figure 2 A crossing chain $(v_0, v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8 = v_0)$ with angles $\theta(v_2) = \theta(v_6) = 180°$ and $\theta(v_i) = 90°$ for $i = 0, 1, 3, 4, 5, 7$.

▶ **Theorem 2** ([8]). *Linked planar 3SAT is NP-complete.*

Proof. (Outline) Pilz gives a reduction from planar 3SAT to linked planar 3SAT[8], which is illustrated in Figure 3. Let ϕ be a given instance of planar 3SAT, and fix an arbitrary straight-line embedding of the planar graph G_ϕ. First we draw the Hamiltonian cycle κ (as on the left) in a large doubled spiral, with the inward spiral being the subpath for clauses and the outward spiral being the subpath for variables (note that we can assume tht κ is $(c_1, c_2, \ldots, c_m, v_1, v_2, \ldots v_n, c_1)$ without loss of generality in Pilz's construction.). Focusing on one half (the dotted square), we obtain a square with alternating horizontal grid lines for clauses and variables.[1] Then we construct a planar drawing of $G_{\tilde{\phi}}$ from G_ϕ with no horizontal edges, all variable vertices on odd grid lines (orange/light), and all clause vertices on even grid lines (purple/dark) as shown in Figure reffig:linked-3sat. This construction

[1] See [8] for the details how to arrange them on a square grid.

of $G_{\tilde{\phi}}$ correctly orders the vertices on κ as on G_ϕ. Finally we replace each edge of this drawing with a sequence of gadgets (as on the right) that duplicates the variable across the horizontal grid lines to reach the clause. Although not notated in our figure, each 4-cycle gadget alternates between positive and negative literals, so that for each variable x and copy x' we have the clauses $\neg x \vee x'$ (i.e., $x \Rightarrow x'$) and $\neg x' \vee x$ (i.e., $x' \Rightarrow x$), which together imply $x = x'$ (i.e., $x \Leftrightarrow x'$). ◀

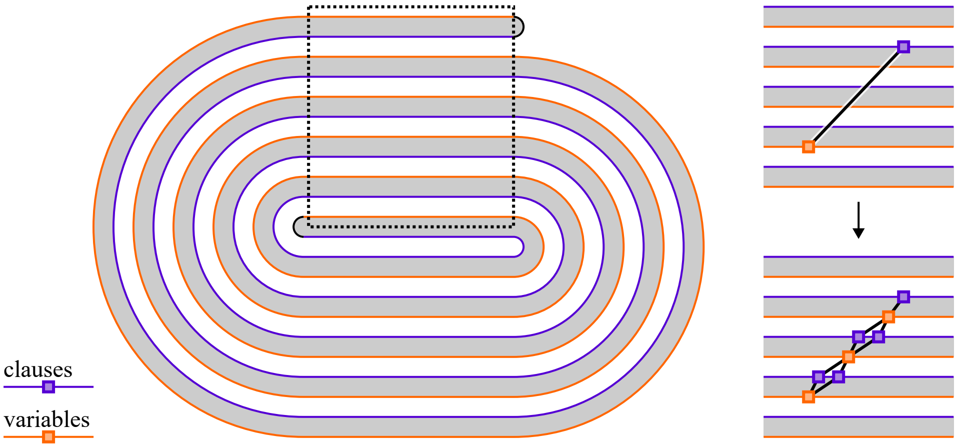

Figure 3 Pilz's reduction from planar 3SAT to linked planar 3SAT [8], based on Figs. 3 and 4 of [8]. (For consistency with later figures, we use colors instead of dashes to distinguish variables and clauses, and rotate the figure 90°.)

3 Embedding fixed-angle orthogonal equilateral closed chains is strongly NP-complete

In contrast to Observation 1, not all fixed-angle orthogonal equilateral *closed* chains are "yes" instances of the embedding problem. In particular, an orthogonal equilateral closed chain must have an even number of edges to have a configuration in 2D. Even with this property, the length-8 chain $(v_0, v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8 = v_0)$ with angles $\theta(v_2) = \theta(v_6) = 180°$ and $\theta(v_i) = 90°$ for $i = 0, 1, 3, 4, 5, 7$ has configurations in 2D but they have crossings at vertices v_2 and v_6 (Figure 2). It is not difficult to show that the embedding problem for fixed-angle orthogonal closed chains is weakly NP-hard by a reduction from the ruler folding problem (see [7, Chap. 2]); this construction requires exponential edge lengths (or equilateral chains with exponentially long straight segments) to be intractable. In this section, we prove that the embedding problem is strongly NP-complete:

▶ **Theorem 3.** *Embedding a fixed-angle orthogonal equilateral closed chain in 2D is strongly NP-complete.*

Section 2.1 argued membership in NP. To show NP-hardness, we mimic the reduction from planar 3SAT to linked planar 3SAT given by Theorem 2. That is, we give a polynomial time reduction from linked planar 3SAT. In particular, assume we have constructed a formula ϕ, the associated graph $G_\phi = (C \cup V, E)$, a Hamiltonian path κ visiting $c_1, c_2, \ldots, c_m, v_1, v_2, \ldots, v_n$ in cyclic order, and a planar embedding of $G_\phi \cup \kappa$ with κ spiraling as in Figure 3. Figure 4 shows an example of a planar 3SAT instance and the result from this transformation, and Figure 5 shows the final result of our reduction.

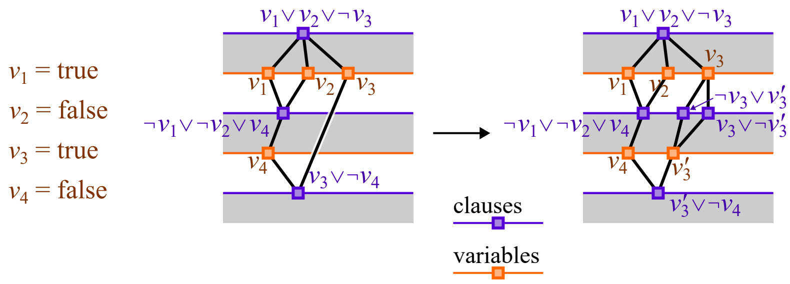

Figure 4 Example planar 3SAT instance and the result of the transformation from Theorem 2.

Figure 6 gives an overview of our construction. In the construction, we have to draw all the gadgets by a fixed-angle orthogonal equilateral chain in one stroke. The chain first draws a double outer frame (shaded gray) and encloses all the gadgets within it. In the frame, rows cycle through sections containing different types of gadgets — "insulation", variables, more insulation, sheaths for clauses, choices for clauses, sheaths for clauses — starting and ending with insulation. The chain thus visits alternating rows of variables and sheaths followed by alternating rows of insulation and clause choices, but it is different from κ in a linked 3SAT instance because it additionally visits rows of insulation and sheaths in between variables and clauses that are connected together. We use the frame and the spiraling of the drawn segments to force the endpoints of each row of gadgets to be as pictured. Then we argue that the insulation gadgets force the locations, endpoints, and bounding boxes of individual variable, sheath, and clause gadgets, which allows us to argue their correctness.

Hereafter, we will give the details of frame gadget, spiral gadget, insulation gadget, choice gadget, clause and sheath gadgets, and variable gadget. Note that the gap in the doubled-frame is different from the other gadgets. This difference in gaps ensures that there is no other way for the frame and other gadgets to fit inside it.

3.1 Frame gadget

We start at the top level with the ***frame gadget***, shown in Figure 7, which surrounds all other gadgets. In fact, the frame gadget works with a $5\times$ scaled version of the rest of the construction, so that we can use congruence arguments (modulo 5).

Formally, suppose we are given a closed chain C, a distinguished edge $\{u, v\}$, an intended bounding box B that the chain C should fold into, and two adjacent points p, q on the bounding box that u, v should fold to respectively. We scale all inputs uniformly by a factor of 5 to produce a new chain $5C$ with a distinguish edge $\{5u, 5v\}$ and a new bounding box $5B$ with distinguished points $5p$ and $5q$ at distance 5 from each other. In particular, $5C$ is the result of replacing each edge of the original chain C with a straight chain of five edges. Thus, in $5C$, all straight chains between $90°$ vertices, which we call ***segments***, have lengths equal to 0 modulo 5. Equivalently, the vertices of $5C$ stay on a scaled square grid with 5×5 cells. We then modify $5C$ by replacing the five-edge chain between $5u$ and $5v$ with the frame gadget F in Figure 7 (i.e., all pictured edges), with the intended folding building most of two rectangles 5 and 6 units away from the bounding box $5B$. Let C' be the resulting chain (the combination of $5C$ and F).

▷ Claim 4. The frame gadget F has a unique folding up to isometries, namely the one shown in Figure 7.

Note that Claim 4 holds for any embedding of the chain.

