FACE FLIPS IN ORIGAMI TESSELLATIONS

Hugo A. Akitaya,[*] *Vida Dujmović,*[†] *David Eppstein,*[‡] *Thomas C. Hull,*[§] *Kshitij Jain,*[¶] *Anna Lubiw*[‖]

Abstract. Given a locally flat-foldable origami crease pattern $G = (V, E)$ (a straight-line drawing of a planar graph on the plane) with a mountain-valley (MV) assignment $\mu : E \to \{-1, 1\}$ indicating which creases in E bend convexly (mountain) or concavely (valley), we may *flip* a face F of G to create a new MV assignment μ_F which equals μ except for all creases e bordering F, where we have $\mu_F(e) = -\mu(e)$. In this paper we explore the configuration space of face flips that preserve local flat-foldability of the MV assignment for a variety of crease patterns G that are tilings of the plane. We prove examples where μ_F results in a MV assignment that is either never, sometimes, or always locally flat-foldable, for various choices of F. We also consider the problem of finding, given two locally flat-foldable MV assignments μ_1 and μ_2 of a given crease pattern G, a minimal sequence of face flips to turn μ_1 into μ_2. We find polynomial-time algorithms for this in the cases where G is either a square grid or the Miura-ori, and show that this problem is NP-complete in the case where G is the triangle lattice.

1 Introduction

An *origami crease pattern* (G, P) is a straight-line drawing of a planar graph $G = (V, E)$ on a region P of $\mathbb{R}^2$, where we allow for the case $P = \mathbb{R}^2$ and G is an infinite graph. A *flat origami* is a function $f : P \to \mathbb{R}^2$ from an origami crease pattern (G, P) to the plane that is continuous, an isometry on each face of G, and non-differentiable on all the edges and vertices of G. If a flat origami function can be found for a given crease pattern (G, P), we say that the origami crease pattern is *phantom foldable*, a term that is meant to indicate that the crease pattern can fold flat ignoring whether or not the region P, thought of as a sheet of paper, would have to intersect itself; this term is adopted from [7]. However, the important object of study is the crease pattern (G, P), since a flat origami function, if it exists, can be generated from (G, P) by fixing a face F_0 of G and then reflecting about the edges in G to determine where all the other faces of G map when all the crease lines are folded (see [6] for details).

We say that an origami crease pattern is *globally flat-foldable* if more conditions are met. Specifically, we record the state of each crease segment with a function $\mu : E \to \{-1, 1\}$,

[*]Tufts University, hugo.alves_akitaya@tufts.edu
[†]University of Ottawa, vida.dujmovic@uottawa.ca
[‡]University of California, Irvine, eppstein@uci.edu
[§]Western New England University,thull@wne.edu
[¶]Borealis AI, k22jain@uwaterloo.ca
[‖]University of Waterloo, alubiw@uwaterloo.ca

where $\mu(e) = -1$ means that the crease e is a *valley crease* (meaning it bends the paper in a concave direction) and $\mu(e) = 1$ means that e is a *mountain crease* (so it bends in a convex direction). We refer to μ as a *mountain-valley (MV) assignment* on G. We then say that a crease pattern is globally flat-foldable if a MV assignment can be found that creates a *superposition order* on the faces of a refinement of the crease pattern that satisfies *non-crossing conditions*, which are meant to model how the folded paper avoids self-intersecting when folded flat; see [16] for details. Determining if a crease pattern is globally flat-foldable is cumbersome and has been shown to be NP-complete [1, 5].

In this paper we will only be concerned with crease patterns (G, P) that are *locally flat-foldable*, which means that for each vertex v of G in the interior of P, the sub-crease pattern (G_v, P_v) made from the faces of G in P that are adjacent to v is globally flat-foldable. As we will see in Section 2, it is easy to determine if a single vertex is globally flat-foldable. However, locally flat-foldable crease patterns possess rich combinatorial structures [6, 14] that have been studied in statistical mechanics [2] and applied to polymer membrane folding [7]. For an example that connects locally flat-foldable crease patterns to more traditional areas of graph theory, consider that counting *locally valid* MV assignments (that make a crease pattern locally flat-foldable) for the family of crease patterns known as the *Miura-ori* has been shown to be equivalent to counting proper 3-colorings of the vertices of grid graphs [10].

A new combinatorial tool that shows promise for helping explore locally valid MV assignments is the face flip. If F is a face in a flat-foldable crease pattern (G, P) and we have a MV assignment μ, then a *face flip of* F *in* (G, P) *under* μ is a new MV assignment μ_F where $\mu_F = \mu$ for all edges in G except those that border F, where we have $\mu_F = -\mu$. That is, we "flip" the creases bordering F from mountain to valley and vice-versa. The fundamental question, then, is whether μ_F will be locally valid given that μ is. If the answer is "yes" then we say that F is a *flippable face* relative to μ.

Face flips seem to have been first introduced by Kyle VanderWerf in [21]. While face flips are otherwise mathematically unexplored, they are related to studies in materials science and mechanical engineering on applied origami (such as [2, 8, 20]). Face flips provide a way of studying the likelihood of a material either being manipulated from one MV assignment to another, or folding to a state that is "close" to the target MV assignment state, where "close" could be interpreted as two vertices close to each other in the origami flip graph configuration space (to be defined in Section 2). Further, if the origami flip graph is disconnected, then its different components could identify folded states that have very little chance of being achieved by an actual folded material with a given target MV assignment.

In this paper, we examine the properties of face flips on flat-foldable crease patterns (G, P) where G is certain regular tilings of the plane. Such flat origamis are also known as *origami tessellations*, and they are of central interest in applications and prior work on flat foldings [2, 7, 8, 10, 18, 20].

Specifically, after setting up background results in Section 2, we will see in Section 3 families of quadrilateral crease patterns where any face flip on a MV assignment will preserve its local validity, another where only certain faces may be flipped, and yet another where no face flip will result in a valid MV assignment. In Section 4 we will prove that any locally-valid MV assignment of the Miura-ori crease pattern can be converted to any other via face flips, thus showing that the configuration space of locally-valid MV assignments of the Miura-ori is connected under face flips. In Section 5 we consider the considerably more complicated case of origami tessellations whose crease pattern is the triangle lattice, showing that its configuration space is also connected. However, we show that, unlike the quadrilateral cases, determining the minimum number of face flips needed to convert one locally-valid MV assignment of the triangle lattice to another is NP-complete using a reduction from minimum vertex cover with maximum degree three in a hexagonal grid.

2 Preliminaries

The most fundamental result of flat-foldability is *Kawasaki's Theorem*:

Theorem 2.1 (Kawasaki). *Let (G, P) be an origami crease pattern where G has only one vertex v in the interior of P and all edges in G are adjacent to v. Let $\alpha_1, \ldots, \alpha_k$ be the sector angles, in order, between the consecutive edges around v. Then (G, P) is globally flat-foldable if and only if $k = 2n$ is even and*

$$\alpha_1 - \alpha_2 + \alpha_3 - \cdots - \alpha_{2n} = 0.$$

See [14] for a proof.

One of the most basic requirements for a MV assignment to be locally valid is for there to be an appropriate number of mountains and valleys at each vertex.

Theorem 2.2 (Maekawa). *Let v be a vertex in a globally flat-foldable crease pattern with MV assignment μ and let A be the set of crease edges adjacent to v. Then*

$$\sum_{e \in A} \mu(e) = \pm 2. \tag{1}$$

This is known as *Maekawa's Theorem*, although it is often written as $M - V = \pm 2$ where M and V are the number of mountain and valley creases, respectively, at v. A proof can be found in [14]; it relies on the fact that the cross section of a physically-flat-folded vertex will be a closed curve with winding number 1 about any point in its interior, giving us that $\pi M - \pi V = \pm 2\pi$.

Kawasaki's and Maekawa's Theorems state necessary conditions for flat-foldability at a vertex. If the sector angles around a vertex are all equal in a single-vertex crease pattern, then Maekawa's Theorem becomes necessary and sufficient.

