Box Pleating is Hard

Hugo A. Akitaya[1], Kenneth C. Cheung[2], Erik D. Demaine[3],
Takashi Horiyama[4], Thomas C. Hull[5], Jason S. Ku[3],
Tomohiro Tachi[6], and Ryuhei Uehara[7]

[1] Tufts University hugo.alves_akitaya@tufts.edu
[2] NASA kenneth.c.cheung@nasa.gov
[3] MIT {edemaine,jasonku}@mit.edu
[4] Saitama University horiyama@al.ics.saitama-u.ac.jp
[5] Western New England Univ. thull@wne.edu
[6] The University of Tokyo tachi@idea.c.u-tokyo.ac.jp
[7] JAIST uehara@jaist.ac.jp

Abstract. Flat foldability of general crease patterns was first claimed to be hard for over twenty years. In this paper we prove that deciding flat foldability remains NP-complete even for box pleating, where creases form a subset of a square grid with diagonals. In addition, we provide new terminology to implicitly represent the global layer order of a flat folding, and present a new planar reduction framework for grid-aligned gadgets.

1 Introduction

In their seminal 1996 paper, Bern and Hayes initiated investigation into the computational complexity of origami [BH96]. They claimed that it is NP-hard to determine whether a given general crease pattern can be folded flat, both when the creases have or have not been assigned crease directions (mountain fold or valley fold). Since that time, there has been considerable work in analyzing the computational complexity of other origami related problems. For example, Arkin et al. [ABD+04] proved that deciding foldability is hard even for simple folds, while Demaine et al. [DFL10] proved that optimal circle packing for origami design is also hard.

While the gadgets in the hardness proof presented in [BH96] for unassigned crease patterns are relatively straightforward, their gadgets for assigned crease patterns are considerably more convoluted, and quite difficult to check. In fact, we have found an error in even their unassigned crossover gadget where signals are not guaranteed to transmit correctly for wires that do not cross orthogonally, which is required in their construction. Part of the reason no one found this error until now is that there was no formal framework in which to prove statements about flat-folded states. We attempt to provide such a framework.

At the end of their paper, Bern and Hayes pose some interesting open questions to further their work. While most of them have been investigated since, two in particular (problems 2 and 3) have remained untouched until now. First, is there a simpler way to achieve a proof for assigned crease patterns (i.e. "without tabs")? Second, their reductions construct creases at a variety of unconstrained angles. Is deciding flat foldability easy under more restrictive inputs? For example, *box pleating* involves folding creases only along on a subset of a square grid and the diagonals of the squares, a special case of particular interest in transformational robotics and self-assembly, with a universality result constructing arbitrary polycubes using box pleating [BDDO10].

In this paper we address both these questions. We prove that deciding flat foldability of box-pleated crease patterns is NP-hard in both the unassigned and assigned cases, using relatively simple gadgets containing no more than 25 layers at any point.

2 Definitions

In general, we are guided by the terminology laid out in [DO07] and [Rob77]. An *isometric flat folding* of a paper P is a function $f : P \to \mathbb{R}^2$ such that if γ is a piecewise-geodesic curve on P parameterized with respect to arc-length, then $f(\gamma)$ is also a piecewise-geodesic curve parameterized with respect to arc-length. It is not hard to show that under these conditions f must be continuous and non-expansive. Let X_f be the boundary of a paper P together with the set of points not differentiable under f. Then one can prove that X_f is a straight-line graph embedded in the paper [Rob77], with vertex set V_f and edge set C_f, the *creases* of our folding f. A vertex or crease in V_f or C_f is *external* if it contains a boundary point of P, and *internal* otherwise. Subtracting X_f from P results in a disconnected set of open polygons F_f we call *faces*. For any face $F \in F_f$, $f(F)$ is either an isotopic transformation in $\mathbb{R}^2$, or the transformation involves a reflection and is anisotopic. Define $u_f : P \setminus X_f \to \{-1, 1\}$ such that $u_f(p) = -1$ if the face containing p is reflected under f and $u_f(p) = 1$ otherwise. We call $u_f(p)$ the *orientation* of the face containing p. Every point in P is in exactly one of V_f, C_f, or F_f. We call this partition of P the *isometrically flat foldable crease pattern* $\Sigma_f = (V_f, C_f, F_f)$ induced by f. We call a folding *box pleating* if every vertex lies on two dimensional integer lattice, and the creases are aligned at multiples of $45°$ to each other.