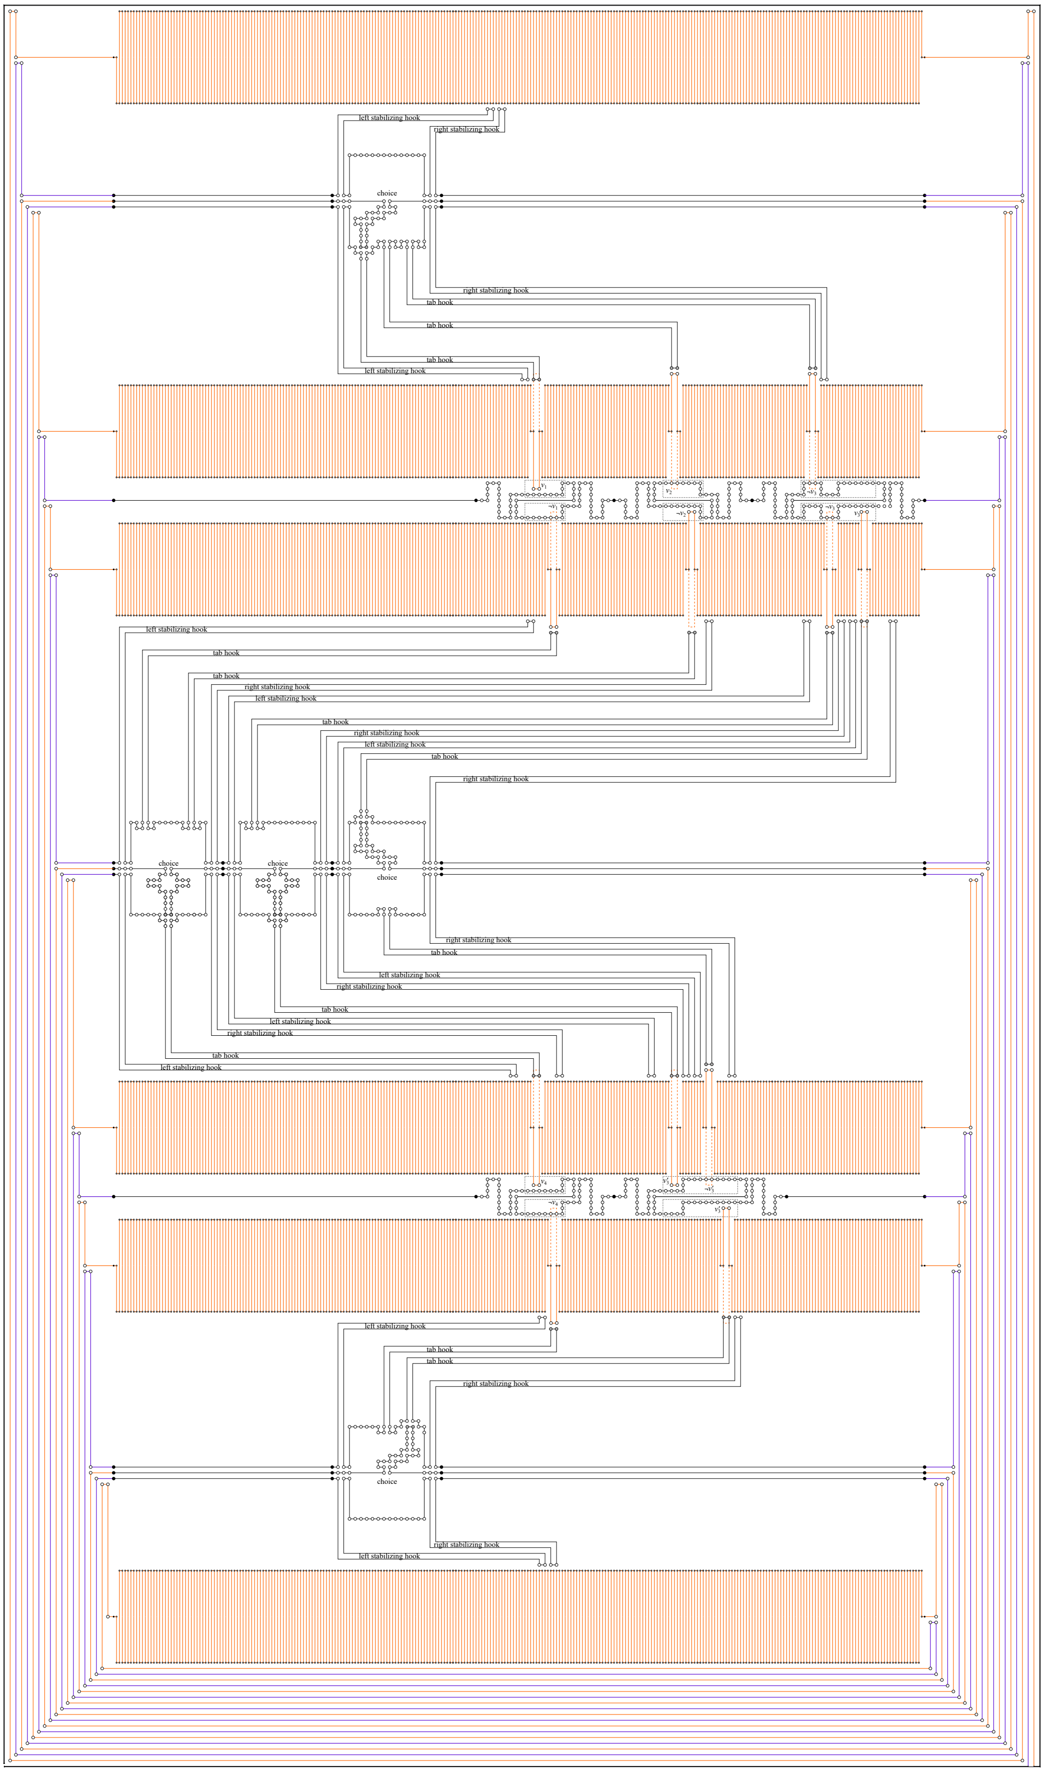

Figure 5 An example of the reduction from the instance in Figure 4, and a solution embedding corresponding to assignment $v_1 =$ true, $v_2 =$ false, $v_3 =$ true, and $v_4 =$ true. (Insulation height and hook segments are not drawn to scale.)

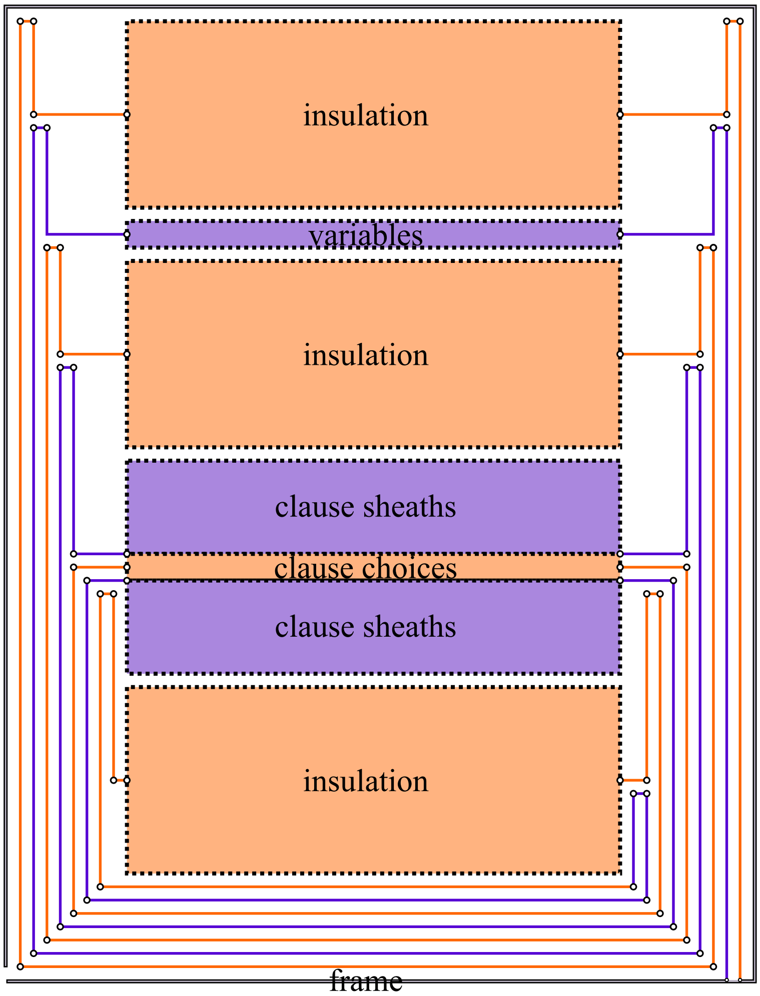

Figure 6 Construction overview, consisting of the frame gadget (shaded gray, on a grid 5× smaller than the rest), spiral gadget (colored segments), and rows of gadgets (dotted boxes) cycling through insulation gadgets, variable gadgets, sheath gadgets, choice gadgets, sheath gadgets, . . . , and ending with insulation gadgets. Insulation rows are extremely thick ($\Theta(n^2)$); sheath rows are medium thickness ($\Theta(n)$); variable rows are thin ($\Theta(1)$); and choice rows are extremely thin (1). (Straight vertices are not drawn to simplify the figure.)

Proof. First observe that, once we embed one edge of the chain, we know which segments are horizontal and vertical. Thus we can assume by isometry that the horizontal/vertical assignment is as in Figure 7.

Because the chain C' forms a cycle, the signed vertical and horizontal distances must sum to zero, and so in particular must sum to 0 modulo 5. The outer frame of the frame gadget consists of three vertical straight chains and three horizontal segments. In each group of three, there are two segments of length 1 modulo 5 (one segment in fact has length exactly 1), and one segment of length 2 modulo 5. All other segments (including those in $5C$) have length 0 modulo 5. The only solutions to $\pm 2 \pm 1 \pm 1 \equiv 0$ modulo 5 are $2 - 1 - 1 \equiv 0$ and $-2+1+1 \equiv 0$. Thus the 2-modulo-5 segment in the group must have the opposite orientation from the two 1-modulo-5 segments in the group. This forces the folding of the frame gadget up to reflections in the coordinate axes.

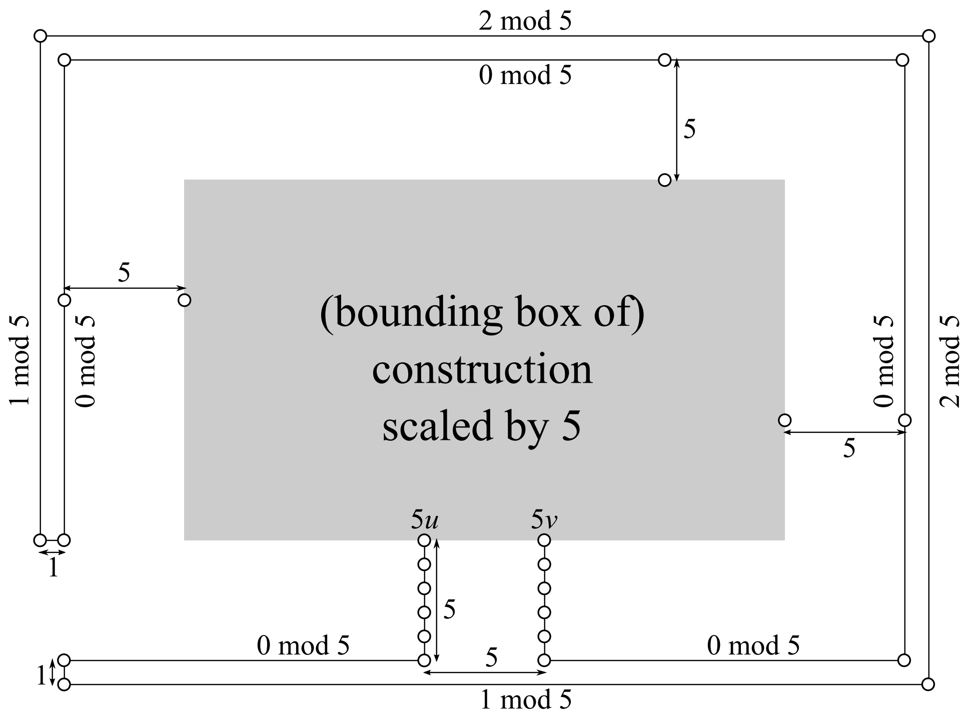

Figure 7 Frame gadget for closed chains, consisting of a frame construction F (all pictured edges) attached to a $5\times$ scaled version of a given construction. (Some straight vertices are not drawn to simplify the figure.)

The rest of the frame gadget has segment lengths that are 0 modulo 5, so it is impossible for it to transition between inside and outside of the outer frame (of the doubled frame) without collisions. (The gap in the lower-left corner does not intersect any points on the scaled square grid.) Thus the rest of the chain must be inside: if the two bottom left unit-length segments next to the outer frame both went outside, then they would immediately collide; and if one went outside and one went inside, then they could never meet. Now the rest of the inner frame has a forced folding because every segment length is at least 5, while the frame border has thickness 1, so each consecutive edge is forced to fold the only direction that remains inside the outer frame. ◁

▷ Claim 5. The framed construction C' has a planar noncrossing embedding if and only if C has a planar noncrossing embedding within B.

Proof. This follows from Claim 4 and that the scaled input chain $5C$ cannot escape the frame (in the lower left) because it remains in a scaled square grid of coordinates 0 modulo 5. ◁

Technically, the frame gadget comes last: the chain C is the rest of the construction (to be specified), and then our overall reduction is the framed chain C'. By Claim 5, we can assume that the rest of the construction C is forced to remain within a desired bounding box B.