Theorem 2.3. *Let v be a vertex in a globally flat-foldable crease pattern where the sector angles between consecutive creases at v are all equal. Then a MV assignment μ will be locally valid at the vertex v if and only if μ satisfies Equation (1) at v.*

Proof. Sufficiency can most easily be seen by an inductive argument, although to do this one must extend the concept of flat foldability to single-vertex *cone folds*, where a crease pattern is embedded on a cone C and the only vertex of the crease pattern in the interior of C is at the cone's apex. The proof Maekawa's Theorem for cone folds (the necessary direction of the current theorem) is the same as it is for flat paper (see [14]). In the other direction, suppose we have a crease pattern (G, P) where G has only one vertex v in the interior of P, the sector angles between the consecutive creases at v are all equal, P is a subset of $\mathbb{R}^2$ or a cone with v at its apex, and $\sum \mu(e) = \pm 2$ where the sum is over all edges e adjacent to v. Then the degree of v must be even (since, if M and V are the number of mountains and valleys, respectively, at v, we have $\deg(v) = M + V = M - V + 2V = 2V \pm 2$). Inducting on n where $\deg(v) = 2n$, we have that the base case has two creases, equal sector angles, and either two mountains or two valleys, which clearly can fold flat. If $n > 2$, then there must exist two consecutive creases (say, going clockwise around v) with different MV parity; these two creases may be folded, and the others left unfolded, to turn (G, P) into a new crease pattern (G', P') with two fewer creases, equal sector angles, smaller cone angle, and $\sum \mu(e) = \pm 2$ in G'. By induction, (G', P') will globally fold flat (which is equivalent to local flat foldability since G has only one interior vertex), implying that (G, P) will as well. □

When the angles between consecutive creases around a flat-folded vertex are not all equal, then Maekawa is only a necessary condition. Nonetheless, other constraints on valid MV assignments for such vertices can be deduced. For example, if we have consecutive sector angles $\alpha_{i-1}, \alpha_i, \alpha_{i+1}$, with α_i between crease edges e_i and e_{i+1}, at a vertex with a valid MV assignment μ, where α_i is strictly smaller than both of the other angles, then we must have $\mu(e_i) = -\mu(e_{i+1})$. This is because if $\mu(e_i) = \mu(e_{i+1})$ then the two sectors with angles α_{i-1} and α_{i+1} would be folded over, and more than cover, the sector with angle α_i on the same side of the paper, causing the sectors of paper with angles α_{i-1} and α_{i+1} to intersect each other. This constraint, where $\mu(e_i) = -\mu(e_{i+1})$ is forced, is sometimes called the *Big-Little-Big Angle Lemma*. This is actually a special case of the following theorem, proved in [14]:

Theorem 2.4. *Let v be a vertex in a globally flat-foldable crease pattern with MV assignment μ, and suppose that we have a local minimum of consecutive equal sector angles between the crease edges $e_i, \ldots, e_{i+k+1}$ at v. That is, $\alpha_i = \alpha_{i+1} = \cdots = \alpha_{i+k}$ where $\alpha_{i-1} > \alpha_i$ and $\alpha_{i+k+1} > \alpha_i$. Then*

$$\sum_{j=i}^{i+k+1} \mu(e_j) = \begin{cases} 0 & \text{if } k \text{ is even,} \\ \pm 1 & \text{if } k \text{ is odd.} \end{cases}$$

Remark on local vs. global flat-foldability: In the Introduction we commented that determining the global flat-foldability of a crease pattern is cumbersome. At the same time, we need to know when a single vertex in a crease pattern will be globally flat-foldable under a given MV assignment in order to determine the local validity of a MV assignment (and thus the local flat-foldability of the crease pattern). Fortunately, determining the global flat-foldability of a single vertex crease pattern is more easily done by modeling how a single vertex folds flat using concepts like the winding number, the alternating sum of the sector angles around the vertex, and the "big-little-big angle" constraint. These tools make the proofs of Theorems 2.1, 2.2, 2.3, and 2.4 rigorous, as seen in [6,14,16]. This is what allows us to speak with confidence about the combinatorial geometry of locally valid MV assignments for flat-foldable crease patterns.

In the Introduction we defined face flips of a flat-foldable crease pattern with a locally valid MV assignment. We also say that two locally valid MV assignments μ_1 and μ_2 are *face-flippable* if there exists a sequence of flippable faces (thus preserving the local validity of the MV assignments along the way) in the crease pattern whose flipping will turn μ_1 into μ_2, or vice-versa.

We now formalize the concept of the configuration space of locally valid MV assignments with a variation of the flip graph from discrete geometry. Given a flat-foldable crease pattern (G, P), define the *origami flip graph* to be the graph whose vertices are locally valid MV assignments of (G, P), and where two MV assignments μ_1 and μ_2 are adjacent in this graph if and only if μ_1 is face-flippable to μ_2 by flipping exactly one face. We then say that the configuration space of MV assignments for (G, P) is connected if its origami flip graph is a connected graph.

For examples of crease patterns with a given MV assignment, see the two shown in Figure 1(a). In this Figure, and throughout the paper, bold lines indicate mountain creases and non-bold lines represent valleys. Both of the MV assignments μ_1 and μ_2 are clearly locally valid (each vertex satisfies Maekawa's Theorem and the sector angles around each vertex are all 90°, so Theorem 2.3 applies). They are also globally flat-foldable with these MV assignments, however. μ_2 can be folded by making the vertical valley creases first, say rolling the paper from left-to-right to get a 3×1 strip, and then folding the horizontal creases. μ_1 is much harder to fold (try folding the entire bottom-most horizontal crease line first).

3 Square and kite tessellations

In this section we will analyze three families of quadrilateral-based origami tessellations to demonstrate different kinds of face flip-behavior and different configuration spaces. Specifically, we will see:

- *Square grid* crease patterns, where any face can be flipped and the configuration space of MV assignments is connected.
- *Huffman grid* tessellations, where no face flips are possible and thus the configuration space is totally disconnected.
- *Square twist* tessellations, where only half of the faces can be flipped but the configuration space is still connected.

3.1 Square grid tessellations

We first consider an $m \times n$ grid of squares, denoted $G_{m,n}$, as our crease pattern, where the region of paper P will be an $m \times n$ rectangle. This will have $(m-1)(n-1)$ vertices in the interior of P, each of which will have degree 4 and 90° angles between the creases. Thus a

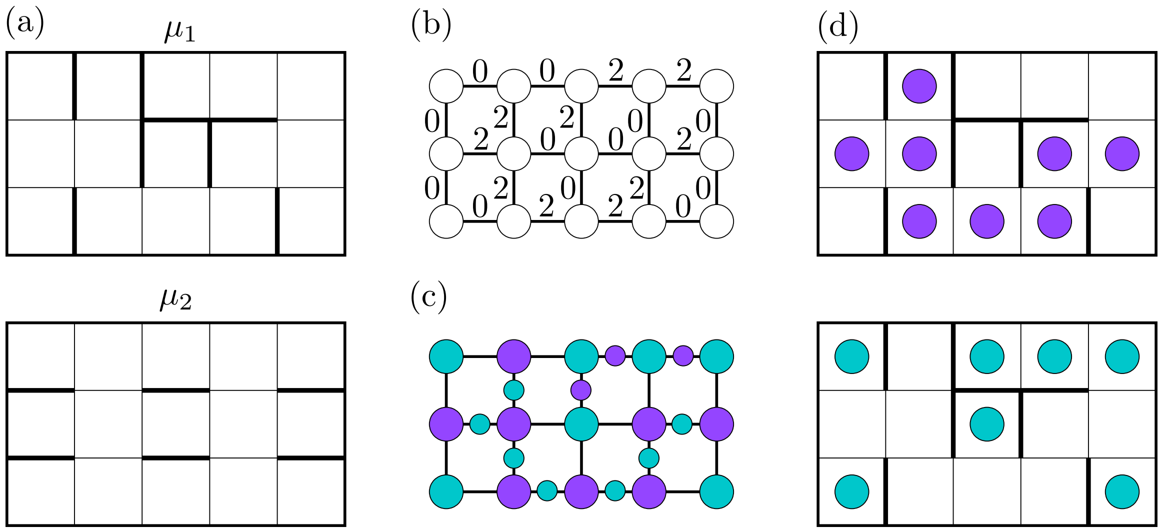

Figure 1: (a) Two locally-valid MV assignments μ_1 and μ_2 of $G_{3,5}$. (b) The weight w on the edges of $G^*(3,5)$. (c) The graph $\overline{G^*}_{3,5}$ with a 2-coloring. (d) The face flip sets.