We say two disjoint simply connected subsets of P are *adjacent* to each other if their closures intersect; we call such an intersection the *adjacency* of the adjacent subsets. We say a simply connected subset of P is *uncreased* under f if f is injective when restricted to the subset. We say two simply connected subsets of P *overlap* under f if the interiors of their images under f intersect. We say two simply connected subsets of P *strictly overlap* under f if their images under f exactly coincide. It is known that the set of creases adjacent to an internal vertex of a crease pattern obey the so called Kawasaki-Justin Theorem: the alternating sum of angles between consecutive creases when cyclically ordered around the vertex equals zero [DO07]. This condition turns out to be necessary sufficient: given a paper P exhaustively partitioned into a set of isolated points V, open line segments C, and open disks F such that every point in V is adjacent to more than two segments in C, then (V, C, F) is an isometrically flat foldable crease pattern induced by a unique isometric flat folding if and only if (V, C, F) obeys the Kawasaki-Justin Theorem.

Let a function $\lambda_f : P \times P \to \{-1, 1\}$ be a *global layer ordering* of an isometric flat folding f if it obeys the following six properties.

Existence: λ_f satisfies *existence* if $\lambda_f(p, q)$ is defined for every distinct pair of points p and q that strictly overlap under f and at least one of p or q is not in X_f; otherwise $\lambda_f(p, q)$ is undefined. Informally, order is only defined between a point on a face and another point overlapping it in the folding.

Antisymmetry: λ_f is *antisymmetric* if $\lambda_f(p, q) = -\lambda_f(q, p)$, where λ_f is defined. Informally, if p is above q, then q is below p.

Transitivity: λ_f is *transitive* if $\lambda_f(p, q) = \lambda_f(q, r)$ implies $\lambda_f(p, r) = \lambda_f(p, q)$, where λ_f is defined. Informally, if q is above p and r is above q, then r is above p.

Consistency (Tortilla-Tortilla Property): For any two uncreased simply connected subsets O_1 and O_2 of P that strictly overlap under f, λ_f is *consistent* if $\lambda_f(p_1, p_2)$ has the same value for all $(p_1, p_2) \in O_1 \times O_2$, where λ_f is defined. See Figure 1. Informally, if two regions completely overlap in the folding, one must be entirely above the other.

Face-Crease Non-crossing (Taco-Tortilla Property): For any three uncreased simply connected subsets O_1, O_2, and O_3 of P such that O_1 and O_3 are adjacent and strictly overlap, and O_2 overlaps the adjacency between O_1 and O_3 under f, λ_f is *face-crease non-crossing* if $\lambda_f(p_1, p_2) = -\lambda_f(p_2, p_3)$ for any points $(p_1, p_2, p_3) \in O_1 \times O_2 \times O_3$, where λ_f is defined. See Figure 1. Informally, if a region overlaps a nonadjacent internal crease, the region cannot be between the regions adjacent to the crease.

Crease-Crease Non-crossing (Taco-Taco Property): For any two adjacent pairs of uncreased simply connected subsets (O_1, O_2) and (O_3, O_4) of P such that every pair of subsets strictly overlap and the adjacency of O_1 and O_2 strictly overlaps the adjacency of O_3 and O_4 under f, λ_f is *crease-crease non-crossing* if either $\{\lambda_f(p_1, p_3),\ \lambda_f(p_1, p_4),\ \lambda_f(p_2, p_3),\ \lambda_f(p_2, p_4)\}$ are all the same or half are $+1$ and half are -1, for any points $(p_1, p_2, p_3, p_4) \in O_1 \times O_2 \times O_3 \times O_4$, where λ_f is defined. See Figure 2. Informally, if two creases overlap in the folding, either the regions incident to one crease lie entirely above the regions incident to

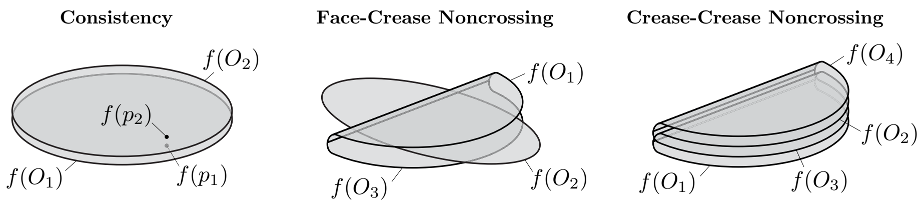

Fig. 1. Topologically different local interactions within an isometric flat folding. Forbidden configurations are shown for Face-Crease and Crease-Crease Non-Crossing.

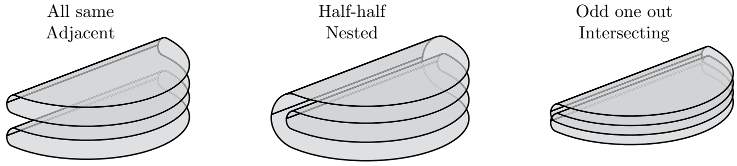

Fig. 2. Local interaction between overlapping regions around two distinct creases.

the other (all same), or the regions incident to one crease nest inside the regions incident to the other (half-half).