3.2 Spiral gadget

The ***spiral gadget*** consists of the colored segments in Figure 6, i.e., all drawn segments other than the frame gadget. Other than the factor-5 scaling from the frame gadget, the segments of the spiral gadget are drawn as tightly as possible, on adjacent (scaled) grid lines.

▷ Claim 6. The spiral gadget has a unique folding, pictured in Figure 6. In particular, the endpoints where dotted boxes meet the spiral gadget have fixed positions relative to the frame.

Proof. Consider the segments $s_1, s_2, \ldots, s_k$ of the spiral gadget, starting with the two extreme segments s_1, s_k that attach to the frame; refer to Figure 8. These segments must both go upward, to avoid intersecting the frame. The outer (right) segment s_1 is maximally right in the bounding box B and goes all the way to the top, while the inner (left) segment s_k is one grid line to the left of s_1 and is a little shorter. For consistent labeling of segments, we assume that s_1 [and s_k] are immediately followed preceded] by a U-turn, but we also describe what modifications are necessary to handle the non-U-turn case.

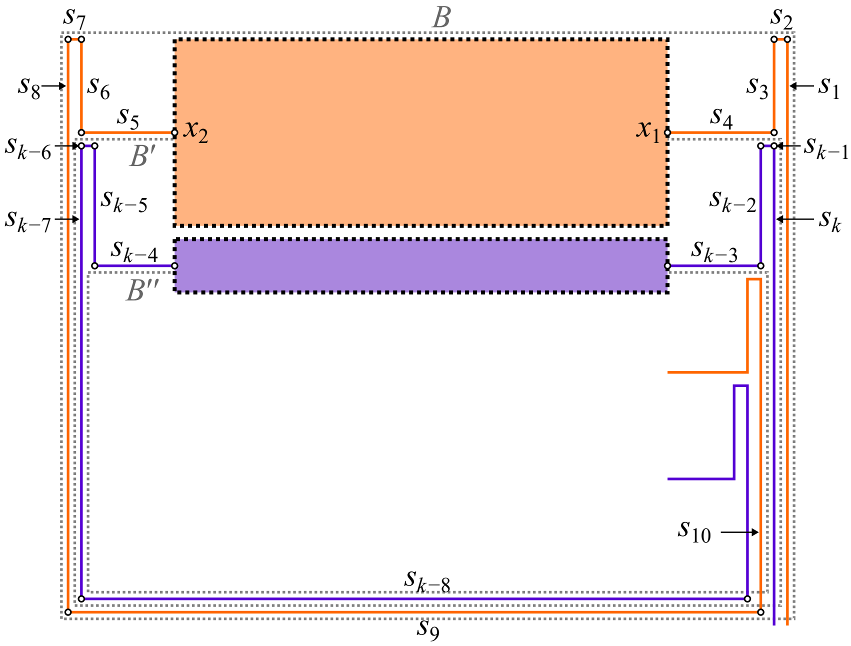

Figure 8 Forced folding of spiral gadget from Figure 6.

We argue that segments s_1, s_2, s_3, s_4 have forced foldings. Because s_1 is maximally right in B, segment s_2 of length 1 must go left to avoid intersecting the right side of the frame. Because s_1 reaches the very top of B, segment s_3 must then go down to avoid intersecting the top side of the frame. Because s_2 has length 1, so s_3 is just one grid line left of s_1, segment s_4 must go left to avoid intersecting s_1. If s_1 is in fact not followed by a U-turn, then we view segments s_2 and s_3 as having length 0, and instead s_4 directly turns; in this case, s_4 must turn left to avoid intersecting the right side of the frame.

Segment s_4 reaches one endpoint x_1 of the topmost clause row. We determine the location of the other endpoint x_2, and the locations of the following segments s_5, s_6, s_7, s_8, s_9, as follows. Let s_5 denote the next segment of the spiral gadget, which is attached to x_2. Segment s_8 is the entire height of the bounding box B, which enforces the set of y coordinates of its ends, but not whether it points up or down. In fact, s_8 must point down: otherwise, segment s_9 would be at the top and go rightward to be adjacent to (but not quite intersecting) s_2, and then segment s_{10} would be forced to go down (to avoid leaving B) and intersect s_4: s_{10} is at least as long as the height of the bottommost dotted box, which we assumed is the tallest dotted box, so is at least as long as s_3, which is half the height of the current top dotted box. Thus s_9 must be at the very bottom of the bounding box B. Segments s_7, s_6, and s_5 then have forced orientations of left, up, and left (from x_2), to avoid local intersection with the left side of the frame, the top side of the frame, and s_8, respectively.

At this point, we have shrunk the effective bounding box B by 1 unit on the right, left, and bottom sides (because of segments s_1, s_8, and s_9) and put the top side just below segments s_4 and s_5. The new bounding shape B' is not exactly a box, because the topmost clause row has an unknown shape, but we can treat it as a box because we will only argue about columns that do not overlap any clause/variable row.

Next we look at the other end of the chain. The process of folding is performed by alternating constraints between the orange and purple gadgets in Figure 6. We first apply the s_1, s_2, s_3, s_4 argument to put the first insulation gadget. Then it leads the foldings of $s_k, s_{k-1}, s_{k-2}, s_{k-3}, whichdeterminesthepositionofthevariablegadget.$ Then we apply the s_5, s_6, s_7, s_8, s_9 argument for the second insulation gadget, which determines the foldings of $s_{k-4}, s_{k-5}, s_{k-6}, s_{k-7}, s_{k-8}$. It leads the first clause sheaths, and so on. In each placement of a gadget, we have shrunk the effective bounding box by 1 unit on the right, left, and bottom sides, and put the top side just below segments.

With the smaller bounding box, we can repeat the same arguments, alternating between analyzing the next nine segments from the front and the previous nine segments from the back. Here we exploit the symmetry of the spiral gadget construction: each round with a smaller box looks just like the previous round. The only special segment is $s_{(k+1)/2}$ which connects the front and back of the chain directly instead of via a variable/clause row, but this only simplifies the argument. In the end, we determine all edges of the spiral gadget. ◀

3.3 Insulation gadget

Figure 9 shows an ***insulation gadget*** which occupies an entire row in Figure 6. This gadget works on a half-grid relative to all other gadgets except the frame; in other words, the spiral, variable, sheath, and clause gadgets are all scaled 2× relative to the insulation gadget, which is scaled 5× relative to the frame gadget, for a total scale of 10× for the spiral, variable, sheath, and clause gadgets. For simplicity, we allow segments of length $\frac{1}{2}$ so that the spiral remains on a 1×1 grid; but in the end everything will be scaled by 10× to ensure all lengths are integers divisible by 5.

The insulation gadget starts and ends with horizontal segments of length $\frac{1}{2}$ that lie on a common horizontal ***axis***. In between, the gadget consists of an alternation between "grilles" and "tabs", starting and ending with a grille, separated by horizontal segments of length $\frac{1}{2}$ that also lie on the axis. A ***grille*** consists of a sequence of segments with lengths in the pattern $h, \frac{1}{2}, 2h, \left(\frac{1}{2}, 2h\right)^{2k}, \frac{1}{2}, h$, where $\cdots^{2k}$ denotes an even number of repetitions, resulting in an integer width $k+1$ (for any desired integer $k \geq 0$). A ***tab*** consists of three segments with lengths $h+2, 1, h+2$, which has width 1. Notably, tabs lie on the integral grid used by all other gadgets other than the frame, and extend 2 units farther than grilles. All grilles and tabs use the same integer parameter h, the ***half-height*** of the insulation gadget.

To force the folding of an insulation gadget, we need the property that h is larger than the height of any possible folding of the adjacent rows of gadgets above or below the insulation gadget. This property is easy to achieve by setting h to be larger than the sum of lengths of all vertical segments in those rows.

▷ Claim 7. In any folding of the entire construction, an insulation gadget must be folded as in Figure 9, with each grille and tab optionally reflected through the axis.

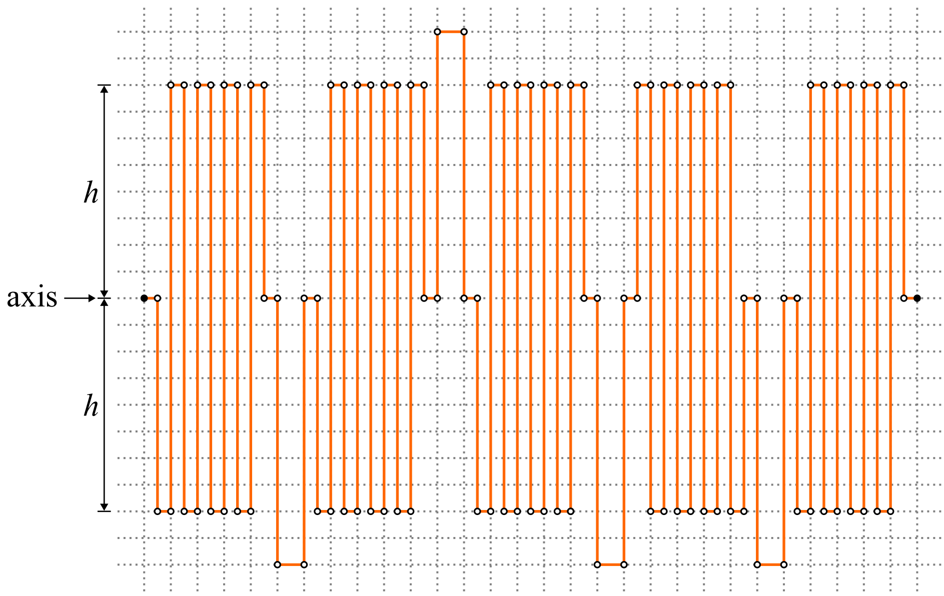

Figure 9 An example insulation gadget of height h. (Straight vertices are not drawn to simplify the figure.)

Proof. Because the endpoints of the insulation gadget (filled in black) are attached to the spiral gadget, their positions are fixed and the incident length-$\frac{1}{2}$ segments must be horizontal by Claim 6. We refer to the extension of these two extreme segments as the ***axis***.

Each grille and tab has an initial choice for its first segment (which has length h) to fold up or down, corresponding to choosing one folding or its reflection through the axis. After this first segment, the chain is at $\pm h$ vertical distance from the axis. The third segment, which has length h or $2h$, must go toward the axis (i.e., in the opposite direction as the first segment); otherwise, the insulation gadget penetrates the adjacent row of gadgets by a vertical distance of at least h, which prevents the chain in that row from connecting its two endpoints without collisions, by our assumption that h is larger than the height of any possible folding of the row. Thus the second segment (which has length $\frac{1}{2}$) must go rightward; otherwise, the third segment would intersect the previous horizontal segment.

For the grille, any remaining segments of length $2h$ must alternate in direction, always going toward the axis, by a symmetric argument, as we remain at $\pm h$ vertical distance from the axis; and all intervening segments of length $\frac{1}{2}$ must go rightward or they would cause collision between the previous and next vertical segments.

For both the grille and tab, the last segment (which has length h) must again go toward the axis, and thereby return to the axis. The next-to-last segment (of length $\frac{1}{2}$ for a grille and 1 for a tab) must again go rightward: for a grille, going left would cause collision between the previous and last vertical segment, and for a tab, going left would cause collision between the previous grille (which exists because the insulation gadget starts with a grille) and the last vertical segment.