MV assignment μ for $G_{m,n}$ will be locally valid if and only if each vertex satisfies Maekawa's Theorem, i.e., has 3 mountains and 1 valley or vice-versa.

Theorem 3.1. *Given a square grid tessellation and a locally-valid MV assignment μ, flipping any face will result in another locally-valid MV assignment.*

Proof. Flipping a face F in $G_{m,n}$ will affect μ for at most four interior vertices in P. Consider one of them, v, and let $e_1, \ldots, e_4$ be the crease edges adjacent to v where e_1 and e_2 border F. If $\mu(e_1) = \mu(e_2)$ then $\sum_{i=1}^4 \mu(e_i) = -\sum_{i=1}^4 \mu_F(e_i)$. If $\mu(e_1) = -\mu(e_2)$ then $\mu_F(e_1) = -\mu_F(e_2)$ as well. In both cases we have that Maekawa's Theorem is still satisfied at v under μ_F, and thus μ_F will be locally valid by Theorem 2.3. □

Lemma 3.2. *Let μ_1 and μ_2 be two locally-valid MV assignments of $G_{m,n}$ and let v be an interior vertex of $G_{m,n}$. Then among the four edges adjacent to v, μ_1 and μ_2 can agree on all four, only 2, or none of the edges, but not on three or one.*

Proof. Let the edges at v be $e_1, \ldots, e_4$, and suppose that μ_1 and μ_2 agree on only one or three of these edges. Then in either case we have that $\prod_{i=1}^4 \mu_1(e_i)\mu_2(e_i) = -1$, since the disagreeing pairs of $\mu_1(e_i)$ and $\mu_2(e_i)$ will each contribute -1 and the agreeing pairs will contribute 1 to the product. However, this product also equals $\prod_{i=1}^4 \mu_1(e_i) \prod_{i=1}^4 \mu_2(e_i) = 1$ since $\prod_{i=1}^4 \mu(e_i) = -1$ for all locally-valid MV assignments μ. This is a contradiction. □

Our goal is to prove that the MV configuration space for $G_{m,n}$ is connected and to devise an algorithm to find the smallest number of face flips needed to flip between two given locally-valid MV assignments μ_1 and μ_2.

To that end, let μ_1 and μ_2 be locally-valid MV assignments of $G_{m,n}$, such as those shown in Figure 1(a). We consider the internal planar dual graph, $G^*_{m,n}$ (that is, the dual of

$G_{m,n}$ ignoring the external face). For every edge e of $G_{m,n}$ denote the corresponding edge in $G^*_{m,n}$ by e^*. Assign a weight function w to the edges in $G^*_{m,n}$ given by $w(e^*) = |\mu_1(e)+\mu_2(e)|$. That is, $w(e^*)$ will equal 0 if $\mu_1(e) \neq \mu_2(e)$ and 2 if $\mu_1(e) = \mu_2(e)$. See the example in Figure 1(b).

Now create a new graph $\overline{G^*}_{m,n}$ made by taking $G^*_{m,n}$ and adding a vertex in the middle of every edge e^* with $w(e^*) = 2$. See Figure 1(c).

Lemma 3.3. *The graph $\overline{G^*}_{m,n}$ is properly 2-vertex colorable.*

Proof. This follows from Lemma 3.2; since the four edges e_i adjacent to each internal vertex v of $G_{m,n}$ have $\mu_1(e_i) = \mu_2(e_i)$ for only four, two, or none of the e_i, we have that each square face of $G^*_{m,n}$ will either remain a square, become a hexagon, or become an octagon in $\overline{G^*}_{m,n}$. Thus $\overline{G^*}_{m,n}$ has only even cycles, which means it is properly 2-vertex colorable. □

Let $c : E(\overline{G^*}_{m,n}) \to \{\text{purple, teal}\}$ be a proper 2-vertex coloring of $\overline{G^*}_{m,n}$. Then c will give us a (most likely not proper) 2-coloring of the vertices of $G^*_{m,n}$.

Theorem 3.4. *Let μ_1 and μ_2 be two locally-valid MV assignments of $G_{m,n}$. Then μ_1 and μ_2 are face-flippable, and this can be achieved by flipping the faces corresponding to all the purple (or all the teal) vertices in $G^*_{m,n}$ under the 2-coloring c described above.*

Proof. Suppose we start with the MV assignment μ_1 and flip all the faces corresponding to the purple vertices in $G^*_{m,n}$. Consider an edge e of $G_{m,n}$ where $\mu_1(e) = \mu_2(e)$. Then $w(e^*) = 2$, and thus the faces in $G_{m,n}$ that border e have corresponding vertices in $G^*_{m,n}$ that are both colored purple or teal. This means that both of these faces were flipped or both were not flipped. In either case, the edge e remains with the same MV assignment after all the purple flips.

Now consider $e \in E(G_{m,n})$ where $\mu_1(e) \neq \mu_2(e)$. Then $w(e^*) = 0$, so the vertices adjacent to e^* in $G^*_{m,n}$ are colored differently under c, which means one of the faces bordering e in $G_{m,n}$ is flipped and the other is not flipped. This implies that the edge e will change its MV assignment after all the purple flips. We conclude that the MV assignment μ_1 will turn into μ_2 after flipping all the faces corresponding to the purple vertices in $G^*_{m,n}$, and the same argument works for the teal vertices. □

It is clear from the construction that the two face flip sets generated by the purple vertices and by the teal vertices in $G^*_{m,n}$ are minimal in that removing any face from them will not result in flipping from μ_1 to μ_2 or vice-versa. Could there be some other minimal set of faces that flips from μ_1 to μ_2? No, since every edge e in $G_{m,n}$ with $\mu_1(e) \neq \mu_2(e)$ requires exactly one of its adjacent faces to be in the flip set, and each of these faces are in a different 2-color class. Furthermore, every edge e with $\mu_1(e) = \mu_2(e)$ will either have both or neither of its adjacent faces in the flip set, which is again determined by the 2-color class. Therefore, picking a face from one of the 2-color classes forces the rest of that color class to also be in the face flip set. We conclude that the smallest number of face flips needed to flip from μ_1 to μ_2 is the smaller of the purple and the teal face flip sets. Further, computing

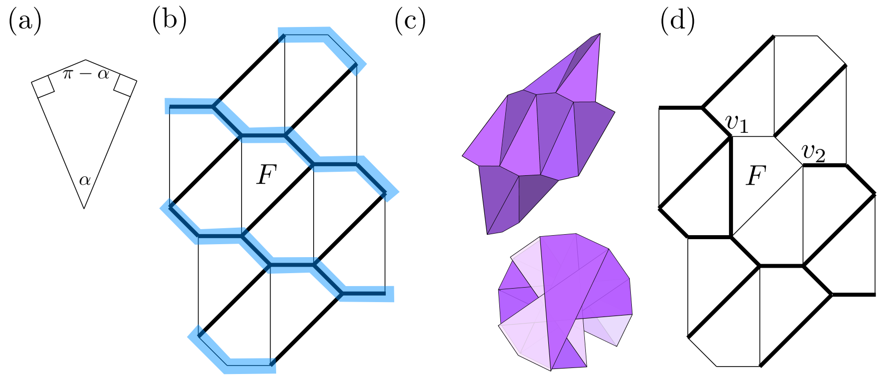

Figure 2: (a) The kite tile for a Huffman grid. (b) The resulting Huffman grid, with the short rows highlighted in blue. (c) The crease pattern partially folded and folded flat.[1] (d) The face F flipped.

this minimal set of face flips can be done in linear time on the number of faces, mn, since each step in the process involves a computation on the vertices or edges of $G_{m,n}$.