If there exists a global layer ordering for a given isometrically flat foldable crease pattern, we say the crease pattern is *globally flat foldable*. Consider an isometrically flat foldable crease pattern Σ_f containing two adjacent uncreased simply connected subsets O_1 and O_2 of P that strictly overlap under f, and let p and q be points in O_1 and O_2 respectively that overlap under f. O_1 and O_2 are subsets of disjoint adjacent faces of the crease pattern mutually bounding a crease. If λ_f is a global flat folding of Σ_f, then it induces a *mountain/valley assignment* $\alpha_{\lambda_f}(c) = u(p)\lambda_f(p, q)$ for each crease point c in the adjacency of O_1 and O_2. This assignment is unique by consistency. We call a crease point c a *valley fold* (V) if $\alpha_{\lambda_f}(c) = 1$ and a *mountain fold* (M) if $\alpha_{\lambda_f}(c) = -1$. In the figures, mountain folds are drawn in red while valley folds are drawn in blue. By convention, if $\lambda_f(p, q) = -1$ we say that p is *above* q, and if $\lambda_f(p, q) = 1$ we say that p is *below* q.

Given an isometrically flat foldable crease pattern Σ_f, the UNASSIGNED-FLAT-FOLDABILITY problem asks whether there exists a global layer ordering for f. Alternatively, given an isometrically flat foldable crease pattern Σ_f and an assignment $\alpha : C_f \to \{M, V\}$ mapping creases to either mountain or valley, the ASSIGNED-FLAT-FOLDABILITY problem asks whether there exists a global layer ordering for f whose induced mountain valley assignment is consistent with α.

We now prove the following implied properties of globally flat foldable crease patterns relating the layer order between points contained in multiple overlapping faces. Informally, *Pleat-Consistency* says if a face is adjacent and overlapping two larger faces, then the creases between them must have different M/V assignment, forming a pleat. *Path-Consistency* says that a face overlapping creases connecting an adjacent sequence of faces is either above or below all of them.

Lemma 1. (Pleat-Consistency) *If Σ_f is a globally flat foldable crease pattern containing disjoint uncreased simply connected subsets O_1, O_2, and O_3 of P with O_2 adjacent to both O_1 and O_3 such that O_2 strictly overlaps subsets $O'_1 \subset O_1$ and $O'_3 \subset O_3$, and the interiors of O_1 and O_3 overlap the adjacencies of O_2, O_3 and O_1, O_2 respectively, then $\lambda_f(p_1, p_2) = \lambda_f(p_2, p_3)$ for any pairwise overlapping points $(p_1, p_2, p_3) \in O_1 \times O_2 \times O_3$.*

Proof. Taco-Tortilla applied to O_3 which overlaps the adjacency of strictly overlapping sets O_2 and O'_1 implies $\lambda_f(p_2, p_3) = -\lambda_f(p_3, p_1)$. Similarly, Taco-Tortilla applied to O_1 which overlaps the adjacency of strictly overlapping sets O_3' and O_2 implies $\lambda_f(p_3, p_1) = -\lambda_f(p_1, p_2)$, so $\lambda_f(p_1, p_2) = \lambda_f(p_2, p_3)$. □

Lemma 2. (Path-Consistency) *If Σ_f is a globally flat foldable crease pattern containing uncreased simply connected subset T of P and a disjoint sequence of adjacent uncreased simply connected subsets $O_1, \ldots, O_n$ of P such that O_i strictly overlaps some subset T_i of T and the interior of O overlaps the adjacency of each pair O_i and O_{i+1} for $i = \{1, \ldots, n-1\}$, then $\lambda_f(t_j, p_j) = \lambda_f(t_k, p_k)$ for any two pairs of overlapping points $(t_j, p_j) \in T_j \times O_j$ and $(t_k, p_k) \in T_k \times O_k$ for $j, k \in \{1, \ldots, n\}$.*

Proof. If some O_i and O_{i+1} overlap, Taco-Tortilla and Consistency ensure that $\lambda_f(t_i, p_i) = \lambda_f(t_{i+1}, p_{i+1})$ for $(t_i, p_i) \in T_i \times O_i$ and $(t_{i+1}, p_{i+1}) \in T_{i+1} \times O_{i+1}$. Alternatively, O_i and O_{i+1} do not overlap and the closure of $O_i \cup O_{i+1}$ is an uncreased region for which $\lambda_f(t_i, p_i) = \lambda_f(t_{i+1}, p_{i+1})$ by consistency. Applying sequentially to each pair of faces proves the claim. □

The proofs in Section 5 and 6 contain many examples of the application of these properties. When proving the existence of a global layer ordering λ_f, it is often impractical to define λ_f between every pair of points. Frequently λ_f is uniquely induced by a M/V assignment, consistency, and transitivity. When it is not, we will provide λ_f between additional point pairs so that it will be. We present crease patterns with this implicit layer ordering information and encourage readers to fold them to reconstruct the unique layer orderings they induce.