Therefore every grille and tab is forced modulo the initial up/down choice. ◁

For understanding the impact of insulation on other gadgets, we can restrict attention to the occupied points of the integral grid in Figure 9. Grilles act as walls: independent of whether they are flipped through the axis, they block the same $(k + 1) \times (2h + 1)$ rectangle of points. Tabs act as local binary choices (wires): they either block a top $2 \times (h + 2)$ rectangle of points above the axis and leave empty the bottom $2 \times (h+2)$ rectangle of points below the axis, or vice versa. In particular, we use the fact that the vertical gap between two consecutive grilles (on the side that does not have the tab) is only two grid points wide. Thus, if an integral chain entered and exited such a gap, it must have exactly two turns separated by a segment of length 1; no more turns are possible.

Tabs in the insulation gadget provide communication wires between gadgets above and below the insulation, while the forced blocking of grilles gives those adjacent gadgets an effective bounding box. The insulation gadget can support any pattern of (integer-aligned) tabs, provided no two tabs are consecutive and there is no tab at the left or right extreme: we simply fill in the remaining space with grilles.

3.4 Choice gadget

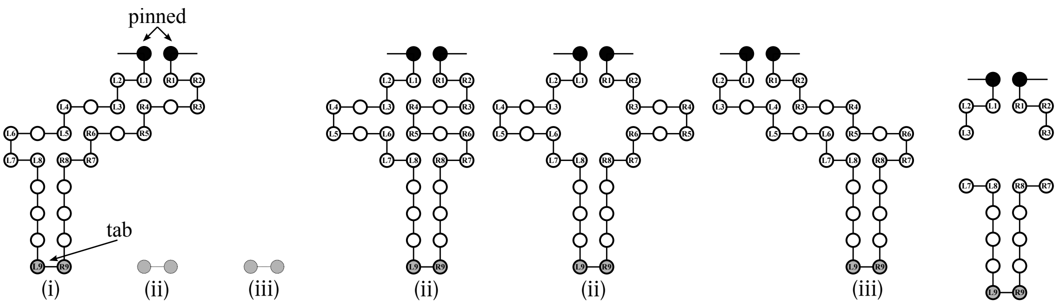

(a) The gray tab can be placed in one of three places below the pinned vertices. **(b)** Representative configurations (modulo reflection). **(c)** Forced choices.

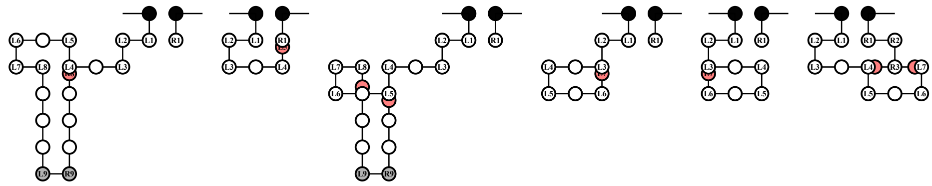

(d) Attempting to turn L4 or L6 upward causes intersection.

Figure 10 Choice gadget, where black vertices are pinned.

Figure 10 defines the ***choice gadget***, which is the central part of a clause gadget. We refer to the two gray vertices as the ***tab*** of this gadget. For now, we assume the two black endpoints are horizontally adjacent; we will effectively pin these endpoints later. We will also assume that the both ends of the chain turn the same direction — either upward or downward — at these endpoints, which will be forced by the long chains attached to the endpoints. Under these assumptions, there are six types of configurations as characterized by the tab locations:

▷ Claim 8. Assuming the endpoints are horizontally adjacent both turn the same direction (upward or downward), the tab of a choice gadget can be placed in exactly six locations (with three different horizontal shifts).

In Figure 10a, three possible tab locations are labeled by (i), (ii), and (iii).

Proof. We claim that there are ten possible embeddings of the choice gadget under the assumption in Claim 8: the three in Figure 10b, their reflections through the vertical line bisecting the endpoints (which adds two more, as the middle diagram is reflectionally symmetric), and the reflections of these five embeddings through the horizontal line connecting the endpoints. We assume by symmetry that the endpoints both turn downward, reducing to the first five embeddings.

To enumerate all embeddings, we consider which way each 90° vertex can turn. As in Figure 10, we label the 90° vertices after the left endpoint by L1, L2, . . . , L9, where L9 is the left tab vertex; and symmetrically label the 90° vertices before the right endpoint by R1, R2, . . . , R9, where R9 is the right tab vertex. These 90° vertices offer a sequence of choices, alternating between turning left vs. right and turning up vs. down.

Some of these choices are immediately forced; refer to Figure 10c. First, L1 and R1 (the two vertices adjacent to the endpoints) must turn away from each other, in order to not immediately intersect each other. Second, L2 and R2 must turn downward, in order to avoid intersecting the other neighbors of the endpoints. Third, L9 and R9 (the tab vertices) must both point upward: if L9 pointed downward, say, then L8 would be so much lower than L9 that the chain could not reach the left endpoint without intersection. Fourth, L8 and R8 must turn away from each other, in order to not immediately intersect each other. Furthermore, L8 must point left and R8 must point right — otherwise, they could not reach their respective endpoints without crossing — so L7 must turn right and R7 must turn left.

Figure 10b and its reflections correspond to always (in particular, L4, R4, L6, R6) choosing to turn down instead of up; making all possible left vs. right choices for L3 and R3; and making all possible left vs. right choices for L5 and R5 that keep the tab vertices in adjacent columns. The remaining cases to consider are when L4, R4, L6, or R6 turn up. By symmetry, it suffices to consider the cases when L4 or L6 turn up. No matter what choices we made for L3 and L5, the gadget causes intersect in these cases as shown in Figure 10d. ◁

3.5 Clause and sheath gadgets

The clause gadget, shown in Figures 11 and 12, consists of three separate chains (with endpoints marked black): one extending the choice gadget, and one "sheath" above and one below the choice chain. On the left and right ends of the choice gadget, we add suitably long horizontal segments. Above and below the choice gadget, we add a ***sheath gadget*** forming a rectangular "container" around the choice gadget, and up to five outward ***hooks*** attached. Each hook is a path of three segments (vertical, horizontal, vertical) doubled to have thickness 1.

At minimum, each sheath of the clause gadget has one ***stabilizing hook*** immediately left and right of the container, and each of these hooks has its end adjacent to a grille of the adjacent insulation gadget. In addition, the two sheaths of a clause have a total of three ***tab hooks*** attached to the container, for connections to variable gadgets either above or below the clause. Tab hooks are attached to the sheath (container) via two length-1 segments on either side, allowing for the hooks to extend or retract their ***tabs*** (gray vertices) vertically with an offset of 2; the attachment is aligned so that the choice gadget's extended tab forces the corresponding sheath tab hook to be extended. The tabs of a tab hook are horizontally aligned with a tab of the adjacent insulation gadget; when the tab hook is retracted, the insulation tab is immediately adjacent (but not intersecting), and when the tab hook is extended, the insulation tab would intersect and so is forced to flip away. The horizontal segments of different hooks are separated enough from each other so that each hook can freely extend or retract.

Figures 11 and 12 show two different versions of the clause gadget. In general, for each connection in G_ϕ from this clause to a variable gadget in the row below, we assign one of

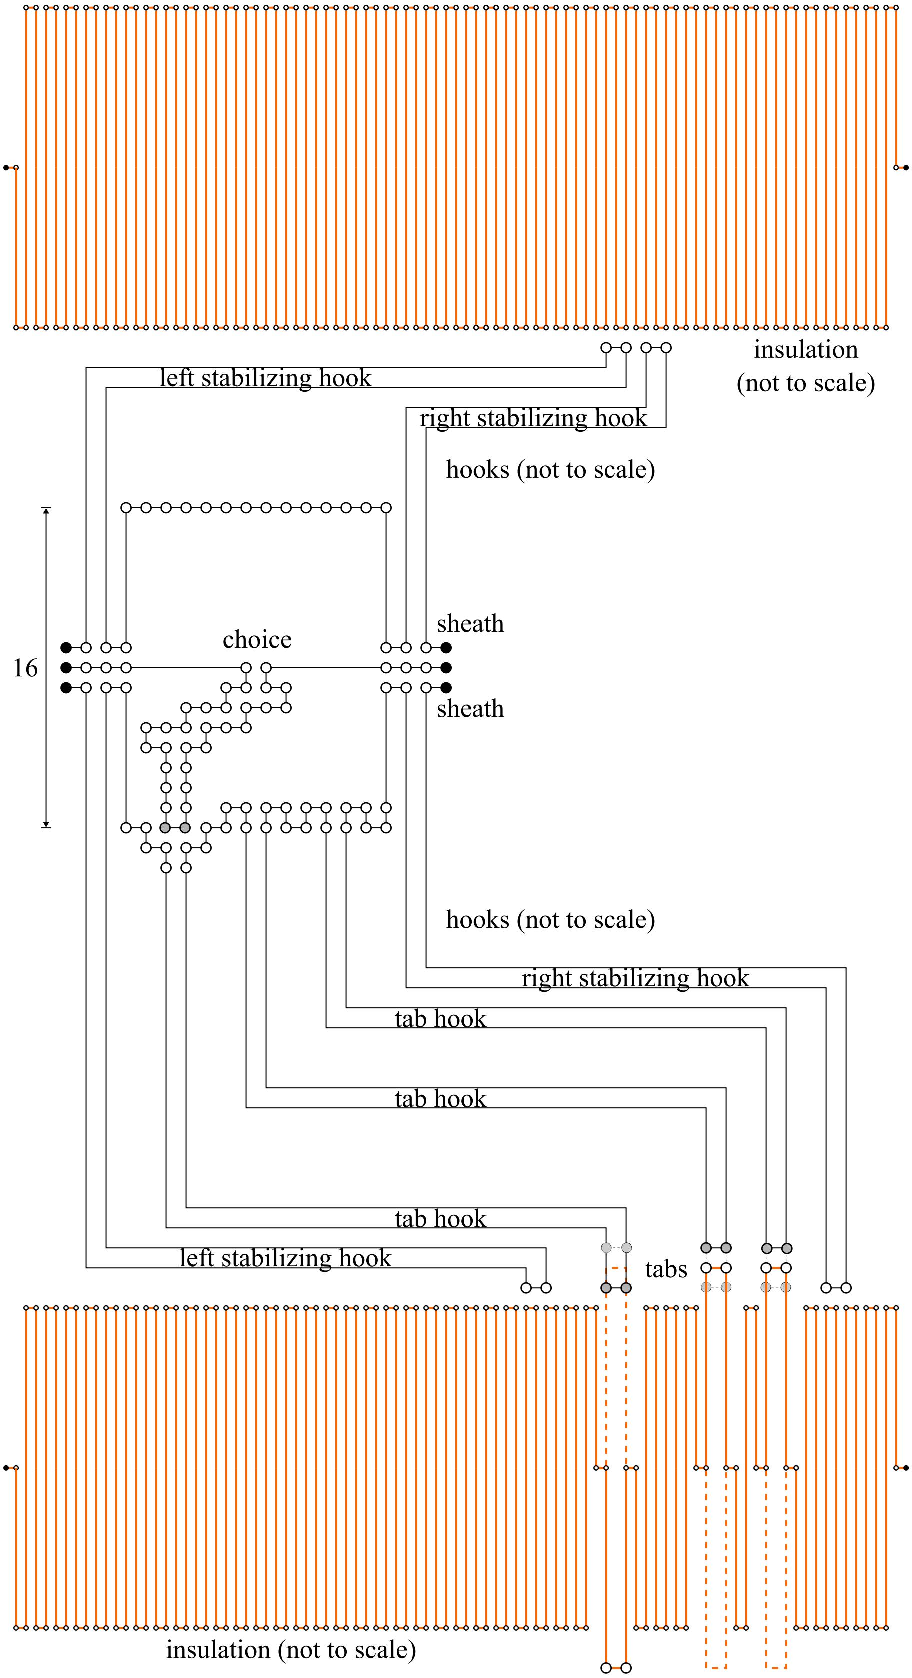

Figure 11 One version of clause gadget (three chains of black segments with black endpoints) and its interaction with the neighboring insulation gadgets (colored segments). This version has three tabs on the bottom, and the leftmost tab extended. Each tab shows the unused alternate state with dashed lines. (Some vertices are not drawn to simplify the figure.)