For a few interesting examples, if $\mu_1 = \mu_2$ then one of the color sets, say purple, will include all the faces of $G_{m,n}$ and the other, teal, will be the empty set. Clearly the empty set is the smaller set, meaning that no flips are needed. If $\mu_1(e) = -\mu_2(e)$ for all $e \in E(G_{m,n})$ then the 2-coloring of $G^*_{m,n}$ will be a simple checkerboard coloring, and if mn is even then the purple and teal flip sets will have equal size.

3.2 Huffman grid tessellations

A *Huffman grid* is a type of monohedral origami tessellation introduced by Huffman in [13]. (See also [8].) The generating tile is a quadrilateral with two opposite corners having right angles, as in Figure 2(a). The other two interior angles are labeled α and $\pi - \alpha$, where we assume $\alpha < \pi/2$. The tiling generated by this tile, shown in Figure 2(b), has vertices that satisfy Kawaski's Theorem, and by the Big-Little-Big Lemma, the creases bordering an angle of α must have different MV parity. The creases that do not border an α angle form zig-zag paths, which we call *short rows*; they are highlighted blue in Figure 2(b). Applying Maekawa's Theorem at each vertex implies that each short row must be either entirely mountain or entirely valley creases. As seen in [8], folding large Huffman grids according to a locally-valid MV assignment will eventually cause the paper to curl up and self-intersect.

Theorem 3.5. *Given the Huffman grid tessellation and a valid MV assignment, any face flip will generate a MV assignment that is not valid.*

Proof. Suppose we flip a face F where we label F's right angle corners v_1 and v_2, as in Figure 2(d). Then before we flipped F, the two creases surrounding the angle α at v_1 had

[1]The images in Figure 2(c) were generated using R. J. Lang's *Tessellatica 11.1* Mathematica code [17].

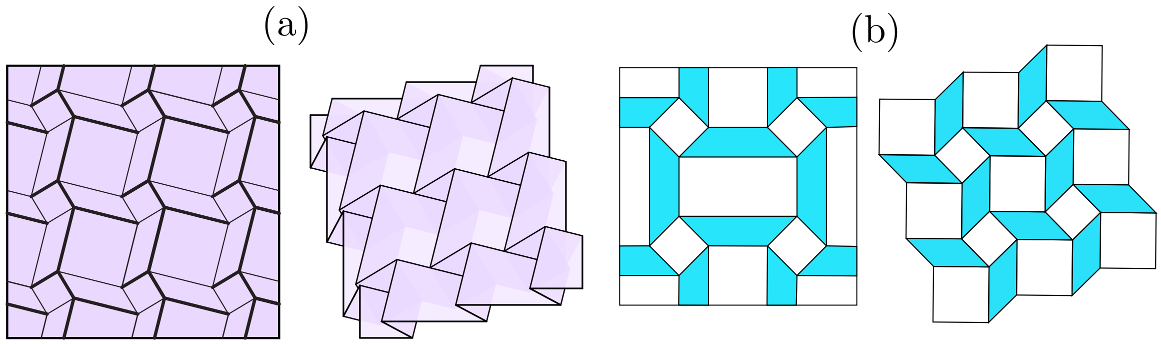

Figure 3: (a) A square twist tessellation with MV assignment and the folded form.[2] (b) A square twist tessellation with the flippable faces colored.

opposite MV parity. After flipping F these creases will have the same MV parity, which violates the Big-Little-Big Lemma at v_1, and the same thing will happen at v_2. Thus if μ was our original locally-valid MV assignment, the flipped assignment μ_F will not be locally valid. □

Theorem 3.5 implies that there are no edges in the face flip graph for Huffman grid tessellations, and thus its configuration space is as disconnected as possible.

3.3 Square twist tessellations

Twist folds are collections of creases that, when folded flat, cause a polygon to rotate from the unfolded to the flat-folded state relative to the rest of the paper. They are of interest for their geometric character [9], applications in origami mechanics [20], and their ability to be used in tessellations [11]. The classic *square twist*, where identical degree-4 vertices with sector angles 45°, 90°, 135°, and 90° are used to twist a square, can tessellate in several different ways; a few examples are shown in Figure 3. (Note that other angles besides the $(45°, 135°)$ pair are possible; see [11, 18].) In all classic square twist tessellations the 90° angles will form square or rectangle faces in the crease pattern tiling, while the 45° and 135° angles will form parallelograms or trapezoids, as shown in Figure 3(b).

The Big-Little-Big Lemma implies that the two creases bordering the 45° angles in the vertices of a square twist tessellation must have different MV assignments. Thus, given any locally valid MV assignment for a classic square twist tessellation, flipping any square or rectangle face will cause a Big-Little-Big Lemma violation at a 45° angle and make the MV assignment no longer locally valid. Therefore square or rectangle faces can never be flipped by themselves. In fact, the only way to flip a square or rectangle face and not violate the Big-Little-Big Lemma somewhere is to flip *all* of the square and rectangle faces at the same time, which is equivalent to flipping all the non-square/rectangle faces.

Therefore the only faces that are safe to flip (and preserve local validity of the MV

[2] The images in Figure 3(a) were generated using R. J. Lang's *Tessellatica 11.1* Mathematica code [17].

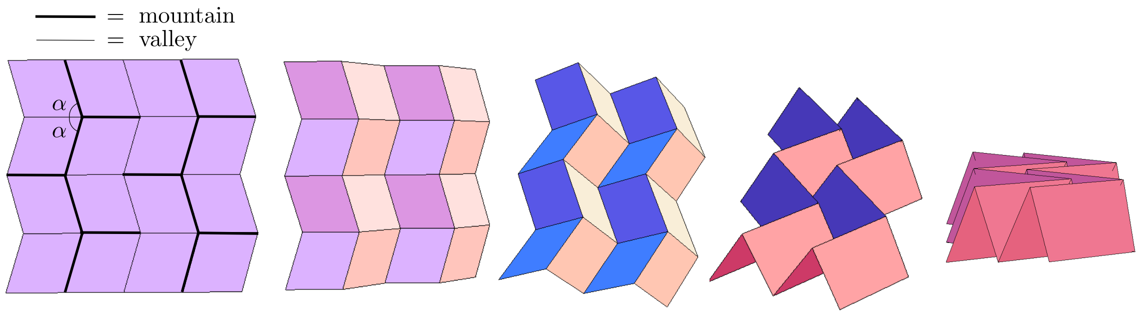

Figure 4: A 4×4 Miura-ori crease pattern with the classical MV assignment, along with how it globally folds flat.

assignment) in a square twist crease pattern are the parallelograms and trapezoids, shown in blue in Figure 3(b). Since all of these parallelograms and trapezoids are edge-disjoint, and these represent the only ways to modify a MV assignment of such patterns, we arrive at the following:

Theorem 3.6. *The MV assignment configuration space of a square twist tessellation crease pattern is connected, and the only faces that are flippable are the parallelograms and trapezoids.*

Determining the minimum number of face flips needed to flip between two locally valid MV assignments μ_1 and μ_2 of a square twist tessellation can be determined in linear time, since all it requires is determining which parallelograms and trapezoids differ between μ_1 and μ_2.

4 The Miura-ori

The Miura-ori [10, 21] is a folding pattern named after Koryo Miura, formed by a sequence of equally spaced parallel lines crossed by zigzag paths that divide the paper into equal parallelograms. These parallelograms tile the plane by reflection across the parallel lines, and by translation in the direction parallel to the parallel lines; see Figure 4. In its classical global flat-folded state, the folds of the Miura-ori alternate between mountain and valley folds along each of the parallel lines, with each zigzag path consisting entirely of mountain folds or entirely of valley folds (in alternation along the sequence of zigzag paths). However, the same folding pattern has many locally flat-folded states. In a locally valid MV assignment of this folding pattern, we have that Maekawa's Theorem must hold and at each vertex the crease bordered by obtuse angles must have the same MV type as the majority of the creases at the vertex (see [10, 14]).