3 Bern and Hayes and k-Layer-Flat-Foldability

Two crossover gadgets are presented in the reduction to Unassigned-Flat-Foldability provided in [BH96]. For each, they claim that the M/V assignment of the crease pair intersecting one edge of the gadget deterministically implies the M/V assignment of the crease pair on the opposite side. This claim is true for their perpendicular crossover gadget, but is unfortunately not true for the other for wires meeting at 45°. The gadget as described requires an exterior 45° angle between incoming wires that is the smallest angle at a four-crease vertex, forbidding the wires to be independently assigned by Pleat-Consistency. For completeness, we have also checked the family of possible gadgets of this form, with a rotated internal parallelogram, and no choice of rotation allows the gadget to function correctly as a crossover for the range of widths of wires that appear in the construction. Our proof to follow only uses the perpendicular crossover, avoiding this complication.

Also in [BH96], they define k-Layer-Flat-Foldability to be the same as Unassigned-Flat-Foldability or Assigned-Flat-Foldability but with the additional constraint that f maps at most k distinct points to the same point. They claim that their reduction implies hardness of Unassigned-k-Layer-Flat-Foldability for $k = 7$. But in fact their perpendicular crossover gadget requires nine points to be mapped to the same point. Our reduction uses

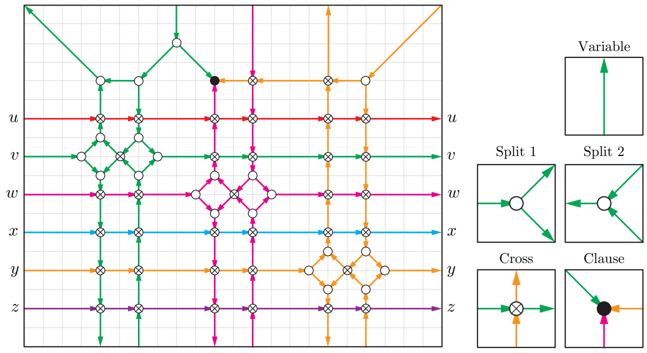

Fig. 3. SCN Gadgets. [Left] A Complex Clause Gadget constructed from the Not-All-Equal clause on variables v, w, and y of a NAE3-SAT instance on six variables. [Right] The five elemental SCN Gadgets.

the same gadget as a crossover, so we reconfirm that Unassigned-k-Layer-Flat-Foldability is NP-complete for $k \geq 9$, even for box pleated crease patterns. Also, because of the complexity of their assigned crease pattern reduction, they were unable to bound the number of layers in their reduction. We explicitly provide gadgets for the assigned case to prove Assigned-k-Layer-Flat-Foldability is NP-complete for $k \geq 25$, even for box pleated crease patterns.

4 SCN-Satisfiability

Our reductions will be from the following NP-complete problem [Sch78].

Problem 1. **(Not-All-Equal 3-SAT)** Given a collection of clauses each containing three variables, Not-All-Equal 3-SAT (NAE3-SAT)[8] asks if variables can be assigned True or False so that no clause contains variables of only one assignment.

We can construct a planar directed graph G embedded in $\mathbb{R}^2$ from an instance $\mathcal{N}$ of NAE3-SAT. For each clause, construct a Complex Clause Gadget as the one shown in Fig. 3. The motivation behind the Complex Clause Gadget is to encode the bipartite graph implicit in $\mathcal{N}$ in a planar grid embedding that can be modularly connected. Each directed edge of the Complex Clause Gadget is associated with a different variable, and we associate a different color with each variable. Some variables do not participate in the clause and simply form a straight chain of directed segments from left to right. However, the three variables participating in the clause are rerouted to intersect at the black dot. We construct a Complex Clause Gadget for each clause in the instance of NAE3-SAT and chain them together side by side, so the arrows exiting the right side of one enter the

[8] This problem is sometimes called ‘positive’ as variables cannot appear negated within clauses, however we follow the naming convention from [Sch78].

left side of another. Graph G has vertices that are adjacent to edges associated with exactly one, two, or three variables. We call these vertices *split*, *cross*, and *clause* vertices respectively. In the figures, they are labeled with white circles, crossed circles, and black circles respectively. We call such a directed graph G a *Split-Cross-Not-All-Equal* (SCN) graph.