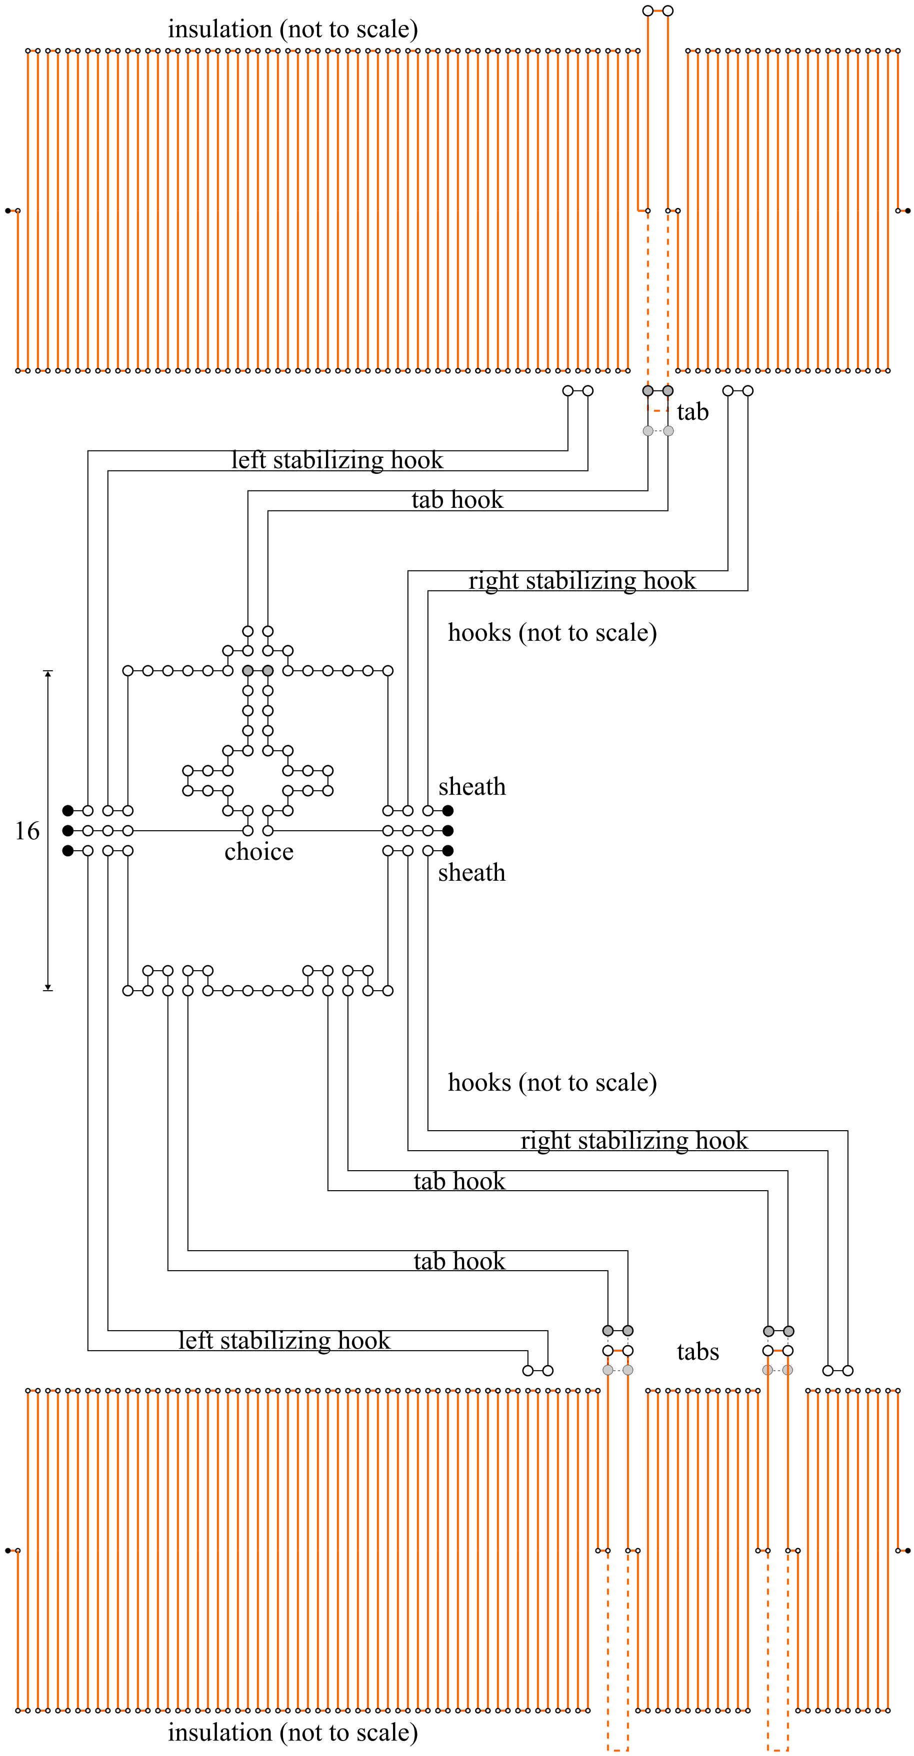

Figure 12 Another version of clause gadget (three chains of black segments with black endpoints) and its interaction with the neighboring insulation gadgets (colored segments). This version has two tabs on the bottom and one option on the top, and the upper tab extended. Each tab shows the unused alternate state with dashed lines. (Some vertices are not drawn to simplify the figure.)

the choice gadget's three possible horizontal shifts of its tab to the bottom side, and add a corresponding tab hook to the bottom sheath that routes the tab to be horizontally aligned with the corresponding variable, along with a corresponding tab to the insulation gadget below. When the choice gadget places its tab in this location, the tab hook must extend (shift down by 2), forcing the insulation tab to flip down. (This will enable communication with the variable below the insulation gadget.) Similarly, for each connection in G_ϕ from this clause to a variable gadget in the row above, we assign one of the choice gadget's three possible horizontal shifts of its tab to the top side, add a corresponding hook tab to the top sheath, and add a corresponding tab to the insulation gadget above the clause.

In Figure 11, the clause is connected to three variables below, and we have extended the leftmost tab and hook; while in Figure 12, the clause is connected to one variable above and two below, and we have extended the upper tab and hook. These two variations and their reflections through a horizontal line are all the cases we need for clauses, as every clause either has all three connections on one side or splits its connections into a group of one and a group of two. (We could additionally re-assign the horizontal shifts of the choice gadget's tabs between up vs. down, but this flexibility is unnecessary.) In each case, the inner choice gadget has exactly three foldings: among the six foldings from Claim 8, three of them intersect a horizontal part of a sheath. For example, in Figure 11, the three upper options for the tab of the choice are prevented by not crossing the upper horizontal sheath.

To force the folding of a clause gadget, we need to specify features of the overall layout (already implicit in Figures 11 and 12). Specifically, the clause gadgets all appear in the leftmost quarter of the construction, and the variable gadgets all appear in the rightmost quarter. Thus the hooks' nonunit horizontal segments have length more than half the width of the construction.

We also constrain the design of hooks as follows. First, we require the first and last extreme vertical segments to be at least 50 longer than the two other vertical segments, so that the hook's initial vertical travel is significantly longer than the second vertical travel. Finally, we require that all vertical segments of a hook have length at least $\ell_{\min} = \max\{50, 16m + 21\}$. (Recall that m is the number of clauses.)

Here we show that each hook gadget should follow the desired configuration as shown in Figures 11 and 12. In order to deal with the desired configuration, we suppose a bounding box of the hook, which is a rectangular region that exactly encompasses a given configuration of the hook. The bounding boxes of two or more hook gadgets can overlap, however, we show that it will cause trouble if a hook strays from the desired configuration.

▷ Claim 9. Consider a downward hook gadget with a fixed starting point in the left quarter of the construction. We assume that the intended foldings descends to 1 above a tab of the insulation gadget below, or lower. Then this hook gadget has a unique folding, provided no vertical segment can go $\ell_{\min}$ above the starting vertex.

Proof. Refer to Figure 13, which enumerates all possible foldings of a downward hook gadget up to the first intersection with a hypothetical rectangular bounding box. For example, folding (a) considers when the initial vertical segment goes up instead of down, while all other foldings consider when it goes down; folding (b) considers when the second segment goes left instead of right; foldings (c–h) consider when the third segment goes up instead of down; foldings (i–j) consider when the fourth segment goes left instead of right; folding (k) considers when the fifth segment goes down instead of up; folding (l) considers when the sixth segment goes right instead of left; and folding (m) considers when the seventh segment goes down instead of up.

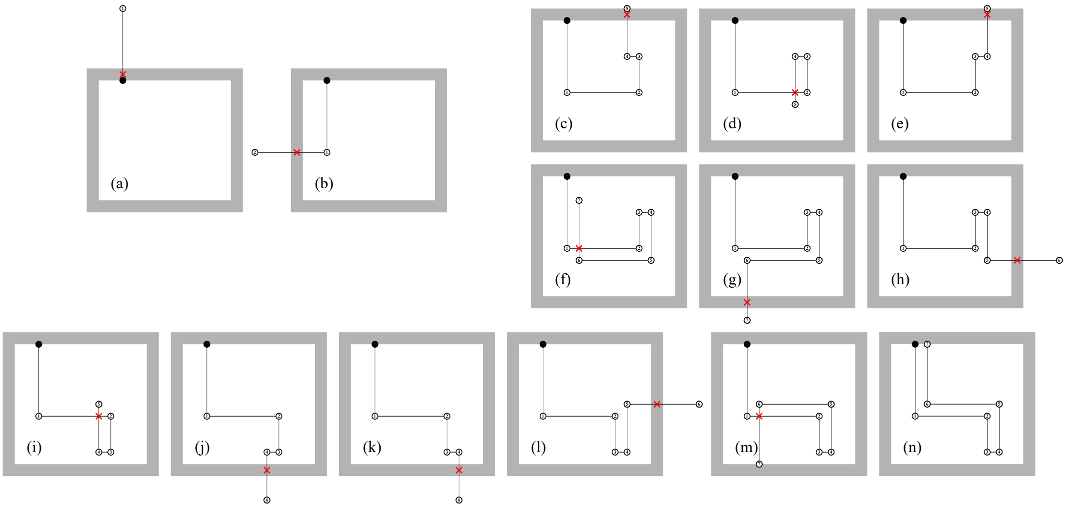

Figure 13 A hook gadget can fold in only one way (n) if it is constrained to lie within a rectangle (gray).