The $m \times n$ parallelograms of the Miura-ori have the same combinatorial structure (although different in their symmetry groups) as the square grid, and it will be helpful for us to use a combinatorial bijection between locally valid MV assignments of the Miura-ori and proper 3-colorings (with one vertex pre-colored) of the vertices of the internal planar

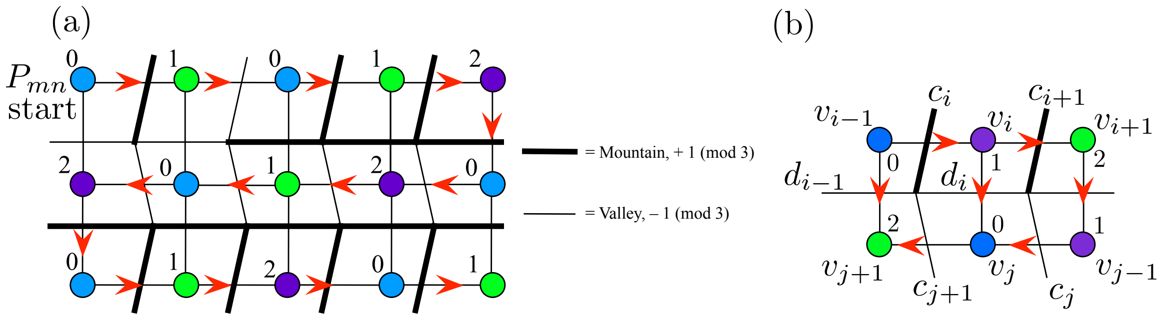

Figure 5: Bijection between local flat foldings of the Miura-ori and 3-colorings of the squares of a grid

dual grid graph $G_{m,n}$, as seen in Section 3.1. To construct the bijection, we position the Miura-ori as in Figure 5(a), with the top-left sector angle being obtuse, and overlay the dual grid graph with one vertex in each parallelogram. Let the three colors of the graph coloring be $\mathbb{Z}_3$, the integers modulo three. We then follow a directed path P_{mn} through the grid graph, starting from the top-left vertex, traveling horizontally to the top-right vertex, then down one edge, traveling horizontally to the left side again, then down another edge, and so on in a bostrophedon pattern through all the vertices in $G_{m,n}$ (this path is indicated by the red arrows in Figure 5(a)). This path determines our vertex coloring; we choose the color corresponding to 0 mod 3 for the first vertex in P_{mn}, and subsequent colors for the vertices on this path are calculated as follows: When the path crosses a mountain fold, we add 1 mod 3 to get the color for the next vertex, and when it crosses a valley fold, we subtract 1 mod 3. (See Figure 5.) More formally, if we let $s : V(G_{m,n}) \to \mathbb{Z}_3$ be a proper 3-coloring of the vertices of $G_{m,n}$, with the top-left vertex colored 0, that corresponds to a locally valid MV assignment μ. We have that for two consecutive vertices v_i, v_{i+1} along the path P_{mn} in $G_{m,n}$ that cross the crease c_k,

$$s(v_{i+1}) - s(v_i) \equiv \mu(c_k) \bmod 3. \tag{2}$$

It can be shown that this method of coloring the vertices of $G_{m,n}$ produces a proper 3-coloring for every locally valid MV assignment of the Miura-ori and, conversely, that every proper 3-coloring of $G_{m,n}$ with the top-left vertex colored 0 comes from a locally valid MV assignment in this way [4, 10].

In order to use this coloring-MV assignment correspondence, we need to know how face flips change the 3-coloring of the corresponding grid graph.

Lemma 4.1. *Suppose we flip a face F of a Miura-ori to change a locally valid MV assignment to another locally valid assignment. Then the corresponding 3-coloring of the grid graph will change only at the vertex that corresponds to F. Conversely, if we change a proper 3-coloring of the grid graph at a single vertex, the corresponding Miura-ori MV assignment will change by flipping a single face.*

Proof. First, we extend the directed path P_{mn} to a direction on the edges in the whole graph $G_{m,n}$ as follows: For each $e \in E(G_{m,n})$ that is in P_{mn}, retain the direction given by P_{mn}.

All edges e not in P_{mn}, will be vertical edges following the depiction of $G_{m,n}$ as shown in Figure 5; have all such edges be oriented in the downward direction. We claim that in this new directed version of $G_{m,n}$, which we denote $G'_{m,n}$, Equation (2) will still hold for the directed edges not in P_{mn}. That is, if v_i is directly above v_j in the grid graph, so (v_i, v_j) is a directed edge in $G'(m,n)$ that is not in P_{mn} and that crosses crease d_i in the Miura-ori crease pattern, we claim that $s(v_j) - s(v_i) \equiv \mu(d_i)$ mod 3 (where s is the proper coloring that corresponds to the locally valid MV assignment μ). The fact that this claim is true is implied in [10], but we will provide an explicit argument here.

Label the vertices and creases near the directed edge (v_i, v_j) as in Figure 5(b), where the red arrows indicate the directed edges of $G'_{m,n}$. We will prove our claim inductively along the path P_{mn}. Skipping, for now, the base case, let us assume that $s(v_{j+1}) - s(d_{i-1}) \equiv \mu(d_{i-1})$ mod 3, and we wish to show that the equivalent equation will be true for $\mu(d_i)$. First note that since the MV assignment of d_{i-1} must equal the majority MV assignment of the creases c_i, d_i, c_{j+1} (from [10, 14]), we have that

$$\mu(d_{i-1}) = \mu(c_i) + \mu(d_i) + \mu(c_{j+1}).$$

Also, by our coloring-MV assignment correspondence, we have

$$s(v_i) - s(v_{i-1}) \equiv \mu(c_i) \bmod 3 \text{ and } s(v_{j+1}) - s(v_j) \equiv \mu(c_{j+1}) \bmod 3.$$

Therefore,

$$\begin{aligned} s(v_j) - s(v_i) &\equiv s(v_{j+1}) - s(v_{i-1}) - \mu(c_{j+1}) - \mu(c_i) \bmod 3 \\ &\equiv \mu(d_{i-1}) - \mu(c_{j+1}) - \mu(c_i) \equiv \mu(d_{i-1}) \bmod 3. \end{aligned}$$

Establishing the base case uses similar computations, but is a bit technical since it requires telescoping a summation of $s(v_i)$ terms along the first two rows of P_{mn} and using an extension of Maekawa's Theorem applied to the first row of vertices in the Miura-ori crease pattern. See Lemmas 1 and 2 of [10] for the full details.

Now, if we flip face F that contains vertex v in $G_{m,n}$, this will change the MV assignments of the creases that border F and thus cause $s(v)$ to change in order to satisfy Equation (2) along the path P_{mn}. But this will not cause any of the other colors $s(u)$ to change for any neighbor u of v in $G_{m,n}$ because such vertices u are all connected to each other in the directed grid graph via paths that that do not include the vertex v (and thus must satisfy their own equations like Equation (2) that do not include v). Conversely, if we change a single grid graph vertex color, the same correspondence shows that the MV assignment cannot change except at the edges surrounding the corresponding face F. □

With this equivalence between flips and vertex recolorings in hand, we can apply known techniques for grid colorings to obtain the corresponding results for Miura-ori face flips.

Theorem 4.2. *Every two locally-valid MV assignments of the Miura-ori crease pattern can be converted to each other via face flips.*

Proof. This follows from Lemma 4.1 and from the already-known fact that every two 3-colorings of a grid can be converted to each other via single-vertex recolorings [12, Lemma 4.4]. More strongly, Goldberg et al. [12, Lemma 4.6] prove that the number of recoloring steps needed, for an $m \times n$ grid with $m \leq n$, is at most $2mn^2$. The same bound holds for face flips of the Miura-ori, where now m and n measure the number of parallelograms in the pattern in either direction. □

Theorem 4.3. *Given any two locally-valid MV assignments of the Miura-ori crease pattern, it is possible to find a minimum-length sequence of face flips that converts one to the other, in polynomial time.*

Proof. The correspondence between MV assignments and grid 3-colorings, described above, can be carried out in linear time, since it only requires a simple calculation for every vertex in $G_{m,n}$. An efficient algorithm for finding the shortest number of vertex-recoloring steps to convert one 3-coloring of any given graph into another (when such a sequence of recoloring steps exists, as we have already proven to be true in this case) was provided by Johnson et al. [15]. □

5 Triangle lattice tessellations

In this section we consider crease patterns that are finite regions of the triangle lattice. In Section 5.1 we show that the configuration space of locally valid MV assignments is connected under face flips, with a linear diameter. In Section 5.2 we show that it is NP-hard to find the minimum number of face flips to reconfigure between two locally valid MV assignments.