Problem 2. **(SCN-Satisfiability)** Given a SCN graph, SCN-SATISFIABILITY asks if variables can be assigned True or False so that no clause vertex is adjacent to edges associated with variables of only one assignment.

The authors introduce SCN-Satisfiability as a useful intermediate problem because it is equivalent to NAE3-SAT but its embedding is planar, lies on a grid, and is constructed only by a small number of local elements. SCN-SATISFIABILITY is equivalent to NAE3-SAT because the bipartite graph connecting SCN variables to clause vertices is exactly the bipartite graph representing $\mathcal{N}$ by construction. However, G has useful structure for many problems. It is planar, the embedding contains edges with only four slopes, and the edges are directed meaning that a variable can be represented locally with respect to that direction. Further G is constructed from only a small number of local elements: a variable gadget, two split gadgets, a cross gadget, and a clause (simple) gadget as shown in Fig. 3. We call these the five *elemental* SCN Gadgets. If we can simulate each of these gadgets in another context, proving that edges of the same color in each gadget must all have the same value, and edges adjacent to a clause vertex do not all have equal value, we can prove other problems NP-hard. This will be our strategy in the following sections.

Theorem 1. *If a problem X can simulate the elemental SCN gadgets such that edges of the same color in each gadget have the same value and edges adjacent to a clause vertex do not all have equal value and if the correspondent gadgets in X can be connected consistently, then X is NP-Hard.*

5 Unassigned Crease Patterns

In this section we present gadgets simulating the elemental SCN gadgets with unassigned crease patterns. They are shown in Fig. 4.

We define a variable gadget to be a pair of parallel creases placed close together having an direction as shown in Fig. 4. By pleat-consistency and transitivity, $\lambda_f(a,b) = \lambda_f(b,c) = \lambda_f(a,c)$ so, local to the gadget, it has exactly two globally flat foldable states. We say the variable is *True* if the face to the right of the variable direction is above the face to left ($\lambda_f(a,c) = 1$), and *False* otherwise.

Lemma 3. *The unassigned crossover gadget is a globally flat foldable crease pattern if and only if opposite variables are equal.*

Proof. Refer to Fig. 4. Assume global flat foldability. Let A, B, C, D, E, F be the maximal subsets of the faces respectively containing points a, b, c, d, e, f such

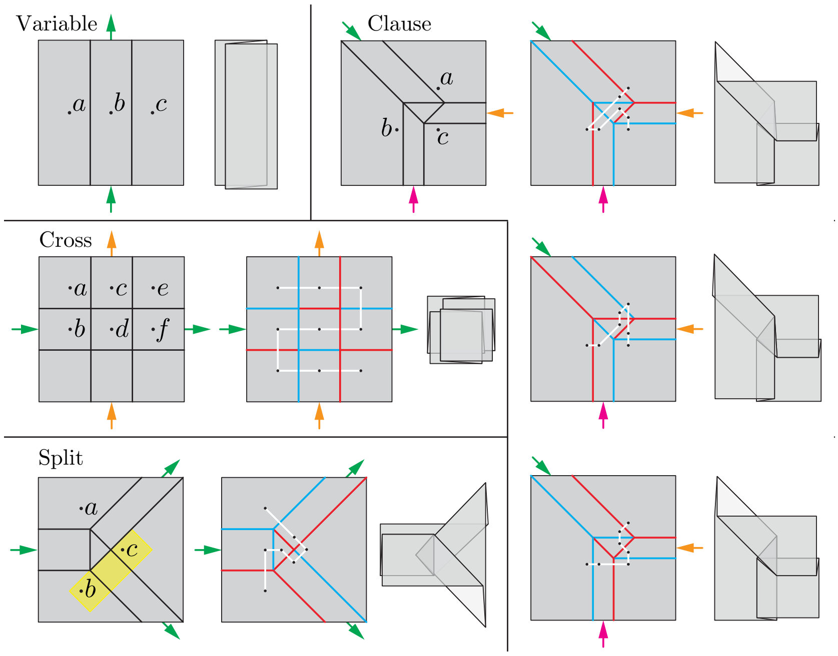

Fig. 4. Elemental SCN Gadgets simulated with unassigned crease patterns.