In each case other than (n), we argue an impossibility as follows. Foldings (d), (f), (i), and (m) have local intersections among segments starting at two vertices of distance 1 from the intersection (because every nonunit segment of the hook gadget has another segment of length ±1 longer). Foldings (a), (c), and (e) go at least $\ell_{\min}$ above the starting vertex, because every vertical segment of the hook gadget is at least $\ell_{\min}$ long. Folding (b) crosses the left side of the frame, because the starting vertex is in the left quarter of the construction, while the second segment of the hook gadget is longer than half the width of the construction. Foldings (h) and (l) cross the right side of the frame, because the second segment going right (as well as the third) puts us in the right quarter of the construction, while the sixth segment of the hook gadget is longer than half the width of the construction. Folding (g) descends at least 50 below the intended bottom of the hook gadget (because the first and last segments are at least 50 longer than the other vertical segments), while foldings (j) and (k) descend at least $\ell_{\min} \geq 50$ below (because the fifth segment has length at least $\ell_{\min}$). In any of these three cases, the hook gadget penetrates the insulation gadget below by at least 50. If it is horizontally aligned with a grille, we get immediate intersection. If it is horizontally aligned with a tab, we get crossings within the next three segments after the hook: (1) if the hook is a tab hook, then the next two segments have length 1, and the turn immediately after crosses into a horizontally adjacent grille; (2) if the hook is a left stabilizing hook, then the next segments have lengths 1 and 7 respectively, so again the following turn immediately crosses into a horizontally adjacent grille; and (3) if the hook is a right stabilizing hook, then the next segment has length > 1, so it immediately crosses into a horizontally adjacent grille. Thus the only remaining folding is the intended folding (n). ◁

We argue that the clause gadgets must be folded as intended in Figures 11 and 12, from left to right within the three rows of gadgets they occupy. The leftmost endpoints of the rows are forced by Claim 6, and the next claim allows us to induct through all the clause gadgets.

▷ Claim 10. Assume the left three black vertices of a clause gadget are placed as in the intended folding (shown in Figures 11 and 12) and the three chains start rightward. Then any folding must place the right three black vertices as in the intended folding; and at least one tab hook must be extended into the corresponding tab of the insulation gadget.

Proof. Define the $y = 0$ line to be the horizontal line through the left endpoint of the clause gadget's choice chain, or equivalently, the left endpoint of the entire row's choice chain.

First we show that no vertical segment of a bottom hook gadget can go up to $y = 16m$, which is useful for applying Claim 9. If it did, then the entire row's choice chain would be unable to connect its two endpoints (where the choice chain connects to the spiral): the total length of its vertical segments is exactly $16m$, so it can reach only $16m$ above the origin, so it could never reach above the wayward vertical segment.

Consider the bottom sheath from left to right. The left stabilizing hook starts at $y = -1$, so Claim 9 guarantees the intended folding, because no vertical segment can go up to $y = -1 + \ell_{\min} \geq 16m$. The next segment (which has length 1) must go right to avoid intersecting the first segment of the sheath. The next segment (which has length 7) must go down, to $y = -8$, to avoid intersecting the choice chain. The next segment must go right to avoid intersecting the first hook.

Now consider the zero, one, two, or three tab hooks on this sheath, one at a time from left to right. Each tab hook is prefixed by a vertical segment of length 1, which can go up or down according to whether the second hook is retracted or extended, and a horizontal segment of length 1, which we will argue must go right. The first tab hook starts at $y = -7$ or $y = -9$, so Claim 9 guarantees the intended folding, because no vertical segment can go up to $y = -9 + \ell_{\min} \geq 16m$. This intended folding implies that the prefix horizontal segment goes right; if it instead went left, then its right endpoint would be where the tab hook gadget ends, causing an intersection. The tab hook is suffixed by a horizontal segment of length 1, which must go right to avoid immediately intersecting the hook, and a vertical segment of length 1, which we will argue must go the opposite direction of the prefix vertical segment. But even if both vertical segments surrounding tab hooks all go down, we reach $y = -10$ after the first tab hook, $y = -12$ after the second tab hook, and $y = -14$ after the third tab hook (if they exist). All of the tab hooks start at $y \geq -12$, so Claim 9 guarantees the intended folding, because no vertical segment can go up to $y = -12 + \ell_{\min} \geq 16m$.

After the hooks, we have a horizontal segment, which must go right to avoid intersecting the previous hook (the left stabilizing hook if there are no bottom tab hooks); a vertical segment of length 7, which we will argue goes up; and a horizontal segment of length 1, which we will argue goes right. Thus we must be at $y \geq -21$ when we reach the right stabilizing hook. Claim 9 guarantees the intended folding, because no vertical segment can go up to $y = -21 + \ell_{\min} \geq 16m$. This intended folding implies that the preceding horizontal segment of length 1 goes right; if it instead went left, then its right endpoint would be where the stabilizing hook gadget ends, causing an intersection. Finally we have a horizontal segment of length 1, which must go right to avoid immediately intersecting the right stabilizing hook.

At this point, we have guaranteed that all bottom hook gadgets fold as intended, and all other horizontal segments of the sheath go right. Thus, we have determined the x coordinates of the entire bottom sheath to be as in the intended foldings of Figures 11 and 12. In particular, the tip of the right stabilizing hook is (as in the intended folding) horizontally aligned with a grille of the insulation gadget below, so to avoid collision the y coordinate must be strictly above the insulation gadget, i.e., no lower than in the intended folding. This guarantees that the starting y coordinate for the right stabilizing hook is no lower than in the intended folding, i.e., at $y \geq -1$. By applying the same argument to the top sheath, its final endpoint is at $y \leq 1$. To leave room for the choice chain to exit on the right (necessary to reach the right endpoint of its row where it meets the spiral), the bottom and top sheaths must in fact end at exactly $y = -1$ and $y = 1$, respectively.

Now we analyze the starting y coordinate for the bottom tab hooks. We have already argued that the leftmost tab hook starts at $y = -7$ (retracted) or $y = -9$ (extended). Now that we have fixed the folding of the right stabilizing hook, we determine the vertical segment of length 7 that precedes it; in particular, it must go up to reach $y = -1$. Thus the rightmost tab hook must start at $y = -7$ or $y = -9$. The remaining case is when there are three bottom tab hooks. If the leftmost and rightmost tab hooks are both retracted ($y = -7$) or extended ($y = -9$), then the middle tab hook could reach $y = -5$ or $y = -11$, respectively. The first case $y = -5$ is impossible, because then all three tab hooks are retracted, so the choice gadget must intersect by Claim 8. The second case $y = -11$ acts the same as $y = -9$: it is still extended into the corresponding tab of the insulation gadget below.

A symmetric argument determines the folding of the top sheath.

Finally we analyze the choice chain. Because there is only a single y coordinate of space between the top and bottom sheaths on the right, namely $y = 0$, the final horizontal segment of length 9 must be at $y = 0$. To avoid intersecting the first horizontal segment of length 9, the x coordinate of the right endpoint of the choice chain must be at or right of where it is in the intended folding (19 right of the left endpoint). Now consider the entire row's choice chain, whose overall width is determined by the spiral gadget. The overall width is the sum of the widths of the individual clauses' choice chains. If any individual clause's choice chain were longer than distance 20 used by the intended folding, then by conservation of the sum, some other clause's choice chain would have to be shorter, which we have argued is impossible. Therefore every clause's choice chain has a width of exactly 20. Thus the endpoints of the choice gadget must be horizontally adjacent as in the intended folding, so Claim 8 guarantees that at least one tab is extended. ◁

3.6 Variable gadget

Figure 14 illustrates the ***variable gadget***. For a variable v that appears in k clauses (as positive literal v or negative literal $\neg v$), the variable gadget consists of two ***zig-zag paths*** occupying a (dotted) rectangle of points of width $3k + 1$ and height 3. The two zig-zag paths are joined by a vertical length-3 ***cap***, followed by a horizontal ***baseline*** of length $3k + 4$, followed by another vertical length-3 cap. In the intended folding, the baseline is in the middle of the available space (aligned with the black endpoints), separating the upper and lower zig-zag paths, which are folded to be identical (so measured along the chain, they are reversals of each other). At the beginning and end of the chain, we add a ***bookend*** consisting of vertical segments of length 3, 6, and 4; horizontal segments of length 2 in between the vertical segments; a horizontal segment of length 2 incident to the black endpoint; and a horizontal segment of length 3 on the other end of the bookend. The two intended solutions of the entire variable gadget are the one shown in Figure 14 (corresponding to setting the variable to true) and its reflection through the baseline (corresponding to setting the variable to false).

Both zig-zag paths contain a horizontal segment of length 3 (connecting four vertices) for each appearance of the variable. The heights of the segments on the upper and lower zig-zag paths have two options depending on whether the corresponding literal uses the variable in its positive or negative form, and on whether the clause that uses the literal is in a row above or below this one. Each pair (h_1, h_2) of heights is measured from the baseline, and is either $h_1 = 3$ and $h_2 = -1$, or $h_1 = 1$ and $h_2 = -3$. The pair of heights is $(1, -3)$ if and only if either the literal is v and the clause is above, or the literal is $\neg v$ and the clause is below; in Figure 14, these are the second, fourth, and fifth pairs of horizontal segments (indicated by ■ in the figure). In the other cases, the heights are $(3, -1)$; in Figure 14, these are the first, third, and sixth pairs of horizontal segments (indicated by ⋄ in the figure). We add

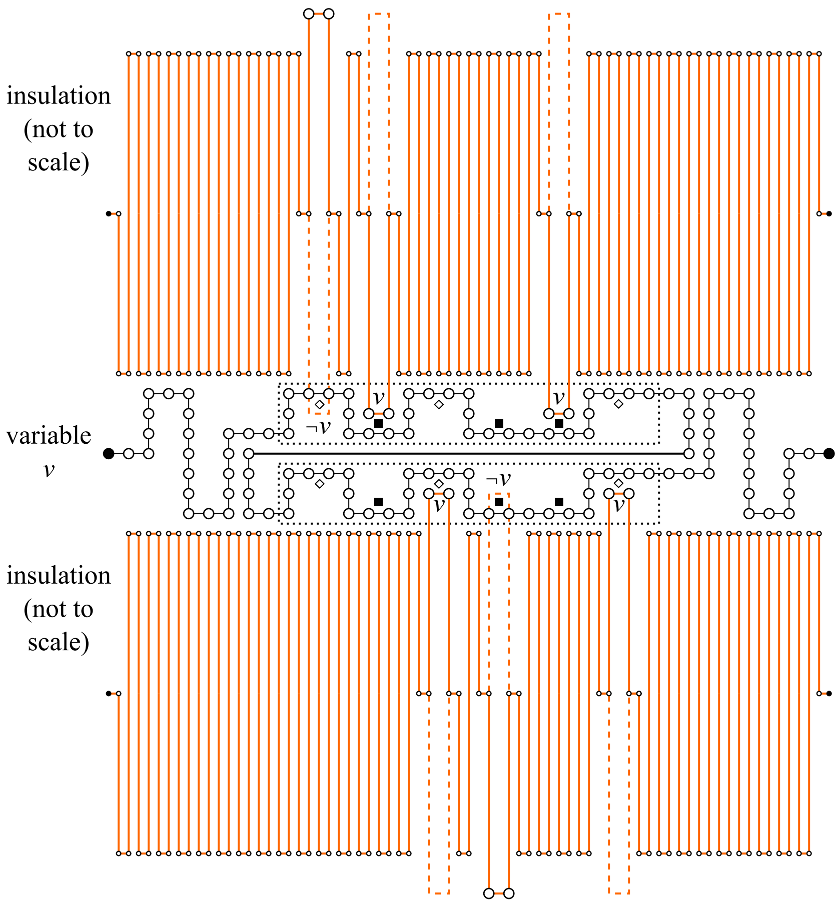

Figure 14 A variable gadget for a variable v that appears six times as v, $\neg v$, $\neg v$, v, $\neg v$, and $\neg v$ coming from tabs above, above, below, below, above, and below, respectively. The dotted areas outline the two zig-zag paths. (Some straight vertices are not drawn to simplify the figure.)

vertical segments of length 2 to transition when necessary between heights 1 and 3 on the upper zig-zag path, and between heights -1 and -3 on the lower zig-zag path. Notably, we do not add vertical segments between length-3 horizontal segments of the matching height; in Figure 14, this occurs between the fourth and fifth pairs of horizontal segments.