We begin with some basic facts about locally valid MV assignments in the triangle lattice. By Theorem 2.3, an MV assignment is locally valid on the triangle lattice if and only if for any vertex v, either v has 4 mountain folds and 2 valley folds, in which case we call it a *mountain vertex*, or v has 4 valley folds and 2 mountain folds, in which case we call it a *valley vertex*.

Consider what happens to a locally valid MV assignment μ at a vertex v if we flip an incident face f. If the two edges of f incident to v have opposite creases (one mountain and one valley), then μ_f will remain locally valid at v by Theorem 2.3—in fact, v retains its mountain/valley designation (Figure 6(a)). If the two edges of f incident to v are mountain creases and v is a mountain vertex, then flipping f changes v to a valley vertex (Figure 6(b)). Similarly, flipping two valley creases at a valley vertex creates a mountain vertex. Finally, if the two edges of f incident to v are mountain creases and v is a valley vertex (Figure 6(c)), or if the two edges are valley creases and v is a mountain vertex (Figure 6(d)), then f cannot be flipped, and we say that v *causes* f to be not flippable. To summarize:

Claim 5.1. *A face can be flipped unless it has 2 mountain creases incident to a valley vertex, or 2 valley creases incident to a mountain vertex.*

In particular, note that a face with a mountain and a valley crease has only one vertex that can potentially cause it to be not flippable.

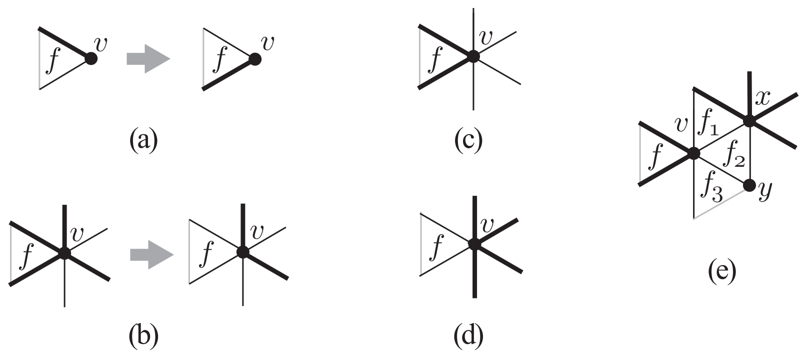

Figure 6: (a–d) Some cases for flipping a triangle f incident to vertex v: (a) one mountain crease (in bold) and one valley crease; (b) two mountain creases at a mountain vertex; (c) valley vertex v causes face f to be not flippable; (d) mountain vertex v causes face f to be not flippable. (e) Illustration for Lemma 5.2.

Lemma 5.2. *Suppose vertex v causes face f to be not flippable. Let f_1, f_2, f_3 be the 3 faces incident to v but not adjacent to f, where f_2 is the* middle *face—the one opposite f. Then at least one of f_1, f_2, f_3 is flippable. Furthermore, if f has both a mountain and a valley edge then after flipping one of f_1, f_2, f_3, face f becomes flippable.*

Proof. Suppose without loss of generality that v is a valley vertex and the two edges of f incident to v are mountain creases. See Figure 6(e). Note that v cannot cause any of f_1, f_2, f_3 to be not flippable. Suppose $f_2 = vxy$ is not flippable. Then this is caused by one of its other vertices, say vertex x, and suppose that x is incident to f_1 (otherwise relabel f_1 and f_3). Since xv is a valley, xy must also be a valley, and x must be a mountain vertex, so its other incident edges must be mountains. Then face f_1 has a mountain edge opposite v and is flippable. □

5.1 Reconfiguring a triangle lattice

To show that any locally valid MV assignment of the triangle lattice can be reconfigured to any other using face flips, we use the standard technique of reconfiguring any locally valid MV assignment to a "canonical" MV assignment. To reconfigure A to B via the canonical configuration C, we reconfigure A to C, and then perform the reverse of the reconfiguration sequence that takes B to C. Our canonical configuration C has mountain folds on the $-30°$ lines of the lattice, and valley folds on the other lines—see Figure 7(a).

Theorem 5.3. *Any locally valid MV assignment of the triangle lattice can be reconfigured to the canonical configuration C using $2n$ flips, where n is the number of faces.*

Proof. We use an iterative process, proceeding from left to right, first fixing the diagonal edges down a column (Figure 7(b)), and then fixing the vertical edges down a line (Figure 7(c)), while always preserving the MV assignment in the "completed" region (shown in the gray shaded areas in the figures). We address these two cases separately.

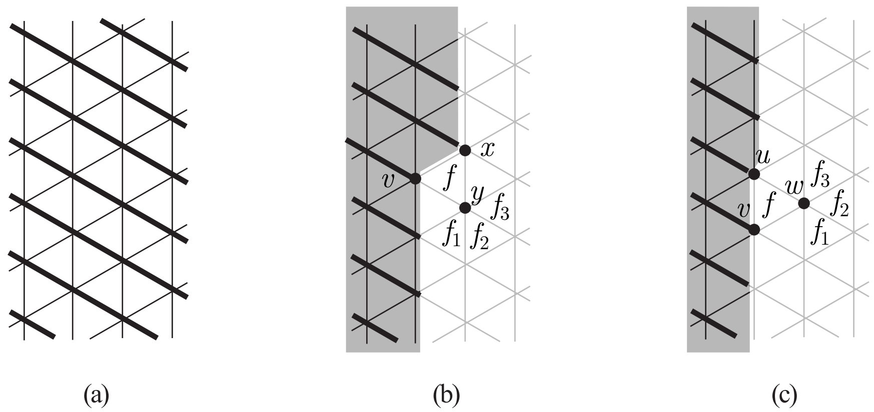

Figure 7: (a) The canonical configuration C. Mountain folds are in bold. (b) Fixing the diagonal edges down a column. The completed region is shaded gray. (c) Fixing the vertical edges down a column.

Case 1. Diagonal edges. Suppose all the diagonal edges down a column match the canonical configuration up until the two edges incident with vertex v on the left-hand side of the column. Let v's neighbors across the column be x and y, with x above y. Because vertex v already has 3 incident valley creases from the completed region, v must be a valley vertex and one of the edges vx, vy must be a mountain and the other a valley.

If edge vx is a valley then vy is a mountain, and the two edges incident to v are correct for the canonical configuration. Thus, we may suppose that vx is a mountain and vy is a valley. Let f be the face vxy. Flipping f corrects the two edges incident to v, so if f is flippable, we are done. Thus, we may suppose that f is not flippable. By Lemma 5.1, f must be unflippable because of vertex x or y.

We first show that x cannot cause f to be unflippable. Suppose it did. Then both edges of f incident to x must be the same; furthermore, they must be mountains because vx is a mountain. Vertex x already had one incident mountain crease from the completed region. But then x has at least 3 incident mountain edges so it must be a mountain vertex, and it cannot prevent f from flipping.

Next, suppose that y causes f to be unflippable. Let f_1, f_2, f_3 be the three faces incident to y and not adjacent to f. By Lemma 5.2 at least one of f_1, f_2, f_3 is flippable. Note that none of the edges of these faces are in the completed region. (See Figure 7(b).) Thus, we can flip one of f_1, f_2, f_3 without disturbing the completed region. Furthermore, f will then be flippable by Lemma 5.2, since f has a mountain and a valley crease.

Case 2. Vertical edges. Suppose that all the vertical edges down a column match the canonical configuration up until the edge uv with u above v. Let the triangle to the right of uv be f, and suppose the f's third vertex is w. If uv is a valley, then it matches the canonical configuration. So, suppose that uv is a mountain. If face f is flippable, that would fix uv, so we may suppose that f is not flippable. We claim that neither u nor v can be the cause. Observe that u and v each have two incident mountain edges. In order for u, say, to cause f to be non-flippable, edge uw must be a mountain like uv. But then u is a mountain vertex so it cannot prevent f from flipping. The same argument applies to v.