that every pair strictly overlap. First assume that $\lambda_f(a, b) = \lambda_f(c, d)$. By Taco-Taco with respect to adjacencies A, C and B, D, $\lambda_f(a, d) = \lambda_f(c, b)$. By Taco-Taco with respect to adjacencies A, B and C, D, $\lambda_f(a, c) = -\lambda_f(b, d)$. By Pleat-Consistency on A, C, E, $\lambda_f(a, c) = \lambda_f(c, e)$. By Pleat-Consistency on B, D, F, $\lambda_f(b, d) = \lambda_f(d, f)$. So $\lambda_f(c, e) = -\lambda_f(d, f)$. By Taco-Taco with respect to adjacencies C, D and E, F, $\lambda_f(c, f) = -\lambda_f(d, e)$. By Taco-Taco with respect to adjacencies C, E and D, F, $\lambda_f(c, d) = \lambda_f(e, f)$. Thus because $\lambda_f(a, b) = \lambda_f(e, f)$, the variable on the left has the same value as the one on the right. Alternatively if $\lambda_f(a, b) = -\lambda_f(c, d)$, the same series of arguments yields that $\lambda_f(c, d) = -\lambda_f(e, f)$, so $\lambda_f(a, b) = \lambda_f(e, f)$. Thus if global flat foldability holds, opposite variables are equal. Now assume that opposite variables are equal. The M/V assignment in Fig. 4 completely induces λ_f, along with consistency and transitivity. The path shown is a linear order on the faces satisfying global layer ordering. Further, every other assignment of variables can be represented by a reflection of this crease pattern. □

Lemma 4. *The unassigned split gadget is a globally flat foldable crease pattern if and only if its three variables are equal.*

Proof. Refer to Fig. 4. Assume global flat foldability. Let A and B be the faces containing points a and b respectively. The region highlighted in the figure and A must satisfy Path-Consistency, so $\lambda_f(a, b) = \lambda_f(a, c)$. Since the crease pattern is symmetric, $\lambda_f(b, a) = \lambda_f(b, c)$. Then, by antisymmetry, $\lambda_f(a, b) = \lambda_f(c, b)$, and therefore all variables are equal. Now assume all variables are equal. The path shown in Fig. 4 is a linear order on the faces satisfying global layer ordering. Further, every other assignment of variables can be represented by a reflection of this crease pattern. □

Lemma 5. *The clause gadget is a globally flat foldable crease pattern if and only if its three variables are not all equal.*

Proof. Refer to Fig. 4. Assume for contradiction the clause gadget is global flat foldable and all variables are equal. By consistency $\lambda_f(a,b) = \lambda_f(b,c) = \lambda_f(c,a)$. By transitivity, $\lambda_f(a,b) = \lambda_f(a,c)$. By antisymmetry, $\lambda_f(a,b) = -\lambda_f(c,a)$, a contradiction. Thus the variables are not all equal. Now assume all variables are not all equal. The paths shown in Fig. 4 are linear orders on the faces satisfying global layer ordering. Further, every other assignment of variables can be represented by the negation of one of these (M/V) assignments. □

Theorem 2. UNASSIGNED-FLAT-FOLDABILITY *is NP-complete, even for box pleated crease patterns.*

Proof. Given λ_f as our certificate, we can check in polynomial time whether it satisfies all conditions for global flat foldability, therefore UNASSIGNED-FLAT-FOLDABILITY is in NP. By Lemma 3, Lemma 4, and Lemma 5, UNASSIGNED-FLAT-FOLDABILITY can simulate the SCN-SATISFIABILITY gadgets. It remains to check if the gadgets can be consistently connected. Let the width of a variable be the distance between its parallel creases. The crossover gadget connects variables of the same width while the clause and split gadgets both connect variables whose ratios differ by a factor of $\sqrt{2}$. Setting the width of one variable in any gadget induces the width of the other variables in the gadget. Fixing the width of one variable in the Complex Clause Gadget (Fig. 3), a consistent unique width for all other variables is induced, resulting in the same width for each variable intersecting a left or right edge. Therefore, by Theorem 1, UNASSIGNED-FLAT-FOLDABILITY is NP-Hard. □

6 Assigned Crease Patterns

In this section we present gadgets simulating the elemental SCN gadgets with assigned crease patterns. They are shown in Fig. 5.

We define a variable gadget as a set of parallel creases placed close together having a direction and a crease assignment as shown in Fig. 5. By Taco-Tortilla, $\lambda_f(a,c) = \lambda_f(b,c) = \lambda_f(a,d) = \lambda_f(b,d)$, so, local to the gadget, it has exactly two globally flat foldable states. We say the variable is *True* if the faces to the right of the variable direction are above the faces to left ($\lambda_f(a,c) = 1$), and *False* otherwise.