For each occurrence of the variable in a clause in a row above, we place a corresponding tab of the insulation gadget immediately above the variable gadget. This tab is horizontally aligned with the middle two vertices of the length-3 horizontal segment of the upper zig-zag path. When the tab is down (corresponding to a clause choosing this variable to satisfy it), it comes down to height 2, which forces the variable to flip so that this horizontal segment has height 1. Similarly, for each occurrence of the variable in a clause in a row below, we place a corresponding tab of the insulation gadget immediately below that is horizontally aligned with the middle two vertices of the length-3 horizontal segment, and which comes up to height -2 when up. This vertical alignment of the insulation gadgets places the grilles of the insulation gadgets above and below at heights 4 and -4 respectively.

▷ Claim 11. For each variable gadget in a row, the only valid foldings are the one discussed above and its reflection through the baseline as shown in Figure 14.

Proof. First we analyze the seven-segment bookend at the beginning and end of each variable gadget's chain. None of the vertical segments could stick into the width-2 gap of a tab in an insulation gadget, because then the incident horizontal segment of length at least 2 would intersect a grille. Thus these vertical segments lie within the height range ± 3 (where height 0 denotes the middle of the space between insulation grilles), so the middle vertical segment of length 6 occupies the full height range, and the first and last vertical segments (of length 3) must go in opposite directions as the middle vertical segment. In particular, the endpoints of the variable gadget must be at height 0. Furthermore, the four horizontal segments of a bookend must all go in the same direction to avoid local intersection among the vertical segments, and this direction must be right to enable eventual connection from the left endpoint of the row to the right endpoint of the row (otherwise the length-6 vertical segment would cut them off).

Now we make a global argument about the variable gadgets in the row, and select a specific variable gadget to consider. Define the ***width*** of a variable gadget to be the signed horizontal distance between its endpoints, which must be positive because of the length-6 vertical segments serving as left-to-right barriers. The total width of the variable gadgets is fixed by the spiral gadget and Claim 6. If any variable gadget had width greater than the width of the intended foldings, then the width of some other variable gadget would be less than intended. Consider a variable gadget whose width is less or equal to the intended width.

We claim that the left endpoint of the baseline is at least 1 unit right of the last vertical segment (of length 4) of the left bookend. First, it must be right of the middle vertical segment (of length 6) of the left bookend, so there are only two intervening x coordinates to consider. In either case, the baseline cannot have height within the height range of the last vertical segment (of length 4) of the left bookend. This leaves just two possible heights: ± 2 and ± 3 (where $\pm$ depends on whether the left bookend is reflected). Assume by symmetry that the left bookend is folded as in Figure 14, so that the baseline has height 2 or 3 if not 0. We use that the baseline has, on both endpoints and hence on the left endpoint, an incident cap segment of length 3. This cap segment must go down from the baseline: going up would intersect a grille, or enter a tab gap which would then cause an intersection because the next horizontal segment has length at least 2. To avoid intersecting the final vertical segment of the left bookend, this forces the left endpoint of the baseline to be in the column in between the middle segment (of length 6) and final segment (of length 4) of the left bookend. This folding wedges the cap segment in between the middle and last vertical segments of the left bookend, so the next horizontal edge intersects one of those vertical segments.

By a symmetric argument, the right endpoint of the baseline is at least 1 unit left of the first vertical segment (of length 4) of the right bookend. These two bounds on the horizontal location of the baseline's endpoints contradict each other if the width of the variable gadget is smaller than needed. Therefore this variable gadget, and all variable gadgets, have the intended width. Furthermore, by the bounds on the endpoints, the baseline's horizontal location is exactly as in the intended foldings.

Next we argue that the baseline has an even height: -2, 0, or 2. We have already fixed the height of the right endpoint of the left bookend to be 1, which is odd. The vertical segments of the intervening zig-zag path are all length 2, which is even. The only odd-length vertical segment between the left bookend and the baseline is the length-3 segment incident to the baseline. Thus the height of the baseline is even.

Now we show that the baseline cannot have height ± 2, leaving only height 0. If the baseline had height -2, then the length-3 cap segment incident to the left endpoint must go up (to avoid the insulation below, as argued earlier) to height 1. Similarly, if the baseline had height 2, then the length-3 cap segment incident to the left endpoint must go down to height -1. In either case, the cap segment intersects the final horizontal segment of the left bookend, which has height ± 1.

At this point, we have determined the horizontal and vertical location of the baseline to be as in the intended foldings. We have also determined the foldings of the left and right bookends, up to reflection through the baseline. The baseline and bookends thus decompose the available space into two disjoint regions, roughly above and below the baseline. The first zig-zag path (attached to the left bookend) must start by going right (to avoid intersecting the middle vertical segment of the left bookend), then go up (to avoid intersecting the baseline), then go right (to avoid cutting off connectivity to the right endpoint), then go down (to avoid intersecting the insulation, which would cause an intersection because the horizontal segments all have length at least 4), and so on. The second zig-zag path is similarly forced, which forces the right bookend to be reflected opposite from the left bookend (as in Figure 14). Therefore we have determined the entire folding of the variable gadget, up to the reflection of the left bookend, which reflects the entire folding through the baseline. ◀

One issue can arise when connecting clause gadgets to variable gadgets: when multiple clauses connect via tab hooks to the same side of a variable gadget, they also need to place two stabilizing hooks against a grille at the transition point. Given that the horizontal alignment of insulation tabs is controlled by the variable gadget, we need to leave horizontal room for such stabilizing hooks. We can do so by adding a "null" occurrence of the variable that is not used in any clauses, and has no corresponding stabilizing hook; instead, it is surrounded by grilles. These null occurrences can have zig-zag path heights of 3 and -1 respectively, or 1 and -3 respectively; it does not matter. For example, in Figure 5, this modification occurs at the bottom of v_3 and at the top of v_3'.

3.7 Putting gadgets together

Figure 5 shows how all the gadgets fit together for an example instance. We join together all upper halves of hook gadgets for $c_1, c_2, \ldots, c_m$; all clause gadgets (and their flaps) for $c_m, c_{m-1}, \ldots, c_1$; all lower halves of the hook gadgets for $c_1, c_2, \ldots, c_m$; and all variable gadgets for $v_1, v_2, \ldots, v_n$, in these orders. Finally, we apply the frame gadget to this closed chain and the intended bounding box, with an edge on a path joining the upper halves of the hook gadgets, or an edge on a path joining the variable gadgets.

Here we claim that we always have such an edge visible from the outside. When the Hamiltonian cycle κ has an edge visible from the outside, we can use any one of them. (In Figure 1, the edges $\{c_1, c_2\}$ and $\{v_1, v_2\}$ are visible.) Otherwise, we have a clause vertex $c \in C$ on the outer boundary. (When κ has no edge visible from the outside, these edges are surrounded by edges joining a clause vertex and a variable vertex. Considering the outmost edge joining a clause vertex c and a variable vertex v, the clause vertex c should be visible from the outside.) Then we can take an edge from the upper half of the hook gadget for the clause c. This completes the construction.

3.8 Correctness

Now we conclude the proof of Theorem 3. This reduction can be done in time polynomial in the size of ϕ: we apply the Pilz reduction to draw ϕ in the spiral fashion with connections only between adjacent rows, push the clauses to the left of each row and push the variables to the right of each row, and route the hook connections in a planar fashion — with the horizontal part of all hooks in a row stacking up vertically, while leaving enough vertical space between them to allow for hooks to freely extend or retract. It remains to show that an instance ϕ of the planar 3SAT is satisfiable if and only if the resulting fixed-angle orthogonal equilateral closed chain has a planar embedding.

When the linked planar 3SAT instance is satisfiable, at least one literal of each clause is satisfied by the assignment. The clause gadget then chooses the corresponding tab location for the choice gadget, and extends the corresponding hook gadget, while retracting the other tabs. The adjacent insulation gadget then flips its corresponding tab to avoid overlap. Reflecting the variable gadget into the assignment corresponding to the literal means that this insulation tab will not intersect the variable gadget. This folding avoids all intersections because the assignment is satisfied.

On the other hand, when the loop has an embedding, the frame gadget folds into the intended bounding box by Claim 5, and the spiral gadget folds as intended by Claim 6. Each insulation gadget folds as intended by Claim 7. Each row of clause gadgets folds as intended by Claim 10, and each row of variable gadgets folds as intended by Claim 11. In particular, by Claim 10, at least one tab hook from each clause must be extended, which forces the corresponding insulation tab to be folded into the corresponding variable gadget. This insulation tab in turn forces the variable to have a satisfying assignment. Once we have an embedding, each clause requires a variable to be true or false, and these requirements on a variable is consistent by the construction of the gadgets. Thus we give an assignment of the variable. When there is a variable that receives no requirement from any clause, we give an arbitrary assignment to the variable. It is easy to see that this is a truth assignment to the instance. Therefore, the instance of planar 3SAT is satisfiable.

4 HP optimal folding a fixed-angle orthogonal equilateral open chain is strongly NP-complete

We now turn to orthogonal equilateral open chains in the HP model, where the vertices are bicolored H or P, and we wish to find a noncrossing configuration in 2D that maximizes the number of H–H contacts. In this section, we prove that this problem is NP-complete, despite the chain being open:

▶ **Theorem 12.** *HP optimal folding of a bicolored fixed-angle orthogonal equilateral open chain is strongly NP-complete, even if the chain has just two H vertices.*

Proof. We use the same reduction in the proof of Theorem 3, except for the frame gadget, which we replace with Figure 15. The differences are that the bottom doubled segment extends very far to the left — more than 10 times the total length L of the given scaled chain $5C$ — and the chain is no longer closed at the left end of the bottom doubled segment. The leftmost two vertices of the bottom doubled segment (the endpoints of the chain) are H, while all other vertices in the chain are P.