Thus, f must be non-flippable because of vertex w. Let f_1, f_2, f_3 be the three faces incident to w and not adjacent to f. By Lemma 5.2 at least one of f_1, f_2, f_3 is flippable. Note that none of the edges of these faces are in the completed region. (See Figure 7(c).) Thus, we can flip one of f_1, f_2, f_3 without disturbing the completed region. Furthermore, f will then be flippable since neither u, v, nor w cause it to be non-flippable. □

5.2 Finding the minimum number of face flips is NP-complete

We show that the problem of finding the sequence of face flips of minimum length between two given crease patterns is NP-complete. For hardness, we reduce from k-Vertex-Cover in max-degree-3 graphs which is NP-complete [19]. Given graph G the problem asks for a subset S of $V(G)$ so that every edge in $E(G)$ is incident to at least one vertex in S, and $|S| \leq k$. We assume that our input is a max-degree-3 graph G embedded in the hexagonal grid. Notice that this does not mean that G is a hexagonal grid graph since the grid is bipartite and admits a polynomial solution for the k-Vertex-Cover problem. Rather, edges in G can be drawn as a path in the grid admitting bends. Such embeddings of polynomial size can be computed in polynomial time [3].

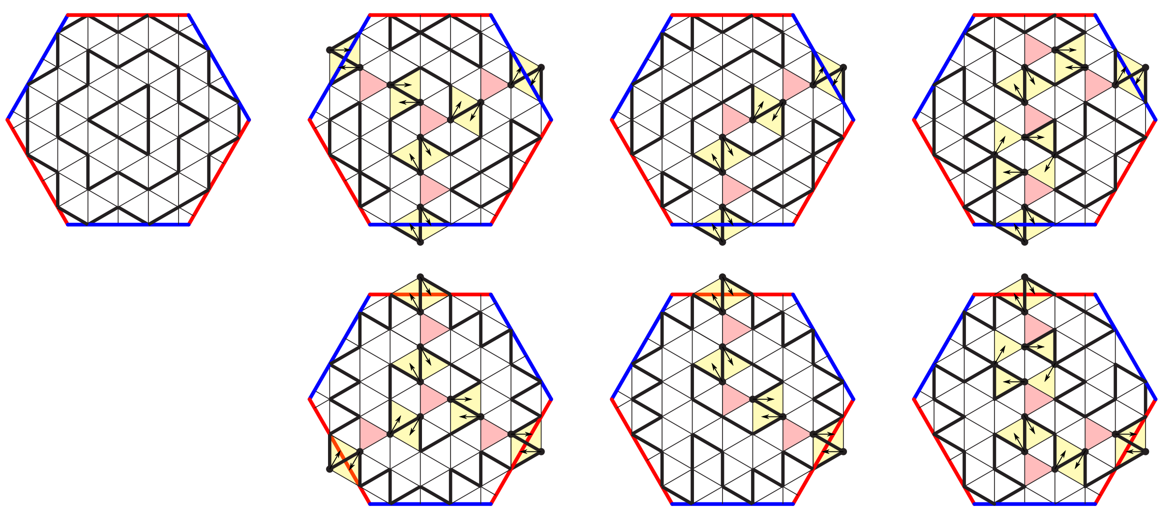

(a) (b) (c) (d)

Figure 8: Reduction from k-Vertex-Cover. (a) Filler, (b) deg-3, (c) deg-2, and (d) bend gadgets. All yellow faces are unflippable and the arrows indicate the vertex that makes it unflippable.

Theorem 5.4. *Given two MV assignments of the triangle lattice A and B, it is NP-complete to decide whether A can be reconfigured into B using at most m face flips.*

Proof. The membership in NP is a consequence of Theorem 5.3. We proceed with the reduction from k-Vertex-Cover. Given a drawing of a graph G in the hexagonal grid, we construct a MV assignment A as follows. We will use the gadgets shown in Figure 8. We first obtain a triangular grid by inserting one *auxiliary* vertex in each hexagon and connecting it to each vertex of the hexagon. We use the triangle grid to tile the plane using hexagons so that each hexagon corresponds to either a vertex of the original hexagonal grid or an auxiliary vertex. If a tile corresponds to a vertex that is not used in the embedding (a degree-3 vertex), we replace it with a *filler* gadget shown in Figure 8 (a) (*deg-3* gadget shown in Figure 8 (b)). If the tile corresponds to a degree-2 vertex (bend), we replace it with the *deg-2* gadget shown in Figure 8 (c) (*bend* gadget shown in Figure 8 (b)) or a 120° rotation. This defines the MV assignment A. B is obtained from A by flipping the assignment of the perimeter of each yellow diamond or triangle as shown in Figure 9. All vertices in the interior of gadgets are locally valid. By construction a blue edge of a gadget will be matched with a red edge. Every vertex on the boundary of a gadget has either zero or an even number of incident mountains and at least one incident valley. A vertex on a red edge is always incident to a mountain crease. Then, the MV assignment A and B are locally valid. We now determine m. Let b be the number of bends in the drawing of G. We set $m = 2k + 8|E(G)| + 14b$.

We now give an informal intuition about the construction and then proceed with the formal proof in the following paragraphs. The assignments A and B only differ on the boundary of the yellow diamonds or triangles. However, by construction, all yellow faces are unflippable. In order to minimize the number of flips we want to flip the minimum number of non-yellow faces so that each vertex that is the tail of an arrow in Figure 8 is adjacent to a flipped face. That allows us to flip the yellow faces. Such a set of non-yellow faces will contain only pink faces. We call the pink central triangle of the deg-3 and deg-2 gadgets *vertex triangles*. We show that whether a vertex triangle is flipped corresponds to whether its corresponding vertex is in a minimum vertex cover.

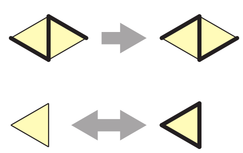

Figure 9: Changes in the MV assignment.

(⇒) Assume that the k-Vertex-Cover instance admits a solution S. We will construct a sequence of at most m flips that brings A into B. Note that all pink triangles are flippable and, because no two are adjacent, flipping any subset of them will not render an unflipped pink triangle not flippable. If a vertex is in S, flip the vertex triangle of its corresponding gadget. Every edge corresponds to a chain of an even number of pink triangles between the vertex triangles of its corresponding endpoints. Since every edge e has at least an endpoint in S, we can flip alternating pink triangles ($1 + 2b_e$ of them where b_e is the number of bends of e) so that every yellow diamond along e has an adjacent flipped triangle. Let P be the set of faces flipped so far. Note that every face colored yellow is initially unflippable, but at the current state at least one face in each diamond or one face in each group of 4 yellow triangles in a bend gadget is flippable. By flipping such a face, an adjacent yellow face becomes flippable. Proceed by flipping all yellow faces. Note that a vertex of a face in P is incident to exactly two yellow faces. Then, every face in P is flippable at the current state. Now, we can flip all faces in P to obtain B. The total number of flips is $2|S| + 8|E(G)| + 14b \leq m$.

($\Leftarrow$) Assume that there exists a sequence F of flips transforming A into B. Notice that it would suffice to flip every yellow face, however, they are unflippable in A. We show that yellow faces are flipped exactly once and we can assume that pink faces are either not flipped, or flipped twice. Assume that F does not flip a yellow face f. Then, by construction of B, there are at least two white or pink faces that share an edge with f that must be flipped an odd number of times. However, flipping those faces create assignments that do not match B, hence requiring further flips. By propagating this argument, we conclude that every white and pink face must be flipped an odd number of times, while yellow faces are either never flipped or flipped an even number of times. The total number of flips is clearly greater then m if $k < |V(G)|$, a contradiction. Therefore, yellow faces are flipped an odd number of times. Since the yellow faces are not flippable in A, at least one face adjacent to each yellow component must be flipped an odd number of times. If there is a white face f_w flipped an odd number of times is adjacent to a pink face f_p that is not flipped, we can change the flipping sequence so that f_p is flipped instead of f_w because the same yellow faces (or a superset) become flippable in both sequences. We can assume that no other faces are flipped since they must be flipped an even number of times and they do not affect whether yellow faces become flippable or not. Except for vertex triangles, flipping a pink face can make a yellow face in at most two different yellow components become flippable. The number of yellow components between vertex triangles corresponding to an edge e is odd and the total number of yellow faces in such components is $6 + 10b_e$ where b_e is the number of bends od e. Then, excluding flips of vertex triangles, at least $8 + 14b_e$ are necessary for each edge e. If an edge achieves such lower bound, then at least one of its endpoints is a vertex triangle that was flipped an even number of times. If more than $8 + 14b_e$ flips were performed for a given edge e, then there is a yellow component adjacent to two pink faces that are flipped at least twice each. We can modify the sequence while not increasing the number of flips so that the only such yellow components are the ones adjacent to vertex triangles. In such a solution, at most k vertex faces can be flipped by the definition of m. Then, we can obtain a solution for the k-Vertex-Cover instance by selecting the vertices corresponding to the vertex faces that are flipped. □

6 Conclusion

We summarize our results in Table 1. The “Flippable Faces" column shows how many of the faces are flippable, where “depends" indicates dependence on the specific locally valid MV assignment being face-flipped.