Lemma 6. *The assigned crossover gadget is a globally flat foldable crease pattern if and only if opposite variables are equal.*

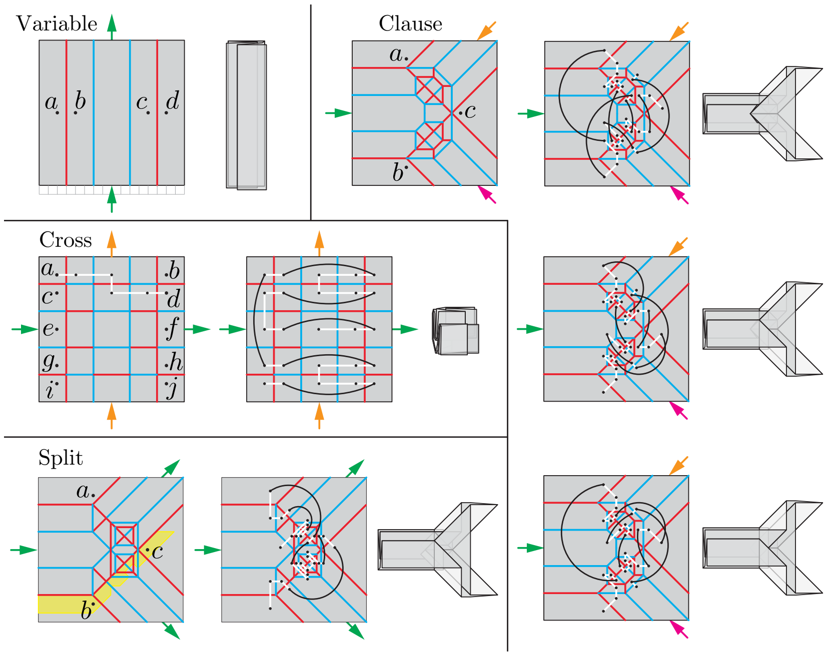

Fig. 5. Elemental SCN Gadgets simulated with assigned crease patterns.

Proof. Refer to Fig. 5. Assume global flat foldability. Let A, B, C, D be the maximal subsets of the faces respectively containing points a, b, c, d such that every pair strictly overlap. By transitivity on subset of λ_f induced by the M/V assignment shown, $\lambda_f(a, d) = \lambda_f(b, c) = -1$. By Taco-Taco with respect to adjacencies A, C and B, D, $\lambda_f(a, b) = -\lambda_f(c, d)$. Repeating this argument for adjacent rows of faces all the way down implies $\lambda_f(a, b) = -\lambda_f(c, d) = \lambda_f(e, f) = -\lambda_f(g, h) = \lambda_f(i, j)$. Thus, the variable on the top edge of the gadget has the same value as the one on the bottom. First assume $\lambda_f(g, a) = \lambda_f(a, b)$. Then previous implications imply $\lambda_f(g, a) = -\lambda_f(g, h)$. By transitivity and antisymmetry, $\lambda_f(g, a) = \lambda_f(h, b)$. Thus, the variable on the left side of the gadget has the same value as the one on the right. Alternatively, assume $-\lambda_f(g, a) = \lambda_f(a, b)$ so $\lambda_f(c, i) = \lambda_f(d, c)$. Then previous implications imply $\lambda_f(c, i) = \lambda_f(i, j)$. By transitivity and antisymmetry, $\lambda_f(c, i) = \lambda_f(d, j)$. Thus, the variable on the left side of the gadget has the same value as the one on the right. So, if globally flat foldable, opposite variables are equal. Now assume that opposite variables are equal. One can fix a unique λ_f by choosing a subset of λ_f in addition to the subset induced by the M/V assignment and consistency. The path shown in Fig. 5 is a linear order on the faces satisfying global layer ordering. Further, every other assignment of variables can be represented by a reflection of this crease pattern. $\square$

Lemma 7. *The assigned split gadget is a globally flat foldable crease pattern if and only if its three variables are equal.*

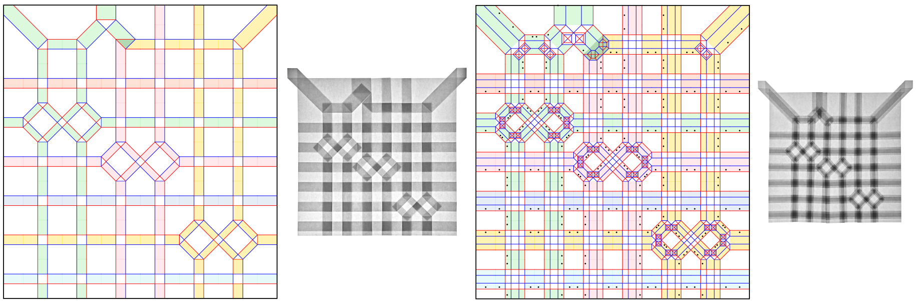

Fig. 6. A folded example of our assigned reduction with two clauses on four variables.