This reduction can be done in polynomial time. Thus it suffices to show that this arrangement of the frame is the only way to obtain the H–H contact at the two H vertices. Because the total length of the given construction inside of the frame is at most L, and the length of each segment of the frame is therefore at most L, the total length of the chain except the bottom doubled segment is at most $9L$. Hence to make the H–H contact between the two H vertices, the two long segments attached to the H vertices must be arranged in parallel as shown in Figure 15: if the two long segments went in opposite directions, or were perpendicular to each other, then the rest of the chain would not be long enough to connect their ends together. Thus the frame construction is forced to act like the closed chain of Figure 7, so it has a unique folding up to isometry by the rest of the proof of Lemma 4. ◀

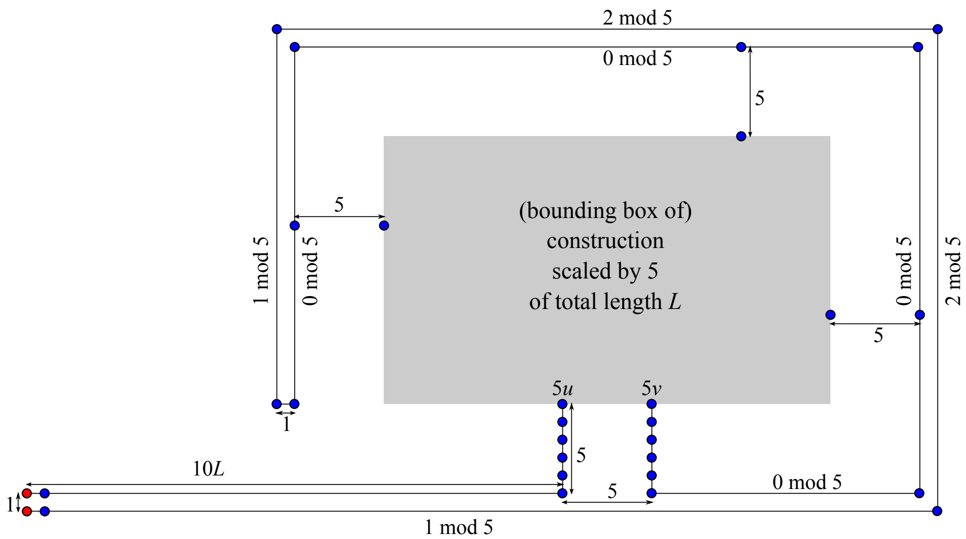

Figure 15 A frame gadget for an HP chain. The two H vertices are drawn red at the far left.

5 Packing fixed-angle orthogonal equilateral open chains into squares is strongly NP-complete

We now address some of the open questions from [1]. First, the authors ask whether a fixed-angle orthogonal equilateral open chain (or in their terminology, an S–T sequence of squares, where each S square must continue straight and each T square must turn left or right) can be packed into a 2D square. Second, they ask whether the problem remains hard when the chain occupies a small fraction of the volume of the target shape. (They ask this question for the 3D version of the problem, but it naturally extends to the 2D version we consider.) We answer both questions by showing that packing a fixed-angle orthogonal equilateral open chain of length $O(s)$ into an $s \times s$ square is strongly NP-complete. This result is tight up to constant factors: if the chain has length $< s$, then it can be packed into an $s \times s$ square via Observation 1.

▶ **Theorem 13.** *Embedding a given fixed-angle orthogonal equilateral open chain into an $s \times s$ square is strongly NP-complete, even if the chain has length $O(s)$.*

Proof. We use the same reduction in the proof of Theorem 12, except for the frame gadget, which we replace with Figure 16.

This frame gadget starts the chain with two straight segments of length $s = 10L + 1$. Any embedding into the $s \times s$ square must place these segments along two boundary edges of the square, say left and bottom as in the figure. The next two segments on the outside of the frame gadget must turn left to remain within the square. At the other end of the chain, we have a vertical (by parity) segment of length $s - 1$, which forces its endpoints to be at the very top and one position up from the bottom (to avoid overlapping the second segment).

Now we make a congruence argument modulo 5. Because the third segment goes up by 2 modulo 5, and the only other vertical travel modulo 5 is by the fifth segment which goes up or down by ± 1, the fifth segment must in fact go down by 1 modulo 5 to reach the position one up from the bottom.

Because the next-to-last segment has length $> 9L$, its right end must be inside the frame. Thus the sixth segment of length 1 must go right to stay inside the frame; otherwise, it could

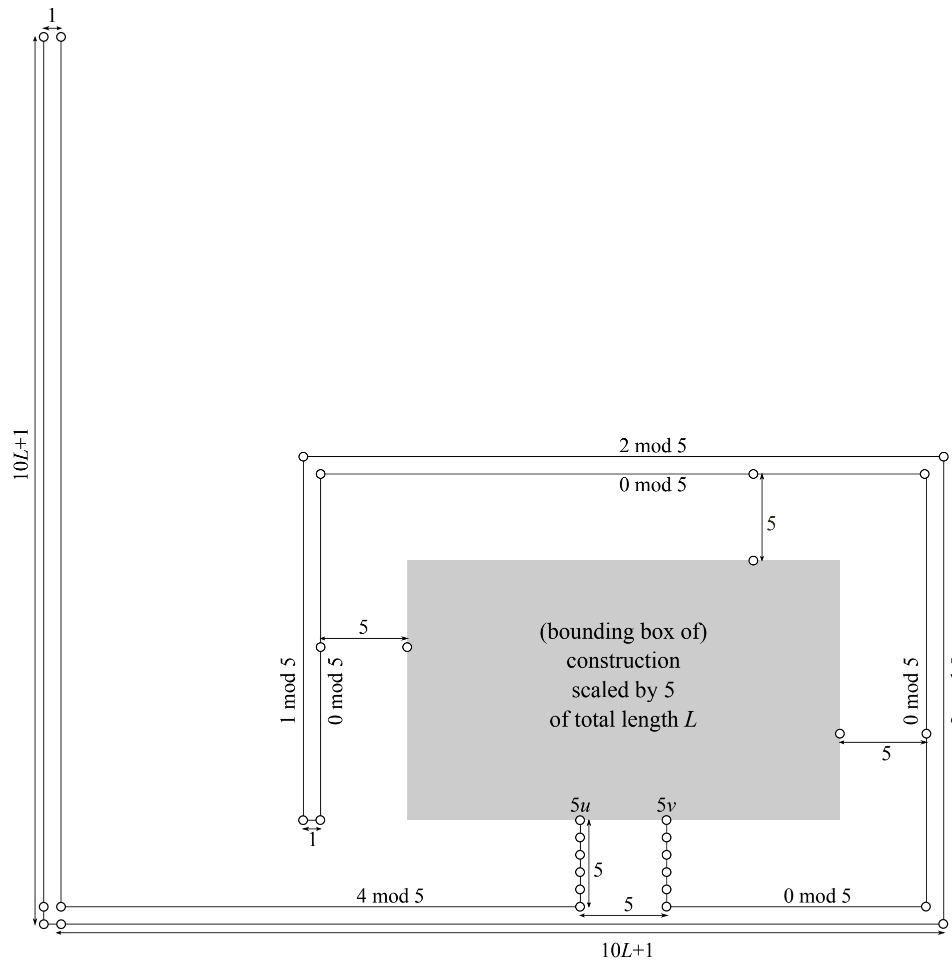

Figure 16 A frame gadget for an open chain which must fit in a $10L + 1$ by $10L + 1$ square.

never connect to the right end of the next-to-last segment. The rest of the segments are then forced to turn as in the figure in order to avoid collisions.

The chain has length at most $48L$ (from the given chain of length L, the smaller frame, and the three long bars). Thus the length is $O(s)$. ◀

It remains open whether the problem of ***densely*** packing a fixed-angle orthogonal equilateral open chain of length s^2 into an $s \times s$ square is NP-complete, meaning that the chain covers *all* grid points in the square. The analogous problem in 3D is strongly NP-complete [1].

References

1 Zachary Abel, Erik D. Demaine, Martin L. Demaine, Sarah Eisenstat, Jayson Lynch, and Tao B. Schardl. Finding a Hamiltonian path in a cube with specified turns is hard. *Journal of Information Processing*, 21(3):368–377, 2013.

2 Zachary Abel, Erik D. Demaine, Martin L. Demaine, Sarah Eisenstat, Jayson Lynch, and Tao B. Schardl. Who needs crossings? Hardness of plane graph rigidity. In *Proceedings of the 32nd International Symposium on Computational Geometry (SoCG 2016)*, pages 3:1–3:15, Boston, Massachusetts, June 2016.

3 Bonnie Berger and Tom Leighton. Protein folding in the hydrophobic-hydrophilic (*HP*) model is NP-complete. *Journal of Computational Biology*, 5(1):27–40, 1998. URL: `https://doi.org/10.1089/cmb.1998.5.27`.

4 Pierluigi Crescenzi, Deborah Goldman, Christos Papadimitriou, Antonio Piccolboni, and Mihalis Yannakakis. On the complexity of protein folding. *Journal of Computational Biology*, 5(3):423–465, 1998. URL: `https://doi.org/10.1089/cmb.1998.5.423`.

5 Erik D. Demaine, Martin L. Demaine, Adam Hesterberg, Quanquan Liu, Ron Taylor, and Ryuhei Uehara. Tangled tangles. In *The Mathematics of Various Entertaining Subjects (MOVES 2015)*, volume 2, pages 141–152. Princeton University Press, 2017.

6 Erik D. Demaine and Sarah Eisenstat. Flattening fixed-angle chains is strongly NP-hard. In *Proceedings of the 12th Algorithms and Data Structures Symposium (WADS 2011)*, pages 314–325, Brooklyn, New York, August 2011.

7 Erik D. Demaine and Joseph O'Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, 2007.

8 Alexander Pilz. Planar 3-SAT with a Clause/Variable Cycle. In *16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018)*, pages 31:1–31:13. LIPIcs, Dagstuhl, 2018.

9 Ashim Garg. New results on drawing angle graphs. *Computational Geometry*, 9(1):43–82, 1998. Special Issue on Geometric Representations of Graphs. URL: `https://www.sciencedirect.com/science/article/pii/S0925772197000163`, `doi:10.1016/S0925-7721(97)00016-3`.

10 Feng Tianfeng, Giovanni Viglietta, and Ryuhei Uehara. Bicolored path embedding problems in protein folding models. *IEICE Trans.*, E105-D(3):623–633, 2022.

11 Giuseppe Di Battista and Luca Vismara. Angles of planar triangular graphs. *SIAM Journal on Discrete Mathematics*, 9(3):349–359, 1996.

12 Susan Landau and Neil Immerman. Embedding linkages on an integer lattice. *Algorithmica*, 32(3):423–436, 2002.

13 Vijayan, V. Geometry of planar graphs with angles. In *Proceedings of the second annual symposium on Computational geometry*, pages 116–124, 1986.

14 Zachary R. Abel, Erik D. Demaine, Martin L. Demaine, David Eppstein, Anna Lubiw, and Ryuhei Uehara. Flat foldings of plane graphs with prescribed angles and edge lengths. *Journal of Computational Geometry*, 9(1):71–91, 2018.