It is interesting to see how different origami tessellations require such different techniques to analyze their origami flip graph structure. This indicates that face flip configuration space graphs are rich in structure among all flat-foldable crease patterns.

Pattern	*Flippable Faces*	*Config Space*	*Min Complexity*
Square grid	all	connected	linear
Huffman grid	none	maximally disconnected	—
Square twist	half	connected	linear
Miura-ori	depends	connected	poly-time
Triangle lattice	depends	connected	NP-hard

Table 1: A summary of the results presented.

Many questions can be posed for further work. For example: Is the NP-hardness of finding a minimal path in the origami flip graph in the triangle lattice due to the fact that the vertices in the crease pattern have degree 6 (aside from the boundary vertices), as opposed to the tessellations in Sections 3 and 4 whose vertices have degree 4? Are other origami tessellations in the engineering literature, such as [8], amenable to the face flip techniques presented here? Specifically, in reference [8] Barreto's "Mars" is a type of square twist tessellation, and the chicken wire tessellation is very similar to the Miura-ori and yields to similar analysis. But the Yoshimura and quadrilateral meshed patterns are unexplored. More broadly, for a general flat-foldable crease pattern it is not known if finding a minimal path in the origami face flip graph is in NP, as it is conceivable that a crease pattern family could exist that requires an exponential number (say in the number of vertices) of face flips to go from a specific locally valid MV assignment to another. Currently no such crease pattern is known.

Acknowledgments

This work was initiated at the 2018 Bellairs Workshop on Computational Geometry, co-organized by Erik Demaine and Godfried Toussaint. We thank the other participants of the workshop for helpful discussions. We also thank Sarah Nash and Natalya Ter-Saakov for helpful comments on an earlier draft of this work. H. A. A. was partially supported by NSF grants CCF-1422311 and CCF-1423615. D. E. was partially supported by NSF grants CCF-1618301 and CCF-1616248. T. C. H. was supported by NSF grant DMS-1906202.

References

[1] Hugo A. Akitaya, Kenneth C. Cheung, Erik D. Demaine, Takashi Horiyama, Thomas C. Hull, Jason S. Ku, Tomohiro Tachi, and Ryuhei Uehara. Box pleating is hard. In Jin Akiyama, Hiro Ito, Toshinori Sakai, and Yushi Uno, editors, *Discrete and Computational Geometry and Graphs*, pages 167–179, Cham, 2016. Springer International Publishing.

[2] Michael Assis. Exactly solvable flat-foldable quadrilateral origami tilings. *Phys. Rev. E*, 98(3):032112, September 2018. doi:10.1103/PhysRevE.98.032112.

[3] Shabnam Aziza and Therese Biedl. Hexagonal grid drawings: Algorithms and lower bounds. In János Pach, editor, *Graph Drawing: 12th International Symposium, GD 2004, New York, NY, USA, September 29–October 2, 2004, Revised Selected Papers*, pages 18–24, Berlin, Heidelberg, 2005. Springer. doi:10.1007/978-3-540-31843-9_3.

[4] Brad Ballinger, Mirela Damian, David Eppstein, Robin Y. Flatland, Jessica Ginepro, and Thomas Hull. Minimum forcing sets for Miura folding patterns. In *Proc. 26th ACM-SIAM Symp. Discrete Algorithms, San Diego, California, 2015*, pages 136–147, 2015. arXiv:1410.2231, doi:10.1137/1.9781611973730.11.

[5] Marshall Bern and Barry Hayes. The complexity of flat origami. In *Proceedings of the 7th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 175–183, Philadelphia, 1996. SIAM.

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

[7] P. Di Francesco. Folding and coloring problems in mathematics and physics. *Bulletin of the American Mathematical Society (N.S.)*, 37(3):251–307, 2000. doi:10.1090/S0273-0979-00-00870-3.

[8] Thomas A. Evans, Robert J. Lang, Spencer P. Magleby, and Larry L. Howell. Rigidly foldable origami gadgets and tessellations. *Royal Society Open Science*, 2(9):150067, 2015. doi:10.1098/rsos.150067.

[9] Thomas A. Evans, Robert J. Lang, Spencer P. Magleby, and Larry L. Howell. Rigidly foldable origami twists. In Koryo Miura et al., editors, *Origami^6: Proceedings of the Sixth International Meeting on Origami Science, Mathematics, and Education*, pages 119–130, Providence, RI, 2015. American Mathematical Society.

[10] Jessica Ginepro and Thomas C. Hull. Counting Miura-ori foldings. *Journal of Integer Sequences*, 17(10):Article 14.10.8, 2014. URL: https://cs.uwaterloo.ca/journals/JIS/VOL17/Hull/hull.html.

[11] Eric Gjerde. *Origami Tessellations: Awe-Inspiring Geometric Designs*. A K Peters, Wellesley, MA, 2008.

[12] Leslie Ann Goldberg, Russell Martin, and Mike Paterson. Random sampling of 3-colorings in $\mathbb{Z}^2$. *Random Structures & Algorithms*, 24(3):279–302, 2004. doi:10.1002/rsa.20002.

[13] David A. Huffman. Curvature and creases: A primer on paper. *IEEE Trans. Comput.*, C-25(10):1010–1019, 1976. doi:10.1109/TC.1976.1674542.

[14] Thomas C. Hull. The combinatorics of flat folds: a survey. In Thomas C. Hull, editor, *Origami^3 (Asilomar, CA, 2001)*, pages 29–38. A K Peters, Natick, MA, 2002. arXiv:1307.1065.

[15] Matthew Johnson, Dieter Kratsch, Stefan Kratsch, Viresh Patel, and Daniël Paulusma. Finding shortest paths between graph colourings. *Algorithmica*, 75(2):295–321, 2016. doi:10.1007/s00453-015-0009-7.

[16] Jacques Justin. Toward a mathematical theory of origami. In Koryo Miura, editor, *Origami Science & Art: Proceedings of the Second International Meeting of Origami Science and Scientific Origami*, pages 15–29, Otsu, Japan, 1997. Seian University of Art and Design.

[17] Robert J. Lang. Tessellatica 11.1. Software, 2018. URL: https://langorigami.com/article/tessellatica/.

[18] Robert J. Lang. *Twists, Tilings, and Tessellations: Mathematical Methods for Geometric Origami*. A K Peters/CRC Press, Boca Raton, FL, 2018.

[19] David Lichtenstein. Planar formulae and their uses. *SIAM Journal on Computing*, 11(2):329–343, 1982. doi:10.1137/0211025.

[20] Jesse L. Silverberg, Jun-Hee Na, Arthur A. Evans, Bin Liu, Thomas C. Hull, Christian D. Santangelo, Robert J. Lang, Ryan C. Hayward, and Itai Cohen. Origami structures with a critical transition to bistability arising from hidden degrees of freedom. *Nature Materials*, 14(4):389–393, 2015. doi:10.1038/nmat4232.

[21] K. VanderWerf. The Physics, Geometry, and Combinatorics of the Miura-ori. Honors thesis, University of Massachusetts, Amherst, 2014.