Proof. Refer to Fig. 5. Assume global flat foldability. Let A and B be the faces containing points a and b respectively. The region highlighted in the figure and A must satisfy Path-Consistency, so $\lambda_f(a,b) = \lambda_f(a,c)$. Since the crease pattern is symmetric, $\lambda_f(b,a) = \lambda_f(b,c)$. Then, by antisymmetry, $\lambda_f(a,b) = \lambda_f(c,b)$, and therefore all variables are equal. Now assume all variables are equal. The path shown in Fig. 5 is a linear order on the faces satisfying global layer ordering. Further, any other assignment of variables can be attained by a reflection. □

Lemma 8. *The assigned clause gadget is a globally flat foldable crease pattern if and only if its three variables are not all equal.*

Proof. Refer to Fig. 5. Assume for contradiction the clause gadget is global flat foldable and all variables are equal. By consistency $\lambda_f(a,b) = \lambda_f(b,c) = \lambda_f(c,a)$. By transitivity, $\lambda_f(a,b) = \lambda_f(a,c)$. By antisymmetry, $\lambda_f(a,b) = -\lambda_f(c,a)$, a contradiction. Thus the variables are not all equal. Now assume all variables are not all equal. The paths shown in Fig. 5 are linear orders on the faces satisfying global layer ordering. Further, any other assignment of variables can be attained by reversing the arrows in the figure. □

Theorem 3. Assigned-Flat-Foldability *is NP-complete, even for box pleated crease patterns.*

Proof. Given λ_f as our certificate, we can check in polynomial time whether it satisfies all conditions for global flat foldability and if it is consistent with the crease assignment, therefore Assigned-Flat-Foldability is in NP. By Lemma 6, Lemma 7, and Lemma 8, Assigned-Flat-Foldability can simulate the SCN-Satisfiability gadgets. It remains to check if the gadgets can be consistently connected. Let the width of a variable be the distance between its two parallel mountain creases. By the same argument as in the proof of Theorem 2, widths of variables can be assigned consistently. Therefore, by Theorem 1, Assigned-Flat-Foldability is NP-Hard. □

7 Conclusion

Table 1 overviews our results and open problems. We proved UNASSIGNED-FLAT-FOLDABILITY and ASSIGNED-FLAT-FOLDABILITY are NP-complete, even for box pleated crease patterns containing no more than 9 and 25 layers respectively. Are these problems still hard for even more restricted inputs? The computational complexity of ASSIGNED-FLAT-FOLDABILITY is still open when the crease pattern is a $m \times n$ square grid called a map [ABD+04]. Orthogonal folding, with crease patterns restricted to orthogonally aligned creases, is also open.

	General	**Box Pleating**	Orthogonal	Map
Unassigned	Hard [BH96]	**Hard (Ours)**	Poly [ABD+04]	Always True
Assigned	Hard [BH96]	**Hard (Ours)**	Open	Open

Table 1. Overview of **our results** and open problems. 'Hard' and 'Poly' designate problems that are NP-complete or solvable in polynomial time respectively.

Acknowledgements: This work was begun at the 2015 Bellairs Workshop on Computational Geometry, co-organized by Erik Demaine and Godfried Toussaint. We thank the other participants of the workshop for stimulating discussions, and to Barry Hayes for helpful comments leading to some simplifications. The research of H. Akitaya was supported by NSF grant CCF-1422311 and Science without Borders. The research of T. Hull and T. Tachi were respectively supported by NSF grant EFRI-ODISSEI-1240441 and the JST Presto Program.

References

[ABD+04] Esther M. Arkin, Michael A. Bender, Erik D. Demaine, Martin L. Demaine, Joseph S. B. Mitchell, Saurabh Sethia, and Steven S. Skiena. When can you fold a map? *Computational Geometry: Theory and Applications*, 29(1):23–46, September 2004.

[BDDO10] Nadia M. Benbernou, Erik D. Demaine, Martin L. Demaine, and Aviv Ovadya. Universal hinge patterns to fold orthogonal shapes. In *Origami*^{5}, pages 405–420. A K Peters, Singapore, July 2010.

[BH96] Marshall Bern and Barry Hayes. The complexity of flat origami. In *Proceedings of the Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '96, pages 175–183, Philadelphia, PA, USA, 1996. Society for Industrial and Applied Mathematics.

[DFL10] Erik D. Demaine, Sándor P. Fekete, and Robert J. Lang. Circle packing for origami design is hard. In *Origami*^{5}, pages 609–626. A K Peters, Singapore, July 2010.

[DO07] Erik D. Demaine and Joseph O'Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra.* Cambridge University Press, July 2007.

[Rob77] S. A. Robertson. Isometric folding of Riemannian manifolds. *Proc. of the Royal Society of Edinburgh*, 79(3-4):275–284, 1977.

[Sch78] Thomas J Schaefer. The complexity of satisfiability problems. In *Proceedings of the tenth annual ACM symposium on Theory of computing*, pages 216–226. ACM, 1978.